E rate is affine in gamma frequency

exp041 · 2 June 2026

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

Abstract

exp037 varied τ_{GABA} at inference and the per-cell E rate tracked the gamma cycle. Does that survive *re-training* at each τ_{GABA}? Yes. Across $\tau_{\mathrm{GABA}} \in \{4.5, 6, 9, 12, 18, 27\}$ ms × 3 seeds, $r_E = 1.15 + 0.183 \cdot f_\gamma$ with $R^2 = 0.994$. The slope is per-cycle E-cell participation; the intercept is a non-rhythmic baseline; accuracy stays at ≈ 86–92% across the sweep. The cycle clock constrains what the network can become.

Method

Sweep. Six τ_{GABA} values $\{4.5, 6, 9, 12, 18, 27\}$ ms × three seeds = 18 networks, trained in the shared hub to the gamma standard (50 epochs on MNIST, Adam at 4×10^{-4}, batch 256, mem-mean readout, no spike budget, $\Delta t = 0.1$ ms, $T = 200$ ms); only τ_{GABA} varies.

Measuring f_γ. For each cell, f_γ is the parabolic-interpolated peak of the Welch PSD on the per-trial population E trace; the fit (Figure 4) uses per-trial peak medians, which avoid the centroid bias of trial-mean PSD peaks. Parabolic interpolation with peak-bin values (y_0, y_1, y_2):

$$f_\gamma = \mathrm{freq[peak]} + \frac{1}{2}\frac{y_0 - y_2}{y_0 - 2y_1 + y_2} \cdot \Delta f, \quad \Delta f = 5 \text{ Hz}.$$

It is needed because the bare 5 Hz bin quantisation would coarsen f_γ across the six conditions; on a well-isolated peak the interpolation error is $O((\Delta f)^3)$.

The predicted law. The shape $r_E = a + p \cdot f_\gamma$ is predicted by cycle dynamics, not curve-fitted. Within one cycle of duration $1/f_\gamma$, a fraction p of E cells emits exactly one spike (those nearest threshold when the I shunt drops); the rest are still recovering, so the cyclic per-cell rate is $p \cdot f_\gamma$. At long τ_{GABA} the I conductance never fully decays and the cycle dissolves into a tonic bath, leaving a feedforward baseline a independent of f_γ:

$$r_E = \underbrace{a}_{\text{feedforward baseline}} + \underbrace{p \cdot f_\gamma}_{\text{cyclic contribution}}.$$

We fit this across the 18 cells.

Convergence. Accuracy plateaus by ≈ epoch 15 while the E rate keeps climbing through training (Figure 1), so the fit uses the final-epoch rates. Fitting $r_E = a + p \cdot f_\gamma$ across the 18 cells is tight, and forcing the intercept through zero barely loosens it, so the law is not an artefact of the free intercept:

