Statistical assessment of the stability of neural movement representations

Ian H. Stevenson[1], Anil Cherian[2], Brian M. London[2], Nicholas A. Sachs[2], Eric Lindberg[2], Jacob Reimer[3], Marc W. Slutzky[1,2,4], Nicholas G. Hatsopoulos[3], Lee E. Miller[1,2,5], and Konrad P. Kording[1,2]

Poisson noise and nonlinear models of tuning
Traditional approaches to modeling cosine tuning assume that spiking noise is Gaussian. However, a popular alternative is to assume Poisson noise.

$$\lambda = b_0 + b_1 \cos(\theta - \theta^*)$$
$$y \sim Poisson(\lambda\, dt)$$

The assumption of Poisson noise is, likely, a better description of discrete spike counts and is consistent with observations that Fano factors tend to be close to one. However, several recent studies have also used a linear-nonlinear-Poisson (LNP) model with an exponential nonlinearity

$$\lambda = \exp(b_0 + b_1 \cos(\theta - \theta^*))$$
$$y \sim Poisson(\lambda\, dt) \qquad .$$

This model accurately captures the nonlinearities observed in some motor cortical neurons (Amirikian and Georgopulos 2000) and falls under the class of generalized linear models, making optimization efficient. In both of these alternative models, rather than minimizing the squared error as one normally would for the Gaussian noise assumption, we maximize the likelihood of this conditional Poisson process (see (Paninski 2004; Truccolo et al. 2005) for details, and (Amirikian and Georgopulos 2000; Swindale 1998) on the choice of tuning function). The maximum likelihood tuning curve estimates for the three models (cosine Gaussian noise, cosine Poisson noise, and LNP) are, in general, similar. However, there are occassional discrepancies in estimates of both preferred direction and modulation (Fig S1).

Bootstrapping as well as adaptive filtering can be applied to either of these two Poisson models as well. Importantly, using these models does not change the conclusions we draw about measurement noise and the stability of preferred directions. Applying the same analysis used in the main text, block-by-block changes in the estimated preferred direction, for the cosine tuning model with Poisson noise, have mean 0.0±1.1 deg and standard deviation 39.3±1.1 deg for blocks of 40 trials and mean 0.0±1.5 deg and standard deviation 22.4±1.3 deg for blocks of 120 trials. Similarly, for the LNP tuning model changes in preferred direction have mean 0.0±1.1 deg and standard deviation 38.9±1.1 deg for blocks of 40 trials and mean 2.4±1.5 deg and standard deviation 22.8±1.5 deg for blocks of 120 trials (Fig S2). In all cases, very few of the PD changes are significant (bootstrapping, α=0.05), and the histograms tend to be, again, almost identical to the histograms of PD changes observed after simulating stable neurons.

Hotelling's T-squared test and bootstrapping
Prevous studies have used a variety of methods to determine whether tuning and changes in tuning are significant. Notably, Rokni et al. (2007) test significance using a bivariate z-test on the parameters c_1 and c_2 (see (Christensen 2001; Rokni et al. 2007) and Methods Eq. (2)). Briefly, this test, based on Hotelling's T-squared statistic, compares a set of samples (c_1 and c_2 pairs) with some covariance to either zero (one-sample test) or another set of samples (two-sample test). Note that this test contains information about both the modulation depth $b_1 = (c_1 - c_2)/(\cos\theta^* - \sin\theta^*)$ and preferred direction $\theta^* = a\tan 2(c_2, c_1)$. Rather than changes in preferred direction alone, this test will, thus, reflect changes in the total cosine

component.

In practice, a one-sample Hotelling's T-Squared test works well for distinguishing tuned from untuned neurons. In the data used here 201 of the 285 recorded neurons were considered tuned ($p<0.05$, samples fit for blocks of 40 trials, all data). This is, in fact, an even more conservative test than the simple (cross-validated) likelihood ratio test where, comparing untuned and cosined tuned models with Gaussian noise, 249 of the 285 recorded neurons were considered tuned (i.e. tuning model has a greater likelihood).

