



Learning to Approximate Functions Using Nb-Doped SrTiO₃ Memristors

Thomas F. Tiotto^{1,2}, Anouk S. Goossens^{1,3}, Jelmer P. Borst^{1,2}, Tamalika Banerjee^{1,3} and Niels A. Taatgen^{1,2*}

¹ Groningen Cognitive Systems and Materials Center, University of Groningen, Groningen, Netherlands, ² Artificial Intelligence, Bernoulli Institute, University of Groningen, Groningen, Netherlands, ³ Zernike Institute for Advanced Materials, University of Groningen, Groningen, Netherlands

Memristors have attracted interest as neuromorphic computation elements because they show promise in enabling efficient hardware implementations of artificial neurons and synapses. We performed measurements on interface-type memristors to validate their use in neuromorphic hardware. Specifically, we utilized Nb-doped SrTiO₃ memristors as synapses in a simulated neural network by arranging them into differential synaptic pairs, with the weight of the connection given by the difference in normalized conductance values between the two paired memristors. This network learned to represent functions through a training process based on a novel supervised learning algorithm, during which discrete voltage pulses were applied to one of the two memristors in each pair. To simulate the fact that both the initial state of the physical memristive devices and the impact of each voltage pulse are unknown we injected noise into the simulation. Nevertheless, discrete updates based on local knowledge were shown to result in robust learning performance. Using this class of memristive devices as the synaptic weight element in a spiking neural network yields, to our knowledge, one of the first models of this kind, capable of learning to be a universal function approximator, and strongly suggests the suitability of these memristors for usage in future computing platforms.

Keywords: neuromorphic computing, supervised learning, interface memristor, Nb-doped SrTiO₃, neural networks, spiking neural network, function approximation

OPEN ACCESS

Edited by:

Xiaobing Yan,
Hebei University, China

Reviewed by:

Jiyong Woo,
Kyungpook National University,
South Korea
Zhongrui Wang,
The University of Hong Kong,
Hong Kong

*Correspondence:

Niels A. Taatgen
n.a.taatgen@rug.nl

Specialty section:

This article was submitted to
Neuromorphic Engineering,
a section of the journal
Frontiers in Neuroscience

Received: 08 November 2020

Accepted: 24 December 2020

Published: 19 February 2021

Citation:

Tiotto TF, Goossens AS, Borst JP,
Banerjee T and Taatgen NA (2021)
Learning to Approximate Functions
Using Nb-Doped SrTiO₃ Memristors.
Front. Neurosci. 14:627276.
doi: 10.3389/fnins.2020.627276

1. INTRODUCTION

The field of Machine Learning is, at its core, concerned with building function approximators from incomplete data samples. The state of the art approach to solving this problem is using artificial neural networks (ANNs), where a large number of real-valued artificial neurons are connected to each other by means of weights. The neurons in such networks are typically arranged into multiple layers, and are therefore referred to as *deep learning*. The optimization process is performed by updating the weight matrices defining the connection weights between pairs of neurons and is guided by learning rules, which are heuristic optimization algorithms capable of iteratively tuning the network weights to minimize some error function. This process is based on either global (as in the classic back-propagation algorithm) or local knowledge (which is more biologically plausible); the typical outcome is an interpolation for the hidden mapping from input samples to observed outputs.

Traditional neural networks usually require considerable power and are not necessarily constrained by physical limitations. One of the main reasons for this shortcoming is that artificial neural networks are run on traditional Von Neumann architectures in which memory and

computation are not *co-located*. The high energy required by deep learning can be ascribed to the fact that an artificial neural network is essentially a non-Von Neumann computational model, where memory and computation are co-located in connection weight matrices, being simulated on hardware that implements a different computational paradigm.

Even though many important advances in deep learning have been “biologically inspired” (e.g., convolutional neural networks), it is unclear how far the current deterministic approach can progress because of energy requirements, architectural complexity, and capacity to generalize.

To reduce energy, alternative approaches to traditional neural networks aim to implement learning in neuromorphic hardware. Memristors are novel fundamental circuit elements that have attracted a great deal of interest for this purpose because they exhibit multilevel conductance states, giving them—for example when arranged in a crossbar array architecture—the ability to carry out parallel vector-matrix multiplication ($C = \mathbf{aB}$), which is the most important mathematical operation underpinning artificial neural network implementations (Xia and Yang, 2019). The ability to carry out this operation, where the input is voltage V , output is current I , and the weight corresponds to conductance G (the inverse of resistance R), follows from Ohm’s law $I = VG$ and Kirchoff’s current law $I_j = \sum I_i$. Due to their non-volatile nature and ability to co-locate memory and computing, memristors can enable parallel physical computation within a low energy envelope.

Memristors have also been used as direct hardware implementations of weights in order to replicate standard deep learning software architectures, such as long short-term memory (LSTM) (Li et al., 2019). Such systems can then be trained using standard gradient-based machine learning algorithms but in doing so limit themselves to the mainstream computational paradigm and do not take full advantage of the material properties of the memristive devices.

Another alternative approach to computation in which the use of memristors is currently being explored is that of *reservoir computing* (Lukoševičius and Jaeger, 2009). A reservoir is a collection of recurrently interconnected units exhibiting short-term memory and non-linear response to inputs; the connections of the reservoir are fixed and the training of the network is focused on a simple feed-forward readout layer thus reducing training time and complexity. Previous literature has shown that memristor-based reservoir computers are capable of excellent performance on a variety of tasks by taking advantage of the intrinsic characteristics of the physical devices. A recurrent connection is a way of adding temporal dynamics to a model, but a memristor already is such a system. Existing memristor-based reservoirs have thus eschewed the recurrent connections between the reservoir units and relied on the fact that programming pulses at different moments in time have varying effect on the state of the device (Du et al., 2017; Midya et al., 2019; Moon et al.,

2019). *Diffusive* memristors—whose memristance is governed by fast diffusive species have also seen use as reservoir units (Wang et al., 2017b), but also as artificial neurons (Wang et al., 2018b), as their dynamics are well-suited to representing the leaky integrate-and-fire model of neuronal computation.

The biological parallel with the brain can be taken further by substituting continuous idealized neurons, found in traditional artificial neural networks, for their spiking equivalent and by using a learning algorithm which updates weights on the basis of local knowledge. For example, considerable interest has been devoted to using memristors to implement Hebbian learning rules which are thought to underpin the capacity of the brain to learn in an unsupervised manner and adapt to novel stimuli. Multiple flavors of spike timing-dependent plasticity (STDP) and hardware architectures have been proposed as ways to make good use of the physical characteristics of memristors, by appropriately shaping voltage pulses (Serrano-Gotarredona et al., 2013; Ambrogio et al., 2016), or by controlling second-order state variables (Kim et al., 2015).

The most widely studied memristors for neuromorphic applications are ones relying on electric field control of nanoscale filaments between two electrodes (Seo et al., 2011; Hu et al., 2018; Moon et al., 2018; Wang et al., 2018a), phase transitions (Kuzum et al., 2012; Ambrogio et al., 2016; Wang et al., 2017a), or voltage induced control of the ferroelectric configuration (Nishitani et al., 2012; Kim and Lee, 2019; Oh et al., 2019). A less studied class of materials for synaptic device applications are interface-type memristors where electric field controlled resistive switching results from changes occurring at interfaces. While the resistive switching in for example metal/Nb-doped SrTiO₃ (Nb:STO) Schottky junctions, is well-documented in literature, there are not many reports in which they are considered as individual neuromorphic components (Yin et al., 2016; Jang et al., 2018; Zhao et al., 2019), and to the best of our knowledge, no reports of their use as components in neural networks. Often, memristive devices require forming processes, which can be unfavorable for device performance and network integration (Amer et al., 2017; Kim et al., 2019). An attractive feature of Nb:STO memristors is that the switching behavior is present in the as-fabricated device. In addition, they also show reproducible multi-level resistive switching and low reading currents at room temperature, making them potentially ideal as components in neuromorphic computers. Such computers need to be able to deal and operate within the physical constraints of the particular class of memristor they utilize as hardware substrate. The heterogeneous characteristics of memristive devices also entails that a neuromorphic computer be aware of their physical characteristics and responses, in order to be able to utilize them to their full potential.

Here, we utilized Nb:STO memristors as synapses in a simulated spiking neural network to explore the suitability of this class of memristors as computing substrate. Measurements were conducted on physical devices, in which the change in resistance in response to a series of voltage pulses was monitored. We found the resistance values in response to forward voltage pulses to follow a power law, whereby each pulse gives rise to an amplitude-dependent decrease in the

Abbreviations: LIF, leaky integrate-and-fire (neurons); mPES, memristor prescribed error sensitivity; MSE, mean squared error; Nb, Niobium; Nb:STO, Nb-doped SrTiO₃; NEF, neural engineering framework; SNN, spiking neural network; SrTiO₃, strontium titanate; ρ , Spearman correlation coefficient.

resistance. Next, we built a simulated model using the Nengo Brain Builder framework, consisting of pre- and post-synaptic neuronal ensembles arranged in a fully-connected topology. Each artificial synapse was composed of a *differential synaptic pair* of memristors, one designated as “positive” and one as “negative,” and the weight of the connection was given by the difference in normalized conductance values between the two paired memristors. Using a training process during which discrete voltage pulses were applied to one of the two memristors in each pair, we showed that our model was capable of learning to approximate transformations from randomly and periodically time-varying signals, suggesting that it could operate as a general function approximator. The initial state of the memristive devices was unknown, as were their exact operating parameters due to the addition of noise; this simulated the same kind of constraints faced in hardware, making it likely that the learning model could be ported to a physical implementation. Most importantly, robust learning performance was proven using only discrete updates based on local information.

2. MATERIALS AND METHODS

2.1. Device Fabrication

Memristive devices are material systems which undergo reversible physical changes under the influence of a sufficiently high voltage, giving rise to a change in resistance (Strukov et al., 2008; Brivio et al., 2018). Here, metal/Nb-doped SrTiO₃ devices were used, where resistive switching results from changes occurring at the Ni/Nb:STO interface (Sim et al., 2005; Seong et al., 2007; Sawa, 2008; Rodenbücher et al., 2013; Mikheev et al., 2014; Goossens et al., 2018). We fabricated devices on (001) SrTiO₃ single crystal substrates doped with 0.01 wt% Nb in place of Ti (obtained from Crystec Germany). To obtain a TiO₂ terminated surface, a wet etching protocol was carried out using buffered hydrofluoric acid, this was followed by *in-situ* cleaning in oxygen plasma. To fabricate the Schottky junctions, a layer of Ni (20 nm), capped with Au (20 nm) to prevent oxidation, was grown by electron beam evaporation at a room temperature at a base pressure of $\approx 10^{-6}$ Torr. Junctions with areas between 50×100 and $100 \times 300 \mu\text{m}^2$ were defined using UV-photolithography and ion beam etching. Further UV-photolithography, deposition and lift-off steps were done to create contact pads and finally wire bonding was carried out to connect the contacts to a chip carrier. Current-voltage measurements were performed using a Keithley 2410 SourceMeter in a cryostat under vacuum conditions (10^{-6} Torr). All electrical measurements were done using three-terminal geometry with a voltage applied to the top Ni electrode. The measurements in the present work are performed on junctions of $20,000 \mu\text{m}^2$. More details on the measurement techniques and device characterization can be found in Goossens et al. (2018). In light of practical applications and due to the attractive lowering of reverse bias current with increasing temperatures (Susaki et al., 2007; Goossens et al., 2018), all measurements were done at room temperature.

