

spike data. The correlated Poisson distribution offers significant improvements in predictive likelihood for GLMs and R-GLMs. We evaluate the performance of the models on a dataset recorded from a Utah array implanted into motor areas of a macaque monkey during a delayed reaching task. We report that the R-GLM consistently finds long timescales (of up to several seconds) of correlated activity similar to those found by LDS and longer than the timescales learnt by standard GLMs (up to 400 ms). Like all GLMs, the proposed model can be used with any link function and any output distribution. This is unlike models based on LDS which require careful approximations to be trained with Poisson outputs.

III-11. Distinct coherent ensembles reflect working memory processes in primate PFC

David Markowitz
Bijan Pesaran
New York University

DAVID.MARKOWITZ@NYU.EDU
BP31@NYU.EDU

Coherent ensemble spiking activity has been linked with attention, coordination and decision-making processes. Computational modeling studies have suggested that coherent activity may be an important property of working memory (WM) networks, as well. Although recent experimental work has identified beta (11-30 Hz) frequency LFP oscillations as a prominent signature of WM in prefrontal cortex (PFC), the origin of these oscillations and their relationship to spiking activity in PFC during WM remain unknown. To address this question, we chronically implanted two monkeys with a 32-channel movable electrode array microdrive over PFC and recorded from ensembles of isolated units at each depth of cortex while the animals performed memory- and visually-guided delayed saccades. We find that distinct cellular response classes are revealed when memory-tuned units (391/706 cells) are grouped by recording depth (above/below 1 mm) and the presence of significant beta frequency spike-field coherency. Superficial coherent neurons ($n=74$) respond in a memory-specific manner, with persistently elevated rates during memory delays and baseline firing rates during visual delays. Deep coherent neurons ($n=75$) do not respond in a memory-specific manner, and exhibit the same persistently elevated rates during both memory and visual delays. These findings suggest that populations of superficial beta-coherent neurons specifically support memory maintenance through cross-columnar interactions, consistent with theoretical predictions. To determine if such interactions lead to firing rate covariation within coherent ensembles, we estimated noise correlations during the memory delay for $n=77$ superficial coherent pairs separated by 1.5 – 8 mm, and find that this is not significantly different from zero (0.01 *pm* 0.01 s.e.m.). Based on these findings, we propose a computational framework in which working memory processes are supported by distinct ensembles of coherent neurons at superficial depths (< 1 mm) within the prefrontal cortical sheet.

III-12. Dimensionality, dynamics, and correlations in the motor cortical substrate for reaching

Peiran Gao¹
Eric Trautmann¹
Byron Yu²
Gopal Santhanam
Stephen Ryu¹
Krishna Shenoy¹
Surya Ganguli¹

PRGAO@STANFORD.EDU
ETRAUT@STANFORD.EDU
BYRONYU@CMU.EDU
GOPAL@NERUR.COM
SEOULMANMD@GMAIL.COM
SHENOY@STANFORD.EDU
SGANGULI@STANFORD.EDU

¹Stanford University
²Carnegie Mellon University

How coherent motor behavior emerges from large populations of neurons constitutes a fundamental question in

neuroscience. Dimensionality reduction and correlation analysis are often used to address this question. Interestingly, neuronal data's effective dimensionality is often far smaller than the number of recorded neurons [Yu 2009], and recent work has shown that low-dimensional neuronal dynamics exhibits rotational structure [Churchland 2012]. Moreover, despite weak pairwise correlations, one can accurately predict a neuron's spiking from only $O(100)$ other neurons' activities [Truccolo 2005]. These various low dimensional views leave open several important questions: What determines the effective dimensionality of neuronal activity? What creates rotational structure? And how do spatiotemporal correlations enable spiking prediction? By analyzing 109 simultaneously recorded PMd neurons from monkeys performing an eight direction delayed reach task [Yu 2007], we find a simple view that answers these questions. Using Gaussian mixture models fitted to trial averaged activity, we find that most neurons exhibit a sharp, monophasic activity peak during movement [but see Churchland...2007]. Each peak's timing, but not amplitude, is largely independent of reach angle. This sparse wave of neural activity comprises a nonlinear manifold, which does not lie within a single low dimensional linear space, and evolves through different dimensions over time. We show empirically and analytically that: (a) the dimensionality of the smallest linear subspace containing this manifold is near an upper bound estimated by task complexity and network correlation time; (b) when projected onto a lower dimension, this manifold exhibits rotational structure; (c) inter-trial variability concentrates along the manifold; and (d) sparsity of activity underlies our ability to predict single neuron behavior from the ensemble. This work unifies and explains, through a single high-dimensional perspective, disparate phenomena previously viewed through different low-dimensional lenses, and suggests new analysis methods for finding rich neural structures that may be missed by time-independent dimensionality reduction.

III-13. Optimal neural tuning for arbitrary stimulus priors with Gaussian input noise

Zhuo Wang¹
 Alan Stocker¹
 Haim Sompolinsky²
 Daniel Lee¹

WANGZHUO@SAS.UPENN.EDU
 ASTOCKER@SAS.UPENN.EDU
 HAIM@FIZ.HUJI.AC.IL
 DDLEE@SEAS.UPENN.EDU

¹University of Pennsylvania

²The Hebrew University

Sensory neurons represent input information via their tuning curves. The optimal tuning curve can be determined by many factors, e.g. optimal criteria, stimulus prior distribution, output noise model and encoding time (Bethge et al. 2002, Ganguli & Simoncelli 2010). The output noise, i.e. the variability in spike-count over time, has been well investigated. Previous theoretical work (Nadal & Parga 1994) has provided some insights into the role input noise can play in optimal encoding, where input noise is considered any noise that affects the input of the neuron. Here we argue that by analyzing the input - output noise pair, we can have a better understanding of the neurons in deeper layers of a neural network since those neurons can only access the noisy output received from upstream neurons. In this work, we assume a Poisson output noise model and generalize our long encoding time limit result (Wang, Stocker & Lee 2012) to the case with the presence of an additive Gaussian input noise. Having the new Fisher information computed, various objective functions, e.g. mutual information and mean asymptotic Lp loss, can be calculated via Cramer-Rao lower bound and Mutual Information Bound (Brunel & Nadal 1998). We analytically solved the optimization problem and derived the optimal condition for the tuning curve. Our work answers the question how neurons should optimally combat with the input noise. Compare to other biophysiological approach (Chance et al. 2002, Lundstrom et al. 2009), our result directly links the phenomena - neuron's adaptation to synaptic noise - to the underlying motivation i.e. optimality criteria. Furthermore, our model can be used to better understand the loss function a neuron is trying to optimize when encoding information. By applying a small input perturbation and observe the change of the tuning curve, the actual optimality criterion can be identified.