

Statistical assessment of the stability of neural movement representations

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

J Neurophysiol 106:764-774, 2011. First published 25 May 2011; doi:10.1152/jn.00626.2010

You might find this additional info useful...

Supplemental material for this article can be found at:

<http://jn.physiology.org/content/suppl/2011/06/30/jn.00626.2010.DC1.html>

This article cites 52 articles, 28 of which can be accessed free at:

<http://jn.physiology.org/content/106/2/764.full.html#ref-list-1>

Updated information and services including high resolution figures, can be found at:

<http://jn.physiology.org/content/106/2/764.full.html>

Additional material and information about *Journal of Neurophysiology* can be found at:

<http://www.the-aps.org/publications/jn>

This information is current as of August 12, 2011.

Statistical assessment of the stability of neural movement representations

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

Departments of ¹Physical Medicine and Rehabilitation, ²Physiology, and ⁵Neurology, Feinberg School of Medicine, and ³Department of Biomedical Engineering, Northwestern University, Chicago, Illinois; and ⁴Department of Organismal Biology and Anatomy, University of Chicago, Chicago, Illinois

Submitted 15 July 2010; accepted in final form 18 May 2011

Stevenson IH, Cherian A, London BM, Sachs NA, Lindberg E, Reimer J, Slutzky MW, Hatsopoulos NG, Miller LE, Kording KP. Statistical assessment of the stability of neural movement representations. *J Neurophysiol* 106: 764–774, 2011. First published May 25, 2011; doi:10.1152/jn.00626.2010.—In systems neuroscience, neural activity that represents movements or sensory stimuli is often characterized by spatial tuning curves that may change in response to training, attention, altered mechanics, or the passage of time. A vital step in determining whether tuning curves change is accounting for estimation uncertainty due to measurement noise. In this study, we address the issue of tuning curve stability using methods that take uncertainty directly into account. We analyze data recorded from neurons in primary motor cortex using chronically implanted, multi-electrode arrays in four monkeys performing center-out reaching. With the use of simulations, we demonstrate that under typical experimental conditions, the effect of neuronal noise on estimated preferred direction can be quite large and is affected by both the amount of data and the modulation depth of the neurons. In experimental data, we find that after taking uncertainty into account using bootstrapping techniques, the majority of neurons appears to be very stable on a timescale of minutes to hours. Lastly, we introduce adaptive filtering methods to explicitly model dynamic tuning curves. In contrast to several previous findings suggesting that tuning curves may be in constant flux, we conclude that the neural representation of limb movement is, on average, quite stable and that impressions to the contrary may be largely the result of measurement noise.

neurons; sensory stimuli; spatial tuning curves; Poisson noise; tuning curves; primary motor cortex

MANY EXPERIMENTS IN NEUROSCIENCE aim at measuring how experimental manipulations affect the tuning properties of neurons. Tuning curves typically characterize how firing rates depend on a single relevant property of a stimulus or movement. For example, visual neurons are often characterized by their dependence on stimulus orientation, auditory neurons by their dependence on pitch, and motor cortical neurons by their dependence on the direction of movement of an animal's hand. A common method for studying how these coding properties change is to characterize the tuning curve, introduce an experimental manipulation, and characterize the tuning curve again to assess if changes have occurred (Cronin et al. 2010; Li et al. 2001; Rokni et al. 2007; Schummers et al. 2007). Tuning curves, in a variety of brain areas, have been shown to change during adaptation and learning. In the absence of an experi-

mental manipulation, we can ask how stable tuning curves are and whether they drift due to the passage of time. Importantly, determining whether changes in tuning are statistically significant depends on our ability to take into account the effects of noise and limited data.

Since noise in the recorded data (e.g., Poisson noise in spike counts) will affect tuning curve estimates, it is important to characterize any tuning curve parameters with confidence bounds. Without such bounds, it is impossible to determine whether an observed change is due to actual changes in tuning or simply measurement noise and uncertainty in the estimation. We will tend to overestimate the average (absolute) magnitude of any existing tuning curve changes, since the measured change consists of the real change plus apparent changes due to measurement noise. Therefore, measurement noise must be considered to obtain reasonable estimates of the magnitude of changes in tuning curves.

Here, we focus on the role of measurement uncertainty in estimating the tuning curves of neurons in primary motor cortex (M1). In particular, we examine the discharge properties of neurons during a well-learned, center-out task, where a monkey is trained to reach from the workspace center to several peripheral target locations (Georgopoulos et al. 1982; Kakei et al. 1999; Kalaska and Hyde 1985; Morrow et al. 2007; Scott and Kalaska 1997). Typically, tuning curves that characterize the discharge of neurons in M1, as a function of reaching direction, are well fit by a cosine function (Georgopoulos et al. 1982, 1988). The peak of this cosine indicates the direction of movement for which the neuron's firing rate is maximal, the "preferred direction" (PD).

A central question in the neural control of movement is whether the PDs of M1 neurons are stable during normal behavior, from reach to reach. If tuning curves fluctuate from reach to reach, the motor system must either be redundant, such that fluctuations do not affect behavior (Pohlmeier et al. 2007; Rokni et al. 2007), or must rapidly adapt to allow for stable movement. By recording many trials from the same neurons over the course of several days, a number of studies have shown that neuronal tuning appears to be relatively stable (Chestek et al. 2007; Dickey et al. 2009). However, there is also some contrasting evidence to suggest that neurons fluctuate rapidly on timescales on the order of minutes to hours (Rokni et al. 2007). One recent study found that the PDs of movement-related neurons in M1 varied substantially over time with a majority of neurons shifting as much as 30° within 15 min (Rokni et al. 2007). These results suggested that neurons may be quite unstable and that there must be substan-

Address for reprint requests and other correspondence: K. P. Kording, Dept. of Physical Medicine and Rehabilitation, Northwestern Univ., Feinberg School of Medicine, 345 East Chicago Ave., Attn: ONT-931, Chicago, IL 60611 (e-mail: kk@northwestern.edu).

tial redundancy in the nervous system and downstream processing to ensure stable behavior. However, these estimates were obtained with methods that did not explicitly correct for measurement noise.

In this study, we address these issues using methods that take measurement uncertainty and potential instability directly into account. We analyze data recorded from neurons in M1 using chronically implanted multielectrode arrays. With the use of simulations, we demonstrate that under typical experimental conditions, the effect of neuronal noise on estimated PD can be quite large. With the use of experimental data, we find that after taking uncertainty into account, the average patterns of neuronal activity expressed by tuning curves are quite stable on a timescale of minutes to hours. We then explore ways of explicitly modeling dynamic tuning curves using adaptive filtering methods. These methods reveal that there may be tiny but real fluctuations of the tuning properties of neurons but that such changes are very hard to detect robustly. We conclude that the neural representation of limb movement is actually quite stable and that impressions to the contrary may be the result largely of measurement noise.

METHODS

Tasks. Four monkey subjects (designated F, K, C, and R) were trained on an eight-target center-out reaching paradigm. Monkeys F, K, and C were seated in a primate chair with movement constrained to a horizontal plane. The monkey grasped the handle of a two-link planar manipulandum that moved within a 20-cm by 20-cm workspace. Feedback about movement was given on a computer screen in front of the monkey. Handle position was displayed as a circular cursor, 1–2 cm in diameter. The experiments with monkey R used a KINARM device (BKIN Technologies, Kingston, ON, Canada), in which the monkey's arm rested on cushioned troughs secured to links of a two-joint robotic arm (Scott 1999). The shoulder joint was abducted 90°, such that shoulder and elbow flexion and extension movements were made in the horizontal plane. A cursor coincident with the handle position of the robotic arm was projected onto a horizontal screen placed above the monkey's arm. All trials began with the acquisition of a square center target, which the monkey was required to hold for 0.3–1.1 s (0.6–1.1 s for monkey F, 0.5–0.6 s for monkey C, 0.3–0.5 s for monkey K). After this hold period, subjects reached to one of eight equally spaced peripheral targets. Subjects had ~1.25 s to acquire the outer target and were required to hold this outer target for at least 0.2–0.5 s. Each success was rewarded with juice or water.

Surgery. Once the monkey was able to perform the center-out task satisfactorily, we performed a surgery to implant a 100-electrode intracortical array (Blackrock Microsystems, Salt Lake City, UT). We made a craniotomy centered above the arm area of M1. After opening the dura, we identified the area on the crown of the precentral gyrus, just medial to the spur of the arcuate sulcus. In some cases, we stimulated within this area using a ball electrode array (monopolar, biphasic, 50 Hz, 100- μ s pulse width, <6 mA) to locate proximal limb movements. All surgery was performed under isoflurane gas anesthesia except during intraoperative stimulation. To increase cortical excitability, ~30 min prior to stimulation, we began infusing remifentanyl (0.4 μ g/kg/min iv), while gradually reducing the concentration of isoflurane to 0.25%. We positioned the array within the identified area, taking care to avoid any major surface vessels, and inserted it rapidly with a pneumatic inserter (supplied by Blackrock Microsystems). The dura was closed over the array with a piece of artificial pericardium (Preclude expanded polytetrafluoroethylene membrane, Gore and Associates, Flagstaff, AZ) under the dura to prevent the back of the array from adhering to the dura. Another piece of Preclude

was placed over the dura. The original bone flap was thinned and replaced, and the skin was closed over the craniotomy. All leads from the array were routed to a percutaneous connector secured to the monkey's skull (Nordhausen et al. 1996).

