Credit Bandwidth Lower Bounds for Diffusive Cortical Learning

Aarav Sinha

aarav@syrg.org

Independent Researcher, SYRG https://orcid.org/0009-0001-8238-9309

Research Article

Keywords: credit assignment, recurrent neural networks, graph signal processing, graph Sobolev spaces, neuromodulation, biologically plausible learning, feedback bandwidth, cortical learning, constrained optimization, spectral graph theory

Posted Date: May 29th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9841779/v1

Additional Declarations: The authors declare no competing interests.

Credit Bandwidth Lower Bounds for Diffusive Cortical Learning

Aarav Sinha
aarav@syrg.org

Abstract

Exact recurrent learning assigns neuron-specific adjoint credit fields, whereas cortical plasticity appears to combine local eligibility traces with low-bandwidth diffuse and cell-type-specific feedback. We formalize this mismatch as communication-constrained credit assignment on a cortical graph. A K-source, S-tap modulatory architecture induces an admissible feedback subspace of effective dimension $D \leq KS$. For graph-Sobolev adjoint fields of smoothness s, the minimax squared credit error of any D-dimensional rule is exactly $R^2(1+\mu_{D+1})^{-s}$, achieved by the first D graph eigenmodes. We then lift this width law to recurrent learning: in a stable graph-filter RNN driven by stationary stochastic input, states have uniformly bounded second moment, all T recurrent gradient slices are nonzero, and the optimal constrained full-gradient error is $A_T R^2(1+\mu_{D+1})^{-s}$ with $A_T = \Theta(T)$. The obstruction persists under trajectory-dependent modulatory coefficients and finite code repertoires; in a formal model of q noiseless scalar side observations, it shifts by at most q graph frequencies. A measured-adjoint experiment supports the diagnostic prediction: graph-spectral tails, unlike mismatched random tails, predict constrained-feedback loss gaps. The theory yields a falsifiable channel-budget rule: estimating graph-Fourier adjoint smoothness s and effective feedback dimension D_{eff} predicts which tasks incur limited-feedback loss floors. On d-dimensional manifold-like graphs, the squared credit, state-normalized gradient, and PL optimization-floor errors scale as $D^{-2s/d}$.

1 Introduction

Credit assignment in recurrent neural systems has two incompatible-looking descriptions. In machine learning, backpropagation through time (BPTT) and real-time recurrent learning (RTRL) compute exact gradients by transporting neuron-specific adjoint information through space and time [Werbos, 1990, Williams and Zipser, 1989]. In biological networks, plasticity is local at synapses and is often modeled as an eligibility trace multiplied by a learning or modulatory signal [Bellec et al., 2020, Lillicrap et al., 2020]. Such signals are not private backward channels for each neuron. Neuromodulators and neuropeptides can be spatially diffuse, cell-type-specific, and mediated by receptor expression patterns rather than point-to-point adjoint transport [Smith et al., 2019, Özçete et al., 2024, Liu et al., 2021].

This paper asks a normative question rather than proposing another mechanism. Suppose the only globally available feedback is a small number of diffuse or cell-type-specific modulatory channels. What is the best possible approximation to exact recurrent credit assignment, and when must every such rule fail?

We model neurons as vertices of a weighted cortical graph G with graph Laplacian L_G. Exact recurrent learning produces an adjoint credit field $\lambda_t \in \mathbb{R}^N$ over neurons. A biologically constrained rule may update a synapse using its local eligibility trace, but the postsynaptic credit field must lie in a low-dimensional admissible subspace $\mathcal{H} \subset \mathbb{R}^N$. Diffusion, volume transmission, and cell-type-

Preprint submitted to Research Square.

specific sources all instantiate such subspaces: if P is a local diffusion operator and b_k are source or receptor patterns, then fields generated by K sources and S diffusion taps lie in

$$\mathcal{H}_{K,S}(P,B) = \operatorname{span}\{P^r b_k : 0 \le r < S,\ 1 \le k \le K\}, \qquad \dim(\mathcal{H}_{K,S}) \le KS. \tag{1}$$

We therefore distinguish biological source count K from effective feedback dimension $D = \dim(\mathcal{H})$. This distinction is essential: a theorem stated in terms of K alone is wrong once each channel can contribute multiple independent diffusion taps.

The design principle is that recurrent credit assignment has a feedback-bandwidth budget. Diffuse low-dimensional feedback is well matched to tasks whose exact adjoints are graph-low-frequency; it fails on tasks whose credit alternates across nearby neurons, cell types, or cortical locations. Escaping the failure mode requires increasing feedback dimension, increasing independent source/filter patterns, or learning a representation in which the same task has smoother credit. Biologically, the prediction is that low-dimensional modulatory systems should work best when task credit is smooth across maps or cell classes, and should struggle when nearby sites need opposite credit signs unless an additional pathway, cell-type code, or representational remapping carries that distinction.

Contributions. The first part of the paper isolates the signal class. Proposition 3 shows that K modulatory sources with S fixed diffusion, temporal, or cell-type taps induce a fixed code of dimension at most KS. Theorem 1 is the recurrent credit-factorization step: for the actual recurrent-weight gradient, the best constrained rule is to project the postsynaptic adjoint factor onto this code; cross-time hidden-state correlations cannot turn an omitted credit component into an admissible one.

The second part proves the sharp obstruction. Theorem 2 gives the graph-Sobolev width $R^2(1+\mu_{D+1})^{-s}$. Theorem 3 lifts it to a genuinely multi-step stable RNN: a stationary graph-filtered input drives bounded-energy states at every time, exact adjoints persist across the horizon, the gradient contains T nonzero rank-one eligibility slices, and the optimal full-gradient lower bound is $A_T R^2(1+\mu_{D+1})^{-s}$ with $A_T = \Theta(T)$. Thus the recurrence is not decorative; the spectral tail is accumulated by recurrent eligibility over time.

The third part turns the bound into predictions. Adaptive-feedback lower bounds show that a finite repertoire of M codes is limited by the $(MD+1)$st frequency; in a deliberately strong model, q noiseless scalar observations of the credit field shift the obstruction by at most q frequencies. On manifold-like graphs, $\mu_D \asymp D^{2/d}$ gives the channel-budget law $D^{-2s/d}$. Section 7 states a measurable protocol: fit s from graph-Fourier adjoint spectra, estimate an effective feedback dimension D_{eff}, and test whether the loss floor follows the predicted spectral tail rather than a dimension-only feedback-alignment baseline.

2 Related work and gap

Backpropagation, BPTT, and RTRL give exact gradients but require backward-in-time storage or nonlocal communication [Rumelhart et al., 1986, Werbos, 1990, Williams and Zipser, 1989]. Biologically plausible learning work studies alternative feedback pathways, dendritic microcircuits, predictive coding, feedback alignment, e-prop, and cell-type-specific neuromodulation [Lillicrap et al., 2016, Nøkland, 2016, Sacramento et al., 2018, Whittington and Bogacz, 2017, Meulemans et al., 2021, Bellec et al., 2020, Liu et al., 2021, 2022, Kaleb et al., 2024, Ellenberger et al., 2026]. These works motivate our admissible class, but they do not characterize the minimax-optimal credit field under a finite feedback-bandwidth constraint.

Our graph-spectral analysis uses standard facts from spectral graph theory and graph signal processing [Chung, 1997, Shuman et al., 2013]. The minimax proof is a finite-dimensional Kolmogorov-width calculation over a graph-Sobolev ellipsoid [Pinkus, 1985]. By itself, that calculation says only that arbitrary smooth graph signals have an optimal D-dimensional spectral compression. The recurrent contribution is to locate this compression bottleneck inside the exact three-factor gradient: admissible diffuse feedback restricts only the postsynaptic adjoint factor, projection remains optimal for the actual Frobenius recurrent-gradient error despite cross-time hidden-state correlations, and bounded stable graph-local RNNs realize the omitted modes as genuine adjoints. The adaptive-feedback and optimization results then show how this factor bottleneck changes under finite repertoires, explicitly modeled side-channel observations, and local parameter-space progress.

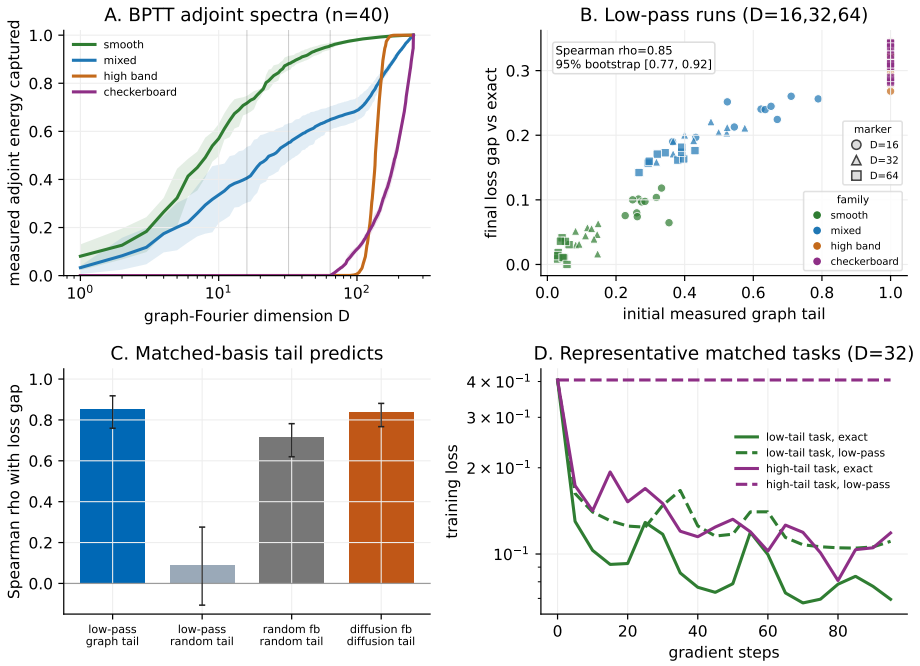

Figure 1: Measured-adjoint learning test of the credit-bandwidth prediction. On a 16×16 grid graph, we generate 40 matched-energy graph-sequence tasks with losses at every time step, measure the exact BPTT adjoints on the initial trajectory, and train the same RNN with exact BPTT, graph-low-pass feedback, random same-dimensional feedback, and K-source diffusion feedback for $D \in \{16, 32, 64\}$. Measured graph-Fourier adjoint tails separate smooth, mixed, high-band, and checkerboard task families and predict the final low-pass loss gap versus exact BPTT (Spearman $\rho = 0.85$, 95% bootstrap interval $[0.77, 0.92]$). Panel C compares predictors: tail measured in the feedback basis has high predictive power, while tail measured in a mismatched random basis is weak for the same low-pass runs. Panel D shows representative matched low-tail and high-tail tasks, not averages over tasks.

