CONTINUOUS-TIME HETEROASSOCIATIVE MEMORY AT BIOLOGICAL TIMESCALES

Marc Gong Bacvanski
MIT, NTT Research
marcbac@mit.edu

Liu Ziyin
NTT Research, MIT
ziyinl@mit.edu

Tomaso Poggio
MIT
tp@ai.mit.edu

ABSTRACT

Associative learning in biological systems unfolds continuously in time, yet most models implicitly assume synchronized, discrete updates that instantaneously deliver the correct teaching signal to each synapse. We study a model of continuous-time heteroassociative memory neural networks in which inference and learning coevolve under coupled ODEs, and where different error-propagation topologies determine how error signals reach synapses. Longer error pathways result in propagation delays of error signals, while deeper networks result in propagation delays of cue signals. We experimentally investigate the boundary at which learning begins to fail under different combinations of propagation time constants, synaptic plasticity time constants, error signal delay Δ, and cue duration T. We find that the plasticity time constant must substantially exceed the cue duration ($\tau_{\text{plas}}/T \approx 10$–$50$ in our settings), a regime aligned with biological evidence. Together, these results yield testable predictions for associative memories in neuroscience, and practical design guidance for hardware implementations.

1 INTRODUCTION

Associative networks learn mappings from cues to desired outputs. In most algorithmic treatments, these associations are learned through discrete optimization steps that assume synchronized availability of forward activity and error signals used for weight updates. Biological learning systems, however, operate with finite propagation and integration speeds, and without separation of phases. In such settings, both cue-driven activity and error signals necessarily arrive at synapses with delays that depend on network architecture and routing pathways. Whether a network can function as an associative system therefore depends not only on the structure of its learning rule, but on the temporal alignment between cue-driven activity and routed error signals throughout the network.

In this work, we study continuous-time heteroassociative memory neural networks with local error-modulated synaptic dynamics. We model neural states and synaptic weights as jointly evolving under coupled differential equations. Learning and inference are not separated into distinct phases; instead, both proceed concurrently as signals propagate through the network and error information is routed back to synapses with nonzero delay. A cue is presented to the input layer for a finite duration, while an error signal (derived from the current output and target) is routed through a separate pathway and reaches synapses with an architecture-dependent delay. Different error propagation topologies specify how error signals reach synapses, but do not alter the local form of the learning dynamics. This formulation allows us to isolate a fundamental timing constraint that governs associative learning in continuous time.

2 COUPLED ODE DYNAMICS FOR HETEROASSOCIATION

In discrete-time associative learning, each update step implicitly assumes three sequential operations occur instantaneously: (i) the cue propagates through the network, (ii) the error is computed at the output and routed back to every synapse, and (iii) all weights update simultaneously using the correct cue–error pair. Moving to continuous time breaks all three assumptions. Neuronal states evolve with finite time constants, so cue-driven activity reaches deeper layers only after a propagation delay.

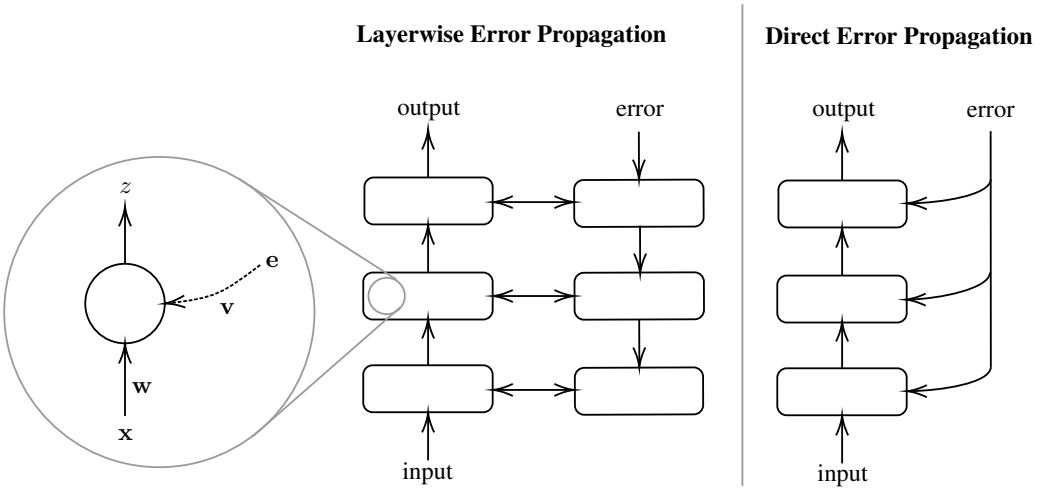

Figure 1: Neuron design with two archetypal topologies. **At left (zoom):** continuous-time heterosynaptic neuron model. Neurons receive forward input $\mathbf{x}$ and produce activated output z with weights $\mathbf{w}$. Error signals $\mathbf{e}$ enter through modulatory weights $\mathbf{v}$ and drive plasticity of $\mathbf{w}$. **In center:** layerwise error propagation topology. **At right:** direct error propagation topology.

