Somato-Dendritic Innovation Learning

Local credit assignment from soma-unpredicted dendritic signals

Inspired by the neuron-specific somato-dendritic residuals reported by Francioni et al., Nature 2026

Question: should a local learner use apical activity itself, or the component that is unexpected from the same neuron's somatic state?

Project overview · August 2026

From the biological observation to a local learning rule

The Harnett-lab study fits each neuron's normal soma–dendrite relationship. The residual is decorrelated from soma and carries outcome and signed task information.

$$a_l = s_l + n_l$$

mixed apical activity

$$\hat{a}_l = P_l(h_l)$$

neutral-period prediction

$$r_l = a_l - \hat{a}_l$$

somato-dendritic innovation

$$\Delta W_l = \eta \, (r_l \odot \varphi'(u_l)) \, h_{l-1}^{T}$$

presynaptic activity × local postsynaptic gain × dendritic innovation

- Training uses locally available quantities, forward observations, and manual synaptic updates.
- Node-perturbation feedback supports the small-network experiments; reciprocal Kolen–Pollack plasticity supports the standard ResNets.
- The inherited credit path transports layer-specific directions. SDIL removes soma-predictable traffic from those directions.

Sources: Francioni et al. (Nature 2026); Lansdell et al. (ICLR 2020); Akrout et al. (NeurIPS 2019).

