Linear-filter analysis of sparse conductance-driven pixel features

exp081 · 10 August 2026 · Draft

Abstract

We tested what stationary linear-filter theory explains about a sparse conductance-driven pixel feature, and where the approximation fails. We simulated how an AMPA synapse, a conductance-based membrane, and finite-time averaging transform a fully active pixel driven from 0 to 25 Hz. The model explains the flat response to slow input fluctuations, attenuation of fast fluctuations, and lobes caused by temporal averaging.

The model does not accurately predict feature magnitude: it overpredicts the mean, and its standard deviation peaks too early and has the wrong shape. Sparse responses depend strongly on the discrete number and timing of spikes, violating the model's assumption of small, stationary fluctuations. The theory explains the filtering, but direct simulation remains necessary for quantitative predictions.

Methods

1. **Simulate the finite-window feature.**

   We fixed the normalized pixel intensity at $x = 1$ and varied only its input rate r from 0 to 25 spikes/s. At each timestep, the pixel generated an event with probability

   $$p_{\text{event}} = \frac{r\Delta t}{1000}. \quad (1)$$

   Here r is the input rate in spikes/s and $\Delta t = 0.1$ ms. During a presentation of duration $T = 200$ ms, the expected number of events is $r\frac{T}{1000}$. Expected event count is therefore a consequence of the chosen rate, not a separate experimental variable.

   Conductance and membrane voltage followed

   $$g(t) = \beta g(t - \Delta t) + wS(t), \quad (2)$$

   $$\beta = \exp\left(-\Delta \frac{t}{\tau_{\text{AMPA}}}\right), \quad (3)$$

   $$C\frac{dv}{dt} = g_L(E_L - v) + g(t)(E_e - v). \quad (4)$$

   Here $S(t)$ is one when an event occurs and zero otherwise, $g(t)$ is AMPA conductance, w is the conductance increment per event, β is the per-timestep decay factor, and τ_{AMPA} is the AMPA decay time constant. The membrane voltage is $v(t)$, C is membrane capacitance, g_L is leak conductance, E_L is the leak reversal potential, and E_e is the excitatory reversal potential.

We used $C = 1$ nF, $g_L = 0.05$ µS, $E_L = -65$ mV, $E_e = 0$ mV, $\tau_{\text{AMPA}} = 2$ ms, and independent probe conductances $w \in \{0.6, 1.2, 2.4\}$ µS. Every presentation began from $g(0) = 0$ and $v(0) = E_L$. Its scalar feature was

$$z = \frac{1}{T}\int_0^T (v(t) - E_L)\,\mathrm{d}t. \quad (5)$$

Here z is the baseline-subtracted mean voltage and T is presentation duration.

For the mean and SD estimates in Figure 1, we generated 512 new presentations at each of 101 input rates and each conductance. For the full response distributions in Figure 2, we generated 4096 new presentations at each selected rate. Direct simulation retains the nonstationary finite-window behaviour of filtered conductance shot noise [**2**].

2. **Calculate the stationary operating point.**

   For stationary Poisson rate r, filtered-shot-noise theory [**1**] gives the mean AMPA conductance

   $$\bar{g}_r = rw\frac{\tau_{\text{AMPA}}}{1000}. \quad (6)$$

   Replacing the fluctuating conductance by this mean gives the operating-point voltage

   $$\bar{v}_r = \frac{g_L E_L + \bar{g}_r E_e}{g_L + \bar{g}_r}. \quad (7)$$

   Because this operating point is constant across the window, its predicted mean feature is

   $$\mu_{\text{linear}}(z) = \bar{v}_r - E_L. \quad (8)$$

   An overbar denotes a stationary mean, $\bar{g}_r$ is mean conductance at rate r, $\bar{v}_r$ is the corresponding mean voltage, and $\mu_{\text{linear}(z)}$ is the resulting linear prediction for mean feature value. Appendix A.1 derives Equations 6–8 from the continuous conductance and membrane equations.