3 Setup: exact adjoints and admissible modulatory feedback

3.1 Cortical graph and graph-Sobolev credit fields

Let $G = (V, E, w)$ be an undirected weighted graph with $|V| = N$ neurons. Let $L_G \in \mathbb{R}^{N \times N}$ be a symmetric graph Laplacian with orthonormal eigenvectors $u_1, \dots, u_N$ and eigenvalues

$$0 = \mu_1 \le \mu_2 \le \cdots \le \mu_N, \qquad L_G u_\ell = \mu_\ell u_\ell. \tag{2}$$

For $s \ge 0$, define

$$\|f\|_{H_G^s}^2 = \sum_{\ell=1}^{N} (1 + \mu_\ell)^s |\langle f, u_\ell \rangle|^2, \qquad \mathcal{B}_s(R) = \{f \in \mathbb{R}^N : \|f\|_{H_G^s} \le R\}. \tag{3}$$

Large s means that admissible exact credit fields have little energy in high graph frequencies. This is an assumption on a task family, not a claim that all cortical adjoints are smooth. The assumption is checkable, however: graph-smooth residuals remain graph-smooth when propagated backward through stable graph-filter dynamics.

Proposition 1 (Smoothness inheritance under graph-filter dynamics)**.** *Suppose $W = p(L_G)$ is a symmetric graph filter with $\sup_\ell |p(\mu_\ell)| \le \gamma < 1$. Let*

$$\lambda_t = z_t + W\lambda_{t+1}, \qquad \lambda_{T+1} = 0, \tag{4}$$

where $z_t = C^\top r_t$ is the instantaneous output-residual drive. Then, for every $s \ge 0$,

$$\|\lambda_t\|_{H_G^s} \le \sum_{q=t}^{T} \gamma^{q-t} \|z_q\|_{H_G^s}. \tag{5}$$

If $\|z_q\|_{H^s_G} \le R_z$ *for all* q*, then* $\|\lambda_t\|_{H^s_G} \le R_z/(1-\gamma)$.

Thus the Sobolev credit class follows, for example, from smooth targets or smooth local residuals combined with contractive graph-filter recurrence. The lower bounds ask what finite-dimensional diffuse feedback can guarantee once the exact adjoints are in this natural inherited class; they do not assume that arbitrary tasks are smooth. The proof is in Appendix A.5.

Proposition 2 (Sobolev-stable nonlinear adjoints)**.** *Consider any differentiable or surrogate-differentiable recurrent model whose adjoint has the form*

$$\lambda_t = z_t + J_t^\top \lambda_{t+1}, \qquad \lambda_{T+1} = 0, \tag{6}$$

where J_t *is the forward Jacobian along the realized trajectory. If the backward Jacobians are graph-Sobolev stable,*

$$\left\|J_t^\top f\right\|_{H^s_G} \le \gamma \|f\|_{H^s_G} \quad \text{for all } t, f, \text{ with } \gamma < 1, \tag{7}$$

then the bound in (5) *holds with* W^{q-t} *replaced by* $J_t^\top \cdots J_{q-1}^\top$.

The graph-filter assumption in Proposition 1 is therefore a transparent sufficient condition, not the boundary of the theory. For a rate or ReLU RNN, condition (7) is the finite, auditable matrix norm $\|(I+L_G)^{s/2} J_t^\top (I+L_G)^{-s/2}\|_{\mathrm{op}} \le \gamma$ along the trajectory. It allows diagonal cell-type gains, gates, and post-activation factors when their measured action does not amplify graph roughness; bounded weights alone do not guarantee this, which is why the condition is stated explicitly.

3.2 Exact recurrent credit

Consider a finite-horizon linear RNN

$$h_{t+1} = Wh_t + Bx_t, \qquad y_t = Ch_t, \qquad t = 0, \dots, T-1, \tag{8}$$

with squared loss

$$\mathcal{L}(W) = \frac{1}{2} \sum_{t=1}^{T} \|Ch_t - y_t^\star\|_2^2 . \tag{9}$$

We call the recurrent dynamics stable when $\rho(W) < 1$. Stability is not needed for the finite-horizon adjoint identity, but it rules out explosive recurrent hard instances.

Let $r_t = Ch_t - y_t^\star$. The exact adjoint fields are

$$\lambda_T = C^\top r_T, \qquad \lambda_t = C^\top r_t + W^\top \lambda_{t+1}, \quad t = T-1, \dots, 1. \tag{10}$$

The recurrent-weight gradient factors as

$$\nabla_W \mathcal{L} = \sum_{t=0}^{T-1} \lambda_{t+1} h_t^\top . \tag{11}$$

Equation (11) is the exact three-factor form: the presynaptic factor is $h_{t,j}$, the postsynaptic credit is $\lambda_{t+1,i}$, and the synaptic update for W_{ij} is their product.

3.3 Admissible modulatory rules

Definition 1 (Admissible feedback code)**.** *A* D*-dimensional admissible feedback code is a fixed subspace* $\mathcal{H} \subset \mathbb{R}^N$ *with* $\dim(\mathcal{H}) \le D$*. A modulatory three-factor rule using this code has updates*

$$\Delta W_{ij} = -\eta \sum_{t=0}^{T-1} e_{ij,t}\, m_{t+1,i}, \qquad m_{t+1} \in \mathcal{H}, \tag{12}$$

where $e_{ij,t}$ *is a synapse-local eligibility factor. For the linear RNN above, the exact eligibility for* W_{ij} *is* $e_{ij,t} = h_{t,j}$*, so the rule estimates the gradient by*

$$\widehat{\nabla_W \mathcal{L}} = \sum_{t=0}^{T-1} m_{t+1} h_t^\top , \qquad m_{t+1} \in \mathcal{H}. \tag{13}$$

The subspace $\mathcal{H}$ is fixed by the feedback architecture: available sources, receptor fields, and diffusion dynamics. The coefficients of m_t may vary with task and time. Nonlinear receptor dose-response can change the map from source amplitude to coefficient, but it does not by itself create additional spatial degrees of freedom unless it creates additional independent spatial patterns.

Proposition 3 (Diffusion and cell-type sources induce finite-dimensional codes)**.** *Let P be any local linear diffusion operator on G, and let $b_1, \dots, b_K \in \mathbb{R}^N$ be source, receptor, or cell-type injection patterns. Any rule whose modulatory field is a linear combination of S diffusion taps,*

$$m_t = \sum_{k=1}^{K} \sum_{r=0}^{S-1} a_{k,r,t} P^r b_k, \tag{14}$$

uses the admissible code

$$\mathcal{H}_{K,S}(P, B) = \operatorname{span}\{P^r b_k : 0 \le r < S,\ 1 \le k \le K\}, \qquad \dim(\mathcal{H}_{K,S}) \le KS. \tag{15}$$

Conversely, every field in $\mathcal{H}_{K,S}$ can be implemented by appropriate coefficients $a_{k,r,t}$. Therefore every lower bound for effective dimension $D = \dim(\mathcal{H}_{K,S})$ applies to all such K-source, S-tap diffusive modulatory rules.

All minimax results below are stated in terms of D. Setting $D = K$ gives the one-pattern-per-channel case. Setting $D \le KS$ gives the source/tap version. This resolves the common ambiguity between number of biochemical sources and number of independent feedback degrees of freedom.

A (K, S)-signal-limited three-factor rule is one whose nonlocal postsynaptic signal lies in a fixed span of at most KS spatial, temporal, or cell-type patterns; the formal definition is in Appendix A.3. This covers output-feedback e-prop-like rules and ModProp-like cell-type modulatory rules.

4 Optimal constrained recurrent credit is projection

For a fixed admissible code $\mathcal{H}$, let $\Pi_{\mathcal{H}}$ denote the Euclidean orthogonal projector onto $\mathcal{H}$.

Theorem 1 (Credit projection theorem)**.** *Fix a trajectory $h_0, \dots, h_{T-1}$ and exact adjoints $\lambda_1, \dots, \lambda_T$. Among all constrained gradient estimators of the form* (13) *with $m_{t+1} \in \mathcal{H}$, one minimizer of the actual recurrent gradient error*

$$\left\| \sum_{t=0}^{T-1} \lambda_{t+1} h_t^\top - \sum_{t=0}^{T-1} m_{t+1} h_t^\top \right\|_F^2 \tag{16}$$

is

$$m_{t+1}^\star = \Pi_{\mathcal{H}} \lambda_{t+1}, \qquad t = 0, \dots, T-1. \tag{17}$$

The resulting minimum error is

$$\left\| \sum_{t=0}^{T-1} (I - \Pi_{\mathcal{H}}) \lambda_{t+1} h_t^\top \right\|_F^2. \tag{18}$$

If the hidden-state factors $h_0, \dots, h_{T-1}$ are linearly independent, this minimizer is unique.

The theorem is the core recurrent factorization link. It is not approximating an arbitrary graph signal after the fact; it identifies the nonlocal postsynaptic factor inside the exact eligibility outer product and shows that the optimal constrained recurrent gradient is obtained by projecting that adjoint factor. Cross-time correlations among hidden states do not change the projected adjoints because every residual column lies in $\mathcal{H}^\perp$ and every admissible correction column lies in $\mathcal{H}$.

For later use, the low-frequency code gives the trajectory upper bound

$$\left\| \nabla_W \mathcal{L} - \widehat{\nabla_W \mathcal{L}}_D \right\|_F^2 \le \left(\sum_{t=0}^{T-1} \|h_t\|_2^2 \right) \left(\sum_{t=0}^{T-1} \left\| (I - \Pi_{\mathcal{H}_D^\star}) \lambda_{t+1} \right\|_2^2 \right), \tag{19}$$

where $\mathcal{H}_D^\star = \operatorname{span}\{u_1, \dots, u_D\}$.