All animal use procedures were approved by the Institutional Animal Care and Use Committees at Northwestern University (Chicago, IL; datasets F, K, and C) or the University of Chicago (Chicago, IL; dataset R).

Tuning curve estimation and simulations. Spike trains from each subject were recorded during center-out reaching for 30–40 min, resulting in at least 300 successful trials for datasets F, R, and K and 290 trials for dataset C. Neural signals were classified as single units based on action-potential shape and minimum interspike intervals of 1.6 ms. Spike sorting was performed offline by manual cluster cutting. Trial-by-trial spike counts from 100 ms prior to movement onset until 300 ms after movement onset are used throughout the analysis.

Following the conventional cosine-tuning model, we assume that the firing rate of each neuron depends on the direction of hand movement, θ as

$$y = b_0 + b_1 \cos(\theta - \theta^*) \quad (1)$$

where the parameters b_0 , b_1 , and θ^* denote the baseline firing rate, modulation, and the PD (Georgopoulos et al. 1986). Traditionally, these parameters are estimated by minimizing the squared error between predicted and actual firing rates (Swindale 1998). To simplify the optimization, we can use the sum-difference formula for cosines to rewrite Eq. 1 as

$$y = b_0 + c_1 \cos(\theta) + c_2 \sin(\theta) \quad (2)$$

giving $\theta^* = \text{atan}(c_2, c_1)$ and $b_1 = (c_1, c_2)/(\cos\theta^* - \sin\theta^*)$. In this form, we can efficiently estimate the parameters using linear regression. Note that since minimizing the squared error is equivalent to maximizing the log likelihood with a Gaussian noise model, this optimization implicitly assumes Gaussian noise. In the low firing rate limit, other noise models, such as Poisson, may be more appropriate (Cronin et al. 2010), and several studies have suggested that exponential cosine-tuning functions may be a better description of neurons in M1 (Amirikian and Georgopoulos 2000; Hatsopoulos et al. 2007). An analysis of these alternative tuning models is presented (see Supplemental Material). However, for the main analysis, we assumed cosine tuning with Gaussian noise as a standard model.

For simulations, the firing rate is assumed to be constant over the observation window, drawn from a Poisson distribution with a firing rate given by Eq. 1. For the example neurons, we fixed the parameters by hand to a physiologically plausible range, and for a more complete analysis, we incrementally varied the modulation depth. For comparisons with recorded data, the parameters b_0 , b_1 , and θ^* were estimated from the entire recording session for each recorded neuron. We then simulated “stable” neurons with the same tuning properties and compared these results, where tuning curves are fixed, with the observed data. This allowed us to test whether results, which previously seemed to suggest tuning curve drift, might be consistent with noisy, stable neurons. It is important to note that although tuning curve estimation assumed Gaussian noise, the simulations used here assumed Poisson noise. This is somewhat nonstandard from the point of view of statistical inference, but we have used this approach since the Gaussian noise assumption is so commonly used for estimating motor tuning curves.

Bootstrapping. To estimate the uncertainty associated with each of the parameters in Eq. 1, we used bootstrapping (Davison and Hinkley 1997; Efron and Tibshirani 1997). Briefly, we randomly resampled the original data with replacement and estimated the parameters from each of these resampled datasets. Bootstrapping creates a distribution of parameters consistent with the data, from which confidence intervals can be estimated, and significance can be assessed. In practice, the results are robust to the choice of number of bootstrap samples and

in this case, were nearly identical for >500 samples. In the following analyses, we used 1,000 bootstrap samples.

The variability of the bootstrap samples gives an indication of how uncertain the parameter estimates are. Confidence intervals for each of the parameters can then be assessed using the quantiles of the bootstrapped distributions over the parameters. For example, 95% confidence intervals are estimated by the 2.5 and 97.5 percentiles of the distribution. After sorting the samples, the lowest and highest 2.5% of the samples are discarded. The range of the remaining samples then defines the error margin. For 1,000 bootstrap samples, these would be the 26th and 975th samples from the ranked list. Importantly, this method does not need to make the assumption that the distribution is Gaussian, symmetric, or even unimodal.

Since the PD parameter is in circular coordinates, special care is needed to estimate the confidence interval. A number of methods exist for dealing with small (<100) sample sizes (Fisher 1996; Otieno and Anderson-Cook 2008). However, in our case (~1,000 bootstrap samples from well-tuned neurons), it suffices first to center the samples by the median and then, to apply the quantile procedure described above.

Estimating confidence intervals allows us to assess statistical significance. Mainly, we want to determine whether the change in PD from one block of trials to the next is statistically significant. In this case, we examined the distribution of differences between the bootstrap samples for PD from each block. If the 95% confidence interval for this distribution of differences did not contain 0, then we reject the null hypothesis (that the difference actually is 0) at the 5% significance level. For a previous application of this method to tuning curves, see Churchland and Shenoy (2007).

In general, the size of confidence intervals is determined by the number of spike observations. For example, the SEM for a Poisson distributed random variable scales as $1/\sqrt{N}$, where N is the number of observations. For cosine-tuned neurons, the confidence interval about the PD is the main quantity of interest. For simplicity, we can reparameterize the tuning curves (Eq. 1) in terms of the total number of spikes emitted by a neuron and the depth of modulation relative to baseline ($M = b_1/b_0$). Given the assumptions of Poisson noise and cosine tuning, uncertainty about PD is affected only by these two factors.

To assess how these factors influence the measurement uncertainty, we parametrically varied the total expected number of spikes as well as the modulation depth. For each simulated neuron, we sampled spikes from a Poisson distribution in eight reach directions with the firing rate given by a cosine-tuning curve. Without loss of generality, we can assume that the PD is 0° . We then used the bootstrapping approach described above to generate bootstrap PD distributions for each modulation-depth and spike-count combination and determined the 95% confidence bounds of the distribution.

To compare our results with previously published data, we read out the maximal and minimal firing rates across direction from the graphs published in several papers (Kalaska et al. 1989; Rokni et al. 2007; Wise et al. 1998). In these cases, we estimated modulation depth from the tuning curves as $\hat{M} = (\max - \min)/((\max + \min)/2)$, and the total number of spikes was estimated as $\hat{N} = Tw(\max + \min)/2$ where T is the number of trials, and w is the integration window.

Corrected estimate of the variance of changes in PD. In the presence of measurement noise, the differences that we observe in PD between blocks will be partially due to a true (hidden) drift and partially due to the measurement noise. Therefore, we will tend to overestimate fluctuations in PD. With the use of PD to denote the true, but unobserved, PD and PD_{est} to denote the estimated PD

$$\text{var}(PD_{1,est} - PD_{2,est}) = \text{var}((PD_1 + \varepsilon_1) - (PD_2 + \varepsilon_2)) \quad (3)$$

where ε_1 and ε_2 denote instantiations of the measurement noise, which is assumed Gaussian. Since the variance of a sum of indepen-

dent random variables is equal to the sum of the variances, the corrected estimate of the fluctuations in PD is given by

$$\text{var}(PD_1 - PD_2) = \text{var}(PD_{1,est} - PD_{2,est}) - \text{var}(\varepsilon_1) - \text{var}(\varepsilon_2) \quad (4)$$

which captures the excess variance (above the measurement noise) that may be attributed to an actual change in PD. We estimated $\text{var}(\varepsilon)$ empirically by using the variance of the bootstrap samples (mean across the population of neurons). In some cases, the corrected value $\text{var}(\Delta PD)$ was negative. These values were set to 0 to calculate summary statistics and suggest that there was no detectable PD drift after accounting for measurement noise. The mean and difference statistics reported in the text were all calculated using circular statistics (Berens 2009). However, the equations above use a linear rather than circular coordinate system for correcting the variance. In the results presented here, the distribution of ΔPD is narrow enough that Eqs. 4 and 5 provide an accurate approximation to the fully circular variance correction.