3. **Derive the local frequency response.**

   Linearizing Equations 2–4 around the operating point gives the response from input-rate fluctuation to voltage fluctuation,

   $$G_r(\omega) = \frac{w}{i\omega + \frac{1}{\tau_{\text{AMPA}}}} \cdot \frac{E_e - \bar{v}_r}{i\omega C + g_L + \bar{g}_r}. \quad (9)$$

   Here ω is angular frequency in rad/ms and i is the imaginary unit. The transfer function $G_{r(\omega)}$ maps a small input-rate fluctuation to its voltage response around the operating point. Appendix A.2 derives Equation 9 by linearizing the synaptic and membrane equations.

   Averaging voltage over the finite presentation contributes

   $$A_T(\omega) = \frac{1 - \exp(-i\omega T)}{i\omega T}, \quad (10)$$

   so the complete input-to-feature response is

   $$H_r(\omega) = A_T(\omega)G_r(\omega). \quad (11)$$

Here $A_{T(\omega)}$ is the transfer function of averaging over duration T, and $H_{r(\omega)}$ is the complete input-to-feature transfer function. Appendix A.3 derives Equations 10 and 11.

To generate Figure 3, we evaluated Equations 9 and 11 from 0.1 to 200 Hz at 0.25, 3, 25 spikes/s for the nominal 1.2 µS probe. Frequency in Hz was converted using $\omega = 2\pi \frac{f}{1000}$. The Bode magnitude was reported relative to the low-drive DC response,

$$M_X(f) = 20 \log_{10} \left(\frac{\left| X_r \left(2\pi \frac{f}{1000} \right) \right|}{\left| X_{r_{\text{low}}}(0) \right|} \right). \quad (12)$$

Here $M_{X(f)}$ is Bode magnitude in dB, $X_r = G_r$ for Figure 3A, $X_r = H_r$ for Figure 3B, f is modulation frequency in Hz, and r_{low} is the lowest operating rate used as the DC reference.

4. **Calculate the linearized feature variance.**

   The centred ideal Poisson input has a white two-sided spectrum on the millisecond time base,

$$S_{\text{in}}(\omega) = \frac{r}{1000}. \quad (13)$$

The feature spectrum and variance are therefore

$$S_z(\omega) = |H_r(\omega)|^2 S_{\text{in}}(\omega), \quad (14)$$

$$\text{Var}_{\text{linear}}(z) = \frac{1}{2\pi} \int_{-\infty}^{\infty} |H_r(\omega)|^2 S_{\text{in}}(\omega) \, \mathrm{d}\omega. \quad (15)$$

Here $S_{\text{in}(\omega)}$ is the two-sided input power spectrum, $S_{z(\omega)}$ is the feature power spectrum, and $\text{Var}_{\text{linear}(z)}$ is the predicted feature variance. Appendix A.4 derives Equations 13–15 from the Poisson spectrum and the linear-filter variance identity.

We integrated Equation 15 numerically on a logarithmic frequency grid. A second calculation with half as many points measured quadrature refinement.

Results

Empirical finite-window response

Mean response increased with input rate and event strength. Variability first increased as spike count and timing diversified, then flattened or declined as relative count fluctuations, shunting, and voltage saturation became more important.

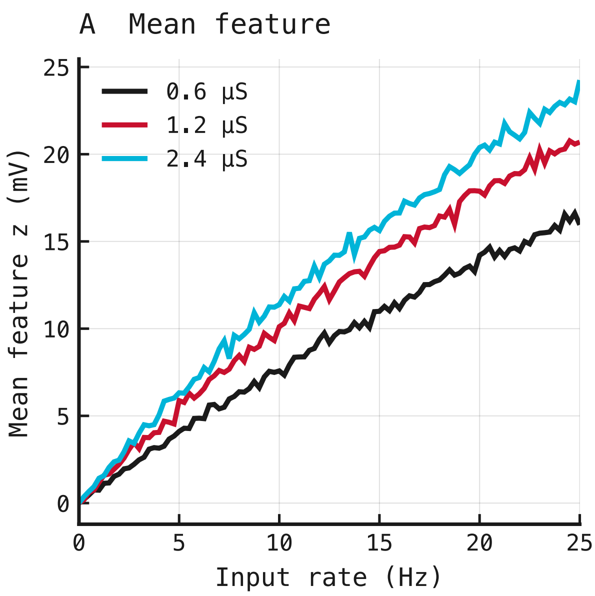

Figure 1: Empirical finite-window response of a fully active pixel over the input-rate grid. Both horizontal axes show input rate in spikes/s. Panel A shows mean feature in mV; Panel B shows feature SD in mV. Each coloured curve contains 101 rate conditions. At every rate, the plotted mean or SD summarizes 512 independently simulated presentations; the figure does not display those individual presentations as points. Black, red, and cyan denote 0.6, 1.2, 2.4 µS conductance increments. Mean response rises with rate, while variability reaches a broad maximum or plateau.