5 Tight graph-spectral credit-bandwidth law

Theorem 2 (Kolmogorov width of graph-Sobolev credit fields)**.** *For every $s \geq 0$, $R > 0$, and $0 \leq D < N$,*

$$\inf_{\dim(\mathcal{H}) \leq D} \sup_{f \in \mathcal{B}_s(R)} \|(I - \Pi_{\mathcal{H}})f\|_2^2 = R^2(1 + \mu_{D+1})^{-s}. \tag{20}$$

The infimum is achieved by the low-frequency graph subspace

$$\mathcal{H}_D^\star = \operatorname{span}\{u_1, \ldots, u_D\}. \tag{21}$$

Theorem 2 is only a graph approximation theorem until we embed it in a recurrent trajectory. We use a stable graph-local recurrence driven at every step by stationary input. Fix $T \geq 2$, $\beta \in (0, 1)$, and

$$0 < \tau \leq \frac{1}{2\max\{1, \mu_N\}}. \tag{22}$$

Let

$$W_{\beta,\tau} = \beta(I - \tau L_G), \tag{23}$$

and write $W_{\beta,\tau} u_\ell = \kappa_\ell u_\ell$, so

$$\frac{\beta}{2} \leq \kappa_\ell = \beta(1 - \tau\mu_\ell) \leq \beta < 1. \tag{24}$$

The recurrent dynamics are nonzero, stable, and graph-local. For $B > 0$, set $\sigma^2 = B^2/N$ and draw a stationary Gaussian graph-filtered state process

$$h_{t+1} = W_{\beta,\tau} h_t + \xi_t, \qquad h_0 \sim \mathcal{N}(0, \sigma^2 I), \qquad \xi_t \sim \mathcal{N}(0, \sigma^2(I - W_{\beta,\tau}^2)), \tag{25}$$

with independent innovations. Then $\mathbb{E} h_t h_t^\top = \sigma^2 I$ and $\mathbb{E}\|h_t\|_2^2 = B^2$ for every t; graph mode u_ℓ is an AR(1) mode with coefficient κ_ℓ. For each $f \in \mathcal{B}_s(R)$, choose residual drives $z_T = f$ and $z_t = (I - W_{\beta,\tau})f$ for $1 \leq t < T$ (equivalently $C = I_N$ and targets $y_t^\star = h_t - z_t$ at this parameter point). The adjoint recursion gives $\lambda_t = f$ for all t, and hence

$$\nabla_W \mathcal{L}_f = \sum_{t=0}^{T-1} f h_t^\top, \tag{26}$$

which has T nonzero rank-one recurrent eligibility slices almost surely whenever $f \neq 0$.

Define

$$A_T = \mathbb{E}\left\|\sum_{t=0}^{T-1} h_t\right\|_2^2 = \frac{B^2}{N}\sum_{\ell=1}^{N}\left[T + 2\sum_{r=1}^{T-1}(T - r)\kappa_\ell^r\right]. \tag{27}$$

It satisfies $B^2 T \leq A_T \leq B^2 T(1 + \beta)/(1 - \beta)$, so $A_T = \Theta(T)$ for fixed β. For an admissible code $\mathcal{H}$, let

$$\mathfrak{C}_{\mathrm{st}}(\mathcal{H}; R) = \sup_{f \in \mathcal{B}_s(R)} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E} \min_{m_t \in \mathcal{H}} \|\lambda_t - m_t\|_2^2, \tag{28}$$

$$\mathfrak{G}_{\mathrm{st}}(\mathcal{H}; R, B, T) = \sup_{f \in \mathcal{B}_s(R)} \mathbb{E} \inf_{m_1, \ldots, m_T \in \mathcal{H}} \left\|\sum_{t=0}^{T-1}(f - m_{t+1}) h_t^\top\right\|_F^2. \tag{29}$$

Theorem 3 (Persistent recurrent credit-bandwidth theorem)**.** *For every $0 \leq D < N$,*

$$\inf_{\dim(\mathcal{H}) \leq D} \mathfrak{C}_{\mathrm{st}}(\mathcal{H}; R) = R^2(1 + \mu_{D+1})^{-s}, \tag{30}$$

$$\inf_{\dim(\mathcal{H}) \leq D} \mathfrak{G}_{\mathrm{st}}(\mathcal{H}; R, B, T) = A_T R^2(1 + \mu_{D+1})^{-s}. \tag{31}$$

Both optima are achieved by $\mathcal{H}_D^\star = \operatorname{span}\{u_1, \ldots, u_D\}$. The lower bound allows m_t to depend on f and on the whole realized trajectory; the obstruction survives because all admissible left factors remain in $\mathcal{H}$, while the omitted credit component remains in $\mathcal{H}^\perp$. The constant-adjoint presentation is not essential: Lemma 1 in the appendix gives the same construction for $\lambda_t = a_t f$ and finite omitted-mode mixtures, with A_T replaced by a positive temporal-energy quadratic form.

Corollary 1 (e-prop/ModProp-like signal-class lower bound)**.** *Any (K, S)-signal-limited three-factor rule in Appendix A.3 has effective nonlocal dimension $D = \dim(\mathcal{H}_{\mathrm{sig}}) \le KS$. On the persistent stochastic class above, even trajectory-dependent coefficients incur*

$$\mathfrak{C}_{\mathrm{st}}(\mathcal{H}_{\mathrm{sig}}; R) \ge R^2(1+\mu_{D+1})^{-s}, \qquad \mathfrak{G}_{\mathrm{st}}(\mathcal{H}_{\mathrm{sig}}; R, B, T) \ge A_T R^2(1+\mu_{D+1})^{-s}. \tag{32}$$

Thus output-feedback e-prop-like rules, random feedback with D fixed channels, and ModProp-like cell-type modulatory rules with K modulators and S predetermined taps are all limited by $D \le KS$.

Corollary 2 (Credit-bandwidth scaling on manifold-like graphs)**.** *If $1+\mu_D \asymp D^{2/d}$ in the relevant bandwidth range, as for graph Laplacians converging to a d-dimensional Laplace–Beltrami operator [Rosenberg, 1997, Belkin and Niyogi, 2007], then the terminal-credit error, the state-normalized persistent-gradient error $\mathfrak{G}_{\mathrm{st}}/A_T$, and the PL optimization floor below scale as $D^{-2s/d}$.*

The channel-budget form follows immediately: reducing a floor $E(D) \asymp cD^{-2s/d}$ by a factor m requires multiplying feedback dimension by $m^{d/(2s)}$. Appendix F visualizes this implication.

5.1 Adaptive and trajectory-dependent feedback

A finite repertoire of M candidate D-dimensional codes is no stronger than their MD-dimensional span, giving the $(MD+1)$st-frequency lower bound (Appendix A.8). A stronger adaptive model may choose a continuum of subspaces after reading q scalar observations of the exact credit field.

Theorem 4 (Information-limited adaptive feedback)**.** *Let $q \ge 0$ and $D+q < N$. For any linear map $A : \mathbb{R}^N \to \mathbb{R}^q$ and arbitrary selector $\Phi(y)$ with $\dim \Phi(y) \le D$,*

$$\sup_{f \in \mathcal{B}_s(R)} \left\| (I - \Pi_{\Phi(Af)}) f \right\|_2^2 \ge R^2(1+\mu_{D+q+1})^{-s}. \tag{33}$$

Consequently, on the persistent recurrent class,

$$\sup_{f \in \mathcal{B}_s(R)} \mathbb{E} \inf_{m_t \in \Phi(Af)} \left\| \sum_{t=0}^{T-1} (f - m_{t+1}) h_t^\top \right\|_F^2 \ge A_T R^2 (1+\mu_{D+q+1})^{-s}. \tag{34}$$

This is an intentionally strong side-channel model: A is noiseless, real-valued, and applied directly to the exact credit field. The result should therefore be read as accounting for q extra communicated credit coordinates, not as ruling out rich adaptation; communicating all high-frequency coefficients simply moves the bandwidth burden into the adaptation channel.

5.2 General trajectory upper bound

For any trajectory class

$$\mathcal{C}_s^{(T)}(R, H) = \left\{ \mathcal{P} : \sum_{t=0}^{T-1} \|\lambda_{t+1}\|_{H_G^s}^2 \le R^2, \quad \sum_{t=0}^{T-1} \|h_t\|_2^2 \le H \right\}, \tag{35}$$

Theorems 1 and 2 give

$$\sup_{\mathcal{P} \in \mathcal{C}_s^{(T)}(R,H)} \left\| \nabla_W \mathcal{L} - \widehat{\nabla_W \mathcal{L}}_D \right\|_F^2 \le HR^2(1+\mu_{D+1})^{-s}. \tag{36}$$

Theorem 3 shows that the state-energy factor is real: in a stable recurrent process H is order TB^2, and the matched lower bound is order $TR^2(1+\mu_{D+1})^{-s}$.

6 Realizable high-frequency recurrent hard instances

The persistent theorem has an explicit omitted-mode interpretation. Choose $f = au_\ell$ with $\ell > D$ and $a = R(1+\mu_\ell)^{-s/2}$. Then $\lambda_t = au_\ell$ at every time, every recurrent slice $au_\ell h_t^\top$ is nonzero almost surely, and every D-dimensional code orthogonal to u_ℓ incurs

$$\mathbb{E} \inf_{m_t \in \mathcal{H}} \left\| \sum_{t=0}^{T-1} (au_\ell - m_{t+1}) h_t^\top \right\|_F^2 = A_T a^2. \tag{37}$$

Thus the omitted high-frequency credit is not a terminal artifact: stable graph-filter recurrence repeatedly exposes the same missing postsynaptic direction to time-distributed eligibility traces. The proof also shows why cross-time cancellation cannot help. All admissible corrections have left factor in $\mathcal{H}$, whereas the omitted mode lies in $\mathcal{H}^{\perp}$; hidden-state correlations only change the scalar factor A_T.