Error signals, routed through feedback pathways, arrive at synapses with additional architecture-dependent lag. Finally, weights change continuously while these signals are still in transit, meaning that at any given moment, a synapse may be integrating a teaching signal that corresponds to the previous cue, not the current one. Whether useful learning still occurs therefore depends on whether correct cue–error pairs overlap in time at each synapse, a constraint that has no analogue in the discrete-time formulation.

In this work, we study a model of continuous-time heteroassociative neural networks trained with heterosynaptic two-signal learning rules. Our goal is to characterize when such systems can learn stable cue-to-output associations. The dynamics introduced in this section are continuous-time versions of the discrete-time heterosynaptic two-signal rules examined in Ziyin et al. (2025), which subsume the special cases studied in Lillicrap et al. (2016), Nøkland (2016), and Akrout et al. (2019). Our contributions lie in the continuous-time, delay-limited formulation, and so we will not restate or reprove the discrete-time equivalence results here.

Figure 1 illustrates the continuous-time heterosynaptic neuronal systems we study. The forward pathway, parameterized by $\mathbf{w}$ and with activation σ, maps inputs $\mathbf{x}$ to neuronal output z. The error pathway, parameterized by $\mathbf{v}$, receives an error signal $\boldsymbol{\epsilon}$ and provides a modulatory influence on learning. Learning is heterosynaptic: the forward weights $\mathbf{w}$ update in proportion to the modulatory drive $(\mathbf{v}^\top \boldsymbol{\epsilon})$, while the error weights $\mathbf{v}$ update in proportion to the forward drive $(\mathbf{w}^\top \mathbf{x})$. Let layer l have width d_l. Stacking the per-neuron variables into matrices, we can use the following notation: $\mathbf{z}_{l-1} \in \mathbb{R}^{d_{l-1}}$, $\mathbf{z}_l \in \mathbb{R}^{d_l}$, $W_l \in \mathbb{R}^{d_l \times d_{l-1}}$, and error source $\boldsymbol{\epsilon}_l$. Define

$$\dot{\mathbf{z}}_l = \frac{-\mathbf{z}_l + \sigma_l(W_l \mathbf{z}_{l-1})}{\tau_{\text{prop}}}, \quad \dot{W}_l = -\frac{W_l}{\tau^W_{\text{dec}}} + \frac{(V_l^\top \boldsymbol{\epsilon}_l)\, \mathbf{z}_{l-1}^\top}{\tau^W_{\text{plas}}}, \quad \dot{V}_l = -\frac{V_l}{\tau^V_{\text{dec}}} + \frac{(W_l \mathbf{z}_{l-1})\, \boldsymbol{\epsilon}_l^\top}{\tau^V_{\text{plas}}}. \quad (1)$$

See Appendix B for the per-neuron update rules. Here τ_{prop} is the neuronal propagation time constant, setting how quickly $\mathbf{z}_l$ relaxes to its driven input. The constants $\tau^W_{\text{plas}}, \tau^V_{\text{plas}}$ control how rapidly synaptic plasticity occurs when presynaptic and modulatory drives coincide. The constants $\tau^W_{\text{dec}}, \tau^V_{\text{dec}}$ govern passive weight decay, setting the forgetting timescale. Inputs to the network are driven by the dataset's input signals, and the error at the output neurons in layer L are $\mathbf{e} = \partial \mathcal{L} / \partial \mathbf{z}_L$. We emphasize that these equations describe rate-based dynamics, not spiking dynamics: the neuronal state $z(t)$ evolves continuously without threshold-triggered resets or discrete spikes (Hopfield, 1984; Wilson & Cowan, 1972).

During evaluation, parameters are held fixed and plasticity and decay dynamics are frozen, such that only inference dynamics $\dot{\mathbf{z}}$ are active: $\dot{W} = \dot{V} = 0$ and $\boldsymbol{\epsilon} \equiv 0$. The model's heteroassociative output is read out at the end of the input presentation window, immediately before the next cue is presented.

3 Timing Robustness

A central question is: under what conditions does a synapse receive an informative update? From the continuous-time dynamics in equation 1, the net weight change accumulated on a single synapse over a cue window of duration T can be written as a temporally filtered correlation between presynaptic activity and routed error drive:

$$\Delta W_l \;\propto\; \int_0^T \mathbf{z}_{l-1}(t)\left(V_l^\top \boldsymbol{\epsilon}_l(t)\right)^\top k_{\tau_{\text{plas}}}(t)\,dt, \qquad k_{\tau_{\text{plas}}}(t) = \exp\big(-(T-t)/\tau_{\text{plas}}\big) \tag{2}$$

The plasticity kernel $k_{\tau_{\text{plas}}}$ arises from the first-order plasticity dynamics that produce a low-pass filtering effect over the learning signal. Learning depends only on the temporal alignment between cue-driven activity and routed error signals.

When $T \ll \tau_{\text{plas}}$, the plasticity kernel $k_{\tau_{\text{plas}}}$ approaches a constant value of 1 over the duration of the cue presentation. Suppose cues are piecewise-constant with duration T, and the routed error signal to that synapse is shifted by a delay Δ relative to the onset of the cue. In this setting, equation 2 reduces to

$$\mathbb{E}[\Delta W] \;\propto\; (T - |\Delta|)_+ \tag{3}$$

where $(x)_+ = \max(x, 0)$. Therefore, the magnitude of the informative update is proportional to the temporal overlap between cue-driven activity and routed error signals. Learning fails when $|\Delta| \geq T$, where there is no overlap and updates are dominated by mismatched cue–error pairs. This predicts the delay sensitivity observed across routing topologies and network depths in the experiments following. For derivation details, see Appendix C.

