A MINIMAL COMPUTATIONAL ASTROCYTE AS A LOCAL SPATIOTEMPORAL OPERATOR

Anonymous authors
Paper under double-blind review

ABSTRACT

Astrocytes integrate neural activity across space and time and modulate nearby synapses, yet it remains unclear which computational primitive is provided by this additional dynamical state. We introduce a minimal computational astrocyte defined by a restricted interface: local event sensing, leaky temporal integration, threshold-triggered astrocytic events, and persistent local modulation of synaptic transmission and plasticity. The resulting mechanism decomposes into spatial pooling, temporal filtering, nonlinear gating, and local feedback modulation. For the deterministic mean-field trajectory under stationary input, we derive an exact threshold-crossing law and a critical input rate; for Poisson spikes, the event-rate law is an approximation. At finite observation times, spatial pooling reduces estimation variance while introducing spatial bias, yielding a predictable operating regime.

We embed the mechanism in a canonical unsupervised Diehl–Cook spiking network. Astrocytic integration parameters are calibrated once from input statistics and then frozen across corruption levels. At the nominal $350\,\mathrm{ms}$ presentation duration, the astrocytic network achieves 78.1% mean accuracy over four seeds under 75% dynamic salt-and-pepper corruption, while the unmodulated baseline reaches 9.8%. The advantage persists across a broad corruption range and then collapses near the analytically predicted loss of mean-field rate separation. Tests on native event streams probe the same local interface in frozen pretrained networks. These results identify a compact computational role for astrocyte-like state: local stateful gating of event-driven information flow.

1 INTRODUCTION

Astrocytes are increasingly viewed as active participants in neural information processing rather than purely supportive cells. They sense local synaptic and neuronal activity, integrate signals over spatially extended domains and comparatively long time scales, and can modulate nearby synapses and neurons (Santello et al., 2019; Kastanenka et al., 2020; Murphy-Royal et al., 2023). These properties make astrocytes computationally interesting because they add a second dynamical layer to neuronal networks: information can be accumulated outside the neuron and returned as a local modulatory signal.

The modeling consequence is less clear. Existing neuron–astrocyte models range from detailed intracellular descriptions to strongly reduced phenomenological units, and similar network-level effects can often be implemented at several levels of abstraction. This creates a basic computational question: *what should remain when an astrocyte model is reduced as far as possible while preserving a distinct computational role?* Without an explicit restriction, this question is ill posed. If neurons or synapses are allowed arbitrary hidden states and interactions, the same computation could simply be absorbed into an enlarged neuronal model.

We therefore take a functional, interface-based view. A computational astrocyte is not defined by reproducing all intracellular biophysics, but by preserving a restricted pattern of interaction: a separate state senses local neural activity, integrates evidence across space and/or time, and feeds a local modulatory signal back to neurons or synapses. Within this interface we study a near-minimal event-driven realization with only two dynamical stages: a leaky local accumulator followed by a

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

persistent modulatory state. The resulting model is simple enough to analyze exactly in mean field, yet retains the spatial, temporal, and feedback structure that motivates an astrocytic interpretation.

We define the interface and derive its stationary and finite-time operating regimes, then evaluate adaptation to switching inputs. In a trained Diehl–Cook network on MNIST (Diehl & Cook, 2015), we separate input and plasticity modulation against filtering controls. Frozen pretrained networks test transfer to native N-MNIST and DVS Gesture events; structured N-MNIST corruption probes a failure mode of the original local accumulator.

2 Related work

Astrocytes as computational elements. Experimental and conceptual work emphasizes that astrocytes integrate and modulate neural activity across multiple spatial and temporal scales (Santello et al., 2019; Kastanenka et al., 2020; Murphy-Royal et al., 2023). Computational studies have used this principle in several forms. Reduced neuron–astrocyte models describe bidirectional regulation (Gordleeva et al., 2012), spatiotemporal detection of neuronal activity (Gordleeva et al., 2019), persistent effects that support working-memory dynamics (Gordleeva et al., 2021), dynamic stimulus representation (Stasenko & Kazantsev, 2023), and image recognition in spiking neuron–astrocyte networks (Lorenzo et al., 2025). These results show that detailed biochemical dynamics are not always required for useful network computation.

We use a restricted interface to isolate spatial pooling, temporal integration, thresholded events, and persistent feedback without attributing any one operation uniquely to astrocytes.

Filtering and robustness in spiking systems. Noise suppression can also be achieved without astrocytes. Spatial denoising, temporal filtering, adaptive thresholds, and preprocessing of event streams are standard tools in signal processing and neuromorphic systems (Hartbauer, 2020; Marchisio et al., 2021; Xu et al., 2026), within the broader context of spike-based machine intelligence (Roy et al., 2019). This is an important comparison class because our minimal model itself contains familiar operations: local averaging, leaky integration, thresholding, and gain modulation.

We compare against global attenuation, static spatial preprocessing, hard local gating, removal of spatial pooling, and an online event filter on the same realized spikes to isolate the effect of the full local feedback loop.

3 What constitutes a computational astrocyte?

Our interface abstracts recurring properties of astrocyte models and experimental descriptions: local sensing of neuronal activity, integration across space and time, and feedback onto nearby neural or synaptic processes (Gordleeva et al., 2012; 2019; Kastanenka et al., 2020; Murphy-Royal et al., 2023). We require an astrocytic unit to satisfy five properties:

- **Separate state:** at least one dynamical variable distinct from the neurons and synapses it interacts with.
- **Local sensing:** input from a spatially or structurally restricted set of neuronal or synaptic variables.
- **Spatiotemporal integration:** integration over time, multiple local inputs, or both.
- **Local feedback modulation:** modulation of neuronal or synaptic variables within the associated local domain, producing the interaction $N/S \to A \to N/S$.
- **Interpretability:** explicit meanings for state variables and parameters such as integration time, threshold, receptive-field scale, and modulation strength.

These constraints deliberately limit expressive power: without them, an arbitrary hidden dynamical system could be labeled an astrocyte. Our goal is not a biological definition, but a computational interface that preserves the astrocyte as a distinct, local, and interpretable component. An abstract slow local interneuron could implement the same input–output operator; these functional criteria alone do not establish cell identity.

The persistent state may be cached at affected synapses; its storage location does not change the restricted sensing and feedback interface. Biochemical fidelity is evaluated separately in Appendix C.

4 MINIMAL ASTROCYTE MODEL

Motivated by reduced astrocyte models based on local activity integration and threshold-triggered modulation (Gordleeva et al., 2012; 2019), each astrocyte a receives spikes from a local domain $\mathcal{N}_a$ and integrates them through a single leaky state g_a,

$$\frac{dg_a}{dt} = -\frac{g_a}{\tau_g} + \frac{A_g}{|\mathcal{N}_a|} \sum_{i \in \mathcal{N}_a} \sum_k \delta(t - t_i^k), \tag{1}$$

so one input spike increments g_a by $A_g/|\mathcal{N}_a|$; for a full 3×3 domain this is $A_g/9$. Here δ has units of inverse time, A_g has units of g per event, and $\mathcal{N}_a$ denotes the input sites actually included in the local average. At a boundary, the same cardinality $|\mathcal{N}_a|$ must be used in both the spike update and the spatial average below. On threshold crossing we use a subtractive reset,

$$g_a \leftarrow g_a - \theta_g, \qquad g_a \geq \theta_g, \tag{2}$$

which preserves threshold overshoot.