The same obstruction can arise from a local loss rather than from a prescribed residual drive. Let ∇_G be a weighted incidence matrix with $\nabla_G^{\top}\nabla_G = L_G$. For $\mu_\ell > 0$, the no-input trajectory $h_t = (a/\mu_\ell)\kappa_\ell^{t-T}u_\ell$ under $W_{\beta,\tau}$ and the terminal edge loss

$$\mathcal{L} = \frac{1}{2}\left\|\nabla_G h_T\right\|_2^2 = \frac{1}{2}\sum_{(i,j)\in E} w_{ij}(h_{T,i} - h_{T,j})^2 \tag{38}$$

have terminal adjoint $\lambda_T = L_G h_T = a u_\ell$ and gradient

$$\nabla_W \mathcal{L} = \frac{Ta^2}{\mu_\ell \kappa_\ell} u_\ell u_\ell^{\top}. \tag{39}$$

If $u_\ell \perp \mathcal{H}$, every admissible estimate is Frobenius-orthogonal in its left factor and the squared error is the squared norm of this matrix. This deterministic variant is not used for the matched stochastic minimax value; it shows that graph-local residuals also generate rough adjoints.

7 Quantitative predictions and learning consequences

The theory is falsifiable because s, d, and D_{eff} can be estimated rather than asserted. For a trained RNN or neural recording embedded in an anatomical, functional, or product graph, compute BPTT adjoints λ_t and fit the graph-Fourier slope

$$\log\left(\frac{1}{|\mathcal{B}_j|T}\sum_{t,\ell\in\mathcal{B}_j}|\langle\lambda_t, u_\ell\rangle|^2\right) = c - \widehat{s}\log(1+\bar{\mu}_j) + \epsilon_j \tag{40}$$

over spectral bins $\mathcal{B}_j$. Smooth cortical-sheet or cell-type embeddings should give $\widehat{s} > 0$, while shuffled embeddings should drive $\widehat{s}$ toward zero. The prediction is then numerical: on a two-dimensional sheet, $E(D) \propto D^{-s}$. For broad neuromodulatory or volume-transmission systems [Avery and Krichmar, 2017, Özçete et al., 2024], if four monoamine/cholinergic systems (DA, NE, 5HT, ACh) are assigned $S = 10$ independent spatial or temporal taps, the illustrative budget is $D_{\text{eff}} \le 40$; a tenfold smaller floor requires $D \approx 400$ when $s = 1$, $D \approx 126$ when $s = 2$, and $D \approx 4000$ when $s = 1/2$. Neuropeptide signaling can raise the effective budget only to the extent that receptor/source expression supplies independent patterns; Smith et al. report 18 neuropeptide precursor genes, 29 neuropeptide-selective GPCR genes, and 37 cognate pairs in mouse cortex [Smith et al., 2019].

This yields falsifiable outcomes rather than a generic "more channels help" claim. Working-memory, evidence-accumulation, and motor-planning tasks should have positive slopes on the appropriate sheet, map, or cell-type graph ($\widehat{s} \approx$ 1–2 would make $D_{\text{eff}} = 40$ useful), but the same adjoints on a randomly permuted graph should have $\widehat{s} \approx 0$. Tasks requiring opposite credit for adjacent cells should leave a loss gap under diffuse feedback even when exact BPTT succeeds. Adding an independent cell-type or peptide-receptor axis should reduce the gap only when it removes the high-frequency tail in the corresponding product graph.

A practical protocol is to report, for each task and embedding, the fitted slope $\widehat{s}$, the cumulative tail $\sum_{\ell > D_{\text{eff}},t}|\langle\lambda_t, u_\ell\rangle|^2 / \sum_{\ell,t}|\langle\lambda_t, u_\ell\rangle|^2$, and the constrained-feedback loss gap. This separates the spectral theory from dimension-only baselines: design tasks with equal Euclidean credit norm and equal random-feedback dimension but different graph roughness. Figure 1 instantiates the protocol with measured multi-time BPTT adjoints rather than prescribed terminal credits; across $D \in \{16, 32, 64\}$, graph-spectral tails order the low-pass loss gaps, while a random same-dimensional tail is a weak predictor for those same low-pass runs. Appendix Figure 2 repeats the diagnostic on rule-generated delayed-response tasks whose targets are not specified by graph-Fourier coefficients.

For a measured trajectory, the spectral prediction is

$$\widehat{\Delta}_D = \frac{1}{2\alpha} \sum_{t=0}^{T-1} \|h_t\|_2^2 \sum_{\ell > D} |\langle \lambda_{t+1}, u_\ell \rangle|^2, \tag{41}$$

up to the local PL constant and the Cauchy–Schwarz looseness in (36). A Haar-random D-subspace of the same dimension has expected omitted energy $(1 - D/N) \sum_{\ell,t} |\langle \lambda_t, u_\ell \rangle|^2$, independent of graph frequency. Thus the discriminative outcome is the ordering across roughness at fixed dimension and fixed Euclidean credit energy, not merely that constrained feedback is worse than BPTT. Appendix D gives the real-data protocol and diagnostics for estimating D_{eff} and d.

The channel budget must also be audited. Given candidate source or receptor patterns b_k and taps P^r, form $B_{\text{code}} = [P^r b_k]_{k,r}$ and report numerical rank or stable rank; correlated receptor fields, redundant diffusion taps, and saturated nonlinear dose-response curves reduce D_{eff}. Estimate d by the low-frequency slope of $\log(1 + \mu_\ell)$ against $\log \ell$ over the feedback-relevant range. A motor-map task should have a smaller tail on a somatotopic graph than after shuffling; a context task asking neighboring units for opposite updates should show the reverse unless a cell-type/product graph separates contexts. Appendix D expands this protocol.

The optimization bridge is standard but important. If training trajectories obey the class condition in (36) and the local objective is L-smooth and satisfies the α-PL inequality, then inexact-gradient descent with the low-frequency constrained gradient has

$$\limsup_{n \to \infty} \big(F(W_n) - F^\star\big) \le \frac{HR^2(1 + \mu_{D+1})^{-s}}{2\alpha}. \tag{42}$$

Thus the omitted graph spectrum is not only a signal-approximation error; in local error-bound regimes it sets the residual loss scale.

8 Limitations and broader implications

Extensions to product graphs and to nonlinear/spiking systems with three-factor adjoints are in Appendices B–C; local diffusion depth versus feedback dimension is treated in Appendix A.12. The main lower bound is worst-case and local in parameter space: it constructs targets/residual drives at the stable graph-filter point, and it does not claim that cortex computes Laplacian eigenvectors or that all cortical adjoints are smooth. Natural tasks with compressible adjoints can perform better, while rough tasks may become smooth after representational learning. The biological numbers above are illustrative effective-dimension calculations, not a claim that each molecule is an independent channel.

9 Conclusion

Biologically plausible credit assignment is often framed as a search for local mechanisms approximating backpropagation. This paper instead gives a communication-constrained lower bound and a design rule: finite-dimensional diffuse feedback must compress exact recurrent adjoints, and the right compression is graph low-pass credit. The first excluded graph frequency sets the unavoidable credit error, the recurrent-gradient bias accumulated over time, and the PL-regime optimization floor. Feedback bandwidth is therefore a quantitative learning resource, analogous to model dimension or sample size, with a measurable channel budget.

References

Michael C. Avery and Jeffrey L. Krichmar. Neuromodulatory systems and their interactions: A review of models, theories, and experiments. *Frontiers in Neural Circuits*, 11:108, 2017. doi: 10.3389/fncir.2017.00108.

Mikhail Belkin and Partha Niyogi. Convergence of laplacian eigenmaps. In *Advances in Neural Information Processing Systems 19*, pages 129–136. MIT Press, 2007.

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(1):3625, 2020. doi: 10.1038/s41467-020-17236-y.

Fan R. K. Chung. *Spectral Graph Theory*, volume 92 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, RI, 1997.

Luke Eilers, Raoul-Martin Memmesheimer, and Sven Goedeke. A generalized neural tangent kernel for surrogate gradient learning. In *Advances in Neural Information Processing Systems 37*. Curran Associates, Inc., 2024. doi: 10.52202/079017-0287.

Benjamin Ellenberger, Paul Haider, Federico Benitez, Jakob Jordan, Kevin Max, Ismael Jaras, Laura Kriener, and Mihai A. Petrovici. Backpropagation through space, time and the brain. *Nature Communications*, 17:66, 2026. doi: 10.1038/s41467-025-66666-z.

Christian Holberg and Cristopher Salvi. Exact gradients for stochastic spiking neural networks driven by rough signals. In *Advances in Neural Information Processing Systems 37*. Curran Associates, Inc., 2024. doi: 10.52202/079017-1003.

Klara Kaleb, Barbara Feulner, Juan A. Gallego, and Claudia Clopath. Feedback control guides credit assignment in recurrent neural networks. In *Advances in Neural Information Processing Systems 37*, pages 5122–5144. Curran Associates, Inc., 2024. doi: 10.52202/079017-0166.

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: 13276, 2016. doi: 10.1038/ncomms13276.

Timothy P. Lillicrap, Adam Santoro, Luke Marris, Colin J. Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, 21:335–346, 2020. doi: 10.1038/s41583-020-0277-3.

Yuhan Helena Liu, Stephen J. Smith, Stefan Mihalas, Eric Shea-Brown, and Uygar Sümbül. Cell-type-specific neuromodulation guides synaptic credit assignment in a spiking neural network. *Proceedings of the National Academy of Sciences of the United States of America*, 118(51): e2111821118, 2021. doi: 10.1073/pnas.2111821118.

Yuhan Helena Liu, Stephen Smith, Stefan Mihalas, Eric Shea-Brown, and Uygar Sümbül. Biologically-plausible backpropagation through arbitrary timespans via local neuromodulators. In *Advances in Neural Information Processing Systems 35*. Curran Associates, Inc., 2022.

Alexander Meulemans, Matilde Tristany Farinha, Javier García Ordóñez, Pau Vilimelis Aceituno, João Sacramento, and Benjamin F. Grewe. Credit assignment in neural networks through deep feedback control. In *Advances in Neural Information Processing Systems 34*, pages 4674–4687. Curran Associates, Inc., 2021.

Arild Nøkland. Direct feedback alignment provides learning in deep neural networks. In *Advances in Neural Information Processing Systems 29*. Curran Associates, Inc., 2016.

Özge D. Özçete, Aditi Banerjee, and Pascal S. Kaeser. Mechanisms of neuromodulatory volume transmission. *Molecular Psychiatry*, 29:3680–3693, 2024. doi: 10.1038/s41380-024-02608-3.

