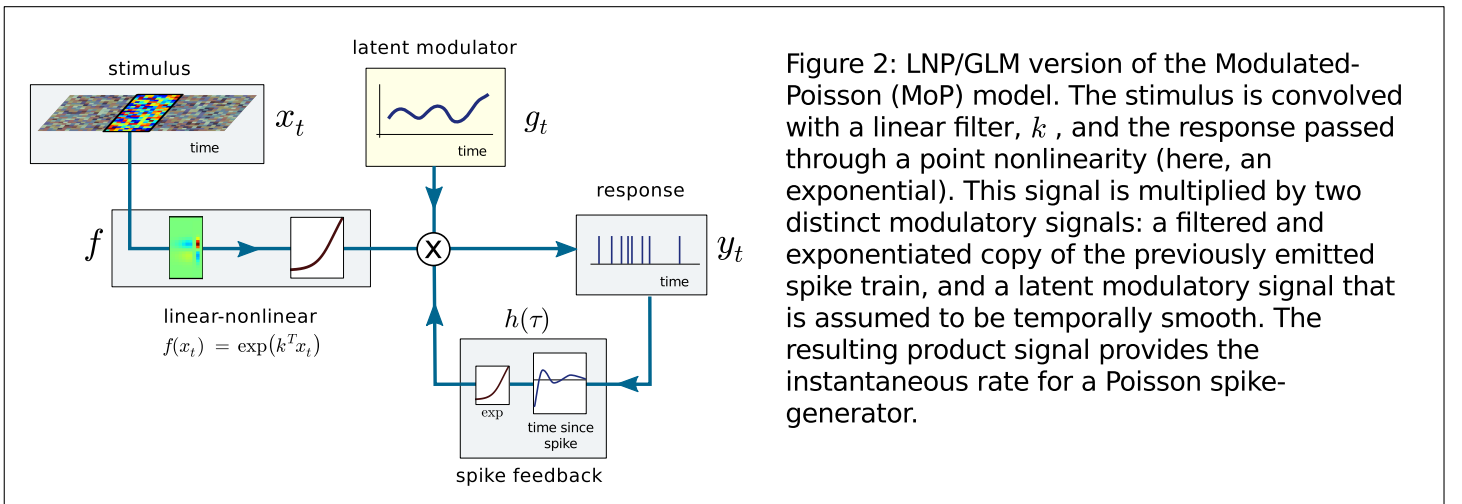
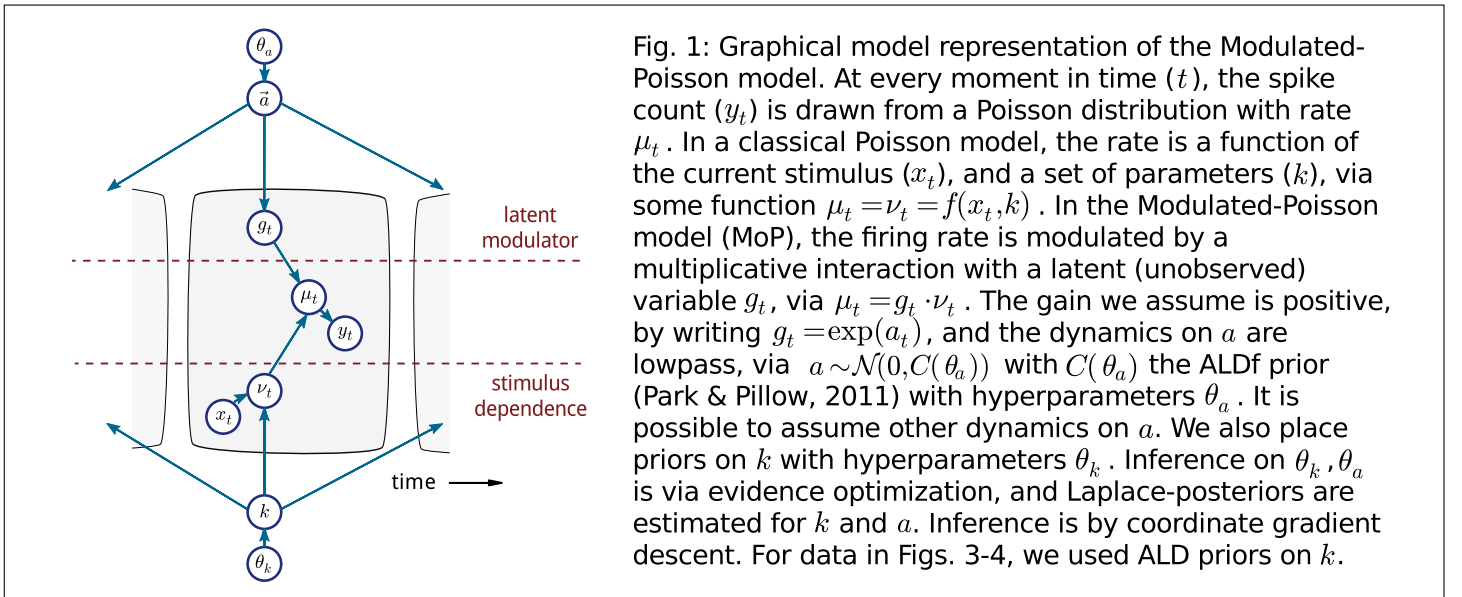