An astrocytic event updates a modulatory state associated with affected neurons,

$$\nu_i \leftarrow \nu_i + A_\nu, \qquad i \in \mathcal{O}_a, \qquad \tau_\nu \frac{d\nu_i}{dt} = -\nu_i. \tag{3}$$

Overlapping output domains add their contributions to the same local modulatory state. The connectivity geometry is otherwise unrestricted: local domains may be spatial, graph-based, or defined by another interpretable notion of locality.

The model therefore consists of four operations,

$$\text{spatial pooling} \rightarrow \text{temporal filtering} \rightarrow \text{nonlinear gating} \rightarrow \text{persistent local modulation}. \tag{4}$$

Several familiar mechanisms appear as limiting cases: small τ_g approaches coincidence detection, $\tau_g \rightarrow \infty$ approaches a persistent accumulator, $|\mathcal{N}_a| = 1$ removes population-level spatial pooling, and suppressing thresholded event generation produces a continuous local filter. These are useful as ablations rather than as uniquely astrocytic operations.

The two time constants have distinct roles. The integration constant τ_g determines how long presynaptic evidence remains available before an astrocytic event is generated; τ_ν determines how long the *consequence* of that event persists after threshold crossing. Likewise, θ_g controls activity selectivity, whereas A_ν and the downstream modulation scale control the magnitude of the returned signal. Keeping these roles separate is useful experimentally because detector selectivity can be calibrated from input statistics without simultaneously fixing the strength of network-level modulation.

We cache ν at affected input sites for causal online modulation.

5 ANALYSIS OF THE MINIMAL OPERATOR

5.1 ACTIVITY SELECTION AND STATIONARY OPERATOR

Let $N_i(t)$ count input spikes at site i. Between resets, Eq. 1 is equivalently $dg_a = -(g_a/\tau_g)\,dt + (A_g/|\mathcal{N}_a|)\sum_{i \in \mathcal{N}_a} dN_i$. If the stationary rates are λ_i, then $\mathbb{E}[dN_i] = \lambda_i\,dt$ and the corresponding deterministic mean-field drift is

$$\frac{dg_a}{dt} = -\frac{g_a}{\tau_g} + A_g(B_K\lambda)_a, \qquad (B_K\lambda)_a = \frac{1}{|\mathcal{N}_a|} \sum_{i \in \mathcal{N}_a} \lambda_i.$$

For a homogeneous local rate, $(B_K\lambda)_a = \lambda$. Starting from $g = 0$, the threshold-crossing time and corresponding astrocytic event rate are

$$t_{\text{fire}} = \tau_g \log\left(\frac{\lambda A_g \tau_g}{\lambda A_g \tau_g - \theta_g}\right), \qquad f_A(\lambda) = t_{\text{fire}}^{-1}, \tag{5}$$

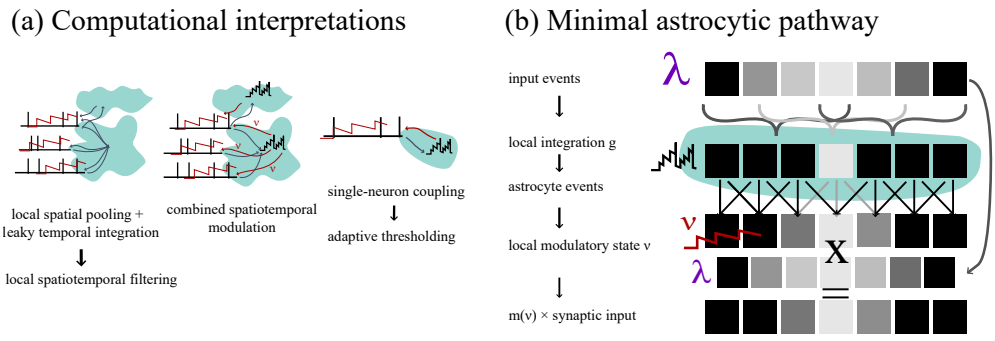


Figure 1: Minimal astrocytic computation. (a) Limiting or ablated configurations recover familiar primitives such as local spatiotemporal filtering and single-channel adaptive gating. (b) The event-driven pathway: local input events are integrated in g, threshold crossings generate astrocytic events, and these update the persistent modulatory state ν used to gate nearby synaptic input.

defined only for

$$\lambda > \lambda_c = \frac{\theta_g}{A_g \tau_g}. \tag{6}$$

Thus g performs activity selection, whereas ν controls persistence: for a fixed event rate,

$$\langle \nu \rangle = A_\nu \tau_\nu f_A. \tag{7}$$

Let $\lambda(\mathbf{x})$ denote a stationary rate field; B_K is the normalized input average just defined. Define the output backprojection by $(R_K z)_i = \sum_{a: i \in \mathcal{O}_a} z_a$: each overlapping astrocyte contributes its full modulatory increment to neuron i. In particular, R_K is an unnormalized incidence backprojection rather than the matrix transpose of the normalized input averaging operator B_K. This follows from Eq. 3: an astrocytic event adds A_ν to *each* output site, without a factor of $1/|\mathcal{O}_a|$. The stationary mean-field action before saturating the downstream readout is

$$\boxed{\mathcal{A}_K[\lambda] = \lambda \odot R_K \left[A_\nu \tau_\nu f_A(B_K \lambda) \right].} \tag{8}$$

This exposes the computation as local aggregation, nonlinear activity selection, spatial backprojection, and multiplicative modulation. The critical rate λ_c is particularly important because it turns a continuous input-rate field into a sparse field of astrocytic events without requiring a separate hand-designed threshold at the neuronal output. Above threshold, the exact logarithmic dependence in Eq. 5 controls how rapidly event rate grows with local drive. Stochastic event-driven simulations closely follow this mean-field operator away from the threshold boundary (Fig. 2a,b). The first-passage law is exact for the deterministic mean-field trajectory starting from zero and resetting to zero at threshold. For the discrete spike-driven process, Poisson fluctuations can produce crossings below λ_c even at long observation times, and subtractive reset preserves overshoot. Thus the law approximates, rather than exactly equals, the stationary event rate of the stochastic implementation.

5.2 Finite-time regime

For a Poisson source observed for duration T,

$$\hat{\lambda}_i(T) = \frac{N_i(T)}{T}, \qquad \mathbb{E}\hat{\lambda}_i = \lambda_i, \qquad \mathrm{Var}\hat{\lambda}_i = \frac{\lambda_i}{T}. \tag{9}$$

For a d-dimensional spatial domain (here $d = 2$), a full neighborhood contains K^d independent sources; near boundaries it contains only $|\mathcal{N}_a|$ sources. Their locally averaged rate estimator has

$$\mathrm{Var}\hat{\lambda}_a(T) = \frac{\bar{\lambda}_a}{|\mathcal{N}_a| T}, \qquad |\mathcal{N}_a| = K^d \text{ for full domains.} \tag{10}$$

Larger domains reduce sampling variance but blur spatial structure. Longer integration reduces stochastic variability while delaying responses to changes in the underlying signal. The useful regime is therefore finite in both spatial scale and integration time.