Residualization is load-bearing under predictable traffic

Controlled MNIST intervention · depth 3 · width 256 · five paired seeds

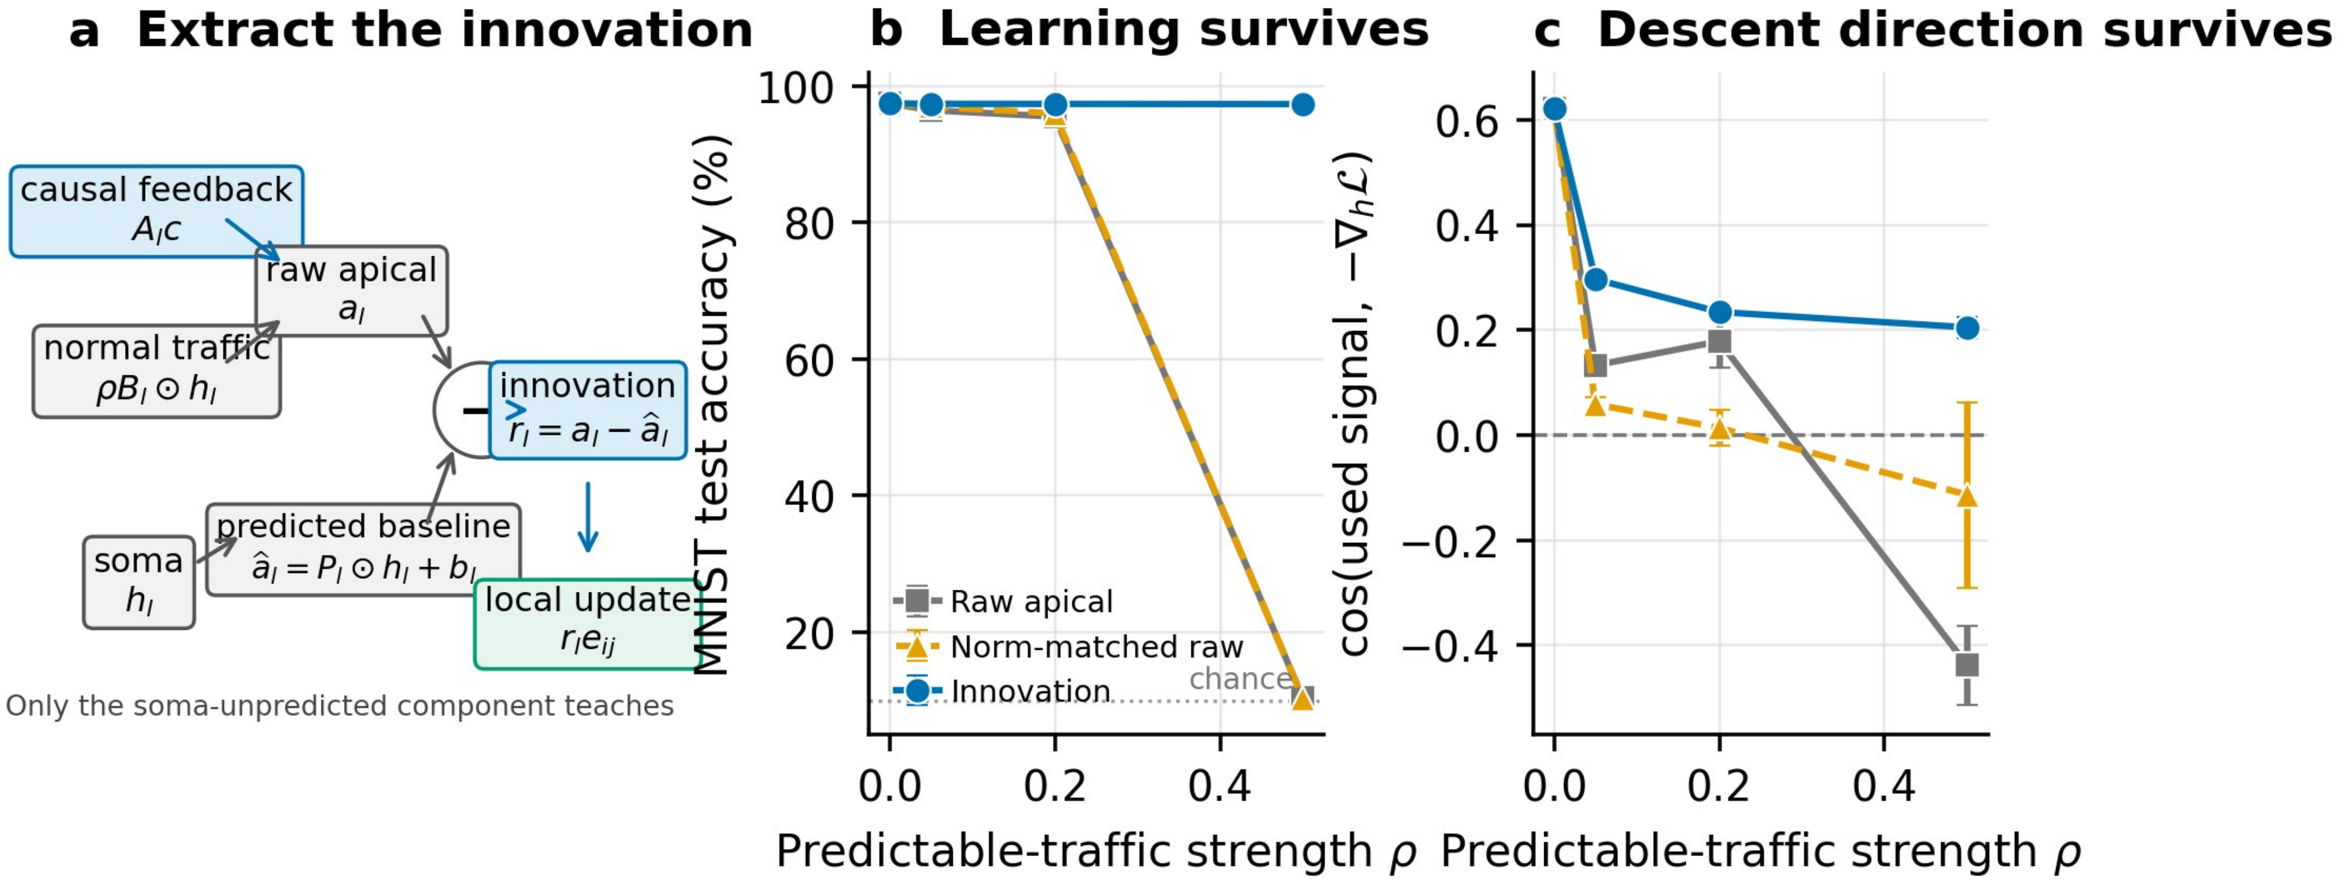

At traffic strength ρ = 0.5: raw 10.38%, norm-matched raw 10.31%, innovation 97.35%.

