

Memristive STDP Synapses and Unsupervised Learning

Tabeeb Rahman, Bsc Mathematics and Physics, Department of Physics, University of Bath, Bath UK

Abstract—In the field of machine learning, Artificial Neural Networks (ANNs of the generalised perceptron type [1]) model synapses as scalars (for a generic synaptic strength variable), and typically relies on updating the value of the synapse (synaptic plasticity) in a globally coordinated (synchronous), iterative process (Gradient Descent [2]) to train the model. In the currently prevalent Von Neumann computer hardware paradigm, such computations are energy expensive. Motivated by the energy efficiency of biological brains, the spike-timing-dependent plasticity (STDP) rule supposes synaptic plasticity be modelled asynchronously for a biologically inspired autonomous synapse, with strength change dependent on the relative spike timing between a pair of connected biologically inspired spiking neurons. These neurons differ from fixed scalar representations of neurons in ANNs by emitting discrete voltage spikes when sufficiently stimulated, and networks of such are called Spiking Neural Networks (SNNs).

This project computationally simulated the AIST (Ag-In-Sb-Te) memristor to investigate the claim that it can act as an artificial STDP synapse to support learning inside of an SNN. Memristors are compact resistive devices that can have their conductances (storing the STDP synaptic strength variable) programmed with external voltages, known as pulse programming. A pair of spiking-type neurons were situated pre- and post- of the AIST memristor, forming a pulse programming scheme. Ideal STDP and the obtained pulse programming updates agreed strongly, with $1 - |r|$ (Pearson's r) found to be on the order of 10^{-4} (the closer to 0 the stronger the fit), supporting the claim. Challenges remain as fabricated memristors suffer from stochasticity and device-to-device variations that affect STDP reproduction. These non-idealities were not used in the AIST simulation model for simplicity. For this reason, hardware implementations of large memristive SNNs remain elusive.

The project also assessed the claim that the STDP rule is sufficient to effectively train SNNs with unlabelled training data (unsupervised learning). A full-scale ideal STDP SNN with 100-, 400- and 1600-neuron SNNs were trained on monochrome images of handwritten digits to classify the correct characters, with the 1600 neuron SNN reaching a peak 98.0% and mean 93.544% test accuracy. This strong performance supports the claim, though falls short of the 99%+ accuracy of Gradient Descent used in traditional ANNs.

I. INTRODUCTION

Biological neural systems continue to motivate the search for more energy-efficient Machine learning approaches. Many conventional machine learning methods including Gradient Descent are made to run iterative plasticity updates that require global coordination with the rest of the network, in order to adhere to the in order (serial) nature of the Von Neumann computational paradigm that most computers since the dawn of the transistor age uses. The bottlenecks found in the Von Neumann architecture is fundamental to its synchrony

and serial command execution nature. Iterations must be synchronised and carried out serially. Results of prior iterations are shuttled serially to and from the place where results are physically stored (memory) to the circuits responsible for mathematical computation and program execution, via serial data buss circuits. Each result in an iterative approach relies on the accuracy of prior results, so the signal integrity needed over these data busses is high. Additionally, larger networks require more data to be shuttled and synchronised over these serial busses than current data busses can handle. This Von neumann Bottleneck is fundamental to machine learning tasks, since in the context of neural networks, they are represented by graphs of edge (representing the synapse and associated strength) connected nodes (neurons). In densely connected graphs adding nodes makes the number of synapses grow quadratically, whereas scaling data throughput through serial data busses demands exponentially higher signal-to-noise ratios [3], requiring an increase in the number of more inefficient higher accuracy signal amplifiers. Also, Gradient Descent allows for synaptic updates to be done independently of one another (in parallel) on a proper subset of neurons and synapses at once, with each set being called a layer of neurons. Doing layer level computations simultaneously is poorly suited for the serial nature of these data busses, and is hard to scale for layers with more neurons in them. Adding more busses in parallel only grows the data throughput linearly but is prohibitively expensive to manufacture since adding more data busses takes up die space on silicon wafers, that would otherwise have been populated with mathematical logic circuits.

By contrast, biological brains are thought to implement synaptic plasticity by autonomous self-updating synapses acting locally on themselves, allowing for parallel computation at a per neuron-pair synapse level, vastly increasing the scalability of neural computations compared to Von Neumann based Gradient Descent approaches. Biological brains are also far more noise tolerant, with each synapse in, for example a human brain existing in a very noisy environment. Humans still consistently outperform state of the art Reasoning Language models [4] whilst consuming just 20 Watts of power [5]. The high noise tolerance and parallel native architecture of human and biological brains are thought to be primarily responsible for the currently inaccessible task-performance per watt capabilities they possess. This is the primary motivation for research into neuromorphic (inspired by neuroscience) hardware.

For the purposes of the SNN section and the pulse programming used in Memristive STDP synapse simulations, Various biological neuron models were researched. Hodgkin Huxley

(HH), Morris-Lecar [6] and Leaky Integrate and Fire (LIF) [7] models were used extensively here due to them being standard models in the field of computational neuroscience. All share the feature that with sufficient stimulation, these neurons produce voltage spikes, that transmit onto other neurons via synapses whose strength dictates the amount of stimulation neurons downstream of them experience. Neurons with greater stimulation have a higher chance of themselves firing. Neurons that provide positive stimulus to neurons downstream of it are known as excitatory neurons, and unless stated otherwise, all neurons going forward should be assumed as such.

Following the work of Lindner and Schimansky-Geier [8], it was found that stimulating certain neurons with a Gaussian noise-voltage waveform of a particular optimum amplitude produced the strongest periodicity in the emitted spike timings of the neuron. This is known as Coherence Resonance (CR) and is explained by Lindner as the intensity of noise enhancing the excitability of the neuron, as higher intensities increases the probability for a transient voltage being sufficiently high to illicit a spike. Too little amplitude and the chance of the noise producing enough voltage for the neuron to fire is unlikely. Too much noise amplitude and the neuron fires too often and stochastically. This forms the basis of the related phenomena known as Stochastic Resonance (SR), where in certain systems there exists an ideal amount of noise that maximises the correlation between upstream and downstream neurons spiking. This is important because spikes encode information in an snn so this corresponds to maximising the distance information encoded with spikes, can travel through the network. This motivated the use of some form of intrinsic noise inside of the SNN architecture by encoding intensities of input image data with frequency modulation, but instead of a deterministic frequency, spikes occurred randomly being Poisson distributed about the interval between a deterministic frequency signal. Instead of explicitly computing the optimum amount of noise (prohibitively expensive for large networks), the intensity of the noise was increased by shifting the entire frequency modulation range up by 5Hz any time fewer than 5 neurons in the network spiked during a training example [9]. An analytical treatment of CR and SR is provided in their original paper [8] for a noisy excitable FitzHugh-Nagumo type neuron system. This project primarily deals in demonstrating CR on the simplified Leaky Integrate Fire model [7] (the neuron used in the SNN section) for computational efficiency. Additionally Bio-plausible deterministic yet chaotic sources of pseudorandom aperiodic drive were also demonstrated using a ring of three coupled oscillators, made out of Morris-Lecar style neurons.

Device-level and simulation-level work on memristive synaptic circuits [10] in the literature describes AIST-type threshold memristors as two-terminal resistive devices that are plausible candidates for artificial STDP synapses [11]. Their conductance depends on non-volatile (persistent even when powered off) internal state variables that evolve under electrical stimulation. They possess two suprathreshold voltages, V_{set} or equivalently V_T^+ , and V_{reset} or equivalently V_T^- . Only when the voltage applied across the terminals exceeds either threshold, there is a change of opposing sign in the

observed conductance of the memristor. Zhang et. al specifies an explicit voltage threshold model to explain the empirical findings of a fabricated AIST memristor by Y li et al. [10]. This group also reproduced STDP directly on the fabricated memristor, agreeing with STDP on average with non negligible stochasticity for any isolated update. This forms the backbone of the pulse programming scheme described in the memristor STDP synapse section.

Finally, the STDP LIF SNN section based its work on that done by Diehl and Cook in 2015 [9]. They were the first to demonstrate, using pure unmodified STDP, higher accuracy than older approaches with simplified STDP rule proxies. Crucially, this architecture was fully unsupervised, marking a genuine milestone in neuromorphic machine learning.

This project computationally simulated the AIST memristor to investigate the claim that it can act as an artificial STDP synapse to support learning inside of an SNN. The project also assessed the claim that the STDP rule was sufficient to perform unsupervised learning on an SNN.

This paper is structured as follows: Firstly STDP is defined, moving on to the memristor simulation model used, where pulse programming and STDP is reproduced. Following that, an exploration of legacy biological neuron models that can both oscillate and fire only when sufficiently externally stimulated, are explored. These models explain the simplifications made in the more contemporarily relevant LIF non-oscillatory neuron model used in the SNN section. Then CR was investigated on an LIF neuron, with a brief explanation of SR followed by simulating three neuron ring oscillators using Morris-Lecar neurons, to produce aperiodic and synchronised oscillations between them. Finally the STDP SNN section is presented, incorporating all sections before it: STDP synapses, LIF Neurons and SR. 100-, 400- and 1600-neuron SNNs were trained and evaluated. Network specific mechanisms such as sparsity and homeostasis were also introduced, and explored. These findings are discussed in the broader context of memristor and neuromorphic literature, before the main conclusions and possible directions for future work are summarised.

II. SIMULATING STDP

A. Biphasic Exponential STDP

Let $\Delta t = t_{\text{post}} - t_{\text{pre}}$ denoting the time difference between a post-synaptic (post) and pre-synaptic (pre) neuron spike pair (top of Fig. 5, memristor symbol used in place of a biological synapse). When $\Delta t > 0$, the pre neuron fires before the post neuron and is called causal firing. In this case the synapse experiences what is called Long Term Potentiation (LTP). This is interpreted as reinforcing synapses when the pre neurons activity stimulated and thus is a candidate for being the proximate cause of the later firing of the post neuron. By contrast, when the post neuron fires before the pre neuron ($\Delta t < 0$), the synapse experiences what is called Long Term Depression, LTD. This penalises cases where the pre neuron's spike could not have been the proximate cause for the post neuron's firing, and is thusly named anti-causal firing.

The explicit amount to update the synaptic strength w , Δw is given by the following exponential update law:

$$\Delta w(\Delta t) = \begin{cases} A_+ \exp\left(-\frac{\Delta t}{\tau_+}\right) & \text{for } \Delta t > 0 \\ -A_- \exp\left(\frac{\Delta t}{\tau_-}\right) & \text{for } \Delta t < 0 \end{cases} \quad (1)$$

where A_+ and A_- set the amplitudes of potentiation and depression, and τ_+ and τ_- set the corresponding time constants of LTP and LTD.

This dependence on spike ordering gives rise to a characteristic STDP curve as seen in Fig. 1 with two lobes (**biphasic**) of opposite sign with a discontinuity at $\Delta t = 0$. As expected from (1), there is a positive lobe for $\Delta t > 0$, corresponding to LTP, and a negative lobe for LTD when $\Delta t < 0$. Another consequence of (1) is that spike pairs with $|\Delta t| \ll 1$ have the highest update amplitude. STDP is often used in training SNNs as a local learning rule because it biases synaptic updates toward repeatedly correlated spike pairs while weakening anti-causal pairings. In this sense it provides a task-agnostic mechanism by which temporally consistent spike relationships may be reinforced over repeated stimulus presentations. This is explored in the discussion.




Fig. 1. Idealized Biphasic STDP curve showing LTP and LTD regions.

