

Variable-rate streaming with a spike-rate readout

exp082 · 10 August 2026 · Draft

Abstract

This experiment will test whether PING networks trained across input rates can classify continuous MNIST streams without the fixed-rate and mean-membrane mismatch inherited by exp048. Excitatory spikes drive ten spiking output LIF neurons, and each class logit is the corresponding output neuron’s firing rate over a window equal to the current digit’s presentation duration. The planned results comprise matched-condition streaming, variable-rate and variable-duration streaming, a 200 ms rate psychometric, and a duration-by-rate accuracy map. Results remain pending until exp022 produces all three variable-rate checkpoints.

Methods

1. Training source

The experiment loads `ping__variable_rate__seed42`, `ping__variable_rate__seed43`, and `ping__variable_rate__seed44` from exp022. Each PING network is trained by sampling one maximum-pixel Poisson rate independently per image presentation from 0.5, 0.75, 1, 1.5, 2, 3, 5, 7.5, 10, 15, and 25 Hz. The recurrent weights and classifier are frozen throughout this experiment.

Note. The exp022 cells are registered and `tools/snn` implements the required output-LIF `spike-rate` training, restore, and output-raster path. Presentation boundaries reset only the output LIF and its counter; the hidden PING state remains continuous. This entry still fails loudly if the variable-rate checkpoint bank is absent or uses another readout.

2. Spike-rate readout

At each timestep, the 1024-element excitatory spike vector $\mathbf{s}^{E(t)}$ drives a trained 1024-by-10 projection W_{out} into ten output LIF neurons. For a presentation beginning at timestep a and ending before timestep b , class evidence is

$$z_c = \frac{1}{T} \sum_{t=a}^{b-1} s_c^{\text{out}(t)}, \quad c \in \{0, \dots, 9\}, \quad T = (b - a) \Delta \frac{t}{1000}. \quad (1)$$

Here $s_c^{\text{out}(t)}$ is the spike emitted by output LIF neuron c at timestep t , T is the matched presentation duration in seconds, and $\mathbf{z}$ contains the ten output firing-rate logits in Hz. The predicted digit is the index of the largest rate. The output neurons therefore have membrane state, leak, threshold, spikes, and reset; unlike `mem-mean`, their membrane voltages are not the logits.

3. Matched inference

Every digit is read from only the spikes generated during its own presentation:

$$T_{\text{readout}} = T_{\text{presentation}} \cdot (2)$$

Here T_{readout} is the spike-summation window and $T_{\text{presentation}}$ is the duration of the corresponding digit. The first protocol holds both at 200 ms and presents a five-digit stream. This checks the trained duration and readout before either input rate or duration is varied.

4. Variable-rate and variable-duration inference

A second five-digit stream changes both presentation duration and maximum-pixel Poisson rate at digit boundaries. Each digit retains its own matched readout window. A factorial evaluation then crosses presentation durations 25, 50, 100, and 200 ms with the eleven training rates. Every cell is evaluated independently for seeds 42, 43, and 44.

5. Psychometric protocol

The rate psychometric fixes presentation and readout duration at 200 ms, varies the maximum-pixel rate across the eleven training values, and reports held-out classification accuracy across the three trained seeds. This isolates rate robustness from shortened evidence windows.

Results

1. Matched-condition stream

TODO. Reproduce the structure of exp048’s streaming figure: excitatory, inhibitory, and output rasters; digit boundaries; and the ten online class-evidence traces. Unlike exp048, every trace will be the cumulative output-spike rate from the start of the current digit: spikes observed so far divided by elapsed time in the matched window.

2. Variable conditions

TODO. Reproduce exp048’s joint temporal and encoding-rate summary using the variable-rate-trained weights. Pair the duration-by-rate accuracy map with the 200 ms psychometric curve, so the effects of rate and available integration time remain distinguishable.

3. Comparison with exp048

TODO. Compare the completed results with exp048. The comparison must be interpretive rather than a silent replacement. Exp048 uses fixed-25-Hz training and reconstructs a sliding mem-mean output. This experiment uses mixed-rate training and the trained spike-rate readout. Any accuracy difference therefore reflects the combined training-distribution and readout correction, not a single isolated intervention.