# Multiscale Astrocyte Network Calcium Dynamics for Biologically-Informed Intelligence in Anomaly Detection

Berk Iskar, and Michael T. Barros, *Senior Member, IEEE*

***Abstract*—Network anomaly detectors trained offline can lose performance as traffic patterns and attack distributions change over time. We introduce an astrocyte-informed $Ca^{2+}$-modulated learning framework that couples a multicellular $Ca^{2+}$ simulator to a deep neural network (DNN). The simulator captures $IP_3$-mediated $Ca^{2+}$ release, SERCA uptake, and conductance-aware diffusion through gap junctions. Its output contributes to a metaplastic learning rule in which local, heterosynaptic, postsynaptic, supervisory, and multicellular $Ca^{2+}$ signals regulate gradient updates through an adaptive threshold. We evaluate the framework on CTU-13 Neris traffic using conventional and leakage-controlled chronological settings, strong adaptive DNN baselines, and five-seed ablations. Under temporal distribution shift, the full $Ca^{2+}$ model gives the strongest mean performance among the evaluated adaptive baselines and reduces missed detections relative to the matched DNN. Factorial ablation identifies $Ca^{2+}$ gating as the main contributor to this gain: Laplacian regularisation alone provides little improvement, whereas removing the multicellular $Ca^{2+}$ term substantially reduces performance. Randomised, shuffled, and constant-signal controls further show that the benefit is linked to $Ca^{2+}$-dependent metaplastic regulation, while the precise temporal waveform is not uniquely required. These results support astrocyte-informed $Ca^{2+}$ signalling as a mechanism for regulating neural-network adaptation under changing data distributions.**

***Index Terms*—Astrocyte networks, $Ca^{2+}$ signaling, bio-informed machine learning, anomaly detection, deep neural networks.**

## I. INTRODUCTION

Modern cyber threats present an escalating challenge for network security as intrusion and malware campaigns now evolve rapidly [1]–[3]. Static machine-learning (ML) detectors, typically trained offline on historical data, struggle to adapt to these [4]–[6]. They suffer from concept drift as attack signatures evolve, generate elevated false alarms when encountering novel patterns, and exhibit poor generalization across different network environments or attack scenarios. While periodic retraining offers a potential solution, it remains computationally expensive and prohibitively slow for real-time threat response. Even sophisticated regularization techniques fail to provide the context-dependent, adaptive plasticity that operational security environments demand [2], [3].

In biological neural circuits, learning emerges from activity-dependent plasticity mechanisms that extend beyond neurons to include astrocytes, whose intracellular $Ca^{2+}$ dynamics coordinate synaptic efficacy across both spatial and temporal domains. Astrocytes form tripartite synapses with pre- and post-synaptic neurons, integrate diverse neuromodulatory signals, and communicate through gap junction networks [7]. Contemporary models of $Ca^{2+}$-controlled plasticity demonstrate that synaptic weight changes can be dynamically gated by local $Ca^{2+}$ concentrations relative to slowly adapting cellular thresholds [8]. These mechanisms are biologically plausible at the level of the modeled biochemical sub-cellular $Ca^{2+}$ processes. The learning framework constructed from them is biologically informed rather than a literal biological implementation. They may provide methods for adaptive machine learning systems that maintain data efficiency and robustness [9]–[11].

Astrocyte $Ca^{2+}$ signaling is being incorporated into artificial neural and neuromorphic models. Recent studies report astrocyte-modulated spiking networks [12], astrocyte-enabled transformer components [13], and astrocyte-augmented large language models [14]. Across these settings, astrocyte-like units have been used to gate plasticity, maintain context over longer horizons, and inject biologically motivated constraints, resulting in improvements in temporal predictability, context-dependent adaptation, and (to a degree) interpretability [15]–[17]. This suggests that glial-informed controllers can complement neuronal dynamics, motivating learning rules that couple task signals to biophysically grounded $Ca^{2+}$ modulation, an approach we pursue here.

Neuron-astrocyte interactions in the brain facilitate learning across multiple time-scales while conveying task-relevant contextual information to circuit dynamics [18]. Within the molecular communications (MC) paradigm, astrocytic $Ca^{2+}$ signaling has been modeled explicitly as a communication substrate across scales. [19] developed an end-to-end MC analysis for $Ca^{2+}$-signaling in biological tissues. Building on this, [20] showed how control strategies shape reliability and responsiveness of intercellular signaling. Disease-linked remodeling has also been explored, [21] examined how Alzheimer's pathology alters astrocyte network topology and, consequently, MC capacity across scales. Lenk *et al.* demonstrates how connection radius influences hub-astrocyte emergence in 3D [22]. Complementing these, Hyttinen *et al.* survey and demonstrate astrocytes' modulatory roles over subcellular, cellular, and intercellular communication channels [23]. These studies provide quantifiable $Ca^{2+}$ dynamics that inform the design of adaptive, bio-grounded learning signals used in this work.

We propose *$Ca^{2+}$-modulated learning for anomaly detection*. In our framework, we first simulate realistic $Ca^{2+}$ dynamics across a network of connected astrocyte cells based on previous work [16], [19]. We then map these $Ca^{2+}$ signals to individual synapses in a DNN and smooth them over time to reduce noise. The resulting $Ca^{2+}$-based signals act as gates that control when and how strongly synaptic weights

B. Iskar and M. T. Barros are with the School of Computer Science and Electronic Engineering, University of Essex, UK.

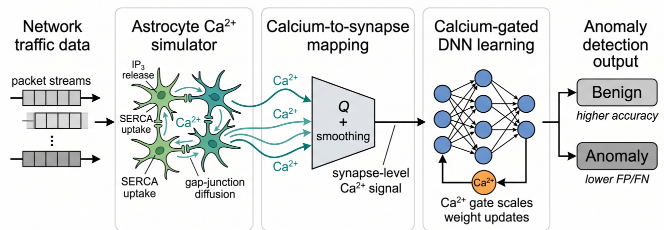


Fig. 1: Simplified overview of the proposed framework. Astrocyte $Ca^{2+}$ dynamics are simulated, mapped to synaptic sites, and used to define a metaplastic gate that modulates DNN training. The trained model then produces binary anomaly predictions, evaluated on the CTU-13 dataset.

are updated during learning. Specifically, weight updates are scaled based on whether local $Ca^{2+}$ levels exceed an adaptive threshold, creating a biologically-informed form of *metaplasticity*. Additionally, we apply a regularization technique that encourages neighbouring synapses to adapt in coordinated ways, promoting stable learning even when data patterns shift over time. The contributions of our paper are as follows: The mesoscopic reaction–diffusion model (Eqs. (5)–(10)), including $IP_3$-mediated CICR, SERCA pump uptake, and ER/$IP_3$ pool kinetics, follows the astrocyte $Ca^{2+}$ signaling framework of Barros *et al.* [16], [19]. The gap-junction Markov model (Eqs. (13)–(14)) and conductance-weighted Laplacian (Eq. (18)) extend the same framework with $Ca^{2+}$-dependent hemichannel gating informed by connexin physiology [24], [25]. The Hill-type kinetics and parameter ranges are standard in computational astrocyte modeling [7].

The following elements are *new* to this study:

1) *$Ca^{2+}$-gated DNN learning rule.* We formulate a supervised plasticity rule that combines local input, heterosynaptic drive, somatic/BAP activity, supervision, and multicellular $Ca^{2+}$ into a total $Ca^{2+}$ drive (Eq. (22)), gated through an adaptive metaplastic threshold (Eqs. (23)–(24)). This learning-rule integration is original.
2) *Mass-preserving synapse–cell map and exponential smoothing.* The linear map $Q$ (Eq. (4)) with column-stochastic mass preservation and exponential smoothing (Eq. (19)) are new interface components coupling the biophysical simulator to gradient-based training.
3) *Heterosynaptic Laplacian regularization.* The synapse-graph Laplacian coupling in Eq. (27) and $Ca^{2+}$ smoothness penalty are new regularization terms designed for this framework.
4) *Empirical validation on CTU-13 (Neris).* We add strong classical and adaptive baselines, a five-seed chronological evaluation across three Neris scenarios, leakage-controlled feature modes, factorial $Ca^{2+}$/Laplacian controls, biophysical mechanism ablations, and runtime profiling.

