

Emergence Requires No Miracle: The Eligibility Trace as the Necessary Engine of Reward-Driven Differentiation in Spiking Networks

Yahua Ruan

yahua.ruan@gientech.com

GienTech Research Institute

Research Article

Keywords: Spiking neural network, Emergence, Eligibility trace, Reward-modulated STDP, Sparsification, Temporal credit assignment, Ablation, Scaling law

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11433074/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Emergence Requires No Miracle: The Eligibility Trace as the Necessary Engine of Reward-Driven Differentiation in Spiking Networks

Yahua Ruan

GienTech Research Institute, Beijing 100086, China

yahua.ruan@gientech.com

ORCID: 0009-0009-9760-1542

Abstract

Emergent differentiation in reward-driven learning requires local events to persist long enough for global evaluation to bind them. In a spiking neural network trained with reward-modulated STDP ($N = 800$ – $80,000$ neurons, up to 128M synapses), five-condition ablation shows that removing the eligibility trace eliminates positive sparsification at every scale on the post-warmup baseline ($\Delta\text{Gini} = -0.006 \pm 0.001$ at $N=800$; -0.005 ± 0.003 at $N=80,000$); measured from random initialization, the traceless condition still drifts upward at small scale ($+0.088$ at $N=800$), so we state the necessity claim on the post-warmup baseline and report both baselines. Removing the reward prediction error or the logistic brake amplifies sparsification but destabilizes the system (R locks at 0.31 – 0.35 and never self-arrests; winners hit the weight cap); removing multiplicative potentiation is free at $N=800$ (99% of full effect) but becomes necessary at $N \geq 8,000$ (16–29%). A scaling law follows: the braking timescale must scale as $\tau_{\text{exp}} \propto \sqrt{N}$; extending τ from 20,000 to 200,000 at $N = 80,000$ raises the differentiation equilibrium from 0.365 (three-seed mean at default τ) to 0.561 ± 0.024 (three seeds at $\tau = 200,000$; +54% on the mean, near-converged at 500k steps), surpassing the $N = 800$ equilibrium (0.509). In conditions retaining both the trace and RPE, $R \rightarrow 0$ (self-arrest) and $\cos(\mathbf{w}_t, \mathbf{w}_{t-1}) \rightarrow 1.0$ (attractor convergence). We propose that temporal credit assignment is the necessary engine of reward-driven emergent differentiation in this system, with partial structural analogies to SGD momentum and evolution; direct experimental verification is limited to synthetic spiking-network tasks,

and sufficiency is not tested.

Keywords: Spiking neural network, Emergence, Eligibility trace, Reward-modulated STDP, Sparsification, Temporal credit assignment, Ablation, Scaling law

1 Introduction

Emergence—the spontaneous appearance of system-level order from local interactions—is typically explained by invoking complex mechanisms: positive feedback loops Camazine et al. [2001], competitive dynamics Song et al. [2000], Gilson and Fukai [2011], carefully designed schedules Bengio et al. [2009], or architectural constraints. In large language models, emergent abilities appear suddenly at scale Wei et al. [2022] and provoke measurement debates Schaeffer et al. [2023], Barak et al. [2022], Power et al. [2022], Olsson et al. [2022]. In reward-driven spiking networks, the prerequisite is theoretically established at the level of three-factor learning rules Klopff [1982], Legenstein et al. [2008], Frémaux and Gerstner [2016], Gerstner et al. [2018]: without a decaying memory of recent synaptic events, a delayed global reward cannot modulate STDP at all. What remains open is the quantitative anatomy—which components drive, which bound, and how the answers shift with scale. We answer this by controlled ablation.

The claim is measured. In our spiking network, removing the eligibility trace (a 20-step exponential memory of recent synaptic events) eliminates positive sparsification ($\Delta\text{Gini} = -0.006 \pm 0.001$ at $N=800$, own-start baseline). Removing the brakes (reward prediction error, logistic saturation) *amplifies* the phenomenon (+308% to +1004% across scales) while destabilizing the system; removing multiplicative STDP preserves it at $N=800$ (99%) but not at $N \geq 8,000$ (16–29%). The eligibility trace is the only component whose removal eliminates positive sparsification at every scale tested.

Why this matters. The finding restructures the theoretical narrative of emergent differentiation:

- **Engine vs. brake.** Prior work conflates mechanisms that *drive* differentiation with mechanisms that *limit* it. Our ablation cleanly separates them: the eligibility trace drives; RPE and logistic bound; multiplicative STDP accelerates (at scale).

- **Simplicity.** The necessary engine is: local events persist in a decaying trace; a global reward multiplies the trace; weights update. Three lines of pseudocode. (We test necessity by removal; sufficiency—a minimal trace-plus-reward construction—is future work.)
- **Cross-substrate generality.** A partial structural analogy appears in SGD (momentum first-moment EMA as eligibility-like memory) and evolution (genotype as persistence memory) Elsayed et al. [2024]—an analogy, not an isomorphism, as we discuss in §8.1.

Contributions.

1. **Ablation discovery** (§4): Systematic 5-condition $\times$ 3-seed ablation revealing the engine/brake/accelerant structure, with the eligibility trace as the only component whose removal eliminates positive sparsification at every scale.
2. **Scale invariance** (§5): Replication at $N = 800$, 8,000, and 80,000 (128M synapses, 500k steps), showing ΔGini is preserved across two orders of magnitude while the role of multiplicative STDP transitions from accelerant to necessity.
3. **τ -scaling law** (§5.4): Discovery that the braking timescale must scale as $\tau_{\text{exp}} \propto \sqrt{N}$; verified across three seeds by raising $N = 80,000$ Gini from 0.365 to 0.561 ± 0.024 at 500k steps (surpassing $N = 800$'s 0.509) with $\tau = 200,000$.
4. **Robustness** (§6): Five-seed validation with $\text{Gini} = 0.509 \pm 0.020$, $R \rightarrow 0.001$ (self-arrest), and $\cos \rightarrow 1.0$ (attractor convergence) in all seeds.
5. **Theoretical unification** (§8.1): Proposal that temporal credit assignment is the necessary engine of emergent differentiation under delayed global feedback, with partial cross-substrate analogies and testable predictions.