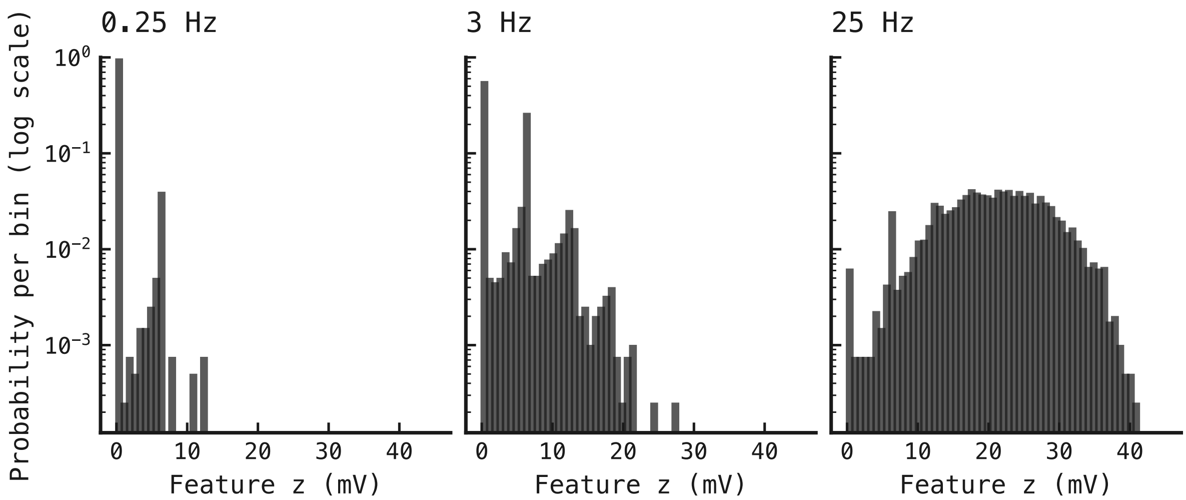

Figure 2: Empirical feature distributions at 0.25, 3, 25 spikes/s for the nominal 1.2 µS probe. Each panel contains 4096 new presentations. The horizontal axes show feature value in mV. Bars report probability per common fixed-width bin on a shared logarithmic vertical axis. Sparse input produces an atom at rest and separated timing-dependent responses; increasing rate produces a smoother distribution.

Input rate fixes only the average number of events. Individual presentations can still contain no events, one event, or several events. Even when two presentations contain the same number, different event times change the finite-window average.

Analytical Bode-magnitude response

The synapse and membrane passed slow modulation but attenuated fast modulation. Finite-window averaging superimposed regularly spaced zeros and lobes on that falling response.