2.2. Device Experimental Evaluation

Biological synapses strengthen and weaken over time in a process known as *synaptic plasticity*, which is thought to be one of the most important underlying mechanisms enabling learning and memory in the brain. For a device to be useful as a neural network component, its resistance should thus be controllable through an external parameter, such as voltage pulses—so that it may take on a range of values (within a certain window). The as-fabricated Ni/Nb-doped SrTiO₃ memristive devices showed stable and reproducible resistive switching without any electroforming process.

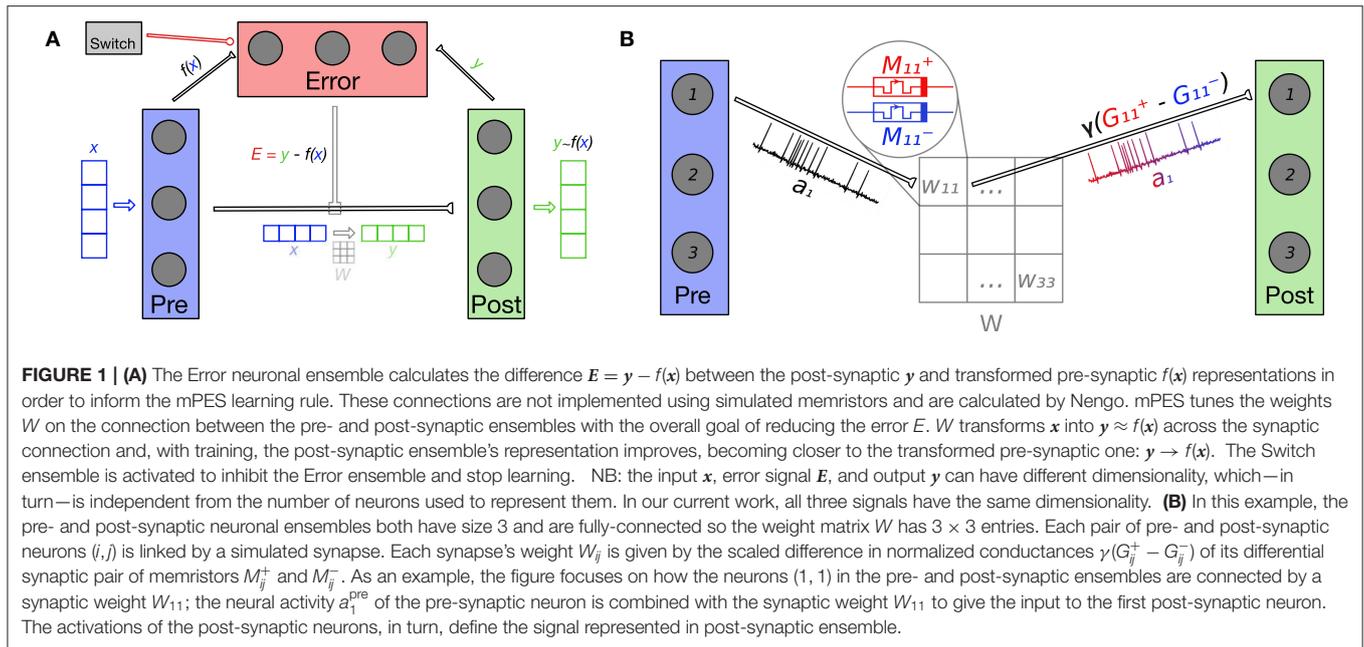
Hence, to investigate the effect of applying pulses across the interface, we conducted a series of measurements in which a device was subjected to a SET voltage of +1 V for 120 s to bring it to a low resistance state and reduce the influence of previous measurements. Then, 25 RESET pulses of negative polarity (−4 V) were applied and the current was read after each pulse. A negative read voltage was chosen because significantly larger hysteresis is observed in reverse bias compared to forward bias. Because of differences in the charge transport mechanisms under forward and reverse bias, a much smaller current flows in this regime: to ensure the measured current is sufficiently large to be read without significant noise levels, a reading voltage of −1 V was chosen. The RESET pulse amplitude was varied from −2 to −4 V. This procedure was repeated several times for each amplitude sequentially. The pulse widths were ≈ 1 s.

Next, we performed a set of measurements in which devices were SET to a low resistance state by applying +1 V for 120 s before applying 10 RESET pulses of −4 V to bring devices to a depressed state. This was followed by applying a series of potentiation pulses ranging from +0.1 to +1 V.

2.3. Nengo Brain Maker Framework

The aforementioned devices were integrated into a simulated spiking neural network (SNN); this was implemented using the Nengo Brain Maker Python package (Bekolay et al., 2014), which represents information according to the principles of the Neural Engineering Framework (NEF) (Eliasmith and Anderson, 2003). Nengo was chosen because its two-tier architecture enables it to run on a wide variety of hardware backends—including GPUs, SpiNNaker, Loihi—with minimal changes to user-facing code, opening up the possibility of running models on a variety of neuromorphic hardware (Furber et al., 2014; Davies et al., 2018). The results of our current work are framework-independent and would apply to any computational model congruent to the one we simulated in Nengo.

In Nengo, information is represented by real-valued vectors which are *encoded* into neural activity and *decoded* to reconstruct the original signal (Bekolay et al., 2014). The representation is realized by *neuronal ensembles*, which are collections of neurons representing a vector; the higher the number of neurons in the ensemble, the better the vector can be encoded and decoded/reconstructed. Each neuron i in an ensemble (which in our case are leaky integrate-and-fire (LIF) neurons) has its own *tuning curve* that describes how strongly it spikes in response to particular inputs (*encoding*). The collective spiking activity of the neurons in an ensemble represents the input vector, which can



be reconstructed in output by applying a temporal filter to each spike train and linearly combining these individual contributions (*decoding*). Thus, each neuronal ensemble represents a vector through its unique neuronal activation pattern through time.

Ensembles are linked by *connections* which transform and transmit the encoded vector represented in the *pre-synaptic ensemble* by means of a *weight matrix*. The *post-synaptic ensemble* will thus be trying to render a *transformation* of the vector represented in the pre-synaptic neuronal ensemble; in the default case Nengo attempts offline optimization to find a series of optimal weights to approximate the function along the connection.

Learning rules can be applied to the connection between two ensembles in order to iteratively modify its weighting and thus find the correct transformation in an *online* manner.

2.4. Simulated Model

The Nengo model used to evaluate the usage of our memristive devices in a neuromorphic learning setting, consisted of three neuronal ensembles: pre- and post-synaptic, and one for calculating an error signal E . The pre- and post-synaptic ensembles were linked by a connection which resulted in a fully-connected topology between their constituent neurons; this is equivalent to the physical arrangement of memristors into *crossbar arrays* (Xia and Yang, 2019) used in many previous works to realize efficient brain-inspired computation (Hu et al., 2014; Li et al., 2019). The specific topology, which is crucial in enabling the model to learn, is shown in **Figure 1A**.

The pre-synaptic ensemble had as input signal a d -dimensional vector x and the connection between the ensembles was tuned in order for the post-synaptic ensemble to represent a transformation/function $f: x \rightarrow y$ of the vector encoded into the pre-synaptic neuronal population. This time-varying input

signal consisted of either d uniformly phase-shifted sine waves x_d described by:

$$x_d = \sin\left(\frac{1}{4}2\pi t + i\frac{2\pi}{d}\right), i \in [0, d] \quad (1)$$

or of d white noise signals low-filtered by a 5 Hz cutoff. The latter signal was generated by using the `nengo.processes.WhiteSignal` class and is naturally periodic but, in order to be able to test network on unseen data, the period was chosen to be double the total simulation time of 30 s.

The third neuronal ensemble, dedicated to calculating the error signal, was connected to the pre- and post-synaptic ensembles using standard Nengo connections whose weights were pre-calculated by the framework. The connection from the post-synaptic ensemble implemented the identity function—simply transferring the vector—while the one from the pre-synaptic ensemble calculated $-f(x)$ with f being the transformation of the input signal we sought to teach the model to approximate. The connection from pre to error ensembles did indeed already represent the transformation $f(x)$ that the model needed to learn, but the optimal network weights for this connection were calculated by Nengo before the start of the simulation based on the response characteristics of the neurons in the connected ensembles. Therefore, the weights realizing the transformation on this connection were not learned from data and did not change during the simulation. Given these connections from pre and post-ensembles, the error ensemble represented a d -dimensional vector $E = y - f(x)$ which was used to drive the learning in order to bring the post-synaptic representation y as close as possible to the transformed pre-synaptic signal $f(x)$, the *ground truth* in this *supervised learning* context.

The learning phase was stopped after 22 s of simulation time had passed; this was done by activating a fourth neuronal ensemble which had strong inhibitory connections to the error ensemble. Suppressing the firing of neurons in the error ensemble stopped the learning rule on the main synaptic connection from modifying the weights further and thus let us test the quality of the learned representation in the remaining 8 s of simulation time. An example simulation run is shown in **Supplementary Figures 2, 3**.

In separate experiments, the transformation f for the model to learn was either $f(x) = x$ or $f(x) = x^2$. Learning the square function was expected to be a harder problem to solve as $f(x) = x^2$ is not an *injective* function; i.e., different inputs can be mapped to the same output [$(-1)^2 = 1^2 = 1$, for example]. Even learning the identity function $f(x) = x$ is not a trivial problem in this setting, as the pre- and post-synaptic ensembles are not congruent in how they represent information, even when they have the same number of neurons. Each of the ensembles' neurons may respond differently to identical inputs, so the same vector decoded from the two ensembles will be specified by different neuronal activation patterns. Therefore, the learning rule has to be able to find an optimal mapping between two potentially very different vector spaces even when the implicit transformation between these is the identity function.

The quality of the learned representation was evaluated on the last 8 s of neuronal activity by calculating the mean squared error (MSE) and Spearman correlation coefficient ρ between the ground truth given by the transformed pre-synaptic vector $f(x)$, and the y vector represented in the post-synaptic ensemble.

At a lower level, each artificial synapse—one to connect each pair of neurons in the pre- and post-synaptic ensemble—in our model was composed of a “positive” M^+ and a “negative” M^- simulated memristor (see **Figure 1B**). The mapping of the current resistance state R^\pm of a pair of memristors to the corresponding synaptic network weight ω was defined as:

$$\begin{aligned}\omega &= \gamma \left[\left(\frac{\frac{1}{R^+} - \frac{1}{R_1}}{\frac{1}{R_0} - \frac{1}{R_1}} \right) - \left(\frac{\frac{1}{R^-} - \frac{1}{R_1}}{\frac{1}{R_0} - \frac{1}{R_1}} \right) \right] \\ &= \gamma \left[\left(\frac{G^+ - G_0}{G_1 - G_0} \right) - \left(\frac{G^- - G_0}{G_1 - G_0} \right) \right]\end{aligned}\quad (2)$$

Note that the relationship between conductance and resistance was specified by $G = \frac{1}{R}$; thus, when normalizing, the maximum conductance G_1 was given by the inverse of the minimum resistance $\frac{1}{R_0}$, and vice-versa. So, the network weight of a synapse at each timestep was the result of the difference between its current memristor conductances G^\pm , which were normalized in the range $[0, 1]$, and then multiplied by a gain factor γ .