4 Experiments

We simulate the continuous-time dynamics described in Section 2 using ODE integrators (details in Appendix D). For each cue of duration T, neural states and synaptic weights evolve jointly under the coupled differential equations, and error signals are routed according to the chosen topology. Performance is evaluated with weights frozen and only state dynamics active.

We evaluate three experimental regimes. Section 4.1 demonstrates that networks can robustly learn so long as error signals overlap with inputs. Section 4.2 demonstrate that deeper networks accumulate propagation lag, making them less tolerant to error delays and requiring longer sample times for stable learning. Finally, Section 4.3 shows that our model operates effectively under synaptic, plasticity, and decay timescales that align with known cortical physiology.

4.1 Direct Error Routing

In the direct routing topology, each layer receives the global output error, $\boldsymbol{\epsilon}_l = \mathbf{e}_L$. Figure 2 shows accuracy on 7×7 MNIST as cue duration T and cue–error delay Δ are varied. Learning remains robust as long as $|\Delta| < T$, and fails once the delay approaches or exceeds the cue duration. When $|\Delta| \geq T$, cue and error no longer overlap at each synapse and updates are dominated by mismatched cue–label correlations (accuracy $\sim 10\%$). Intermediate regions correspond to partial overlap, producing mixed correct and incorrect updates. With $\tau_{\text{plas}} = 10$s and $T \in [0.01, 0.25]$s, the system operates in the flat-kernel regime, yielding near-symmetric sensitivity to early and late error signals.

4.2 Layerwise Error Routing

In the layerwise routing topology, error is propagated sequentially from the output layer via $V_l^\top \boldsymbol{\epsilon}_l$. Figure 3 shows results on a nonlinear circles task whose degree of nonlinearity requires deeper networks. Deeper networks exhibit stricter timing constraints with layerwise error routing than with direct error routing. Because error must propagate through multiple layers, the effective error

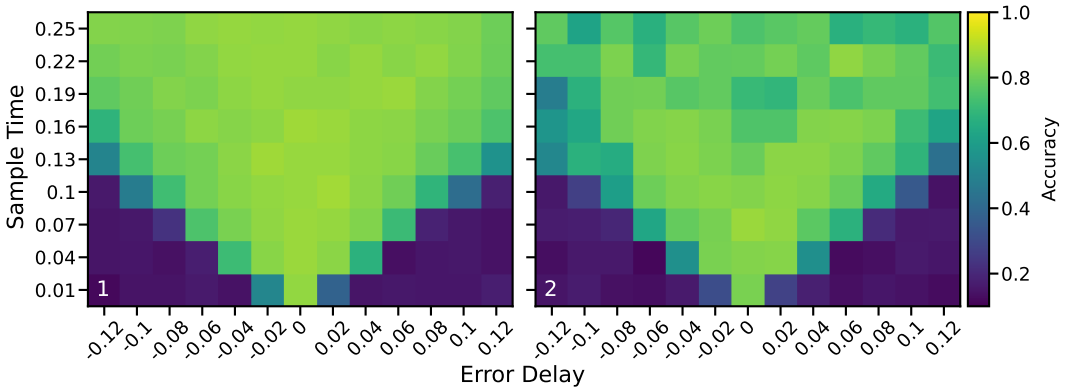

Figure 2: Evaluation of the direct error routing topology. The x-axis denotes the temporal delay between input signal and label: negative values indicate the label arrives before the input, while positive values indicate the label arrives after the input. At left is a network with 1 hidden layer of 49 neurons, and at right is a network with 2 hidden layers of 49 and 32 neurons. Learning is fairly robust until the delay exceeds the sample time.

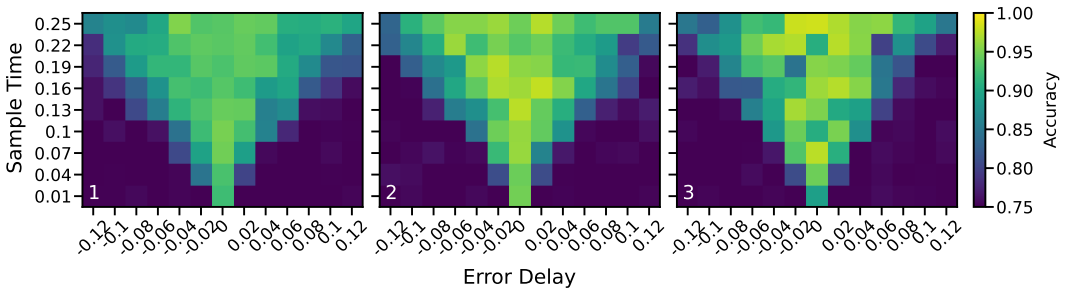

Figure 3: Evaluation of the layerwise error routing topology on the circle dataset as error signal delay and sample duration are swept. The network topologies increase in the number of hidden layers, where each added hidden layer has 24 neurons. From left to right, the networks have one hidden layer of 24 ReLU neurons, two hidden layers of 24 ReLU neurons, and three hidden layers of 24 ReLU neurons. Layerwise error routing imposes stricter requirements on delay due to longer error propagation paths.

signal delay at synapses accumulates, reducing cue–error overlap particularly in early layers. As predicted by the overlap law equation 3, increasing cue duration T mitigates this effect by enlarging the overlap window, whereas a large error delay $|\Delta|$ eliminates overlap and causes complete failure.