For detecting changes between blocks of trials, the two-sample Hotelling's T-Squared test has several features that complicate the process of determining whether preferred directions are stable. 1) Changes in preferred direction cannot be separated from changes in modulation. As above this test works on $c_1 c_2$ pairs. 2) The test requires an estimate of the a sample covariance. In general, testing changes between blocks will require a *set* of samples for each block (~5-10). It is not obvious how to select this sub-block size, and the amount of data used to generate each $c_1 c_2$ sample will affect the covariance estimates and thus the test itself. These properties of Hotelling's T-Squared statistic suggest that bootstrapping may be a better, more reliable way to test for changes in individual tuning parameters.

A note on fully Bayesian methods and bootstrapping
Bootstrapping is a fast, easy to implement approach to estimating parameters and confidence intervals. When there is very little data, however, bootstrapping can be noisy and special care must be taken to smooth or correct bootstrap estimates. One solution to dealing with very small amounts of data is to use fully Bayesian methods (Cronin et al. 2010). However, for the amount of data used here (>40 trials) bootstrapping and fully Bayesian methods give nearly identical results. Since we are not implementing informative priors, bootstrapping is used here for its simplicity.
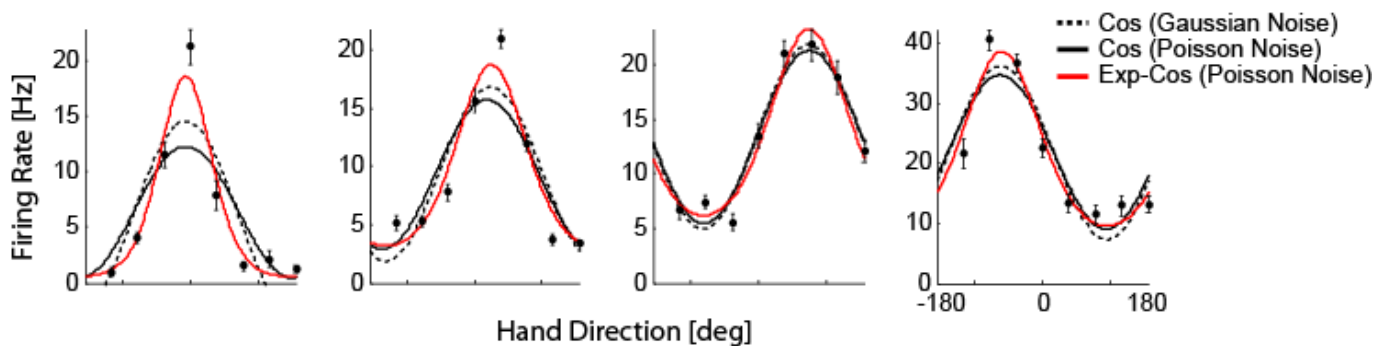