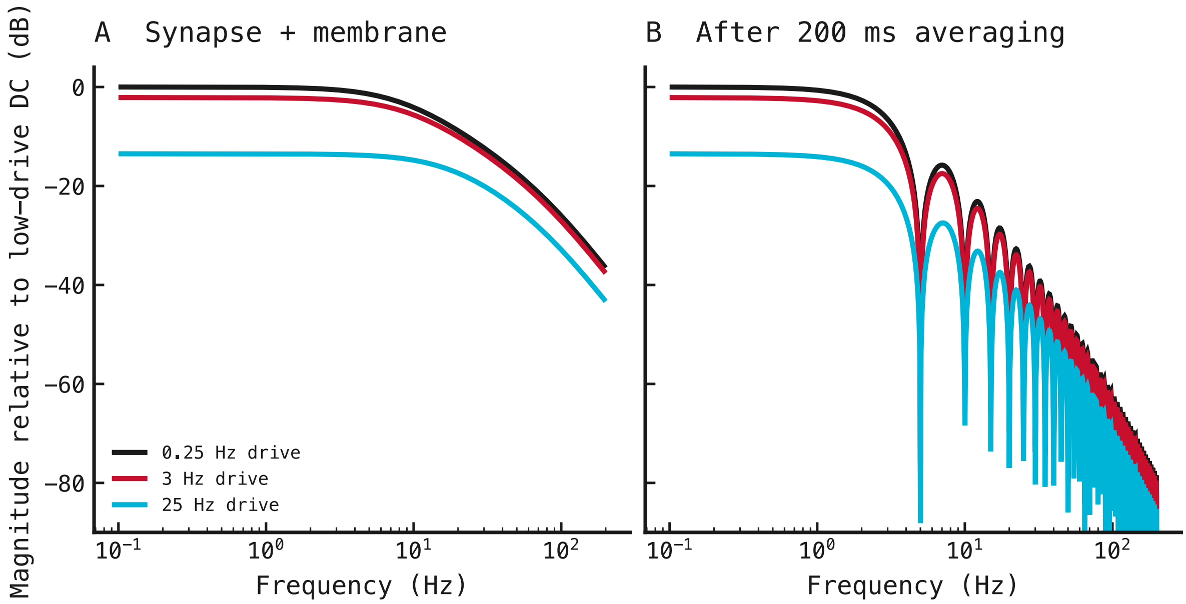

Figure 3: Analytical Bode-magnitude plots at the nominal 1.2 μS probe. Black, red, and cyan denote operating rates 0.25, 3, 25 spikes/s. The horizontal axis is modulation frequency in Hz and the vertical axis is gain relative to the low-drive direct-current response in decibels. Panel A shows Equation 9; Panel B shows the complete 200 ms window-averaged response from Equation 11.

In Panel A, each curve is flat at low frequency because the AMPA conductance and membrane voltage can follow a slowly varying input quasi-statically. Greater drive lowers that plateau by shunting the membrane and reducing the excitatory driving force. The curves bend downward once modulation becomes too fast for the synaptic and membrane timescales to follow.

Panel B is the same synapse-plus-membrane response multiplied by the rectangular-window response. With $T_s = \frac{T}{1000}$ s,

$$\left| A_T \left(2\pi \frac{f}{1000} \right) \right| = \left| \frac{\sin(\pi f T_s)}{\pi f T_s} \right|. \quad (16)$$

Its zeros at $f = \frac{n}{T_s}$ cancel modulation containing an integer number of cycles within the presentation. Here n is any nonzero integer. The absolute sinc response creates the lobes, while the Panel A low-pass response supplies their falling envelope. Appendix A.3 derives Equation 16 from the finite-window averaging kernel.

Analytical and empirical moments

The stationary model reproduced the mean curve's broad shape but not its magnitude, and it failed to reproduce the shape of the empirical standard deviation.