Allan Pinkus. *n-Widths in Approximation Theory*, volume 7 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer-Verlag, Berlin, 1985. doi: 10.1007/978-3-642-69894-1.

Steven Rosenberg. *The Laplacian on a Riemannian Manifold: An Introduction to Analysis on Manifolds*, volume 31 of *London Mathematical Society Student Texts*. Cambridge University Press, 1997. doi: 10.1017/CBO9780511623783.

David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986. doi: 10.1038/323533a0.

João Sacramento, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. Dendritic cortical microcircuits approximate the backpropagation algorithm. In *Advances in Neural Information Processing Systems 31*. Curran Associates, Inc., 2018.

David I. Shuman, Sunil K. Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98, 2013. doi: 10.1109/MSP.2012.2235192.

Stephen J. Smith, Uygar Sümbül, Lucas T. Graybuck, Forrest Collman, Sharmishtaa Seshamani, Rohan Gala, Olga Gliko, Leila Elabbady, Jeremy A. Miller, Trygve E. Bakken, Jean Rossier, Zizhen Yao, Ed Lein, Hongkui Zeng, Bosiljka Tasic, and Michael Hawrylycz. Single-cell transcriptomic evidence for dense intracortical neuropeptide networks. *eLife*, 8:e47889, 2019. doi: 10.7554/eLife.47889.

Paul J. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990. doi: 10.1109/5.58337.

James C. R. 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. doi: 10.1162/NECO_a_00949.

Ronald J. Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1(2):270–280, 1989. doi: 10.1162/neco.1989.1.2.270.

A Proofs

A.1 Adjoint factorization

Proof of Equation (11). Let δh_t be the first-order perturbation induced by δW. The dynamics give

$$\delta h_{t+1} = W\delta h_t + \delta W h_t, \qquad \delta h_0 = 0. \tag{43}$$

The loss variation is

$$\delta\mathcal{L} = \sum_{t=1}^{T} r_t^\top C\delta h_t. \tag{44}$$

With $\lambda_T = C^\top r_T$ and $\lambda_t = C^\top r_t + W^\top \lambda_{t+1}$, summation by parts yields

$$\delta\mathcal{L} = \sum_{t=1}^{T} \lambda_t^\top \delta h_t - \sum_{t=1}^{T-1} \lambda_{t+1}^\top W\delta h_t \tag{45}$$

$$= \sum_{t=0}^{T-1} \lambda_{t+1}^\top (\delta h_{t+1} - W\delta h_t) \tag{46}$$

$$= \sum_{t=0}^{T-1} \lambda_{t+1}^\top \delta W h_t = \left\langle \sum_{t=0}^{T-1} \lambda_{t+1} h_t^\top, \delta W \right\rangle_F. \tag{47}$$

Since this holds for every δW, the gradient is (11). □

A.2 Admissible diffusion codes

Proof of Proposition 3. Each field m_t is a finite linear combination of the vectors $P^r b_k$, hence belongs to their span. The span has at most KS generators, so its dimension is at most KS. Conversely, by definition of span, every vector in $\mathcal{H}_{K,S}$ is a finite linear combination of those generators, and the corresponding coefficients are precisely the channel/tap amplitudes $a_{k,r,t}$. □

A.3 e-prop/ModProp-like signal class

Definition 2 (e-prop/ModProp-like signal class)**.** *A three-factor rule is called* (K, S)-signal-limited *if, for every task instance and time* t*, its nonlocal postsynaptic learning signal can be written as*

$$m_t = \sum_{k=1}^{K} \sum_{r=0}^{S-1} a_{k,r,t}\, \psi_{k,r}, \qquad \psi_{k,r} \in \mathbb{R}^N \text{ fixed before the instance}, \tag{48}$$

while the synapse-specific factor is local. The vectors $\psi_{k,r}$ *may represent fixed random feedback columns, output-error feedback columns, cell-type receptor fields, local neuromodulatory source patterns, temporal-filter taps, or graph-diffusion taps. The resulting admissible code is*

$$\mathcal{H}_{\mathrm{sig}} = \mathrm{span}\{\psi_{k,r} : 1 \le k \le K,\, 0 \le r < S\}, \qquad \dim(\mathcal{H}_{\mathrm{sig}}) \le KS. \tag{49}$$

Thus output-feedback e-prop-like rules with K *learning-signal channels have* $D \le K$ *when there is one fixed spatial pattern per channel, and* $D \le KS$ *when each channel is combined with* S *fixed filtering taps. ModProp-like rules with* K *cell-type-specific modulators and* S *predetermined filtering taps are included by taking* $\psi_{k,r} = P^r b_k$*.*

A.4 Projection theorem

Proof of Theorem 1. Let $\Pi = \Pi_{\mathcal{H}}$. For arbitrary $m_{t+1} \in \mathcal{H}$,

$$\sum_t \lambda_{t+1} h_t^\top - \sum_t m_{t+1} h_t^\top = \sum_t (I - \Pi)\lambda_{t+1} h_t^\top + \sum_t (\Pi\lambda_{t+1} - m_{t+1}) h_t^\top. \tag{50}$$

The left factors in the first sum lie in $\mathcal{H}^\perp$, while the left factors in the second sum lie in $\mathcal{H}$. Therefore the two matrices are orthogonal in Frobenius inner product:

$$\left\langle \sum_t (I - \Pi)\lambda_{t+1} h_t^\top, \sum_s (\Pi\lambda_{s+1} - m_{s+1}) h_s^\top \right\rangle_F \tag{51}$$

$$= \sum_{t,s} h_t^\top h_s \, \langle (I - \Pi)\lambda_{t+1}, \Pi\lambda_{s+1} - m_{s+1} \rangle = 0. \tag{52}$$

Thus the squared error decomposes as

$$\left\| \sum_t (I - \Pi)\lambda_{t+1} h_t^\top \right\|_F^2 + \left\| \sum_t (\Pi\lambda_{t+1} - m_{t+1}) h_t^\top \right\|_F^2. \tag{53}$$

The first term is independent of m_t and the second is minimized by $m_{t+1} = \Pi\lambda_{t+1}$ for every t. If the h_t are linearly independent, the second term is zero only when every coefficient vector $\Pi\lambda_{t+1} - m_{t+1}$ is zero, giving uniqueness. □

A.5 Smoothness inheritance

Proof of Proposition 1. Because $W = p(L_G)$, it is diagonal in the graph Fourier basis. For $f = \sum_\ell f_\ell u_\ell$,

$$\|Wf\|_{H_G^s}^2 = \sum_\ell (1 + \mu_\ell)^s |p(\mu_\ell)|^2 |f_\ell|^2 \le \gamma^2 \|f\|_{H_G^s}^2. \tag{54}$$

Unrolling the adjoint recursion gives $\lambda_t = \sum_{q=t}^{T} W^{q-t} z_q$. The triangle inequality and the contraction bound imply

$$\|\lambda_t\|_{H_G^s} \le \sum_{q=t}^{T} \gamma^{q-t} \|z_q\|_{H_G^s}. \tag{55}$$

The uniform bound follows by summing the geometric series. □

Proof of Proposition 2. Unrolling the adjoint recursion gives

$$\lambda_t = \sum_{q=t}^{T} J_t^\top J_{t+1}^\top \cdots J_{q-1}^\top z_q, \tag{56}$$

with the empty product equal to the identity. Applying (7) to each factor and then the triangle inequality yields

$$\|\lambda_t\|_{H_G^s} \le \sum_{q=t}^{T} \gamma^{q-t} \|z_q\|_{H_G^s}, \tag{57}$$

which is the claimed inheritance bound. □

A.6 Graph-Sobolev minimax theorem

Proof of Theorem 2. For the upper bound, choose $\mathcal{H}_D^\star = \mathrm{span}\{u_1, \dots, u_D\}$. Writing $f_\ell = \langle f, u_\ell \rangle$,

$$\left\|(I - \Pi_{\mathcal{H}_D^\star}) f\right\|_2^2 = \sum_{\ell > D} |f_\ell|^2 \tag{58}$$

$$\le (1 + \mu_{D+1})^{-s} \sum_{\ell > D} (1 + \mu_\ell)^s |f_\ell|^2 \tag{59}$$

$$\le R^2 (1 + \mu_{D+1})^{-s}. \tag{60}$$

For the lower bound, let $\mathcal{H}$ be any subspace with $\dim(\mathcal{H}) \le D$. The $(D+1)$-dimensional space $E_{D+1} = \mathrm{span}\{u_1, \dots, u_{D+1}\}$ has a nonzero vector $g \in E_{D+1} \cap \mathcal{H}^\perp$. Normalize $\|g\|_2 = 1$. Since g is supported only on the first $D+1$ eigenvectors,

$$\|g\|_{H_G^s}^2 \le (1 + \mu_{D+1})^s. \tag{61}$$

Set $f = R(1 + \mu_{D+1})^{-s/2} g$. Then $f \in \mathcal{B}_s(R)$ and, because $f \perp \mathcal{H}$,

$$\|(I - \Pi_{\mathcal{H}}) f\|_2^2 = \|f\|_2^2 = R^2 (1 + \mu_{D+1})^{-s}. \tag{62}$$

This proves equality and optimality of $\mathcal{H}_D^\star$. □

A.7 Persistent recurrent minimax theorem

Proof of Theorem 3. Stationarity follows because $W_{\beta,\tau}$ is symmetric with spectrum in $(0, 1)$ and

$$W_{\beta,\tau}(\sigma^2 I) W_{\beta,\tau} + \sigma^2 (I - W_{\beta,\tau}^2) = \sigma^2 I. \tag{63}$$

Moreover $\mathbb{E} h_t h_s^\top = \sigma^2 W_{\beta,\tau}^{|t-s|}$. Therefore

$$\mathbb{E} \left\| \sum_{t=0}^{T-1} h_t \right\|_2^2 = \sum_{t,s=0}^{T-1} \mathrm{tr}\, \mathbb{E}[h_t h_s^\top] \tag{64}$$

$$= \sigma^2 \sum_{\ell=1}^{N} \left[T + 2 \sum_{r=1}^{T-1} (T - r) \kappa_\ell^r \right] = A_T. \tag{65}$$

Since $0 \le \kappa_\ell \le \beta$, $B^2 T \le A_T \le B^2 T (1 + \beta)/(1 - \beta)$.