To simulate device-to-device variation and hardware noise we elected to introduce randomness into the model in two ways. Firstly, we initialized the memristor resistances to a random high in the range $[10^8 \Omega \pm 15\%]$ in order to ensure that all the weights were not 0 at the start of the simulation. This value was chosen on the basis of the experimental measurements shown in **Figure 4**, it being a high-resistance state the devices converged to after the

application of -4 V RESET pulses. We can therefore imagine the simulated memristors being brought to their initial resistances by the application of several -4 V RESET pulses prior to the start of the training phase. As each entry in weight matrix W was given by the difference in normalized conductances of a pair of memristors, initializing them all to the same resistance $10^8 \Omega$ would have led them all to necessarily be equal to 0, as per Equation (2). In Machine Learning parlance, the act of assigning a random initial value network parameters is known as “symmetry breaking.” Secondly, we perturbed, via the addition of 15% Gaussian noise, the parameters in Equation (8) used to calculate the updated resistance states of the simulated devices. This was done independently for every simulated memristive device.

2.5. Learning Rule

The process through which learning is effected in an artificial neural network is the iterative, guided modification of the network weights connecting the neurons. To this end, we designed a neuromorphic learning rule that enacted this optimization process by applying SET pulses to our memristive devices.

In order to learn the desired transformation from the pre-synaptic signal, we iteratively tuned the model parameter matrix W , which represented the transformation across the synaptic connection between the pre- and post-synaptic ensembles. This was done by adjusting the resistances R^\pm of the memristors in each differential synaptic pair through the simulated application of $+0.1$ V SET learning pulses guided - at each timestep of the simulation - by a modified version of the supervised, pseudo-Hebbian prescribed error sensitivity (PES) learning rule (MacNeil and Eliasmith, 2011); this extension of PES to memristive synapses will be referred to as *mPES* from here on.

The PES learning rule accomplishes online error minimization by applying at every timestep the following equation to each synaptic weight ω_{ij} :

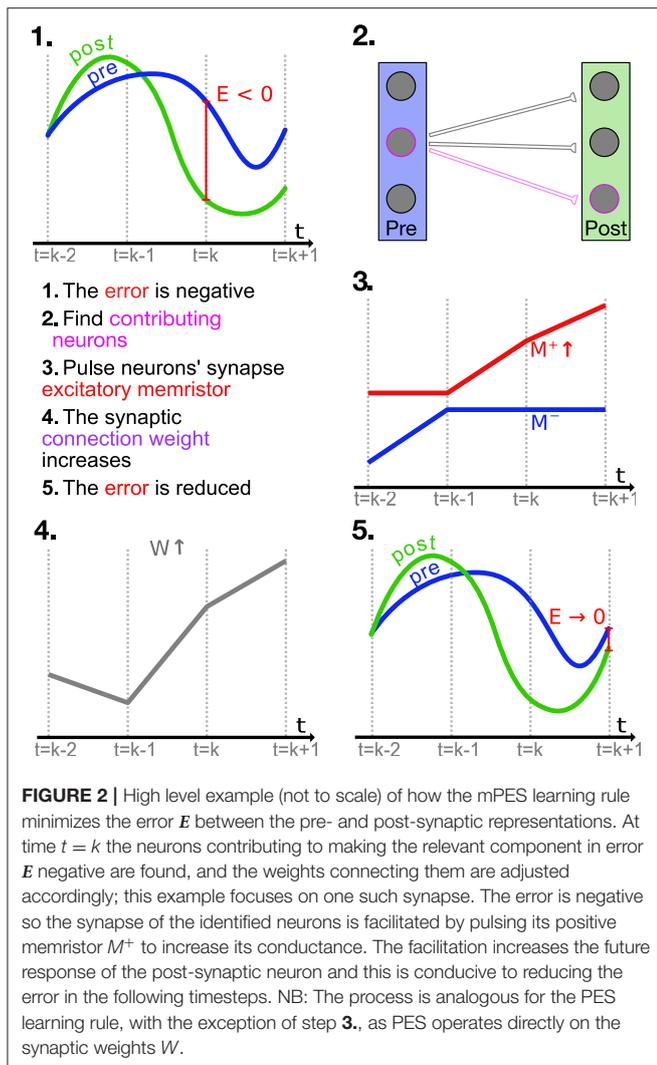
$$\Delta\omega_{ij} = \kappa\alpha_j e_j \mathbf{E} a_i \quad (3)$$

where ω_{ij} is the weight of the connection between pre-synaptic neuron i and post-synaptic neuron j , κ is the learning rate, α_j the gain factor for neuron j , e_j the encoding vector for neuron j , \mathbf{E} the global d -dimensional error to minimize, and a_i the activity of pre-synaptic neuron i . α_j and e_j are calculated for each neuron in the post-synaptic ensemble by the Nengo simulator such that:

$$a_j = G(\alpha_j e_j \cdot \mathbf{x} + b_j) \quad (4)$$

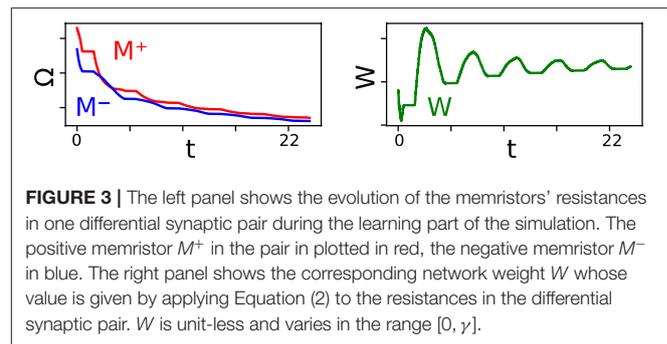
with G the neuronal non-linearity, \mathbf{x} the input vector to the ensemble, and b_j a bias term.

The factor $\epsilon = \alpha_j e_j \mathbf{E}$ is analogous to the *local error* in the backpropagation algorithm widely used to train artificial neural networks (Le Cun, 1986). The overview shown in **Figure 2** applies to both the PES and mPES learning rules, with the exception of step 3 that is characteristic to mPES; in this high-level depiction we can see how the neuronal pairs i, j whose $\omega_{ij} > 0$ are the ones which contribute to the post-synaptic representation being lower than the pre-synaptic signal and



should thus be the ones potentiated in order to bring the post-synaptic representation closer to the pre-synaptic one. The converse is also true, so the synapses whose corresponding $\omega_{ij} < 0$ should be depressed in order to reduce E . Thus, the neurons in the post-synaptic ensemble that are erroneously activating will be less likely to be driven to fire in the future, changing the post-synaptic representation toward a more favorable one.

mPES is a novel learning rule developed to operate under the constraints imposed by the memristive devices and essentially behaves as a discretized version of PES. This discretization is a natural consequence of the models it operates on not having ideal continuous network weights but instead, by these being given by the state of memristors, which were updated in a discrete manner through voltage pulses. The gradual change in the weights W between pre- and post-synaptic neuronal ensembles in the model can be classified as a *self-supervised learning* optimization process, as the learning was informed by the error E vector calculated by the third neuronal ensemble—which was also a part



of the model. Our learning rule had to be able to operate in a very restricted and stochastic environment given that we only supplied $+0.1$ V SET potentiation pulses to our memristive devices during the online training process, and that we introduced uncertainty on the initial resistance and update magnitude of the memristors. Furthermore, at most one memristor in each synapse could have its resistance updated (either M^+ or M^-) and the exact result of this operation was also uncertain due to the fact that we injected noise into the parameters in Equation (8) governing the update. **Figure 3** shows how the resistances of the memristors in one such differential synaptic pair are modulated during training, and how the corresponding weight evolves over time. An interesting characteristic stemming from the power-law memristance behavior is that each SET pulse has a smaller effect than the one preceding it; the effect this has on the network weight is to help it converge to an optimum in a manner very similar to the cooling in *simulated annealing* (Kirkpatrick et al., 1983), as can be seen for W in **Figure 3**. This kind of behavior was seen for most—if not all—of the synapses in our simulations.

The learning rule is gated, pseudo-Hebbian, and uses discrete updates and local information, therefore making it biologically plausible. Our claim that the learning rule is local is substantiated by the fact that only information available to post-synaptic neuron j is used when calculating the weight update $\Delta\omega_{ij}$. The key point is that each neuron has a different random encoder and is therefore responsive to a different portion of the error vector space. Each update is thus optimizing a separate local error ϵ which depends on the neural state of j ; there is no collaboration between units toward optimizing a global error function as in classical gradient methods. The algorithmic approach is akin to that presented in, e.g., Mostafa et al. (2018). Specific to our current setup is also the fact that we have no *hidden layers* in the network and therefore all information is necessarily local, there being no intermediaries. The characterization as a gated learning rule derives from the fact that when $E = 0$ learning does not occur, while the pseudo-Hebbian element is given by the post-synaptic activity being represented indirectly by $\alpha_j e_j$ instead of explicitly by a_j .

mPES required the definition and tuning of a hyperparameter γ , which is analogous to the *learning rate* found in most machine learning algorithms. γ represents the gain factor used to scale the memristors' conductances G^\pm to network weights W whose magnitude was compatible with the signals the model needs

to represent. A larger γ made each pulse to the memristors have an increased effect on the network weights but, unlike a learning rate κ , the magnitude of each update on the underlying objects was not changed by γ , affecting only the translation to network weights.

Our learning rule also presented an *error threshold* θ_ϵ hyperparameter whose role was to guarantee that a very small error E would not elicit updates on the memristors, given that this was seen to negatively impact learning. At this stage, the value for the error threshold was experimentally determined but kept constant and considered a static part of the model rather than a hyperparameter.

The mPES learning rule went through the following stages at each timestep of the simulation, in order to tune the resistances of each memristor in every differential synaptic pair with the overall goal of minimizing the global error E :

1. The *global error* signal vector E was taken as input to the learning rule.
2. A *local error* ϵ vector was calculated as $\epsilon = -e_j E$. ϵ is a projection of E onto each neuron's response and, consequently, encodes the contribution of each post-synaptic neuron to the global error E (Note that the sign of E is inverted).
3. If no value in the ϵ vector was greater than an experimentally determined *error threshold* $\theta_\epsilon = 10^{-5}$ the updates in this timestep were skipped. The reason for this check was 2-fold. Firstly, if we didn't filter for local errors very close to 0, the learning performance would suffer. This is because a very small error may have no consequence in a continuous-update setting, as that where original PES is applied, but in our setup it could trigger a voltage pulse on the corresponding memristor and bring it to a new resistance state. As memristors move between discrete resistance states, such an update would end up having an outsized effect. Secondly, inhibiting the error ensemble to terminate the learning phase did not completely shut off its activity so E would never be exactly $\mathbf{0}$ and the adjustments to the resistances would not stop, thus leading to the learning phase going on indefinitely.
4. A *delta* matrix $\Delta = -\epsilon \otimes \mathbf{a}^{\text{pre}}$, whose entries were the equivalent to the PES $\Delta\omega_{ij}$ in Equation (3), was calculated. Each entry in Δ has a one-to-one correspondence to a synapse in the model and encodes the participation of each pair of pre- and post-synaptic neurons (i, j) to the error E . Therefore, the sign of each entry Δ_{ij} determines if the corresponding synapse should be potentiated ($\Delta_{ij} > 0$) or depressed ($\Delta_{ij} < 0$).
5. The \mathbf{a}^{pre} vector of pre-synaptic activations' elements were discretized into binary values 0 and 1, as the only information of interest was if a pre-synaptic neuron had spiked, not the intensity of its activity.
6. The update direction V for each synapse was calculated as $V = \text{sgn}(\Delta)$, in order to determine which memristor in each pair needed adjusting:
 - a. The positive memristor M_{ij}^+ in a differential synaptic pair was pulsed when the corresponding term $V_{ij} > 0$.
 - b. The negative memristor M_{ij}^- was pulsed when $V_{ij} < 0$.
7. Finally, Equation (2) was applied to each synapse in order to convert from M^\pm 's resistances to network weights W .