4.3 Biological Timescales

A central motivation for studying continuous-time associative networks is that biological neural circuits operate across multiple nested timescales: millisecond synaptic transmission, second-scale plasticity processes, and slower homeostatic decay. Here, we evaluate our model using timescale parameters reflective of the ranges found in biology.

Figure 4 sweeps the plasticity time constant τ_{plas} while holding the cue duration fixed at $T = 50$ ms. Learning becomes robust once $\tau_{\text{plas}} \gtrsim 2$ s. This corresponds to a ratio $\tau_{\text{plas}}/T \approx 40$, indicating that effective learning requires plasticity dynamics that substantially outlast cue presentation.

These results align with the hierarchy $\tau_{\text{prop}} \ll \tau_{\text{plas}} \ll \tau_{\text{dec}}$ observed in cortical circuits. In this regime, fast state dynamics track the cue, plasticity integrates cue–error coincidence over seconds, and slow decay sets a baseline weight timescale without disrupting learning.

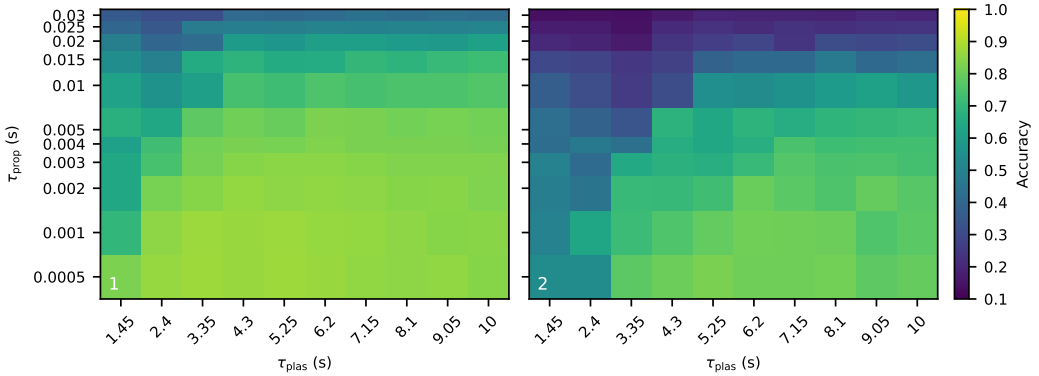

Figure 4: Evaluation of a layerwise error routing network on the 7×7 MNIST dataset. **Left:** 1 hidden layer (49 neurons). **Right:** 2 hidden layers (49 and 32 neurons). Learning is unstable when the plasticity timescale τ_{plas} is comparable to the presentation window ($T = 50\,\text{ms}$), but becomes robust only once τ_{plas} exceeds ~ 2 s. This corresponds to $\tau_{\text{plas}}/T \approx 40$, highlighting the requirement that plasticity persist far longer than the input presentation time.

5 Discussion

We analyzed a model of continuous-time heteroassociative networks in which inference and learning coevolve under coupled ODEs. In this setting, synaptic change reduces to a temporally filtered cross-correlation between presynaptic behavior and a routed error signal. In the continuous-time setting, different "credit assignment" architectures differ primarily in how much temporal overlap they preserve across layers. From an associative memory viewpoint, this reframes depth not as an abstract optimization difficulty, but as a physical timing budget: deeper memories require either longer cue windows, faster signal and error pathways, or architectural shortcuts that reduce effective lag.

The overlap perspective also isolates a timescale requirement for weight potentiation required for stable association learning. The synaptic update is weighted by a plasticity kernel due to the low pass dynamics of the time constant τ_{plas}; if this kernel varies substantially over a single cue's presentation, such as if τ_{plas} is on the order of the cue presentation duration T, then early and late portions of the cue–error coincidence are differentially weighted, increasing sensitivity to timing jitter and mismatch. Robust learning therefore requires weight plasticity dynamics to run much slower than signal and error propagation and cue duration. Empirically, we observe reliable learning only once τ_{plas}/T reaches tens (roughly 10-50 in our settings), for which cortical-like T in the tens of milliseconds places the effective eligibility regime in the seconds range. This is consistent with a broad class of biochemical and synaptic-tag processes that integrate coincidence over seconds, and suggests a simple interpretation: continuous-time heteroassociative learning needs a slow memory of recent coincidence to average away propagation and integration delays.

Finally, these results suggest a compact design principle for both neuroscience and hardware: the key requirement is not weight symmetry or a particular feedback rule, but the preservation of temporal overlap between cue-related activity and teaching signals at each synapse. A hierarchy of timescales $\tau_{\text{prop}} \ll \tau_{\text{plas}} \ll \tau_{\text{dec}}$ supports this relationship. This principle also points to the role of long-range feedback and shortcut pathways as mechanisms for maintaining temporal overlap in deep networks.