## II. Model and Methodology

The main model framework lies in coupling a mesoscopic astrocyte $Ca^{2+}$ field to a feed-forward DNN for binary anomaly detection on network flows. First, a 3D astrocyte lattice (from [16], [19]) is simulated using a mesoscopic reaction–diffusion model with stochastic updates: $IP_3$-mediated $Ca^{2+}$ release, SERCA uptake, and conductance-aware diffusion through gap junctions whose open states evolve via a $Ca^{2+}$-dependent Markov process. The resulting per-cell $Ca^{2+}$ trajectories are linearly mapped to synapse-specific signals and temporally smoothed to yield a multicellular modulation term. These signals are fused with local presynaptic activity, heterosynaptic drive, somatic/back-propagating action potential (BAP) activity, and a supervisory term to construct a total $Ca^{2+}$ drive that gates weight updates through a fast metaplastic modulator around a slowly adapting $Ca^{2+}$ threshold. Learning proceeds by standard backpropagation whose gradients are multiplicatively scaled by the $Ca^{2+}$ gate and complemented with a Laplacian heterosynaptic coupling regularizer. A $Ca^{2+}$ smoothness penalty promotes spatial coherence of the $Ca^{2+}$ field. Our final analysis evaluates the resulting learning rule directly on CTU-13 Neris under strong baseline calibration, chronological distribution shift, leakage-controlled feature sets, and matched component ablations.

### A. Framework Overview

**Task and data:** Let $\mathcal{D} = \{(x_s, y_s)\}_{s=1}^{T}$ denote $T$ network-flow samples, where $x_s \in \mathbb{R}^d$ is a $d$-dimensional feature vector and $y_s \in \{0,1\}$ is a binary label. We reserve $\tau \geq 0$ for simulation time in the astrocyte model (to avoid overloading the index $s$).

**Classifier:** Our DNN is chosen for its ability to extract statistical relationships from data for binary classification. We use a feedforward topology, which is a reasonable choice when considering bio-informed architectures that balance expressivity and tractability. We define $f_\theta : \mathbb{R}^d \to [0,1]$ with parameters $\theta = \{W^{(\ell)}, b^{(\ell)}\}_{\ell=1}^{L}$ producing

$$\hat{y}_s = f_\theta(x_s) = \sigma(z_s^{(L)}), \qquad z_s^{(L)} = W^{(L)} h_s^{(L-1)} + b^{(L)}, \tag{1}$$

*where* $L$ is the number of layers; $W^{(\ell)} \in \mathbb{R}^{n_\ell \times n_{\ell-1}}$ and $b^{(\ell)} \in \mathbb{R}^{n_\ell}$ are the weights and biases; $n_0 = d$ and $n_L = 1$; $h_s^{(0)} = x_s$ and $h_s^{(\ell)} = \phi^{(\ell)}(W^{(\ell)} h_s^{(\ell-1)} + b^{(\ell)})$ for $1 \leq \ell \leq L-1$; $\sigma(u) = 1/(1+e^{-u})$.

**Loss function:** We chose a loss function that is able to be sensitive to weight adaptation effectively across the network, when scaled, to ensure comparability. We thus used a Bernoulli negative log-likelihood [26], yielding statistically consistent estimators and stable gradients, defined by:

$$\mathcal{L}_{\text{base}}(\theta) = -\frac{1}{T} \sum_{s=1}^{T} \left( y_s \log \hat{y}_s + (1-y_s) \log(1-\hat{y}_s) \right). \tag{2}$$

*where* $T$ is the number of samples, $y_s \in \{0,1\}$ is the ground-truth label for sample $s$, and $\hat{y}_s \in (0,1)$ is the predicted probability.

**$Ca^{2+}$ field:** As our multicellular astrocyte model is not integrated with the DNN in an operational fashion, we define a $Ca^{2+}$ field to be able to modularly match synaptic weight operations as

$$c(\tau) = \begin{bmatrix} c_1(\tau) & \cdots & c_{N_c}(\tau) \end{bmatrix}^\top \in \mathbb{R}_{\geq 0}^{N_c}. \tag{3}$$

*where* $c_i(\tau)$ is the intracellular $Ca^{2+}$ concentration (units: $\mu$M) in cell $i$ at simulation time $\tau \geq 0$, and nonnegativity reflects physical constraints. More details in the next section.

**Astrocyte multicellular lattice model:** We model the multicellular astrocyte network as an undirected graph $G = (\mathcal{V}, \mathcal{E})$ with $|\mathcal{V}| = N_c$ cells embedded on a 3D lattice (voxel spacing $h > 0$). The adjacency is $A \in \{0,1\}^{N_c \times N_c}$, the degree matrix $D = \mathrm{diag}(A\mathbf{1})$, and the graph Laplacian $L = D - A$ encodes diffusive coupling along gap junctions. The Laplacian enforces two key principles: locally, flux flows proportionally to differences between adjacent cells, and globally, the total amount is conserved except for boundary conditions and leak terms.

**Synapse–cell mapping:** From the $Ca^{2+}$ field we define a linear, nonnegative, mass-preserving map to obtain the necessary position on the DNN where $Ca^{2+}$ will affect network weights. We define this as:

$$\bar{C}(\tau) = Q\, c(\tau) \in \mathbb{R}^{N_s}_{\geq 0}, \qquad \mathbf{1}^\top Q = \mathbf{1}^\top. \tag{4}$$

*where* $Q \in \mathbb{R}^{N_s \times N_c}_{\geq 0}$ maps cell $Ca^{2+}$ to $N_s$ synaptic sites; $\bar{C}_i(\tau)$ is the synapse-$i$ $Ca^{2+}$ drive; $\mathbf{1}$ is the all-ones vector, and the column-stochastic constraint $\mathbf{1}^\top Q = \mathbf{1}^\top$ ensures mass preservation ($\mathbf{1}^\top \bar{C} = \mathbf{1}^\top c$). This is the minimal assumption consistent with superposition of weak astrocytic influences and compatibility with gradient-based learning.

### B. Multicellular $Ca^{2+}$ Field (Mesoscopic Model)

We model intracellular $Ca^{2+}$ as a stochastic reaction–diffusion process on the astrocyte graph $G$ with discrete simulation step $\Delta t > 0$:

$$\begin{aligned} c(\tau+\Delta t) = c(\tau) + \Delta t\Big(&J_{\mathrm{in}}(\tau) - J_{\mathrm{out}}(\tau) \\ &- \underbrace{\kappa_D\, L\, c(\tau)}_{\text{gap-junction diffusion}}\Big) + \sqrt{\Delta t}\,\zeta(\tau). \end{aligned} \tag{5}$$

*where* $c(\tau) \in \mathbb{R}^{N_c}_{\geq 0}$ stacks cytosolic $Ca^{2+}$ concentrations ($\mu$M) for $N_c$ cells at time $\tau \geq 0$; $J_{\mathrm{in}}(\tau), J_{\mathrm{out}}(\tau) \in \mathbb{R}^{N_c}_{\geq 0}$ are influx/efflux flux vectors ($\mu$M/ms); $L = D - A$ is the graph Laplacian induced by gap-junction connectivity; $\kappa_D > 0$ is an effective conductance/diffusion coefficient ($\mathrm{ms}^{-1}$); and $\zeta(\tau) \sim \mathcal{N}(0, \Sigma_\zeta)$ is zero-mean Gaussian noise with covariance $\Sigma_\zeta \succeq 0$, capturing mesoscopic stochasticity from channel opening and finite-molecule effects.

