

UCSF

UC San Francisco Previously Published Works

Title

Motor Variability Arises from a Slow Random Walk in Neural State

Permalink

<https://escholarship.org/uc/item/37959685>

Journal

Journal of Neuroscience, 34(36)

ISSN

0270-6474

Authors

Chaisanguanthum, Kris S
Shen, Helen H
Sabes, Philip N

Publication Date

2014-09-03

DOI

10.1523/jneurosci.3001-13.2014

Peer reviewed

Motor Variability Arises from a Slow Random Walk in Neural State

Kris S. Chaisanguanthum,^{1,2} Helen H. Shen,¹ and Philip N. Sabes^{1,2}

¹Center for Integrative Neuroscience and Department of Physiology and ²Sloan-Swartz Center for Theoretical Neurobiology, University of California, San Francisco, San Francisco, California 94143

Even well practiced movements cannot be repeated without variability. This variability is thought to reflect “noise” in movement preparation or execution. However, we show that, for both professional baseball pitchers and macaque monkeys making reaching movements, motor variability can be decomposed into two statistical components, a slowly drifting mean and fast trial-by-trial fluctuations about the mean. The preparatory activity of dorsal premotor cortex/primary motor cortex neurons in monkey exhibits similar statistics. Although the neural and behavioral drifts appear to be correlated, neural activity does not account for trial-by-trial fluctuations in movement, which must arise elsewhere, likely downstream. The statistics of this drift are well modeled by a double-exponential autocorrelation function, with time constants similar across the neural and behavioral drifts in two monkeys, as well as the drifts observed in baseball pitching. These time constants can be explained by an error-corrective learning processes and agree with learning rates measured directly in previous experiments. Together, these results suggest that the central contributions to movement variability are not simply trial-by-trial fluctuations but are rather the result of longer-timescale processes that may arise from motor learning.

Introduction

Although professional athletes, such as 2013 Major League Baseball (MLB) National League Cy Young winner Clayton Kershaw, can perform highly trained movements with remarkable precision, they still exhibit movement variability (Fig. 1*a*). This inevitable variability is thought to reflect a fundamental limit on movement precision attributable to sensorimotor “noise.” This noise could arise at any stage of sensorimotor processing, from perception of the spatial cues for movement to motor preparation and execution (Harris and Wolpert, 1998; Todorov and Jordan, 2002; Osborne et al., 2005; Churchland et al., 2006; Faisal et al., 2008). Here we focus on cortical contributions to movement variability, examining the patterns of activity in populations of neurons in macaque dorsal premotor cortex (PMd) and primary motor cortex (M1) recorded while animals performed well practiced reaching movements. These brain areas represent the output of the cortical circuits for reach planning and execution, and neural activity in these areas is known to account for a significant amount of motor variation (Churchland et al., 2006). However, the character of this centrally generated motor variation, or indeed motor variation as a whole, is not well understood.

Here, we show that central contributions to movement variability do not resemble simple sensory or neural noise but may

instead reflect slowly drifting neural states. We make this argument through a series of analyses. First, we show that, in macaque monkeys making center-out reach movements—as well as in professional baseball pitchers—a substantial fraction of the movement-by-movement variability is attributable to a slow drift in the mean movement. The preparatory activity of PMd/M1 neurons in monkey exhibits an analogous slow drift in firing rates over time. Thus, variability in both behavior and neurons can be decomposed into two statistical components: (1) the slow drift; and (2) fast trial-by-trial fluctuations about the drifting mean. Because preparatory neural activity predicts a significant amount of movement variability, at least one of the following must be true: (1) the fast fluctuations in the neural activity are related to the fast fluctuations in behavior; and/or (2) the drifts in the neural activity and in behavior are related. The second set of analyses tests the former hypothesis and shows the surprising result that there is little correlation between the fast fluctuations in central neural activity and behavior. Therefore, the fast trial-by-trial variability must arise primarily downstream of PMd/M1 or otherwise independently. The third set of analyses tests the latter hypothesis and shows that drift in neural activity is related with drift in behavior: we found that the magnitude of the neural drifts correlate with the magnitude of the behavioral drift across sessions. We also found that both neural and behavioral drifts are well described by similar autocorrelation time constants, which is additional evidence that they are of common origin. Finally, we present evidence that the drift arises from an underlying error-corrective learning process that is a fundamental feature of practiced movements in primates.