Adaptive filtering. Although the methods described above for assessing and describing drift are fairly robust, one potential issue with these approaches is that they assume that tuning curves are stable within each block and that each trial is independent. Ideally, we should use a method that explicitly assumes that tuning curve parameters are drifting over time from trial to trial. Adaptive filtering is one approach that makes this assumption explicit. Several recent studies have used point-process adaptive filtering to describe time-varying tuning of place cells in hippocampus (Brown et al. 2001; Eden et al. 2004). Here, for consistency with the initial analysis, we analyze trial-by-trial spike counts, assuming cosine tuning with Gaussian noise, and use least-mean-squares (LMS) or steepest-descent adaptive filtering (Haykin 1996). Briefly, the parameters of the tuning model β are updated every trial following

$$\hat{\beta}_k = \hat{\beta}_{k-1} - \varepsilon \frac{\partial l_k(\beta)}{\partial \beta} \Big|_{\beta=\hat{\beta}_{k-1}} \quad (5)$$

where $l_k(\beta)$ denotes the instantaneous, negative log likelihood (squared error), and ε denotes a learning rate that must be optimized separately. Assuming Gaussian noise and writing (Eq. 2) in matrix form, such that $y = X\beta$, where $X_k = [1 \cos(\theta_k) \sin(\theta_k)]$, and $\beta = [b_0 \ c_1 \ c_2]^T$, we have the updates

$$\hat{\beta}_k = \hat{\beta}_{k-1} + \varepsilon(y_k - X_k \hat{\beta}_{k-1})X_k^T \quad (6)$$

The aim of this analysis is to update the parameters at each time-step to minimize prediction errors. The learning rate for each neuron and parameter is chosen to maximize the total likelihood of the time-varying parameter estimates. The final time-varying parameter estimates then allow us to extract the per-trial drift rate and fluctuations for each parameter. For the analyses presented here, we focus solely on the dynamics of the PD. In this case, the tuning curve is parameterized following Eq. 1, and the baseline and modulation are held constant, whereas the PD is updated as

$$\hat{\theta}_k^* = \hat{\theta}_{k-1}^* + \varepsilon(y_k - (b_0 + b_1 \cos(\theta_k - \hat{\theta}_{k-1}^*))) * \sin(\theta_k - \hat{\theta}_{k-1}^*). \quad (7)$$

The LMS algorithm works well for processes where the parameters follow a Gaussian random walk. To examine the performance of adaptive filtering, we simulated cosine-tuned Poisson neurons with fixed baseline firing rate and modulation, but where the PD followed a random walk

$$\theta_k^* = \theta_{k-1}^* + \eta \quad \eta \sim N(\mu, \sigma). \quad (8)$$

In this case, the PD changes by a random amount η each trial, drawn from a Gaussian distribution with mean μ and SD σ .

Time-varying tuning can lead to surprising results. For instance, a neuron that appears very broadly tuned using traditional block-

estimation techniques may, in reality, be a sharply tuned but drifting neuron (Brown et al. 2001). Finally, it is important to note that when the parameters do not have linear dynamics but instead, change suddenly, recursive methods may be more appropriate to estimate these changes accurately [see Eden et al. (2004)]. Here, we used the more basic LMS algorithm as a first step in analyzing trial-by-trial fluctuations in tuning parameters.

RESULTS

Measuring uncertainty in estimates of PD. To illustrate how measurement noise affects estimates of neuronal PD, we simulated the discharge of a typical cosine-tuned Poisson neuron (Fig. 1A) during repeated reaches to eight targets. We assume the simulated neuron has a movement-related discharge with a baseline firing rate of 20 Hz and a modulation of 5 Hz. Without loss of generality, we assume that the neuron has an actual PD of 180° . Due to Poisson noise in the simulation, the number of spikes generated during different reaches to the same target differed considerably (Fig. 1B). This type of variability is typical of experimental recordings. To characterize measurement uncertainty in this situation, we used the well-established technique of bootstrapping. Essentially, bootstrapping is a procedure that produces alternative fits, which could have been obtained if one could repeat a given experiment many times. In the case of the simulated neuron in Fig. 1, we initially simulated 40 reaches. With bootstrapping, we resampled from these data with replacement, each time computing a tuning curve and PD (Fig. 1C). Ultimately, this process can be used to obtain a distribution of possible PDs that is compatible with the original spike data and that can be used to estimate how reliable each tuning parameter is (Fig. 1D). In this example, the 95% confidence interval spanned $\sim 40^\circ$. The parameters that we used were well within the range of typical M1 experiments (Georgopoulos et al. 1982; Kalaska et al. 1989; Morrow et al. 2007; Wise et al. 1998); however, even after 40 trials, there was considerable uncertainty in the PD, and this uncertainty directly determines whether a change in tuning can be detected.

To determine whether the tuning curve of a neuron has changed, we can use this same bootstrapping approach on two successive blocks of data. For illustration, we simulated a hypothetical neuron with a tuning curve that changed following an experimental manipulation (Fig. 2A). We again assumed that this neuron has a baseline firing rate of 20 Hz and a modulation of 5 Hz. In this case, we assumed that the PD changes from 90° in the first condition to 135° in the second condition. After simulating 40 reaches for each condition (Fig. 2B), we calculated distributions over potential PDs for each of the two conditions using bootstrapping (Fig. 2C). The two resulting distributions of PDs overlapped only slightly, and the distribution of differences between the PDs was significantly different from 0. For this simulation, bootstrapping allowed us to conclude that the PD changed as a result of the experimental manipulation. However, the width of the resulting histograms ($\sim 40^\circ$) indicates that changes in PD for a given neuron need to be rather large to be visible with this amount of data.

It is important to understand how measurement uncertainty varies as a function of these simulation parameters. Given cosine tuning and Poisson noise, there are two relevant parameters that determine the precision of PD estimates: the total number of spikes and the modulation depth relative to baseline (see METHODS). We used the bootstrapping methods described

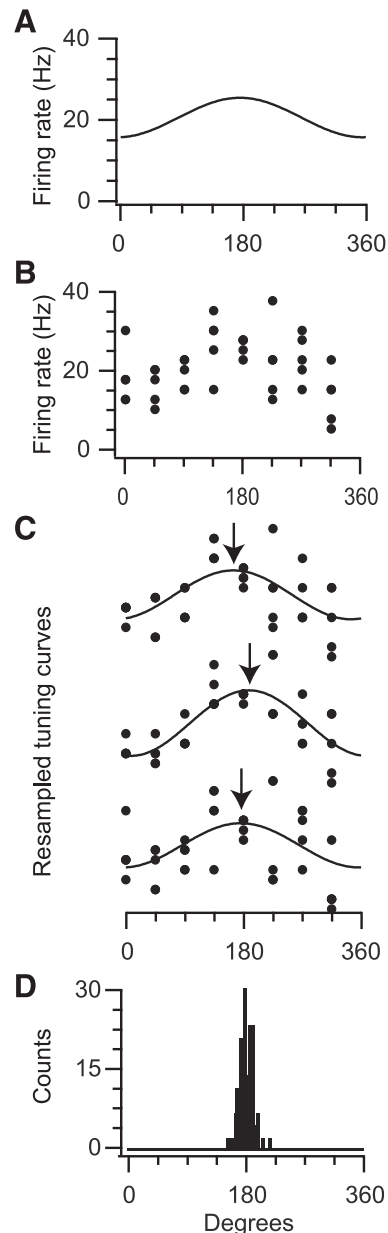


Fig. 1. Measurement noise for a simulated neuron with known tuning. A: idealized cosine-tuning function of a neuron with typical movement-related discharge (20 Hz) and a modulation (5 Hz). B: 40 trials simulated during movement in 8 directions with Poisson spike noise. C: 3 example bootstrap samples from the initial set of observations in B. Note that there is substantial variability in the estimated preferred direction (PD; arrows). D: using many bootstrap samples, we build a distribution of PDs that allows estimation of confidence intervals. In this case, the 95% confidence interval spans $\sim 40^\circ$.

above to calculate the average confidence interval size as a function of these two parameters (Fig. 3). In this case, a modulation depth of 1 corresponded to a cosine between 0 and the maximal firing rate, and a modulation depth of 0 corresponded to an untuned cell. In general, more strongly modulated neurons yield more precise PD estimates, and as more data are observed, the estimates of all tuning curve parameters become more precise.

Based on these simulated results, the precision of PD estimates in typical studies is likely to vary substantially. Here, we have shown total spike numbers and modulation depths, for

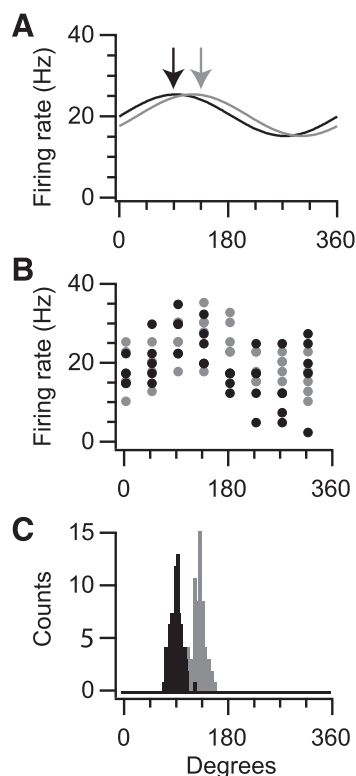


Fig. 2. Measurement noise affects the ability to detect PD changes. *A*: cosine-tuning curves for a pair of simulated neurons having PDs of 90° (black) and 135° (gray). *B*: 40 trials simulated from each of these tuning curves during movement in 8 directions with Poisson spike noise. *C*: bootstrap distributions for the PD based on the data in *B*. Note that there is overlap between these distributions, and the differences between the two distributions allow us to assess whether there is a significant tuning difference between the two conditions.

example, neurons in several published studies (Kalaska et al. 1989; Rokni et al. 2007; Wise et al. 1998), as well as our own data (Fig. 3*B*). Some of these studies allowed PDs to be estimated with a precision of $\sim 10^\circ$, whereas others were probably precise only to 50° or 60° . For our data, the average modulation depth was 0.49 ± 0.02 , with 6.23 ± 0.39 spikes/trial on average (400 ms integration window). Mapping these values onto the simulation results suggested an average confidence interval size of 33.5° for blocks of 40 trials, 20.6° for blocks of 120 trials, and 16.9° for blocks of 240 trials.