Modeling neural responses in the presence of unknown modulatory inputs

Neil C. Rabinowitz, Robbe L. Goris, Johannes Ballé, & Eero P. Simoncelli

Neurons transmit information with spike trains that differ across repeated measurements. The origin of this variability is unknown, but it is common to describe spike count distributions as Poisson, despite the fact that their variance generally exceeds that expected of a Poisson process. This is likely because neurons' firing rates are also at the mercy of numerous uncontrolled and/or unobserved modulatory factors that alter their gain, including the influence of recently emitted spikes, locally-generated gain control, top-down signals (e.g. attention, arousal, motivation), and physiological conditions (e.g. metabolic resource availability). Regardless of their origin, fluctuations in these signals can confound or bias the inferences that one derives from spiking responses.

These effects can be captured by a modulated Poisson model, whose rate is the product of a stimulus-driven response function and an unknown modulatory signal (Goris, Movshon, Simoncelli, 2013). Here, we extend this model, by including modulatory elements that are known (specifically, spike-history dependence, as in previous GLM models, Pillow et al, 2008), and by constraining the remaining latent modulatory signals to be smooth in time. We fit the entire model, including hyperparameters, via evidence optimization (Park & Pillow, 2011), to the responses of ferret auditory midbrain and cortical neurons to complex sounds. Integrating out the latent modulators yields more readily-interpretable receptive field estimates than a standard Poisson model. Conversely, integrating out the stimulus dependence yields estimates of the slowly-varying latent modulators. For example, when applied to array recordings of macaque V1, we find complex spatial patterns of correlation amongst the latent modulators, including clusters of co-modulated units. In sum, use of the modulated Poisson model improves inference, and enables the study of signals underlying non-stationarities in neural responses.