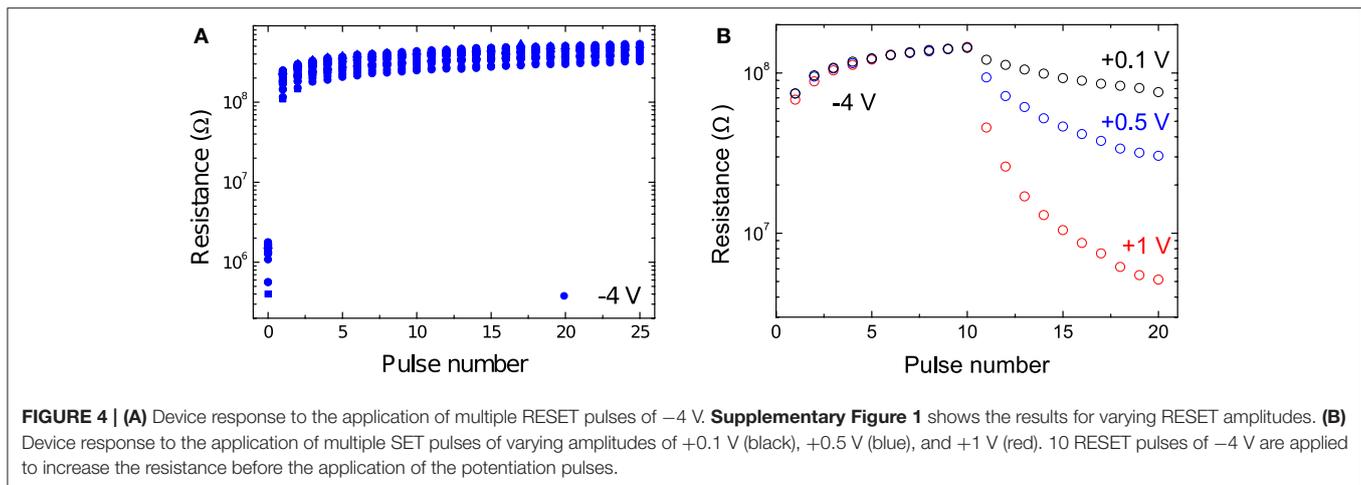
2.6. Optimization Experiments

In order to assess the effect that the gain γ hyperparameter had on the learning capacity of the model, we set up a series of eight experiments, one for each combination of the three binary parameters: input function (sine wave or white noise), learned transformation ($f(\mathbf{x}) = \mathbf{x}$ or $f(\mathbf{x}) = \mathbf{x}^2$), and ensemble size (10 or 100 pre- and post-synaptic neurons). In each experiment we ran 100 randomly initialized models for each value of γ logarithmically distributed in the range $[10, 10^6]$ and recorded the average mean squared error and Spearman correlation coefficient measured on the testing phase of each simulation, during the final 8 s.

2.7. Learning Experiments

To assess the learning capacity of our model and learning rule (mPES) we compared them against a network using prescribed error sensitivity (PES) as learning algorithm and continuous synapses together with noiseless weight updates. The default model with mPES used to generate our results was run with 15% of injected Gaussian noise on both the memristor parameters and the initial resistances, as previously described. To show that our model was not just learning the transformed input signal $f(\mathbf{x})$ but the actual function f , we tested both on the same signal that had been learned upon, and on the other. In other words, the model learned the $f(\mathbf{x})$ transformations from both a sine wave and a white noise signal, using either 10 or 100 neurons, and the quality of the learning was then tested on both functions. Both input learning and testing signals were 3-dimensional in all cases.

To have a quantitative basis to compare the two models, we measured the mean squared error (MSE) together with Spearman correlation ρ between the pre- and post-synaptic representations in the testing phase of the simulation (corresponding to the last 8 of the total 30 s). We took PES as the "golden standard" to compare against, but also looked at what the error would have been when using the memristors without a learning rule modifying their resistances. A lower mean squared error indicates a better correspondence between the pre- and post-synaptic signals, while a correlation coefficient close to +1 describes a strong monotonic relationship between them. A higher MSE-to- ρ ($\frac{\rho}{\text{MSE}}$) ratio suggests that a model is able to learn to represent the transformed pre-synaptic signal more faithfully.



2.8. Other Experiments

We also tested the learning performance of the network for varying amounts of noise, as defining models capable of operating in stochastic, uncertain settings is one of the fundamental goals of the field on Neuromorphic Computing. To this end, we defined the simplest model possible within our learning framework (10 neurons, learning identity from sine input) and used it to evaluate the learning performance for varying amounts of Gaussian noise and for different values of the c parameter in the resistance update Equation (8).

We ran 10 randomly initialized models for each of 100 levels of noise—expressed as coefficient of variation $\frac{\sigma}{\mu}$ —in the range $[0, 100\%]$ and reported the averaged MSE-to- ρ ratio for each noise percentage. The Gaussian noise was added in the same amount to the R_0 , R_1 , and c parameters in Equation (8) as well as to the initial resistances of the memristors. The methodology we followed closely resembled that reported in Querlioz et al. (2013), but it should be noted that, in doing so, the magnitude of the uncertainty of each memristance update depends on the devices' current resistance state; it is presently unclear if this method best models the actual behavior that would be seen in a physical device, but it suffices for the scope of this work.

Using the same methodology, we also ran a *parameter search* for the exponent c in Equation (8) in order to find the SET voltage magnitude best able to promote learning. To realize this, we evaluated the MSE-to- ρ ratio of randomly initialized models with 15% noise for 100 values of c uniformly distributed in the range $[-1, -0.0001]$.

3. RESULTS

3.1. Device Experimental Evaluation

We investigated the effect of applying pulses across the interface by administering a long SET pulse followed by 25 RESET pulses. Due to the resistance's dependence on the memristor's previous history, there is some variation in the starting state, shown by the measurement at pulse number 0. **Figure 4A** shows the results for pulses of -4 V in amplitude. Results for other pulse amplitudes can be found in **Supplementary Figure 1**. The first pulse gave rise

to the largest increase in resistance, with subsequent ones having a much smaller effect. The change in resistance quickly leveled off and the influence of subsequent pulses was significantly smaller. The application of a RESET pulse resulted in a switch to a high resistance state that depended strongly on the amplitude of the applied pulse, but not particularly on the number of pulses of that amplitude that were applied.

We also explored the devices' response to SET voltages. It should be pointed out that in **Figure 4B** the initial state (and hence also the large difference induced by the first RESET pulse) after applying a long SET voltage, which would correspond to pulse number 0, is not shown. Note that different pulse amplitudes were chosen for SET and RESET pulses because the asymmetric nature of the charge transport results in much larger currents in forward bias (Goossens et al., 2018). With the application of positive pulses, the device saw its resistance gradually decrease, i.e., it was potentiated. It is clear that in this case both the pulse amplitude and the number of applied pulses had a great impact on the resistance state of the device. The larger the amplitude of the SET pulse, the greater the induced difference to the resistance with each applied pulse. Interesting, compared to applying RESET pulses, the difference between the change induced by the first and later pulses was not severe. While 10 applied RESET pulses gave rise to close to saturated high resistance states, this was not observed when positive pulses were applied.

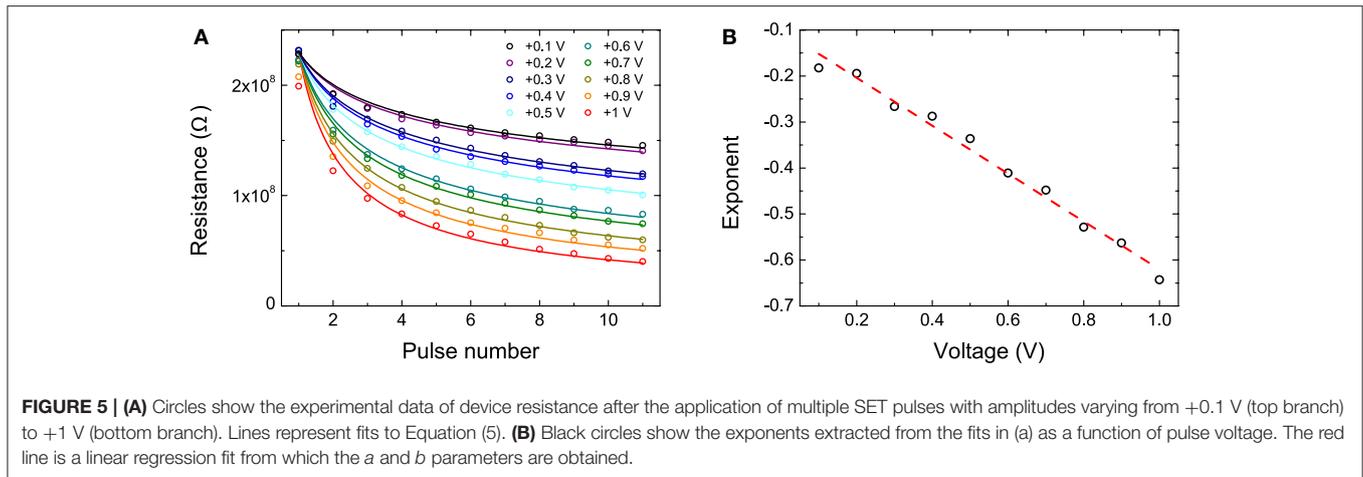
3.2. Device Memristive Response

In order to simulate our memristive devices in Nengo, we modeled their behavior in response to SET pulses on the basis of the experimental measurements we carried out.

The device behavior when applying a series of SET pulses, i.e., its *forward bias* response, was seen to be well-described by an exponential equation of the form:

$$R(n, V) = R_0 + R_1 n^{a+bV} \quad (5)$$

in which V represented the amplitude of the SET pulse, n the pulse number, R_0 the lowest value that the resistance $R(n, V)$



could reach, and $R_0 + R_1$ the highest value. One of the reasons we chose a fit of this form was because of the parallel we saw with the classic *power law of practice*, a psychological and biological phenomenon by virtue of which improvements are quick at the beginning but become slower with practice. In particular, skill acquisition has classically been thought to follow a power law (Newell and Rosenbloom, 1981; Hill et al., 1991).

