

Comparing integrate-and-fire models estimated using intracellular and extracellular data¹

Liam Paninski^{a,b,2} Jonathan Pillow^b Eero Simoncelli^b

^a *Gatsby Computational Neuroscience Unit, University College London*

^b *Howard Hughes Medical Institute, Center for Neural Science, NYU*
{pillow, eero}@cns.nyu.edu

Abstract

We have recently developed a maximum-likelihood (ML) method for estimating integrate-and-fire-based stimulus encoding models that can be used even when only extracellular spike train data is available. Here we derive the ML estimator given the full intracellular voltage trace and apply both the extracellular-only and intracellular method to responses recorded *in vitro*, allowing a direct comparison of the model fits within a unified statistical framework. Both models are able to capture the behavior of these cells under dynamic stimulus conditions to a high degree of temporal precision, although we observe significant differences in the stochastic behavior of the two models.

Key words: Integrate-and-fire, stimulus-response encoding, noise

A central issue in systems neuroscience is the characterization of the functional relationship between sensory stimuli and neural spike trains. A common model for this relationship consists of linear filtering of the stimulus, followed by a nonlinear, probabilistic spike generation process. The linear filter is typically interpreted as the neuron's "receptive field," while the spiking mechanism accounts for simple nonlinearities like rectification and response saturation. Given a set of stimuli and observed spike times, the characterization problem consists of estimating both the linear filter and the parameters governing the spiking mechanism.

We have recently developed methods for estimating a model of this type (given only extracellular spike train data), in which the nonlinear spike generation step is implemented by a noisy integrate-and-fire mechanism; see e.g. (2-4) for detailed

¹ JP was partially supported by an NYU Dean's Dissertation fellowship. The data analyzed here were collected jointly with B. Lau in the lab of A. Reyes. We thank T. Lewis for an advance copy of (1).

² Contact author: liam@gatsby.ucl.ac.uk; <http://www.gatsby.ucl.ac.uk/~liam>.

arguments on why models of this type are of interest. On the other hand, the problem of estimating this class of model given the full intracellular voltage trace can be solved more easily and has been previously studied (1, 5). Here we examine the performance of the two techniques on *in vitro* data (6), allowing a detailed comparison of the methods. Our results serve both to validate the extracellular method and, more importantly, to illuminate the statistical differences in the models.

Methods

Experimental procedures. See (6) for physiological details of the *in vitro* recordings analyzed here. Briefly, dual-electrode whole-cell recordings were made from pyramidal cells from layers III and V in slices from sensorimotor cortex of rats aged P14-24; Gaussian noise current stimuli were delivered through one electrode, while voltage responses were recorded through the other electrode. The current was not “frozen,” that is, a new Gaussian current was drawn independently for each trial.

Intracellular model. We use a slightly different estimation method than either (5) or (1), which is worth describing here. One major difference is that we pose the problem in an explicitly probabilistic setting, allowing us to systematically examine the noise properties of the model (i.e., the variability of response to a fixed stimulus).

We develop the maximum likelihood estimator (MLE) for the following integrate-and-fire-based model, a generalized form of the “spike-response” model of (2). In discrete time, the voltage V evolves according to

$$V(1) = V(0) + dt \left(- \sum_{i=-m}^0 g_i(-t_0) V(i) + \sum_l a_l f_l(V(0)) + \vec{k} \cdot \vec{x}(0) + \sum_{j=-\infty}^0 h(-t_j) \right) + \sigma \sqrt{dt} N.$$

The tap weights g_i implement an I.I.R. filter whose coefficients can vary as a function of the time t_0 since the last spike (if $m = 0$, $g_0 = g$ is simply the membrane conductance); the functions f_l allow us to model nonlinear subthreshold dynamics; \vec{k} summarizes the linear filtering properties of the cell given the input \vec{x} (here just the scalar stimulus current). On short time scales, the current h induces the spike shape; on longer time scales, h can model burstiness, refractoriness, or adaptation effects, depending on its form and magnitude (3). The cell spikes whenever V reaches a threshold potential V_{th} . Finally, we take the noise term N to be Gaussian and i.i.d.; this means that the MLE for the model parameters has the usual mean-square formulation, which we now describe.

We assume the spike times $\{t_j\}$ have been detected (typically via automatic thresholding, assuming sufficiently well-isolated recordings). Then the model parameters $\{g_i, a_l, \vec{k}, h\}$ act linearly on $\{V(i), f_l(V(0)), \vec{x}(0), \sum \delta(t_j)\}$ in the model equation to produce the voltage differences $(V(1) - V(0)) / dt$ corrupted by additive Gaussian noise. The ML estimate for these parameters is therefore given by least-squares linear regression, and the noise scale σ is then simply estimated as the root-mean-square of the residuals.