fit	a (Hz)	p (Hz/Hz)	R^2
affine	1.15	0.183	0.994
through origin	0	0.213	0.966

Results

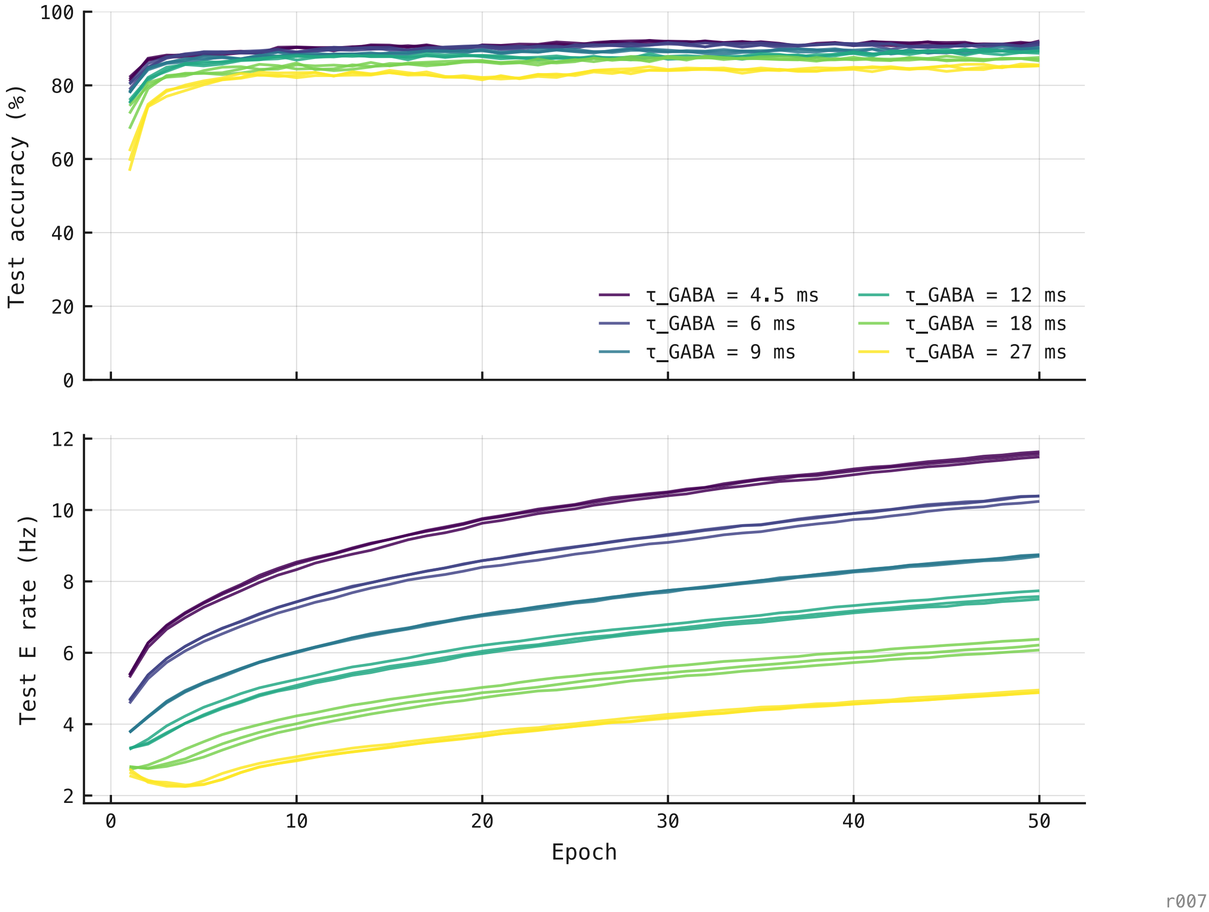

Figure 1: Per-cell accuracy (top) and E rate (bottom) over training, one line per cell. Accuracy plateaus by ≈ epoch 15; the E rate keeps climbing through the 50 epochs, so the final-epoch rates are the ones fit.