B. Memristor model

The memristor that was investigated was modelled using the voltage-controlled threshold formulation proposed by Zhang, Wang [11], to fit empirically observed AIST memristor behaviour by Li [11]. This model introduces voltage thresholds, and defines suprathreshold (exceeding either threshold) conductance change dynamics that taper as either end of the range of conductances a given memristor chemistry can possess, is approached, known as saturation. This model remains compact enough for repeated device-level and network-level simulation, but omits device stochasticity for simplicity.

a) Suprathreshold Voltages and the state variable:

On a microscopic level, the AIST memristor contains silver electrodes for its two terminals, with a solid polycrystalline Ag-In-Sb-Te substrate sandwiched in the gap between the terminals (terminal gap). Ag-In-Sb-Te has two separate solid crystal phases with distinct resistances (and thus conductances). $R_{\text{on}} = 10 \Omega$ and $R_{\text{off}} = 420 \Omega$ correspond to the

filamentary and base phase resistances if either were to fill the entire terminal gap of width $D = 3 \text{ nm}$ respectively. The memristor can admit any intermediary state with a filamentary phase with an effective width across the terminal gap w (nm) (not to be confused with the synaptic strength variable w in II.A.). The rest of the gap is occupied by the base phase. Thus the overall resistance R is given by the linear interpolation (2) where t is a time variable (see next section).

$$R(t) = R_{\text{on}} \left(\frac{w(t)}{D} \right) + R_{\text{off}} \left(1 - \frac{w(t)}{D} \right). \quad (2)$$

b) Thresholded state dynamics:

When the voltage across the terminal at a time t exceeds $V_T^+ = 0.37 \text{ V}$ the resistance changes in one direction (increases). This corresponds to the filamentary phase dissolving, leaving more of the base phase along the terminal gap, when the applied voltage is less than $V_T^- = -0.19 \text{ V}$ more filamentary phase crystals start growing, reducing resistance as seen in (3):

$$\frac{dw}{dt} = \begin{cases} \mu_v \left(\frac{R_{\text{on}}}{D} \right) \left(\frac{i_{\text{off}}}{i(t) - i_0} \right) f(w(t)) & \text{if } v(t) > V_T^+, \\ 0 & \text{if } V_T^- \leq v(t) \leq V_T^+, \\ \mu_v \left(\frac{R_{\text{on}}}{D} \right) \left(\frac{i(t)}{i_{\text{on}}} \right) f(w(t)) & \text{if } v(t) < V_T^-. \end{cases} \quad (3)$$

$$\mu_v = 1.6 \times 10^{-12}, \quad i_{\text{on}} = 1, \quad i_{\text{off}} = 10^{-5}, \quad i_0 = 10^{-3}$$

μ_v is a mobility parameter with units of $\frac{\text{m}^2}{\text{s}} \Omega$. i_0 , i_{off} , and i_{on} are current parameters with units of A. These are all fitted constants proposed by Zhang et al. [11] to control the switching rate to match Li's empirical data [10], with units dictating the symbols used. They shouldn't be interpreted as corresponding directly to any specific microscopic phenomena. The interval $V_T^- \leq v(t) \leq V_T^+$ defines a dead zone in which the device state remains unchanged, providing the non-volatile behaviour necessary for use in synaptic circuits. Positive and negative suprathreshold voltages therefore drive the state in opposite directions, allowing independent directional control of the memristors resistance (memristance) during pulse programming. And will form the basis of producing Biphasic lobes in the memristor STDP section.

c) Window function:

To suppress unphysical state evolution near the boundaries $w = 0$ and $w = D$, the state equation was multiplied by a nonlinear window function,

$$f(w(t)) = 1 - \left(2 \frac{w(t)}{D} - 1 \right)^{2p}, \quad (4)$$

where $p = 2$ is a fitting parameter chosen [11]. This term reduces the rate of change of w as the device approaches either terminal state and therefore enforces smooth saturation.

RK4 Numerical integration was used on (3), applying (4) and then substituting w into (2), to get a final resistance that is used by all the AIST memristor simulations to follow. Note that resistance is sometimes used in this section because Zhang et. al uses memristance, however in the STDP memristor section, conductances will be used exclusively. Conductance is the reciprocal of resistance.

C. Pulse Programming Results

a) Pulse Programming:

In Fig. 2 the conductance is dictated by the slope of the line from the origin to the point of intersection of the currently applied voltage and the IV curve. The two saturated resistances R_{on} (shallower line) and R_{off} (steeper line) are seen as lines going through the origin, thus any points along it at any subthreshold voltages on the I-V sweep have the same conductance and resistance. It's only in the parts of I-V sweep where the voltage exceeds V_T^+ and V_T^- , does the curve evolve into a line that doesn't intersect the origin, thus describing an evolving conductance as suprathreshold memristance programming is happening. The top right curve demonstrates an increasing conductance and therefore a decreasing resistance, which is associated with filament formation. By the same argument the bottom left most non origin intersecting segment of the I-V curve represents filament dissolution, as the resistance is increasing. This paragraph assumes a counter clockwise rotation around the curve.




Fig. 2. A sub-volt sinusoidal sweep produces the pinched hysteresis expected from the AIST threshold model. AIST-style memristor simulation driven by a ± 0.5 V sinusoid at 500 Hz integration timestep used was $\Delta t = 0.1 \mu\text{s}$, with $V_T^+ = 0.37$ V, and $V_T^- = -0.19$ V thresholds. The solid black loop is the simulated I - V trajectory in amperes. The slope of the line from the origin to a point on the curve represents the conductance. Note that when the suprathreshold voltages are exceeded, the slope changes, reflecting the changing internal state of the AIST memristor.

Fig. 3 shows the programming input and terminal response of the threshold memristor. The device is driven by a ± 0.6 V, 2 kHz square wave over 0–1 ms with a $0.1 \mu\text{s}$ integration step, using the memristor parameters stated above. The current response is not simply a rectangular copy of the applied voltage. Current is low initially initially since resistance is high. As the filament grows as seen in Fig. 4 the filamentary phase occupies a larger proportion of the terminal gap D which has the lower resistance $R_{\text{on}} = 10 \Omega$ (see (2)). As a result the resistance decreases and so the current slowly increases (by Ohm's law since voltage is constant on the positive half of the square wave). The converse is true during the negative edge of the square voltage, where the opposite threshold is exceeded, resulting in an increasing resistance approaching

$R_{\text{off}} = 420 \Omega$. It is clear from the curves seen in Fig. 2, Fig. 4 and Fig. 3 that the suprathreshold conductance updates are non linear. featuring fast initial decays in currents for the filament decay regime. In the filamentary growth phase, growth is initially slow, and then increases till it almost rises linearly to the filamentary phase saturation resistance R_{on} . It is clear the filament formation is slower than filament dissolution.




Fig. 3. Square-wave pulse programming exceeds the update thresholds and produces state-dependent current peaks. Threshold-memristor simulation with $R_{\text{on}} = 10 \Omega$, $R_{\text{off}} = 420 \Omega$, $V_T^+ = 0.37$ V, and $V_T^- = -0.19$ V, driven by a ± 0.6 V square wave at 2 kHz over 0–1 ms using a fixed $0.1 \mu\text{s}$ integration resolution. Panel (a) shows the applied voltage and panel (b) shows device current; horizontal guides mark V_+ and V_- . Note the asymmetric current excursions and the gradual change in peak current magnitude across successive pulses.




Fig. 4. The bounded internal state integrates pulse history into a cumulative conductance trajectory. Same pulse-programming run as Fig. 3, with $R_{\text{on}} = 10 \Omega$, $R_{\text{off}} = 420 \Omega$, $V_T^+ = 0.37$ V, and $V_T^- = -0.19$ V, and a ± 0.6 V drive at 2 kHz. Solid black trace is the normalised state $\frac{w}{D}$, representing the proportion of the terminal gap occupied by the filamentary phase. dashed grey trace is conductance $G = \frac{1}{M}$ in mS. Note the stepwise updates in $\frac{w}{D}$ and the corresponding nonlinear change in G .

Fig. 4 normalises the internal state variable, showing the changes that occur during supra-threshold parts of the stimulating square wave changing the filament width and thus resistance. The slope of the changing filaments begin to flatten as it approaches the bounded switching limits. The nonlinearity and asymmetry in the suprathreshold memristance update explains the pinch seen on the left of the I-V curve.

D. Implementing STDP on a memristor Results

In the previous section pulse programming was introduced, demonstrating the memristors conductance G changes when suprathreshold voltages are exceeded. Thus the Δw variable seen in (1) is set as ΔG according to the pulse programming scheme shown in Fig. 5.

The neuron voltage-waveform (spike kernel) used was that described in the *Model Agnostic Spike Kernel* section with 50 ms and 150 ms time constants for depolarisation and hyperpolarisation respectively Fig. 14. The pre- and post neurons' spike kernels were applied to either terminal of the memristor according to $V_{post} - V_{pre}$, known as diffusive coupling. With the specified spike kernel, diffusive coupling voltage produces a peak voltage exceeding the positive V_{post} action potential during causal firing. During anti-causal firing, Diffusive coupling dips below the negative $-V_{pre}$ action potential. In both cases, as the $|\Delta t|$ approaches 0, the diffusive coupling voltage stays exceeding the signed action potential for longer. Combining this with the fact that memristor's update their conductance depending on if and for how long their suprathreshold voltage is exceeded, if the action potential is scaled to lie just below the positive update threshold V_T^+ by some small δV , then the diffusive coupling voltage becomes a natural LTP and LTD adjacent pulse-programming scheme that is very simple to implement in hardware (just subtract the voltages). Since The AIST memristor has $|V_T^-| < |V_T^+|$, the parts of the waveform with $V < 0$, had to be scaled by an additional factor $\frac{|V_T^-|}{|V_T^+|}$ in order to make the potentiation and depression lobes symmetric at the programming level. This also motivated the choice of using exponentials to define the kernel, since the duration of time for which the diffusive overlap remains surpassing either threshold voltage is invariant if both pre- and post- synaptic spikes are scaled by the same factor on either half of the voltage time plane.

Fig. 6 shows the conductance update produced by evaluating ΔG across three separate initial conductances, as the Δt between the Pre and Post spike kernels was swept from -100 ms to $+100$ ms. The timing dependence was approximately exponential on either side of the central sign-change region. As expected positive timing offsets produced potentiation, while negative timing offsets produced depression, demonstrating distinct LTP and LTD lobes.

The behaviour very close to $|\Delta t| = 0$ (< 7 ms) is the transition region where the causal and anti-causal pulse overlaps exchange sign, so the fitted exponential lobes describe the separate positive and negative branches away from the origin rather than during this cusp region itself. This can be interpreted as the discontinuity at $\Delta t = 0$ defined in (1), but since

such discontinuities is un-physical, a sign transition occurs over a narrow range.




Fig. 5. Pre/post spike overlap converts spike timing into a memristor programming voltage. The presynaptic and postsynaptic spike kernels are applied across opposite memristor terminals, so the voltage produced by diffusive coupling $V_{post} - V_{pre}$. Note that the negative coefficient for V_{pre} has already been used to invert the waveform, so the dotted curve shows the sum, for Causal and anti Causal firing. Changing the relative timing and sign of Δt changes how long this differential waveform exceeds the positive or negative update threshold, providing the pulse-programming mechanism that generates LTP and LTD.