At very short durations there is insufficient evidence; at intermediate durations local pooling can recover structure from fewer samples; and at long durations the raw spike histogram itself converges to the rate field. The simulation in Fig. 2c illustrates a bounded region in the (K, T) plane where astrocytic modulation improves reconstruction. We define the panel's correlation improvement against the clean underlying rate image λ^{clean} as

$$\Delta\rho(K,T) = \rho\Big(\widehat{\lambda}(T) \odot m(\nu_{K,T}), \lambda^{\mathrm{clean}}\Big) - \rho\Big(\widehat{\lambda}(T), \lambda^{\mathrm{clean}}\Big), \qquad \widehat{\lambda}_i(T) = N_i(T)/T, \tag{11}$$

where ρ is the Pearson correlation across image pixels (defined when both spatial fields have nonzero variance) and $m(\nu) = \min(1, \nu)$ is the model's saturating local readout. This observation is important for the end-to-end experiment below: the astrocytic pathway is expected to be most useful when the neuronal network must act before independent spike counts have accurately resolved the local rate structure.

5.3 Adaptation to switching spatial inputs

The stationary operator does not determine how quickly the persistent state represents a *new* input after a switch. We alternate two complementary 28×28 spatial patterns with pixel rates of 22.5 and $37.5\,\mathrm{Hz}$. Each pattern is shown for dwell time B before switching to the other. Neither the detector state g nor the modulatory state ν is reset between switches. This controlled experiment uses synthetic spatial patterns rather than successive MNIST samples, isolating adaptation of the local operator from classification.

Reorientation after an image switch. At $\tau_\nu = 200\,\mathrm{ms}$ we vary B and measure Pearson correlation of the spatial ν field with the *current* pattern, both averaged during the dwell and at its end. We also measure the first time after each switch at which this correlation becomes nonnegative. These readouts test whether the previous image persists long enough to delay representation of the new one.

Persistence-time sweep. We cross dwell times $B \in \{20, 50, 75, 100, 150, 200, 300, 400, 600, 800\}\,\mathrm{ms}$ with $\tau_\nu \in \{25, 50, 100, 200, 400, 800\}\,\mathrm{ms}$. The detector is held at $K = 3$, $\tau_g = 105\,\mathrm{ms}$, and $\theta_g = 3.35$, with $A_\nu = 1$ fixed. For each seed and B, the same Poisson input and astrocytic event stream is reused across all τ_ν. Six alternating segments establish the state; the following twelve are measured. Figure 3b displays their mean within-dwell correlation with the current pattern.

6 Running examples and end-to-end evaluation

We use three distinct evaluation settings. Static MNIST provides a controlled Poisson-rate task in which we train an unsupervised Diehl–Cook network and can separate modulation of incoming synaptic transmission from modulation of the plasticity rate. N-MNIST and DVS Gesture supply native event streams and separate pretrained spiking networks; in those settings we freeze neuronal weights and insert the astrocytic operator at the input, without retraining or a plasticity-rate intervention. These settings do not share a neuronal architecture or an input-noise model.

6.1 MNIST: training and the two modulation pathways

We embed the minimal module in the canonical unsupervised Diehl–Cook network (Diehl & Cook, 2015). Each MNIST pixel (LeCun, 1998) drives a Poisson input neuron with maximum rate $60\,\mathrm{Hz}$. Astrocytes use overlapping normalized 3×3 input domains, so each presynaptic event in a full nine-site domain contributes $A_g/9$ to g, and astrocytic state is reset between examples. Simulations are implemented in Brian 2 (Stimberg et al., 2019).

For salt-and-pepper corruption probability p, we use the marginal corrupted rate

$$r_{\mathrm{noisy}} = (1-p)r + \frac{p}{2} r_{\max}, \qquad r_{\max} = 60\,\mathrm{Hz}, \tag{12}$$

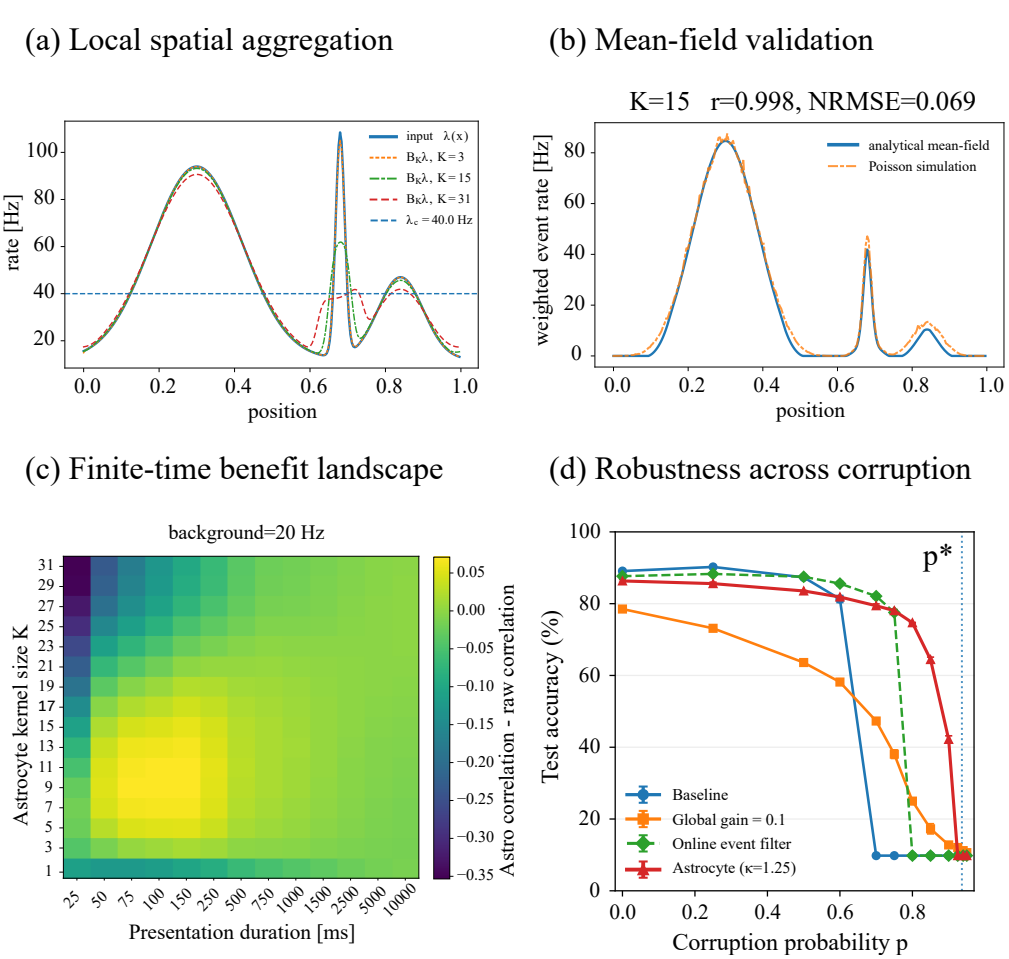


Figure 2: Spatial, finite-time, and end-to-end behavior. (a) Local aggregation of the input rate field for several domain sizes K. (b) Analytical mean-field operator versus stochastic event-driven simulation. (c) Difference in pixelwise Pearson correlation with the clean rate image between the locally modulated and raw finite-time rate fields, defined in Eq. 11, versus K and duration T. (d) Diehl–Cook test accuracy versus dynamic salt-and-pepper corruption for the unmodified baseline, fixed global-gain control, online event-domain spatiotemporal filter, and minimal astrocyte. The astrocytic curve uses modulation scale $\kappa = 1.25$. Points show mean $\pm$ s.d. over four seeds on the matched corruption grid. The dashed boundary marks $p^\star$, the predicted loss of mean-field rate separation.