Norm matching equalizes per-example update magnitude. The remaining difference comes from teaching direction.

Dynamic innovation survives a standard ResNet-20 test

CIFAR-10 · 200 epochs · five untouched paired seeds · four-times-RMS soma-predictable traffic

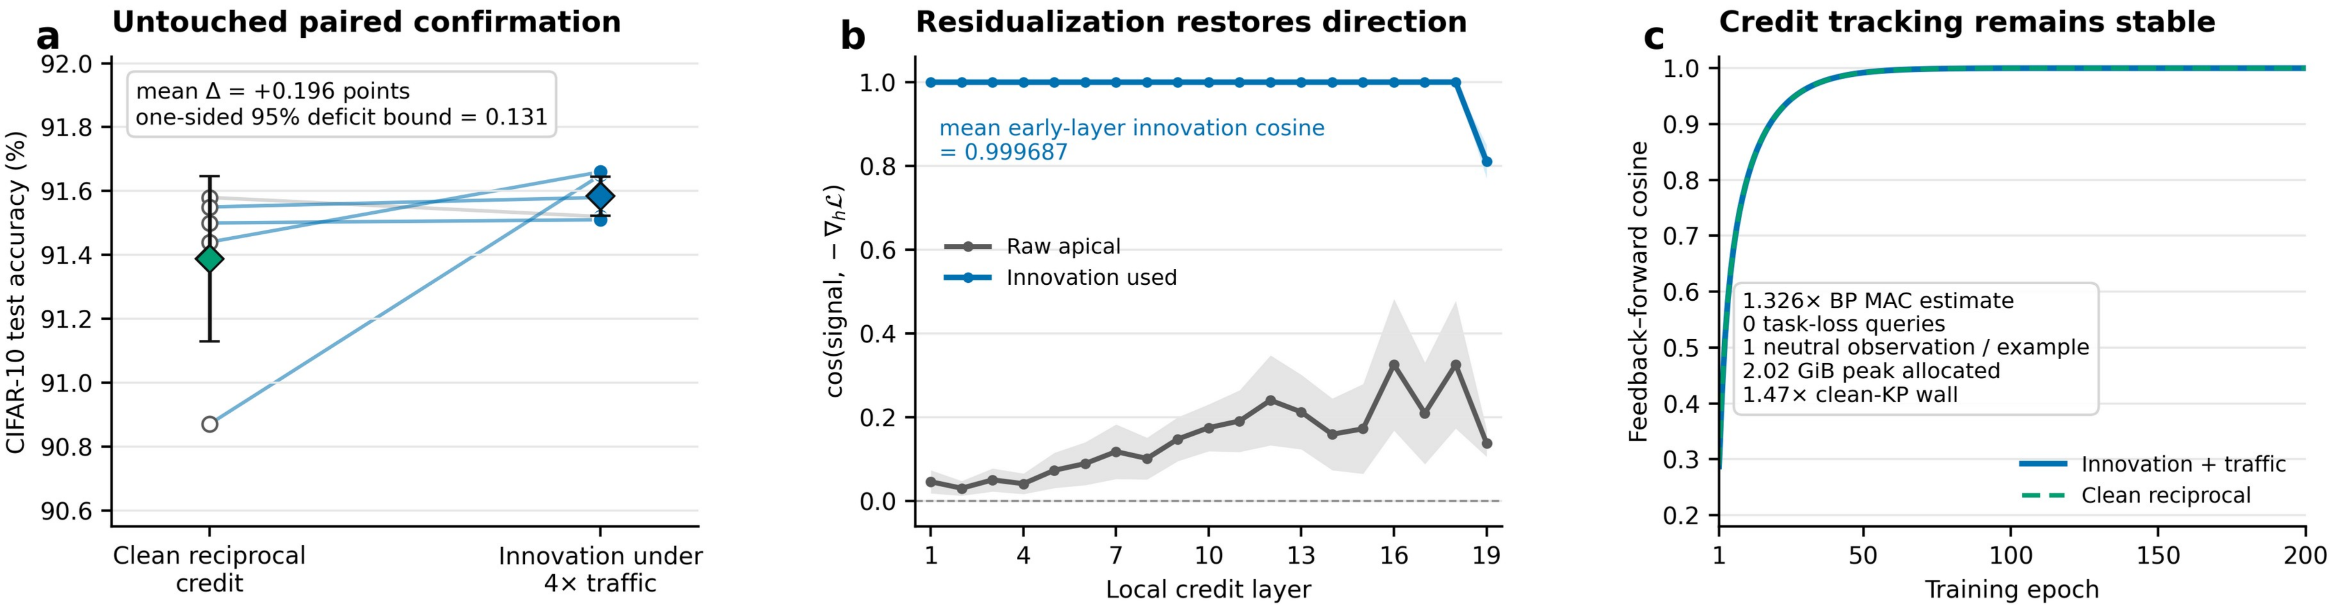

- Dynamic SDIL under traffic: 91.584% mean test accuracy; clean reciprocal KP: 91.388%.
- Mean early-layer teaching-signal cosine: 0.999687.
- Cost: 1.326× the matched BP MAC estimate, zero task-loss queries, and one neutral observation per example.

Exact gradients serve diagnostic measurement; local signals drive every learning update.

Scaling from ResNet-20 to ResNet-56

Complete 60-endpoint CIFAR-10 panel · five seeds per method and depth · shared hyperparameters