References

Mohamed Akrout, Collin Wilson, Peter Humphreys, Timothy Lillicrap, and Douglas B Tweed. Deep learning without weight transport. *Advances in neural information processing systems*, 32, 2019.

Atilim Gunes Baydin, Barak A Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey. *Journal of machine learning research*, 18(153):1–43, 2018.

Yoshua Bengio and Asja Fischer. Early inference in energy-based models approximates back-propagation. *arXiv preprint arXiv:1510.02777*, 2015.

Zuzanna Brzosko, Wolfram Schultz, and Ole Paulsen. Retroactive modulation of spike timing-dependent plasticity by dopamine. *elife*, 4:e09685, 2015.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

Alain Destexhe, Zachary F Mainen, Terrence J Sejnowski, et al. Kinetic models of synaptic transmission. *Methods in neuronal modeling*, 2:1–25, 1998.

Michael D Ehlers. Reinsertion or degradation of ampa receptors determined by activity-dependent endocytic sorting. *Neuron*, 28(2):511–525, 2000.

Michael D Ehlers. Activity level controls postsynaptic composition and signaling via the ubiquitin-proteasome system. *Nature neuroscience*, 6(3):231–242, 2003.

Benjamin Ellenberger, Paul Haider, Jakob Jordan, Kevin Max, Ismael Jaras, Laura Kriener, Federico Benitez, and Mihai A Petrovici. Backpropagation through space, time, and the brain. *arXiv preprint arXiv:2403.16933*, 2024.

Jörg RP Geiger, Joachim Lübke, Arnd Roth, Michael Frotscher, and Peter Jonas. Submillisecond ampa receptor-mediated signaling at a principal neuron–interneuron synapse. *Neuron*, 18(6): 1009–1023, 1997.

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.

Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256. JMLR Workshop and Conference Proceedings, 2010.

Paul Haider, Benjamin Ellenberger, Laura Kriener, Jakob Jordan, Walter Senn, and Mihai A Petrovici. Latent equilibrium: A unified learning theory for arbitrarily fast computation with arbitrarily slow neurons. *Advances in neural information processing systems*, 34:17839–17851, 2021.

John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.

John J Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the national academy of sciences*, 81(10):3088–3092, 1984.

A Kapur, RA Pearce, WW Lytton, and LB Haberly. Gabaa-mediated ipscs in piriform cortex have fast and slow components with different properties and locations on pyramidal cells. *Journal of neurophysiology*, 78(5):2531–2545, 1997.

Patrick Kidger. *On Neural Differential Equations*. PhD thesis, University of Oxford, 2021.

David Koplow, Tomaso Poggio, and Liu Ziyin. Emergence of hebbian dynamics in regularized non-local learners. *arXiv preprint arXiv:2505.18069*, 2025.

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

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 2002.

Qianli Liao, Liu Ziyin, Yulu Gan, Brian Cheung, Mark Harnett, and Tomaso Poggio. Self-assembly of a biologically plausible learning circuit. *arXiv preprint arXiv:2412.20018*, 2024.

Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7(1): 13276, 2016.

Timothy P Lillicrap, Adam Santoro, Luke Marris, Colin J Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6):335–346, 2020.

Roger A Nicoll and Howard Schulman. Synaptic memory and camkii. *Physiological reviews*, 103 (4):2897–2945, 2023.

Arild Nøkland. Direct feedback alignment provides learning in deep neural networks. *Advances in neural information processing systems*, 29, 2016.

Richard J O'Brien, Sunjeev Kamboj, Michael D Ehlers, Kenneth R Rosen, Gerald D Fischbach, and Richard L Huganir. Activity-dependent modulation of synaptic ampa receptor accumulation. *Neuron*, 21(5):1067–1078, 1998.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.

João Sacramento, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. Dendritic cortical microcircuits approximate the backpropagation algorithm. *Advances in neural information processing systems*, 31, 2018.

Benjamin Scellier and Yoshua Bengio. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in computational neuroscience*, 11:24, 2017.

Ch Tsitouras. Runge–kutta pairs of order 5 (4) satisfying only the first column simplifying assumption. *Computers & mathematics with applications*, 62(2):770–775, 2011.

Gina G Turrigiano, Kenneth R Leslie, Niraj S Desai, Lana C Rutherford, and Sacha B Nelson. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391(6670):892–896, 1998.

James CR Whittington and Rafal Bogacz. An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. *Neural computation*, 29(5): 1229–1262, 2017.

James CR Whittington and Rafal Bogacz. Theories of error back-propagation in the brain. *Trends in cognitive sciences*, 23(3):235–250, 2019.

Hugh R Wilson and Jack D Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal*, 12(1):1–24, 1972.

Xiaohui Xie and H Sebastian Seung. Equivalence of backpropagation and contrastive hebbian learning in a layered network. *Neural computation*, 15(2):441–454, 2003.

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

Liu Ziyin, Isaac Chuang, and Tomaso Poggio. Heterosynaptic circuits are universal gradient machines. *arXiv preprint arXiv:2505.02248*, 2025.

