PING tolerates 80% dropped spikes but collapses on added noise

exp037 · 30 May 2026

The trained networks this entry uses are produced once in the shared training hub, exp022 (Training), and reused here rather than retrained.

Abstract

Perturbs the hidden spike stream of trained PING and COBA networks at inference (drop spikes, add spikes) to ask whether the PING rate floor is dynamical or informational. PING tolerates dropping ≈ 80% of its emitted spikes (accuracy 85% at that level, from 91% unperturbed) yet collapses once added Poisson noise passes ≈ 81% of its own baseline rate. COBA is roughly flat to both, holding 89% across the whole add sweep (which reaches only ≈ 22% of its far higher baseline). The asymmetry, drops forgiven while adds break the gating, is the gamma cycle made visible: a structural feature of the architecture, not a readout-side trade-off.

Method

Parameter	Value
Integration timestep Δt	0.1 ms
Trial duration T	200 ms
Held-out MNIST test samples	500
Baseline training epochs (in exp022)	30

The PING and COBA baseline definitions and training recipe are in exp025, where the rate-floor mechanism is also worked out. This entry tests whether the floor is **dynamical** (locked by the cycle period) or **informational** (locked by the readout's spike-count requirement) by perturbing the hidden spike stream of the trained networks at inference.

The perturbation. A per-step callback on the COBANet's *_hidden_perturb_fn* slot fires every timestep of every trial, with no warm-up, schedule, or exclusion. Trials are $T = 200$ ms at $\Delta t = 0.1$ ms → 2000 fires per trial, applied across the held-out test set against the same trained network. At each timestep t:

1. *update conductances* from the previous step's spikes: $g_e^{(t)} \leftarrow g_e^{(t-1)} \cdot d_{\text{AMPA}} + \tilde{\boldsymbol{s}}^{E,(t-1)} W_{ee} + \text{input}^{(t)} W_{\text{in}}$, with the analogous update for $g_i^{(t)}$ from $\tilde{\boldsymbol{s}}^{I,(t-1)} W_{ie}$ and the E→I conductance from $\tilde{\boldsymbol{s}}^{E,(t-1)} W_{ei}$;
2. *LIF step* integrates $V^{(t)}$ and emits raw spike vectors $\boldsymbol{s}^{E,(t)} \in \{0,1\}^{B \times N_E}$, $\boldsymbol{s}^{I,(t)} \in \{0,1\}^{B \times N_I}$;
3. *perturbation callback* rewrites the raw vectors → $\tilde{\boldsymbol{s}}^{E,(t)}, \tilde{\boldsymbol{s}}^{I,(t)}$;
4. *record* the perturbed vectors into the spike buffer;

5. *readout accumulator* adds $\tilde{\boldsymbol{s}}^{E,(t)} W_{\text{out}}$ to the mem-mean integrator.

Step 1 of timestep $t+1$ consumes the perturbed spikes through W_{ee}, W_{ei}, W_{ie}: a dropped E spike fails to drive I next step and contributes nothing to the readout; an injected I spike adds inhibition next step and counts in the rate metric. E and I get the same mode and level with independent draws.

Drop mode. For each (batch, neuron, timestep) slot i, draw $u_i \sim \text{Uniform}(0,1)$ i.i.d. and keep each emitted spike with probability $1 - p_{\text{drop}}$:

$$\tilde{s}_i = s_i \cdot \mathbb{1}\left[u_i \geq p_{\text{drop}}\right].$$

Drop thins the count fed to both the readout and the next-step conductance update while leaving the E→I→E loop intact. Sweep $p_{\text{drop}} \in \{0.0, 0.1, ..., 1.0\}$, 11 levels.

Add mode. For each slot draw $u_i \sim \text{Uniform}(0,1)$ i.i.d. and flip silent slots on at Poisson statistics, phase-independent:

$$\tilde{s}_i = \min(s_i + \mathbb{1}[u_i < r_{\text{add}} \cdot \Delta t / 1000], 1).$$

Sweep $r_{\text{add}} \in \{0, 2, ..., 40\}$ Hz per neuron, 21 levels, applied equally to E and I. Because the two architectures sit at very different baselines (COBA ≈ 181 Hz per E cell versus PING ≈ 12 Hz, a factor of ≈ 15), a fixed added rate is a much larger relative insult to PING than to COBA, so the right panel of Figure 1 expresses the added rate as a percentage of each model's own baseline E rate (the architecture-fair view). The per-step RNG is seeded separately from the input encoder, so the Poisson input stream matches the unperturbed baseline. Total: 2 models × (11 drop + 21 add) = 64 forward passes.

