



Biases in white noise analysis due to non-Poisson spike generation

Jonathan W. Pillow^{a,*}, Eero P. Simoncelli^b

^aCenter for Neural Science, New York University, 4 Washington Pl 809, New York 10003, USA

^bHoward Hughes Medical Center, and Center for Neural Science, New York University,
New York, USA

Abstract

White noise analysis methods for characterizing neurons typically ignore the dynamics of neural spike generation, assuming that spikes arise from an inhomogeneous Poisson process. We show that when spikes arise from a leaky integrate-and-fire mechanism, a classical white noise estimate of a neuron's temporal receptive field is significantly biased. We develop a modified estimator for linear characterization of such neurons, and demonstrate its effectiveness in simulation. Finally, we apply it to physiological data and show that spiking dynamics may account for changes observed in the receptive fields measured at different contrasts.

© 2003 Elsevier Science B.V. All rights reserved.

Keywords: White noise analysis; Spike generation

White noise analysis has become a widely used technique for characterizing response properties of spiking neurons in sensory systems. A sequence of stimuli are drawn randomly from an ensemble and presented in rapid succession, and one examines the stimuli that elicit action potentials. In the most widely used form of this analysis, one estimates a linear approximation to the receptive field (i.e. first-order Wiener kernel) by computing the spike-triggered average (STA); that is, the average stimulus preceding a spike (e.g. Refs. [6,8]). Under the assumption that spikes are generated by a Poisson process with instantaneous rate determined by linear projection onto a kernel followed by a static nonlinearity, the STA provides an unbiased estimate of the underlying kernel [5].

* Corresponding author.

E-mail addresses: pillow@cns.nyu.edu (J.W. Pillow), eero.simoncelli@nyu.edu (E.P. Simoncelli).

The white noise approach is considered to have several advantages over traditional characterization approaches, including the ability to explore a large portion of the input space, and receptive field estimation that is robust to drift or fluctuation in the responsiveness of a neuron. Despite these advantages, it has also become clear that there are drawbacks to the standard characterizations obtained with white noise methods. One such shortcoming is the well-known phenomenon that the shape of the STA varies with the amplitude (e.g. contrast) of the white noise stimuli. (e.g. Refs. [4,9,13]). This type of change cannot be explained by a linear model followed by a static nonlinearity and Poisson spike generation (the L–N–P model), since it implies a change in the linear front end. We have previously shown that nonlinear suppressive interactions such as those found in cortical neurons can explain biases in the STA, that a spike-triggered covariance analysis can be used to characterize these suppressive interactions, and that the resulting corrected model can account for the changes of STA with contrast [10].

Here, we explore another potential source of failure in white noise characterization: the assumption of Poisson spike generation. The significance of temporal dynamic (i.e. non-Poisson) properties of biological spike generation for white noise characterization of neurons has not been thoroughly analyzed¹. However, we show that in simulated white noise experiments, a linear model which drives an integrate-and-fire spiking mechanism is inaccurately characterized by the STA. Furthermore, we show that the integrative behavior of this model can account for some of the changes in the STA estimated at different stimulus amplitudes in real neurons. Finally, we propose a new method for recovering the linear temporal filter governing neural response. We demonstrate through simulation that this approach can correctly estimate the linear kernel of a model neuron, and we also apply our method to real neural data, demonstrating that the recovered linear kernel is fairly stable with changes in stimulus contrast. We thus conclude that the recovered linear kernel may provide a more fundamental functional description of neural behavior, and might well be more directly related to the mechanisms underlying neural response.

1. Leaky integrate-and-fire model

Our analysis is based on a leaky integrate-and-fire (LIF) model. The input is convolved with a linear filter K , and this response drives a leaky integrator. When the level of this integrator reaches a threshold value, the neuron fires a spike and the integrator is reset to zero. The time evolution of the model membrane potential $V(t)$ is characterized by a single differential equation

$$\frac{dV}{dt} = -\frac{1}{\tau} V(t) + I(t), \quad (1)$$

¹ However, Arcas et al. have recently examined the behavior of a Hodgkin–Huxley model under white noise stimulation, and have made a number of interesting observations regarding the spike-triggered stimulus ensemble [2].