and generate Poisson events online from this rate field. The nominal presentation duration is $T = 350\,\mathrm{ms}$. We train the Diehl–Cook network with unsupervised plasticity on corrupted training presentations; corruption is therefore present during learning, not introduced solely at test time. Following the Diehl–Cook readout (Diehl & Cook, 2015), each excitatory neuron receives one digit label after training, selected by its highest average response to labeled examples of that digit. Multiple neurons may receive the same label; classification averages the responses of neurons assigned to each digit and selects the largest class average. Matched astrocytic and non-astrocytic runs use the same neuronal architecture, training protocol, presentation schedule, and random seeds. The astrocytic state is reset between examples so that modulation reflects evidence from the current stimulus rather than memory carried across samples.

We use corruption as a probe rather than as a claim of state-of-the-art denoising. As p increases, contrast between low- and high-rate pixels is compressed toward the common midpoint of $30\,\mathrm{Hz}$,

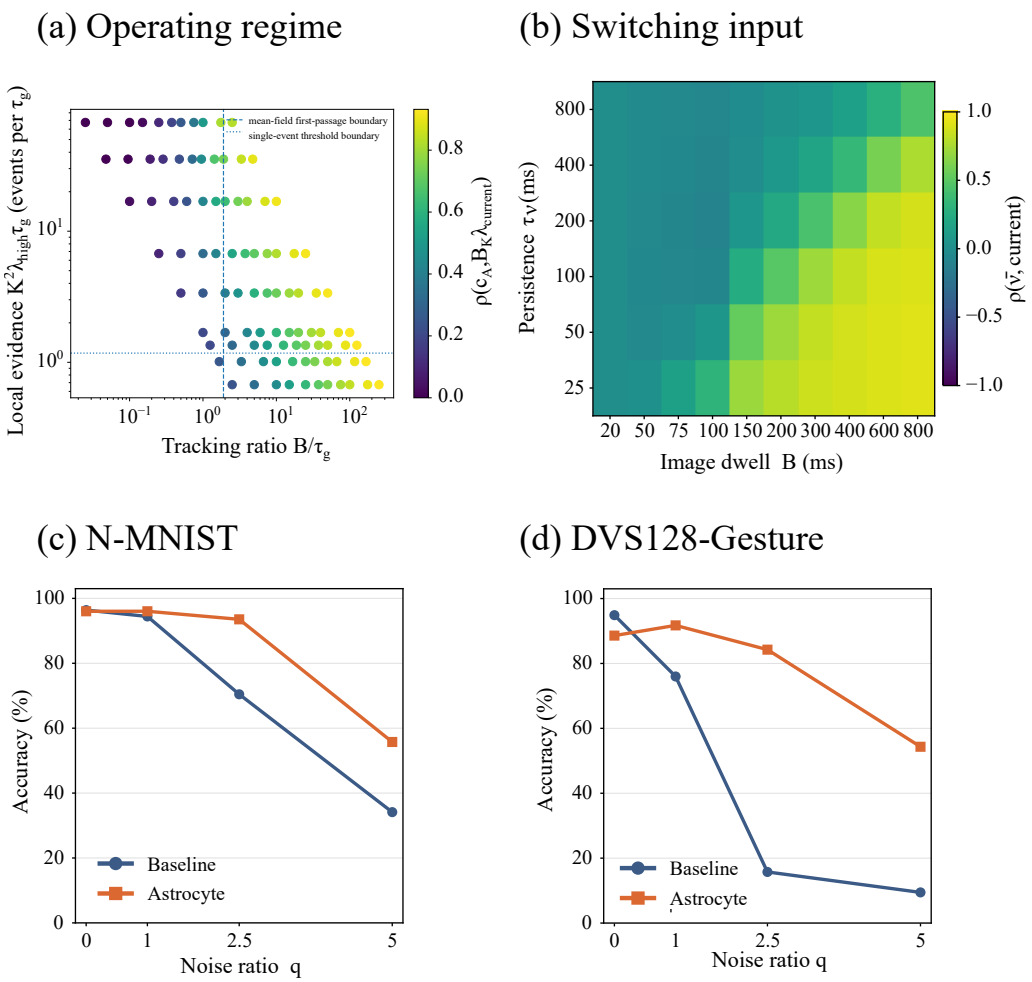


Figure 3: Operating regimes, switching dynamics, and native-event transfer. (a) Correlation across local evidence and the input tracking ratio; dashed lines mark the mean-field first-passage and single-event boundaries. (b) Mean within-dwell spatial correlation of ν with the current synthetic pattern versus dwell time B and persistence τ_ν; neither g nor ν is reset at a switch. (c) N-MNIST and (d) DVS128-Gesture test accuracy with and without astrocytic input modulation. In (c,d), q denotes the expected ratio of injected to clean sensor events; each paired comparison uses the same corrupted event stream. The pretrained neuronal weights remain fixed.

creating a controlled loss of rate separability. This makes it possible to connect a network-level failure regime to the detector-level threshold derived above.

Calibration. We calibrate τ_g and θ_g once from input statistics rather than downstream classification accuracy. At the nominal condition $p = 0.75$, inactive and maximally active clean rates map to 22.5 and 37.5 Hz. For the finite presentation window we require a first-passage false-crossing probability of at most approximately 2.3% and a true-crossing probability of at least approximately 97.7%, using $M(T) = \max_{0<t<T} g(t)$. This gives an operating region around $\tau_g = 100$–110 ms and $\theta_g = 3.2$–3.5 for $A_g = 1$. We use

$$\tau_g = 105\,\text{ms}, \qquad \theta_g = 3.35, \qquad \lambda_c \simeq 31.90\,\text{Hz},$$

and freeze these values for all corruption levels.

The modulatory state uses $A_\nu = 1$ and $\tau_\nu = 200$ ms with saturating readout $m(\nu) = \min(1, \nu)$. In the Diehl–Cook network the same local readout modulates synaptic transmission and the local

plasticity rate,

$$\Delta g_e = \kappa\, m(\nu_i) w, \qquad \eta_i = \kappa\, m(\nu_i)\eta_0,$$

In Fig. 2d we use $\kappa = 1.25$. This scalar changes the downstream modulation strength but not λ_c or the astrocytic event boundary. Under clean input at this operating point, mean excitatory activity is 33.06 ± 0.12 versus 32.80 ± 0.35 spikes/image for the unmodified network (four seeds).