Fig. 6. Spike-timing offset produces a biphasic conductance update with approximately exponential tails. AIST threshold-memristor model with $R_{on} = 10\Omega$, $R_{off} = 420\Omega$, $V_T^+ = 0.37$ V, and $V_T^- = -0.19$ V, driven by a pre/post superposition built from a piecewise exponential spike kernel. The timing offset is swept over $\Delta t \in [-100, 100]$ ms. Solid traces show the full simulated $\Delta G(\Delta t)$ waveform in mS for three representative initial conductances, including the central sign-switching cusp. Thin dotted overlays show shifted exponential fits applied only to the outer lobes.

The STDP reproduction scheme was applied on a sweep of 50 initial conductances G_0 to investigate the impact starting at a particular conductance value the memristor can possess (the conductance band), has on the amplitude of ΔG and time constant of the reproduced update rule.

In Fig. 7, ideal STDP and the obtained pulse programming updates are shown to agree strongly. In this fit, the amplitude parameter sets the scale of the conductance update, while the time constant described how quickly $|\Delta G|$ decayed as $|\Delta t|$ increased. Taking the logarithm of $|\Delta G|$ gives an approximately linear relation for an ideal exponential lobe, so the Pearson coefficient r of this linearised fit was used as a compact measure of fit quality. A strong exponential fit would give $|r|$ close to 1, and therefore $1 - |r|$ close to zero. $1 - |r|$ was found to be on the order of 10^{-4} , ranging from 1.143×10^{-4} to 1.334×10^{-4} for the positive lobe and from 1.145×10^{-4} to 1.336×10^{-4} for the negative lobe with the positive and negative lobes fitted separately to the non-transitional points of the G_0 -centred STDP sweeps. Here, $1 - |r|$ remained close to 10^{-4} . The exponential description was therefore strong, with only a small deviation across G_0 .

The fit quality also remained consistent across conductance because the small suprathreshold perturbation kept the window-function scaling approximately separable from the spike-overlap shape. Across much of the conductance range, changing G_0 mainly changed the magnitude of the fitted conductance update, while the fitted exponential decay with Δt remained nearly unchanged. The state dependence therefore appeared primarily in the A_{\pm} amplitude scale, rather than in a strong coordinate transform of τ or equivalently Δt . Fig. 9 shows the linear fits of τ_+ & τ_- . It give gradients of $+2.055 \times 10^{-7} \text{mS}^{-1}$ and $-2.056 \times 10^{-7} \text{mS}^{-1}$ respectively. The fit is very strong with $|r| > 0.99$ in both cases. The small gradients also inform that the time constants of the produced update rule barely depends on G_0 at all, even during saturation.

Fig. 8 shows the fitted exponential amplitude and time-scale parameters extracted from the same G_0 sweep. The conductance range of the AIST memristor was approximately 2.38 to 100 mS, corresponding to the 420 Ω and 10 Ω resistance limits respectively. The fitted exponential amplitude fell toward zero at both ends of this range, indicating that the window function attenuated conductance updates as the device approached either limiting state. This confirmed that the model preserved bounded conductance while still permitting STDP-like plasticity within the interior of the conductance band.

The same fitted-amplitude curve in Fig. 8 also quantified the usable update scale. The largest conductance updates occurred around the middle of the conductance band, where the attenuation from either boundary was weakest. Using the full 2.381 to 100 mS conductance range, whose width is 97.619 mS, the maximum fitted conductance update was 0.3296 mS, corresponding to approximately 0.338% of the full conductance-band width.

If attention is restricted to the broad central update region, roughly 30 to 78 mS, whose width was about 48 mS, the same update corresponded to approximately 0.687% of that effective

range. This sub-percent scale is consistent with multi-level memristive synapse simulations that assume hundreds of resolvable conductance states; for example, 256 states imply a characteristic update quantum of about 0.39% of the available range [12].




Fig. 7. Exponential-fit quality is approximately constant across initial conductance in the single-pair regime. For each G_0 , Pearson's r is computed between the shifted timing predictor and the log-transformed response $\ln|\Delta G|$ for each outer lobe after excluding the central cusp. The plotted metric is $1 - |r|$, where 0 indicates a perfect linearised exponential. Open circles are the positive lobe; open squares are the negative lobe. The absolute deviation remains small, but the non-zero opposite-sign gradients show a systematic conductance-dependent drift.




Fig. 8. Fitted STDP lobe magnitudes vary with initial conductance, while the fitted time constants have only a weak linear drift. Shifted exponentials are fitted independently to the positive and negative outer lobes of $\Delta G(\Delta t)$ for 50 initial conductances G_0 sampled over $[\frac{1}{R_{off}}, \frac{1}{R_{on}}]$. The cusp/transition region between the lobe extrema is excluded before fitting. Open circles denote the positive-lobe fit; open squares denote the negative-lobe fit. Note the suppression of $|A|$ (see (1)) near the conductance extremes.

Fig. 10 shows the relative lobe-amplitude imbalance $\frac{|A_+ - A_-|}{A_+ + A_-}$ across the same 50-point initial conductance sweep. The potentiation and depression amplitudes were very closely matched. The imbalance reached its minimum near the middle of the conductance band and increased smoothly toward the endpoints, but it remained small across the sweep (10^{-3} order of magnitude). This suggested that the reproduced STDP rule was relatively balanced and did not strongly favour either potentiation or depression. This balance was important because

a strong systematic bias in either direction would have made learning unstable.




Fig. 9. Positive and negative fitted STDP time constants track each other with a weak conductance dependence. Time constants τ_+ and τ_- are extracted from shifted exponential fits to the positive and negative outer lobes of $\Delta G(\Delta t)$ for 50 uniformly spaced G_0 samples, with the central cusp/transition interval omitted. Open circles denote τ_+ ; open squares denote τ_- . The fitted values remain close to the timing-dominated kernel scale.




Fig. 10. Potentiation/depression amplitude imbalance is small over most of the conductance range. From the same 50-sample STDP sweep, the plotted metric is $\frac{|A_+ - A_-|}{A_+ + A_-}$ using the fitted outer-lobe amplitudes after excluding the central cusp region. Values near zero indicate nearly symmetric lobe magnitudes, while peaks indicate conductance regions where set/reset nonlinearity breaks symmetry.

III. ANATOMY OF A NEURON

In the previous section, STDP was demonstrated on a memristor using a *Model Agnostic spike Kernel* for pre- and post-neuron spike voltage waveforms shown in Fig. 14, applied over a sweep of discrete relative spike times. This section aims to motivate the usage of Fig. 14, by investigating the qualitative shapes of the spikes produced by various Dynamical Neuron models. By analysing the stability of the stable points of a simplified 2D Dynamical Morris-Lecar Neuron, these qualitative voltage waveform characteristics were demonstrated, to justify Fig. 14. The main points demonstrated will be that neurons can be assumed to be charge

leaking capacitors (equivalent to RC circuits), and the firing of a neuron corresponds to a sudden discharging of this capacitor and a part of the energy inside of a chemical battery store. The stable points of these dynamical systems dictate that such spikes are triggered by the neuron’s capacitor voltage exceeding a threshold (suprathreshold voltage crossing). The accumulation of charge in a neuron that slowly leaks is known as leaky integration. The Leaky Integrate and Fire (LIF) neuron is then introduced, which simplifies such models by forgoing multidimensional dynamics and instead modelling each qualitative behavior of spiking neurons independently, switching between regimes when appropriate. Then CR is demonstrated on LIF neurons, discussing its implications for using stochasticity in the SNN section and aperiodic oscillators are shown as an alternative non-stochastic form of aperiodic drive.

Neuron models span a broad range of abstraction. At one extreme are biophysically detailed models, such as Hodgkin–Huxley (HH) type formulations, which reproduce electrophysiological behaviour with high fidelity but are comparatively costly to analyse and simulate. At the other extreme lies the artificial neurons which are modeled as simply the sum of a linear combination of scalar neuron values feeding into it in ANNs as described in the abstract. Between these ends sits the leaky integrate-and-fire (LIF) neuron, which retains much of the qualitative behaviour of biologically faithful neuron models while remaining relatively computationally efficient, and easy to implement in hardware. This balance made it a suitable candidate for the present work, where spike timing and synaptic interaction were of greater interest. However, starting with the detailed differential equation geometry that governs neurotransmitter and ionic channel kinetics, in more biologically accurate models, explains the pragmatic choices that inspire the LIF model.

a) Conductance-based neuron background:

A biological neuron is an electrically excitable cell that communicates by emitting short voltage pulses called spikes, with a peak amplitude referred to as an action potential. To understand how all of this arises, first consider that the ion permeable insulating lipid membrane of a neuron acts as the plates of a capacitor that contain gated holes, that form resistive ion flow channels between the conductive charged intracellular and extracellular solutions, providing a resistive charge leakage pathway. This reduces A neuron to a Capacitor Resistor circuit. A change in the balance of ions across the membrane therefore changes the membrane capacitor voltage. The charge imbalance between the solutions also forms an electro-chemical battery in series with these resistive ion channels.

The HH model is the benchmark for the most biologically realistic mechanistic neuron model in neuroscience, given by (5) describing the specifics of this leaky capacitor behaviour, with the opening and closing of three primary ion channel gating proteins modelled as voltage dependent variable resistors.

$$C_m \frac{dV}{dt} = g_{Na} m^3 h (E_{Na} - V) + g_K n^4 (E_K - V) + g_L (E_L - V) + I_{ext}. \quad (5)$$




Fig. 11. The Hodgkin–Huxley circuit represents spikes as voltage-gated ionic current balance. Equivalent circuit for a conductance-based membrane with capacitance C_m in parallel with sodium, potassium, and leak branches. Each branch contains a conductance, g_{Na} , g_K , and g_L , and a reversal potential, E_{Na} , E_K , or E_L . Note that this circuit retains biophysical channel structure that the LIF abstraction removes.

Referring to Fig. 11, C_m denotes the Membrane capacitance with a membrane voltage V , with several proteins that gate the ion channels in a voltage dependent way. These channels possess intrinsic conductances g . An external bias current I_{ext} is used as a coarse grain approximation of the ambient input current a neuron experiences from the local activity of neighbouring neurons. Positive charge flowing into the cell is called an inward current, increasing V . Outward current is the converse.

The inward current is modelled through positive sodium ions flowing into the cell via sodium channels which are gated by two distinct m and h proteins. The dominant outward current is the outward flow of positive potassium ions flowing out of the potassium channels, gated by a third n type protein. In (5) The variables m , h and n represent the proportion of all of each type of protein gate that is currently open. A small non gated leak current I_L provides passive relaxation toward a resting potential and represents the passive permeability of the lipid membrane outside of the specialised localised ion channel sites.

E_{Na} , E_K , and E_L are the reversal potentials of the sodium, potassium, and leak channels signifying the voltages toward which each channel would drive towards if it acted alone, dictated by the concentrations of the ions in and around the cell.