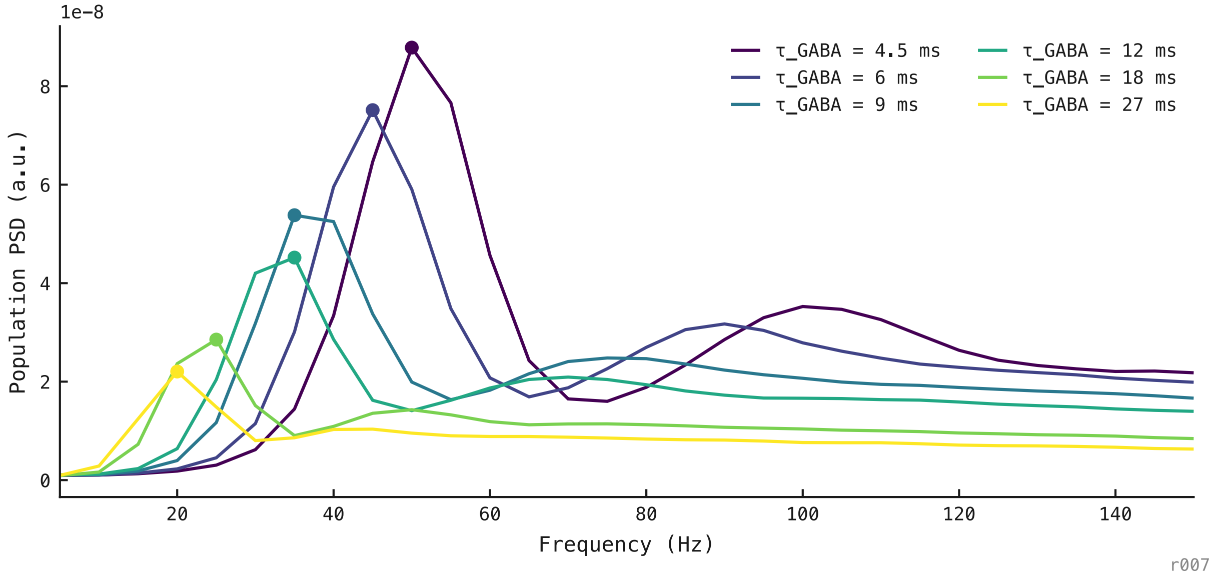

Figure 2: Trial-mean Welch PSDs by τ_{GABA}; dots mark the parabolic-interpolated peak. The peak shifts cleanly from ≈ 20 Hz at $\tau_{\text{GABA}} = 27$ ms to ≈ 51 Hz at $\tau_{\text{GABA}} = 4.5$ ms, with no overlap between adjacent conditions.