where τ is the time constant governing decay of the membrane potential, and $I(t)$ is the input current, generated by convolving the input signal $S(t)$ with the fixed kernel K :

$$I(t) = K * S(t) = \int_{-\infty}^0 K(u)S(t - u) du. \tag{2}$$

This model has an analytical solution relative to the time of the most recent spike:

$$V(t) = \int_{t^-}^t I(u)e^{(u-t)/\tau} du, \tag{3}$$

where t^- is the time of occurrence of the last spike before t . This dependence on the time of the previous spike (and past input to the integrator) represents a fundamental departure from L–N–P model described earlier, where the probability of firing a spike is an instantaneous function of the projection of the stimulus onto K .

2. Simulation results and comparison

We simulated a white noise analysis experiment with the model described above. In our simulations, the kernel K was chosen to be a 32-sample function whose shape loosely resembles temporal kernels measured in retinal ganglion cells. As in classical white noise experiments, we generated a random discrete stimulus $S(t)$ that was temporally white, drawing the stimulus intensity as an independent Gaussian random variable in each time step. We computed the STA as the average stimulus in the 32 time bins preceding each spike.

Fig. 1 (left) shows a plot of the actual kernel K superimposed on the STA for three different values of the membrane time constant τ . First, note that in all three cases, the

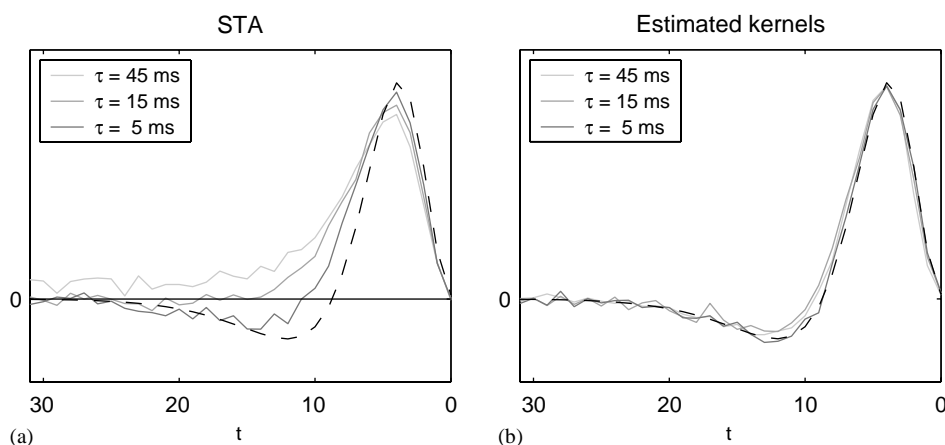


Fig. 1. Simulation of integrate-and-fire neuron. Left: STA kernels retrieved for three different contrast levels (solid lines), plotted along with the true model kernel (dashed curve). Right: Kernels recovered using our algorithm.

STA differs significantly from K . This bias reflects the integrative spiking mechanism of the LIF model, as the STA is quite close to K if the same input were given to an L–N–P model [5]. Furthermore, the discrepancy between K and the STA depends on τ . For small τ (i.e. rapid decay of V), the STA more closely resembles K , whereas larger τ (slower decay) gives rise to an STA which is smoother and more biased away from the true K . Note that although this basic effect is unsurprising, it is *not* the case that the STA shape arises simply from a low-pass filtering of K with an exponential filter. Specifically, the STA measured for a stand-alone LIF spike generator is decidedly non-exponential.

Physiological evidence indicates that at higher firing rates, the membrane conductance of neurons increases, which corresponds to a decrease in membrane time constant τ [1,3,7]. Moreover, STAs measured in real neurons at high contrast tend to be narrower than those measured at low contrasts. This suggests that an integrative spiking mechanism with time constant that depends on firing rate is at least consistent with contrast-dependent changes in the STA of real neurons.