Each gate variable evolves according to a first-order equation of the form

$$\frac{dx}{dt} = \frac{x_\infty(V) - x}{\tau_x}(V) = \alpha_x(V)(1 - x) - \beta_x(V)x, \quad (6)$$

where $x \in \{m, h, n\}$ and $\alpha_x(V)$ and $\beta_x(V)$ are the voltage-dependent opening and closing rates of the corresponding channel-gate proteins. These competing rates drive the gate toward the steady-state equilibrium

$$x_\infty(V) = \frac{\alpha_x(V)}{\alpha_x(V) + \beta_x(V)}. \quad (7)$$

This makes the HH model a 4D system, that is hard to model and computationally expensive to integrate. A simplified Morris-Lecar style 2D approximation is used [6]. $m \approx m_\infty(V)$, $h \approx 1$, since empirically m reaches equilibrium much faster than the other gates, and n gates dominate potassium channel gating behaviour. This gives the Morris Lecar model:

$$C \frac{dV}{dt} = I_{in} + I_{out} + I_{ext} = g_{Na} m_\infty(V)(E_{Na} - V) + g_K n(E_K - V) + g_L(E_L - V) + I_{ext}, \quad (8)$$

$$\frac{dn}{dt} = \frac{n_\infty(V) - n}{\tau_n(V)} = \alpha_n(V)(1 - n) - \beta_n(V)n. \quad (9)$$

The opening and closing rates were approximated by a phenomenological Eyring-type exponential law [13], [14] in which membrane voltage linearly biases the proportion of open gates (effective activation barrier) and therefore makes the gate closing β_n , opening α_n transition rates exponential functions of V :

$$\alpha_n(V) = A_n \exp\left(\frac{V - V_n}{k_n}\right), \quad (10)$$

$$\beta_n(V) = A_n \exp\left(-\frac{V - V_n}{k_n}\right),$$

where V_n is a chosen parameter that gives $n_\infty(V_n) = \frac{1}{2}$. Now the stable points of (8) are investigated, by finding the curves of $\frac{dV}{dt} = 0$ and $\frac{dn}{dt} = 0$ (the n and V nullclines). This gives $\frac{dn}{dV} = 0$ when both curves intersect, producing a stable point of the system, and can be graphed and visualised.

By inferring the expression for $n_\infty(V)$ from (7), substituting (10) and setting $\frac{dn}{dt} = 0$, (9) produces the n -nullcline:

$$n = \frac{\alpha_n(V)}{\alpha_n(V) + \beta_n(V)} = \frac{1}{2} \left(1 + \tanh\left(\frac{V - V_n}{k_n}\right) \right). \quad (11)$$

The right-hand side is a logistic sigmoid, written here in its equivalent hyperbolic-tangent form.

Setting $\frac{dV}{dt} = 0$ in (8) gives the V -nullcline. Solving for n yields:

$$n = \frac{I_{ext} + g_L(E_L - V) + g_{Na} m_\infty(V)(E_{Na} - V)}{g_K(V - E_K)}. \quad (12)$$

In the following section the flow of n w.r.t V given by the Morris-Lecar dynamical system (9) is plotted alongside their nullcline curves given by (11) & (12) respectively.

A. Monostable Dynamical Neuron Results

Depending on the set parameters, the nullclines describe different neuron types. The two that will be the focus of this thesis are shown in Fig. 12 and Fig. 19. the latter will be addressed separately in its own section. For suitable parameter values, V and n nullclines intersect at three points as seen in Fig. 12. In that regime, the leftmost intersection is a stable node denoting the resting potential, the middle node is a saddle node, and the rightmost node is unstable. This is known as the excitable regime. Since n evolves more slowly than V , a brief positive voltage perturbation initially moves the state almost horizontally to the right in phase space (see grey circles and yellow triangle). The stimulation the neuron

faces is modelled by transient voltage perturbations from the resting state.

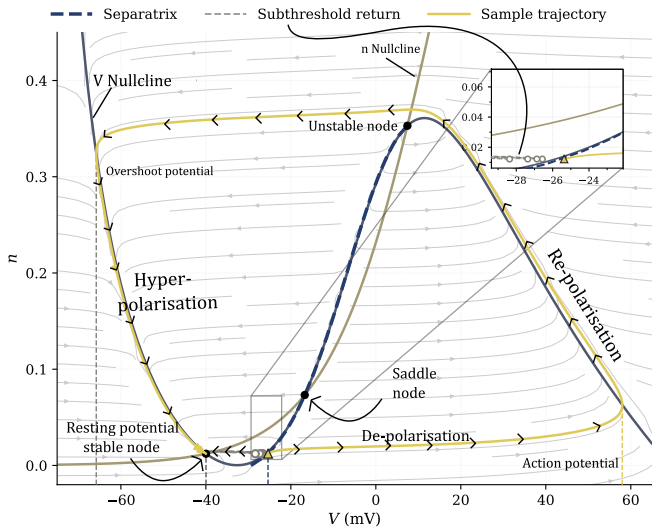


Fig. 12. The stable manifold of the saddle forms a separatrix that defines the excitability threshold. Reduced conductance-based neuron model with slow recovery is shown in the (V, n) phase plane. The separatrix (purple dash) separates subthreshold return-to-rest trajectories from a suprathreshold trajectory undergoing a spiking voltage large excursion, giving a geometric interpretation of the firing threshold voltage of excitable neuronal cells. $I_{\text{ext}} = 34.5 \mu\text{A}/\text{cm}^2$, $g_K = 8.0 \text{ mS}/\text{cm}^2$, $E_{K^-} = -84.0 \text{ mV}$, $g_{\text{Na}} = 4.4 \text{ mS}/\text{cm}^2$, $E_{\text{Na}} = 120.0 \text{ mV}$, $g_L = 2.0 \text{ mS}/\text{cm}^2$, $E_L = -60.0 \text{ mV}$, $\tau_n = 10.0 \text{ ms}$.

The stability of the saddle node forms the separatrix, which acts to set a firing threshold voltage. If the perturbation remains on the resting potential side of this separatrix, the trajectory decays back to it. The accumulation of subsequent transient voltage perturbations, that slowly decay back to the resting potential is known as leaky integration. If the amplitude and frequency of these perturbations are sufficient to cross the separatrix, the trajectory cannot return directly via the decay path and instead makes a large excursion through phase space, which can be seen in the yellow sample trajectory in Fig. 12. The trajectory flows to the right and up, where voltage rise rapidly past the threshold voltage, first arriving at the action potential. This step is called depolarisation, and is the initial voltage rise seen in Fig. 13. The trajectory then traverses up and then left very quickly, reducing the voltage to an overshoot potential, in a process known as repolarisation. The slower outward recovery current pulls the trajectory back down to the resting potential, known as hyperpolarisation. This part of the trajectory hugs the much slower V nullcline, and thus hyperpolarisation is always slower than re- and de- polarisation. Note that any perturbation during the large voltage excursion, doesn't meaningfully change the qualitative behaviour of the trajectory until the resting potential is reached again. This can be seen from the flow lines where all excursion trajectories flow alongside each other towards rest. This makes the duration of spiking, a refractory period, where excitations can't prematurely trigger another action potential, before the resting potential is returned to first. Such a spiking regime is

known as burst firing and Neurons that are current biased to behave like this are known as monostable neurons.

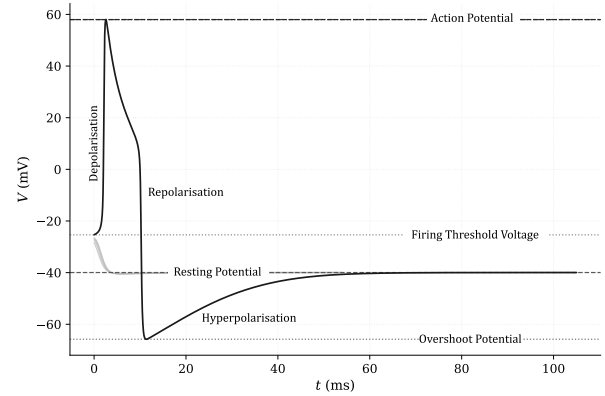


Fig. 13. In the excitable regime, only suprathreshold perturbations trigger a full spike-and-recovery excursion. Voltage time series are shown for initial conditions chosen relative to the phase-plane threshold boundary. Light grey traces show subthreshold trajectories that return to rest; the solid black trace is a suprathreshold excursion producing a typical biological neuron spike kernel.

B. Model Agnostic Spike Kernel

The spike waveform shown in Fig. 14 was modelled by a piecewise exponent kernel intended to approximate and capture the qualitative regions found in the voltage spikes of a diverse range biological neuron models. This follows on from the standard spike-response-model approach seen in Gerstner's work [15], [16]. For $t < 0$, the kernel rose exponentially with a short time constant to the action potential, representing depolarisation. At $t = 0$, it switched discontinuously to a negative overshoot potential, representing a very fast repolarisation. For $t > 0$ a slower negative exponential decay back to the resting potential V_{rest} (not to be confused with V_{reset}) was used to represent repolarisation. Each side of the kernel was truncated at five time constants. For any given time τ that a neuron fires, such a spike kernel may be used with the time translation $t - \tau$.

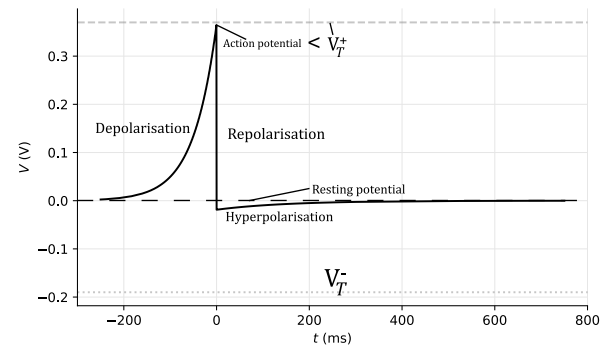


Fig. 14. The artificial spike kernel concentrates threshold-crossing drive near $t = 0$ while retaining an exponential tail. Piecewise exponential kernel used to build pre/post waveforms for the STDP sweeps, with $\tau_{\text{spike}} = 50 \text{ ms}$ and $\tau_{\text{tail}} = 150 \text{ ms}$. Solid black curve is the kernel voltage; dashed and dotted horizontal guides indicate the memristor thresholds $V_T^+ = 0.37 \text{ V}$ and $V_T^- = -0.19 \text{ V}$. Note that time stretching increases the duration of above-threshold overlap.

C. The LIF Model

The preceding sections introduced mechanistic conductance gated models for the membrane voltage, of a continuous dynamical system neuron. Leaky integration, threshold potential spiking and Spike kernels were established. In this section the multidimensional dynamical systems seen prior will be reduced to produce a monostable neuron, known as the Leaky Integrate and Fire (LIF) model first introduced in 1907 by Lapique [7].




Fig. 15. The LIF circuit reduces membrane dynamics to leakage, capacitance, threshold, and reset. Equivalent circuit for the leaky integrate-and-fire abstraction used in the coherence-resonance simulations. The input current charges C_m in parallel with R_m and a leak reversal E_L ; the comparator marks threshold V_θ and the reset label marks the post-spike return voltage.

a) Membrane Leaky Integration:

The equivalent LIF circuit is shown in Fig. 15. The membrane is represented by a capacitance C_m just like in HH and Morris-Lecar. The passive ionic leakage is represented by a resistance R_m , and the total synaptic and external drive was represented by an input current $I_{in}(t)$ corresponding to I_{ext} in (5) & (8). Let $V(t)$ be the membrane voltage relative to the resting potential V_{rest} . Current conservation gives

$$I_{in}(t) = I_R(t) + I_C(t). \quad (13)$$

Using Ohm's law and capacitor voltage charge equation, the passive membrane equation becomes

$$I_{in}(t) = \frac{V(t) - V_{rest}}{R_m} + C_m \frac{dV}{dt} \quad (14)$$

Multiplying by R_m and defining $\tau_m = R_m C_m$ gives

$$\tau_m \frac{dV}{dt} = V_{rest} - V(t) + R_m I_{in}(t). \quad (15)$$

If the input is written as an effective drive voltage,

$$V_{drive}(t) = R_m I_{in}(t), \quad (16)$$

then

$$\tau_m \frac{dV}{dt} = V_{rest} - V(t) + V_{drive}(t). \quad (17)$$

For a conductance-based synaptic input, the input current can instead be written as

$$I_{syn}(t) = g_{syn}(t)(E_{syn} - V(t)). \quad (18)$$

Substituting this into the LIF equation gives

$$\tau_m \frac{dV}{dt} = V_{rest} - V(t) + R_m g_{syn}(t)(E_{syn} - V(t)). \quad (19)$$

Defining the dimensionless conductance drive

$$q(t) = R_m g_{syn}(t), \quad (20)$$

this becomes

$$\tau_m \frac{dV}{dt} = -(V(t) - V_{rest}) + q(t)(E_{syn} - V(t)). \quad (21)$$

Crucially, the complicated voltage dependence and time evolution of the gating proteins has been omitted leaving just the passive membrane decay pathway corresponding to g_L in (5) & (8) and the subthreshold decay dynamics of the other active gates. This only models integrating the subthreshold stimulation, decaying back to the resting potential. Only this Leaky Integration phenomena is being modelled by (15).

b) Threshold Firing and reset:

To extend and include the qualitative spiking behaviour of neurons, the LIF neuron models a binary event variable

$$S[t] = \begin{cases} 1 & \text{if } V[t] \geq V_{thr} + \theta[t] \\ 0 & \text{otherwise} \end{cases} \quad (22)$$

where V_{thr} is the firing threshold potential and $\theta[t]$ is an optional adaptive threshold. When a suprathreshold crossing event ($V > V_{thr} + \theta[t]$) occurs at time t , a spike is recorded at that time, and the V of the system gets subtracted by a reset voltage V_{rst} :

$$V[t] \rightarrow V[t] - S[t]V_{rst}. \quad (23)$$

This discontinuous subtraction is a coarse-grained representation of post-spike refractoriness, and return to resting potential. It is often set to V_{thr} or just below to simulate repolarisation overshoot potentials.