The chosen residual drives give $\lambda_T = z_T = f$ and, for $t < T$, $\lambda_t = z_t + W_{\beta,\tau} \lambda_{t+1} = (I - W_{\beta,\tau}) f + W_{\beta,\tau} f = f$. Thus the exact gradient is (26). For a fixed trajectory and any admissible $m_t \in \mathcal{H}$, Theorem 1 implies that $m_t = \Pi_{\mathcal{H}} f$ is optimal. The residual gradient is

$$\sum_{t=0}^{T-1} (I - \Pi_{\mathcal{H}}) f \, h_t^\top = (I - \Pi_{\mathcal{H}}) f \left(\sum_{t=0}^{T-1} h_t \right)^\top, \tag{66}$$

so its expected squared Frobenius norm is

$$\|(I - \Pi_{\mathcal{H}}) f\|_2^2 \, A_T. \tag{67}$$

Similarly, because every $\lambda_t = f$, the normalized trajectory-credit risk is exactly $\sup_{f \in \mathcal{B}_s(R)} \|(I - \Pi_{\mathcal{H}}) f\|_2^2$. Taking the infimum over all D-dimensional codes and applying Theorem 2 proves both equalities. The statement for signal-limited rules follows by applying the same argument to the fixed code $\mathcal{H}_{\mathrm{sig}}$ from Definition 2. □

Lemma 1 (Time-varying omitted adjoints)**.** *Use the stationary process and graph-filter recurrence from Theorem 3, and fix any admissible code $\mathcal{H}$. Let $g_1, \dots, g_q$ be orthonormal vectors in $\mathcal{H}^\perp$ lying in a $W_{\beta,\tau}$-invariant graph-Fourier subspace, and let $a^{(j)} = (a_0^{(j)}, \dots, a_{T-1}^{(j)}) \in \mathbb{R}^T$. There are residual drives and targets for which*

$$\lambda_{t+1} = \sum_{j=1}^{q} a_t^{(j)} g_j, \qquad t = 0, \dots, T-1. \tag{68}$$

For these tasks,

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \min_{m_{t+1} \in \mathcal{H}} \|\lambda_{t+1} - m_{t+1}\|_2^2 = \frac{1}{T} \sum_{j=1}^{q} \left\| a^{(j)} \right\|_2^2, \tag{69}$$

$$\mathbb{E} \inf_{m_1, \dots, m_T \in \mathcal{H}} \left\| \sum_{t=0}^{T-1} (\lambda_{t+1} - m_{t+1}) h_t^\top \right\|_F^2 = \sum_{j=1}^{q} (a^{(j)})^\top \Gamma_T a^{(j)}, \tag{70}$$

where $(\Gamma_T)_{ts} = \sigma^2 \mathrm{tr}(W_{\beta,\tau}^{|t-s|})$ for $0 \le t, s < T$. Moreover

$$B^2 \frac{1-\beta}{1+\beta} \sum_{j=1}^{q} \left\| a^{(j)} \right\|_2^2 \le \sum_{j=1}^{q} (a^{(j)})^\top \Gamma_T a^{(j)} \le B^2 \frac{1+\beta}{1-\beta} \sum_{j=1}^{q} \left\| a^{(j)} \right\|_2^2. \tag{71}$$

Thus replacing the persistent credit $\lambda_t = f$ by a time-varying amplitude $a_t f$ or by a finite omitted-mode mixture changes only the temporal energy factor; sign changes or smooth temporal modulation do not remove the feedback-dimension obstruction.

Proof. Choose the desired adjoint sequence above and set

$$z_T = \lambda_T, \qquad z_t = \lambda_t - W_{\beta,\tau} \lambda_{t+1} \quad (1 \le t < T). \tag{72}$$

With $C = I_N$ and targets $y_t^\star = h_t - z_t$, the adjoint recursion gives these λ_t exactly. The $W_{\beta,\tau}$-invariant graph-Fourier assumption ensures the residual drives remain in the same structured spectral subspace.

Because every λ_t lies in $\mathcal{H}^\perp$, Theorem 1 gives the optimal admissible signal $m_t = 0$. Orthonormality of the g_j gives the normalized credit identity. For the gradient term,

$$\mathbb{E} \left\| \sum_{t=0}^{T-1} \lambda_{t+1} h_t^\top \right\|_F^2 = \sum_{j=1}^{q} \sum_{t,s=0}^{T-1} a_t^{(j)} a_s^{(j)} \, \mathbb{E}[h_t^\top h_s] \tag{73}$$

$$= \sum_{j=1}^{q} (a^{(j)})^\top \Gamma_T a^{(j)}, \tag{74}$$

since $\mathbb{E}[h_t^\top h_s] = \sigma^2 \mathrm{tr}(W_{\beta,\tau}^{|t-s|})$.

It remains to bound Γ_T. In the graph-Fourier basis,

$$\Gamma_T = \sigma^2 \sum_{\ell=1}^{N} K_{\kappa_\ell}, \qquad (K_\kappa)_{ts} = \kappa^{|t-s|}. \tag{75}$$

For finite a, the AR(1) Toeplitz matrix has the Fourier representation

$$a^\top K_\kappa a = \frac{1}{2\pi} \int_{-\pi}^{\pi} \left| \sum_{t=0}^{T-1} a_t e^{\mathrm{i}t\theta} \right|^2 \frac{1-\kappa^2}{1 - 2\kappa\cos\theta + \kappa^2} \, d\theta. \tag{76}$$

For $0 \le \kappa \le \beta < 1$, the spectral density is bounded between $(1-\beta)/(1+\beta)$ and $(1+\beta)/(1-\beta)$. By Parseval,

$$\frac{1-\beta}{1+\beta} \|a\|_2^2 \le a^\top K_\kappa a \le \frac{1+\beta}{1-\beta} \|a\|_2^2. \tag{77}$$

Summing over ℓ and using $\sigma^2 = B^2/N$ proves the bounds. □

Deterministic slice variant. For reference, the earlier one-slice construction is recovered by choosing any v with $\|v\|_2 = V$, setting $h_0 = \cdots = h_{T-2} = 0$, $h_{T-1} = v$, $h_T = f$, and choosing inputs $x_{T-2} = v$, $x_{T-1} = f - W_{\beta,\tau} v$, all others zero. With terminal-only loss $\frac{1}{2}\|h_T\|_2^2$, one has $\lambda_T = f$, $\nabla_W \mathcal{L} = f v^\top$, and state/input bounds $\max_t \|h_t\|_2 \le \max\{R, V\}$ and $\max_t \|x_t\|_2 \le \max\{V, R + \beta V\}$. This yields the same width multiplied by V^2, but the main theorem uses the persistent stochastic class because the eligibility energy is distributed across time.

A.8 Adaptive feedback lower bounds

Proposition A.1 (Finite-repertoire adaptive codes)**.** *Let $\mathcal{H}_1, \dots, \mathcal{H}_M \subset \mathbb{R}^N$ be any repertoire of possible feedback codes with $\dim(\mathcal{H}_j) \le D$, and assume $MD < N$. Even if the rule may choose the best code in the repertoire after seeing the task instance,*

$$\sup_{f \in \mathcal{B}_s(R)} \min_{1 \le j \le M} \left\|(I - \Pi_{\mathcal{H}_j}) f\right\|_2^2 \ge R^2 (1 + \mu_{MD+1})^{-s}. \tag{78}$$

The same lower bound holds for normalized credits over the persistent stochastic recurrent class, and the corresponding expected full-gradient lower bound is multiplied by A_T.

Proof of Proposition A.1. Let $\mathcal{S} = \mathcal{H}_1 + \cdots + \mathcal{H}_M$. Then $\dim(\mathcal{S}) \le MD$. The space $E_{MD+1} = \operatorname{span}\{u_1, \dots, u_{MD+1}\}$ has dimension $MD+1$, so there is a unit vector $g \in E_{MD+1} \cap \mathcal{S}^\perp$. Then g is orthogonal to every $\mathcal{H}_j$ and $\|g\|_{H_G^s}^2 \le (1 + \mu_{MD+1})^s$. Set $f = R(1 + \mu_{MD+1})^{-s/2} g$. For every j,

$$\left\|(I - \Pi_{\mathcal{H}_j}) f\right\|_2^2 = \|f\|_2^2 = R^2 (1 + \mu_{MD+1})^{-s}. \tag{79}$$

For the persistent recurrent class, $\lambda_t = f$ for every t and the proof of Theorem 3 shows that full-gradient risk is exactly A_T times the vector-credit error. □

Proof of Theorem 4. Let $E_{D+q+1} = \operatorname{span}\{u_1, \dots, u_{D+q+1}\}$. Consider the zero observation $y = 0$ and write $\mathcal{H}_0 = \Phi(0)$, so $\dim(\mathcal{H}_0) \le D$. Since A has rank at most q,

$$\dim\left(E_{D+q+1} \cap \ker A \cap \mathcal{H}_0^\perp\right) \ge (D + q + 1) - q - D = 1. \tag{80}$$

Choose a unit vector g in this intersection. Then $Ag = 0$, $g \perp \mathcal{H}_0$, and g is supported only on the first $D + q + 1$ graph eigenvectors, so

$$\|g\|_{H_G^s}^2 \le (1 + \mu_{D+q+1})^s. \tag{81}$$

Set $f = R(1 + \mu_{D+q+1})^{-s/2} g$. Then $f \in \mathcal{B}_s(R)$, $Af = 0$, the adaptive rule selects $\Phi(Af) = \Phi(0) = \mathcal{H}_0$, and $f \perp \mathcal{H}_0$. Therefore

$$\left\|(I - \Pi_{\Phi(Af)}) f\right\|_2^2 = \|f\|_2^2 = R^2 (1 + \mu_{D+q+1})^{-s}. \tag{82}$$

For the persistent recurrent class, $\lambda_t = f$ for all t, and the expected full-gradient error equals $A_T \left\|(I - \Pi_{\Phi(Af)}) f\right\|_2^2$ by the same calculation as in Theorem 3. □

A.9 Trajectory upper bound

Let $q_{t+1} = (I - \Pi_{\mathcal{H}_D^\star}) \lambda_{t+1}$. By Cauchy–Schwarz,

$$\left\|\nabla_W \mathcal{L} - \widehat{\nabla_W \mathcal{L}}_D\right\|_F = \left\|\sum_{t=0}^{T-1} q_{t+1} h_t^\top\right\|_F \tag{83}$$

$$\le \sum_{t=0}^{T-1} \|q_{t+1}\|_2 \, \|h_t\|_2 \tag{84}$$

$$\le \left(\sum_{t=0}^{T-1} \|h_t\|_2^2\right)^{1/2} \left(\sum_{t=0}^{T-1} \|q_{t+1}\|_2^2\right)^{1/2}. \tag{85}$$

The spectral tail bound gives

$$\sum_t \|q_{t+1}\|_2^2 \le (1 + \mu_{D+1})^{-s} \sum_t \|\lambda_{t+1}\|_{H_G^s}^2, \tag{86}$$

which proves (36).