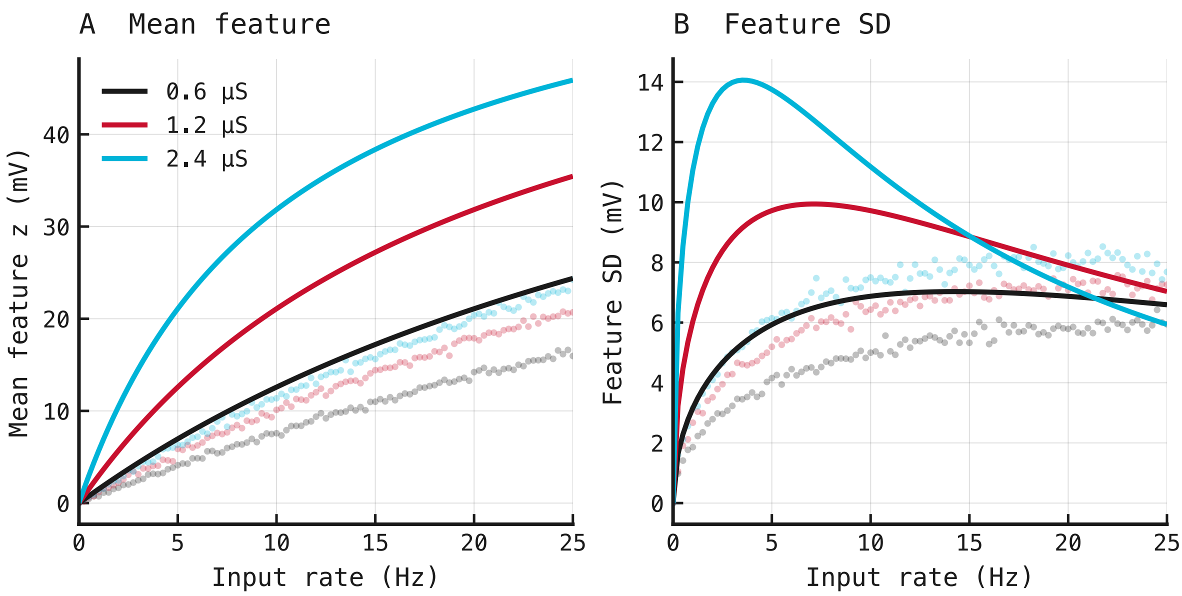

Figure 4: Analytical and empirical moments over input rate for a fully active pixel. The horizontal axis is input rate in spikes/s. Solid curves are stationary predictions from Equations 8 and 15; faint points are independently simulated finite-window estimates. Panel A shows mean feature in mV and Panel B feature SD in mV. The stationary theory overpredicts the mean and places the SD maximum too early.

The analytical mean preserved the broad curvature and conductance ordering but exceeded the empirical magnitude. Its median predicted-to-empirical ratio was 1.972. The stationary calculation evaluates the response at mean conductance, $z(E[g])$, whereas simulation estimates the response over random conductance paths, $E[z(g)]$. Saturation makes these unequal, approximately

$$E[z(g)] < z(E[g]).$$

The stationary input also acts throughout the window; a real late event contributes only briefly to Equation 5.

The SD mismatch is more fundamental. The empirical distribution is a spike-count and spike-time mixture,

$$p_Z(z) = \sum_{n=0}^{\infty} P(N = n) p_{Z|N}(z|n). \quad (17)$$

Here Z is the random feature value, N is the event count, $P(N = n)$ is the probability of observing n events, and $p_{Z|N}(z|n)$ is the feature distribution conditional on that count.

Equation 15 instead treats input as a small stationary fluctuation around one operating point. At low rates, zero-event trials remain exactly at rest, one large event produces a timing-dependent response, and multiple events interact through shunting and saturation. That non-Gaussian mixture explains why the analytical SD rises too sharply and peaks too early, especially for the largest conductance increment.

Conclusion

Stationary linear-filter theory correctly explains the low-frequency plateau, high-frequency roll-off, averaging-window zeros, and qualitative operating- point dependence. It does not quantitatively predict finite-window moments in the sparse, large-jump regime. Direct simulation is therefore required for empirical rate selection, while the analytical model remains useful for explaining the filter's structure.

Limitations

The empirical moment grid used 512 presentations per condition; it was designed to resolve the qualitative theoretical comparison, not to estimate extreme distribution tails. The analytical variance assumes ideal continuous-time Poisson drive and a local stationary linearization, whereas simulation uses discrete Bernoulli events and begins from rest.

Reproducibility

`uv run python experiments/exp081.py` regenerates every sample, summary, and figure. The runner contains a complete independent physical specification and consumes no artifacts or code from another experiment.

Appendix A: derivation of the analytical filter

A.1 Stationary conductance and voltage

Write the continuous AMPA equation driven by an event train $s(t) = \sum_k \delta(t - t_k)$ as

$$\frac{dg}{dt} = -\frac{g}{\tau_{\text{AMPA}}} + ws(t). \quad \text{(A1)}$$

Here $s(t)$ is the impulse train, t_k is the time of event k, and $\delta(t - t_k)$ is a Dirac impulse at that time.