By solving Equation (5) for n , we can calculate the pulse number from the current resistance $R(n, V)$ by:

$$n = \left(\frac{R(n, V) - R_0}{R_1} \right)^{\frac{1}{a+bV}} \quad (6)$$

Parameters a and b were found based on the data measured by applying SET voltages between +0.1 and +1 V to the memristive devices, as shown in **Figure 5A**. The exponents of the curves best describing the memristors' behavior were then fitted with linear regression on the log-transformed pulse numbers and resistances, shown in **Figure 5B**, in order to obtain a linear expression $a + bV$. This process yielded an estimated best fit for the memristor behavior of:

$$R(n, V) = 200 + 2.3 \times 10^8 n^{-0.093-0.53V} \quad (7)$$

In our current work we always supplied SET pulses of +0.1 V so we could subsume the exponential term $a + bV$ into a single parameter c :

$$\begin{aligned} R(n) &= R_0 + R_1 n^c \\ &= 200 + 2.3 \times 10^8 n^{-0.093-0.53 \times 0.1} \\ &= 200 + 2.3 \times 10^8 n^{-0.146} \end{aligned} \quad (8)$$

3.3. Optimization Experiments

The gain factor γ in Equation (2) was analogous to the *learning rate* κ in Equation (3)—present in most machine learning algorithms—in that it defined how big an effect each memristor conductance update had on the corresponding network weight.

The rationale behind this experiment was to execute an equivalent of hyperparameter tuning—routinely done for

artificial neural networks—as we had realized that γ was homomorphous to a learning rate. Therefore searching if there existed a “best” value of γ was integral in helping mPES obtain good learning performance from the models.

The results for this experiment are shown in **Table 1**, with the optimal value of the hyperparameter γ that was found for each combination of factors highlighted in bold. This “best” gain value was selected as the one giving the highest mean squared error-to-Spearman correlation coefficient ratio ($\frac{\rho}{\text{MSE}}$) across the various models. In our case this value was $\gamma = 10^4$, which would be the hyperparameter used in our subsequent experiments to evaluate the learning performance of our models in greater depth. Most modern Machine Learning algorithms employ some form of scheduling of the learning rate in order to help convergence. We had considered adding a decay to our gain factor but elected against this because, as previously stated, our eventual goal is to move our memristor-based learning model from simulation to physical circuit implementation and adding a schedule to our gain factor would entail a more complex CMOS logic. It should also be noted that the memristors themselves naturally implement a decaying learning schedule as their memristance behavior is described by a power law (Equation 5); each subsequent SET learning pulse will have a smaller effect on the resistance than the previous one.

3.4. Learning Experiments

The results for the experiments comparing learning a function using PES or mPES are shown in **Table 2**, where it can be immediately seen that our mPES learning rule is competitive with PES by noticing that the statistics encoding the quality of the learning, the mean squared error (MSE) and Spearman correlation coefficient (ρ), are consistently more favorable to mPES—or when not, very close to equal—across the spectrum of tested models.

Figures 6A–D show a selection of results from our simulations for various combinations of neuronal ensemble sizes, desired transformations f and input signals. The plots show the decoded output for the post-synaptic neuronal ensemble (filled colors), overlaid to the signal represented in the

TABLE 1 | Influence of the γ hyperparameter on the learning performance, expressed as mean squared error (MSE), Spearman correlation coefficient (ρ), and MSE-to- ρ ratio - measured in the testing phase - of the model for different combinations of ensemble size, input signal, and learned transformation.

10 neurons, sine, $f(x) = x$				10 neurons, sine, $f(x) = x^2$			
Gain γ	MSE	ρ	$\frac{\rho}{\text{MSE}}$	Gain γ	MSE	ρ	$\frac{\rho}{\text{MSE}}$
10^1	0.3058	0.0422	0.1381	10^1	0.1710	0.0061	0.0356
10^2	0.2737	0.3234	1.1814	10^2	0.1389	0.0364	0.2620
10^3	0.1416	0.7997	5.6478	10^3	0.1188	0.1745	1.4696
10^4	0.1183	0.8674	7.3303	10^4	0.1346	0.2306	1.7129
10^5	0.2190	0.7670	3.5016	10^5	0.2752	0.1161	0.4219
10^6	0.3177	0.6802	2.1408	10^6	0.4189	0.0747	0.1784
10 neurons, white, $f(x) = x$				10 neurons, white, $f(x) = x^2$			
Gain γ	MSE	ρ	$\frac{\rho}{\text{MSE}}$	Gain γ	MSE	ρ	$\frac{\rho}{\text{MSE}}$
10^1	0.1733	0.0107	0.0618	10^1	0.0863	0.0030	0.0349
10^2	0.1631	0.1070	0.6563	10^2	0.0752	0.0280	0.3722
10^3	0.1352	0.5313	3.9290	10^3	0.0869	0.0882	1.0149
10^4	0.1287	0.7599	5.9052	10^4	0.1298	0.1542	1.1882
10^5	0.1830	0.7308	3.9923	10^5	0.1831	0.1332	0.7276
10^6	0.2302	0.6136	2.6655	10^6	0.2866	0.0936	0.3268
100 neurons, sine, $f(x) = x$				100 neurons, sine, $f(x) = x^2$			
Gain γ	MSE	ρ	$\frac{\rho}{\text{MSE}}$	Gain γ	MSE	ρ	$\frac{\rho}{\text{MSE}}$
10^1	0.3049	0.0357	0.1172	10^1	0.1817	0.0065	0.0360
10^2	0.2740	0.3171	1.1571	10^2	0.1439	0.0452	0.3144
10^3	0.1461	0.7886	5.3989	10^3	0.1169	0.1511	1.2920
10^4	0.1239	0.8741	7.0567	10^4	0.1924	0.2143	1.1138
10^5	0.2341	0.7785	3.3254	10^5	0.2710	0.1426	0.5261
10^6	0.3649	0.6649	1.8220	10^6	0.4395	0.0843	0.1918
100 neurons, white, $f(x) = x$				100 neurons, white, $f(x) = x^2$			
Gain γ	MSE	ρ	$\frac{\rho}{\text{MSE}}$	Gain γ	MSE	ρ	$\frac{\rho}{\text{MSE}}$
10^1	0.1761	0.0130	0.0736	10^1	0.0790	0.0031	0.0389
10^2	0.1667	0.1001	0.6005	10^2	0.0710	0.0255	0.3590
10^3	0.1413	0.5155	3.6488	10^3	0.0928	0.0797	0.8583
10^4	0.1219	0.7689	6.3070	10^4	0.1255	0.1593	1.2694
10^5	0.1769	0.7292	4.1219	10^5	0.1797	0.1373	0.7642
10^6	0.2225	0.6467	2.9062	10^6	0.2844	0.0767	0.2696

The bolded entries highlight the optimal value found for each statistic, and the best γ chosen for each model. NB: the γ chosen for the [100 neurons, white, $f(x) = x$] model was $\gamma = 10^4$ even though the $\frac{\rho}{\text{MSE}}$ was lower for $\gamma = 10^3$.

pre-synaptic ensemble (faded colors) during the testing phase of the simulation. The x-axis shows the final 8 s of the simulation, while the y-axis is fixed to $[-1, +1]$ to show the value of the signals. In all cases the training input and testing signals were 3-dimensional, with the first dimension plotted in blue, the second in orange, and the third in green. These results were chosen to give an intuition on the outcome of learning from a qualitative point of view, given that mere statistics might not encode the notion of what an observer would deem a “good” fit.

Figure 6A shows how pre, post and error ensembles of 10 neurons each, after the learning period, were able to

learn the identity function from a 3-dimensional sine wave. Increasing the model’s ensembles to 100 neurons (see **Figure 6B**) gave a much cleaner output that better approximated the pre-synaptic signal. It should also be noted that the signal decoded from the smaller neuronal ensembles is much noisier, apart from having a worse average fit (i.e., has both higher MSE and lower ρ). **Figure 6C** shows that 100-neuron ensembles could learn the identity function after being trained on a white noise signal. Finally, **Figure 6D** indicates that neuronal ensembles of 100 neurons each could also learn to approximate the square from an input sine wave, with the

TABLE 2 | Mean squared error (MSE), Spearman correlation coefficient (ρ), and MSE-to- ρ —averaged over 100 runs—for models trying to learn either $y = x$ or $y = x^2$ using PES or mPES as learning rule, from a 3-dimensional learning input signal.

Neur.	Learn x	$f(x)$	Test x	PES			mPES			No learning		
				MSE	ρ	$\frac{\rho}{\text{MSE}}$	MSE	ρ	$\frac{\rho}{\text{MSE}}$	MSE	ρ	$\frac{\rho}{\text{MSE}}$
10	Sine	x	Sine	0.2088	0.8283	3.9673	0.1283	0.8719	6.7957	0.5675	0.0511	0.0900
			White	0.2800	0.5526	1.9735	0.1822	0.6268	3.4412	0.4044	-0.0128	-0.0317
		x^2	Sine	0.2048	0.2132	1.0411	0.1673	0.2032	1.2146	0.1849	-0.0075	-0.0406
			White	0.1783	0.0390	0.2189	0.1322	0.0992	0.7499	0.1005	0.0153	0.1524
	White	x	Sine	0.1315	0.8708	6.6223	0.1712	0.8058	4.7077	0.8719	0.0222	0.0255
			White	0.1515	0.8322	5.4928	0.1307	0.7719	5.9074	0.8021	0.0006	0.0008
		x^2	Sine	0.1561	0.2786	1.7854	0.1766	0.2024	1.1465	0.1676	0.0007	0.0042
			White	0.1586	0.2266	1.4288	0.1247	0.1374	1.1013	0.0002	-0.0009	-5.6996
100	Sine	x	Sine	0.1385	0.8812	6.3601	0.1197	0.9421	7.8731	0.4208	-0.0305	-0.0724
			White	0.1905	0.6810	3.5742	0.1847	0.7344	3.9758	0.2204	-0.0179	-0.0811
		x^2	Sine	0.0816	0.5956	7.3007	0.1316	0.4796	3.6441	0.2579	-0.0322	-0.1249
			White	0.0997	0.1739	1.7445	0.1724	0.1215	0.7047	0.1222	-0.0139	-0.1140
	White	x	Sine	0.0176	0.9865	56.1002	0.0867	0.9614	11.0912	0.4653	-0.0152	-0.0327
			White	0.0162	0.9784	60.2713	0.0654	0.9554	14.6141	0.2508	0.0186	0.0742
		x^2	Sine	0.0385	0.8384	21.7865	0.1405	0.5152	3.6670	0.2465	-0.0008	-0.0031
			White	0.0248	0.7957	32.0590	0.1013	0.5604	5.5338	0.1122	-0.0014	-0.0129

The models are given by a combination of number of neurons in their ensembles, training input signal, and function learned. Each model was tested on both kinds of testing input. The results where mPES is run with no learning time allotted are also included in the “No learning” column.

caveat being that the quality of the learned transformation was noticeably lower—which is also reflected by the statistics in **Table 2**.

An argument in support of reporting the notion of qualitative fit is given by analysing the results in **Table 2** for the models in **Figures 6B,D**. Here we see how these particular models perform similarly in regards to the measured mean squared error when trying to learn the identity and square function from a sine wave, but how the Spearman correlation coefficient is two times worse when learning the latter. The answer to the dilemma can be found by looking at the qualitative fits of the simulations (**Figures 6B,D**) and noticing that, even though the pre- and post-synaptic signals are—on average—close in both cases, when trying to learn the square the post-synaptic signal is much noisier. A noisy signal negatively impacts the correlation coefficient as the pre- and post-synaptic signals are nearly never moving in tandem, even when the oscillations are around the correct value. This could be ascribed to the square being a harder function to learn than the identity, as already noted throughout, which could lead to a model being “overpowered,” in the sense that its representational power is allotted to trying to learn the harder function, decreasing the resources available to represent the signal.

The sizes of the neuronal ensembles were not fine-tuned for compactness so it may well be possible to learn representations of similar quality using smaller neuronal ensembles.