A.10 Additional deterministic high-frequency instances

The stochastic theorem is the main recurrent lower bound. Two deterministic variants are useful for interpreting the omitted graph mode. First, for any eigenmode u_ℓ and scalar a, the slice construction in Appendix A.7 with $f = au_\ell$ gives $\lambda_T = au_\ell$ and $\nabla_W \mathcal{L} = au_\ell v^\top$. If $u_\ell \perp \mathcal{H}$, every $\mathcal{H}$-admissible rule incurs terminal-credit error a^2 and full-gradient error $a^2 \|v\|_2^2$.

Second, the obstruction can be generated by a graph-local residual rather than by an arbitrary terminal target. Let ∇_G be any weighted signed incidence matrix with $\nabla_G^\top \nabla_G = L_G$. For $\mu_\ell > 0$, set

$$h_t = \frac{a}{\mu_\ell} \kappa_\ell^{t-T} u_\ell, \qquad t = 0, \dots, T, \tag{87}$$

with no inputs and terminal loss $\mathcal{L} = \frac{1}{2} \|\nabla_G h_T\|_2^2$. Then $\lambda_T = L_G h_T = au_\ell$ and

$$\nabla_W \mathcal{L} = \sum_{t=0}^{T-1} \lambda_{t+1} h_t^\top = \frac{Ta^2}{\mu_\ell \kappa_\ell} u_\ell u_\ell^\top. \tag{88}$$

If $u_\ell \perp \mathcal{H}$, the constrained estimate is Frobenius-orthogonal to this matrix in its left factor, so the squared gradient error is its squared Frobenius norm. Choosing $a = R(1+\mu_\ell)^{-s/2}$ places the terminal adjoint on the Sobolev sphere.

A.11 Optimization theorem

The main-text optimization-floor consequence follows from the trajectory spectral bound (36) and the following standard inexact-gradient statement.

Theorem 5 (Inexact-gradient descent with spectral credit error)**.** *Assume F is L-smooth and satisfies the α-PL inequality*

$$\frac{1}{2} \|\nabla F(W)\|_F^2 \geq \alpha\big(F(W) - F^\star\big). \tag{89}$$

Run $W^+ = W - L^{-1}\widehat{g}_D(W)$, where $\widehat{g}_D(W) = g(W) + e_D(W)$. If $\|e_D(W)\|_F \leq \delta_D$ for all iterates, then

$$F(W^+) - F^\star \leq \left(1 - \frac{\alpha}{L}\right)(F(W) - F^\star) + \frac{\delta_D^2}{2L}, \tag{90}$$

and consequently

$$\limsup_{n\to\infty} \big(F(W_n) - F^\star\big) \leq \frac{\delta_D^2}{2\alpha}. \tag{91}$$

If $\|e_D(W)\|_F \leq \rho \|g(W)\|_F$ with $0 \leq \rho < 1$, then

$$F(W^+) - F^\star \leq \left(1 - \frac{\alpha}{L}(1-\rho^2)\right)(F(W) - F^\star). \tag{92}$$

If the exact adjoints and states satisfy the class condition defining $\mathcal{C}_s^{(T)}(R, H)$, then the low-frequency code obeys

$$\delta_D^2 \leq HR^2(1+\mu_{D+1})^{-s}. \tag{93}$$

Proof of Theorem 5. Let $g = \nabla F(W)$, $\widehat{g} = g + e$, and $W^+ = W - L^{-1}\widehat{g}$. Smoothness gives

$$F(W^+) \leq F(W) - \frac{1}{L}\langle g, g+e\rangle + \frac{1}{2L}\|g+e\|_F^2 \tag{94}$$

$$= F(W) - \frac{1}{2L}\|g\|_F^2 + \frac{1}{2L}\|e\|_F^2. \tag{95}$$

If $\|e\|_F \leq \delta_D$, the PL inequality implies

$$F(W^+) - F^\star \leq \left(1 - \frac{\alpha}{L}\right)(F(W) - F^\star) + \frac{\delta_D^2}{2L}. \tag{96}$$

Iterating this affine recursion gives the stated limit superior. If $\|e\|_F \leq \rho\|g\|_F$, the same display gives

$$F(W^+) - F(W) \leq -\frac{1-\rho^2}{2L}\|g\|_F^2 \leq -\frac{\alpha(1-\rho^2)}{L}(F(W) - F^\star), \tag{97}$$

which proves the relative-error rate. The spectral bound (93) is exactly the trajectory upper bound (36). □

A.12 Polynomial graph filters

Proposition A.2 (Local polynomial graph filters)**.** *Let $p(\mu) = \sum_{r=0}^{S} c_r \mu^r$ be a degree-S polynomial. Then $p(L_G)b$ can be computed from a source pattern b by S rounds of local graph communication. If, for a cutoff D,*

$$\max_{\ell \le D} |1 - p(\mu_\ell)| \le \varepsilon, \qquad \max_{\ell > D} |p(\mu_\ell)| \le \varepsilon, \tag{98}$$

then $\left\| p(L_G) f - \Pi_{\mathcal{H}_D^\star} f \right\|_2 \le \varepsilon \|f\|_2$ for all $f \in \mathbb{R}^N$. If $\mu_D < \mu_{D+1}$, exact spectral projection is obtained on a finite graph by a polynomial of degree at most $N - 1$.

Proof of Proposition A.2. Because L_G has nonzero entries only on vertices at graph distance at most one, $L_G^r b$ depends only on values within r graph hops of the support of b. Therefore $p(L_G)b = \sum_{r=0}^{S} c_r L_G^r b$ is computable by S rounds of local communication. For the approximation statement, write $f = \sum_\ell f_\ell u_\ell$. Then

$$p(L_G) f - \Pi_{\mathcal{H}_D^\star} f = \sum_{\ell \le D} (p(\mu_\ell) - 1) f_\ell u_\ell + \sum_{\ell > D} p(\mu_\ell) f_\ell u_\ell. \tag{99}$$

Orthogonality of the eigenvectors and the assumed spectral bounds give

$$\left\| p(L_G) f - \Pi_{\mathcal{H}_D^\star} f \right\|_2^2 \le \varepsilon^2 \sum_\ell |f_\ell|^2 = \varepsilon^2 \|f\|_2^2 . \tag{100}$$

If $\mu_D < \mu_{D+1}$, the cutoff does not split an eigenspace. A Lagrange interpolation polynomial on the distinct eigenvalues can then match the ideal values 1 on all eigenvalues at or below the cutoff and 0 on all eigenvalues above the cutoff, with degree at most the number of distinct eigenvalues minus one, hence at most $N - 1$. If the cutoff splits a repeated eigenvalue, no polynomial in L_G can distinguish the equal-eigenvalue directions; in that case the exact polynomial projector is the projector onto whole eigenspaces, while the minimax theorem may choose any D-dimensional subspace inside the degenerate eigenspace. □

B Product-graph extension

Let $L_x u_i = \mu_i u_i$ and $L_c v_j = \nu_j v_j$. The Kronecker-sum Laplacian $L_{x \times c} = L_x \otimes I + \gamma I \otimes L_c$ satisfies

$$L_{x \times c}(u_i \otimes v_j) = (\mu_i + \gamma \nu_j)(u_i \otimes v_j). \tag{101}$$

Since the proofs use only the spectral decomposition of a symmetric positive semidefinite Laplacian, every theorem holds with the sorted product eigenvalues $\mu_i + \gamma \nu_j$ and the corresponding product Sobolev norm. This gives the formal version of the space-by-cell-type extension mentioned in the main text.

C Nonlinear and spiking adjoint scope

Nonlinear rate networks and spiking networks require their own exact or surrogate adjoint definitions. Once a gradient admits a three-factor form with a postsynaptic credit field, however, the communication argument is unchanged: replacing that credit field by a D-dimensional broadcast/modulatory code creates the same projection problem. Surrogate-gradient and event-driven theories provide such adjoint notions in several spiking settings [Eilers et al., 2024, Holberg and Salvi, 2024]. For these models, the present theorems should be read as conditional statements: if the relevant adjoints inherit graph-Sobolev smoothness and the nonlocal learning signal lies in a D-dimensional code, then the same spectral compression limit applies.

D Empirical protocol details

For real data, the adjoint-spectrum analysis can be run retrospectively. Train a small BPTT RNN on a cognitive task, build an anatomical sheet graph or a functional correlation graph over units, compute adjoints during training, and report $\widehat{s}$ and the tail predictor (41). A motor-planning task

with a smooth somatotopic readout should have a small tail on a motor-map graph and a larger tail after shuffling. A context-dependent task that asks neighboring units to update in opposite directions should show the reverse: exact BPTT can solve it, but low-pass feedback should leave a gap unless a cell-type/product graph separates the contexts. An evidence-accumulation task with a global scalar decision variable is the easy case, because the adjoint is nearly constant across the graph and the first few modes capture most of the energy. The strong biological-alignment story is the specific pattern $\widehat{s}_{\text{sheet}} > 0$ and $\widehat{s}_{\text{shuffle}} \approx 0$, not merely good performance of a low-dimensional feedback rule.