The Adaptive Threshold $\theta[t]$ was a homeostasis variable used in the SNN, which increased after every spike and decayed otherwise. A simple discrete form with time steps $n, n + 1$ is

$$\theta[t_n + 1] = \lambda_\theta(\theta[t_n] + S[t_n]), \quad (24)$$

with $0 < \lambda_\theta < 1$. and S as before from (22). Recently active neurons therefore became temporarily harder to excite, helping prevent a small subset of excitatory neurons from dominating the competitive network. In hardware, again θ would be implemented as the voltage of a capacitor in an RC circuit with a much longer time constant and lower peak voltage compared to the Membrane capacitor, feeding into the circuit block responsible for classifying a suprathreshold crossing event.

c) spike shaping:

Spike times are recorded, but to produce an actual spike kernel, LIF is optionally extended with an external pulse shaping circuit where the spike kernel is free to be chosen, such as Fig. 14 used in the memristor STDP section.

d) Many to One Neuron coupling:

Many pre-synaptic neurons may share the same post-synaptic neuron so the input current to post synaptic neuron i was written as a weighted sum of the currents caused by the activity of all the connected pre-synaptic neurons:

$$I_{in}[t] = \sum_j w_{ij} x_j[t], \quad (25)$$

where the weight w_{ij} is the conductances of the memristor STDP synapse or some other synaptic proxy, between the post synaptic neuron i and the pre synaptic neurons j . $x_j[t]$ is the spike voltage waveform of each pre-synaptic neuron j .

D. Coherence resonance

In this section Coherence Resonance in the LIF neuron is demonstrated motivating the use of Poisson spike encoding in the latter SNN section. To quantify coherence resonance, the LIF neuron was driven by a Gaussian stimulation voltage centered around 0V of intensity σ . For each σ , 12 independent trials were simulated and the regularity of the resulting spike train was measured using the normalised inter-spike-interval (ISI) variance (see Fig. 16):

$$C_{\text{ISI}} = \frac{\text{Var}(\text{ISI})}{\langle \text{ISI} \rangle^2} \quad (26)$$

This produces a decaying random walk of the LIF Membrane Voltage, occasionally crossing the firing threshold (23). Following the parameter guidelines outlined by Morris Lecar, The LIF parameters used were: $V_{\text{rest}} = -65$ mV, $V_{\text{rst}} = -20$ mV (making post firing voltage 70 mV), $V_{\text{thr}} = -50$ mV, $\tau_m = 20$ ms (note this eliminates the need to specify R_m, C_m explicitly, with a current density being sufficient for any R in the system, see (15)), integration step $\Delta t = 0.1$ ms, and constant subthreshold drive $V_{\text{in}} = 13.7$ mV [6]. this is added to V_{rest} in (15), eliminating the need to specify I_{in} . Thus we should expect to see the resting voltage to be -51.3 mV, just below the -50 mV firing threshold potential, which is exactly what we see in the voltage traces most clearly seen in Fig. 17 a).

E. Coherence Resonance Results

e covariance




Fig. 16. Intermediate noise minimises LIF spike-timing covariance, demonstrating coherence resonance. Noise intensity σ is swept in a subthreshold LIF neuron, and the mean normalised ISI covariance $C_{\text{ISI}} = \text{Var}(\text{ISI}) * (\text{ISI})^{-2}$ is computed over 12 trials that produced enough spikes for an ISI statistic. The dashed vertical guide marks the minimum at $\sigma^* = 0.665$ mV. P_{ISI} denotes the proportion of 12 trials that produced at least 5 spikes. $\sigma < 0.36$ mV failed to reach $P_{\text{ISI}} = 1$




Fig. 17. Voltage traces confirm that the Normalised ISI Variance C_{ISI} minimum corresponds to the most regular noise-driven spiking. Representative LIF simulations are shown for noise below, at, and above the optimum from the sweep: $\sigma_{\text{low}} = 0.150$ mV, $\sigma^* = 0.665$ mV, and $\sigma_{\text{high}} = 2.044$ mV. Spike rasters are aligned below each voltage trace. The middle trace has the most evenly spaced spikes.

The plotted values in Fig. 16 is the 12 run average of C_{ISI} statistic. When σ was below 0.35 mV, less than 80% of trials yielded more than 5 spikes, thus increasing the error bounds on those values of C_{ISI} significantly. Yet still the V shape is clear with a minimum reached at $\sigma^* = 0.665 \pm 0.0001$ mV, indicating the lowest variance in the Inter spike timings. Any deviation from σ^* in either direction vastly increases the variance in the inter spike timing intervals.

The representative traces in Fig. 17 compare noise levels below, at, and above σ^* . Each panel shows the membrane voltage waveform with the firing threshold marked by a dashed horizontal line. The spike raster below each voltage trace uses vertical ticks to mark the detected spike times. These ticks should not be interpreted as a separate physical voltage; they indicate the time instants at which a downstream pulse-shaping circuit could be triggered by suprathreshold LIF firing events. Panel (b), corresponding to σ^* , shows the most regular spacing between events. Too little σ amplitude is shown in a) where the chance of the noise producing a high enough transient voltage for the neuron to fire is unlikely. Too much noise is shown in c) where the random walk exceeds the firing threshold too often and thus the stochastic nature of the noise signal shows up as highly aperiodic spike timings with a high inter spike time variance.

This forms the basis of the related phenomena known as Stochastic Resonance (SR), where in certain systems there exists an ideal amount of noise that maximises the correlation between upstream and downstream neurons spiking. This corresponds to maximising the distance spikes and therefore information, can travel through the SNN. CR demonstrates that noise minimises the variance in the firing of an LIF Neuron. The lower the stochasticity of the firing of a neuron, the more different it appears to downstream neurons that otherwise are themselves stimulated by ambient stochastic

noise. This means a lower C_{ISI} corresponds to a higher chance that spike signal gets distinguished from ambient noise by downstream neurons.

F. Oscillators and Aperiodic Spike Circuits

Having an aperiodic form of stimulation whether stochastic or otherwise is useful for Stochastic Resonance. This is explored further in the discussion section. Here, an aperiodic chaotic oscillator is shown. LIF neurons are monostable only, so this section returns to Morris-Lecar neurons.

Fig. 18 shows Morris-Lecar neurons $i, i \in \{1, 2, 3\}$ as circle nodes with membrane voltages Φ_i , where 1 drives 2, 2 drives 3 and 3 drives 1, each with directional conductances g_{12} , g_{23} and g_{31} , respectively [6]. By specifying particular parameters for g_{ij} and $I_{\text{ext } i}$ (see (8)), this network can produce both self-sustained oscillations with a fixed constant phase offset and also aperiodic firing, with the former being referred to as a Central Pattern Generator (CPG). CPGs are often found in the motor cortex of animals for autonomic nervous system functions that require rhythmic alternating motions such as breathing, and walking. Self sustained spiking is called tonic firing, with neurons capable of such firing are described as oscillator neurons. Referring back to the nullcline from the Morris-Lecar section, as the external bias current density $I_{\text{ext } i}$ is increased, the V nullcline in Fig. 19 is translated upwards until a saddle-node-on-invariant-circle (SNIC) bifurcation occurs, which produces an oscillatory limit cycle. Note how there is only one point of intersection between the nullclines. This is how the same model can describe both monostable neurons and also produce oscillator neurons, furthering the idea that such multidimensional dynamical models encode a large number of important distinct qualitative features in a very compact yet computationally expensive way.




Fig. 18. A three-neuron ring circuit provides a minimal directionally coupled oscillator network. Each neuron projects forward to the next neuron in a closed loop, so activity circulates around the ring rather than terminating at a feedforward output layer. The directed coupling defines the firing order and supports phase-structured dynamics, including synchronised circulation, stable phase offsets, and aperiodic timing structure in the oscillator simulations. The couplings are excitatory.

G. Oscillator and Aperiodic Spike Circuit Results

Fig. 19 shows the post SNIC bifurcation flow diagram of the dynamical system just like in Fig. 19, where the saddle and stable nodes got closer together till they meet and annihilate leaving just one unstable node remaining, as denoted by the intersection point of both nullclines. Once this happens, there is no longer a stable resting state to trap the dynamics, and the trajectory instead evolves onto a stable limit cycle, producing repetitive oscillatory spiking as can be seen in the membrane voltage time evolution in Fig. 20. These limit cycles possess basins of stability, as seen in the flow, where all trajectories approach the limit cycle, in the long term.




Fig. 19. Increasing bias current from $35\text{mA}/\text{cm}^2$ in Fig. 12 to $47\text{mA}/\text{cm}^2$ moves the system from an excitable separatrix geometry to a tonic-spiking limit cycle. The phase-plane is of the same reduced conductance model as shown in the monostable nullcline in Fig. 12 showing a post-SNIC bifurcation oscillatory regime. Note that now there is only one unstable point of intersection between the nullclines. The flow paths of non limit cycle trajectories also corroborates the persistent limit cycle.




Fig. 20. Post-SNIC bias produces tonic spiking with a locked constant interspike period. Voltage time series in the oscillatory regime show repeated spikes after transient decay. The absence of convergence to a resting potential and the near-regular repetition are consistent with a stable limit cycle with frequency $\omega = 0.550349 \frac{\text{rad}}{\text{s}}$.

The following parameters produced synchronized oscillations Fig. 21, with stable $2\frac{\pi}{3}$ phase offsets and a 1–2–

3 firing order. $g_{12} = g_{23} = g_{31} = 0.005$ mS, with $I_{\text{ref } i} = 10.0 \mu\text{A}/\text{cm}^2$ (for all i), $\omega_{\text{ref}} = 0.550349$ rad/ms (the base frequency of the oscillators), Due to the symmetry of the three-neuron chain, any of the $3! = 6$ possible firing permutations can be achieved by permuting the bias currents and coupling conductances across the labels. Fig. 21, Clearly shows the expected $\varphi_{12} = 2\frac{\pi}{3}$, $\varphi_{13} = 4\frac{\pi}{3}$.