The influx/efflux terms aggregate channel and pump currents via Hill-type nonlinearities:

$$\begin{aligned} \big[J_{\mathrm{in}}(\tau)\big]_i = V_{\mathrm{IP3}} \quad [t]\, &\underbrace{\frac{c_i(\tau)^n}{K_1^n + c_i(\tau)^n} \cdot \frac{\mathrm{IP3}_i(\tau)^m}{K_I^m + \mathrm{IP3}_i(\tau)^m}}_{\mathrm{IP}_3\text{-mediated CICR}} \\ &\cdot \big(E_i(\tau) - c_i(\tau)\big), \end{aligned} \tag{6}$$

$$\big[J_{\mathrm{out}}(\tau)\big]_i = V_{\mathrm{SERCA}}\, \frac{c_i(\tau)^p}{K_2^p + c_i(\tau)^p} \;+\; \kappa_o\, c_i(\tau). \tag{7}$$

*where* $i \in \{1, \dots, N_c\}$ indexes cells; $V_{\mathrm{IP3}} > 0$ is the maximal $\mathrm{IP}_3$-receptor-mediated release rate; $K_1, K_I > 0$ are half-saturation constants for cytosolic $Ca^{2+}$ and $\mathrm{IP}_3$, respectively; $n, m \geq 1$ are Hill exponents capturing cooperativity; $E_i(\tau) \geq 0$ is ER $Ca^{2+}$ (units: $\mu$M) providing the store gradient $\big(E_i - c_i\big)$; $V_{\mathrm{SERCA}} > 0$ is the SERCA pump capacity with half-saturation $K_2 > 0$ and exponent $p \geq 1$; $\kappa_o > 0$ lumps plasma-membrane extrusion (PMCA/NCX).

Auxiliary pools follow first-order kinetics:

$$E_i(\tau+\Delta t) = E_i(\tau) + \Delta t\Big(J_{\mathrm{out},i}(\tau) - J_{\mathrm{in},i}(\tau) \tag{8}$$

$$- \kappa_f\big(E_i(\tau) - c_i(\tau)\big)\Big), \tag{9}$$

$$\mathrm{IP3}_i(\tau+\Delta t) = \mathrm{IP3}_i(\tau) + \Delta t\Big(V_{\mathrm{PLC}}\, \frac{c_i(\tau)^2}{K_p^2 + c_i(\tau)^2} \tag{10}$$

$$- \kappa_d\, \mathrm{IP3}_i(\tau)\Big). \tag{11}$$

*where* $E_i(\tau)$ is ER $Ca^{2+}$; $\kappa_f > 0$ is an ER leak/forward-exchange rate; $\mathrm{IP3}_i(\tau) \geq 0$ is inositol trisphosphate concentration (arbitrary units or $\mu$M if calibrated); $V_{\mathrm{PLC}} > 0$ sets phospholipase C production strength with half-saturation $K_p > 0$; and $\kappa_d > 0$ is the $\mathrm{IP}_3$ degradation rate.

**Corollary:** A useful conservation identity follows from $\mathbf{1}^\top L = 0$:

$$\begin{aligned} \mathbf{1}^\top c(\tau+\Delta t) = \mathbf{1}^\top c(\tau) + \Delta t\, \mathbf{1}^\top\big(J_{\mathrm{in}}(\tau) - J_{\mathrm{out}}(\tau)\big) \\ + \sqrt{\Delta t}\, \mathbf{1}^\top \zeta(\tau), \end{aligned} \tag{12}$$

*which shows* diffusion is mass-conservative at the network level; only channel/pump fluxes and noise change total cytosolic $Ca^{2+}$.

**Numerical and physical constraints.** (i) *Nonnegativity:* after each step we project to the nonnegative orthant, $c_i \leftarrow \max\{0, c_i\}$ and similarly for $E_i$, $\mathrm{IP3}_i$, preserving physical meaning. (ii) *Stability:* choose $\Delta t$ to satisfy a diffusive CFL-type bound, e.g., $\Delta t \lesssim \big(\kappa_o + V_{\mathrm{SERCA}} + \kappa_D\, \lambda_{\max}(L)\big)^{-1}$, to keep the explicit scheme stable (empirically tuned in our simulator). (iii) *Boundary conditions:* encoded through $A$ (and thus $L$); no-flux boundaries correspond to removing exterior edges, while anisotropy is captured by weighted edges if needed. (iv) *Units/scale:* parameters $(V_{\mathrm{IP3}}, V_{\mathrm{SERCA}}, \kappa_o, \kappa_f, V_{\mathrm{PLC}}, \kappa_d, \kappa_D)$ are in $\mathrm{ms}^{-1}$ (or per-step units) when concentrations are in $\mu$M.

### C. Gap-Junction State Dynamics and Conductance

We model each gap junction between neighboring cells $(i, j) \in \mathcal{E}$ as a three-state process with hemichannel configurations $s_{ij}(\tau) \in \{\mathrm{HH}, \mathrm{HL}, \mathrm{LH}\}$ (open–open, open–closed, closed–open). The state evolves as a $Ca^{2+}$-dependent discrete-time Markov chain:

$$\mathbb{P}[s_{ij}(\tau+\Delta t) = \sigma \mid s_{ij}(\tau) = \sigma'] = \Pi_{ij,\sigma' \to \sigma}(\tau). \tag{13}$$

*where* $\Delta t > 0$ is the simulation step; $\Pi_{ij}(\tau) \in \mathbb{R}^{3\times 3}$ is a row-stochastic transition matrix with entries $\Pi_{ij,\sigma'\to\sigma}(\tau) \in [0,1]$; and $\sigma', \sigma \in \{\mathrm{HH}, \mathrm{HL}, \mathrm{LH}\}$.

We summarize $Ca^{2+}$ dependence through an *open propensity*, meaning $|c_i - c_j|$ captures transjunctional drive, while $(c_i + c_j)$ captures local facilitation/inactivation. We define this by:

$$p_{ij}^{\mathrm{open}}(\tau) = \sigma\Big(a_0 + a_1\, |c_i(\tau) - c_j(\tau)| + a_2\, (c_i(\tau) + c_j(\tau))\Big), \tag{14}$$

*where* $\sigma(u) = \frac{1}{1+e^{-u}}$ and $p_{ij}^{\mathrm{open}}(\tau) \in (0,1)$ is the probability a hemichannel is open at time $\tau$; $c_i(\tau)$ and $c_j(\tau)$ are cytosolic $Ca^{2+}$ concentrations (units: $\mu$M) from (3); $a_0, a_1, a_2 \in \mathbb{R}$ are logistic sensitivities (dimensionless); and $\sigma(\cdot)$ is the logistic function.

Assuming independent, symmetric hemichannels with the same open propensity, the instantaneous state probabilities factor as

$$\pi_{ij}^{\mathrm{HH}}(\tau) = \big(p_{ij}^{\mathrm{open}}(\tau)\big)^2, \qquad \pi_{ij}^{\mathrm{LH}}(\tau) = \big(1 - p_{ij}^{\mathrm{open}}(\tau)\big)p_{ij}^{\mathrm{open}}(\tau). \tag{15}$$

*where* $\pi_{ij}^{\sigma}(\tau) \in [0,1]$ and $\sum_\sigma \pi_{ij}^{\sigma}(\tau) = 1$. *Rationale:* independence yields the product form; symmetry gives $\pi_{ij}^{\mathrm{HL}} = \pi_{ij}^{\mathrm{LH}}$.