2 Background

Spiking neural networks Maass [1997], Gerstner and Kistler [2002], Pfeiffer and Pfeil [2018]—realized on neuromorphic hardware Mead [1990], Merolla et al. [2014], Indiveri et al. [2011] and studied as liquid state machines Maass et al. [2002], Jaeger [2001]—provide the substrate of the present

work. Their plasticity is governed by spike-timing-dependent plasticity Bi and Poo [1998, 2001], Gerstner et al. [1996] and its reward-modulated variants Florian [2007], Izhikevich [2007]. This research area remains highly active: eligibility-propagation schemes now train deep recurrent spiking networks Bellec et al. [2020], with recent extensions to neuromorphic hardware and learnable delays Quintana et al. [2024], Vassallo and Taherinejad [2026]; behavioral-timescale eligibility has been identified in hippocampal and striatal plasticity Milstein et al. [2021], Yagishita et al. [2014]; and cortical gamma synchrony has been shown to couple reinforcement-learning signals to long-term memory formation in the mammalian brain Yun et al. [2026]; surrogate-gradient theory continues to mature Zenke and Vogels [2021]; the conditions for emergent structure in plastic spiking networks are under active investigation Bergoin et al. [2025]; and current reviews attest that the eligibility trace remains a live conceptual substrate for learning Shouval and Kirkwood [2025].

2.1 Leaky Integrate-and-Fire Network

We simulate N neurons with LIF dynamics:

$$\tau_m \frac{dv_i}{dt} = -(v_i - v_{\text{rest}}) + I_i(t), \quad (1)$$

with $v_{\text{rest}} = 0$, $v_{\text{thresh}} = 0.5$, $v_{\text{reset}} = -0.3$, $\tau_m = 20$ steps, and refractory period of 2 steps. Neurons are connected via $S \approx N^2 p$ synapses ($p = 0.02$), with initial weights $w \sim \mathcal{N}(0, 0.1)$ clamped to $[-1, 1]$ and transmission delays $d \in \{1, 2, 3\}$.

2.2 Reward-Modulated STDP (R-STDP)

Standard STDP modifies weights based on pre/post spike timing:

$$\Delta w_{\text{STDP}} = \begin{cases} A^+ \cdot f(w) \cdot e^{-\Delta t / \tau_+} & \text{if } \Delta t > 0 \text{ (LTP)} \\ -A^- \cdot e^{-|\Delta t| / \tau_-} & \text{if } \Delta t < 0 \text{ (LTD)} \end{cases} \quad (2)$$

where $\tau_+ = \tau_- = 20$ steps, $A^- = 0.0105$ (additive LTD).

In R-STDP Florian [2007], STDP updates do not directly modify weights. Instead, they accu-

95 modulate in an **eligibility trace**:

$$e_i(t+1) = \lambda_e \cdot e_i(t) + \Delta w_{\text{STDP},i}(t), \quad \lambda_e = e^{-1/\tau_e}, \quad \tau_e = 20 \text{ steps} \quad (3)$$

96 Weight updates occur only when a global reward signal $R(t)$ arrives:

$$\Delta w_i = R(t) \cdot \alpha \cdot e_i(t), \quad \alpha = 0.01 \quad (4)$$

97 2.3 The Four Components Under Test

98 Our full system combines four components, each previously argued to be necessary for emergent
99 sparsification:

- 100 1. **Eligibility trace** ($\tau_e = 20$): temporal credit assignment—bridges local STDP events to
101 delayed global reward.
- 102 2. **Multiplicative LTP**: $f(w) = \max(w, w_{\text{floor}}) \cdot (1 - w/w_{\text{max}})$ —“rich get richer” amplification
103 with soft saturation.
- 104 3. **Logistic saturation**: the $(1 - w/w_{\text{max}})$ factor bounds potentiation near w_{max} .
- 105 4. **Reward prediction error (RPE)**: $R(t) = \text{PA}(t) - \mathbb{E}[\text{PA}]$, where $\mathbb{E}[\text{PA}]$ is an EMA with
106 $\tau = 20,000$ —endogenous self-braking.

107 3 Experimental Setup

108 3.1 Ablation Design

109 Five conditions: the full model plus four ablations, each removing exactly one component:

Table 1: Ablation conditions.

Condition	Removed	Replacement
FULL	—	—
no.RPE	RPE reward	State reward: +1 if PA > mean, −0.3 otherwise
no.MULT	Multiplicative LTP	Additive LTP: $f(w) = A_{\text{add}}^+ = 0.02$
no.ELIG	Eligibility trace	Immediate modulation: $\Delta w = R \cdot \alpha \cdot \Delta w_{\text{STDP}}$
no.LOG	Logistic factor	$f(w) = \max(w, w_{\text{floor}})$ (no saturation)

3.2 Protocol

Each condition runs for 100,000 steps at $N \leq 8,000$ and 500,000 steps at $N = 80,000$:

- **Warmup** (300 steps): additive STDP without reward modulation, establishing baseline activity.
- **Main loop** (100k steps): R-STDP with condition-specific components.
- **Measurement**: Gini coefficient of weight distribution, reward R , maximum weight $w_{\max}$, and weight-direction cosine similarity $\cos(\mathbf{w}_t, \mathbf{w}_{t-\Delta})$ recorded every 5,000 steps.

3.3 Metrics

Gini coefficient: measures inequality of the weight distribution. Rising Gini = increasing sparsification (few synapses dominate). Computed on $|w|$ with correction for padded entries in GPU implementation.

Self-arrest: $R \rightarrow 0$ indicates the system has reached a state where performance matches expectation—no further learning signal. This is the RPE brake engaging.

Attractor convergence: $\cos(\mathbf{w}_t, \mathbf{w}_{t-\Delta}) \rightarrow 1.0$ indicates the weight vector has stopped changing direction—the system has converged to a fixed-point attractor.

4 Results I: The Ablation

4.1 $N = 800$ (12,800 synapses, 3 seeds)

Fig. 1 shows the Gini trajectories for all five conditions. Table 2 reports the final values.

