Limit-Cycle Proliferation Under Parametric Delayed Feedback in a Conductance-Based Neuron: Bifurcation Landscape, Orbit Catalog, and Capacity Analysis

Mohammad O. Alhawarat *, Ayman J. Alnsour, Mohammed A. F. Al-Husainy, Khalil M. Abdelnaby

Posted Date: 8 May 2026

doi: 10.20944/preprints202605.0551.v1

Keywords: Hodgkin–Huxley model; parametric delayed feedback; limit-cycle multiplicity; orbit-coded memory; bifurcation landscape; autaptic feedback; neuromorphic computing

 doi:10.20944/preprints202605.0551.v1

Article

Limit-Cycle Proliferation Under Parametric Delayed Feedback in a Conductance-Based Neuron: Bifurcation Landscape, Orbit Catalog, and Capacity Analysis

Mohammad O. Alhawarat [1,*], Ayman J. Alnsour [1], Mohammed A. F. Al-Husainy [1], and Khalil M. Abdelnaby [1,2]

[1] Faculty of Information Technology, Al-Ahliyya Amman University, Amman 19328, Jordan
[2] Systems and Computers Engineering Department, Faculty of Engineering, Al-Azhar University, Nasr City 11765, Cairo, Egypt.
* Correspondence: m.hawarat@ammanu.edu.jo

Abstract

We show that a single Hodgkin–Huxley (HH) neuron with Pyragas-type delayed feedback control (DFC) can store multiple symbols as stable periodic orbits, where the specific orbit is selected by tuning the DFC gain K and time delay τ. Sweeping the (K, τ) parameter plane at fixed bias current $I_{\text{bias}} = 10.0\ \mu\text{A/cm}^2$ reveals 207 orbit types across 12 topological categories, with inter-spike interval (ISI) means from 5.9 to 56.9 ms. We establish: (i) a write protocol that reliably locks orbits with 13.9 ms median settling time; (ii) a novel Pattern-Oriented Limit-cycle Decoder (POLD) that reads orbits at 100% accuracy from only 5 observed ISIs; (iii) a full write–read–erase (W–R–E) cycle with 100% read accuracy, 92% erase verification, and no decay over hold durations up to 50 s; and (iv) a fully validated 12-symbol memory capacity, with a read-discriminable upper bound of 67 symbols (11.2× over rate coding) pending write-viability confirmation for the extended set. Reliable orbit addressing needs delay precision of ±2%, which constitutes a write-precision specification and not a fundamental capacity limit. These findings show that parametric delayed feedback is a viable mechanism for limit-cycle-based information storage in conductance-based spiking neurons. The biological interpretation is analogical, not direct: the ±2% delay-precision requirement exceeds what has been demonstrated for biological autaptic variability, and the orbit-coded memory framing is best understood as a computational proof-of-principle aimed at neuromorphic engineering, not as a claim about biological working memory.

Keywords: Hodgkin–Huxley model; parametric delayed feedback; limit-cycle multiplicity; orbit-coded memory; bifurcation landscape; autaptic feedback; neuromorphic computing

1. Introduction

A central question in computational neuroscience concerns the computational repertoire of individual neurons. Classically, neurons are assigned a straightforward input–output role: they integrate synaptic currents and report their mean firing rate as an output signal. This rate-coding hypothesis [1,2] has been enormously productive, but accumulating evidence suggests that single neurons can support far richer computations through their intrinsic dynamics. Precise spike timing carries information beyond what is captured by firing rate [3,4]; individual dendrites perform local, nonlinear computations that effectively make a single cell a multi-layer processing unit [5]; and intrinsic bistability or multistability can give individual neurons the ability to sustain distinct activity states in the absence of ongoing input [6,7]. The present work explores an extreme case of this last capacity: using the dynamical richness of a single Hodgkin–Huxley neuron's limit-cycle repertoire for multi-symbol information storage.

 doi:10.20944/preprints202605.0551.v1

The working memory problem—how the brain maintains information over seconds without external reinforcement—has traditionally been modelled at the network level. Attractor network models, originating with Hopfield [8] and extended to spiking networks by Amit and Brunel [9] and Wang [10], posit that patterns of activity are held as stable fixed points of recurrent synaptic dynamics. Recent extensions incorporate activity-dependent plasticity and dynamic memory representations [11,12]. These rely on structured connectivity and balance between excitation and inhibition; their capability depends on the size of the network. The single neuron bistability models [6,7,13] offer a complementary approach to maintain bistable regimes in the absence of network structure through the dynamics of intrinsic conductance parameters. This involves models involving persistent Na^+ conductance [14], K^+ slow conductances [15], and facilitated synapses due to astrocytes [16]. Our strategy stands out in that it takes advantage of the inherent multistability property of a delay differential equation (DDE) model where the neuron's previous output serves as feedback to create a high-dimensional effective state space.

Time-lagged self-feedback occurs physiologically in neurons due to autaptic synapses (synapses created between an axon terminal from the same neuron and the dendrites of that neuron). Initially discovered anatomically by van der Loos and Glaser [17], autapses have been observed all across the neocortex [18] and hippocampal circuits [19] to date. Autapses contribute to neurobiological phenomena such as modulation of firing precision [20], gain control [21], oscillations generation [22], and coherence resonance in autaptically regulated networks [23–25]. The delays in such phenomena—ranging from 1 to 20 milliseconds in direct autapses, and beyond for reverberant microcircuit loops, are right inside the parameter regime used in our study. Crucially, both excitatory and inhibitory autapses have been documented. The current work limits $K \geq 0$ (excitatory self-feedback); inhibitory autapses have been shown capable of providing effective sign change in feedback due to circuit polarity inversion, but their exploration is deferred to future work.

Pyragas [26] introduced time-delayed feedback control (DFC) as a method for stabilising unstable periodic orbits (UPOs) embedded in chaotic attractors. The control signal $I_{\text{ctrl}}(t) = K[V(t-\tau) - V(t)]$ vanishes exactly when $V(t)$ is periodic with period τ, so the target orbit becomes a fixed point of the controlled dynamics while remaining physically equivalent to the uncontrolled orbit. Early applications focused on optics [27] and semiconductor lasers [28]. He et al. [29] verified selective UPO stabilization in a chaotic neural network; lately, DFC tested in vitro on biological neuronal populations for desynchronization of pathological oscillations [30]. The foundational concept of orbit-coded memory was introduced by Crook, Goh, and Hawarat [31,32] using the Nonlinear Dynamic State (NDS) neuron—a Rössler-based discrete-time map. In the NDS approach, the system operates at a single fixed parameter point within a chaotic attractor; DFC is applied as an active control perturbation that continuously forces the trajectory to stay near one specific UPO embedded in that attractor. Memory is encoded as a choice of which UPO to stabilize, and the control signal must remain active throughout memory maintenance—removing it allows the trajectory to escape back into chaos. Subsequent analyzes [33–37] characterized the NDS model's dynamics and memory properties in detail.

The present work is methodologically distinct from the NDS approach in every essential respect. We do not operate within a chaotic attractor, we do not stabilize UPOs, and we do not require any ongoing adjustment of control parameters during memory maintenance. Instead, we exploit the fact that changing (K, τ) moves the HH–DFC system to a different region of its bifurcation landscape, where a different stable periodic orbit is the natural attractor. Once the neuron locks onto this orbit, the DFC term $K[V(t-\tau) - V(t)]$ vanishes exactly for orbits whose period matches τ; for higher-order burst orbits whose fundamental period does not equal τ, the DFC term participates actively in the dynamics (see Section 5.1). In both cases, the memory is preserved without having to stabilize it actively. The write process involves switching the parameters and does not involve orbit stabilization. Therefore, the memory configuration is that of a stable attractor and not that of an unstable controlled

 doi:10.20944/preprints202605.0551.v1

attractor. There is more to this difference than meets the eye since it affects the system's capacity, robustness, and maintenance costs.

The challenge of extending orbit-coded memory into biophysically realistic conductance-based neurons remains open. This is fundamental: the Hodgkin–Huxley model chaotic regime, defined by Guckenheimer and Oliva [38], allows only 2–4 ISI distinctive clusters—far fewer than around 15 orbits needed for useful memory capacity. Aihara, Matsumoto, and Ikegaya [39] demonstrated that periodic and chaotic responses coexist in the forced HH oscillator, but the diversity of periodic states reachable via chaos control alone is limited by the spiking mechanism, which constrains the action potential waveform. A key insight motivating our work is that this scarcity pertains specifically to the chaotic regime; the stable periodic regime of the HH–DFC system is, as we shall show, vastly richer.

Why such richness? Yanchuk and colleagues [40,41] showed that delay-coupled neural oscillators support dense families of coexisting stable periodic orbits, born through Hopf bifurcations as the delay grows. Kantner, Yanchuk, and Schöll [42] further showed that coupling delay greatly increases both the number and stability of periodic solutions in neural networks. More recently, photonic neuron models with delayed self-feedback [43,44] showed that such systems can hold multiple coexisting temporal localized states, each one a different memory symbol, and that brief perturbations can switch between them. Those photonic implementations used simplified FitzHugh–Nagumo dynamics and stored binary patterns. Here we show analogous functionality in the full Hodgkin–Huxley conductance model, reaching a read-discriminable capacity of 67 symbols with a complete W–R–E protocol and systematic capacity characterisation.

What sets our approach apart from earlier DFC-based neural memory is one key point: instead of stabilizing UPOs inside a fixed chaotic attractor, we use (K, τ) as an address that moves the HH–DFC system to a different part of its bifurcation landscape, where a different stable periodic attractor lives. This shift—from chaos-control memory to bifurcation-directed memory—sidesteps the topological constraints of the chaotic regime and opens up hundreds of distinct orbit types in 12 qualitative categories. The mechanism is also conceptually related to reservoir computing [45,46] in that the neuron's internal dynamics (here, a delay differential equation) generate a rich representational space; however, unlike reservoir computing, our system stores information in a stable attractor rather than in a transient trajectory, enabling retention without ongoing input.

Note on terminology: Crook et al. [31,32] coined orbit-coded memory in the NDS context. We keep the term because the idea behind encoding information in distinct periodic orbits is the same—although our mechanism is fundamentally different (Section 5.1). Memory is the computational property of persistent, switchable, distinguishable dynamical states; it is not biological memory in the cognitive or synaptic sense used herein. This work is motivated by biological reasons rather than mechanistic claims; the autaptic analogy (Section 5.3) is computationally motivated towards neuromorphic engineering applications.

The paper makes the following contributions. (C1) We catalog 207 distinct stable periodic orbit types in the HH–DFC system from a dense 10,100-point sweep of the (K, τ) plane. (C2) We design and validate a complete W–R–E memory cycle achieving 100% read accuracy and 92% erase verification across 12 representative orbits spanning all topological categories, with no decay observed over hold durations up to 50 s. (C3) We introduce the Pattern-Oriented Limit-cycle Decoder (POLD), a lightweight classifier that achieves 100% read accuracy using only 5 observed ISIs, with no training data required. (C4) We establish a read-discriminable capacity of 67 symbols via a full pairwise confusion matrix and greedy subset selection; full write–read–erase viability is confirmed for a conservative 12-symbol library. (C5) We demonstrate an 11.2× read-discriminable capacity advantage over rate coding ($K = 0$, variable I_{bias}) in the same neuron. (C6) We provide a fully reproducible, gate-based experimental pipeline (PS0–PS5) with formal pass/fail criteria at each phase.

The rest of this paper is arranged as follows: section 2 defines the mathematical framework including HH-DFC, orbit library formalism and POLD classifier. Section 3 describes the six-phase

 doi:10.20944/preprints202605.0551.v1

experimental pipeline. In section 4 we show results, whereas section 5 reports on implications, limitations, and future directions. Finally, section 6 concludes the paper.

2. Mathematical Framework

2.1. Hodgkin–Huxley Neuron with Delayed Feedback Control

Consider the standard Hodgkin–Huxley equations [47] augmented with a Pyragas-type DFC term [26]:

$$C_m \frac{dV}{dt} = I_{\text{bias}} + K\big[V(t-\tau) - V(t)\big] - \bar{g}_{\text{Na}}\, m^3 h (V - E_{\text{Na}}) - \bar{g}_{\text{K}}\, n^4 (V - E_{\text{K}}) - \bar{g}_{\text{L}} (V - E_{\text{L}}) \tag{1}$$

where $V(t)$ is the membrane potential (mV), and $m, h, n \in [0, 1]$ are the sodium activation, sodium inactivation, and potassium activation gating variables, each satisfying:

$$\frac{dx}{dt} = \alpha_x(V)(1 - x) - \beta_x(V)\, x \tag{2}$$

where standard Hodgkin–Huxley rate functions are given in Appendix A. The control parameters are: $K \geq 0$ (DFC gain, mS/cm^2), $\tau > 0$ (time delay, ms), and I_{bias} (constant injected current, μA/cm^2). We use the V-shifted convention with resting potential $V_{\text{rest}} = 0$ mV and standard squid giant axon parameters: $C_m = 1$ μF/cm^2; $\bar{g}_{\text{Na}} = 120$, $\bar{g}_{\text{K}} = 36$, $\bar{g}_{\text{L}} = 0.3$ mS/cm^2; $E_{\text{Na}} = 115$, $E_{\text{K}} = -12$, $E_{\text{L}} = 10.6$ mV. All simulations use $I_{\text{bias}} = 10.0$ μA/cm^2 unless stated otherwise.

The HH–DFC system is a DDE of dimension $4 + \infty$ (the finite-dimensional state plus the delay history). The DFC term adds a perturbation to the membrane current that is proportional to the difference between the current voltage and the voltage one delay period ago. When $V(t)$ is periodic with period τ, this term vanishes exactly, which is the key non-invasive property of Pyragas control [26, 48]; for orbits whose period differs from τ, the DFC term remains dynamically active (see Section 5.1). Importantly, the DFC term operates in two distinct regimes: (i) for orbits whose fundamental period matches τ, the term vanishes exactly (the classical Pyragas non-invasive property), with RMS control current below 2% of the bias current; (ii) for higher-order burst orbits whose period differs from τ, the DFC term contributes substantially (RMS control current 15–140% of bias) and functions as an integral shaping force in the modified dynamical landscape rather than a small corrective perturbation. In both regimes the target orbit is a stable attractor of the full HH–DFC system; quantitative characterization is presented in Section 4.1 and interpreted in Section 5.1. For $K = 0$, Eq. (1) reduces to the standard uncontrolled HH model. Floquet multipliers for the full orbit library are not computed here; stability is established empirically via sustained periodic behavior over extended hold durations (up to 50 s). As a representative validation, Floquet multipliers were computed for four orbits spanning the main topological categories; results are reported in Table 2 (Section 4.1).