Materials and Methods

MLB data. Baseball pitching data were taken from the PITCH/x database, created and maintained by Sportvision and publicly available

Received July 15, 2013; revised June 10, 2014; accepted July 25, 2014.

Author contributions: K.S.C., H.H.S., and P.N.S. designed research; K.S.C., H.H.S., and P.N.S. performed research; K.S.C., H.H.S., and P.N.S. analyzed data; K.S.C. and P.N.S. wrote the paper.

This research was supported by National Institutes of Health (NIH)/National Institute of Mental Health Grant P50-MH-077970, NIH/National Eye Institute Grant EY-015679, Defense Advanced Research Projects Agency Reorganization and Plasticity to Accelerate Injury Recovery Grant N66001-10-C-2010, and the Swartz Foundation.

Correspondence should be addressed to Dr. Philip N. Sabes, University of California, San Francisco, 675 Nelson Rising Lane, #514G, San Francisco, CA 94143-0444. E-mail: sabes@phy.ucsf.edu.

DOI:10.1523/JNEUROSCI.3001-13.2014

Copyright © 2014 the authors 0270-6474/14/3412071-10\$15.00/0




Figure 1. Examples of movement variability. *a*, Spread of release points for all four-seam fastball pitches thrown by Los Angeles Dodgers pitcher Clayton Kershaw versus the New York Mets on May 8, 2011. The location is defined in the frontal plane of the pitcher but in coordinates fixed relative to the field. Baseball circumference shown for scale. *b*, Reach paths from all tuning block trials in a single experimental session (monkey E, color coded by target).

through MLB. PITCHf/x measures the release point and trajectory of pitches via video tracking using a multicamera setup in each MLB stadium. More detailed descriptions of the dataset have been described previously (Albert, 2010). PITCHf/x release point has an SD of <1.3 cm (Sportvision).

We analyzed release-point location and initial speed. The release point is measured as the transverse location of the baseball when it crosses the plane 50 feet (15.2 m) away from home plate. Because home plate is 60.5 feet (18.4 m) away from the pitching rubber and pitches are typically released with the pitcher near full forward extension, this measured location is a good proxy for where the pitcher released the ball. We chose to analyze release point because it is the endpoint of the trajectory of the ball that is directly under the pitcher's control; for example, it is not confounded by ball spin. We also focused our analysis on a single pitch type, four-seam fastballs, a common pitch thrown with high velocity, because these pitches follow the straightest paths. Nonetheless, release-point drift is apparent across all pitch types.

Note that not all of the variability in pitcher release point need be "sensorimotor," i.e., some of it may, in fact, be intentional, reflecting strategy of play. However, such intentional variability would not be expected to drift on long timescales. Thus, we expect that, although such intentional motor variability might affect the relative contributions of the two distinct types of variability, it should not affect the character of the drift.

Macaque reaching experiments. All procedures were approved by the University of California, San Francisco Institutional Animal Care and Use Committee and followed National Institutes of Health Guide for the Care and Use of Laboratory Animals. Data were collected in 12 and 14 experimental sessions for monkeys D and E, both male, respectively. Sessions consisted of a series of trials in which the animals made reaching movements on a horizontal surface, located just below shoulder height. Direct view of the hand and arm were blocked, with visual targets and visual feedback of the hand displayed via a mirror and rear projection screen such that objects appeared coplanar with the reaching hand (McGuire and Sabes, 2011). Hand position was measured at the location of the index fingertip using an electromagnetic tracking device (Liberty; Polhemus), and visual feedback was given as a small disk of light at that location. Eye position was also monitored using an infrared eye tracker (Iscan).