We assign conductance to each configuration and define the instantaneous junction conductance:

$$g_{ij}(\tau) = g_{\max}\Big(\mathbf{1}\{s_{ij}(\tau) = \mathrm{HH}\} + \rho\,\mathbf{1}\{s_{ij}(\tau) \in \{\mathrm{HL}, \mathrm{LH}\}\}\Big). \tag{16}$$

*where* $g_{ij}(\tau) \geq 0$ is the conductance of edge $(i,j)$ at time $\tau$; $g_{\max} > 0$ is the fully open (HH) conductance; and $\rho \in [0,1)$ is the relative conductance when one hemichannel is closed (HL/LH). This is because we known that permeability is reduced but nonzero for half-open junctions and $\rho = \frac{1}{2}$ recovers the heuristic used in our simulator.

For diffusion between astrocytes $i$ and $j$, we use the conditional expectation given $Ca^{2+}$:

$$\begin{aligned} \bar{g}_{ij}(\tau) &= \mathbb{E}[g_{ij}(\tau) \mid c(\tau)] \\ &= g_{\max}\big(p^2 + 2\rho\, p(1-p)\big), \end{aligned} \tag{17}$$

*where* $p = p_{ij}^{\mathrm{open}}(\tau)$ for brevity.

To model real astrocyte networks accurately, we must account for heterogeneous gap junction conductances. Standard Laplacians assume uniform coupling, which is biologically unrealistic. Weighting the Laplacian by conductance captures the actual strength of intercellular connections [24], [25]. We finally build a conductance-weighted Laplacian from $\bar{g}_{ij}(\tau)$:

$$G_{ij}(\tau) = \begin{cases} \bar{g}_{ij}(\tau), & (i,j) \in \mathcal{E},\ i \neq j, \\ 0, & i = j \text{ or } (i,j) \notin \mathcal{E}, \end{cases}$$

$$D_g(\tau) = \mathrm{diag}\Big(\sum_j G_{ij}(\tau)\Big), \tag{18}$$

$$\tilde{L}(\tau) = D_g(\tau) - G(\tau).$$

*where* $G(\tau) \in \mathbb{R}^{N_c \times N_c}$ stores edge conductances, $D_g(\tau) \in \mathbb{R}^{N_c \times N_c}$ is the conductance degree matrix, and $\tilde{L}(\tau) \succeq 0$ is the conductance-weighted graph Laplacian.

### *D. From $Ca^{2+}$ Field to Learning Signals*

We aggregate cell $Ca^{2+}$ to synapses via (4) and apply exponential smoothing (low-pass filtering) to reduce sampling noise:

$$\tilde{C}(\tau) = (1-\phi)\,\tilde{C}(\tau - \Delta t) + \phi\,\bar{C}(\tau), \qquad \phi \in (0,1]. \tag{19}$$

*where* $\bar{C}(\tau) \in \mathbb{R}_{\geq 0}^{N_s}$ is the synapse-level $Ca^{2+}$ drive from (4) (units: $\mu$M), $\tilde{C}(\tau) \in \mathbb{R}_{\geq 0}^{N_s}$ is its smoothed version (units: $\mu$M), $\Delta t > 0$ is the simulation step, and $\phi$ is the smoothing gain. For interpretability we parameterize

$$\phi \;=\; 1 - \exp\!\Big(-\frac{\Delta t}{\tau_s}\Big), \qquad \tau_s > 0, \tag{20}$$

*where* $\tau_s$ is the smoothing time constant (ms): larger $\tau_s$ produces heavier smoothing.

To stabilize scales across synapses we perform per-synapse z-normalization using exponentially weighted moments [26]:

$$\begin{aligned} \hat{C}_i(\tau) &= \frac{\tilde{C}_i(\tau) - \mu_i(\tau)}{\sigma_i(\tau) + \epsilon_{\mathrm{num}}}, \\ \mu_i(\tau + \Delta t) &= (1-\rho)\mu_i(\tau) + \rho\,\tilde{C}_i(\tau), \\ \sigma_i^2(\tau + \Delta t) &= (1-\rho)\sigma_i^2(\tau) + \rho\big(\tilde{C}_i(\tau) - \mu_i(\tau)\big)^2. \end{aligned} \tag{21}$$

*where* $i \in \{1, \ldots, N_s\}$ indexes synapses; $\hat{C}_i(\tau) \in \mathbb{R}$ is the standardized, *dimensionless* $Ca^{2+}$ signal used for learning; $\mu_i(\tau)$ and $\sigma_i(\tau) \geq 0$ are exponentially weighted mean and standard deviation estimates (units: $\mu$M); $\epsilon_{\mathrm{num}} > 0$ is a small numerical constant preventing division by zero (e.g., $\epsilon_{\mathrm{num}} = 10^{-6}\,\mu$M); it is distinct from the biological coefficient $\varepsilon$ multiplying the multicellular term in Eq. (22); and $\rho \in (0,1]$ is the moment-update gain. This method of z-normalization removes baseline and scale heterogeneity across synapses, making $Ca^{2+}$ contributions comparable. As with (20), we tie $\rho$ to a time constant similarly to $\phi$ in Eq. (20).

### *E. $Ca^{2+}$-Weighted Plasticity With Adaptive Thresholds*

We combine five signals to decide when a synapse should strengthen or weaken [15], [27], [28]: (i) the input arriving at that synapse ($x_i$; *local evidence*), (ii) the activity of nearby inputs summarized through the current weights ($[Wx]_i$; *neighborhood context*), (iii) the neuron's own output ($\hat{y}$; *did the cell fire?*), (iv) a label-derived cue ($Z_i$; *teacher signal*), and (v) a field from surrounding astrocytes ($\hat{C}_i$; *multicellular context*). In short, plasticity is strongest when local input coincides with postsynaptic firing and supportive context, and it switches sign when the context contradicts the local input [29], [30]. We define the $Ca^{2+}$ control of synaptic plasticity hypothesis by:

$$\begin{aligned} C_i(t) = \underbrace{\alpha\, x_i(t)}_{\text{local}} + \underbrace{\beta\,[Wx(t)]_i}_{\text{heterosynaptic}} + \underbrace{\gamma\,\hat{y}(t)}_{\text{BAP/somatic}} \\ + \underbrace{\delta\, Z_i(t)}_{\text{supervisor}} + \underbrace{\varepsilon\,\hat{C}_i(t)}_{\text{multicellular astrocyte}}. \end{aligned} \tag{22}$$

*where* $t \in \mathbb{N}$ indexes *learning updates* (distinct from simulator time $\tau$); $i$ indexes a synapse onto a given postsynaptic unit in the network; $x_i(t) \in \mathbb{R}$ is the presynaptic activation feeding that synapse at step $t$ (e.g., $h_i^{(\ell-1)}$); $W$ is the weight row-vector into the postsynaptic unit, $x(t)$ is the full presynaptic activity vector, and $[Wx(t)]_i$ denotes the heterosynaptic drive attributed to synapse $i$ (e.g., a local share of the postsynaptic current $\sum_j w_j x_j$); $\hat{y}(t) \in [0,1]$ is the postsynaptic output (e.g., neuron or network output probability at $t$); $Z_i(t) \in \{-1, +1\}$ is a supervision signal (e.g., $Z_i(t) = 2y - 1$ on output synapses and 0 otherwise); $\hat{C}_i(t) \in \mathbb{R}$ is the *dimensionless*, z-normalized multicellular $Ca^{2+}$ signal from (21). The nonnegative coefficients $\alpha, \beta, \gamma, \delta, \varepsilon \geq 0$ weight the relative contributions and place all terms on a comparable scale.