Results

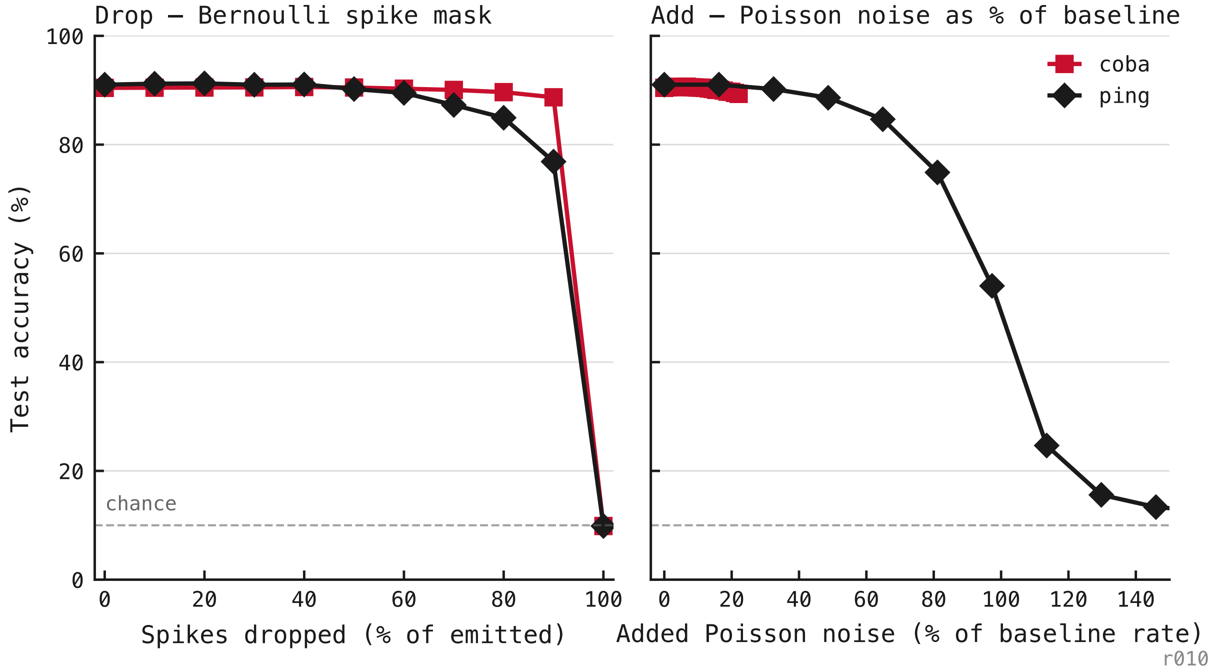

Figure 1: **Left (drop):** Bernoulli mask, percent of emitted spikes dropped. **Right (add):** Poisson injection as a percent of each model's own baseline E rate, so the two architectures are insult-matched. The asymmetry is direct: PING (black) holds accuracy to ≈ 85% at 80% drop but collapses once added noise passes ≈ 81% of its baseline, while COBA (red) stays at 89% across its whole sweep, which reaches only ≈ 22% of its far higher baseline. The dashed line marks chance.

