Low-dimensional models of neural population recordings with complex stimulus selectivity

Evan Archer, Jonathan W. Pillow, & Jakob H. Macke

Abstract

Modern experimental technologies such as multi-electrode arrays and 2-photon population calcium imaging make it possible to record the responses of large neural populations (up to 100's of neurons) simultaneously. These high-dimensional data pose a significant challenge for analysis. Recent work has focused on extracting low-dimensional dynamical trajectories that may underlie such responses. These methods enable visualization of high-dimensional neural activity, and may also provide insight into the function of underlying circuitry. Previous work, however, has primarily focused on models of a population's intrinsic dynamics, without taking into account any external stimulus drive.

We propose a new technique that integrates linear dimensionality reduction of stimulus-response functions (analogous to spike-triggered average and covariance analysis) with a latent dynamical system (LDS) model of neural activity. Under our model, the population response is governed by a low-dimensional dynamical system with nonlinear (quadratic) stimulus-dependent input. Parameters of the model can be learned by combining standard expectation maximization for linear dynamical system models with a recently proposed algorithms for learning quadratic feature selectivity [1]. Unlike models with all-to-all connectivity, this framework scales well to large populations since, given fixed latent dimensionality, the number of parameters grows linearly with population size.

Simultaneous modeling of dynamics and stimulus dependence allows our method to model correlations in response variability while also uncovering low-dimensional stimulus selectivity that is shared across a population. Because stimulus selectivity and noise correlations both arise from coupling to the underlying dynamical system, it is particularly well-suited for studying the neural population activity of sensory cortices, where stimulus inputs received by different neurons are likely to be mediated by local circuitry, giving rise to both shared dynamics and substantial receptive field overlap.

Additional Information

Model formulation: The dynamics of the n latent units x_t are given by

$$\mathbf{x}_t = A\mathbf{x}_{t-1} + f_B(h_t) + \epsilon_t,\ \ \epsilon_t \sim \mathcal{N}(0, Q),$$

where A is the $n \times n$ dynamics matrix, ϵ is Gaussian innovation noise with covariance matrix Q. Each dimension of x_t is influenced by the stimulus h_t via a quadratic stimulus model:

$$f_{B,i}(h_t) = a_i \left({w_i}^\top h_t\right)^2 + b_i \left({w_i}^\top h_t\right) + c_i,$$

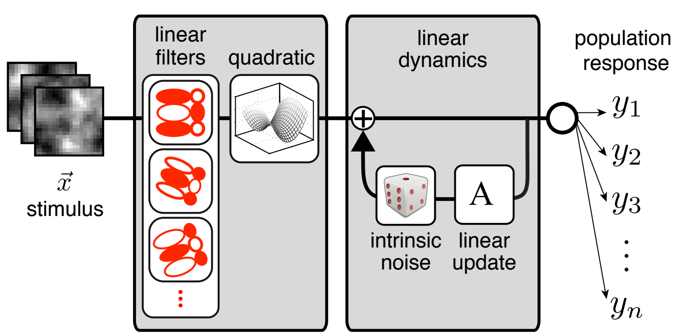

Figure 1: **Model schematic illustrating the ingredients of QLDS.** The sensory stimulus is processed by multiple units with quadratic stimulus selectivity (only one of which is shown). These units then provide input into a multi-dimensional linear dynamical system model. Finally, each neuron y_i in the population is influenced by the dynamical system via a linear readout. QLDS therefore models both the stimulus selectivity as well as the spatio-temporal correlations of the population.

where parameters $B = \{w_i, a_i, b_i, c_i : i \in 1, \ldots, m\}$ include multiple stimulus filters w_i as well as quadratic parameters (a_i, b_i, c_i), as discussed in [1]. The activity y_t of the m observed neurons is modelled as

$$y_t = Cx_t + d + \eta_t,\ \ \eta_t \sim \mathcal{N}(0, R),$$

i.e., each neuron has a linear readout of the latent dynamics, and single-neuron noise is additive Gaussian with (diagonal) covariance R. Variability-correlations across neurons arise from the fact that all neurons are coupled to the same underlying dynamical system. All parameters of the system are learned via expectation maximization. In the M-step, the objective function maximized to obtain the parameters B is given by

$$g(B) = -\frac{1}{2}\sum_{t=2}^{T}\left[(\mu_t - A\mu_{t-1} - f_B(h_t))^\top Q^{-1}(\mu_t - A\mu_{t-1} - f_B(h_t))\right],\ \text{where}\ \mu_t = \mathbb{E}[x_t|y_{t-1}, h_{t-1}].$$

We obtain the update for B numerically, iterating with closed-form updates for A and Q (similar to those derived in [1]), and making use of the gradient

$$\frac{\partial g(B)}{\partial B} = -Q^{-1}\sum_{t=2}^{T}(\mu_t - A\mu_{t-1} - f_B(h_t))\frac{\partial f(h_t)}{\partial B}.$$

Fig 2 shows an application of this method to a simulated neural dataset, illustrating that a low-dimensional LDS with quadratic inputs can describe the response properties of a heterogeneous collection of complex cells.