A slowly adapting $Ca^{2+}$ threshold implements homeostatic metaplasticity:

$$\theta_i(t+1) = (1-\eta_\theta)\,\theta_i(t) + \eta_\theta\, C_i(t). \tag{23}$$

*where* $\theta_i(t) \in \mathbb{R}$ is the running $Ca^{2+}$ threshold at synapse $i$, and $\eta_\theta \in (0,1)$ is the adaptation gain. (23) makes potentiation/depression history-dependent, preventing runaway growth. For interpretability we may set $\eta_\theta = 1 -$

$\exp(-\Delta t_{\text{learn}}/\tau_\theta)$ with time constant $\tau_\theta > 0$ and learning-step duration $\Delta t_{\text{learn}} > 0$.

The potentiation–depression modulator gates plasticity around the adaptive threshold:

$$m_i(t) = 2\,\sigma\big(k\,(C_i(t) - \theta_i(t))\big) - 1, \qquad \sigma(u) = \frac{1}{1+e^{-u}}. \tag{24}$$

*where* $m_i(t) \in (-1, 1)$ is a signed gain that is positive for potentiation and negative for depression; $k > 0$ is the steepness parameter of the logistic $\sigma(\cdot)$ controlling sensitivity to the deviation $C_i - \theta_i$. *Rationale:* the symmetric range $(-1, 1)$ yields balanced LTP/LTD, while $k$ tunes how sharply the rule switches with $Ca^{2+}$.

### F. $Ca^{2+}$-Gated Weight Dynamics on a DNN

Let $h_j^{(\ell-1)}(t)$ be the presynaptic activation feeding weight $w_{ij}^{(\ell)}(t)$, and let $z_i^{(\ell)}(t) = \sum_j w_{ij}^{(\ell)}(t)\,h_j^{(\ell-1)}(t) + b_i^{(\ell)}(t)$ be the pre-activation. For the Bernoulli negative log-likelihood in Eq. (2) combined with the sigmoid output in Eq. (1), the derivative with respect to the output logit simplifies to

$$\delta^{(L)}(t) = \frac{\partial \mathcal{L}_{\text{base}}}{\partial z^{(L)}} = \hat{y}_t - y_t. \tag{25}$$

Here $\hat{y}_t \in (0, 1)$ is the predicted anomaly probability and $y_t \in \{0, 1\}$ is the label. This form avoids multiplying by the sigmoid derivative twice when binary cross-entropy is differentiated through the sigmoid output.

For a hidden layer $(1 \le \ell < L)$ with activation $\phi^{(\ell)}$, the standard backprop recursion is

$$\delta_i^{(\ell)}(t) = \phi^{(\ell)\prime}\big(z_i^{(\ell)}(t)\big) \sum_{k=1}^{n_{\ell+1}} w_{ki}^{(\ell+1)}(t)\,\delta_k^{(\ell+1)}(t), \tag{26}$$

*where* $\phi^{(\ell)\prime}(\cdot)$ is the derivative of the hidden activation, $n_{\ell+1}$ is the width of layer $\ell{+}1$, and $w_{ki}^{(\ell+1)}(t)$ is the weight from unit $i$ in layer $\ell$ to unit $k$ in layer $\ell{+}1$.

We propose a $Ca^{2+}$-gated, heterosynaptically regularized update:

$$\begin{aligned} \Delta w_{ij}^{(\ell)}(t) = &-\eta\left(1 + \lambda_m\, m_i^{(\ell)}(t)\right) \delta_i^{(\ell)}(t)\, h_j^{(\ell-1)}(t) \\ &- \lambda_w\, w_{ij}^{(\ell)}(t) \\ &- \xi \sum_{u=1}^{n_\ell} L_{iu}^{(\ell)}\, w_{uj}^{(\ell)}(t) \\ &+ \mu\, \Delta w_{ij}^{(\ell)}(t-1). \end{aligned} \tag{27}$$

*where* $\eta > 0$ is the base learning rate; $m_i^{(\ell)}(t) \in (-1, 1)$ is the $Ca^{2+}$ modulator from (24) computed for synapse-index $i$ on layer $\ell$ (using $C_i(t)$ and $\theta_i(t)$ from (22)–(23)); $\lambda_m \ge 0$ scales the strength of $Ca^{2+}$ gating (typically $\lambda_m \in [0, 1]$ to keep $1{+}\lambda_m m_i^{(\ell)}(t) > 0$); $\delta_i^{(\ell)}(t)$ is the backprop error from (25)–(26); $h_j^{(\ell-1)}(t)$ is the presynaptic activation into weight $w_{ij}^{(\ell)}$; and $\lambda_w \ge 0$ is the $\ell_2$ weight-decay coefficient. The coefficient $\xi \ge 0$ weights the positive quadratic synapse-graph regularizer $\frac{\xi}{2}\sum_j (w_{\cdot j}^{(\ell)})^\top L^{(\ell)} w_{\cdot j}^{(\ell)}$. Because Eq. (27) is written as a weight increment under gradient descent, its Laplacian contribution is $-\xi L^{(\ell)} W^{(\ell)}$. Here $L^{(\ell)} \in \mathbb{R}^{n_\ell \times n_\ell}$ is symmetric and positive semidefinite with $\mathbf{1}^\top L^{(\ell)} = 0$, and $\mu \in [0, 1)$ is the momentum coefficient. The weights are updated as $w_{ij}^{(\ell)}(t{+}1) = w_{ij}^{(\ell)}(t) + \Delta w_{ij}^{(\ell)}(t)$.

### G. Evaluation Protocol

To address baseline calibration, component attribution, and temporal generalisation, we added a second evaluation protocol focusing on stricter evaluation. Three CTU-13 captures generated with the Neris botnet are used to construct a temporal sequence: Scenario 1 (CTU-Malware-Capture-Botnet-42), Scenario 2 (Botnet-43), and Scenario 9 (Botnet-50), corresponding to captures from 10, 11, and 17 August 2011, respectively. We treat these as sequential environments A, B, and C. Flow order is preserved within each environment. The primary leakage-controlled feature mode removes absolute start time, source and destination addresses, and source and destination ports (`StartTime`, `SrcAddr`, `DstAddr`, `Sport`, `Dport`). A traffic-only mode removes categorical protocol/direction/state variables as an additional stress test. Feature transformations are fitted without using future environments.

The principal comparison uses five random seeds (1, 7, 42, 123, and 2026) and identical data order, model initialisation, revealed-label indices, and update budget within each paired seed. We compare the proposed model against frozen, online, class-weighted online, adaptive-learning-rate, replay, and elastic-weight-consolidation (EWC) DNNs. Final runs use minibatches of 256 and 20 initial Scenario-A training epochs. The primary sequential condition uses a 10% revealed-label budget for online updates. Classification thresholds and all model-selection choices are fixed from pre-B/C training/-validation information and are not retuned on environments B or C. Means and uncertainty intervals in the reviewer-grade figures are calculated from the five seed-level outcomes; batches and individual flows are not treated as statistical replicates.

For baseline calibration on the original random all-feature capture, the 2,087,508 flows are partitioned 70/15/15 into train/validation/test sets, with an anomaly prevalence of approximately 8.86%. The validation partition is used for model/threshold selection where applicable, and the class-weighted DNN explicitly tests whether imbalance handling explains the result. The held-out test partition is used only for the reported baseline-calibration metrics.

Component attribution uses a matched factorial design: baseline DNN; $Ca^{2+}$ gate only; Laplacian only; $Ca^{2+}$ gate plus Laplacian; full model; $\varepsilon = 0$; random amplitude-matched, time-shuffled, phase-shifted, and constant-mean $Ca^{2+}$ controls; fixed-threshold and no-metaplasticity controls; gap-junction, CICR, and SERCA ablations; and a shuffled synapse–cell map $Q$. The logged reviewer runs use $\alpha = 1$, $\beta = 0.2$, $\gamma = 0.7$, $\delta = 1$, $\varepsilon = 0.75$, $\lambda_m = 0.5$, and Laplacian coefficient $10^{-4}$; baseline and single-component controls switch only the mechanism under test. A separate feature/split analysis compares random/all-feature, chronological/all-feature, chronological/no-address-no-port-no-time, and chronological/traffic-only conditions.