4.2 The Engine/Brake Structure

The ablation reveals a clean three-tier structure:

Tier 1: The Engine (necessary). Removing the eligibility trace *abolishes* sparsification ($\Delta\text{Gini} = -0.006$, negative). Without temporal credit assignment, local STDP events are isolated—each occurs and evaporates before the global reward arrives. No binding, no accumulation, no differentiation. The weight distribution becomes *more* uniform, not less.




Figure 1: Gini trajectories for all five conditions at $N = 800$ (3 seeds). The eligibility trace (FULL) drives sparsification; removing it (no_ELIG) flatlines below its post-warmup start.

Tier 2: The Brakes (limiting). Removing RPE or logistic *amplifies* sparsification:

- no_RPE: Gini reaches 0.82 (vs. 0.482 in FULL), but R locks at 0.31—the system never self-arrests. Sparsification proceeds unchecked by prediction error.
- no_LOG: Gini reaches 0.79, but $w_{\max}$ hits the hard cap (50.0)—winners grow without bound. The logistic factor was *containing* sparsification, not causing it.

Tier 3: The Accelerant (scale-dependent). At $N = 800$, removing multiplicative LTP preserves 99% of the effect ($\Delta\text{Gini} = +0.070$ vs. $+0.071$). Additive LTP alone, combined with eligibility and RPE, produces equivalent sparsification. The multiplicative rule accelerates early differentiation but is not the driver.

4.3 Judgment Criteria

We define component necessity by the ratio $\Delta\text{Gini}_{\text{ablated}}/\Delta\text{Gini}_{\text{FULL}}$ (own-start baselines):

Table 2: Ablation results at $N = 800$ (V100 GPU). Mean $\pm$ std over 3 seeds (42, 7, 123). $\Delta\text{Gini} = \text{Gini}_{\text{final}} - \text{Gini}_{\text{start}}$ (post-warmup own-start; for no_LOG, warmup drives Gini to ~ 0.81 within $< 1\text{k}$ steps, so $\Delta \approx -0.02$ reflects post-explosion drift—see Table 5 for the random-init comparison). FULL here is the 3-seed V100 mean (0.482); the independent 5-seed CPU robustness gives 0.509 (Table 6).

Condition	Gini (final)	ΔGini	R (final)	$w_{\max}$	$\cos_t$
FULL	0.482 ± 0.012	$+0.071 \pm 0.007$	0.001	1.08	1.000
no_RPE	0.819 ± 0.002	$+0.219 \pm 0.006$	0.312	2.20	0.9999
no_MULT	0.505 ± 0.019	$+0.070 \pm 0.014$	0.001	3.44	1.000
no_ELIG	0.411 ± 0.003	-0.006 ± 0.001	0.003	1.00	1.000
no_LOG	0.790 ± 0.007	-0.019 ± 0.003	0.002	50.0	1.000

• < 0 : **absolutely necessary** (removal reverses the effect)

• 0–0.3: **necessary**

• 0.3–0.7: **partially necessary**

• > 0.7 : **non-essential**

• > 1 : removal *amplifies* the effect — a brake, not a necessity

One caveat: a condition whose brake acts during warmup (no_LOG: weights explode pre-baseline, Table 5 note) has a negative own-start Δ from post-explosion drift, not from stability; we classify it by the brake role, not the ratio.

At $N = 800$: no_ELIG = -8% (absolutely necessary); no_LOG = -27% (brake acting during warmup); no_MULT = 99% (non-essential); no_RPE = 308% (brake removed; amplifies).

5 Results II: Scale Invariance

5.1 $N = 8,000$ (1.28M synapses, 3 seeds)

5.2 The Scaling Law of Component Necessity

Key finding: The eligibility trace is *scale-invariantly nullifying*—its removal reduces ΔGini to zero or below at every scale tested (-0.006 ± 0.001 at $N = 800$; -0.001 ± 0.002 at $N = 8,000$; note the latter’s confidence interval includes zero, but the direction is consistently non-positive and the magnitude is two orders below FULL’s $+0.097$). The multiplicative rule undergoes a *necessity*




Figure 2: Engine/brake structure: ΔGini at $N = 800$ for each ablation. Only no_ELIG abolishes sparsification ($\Delta < 0$); removing brakes (no_RPE, no_LOG) amplifies it.

transition: at $N = 800$, additive LTP suffices (the fixed increment is large relative to the weight range); at $N \geq 8,000$, the weight distribution is denser and additive increments are drowned in noise—multiplicative amplification (“the rich get richer”) becomes necessary to pull winners above the bulk. (So the eligibility trace is the engine at every scale, while the turbocharger’s necessity is scale-dependent.)

Interpretation: The engine (eligibility) is universal. The turbocharger (multiplicative) is needed only when the engine’s raw power is insufficient for the load. Larger networks carry more noise; they need the turbocharger.




Figure 3: Scale invariance: ΔGini across $N = 800$ to $80,000$. The eligibility trace nullifies at every scale; multiplicative LTP transitions from accelerant to necessity.

5.3 $N = 80,000$ (128M synapses, 3 seeds, MetaX C550 GPU)

Convergence: FULL reaches equilibrium at step $\approx 150,000$ (Gini = 0.365, unchanged for the remaining 350k steps; $R = 10^{-5}$, $\cos_t = 1.0$). The 500k-step protocol provides $> 2\times$ the convergence time, confirming the result is a true fixed point.

Verdict preservation: The engine/brake structure is fully preserved at $N = 80,000$:

- no_ELIG: $\Delta\text{Gini}_{\text{own}} = -0.005$ (nullified, consistent with $N = 800$ and $8,000$).
- no_RPE: $\Delta\text{Gini}_{\text{own}} = +0.442$ ($10\times$ FULL; brake removed, runaway amplification).
- no_LOG: $\Delta\text{Gini}_{\text{own}} = -0.009$, but $\text{Gini}(\text{start}) = 0.78$ and $w_{\text{max}} = 50.0$ —the brake failed during warmup (weight explosion); the negative own-start reflects post-explosion drift, not stability.
- no_MULT: $\Delta\text{Gini}_{\text{own}} = +0.013$ (29% of FULL; multiplicative now necessary at this scale).

Table 3: Ablation results at $N = 8,000$ (V100 GPU). Mean $\pm$ std over 3 seeds. ΔGini is post-warmup own-start.