Here, we have assumed that the only source of variability, in addition to directional tuning, is Poisson noise. Because there are certainly other noise sources, the values reported in Fig. 3 should be seen as a lower bound on the uncertainty. However, this figure can be used to estimate measurement uncertainty, both for data that are already published and also to provide a type of power analysis for future studies. Given a baseline firing rate and modulation depth, this figure provides an approximate number of trials necessary to achieve a desired precision in PD estimates. For the purposes of this study, it is important to note that uncertainty in estimates of PD can be of the same order of magnitude as the size of changes reported in previous papers—changes that have been attributed to either experimental manipulations or instability.

Assessing tuning curve stability, taking uncertainty into account. Whereas examples and power analysis illustrate how bootstrapping can be used to assess measurement uncertainty

in the presence of simulated Poisson noise, bootstrapping is also a powerful technique for the analysis of actual data. Several recent studies have suggested that the PDs of neurons in M1 may be drifting rapidly over the course of tens of minutes (Carmena et al. 2005; Rokni et al. 2007). However, in describing changes in PDs, these studies did not explicitly account for measurement uncertainty. As noted above, estimated changes in PD can easily be inflated, since they include effects due to both any actual changes and measurement noise. Here, we used bootstrapping to assess the stability of M1 neurons during well-learned, center-out reaching.

To address these issues experimentally, we analyzed how the properties of recorded neurons in M1 from four animal subjects evolved over the course of a recording session. We recorded four datasets: C (79 neurons), F (36 neurons), K (92 neurons), R (78 neurons), each consisting of at least 290 successful trials with sessions lasting 30–40 min. After collecting spike counts from 100 ms prior to movement onset to 300 ms following movement onset, 70.5% of the recorded neurons (201/285 total—38/79, 34/36, 72/92, and 57/78 in datasets C, F, K, and R, respectively) showed significant cosine tuning across all trials ($\alpha = 0.05$, Hotelling's T^2 -test, see Supplemental Note). In the following sections, we limited our analysis to these significantly tuned neurons.

Following previous approaches (Rokni et al. 2007), we divided our dataset of center-out movements into blocks of trials. For each block, we estimated the neuron's PD, baseline firing rate, and modulation (see METHODS) and determined the 95% confidence interval for the PD estimates using bootstrapping. To compare these empirical results with the previous power analysis, we first varied the number of trials included in each block (block size) and noted the average confidence interval size. As suggested by the power analysis, the size of the confidence interval decreases as the block size increases (Fig. 4). In this case, where modulation is measured from the data, confidence interval size dropped $\sim 1/\sqrt{\text{number of trials}}$. In close approximation to the power analysis estimates, we found that in real data, the average confidence interval size is 43.7° for blocks of 40 trials, 24.4° for blocks of 120 trials, and 16.3° for blocks of 240 trials. Simulated Poisson neurons that were stable and matched for tuning showed a very similar decrease in confidence interval size as a function of the number of trials (Fig. 4, red curves).

Keeping in mind that measurement uncertainty (confidence interval size) is related to the block size, we then assessed stability in PD and examined changes in tuning between blocks. With the use of bootstrapping analysis and examining successive blocks of 120 trials, we found that many neurons appear to be quite stable (Fig. 5*A*). However, there were neurons in each of the four datasets that appeared to drift and have relatively large changes in PD, modulation, and baseline firing between successive blocks (Fig. 5*B*). To determine whether these changes are real effects or simply artifacts of high measurement uncertainty, it is important to consider whether the observed changes are statistically significant. Greater drift in PD is correlated with smaller baseline firing rates (Fig. 5*C*), as well as smaller modulation (Fig. 5*D*). In general, the neurons with larger apparent changes were the ones that also had large confidence intervals and high uncertainty about PD (Fig. 5*E*).

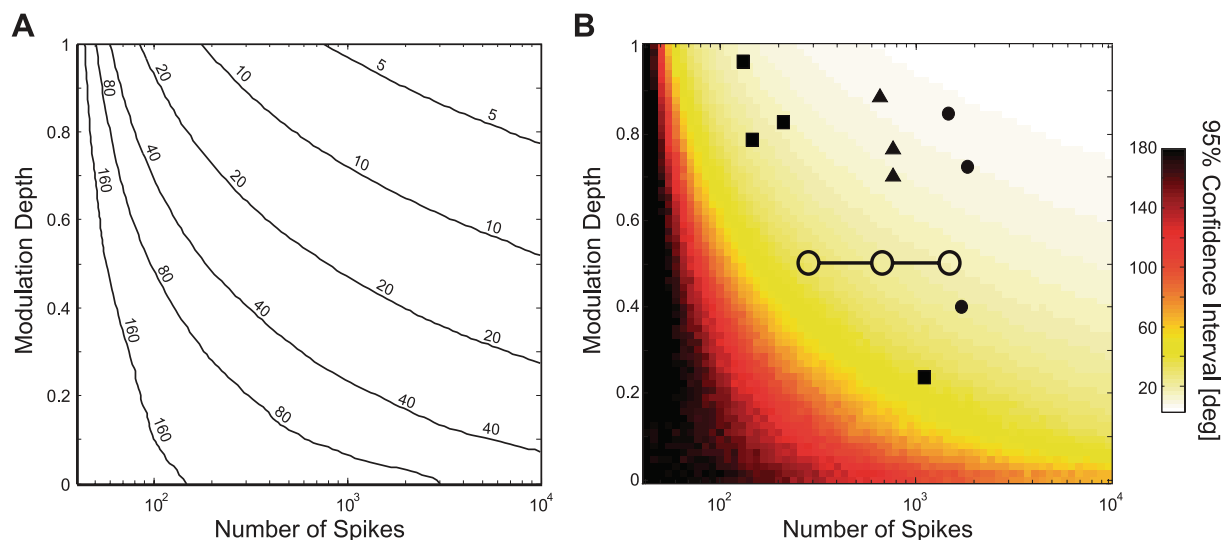


Fig. 3. Uncertainty in PD as a function of the number of spikes and modulation depth of a simulated neuron. Uncertainty about the PD decreases both with increasing number of observations and increasing modulation. **A:** contour plots illustrating the width of the 95% confidence intervals. Levels denote 1-sided 95% confidence intervals in degrees. For instance, for 0 observations or a modulation of 0, the confidence intervals span $\pm 180^\circ$. **B:** color plot of same values with representative examples from 3 published studies: ■, Rokni et al. (2007); ●, Kalaska et al. (1989); ▲, Wise et al. (1998). The average confidence intervals for the data used here are shown as open circles for blocks of 40, 120, and 240 trials (from left to right).

In all datasets, we found that the average change in PD from one block to the next was small. For a block size of 40 trials, the mean change was $-0.8 \pm 1.1^\circ$ and for blocks of 120 trials, $2.3 \pm 1.5^\circ$ for all neurons. Despite the fact that the average PD change is close to 0, there was substantial variability in the changes, which at first glance, appears to be an indication that

PDs fluctuate (Fig. 6). The SD of PD changes was $39.0 \pm 2.2^\circ$ for blocks of 40 trials and $22.9 \pm 3.2^\circ$ for blocks of 120 trials. However, based on the bootstrapping analysis, only a small number of these changes were significant (5.8% for 40 trials and 5.5% for 120 trials; bootstrapping test, $\alpha = 0.05$, not correcting for multiple comparisons). In fact, simulated Poisson neurons that were stable and matched for tuning had a distribution of changes from one block to the next, which was almost identical to the distribution of observed PD changes. These results suggest that the majority of neurons are stable on timescales on the order of minutes or if they are unstable, drift at a rate that is undetectable given the amount of measurement noise ($<43.7^\circ$ over 40 trials and $<16.3^\circ$ over 120 trials).

For the population, we can attempt to estimate the true variability in PD changes using a corrected measure of SD that attributes the observed fluctuations to a combination of real changes in PD and measurement noise (see METHODS). With the use of this method, we found that the corrected SD for blocks of 40 trials was $0.4 \pm 1.1^\circ$ and for blocks of 120 trials, was $1.8 \pm 2.5^\circ$, compatible with absolutely stable PDs. Our findings are thus consistent with the results from several past studies that incorporated enough data to reduce measurement errors (Chestek et al. 2007; Ganguly and Carmena 2009) and provide statistical insights into previous findings of apparent instabilities.