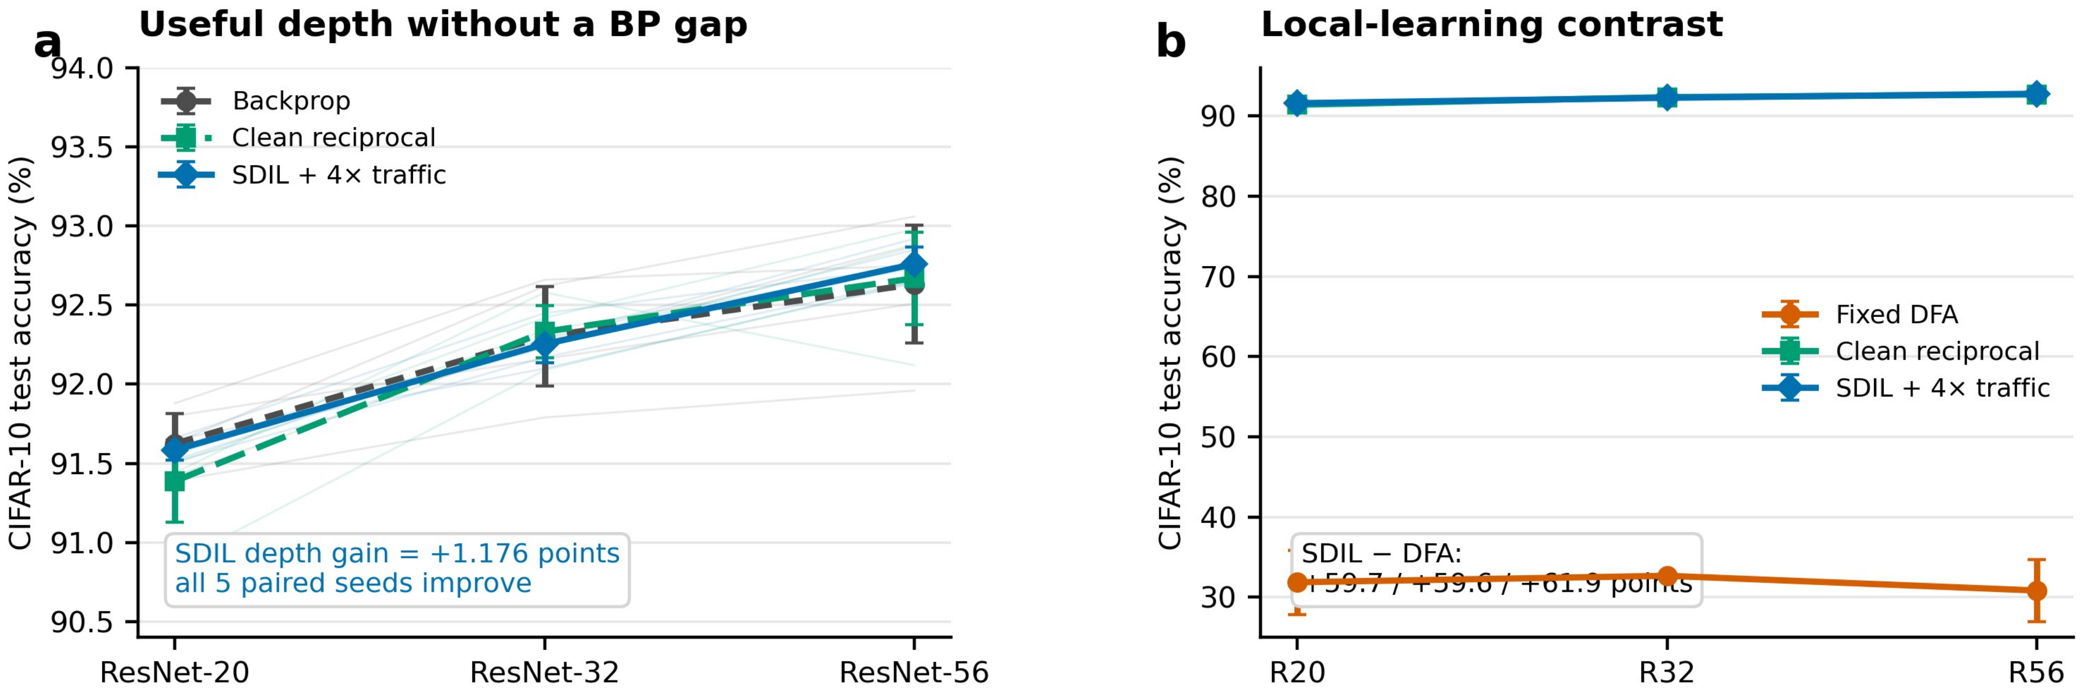

SDIL under 4× traffic: 91.584% → 92.254% → 92.760%; all five R20/R56 pairs improve.

Reciprocal KP provides the scalable credit path. SDIL preserves this path under mixed apical traffic.

BP and clean reciprocal KP follow the same depth trend; fixed DFA remains near 31%.