Condition	Gini (final)	ΔGini	R (final)	$w_{\max}$	$\cos_t$
FULL	0.492 ± 0.016	$+0.097 \pm 0.011$	0.0005	1.00	1.000
no.RPE	0.850 ± 0.002	$+0.391 \pm 0.028$	0.350	2.41	0.9999
no.MULT	0.416 ± 0.012	$+0.016 \pm 0.005$	-0.0001	1.67	1.000
no.ELIG	0.388 ± 0.009	-0.001 ± 0.002	0.0003	1.00	1.000
no.LOG	0.804 ± 0.023	-0.032 ± 0.026	0.002	50.0	0.9999

Table 4: Component necessity across scales (ΔGini as % of FULL, own-start baseline). All values are post-warmup own-start ($\text{Gini}_{\text{final}} - \text{Gini}_{\text{start}}$ per condition), isolating training-phase effects; the random-init (0.323) comparison is in Table 5. Logistic saturation’s brake acts during warmup (preventing weight explosion), so its training-phase Δ is negative despite the large final Gini (0.77–0.78).

Component	$N = 800$	$N = 8,000$	$N = 80,000$	Transition
Eligibility trace	−8%	−1%	−11%	Always nullifying (engine)
Logistic saturation	−27%	−33%	−20%	Brake acts in warmup
Multiplicative LTP	99%	16%	29%	Non-essential $\rightarrow$ necessary
RPE	308%	403%	1004%	Brake removed (amplifies)

5.4 The τ -Scaling Hypothesis

The equilibrium Gini at $N = 80,000$ (0.365) is lower than at $N = 800$ (0.509). This is *not* a weakening of the mechanism. The root cause is that population-activation noise scales as $1/\sqrt{N}$: larger networks produce smoother PA signals, so the per-step RPE magnitude shrinks as $1/\sqrt{N}$ while the braking window (τ_{exp}) stays fixed. Cumulative differentiation scales roughly as (per-step signal) $\times$ (effective window); holding it constant gives $\tau_{\text{exp}} \propto \sqrt{N}$.

Prediction: If τ_{exp} is scaled by $\sqrt{N}$ (from 20,000 to 200,000 for $N = 80,000$), the equilibrium Gini should rise toward—or exceed—the $N = 800$ value (0.509).

Verification (confirmed, three seeds {7, 42, 123}): With $\tau_{\text{exp}} = 200,000$ at $N = 80,000$, the equilibrium Gini rises from the default- τ three-seed mean 0.365 to 0.561 ± 0.024 (three-seed mean; per-seed finals 0.536/0.563/0.584) — a +54% rise that surpasses the $N = 800$ equilibrium (0.509).

One seed’s trajectory (seed 42):

- Step 50k: Gini = 0.410, already exceeding the standard equilibrium (0.365).
- Step 175k: Gini = 0.520, *exceeding* the $N = 800$ equilibrium (0.509).
- Step 275k: Gini = 0.554, with $R = 0.004$ still positive.

Table 5: Ablation results at $N = 80,000$ (500k steps). Mean $\pm$ std over 3 seeds (42, 7, 123). Device: MetaX C550. $\Delta\text{Gini}^{\text{own}}$ is the post-warmup own-start increment ($\text{Gini}_{\text{final}} - \text{Gini}_{\text{start}}$ per condition); Gini(start) is shown for full transparency. For reference, ΔGini from the random-init baseline (0.323): FULL +0.042, no_RPE +0.510, no_MULT +0.020, no_ELIG -0.007, no_LOG +0.448. [†]no_LOG: without logistic saturation, weights explode during warmup ($<5\text{k}$ steps, Gini $\rightarrow \sim 0.78$); own-start Δ (-0.009) reflects post-explosion drift, not stability.

Condition	Gini (start)	Gini (final)	$\Delta\text{Gini}^{\text{own}}$	R (final)	w_{max}	cos_t
FULL	0.321	0.365 ± 0.002	$+0.044 \pm 0.002$	0.00001	1.00	1.000
no_RPE	0.391	0.833 ± 0.001	$+0.442 \pm 0.001$	0.350	2.82	1.000
no_MULT	0.330	0.343 ± 0.002	$+0.013 \pm 0.003$	-0.0001	2.09	1.000
no_ELIG	0.321	0.316 ± 0.004	-0.005 ± 0.003	0.0003	1.00	1.000
no_LOG	0.780	0.771 ± 0.034	-0.009 ± 0.034	0.003	50.0	1.000

- Step 300k: Gini = 0.560, with $R = 0.003$ still positive.
- Step 325k: Gini = 0.565, with $R = 0.003$.
- Step 400k: Gini = 0.576, with $R = 0.002$.
- Step 500k: Gini = 0.584, with $R = 0.001$ —near-converged (residual drift +0.008 over the last 100k).

The trajectory converges at $\approx 500\text{k}$ steps (vs. the standard $\tau = 20,000$ run which converged at 150k), confirming that the extended braking timescale successfully prolongs the effective learning window. Across all three seeds, the equilibrium under $\tau = 200,000$ (0.561 ± 0.024) exceeds the $N = 800$ equilibrium (0.509); the trajectories were still rising slowly at 500k, so this value is a conservative lower bound.

Conclusion: The equilibrium Gini is a function of the *parameter-scale match*, not an intrinsic limitation of the mechanism at large N . The braking timescale should scale as:

$$\tau_{\text{exp}} \propto \sqrt{N} \quad (5)$$

to maintain a constant effective learning window across network sizes. This is a testable scaling law for reward-modulated plasticity: *the brake must slow down as the network gets quieter.*

6 Results III: Robustness

Five independent seeds (42, 7, 123, 999, 2024) under the FULL condition at $N = 800$ (Fig. 4):




Figure 4: Five-seed robustness at $N = 800$, FULL condition. All seeds converge to $\text{Gini} \approx 0.51$ with $R \rightarrow 0$ (self-arrest) and $\cos \rightarrow 1.0$ (attractor convergence).

Table 6: Five-seed robustness ($N = 800$, FULL condition, 100k steps).