## III. Results

The Results focus on the experiments that directly test the proposed $Ca^{2+}$-modulated learning, including the five-seed chronological comparison, factorial ablations, and matched

TABLE I: Calcium simulator per-cell state vector (30 slots). Units are $\mu$M/ms where applicable.

| Idx | Code | Quantity | Units | Used in ML | Brief role in CalComSim |
|---|---|---|---|---|---|
| 1 | $E$ | ER $Ca^{2+}$ | $\mu$M | No | Store $Ca^{2+}$; source/sink in ER–cytosol exchange ((8)). |
| 2 | $I$ | $IP_3$ | $\mu$M (a.u.) | No | Second messenger driving CICR; evolves via ((10)). |
| 3 | $V_{\mathrm{IP3}}$ | $IP_3$R max flux | $\mathrm{ms}^{-1}$ | No | Scales $IP_3$-mediated release in ((6)). |
| 4 | **C** (4C) | Cytosolic $Ca^{2+}$ | $\mu$M | **Yes** | Primary signal; mapped to synapses ((4)), smoothed ((19)). |
| 5 | $\kappa_f$ | ER leak/forward exchange | $\mathrm{ms}^{-1}$ | No | Exchange term in ((8)). |
| 6 | $K_2$ (6K2) | SERCA half-saturation | $\mu$M | No | Nonlinearity of uptake in ((7)). |
| 7 | $K_1$ | $Ca^{2+}$ half-saturation ($IP_3$R) | $\mu$M | No | Nonlinearity of release in ((6)). |
| 8 | $k_{\mathrm{out}}$ (8kout) | PMCA/NCX extrusion | $\mathrm{ms}^{-1}$ | No | Linear efflux (lumped as $\kappa_o$) in ((7)). |
| 9 | $S$ (9S) | SERCA capacity factor | a.u. | No | Scales $V_{\mathrm{SERCA}}$ in ((7)). |
| 10 | $k_f$ (10kf) | Forward channel rate | $\mathrm{ms}^{-1}$ | No | Auxiliary gating for CICR. |
| 11 | $k_p$ (11kp) | Pump rate | $\mathrm{ms}^{-1}$ | No | Auxiliary pump kinetics. |
| 12 | $k_{\mathrm{deg}}$ (12kdeg) | $IP_3$ degradation | $\mathrm{ms}^{-1}$ | No | Decay term in ((10)). |
| 13 | $K_I$ | $IP_3$ half-saturation | $\mu$M | No | Nonlinearity of release in ((6)). |
| 14 | $n$ | Hill exponent ($Ca^{2+}$) | — | No | Cooperativity in ((6)). |
| 15 | $m$ | Hill exponent ($IP_3$) | — | No | Cooperativity in ((6)). |
| 16 | $p$ | Hill exponent (SERCA) | — | No | Cooperativity in ((7)). |
| 17 | $V_{\mathrm{SERCA}}$ | SERCA capacity | $\mu\mathrm{M}\,\mathrm{ms}^{-1}$ | No | Uptake term in ((7)). |
| 18 | $V_{\mathrm{PLC}}$ | PLC production rate | $\mu\mathrm{M}\,\mathrm{ms}^{-1}$ | No | Production term in ((10)). |
| 19 | $K_p$ | PLC half-saturation | $\mu$M | No | Nonlinearity in ((10)). |
| 20 | $\kappa_d$ | $IP_3$ degradation rate | $\mathrm{ms}^{-1}$ | No | Decay in ((10)). |
| 21 | $\kappa_D$ | Diffusion coefficient | $\mathrm{ms}^{-1}$ | No | Scales Laplacian in ((5)). |
| 22 | $a_0$ | Junction logistic bias | — | No | Baseline in open-probability ((14)). |
| 23 | $a_1$ | Sensitivity to $\lvert c_i - c_j \rvert$ | — | No | Slope for transjunctional drive in ((14)). |
| 24 | $K$ (24K) | Generic dissociation | $\mu$M | No | Channel/pump tuning constant (implementation slot). |
| 25 | $a_2$ | Sensitivity to $(c_i + c_j)$ | — | No | Slope for local facilitation/inactivation in ((14)). |
| 26 | $g_{\max}$ | Max junction conductance | a.u. | No | Conductance for HH state in ((16)). |
| 27 | $p_{\mathrm{HH}}$ (27phh) | Open–open probability | [0,1] | No | Junction state; contributes to $\bar{g}_{ij}$ ((17)). |
| 28 | $p_{\mathrm{HL}}$ (28phl) | Open–closed probability | [0,1] | No | Junction state; contributes to $\bar{g}_{ij}$ ((17)). |
| 29 | $p_{\mathrm{LH}}$ (29plh) | Closed–open probability | [0,1] | No | Junction state; contributes to $\bar{g}_{ij}$ ((17)). |
| 30 | $\rho$ | Half-open conductance fraction | — | No | Scales HL/LH conductance in ((16)). |

TABLE II: Single-term coefficient isolation retained from the original study. The astrocyte-derived term $\varepsilon\hat{C}_i(t)$ is held fixed.

| Active term | $\alpha$ | $\beta$ | $\gamma$ | $\delta$ | Acc. (%) |
|---|---|---|---|---|---|
| Reference | 0 | 0 | 0 | 0 | 27.50 |
| Local | 2 | 0 | 0 | 0 | 50.06 |
| Heterosynaptic | 0 | 0.2 | 0 | 0 | 72.05 |
| Somatic/BAP | 0 | 0 | 0.7 | 0 | 80.00 |
| Supervisor | 0 | 0 | 0 | 1 | 73.95 |

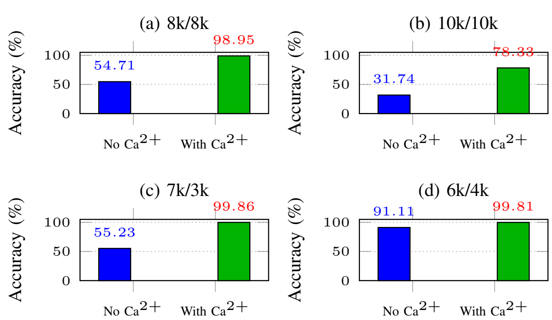


Fig. 2: Fixed-size split diagnostic, retained only for continuity with the submission. The revised claims use the stronger calibration and chronological analysis below.

controls.

### A. Coefficient and Legacy Diagnostics

The single-term study isolates the activity-dependent components of Eq. (22) while holding the multicellular term fixed (Table II). Each term changes learning, with the largest isolated effect produced by the somatic/BAP component.

The fixed-size splits showed large apparent gains against an untuned architecture-matched DNN (Fig. 2). A separate model diagnostic also showed that, from 10 to 30 epochs, TP increased from 1924 to 2159 while FP fell from 1072 to 613 and FN from 423 to 188. Component attribution is instead based on the matched factorial experiment in the following.

### B. Leakage-Controlled Temporal Evaluation

The primary evaluation uses the three temporally separated CTU-13 Neris captures as sequential environments and preserves their chronological order. To reduce reliance on dataset-specific identifiers, source and destination addresses, ports, and absolute time are removed from the input representation. This setting therefore evaluates whether the models can adapt to changes in traffic and attack distributions using traffic-derived features rather than memorising persistent identifiers.