To separate the two feedback routes, matched MNIST ablations train with input-transmission modulation only, plasticity-rate modulation only, both pathways, or neither. The detector and the neuronal training protocol are held fixed across conditions.

Predicted corruption boundary. For the limiting clean rates 0 and 60 Hz, the marginal corruption model gives

$$\lambda_-(p) = 30p, \qquad \lambda_+(p) = 60 - 30p.$$

Mean-field separation requires $\lambda_-(p) < \lambda_c < \lambda_+(p)$. For the fixed operating point, the upper condition is lost at

$$p^\star = \frac{60 - \lambda_c}{30} \simeq 0.9365. \tag{13}$$

This is a predicted loss of detector-level rate separation, not a prediction of the exact downstream classification transition.

Robustness across corruption. The curve in Fig. 2d uses $\kappa = 1.25$ (four seeds). At $p = 0$, the baseline and astrocytic networks achieve 89.08% and 86.33% mean accuracy, respectively. The baseline remains accurate through moderate corruption but collapses between $p = 0.6$ and 0.7. The astrocytic network degrades smoothly, achieving 79.42% at $p = 0.7$, 78.06% at $p = 0.75$, 74.68% at $p = 0.8$, and 42.17% at $p = 0.9$. It reaches chance-level 9.8% at $p = 0.925$ as corruption approaches the mean-field boundary in Eq. 13.

The online-filter control integrates the same realized spikes over a 3×3 domain with the same τ_g, reading out continuously without an astrocytic threshold, reset, or persistent ν. Additional duration, spatial and modulation-strength controls appear in Appendix A.

6.2 N-MNIST and DVS Gesture: pretrained native-event networks

For N-MNIST (Orchard et al., 2015), we evaluate a separately pretrained HM2-BP spiking network (Jin et al., 2018) on the native event stream; its weights are from the released ETRI NEST-SNN implementation (`https://github.com/etri/nest-snn`). The astrocytic detector processes events causally before the network's sensor-motion stabilization and $600\,\mu$s temporal bucketing; its pre-event local state scales input transmission to the first pretrained layer. All neuronal weights remain frozen. The paired baseline and modulated branches receive the same corrupted events, and no learning-rate modulation is applied during inference. Figure 3c compares baseline and astrocytic test accuracy as a function of the uniform native-event noise ratio q.

For DVS128 Gesture (Amir et al., 2017), we use the pretrained, post-training-quantized convolutional SNN from TU Dresden's PTQ release (Arfa et al., 2025), implemented in snnTorch (Eshraghian et al., 2023). We retain its event-to-frame inference pipeline. Noise is added to the native 128×128 sensor stream. Local astrocytic weights are computed on that stream before denoising, spatial downsampling, and framing, and travel with surviving events to modulate the input to the frozen network. The baseline and astrocytic branches again share each corrupted recording. Figure 3d shows the paired accuracy sweep for this network. These native-event studies test input insertion into fixed networks; the MNIST experiment also tests training. Their noise axes are not pooled.

Structured noise on N-MNIST. The uniform-noise sweep does not test concentrated or bursty interference. Using the same frozen N-MNIST network, we therefore compare the baseline, the original minimal astrocyte, and an alternative local astrocyte that thresholds the *fraction of distinct channels with recent glutamate*. We evaluate clean input once, and temporal-burst and hot-pixel corruption at $q \in \{5, 10, 20\}$, where q is the expected injected-to-clean event ratio. All three branches process each identical realized event stream; their states reset between samples. Each astrocyte keeps its own frozen parameters, and no weights are retrained. On the pilot-excluded set ($n = 9880$), the

active-channel detector is consistently worse than both comparators under temporal bursts, but under hot-pixel noise it reaches 90.29–90.45% accuracy versus 82.62–83.01% for the unmodulated baseline and 77.45–77.82% for the original minimal detector. The exact dynamics, noise protocol, and paired comparisons are given in Appendix B. Thus the extension changes the local input statistic rather than providing a uniformly stronger denoiser; it is not a calcium-dynamics comparison.

7 Discussion

The restricted interface combines pooling, integration, thresholded events, and persistent feedback in a local loop. Its detector state g controls activity selection, the critical rate λ_c describes the stationary regime, and ν determines feedback persistence. Added state and event processing are quantified in Appendix A; the distinct boundary of the biological comparison is in Appendix C.

Pooling across K^d independent sources reduces Poisson sampling variance as $1/K^d$ but blurs spatial structure; longer integration improves rate estimation but delays response. The duration sweep in Appendix A probes this finite-time trade-off.

In the matched pathway ablation (Table 4), transmission or plasticity-rate modulation alone improves performance at $p = 0.75$, while their combination gives the highest accuracy. The table does not establish statistical independence of the two routes.

The theoretical corruption boundary should be interpreted at the same level. Equation 13 predicts when the limiting low- and high-rate regimes can no longer be separated by the mean-field detector. It does not predict an exact classification accuracy or require the downstream network to fail precisely at $p^\star$. At $\kappa = 1.25$, accuracy collapses earlier, at $p = 0.925$, while remaining useful through $p = 0.9$. We therefore treat the agreement as an operating-range correspondence between detector theory and network behavior, not as a parameter-free prediction of the classifier.

Several limitations remain. The stationary Poisson calculation assumes independent inputs, while real sensor events can be correlated and bursty. The native-event experiments test transfer into frozen networks, not the stationary mean-field law. Equal expected injected-event ratios can yield different spatial concentrations: Table 5 shows that the active-channel extension helps under hot-pixel corruption but hurts under temporal bursts. This statistic-specific trade-off limits any claim of uniform robustness. The present model also omits intracellular calcium pathways, receptor-specific effects, and morphology. In an isolated Postnov-derived calcium-compartment comparison, the two local responses can be spatially similar under a short stationary input, but sustained input produces a regenerative calcium spike absent from the minimal state (Appendix C). We therefore interpret the minimal unit as a functional approximation to graded subthreshold integration, not to calcium-spike dynamics.

These limitations suggest a natural next step. The restricted interface can be kept fixed while the task changes. If a task requires signed integration, multiple interacting time scales, or state-dependent feedback, the minimal unit can be enriched one mechanism at a time. This provides a controlled way to ask not whether astrocytes are "necessary" for a computation, but which astrocyte-like dynamical resources are sufficient under a specified neural interface and temporal protocol.

AI use statement

In this work, generative AI tools were used to aid and polish writing; for retrieval and discovery of related work; for research ideation and execution, including experimental design and code assistance; to draft portions of the manuscript; and to assist with mathematical derivations. AI tools were not used to generate synthetic datasets. All AI-assisted mathematical claims, code, experimental results, citations, and manuscript text were reviewed and verified by the authors. The authors take responsibility for the final content of the paper.

Reproducibility statement

The model equations, calibration rule, corruption protocol, and experimental controls are specified in the main text. Full configuration files, per-seed results, plotting scripts, and implementation code will be provided in the supplementary material. Additional ablations are reported below.