3.5. Other Experiments

3.5.1. Increasing Noise Experiments

As previously mentioned, the impact of noise is not linear across the resistance range, with memristors in a HRS seeing their response to pulses having a higher uncertainty. This is due to the

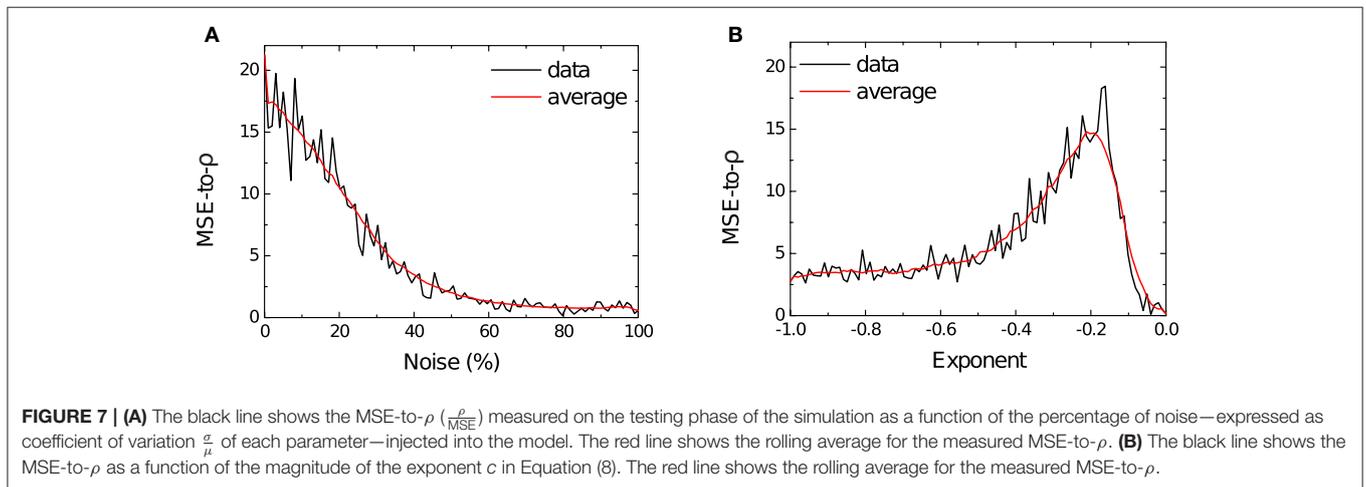
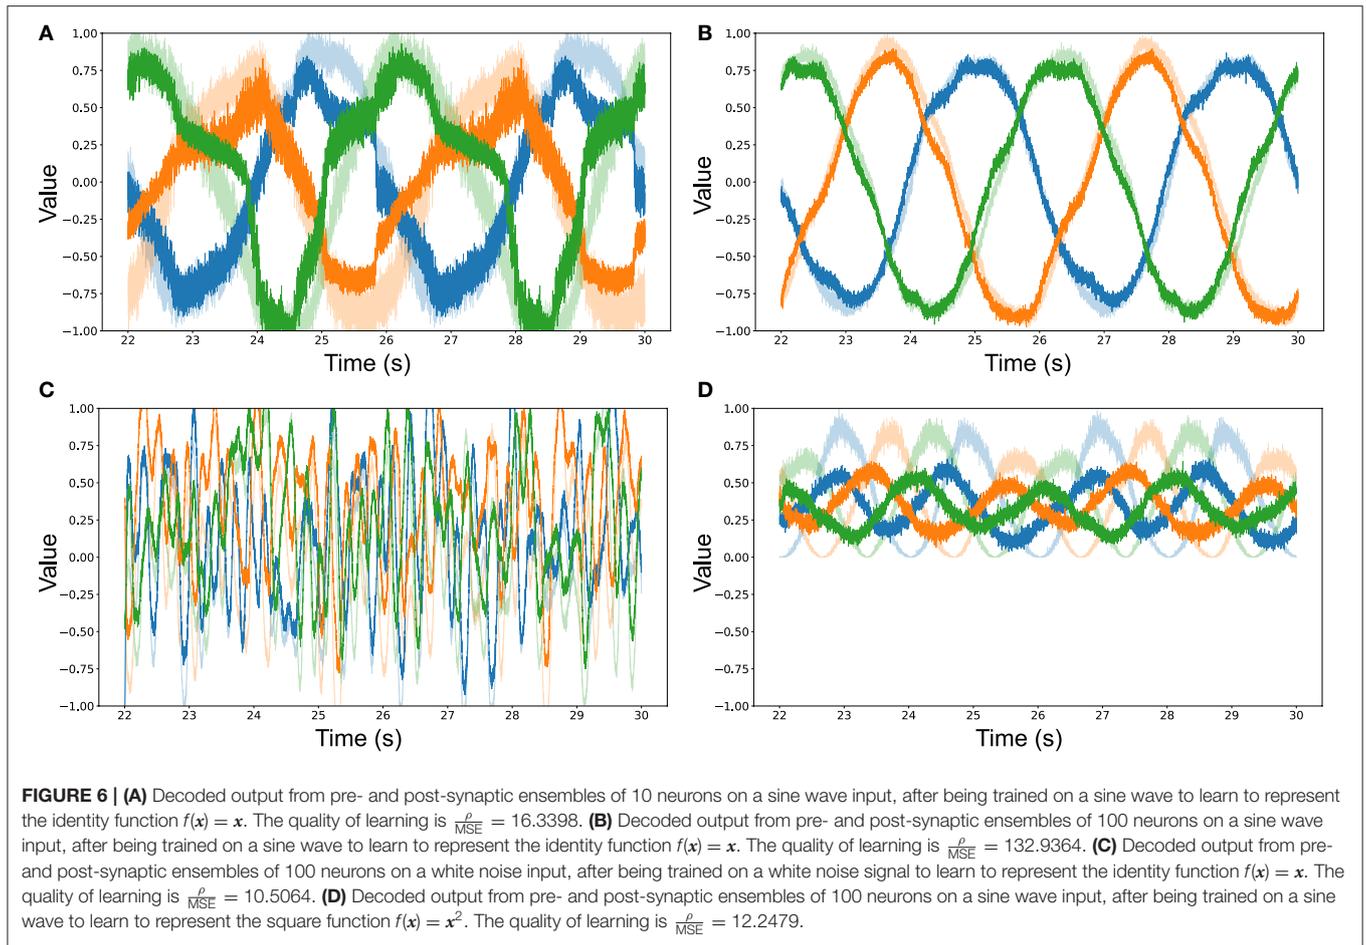
fact that we were introducing variation on the update equation itself and not on the resulting resistance state.

As can be seen in **Figure 7A**, learning performance decreased for high levels of noise and this degradation was graceful, as the MSE-to- ρ monotonically decreased toward 0. As previously stated: a lower MSE-to- ρ ratio ($\frac{\rho}{\text{MSE}}$) indicates a model that cannot faithfully represent the transformed pre-synaptic signal $f(x)$. Performance was well-maintained up until $\approx 15\%$ of variation of the parameters and smoothly decreased until stabilizing around 0 for noise $> 60\%$. An example of how the memristor resistances behave in the presence of various amounts of noise is shown in **Supplementary Figure 4**.

3.5.2. Exponent Search Experiments

The results for the exponent parameter search are shown in **Figure 7B** and here we can see how strongly the performance of the model depends on the magnitude of the exponent in the power law (Equation 8). The quality of the fit is in direct correspondence with the value of c up until $c \approx -0.162$, with a rapid decrease in learning performance thereafter. The highest performance is found for $c \in [-0.17, -0.16]$. Our choice of supplying learning SET pulses of $+0.1$ V implied $c = -0.146$ (Equation 8), suggesting that a slightly higher tension could have led to better performance in the task at hand.

The decrease in MSE-to- ρ for the smallest values of the exponent is most likely the result of the memristors having too small of a response for the training time supplied (which was 24 s in all out experiments) to be able to learn the function. This is the equivalent to having too small a learning rate in a traditional machine learning model.



4. DISCUSSION

Even with the restrictions imposed by the devices’ physical characteristics, our experiments showed that we could enable the post-synaptic neuronal ensemble to well-represent

transformations of the pre-synaptic signal by applying our novel mPES learning rule to synapses based on Nb:STO memristive devices. This process leads to a set of memristor resistances that can be mapped to network weights, which transform the supplied input into the desired output by encoding the learned

function. Crucially, our results are not limited to our Nb:STO memristors but generalize to any memristive device whose memristance follows a power law. We showed that our model could learn to approximate both the identity $f(x) = x$ and square $f(x) = x^2$ functions—the latter being a qualitatively harder problem to solve—with no theoretical reason limiting it to only representing these transformations. Learning was effective both when the input signal was a periodic sine wave and also when it was random, as the case where the pre-synaptic ensemble was fed white noise, or was switched between learning and testing phases; this second result is important in that it shows that the model was able to learn the transformation independently of the regularity of the supplied training data and of the input signal itself. It had already been suggested that the original PES learning rule, on which our mPES algorithm is based, would be able to learn arbitrary vector functions (Bekolay, 2010; MacNeil and Eliasmith, 2011). The fact that our learning rule is derived from one already capable of approximating arbitrary functions leads further credence to our conclusion that our model is capable of learning any transformation of n -dimensional vectors. Critically, as our model aimed to present strong neuromorphic credentials, these results are robust to the presence of noise on the updates and on the initial state of the memristors' resistances.

The results in the learning experiments—during which we compared the quality of the representations in the pre- and post-synaptic ensembles—were competitive with the original PES learning rule even though the latter had the advantage of operating on ideal, continuous network weights, with no uncertainty on their updates. Our mPES learning rule showed a post-training MSE-to- ρ ratio that was consistently close to that of PES; the results were especially favorable when trying to learn the identity function $f(x) = x$. In general, the MSE-to- ρ ratio was consistently lower for both algorithms when trying to learn the square transformation $f(x) = x^2$ compared to the $f(x) = x$ case, but the drop in performance for mPES was more marked. These observations suggest that the square function is a harder problem to learn and that the gap in performance between PES and mPES when trying to learn this transformation could be bridged by using a higher number of neurons in conjunction with our learning rule. Both learning rules exhibited higher performance when trained on a white noise input signal compared to a sine wave; we attribute this difference to the dissimilarity in how the input spaces are sampled but have no grounded theoretical explanations to explain this influence. Most importantly, we did see a decline in learning performance when switching to a different input signal between learning and testing phases, but this effect was limited; this lends credence to the belief that our model is actually able to generalize and thus is learning the function f , not just the transformed input signal $f(x)$.

The evaluations on the learning performance for increasing amounts of noise were done with models composed of a very small number of neurons, which could potentially predispose the network to be more susceptible to the injected randomness; this suggests that using our memristive devices and learning rule can lead to robust performance. As previously reported, the model maintained its best performance up to a moderate amount of Gaussian noise on the initial memristor resistances and on the

parameters of the update equation (Equation 8). The measured performance in the presence of increasing noise is less than that reported in previous work, but in our case we did not have a well-defined metric to measure performance against, for example recognition rate as in Querlioz et al. (2013). Again we raise the question of what constitutes a “good enough fit,” reiterating that—in our view—this is entirely dependent on the observer and the task at hand. The results also suggest that noise aids the model to a certain degree, given that the act of initializing each network weight to a different value is, in Machine Learning terminology, known as “symmetry breaking”; failing to do so can make a model difficult or impossible to train. Therefore, the noise on the initial resistance state could be effectively “breaking the symmetry” of the weights.