We compare the proposed $Ca^{2+}$-modulated model with matched adaptive DNN baselines under the same architecture, data order, label availability, and optimisation budget. Results are reported across five random seeds. As shown in Fig. 3, the full $Ca^{2+}$ model achieves the highest mean environment PR-AUC, F1 score, and balanced accuracy, while also producing the lowest mean cumulative false-negative count among the principal adaptive comparators. Class-weighted learning and EWC remain competitive, whereas replay exhibits greater variability across seeds. The differences indicate that the benefit of $Ca^{2+}$ modulation lies in improving adaptation under temporal distribution shift rather than simply increasing performance.

### C. Factorial and Mechanistic Attribution

The factorial decomposition in Fig. 4 separates the contribution of $Ca^{2+}$-dependent gating from conventional regularisation. Laplacian regularisation alone remains close to the baseline, whereas $Ca^{2+}$ gate-only, $Ca^{2+}$ plus Laplacian, and the full model consistently achieve higher performance. Setting the multicellular $Ca^{2+}$ coefficient to zero ($\varepsilon = 0$)

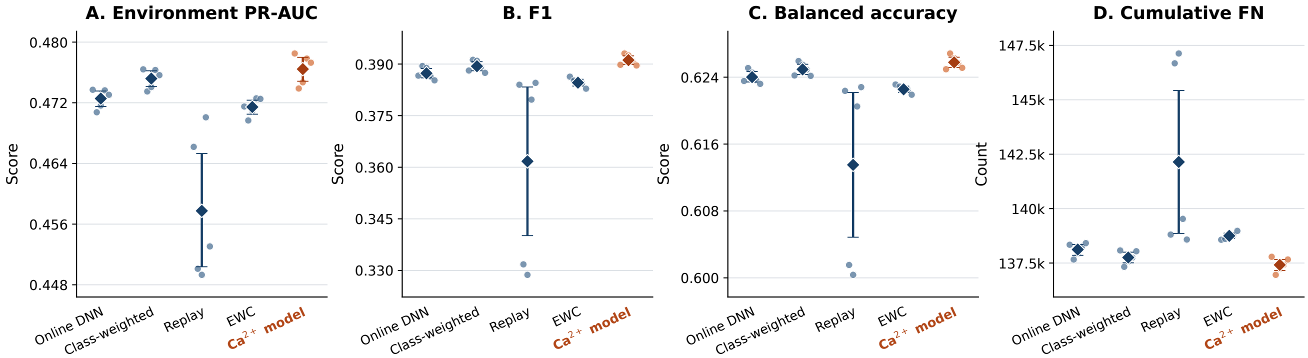


Fig. 3: Five-seed comparison under the leakage-controlled chronological protocol, with address-, port-, and absolute-time features removed. Points denote individual seeds, red markers indicate means, and black intervals denote uncertainty.

shifts performance back toward baseline. Random amplitude-matched, time-shuffled, and constant-mean $Ca^{2+}$ controls retain part of the performance gain, indicating that the benefit is associated more strongly with $Ca^{2+}$-dependent metaplastic regulation than with the exact temporal waveform of the simulated $Ca^{2+}$ signal.

Biophysical and structural ablations further clarify where this sensitivity arises. Removing gap-junction coupling, CICR, or SERCA produces comparatively small changes in adaptation PR-AUC, whereas removing metaplasticity or shuffling the astrocyte-to-synapse mapping $Q$ leads to the largest increases in cumulative false negatives. Representative dynamics show that $C_i - \theta_i$, the signed modulator, and the gradient multiplier $1 + \lambda_m m_i$ vary across environment transitions, confirming that the $Ca^{2+}$ pathway provides state-dependent modulation rather than a constant rescaling of the learning rate.

**Runtime.** On the same workload and machine, the baseline DNN requires 13.82 s wall time and 1.93 s user CPU time; executing $Ca^{2+}$ dynamics during training requires 24.20 s and 3.86 s, respectively. The present prototype therefore incurs $1.75\times$ wall-clock, $2.00\times$ user-CPU, and $2.99\times$ profiler-reported peak-memory overhead.

## IV. Discussion

The results show that astrocyte-informed $Ca^{2+}$ modulation can improve adaptation under temporal distribution shift when compared with matched adaptive DNN baselines. Under the leakage-controlled chronological protocol, the full $Ca^{2+}$ model achieves the strongest mean performance across the primary metrics, including PR-AUC, F1, balanced accuracy, and cumulative false negatives. The factorial analysis attributes most of this improvement to $Ca^{2+}$-dependent gating rather than Laplacian regularisation alone, while suppressing the multicellular contribution ($\varepsilon = 0$) shifts performance back toward the baseline. At the same time, random amplitude-matched, time-shuffled, and constant-mean $Ca^{2+}$ controls retain part of the performance gain. The results support a broader role for $Ca^{2+}$-dependent metaplastic regulation in controlling learning.

The mechanistic ablations further indicate that the main task-level sensitivity lies at the interface between multicellular $Ca^{2+}$ state and synaptic plasticity. Removing CICR, SERCA, or gap-junction coupling alters the simulated biological dynamics but produces comparatively small changes in aggregate PR-AUC. In contrast, removing metaplasticity or disrupting the astrocyte-to-synapse mapping $Q$ produces larger increases in missed detections. Together with the observed variation of $C_i - \theta_i$, the signed modulator, and the gradient multiplier across environment transitions, this suggests that the computational effect arises primarily from the state-dependent conversion of $Ca^{2+}$ activity into local plasticity control.

This interpretation also defines the biological scope of the model. $IP_3$-mediated release, SERCA uptake, and gap-junction diffusion are based on established astrocytic mechanisms, whereas the fixed mapping $Q$, per-synapse normalisation, and supervisory contribution are engineering abstractions introduced to couple the biological model to the DNN.

The additional biological dynamics introduce a measurable computational cost. In the current implementation, $Ca^{2+}$-coupled training increases wall-clock time by $1.75\times$ and peak memory use by $2.99\times$ relative to the matched DNN implementation. Because the present implementation is a research prototype, there is also scope for reducing this cost through implementation-level optimisation, including precomputation or caching of $Ca^{2+}$ trajectories where closed-loop coupling is not required. Future work should therefore examine external datasets, alternative architectures, and more tightly closed-loop coupling between neural activity and astrocyte dynamics.

## V. Conclusion

We introduced an astrocyte-informed learning framework that couples multicellular $Ca^{2+}$ dynamics to metaplastic gradient modulation in a DNN. Under a leakage-controlled chronological evaluation of CTU-13 Neris traffic, the full $Ca^{2+}$ model achieves the strongest mean performance among the evaluated adaptive baselines and reduces missed detections under temporal distribution shift. Factorial ablations show that $Ca^{2+}$-dependent gating, rather than Laplacian regularisation alone, accounts for most of the observed improvement, while removing the multicellular $Ca^{2+}$ contribution substantially reduces the benefit. Matched signal controls show that the metaplasticity and the astrocyte-to-synapse mapping have the strongest effects on missed-attack performance. These results support a cross-scale computational principle in which multi-

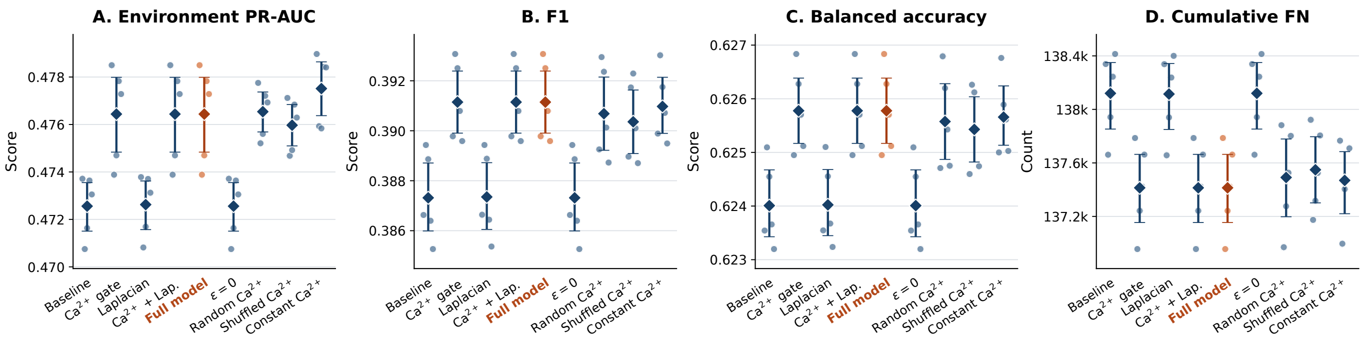


Fig. 4: Five-seed factorial decomposition under the leakage-controlled chronological protocol.

cellular state regulates artificial plasticity, while broader generalisation and more biologically faithful closed-loop coupling remain open directions.