Explicitly modeling dynamic tuning curves with adaptive filtering. In the methods described above, tuning curves are assumed to be static within each block, with changes occurring only between blocks. So far, our results have suggested that any changes in PD that do exist are below the level of measurement uncertainty (16.3° over 120 trials on average). These block-by-block methods gain statistical power by incorporating more trials, but at the same time, by including so many trials, they average over any potential fluctuations that might exist. Ideally, we would like to use a method that explicitly models dynamic tuning curves, which change from trial to trial. Adaptive filtering is one approach for modeling exactly this type of relationship, and several recent studies have used adaptive

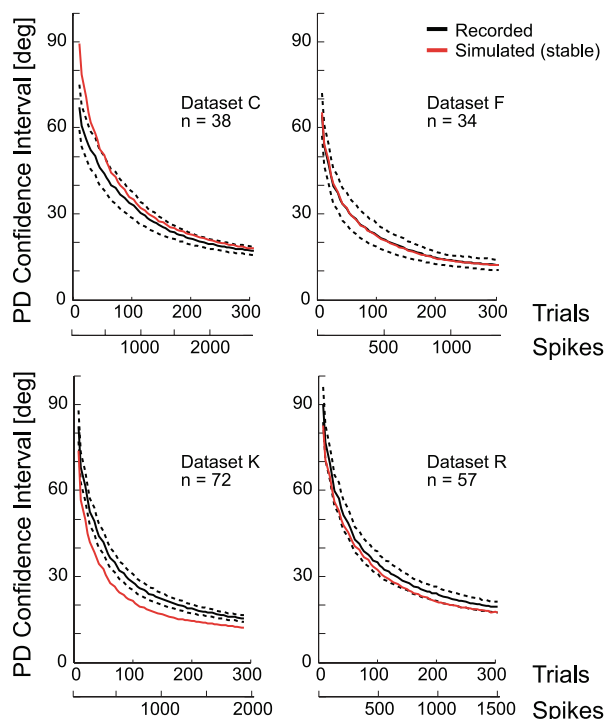


Fig. 4. Uncertainty in PD with increasing numbers of trials for 4 sets of data collected from 4 different monkeys. Average size of the 95% confidence interval for populations of actual (black) and simulated neurons (red). Simulated data were generated by estimating stable, nondrifting tuning curves from the recorded neurons and simulating spiking with Poisson noise. The average number of spikes/neuron is shown to relate these results directly to Fig. 3. Dashed lines denote \pm SEM across neurons.

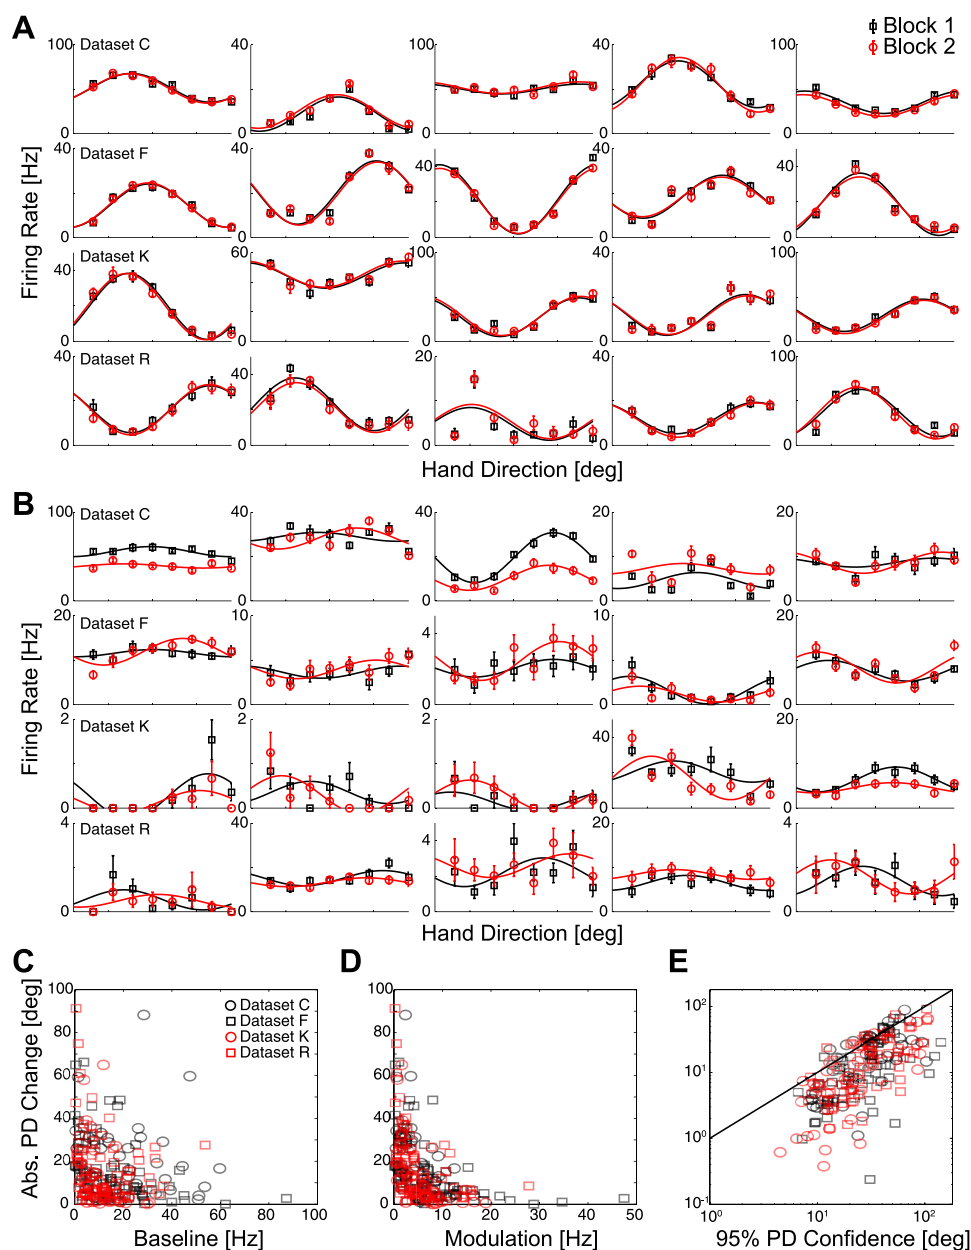


Fig. 5. Changes in tuning across blocks of 120 trials. Tuning curves for the 1st (black) and 2nd (red) blocks of 120 trials for the 5 most stable neurons (A) and the 5 least stable (B) from each of the 4 datasets. Error bars denote mean firing rate \pm SEM; solid lines denote cosine fits. The size [absolute (Abs.) value] of the PD change is negatively correlated with the baseline firing rate (C) as well as modulation (D). Higher uncertainty in the estimate of PD is correlated with larger apparent drift (E). These correlations suggest that estimated changes in PD could be due to measurement noise.

filtering to model dynamic place fields in hippocampal data (Brown et al. 2001; Eden et al. 2004). Here, we use a basic version of adaptive filtering techniques to assess dynamic tuning curves in M1.

Rather than look for differences in tuning between blocks of trials, adaptive filtering methods update the tuning curve parameters after each observation. At each time-step, the parameters are adjusted to reduce the size of prediction errors. Here, we used the LMS based on steepest-descent adaptive filtering. Larger errors result in larger parameter changes, and the size of the parameter updates is determined by a learning rate that is optimized separately. By adjusting the tuning parameters dynamically, adaptive filtering allows potential changes in tuning to be tracked from trial to trial. Here, we assumed that only the PD changes from trial to trial, whereas the baseline firing rate and modulation are constant (see METHODS for details).

Adaptive filtering is able to track changes in the PDs of simulated cosine-tuned neurons. The analysis is built on the assumption that the process of spike generation is noisy and that neurons undergo both drift (fixed changes) and diffusion or fluctuation (random changes). In simulation, cosine-tuned neurons with a constant drift in PD were accurately tracked for a range of drift rates (Fig. 7A). For a drift rate in the range $\pm 2^\circ/\text{trial}$ and no fluctuations, the estimated drift rate had a root-mean-square (RMS) error of only $0.07^\circ/\text{trial}$ (Fig. 7B). Fluctuations in PD could also be tracked, albeit less accurately (Fig. 7C). For fluctuations with a SD between 0 and $20^\circ/\text{trial}$, the estimated fluctuation had a RMS error of $2.5^\circ/\text{trial}$ (Fig. 7D). It is important to note that while adaptive filtering excels at tracking large changes in the parameters, estimates of small fluctuations are substantially less accurate and sensitive to the learning rate. With a RMS error of $2.5^\circ/\text{trial}$, it will be difficult to detect the $0.16^\circ/\text{trial}$ variation ($1.8 \pm 2.5^\circ$ over 120 trials), sug-