Seed	Gini (final)	ΔGini	R (final)	$\cos_t$
42	0.538	+0.108	0.001	1.000
7	0.499	+0.066	0.001	1.000
123	0.499	+0.072	0.001	1.000
999	0.484	+0.058	0.001	1.000
2024	0.527	+0.089	0.001	1.000
Mean $\pm$ std	0.509 $\pm$ 0.020	+0.079 $\pm$ 0.019	0.001	1.000

Attractor signature: All five seeds converge to $R \rightarrow 0.001$ and $\cos \rightarrow 1.000$ regardless of initial conditions. The low variance ($\sigma = 0.020$) relative to the effect size ($\Delta = 0.079$) gives a signal-to-noise ratio of ≈ 4 , confirming the mechanism is robust to initialization.

7 Mechanism: Temporal Credit Assignment

7.1 Why the Eligibility Trace Is the Engine

The eligibility trace solves a *temporal binding problem*: STDP events are local and instantaneous, but the reward signal is global and delayed. Without a memory bridge, these two cannot interact:

- **With trace:** A neuron fires at $t = 100$, creating a pre/post spike pair and hence Δw_{STDP} at the relevant synapse. This enters the trace: $e(100) = \Delta w$. At $t = 105$, the trace has decayed to $e(105) = \lambda^5 \cdot \Delta w \approx 0.78 \cdot \Delta w$ —still alive. When reward arrives at $t = 105$, it multiplies the trace: $\Delta w_{\text{final}} = R \cdot \alpha \cdot 0.78 \cdot \Delta w_{\text{STDP}}$. The local event has been *evaluated* by the global signal.
- **Without trace (no_ELIG):** reward modulates only the *instantaneous* STDP increment, whose sign and magnitude fluctuate from step to step. With no memory of the recent direction of activity, the global signal is multiplied by noise rather than by a persistent trend: nothing accumulates, and the distribution drifts back toward uniformity ($\Delta \text{Gini} = -0.006$). The trace is what aligns the global signal with the recent direction of local change.

7.2 The One-Sentence Law

Emergence is causality bridged by time: the system remembers what it just did, and thereby local actions acquire global consequences.

The eligibility trace *is* the memory. The reward *is* the global consequence. Their product *is* the binding. Remove the memory, and causality is severed—local actions have no global consequences, and no differentiation can accumulate.

7.3 Why Brakes Are Necessary (But Not Sufficient)

Without RPE, the system still differentiates—but it never *stops*. R locks at 0.31 (perpetual “better than average” signal), and Gini climbs to 0.82+. Without logistic, winners hit the hard cap ($w = 50$). The system differentiates maximally but pathologically.

The brakes do not cause differentiation; they *bound* it. FULL’s Gini of 0.51 represents the *equilibrium* between the engine (eligibility $\times$ reward) and the brakes (RPE self-arrest + logistic containment). This equilibrium is the attractor.

8 Discussion

8.1 Cross-Substrate Unification: Temporal Credit Assignment as Universal Engine

The ablation reveals more than a circuit-level mechanism—it suggests a *structural analogy* across adaptive systems. We propose that temporal credit assignment in different substrates shares a common logical structure, instantiated in different physics. The analogy is partial, not an isomorphism: in SGD, momentum accumulates gradients unconditionally, while eligibility is gated by the reward signal and settled per decision Elsayed et al. [2024]; we mark the divergence where it matters.

8.1.1 The Isomorphism

Table 7 maps seven functional roles across three substrates. The critical row is *eligibility*: in each case, a local event must **persist in a decaying memory** long enough for a global signal to evaluate it. The persistence timescale differs—20 steps in our network, a momentum horizon in SGD, one generation in evolution—but the *logical structure* is shared:

$$\text{differentiation} = \underbrace{\text{global signal}}_{\text{evaluation}} \times \underbrace{\text{decaying memory of local events}}_{\text{eligibility}} \quad (6)$$

Table 7: Cross-substrate structural isomorphism. Each row is a functional role; each column is a physical instantiation. The logical structure (Eq. 6) is invariant across columns.

Role	Spiking network	SGD
Local event	STDP (Δw)	Per-sample gradient g_t
Eligibility (memory)	$e(t) = \lambda e + \Delta w$	Momentum EMA: $m_t = \beta_1 m_{t-1} + (1-\beta_1)g_t$
Global signal	Reward $R(t)$	Loss signal $-\nabla L$
Binding	$R \cdot \alpha \cdot e$	m_t modulated by $-\nabla L$
Differentiation	Weight sparsification	Parameter specialization
Brake	RPE ($R \rightarrow 0$)	Convergence ($\ \nabla\ \rightarrow 0$)
Accelerant	Multiplicative LTP	Adam (2nd-moment adaptation)

Evolution (conjectured isomorphism): local event = phenotypic variation; eligibility = genotype (fitness memory); global signal = selection pressure; binding = fitness $\times$ variation; differentiation = species divergence; brake = carrying capacity; accelerant = sexual recombination.

8.1.2 Evidence from Scale Invariance

Our data provide three independent lines of evidence that the mechanism is *substrate-level* rather than *implementation-level*:

1. **Network scale invariance** (§5): The verdict “eligibility nullifying” holds at $N = 800$ ($\Delta\text{Gini} = -0.006 \pm 0.001$), $N = 8,000$ (-0.001 ± 0.002), and $N = 80,000$ (-0.005 ± 0.003) on the post-warmup baseline (from random initialization the traceless $N=800$ run drifts $+0.088$; see Limitations). The *magnitude* of ΔGini in FULL ($+0.04$ – 0.10) is preserved across two orders of magnitude. The engine does not care about the size of the network—only about whether temporal binding is intact.
2. **Implementation invariance**: The CPU implementation (sequential, braincore library, Python loops) and the GPU implementation (batched, vectorized, `scatter_add_`, int32 indexing) produce the same verdict for the headline condition: $\Delta\text{Gini}_{\text{CPU}} = -0.006 \pm 0.001$ vs. $\Delta\text{Gini}_{\text{GPU}} = -0.006 \pm 0.001$ for no_ELIG at $N = 800$. Two completely different codebases, zero shared code, same verdict to three decimal places on the condition that carries the claim (the FULL condition differs slightly, 0.482 vs. 0.509; see Limitations). The mechanism lives in the *mathematics*, not the implementation.
3. **Component necessity transition**: The multiplicative rule transitions from non-essential (99%) at $N = 800$ to necessary (16%) at $N = 8,000$. This mirrors the SNR threshold in gradient-based learning Ruan [2026a]: when noise scales with system size, the “turbocharger” becomes necessary to maintain signal above the noise floor. The engine (eligibility) never transitions—it is nullifying at all scales. Only the amplification requirement is scale-dependent.