It is notable that the learning performance was maintained even though our mPES learning rule had absolutely no knowledge of the magnitudes of the updates and that the only hyperparameter tuned was the gain γ , whose effect was analogous to a learning rate. The learning rule was also only able to adjust the memristors toward lower resistance states, potentially making it harder to backtrack on any error. We experimentally found, as reported in **Table 1**, that a value of $\gamma = 10^4$ led to the best learning performance—defined as balance between testing error and correlation MSE-to- ρ ratio $\frac{\rho}{\text{MSE}}$ —for all models given by combinations of neuronal ensemble size, training input signal, and desired transformation. The best values for mean squared error (MSE) and Spearman correlation coefficient (ρ) were both found for the same value of $\gamma = 10^4$ in all the instances where the models were instructed to learn the identity function $f(x) = x$. Instead, when the models were attempting to learn the square function $f(x) = x^2$ there was a discrepancy between the values of γ leading to the lowest MSE and to the highest ρ . The two models learning the square from the regular sine wave input still had best ρ for $\gamma = 10^4$, but showed the lowest error for a smaller gain factor of $\gamma = 10^3$; the combined MSE-to- ρ statistic still indicated $\gamma = 10^4$ as the best choice for this class of model as the positive gains in correlation counterbalanced the higher error. The models learning the square transformation from the random white noise input followed the same pattern as they also showed highest correlation for $\gamma = 10^4$ and required a smaller gain factor to exhibit the lowest error: $\gamma = 10^2$ in their case. The measured $\frac{\rho}{\text{MSE}}$ for this last class of models was close for both $\gamma = 10^3$ and $\gamma = 10^4$ as the correlation ρ did not suffer as much as when trying to learn the square from the sine wave input.

In the [100 neurons, sine, $f(x) = x^2$] instance the best MSE-to- ρ ratio was found for $\gamma = 10^3$ but, given that the result was very close to the one obtained with $\gamma = 10^4$, the latter was chosen for consistency with all other cases.

The magnitude of γ was crucial to enable learning as it brought the normalized conductances calculated in Equation (2) to a scale compatible with the model and the input signal. The need to precisely tune the learning rate when using memristors presenting non-linear and asymmetric response has already been recognized (Burr et al., 2015) and this effect is somewhat confirmed here, as a too small or too large γ completely disabled the learning capacity of the model, while a sub-optimal value hindered the learning performance. Having to tune

hyperparameters in order to achieve better learning performance is, unfortunately, the norm in Machine Learning so it is not surprising that we ran into the same requirement. The search was not fine-grained, but gave a good indication of the magnitude of γ we could expect to lead to reasonable performance. It could be argued that, in a neuromorphic setting, it is more advantageous to have a “rough” estimate for the hyperparameter— $\gamma = 10^4$ in our case—leading to “robust” learning in a variety of situations, rather than focusing on finding the optimal value for a specific setup. Future work should focus on establishing a theoretical link between the model setup, including input signal characteristics and magnitude, and the expected range for finding a γ leading to sufficiently robust learning performance.

The memristive response of our device had notable implications on its use as a substrate for learning (Burr et al., 2015; Sidler et al., 2016). There is contention regarding the best characteristics that devices’ memristance should have in order to effectively support learning. Recent works highlight how there is currently no consensus on if a linear response (Krishnaprasad et al., 2019) or a non-linear (Brivio et al., 2018; Frascaroli et al., 2018) behavior of memristive synapses is more advantageous. On our side, we can add to the discourse by pointing out that our learning experiments (results in **Table 2**) showed that our mPES algorithm, modulating the resistance of non-linear Nb:STO memristors, was competitive with the original PES learning rule, which operated on linear, ideal synaptic weights. **Figure 3** adds some intuition to why a power-law update would be beneficial for learning compared to a linear synaptic response: the decreasing effect of each training pulse effectively mimics the cooling in simulated annealing and therefore helps the weights converge to an optimal value.

In our current work we only used SET pulses of +0.1 V while negative RESET pulses were largely ignored as the effect of sequential negative pulses was seen to quickly diminish with pulse number. Our initial choice to apply positive learning pulses of +0.1 V implied having an exponent $c = -0.146$ for the power law governing the memristor updates, which put the model close to the region exhibiting best learning performance. From a physical point of view, the magnitude of the SET pulses could be set very low as the current-voltage characteristics of our memristors are continuous, without any step-like features reminiscent of a switching threshold voltage. Hence, in principle we expect even small pulse magnitudes to influence the resistive state of the device. In addition, by utilizing a read voltage of a different polarity than that of the SET pulses, we created an asymmetry between the writing and reading processes that made it possible to use very small voltage amplitudes for the SET pulses. Thus it should be possible, in principle, to use physical voltage pulses as small as needed to guarantee the best learning performance. Different mechanisms govern charge transport in forward and reverse bias; this leads to asymmetries in current-voltage characteristics and may also result in differences between the SET and RESET processes.

The disadvantage of only focusing on the forward bias is that long training regimes could potentially lead to the memristors’ resistances to saturate. That is, using only positive pulses progressively brings the devices to a low resistance state with

no way to backtrack. This issue could be alleviated by a careful pulsing scheme where memristor close to saturating were “re-initialized” to a high resistance state by supplying RESET voltage pulses, at the cost of losing some precision in the represented network weights (Burr et al., 2015). This could be done by re-initializing both memristors in a pair to their initial resistances with RESET pulses and then supplying a number of SET pulses proportional to the magnitude of the weight before the reset, on either the positive or negative memristor depending on the polarity of the weight. After the resetting, a highly positive weight would thus see its memristor M^+ receiving more pulses than a lightly positive one, as a moderately negative weight might have its memristor M^- receive fewer SET pulses than a highly negative one. This approach could have the disadvantage of requiring quite a complex meta-architecture, but it is possible that the current learning approach may already be robust enough to deal gracefully with the resetting of a small percentage of its synapses.

Another path that could improve the learning performance could be to grade the number, the duration, or the magnitude of pulses applied to each memristor, based on the participation of its pre- and post-synaptic neurons to the global error. This could prove to be a less demanding undertaking than shaping the pulses based on the memristors’ current state, while still resulting in positive gains to learning performance.

5. CONCLUSIONS

In this work we fabricated memristive devices based on Ni/Nb-doped SrTiO₃ and found that their memristance followed a power law. These memristive devices were used as the synaptic weight element in a spiking neural network to simulate, to our knowledge, one of the first models of this kind capable of learning to be a universal function approximator. The performance was tested in the Nengo Brain Builder framework by defining a simple network topology capable of learning transformations from multi-dimensional, time-varying signals. The network demonstrated good learning performance with robustness to moderate noise levels of up to $\approx 15\%$, showing that this class of memristive devices is apt to being used as component of a neuromorphic architecture. It is worth restating that the network weights were found using only discrete updates to the memristors, based on knowledge local to each pre- and post-synaptic neuron pair; this means that our learning model presents many more neuromorphic characteristics than, for example, previous works where memristors are used as mere hardware implementations for artificial neural network weights (for example, see Li et al., 2019).

Using memristive devices as weights can enable efficient computing—which is also one of the cornerstones of the neuromorphic approach—but not pairing them with a neuromorphic learning algorithm and a spiking neural network imposes quite stringent requirements on the physical characteristics of the device. For example, a memristor simply used as a physical implementation of a network weight works best if it exhibits reliable, linear, and symmetric conductance

response in order to approximate the idealized weight it stands in for Li et al. (2018). Present-day machine learning models suffer from “brittleness” (i.e., small changes in input can give rise to large errors) and do not seem well-positioned in being able to match human adaptability across a wide variety of tasks. It is becoming accepted that stochasticity is itself a computational resource (Xia and Yang, 2019)—even though it is still not formally proven as such—and the fact that our model seems to perform best with a moderate amount of noise could be a result of this effect. It had already been proposed that networks using memristors as synapses could match the learning performance of traditional artificial neural networks (Burr et al., 2015) and our results are in line with this conjecture.

CODE AVAILABILITY

All code used in this study is publicly available on GitHub at <https://github.com/Tioz90/Learning-to-approximate-functions-using-niobium-doped-strontium-titanate-memristors>.

DATA AVAILABILITY STATEMENT

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

REFERENCES

- Ambrogio, S., Ciochini, N., Laudato, M., Milo, V., Pirovano, A., Fantini, P., et al. (2016). Unsupervised learning by spike timing dependent plasticity in phase change memory (PCM) synapses. *Front. Neurosci.* 10:56. doi: 10.3389/fnins.2016.00056
- Amer, S., Rose, G. S., Beckmann, K., and Cady, N. C. (2017). “Design techniques for in-field memristor forming circuits,” in *2017 IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS)* (Boston, MA: IEEE), 1224–1227. doi: 10.1109/MWSCAS.2017.8053150
- Bekolay, T. (2010). *Learning Nonlinear Functions on Vectors: Examples and Predictions*. Technical report, CTN-TR-20101217-010, Centre for Theoretical Neuroscience, Waterloo.
- Bekolay, T., Bergstra, J., Hunsberger, E., DeWolf, T., Stewart, T. C., Rasmussen, D., et al. (2014). Nengo: a python tool for building large-scale functional brain models. *Front. Neuroinform.* 7:48. doi: 10.3389/fninf.2013.00048
- Brivio, S., Conti, D., Nair, M. V., Frascaroli, J., Covi, E., Ricciardi, C., et al. (2018). Extended memory lifetime in spiking neural networks employing memristive synapses with nonlinear conductance dynamics. *Nanotechnology* 30:015102. doi: 10.1088/1361-6528/aae81c
- Burr, G. W., Shelby, R. M., Sidler, S., Nolfo, C., d., Jang, J., et al. (2015). Experimental demonstration and tolerancing of a large-scale neural network (165,000 synapses) using phase-change memory as the synaptic weight element. *IEEE Trans. Electron Devices* 62, 3498–3507. doi: 10.1109/TED.2015.2439635
- Davies, M., Srinivasa, N., Lin, T.-H., Chinya, G., Cao, Y., Choday, S. H., et al. (2018). Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38, 82–99. doi: 10.1109/MM.2018.112130359
- Du, C., Cai, F., Zidan, M. A., Ma, W., Lee, S. H., and Lu, W. D. (2017). Reservoir computing using dynamic memristors for temporal information processing. *Nat. Commun.* 8:2204. doi: 10.1038/s41467-017-02337-y
- Eliasmith, C., and Anderson, C. H. (2003). *Neural engineering: Computation, Representation, and Dynamics in Neurobiological Systems*. Cambridge, MA; London: MIT Press.

AUTHOR CONTRIBUTIONS

TT, AG, JB, TB, and NT developed the ideas and wrote the paper. AG conducted the experimental work. TT, AG, and NT derived the mathematical results. TT designed and ran the simulations. All authors contributed to the article and approved the submitted version.

ACKNOWLEDGMENTS

This manuscript has been released as a pre-print at arXiv.org (arXiv:2011.02794) (Tiotto et al., 2020).