Fig. 21. Identical Morris-Lecar neuron oscillators provide a 120° CPG reference case. The reference waveform is constructed from the same Morris-Lecar limit cycle and repeated with phase offsets $\varphi = 0$, $\varphi = 2\frac{\pi}{3}$, and $\varphi = 4\frac{\pi}{3}$. This is the zero-detuning reference condition. $I_i = 10 \mu\text{A}/\text{cm}^2$




Fig. 22. Persistent phase slippage rather than phase locking in a three-neuron diffusive network is demonstrated. Three reduced Morris-Lecar neurons with weak diffusive coupling are simulated under the current triple $(I_1, I_2, I_3) = (8, 10, 14) (\mu\text{A}/\text{cm}^2)$, and coupling constants $g = 0.005$ mS. The repeated near-alignments and separations of spikes are consistent with ongoing phase drift rather than a fixed phase relationship.

Following the literature aperiodic firing parameters were identified as $(I_1, I_2, I_3) = (8, 10, 14) (\mu\text{A}/\text{cm}^2)$, with the same coupling conductances as before [6]. The phase offset between these neurons do not settle because their oscillation frequencies are sufficiently different and the coupling strength is too weak to produce synchrony. Fig. 23 corroborates the expectation that the phases diverge monotonically. This is

further seen in the voltage waveform Fig. 22, where the firing phase and order appear chaotic. Simulation times of 4700 ms and 1000 ms are shown respectively to highlight that this aperiodicity isn't just a transient behaviour before the system settles to synchrony.




Fig. 23. Unwrapped phase differences drift approximately linearly, confirming absence of stable phase locking. Same three-neuron network as Fig. 22; phases are extracted from late-time voltages and differences are unwrapped. Both differences continue to increase without settling, which is direct evidence of persistent phase drift.

IV. FEED-FORWARD STDP SNN

The previous sections introduced the concept of an STDP synapse, Then AIST memristors were shown to reproduce this rule well, when spike shaped LIF neurons were applied across it with a diffusive coupling voltage.

The LIF neuron was also introduced and investigated including under stochastic stimulation. This section focusses on highlighting the value of reproducing STDP, with its implications for training SNNs being explored. The LIF Neuron, STDP model and Poisson based stochastic aperiodic drive will be incorporated into an SNN trained to recognise handwritten digits. Some new network level ideas will be introduced, such as homeostasis mechanisms which act to prevent a subset of neurons from firing too often, and sparsity, which is any mechanism that promotes a network to store learned representations of training data on as few synapses in the network as possible. This maximises representational efficiency given a finite number of synapses in an SNN.

a) Synaptic Plasticity Update:

The symmetric biphasic STDP (1) rule was used

b) Network Dynamics and Synaptic Integration:

Uptill this point only excitatory coupling between neurons have been considered, where the pre-synaptic Neuron drives the post synaptic neuron, increasing its stimulation and thus its chance of firing. To implement a functional competitive in the network, The architecture relies on the interplay between excitatory (E) neurons and inhibitory (I) neurons. These Neurons actively discharge the membrane capacitors of the post synaptic neurons they feed into, reducing their chance of firing. Thus the LIF model used is identical to (21) with the addition of k inhibitory and j excitatory pre synaptic neurons feeding into a post synaptic neuron i :

$$\tau \dot{V}_i = (E_{\text{rest}} - V_i) + \sum_j g_{ij}^E (E_{\text{exc}} - V_i) + \sum_k g_{ik}^I (E_{\text{inh}} - V_i) \quad (27)$$

As before τ represents the membrane time constant, which dictates the charge leak & integration rates of the neuron.

For each presynaptic spike, the corresponding synaptic ion channel conductance (g_{ij}^E and g_{ik}^I) undergoes an instantaneous jump followed by an exponential decay:

$$\tau_{g_e} \dot{g}_{ij}^E = -g_{ij}^E \quad \text{and} \quad \tau_{g_i} \dot{g}_{ik}^I = -g_{ik}^I \quad (28)$$

with $g_{ij}^E(t=0) = g_{ij}$, $g_{ik}^I(t=0) = g_{ik}$,

where g_{ik} & g_{ij} are the learned conductances using the STDP rule. The product of these exponentially decaying conductances and their respective fixed channel potentials results in an exponential decaying current flow into the post neuron. In contrast, the STDP Memristor picture had a persistent non zero synaptic conductance g_{xy} from pre x to post y neurons and a varying voltage spike kernel. This means for the exponentially decaying current proposed by Diehl Cook to be equivalent to the LIF STDP memristor synapse, the spike kernel needs to contain an exponential voltage rather than a fixed gate potential seen in LIF and Morris Lecar Neuron models. This is another motivation behind using exponential voltages in Fig. 14.

The distinction between E and I neuron populations lies in their reversal potentials. While E_{exc} drives V_{post} upward, the inhibitory reversal potential (E_{inh}) is set to a value significantly more negative than the resting potential (E_{rest}). Consequently, the driving force ($E_{\text{inh}} - V_i$) is strongly negative, ensuring the inhibitory current kernel acts with an opposite sign to excitatory input. Rather than simply opposing excitation, when inhibitory Neurons fire they accelerate the removal of charge from the membrane capacitor of the post-synaptic neuron, discharging it at a rate exceeding its passive leakage.

c) Training Data:

Input images were 28×28 monochrome images of a curated set of 420000 human handwriting examples of the digits 1 to 9, by the Modified National Institute of Standards and Technology (MNIST) [17], encoded as Poisson spike trains with firing rates proportional to pixel intensity. Each image was presented for 350 ms. Following the original procedure, the maximum pixel intensity, out of 255 equally spaced possible intensities, was divided by 4 and converted into frequencies, giving initial input rates between 0 and 63.75 Hz. network input spikes were Poisson distributed about repeated time instances occurring with a period given by each pixel's intensity encoded into a frequency. If fewer than five spikes were emitted by the entire excitatory population of the SNN during presentation of a particular example, the maximum input rate was increased 5Hz and the same image was presented again until sufficient activity had been elicited. This harks back to the CR section, which inspires the use of stochastic Poisson input timings, to improve network accuracy. Unlike the simple case of a single Neuron, finding σ^* directly is difficult. Instead the adaptive Poisson frequency range as a variable σ proxy.

Between successive images, a 150 ms interval with no input was inserted so that fast state variables could decay back toward baseline.

d) Network Connectivity:

The reproduced network consisted of three layers, as shown in Fig. 24. The first was an input layer of $28 \times 28 = 784$ neurons, one for each (Poisson encoded) pixel intensities of the MNIST image. The second was the main learning excitatory layer. All input pixel spike trains were fed into every Excitatory neuron, with STDP synapses in between. The pattern of input pixels to which a particular excitatory neuron has become most sensitive is known as the receptive field of each excitatory neuron; because all pixels are fed into each excitatory neuron, the synapse strengths for each input pixel can be visualised like a 28×28 MNIST image. There is an equal number of inhibitory neurons in the third layer. Each inhibitory Neuron is exclusively assigned its own excitatory neuron from the previous layer that drives it with a fixed non trainable synapse. Every Inhibitory neuron connects backwards all-to-all except to that the neuron driving it in the excitatory layer.




Fig. 24. Diehl–Cook network architecture used for unsupervised MNIST learning. Pixel intensities from a 28×28 MNIST image are converted to Poisson spike trains that project all-to-all to the excitatory population. Each excitatory neuron drives one inhibitory neuron, and each inhibitory neuron projects back to all excitatory neurons except its source neuron, creating lateral inhibition and a soft winner-take-all dynamic.

A. Competition and Homeostasis

a) Lateral Inhibition:

Competition is a type of sparsity mechanism. Neurons react to inputs and compete, with whichever neuron emitting the most charge, having its synapse strengthened more than others reacting to the same thing. It was implemented on the excitatory neurons through a dedicated inhibitory population, using the connectivity shown in Fig. 24. Each excitatory neuron projected to exactly one inhibitory neuron, and each inhibitory neuron projected back to all excitatory neurons except the one from which it received input. This connectivity created lateral inhibition and hence a soft winner-take-all dynamic. The

strongest neuron most strongly inhibited all others. Because it reacted most strongly, the lateral inhibition it experiences is also less than what all other excitatory neurons experience. This made it more likely that only a small subset of excitatory neurons reacts most strongly to a given training example, thus promoting sparse representations in the excitatory neurons.

Note that excitatory neurons are given a longer time constant than inhibitory neurons [9]. Input digits are encoded as Poisson spike rates so a larger excitatory τ allows the membrane to average over many stochastic spike arrivals, better gauging the true pixel intensity, and averages pixel information across multiple images providing more stable learning, whilst inhibition can still act quickly on a single training example providing a per image winner neuron. This forces the inhibitory neurons to only provide competition, and does not meaningfully participate in building a rich learnt representation of the training data.

b) *Adaptive Thresholding:*

To prevent a small subset of excitatory neurons from dominating the response, the network also implemented a homeostasis mechanism through an adaptive threshold. This was identical to that described in equation (24). A neuron that had recently been highly active becomes temporarily harder to excite, redistributing competition across the population and encouraging receptive fields to differentiate, and to get excitatory neurons to detect different digits from one another.

B. *Training*

Training was unsupervised. Synaptic weights were learned from presentations of the MNIST training set without class labels. Labels were introduced only after learning had finished. At that stage, learning was disabled, neuron thresholds were fixed, and each excitatory neuron was assigned to the digit class for which it showed the highest average response over the training set. During evaluation, the predicted digit was taken to be the class whose assigned neurons exhibited the highest mean firing response. No separate classifier was placed on top of the spiking network. An Nvidia RTX 4090 graphics processing unit (GPU) was used to simulate networks with 100, 400, and 1600 excitatory neurons.

C. *Training results*

Fig. 26 shows a clear population-size effect. The $N = 100$ run reaches a final training-window accuracy of 74.8% after 40000 images, the $N = 400$ run reaches 94.4% after 180000 images, and the $N = 1600$ run reaches 98.0% after 420000 images. The larger model does not merely improve the final checkpoint: it also spends a much larger fraction of training in the high-accuracy regime, consistent with the original Diehl–Cook observation that larger excitatory populations provide more feature-specialised units.

The same scaling trend is visible on the held-out test set in Fig. 25. $N = 25$ results was from a trial run during the process of getting the simulation working. It is expectedly weak in performance, with too few neurons to learn what handwritten characters correspond to which digit. Moving from $N = 100$ to $N = 1600$ raises the mean test accuracy from 79.398% to 93.544%. This is the strongest network-

level result in the simulation set: the architecture does not only fit its training-window metric, but retains substantially improved performance on the test data. For this graph, the 400 neuron result was omitted because the simulation crashed before those figures could be obtained.




Fig. 25. Test accuracy improves strongly between the pilot simulation and the 1600-neuron model. Bars show mean test accuracy for available population-size simulations with $N = 25$, $N = 100$, and $N = 1600$ excitatory neurons; error bars show the standard deviation over repeated test evaluations where available. The $N = 1600$ model reaches 93.544% mean test accuracy, compared with 79.398% for the $N = 100$ run and 35.938% for the local pilot simulation.

The confusion matrix in Fig. 27 shows where the remaining errors occur. Highest probability (yellow colour intensity) lies on the diagonal, confirming that the output assignment rule is meaningful across all ten classes. However, the diagonal is not perfectly uniform. The diagonal mean of approximately 93.55%, giving a 7% error rate and agrees with the repeated-test mean accuracy. The final network is better interpreted as a successful qualitative reproduction of unsupervised learning rather than as a fully optimised state of the art classifier, and there are certainly many ANNs that can outperform it (Gradient Descent exceeds 99% accuracy as noted by Diehl & Cook [9]).