REFERENCES

Arnon Amir, Brian Taba, David Berg, Timothy Melano, Jeffrey McKinstry, Carmelo Di Nolfo, Tapan Nayak, Alexander Andreopoulos, Guillaume Garreau, Marcela Mendoza, et al. A low power, fully event-based gesture recognition system. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7388–7397, 2017. doi: 10.1109/CVPR.2017.781.

Sirine Arfa, Bernhard Vogginger, Chen Liu, Johannes Partzsch, Mark Schöne, and Christian Mayr. Efficient deployment of spiking neural networks on SpiNNaker2 for DVS gesture recognition using neuromorphic intermediate representation, 2025. arXiv:2504.06748.

Evan Cresswell-Clay, Nathan Crock, Joël Tabak, and Gordon Erlebacher. A compartmental model to investigate local and global Ca2+ dynamics in astrocytes. *Frontiers in Computational Neuroscience*, 12:94, 2018. doi: 10.3389/fncom.2018.00094.

Peter U Diehl and Matthew Cook. Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Frontiers in Computational Neuroscience*, 9:99, 2015.

Jason K. Eshraghian, Max Ward, Emre Neftci, Xinxin Wang, Gregor Lenz, Girish Dwivedi, Mohammed Bennamoun, Doo Seok Jeong, and Wei D. Lu. Training spiking neural networks using lessons from deep learning. *Proceedings of the IEEE*, 111(9):1016–1054, 2023.

Susanna Yu Gordleeva, Sergey V Stasenko, Alexey V Semyanov, Alexander E Dityatev, and Victor B Kazantsev. Bi-directional astrocytic regulation of neuronal activity within a network. *Frontiers in Computational Neuroscience*, 6:92, 2012.

Susanna Yu Gordleeva, Anastasia V Ermolaeva, Innokentiy A Kastalskiy, and Victor B Kazantsev. Astrocyte as spatiotemporal integrating detector of neuronal activity. *Frontiers in Physiology*, 10: 294, 2019.

Susanna Yu Gordleeva, Yuliya A Tsybina, Mikhail I Krivonosov, Mikhail V Ivanchenko, Alexey A Zaikin, Victor B Kazantsev, and Alexander N Gorban. Modeling working memory in a spiking neuron network accompanied by astrocytes. *Frontiers in Cellular Neuroscience*, 15:631485, 2021.

Manfred Hartbauer. From insect vision to a novel bio-inspired algorithm for image denoising. In *Biomimetics*. IntechOpen, 2020.

Yingyezhe Jin, Wenrui Zhang, and Peng Li. Hybrid macro/micro level backpropagation for training deep spiking neural networks. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Ksenia V Kastanenka, Rubén Moreno-Bote, Maurizio De Pittà, Gertrudis Perea, Abel Eraso-Pichot, Roser Masgrau, Kira E Poskanzer, and Elena Galea. A roadmap to integrate astrocytes into systems neuroscience. *Glia*, 68(1):5–26, 2020.

Yann LeCun. The MNIST database of handwritten digits. *http://yann.lecun.com/exdb/mnist/*, 1998.

Jhunlyn Lorenzo, Juan-Antonio Rico-Gallego, Stéphane Binczak, and Sabir Jacquir. Spiking neuron-astrocyte networks for image recognition. *Neural Computation*, 37(4):635–665, 2025.

Alberto Marchisio, Giacomo Pira, Maurizio Martina, Guido Masera, and Muhammad Shafique. R-snn: An analysis and design methodology for robustifying spiking neural networks against adversarial attacks through noise filters for dynamic vision sensors. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6315–6321. IEEE, 2021.

Ciaran Murphy-Royal, ShiNung Ching, and Thomas Papouin. A conceptual framework for astrocyte function. *Nature Neuroscience*, 26(11):1848–1856, 2023.

Garrick Orchard, Ajinkya Jayawant, Gregory K Cohen, and Nitish Thakor. Converting static image datasets to spiking neuromorphic datasets using saccades. *Frontiers in Neuroscience*, 9:437, 2015. doi: 10.3389/fnins.2015.00437.

Dmitry E Postnov, Lyudmila S Ryazanova, and Olga V Sosnovtseva. Functional modeling of neural-glial interaction. *BioSystems*, 89(1–3):84–91, 2007. doi: 10.1016/j.biosystems.2006.04.012.

Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575(7784):607–617, 2019.

Mirko Santello, Nicolas Toni, and Andrea Volterra. Astrocyte function from information processing to cognition and cognitive impairment. *Nature Neuroscience*, 22(2):154–166, 2019.

Sergey V Stasenko and Victor B Kazantsev. Dynamic image representation in a spiking neural network supplied by astrocytes. *Mathematics*, 11(3):561, 2023.

Marcel Stimberg, Romain Brette, and Dan FM Goodman. Brian 2, an intuitive and efficient neural simulator. *eLife*, 8:e47314, 2019.

Yue Xu, Ye Zhao, Yumeng Ren, Long Chen, Liang Chen, Yulin Zhang, and Shushan Qiao. Aefsnn: Adaptive filtering spiking neural network for event-based sensors. *Applied Sciences*, 16(4):2073, 2026.

A Additional ablations

Computational overhead of the local operator. Let M be the number of local detectors, N the number of modulated input sites, and E_{in} and E_{out} their respective neighborhood incidences. The two dynamical arrays g and ν require $M + N$ scalar states; a sparse event implementation additionally stores $M+N$ last-update times and neighborhood indices, for $O(M+N+E_{\text{in}}+E_{\text{out}})$ memory. With square $K\times K$ domains, one input event visits at most K^2 detectors, and each detector that fires updates at most K^2 output sites. Thus its local processing cost is $O(K^2+FK^2)$, where F is the number of detectors firing in response to the event; batching simultaneous events also requires grouping their input sites. As an illustrative 28×28 implementation with one detector and one modulated site per pixel, $K = 3$ entails 784 detector states and 784 modulatory states, up to nine detector visits per isolated input event and up to nine output visits per firing detector. The specific native-event implementation uses two float64 state arrays and two int64 timestamp arrays (24.5 KiB of raw state/timestamp payload for $M = N = 784$), in addition to neighborhood indices, temporary buffers, and the pretrained neuronal model. These operation counts quantify the added state and local work; they do not measure wall-clock acceleration or energy use.

Presentation duration. As a sensitivity analysis, we varied presentation duration without retuning the astrocytic integration dynamics. The sweep below was run with unit downstream modulation scale $\kappa = 1.0$ and is therefore reported as an ablation rather than as the final operating point used in Fig. 2d.

Table 1: Presentation-duration sensitivity at $\kappa = 1.0$. Values are mean accuracy over four matched seeds.

	Clean ($p = 0$)			Corruption ($p = 0.75$)		
T (ms)	Baseline	Astro	Δ	Baseline	Astro	Δ
250	87.87	79.33	-8.54	9.80	42.23	$+32.43$
300	88.45	83.00	-5.45	9.80	61.01	$+51.21$
350	89.08	85.00	-4.08	9.80	73.06	$+63.26$
400	89.79	86.79	-3.01	9.80	79.18	$+69.38$
450	90.43	87.73	-2.70	9.80	82.31	$+72.51$

The clean accuracy cost decreases with presentation duration, while the high-noise advantage increases, consistent with a robustness–latency trade-off.