## REFERENCES


[1] M. A. I. Mallick and R. Nath, "Navigating the cyber security landscape: A comprehensive review of cyber-attacks, emerging trends, and recent developments," *World Scientific News*, vol. 190, no. 1, pp. 1–69, 2024.

[2] Y. Xin *et al.*, "Machine learning and deep learning methods for cybersecurity," *IEEE Access*, vol. 6, no. 6, pp. 35 365–35 381, 2018.

[3] S. García, M. Grill, J. Stiborek, and A. Zunino, "An empirical comparison of botnet detection methods," *Computers & Security*, vol. 45, pp. 100–123, 2014.

[4] D. N. Katiyar, M. S. Tripathi, M. P. Kumar, M. S. Verma, D. A. K. Sahu, and D. S. Saxena, "Ai and cyber-security: Enhancing threat detection and response with machine learning," *Educational Administration: Theory and Practice*, vol. 30, no. 4, pp. 6273–6282, 2024.

[5] J. N. Chukwunweike, A. Praise, and B. A. Bashirat, "Harnessing machine learning for cybersecurity: How convolutional neural networks are revolutionizing threat detection and data privacy," *International Journal of Research Publication and Reviews*, vol. 5, no. 8, 2024, dOI not available.

[6] J. P. Pramod, B. G. Yadav, and S. Fatima, "Machine learning meets cybersecurity," *International Journal of Scientific Research in Science, Engineering and Technology*, vol. 11, no. 6, pp. 17–24, 2024.

[7] S. Mederos, C. González-Arias, and G. Perea, "Astrocyte–neuron networks: A multilane highway of signaling for homeostatic brain function," *Frontiers in Synaptic Neuroscience*, vol. 10, no. 45, 2018.

[8] T. Moldwin, L. S. Azran, and I. Segev, "The calcitron: A simple neuron model that implements many learning rules via the calcium control hypothesis," *PLOS Computational Biology*, vol. 21, no. 1, p. e1012754, 2025.

[9] S. Balasubramaniam, S. Kadry, M. K. TK, and K. S. Kumar, *Bio-inspired Algorithms in Machine Learning and Deep Learning for Disease Detection*. CRC Press, 2025.

[10] M. Aslam, D. Ye, M. Hanif, and M. Asad, "Adaptive machine learning: A framework for active malware detection," in *2020 16th International Conference on Mobility, Sensing and Networking (MSN)*. IEEE, 2020, pp. 57–64.

[11] C. T and D. V. Jose, "Future prospects of bio-inspired algorithms based ml/dl models in healthcare," in *2024 2nd International Conference on Computing and Data Analytics (ICCDA)*, 2024, pp. 1–6.

[12] J. Lorenzo, J.-A. Rico-Gallego, S. Binczak, and S. Jacquir, "Spiking neuron-astrocyte networks for image recognition," *Neural Computation*, vol. 37, no. 4, pp. 635–665, 2025.

[13] L. Kozachkov, K. V. Kastanenka, and D. Krotov, "Building transformers from neurons and astrocytes," *Proceedings of the National Academy of Sciences*, vol. 120, no. 34, p. e2219150120, 2023.

[14] G. Shen, D. Zhao, Y. Dong, Y. Li, J. Li, K. Sun, and Y. Zeng, "Astrocyte-enabled advancements in spiking neural networks for large language modeling," *arXiv preprint arXiv:2312.07625*, 2023.

[15] C. S. Yang, S. J. Gates III, D. De Zoysa, J. Choe, W. Losert, and C. B. Hart, "Characterizing learning in spiking neural networks with astrocyte-like units," *arXiv preprint arXiv:2503.06798*, 2025.

[16] M. T. Barros, P. Doan, M. Kandhavelu, B. Jennings, and S. Balasubramaniam, "Engineering calcium signaling of astrocytes for neural–molecular computing logic gates," *Scientific reports*, vol. 11, no. 1, p. 595, 2021.

[17] C. M. Schöfmann, M. Fasli, and M. T. Barros, "Investigating biologically plausible neural networks for reservoir computing solutions," *IEEE Access*, vol. 12, pp. 50 698–50 709, 2024.

[18] L. Gong, F. Pasqualetti, T. Papouin, and S. Ching, "Astrocytes as a mechanism for contextually-guided network dynamics and function," *PLoS computational biology*, vol. 20, no. 5, p. e1012186, 2024.

[19] M. T. Barros, S. Balasubramaniam, and B. Jennings, "Comparative end-to-end analysis of $ca^2+$-signaling-based molecular communication in biological tissues," *IEEE Transactions on Communications*, vol. 63, no. 12, pp. 5128–5142, 2015.

[20] M. T. Barros and S. Dey, "Feed-forward and feedback control in astrocytes for $ca^2+$-based molecular communications nanonetworks," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 17, no. 4, pp. 1174–1186, 2018.

[21] M. T. Barros, W. Silva, and C. D. M. Regis, "The multi-scale impact of the Alzheimer's disease on the topology diversity of astrocytes molecular communications nanonetworks," *IEEE Access*, vol. 6, pp. 78 904–78 917, 2018.

[22] K. Lenk *et al.*, "Larger connection radius increases hub astrocyte number in a 3d neuron–astrocyte network model," *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, vol. 7, no. 2, pp. 83–88, 2021.

[23] J. A. K. Hyttinen *et al.*, "Astrocytes in modulating subcellular, cellular and intercellular molecular neuronal communication," in *Proceedings of the 8th ACM International Conference on Nanoscale Computing and Communication (NanoCom '21)*, 2021, pp. 1–6.

[24] A. Araque, V. Parpura, R. P. Sanzgiri, and P. G. Haydon, "Tripartite synapses: glia, the unacknowledged partner," *Trends in Neurosciences*, vol. 22, no. 5, pp. 208–215, 1999.

[25] C. Eroglu and B. A. Barres, "Regulation of synaptic connectivity by glia," *Nature*, vol. 468, no. 7321, pp. 223–231, 2010.

[26] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *Journal of the American statistical Association*, vol. 102, no. 477, pp. 359–378, 2007.

[27] A. Citri and R. C. Malenka, "Synaptic plasticity: Multiple forms, functions, and mechanisms," *Neuropsychopharmacology*, vol. 33, no. 1, pp. 18–41, 2007.

[28] G. J. Stuart and B. Sakmann, "Active propagation of somatic action potentials into neocortical pyramidal cell dendrites," *Nature*, vol. 367, no. 6458, pp. 69–72, 1994.

[29] M. Graupner and N. Brunel, "Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location," *PNAS*, vol. 109, no. 10, pp. 3991–3996, 2012.

[30] C. M. Tigaret, V. Olivo, J. H. L. P. Sadowski, M. C. Ashby, and J. R. Mellor, "Coordinated activation of distinct ca2+ sources and metabotropic glutamate receptors encodes hebbian synaptic plasticity," *Nature Communications*, vol. 7, no. 1, 2016.