8.1.3 Connection to the SNR Threshold and the “Dead Phase”

In companion work on linear attention scaling Ruan [2026a] (data publicly archived, Zenodo DOI 10.5281/zenodo.21875430), we observed that $\text{SNR} \geq 1$ (target-key gradient signal-to-noise ratio) accompanied long-range retrieval breakthrough in 14 runs with zero exceptions, and that *direction alignment*, not rank, tracks whether the accumulated signal breaks through Ruan [2026b]. These

companion results are cited as motivation, not as load-bearing evidence.

These findings connect to the present work as follows:

- **The “dead phase” is trace accumulation below noise floor.** In attention scaling, a long phase of apparent stagnation (loss not decreasing, rank not changing) precedes sudden breakthrough. We *conjecture* that during this phase gradient direction accumulates temporal coherence— $\cos(\mathbf{g}_t, \mathbf{g}_{t+1})$ rising while magnitude remains below threshold—corresponding to the eligibility trace accumulating signal below the $\text{SNR} = 1$ boundary. We state this as a target for ongoing experiments rather than an established result; preliminary direction-coherence traces have not yet confirmed it cleanly.
- **Breakthrough is the trace crossing threshold.** When accumulated coherence exceeds noise, the self-reinforcing phase transition ignites—rank collapses, loss drops, specialization emerges. In our spiking network, this corresponds to Gini crossing above the random-init baseline and beginning its monotonic rise.
- **Direction > magnitude.** The LRR scaling result (direction alignment cosine 0.64–0.81 for success vs. 0.14 for failure) parallels our finding that the *temporal structure* of the trace matters, not its amplitude. Removing eligibility (destroying temporal coherence) kills emergence; removing multiplicative amplification (reducing amplitude) does not—at small scale.

The three papers thus form a unified narrative: the companion identifies the *threshold* ($\text{SNR} \geq 1$), the LRR study identifies the *quality metric* (direction alignment), and the present paper identifies the *mechanism* (temporal credit assignment). The mechanism generates the quality; the quality determines whether the threshold is crossed.

8.1.4 Testable Predictions

The cross-substrate isomorphism yields three falsifiable predictions:

1. **SGD:** Destroying gradient temporal coherence—by randomizing gradient directions across time steps while preserving magnitude—should prevent parameter specialization (no emergent feature selectivity in hidden layers), even with correct learning rate, architecture, and data. Prediction: final-layer Gini of neuron activation norms remains at random-init level.

2. **Evolution:** Eliminating heritability—the genotype-to-phenotype memory bridge—should prevent speciation regardless of selection pressure. This is partially confirmed by artificial life experiments where “Lamarckian” inheritance (no persistent genotype) produces adaptation but not divergence.

3. **Scale-dependent accelerant:** In SGD, momentum/Adam should transition from non-essential to necessary as model size increases—precisely mirroring our multiplicative STDP result. It is widely observed that large language models ($> 10^9$ parameters) require adaptive optimizers (Adam-family Kingma and Ba [2015]), while small models ($< 10^6$) can train with vanilla SGD—consistent with a scale-dependent accelerant transition.

8.1.5 The Universal Law

We can now state the law in substrate-independent language:

In systems adapting under delayed global feedback, emergent differentiation occurs when local events persist in a decaying memory readable by the global evaluation signal, and the persistence window exceeds the action-to-evaluation delay. We claim necessity within this class and test it by removal; self-organizing systems without any global evaluation signal (competitive STDP, reaction–diffusion patterning) exist and lie outside this scope.

This is not a metaphor. It is a falsifiable, quantitative statement. In our network the reward is computed from population activation at every step, and the trace ($\tau_e = 20$ steps) spans the window over which recent activity remains readable by it; the window-over-delay condition is tested here at the zero-window corner (no_ELIG), and a window scan is future work. In SGD with momentum, the first-moment EMA ($m_t = \beta_1 m_{t-1} + (1-\beta_1)g_t$) plays the analogous memory role, with $\lambda = \beta_1$; in evolution, one generation provides the eligibility window. In each case, *shortening the window below the evaluation delay abolishes emergence*—our no_ELIG condition is the limiting case where the window is zero.

8.2 Relation to Synaptic Tagging

Our eligibility trace has a direct biological analog: *synaptic tagging and capture* Frey and Morris [1998]. In biological neurons, a local “tag” is set at active synapses; a delayed global signal

(dopamine, representing reward prediction error Schultz et al. [1997]) then selectively consolidates tagged synapses. Our finding suggests this is not merely *one* mechanism for plasticity—it is the *necessary* engine for emergent differentiation under delayed reward in our system. The tag (eligibility) and the capture signal (reward) are not two components; they are the two halves of a single temporal bridge.

8.3 Relation to competitive self-organization

Unrewarded STDP alone can produce heavy-tailed weight distributions and neuronal specialization Song et al. [2000], Gilson and Fukai [2011], and inhibitory Hebbian/anti-Hebbian variants grow modular structure without any global reward Bergoin et al. [2025], and sparse codes can be learned without any reward at all Olshausen and Field [1996]. Our claim is therefore scoped: the eligibility trace is necessary for *reward-driven* emergent differentiation under a delayed global signal—not for structure formation in general. Whether our additive-STDP parameter region overlaps the regimes in which plain competitive STDP self-organizes is a question we flag for future work.

8.4 Relation to Reinforcement Learning

The eligibility trace in R-STDP is structurally identical to the eligibility trace in temporal-difference learning Sutton [1988], Sutton and Barto [2018], a foundational idea with roots in both reinforcement learning Minsky [1961] and the hedonistic-neuron hypothesis Klopff [1982]: both solve the temporal credit assignment problem (which past actions deserve credit for the current reward?). That such a trace is required when a delayed global reward must modulate local plasticity is theoretically established in the three-factor-rule literature Legenstein et al. [2008], Frémaux and Gerstner [2016], Gerstner et al. [2018]. Our contribution is the quantitative anatomy: a controlled ablation showing which components drive, which bound, and how the boundary shifts with scale—with the eligibility trace uniquely necessary at every scale tested.

