# Empirical input-rate calibration for variable-rate PING training

exp080 · 10 August 2026


## Abstract

This experiment asks which input-rate interval preserves enough digit information after synaptic and membrane filtering to justify using it in a variable-rate PING training run. We trained a nonlinear decoder on freshly simulated MNIST features spanning eight rates, then evaluated the frozen decoder on held-out images at each rate. The selected interval is 0.5 to 25 Hz. Its lower edge is the first tested rate at which all three decoders meet or exceed 50% held-out accuracy. The result calibrates this feature representation and decoder; it does not measure PING-network accuracy.


## Methods

1. **Partition the MNIST dataset.** The official MNIST training partition contains 60000 images. We assigned its first 20000 images to decoder training and the next 5000 images to checkpoint selection. The remaining 35000 images were not used. Final evaluation used the first 5000 images from the separate official test partition. Training, validation, and test images therefore did not overlap.
2. **Simulate filtered image features.**

   Each normalized pixel intensity $x_i$ generated an independent binary event at every $\Delta t = 0.1$ ms timestep,

$$S_i(t) \sim \text{Bernoulli}\left(r x_i \Delta \frac{t}{1000}\right). \quad (1)$$

   Here $r$ is the maximum-pixel encoding rate in spikes/s. Conductance followed

$$g_i(t) = \exp\left(-\Delta \frac{t}{\tau_{\text{AMPA}}}\right) g_i(t - \Delta t) + w S_i(t), \quad (2)$$

   with $\tau_{\text{AMPA}} = 2$ ms and $w = 1.2$ µS. The non-spiking conductance-based membrane obeyed

$$C_E \frac{dv_i}{dt} = g_{\text{L,E}}(E_L - v_i) + g_i(t)(E_e - v_i), \quad (3)$$

   with $C_E = 1$ nF, $g_{\text{L,E}} = 0.05$ µS, $E_L = -65$ mV, and $E_e = 0$ mV. Conductance began at zero and voltage at $E_L$. The decoder feature was the baseline-subtracted voltage averaged over the complete 200 ms presentation,

$$z_i = \frac{1}{T} \int_0^T (v_i(t) - E_L)\, \mathrm{d}t. \quad (4)$$

Every training, validation, illustration, and test feature used a newly drawn spike train and direct evaluation of Equations 1–4. The random spikes form *shot noise*: each spike produces a discrete jump in conductance, followed by the exponential AMPA decay in Equation 2. These conductance pulses alter both the voltage toward which the membrane moves and the speed at which it moves there. The effect of a spike therefore depends on the voltage and conductance left by earlier spikes, rather than adding a fixed voltage increment.

Direct simulation is particularly important at low input rates. A finite presentation may contain no spikes, one spike, or a few spikes arriving at different times. The response statistics can consequently change during the presentation (*nonstationary*) and the distribution across presentations can be asymmetric or concentrated around a few distinct outcomes (*non-Gaussian*) [**1**]. Evaluating Equations 1–4 for each presentation preserves these effects instead of replacing them with a steady, bell-shaped approximation.

3. **Train a mixed-rate decoder.**

   The official MNIST training partition supplied the first 20000 images for training and the next 5000 for validation. Every presentation sampled one of the eight rates 0.1, 0.25, 0.5, 1, 2, 5, 10, 25 Hz uniformly and independently. A 784–1024–10 ReLU decoder was trained with cross-entropy and Adam for 10 epochs. Seeds 42, 43, 44 defined independent initializations, rate assignments, and spike trains. Validation accuracy selected one checkpoint per seed.