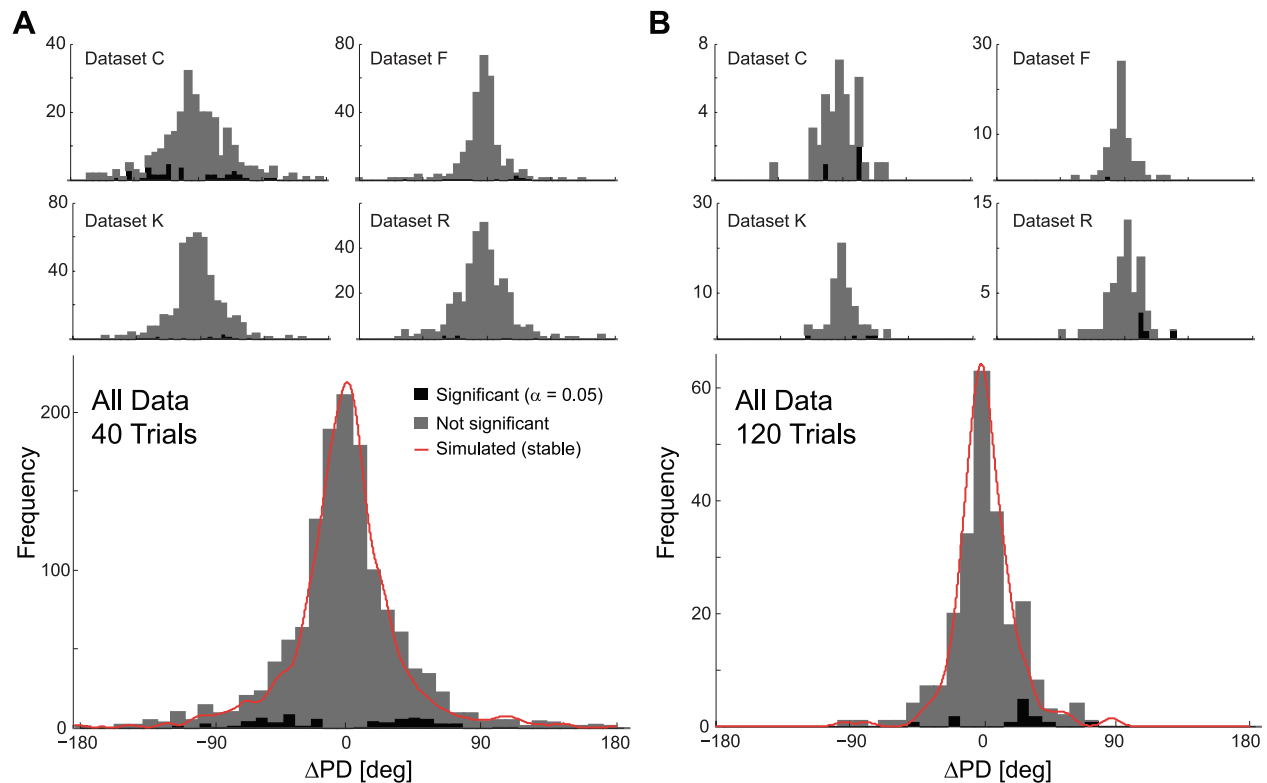


Fig. 6. Estimated changes in PD. A: histograms of PD changes for each of the datasets separately and for all neurons for a block size of 40 trials. B: PD changes for a block size of 120 trials. Black blocks denote changes that were significant at the 95% level by bootstrapping. Note that for blocks of 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. Red curves denote the PD changes observed when simulating stable Poisson neurons that were matched to have the same tuning as the observed data.

gested by the corrected SD analysis above. However, these simulation results suggest that if neurons in M1 are truly unstable with large, random fluctuations in PD, adaptive filtering methods may reveal this instability.

In actual data, adaptive filtering revealed some degree of fluctuation in the PDs of a subset of M1 neurons. Of the 201 significantly tuned neurons from our four datasets, 177 neurons were best fit by a stable rather than dynamic tuning curve. The remaining 24 neurons ($\sim 12\%$) showed some degree of fluctua-

tion, with a median fluctuation of $2.1 \pm 0.3^\circ/\text{trial}$ (Fig. 8A). Assuming that PDs follow a Gaussian random walk, results from Rokni et al. (2007), as well as the uncorrected block-by-block analysis (Fig. 6), suggest fluctuations on the order of 2.1° – $2.2^\circ/\text{trial}$ on average [$22.9 \pm 3.2^\circ$ over 120 trials for our analysis and $29 \pm 3^\circ$ over 160 trials (Rokni et al. 2007)]. At first glance, this subset of neurons thus appears to be fluctuating consistent with previous results; however, even in this minority of neurons, it is possible that the observed fluctuations are false positives.

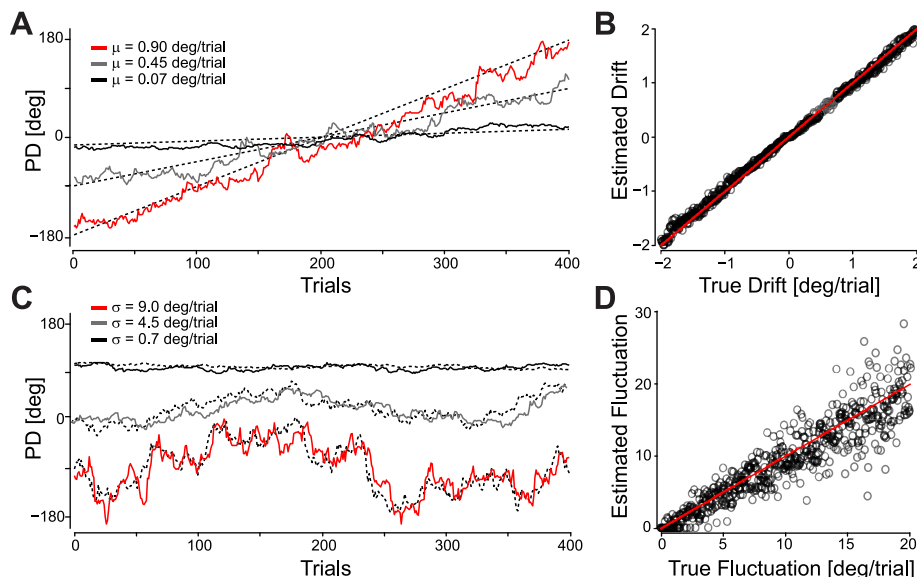


Fig. 7. Adaptive filtering for detecting drift and fluctuation. A: true and estimated PD for 3 simulated neurons with PDs drifting 0.9, 0.45, and 0.07 $^\circ/\text{trial}$. Dashed lines denote the true, underlying PD. B: estimated PD drift as a function of the true drift for 512 simulation runs using data from 400 trials. C: true and estimated PD for 3 simulated neurons with no mean drift, fluctuating 9, 4.5, and 0.7 $^\circ/\text{trial}$. D: estimated PD fluctuation as a function of the true fluctuation for 512 simulation runs. For data from 400 trials, the degree of fluctuation is estimated fairly well but much less accurately than the mean drift.

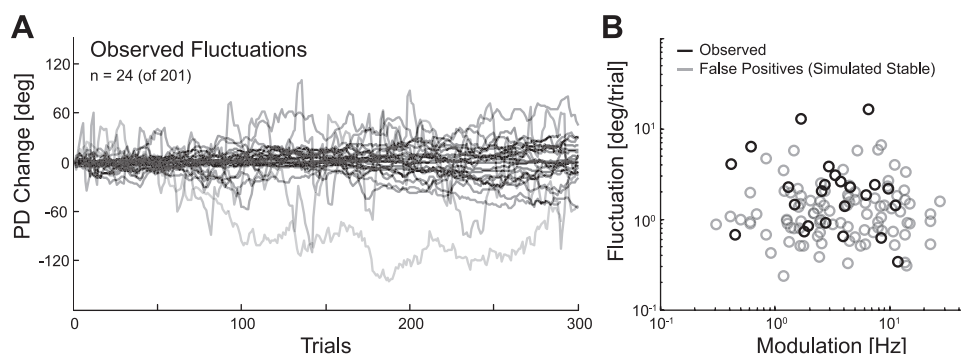


Fig. 8. Adaptive filtering for actual data. For this analysis, we combined neurons from all 4 datasets (C, F, K, and R) and use only the significantly tuned neurons ($n = 201$). Only 12% ($n = 24$) of these neurons showed some degree of fluctuation. *A*: trial-by-trial PD changes estimated by adaptive filtering for the subset of neurons that appeared to be fluctuating. Note that 177 of the 201 neurons were better fit by a stable rather than a dynamic tuning curve. *B*: estimated fluctuation as a function of the modulation for the subset of fluctuating neurons (black) and simulated, stable neurons (gray). The fluctuations revealed by adaptive filtering are consistent with false positives.

The fact that these estimates are close to the RMS error observed when estimating fluctuations in simulated data suggests that some care is needed in interpreting these values. The fluctuations found by adaptive filtering did improve spike prediction. The average log-likelihood ratio for the dynamic tuning model was 8.66 ± 1.6 bits/trial, relative to a homogeneous Poisson process, whereas the average log-likelihood ratio for the static tuning model was 8.56 ± 1.6 bits/trial, relative to a homogeneous Poisson process (both on training data). However, with any finite amount of data, adaptive filtering methods can find small false-positive fluctuations even for simulated, stable neurons.