Fig. 26. Increasing the excitatory population size improves the Diehl–Cook learning trajectory. Training accuracy is plotted against MNIST presentations for reproduced Diehl–Cook simulations with $N = 100$, $N = 400$, and $N = 1600$ excitatory neurons using seed 42. Line styles distinguish network size and markers indicate the final sampled point. Final checkpoint accuracies are 74.8%, 94.4%, and 98.0% respectively.




Fig. 27. The 1600-neuron network classifies most MNIST digits correctly as seen on this confusion-matrix diagonal. Rows are true labels y and columns are predicted labels \hat{y} for the $N = 1600$ simulation evaluated on 10000 MNIST test images with 10 repeated evaluations. Cell shading gives the row-normalised prediction probability and annotations give percentages for visible cells. The mean diagonal probability is 93.55%, consistent with the mean test accuracy.

The mechanistic reason this architecture can classify without labels is visible in the receptive fields of the excitatory neurons. Fig. 29 shows examples of receptive fields that clearly resemble digit strokes or whole digit prototypes. This indicates that lateral inhibition, adaptive thresholds, and STDP-like competition split the excitatory population into neurons responding preferentially to different recurring input features (specialised into detecting different digits).




Fig. 28. Averaging assigned receptive fields reveals the class-level templates learned by the 1600-neuron network. Each panel averages all learned input weight vectors assigned to one MNIST digit in the $N = 1600$ simulation. Panel labels give the assigned digit, and darker pixels indicate larger mean synaptic weights. The class averages are blurrier than individual neurons because they pool multiple stroke styles, but they preserve recognisable digit geometry.

The averaged receptive fields in Fig. 28 provides a more compact view of the same effect. Averaging within each assigned digit naturally blurs individual stroke variation, but the class-level templates remain recognisable. This supports the interpretation that the improved accuracy in Fig. 26 and Fig. 25 is associated with genuine unsupervised feature formation.



Fig. 29. The 1600-neuron network learns digit-selective receptive fields. Each tile shows one input-to-excitatory weight vector reshaped to the 28×28 MNIST pixel grid for the $N = 1600$ simulation. Rows group neurons by their assigned digit label, and columns show the eight most selective neurons within each assignment group; darker pixels indicate stronger learned input weights. The emergence of digit-like templates indicates distributed unsupervised feature specialisation.

Finally, Fig. 30 gives a useful check on the failure mode. These failed classification examples are understandable since they demonstrate genuine digit ambiguity, and poorer handwriting. Despite the successes of the receptive fields in generalising digits, it is not foolproof, and the a digit can still be misclassified by this SNN.




Fig. 30. Residual 1600-neuron errors are concentrated in visually ambiguous MNIST examples. The first 24 misclassified examples from the $N = 1600$ test evaluation are shown. Each label gives true digit y followed by predicted digit \hat{y} . These examples are not a random test subset; they diagnose the failure modes remaining after the model reaches 93.544% mean test accuracy.

V. DISCUSSION

A. Device-Level Plasticity

Initially the STDP rule (1) was defined. STDP can be explained as promoting meaningful computations to take place in an SNN without explicitly defining an algorithm to solve a particular task. Computation in an SNN occurs by manipulating the time distribution of spikes as they travel through the network, from input sensory neurons to output action neurons. When some meaningful coordinated computation is being done by some neurons, then it may be expected that a group of connected neurons will fire close together in time, in the order of the feed forward excitatory influence they have been arranged in. This way LTP reinforces any task agnostic, well coordinated computations that may randomly arise in the SNN when sensory input is fed in. The idea is that applying STDP after every exposure to many different training examples for a particular task, only the subset of well coordinated computations actually relevant to solving the training task, gets reinforced, potentiating the same synaptic pathways consistently. IN contrast, chance based correlations found in noise or unrelated parts of the training data, change which synapses get strengthened every new training example. Thus these irrelevant yet well coordinated computations average out to contribute very little to the overall synaptic structure of the network. What is left may be thought of as groups of neurons forming specialised circuits that perform some task relevant computations that feed into each other to give the desired output. LTD is there to weaken synapses that may happen to exist between a pair of neurons that have settled into different independent specialised circuits. Due to this, Spikes produced by this pair aim to work independently each other, so the downstream neuron may consistently (though not always) fire before the upstream neuron, producing $\Delta t < 0$ (the condition for LTD). thus synapses that fit this case, are weakened in STDP [18], [19], [20].

LTP weakens with $\Delta t \rightarrow \infty$ since long time separations in the firing between excitatory neurons (that feed forward their stimulation), because many pre synaptic neurons can feed into one post synaptic neuron (via a distinct synapse for each pre-synaptic neuron). A larger time gap between a pre and post neuron spiking has a larger chance that a different other pre neuron caused the suprathreshold voltage crossing that triggered the firing of the post synaptic neuron. This ambiguity increases with time so it is sensible for LTP to decay with spike timing separation. By the same argument, LTD also decays when $|\Delta t| \rightarrow \infty$.

The first part of the project established the physical link between an AIST memristor and an STDP-like synapse. The AIST model in Fig. 2 and the pulse-programming response in Fig. 3 and Fig. 4 show the same central mechanism of device conductance evolving with the time integral of a non linear suprathreshold update function (3) when driven by either a positive or negative suprathreshold voltage. This corresponds to the formation and dissolution of a lower resistance filamentary phase in the crystal structure of the AIST substrate. The resistance of such a device depends on its own recent programming history, with a bounded resistance caused by the

fact that the memristor's terminal gap is finite (3nm). When no external voltage is applied the width of the filament phases persist, making its conductance G suitable for storing synaptic strength w from (1) and ΔG as Δw with pulse programming.

The diffusive coupling $V_{\text{post}} - V_{\text{pre}}$ of pre and post synaptic neurons, combined with a specially chosen spike kernel Fig. 14, resulted in an overlap waveform that converted relative spike timing into a suprathreshold pulse duration, signed oppositely between causal and anti causal firing cases. In Fig. 6, causal timing produced LTP and anti-causal timing produced LTD as expected, because the spike kernel under diffusive voltage crosses opposite thresholds depending on spike order as seen in Fig. 5. The full ΔG waveform produced showed strongly exponential biphasic lobes ideal STDP demands. However, the two lobes were separated by a cusp where the sign of the effective programming pulse changed about $|\Delta t| = 0$. This is why the fits in Fig. 8, Fig. 7, Fig. 9, and Fig. 10 agree well with pearson $r > 0.999$ but only outside of this cusp zone ($|\Delta t| < 7\text{ms}$). this has the potential to form weak zones where the synaptic strength update is weaker than expected by Ideal STDP. However the synaptic sign is still correct, so although synaptic updates may not be as strong near zero, the correct type of update still happens. Compared to the approximately 65 ms time constants found in Fig. 8, $1 - e^{-\frac{7}{65}} \approx 10\%$ of the overall Ideal STDP strength is lost by having this cusp region. This is not Ideal for use in training since it further reduces the strength of the updates, potentially making training slower.

The ΔG amplitude also falls to zero near the conductance limits. This prevents runaway potentiation and depression when such a memristor STDP synapse is used inside of an SNN for example. Over the region where the amplitude is the highest and the flattest (30 to 78 mS) the update amplitude percentage is 0.6% of this window which falls in line with estimates made by Kim et. al [12], for sufficient amplitude ratios needed for effectively training an SNN using the STDP rule. All in all this strongly supports the use of such AIST memristors as STDP synapses, fulfilling the first aim of this project. Although there is the caveat from the 10% cusp attenuation from the previous paragraph that needs to be considered. Also the device model used did not take into account any device stochasticity and device-to-device variability that matters in a fabricated SNN array, that affect STDP reproducibility. Approaches such as active spike amplitude normalisation, batch together many synapses and re-normalise the action potentials of neurons explicitly to stabilise the synaptic update amplitude to be more consistent across the conductance range, and to artificially shrink the cusp region. These form a good passage for further exploration.

Then the groundwork for the SNN section was laid with the LIF neuron being derived from more complicated dynamical neuron models, by replicating threshold firing and Laaky integration dynamics of biological ion channels. Then followed Coherence Resonance results for the LIF neuron using the parameters described in the CR section, with the normalised interspike variance C_{ISI} from Fig. 16 was minimised for a particular level of gaussian noise amplitude $\sigma^* = 0.665\text{mV}$. This informed the use of some form of stochastic drive in the

SNN architecture. Other Network mechanisms such as lateral inhibition and adaptive thresholding was used to implement competition and homeostasis respectively.

CR demonstrates that noise increases the excitability of the LIF neuron, so it is easier for a neuron to transmit information to downstream. This is of course also true for just reducing the firing threshold voltage of LIF neurons, but that would also make the SNN less tolerant of external noise. It is by the fact that Gaussian noise is Aperiodic and centered on 0V, that on average the overall increase in the excitability of the neuron is 0. Adding an aperiodic background drive means that connected chains of neurons will only occasionally have a transient boost in the excitability of the downstream neuron favoring the chance that a coincidental upstream neuron firing will trigger a cascade of downstream neurons to fire. Since STDP strengthens synapses when downstream neurons fire shortly after upstream neurons, it asymmetrically favors and strengthens the stimulation the upstream neuron has on the downstream neuron. In the case where the transient gaussian noise voltage waveform made the downstream neuron harder to fire (by being negative), there is a lower chance that the pre synaptic neuron causes the post neuron to have a suprathreshold voltage crossing event to begin with, reducing the chance for there being a spike pair for STDP to react to. this way STDP asymmetrically favours aperiodic voltage sources superimposed into the membrane voltages of spiking neurons in SNNs, increasing the distance information can travel in the network, whilst maintaining its original noise tolerance. Using randomness to achieve this is known as Stochastic Resonance (SR). This motivated investigating aperiodic Chaotic biologically plausible oscillators presented in its own section, and the use of Poisson distributed sensory neuron stimulus in the Diehl Cook SNN.