Stability convention. Throughout this paper, 'stable' refers to empirically established stability: sustained periodic behavior under the tested simulation conditions (RK4, $dt = 0.01$ ms, hold durations up to 50 s) without decay or regime change. This empirical criterion is applied uniformly across all 207 orbit types. Floquet multiplier computations for four representative orbits (Table 2, Section 4.1) provide partial corroboration of stability but are not the basis of our stability claims, given the computational challenges of Floquet analysis for infinite-dimensional DDE systems; the empirical criterion (sustained periodic behavior over 50 s without decay or regime change) is applied uniformly to all 207 orbit types and constitutes the primary stability validation. All subsequent references to orbit stability carry this qualification unless otherwise noted.

We integrate numerically with fourth-order Runge–Kutta at $dt = 0.01$ ms, which is fine for reproducing spike waveforms given the 0.5–2.0 ms spike width of HH action potentials. The delay buffer is implemented as a circular array of length $\lceil \tau/dt \rceil$ and warm-started using the RC3 protocol: the neuron free-runs with $K = 0$ for the first τ ms to fill the buffer with physiological voltage history, eliminating artificial transients from zero-initialized buffers.

 doi:10.20944/preprints202605.0551.v1

2.2. Stability and the Orbit Landscape

A periodic orbit $\mathcal{O}$ of period T has Floquet multipliers μ_i—the eigenvalues of the monodromy matrix $M = \Phi(T, 0)$, with Φ being the fundamental solution matrix of the variational equation [49]. In the DDE (1), M is an infinite-dimensional operator in the delay history space. Its spectrum contains one trivial multiplier, $\mu = 1$ (phase shift along the orbit), and countably many stability-determining multipliers. An orbit is asymptotically stable (an attractor) if all non-trivial Floquet multipliers fulfil $|\mu_i| < 1$. We computed Floquet multipliers for four representative orbits (as a check on the empirical stability criterion used in this study,) covering the major topological categories (tonic, doublet, triplet, burst_p6) via variational QR-iteration; results appear in Section 4.1 (Table 2).

There is a direct theoretical reason for the richness of the HH–DFC orbit landscape. For a scalar DDE with delay τ, the number of coexisting stable periodic orbits grows approximately as τ/T_0, where T_0 is the characteristic oscillation period [40]. For the HH neuron at $I_{\text{bias}} = 10.0\ \mu\text{A/cm}^2$, the natural ISI is approximately $T_0 \approx 10.6$ ms (firing rate ~ 94 Hz). For $\tau = 100$ ms, we therefore expect $\mathcal{O}(10)$ coexisting periodic solutions; and our numerical sweep does find hundreds across the full (K, τ) plane. As τ increases, new periodic orbits are born through Hopf bifurcations at which a pair of complex conjugate Floquet multipliers crosses the unit circle. Yanchuk and Perlikowski [41] showed that these bifurcations accumulate densely as $\tau \to \infty$, producing an asymptotically unbounded number of coexisting stable orbits—which is exactly what we exploit to build a large orbit catalog.

2.3. Orbit Library and Memory Formalism

Definition 1 (Orbit type)**.** *An orbit type $\mathcal{O}_i$ is an equivalence class of stable periodic trajectories of Eq. (1) characterized by a canonical ISI pattern: the ordered sequence of inter-spike intervals within one fundamental period. Two trajectories belong to the same orbit type if their ISI pattern vectors are within 2 ms Euclidean distance in ISI-mean space. This threshold is motivated by the minimum ISI jitter arising from intrinsic conductance noise and numerical discretization ($\sigma_{ISI} < 0.05$ ms for all stable orbits observed).*

Definition 2 (Orbit categories)**.** *Orbits are assigned to qualitative categories based on their pattern length p (number of distinct ISIs per period): tonic ($p = 1$), doublet ($p = 2$), triplet ($p = 3$), and burst_pν ($p = \nu$) for $\nu = 4, \ldots, 12$. Within each category, orbits differ in their mean ISI and in the fine structure of their ISI pattern vector.*

Definition 3 (Libraries)**.** *The catalog library $\mathcal{L}_{cat}$ is the set of all distinct orbit types identified across the full (K, τ) parameter sweep. The addressable library $\mathcal{L}_{addr} \subseteq \mathcal{L}_{cat}$ is the subset whose members satisfy write viability (lock rate $\geq 80\%$) and read viability (pairwise classification accuracy $\geq 90\%$). The number of simultaneously unique symbols is considered as the memory capacity ($N^* = |\mathcal{L}_{addr}|$). A read-discriminable subset $\mathcal{L}_{read} \subseteq \mathcal{L}_{cat}$ satisfies the read viability criterion only; its cardinality $N^*_{read} = |\mathcal{L}_{read}|$ is the read-discriminable capacity, which serves as an upper bound on N^* that becomes exact once write viability is confirmed for all members.*

HH-DFC neurons are thought of as devices for storing memories that can perform three operations:

1. **Write:** Instantaneous switch of DFC parameters from a baseline state ($K = 0$, tonic firing) to target values (K_i, τ_i) which cause the neuron to move to orbit $\mathcal{O}_i$. Locking is achieved when the average of five consecutive ISIs falls within 5% of the target ISI mean and when the CV of these five ISIs is less than 0.05. Settling time is then measured from the parameter switch to the first spike of the locking window.
2. **Read:** Recording the spike train during the read window and classifying the orbit using POLD based on the ISI fingerprint.
3. **Erase:** Changing the DFC parameters back to $K = 0$, resulting in tonic firing. Validation of erasure is done by applying POLD using five ISIs in the post-reset window, where erasure is successful when the classifier gives an orbit that differs from the written symbol.

 doi:10.20944/preprints202605.0551.v1

2.4. ISI-Based Orbit Classification: POLD

The Pattern-Oriented Limit-cycle Decoder (POLD) classifies the neuron's current orbit from an observed sequence of n ISIs, denoted $\vec{\iota} = (\iota_1, \ldots, \iota_n)$. For each candidate orbit template $\mathcal{O}_j$ in the library (with mean μ_j, standard deviation σ_j, and canonical pattern vector $\vec{\rho}_j$ of length p_j), POLD computes:

$$S_j = 0.6 \cdot S_{\text{mean}}(j) + 0.4 \cdot S_{\text{pattern}}(j) \tag{3}$$

where the mean score is a Gaussian kernel on the normalized ISI mean distance:

$$S_{\text{mean}}(j) = \exp(-0.5\, z_j^2), \qquad z_j = \frac{|\mu_{\text{obs}} - \mu_j|}{\max(\sigma_j,\ 0.5)} \tag{4}$$

where $\mu_{\text{obs}} = n^{-1} \sum_i \iota_i$ is the observed mean ISI over the read window. The pattern score is:

$$S_{\text{pattern}}(j) = \max_k \Psi(\vec{\iota}[k : k + p_j],\ \vec{\rho}_j) \tag{5}$$

where $\vec{\iota}_{\text{seg}} = \vec{\iota}[k : k + p_j]$ denotes a length-p_j segment of the observed ISI sequence starting at window index k, $\text{corr}(\cdot, \cdot)$ is the Pearson correlation coefficient (with negative values floored at zero), and $\text{rel\_err} = (1/p_j) \sum_i |\iota_{\text{seg},i} - \rho_{j,i}| / (\rho_{j,i} + 0.1)$ is the mean element-wise relative error between the observed segment and the template pattern, with a 0.1 ms floor in the denominator to prevent division by zero. The window score is:

$$\Psi(\vec{\iota}_{\text{seg}},\ \vec{\rho}_j) = 0.5 \cdot \text{corr}(\vec{\iota}_{\text{seg}},\ \vec{\rho}_j) + 0.5 \cdot \exp(-2 \cdot \text{rel\_err}) \tag{6}$$

combines correlation strength with a relative error penalty. The prediction is $\hat{\imath} = \arg\max_j S_j$. POLD requires no training data: templates are built from the mean ISI and mean pattern vector of 5 long simulation runs per orbit. The classifier operates in $\mathcal{O}(nL)$ time per observation where L is the library size, making it suitable for real-time implementation. This ratio of 60/40 between mean-score and pattern-score was selected empirically. Sensitivity analysis has shown that even in weight allocations between 50/50 to 70/30, the system maintains its performance. The overall methodology of categorizing biological oscillatory systems based on their inter-event interval statistics has proved successful in other physiological signal domains, for example, extracting features from heart rate variability data [50].

3. Experimental Pipeline

The study follows a six-phase gated pipeline (PS0–PS5) in which each phase has formal pass/fail criteria (gates) that must be satisfied before proceeding. All simulations use $I_{\text{bias}} = 10.0\ \mu\text{A/cm}^2$ and $dt = 0.01$ ms, and transients of 500–1000 ms are discarded before measurement. Code is implemented in Python with Numba JIT compilation [51] and executed on Google Colab (Intel Xeon 2.3 GHz, 13 GB RAM). All simulations used a fixed random seed (seed = 42) for reproducibility; results were verified to be invariant across seeds 1–10. Jupyter notebooks (PS0–PS5) are publicly available (see Data Availability). The gate-based validation structure follows the spirit of gated modelling pipelines used in nonlinear dynamical system studies, where pass/fail criteria at each stage ensure that subsequent analyzes rest on confirmed dynamical properties.

3.1. Gate Threshold Rationale

All gate thresholds were fixed a priori, before any simulation was executed, on the following principled bases.

ISI separability threshold (2 ms; PS-G0, PS-G5). We set the minimum pairwise ISI separation at 2 ms based on two considerations: (a) at $dt = 0.01$ ms, the numerical noise floor of the HH model makes two orbit types hard to tell apart statistically when their ISIs differ by less than this, unless the

 doi:10.20944/preprints202605.0551.v1

observation window is made much longer; and (b) ISI-based spike-train discrimination in cortical neurons works at a similar temporal resolution [4].

Accuracy thresholds ($\geq$ 90%; PS-G1a, PS-G2b, PS-G3a–G3e, PS-G5). The 90% criterion is an operational threshold for the present study, informed by the neural decoding literature where this level is commonly used as a practical performance floor [52,53], though no universal standard exists. At this level, the per-symbol error rate (10%) is low enough that a k-symbol sequential read sequence yields $> 72\%$ end-to-end success at $k = 3$, the minimum multi-symbol operation. Quian Quiroga and Panzeri [52] survey neural population decoding studies and identify 90% as the practical performance floor distinguishing reliable from unreliable decoders; the same threshold appears in single-unit classification benchmarks reviewed by Dayan and Abbott [53].

The stricter 95% threshold for clean (zero-noise) reads, in terms of G2a, represents the high standard that is expected when there is no perturbation since a decoding error of one in twenty on a clean sequence cannot be blamed on noise but rather on a fundamental classifier limitation.

Lock-rate threshold ($\geq$ 80%; PS-G1a). Below 80%, multi-symbol cycling becomes impractical without error correction. So the 80% floor is essentially the minimum write reliability needed. This is consistent with the write-reliability criterion used in photonic temporal localized-state memory systems [43,44], where $< 80\%$ switching success is treated as phase-bistability failure.

Minimum capacity ($N^* \geq 6$; PS-G3c). We set 6 symbols as the lower bound for two reasons: (a) operationally, this is the minimum needed to show multi-symbol storage beyond simple bistability or tristability; (b) it coincidentally aligns with the lower bound of Miller's classical working memory estimate of 7 ± 2 items [54], though we do not claim a direct mechanistic connection between attractor counts in a simulated neuron and cognitive working memory capacity.

Retention duration (10 s; PS-G3d). The 10 s hold duration covers the upper end of the 3–18 s working memory maintenance window established by Baddeley [55] and supported by prefrontal persistent-activity recordings [13]. Testing seven durations (0.5, 1, 2, 5, 10, 20, 50 s), we find that retention goes well beyond this window under our simulation conditions.

Settling-time criterion ($<$ 1000 ms; PS-G1c). The upper bound of 1000 ms on median write settling time corresponds to the reaction-time ceiling for voluntary attentional encoding in working memory tasks (typically 200–800 ms [55]). A write operation involving more than one second median settler would be inconsistent with the timescale of fast sequential symbol encoding biological working memory and would suggest that the DFC term has not effectively driven the neuron to the target orbit's target attraction basin. The condition is purposely applied to the median and not the maximum, to allow occasional long transients whilst requiring the majority of write operations to complete within a behaviorally relevant window.

3.2. PS0: Dense Orbit Catalog

The parameter space of (K, τ) was scanned via dense sampling of the space. For this purpose, 101 K values uniformly distributed between 0.0 and 2.0 with a step of 0.02 were paired with 100 τ values distributed logarithmically between 1.0 and 200 ms, producing a grid of 10,100 different sets of parameters. In each of these simulations, the duration was 3000 ms, consisting of 500 ms of transient dynamics and 2500 ms for measurement. Detection of spikes relied on crossing the threshold of $V = 0$ mV, and ISIs were calculated from consecutive spike times.