For stationary input rate r spikes/s, the rate on the millisecond time base is $\frac{r}{1000}$. Taking expectations of Equation A1 and setting the stationary derivative to zero gives

$$0 = -\frac{\bar{g}_r}{\tau_{\text{AMPA}}} + w\frac{r}{1000},$$

hence Equation 6. The stationary membrane equation is

$$0 = g_L(E_L - \bar{v}_r) + \bar{g}_r(E_e - \bar{v}_r), \quad \text{(A2)}$$

and collecting the terms multiplying $\bar{v}_r$ gives Equation 7.

A.2 Synapse-plus-membrane linearization

Decompose each quantity into its operating point and a small fluctuation,

$$g = \bar{g}_r + \delta g, \quad v = \bar{v}_r + \delta v, \quad s = \frac{r}{1000} + \delta s.$$

Substitution into Equation A1 and cancellation of the stationary terms gives

$$\frac{d\delta g}{dt} = -\delta\frac{g}{\tau_{\text{AMPA}}} + w\delta s. \quad \text{(A3)}$$

With Fourier convention $\frac{d}{dt} \to i\omega$,

$$\delta g(\omega) = \frac{w}{i\omega + \frac{1}{\tau_{\text{AMPA}}}}\delta s(\omega). \quad \text{(A4)}$$

Substitute the decompositions into Equation 4. After cancelling Equation A2 and discarding the second-order product $\delta g \delta v$,

$$C\frac{d\delta v}{dt} = -(g_L + \bar{g}_r)\delta v + (E_e - \bar{v}_r)\delta g. \quad \text{(A5)}$$

Fourier transformation and substitution of Equation A4 yields Equation 9. Its two denominator factors are the AMPA and membrane low-pass terms; the numerator contains the conductance increment and local excitatory driving force.

A.3 Finite-window average

Linearizing Equation 5 leaves the average of the voltage fluctuation,

$$\delta z = \frac{1}{T}\int_0^T \delta v(t)\,\mathrm{d}t. \quad \text{(A6)}$$

The averaging kernel is $a_T(t) = \frac{1}{T}$ for $0 \leq t \leq T$ and zero otherwise. Its Fourier transform is

$$A_T(\omega) = \frac{1}{T}\int_0^T \exp(-i\omega t)\,\mathrm{d}t = \frac{1 - \exp(-i\omega T)}{i\omega T},$$

proving Equation 10. Because averaging follows the membrane response, the transfer functions multiply, proving Equation 11. Using $1 - \exp(-ix) = 2i\exp\left(-i\frac{x}{2}\right)\sin\left(\frac{x}{2}\right)$ gives

$$|A_T(\omega)| = \left|\frac{\sin\left(\omega\frac{T}{2}\right)}{\omega\frac{T}{2}}\right|,$$

which becomes Equation 16 after substituting $\omega = 2\pi\frac{f}{1000}$ and $T_s = \frac{T}{1000}$.

A.4 Poisson spectrum and feature variance

For ideal Poisson events with rate $\frac{r}{1000}$ per ms, disjoint intervals have independent counts and count variance equals count mean. The centred impulse train therefore has autocovariance

$$R_{\delta s}(\tau) = \left(\frac{r}{1000}\right)\delta(\tau). \quad \text{(A7)}$$

Here $R_{\delta s}(\tau)$ is the input autocovariance at lag τ, and $\delta(\tau)$ is the Dirac delta function.

Fourier transformation of the delta function gives the constant spectrum in Equation 13. A linear time-invariant filter multiplies the input spectrum by squared transfer magnitude, proving Equation 14. Finally, inverse Fourier transformation at zero lag gives

$$\text{Var}(z) = R_z(0) = \frac{1}{2\pi}\int_{-\infty}^{\infty} S_z(\omega)\,\mathrm{d}\omega,$$

which, after substituting Equation 14, proves Equation 15. The result is exact for the stated stationary linearized model but only approximate for the nonlinear, start-from-rest, finite-event simulation.

References

1. Wolff & Lindner: *Mean, Variance, and Autocorrelation of Subthreshold Potential Fluctuations Driven by Filtered Conductance Shot Noise.* Neural Computation, 2010. doi:10.1162/neco.2009.02-09-958
2. Brigham & Destexhe: *Nonstationary Filtered Shot-Noise Processes and Applications to Neuronal Membranes.* Physical Review E, 2015. doi:10.1103/PhysRevE.91.062102