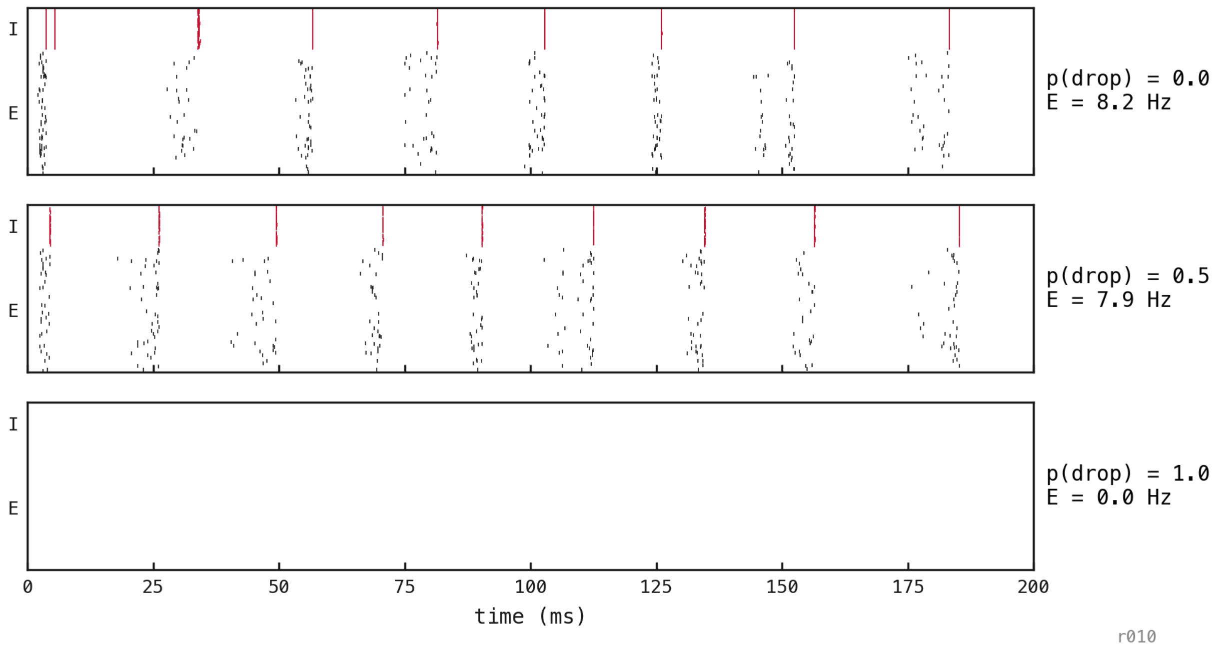

Figure 2: Trained PING replayed on the same MNIST digit 0 trial across three drop levels (0, 50, 100%); E (black) above I (red). The gamma cadence persists as spikes are thinned and only vanishes at total drop: dropping spikes thins the train without injecting phase-incoherent activity.

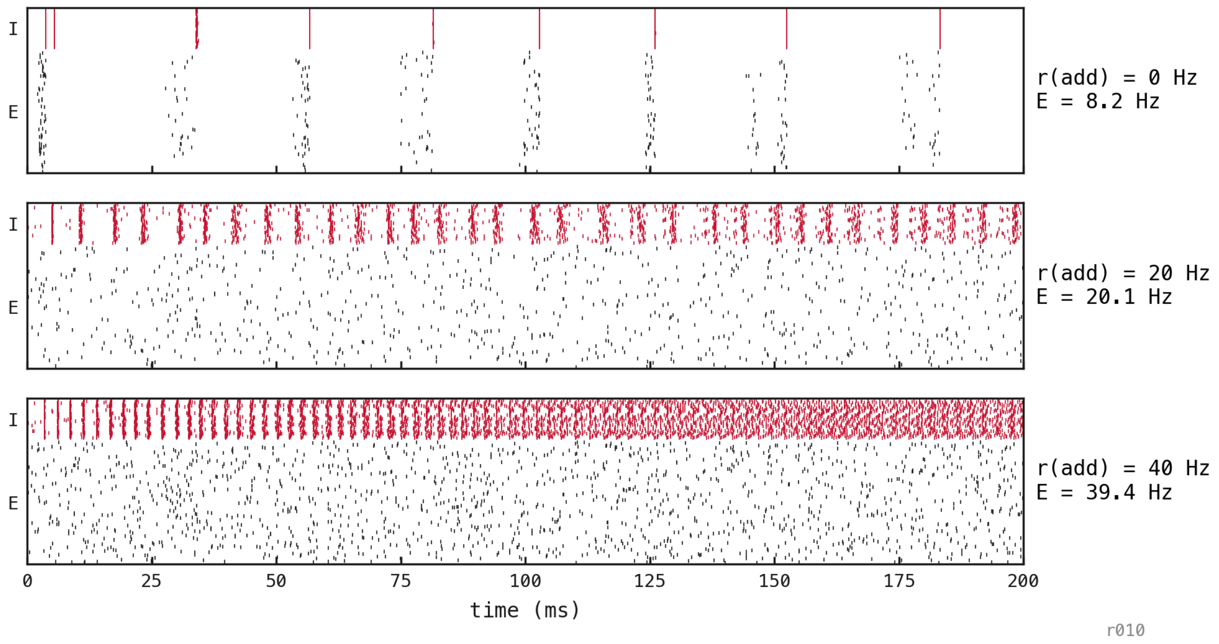

Figure 3: Trained PING replayed across three added-noise levels (0, 20, 40 Hz per neuron). The gamma banding dissolves into asynchronous firing as the added rate rises, matching the accuracy cliff in the previous figure.