AG and TB would like to thank all members of the Spintronics of Functional Materials group at the University of Groningen, in particular Arijit Das for help with device fabrication. Device fabrication was realized using NanoLab NL facilities. AG and TB acknowledge technical support from J. G. Holstein, H. H. de Vries, T. Schouten, and H. Adema. AG and TT were supported by the CogniGron Center, University of Groningen.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnins.2020.627276/full#supplementary-material>

- Frascaroli, J., Brivio, S., Covi, E., and Spiga, S. (2018). Evidence of soft bound behaviour in analogue memristive devices for neuromorphic computing. *Sci. Rep.* 8:7178. doi: 10.1038/s41598-018-25376-x
- Furber, S. B., Galluppi, F., Temple, S., and Plana, L. A. (2014). The SpiNNaker project. *Proc. IEEE* 102, 652–665. doi: 10.1109/JPROC.2014.2304638
- Goossens, A. S., Das, A., and Banerjee, T. (2018). Electric field driven memristive behavior at the Schottky interface of Nb-doped SrTiO₃. *J. Appl. Phys.* 124:152102. doi: 10.1063/1.5037965
- Hill, L. B., Rejall, A. E., and Thorndike, E. L. (1991). Practice in the case of typewriting. *J. Genet. Psychol.* 152, 448–461. doi: 10.1080/00221325.1991.9914706
- Hu, M., Li, H., Chen, Y., Wu, Q., Rose, G. S., and Linderman, R. W. (2014). Memristor crossbar-based neuromorphic computing system: a case study. *IEEE Trans. Neural Netw. Learn. Syst.* 25, 1864–1878. doi: 10.1109/TNNLS.2013.2296777
- Hu, P., Wu, S., and Li, S. (2018). “Synaptic behavior in metal oxide-based memristors,” in *Advances in Memristor Neural Networks-Modeling and Applications* (London: IntechOpen). doi: 10.5772/intechopen.78408
- Jang, J. T., Ko, D., Ahn, G., Yu, H. R., Jung, H., Kim, Y. S., et al. (2018). Effect of oxygen content of the LaAlO₃ layer on the synaptic behavior of Pt/LaAlO₃/Nb-doped SrTiO₃ memristors for neuromorphic applications. *Solid State Electron.* 140, 139–143. doi: 10.1016/j.sse.2017.10.032
- Kim, G. S., Song, H., Lee, Y. K., Kim, J. H., Kim, W., Park, T. H., et al. (2019). Defect-engineered electroforming-free analog HfO_x memristor and its application to the neural network. *ACS Appl. Mater. Interfaces* 11, 47063–47072. doi: 10.1021/acsami.9b16499
- Kim, M. K., and Lee, J. S. (2019). Ferroelectric analog synaptic transistors. *Nano Lett.* 19, 2044–2050. doi: 10.1021/acs.nanolett.9b00180
- Kim, S., Du, C., Sheridan, P., Ma, W., Choi, S., and Lu, W. D. (2015). Experimental demonstration of a second-order memristor and its ability to biorealistically implement synaptic plasticity. *Nano Lett.* 15, 2203–2211. doi: 10.1021/acs.nanolett.5b00697
- Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. (1983). Optimization by simulated annealing. *Science* 220, 671–680. doi: 10.1126/science.220.4598.671

- Krishnaprasad, A., Choudhary, N., Das, S., Dev, D., Kalita, H., Chung, H.-S., et al. (2019). Electronic synapses with near-linear weight update using MoS₂/graphene memristors. *Appl. Phys. Lett.* 115:103104. doi: 10.1063/1.5108899
- Kuzum, D., Jeyasingh, R. G. D., Lee, B., and Wong, H. S. P. (2012). Nanoelectronic programmable synapses based on phase change materials for brain-inspired computing. *Nano Lett.* 12, 2179–2186. doi: 10.1021/nl201040y
- Le Cun, Y. (1986). “Learning process in an asymmetric threshold network,” in *Disordered Systems and Biological Organization* (Berlin; Heidelberg: Springer), 233–240. doi: 10.1007/978-3-642-82657-3_24
- Li, C., Belkin, D., Li, Y., Yan, P., Hu, M., Ge, N., et al. (2018). Efficient and self-adaptive *in-situ* learning in multilayer memristor neural networks. *Nat. Commun.* 9:2385. doi: 10.1038/s41467-018-04484-2
- Li, C., Wang, Z., Rao, M., Belkin, D., Song, W., Jiang, H., et al. (2019). Long short-term memory networks in memristor crossbar arrays. *Nat. Mach. Intell.* 1, 49–57. doi: 10.1038/s42256-018-0001-4
- Lukoševičius, M., and Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* 3, 127–149. doi: 10.1016/j.cosrev.2009.03.005
- MacNeil, D., and Eliasmith, C. (2011). Fine-tuning and the stability of recurrent neural networks. *PLoS ONE* 6:e22885. doi: 10.1371/journal.pone.0022885
- Midya, R., Wang, Z., Asapu, S., Zhang, X., Rao, M., Song, W., et al. (2019). Reservoir computing using diffusive memristors. *Adv. Intell. Syst.* 1:1900084. doi: 10.1016/j.aisy.201900084
- Mikheev, E., Hoskins, B. D., Strukov, D. B., and Stemmer, S. (2014). Resistive switching and its suppression in Pt/Nb:SrTiO₃ junctions. *Nat. Commun.* 5:3990. doi: 10.1038/ncomms4990
- Moon, J., Ma, W., Shin, J. H., Cai, F., Du, C., Lee, S. H., et al. (2019). Temporal data classification and forecasting using a memristor-based reservoir computing system. *Nat. Electron.* 2, 480–487. doi: 10.1038/s41928-019-0313-3
- Moon, K., Lim, S., Park, J., Sung, C., Oh, S., Woo, J., et al. (2018). RRAM-based synapse devices for neuromorphic systems. *Faraday Discuss.* 213, 421–451. doi: 10.1039/C8FD00127H
- Mostafa, H., Ramesh, V., and Cauwenberghs, G. (2018). Deep supervised learning using local errors. *Front. Neurosci.* 12:608. doi: 10.3389/fnins.2018.00608
- Newell, A., and Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. *Cogn. Skills Acquisit.* 1, 1–55.
- Nishitani, Y., Kaneko, Y., Ueda, M., Morie, T., and Fujii, E. (2012). Three-terminal ferroelectric synapse device with concurrent learning function for artificial neural networks. *J. Appl. Phys.* 111:124108. doi: 10.1063/1.4729915
- Oh, S., Hwang, H., and Yoo, I. K. (2019). Ferroelectric materials for neuromorphic computing. *APL Mater.* 7:091109. doi: 10.1063/1.5108562
- Querlioz, D., Bichler, O., Dollfus, P., and Gamrat, C. (2013). Immunity to device variations in a spiking neural network with memristive nanodevices. *IEEE Trans. Nanotechnol.* 12, 288–295. doi: 10.1109/TNANO.2013.2250995
- Rodenbücher, C., Speier, W., Bihlmayer, G., Breuer, U., Waser, R., and Szot, K. (2013). Cluster-like resistive switching of SrTiO₃:Nb surface layers. *New J. Phys.* 15:103017. doi: 10.1088/1367-2630/15/10/103017
- Sawa, A. (2008). Resistive switching in transition metal oxides. *Mater. Today* 11, 28–36. doi: 10.1016/S1369-7021(08)70119-6
- Seo, K., Kim, I., Jung, S., Jo, M., Park, S., Park, J., et al. (2011). Analog memory and spike-timing-dependent plasticity characteristics of a nanoscale titanium oxide bilayer resistive switching device. *Nanotechnology* 22:254023. doi: 10.1088/0957-4484/22/25/254023
- Seong, D., Jo, M., Lee, D., and Hwang, H. (2007). HPHA effect on reversible resistive switching of Pt/Nb-doped SrTiO₃ Schottky junction for nonvolatile memory application. *Electrochem. Solid State Lett.* 10:H168. doi: 10.1149/1.2718396
- Serrano-Gotarredona, T., Masquelier, T., Prodromakis, T., Indiveri, G., and Linares-Barranco, B. (2013). STDP and STDP variations with memristors for spiking neuromorphic learning systems. *Front. Neurosci.* 7:2. doi: 10.3389/fnins.2013.00002
- Sidler, S., Boybat, I., Shelby, R. M., Narayanan, P., Jang, J., Fumarola, A., et al. (2016). “Large-scale neural networks implemented with non-volatile memory as the synaptic weight element: impact of conductance response,” in *2016 46th European Solid-State Device Research Conference (ESSDERC)* (Lausanne: IEEE), 440–443. doi: 10.1109/ESSDERC.2016.7599680
- Sim, H., Choi, H., Chang, D. L. M., Choi, D., Son, Y., Lee, E. H., et al. (2005). *Excellent Resistance Switching Characteristics of Pt/SrTiO₃ Schottky Junction for Multi-bit Nonvolatile Memory Application*. Washington, DC: Technical Digest-International Electron Devices Meeting, IEDM.
- Strukov, D. B., Snider, G. S., Stewart, D. R., and Williams, R. S. (2008). The missing memristor found. *Nature* 453, 80–83. doi: 10.1038/nature06932
- Susaki, T., Kozuka, Y., Tateyama, Y., and Hwang, H. Y. (2007). Temperature-dependent polarity reversal in Au/Nb:SrTiO₃ Schottky junctions. *Phys. Rev. B* 76:155110. doi: 10.1103/PhysRevB.76.155110
- Tiotto, T. F., Goossens, A. S., Borst, J. P., Banerjee, T., and Taatgen, N. A. (2020). Learning to approximate functions using Nb-doped SrTiO₃ memristors. *arXiv* 2011.02794.
- Wang, L., Lu, S.-R., and Wen, J. (2017a). Recent advances on neuromorphic systems using phase-change materials. *Nanoscale Res. Lett.* 12:347. doi: 10.1186/s11671-017-2114-9
- Wang, W., Pedretti, G., Milo, V., Carboni, R., Calderoni, A., Ramaswamy, N., et al. (2018a). Computing of temporal information in spiking neural networks with ReRAM synapses. *Faraday Discuss.* 213, 453–469. doi: 10.1039/C8FD00097B
- Wang, Z., Joshi, S., Savel'ev, S., Song, W., Midya, R., Li, Y., et al. (2018b). Fully memristive neural networks for pattern classification with unsupervised learning. *Nat. Electron.* 1, 137–145. doi: 10.1038/s41928-018-0023-2
- Wang, Z., Joshi, S., Savel'ev, S. E., Jiang, H., Midya, R., Lin, P., et al. (2017b). Memristors with diffusive dynamics as synaptic emulators for neuromorphic computing. *Nat. Mater.* 16, 101–108. doi: 10.1038/nmat4756
- Xia, Q., and Yang, J. J. (2019). Memristive crossbar arrays for brain-inspired computing. *Nat. Mater.* 18, 309–323. doi: 10.1038/s41563-019-0291-x
- Yin, X. B., Yang, R., Xue, K. H., Tan, Z. H., Zhang, X. D., Miao, X. S., et al. (2016). Mimicking the brain functions of learning, forgetting and explicit/implicit memories with SrTiO₃-based memristive devices. *Phys. Chem. Chem. Phys.* 18, 31796–31802. doi: 10.1039/C6CP06049H
- Zhao, L., Xu, J., Shang, X., Li, X., Li, Q., and Li, S. (2019). Synaptic memory devices from CoO/Nb:SrTiO₃ junction. *R. Soc. Open Sci.* 6:181098. doi: 10.1098/rsos.181098

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2021 Tiotto, Goossens, Borst, Banerjee and Taatgen. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.