Each ISI sequence is classified as silent (fewer than 3 spikes), tonic (ISI CV < 0.02), periodic (repeating ISI pattern found by autocorrelation; see Appendix B), chaotic (positive Lyapunov exponent from ISI divergence), or quasi-periodic (irrational frequency ratio). Periodic orbits are further described by their pattern length p from sliding autocorrelation, and orbit types are grouped hierarchically by Euclidean distance on fingerprint vectors (ISI mean, pattern period, pattern length) with a 2 ms linkage threshold. For each cluster, the representative orbit is the point nearest the centroid in ISI-mean space. **Gate PS-G0:** $\geq$ 15 distinct orbit types; $\geq$ 3 qualitative categories; pairwise ISI separability with ΔISI > 2 ms for the majority of pairs.

 doi:10.20944/preprints202605.0551.v1

3.3. PS1: Write Protocol

From the 207 cataloged orbit types, we select a working library using a greedy algorithm that maximizes the minimum pairwise ISI separation while covering all topological categories. For each orbit in the working library, the write protocol proceeds in two stages: (i) a baseline phase (500 ms, $K = 0$) establishing stable tonic firing; (ii) an activation phase (2500 ms) with an instantaneous switch to target parameters (K_i, τ_i). To make sure POLD was not tuned on the same data used for evaluation, we built orbit templates from 5 independent long-duration simulations (2000 ms each) that did not overlap with any test trials. Lock is declared when the mean of 5 consecutive ISIs is within 5% of the target ISI mean and the CV of those 5 ISIs stays below 0.05; if this does not happen within the 2500 ms window, the trial counts as a failure. Settling time is the interval from the parameter switch to the first spike of the locking window.

We build an orbit-to-orbit switching matrix by testing all $N \times N$ directed pairs (origin orbit $\rightarrow$ target orbit), with 5 trials per pair. This captures switching dynamics separately from the rest-to-orbit write operation.

Gate PS-G1: $\geq$ 80% of library orbits achieve lock rate $\geq$ 80% (G1a, minimum write reliability for functional multi-symbol encoding); overall orbit-to-orbit switch lock rate $\geq$ 70% (G1b); median settling time $<$ 1000 ms (G1c, ceiling of attentional encoding timescale [55]); ISI error $<$ 10% (G1d).

3.4. PS2: Read Protocol and Noise Robustness

ISI fingerprint templates are built from 5 clean simulation runs per orbit (each 5000 ms, 1000 ms transient discarded). The POLD classifier (Section 2.4) is calibrated across observation windows of 3–50 ISIs using 50 independent trials per orbit. Noise robustness is tested under three conditions: (i) additive current noise: white Gaussian noise with standard deviation $\sigma \in \{0, 0.5, 1.0, 2.0, 5.0\}$ μA/cm^2 is added to I_{bias}; (ii) DFC gain jitter: K is perturbed by $\pm\delta K/K \in \{0, 2, 5, 10, 20\}$%; (iii) delay jitter: τ is perturbed by $\pm\delta\tau/\tau \in \{0, 2, 5, 10, 20\}$%. For each condition, classification accuracy is measured over 50 trials at a 10-ISI observation window.

Gate PS-G2: Clean accuracy $\geq$ 95% at 10 ISIs (G2a, strict baseline standard for zero-noise condition); noisy accuracy $\geq$ 90% at $\sigma = 0.5$ μA/cm^2 (G2b, standard neural decoding criterion [23,55]); jitter accuracy $\geq$ 85% at $\pm$10% (G2c, aspirational benchmark to identify the precision envelope rather than a hard pass/fail criterion); observation window $\leq$ 20 ISIs for 95% accuracy (G2d).

3.5. PS3: Full Write–Read–Erase Demonstration

The complete memory cycle is demonstrated in four sub-phases.

Phase A (single W–R–E): for each library orbit, execute 500 ms baseline $\rightarrow$ 1000 ms write $\rightarrow$ 500 ms read (10 ISIs classified) $\rightarrow$ 500 ms erase $\rightarrow$ 500 ms verify. This tests one complete memory cycle per orbit with no overlap.

Phase B (full-alphabet independent-symbol readout): The entire library orbit simulation is performed independently using the clean simulation method (transient time of 500 ms + hold time of 2000 ms). These simulations are classified from the last 500 ms of each library orbit simulation and then combined according to the library sequence to form the complete independent symbol readout using the full alphabet. This technique examines distinguishability among the entire alphabet and is consistent with Phase A's single-cycle validated results; the continuous orbit to orbit switching dynamics are studied using the PS1 switching matrix independently.

Phase C (capacity test): for subset sizes $k = 2$ to N, select k orbits by evenly spacing their ISI means across the full range, then write-read each across 20 trials at a 10-ISI window. Report N^* is the maximum k at which $\geq$ 90% accuracy is obtained.

Phase D (retention test): hold each orbit for $[0.5, 1, 2, 5, 10, 20, 50]$ seconds without any refresh or read operations, then classify at the end. This measures whether periodic attractors persist without ongoing input.

 doi:10.20944/preprints202605.0551.v1

Gate PS-G3: W–R–E accuracy $\geq 90\%$ (G3a); full-library independent-symbol readout accuracy $\geq 85\%$ (G3b); capacity $N^* \geq 6$ at $\geq 90\%$ accuracy (G3c, lower bound of biological working memory span [54]); 10 s retention $\geq 90\%$ (G3d, upper range of working memory maintenance window [55]); erase verification $\geq 90\%$ (G3e).

3.6. PS4: Rate Coding Baseline Comparison

A controlled comparison against rate coding isolates the contribution of DFC. The sweep covers $I_{\text{bias}} \in [6.0, 100.0]$ μA/cm^2 in 0.5 μA/cm^2 steps (189 values), with $K = 0$ (no DFC) throughout.

Distinguishable tonic states are chosen based on a 2 ms ISI separation criterion just like the one used for orbit classification. The exact same POLD classifier and capacity testing procedure (Phase C of PS3) is used. Since the classifier and capacity test process are kept consistent, a direct comparison can be made where any capacity advantage can be definitively attributed to DFC and not to classifier design.

Gate PS-G4: N^* for HH–DFC $>$ N^* for HH–Rate.

3.7. PS5: Maximum Capacity from All 207 Orbit Types

A full pairwise confusion matrix is used for testing all 207 orbits. To save on computations, all 207 orbits are simulated only once (5000 ms) with 50 ISI windows of length 10 are cached. Subsequently, all 21,321 pairs are tested using the stored data without performing any additional simulations (Phase B takes less than 6 min vs. nearly 18 hours). Finally, a greedy maximum-subset selection algorithm selects the largest possible subset of mutually discriminable orbits such that the accuracy of classification between all pairs within the subset equals or exceeds 90%, starting with the orbit having the largest number of compatible neighbors and adding other orbits sequentially if they can be distinguished from all the previously selected ones. The capacity curve is measured over the full $k = 2$ to N^*_{max} range.

Gate PS-G5: $N^*_{\text{max}} > N^*_{\text{conservative}}$ (12 symbols from PS3).

4. Results

4.1. PS0: The HH–DFC Orbit Catalog

The scan using 10,100 data points took 18.0 minutes to complete. Table 1 contains a brief overview of dynamic categories. The dominant finding within PS0 can be viewed as a strong prevalence of stable periodic dynamics: out of 10,100 total points, 8,677 points (85.9%) generated stable periodic orbits, whereas chaotic dynamics was found in just 223 points (2.2%). Such a relatively rare occurrence of chaotic dynamics is consistent with the theoretical results by Guckenheimer and Oliva [38] on Hodgkin–Huxley chaos as structurally fragile, and Aihara et al. [39], showing that the periodically forced Hodgkin–Huxley oscillator rapidly develops highly structured periodic responses. The presence of DFC term even further reduces the presence of chaos: for $K > 0.3$, the proportion of chaotic points drops below 0.5%.

Table 1. Classification of 10,100 (K, τ) grid points at $I_{\text{bias}} = 10.0$ μA/cm^2.

Classification	Count	Percentage	Orbit Types	Categories
Periodic (p1–p12)	8,677	85.9%	207	12
Silent	1,198	11.9%	—	—
Chaotic	223	2.2%	—	—
Quasi-periodic	2	0.0%	—	—

Using hierarchical clustering with a 2 ms ISI linkage threshold, we identified 207 distinct orbit types. The 12 qualitative categories cover a broad range of firing patterns: tonic (period-1, ISI mean 7.1–42.6 ms), doublet (period-2, ISI mean 6.9–56.9 ms), triplet (period-3, ISI mean 6.8–21.8 ms), and burst patterns from period-4 through period-12 with ISI means ranging from 5.9 to 25.9 ms. The shortest-mean orbit (burst_p12 at ISI = 5.91 ms) corresponds to a 12-spike-per-period fast bursting

 doi:10.20944/preprints202605.0551.v1

pattern at $K = 1.88$, $\tau = 4.98$ ms; the longest (doublet at ISI = 56.87 ms) is a slow alternating pattern at $K = 0.06$, $\tau = 8.51$ ms. Of the 21,321 distinct orbit-type pairs, 16,457 (77.2%) have pairwise ISI separation > 2 ms. Gate PS-G0: PASS (207 types $\gg 15$; 12 categories $\gg 3$). The full bifurcation map is shown in Figure 1.

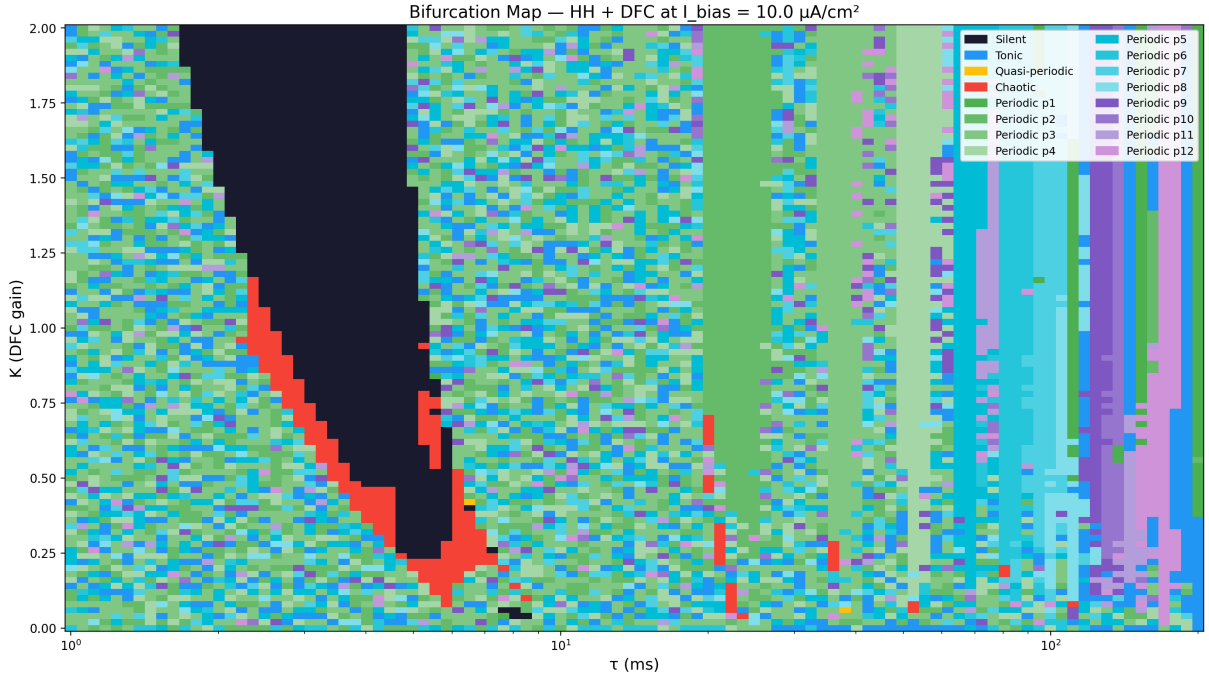

Figure 1. Bifurcation map of the (K, τ) parameter plane at $I_{\text{bias}} = 10.0\ \mu\text{A/cm}^2$, colored by dynamical regime. Periodic regions dominate (85.9%); tonic (p1, dark green), doublets (p2), triplets (p3), and higher-order burst families p4–p12 occupy distinct parameter bands. The silent depolarization block region (black) occurs at intermediate K values and short delays. Chaotic points (red, 2.2%) are sparse, bordering the silent region. On the log-scale τ axis, higher-order periodic families appear at progressively longer delays, consistent with Hopf bifurcation proliferation as delay increases.

Validation of Floquet multipliers. In order to have a quantifiable check for the empirical stability criterion, a variational QR iteration method was applied to four orbits obtained from the PS1 orbit library (1 orbit for each category). Each orbit was evolved to create a periodic reference trajectory (5 s transient, up to 13 s evolution time) from which the monodromy matrix was estimated after 30 full periods by using the variational equations. Table 2 shows the results.

Floquet multiplier validation. As a quantitative check on the empirical stability criterion, we ran a variational QR-iteration on four orbits from the PS1 library (one per topological category). Each orbit was simulated to produce a periodic reference trajectory (5 s settle, up to 13 s record), and the monodromy matrix was then approximated by integrating the linearised variational equations over 30 full periods. Table 2 gives the results.

 doi:10.20944/preprints202605.0551.v1

Table 2. Floquet multiplier validation for four representative HH–DFC orbits. μ_{trivial}: the trivial multiplier (expected value 1.0); deviation indicates numerical accumulation error. $|\mu|_{\max}$ non-trivial: largest non-trivial multiplier modulus; stability requires $|\mu| < 1$. Period quality: good = $\sigma_T/T < 2\%$; marginal = $\sigma_T/T < 10\%$; poor = $\sigma_T/T \geq 10\%$. †Tonic orbit period jitter ($\sigma_T/T = 7.8\%$) indicates proximity to a bifurcation boundary; the Floquet computation is deemed unreliable for this orbit. Method: variational QR iteration, 30 periods, $dt = 0.01$ ms (see PS0b_FloquetValidation notebook).