This determines all model parameters but V_{th} . The likelihood for this parameter is defined as the probability that the Gaussian noise N will push the voltage V over

the threshold V_{th} at spike times and keep the voltage below threshold at non-spike times. This loglikelihood can be written

$$L(V_{th}) \equiv \sum_j \log \left(\int_{V_{th}}^{\infty} G_{\mu_j, \sigma^2} dV \right) + \sum_{j'} \log \left(\int_{-\infty}^{V_{th}} G_{\mu_{j'}, \sigma^2} dV \right),$$

where $G_{\mu, \sigma^2}(V)$ is the Gaussian density of mean μ and variance σ^2 , and μ_j denotes the mean of V at the spike times $\{t_j\}$ under the spike-response model, $\mu_j \equiv$

$$V(s) + dt \left(- \sum_{i=-m}^0 g_i(s - t_{j-1}) V(s + i) + \sum_l a_l f_l(V(s)) + \vec{k} \cdot \vec{x}(s) + \sum_{j_*=-\infty}^{j-1} h(s - t_{j_*}) \right),$$

with $s \equiv t_j - 1$; $\mu_{j'}$ is the mean computed at all other times (at which no spike was observed). One can show that $L(V_{th})$ is concave in V_{th} (7), and hence easy to maximize using ascent algorithms (note also that the derivatives $L^{(i)}$ are easily computed once one has computed L). Note that this approach requires just a single “pass” through the data (we obtain the spike times, then fit the parameters $\{g_i, a_l, \vec{k}, h\}$ by linear least-squares, then maximize $L(V_{th})$ to fit V_{th}); we do not need to generate a new voltage trace for each setting of the model parameters, to numerically optimize the mean-square difference between the model and true voltages (this latter optimization problem is non-concave in general and is therefore more challenging, due to the possible presence of local minima, than the solution of $L(V_{th})$), and thus this likelihood-based approach is faster and somewhat more easily interpretable than the method described in (1).

We have just described the MLE given both the voltage $V(t)$ and the spike times $\{t_j\}$. It is also possible — indeed, somewhat more straightforward — to write down the MLE given only the voltage $V(t)$ (in which case the solution, including V_{th} , is completely given by least-squares). However, our experiments with the latter approach (details not shown) indicate that this method is less robust; the MLE given only $V(t)$ tends to set V_{th} far too high, basically because missing the first part of the spike (due to improperly high V_{th}) costs less in mean square error than does adding a few false threshold crossings (as is inevitable with V_{th} near the true optimal value, if the cell’s threshold is not perfectly sharp).

Extracellular model. The model for extracellular data introduced in (3) is a simplified version of the above spike-response model:

$$V(1) = V(0) + dt \left(-gV(0) + I_{DC} + \vec{k} \cdot \vec{x}(0) + \sum_{j=-\infty}^0 h(-t_j) \right) + \sigma \sqrt{dt} N.$$

We lack the space necessary to describe the MLE for this model, given only the spike train $\{t_j\}$; see (3) for full details. For fair comparison, this is the version of the intracellular model we examine below; that is, we set $m = 0$, $g(t) \equiv g$, and $f_l = 1$, so the extracellular and intracellular models have the same parameter spaces.

Note that, despite the fact that the full $V(t)$ provides a much richer data set than just the spike times $\{t_j\}$, the extracellular model is in some senses easier to estimate,




Fig. 1. Responses to white noise current input. **Top:** Comparison of sample from fit intracellular model to true observed voltage. Horizontal line gives estimated threshold voltage. **Middle:** Comparison of spike rasters. Asterisks indicate actual cell spike times. Rasters for trials 1-49 are example spike trains from the intracellular model, with noise level decreasing linearly from the fit value σ_{intra} down to zero. Rasters for trials 51-100 are example spike trains from the extracellular model, with noise level increasing linearly from zero up to the fit value σ_{extra} . Note that noisy and deterministic rasters in the intracellularly-fit case match fairly well, while in the extracellular case estimated noise has a much greater effect. **Bottom:** Deterministic and random sample from extracellular model; note that V is dimensionless here, with arbitrarily chosen units. Noise induces an increased number of threshold crossings (i.e., higher firing rate); c.f. especially bottom traces in middle panel.

because we don't have to worry about matching V precisely, just the spike times. This means that we only have to deal with voltage up to a scale and offset that we don't need to estimate (in other words, V is dimensionless in the extracellular model). More importantly, if the threshold doesn't depend linearly on just V , but also, say, $\frac{\partial V}{\partial t}$ (2), the intracellular model will do poorly (because this model has to simultaneously fit the subthreshold voltage trace and the spike times), whereas the extracellular model will perform well (since this model doesn't have to fit V , and $\frac{\partial V}{\partial t}$ is a linear functional of the $V(t)$ signal). A possible improvement for the intracellular model, not pursued here, would be to generalize the simple voltage threshold condition to allow linear combinations of $\vec{V} \equiv \{V(i)\}_{m \leq i \leq 0}$, for example allowing a spike to occur when $\vec{V} \cdot \vec{V}_{th} > 1$ for some vector \vec{V}_{th} . This more powerful model remains easily solvable by the global convexity of the corresponding loglikelihood

$L(\vec{V}_{th})$ (as can be shown, again, by the results in (7)).