8.5 Honest Limitations

1. **Toy network:** $N = 800\text{--}80,000$ LIF neurons with random connectivity. Biological networks have structured topology, multiple neuron types, and neuromodulatory complexity we do not model.

2. **Single reward signal:** We use a scalar global reward (population activation). Biological dopamine is region-specific, temporally complex, and multiplexed.
3. **Sparsification $\neq$ function:** We measure weight inequality (Gini), not computational performance. Whether sparsification *improves* information processing is a separate question.
4. **Cross-substrate analogy $\neq$ proof:** The SGD/evolution parallels are structural analogies, not experimental validations. The prediction (disrupt temporal bridge $\rightarrow$ abolish differentiation) is testable but untested in those substrates.
5. **Scale:** $N = 80,000$ (128M synapses, MetaX C550 GPU) is the largest scale tested, with full convergence confirmed at 150k steps. Whether the mechanism persists at 10^9+ synapses (cortical scale) is unknown, though the τ -scaling hypothesis (§5.4) provides a concrete prescription for maintaining effectiveness at larger scales.
6. **Baseline convention:** necessity claims are stated on the post-warmup (own-start) baseline; measured from random initialization, the traceless condition still drifts upward at $N=800$ (+0.088), and we report both baselines. An amplitude-matched immediate-modulation control (removing the trace while preserving update magnitude) is future work.
7. **Control-condition caveats:** the no.RPE replacement reward carries a positive mean ($\mathbb{E}[R] \approx +0.35$), so its amplification is jointly attributable to brake removal and positive bias (a zero-mean control is future work); warmup uses condition-specific dynamics, so we report own-start increments to isolate training-phase effects.
8. **Hardware:** $N=800/8,000$ on V100, $N=80,000$ on MetaX C550; cross-hardware consistency is shown for the headline no.ELIG condition (-0.006 ± 0.001 on both), and FULL differs slightly between implementations (0.482 V100 vs. 0.509 CPU five-seed), which we attribute to implementation-level noise rather than the mechanism.
9. **Convergence wording:** the $\tau=200,000$ runs are near-converged at 500k steps (residual drift +0.008 per 100k); the reported 0.561 ± 0.024 is a conservative lower bound.

9 Conclusion

Emergence requires no miracle. It requires a memory bridge between local action and global evaluation. In our spiking network, this bridge is the eligibility trace ($\tau_e = 20$ steps): remove it, and reward-driven sparsification loses its engine at every scale tested ($\Delta\text{Gini} \leq 0$ on the post-warmup baseline). Remove the brakes, and the engine runs the system off the rails; remove the accelerant, and the engine stalls only at scale.

More precisely: the eligibility trace is the necessary *engine* of reward-driven differentiation in this system. The brakes (RPE, logistic) are needed to keep the engine from destroying the system. The turbocharger (multiplicative STDP) is needed at scale. But the engine itself—the thing that converts local noise into global order—is temporal credit assignment, and nothing else we tested can replace it. Sufficiency is not claimed: we test necessity by removal, and the minimal trace-plus-reward construction remains future work.

The broader claim, scoped to systems adapting under delayed global feedback: such systems need local events to persist long enough for evaluation to bind them. Whether that memory is called an eligibility trace, a momentum buffer, or a genotype is a matter of substrate. We develop the theoretical implications of this structure for models of computation in a companion paper Ruan [2026c].

Acknowledgements

This work builds on the brain-prototype spiking network simulation framework.

Author contributions

Yahua Ruan: Conceptualization, Methodology, Software, Investigation, Formal analysis, Visualization, Writing – original draft, Writing – review & editing.

Competing interests

The author declares no competing interests.

Code availability

Analysis scripts and experimental data are publicly available at Zenodo (DOI: 10.5281/zenodo.21748099).

The simulation code is available from the corresponding author upon reasonable request.

AI-assisted technologies

During the preparation of this work the author used the following AI tools as research assistants: DeepSeek V4-Pro (GPU-vectorized eligibility-trace implementation), Qwen 3.8Max-Preview (five-condition experimental design, ablation taxonomy, and scaling-law analysis), and GLM-5.1 (SNR-threshold formalization and narrative refinement). All substantive scientific decisions, including experimental design, data analysis, and manuscript writing, were made by the human author.

Data availability

Experimental data and scripts deposited at Zenodo (DOI: 10.5281/zenodo.21748099; publicly available). This manuscript is available as a preprint (DOI: 10.5281/zenodo.21754028).

Funding

The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

References

Boaz Barak, Benjamin L Edelman, Surbhi Goel, Sham Kakade, Eran Malach, and Cyril Zhang. Hidden progress in deep learning: Sgd learns parities near the computational limit. *Advances in Neural Information Processing Systems*, 35:21750–21764, 2022.

Guillaume Bellec, Franz Scherr, Anand Subramoney, Elias Hajek, Darjan Salaj, Robert Legenstein, and Wolfgang Maass. A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*, 11:3625, 2020.

- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 41–48, 2009.
- Raphaël Bergoin, Alessandro Torcini, Gustavo Deco, Mathias Quoy, and Gorka Zamora-López. Emergence and maintenance of modularity in neural networks with Hebbian and anti-Hebbian inhibitory STDP. *PLoS Computational Biology*, 21(4):e1012973, 2025.
- Guo-Qiang Bi and Mu-Ming Poo. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience*, 18(24):10464–10472, 1998.
- Guo-Qiang Bi and Mu-Ming Poo. Synaptic modification by correlated activity: Hebb’s postulate revisited. *Annual Review of Neuroscience*, 24:139–166, 2001.
- Scott Camazine, Jean-Louis Deneubourg, Nigel R Franks, James Sneyd, Guy Theraulaz, and Eric Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, 2001.
- Mohamed Elsayed, Elena Sorina Lupu, Gautham Vasan, and A. Rupam Mahmood. Streaming deep reinforcement learning finally works. *arXiv preprint arXiv:2410.14606*, 2024.
- Răzvan V Florian. Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. *Neural Computation*, 19(6):1468–1502, 2007.
- Nicolas Frémaux and Wulfram Gerstner. Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Frontiers in Neural Circuits*, 9:85, 2016. doi: 10.3389/fncir.2015.00085.
- Uwe Frey and Richard GM Morris. Synaptic tagging: implications for late maintenance of hippocampal long-term potentiation. *Trends in Neurosciences*, 21(5):181–188, 1998. doi: 10.1016/S0166-2236(97)01189-2.
- Wulfram Gerstner and Werner M Kistler. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, 2002.