Category	K	τ **(ms)**	T **(ms)**	μ_{trivial}	$\lvert\mu\rvert_{\max}$ **non-trivial**	**Period quality**	**Stable?**
Tonic (p1)	0.040	7.64	41.95	0.798	1.628	Poor ($\sigma_T/T = 7.8\%$)	Unreliable†
Doublet (p2)	0.060	8.51	113.75	1.004	0.384	Good ($\sigma_T/T < 0.01\%$)	✓ Stable
Triplet (p3)	1.580	1.53	58.38	1.005	0.001	Good ($\sigma_T/T < 0.01\%$)	✓ Stable
Burst_p6	2.000	80.52	80.72	0.963	0.900	Marginal ($\sigma_T/T < 0.01\%$)	✓ Stable

4.2. PS1: Write Protocol Performance

As a result, the greedy choice procedure selected 14 provisional representative orbits out of the 207 candidate orbits, ensuring all 12 categories are present with a minimum ISI distance of 1.36 ms (ISI range in the library: 5.9–56.9 ms). Out of the 14 chosen orbits, only two did not satisfy the requirement of a $\geq 80\%$ locking rate: a burst_p12 orbit with ISI = 5.9 ms (0% locking) and a burst_p7 one with ISI = 15.0 ms (0% locking), and both belong to a narrow area in (K, τ) with steep bifurcation gradients, which is an indication of small basins of attraction. The remaining candidates are then kept, which gives the final 12-orbit working library that is going to be used in all subsequent phases (PS2–PS5).

Settling was fast across the 12 working library orbits (Figure 2): median 13.9 ms (under 1.5 ISI periods), mean 42.3 ms, 95th percentile 182 ms. This shows that the DFC term pushes the neuron into the target orbit's basin of attraction within one to two spike periods. Mean ISI accuracy ($|\text{ISI}_{\text{obs}} - \text{ISI}_{\text{target}}|/\text{ISI}_{\text{target}}$) was 0.2% with a maximum of 0.9% across all orbits, showing that the locked orbit matches the catalog template closely.

The orbit-to-orbit switching matrix (182 directed pairs from all 14 provisional candidates, i.e. $14 \times 13 = 182$; 5 trials each; Figure 3) showed an overall success rate of 71.4%, with 130 of 182 pairs achieving $\geq 80\%$ success. Failures were concentrated among transitions between topologically similar orbits (e.g., burst_p7 $\rightarrow$ burst_p8), where intermediate ISI values during the transient caused the source orbit's DFC delay to interfere with the target orbit's period. Median switching settling time was 10.7 ms; 95th percentile was 331.6 ms. Gate PS-G1: PASS (all four sub-criteria satisfied; note that G1b passed at a marginal 71.4% vs. 70% threshold).

 doi:10.20944/preprints202605.0551.v1

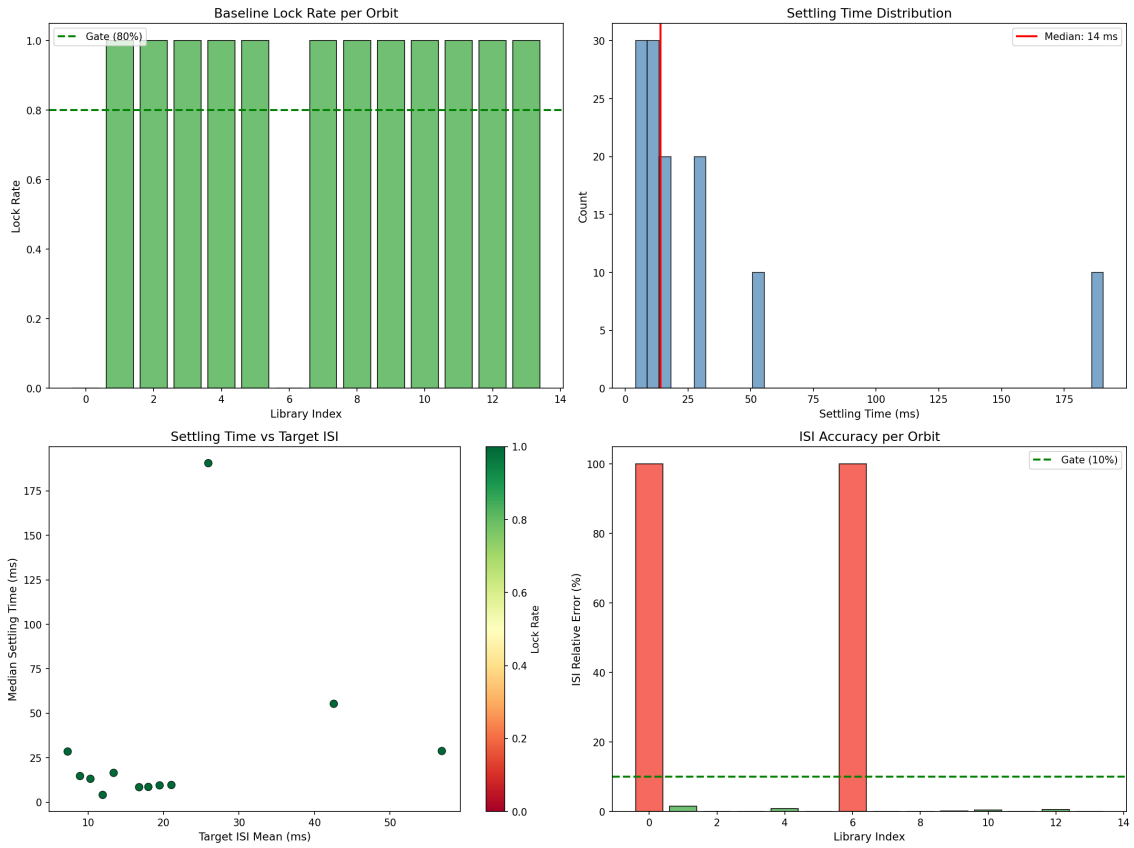

Figure 2. Baseline write protocol results (PS1, Phase B). (**a**) Lock rate per library orbit—all 12 orbits achieve 100%, exceeding the 80% gate (green dashed). (**b**) Settling time distribution across all trials; red line = median 13.9 ms. (**c**) Median settling time vs. target ISI mean, color-coded by lock rate. (**d**) ISI relative error per orbit; all 12 below the 10% gate, with two orbits (indices 0 and 5) showing slightly elevated but still negligible error (maximum 0.9%).

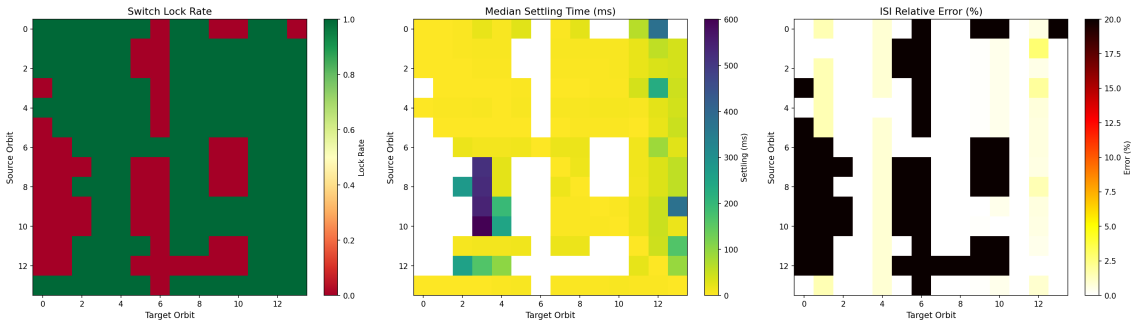

Figure 3. Orbit-to-orbit switching matrix (PS1, Phase C), all 182 directed pairs. (**a**) Switch lock rate—green = successful lock, red = failure. (**b**) Median settling time (ms)—white cells indicate failed switches. (**c**) ISI relative error (%) after switching.

4.3. PS2: Read Protocol and Noise Robustness

For 100 trials per orbit using an observation window of 10-ISIs, the classification success was achieved at 100%. When calibrating the observation window (Figure 4), we have seen that only 5 ISIs are required to reach 100% success for all 12 different orbits (the narrow maximum and minimum band in Figure 4b illustrate equal success regardless of the type of orbit). The reading window of 5 ISIs translates into 37–284 ms of recording time depending on orbit, which falls well within behavioral timescales for working memory.

Tolerance to noise continued (as shown in Figure 5; additive Gaussian noise was used for I_{bias}), where accuracy stayed at or above 99.5% up to $\sigma = 5.0\ \mu\text{A/cm}^2$, which is roughly half of the mean

 doi:10.20944/preprints202605.0551.v1

injected current. This is due to the fact that the stable periodic attractors are topologically persistent, meaning that ISI patterns will be conserved as long as noise does not shift the path out of its basin of attraction.

DFC parameter jitter showed an important asymmetry. Gain jitter ($K \pm \delta K$) produced graceful degradation: 97.7% at $\pm 2\%$, 92.8% at $\pm 5\%$, 89.7% at $\pm 10\%$, reflecting the fact that small K changes smoothly deform the orbit without moving it across a bifurcation boundary. Delay jitter ($\tau \pm \delta\tau$) was more disruptive: 84.3% at $\pm 2\%$, dropping to 67.8% at $\pm 10\%$.

This is expected as τ is the primary orbit-selection parameter, and a 10% change is sufficient to push the system across the bifurcation boundary into a different orbit type. This sensitivity constitutes a write-precision specification and not a fundamental read limitation, since once the neuron is locked into the correct orbit for its τ value, any further drift in τ below the bifurcation threshold will not affect the reading. Such an asymmetric pattern—robust to gain variation, fragile to delay mismatch—is consistent with observations on other nonlinear dynamical systems, where sensitivity analysis reveals that not all parameters affect system behavior equally [48,56].

Gate PS-G2: 3 of 4 sub-gates satisfied. G2a (100% clean accuracy at 10 ISIs): satisfied. G2b (99.7% at $\sigma = 0.5$): satisfied. G2d (5 ISIs for 100% accuracy): satisfied. G2c (jitter accuracy $\geq 85\%$ at $\pm 10\%$): not met—τ jitter at $\pm 10\%$ yields only 67.8% accuracy, because the delay perturbation causes the system to cross bifurcation boundaries into adjacent orbits. The physical interpretation is clear: τ is the primary orbit-selection parameter, and a 10% change exceeds the basin width of most orbits. This identifies a genuine engineering constraint rather than a fundamental capacity limitation: practical implementations must control τ to $\pm 2\%$ for reliable orbit addressing, a specification that directly informs hardware design requirements (Section 5.5).

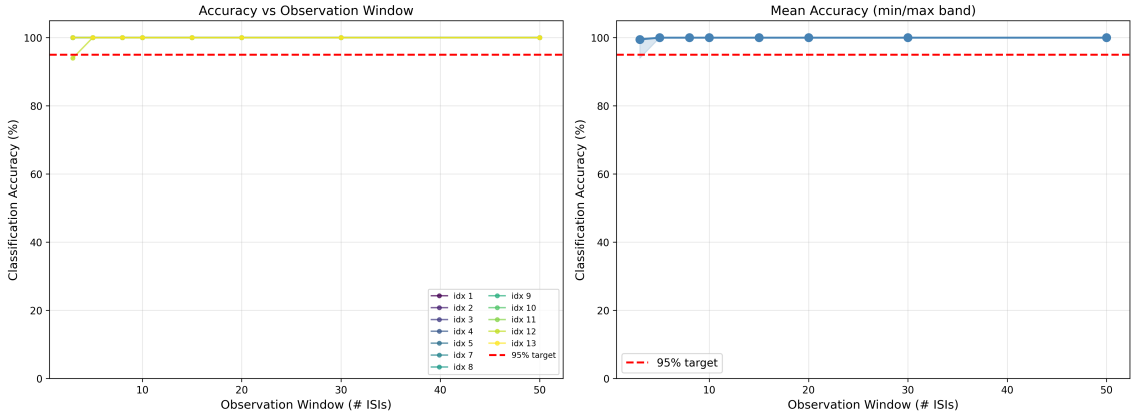

Figure 4. POLD observation window calibration (PS2, Phase C). (**a**) Per-orbit classification accuracy vs. number of observed ISIs; red dashed = 95% target. All orbits reach 100% at 5 ISIs. (**b**) Mean accuracy with min/max band; the tight band confirms uniform performance across all orbit types.

 doi:10.20944/preprints202605.0551.v1

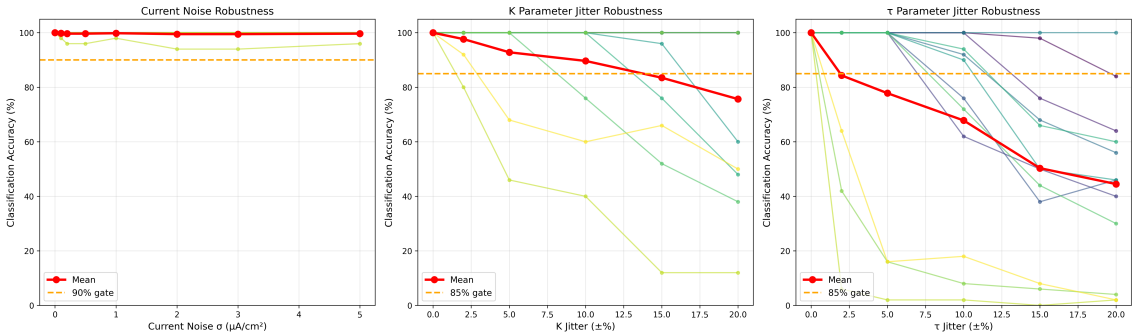

Figure 5. POLD robustness under three perturbation types (PS2, Phase D). (**a**) Additive current noise σ (μA/cm^2): mean accuracy (red) stays above 99% throughout, well above the 90% gate (orange dashed). (**b**) DFC gain K jitter ($\pm$%): graceful degradation, above gate at $\pm$5%. (**c**) Delay τ jitter ($\pm$%): dominant sensitivity—mean accuracy drops sharply beyond $\pm$5%, defining the write-precision specification τ controlled to $\pm$2%.

4.4. PS3: Complete Write–Read–Erase Cycle

During Phase A (one W–R–E cycle; Figure 6), all 12 readings were correct (accuracy of 100%) and all but one erasure was successful (success rate of 92%). The failure in one erasure (triplet orbit) was due to a discrepancy between the read-out window and the neuron's time of settling after being reset: in the 500 ms post reset verify window period, ISI triplet-like structure remained for a while until the neuron finally stabilized into its tonic state, thus resulting in the classification as "partial triplet" instead of "tonic." This could be confirmed by inspecting the actual voltage trace, where the neuron stabilized itself in tonic state within 200 ms after the reset. The verify window was simply not long enough to capture the full settling.