We examined whether the observed fluctuations may have been false positives by again simulating stable cosine-tuned neurons of varying modulation depth. In this case, after applying adaptive filtering, $\sim 4\text{--}5\%$ of the stable neurons are mistakenly identified as fluctuating. With the use of these false positives, we constructed a null distribution and compared this distribution with the subset of observed fluctuating neurons (Fig. 8*B*). The fluctuations observed in 12% of the recorded neurons are consistent with the results from stable neurons, suggesting that they may indeed be false positives. The null results suggest that for typical physiologically realistic tuning curve parameters, fluctuations have to be rather large before they are detectable by adaptive filtering ($\sim 5^\circ/\text{trial}$).

DISCUSSION

We have presented results from two approaches aimed at detecting fluctuations in the PDs of cosine-tuned neurons in M1. With the use of bootstrapping on simulations of stable neurons, we have quantified how measurement uncertainty or confidence interval size is affected by modulation depth and the amount of data available. In real data, bootstrapping allows us to estimate changes between blocks of trials and test for significance in a way that directly captures the measurement uncertainty. Finally, we used adaptive filtering techniques to model trial-by-trial changes in PD explicitly. After comparing the results with those from simulated, stable neurons, we find no evidence for large fluctuations in PD using either bootstrapping or adaptive filtering. Small fluctuations in PD may exist, but detecting these changes is difficult in the presence of spiking noise.

When examining the properties of neural discharge, the experimental manipulations, as well as the statistical ap-

proaches used for analysis, will influence the interpretation of the results. In this report, we have demonstrated that the uncertainty related to estimates of PD is directly related to the total number of spikes, as well as the neuron's modulation depth over different reach directions. The greater the modulation depth and the more data available (number of spikes), the more certain one can be about the tuning of an individual neuron. For typical experiments with limited amounts of data, the uncertainty about PD can be quite large compared with the typical effects of a manipulation—both on the order of a few tens of degrees. Measurement noise is thus of central importance when estimating how tuning curves change over time.

With the use of bootstrapping, we can determine confidence intervals for each of the tuning curve parameters. For the observed neurons, the confidence intervals for PD were rather large— 16.3° over 120 trials on average. This value fundamentally limits how well we can detect changes in PD. For instance, the SD of changes in PD between blocks of 120 trials was $22.9 \pm 3.2^\circ$ —the same order of magnitude as the confidence interval. Comparing these changes with those from stable, simulated neurons with matched tuning curves and using a corrected measure, we find that the SD of changes is likely closer to $1.8 \pm 2.5^\circ$ over 120 trials. Adaptive filtering estimates PD changes in a small subset of neurons on the order of $2.1^\circ/\text{trial}$, but after comparing these results with stable, simulated neurons, we find that these changes are again consistent with stable PDs.

A previous study using large blocks of data reported similar stability of tuning parameters (Chestek et al. 2007). However, this study used free-reaching rather than a manipulandum and analyzed neurons primarily from the dorsal premotor cortex (PMd), leading the authors to speculate that the differences between their results and those of Rokni et al. (2007) could be due to experimental design or the specific population of neurons. The experiments presented here are far closer in design to those by Rokni et al. (2007), and yet, we find no significant PD changes after corrections for measurement noise. Several studies have examined the stability of neural activity in the context of brain-machine interfaces and observe that firing properties may be highly variable (Carmena et al. 2005) or relatively stable (Ganguly and Carmena 2009) depending on training. However, interpreting these results can be difficult, since it is often unclear how a given decoding scheme relates to the tuning properties of individual neurons. Variability in the

decoded behavior tends to be much higher than during normal reaching, and brain-control itself may introduce certain constraints on movement coding.

As a practical matter in data analysis, anything that is not part of the model but affects neural signals is considered noise. We considered only Poisson noise, which is compatible with the Fano factor, typically observed in cortical recordings (Stevens and Zador 1998; Zacksenhouse et al. 2007). However, there are many other factors beyond the monkey's hand direction that potentially influence firing rates (Johnson et al. 2001), including other kinematic variables such as reach speed (Chestek et al. 2007) and limb posture (Caminiti et al. 1990; Scott and Kalaska 1997), as well as added loads (Kalaska et al. 1989), cortical waves (Rubino et al. 2006), neuromodulator concentrations (Ahern et al. 2002), oxygen concentration (Jiang and Haddad 1994), and circadian rhythms (Barnes et al. 1977). Modeling these other sources of variability, whether they are observed (Paninski et al. 2003; Saleh et al. 2010; Truccolo et al. 2005; Wu and Hatsopoulos 2006) or unobserved (Kulkarni and Paninski 2007; Stevenson et al. 2010) would likely improve the estimates of PD stability. These uncontrolled sources of variability could inflate estimates of both real fluctuations and measurement uncertainty, and the results here thus provide only an upper bound for the instability of PDs of neurons in M1. If other sources of measurement noise could be accounted for, the estimated changes in PD may very well be even smaller.

Whereas our results suggest that PDs are substantially more stable during normal reaching than some previous reports, there is also convincing evidence that the relationship between a given neuron's activity and hand direction does change over time. Tuning to hand direction changes on very short timescales due to changing kinematics and dynamics (Churchland and Shenoy 2007; Hatsopoulos et al. 2004; Sergio et al. 2005), as well as over longer timescales during sensorimotor learning (Jarosiewicz et al. 2008; Li et al. 2001; Paz and Vaadia 2004). While several studies have shown that tuning curves are sensitive to the measurement epoch and specific task constraints (Hamel-Pâquet et al. 2006; Sergio et al. 2005), here, we focused on a fixed, specific portion of the reaches (100 ms prior to, through 300 ms after movement onset). This type of analysis ignores the short-timescale kinematics and dynamics of reaching and is aimed to test whether tuning curves are stable on longer timescales during a well-learned task. Additionally, we have focused primarily on the stability and uncertainty in estimates of PD. Both the modulation and baseline firing rate of cosine-tuned neurons may show a higher degree of instability (Chestek et al. 2007).

The stability of neuronal properties is of central importance to many computational theories. If presynaptic neurons change rapidly, then the motor system must either be redundant to the extent that fluctuations do not affect behavior (Rokni et al. 2007), or postsynaptic neurons must adapt to allow for stable movement and decisionmaking. Many behavioral models (e.g., Cheng and Sabes 2007; Wei and Körding 2009) have suggested that learning is an ongoing process, where errors are constantly being corrected, even during apparently stable behavior. This might suggest that the fluctuations in tuning, if they exist, may actually be functional. Rather than being an artifact of redundancy in the cortical representation of movement, fluctuations may be a reflection of ongoing attempts to

correct small-reaching errors. While measurement noise makes it difficult to distinguish between these hypotheses, methods are being developed to isolate the effects of reach errors (Chase et al. 2010; Scheidt et al. 2000) and better understand redundancy in the motor system (Jarosiewicz et al. 2008). Statistical techniques that allow modeling of nonstationary data (Kim et al. 2006; Wu and Hatsopoulos 2008), as well as experimental techniques that allow for longer-term recordings (Dickey et al. 2009; Tolia et al. 2007), should both serve to reveal the more-detailed structure of tuning curve dynamics during learning as well as stable reaching.

ACKNOWLEDGMENTS

Present address of J. Reimer: Department of Neuroscience, Baylor College of Medicine, One Baylor Plaza, S603, Houston, TX 77030.

GRANTS

I. H. Stevenson and K. P. Körding were supported by the Chicago Community Trust and by National Institutes of Health Grants 1R01-NS-063399 and 2P01-NS-044393. A. Cheria, B. M. London, and L. E. Miller were supported by National Institutes of Health Grant R01-NS-048845. N. A. Sachs was supported by the National Center for Research Resources Clinical and Translational Science Awards Program (CTSA), part of the Roadmap Initiative, "Re-Engineering the Clinical Research Enterprise", UL1RR0254741. E. Lindberg and M. W. Slutzky were supported by National Institutes of Health Grant K08-NS-060223. J. Reimer and N. G. Hatsopoulos were supported by National Institute of Neurological Disorders and Stroke Grant R01NS45853.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author(s).