A Related Work

Related families of learning algorithms, such as contrastive Hebbian learning (Xie & Seung, 2003), equilibrium propagation (Scellier & Bengio, 2017), and other energy-based formulations (Hopfield, 1982; Bengio & Fischer, 2015), also instantiate continuous-time learning rules. Latent-equilibrium networks and their extensions derive both neural and synaptic dynamics as gradient flows so that arbitrarily slow neurons can approximate backpropagation (and backpropagation through time) without explicit forward/backward phases (Haider et al., 2021; Ellenberger et al., 2024). These approaches differ from the heterosynaptic two-signal rules that are our focus here. While our model shares the spirit of framing learning as a dynamical process, we restrict our focus in this work to error-propagation style rules (FA, DFA, KP) and their continuous-time realizations.

Whittington and Bogacz (Whittington & Bogacz, 2019) review biologically plausible backpropagation schemes, but emphasize predictive-coding and dendritic-error frameworks rather than weight-transport-free methods such as feedback alignment or Kolen–Pollack. Predictive coding (Whittington & Bogacz, 2017) uses continuous-time ODEs for neuronal states but treats error terms and weight updates as instantaneous algebraic quantities. Similarly, dendritic error models (Sacramento et al., 2018) compute apical errors by algebraic functions of current somatic activity, which is distinct from feedback alignment. Our work is orthogonal to these directions: we analyze feedback alignment type rules (FA, DFA, and KP) in layered feedforward networks, showing that these discrete-time algorithms admit a continuous-time formulation in which neural states and synaptic weights coevolve under coupled ODEs.

Both neural ODEs (Chen et al., 2018) and our neural differential-equation model cast network computation as a continuous-time dynamical system, replacing discrete layers/updates with ODE flows over time. In standard neural ODEs, parameters are fixed during the forward solve, and gradients are typically recovered by integrating an adjoint ODE backward in time (Chen et al., 2018) or by differentiating through the solver (Baydin et al., 2018). By contrast, our model couples learning and inference in one forward-in-time system: both neural states and parameters evolve by ODEs, with weights updated online via locally computed and propagated error terms.

Concurrently, "self-assembling" heterosynaptic circuits specifying four plastic synapses linking forward and feedback streams have been shown to match backpropagation on standard benchmarks (Liao et al., 2024). A key insight from this line of work is that error-propagation rules such as feedback-style algorithms (FA, DFA, KP) can emerge without requiring explicit weight symmetry, and only using local rules. This further motivates our focus on these algorithms as being attractive candidates for biologically plausible and hardware-realizable learning.

Biological relevance of feedback-type algorithms has been argued in various works. For example, Lillicrap et al. (2020) argued how FA may be realizable in the cerebellum, where the Purkinje cells are the trainable neurons, and the one-to-one climbing fibres carry in the feedback error signals. Koplow et al. (2025) computationally showed how these algorithms are consistent with the widely observed Hebbian and anti-Hebbian plasticities in the brain. However, it is not yet clear whether these algorithms can work in continuous-time, and if so, whether they work at biologically reasonable timescales. This is a key question we address.

B Dynamics of a Single Neuron

The single-neuron version of the dynamics presented in equation 1:

$$\dot{z} = \frac{-z + \sigma(\mathbf{w}^\top \mathbf{x})}{\tau_{\text{prop}}}, \quad \dot{\mathbf{w}} = -\frac{\mathbf{w}}{\tau^W_{\text{dec}}} + \frac{(\mathbf{v}^\top \boldsymbol{\epsilon})\,\mathbf{x}}{\tau^W_{\text{plas}}}, \quad \dot{\mathbf{v}} = -\frac{\mathbf{v}}{\tau^V_{\text{dec}}} + \frac{(\mathbf{w}^\top \mathbf{x})\,\boldsymbol{\epsilon}}{\tau^V_{\text{plas}}}.$$

C Temporal Overlap and Informative Updates

A central question is: *under what conditions does a synapse receive a correct update?* Two key considerations are (i) the temporal mismatch between input and error signals, and (ii) the rate at which the input changes. Figure 5 illustrates the dynamics of a final-layer weight when the error signal arrives earlier or later than the corresponding input. During the mismatch period, the instantaneous weight update is incorrect, leading to a cumulative update that deviates from the non-delayed case.

To analyze robustness, we isolate only the part of the weight update that is informative: that which is proportional to the correlation between presynaptic activity and the matching error drive. Vectorized, the informative weight change accumulated over a presentation window of length T is

$$\Delta W_l \;\propto\; \int_0^T \mathbf{z}_{l-1}(t)\left(V_l^\top \boldsymbol{\epsilon}_l(t)\right)^\top k_{\tau_{\text{plas}}}(t)\; dt, \qquad k_{\tau_{\text{plas}}}(t) = \exp\big(-(T-t)/\tau_{\text{plas}}\big), \quad (4)$$

where the causal exponential kernel arises from the low-pass plasticity dynamics and weights more recent coincidence more strongly. For a single synapse $(i \to j)$, this reduces to the scalar form

$$\Delta(W_l)_{ij} \;\propto\; \int_0^T (\mathbf{z}_{l-1}(t))_i \left((V_l)_{:,j}^\top \boldsymbol{\epsilon}_l(t)\right) k_{\tau_{\text{plas}}}(t)\; dt. \quad (5)$$