Rate-domain local controls. We additionally tested static local preprocessing and reduced astrocytic variants at $T = 450\,\text{ms}$ with $\kappa = 1.0$. These controls act on the effective rate field before Poisson generation and should therefore be interpreted as rate-domain baselines rather than exact event-domain replicas.

Static smoothing alone fails under severe corruption. Hard gating and temporal astrocytic dynamics without spatial pooling each recover substantial accuracy but remain below the full $K = 3$ mechanism.

Table 2: Mechanistic controls at $T = 450\,\mathrm{ms}$ and $\kappa = 1.0$. Values are mean $\pm$ s.d. across four matched seeds.

Method	Clean acc. (%)	$p = 0.75$ acc. (%)	Noisy spikes/image
Box 3×3	89.69 ± 0.48	9.80 ± 0.00	144.93 ± 75.77
Gaussian 3×3	90.54 ± 0.31	9.80 ± 0.00	139.55 ± 50.74
Hard local gate	85.90 ± 0.30	75.61 ± 0.68	20.33 ± 1.33
Astrocyte, $K = 1$	87.46 ± 0.26	74.45 ± 0.65	27.63 ± 0.21
Astrocyte, $K = 3$	87.73 ± 0.19	82.31 ± 0.59	29.15 ± 0.19

Sensitivity to the downstream modulation scale. To verify that the main result is not tied to a single hand-selected modulation strength, we repeated the nominal $T = 350\,\mathrm{ms}$ experiment for three values of κ while keeping $\tau_g = 105\,\mathrm{ms}$, $\theta_g = 3.35$, $A_g = 1$, $\tau_\nu = 200\,\mathrm{ms}$, and $A_\nu = 1$ fixed. The detector threshold and therefore λ_c are unchanged across this sweep.

Table 3: Sensitivity to downstream modulation strength in the current model. Values are mean $\pm$ s.d. over four seeds; the unmodulated baseline has no κ parameter.

κ	Clean acc. (%)	Clean spikes/image	$p = 0.75$ acc. (%)	$p = 0.75$ spikes/image
Baseline	89.08 ± 0.57	32.80 ± 0.35	9.80 ± 0.00	77.99 ± 21.22
1.0	85.00 ± 0.36	23.18 ± 0.03	73.06 ± 0.17	16.11 ± 0.08
1.25	86.33 ± 0.35	33.06 ± 0.12	78.06 ± 0.30	22.46 ± 0.06
1.3	86.74 ± 0.27	35.20 ± 0.02	78.61 ± 0.56	23.89 ± 0.11
1.7	87.46 ± 0.25	48.74 ± 0.16	80.91 ± 0.45	34.05 ± 0.15

The trend is smooth rather than sharply tuned: stronger modulation improves accuracy in both clean and corrupted conditions while increasing network activity. At $\kappa = 1.3$, clean input produces 35.2 versus 32.8 spikes/image for the unmodified network. Fig. 2d uses $\kappa = 1.25$; its clean and $p = 0.75$ results are included above. The remaining values show sensitivity to modulation strength.

Modulation-pathway ablation. The matched final-configuration ablation separately enables synaptic transmission modulation and plasticity-rate modulation, yielding transmission-only, plasticity-only, both-pathways, and neither-pathway conditions. The same training and test protocol and matched seeds are used for each. Table 4 reports means and standard deviations across four seeds. At $p = 0.75$, either pathway alone strongly outperforms the unmodulated baseline, and the combined model improves on transmission-only and plasticity-only models by 4.89 and 6.68 percentage points, respectively. The summary contains no per-seed paired differences, so these contrasts are descriptive rather than paired significance tests.

Table 4: Matched modulation-pathway ablation at $T = 350\,\mathrm{ms}$ and $\kappa = 1.25$. Mean $\pm$ standard deviation over four seeds; spikes/image are measured at $p = 0.75$.

Condition	Clean acc. (%)	$p = 0.75$ acc. (%)	Noisy spikes/image
Neither pathway	89.08 ± 0.57	9.80 ± 0.00	77.99 ± 21.22
Input only	86.84 ± 0.23	73.17 ± 0.38	26.75 ± 0.67
Learning rate only	85.66 ± 0.79	71.38 ± 0.93	42.45 ± 1.33
Both pathways	86.33 ± 0.35	78.06 ± 0.30	22.46 ± 0.06

B Structured native-event noise and an active-channel astrocyte

The original N-MNIST detector integrates the total local spike count g_a as in Eq. 1. To distinguish many spikes from one input channel from activity distributed across channels, we also evaluate a local active-channel detector. A channel $c = (x, y, p)$ denotes one pixel and one polarity, with leaky

Table 5: Paired N-MNIST accuracy under structured injected events on the pilot-excluded set ($n = 9880$). One clean row is shared across corruption types. Values are percentages. Both astrocyte parameter sets are fixed across rows; the 300-recording parameter-selection pilot is excluded from all reported accuracies.

Corruption	q	n	Baseline	Minimal	Active-channel
Clean	0	9880	96.39	96.02	93.31
Temporal burst	5	9880	92.66	91.80	87.63
Temporal burst	10	9880	92.54	90.80	85.55
Temporal burst	20	9880	91.83	89.96	83.58
Hot pixel	5	9880	82.62	77.51	90.45
Hot pixel	10	9880	82.72	77.45	90.40
Hot pixel	20	9880	83.01	77.82	90.29

glutamate proxy h_c:

$$\tau_h \dot{h}_c = -h_c \quad \text{between input events}, \qquad h_c(t^+) = h_c(t^-)+1 \quad \text{at a spike on } c, \qquad b_c(t) = \mathbf{1}\{h_c(t) > \theta_h\}. \tag{14}$$

For each astrocyte a with spatial neighborhood $\mathcal{N}_a$, let $\mathcal{C}_a$ contain both polarities of every included pixel and define

$$n_a(t) = \sum_{c \in \mathcal{C}_a} b_c(t), \qquad H_a(t) = \mathbf{1}\{n_a(t) \geq \max(n_{\min}, \lceil f_\theta |\mathcal{C}_a| \rceil)\}. \tag{15}$$

An astrocytic event occurs only on an upward crossing $H_a : 0 \to 1$; the detector rearms when enough flags decay for H_a to return to zero. Each event adds A_ν to every ν_i in its overlapping output neighborhood, which then decays with τ_ν as in Eq. 3. For each sensor event, the input to the frozen network is multiplied by $\kappa \min(1, \max(0, \nu_i(t^-)))$. Same-timestamp events all read the pre-update state. Input-channel threshold crossings and their analytic expiry times are processed in event time, even when no additional input spike occurs at the expiry. The state resets for each recording. This active-channel variant adds one leaky glutamate proxy and one active flag per polarity channel plus local active counts; it is an engineering extension of the computational interface, not a model of a regenerative calcium event.