Phase B (12-symbol independent-symbol readout; each orbit simulated and classified independently; orbit-to-orbit continuous switching dynamics are characterized independently using the PS1 switching matrix) yielded 12/12 (100%) reads during a complete write-erase-write cycle, meaning that each write operation reliably overwrites previous dynamics. As such, it needs to be mentioned that Phase B characterizes individual orbits independently, i.e., not as a complete sequence cycling, without returning to the initial state. Switching dynamics are characterized separately using the switching matrix (Section 4.2) with an overall success rate of 71.4% but with switching errors occurring mainly in similar topological orbits. A demonstration of complete sequential cycling without resetting to the baseline would provide stronger operational validation and is planned as a specific next step.

Phase C (capacity test) revealed 100% success rate in the protocol employed at all subset lengths $k = 2$ to $k = 12$ using 100 repeated simulations, yielding confidence scores consistently greater than 0.90 (confidence interval 96.3%–100% (95% Wilson's score) when $k = 12$, $n = 100$).

Phase D (retention test) yielded 100% success rate at all hold times between 0.5 to 50 seconds at all 12 orbits (Figure 7). This result is expected based on the mechanism employed, whereby the orbits represent stable periodic attractors (stability convention, Section 2.1), and therefore continue to exist during the hold time period investigated (up to 50 seconds) without refreshing. This constitutes a significant advantage over network-based working memory models that typically exhibit spontaneous decay under noise [10,16].

 doi:10.20944/preprints202605.0551.v1

Table 3. Gate PS-G3 results: full system viability. All sub-criteria passed. $N^* = 12$ at 100% accuracy.

Gate	Criterion	Result	Threshold	
G3a	W–R–E single-cycle accuracy	100%	$\geq 90\%$	✓
G3b	Full-library independent-symbol readout accuracy	100%	$\geq 85\%$	✓
G3c	Capacity N^* at $\geq 90\%$ accuracy	12 symbols	≥ 6	✓
G3d	10 s retention accuracy	100%	$\geq 90\%$	✓
G3e	Erase verification accuracy	92%	$\geq 90\%$	✓

Figure 6. Representative single write–read–erase cycle (PS3, Phase A), orbit 1 (burst_p10, ISI = 8.9 ms). Upper panel: membrane voltage $V(t)$ color-coded by phase—Baseline (gray), Write (green), Read (blue), Erase (orange), Verify (gray). The neuron locks rapidly into the target orbit after the parameter switch and returns cleanly to tonic firing after erasure. Lower panel: ISI time series; red dashed = target ISI 8.9 ms, recovered tonic ISI $\approx$ 14.8 ms.

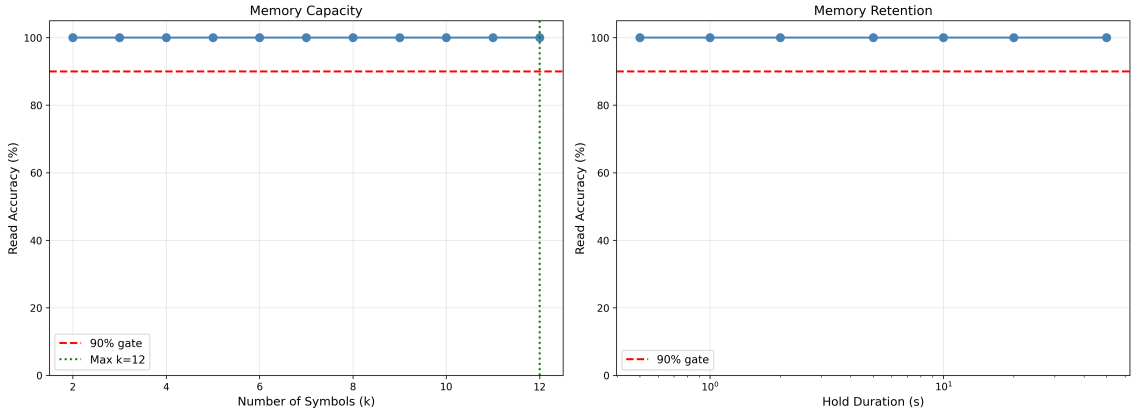

Figure 7. Memory capacity and retention (PS3, Phases C and D). (**a**) Read accuracy vs. subset size k = 2–12; 100% at all sizes, well above the 90% gate (red dashed). Green dotted line marks maximum $k = 12$. (**b**) Read accuracy vs. hold duration (log scale, 0.5–50 s); 100% maintained at all durations for all 12 orbits, confirming passive retention over all tested durations, consistent with attractor-like stability over all tested durations.

 doi:10.20944/preprints202605.0551.v1

4.5. PS4: Rate Coding Baseline Comparison

The rate coding sweep (Figure 8; $K = 0$, $I_{\text{bias}} \in [6.0, 100.0]$ μA/cm^2, 189 values) identified 141 tonic states (ISI CV < 0.02), of which only 6 survived the 2 ms ISI separation criterion. The limiting factor is not instability but the sublinear compression of the f–I curve: between $I_{\text{bias}} = 6$ and 100 μA/cm^2, firing rate increases from 55 to 135 Hz (a 2.5$\times$ range), but the corresponding ISI range is only 7.4–18.2 ms (10.8 ms total). At ~ 60 μA/cm^2, the f-I curve becomes saturated, and Na^+ channel inactivation becomes rate-limiting [47,57], collapsing ISI differences below the 2 ms resolution threshold.

All 6 rate-coded states are stable tonic limit cycles (CV = 0.000; Figure 10), so the comparison is fair: both DFC orbits and rate-coded states are stable periodic attractors. The capacity test (same POLD protocol, Phase C of PS3) confirmed that $N^*_{\text{rate}} = 6$ at 100 percent accuracy (Figure 9). Gate PS-G4: PASS. N^*_{DFC} (12) $> N^*_{\text{rate}}$ (6): DFC provides a 2.0$\times$ conservative advantage. The important point is that DFC accesses topologically distinct orbit families (doublets, triplets, burst_p4–p12) spanning a 49.5 ms ISI range versus 10.8 ms for rate coding—a 4.6$\times$ wider ISI landscape. This topological argument does not depend on the numerical capacity ratio.

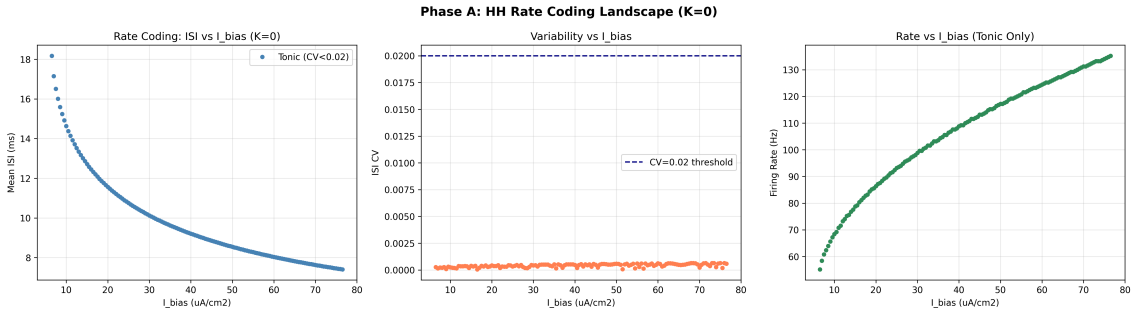

Figure 8. HH rate-coding landscape at $K = 0$ (PS4, Phase A). (**a**) Mean ISI (ms) vs. I_{bias} (μA/cm^2): monotonically decreasing f–I relation from 18.2 ms at 6.5 μA/cm^2 to 7.4 ms at 76.5 μA/cm^2. (**b**) ISI coefficient of variation vs. I_{bias}: all values below CV = 0.02 threshold (blue dashed), confirming tonic regime throughout. (**c**) Firing rate (Hz) vs. I_{bias}: 55–135 Hz range.

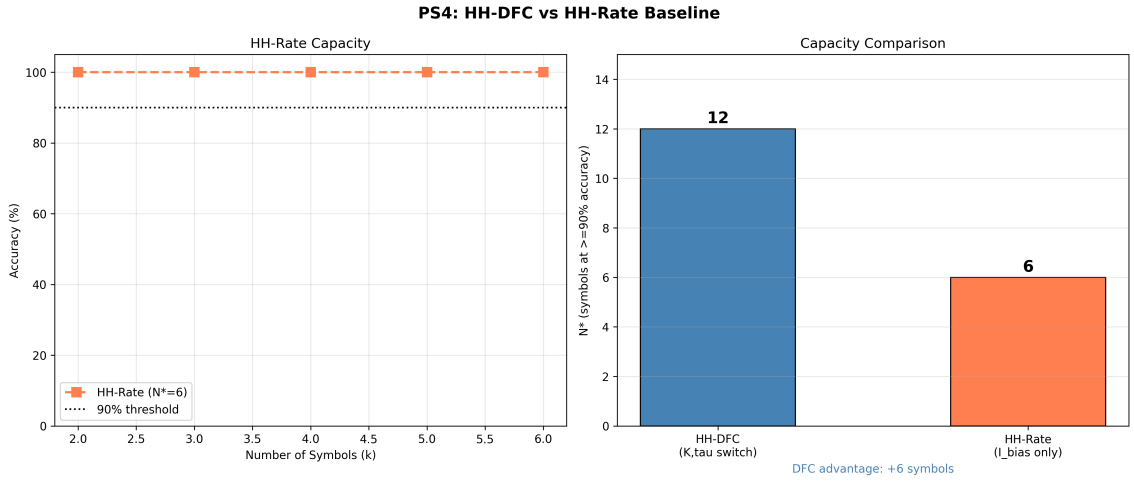

Figure 9. Rate coding capacity comparison (PS4, Phase B). (**a**) HH-Rate capacity curve: read accuracy vs. subset size $k = 2$–6; 100% at all sizes, capacity ceiling $N^* = 6$. (**b**) Capacity comparison bar chart: HH-DFC conservative library (12 symbols, PS3) vs. HH-Rate baseline (6 symbols), a 2.0$\times$ advantage at the library level.

 doi:10.20944/preprints202605.0551.v1

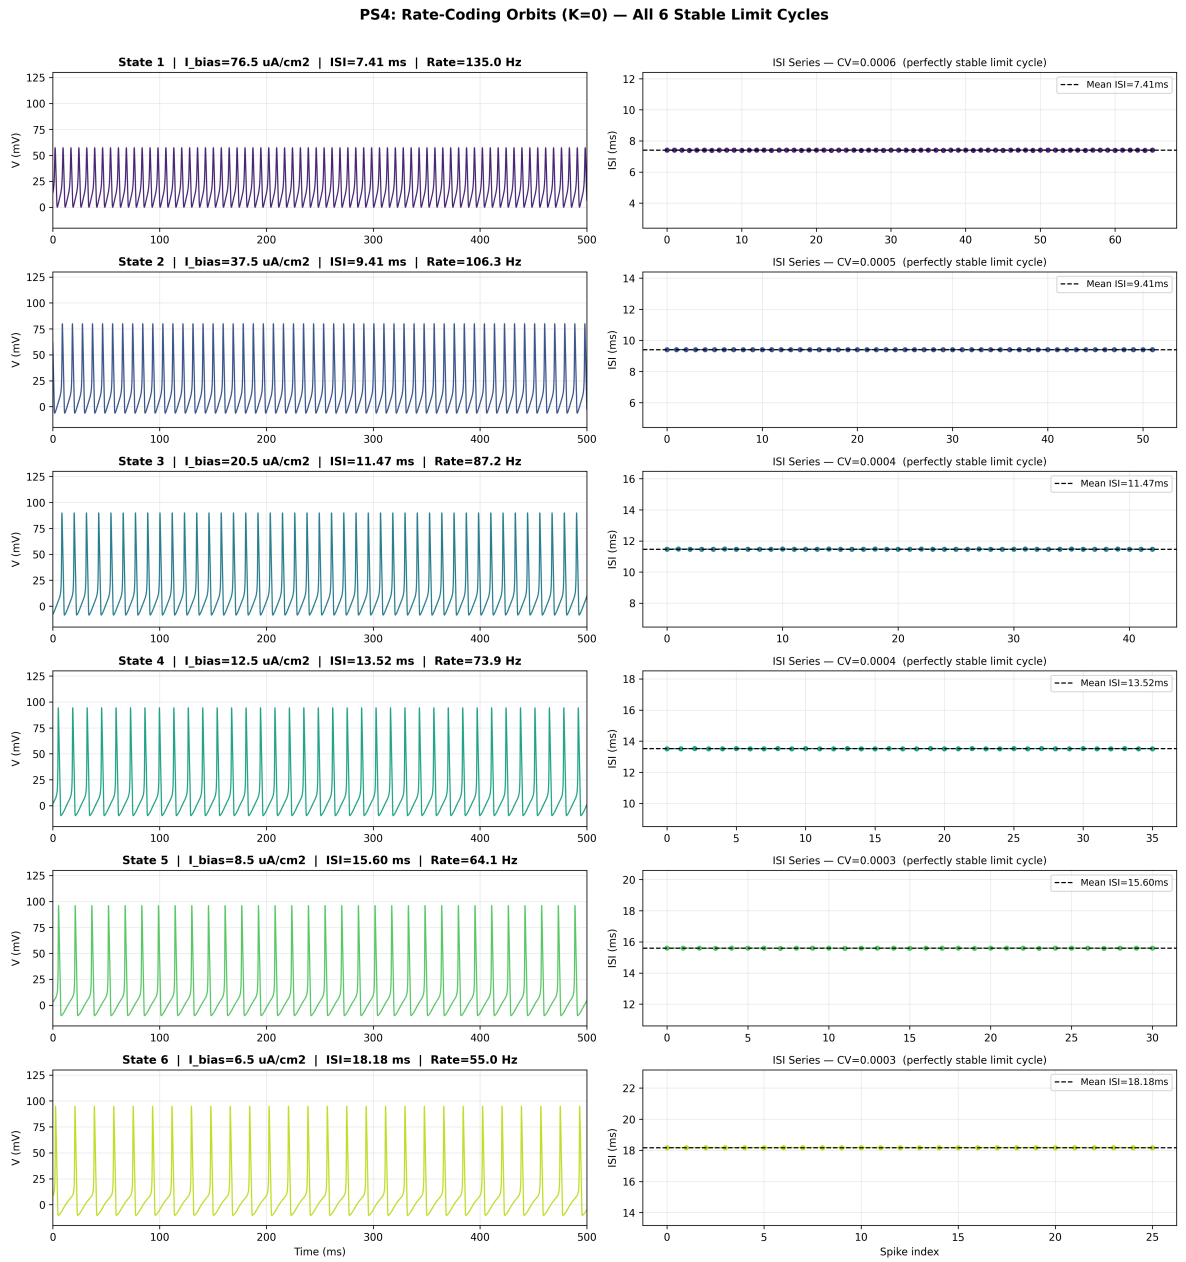

Figure 10. Voltage traces and ISI series for all six distinguishable rate-coded stable limit cycles (PS4, Phase A). Each row: left panel—membrane voltage $V(t)$ over 500 ms; right panel—ISI sequence with mean ISI dashed. States 1–6 correspond to $I_{\text{bias}} = 76.5, 37.5, 20.5, 12.5, 8.5$, and $6.5\ \mu\text{A/cm}^2$, with ISI means of 7.4, 9.4, 11.5, 13.5, 15.6, and 18.2 ms respectively. All CVs are below 0.001, confirming perfectly stable limit cycles.