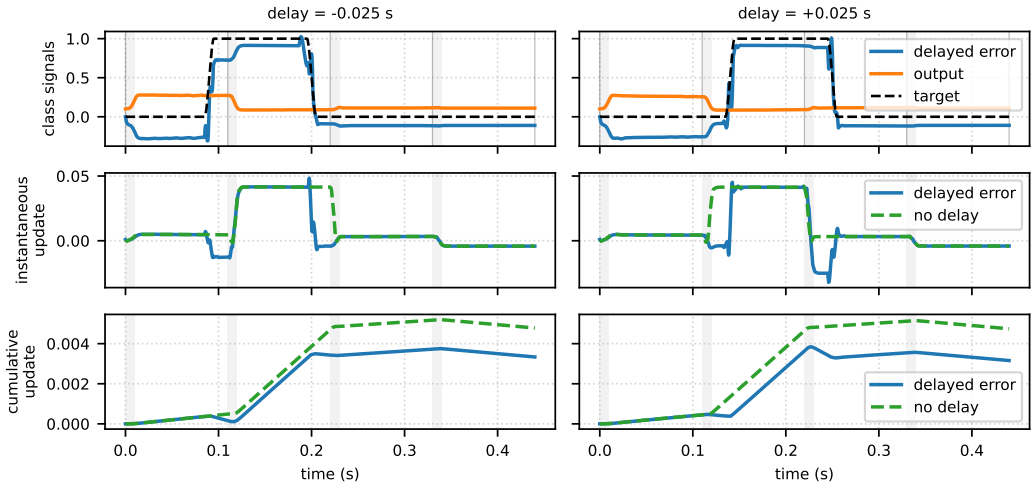

Figure 5: Single-neuron dynamics with different relative timings of the error and input signals. **Left:** error signal arrives early. **Right:** error signal is delayed. The bottom panel of both plots depicts the cumulative weight change and shows that in the presence of delay, the weight accumulates a biased gradient update compared to the case where there was no error delay.

This makes clear that learning depends on the temporal cross-correlation between the presynaptic drive $(\mathbf{z}_{l-1})_i$ and the local modulatory/error drive at neuron j.

Piecewise-constant inputs with delay. Assume $(\mathbf{z}_{l-1})_i(t)$ is active on $[0, T]$ and the error drive is active on $[\Delta, \Delta + T]$ (same duration, delayed by Δ). In the fast-propagation limit $\tau_{\text{prop}} \ll \tau_{\text{plas}}$, the expected update becomes

$$\begin{aligned}\mathbb{E}[\Delta(W_l)_{ij}] &\propto \int_{t_0}^{t_1} \exp\big((t-T)/\tau_{\text{plas}}\big)\, dt \\ &= \tau_{\text{plas}}\Big(e^{-(T-t_1)/\tau_{\text{plas}}} - e^{-(T-t_0)/\tau_{\text{plas}}}\Big) = \tau_{\text{plas}}\, e^{-(T-t_1)/\tau_{\text{plas}}}\Big(1 - e^{-L/\tau_{\text{plas}}}\Big),\end{aligned} \tag{6}$$

where $t_0 = \max(0, \Delta)$, $t_1 = \min(T, \Delta + T)$, and $L = t_1 - t_0 = (T - |\Delta|)_+$.

Flat-kernel limit and prediction. When $T \ll \tau_{\text{plas}}$ (the regime standard in our experiments), $k_{\tau_{\text{plas}}}$ is approximately constant over $[0, T]$, so equation 6 reduces to the symmetric triangular law

$$\mathbb{E}[\Delta(W_l)_{ij}] \;\propto\; (T - |\Delta|)_+. \tag{7}$$

Thus learning succeeds if and only if input and error *overlap in time*, and it degrades sharply as $|\Delta| \to T$. When T approaches τ_{plas}, the exact expression equation 6 predicts a mildly *skewed triangle* that up-weights late-arriving errors (positive Δ) relative to equally early ones; the skew vanishes continuously as $\tau_{\text{plas}}/T \to \infty$.

A fixed overlap budget. For delay Δ, define the correct-overlap set $C(\Delta) = [0, T] \cap [\Delta, \Delta + T]$ of length $L = (T - |\Delta|)_+$ and mismatched set $I(\Delta) = [0, T] \setminus C(\Delta)$. With plasticity kernel $k(t)$, let

$$K_C(\Delta) = \int_{C(\Delta)} k(t)\, dt, \quad K_I(\Delta) = \int_{I(\Delta)} k(t)\, dt,$$

so that $K_C(\Delta) + K_I(\Delta) = K_T := \int_0^T k(t)\, dt$. Thus overlap and mismatch trade off under a fixed budget. In the flat-kernel case $k \equiv 1$, this reduces to $L + (T - L) = T$, yielding $\mathbb{E}[\Delta(W_l)_{ij}] \propto L = (T - |\Delta|)_+$, i.e. accuracy improves with either larger T or smaller $|\Delta|$.

Our analysis predicts and experiments confirm near-symmetry between early and late error in the flat-kernel regime. Incorporating an explicit eligibility gate would recover causal learning behavior without altering our conclusions about temporal overlap and timescale separation; we leave such extensions for future work.