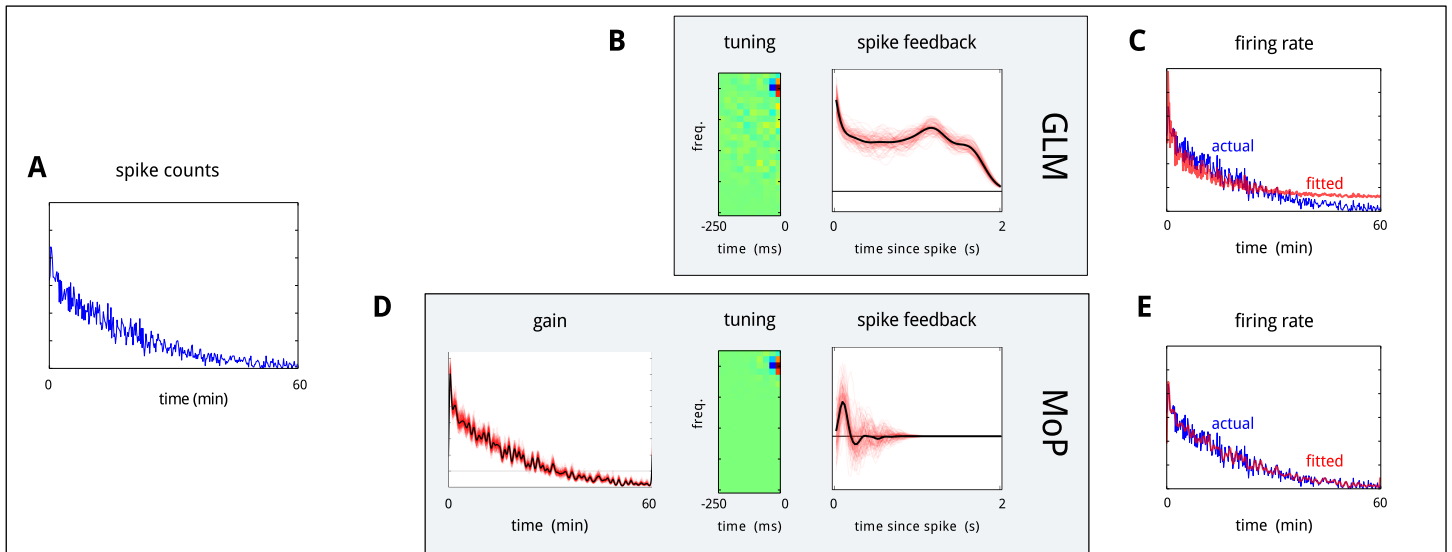


Figure 3: Comparison of GLM to the MoP model, in explaining responses of an auditory midbrain (IC) neuron (data from Rabinowitz et al., 2013). (A) Binned spike counts collected over 1 hr of recording time show a pronounced non-stationarity in firing rate. This could be due to physiological decline, or reduced spike detectability over time. (B) Estimated GLM parameters. Left: spectrotemporal receptive field. Right: spike feedback kernel (black line indicates MAP estimate, red lines are samples from posterior). In order to capture the monotonic decline in firing rate over time, the feedback kernel has a substantial temporal extent, and an unusual shape. (C) Simulated model spike rate compared to measured spike rate. (D) Estimated parameters for the Modulated-Poisson model. Left: estimated time-varying (latent) modulatory signal. Middle/Right: spectrotemporal receptive field and spike feedback kernel (as in GLM model). The latent modulatory input explains the slow non-stationarity in firing rates, freeing the spike feedback term to reveal fast modulations of gain. (E) Simulated model spike rate is improved, relative to the GLM model.




Figure 4: Properties of latent modulators for model fit to array data from anesthetized macaque V1 (data from Graf et al., 2011). Monkeys were shown a sequence of oriented gratings, over a duration of ~ 2.5 hrs. Each grating was shown for 1.28 sec, with 1.28 sec (uniform-contrast) delay between successive stimuli. Spikes were recorded on Utah arrays over a 4mm x 4mm region of visual cortex; from units isolated on unique electrodes, we concurrently estimated both orientation tuning curves, and the latent modulators. (A) Latent modulating signals, for each of the 77 units recorded concurrently. Signals are shown as $\log(\text{gain})$ values, a_t (see Fig. 1), normalized (z-scored) for each unit. The time-scale of the gain fluctuations is a hyperparameter that is learned for each unit. Several units experienced no significant gain modulation, and were best explained as having constant gain throughout (horizontal lines). Most notably, clear concurrent patterns are visible, whereby some (but not all) units participate in nearly identical fluctuations in gain. (B) Cross-correlation between the inferred signals in (A), with cells sorted by a hierarchical clustering algorithm. There are groups of units which undergo highly-correlated modulation (i.e., they have nearly identical latent modulators), and others with modulatory signals that are nearly independent of those of other neurons.

- [1] Goris et al. (2013). *Response variability of visual cortical neurons explained by a modulated...* SfN 2013.
- [2] Pillow et al. (2008). *Spatio-temporal correlations and visual signalling...* Nature, 454(7207).
- [3] Park & Pillow (2011). *Receptive field inference with localized priors.* PLoS computational biology, 7(10).
- [4] Rabinowitz et al. (2013). *Constructing noise-invariant representations of sound...* PLoS Biology, 11(11).
- [5] Graf et al. (2011). *Decoding the activity of neuronal populations in macaque...* Nature Neurosci. 14(2).