REFERENCES

- Ahern GP, Klyachko VA, Jackson MB. cGMP and S-nitrosylation: two routes for modulation of neuronal excitability by NO. *Trends Neurosci* 25: 510–517, 2002.
- Amirikian B, Georgopoulos AP. Directional tuning profiles of motor cortical cells. *Neurosci Res* 36: 73–79, 2000.
- Barnes CA, McNaughton BL, Goddard GV, Douglas RM, Adamec R. Circadian rhythm of synaptic excitability in rat and monkey central nervous system. *Science* 197: 91–92, 1977.
- Berens P. CircStat: a MATLAB toolbox for circular statistics. *J Stat Softw* 31: 1–21, 2009.
- Brown E, Nguyen D, Frank L, Wilson M, Solo V. An analysis of neural receptive field plasticity by point process adaptive filtering. *Proc Natl Acad Sci USA* 98: 12261–12266, 2001.
- Caminiti R, Johnson PB, Urbano A. Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. *J Neurosci* 10: 2039–2058, 1990.
- Carmena J, Lebedev M, Henriquez C, Nicolelis M. Stable ensemble performance with single-neuron variability during reaching movements in primates. *J Neurosci* 25: 10712–10716, 2005.
- Chase S, Schwartz A, Kass R. Latent inputs improve estimates of neural encoding in motor cortex. *J Neurosci* 30: 13873–13882, 2010.
- Cheng S, Sabes P. Calibration of visually guided reaching is driven by error-corrective learning and internal dynamics. *J Neurophysiol* 97: 3057–3069, 2007.
- Chestek CA, Batista AP, Santhanam G, Yu BM, Afshar A, Cunningham JP, Gilja V, Ryu SI, Churchland MM, Shenoy KV. Single-neuron stability during repeated reaching in macaque premotor cortex. *J Neurosci* 27: 10742–10750, 2007.
- Churchland M, Shenoy K. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J Neurophysiol* 97: 4235–4257, 2007.
- Cronin B, Stevenson IH, Sur M, Körding KP. Hierarchical Bayesian modeling and Markov chain Monte Carlo sampling for tuning-curve analysis. *J Neurophysiol* 103: 591–602, 2010.

- Davison A, Hinkley D.** *Reviews of Bootstrap Methods and Their Application.* Cambridge, MA: Cambridge University Press, 1997.
- Dickey A, Suminski A, Amit Y, Hatsopoulos N.** Single-unit stability using chronically implanted multielectrode arrays. *J Neurophysiol* 102: 1331–1339, 2009.
- Eden U, Frank L, Barbieri R, Solo V, Brown E.** Dynamic analysis of neural encoding by point process adaptive filtering. *Neural Comput* 16: 971–998, 2004.
- Efron B, Tibshirani R.** *An Introduction to the Bootstrap.* New York: Chapman & Hall, 1997.
- Fisher N.** *Statistical Analysis of Circular Data.* Cambridge, MA: Cambridge University Press, 1996.
- Ganguly K, Carmena JM.** Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol* 7: e1000153, 2009.
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT.** On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2: 1527–1537, 1982.
- Georgopoulos AP, Kettner RE, Schwartz AB.** Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *J Neurosci* 8: 2928–2937, 1988.
- Georgopoulos AP, Schwartz AB, Kettner RE.** Neuronal population coding of movement direction. *Science* 233: 1416–1419, 1986.
- Hamel-Pâquet C, Sergio LE, Kalaska JF.** Parietal area 5 activity does not reflect the differential time-course of motor output kinetics during arm-reaching and isometric-force tasks. *J Neurophysiol* 95: 3353–3370, 2006.
- Hatsopoulos N, Joshi J, O'Leary JG.** Decoding continuous and discrete motor behaviors using motor and premotor cortical ensembles. *J Neurophysiol* 92: 1165–1174, 2004.
- Hatsopoulos NG, Xu Q, Amit Y.** Encoding of movement fragments in the motor cortex. *J Neurosci* 27: 5105–5114, 2007.
- Haykin S.** *Adaptive Filter Theory.* Englewood Cliffs, NJ: Prentice Hall, 1996.
- Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB.** Functional network reorganization during learning in a brain-computer interface paradigm. *Proc Natl Acad Sci USA* 105: 19486–19491, 2008.
- Jiang C, Haddad GG.** A direct mechanism for sensing low oxygen levels by central neurons. *Proc Natl Acad Sci USA* 91: 7198–7201, 1994.
- Johnson MTV, Mason CR, Ebner TJ.** Central processes for the multiparametric control of arm movements in primates. *Curr Opin Neurobiol* 11: 684–688, 2001.
- Kakei S, Hoffman DS, Strick PL.** Muscle and movement representations in the primary motor cortex. *Science* 285: 2136–2139, 1999.
- Kalaska JF, Cohen DAD, Hyde ML, Prud'homme M.** A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. *J Neurosci* 9: 2080–2102, 1989.
- Kalaska JF, Hyde ML.** Area 4 and area 5: differences between the load direction-discharge variability of cells during active postural fixation. *Exp Brain Res* 59: 197–202, 1985.
- Kim S-P, Wood F, Fellows M, Donoghue JP, Black MJ.** Statistical analysis of the non-stationarity of neural population codes. *IEEE* 811–816, 2006.
- Kulkarni JE, Paninski L.** Common-input models for multiple neural spike-train data. *Network* 18: 375–407, 2007.
- Li CSR, Padoa-Schioppa C, Bizzi E.** Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* 30: 593–607, 2001.
- Morrow MM, Jordan LR, Miller LE.** Direct comparison of the task-dependent discharge of M1 in hand space and muscle space. *J Neurophysiol* 97: 1786–1798, 2007.
- Nordhausen CT, Maynard EM, Normann RA.** Single unit recording capabilities of a 100 microelectrode array. *Brain Res* 726: 129–140, 1996.
- Otieno B, Anderson-Cook C.** *On Bootstrap Confidence Interval Estimation of Preferred Direction for Circular Data* (Paper 805). Allendale, MI: Grand Valley State University, Dept. of Statistics, 2008.
- Paninski L, Fellows MR, Hatsopoulos NG, Donoghue JP.** Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *J Neurophysiol* 91: 515–532, 2003.
- Paz R, Vaadia E.** Learning-induced improvement in encoding and decoding of specific movement directions by neurons in the primary motor cortex. *PLoS Biol* 2: E45, 2004.
- Pohlmeier EA, Solla SA, Perreault EJ, Miller LE.** Prediction of upper limb muscle activity from motor cortical discharge during reaching. *J Neural Eng* 4: 369–379, 2007.
- Rokni U, Richardson AG, Bizzi E, Seung HS.** Motor learning with unstable neural representations. *Neuron* 54: 653–666, 2007.
- Rubino D, Robbins KA, Hatsopoulos NG.** Propagating waves mediate information transfer in the motor cortex. *Nat Neurosci* 9: 1549–1557, 2006.
- Saleh M, Takahashi K, Amit Y, Hatsopoulos NG.** Encoding of coordinated grasp trajectories in primary motor cortex. *J Neurosci* 30: 17079–17090, 2010.
- Scheidt R, Reinkensmeyer D, Conditt M, Rymer W, Mussa-Ivaldi F.** Persistence of motor adaptation during constrained, multi-joint, arm movements. *J Neurophysiol* 84: 853–862, 2000.
- Schummers J, Cronin B, Wimmer K, Stimberg M, Martin R, Obermayer K, Koerding K, Sur M.** Dynamics of orientation tuning in cat v1 neurons depend on location within layers and orientation maps. *Front Neurosci* 1: 145–159, 2007.
- Scott SH.** Apparatus for measuring and perturbing shoulder and elbow joint positions and torques during reaching. *J Neurosci Methods* 89: 119–127, 1999.
- Scott SH, Kalaska JF.** Reaching movements with similar hand paths but different arm orientations. I. Activity of individual cells in motor cortex. *J Neurophysiol* 77: 826–852, 1997.
- Sergio L, Hamel-Paquet C, Kalaska J.** Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. *J Neurophysiol* 94: 2353–2578, 2005.
- Stevens CF, Zador AM.** Input synchrony and the irregular firing of cortical neurons. *Nat Neurosci* 1: 210–217, 1998.
- Stevenson I, Cronin B, Sur M, Kording K.** Sensory adaptation and short term plasticity as Bayesian correction for a changing brain. *PLoS One* 5: e12436, 2010.
- Swindale NV.** Orientation tuning curves: empirical description and estimation of parameters. *Biol Cybern* 78: 45–56, 1998.
- Tollas A, Ecker A, Siapas A, Hoenselaar A, Keliris G, Logothetis N.** Recording chronically from the same neurons in awake, behaving primates. *J Neurophysiol* 98: 3780–3790, 2007.
- Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN.** A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol* 93: 1074–1089, 2005.
- Wei K, Körding K.** Relevance of error: what drives motor adaptation? *J Neurophysiol* 101: 655–664, 2009.
- Wise SP, Moody SL, Blomstrom KJ, Mitz AR.** Changes in motor cortical activity during visuomotor adaptation. *Exp Brain Res* 121: 285–299, 1998.
- Wu W, Hatsopoulos N.** Evidence against a single coordinate system representation in the motor cortex. *Exp Brain Res* 175: 197–210, 2006.
- Wu W, Hatsopoulos N.** Real-time decoding of nonstationary neural activity in motor cortex. *IEEE Trans Neural Syst Rehabil Eng* 16: 213–222, 2008.
- Zacksenhouse M, Lebedev MA, Carmena JM, O'Doherty JE, Henriquez C, Nicolelis MAL.** Cortical modulations increase in early sessions with brain-machine interface. *PLoS One* 2: e619, 2007.