We recompute identical eligible N-MNIST recordings and realized injected events for all three branches. The original minimal detector retains its frozen reference parameters ($K = 3$, $\tau_g = 4\,\text{ms}$, $\theta_g = 0.163155$, $\tau_\nu = 3\,\text{ms}$, $A_\nu = 0.1$, $\kappa = 1$). The active-channel detector uses $K = 3$, $\tau_h = 7\,\text{ms}$, $\theta_h = 0.75$, $f_\theta = 0.304$, $n_{\min} = 3$, $\tau_\nu = 3.755\,\text{ms}$, $A_\nu = 0.376$, and $\kappa = 2.3$. At full 3×3 neighborhoods Eq. 15 requires at least six distinct active polarity channels. These parameters were chosen after examining a 300-recording pilot. That pilot subset (selected with seed 67890 from the same eligible test indices) is excluded from reported accuracies; the full run still processes all eligible recordings. This held-out comparison does not undo the post-hoc choice of detector class.

For each sample with N_{clean} original events, we draw the injected count from $\text{Poisson}(qN_{\text{clean}})$. Temporal-burst noise uses random spatial locations and polarities but clusters event times around three sampled centers; the total nominal burst-window fraction is 0.08 of the recording. Hot-pixel noise distributes event times over the recording and allocates events among three randomly selected spatial sites, with random polarities. Clean input ($q = 0$) is measured once. The exclusion list, seed, pretrained weights, and frozen event-to-network mapping are shared across branches. Table 5 reports percentage accuracy on the pilot-excluded recordings; sample-wise predictions permit paired wins/losses and exact paired tests. No pilot accuracy is substituted for a full-run entry.

The structured-noise result is statistic-specific rather than uniformly beneficial. Under temporal bursts, the active-channel detector loses 5.03–8.25 percentage points relative to the unmodulated baseline and 4.17–6.38 points relative to the original minimal detector. Under hot-pixel corruption, the ordering reverses: the active-channel detector gains 7.29–7.82 points over baseline and 12.47–12.96 points over the minimal detector. These paired contrasts are highly significant by exact paired tests at every structured-noise setting ($p < 10^{-5}$; the largest reported paired p-value is 3.97×10^{-6}).

The local weights expose the corresponding mechanism. For hot-pixel noise, the original minimal detector assigns nearly saturated mean weights to injected events (0.9997–0.9999), whereas the active-channel detector assigns them mean weights of only 0.080–0.081 while retaining mean clean

event weights around 0.555. Thus repeated firing from a few spatial sites looks like strong evidence to total spike-count integration but is rejected by a detector that requires recent activity across distinct channels. The same change in local statistic is detrimental for temporal bursts, which is why we interpret this extension as a targeted inductive bias rather than a generally stronger denoiser.

C BOUNDARY OF THE BIOLOGICAL COMPARISON

As a separate mechanistic illustration, we compared one minimal local unit with an isolated two-pool calcium subsystem of Cresswell-Clay et al. (2018), derived from Postnov et al. (2007). The reference uses the cytosolic and endoplasmic-reticulum calcium states and the leak term in the authors' released implementation; it does not include the complete tripartite synapse, the soma–process diffusion network, or neuronal feedback. The same spike realization drives both units on nonoverlapping 4×4 territories of a single stationary MNIST image. The local unit uses $\tau_g = 250\,\mathrm{ms}$, $\theta_g = 0.75$, $\tau_\nu = 1750\,\mathrm{ms}$, and $A_\nu = 0.00025$; these illustrative parameters are distinct from the trained Diehl–Cook experiment.

Under a $4\,\mathrm{s}$ presentation, we compare the mean final-$500\,\mathrm{ms}$ territory maps for ν and for ΔC_I above the unstimulated calcium trajectory. The single territory followed for $40\,\mathrm{s}$ is chosen by maximum input-image intensity *before* running either model; the shorter stimulus is a prefix of the longer realization. The minimal state shows a graded response to local input. In an illustrative sustained-input run, the isolated calcium core crosses the rising $C_I = 1.2$ event threshold and produces a regenerative event that the minimal state does not reproduce. Biophysical calcium models can also show repeated events; that regime is outside this comparison. The event-rate to calcium coupling and conversion of model time units to milliseconds are experimental choices. Similarity between the short-stimulus maps, if present, identifies only a limited functional resemblance, not biophysical equivalence.

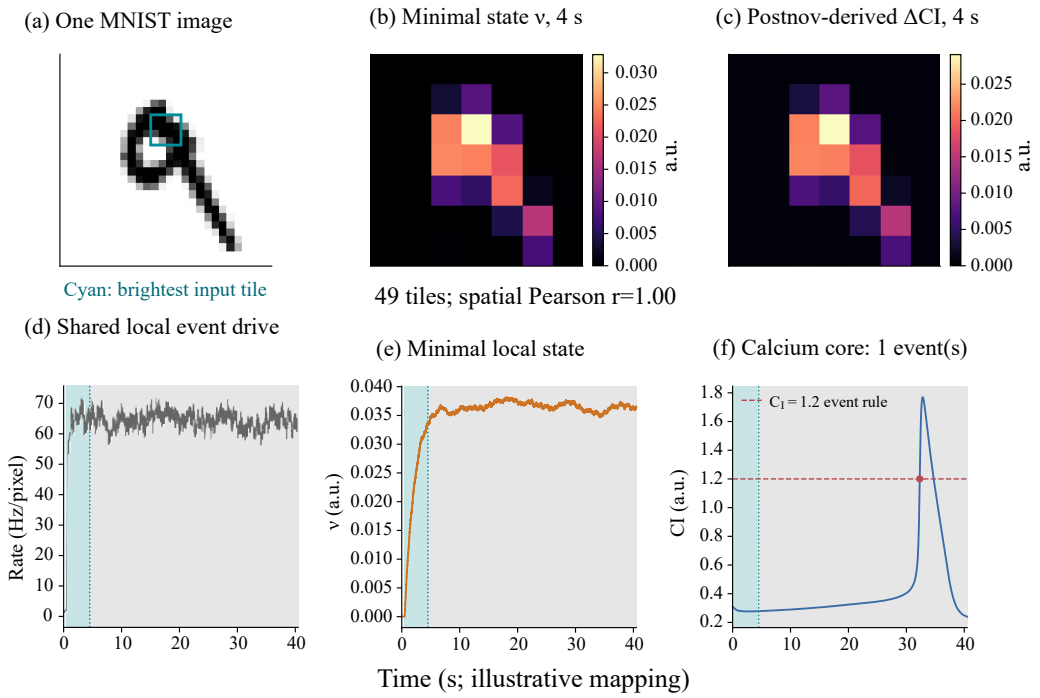


Figure 4: Boundary of a minimal local gate compared with an isolated Postnov-derived two-pool calcium core, driven by identical realized events from one MNIST image at $q = 0$. (a–c) Input image and the two short-input maps from disjoint 4×4 domains; map color scales are separate because the states have different units. (d–f) Shared input rate and model responses in the brightest input territory under prolonged stimulation. The dotted line marks the $4\,\mathrm{s}$ presentation endpoint; the dashed calcium line marks the rising event threshold $C_I = 1.2$. The calcium core can generate a regenerative event that the minimal modulatory state cannot represent. Time conversion and event-rate-to-calcium coupling are illustrative.