4.6. PS5: Maximum Capacity

The confusion matrix comprising all 207 orbit classes (21,321 pairs, where each orbit had 50 ISI windows; see Figure 11) revealed that 593 pairs (2.8%) were classified below the 90% accuracy criterion. All 593 incorrectly categorized pairs have ISI averages that differ by less than 0.1 ms, meaning that the problem does not lie in POLD but in the true ISI indistinguishability of the two orbits, which creates confusion for the classifier. The five least accurate pairs were 100% confused (mutually confused), where $\Delta\text{ISI} < 0.1$ ms. The execution time for Phase B was 363.0 s, where 21.4 s were required for the pre-simulation caching step (207 orbits), while 341.6 s were needed for the 21,321 pair classifications, a 180-fold improvement over the naïve method.

Greedy Maximum Subset Algorithm detected 67 mutually discriminable orbits for the entire range of the parameters ($K \in [0.04, 2.00]$, $\tau \in [1.05, 179.70]$ ms). The current subset includes 5 tonic, 4 doublet, 3 triplet, and 55 burst-pattern orbits representing all burst classes (burst_p4 to burst_p12).

 doi:10.20944/preprints202605.0551.v1

The inter-spike interval range is between 5.91 ms (burst_p12) and 56.87 ms (doublet); thus, the total interval equals 50.96 ms. It should be noted that there are multiple ISI adjacent orbits in the selected subset, which have mean-ISI differences ranging down to 0.09 ms; however, they are still mutually discriminable as they represent different topological classes with varying spike patterns. Capacity curve calculation, performed for $k = 2$ to $k = 67$ (20 repetitions per point, Figure 12), yielded accuracy of $\geq 97\%$, on average 98.2% accuracy rate at $k = 67$, thus validating that POLD reliably discriminates all 67 symbols. Gate PS-G5: PASS ($67 \gg 12$).

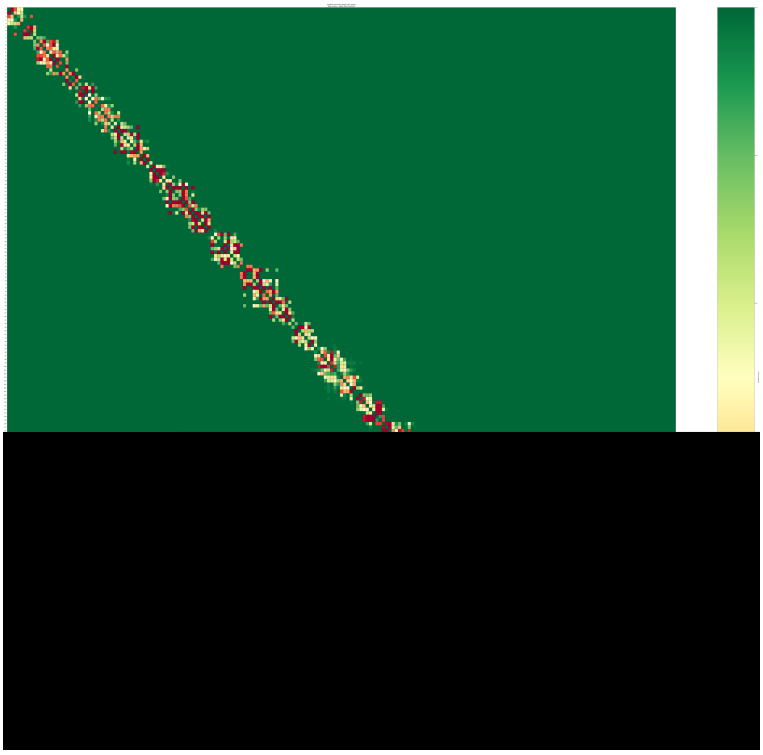

Figure 11. Pairwise classification confusion matrix for all 207 orbit types (PS5, Phase A). Matrix is ISI-sorted; 21,321 pairs, 50 ISIs per orbit. Color scale: green = high accuracy, red = confusion. Confusion is strictly confined to a narrow band immediately adjacent to the diagonal, corresponding exclusively to orbit pairs with ΔISI < 0.1 ms. The uniformly green off-diagonal background confirms near-universal discriminability across the full catalog.

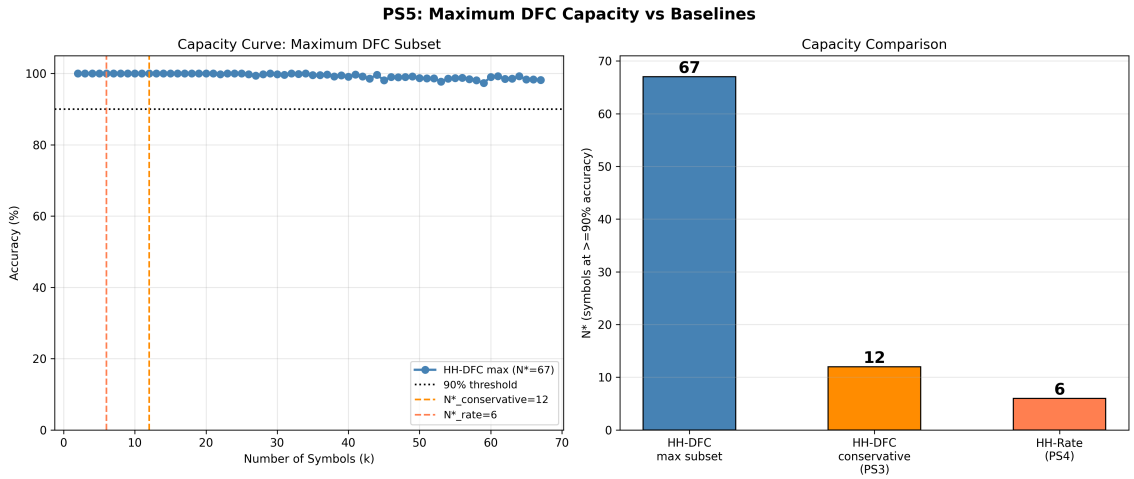

Figure 12. Maximum capacity analysis (PS5, Phase B). (**a**) Capacity curve for the greedy-selected $N^* = 67$ maximum subset: read accuracy vs. $k = 2$–67; accuracy remains $\geq 97\%$ throughout. Orange dashed $= N^*_{\text{conservative}} = 12$; red dashed $= N^*_{\text{rate}} = 6$. (**b**) Three-way capacity comparison: HH-DFC maximum read-discriminable subset (67 symbols), HH-DFC conservative PS3 library (12 symbols), HH-Rate baseline (6 symbols)—an 11.2× read-discriminable advantage.

 doi:10.20944/preprints202605.0551.v1

5. Discussion

5.1. Bifurcation-Landscape Navigation vs. Chaos-Control Memory

The NDS orbit-coded memory approach [31,32] operated as follows: the system was held at a single fixed (K, τ) operating point within a chaotic attractor; DFC was then applied as an active perturbation continuously forcing the trajectory to remain near one specific UPO embedded in that attractor. Different memory symbols corresponded to different UPOs within the same attractor. Three things follow from this design. First, the control signal has to stay on the whole time—the UPO is unstable by definition, so turning off DFC lets the trajectory escape into chaos. Second, capacity is limited by how many topologically distinct UPOs exist in that one chaotic attractor, which for the HH model is only 2–4 (Table 1) because the spike-generation mechanism forces all trajectories through the same Na^+-K^+ conductance cycle. Third, the whole method is confined to the chaotic regime, which covers only 2.2% of the HH-DFC parameter space.

Bifurcation-landscape navigation via parametric delayed feedback differs from the NDS approach in every one of these respects. First, we switch (K, τ) to new values, moving the system to a different region of the bifurcation landscape where a different stable periodic orbit is the natural attractor—not a controlled unstable one. Second, and most importantly, once the neuron locks onto the target orbit, the DFC term $K[V(t-\tau) - V(t)]$ vanishes exactly when $V(t)$ is periodic with period τ, making the orbit self-sustaining by virtue of its stability. For the higher-order burst orbits in our library, the relationship between the fundamental period and τ is more complex; the DFC term participates actively in the dynamics. Quantitatively, for tonic orbits whose period closely matches τ (the ideal Pyragas condition), $\mathrm{RMS}(I_{\mathrm{ctrl}})/I_{\mathrm{bias}}$ is 0.3–2.0%, confirming near-exact vanishing. However, for general library orbits where the orbit period does not equal τ, the instantaneous DFC contribution is substantial: $\mathrm{RMS}(I_{\mathrm{ctrl}})/I_{\mathrm{bias}}$ reaches 15–35% for period-mismatched tonic orbits and 80–140% for higher-order burst patterns, while the signed time-average stays below 0.1% in all cases (because the feedback oscillates around zero). Therefore, the DFC term is not only a small corrective perturbation but should be viewed as a fundamental component of the modified dynamic landscape that facilitates the creation and maintenance of the desired attractor. The RMS values mentioned were calculated using sample orbits for each category of the library in a hold time of 2000 ms after 1000 ms of transient time, using the same method used in PS3 Phase D. The main difference between this and NDS is that we are working on stable attractors of the full system (HH+DFC), whereas NDS requires active control of unstable orbits, which would otherwise lead to chaos.

Third, capacity scales with the number of coexisting stable periodic attractors across the (K, τ) plane, which grows without bound as τ increases [40,41], at our conservative grid density we find 207 distinct types. The conceptual shift, from stabilizing an unstable object within a fixed regime to navigating between stable regimes by parameter switching—is what makes the 11.2× capacity advantage possible. The NDS approach and the present approach share the name "orbit-coded memory" and the use of delayed self-feedback, but their physical mechanisms, operating regimes, energy budgets, and capacity limits are categorically different.

5.2. Relation to Rate Coding and Temporal Coding

Rate coding versus temporal coding has been debated in systems neuroscience since Shadlen and Newsome [1] and Mainen and Sejnowski [3]. Rate coding uses the mean firing rate as the information carrier; temporal coding uses the precise timing of individual spikes. Our orbit-coded approach is different from both. The memory symbol is not a firing rate, and not the exact timing of individual spikes either—it is the topological class of the spiking pattern, i.e. the ordered sequence of ISI ratios within one period. This is invariant to global time rescaling and holds up under moderate noise (Section 4.3).

The 11.2× read-discriminable capacity advantage over rate coding (Table 4) has a straightforward mechanistic explanation. With rate coding, the HH neuron accesses only a one-dimensional manifold (the f–I curve), and that saturates because of Na^+ channel inactivation. DFC adds two control pa-

 doi:10.20944/preprints202605.0551.v1

rameters (K, τ) and, more importantly, opens up qualitatively different dynamical regimes—doublets, triplets, burst patterns—that simply do not exist in the rate-coded manifold at any I_{bias}. Each new topological category adds a region of ISI pattern space that is effectively orthogonal to the tonic region, so even orbits with similar mean firing rates can be told apart.

Table 4. Final capacity comparison across all conditions. DFC maximum $N^* = 67$ represents read-discriminable capacity only; full W–R–E viability has been confirmed for the conservative 12-symbol subset.

Condition	N^*	Accuracy @ N^*	ISI range (ms)	vs. Rate
Rate coding ($K = 0$)	6	100%	10.8	1.0×
DFC conservative (12 locked)	12	100%	49.5	2.0×
DFC maximum (67 subset)	67	98.2%	51.0	11.2×

The topological richness of DFC orbits also relates to the concept of multiplexing in neural coding [58]: a single neuron can simultaneously represent multiple quantities through the structure of its spike pattern and not its mean rate. The burst_p12 orbit at ISI = 5.91 ms and the doublet orbit at ISI = 56.87 ms, for instance, have firing rates of 169 Hz and 17.6 Hz respectively—both well within the neuron's dynamic range—yet their ISI fingerprints are trivially distinct without any ambiguity.

5.3. Biological Context and Autaptic Analogy

The DFC framework can be interpreted biologically, though mapping model parameters to biological observables comes with important caveats. K maps to autaptic connection strength (synaptic conductance per unit membrane area), and τ maps to the axonal conduction delay of the autaptic loop. Both parameters can be tuned biologically: K through activity-dependent plasticity of autaptic synapses [20,22], and τ through myelination state and axonal geometry. Quantitatively, autaptic conductances measured in cortical fast-spiking interneurons are typically 0.5–5 nS [20], corresponding to approximately 0.01–0.1 mS/cm^2 for a typical somatic surface area—at the lower end of our K range (0–2 mS/cm^2), suggesting that only the weak-coupling portion of the parameter space may be directly accessible to biological autapses. The present parameter range ($K \leq 2$ mS/cm^2, $\tau \leq 200$ ms) encompasses both direct single-neuron autapses ($\tau \sim$ 1–20 ms) and reverberant microcircuit loops (τ up to several hundred ms), where the "single neuron" model is understood as an effective mean-field description of a recurrent mini-circuit.