Figure S1: Typical maximum likelihood fits for three different tuning models: linear-nonlinear-Poisson (red), cosine tuned with Poisson noise (black), and cosine tuned with Gaussian noise (dashed) as used in the main text. Dots and errorbars denote mean and SEM for the observed spike counts 100 ms prior through 300 ms after movement onset.
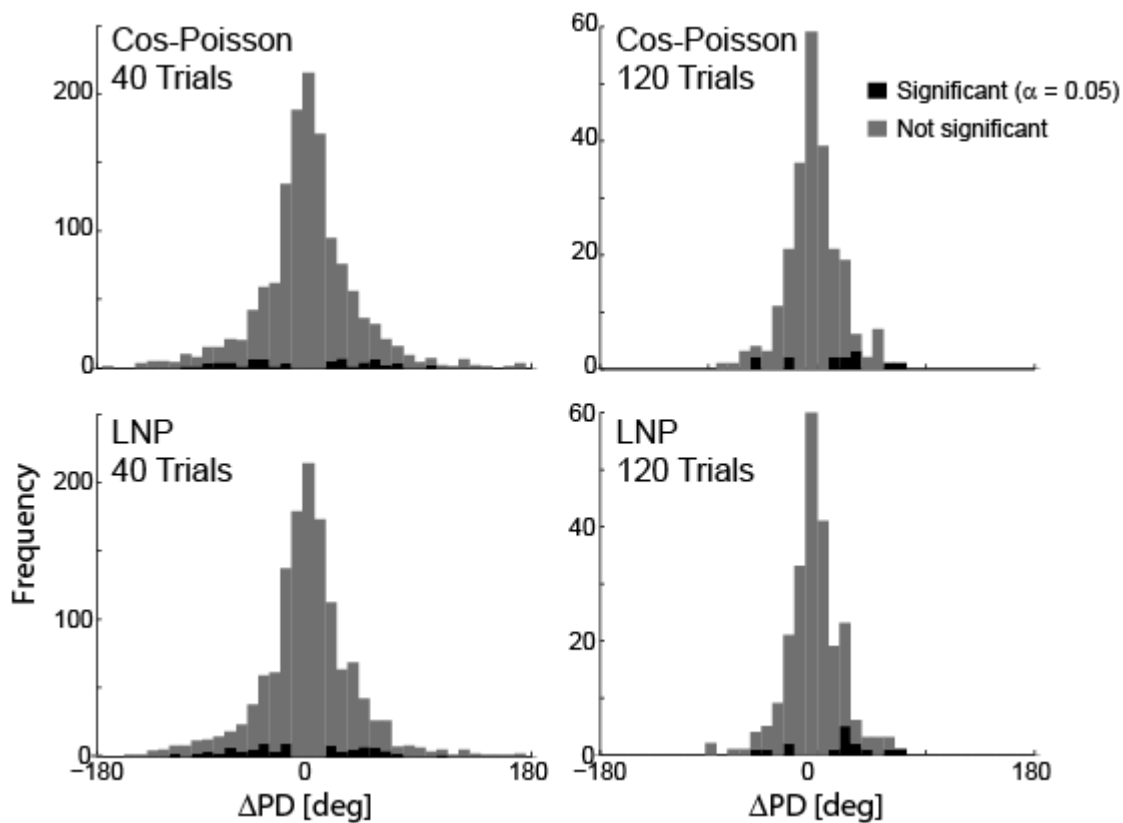
Figure S2: Estimated changes in preferred direction under two alternative tuning models: cosine with Poisson noise (top), and linear-nonlinear-Poisson (bottom). Black blocks denote changes that were significant at the 95% level by bootstrapping. Note that for blocks of 40 trials the variability of changes between blocks is substantially higher than for blocks of 120 trials, suggesting that measurement noise plays a large role in determining the magnitude of PD changes.

Amirikian B, and Georgopulos AP. Directional tuning profiles of motor cortical cells. *Neuroscience Research* 36: 73-79, 2000.

Christensen R. *Advanced linear modeling*. New York: Springer-Verlag, 2001.

Cronin B, Stevenson IH, Sur M, and Kording KP. Hierarchical Bayesian Modeling and Markov Chain Monte Carlo Sampling for Tuning-Curve Analysis. *Journal of Neurophysiology* 103: 591, 2010.

Paninski L. Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems* 15: 243-262, 2004.

Rokni U, Richardson AG, Bizzi E, and Seung HS. Motor learning with unstable neural representations. *Neuron* 54: 653-666, 2007.

Swindale NV. Orientation tuning curves: empirical description and estimation of parameters. *Biological Cybernetics* 78: 45-56, 1998.

Truccolo W, Eden UT, Fellows MR, Donoghue JP, and Brown EN. A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate Effects. *Journal of Neurophysiology* 93: 1074-1089, 2005.