3. Recovering the linear kernel

Assuming that the input to an integrate-and-fire spiking model is determined by projection onto a linear kernel, how can the kernel be recovered from the response to white noise stimuli? Eq. (3) provides a deterministic expression for the voltage at any time since the most recent spike. The voltage at any spike time is therefore given by

$$V(t^+) = V_{\text{th}} = \int_{t^-}^{t^+} [K * S(t)] e^{(t-t^+)/\tau} dt, \quad (4)$$

where V_{th} is threshold, t^- is the time index of the previous spike and t^+ that of the current spike. Using Eq. (2), we can rewrite this (by switching the order of integration):

$$V_{\text{th}} = \int_{-\infty}^0 K(u) \left[\int_{t^-}^{t^+} S(t-u) e^{(t-t^+)/\tau} dt \right] du. \quad (5)$$

Note that, for fixed τ , this equation provides a linear constraint on K , since it expresses V_{th} as the inner product of K with the exponentially weighted S (back to the time of the previous spike). Every spike in the spike train provides one such constraint, so a discretized K can be overconstrained so long as its dimensionality is smaller than the number of spikes collected. K can easily be estimated by finding the least-squares solution to this overconstrained linear system.

In practice, one would like to estimate both τ and K simultaneously, since both are unknown for data collected in real neurons. This can be achieved simply using a nested optimization (a line search algorithm) to find the τ which minimizes the squared error in the least-squares solution for K . This algorithm is guaranteed to converge, and although the solution may be only a local minimum, in simulations it was well-behaved for a wide variety of kernel shapes and a large range of τ values. Fig. 1 (right) shows the kernels estimated for simulations conducted with three different values of τ . Close




Fig. 2. Analysis of in vitro ganglion cell data in monkey retina. Left: STA estimates based on responses recorded at two different input contrast levels. Right: Kernels recovered using our procedure. The associated time constant estimates are 19.1 and 6.5 ms.

estimates of the true values of τ were also obtained. For both graphs in this figure, the stimulus contained 40,000 time samples and approximately 2000 spikes were collected for each τ .

It should be noted, finally, that this estimator for K and τ ignores a huge set of additional constraints—namely, that $V(t)$ be less than threshold at all times within the interspike interval. However, because the problem is already overconstrained by the constraint on $V(t)$ at spike times, and because the additional constraints are much harder to implement, they can be ignored. A significant improvement to the estimator may nevertheless be obtained by considering additional constraints only on the time steps immediately preceding a spike. (This can be implemented by allowing a contribution to the squared error for any pre-spike time bin where V exceeds threshold). Monte Carlo simulations exhibit rapid convergence to the true values of K and τ for this revised estimator.

4. Recovering a kernel from neural data

Our procedure for linear kernel estimation is based on an overly simplistic integrate-and-fire model for neural spike generation. We thus cannot be sure it will be applicable to real neural data. But we note that STA techniques have been used for decades to estimate linear kernels under the assumption of a Poisson spike generator. The integrate-and-fire model incorporates a dependence on the time of the previous spike and is likely to provide a more accurate description of spiking in real neurons.

We have applied our procedure directly to data drawn from a monkey retinal ganglion cell [4]. The data were recorded in vitro, using a stimulus consisting of 80,000 time samples of full-field 120 Hz flickering binary white noise. The stimulus vectors \vec{s} of this sequence are defined over a 25-segment (0.21 s) time window. Two data sets were recorded, at contrasts of 32% and 64%.

Fig. 2 (left) shows STA estimates for both contrast levels. The kernels are quite different; the low-contrast STA is smoother and its peak is shifted earlier in time than the high-contrast STA. Fig. 2 (right) shows the kernels resulting from our estimation procedure. Note that the estimated kernel is now quite stable across different contrasts,

a desirable property for a functional description of neural behavior. The recovered time constants of 19.1 and 6.5 ms are within ranges considered biologically plausible, although their ratio indicates a greater change with amplitude than is commonly reported for cortical neurons (e.g. Refs. [1,3,7]).

5. Discussion

Our results show that spike generation mechanisms can affect the interpretation of results obtained with white noise analysis. In particular, we have shown that even for a simple integrate-and-fire model, the temporal STA does not accurately recover the temporal linear input kernel. For this model, the magnitude of bias in the STA is influenced by the membrane conductance, which is believed to vary with stimulus strength. This amplitude-dependence of the STA mirrors changes in the STA of real neurons measured at different contrasts, and cannot be captured by an L–N–P model.