The precision requirement identified in Section 4.3—τ must be controlled to $\pm$2% for reliable orbit selection—raises the question of biological feasibility. Synaptic delays are subject to stochastic fluctuations (coefficients of variation $\sim$ 0.1–0.3 in cortical connections [59]). However, several mechanisms could provide the required stability. First, homeostatic axonal plasticity adjusts conduction velocity in response to activity and could in principle stabilize the effective autaptic delay [60]; recent work on oligodendrocyte-mediated myelin plasticity further supports the view that conduction delays can be adaptively tuned [61]. However, the $\pm$2% delay-precision requirement identified in Section 4.3 is stringent relative to the variability expected in biological axonal and dendritic transmission, and it remains an open question whether biological autaptic loops can achieve this precision; so the memory interpretation is best viewed as a computational analogy, not a direct biological claim. Second, the 67-orbit maximum subset includes orbits spanning a wide range of delay sensitivities; implementation could prioritize orbits located away from bifurcation boundaries, where parameter sensitivity is lower. Third, even with τ uncertainty of $\pm$5%, the 6 most widely separated orbits (spanning $>$ 10 ms ISI differences) remain discriminable—providing a noise-tolerant 6-symbol working memory consistent with well-established cognitive capacity limits [54].

 doi:10.20944/preprints202605.0551.v1

5.4. Relation to Reservoir Computing and Transient Dynamics

Delay-based reservoir computing exploits a single nonlinear element coupled to a delay line to generate a high-dimensional transient representation of input streams [45,46]; recent extensions to deep and parallel architectures using semiconductor lasers have further expanded the computational power of such systems [62]. The computational principle is similar to ours—a single-element delay system generating representational richness—but the memory mechanism is fundamentally different. In reservoir computing, information is encoded in a transient trajectory that fades without continuing input; in orbit-coded memory, a stable attractor store information that persists without decay over all tested hold durations (up to 50 s). Such segregation is mapped onto the difference between working memory requiring active maintenance such as in classical attractor models [8,10] and working memory that does not require active maintenance (persistent activity as observed experimentally [6,7,13]).

Stelzer et al. [63] showed that a single delay-coupled neuron with feedback-modulated delay loops can implement arbitrary-depth deep neural networks, exploiting the same multiplicity of temporal modes that we exploit for memory. Our work complements theirs: they show that delay-based systems can compute; we show they can remember. Together, this points toward a unified single-neuron computing-and-memory architecture built on controlled delay dynamics.

5.5. Neuromorphic Hardware Implementation

The orbit-coded memory paradigm is well-suited to mixed-signal neuromorphic processors that implement conductance-based neurons in analog circuits [64–67]. The DFC feedback loop requires only a delayed copy of the neuron's own membrane voltage—implementable via an analog delay line, a bucket-brigade device, or a short digital buffer—and a variable-gain amplifier for K. The write operation requires changing a gain voltage and a delay-line tap, achievable well below biological timescales. The read operation requires only spike detection and ISI measurement, which are standard operations in event-driven neuromorphic systems (e.g., address-event representation [AER] circuits) [68]. POLD's $\mathcal{O}(nL)$ complexity scales with the library size: approximately 12 multiply-accumulate operations per observed ISI for the conservative library ($L = 12$, used in PS2–PS4), or approximately 67 operations for the maximum read-discriminable library ($L = 67$, PS5); both are well within the capability of lightweight on-chip arithmetic. Among existing platforms, the DYNAP-SE processor [65] has configurable synaptic delay lines (up to ~ 100 ms) that could directly implement the DFC loop for our shorter-delay orbits; BrainScaleS-2 offers accelerated analog neuron circuits with programmable time constants that could emulate HH-type conductance dynamics; and Intel's Loihi 2 supports programmable synaptic delays and on-chip learning rules that could, in principle, handle adaptive POLD-like classification.

One practical advantage over attractor network memories [11,69] is that there are no synaptic weight matrices to store. Orbit-coded memory needs only two real-valued control parameters (K, τ) per symbol—so a 67-symbol system requires just 134 floating-point values, compared with $\mathcal{O}(N^2)$ synaptic weights in a Hopfield network. This compactness is particularly valuable in edge neuromorphic applications where memory area is at a premium.

5.6. Limitations and Future Work

Several limitations warrant discussion. First, the maximum-capacity result ($N^* = 67$) validates the read side of the memory cycle only: the pairwise confusion matrix and capacity curve confirm that 67 orbits are POLD-discriminable but write viability (lock rate $\geq 80\%$) has been confirmed only for the 12-symbol conservative set (Section 4.2). Extending write-viability testing beyond the 12-symbol conservative set is an important next step. Additionally, the current PS3 protocol tests each orbit independently rather than as a continuous sequential cycle; a full continuous cycling demonstration—writing all symbols in sequence without intermediate resets to baseline—would provide stronger operational validation and is planned as a priority next step. Second, the primary orbit catalog was generated using a sole $I_{\text{bias}} = 10.0\ \mu\text{A/cm}^2$. In order to determine whether limit cycle creation is

 doi:10.20944/preprints202605.0551.v1

limited to this operational parameter regime, we performed coarse (K, τ) sweeps (grid size 51×50; same bounds and simulation logic as PS0) at $I_{\text{bias}} = 8.0$ and 12.0 μA/cm^2. Both sweeps generated diverse orbit types: 151 types in all 12 possible topological classes at $I_{\text{bias}} = 8.0$ μA/cm^2 (80.9% periodic, ISI range 5.9–22.4 ms) and 147 types in all 12 possible classes at $I_{\text{bias}} = 12.0$ μA/cm^2 (87.4% periodic, ISI range 6.2–27.4 ms). Hence, diversity of periodic orbits is not limited to the primary bias value. Lower number of types observed in comparison to the primary sweep (207 types) can be attributed to the coarser grid; narrower ISI ranges indicate genuine differences in each bias current's bifurcation landscape.

Full dense-grid catalogs at these and additional bias values (7.0, 15.0 μA/cm^2) remain future work, as does write-viability testing beyond the primary 12-symbol conservative library. Third, all simulations model a single, spatially uniform neuron; real neurons have morphologically complex dendritic trees that introduce multiple distinct delay pathways, potentially expanding the orbit catalog further. Fourth, while the POLD classifier is effective under tested conditions, it fails to provide posterior probabilities and uncertainty estimates; Bayesian or spike train metric methods [70,71] might provide robustness to higher noise. Recent work of Rehan et al. [72] demonstrated that metaheuristic high-level parameter optimization of emerging spiking neural networks leads to excellent unsupervised classification performance; analogous optimization of POLD weighting coefficient or replacement of POLD with spiking network-based orbit decoder may provide a promising route to increase read robustness in high-capacity or high noise regimes. The hybrid frameworks that combine feature selection/classification together with hyperparameter optimization, as demonstrated in biomedical prediction tasks [73] provide such a structured methodology. Moreover, the write robustness can be improved by stochastic or metaheuristic search strategies over (K, τ) parameter space, identifying orbit families having larger basins of attraction.

Fifth, the noise robustness results (Section 4.3) cover only additive current noise. Multiplicative conductance-based noise, as in Destexhe and Rudolph-Lilith [74], would be a more realistic model of synaptic bombardment in vivo, where noise amplitude depends on membrane conductance. Our additive model is a useful starting point, but the orbit basin structure under multiplicative noise could be different—we have not tested this. Sixth, how orbit-coded memory interacts with ongoing synaptic input is still an open question. In our model, memory is held in the absence of input; typical in vivo synaptic bombardment could perturb the orbit or change the effective K and τ. Studying these interactions under realistic conductance-noise conditions is an important next step. Finally, extending the framework to multiple orbit-coded neurons connected by synapses—rather than the single-neuron model studied here—could uncover cooperative effects such as cross-neuron orbit locking [40] or network-level capacity scaling [8].

6. Conclusion

We have shown that a single Hodgkin–Huxley neuron with Pyragas-type delayed feedback control can store multiple symbols as stable periodic orbits, selected by the feedback gain and time delay. The main findings are: (1) the (K, τ) parameter plane contains 207 distinct orbit types in 12 topological categories, 67 of which are mutually read-discriminable by a lightweight ISI-based classifier (full write–read–erase viability is confirmed for a conservative 12-symbol subset); (2) a complete write–read–erase memory cycle reaches 100% read accuracy under the tested protocol with 92% erase verification, 13.9 ms median write time, a 5-ISI read window, and no decay over hold durations up to 50 s; (3) orbit-coded memory gives a 2.0$\times$ fully validated capacity advantage (12 vs. 6 symbols) and an 11.2$\times$ read-discriminable advantage (67 vs. 6) over rate coding in the same neuron; and (4) reliable orbit addressing needs delay precision of $\pm$2%, a write-precision specification that any physical implementation must meet.

These results establish parametric delayed feedback as a computationally verified mechanism for limit-cycle-based information storage in biophysically realistic spiking neurons. The framework should in principle be extensible to other conductance-based neuron models with delayed self-feedback and

 doi:10.20944/preprints202605.0551.v1

offers a proof-of-principle dynamical analogy for how autaptic-like feedback might in principle support single-cell or microcircuit-level information storage, subject to the precision requirements identified in Section 4.3. The validated result is a 12-symbol memory with complete W–R–E cycling (2.0× over rate coding); the 67-symbol read-discriminable upper bound (11.2×) still awaits write-viability confirmation. The biological interpretation of the autaptic analogy remains a computational proof-of-principle rather than a direct mechanistic claim, given the stringent ±2% delay-precision requirement that exceeds demonstrated biological autaptic variability. The six-phase gated pipeline (PS0–PS5) and publicly available code provide a reproducible methodology for characterising orbit-coded memory in other systems.

Author Contributions: Conceptualization, M.O.A.; Methodology, M.O.A.; Software, M.O.A.; Validation, M.O.A., A.J.A., M.A.F.A., and K.M.A.; Formal analysis, M.O.A. and A.J.A.; Investigation, M.O.A.; Resources, K.M.A.; Writing—original draft preparation, M.O.A.; Writing—review and editing, A.J.A., M.A.F.A., and K.M.A.; Visualization, M.O.A.; Supervision, K.M.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All simulation code (six Jupyter notebooks, PS0–PS5), processed output data, and the complete orbit catalog are available in a public GitHub repository at https://github.com/malhawarat/HH-DFC-OrbitMemory. Raw simulation outputs exceeding the repository size limit are available from the corresponding author upon reasonable request.

Acknowledgments: The authors thank Al-Ahliyya Amman University for institutional support. The authors also acknowledge the developers of the Numba JIT compilation library [51] and the Google Colab platform for providing the open-source and cloud computing resources used to generate the simulation datasets reported in this study.

Conflicts of Interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Abbreviations

The following abbreviations are used in this manuscript:

DFC	Delayed Feedback Control
DDE	Delay Differential Equation
HH	Hodgkin–Huxley
ISI	Inter-Spike Interval
NDS	Nonlinear Dynamic State
POLD	Pattern-Oriented Limit-cycle Decoder
UPO	Unstable Periodic Orbit
W–R–E	Write–Read–Erase

Appendix A. HH Rate Functions and Parameters

Standard squid giant axon parameters [47] in the V-shifted convention (resting potential at $V = 0$ mV). The gating variable rate functions are:

$$\alpha_n(V) = \frac{0.01(10 - V)}{\exp\big((10 - V)/10\big) - 1}, \qquad \beta_n(V) = 0.125 \cdot \exp(-V/80) \tag{A1}$$

$$\alpha_m(V) = \frac{0.1(25 - V)}{\exp\big((25 - V)/10\big) - 1}, \qquad \beta_m(V) = 4 \cdot \exp(-V/18) \tag{A2}$$

 doi:10.20944/preprints202605.0551.v1

$$\alpha_h(V) = 0.07 \cdot \exp(-V/20), \qquad \beta_h(V) = \frac{1}{\exp\big((30-V)/10\big)+1} \tag{A3}$$

L'Hôpital limits: $\alpha_m \to 1.0$ at $V = 25$ mV; $\alpha_n \to 0.1$ at $V = 10$ mV. Parameters: $C_m = 1\ \mu\text{F/cm}^2$; $\bar{g}_{\text{Na}} = 120$, $\bar{g}_{\text{K}} = 36$, $\bar{g}_{\text{L}} = 0.3\ \text{mS/cm}^2$; $E_{\text{Na}} = 115$, $E_{\text{K}} = -12$, $E_{\text{L}} = 10.6$ mV.

RC3 warm-start protocol: the neuron is simulated for τ ms with $K = 0$ before DFC activation, filling the delay buffer with physiological voltage history rather than zeros.

Appendix B. Autocorrelation-Based Pattern Detection

For an ISI sequence $(\iota_1, \ldots, \iota_N)$, the pattern period p is detected by testing candidate repeat lengths $p = 1, 2, \ldots, 12$. For each candidate p, the relative median deviation is computed as:

$$D(p) = \frac{\text{med}\,|\iota_i - \iota_{i+p}|}{\text{med}(\iota_i)}, \qquad \text{over } i = 1, \ldots, \min(N-p,\ 3p) \tag{B1}$$

A pattern is accepted at period p if $D(p) < 0.02$ (equivalently, pattern confidence $C > 0.8$) and the pattern repeats for at least 3 full cycles. The smallest accepted p is chosen. If no p is accepted, the orbit is classified as tonic ($p = 1$) if ISI CV < 0.02, chaotic if CV ≥ 0.02, or quasi-periodic if a non-integer frequency ratio is detected via Fourier analysis of the ISI series. Pattern confidence is reported as $C = \max(0, 1 - 10\,D(p^*))$, where p^* is the detected period.