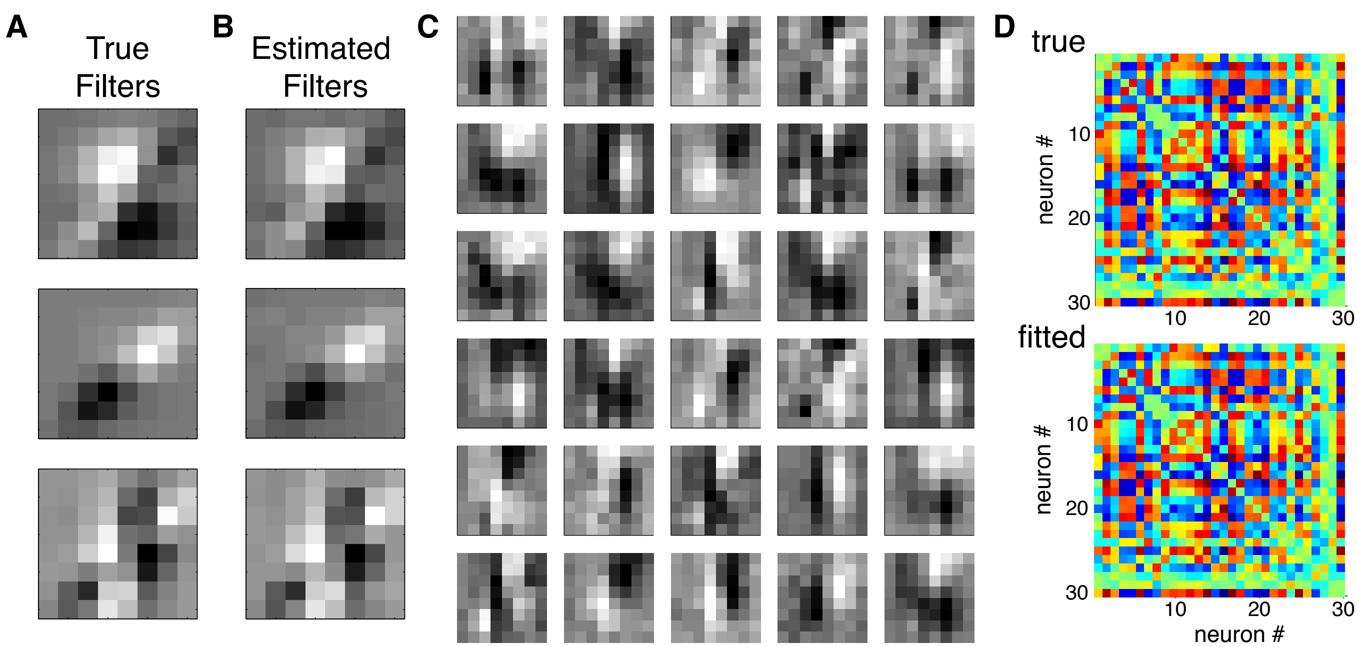

Figure 2: **Illustration of model on simulated data with $m = 30$ neurons A)** Stimulus filters of ground-truth model with $n = 4$ dimensional latent dynamics which has 3 stimulus selective inputs with quadratic nonlinearity ('energy units') **B)** First three stimulus filters recovered by QLDs fit to a simulation of length $T = 10000$ samples (recovered filters rotated to allow comparison with original filters). In this simulation, the fourth estimated dimension was weaker but still carried stimulus selectivity. **C)** Receptive fields of the 30 neutrons estimated by a conventional spike-trigging method. Because of the 'mixing' of stimulus-selectivity by latent dynamics, single-neuron selectivity appears much more complex and belies the simplicity of the underlying generating model. **D)** QLDS also recovers the cross-correlations of the population (top ground truth, bottom from fitted model).

References: [1]: Park, Archer, Priebe, & Pillow, NIPS 2013.