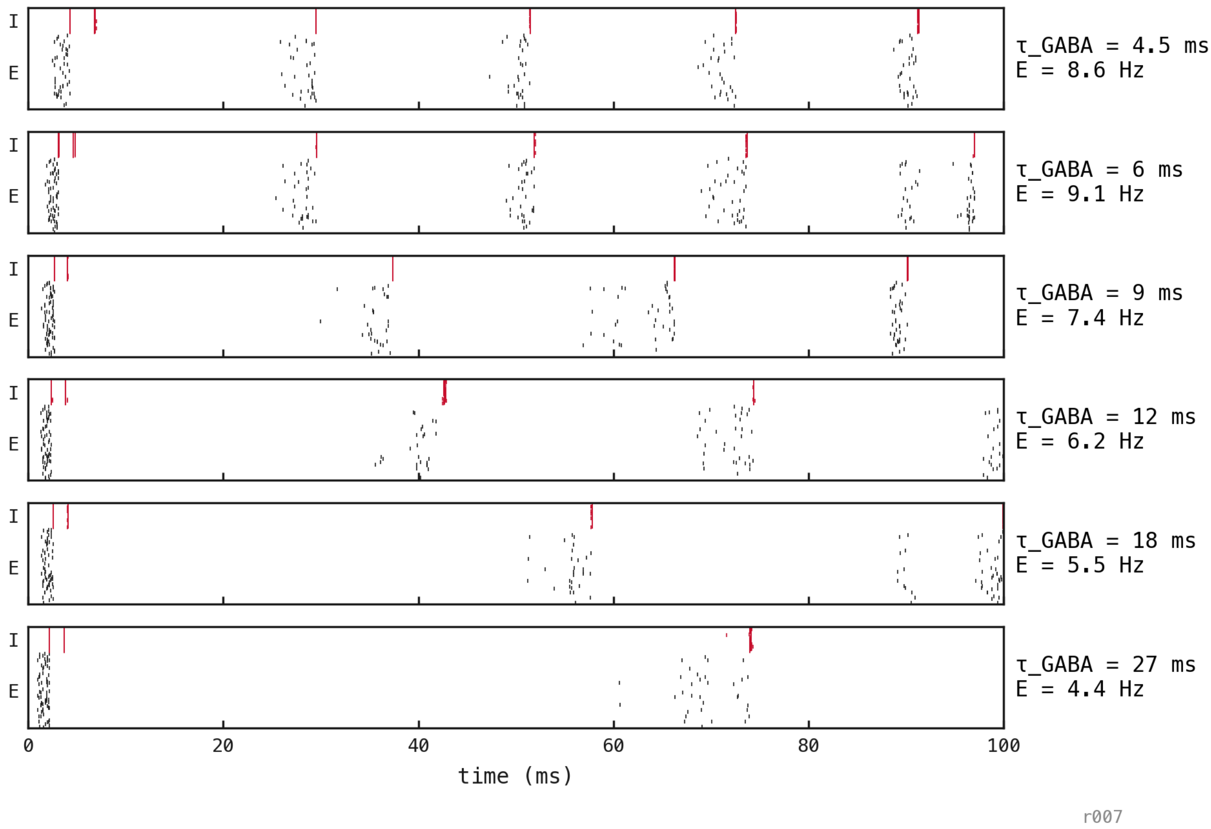

Figure 3: One MNIST trial through each network. The cycle period stretches from ≈ 20 ms ($f_\gamma \approx 51$ Hz) at short τ_{GABA} to ≈ 50 ms ($f_\gamma \approx 20$ Hz) at long, the eye and the spectrum agreeing.

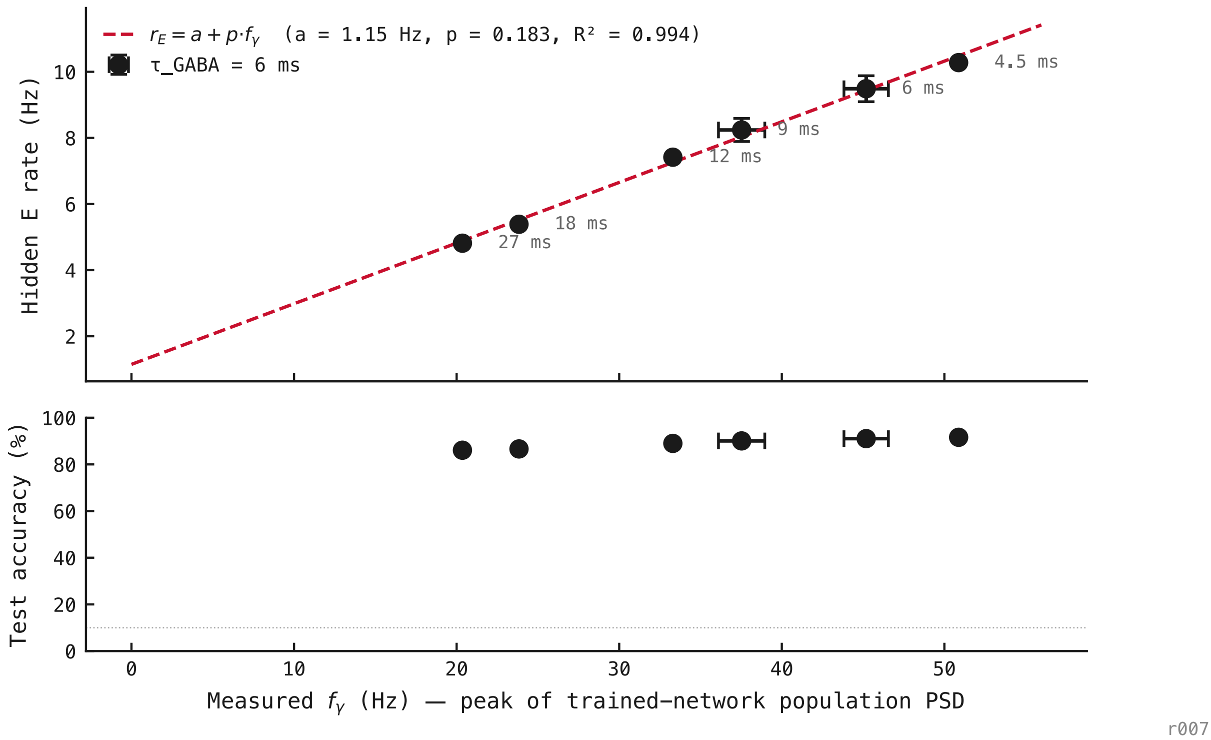

Figure 4: The law itself. Top: mean post-training E rate vs f_γ, six clusters × three seeds, error bars from seed variance; the affine fit passes through every error bar. Bottom: per-cluster accuracy is flat, so the rate change is not paid in classification.