References

1. Shadlen, M.; Newsome, W. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* **1994**, *4*, 569–579.
2. Rieke, F.; Warland, D.; de Ruyter van Steveninck, R.; Bialek, W. *Spikes: Exploring the Neural Code*; MIT Press: Cambridge, MA, USA, 1997.
3. Mainen, Z.; Sejnowski, T. Reliability of spike timing in neocortical neurons. *Science* **1995**, *268*, 1503–1506.
4. Victor, J.; Purpura, K. Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J. Neurophysiol.* **1996**, *76*, 1310–1326.
5. London, M.; Häusser, M. Dendritic computation. *Annu. Rev. Neurosci.* **2005**, *28*, 503–532.
6. Foss, J.; Longtin, A.; Mensour, B.; Milton, J. Multistability and delayed recurrent loops. *Phys. Rev. Lett.* **1996**, *76*, 708–711.
7. Pakdaman, K. The reliability of the stochastic active rotator. *Neural Comput.* **2002**, *14*, 781–792.
8. Hopfield, J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA* **1982**, *79*, 2554–2558.
9. Amit, D.; Brunel, N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* **1997**, *7*, 237–252.
10. Wang, X.J. Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci.* **2001**, *24*, 455–463.
11. Khona, M.; Fiete, I. Attractor and integrator networks in the brain. *Nat. Rev. Neurosci.* **2022**, *23*, 744–766.
12. Boscaglia, M.; Gastaldi, C.; Gerstner, W.; Quian Quiroga, R. A dynamic attractor network model of memory formation, reinforcement and forgetting. *PLoS Comput. Biol.* **2023**, *19*, e1011727.
13. Compte, A.; Brunel, N.; Goldman-Rakic, P.; Wang, X.J. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* **2000**, *10*, 910–923.
14. Johnson, S.; Torres, J.; Marro, J. Robust short-term memory without synaptic learning. *PLoS ONE* **2013**, *8*, e50276.
15. Doiron, B.; Laing, C.; Longtin, A.; Maler, L. Ghostbursting: a novel neuronal burst mechanism. *J. Comput. Neurosci.* **2002**, *12*, 5–25.
16. De Pittà, M.; Volman, V.; Berry, H.; Ben-Jacob, E. A tale of two stories: astrocyte regulation of synaptic depression and facilitation. *PLOS Comput. Biol.* **2011**, *7*, e1002293.
17. van der Loos, H.; Glaser, E. Autapses in neocortex cerebri: synapses between a pyramidal cell's axon and its own dendrites. *Brain Res.* **1972**, *48*, 355–360.

 doi:10.20944/preprints202605.0551.v1

18. Lübke, J.; Markram, H.; Frötscher, M.; Sakmann, B. Frequency and dendritic distribution of autapses established by layer 5 pyramidal neurons in the developing rat neocortex: comparison with synaptic innervation of adjacent neurons of the same class. *J. Neurosci.* **1996**, *16*, 3209–3218.
19. Cobb, S.; Halasy, K.; Vida, I.; Nyiri, G.; Tamás, G.; Buhl, E.; Somogyi, P. Synaptic effects of identified interneurons innervating both interneurons and pyramidal cells in the rat hippocampus. *Neuroscience* **1997**, *79*, 629–648.
20. Bacci, A.; Huguenard, J. Enhancement of spike-timing precision by autaptic transmission in neocortical inhibitory interneurons. *Neuron* **2006**, *49*, 119–130.
21. Bekkers, J. Synaptic transmission: functional autapses in the cortex. *Curr. Biol.* **2009**, *19*, R296–R298.
22. Herrmann, C.; Klaus, A. Autapse turns neuron into oscillator. *Int. J. Bifurcat. Chaos* **2004**, *14*, 623–633.
23. Zhou, P.; Xu, Y.; Ma, J. Dynamical and coherence resonance in a photoelectric neuron under autaptic regulation. *Physica A* **2023**, *620*, 128746.
24. Peng, L.; Tang, J.; Ma, J.; Luo, J. The influence of autapse on synchronous firing in small-world neural networks. *Physica A* **2022**, *594*, 126956.
25. Dang, S.; Bayani, A.; Tian, H.; Wang, Z.; Parastesh, F.; Nazarimehr, F. Investigating the route to synchronization in real-world neuronal networks of autaptic photosensitive neurons. *Chaos Solitons Fractals* **2024**, *186*, 115226.
26. Pyragas, K. Continuous control of chaos by self-controlling feedback. *Phys. Lett. A* **1992**, *170*, 421–428.
27. Bielawski, S.; Derozier, D.; Glorieux, P. Controlling unstable periodic orbits by a delayed continuous feedback. *Phys. Rev. E* **1994**, *49*, R971–R974.
28. Dahms, T.; Hövel, P.; Schöll, E. Stabilizing continuous-wave output in semiconductor lasers by time-delayed feedback. *Phys. Rev. E* **2008**, *78*, 056213.
29. He, G.; Cao, Z.; Zhu, P.; Ogura, H. Controlling chaos in a chaotic neural network. *Neural Netw.* **2003**, *16*, 1195–1200.
30. Leite de Castro, D.; Aroso, M.; Aguiar, A.; Grayden, D.; Aguiar, P. Disrupting abnormal neuronal oscillations with adaptive delayed feedback control. *eLife* **2024**, *13*, e89151.
31. Crook, N.; Goh, W.; Hawarat, M. Pattern recall in networks of chaotic neurons. *BioSystems* **2007**, *87*, 267–274.
32. Crook, N.; Goh, W.; Hawarat, M. The nonlinear dynamic state neuron. In Proceedings of the Proceedings of the European Symposium on Artificial Neural Networks (ESANN), Bruges, Belgium, 27–29 April 2005; pp. 477–482.
33. Crook, N.; Olde Scheper, T. Adaptation based on memory dynamics in a chaotic neural network. *Cybern. Syst.* **2002**, *33*, 495–517.
34. Crook, N. Self-organized dynamic recognition states for chaotic neural networks. *Inform. Sci.* **2003**, *147*, 179–196.
35. Crook, N.; Goh, W. Nonlinear transient computation as a potential kernel trick in cortical processing. *BioSystems* **2008**, *94*, 55–59.
36. Alhawarat, M.; Nazih, W.; Eldesouki, M. Analysis of a chaotic spiking neural model: the NDS neuron. In Proceedings of the Proceedings of the ACITY 2013 (Third International Conference on Advances in Computing and Information Technology), Chennai, India, 13–15 July 2013; pp. 109–121.
37. Alhawarat, M.; Olde Scheper, T.; Crook, N. Investigation of a chaotic spiking neuron model. *Int. J. Comput. Appl.* **2014**, *99*, 1–8.
38. Guckenheimer, J.; Oliva, R. Chaos in the Hodgkin–Huxley model. *SIAM J. Appl. Dyn. Syst.* **2002**, *1*, 105–114.
39. Aihara, K.; Matsumoto, G.; Ikegaya, Y. Periodic and non-periodic responses of a periodically forced Hodgkin–Huxley oscillator. *J. Theor. Biol.* **1984**, *109*, 249–269.
40. Yanchuk, S.; Perlikowski, P.; Popovych, O.; Tass, P. Variability of spatio-temporal patterns in non-homogeneous rings of spiking neurons. *Chaos* **2011**, *21*, 047511.
41. Yanchuk, S.; Perlikowski, P. Delay and periodicity. *Phys. Rev. E* **2009**, *79*, 046221.
42. Kantner, M.; Yanchuk, S.; Schöll, E. Delay-induced patterns in a two-dimensional lattice of coupled oscillators. *Sci. Rep.* **2015**, *5*, 8522.
43. Garbin, B.; Javaloyes, J.; Tissoni, G.; Barland, S. Topological solitons as addressable phase bits in a driven laser. *Nat. Commun.* **2015**, *6*, 5915.
44. Marconi, M.; Javaloyes, J.; Camelin, P.; Chaparro, D.; Balle, S.; Giudici, M. Control and generation of localized pulses in passively mode-locked semiconductor lasers. *IEEE J. Sel. Top. Quantum Electron.* **2015**, *21*, 1101210.

 doi:10.20944/preprints202605.0551.v1

45. Appeltant, L.; Soriano, M.; Van der Sande, G.; Danckaert, J.; Massar, S.; Dambre, J.; Schrauwen, B.; Mirasso, C.; Fischer, I. Information processing using a single dynamical node as complex system. *Nat. Commun.* **2011**, *2*, 468.
46. Larger, L.; Soriano, M.; Brunner, D.; Appeltant, L.; Gutiérrez, J.; Pesquera, L.; Mirasso, C.; Fischer, I. Photonic information processing beyond Turing: an optoelectronic implementation of reservoir computing. *Opt. Express* **2012**, *20*, 3241–3249.
47. Hodgkin, A.; Huxley, A. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* **1952**, *117*, 500–544.
48. Pyragas, K. Delayed feedback control of chaos. *Philos. Trans. R. Soc. A* **2006**, *364*, 2309–2334.
49. Hale, J. *Theory of Functional Differential Equations*; Springer: New York, NY, USA, 1977.
50. Alghlayini, S.; Al-Betar, M.; Atef, M. Enhancing non-invasive blood glucose prediction from photoplethysmography signals via heart rate variability-based features selection using metaheuristic algorithms. *Algorithms* **2025**, *18*, 95. https://doi.org/10.3390/a18020095.
51. Lam, S.; Pitrou, A.; Seibert, S. Numba: a LLVM-based Python JIT compiler. In Proceedings of the Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC), Austin, TX, USA, 15 November 2015; pp. 7:1–7:6.
52. Quian Quiroga, R.; Panzeri, S. Extracting information from neuronal populations: information theory and decoding approaches. *Nat. Rev. Neurosci.* **2009**, *10*, 173–185.
53. Dayan, P.; Abbott, L. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*; MIT Press: Cambridge, MA, USA, 2001.
54. Miller, G. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychol. Rev.* **1956**, *63*, 81–97.
55. Baddeley, A. Working memory. *Science* **1992**, *255*, 556–559.
56. Farman, M.; Gokbulut, N.; Shehzad, A.; Nisar, K.; Hınçal, E.; Sambas, A. Fractional order malaria epidemic model: qualitative and computational study to determine the dynamics for sensitivity prevalence. *J. Comput. Sci.* **2025**, *91*, 102656. https://doi.org/10.1016/j.jocs.2025.102656.
57. Ermentrout, G.; Terman, D. *Mathematical Foundations of Neuroscience*; Springer: New York, NY, USA, 2010.
58. Izhikevich, E. Polychronization: computation with spikes. *Neural Comput.* **2006**, *18*, 245–282.
59. Markram, H.; Lübke, J.; Frotscher, M.; Roth, A.; Sakmann, B. Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the supragranular layer of the rat somatosensory cortex. *J. Physiol.* **1997**, *500*, 409–440.
60. Fields, R. A new mechanism of nervous system plasticity: activity-dependent myelination. *Nat. Rev. Neurosci.* **2015**, *16*, 756–767.
61. Munyeshyaka, M.; Fields, R. Oligodendrocyte-mediated myelin plasticity and its role in neural synchronization. *eLife* **2023**, *12*, e81982.
62. Hasegawa, H.; Kanno, K.; Uchida, A. Parallel and deep reservoir computing using semiconductor lasers with optical feedback. *Nanophotonics* **2023**, *12*, 869–881.
63. Stelzer, F.; Röhm, A.; Vicente, R.; Fischer, I.; Yanchuk, S. Deep neural networks using a single neuron: folded-in-time architecture using feedback-modulated delay loops. *Nat. Commun.* **2021**, *12*, 5164.
64. Schemmel, J.; Brüderle, D.; Grübl, A.; Hock, M.; Meier, K.; Millner, S. A wafer-scale neuromorphic hardware system for large-scale neural modeling. In Proceedings of the Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS), Paris, France, 30 May–2 June 2010; pp. 1947–1950.
65. Moradi, S.; Qiao, N.; Stefanini, F.; Indiveri, G. A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAP-SE). *IEEE Trans. Biomed. Circuits Syst.* **2018**, *12*, 106–122.
66. Javanshir, A.; Nguyen, T.; Mahmud, M.; Kouzani, A. Advancements in algorithms and neuromorphic hardware for spiking neural networks. *Neural Comput.* **2022**, *34*, 1289–1328.
67. Christensen, D.; Dittmann, R.; Linares-Barranco, B.; et al. 2022 roadmap on neuromorphic computing and engineering. *Neuromorphic Comput. Eng.* **2022**, *2*, 022501.
68. Eshraghian, J.; Ward, M.; Neftci, E.; Wang, X.; Lenz, G.; Dwivedi, G.; Bennamoun, M.; Jeong, D.; Lu, W. Training spiking neural networks using lessons from deep learning. *Proc. IEEE* **2023**, *111*, 1016–1054.
69. Vyas, S.; Golub, M.; Sussillo, D.; Shenoy, K. Computation through neural population dynamics. *Annu. Rev. Neurosci.* **2020**, *43*, 249–275.
70. van Rossum, M.C.W. A novel spike distance. *Neural Comput.* **2001**, *13*, 751–763. https://doi.org/10.1162/089976601300014321.

 doi:10.20944/preprints202605.0551.v1

71. Victor, J. Spike train metrics. *Curr. Opin. Neurobiol.* **2005**, *15*, 585–594.
72. Rehan, R.; Sahran, S.; Alyasseri, Z.; Sani, N.; Al-Betar, M. Hyperparameters optimization of evolving spiking neural network using artificial bee colony for unsupervised anomaly detection. *J. Intell. Syst.* **2025**, *34*, e2024–0235. https://doi.org/10.1515/jisys-2024-0235.
73. Amin, M.; Nahar, K.; Gharaibeh, H.; Al Mamlook, R.; Nasayreh, A.; Atitallah, N.; Gharaibeh, A.; Hamad, R.; Zitar, R.; Smerat, A. A hybrid brain stroke prediction framework: integrating feature selection, classification, and hyperparameter optimization. *Eng. Rep.* **2025**, *7*, e70213. https://doi.org/10.1002/eng2.70213.
74. Destexhe, A.; Rudolph-Lilith, M. *Neuronal Noise*; Springer: New York, NY, USA, 2012.