Based on this simple LIF model, we have developed a new method for the recovery of the linear kernel and integration time constant from responses to white noise stimuli. To our surprise, this kernel estimation procedure recovers a stable linear kernel when applied to data recorded from monkey retinal ganglion cells, and the associated estimates of membrane conductance are within a biologically plausible range. Finally, while not discussed here, our technique also appears to be quite robust to the presence of noise in the membrane potential.

We are currently exploring the generalization of these results to more realistic models. In particular, we have found that the incorporation of a voltage floor in the model (corresponding to an ionic reversal potential) produces an STA which is sharper and closer to the true input kernel at high contrast, independent of any changes in membrane conductance. The significance of this phenomenon, along with that of other nonlinearities associated with spike generation, remains to be analyzed.

Our results suggest a mechanistic explanation of the behaviors captured by current functional models of retinal ganglion cells (e.g. Ref. [12]), in which a nonlinear feedback signal is used to adjust the gain of the neuron. We have also previously shown that nonlinear gain control operations might account for a variety of apparent changes in receptive field properties at different contrast levels [11]. The results presented in this paper suggest that some such changes might be due to intracellular mechanisms of spike generation. It would be interesting to test such hypotheses against intracellular measurements.

Acknowledgements

We thank E.J. Chichilnisky and Divya Chander for providing electrophysiological data, and Liam Paninski and Michael Shadlen for helpful discussions. This research was supported by the Sloan–Swartz Theoretical Visual Neuroscience Center at NYU, and by the Howard Hughes Medical Institute.

References

- [1] J.S. Anderson, M. Carandini, D. Ferster, Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex, *J. Neurophysiol.* 84 (2000) 909–926.
- [2] B.A.Y. Arcas, A.L. Fairhall, W. Bialek, What can a single neuron compute? in: T. Leen, T. Dietterich, V. Tresp (Eds.), *Advances in Neural Information Processing Systems*, Vol. 13, 2000, MIT Press, Cambridge, MA, pp. 75–81.
- [3] L. Borg-Graham, C. Monier, Y. Fregnac, Visual input evokes transient and strong shunting inhibition in visual cortical neurons, *Nature* 393 (1998) 369–373.
- [4] D. Chander, E.J. Chichilnisky, Adaptation to temporal contrast in primate and salamander retina, *J. Neurosci.* 21 (24) (2001) 9904–9916.
- [5] E.J. Chichilnisky, A simple white noise analysis of neuronal light responses, *Network: Comput. Neural Syst.* 12 (2) (2001) 199–213.
- [6] E. deBoer, P. Kuyper, Triggered correlation, *IEEE Trans. Biomed. Eng.* 15 (1968) 169–179.
- [7] J. Hirsch, J.M. Alonso, C.R. Reid, L. Martinez, Synaptic integration in striate cortical simple cells, *J. Neurosci.* 15 (1998) 9517–9528.
- [8] J.P. Jones, L.A. Palmer, The two-dimensional spatial structure of simple receptive fields in the cat striate cortex, *J. Neurophysiol.* 58 (1987) 1187–11211.
- [9] K.J. Kim, F. Rieke, Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells, *J. Neurosci.* 21 (1) (2001) 287–299.
- [10] O. Schwartz, E.J. Chichilnisky, E.P. Simoncelli, Characterizing neural gain control using spike-triggered covariance, in: T.G. Dietterich, S. Becker, Z. Ghahramani (Eds.), *Adv. Neural Information Processing Systems*, Vol. 14, MIT Press, Cambridge, MA, 2002, pp. 269–276.
- [11] O. Schwartz, E.P. Simoncelli, Natural signal statistics and sensory gain control, *Nat. Neurosci.* 4 (8) (2001) 819–825.
- [12] R.M. Shapley, J.D. Victor, How the contrast gain control modifies the frequency response of cat retinal ganglion cells, *J. Physiol. (London)* 318 (1981) 161–171.
- [13] S.M. Smirnakis, M.J. Berry, D.K. Warland, W. Bialek, M. Meister, Adaptation of retinal processing to image contrast and spatial scale, *Nature* 386 (1997) 69–73.