The effective dimension should be measured, not equated with a molecule count. Given candidate source or receptor patterns b_k and taps P^r, form the basis matrix $B_{\text{code}} = [P^r b_k]_{k,r}$ and report its numerical rank or stable rank. Highly correlated receptor fields, redundant diffusion taps, or saturated nonlinear dose-response curves reduce D_{eff}; distinct receptor patterns that tile cell classes increase it. The theorem is stated in terms of this rank because two biochemical channels with nearly identical spatial action are one feedback dimension for credit assignment, whereas one molecule with multiple independent receptor-defined patterns can contribute more than one dimension.

The graph dimension d is likewise empirical. On a cortical sheet the low-frequency counting law should be close to $d = 2$ over a finite bandwidth, but a product graph over space and cell type can have a larger effective dimension, and a functional correlation graph can have a different spectral growth. Regressing $\log(1 + \mu_\ell)$ on $\log \ell$ over the feedback-relevant spectral range estimates $2/d$. Reporting both $\widehat{s}$ and $\widehat{d}$ makes the predicted exponent $2\widehat{s}/\widehat{d}$ auditable.

The protocol has three levels. A signal-level test asks whether measured adjoint tails follow the fitted Sobolev envelope. A learning-level test asks whether constrained-feedback loss gaps follow (41). A biological-alignment test asks whether the anatomical or cell-type graph gives a substantially larger $\widehat{s}$ than shuffled or random graphs. Passing only the first test supports the mathematical compression model; passing all three supports the stronger claim that the biological embedding is aligned with credit structure.

E Rule-task measured-adjoint experiment

Figure 2 repeats the adjoint-spectrum diagnostic on delayed spatial-memory rules rather than spectral hard families. Each task is a batch of cue-location trials on a 14×14 grid: smooth recall, shifted recall, two-cue binding, or checker-context recall. Targets are generated by these rules and are never chosen by specifying graph-Fourier coefficients. We train the same graph-filter RNN with exact BPTT and constrained feedback, measure the exact BPTT adjoints on the trajectory, and ask whether their graph tail predicts the fraction of exact-BPTT progress lost by low-pass feedback.

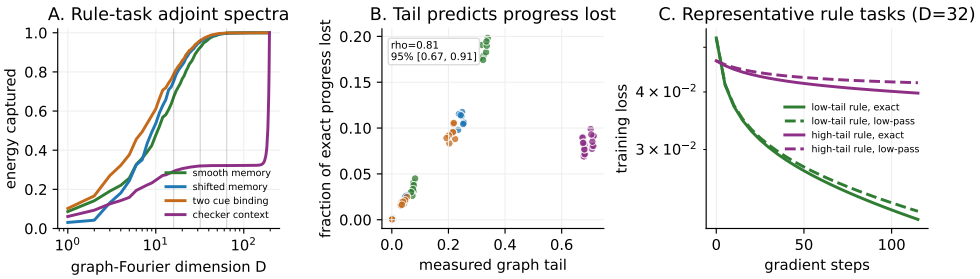

Figure 2: Rule-driven measured-adjoint check. Across 32 delayed-response tasks, batch size 18, horizon 12, and $D \in \{16, 32, 64\}$, measured graph-spectral adjoint tail predicts the fraction of exact-BPTT progress lost by graph-low-pass feedback (Spearman $\rho = 0.81$, 95% bootstrap interval $[0.67, 0.91]$). This experiment is not a biological benchmark; it is a non-hard-instance diagnostic showing that the spectral-tail predictor also appears when targets are generated by simple task rules rather than by prescribing residual graph spectra.

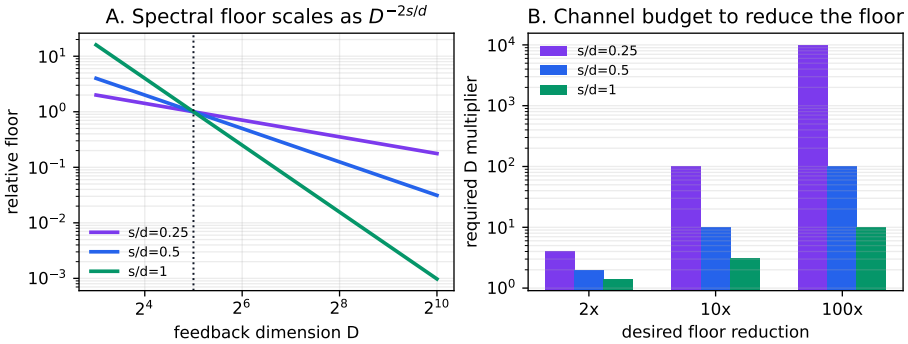

Figure 3: Channel-budget interpretation of the spectral credit-bandwidth law. Left: the relative credit, gradient, or PL optimization floor follows $E(D)/E(D_0) = (D/D_0)^{-2s/d}$ on manifold-like graphs. Right: reducing the floor by a factor m requires multiplying feedback dimension by $m^{d/(2s)}$. The theorem therefore predicts not only that high-frequency credit is hard, but how expensive it is to make it learnable with additional diffuse feedback channels.

F Channel-budget visualization

G Large-scale numerical stress test

Figure 4 complements the learning experiment in Figure 1 with a scalability check. For a two-dimensional torus grid, graph-Laplacian eigenvectors are Fourier modes and the eigenvalues are analytic, so omitted low-pass credit energy can be measured by FFT without dense eigendecomposition. We used this structure to evaluate grids with 256^2, 512^2, and $1024^2 = 1{,}048{,}576$ nodes. Smooth heat-filtered credit fields are highly compressible, whereas high-pass, checkerboard, and hard Fourier-mode credit fields remain almost entirely outside a $D \approx 65{,}000$ low-pass feedback subspace on the million-node graph.

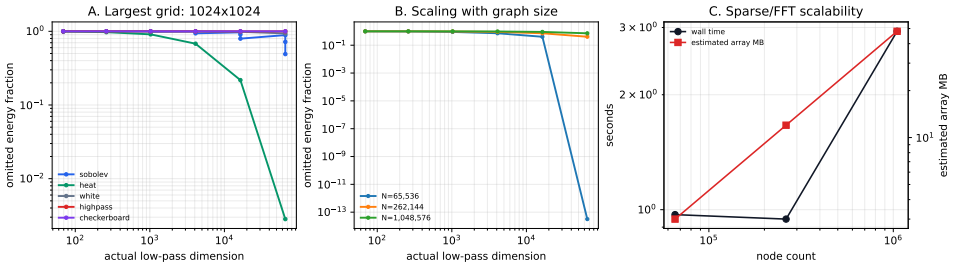

Figure 4: Large-scale FFT stress test on torus-grid graphs up to $1{,}048{,}576$ nodes. Low-pass feedback captures smooth heat-filtered credit fields with small omitted energy, but high-frequency fields remain invisible even at large feedback dimension. The right panel reports the CPU JAX run-time and estimated array footprint for the sparse/FFT implementation, illustrating that the spectral credit-bandwidth phenomenon can be evaluated at million-node scale without full eigendecomposition.

H NeurIPS paper checklist

1. **Claims.** The abstract and introduction state the paper's main claims: smoothness inheritance for graph-filter and Sobolev-stable nonlinear adjoints, an admissible-class theorem, a projection theorem for actual recurrent gradient error, a tight graph-Sobolev width law, a persistent stationary recurrent minimax theorem with $A_T = \Theta(T)$ full-gradient error, finite-repertoire and information-limited adaptive lower bounds, a formal e-prop/ModProp-like signal-class corollary, deterministic high-frequency residual variants, product-graph and nonlinear/spiking scope statements, quantitative channel-budget predictions, and a spectral optimization-floor consequence. The formal assumptions are stated before each theorem.

2. **Limitations.** Limitations are discussed in Section 8. The core minimax theorem is for stable linear RNNs driven by a stationary Gaussian input and fixed, finite-repertoire, or information-limited admissible feedback subspaces. Nonlinear or spiking applications require a separately justified exact or surrogate adjoint and the stated Sobolev-stability/smoothness condition.
3. **Theory assumptions and proofs.** All theorem assumptions are stated explicitly. Full proofs are included in the appendix. The K-source versus effective-dimension accounting is handled by Proposition 3; all lower bounds are in terms of $D = \dim(\mathcal{H})$.
4. **Experiments.** This is primarily a theory paper. Figure 3 visualizes the theorem's channel-budget scaling. Figure 1 gives a JAX measured-adjoint learning experiment showing that omitted graph-spectral credit energy predicts constrained-feedback loss gaps across task families and feedback dimensions. Appendix Figure 2 adds a rule-generated delayed-response check whose targets are not specified spectrally. Figure 4 adds a large-scale FFT stress test on torus-grid graphs up to $1{,}048{,}576$ nodes.
5. **Datasets, code, and compute.** No external datasets, learned models, or benchmarks are used. The channel-budget figure is generated by `scripts/render_credit_budget_figure.py`. The measured-adjoint learning experiment is generated by `scripts/run_jax_measured_adjoint_rnn_experiment.py`; it uses a procedurally generated 16×16 grid, 40 matched-energy sequence tasks, horizon 10, $D \in \{16, 32, 64\}$, exact BPTT, graph-low-pass feedback, random same-dimensional feedback, and K-source diffusion feedback. The rule-task check is generated by `scripts/run_jax_rule_task_adjoint_experiment.py`; it uses a 14×14 grid, 32 delayed-response tasks, batch size 18, horizon 12, and the same feedback dimensions. The large-scale experiment uses procedurally generated torus grids with 256^2, 512^2, and 1024^2 nodes and FFT-based graph Fourier transforms. The JAX runs are CPU-only and complete in seconds to a few minutes on a laptop-scale machine.
6. **Reproducibility.** The formal claims are reproducible from the theorem statements and proofs. The numerical figures use fixed procedural graph constructions, fixed seeds, stated feedback dimensions, and saved diagnostic quantities; no training data or learned model checkpoint is involved.
7. **Ethics and broader impacts.** The work is theoretical and does not introduce deployed systems, human-subject data, or safety-critical applications. The broader impact is a better understanding of limits on biologically plausible learning and communication-constrained optimization.
8. **Licenses and assets.** No external datasets or pretrained assets are introduced. Figures 3, 1, 2, and 4 are generated from procedural data or theorem-implied analytic curves produced by the authors.
9. **Responsible reviewing.** The authors acknowledge the responsible reviewing obligations associated with NeurIPS submissions.