A local actor–critic produces temporal outcome signals

Synthetic BCI task · six task clusters × five model seeds · manual local updates

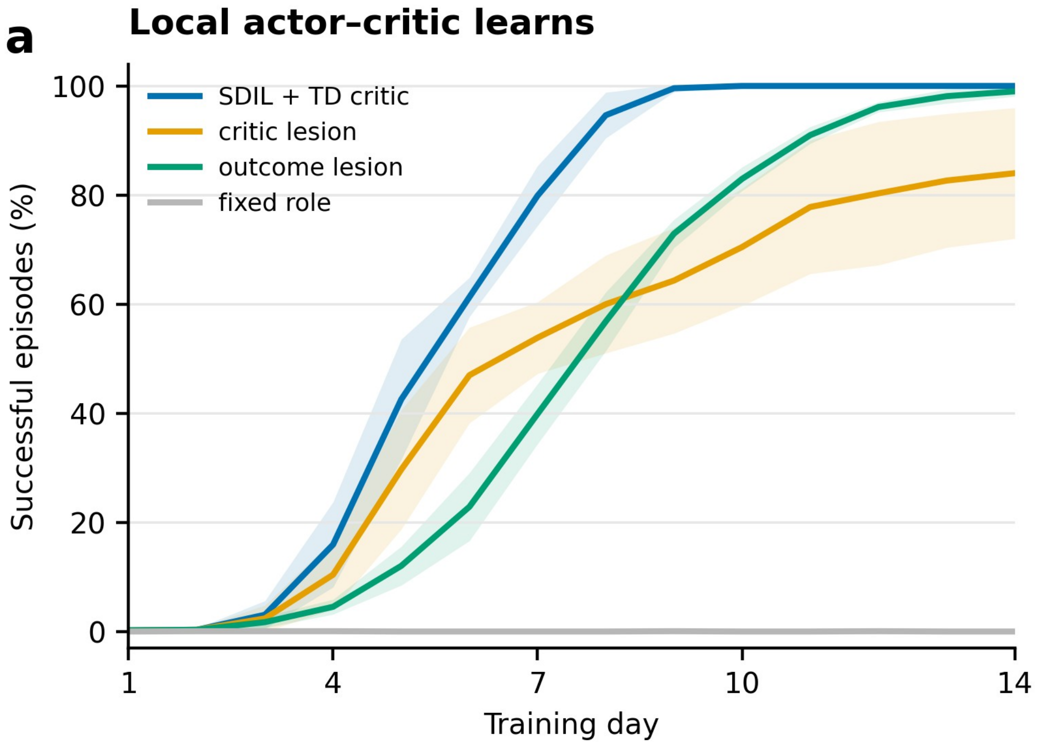

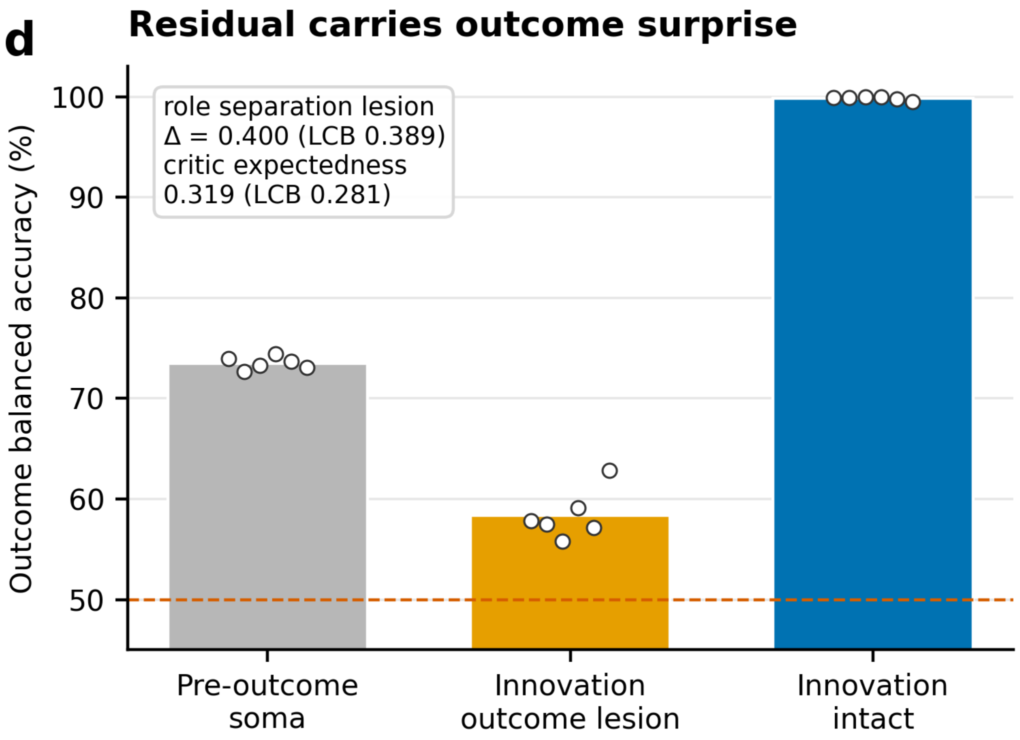

Final success: 100% · terminal outcome decoding: 99.83% · outcome-lesion effect: 0.400.

The synthetic task supplies terminal reward and tests neuron-specific outcome surprise with causal lesions.

Outcome labels and exact causal roles serve evaluation.

Current scientific position

Established

- Soma-predictable apical traffic can rotate a raw local teaching signal.
- Per-neuron innovation preserves learning direction under controlled traffic.
- Dynamic innovation works with a scalable BP-free reciprocal credit path on ResNet-20/32/56.
- A separate local actor–critic reproduces neuron-specific outcome signaling in a controlled dynamical task.

Attribution

- Reciprocal KP provides clean-setting scale and feedback tracking.
- SDIL provides soma-conditioned traffic removal and the associated mechanism tests.

Next decisive experiment

Evaluate raw reciprocal credit and SDIL under naturally generated, state-dependent mixed traffic with component-level bias hidden from the learner.

Target claim: scalable local credit assignment with adaptive removal of predictable feedback contamination.