4. **Evaluate held-out accuracy and select the interval.**
   1. **Select the held-out images.** We took the first 5000 images from the official MNIST test partition. These images were not used for training or checkpoint selection.
   2. **Simulate shared test features.** We simulated every held-out image once at each registered input rate. All three validation-selected decoders received the same simulated feature for a given image and rate, so differences among decoders were not caused by different spike realizations.
   3. **Measure each decoder's accuracy.** For each input rate and decoder, we divided the number of correctly classified test images by 5000. This produced one held-out accuracy per decoder at each rate.
   4. **Select the interval.**

      The practical floor was the lowest tested rate at which every decoder met or exceeded 50% accuracy,

$$r_{\text{train}} = \min\Big\{r \in \mathcal{R} : \min_{s \in \mathcal{S}} A_s(r) \geq 0.5\Big\}. \quad (5)$$

      Here $\mathcal{R}$ is the set of tested rates, $\mathcal{S}$ is the set of decoder seeds, $A_{s(r)}$ is held-out accuracy for decoder $s$ at rate $r$, and $r_{\text{train}}$ is the selected training floor. The upper edge was the highest tested rate. We did not interpolate between tested rates.

## Results

### Decoder training

Validation accuracy improved across training for all three decoder seeds.

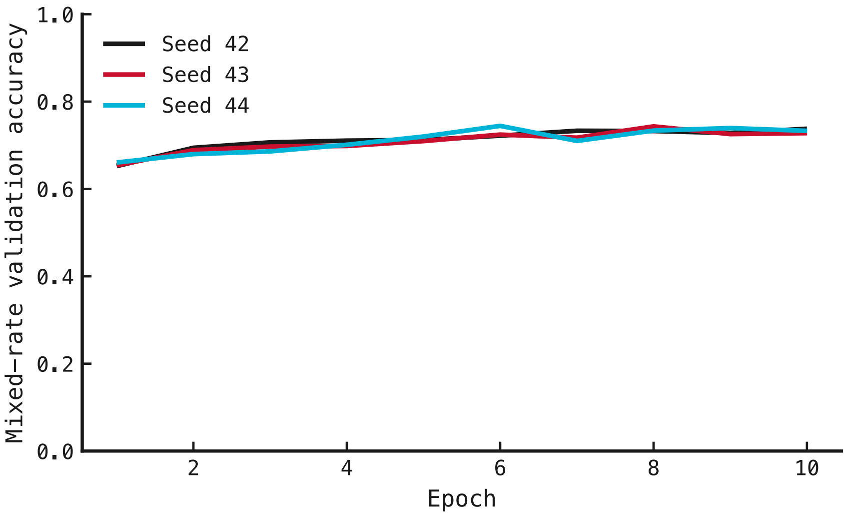


Figure 1: Mixed-rate validation accuracy across training. The horizontal axis is epoch and the vertical axis is validation accuracy. Each curve is one independently trained nonlinear decoder, and every epoch uses fresh direct feature simulations. All three decoders improve before checkpoint selection.

**What the decoder saw**

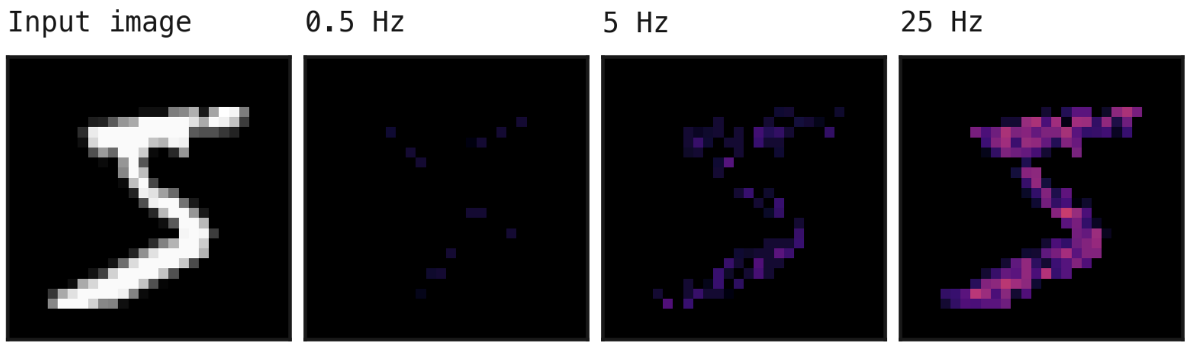


Figure 2: An MNIST input image and its directly simulated filtered features. The left panel shows the normalized input. The remaining panels show independent spike realizations at maximum-pixel encoding rates of 0.5 Hz, 5 Hz, and 25 Hz. Each simulation uses a 1.2 µS synaptic conductance and a 200 ms presentation. Greater input rate preserves more of the digit's spatial structure.