The Diehl–Cook unsupervised SNN reproduction combined the preceding pieces into an explicit network-level test of whether the local STDP rule could facilitate unsupervised learning. The architecture in Fig. 25 highlights a crucial mechanism. All input pixels project to all excitatory neurons, each excitatory neuron projects to one of its own inhibitory neuron, and inhibitory feeds backwards suppresses every other competing excitatory neurons in the SNN. This is known as Lateral inhibition and creates a soft winner-take-all mechanism. When one neuron responds strongly to an input digit, lateral inhibition prevents all neighbouring neurons from learning the same pattern at the same time, thus making learned representations to be more efficiently stored on a minimal number of synapse weights. STDP then strengthens the weights between the poisson coded active input pixel neurons and the winning neuron, making the neuron more selective for that feature in future presentations. Homeostasis is a complementary mechanism. The adaptive threshold described in the LIF section makes recently active excitatory neurons temporarily harder to activate, so competition does not collapse onto the first few neurons that happen to react strongly to something. This allows for more training examples to have a chance to influence the learned features, without the SNN getting too hung up on some correlation that may occur in early examples that might have just been noise. This

explains why the receptive fields of each neuron specialises for a specific digit and in Fig. 29 and Fig. 30 appear to visually encode the average of many different common handwriting variations, demonstrating the ability of the network to learn the general concepts behind the features present in the training data, without the use of any image labels, or given any external feedback on how accurate its predictions were. Despite this, the SNN failed on around 7% of examples on average with the final tested accuracy being 98%. In summary lateral inhibition strengthens transient correlations found in the most recent training examples, while homeostasis slowly redistributes the incremental knowledge gained each training example, whilst protecting previously learned knowledge from being overwritten by strong transient updates. This pairing is why the network resembles vector quantisation or self-organising maps seen in traditional ANNs that require expensive synchronised global weight matrix updates. The accuracy results are strong in the broader context of Neuromorphic computing. Diehl and Cook report average MNIST test accuracies of 82.9%, 87.0%, 91.9%, and 95.0% for $N = 100, 400, 1600,$ and 6400 excitatory neurons respectively [3]. The reproduced $N = 1600$ model reached 98.0% final training-window accuracy in Fig. 26 and 93.544% mean test accuracy in Fig. 27, which is within margin of error of the original author’s results. Note that the 6400 neuron simulation was not reproduced due to restricted funding for compute resources compared to the Authors. The confusion matrix in Fig. 28 and the a sample of the SNN’s failures in Fig. 31 show its limitations. Ambiguous handwritten examples can still activate the wrong assigned population. Such failure modes are expected from a shallow unsupervised architecture, since it learns useful stroke and digit templates, but it does not have the hierarchical abstraction of more layered networks, that allow them to abstract and handle ambiguity better. Diehl and Cook make the same broader comparison in their benchmark table, where supervised or less biologically constrained systems performed worse, but state of the art conventional Gradient Descent methods can reach higher accuracies (99%). Diehl & Cook’s work remain valuable because of its simplicity, unsupervised nature and class leading performance amongst biologically inspired synaptic update/ learning rules. Aside from some ANNs getting higher accuracies, the 98% peak accuracy obtained still strongly supports the conclusion that STDP is sufficient to perform very well as a training rule for an SNN, answering the second aim of this project. Further work could involve modifying STDP to perform even better, to try and reach parity with ANNs and to investigate ways in which much larger SNNs with more neuron layers can be implemented.

B. Overall Assessment

Taken together, the results support the project aim: a compact AIST-style memristor model can reproduce a local, timing-dependent conductance update, and a Diehl–Cook style STDP SNN can learn recognisable MNIST structure with strong accuracy. The strongest conceptual chain is that pulse timing becomes AIST memristor conductance updates, memristor conductance updates mimics the STDP synaptic

plasticity rule in an SNN, lateral inhibition turns plasticity into competition inside of an SNN, and homeostasis prevents competition from degenerating into a few neurons monopolising SNN behaviour. This chain gives a coherent neuromorphic story from device physics to network function, motivated by the energy efficiency of neuromorphic systems seen in the literature.

The main caveat is that the results in this paper do not directly demonstrate the hardware efficiency that motivates neuromorphic computing. Training was slow and inefficient because event-driven STDP still had to be simulated sequentially on conventional synchronous hardware. A physical memristor crossbar, by contrast, would update many synapses asynchronously in parallel through local device physics, which is the reason memristive synapses and spiking neuromorphic hardware remain attractive. Existing neuromorphic processors already show the appeal of this direction: TrueNorth demonstrated large-scale spiking networks performing tasks successfully at milliwatt power levels [21], with Loihi demonstrating similar energy gains on programmable Spiking Neural architecture [22]. No power measurement was performed in this project, so those references should be treated as motivation rather than evidence in favour of the presented SNN compared to traditional ANNs.

VI. CONCLUSION

This project demonstrated a coherent simulation pathway from device-level memristor dynamics to network-level unsupervised learning. The AIST threshold-memristor model reproduced the core ingredients required for a synaptic device: non-volatile state-dependent conductance, threshold-ed pulse programming, bounded saturation, and polarity-dependent set/reset behaviour. When driven by overlapping pre- and postsynaptic spike kernels, the model produced a biphasic STDP-like conductance update, with causal timing producing LTP and anti-causal timing producing LTD. The fitted outer lobes were approximately exponential across the usable conductance range, while the update amplitude correctly fell toward zero near the conductance limits. The conductance update was weak when the spike timing magnitude fell below 7ms, and the lack of control in the strength of the updates is a drawback of using AIST memristors as STDP synapses. Further work could explore the many ways these drawbacks are overcome by supplementing AIST memristors with additional circuitry. Overall this still supports the use of AIST memristors for STDP synapses well and shows promise, thus the first aim of the project was successful.

The neuron-model simulations helped justify the abstractions used in the network section. Conductance-based and reduced phase-plane models showed how threshold crossing and spike generation arise from continuous membrane dynamics, while the LIF coherence-resonance simulations showed that noise can improve spike-timing regularity at an intermediate noise strength. The oscillator and ring-circuit simulations further illustrated how structured or aperiodic timing sources could be used to as a pseudorandom noise source.

At network scale, the reproduced Diehl–Cook style architecture showed strong unsupervised learning behaviour [9]. Lateral inhibition created competition between excitatory neurons, while adaptive thresholds provided homeostasis so that learning was distributed across the population rather than monopolised by a small number of early winners. The resulting receptive fields became digit-like, and the $N = 1600$ model achieved 98.0% final training-window accuracy and 93.544% mean test accuracy. This result compares favourably with the original Diehl & Cook paper. Compared to previous neuromorphic architectures such, this marks a step up in performance whilst being unsupervised. For example The work done by Hussain et al. [23] just a year before, using the supervised Morphology rule only achieved an accuracy of 90.3% [9]. The use of STDP in this SNN provides an elusive pathway to implement such devices with memristive STDP synapses providing low energy spike circuits and asynchronous local plasticity, where autonomous synaptic strength updates occur naturally from device physics, allowing for massively parallel computation at a per synapse level. The two aims of the project are therefore complementary, as the AIST memristor provides a very compact STDP synapse implementation (the terminal gap is just 3nm), providing a proof of concept to implementing low energy artificial STDP synapses in hardware, and the STDP SNN gives STDP synapses a meaningful role in facilitating highly parallel asynchronous synaptic updates, that this SNN demands. Future work should therefore focus on moving beyond the concept level, and towards tackling the device non-idealities present in fabricated arrays, increasing the depth of the networks (currently there's only 3 layers) and size of the networks, to be able to take on more difficult and practical real world machine learning tasks. Put together, the second aim of the project was also successfully carried out. All together both aims were thoroughly investigated, and plans for future work have been outlined.

VII. ACKNOWLEDGEMENTS:

All reproduction graphs were reproduced independently by Tabeeb Rahman. Fig. 24 comes from [9].

REFERENCES

- [1] F. Rosenblatt, "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain," *Psychological Review*, vol. 65, no. 6, pp. 386–408, 1958, doi: 10.1037/h0042519.
- [2] B. Widrow and M. E. Hoff, "Adaptive Switching Circuits," *1960 IRE WESCON Convention Record*, vol. 4, pp. 96–104, 1960.
- [3] C. E. Shannon, "A Mathematical Theory of Communication," *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948, doi: 10.1002/j.1538-7305.1948.tb01338.x.
- [4] F. Chollet, M. Knoop, G. Kamradt, B. Landers, and H. Pinkard, "ARC-AGI-2: A New Challenge for Frontier AI Reasoning Systems," *arXiv preprint arXiv:2505.11831*, 2025, [Online]. Available: <https://arxiv.org/abs/2505.11831>
- [5] V. Balasubramanian, "Brain power," *Proceedings of the National Academy of Sciences*, vol. 118, no. 32, p. e2107022118, 2021, doi: 10.1073/pnas.2107022118.
- [6] C. Morris and H. Lecar, "Voltage Oscillations in the Barnacle Giant Muscle Fiber," *Biophysical Journal*, vol. 35, no. 1, pp. 193–213, July 1981, doi: 10.1016/S0006-3495(81)84782-0.

- [7] L. Lapicque, "Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation," *Journal de Physiologie et de Pathologie Générale*, vol. 9, pp. 620–635, 1907.
- [8] B. Lindner and L. Schimansky-Geier, "Analytical approach to the stochastic FitzHugh-Nagumo system and coherence resonance," *Physical Review E*, vol. 60, no. 6, pp. 7270–7276, 1999, doi: 10.1103/PhysRevE.60.7270.
- [9] P. U. Diehl and M. Cook, "Unsupervised learning of digit recognition using spike-timing-dependent plasticity," *Frontiers in Computational Neuroscience*, vol. 9, p. 99, 2015, doi: 10.3389/fncom.2015.00099.
- [10] Y. Li *et al.*, "Activity-dependent synaptic plasticity of a chalcogenide electronic synapse for neuromorphic systems," *Scientific Reports*, vol. 4, p. 4906, May 2014, doi: 10.1038/srep04906.
- [11] Y. Zhang, X. Wang, Y. Li, and E. G. Friedman, "Memristive Model for Synaptic Circuits," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 64, no. 7, pp. 1–5, 2017, doi: 10.1109/TCSII.2016.2605069.
- [12] H. Kim *et al.*, "Spiking Neural Network (SNN) With Memristor Synapses Having Non-linear Weight Update," *Frontiers in Computational Neuroscience*, vol. 15, p. 646125, 2021, doi: 10.3389/fncom.2021.646125.
- [13] D. Sigg, "Modeling ion channels: Past, present, and future," *Journal of General Physiology*, vol. 144, no. 1, pp. 7–26, 2014, doi: 10.1085/jgp.201311130.
- [14] L. S. Milesu, T. Yamanishi, K. Ptak, M. Z. Mogri, and J. C. Smith, "Real-time kinetic modeling of voltage-gated ion channels using dynamic clamp," *Biophysical Journal*, vol. 95, no. 1, pp. 66–87, 2008, doi: 10.1529/biophysj.107.118190.
- [15] W. Gerstner and W. M. Kistler, *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge: Cambridge University Press, 2002.
- [16] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge: Cambridge University Press, 2014. [Online]. Available: <https://neurondynamics.epfl.ch/>
- [17] Y. LeCun, C. Cortes, and C. Burges, "MNIST handwritten digit database," *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, vol. 2, 2010.
- [18] S. Song, K. D. Miller, and L. F. Abbott, "Competitive Hebbian Learning through Spike-Timing-Dependent Synaptic Plasticity," *Nature Neuroscience*, vol. 3, no. 9, pp. 919–926, 2000, doi: 10.1038/78829.
- [19] T. Masquelier and S. J. Thorpe, "Unsupervised Learning of Visual Features through Spike Timing Dependent Plasticity," *PLoS Computational Biology*, vol. 3, no. 2, p. e31, 2007, doi: 10.1371/journal.pcbi.0030031.
- [20] T. Masquelier, R. Guyonneau, and S. J. Thorpe, "Spike Timing Dependent Plasticity Finds the Start of Repeating Patterns in Continuous Spike Trains," *PLoS ONE*, vol. 3, no. 1, p. e1377, 2008, doi: 10.1371/journal.pone.0001377.
- [21] P. A. Merolla *et al.*, "A million spiking-neuron integrated circuit with a scalable communication network and interface," *Science*, vol. 345, no. 6197, pp. 668–673, 2014, doi: 10.1126/science.1254642.
- [22] M. Davies *et al.*, "Loihi: A Neuromorphic Manycore Processor with On-Chip Learning," *IEEE Micro*, vol. 38, no. 1, pp. 82–99, 2018, doi: 10.1109/MM.2018.112130359.
- [23] S. Hussain, S.-C. Liu, and A. Basu, "Improved Margin Multi-Class Classification Using Dendritic Neurons with Morphological Learning," in *2014 IEEE International Symposium on Circuits and Systems (ISCAS)*, Melbourne, VIC, Australia: IEEE, 2014, pp. 2640–2643. doi: 10.1109/ISCAS.2014.6865683.