D DATASETS & METHODOLOGY

We use two datasets in our evaluation of these models. The 7×7 downsampled MNIST dataset is the standard MNIST (LeCun, 1998) dataset that has been downsampled with 4×4 average pooling. The circles dataset is the `make_circles` dataset from scikit-learn (Pedregosa et al., 2011). It is well established that simple linear classifiers achieve surprisingly high accuracy on MNIST, with only modest improvements from deeper architectures (LeCun et al., 2002). We find a linear softmax regressor on 7×7 downsampled MNIST achieves over 89% test accuracy. In contrast, a logistic regressor on the 2-circle concentric-rings dataset achieves only 75% test accuracy, which corresponds to the class imbalance itself (75% outer ring, 25% inner ring).

W are initialized according Xavier normalization (Glorot & Bengio, 2010). V are initialized to a fixed constant 0.1. Classification decisions are read out from output neurons at the very end of each input sample's presentation. Evaluation is done on frozen W and V dynamics. Each heatmap data point is the average of 3-5 runs, depending on the experiment.

We integrate all continuous-time dynamics using Diffrax's (Kidger, 2021) Tsit5 solver (Tsitouras, 2011) (a fifth-order explicit Runge–Kutta method with an embedded fourth-order error estimate) equipped with a PID adaptive step-size controller ($\texttt{rtol} = 2 \times 10^{-3}$, $\texttt{atol} = 10^{-5}$). This setup allows the solver to take large steps during slowly varying segments of the dynamics while automatically refining steps around rapid transients induced by input switches and error onsets. Although the timescales of our systems span several orders of magnitude, we empirically find that they are not so stiff as to require an implicit method; Tsit5 remains stable and efficient under these tolerances. Compared to a fixed-step forward Euler integrator tuned to resolve the fastest timescale, this adaptive scheme reduces wall-clock time by several orders of magnitude while producing indistinguishable learning curves and dynamics.

E BIOLOGICAL TIMESCALES

We parameterize our continuous-time networks with biologically motivated constants and show that they learn effectively on timescales observed in the brain. Of the three time constants relevant in our model, τ_{plas} is the least constrained biologically. Our results therefore provide a new theoretical prediction that narrows its plausible functional range.

A central result of our study is that robust learning requires plasticity windows that outlast the stimulus duration by at least an order of magnitude, placing τ_{plas} firmly in the few-second range. This prediction is biologically plausible and experimentally testable, and it narrows the functional range of τ_{plas}, which has so far been the least constrained timescale in biology. Our simulations demonstrate that the hierarchy $\tau_{\text{prop}} \ll \tau_{\text{plas}} \ll \tau_{\text{dec}}$, characteristic of cortical tissue, is sufficient to support effective learning.

In our model, τ_{prop} corresponds to the dominant synaptic conductance time constant, τ_{plas} to the biochemical induction gate during which coincident presynaptic drive and modulatory/error input can trigger plasticity, and τ_{dec} to slow synaptic weakening. Once $\tau_{\text{dec}} \gg \tau_{\text{prop}}, \tau_{\text{plas}}$, it primarily sets a slow baseline for weight decay rather than shaping learning dynamics. Table 1 summarizes these biophysical processes.

Table 1: Biophysical timescales in continuous-time network models. Time constants correspond to distinct molecular/cellular mechanisms: rapid receptor/channel kinetics (τ_{prop}), seconds-scale intracellular signaling and receptor trafficking (τ_{plas}), and slow homeostatic regulation (τ_{dec}).

	Corresponding Biophysical Process	**Biologically Plausible Range**	**Successful Range (Fig 4)**
τ_{prop}	Fast synaptic transmission determined by receptor/channel kinetics and membrane RC filtering such as AMPA/$GABA_A$ receptor conductance decay after vesicular glutamate/GABA release (Destexhe et al., 1998; O'Brien et al., 1998).	2–30ms in cortex (Kapur et al., 1997); sub-ms in auditory brainstem synapses (Geiger et al., 1997).	Swept over 5–30ms, best learning with $\tau_{\text{prop}} < 20$ms.
τ_{plas}	Coincidence-gated plasticity via second-messenger cascades (dopamine D1/D2→cAMP/PKA, Ca^{2+} →CaMKII) (Nicoll & Schulman, 2023), regulating AMPAR phosphorylation and trafficking. Defines the biochemical "induction gate" during which pre/post and modulatory signals interact (Yagishita et al., 2014; Gerstner et al., 2018).	∼0.3–10s depending on circuit (striatal vs. cortical/hippocampal). Some studies find an effective eligibility trace for dopamine-driven LTP lasts 1–2 minutes, while weaker effects can persist up to 10 minutes (Brzosko et al., 2015).	1.45s to 10s, with best learning in $\tau_{\text{plas}} \gtrsim 2$s.
τ_{dec}	Slow synaptic weakening via protein turnover, phosphatase activity, and homeostatic scaling (Turrigiano et al., 1998; Ehlers, 2003) (e.g., AMPAR endocytosis (Ehlers, 2000), transcriptional regulation).	Minutes to tens of minutes.	$\sim$ 20 minutes. We find that as long as τ_{dec} is sufficiently larger than $\tau_{\text{prop}}, \tau_{\text{plas}}$, it plays little role in the learning dynamics.