- Wulfram Gerstner, Richard Kempter, J Leo van Hemmen, and Hermann Wagner. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383(6595):76–78, 1996.
- Wulfram Gerstner, Marco Lehmann, Vasiliki Liakoni, Dane Corneil, and Johanni Brea. Eligibility traces and plasticity on behavioral time scales: Experimental support of NeoHebbian three-factor learning rules. *Frontiers in Neural Circuits*, 12:53, 2018. doi: 10.3389/fncir.2018.00053.
- Matthieu Gilson and Tomoki Fukai. Stability versus neuronal specialization for STDP: Long-tail weight distributions solve the competition. *PLoS ONE*, 6(10):e25339, 2011. doi: 10.1371/journal.pone.0025339.
- Giacomo Indiveri, Bernabe Linares-Barranco, Tara Julia Hamilton, Andre Van Schaik, Ralph Etienne-Cummings, et al. Neuromorphic silicon neuron circuits. *Frontiers in Neuroscience*, 5:73, 2011.
- Eugene M Izhikevich. Solving the distal reward problem through linkage of stdp and dopamine signaling. *Cerebral Cortex*, 17(10):2443–2452, 2007.
- Herbert Jaeger. The “echo state” approach to analysing and training recurrent neural networks. *GMD Report 148, German National Research Center for Information Technology*, 2001.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2015.
- A. Harry Klopf. *The Hedonistic Neuron: A Theory of Memory, Learning, and Intelligence*. Hemisphere Publishing, 1982.
- Robert Legenstein, Dejan Pecevski, and Wolfgang Maass. A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. *PLoS Computational Biology*, 4(10):e1000180, 2008. doi: 10.1371/journal.pcbi.1000180.
- Wolfgang Maass. Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, 10(9):1659–1671, 1997.
- Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable

states: A new framework for neural computation based on perturbations. *Neural Computation*, 14(11):2531–2560, 2002.

Carver Mead. Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10):1629–1636, 1990.

Paul A Merolla, John V Arthur, Rodrigo Alvarez-Icaza, Andrew S Cassidy, Jun Sawada, et al. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345(6197):668–673, 2014.

Aaron D. Milstein, Yiding Li, Katie C. Bittner, Christine Grienberger, Ivan Soltesz, Jeffrey C. Magee, and Sandro Romani. Bidirectional synaptic plasticity rapidly modifies hippocampal representations. *eLife*, 10:e73046, 2021.

Marvin Minsky. Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1):8–30, 1961.

Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.

Catherine Olsson, Nelson Elhage, Neel Nanda, Nova Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. In-context learning and induction heads. *Transformer Circuits Thread*, 2022.

Michael Pfeiffer and Thomas Pfeil. Deep learning with spiking neurons: Opportunities and challenges. *Frontiers in Neuroscience*, 12:774, 2018.

Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*, 2022.

Fernando M. Quintana, Fernando Perez-Peña, Pedro L. Galindo, Emre O. Neftci, Elisabetta Chicca, and Lyes Khacef. ETLF: Event-based three-factor local plasticity for online learning with neuromorphic hardware. *Neuromorphic Computing and Engineering*, 4(3):034006, 2024.

Yahua Ruan. A signal-to-noise threshold and companion mechanism in linear attention. *Neural Networks*, 2026a. Under review.

509 Yahua Ruan. Subspace direction, not rank: Scaling linear attention via sweet-spot init transfer.
510 *Neural Networks*, 2026b. Under review; preprint DOI 10.5281/zenodo.21590345.

511 Yahua Ruan. Time as a computational substrate: The eligibility trace and the first challenge to
512 the von neumann architecture. *Zenodo preprint*, 2026c. DOI 10.5281/zenodo.21754002.

513 Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. Are emergent abilities of large language
514 models a mirage? *Advances in Neural Information Processing Systems*, 36:55565–55581, 2023.

515 Wolfram Schultz, Peter Dayan, and P Read Montague. A neural substrate of prediction and reward.
516 *Science*, 275(5306):1593–1599, 1997.

517 Harel Z. Shouval and Alfredo Kirkwood. Eligibility traces as a synaptic substrate for learning.
518 *Current Opinion in Neurobiology*, 91:102978, 2025.

519 Sen Song, Kenneth D. Miller, and Larry F. Abbott. Competitive Hebbian learning through spike-
520 timing-dependent synaptic plasticity. *Nature Neuroscience*, 3:919–926, 2000. doi: 10.1038/78829.

521 Richard S Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*,
522 3(1):9–44, 1988.

523 Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT Press,
524 2018.

525 Luke Vassallo and Nima Taherinejad. Three factor delay learning rules for spiking neural networks.
526 *Frontiers in Neuroscience*, 20:1814505, 2026.

527 Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yo-
528 gatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language
529 models. *Transactions on Machine Learning Research*, 2022.

530 Sho Yagishita, Akiko Hayashi-Takagi, Timothy J. R. Ellis-Davies, Hidetoshi Urakubo, Shin Ishii,
531 and Haruo Kasai. A critical time window for dopamine actions on the structural plasticity of
532 dendritic spines. *Science*, 345(6204):1616–1620, 2014.

533 Di Yun, Zheng Wang, Shenglin Zhao, Zhenjie Wang, Huan Ma, Fei He, Jun Lu, Yuanning Li,
534 Hong Xie, and Ji-Song Guan. Distributed cortical learning through LEC-mediated γ -synchrony.
535 *Nature Communications*, 17:10396, 2026. doi: 10.1038/s41467-026-77266-w.

536 Friedemann Zenke and Tim P. Vogels. The remarkable robustness of surrogate gradient learning
537 for instilling complex function in spiking neural networks. *Neural Computation*, 33(4):899–925,
538 2021.