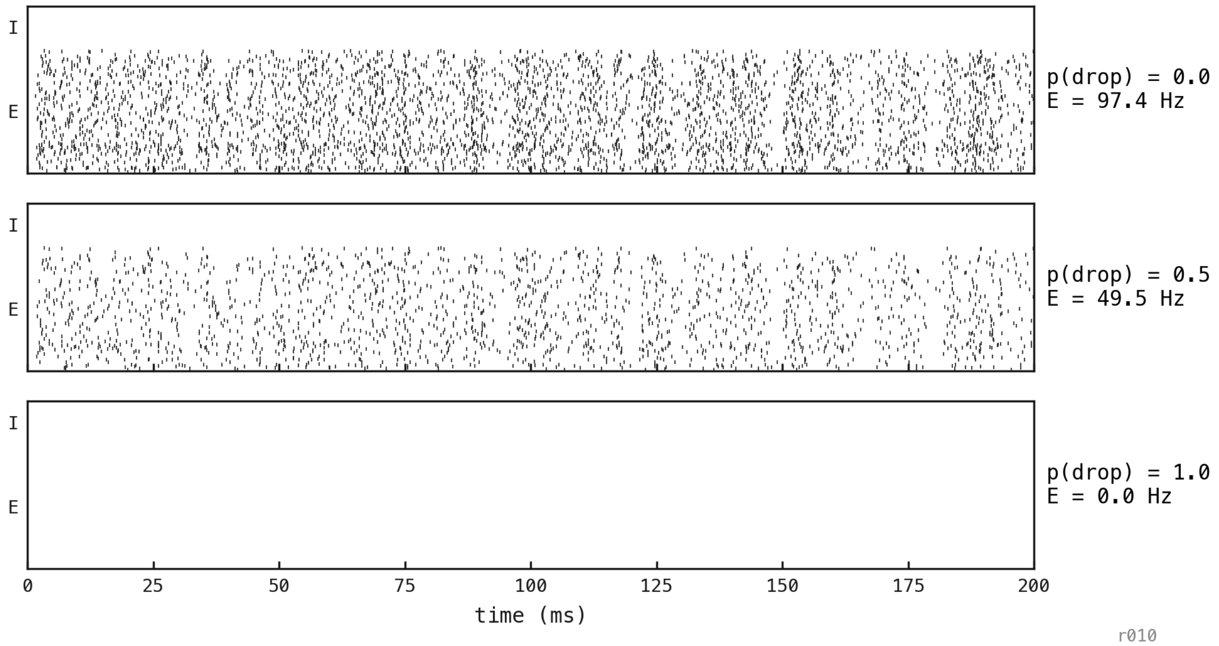

Figure 4: Trained COBA replayed across the same drop sweep. With no cycle to preserve, dropping spikes just thins a uniform asynchronous mean.

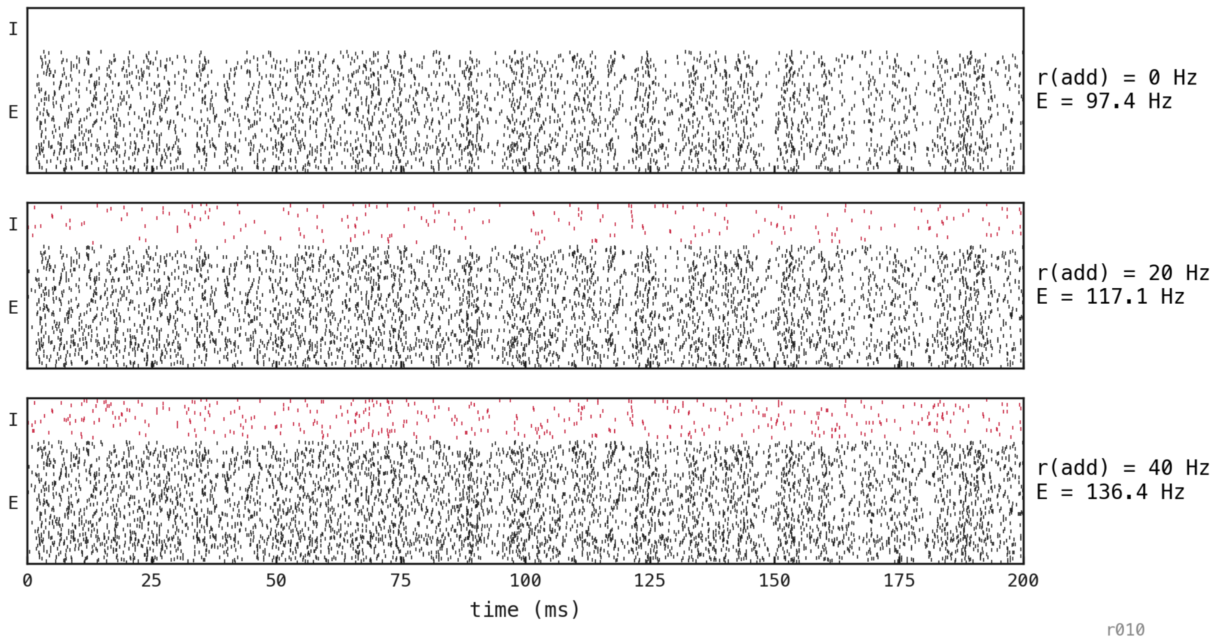

Figure 5: Trained COBA replayed across the same add sweep. Added spikes blend into COBA's asynchronous mean: no temporal structure to corrupt, just a higher mean rate.