Sparse presentations retained fragments of the digit rather than a uniformly attenuated image. Increasing rate filled in the spatial pattern and reduced the importance of individual event times.

**Empirical rate selection**

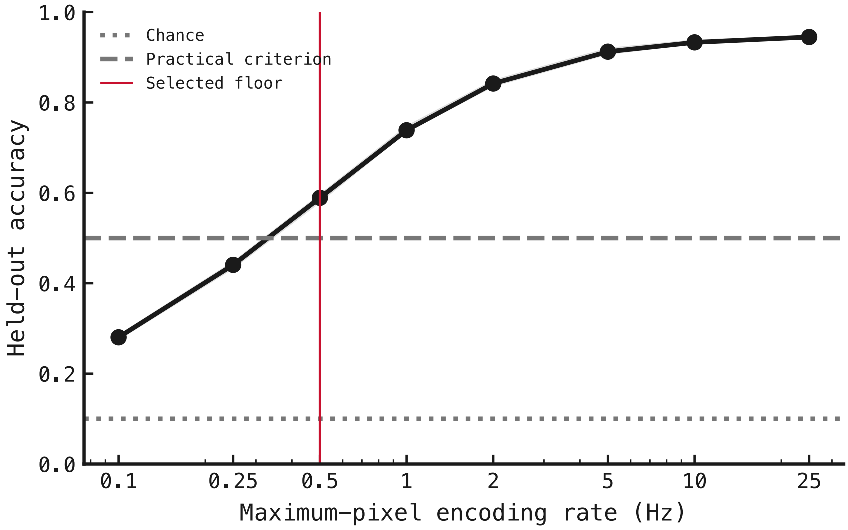


Figure 3: Held-out nonlinear-decoder accuracy against maximum-pixel encoding rate. Points average the official test images and three independently trained decoders. The band spans the lowest and highest decoder accuracy at each rate. Horizontal rules mark chance and the 50% practical criterion. The red vertical rule marks the first tested rate at which all three decoders meet that criterion, which defines the selected floor.

All three decoders first met the practical 50% criterion at 0.5 Hz. Mean accuracy at that condition was 58.9%. We therefore select 0.5 to 25 Hz for later variable-rate PING training.

## Conclusion

The filtered MNIST representation retained usable digit information from 0.5 Hz upward. All three independently trained decoders met the 50% held-out accuracy criterion at the selected lower bound, and performance continued to improve across the tested range. We therefore carry 0.5 to 25 Hz forward as the empirical input-rate interval. This is a decoder-based calibration, not a measurement of PING-network performance, so the interval must still be checked in the network for which it was selected.

## Relation to prior work

Neural decoding measures information accessible to a specified readout, not an absolute information content or a mechanistic account of the encoded population [**2**]. We therefore interpret the ANN psychometric curve only as a decoder-relative calibration. Nonstationary filtered-shot-noise theory motivates direct simulation of the conductance and membrane dynamics [**1**].

## References

1. Brigham & Destexhe: *Nonstationary Filtered Shot-Noise Processes and Applications to Neuronal Membranes*. Physical Review E, 2015. doi:10.1103/PhysRevE.91.062102
2. Quian Quiroga & Panzeri: *Extracting Information from Neuronal Populations: Information Theory and Decoding Approaches*. Nature Reviews Neuroscience, 2009. doi:10.1038/nrn2578