Results

Our findings are illustrated for a single cell in Fig. 1 and summarized in Fig. 2. Both models predict the responses to dynamic input fairly well. A large proportion of spikes are predicted with millisecond accuracy; “mistaken” spikes — that is, spikes which occur in the model but not in the data — typically appear at plausible times (i.e., near peaks in the true voltage). The intracellular model also captures the subthreshold responses well (as previously emphasized in (1, 5)).

In particular, the extracellular model estimation procedure performs almost as well as the intracellular procedure, even though the latter is given a much more complete representation of the subthreshold voltage trace (Fig. 2). The one exception to this comes when we consider the estimate of the noise value σ : somewhat unexpectedly, we find that the maximum likelihood estimate for the extracellular model actually adds noise by increasing σ_{extra} beyond the value obtained by the intracellular method. This noise increase is formally identical to a shrinkage of $\{\vec{k}, I_{DC}, h\}$ and is a kind of model selection phenomenon (8, 9): roughly stated, for robust fitting of this model, it is better to be uncertain than certain and wrong. In this case, model error — the inability of the extracellular model to capture every spike time — is represented by the ML estimate as increased uncertainty, in the form of the noise parameter σ_{extra} . This has the effect of increasing the extracellular model’s simulated firing rate and decreasing its selectivity (Fig. 1), even though the noiseless predictions of the extracellular model are comparable in accuracy to those of the intracellular model (Fig. 1, center rasters, and Fig. 2). This decreased selectivity, in turn, reduces the accuracy of the extracellular model’s predictions in the nonzero noise case (Fig. 2). Of course, to predict the spike train given the stimulus, it makes sense to restrict our attention to the most likely spike train given the parameters and stimulus, which (since we assumed N to be Gaussian, with a mode at 0) corresponds exactly to the (much more accurate) $\sigma = 0$ case.

We are currently examining the performance of the model in predicting sensory and motor responses in intact preparations. We close by mentioning one additional interesting direction for future work: we found that the model fits are stable over different input variance (i.e., firing rate) regimes (data not shown), but that the estimated g consistently increases with firing rate (see (10) for another example of this). To model, e.g., variance (“contrast”) adaptation (6) accurately, it may be necessary to capture this effect. One way to do this is to allow g to vary with time, as mentioned above and discussed in more detail in (1, 5); since $g(t)$ is typically found to be largest just after a spike, increasing the spike rate increases the average g . It turns out that the results in (3) on the estimation of the extracellular model may be extended to this time-varying g case; future work will examine the relevance of this time-varying conductance more systematically.




Fig. 2. Summary of fit results. 23 experiments (25-100 sec each) in ten different cells examined; variance of input current varied systematically to explore different firing rate regimes. Only experiments in which ≥ 300 spikes were elicited are examined here; experiments with fewer spikes tended to lead to less accurate estimates. **Left:** Example cross-correlation density (standard cross-correlation normalized by width of time bin used to discretize spike train, .125 ms here) between real spike train and sample spike trains from estimated models. Note that the noiseless extracellular and intracellular fits are about equally accurate, and that the MLE overestimates the noise in the extracellular case, leading to a greatly reduced accuracy. **Middle:** Cumulative distribution of observed cross-correlations (computed as integral of cross-correlation density between $+/- 2$ ms), showing similar results across the population. **Right:** Scatterplot of cross-correlations for extracellular and intracellular fits. Identity line plotted for comparison.

References

1. R. Jolivet, T. Lewis, W. Gerstner, *Journal of Neurophysiology* **92**, 959 (2004).
2. W. Gerstner, W. Kistler, *Spiking Neuron Models: Single Neurons, Populations, Plasticity* (Cambridge University Press, 2002).
3. L. Paninski, J. Pillow, E. Simoncelli, *Neural Computation* **In press** (2004).
4. J. Keat, P. Reinagel, R. Reid, M. Meister, *Neuron* **30**, 803 (2001).
5. C. Stevens, A. Zador, *Proc. 5th joint symp. neural computation, UCSD* (1998).
6. L. Paninski, B. Lau, A. Reyes, *Neurocomputing* **52**, 877 (2003).
7. Y. Rinott, *Annals of Probability* **4**, 1020 (1976).
8. W. James, C. Stein, *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability* **1**, 361 (1960).
9. M. Sahani, J. Linden, *NIPS* **15** (2003).
10. J. Pillow, E. Simoncelli, *Neurocomputing* **52**, 109 (2003).