Each trial began with the appearance of a circular start target at the initial position, which was located sagittally along the midline and was fixed throughout the experiment. Animals then moved their hand to the initial position with visual feedback. After a hold period of 300 ms, a reach target appeared. There were eight potential target locations, arrayed evenly around a circle, 7 cm from the initial position ("center-out" reaches). Animals were then required to visually fixate the target, and fixation was enforced for the remainder of the trial. After 300 ms of visual fixation, a variable delay began, with duration sampled uniformly between 500 and 1000 ms. An audible "go" cue was then played, and the

start target and feedback disappeared, cuing the animal to move to the reach target. Reaches were made in the absence of visual feedback. The trial ended when the arm came to a stop and remained stationary for 200 ms. Visual feedback was then turned back on, and fluid reward was delivered in a graded manner based on final endpoint error: reward fell linearly with the distance from the target, and trials with endpoint errors >2.5 cm were not rewarded. The average endpoint error was ~5 mm from the center of the target.

Extracellular signals were recorded from the PMd of both animals and from M1 of monkey E using chronically implanted 96-microelectrode arrays (Blackrock Microsystems). Single units were identified by spike sorting with Plexon Offline Sorter, using a combination of automatic (*T*-Dist E-M algorithm) and manual processing. Identified units had an average signal-to-noise ratio of 5.3 for monkey D and 3.9 for monkey E. Mean firing rates were calculated during the instructed delay period (between target presentation and the go cue), which was of variable duration, between 800 and 1300 ms. We have also repeated the analyses presented here using firing rates calculated from different epochs of the trial, both before and after the delay period, as well from subsets of the delay period; the results presented here are not sensitive to the choice of time window.

There were two types of trial blocks, which alternated throughout the experimental session: (1) 40-trial "tuning" blocks (five trials randomly interleaved to each of the eight targets); and (2) 210-trial "three-target" blocks (70 trials randomly interleaved to each of three targets in a quadrant, e.g., at 0°, 45°, and 90°; quadrant was fixed within a session but randomized across sessions). For this study, the only purpose of the tuning block data is to confirm that the responses of the recorded neurons are indeed tuned to the location of the reach target cue. Using the tuning block data, many recorded units exhibit "cosine tuning" in target angle (Georgopoulos et al., 1981), and the location of the reach target can be decoded from the vector of delay-period firing rates of the neurons, with a cross-validated R^2 of 84.6% (averaged across all experimental sessions). Hereafter, the data we present are from the three-target blocks, which, for each session, consist of several hundred reaches to each target. The large number of repeated reaches allowed us to observe and measure variability in neural activity and behavior, including slow changes over time.

Optimal linear estimators. We are interested in relating variations in neural activity to variations in behavior, e.g., the speed or direction of a reaching movement. This is often done with "population decoders" that estimate the behavioral variable from recorded neural activity using knowledge about the tuning curve of each individual neuron (and possibly some model of neural noise; Georgopoulos et al., 1986). We could estimate such tuning curves (e.g., rate vs. movement direction) using data from the tuning blocks. However, because the fine structure of the tuning curves may not match simple parametric models such as cosine tuning and because even cosine tuning curves are difficult to measure with precision (Stevenson et al., 2011), the results of this approach would be too biased and/or noisy to relate the relatively small variations in neural activity and behavior that are observed in repeated reaches to a nominally identical target.

Instead, we designed the three-target block structure: with a large series of reaches to a single target, we can determine the relationship between neural variation and behavioral variation in a way that is agnostic of the tuning properties of the neurons. For a set of reaches to a single target, we first subtracted means from neural activity and behavioral metrics. This left only the target-independent variability in the system, allowing us to determine the relationship between the neural and behavioral variability in a data-driven manner: we found the linear estimator that best predicted each behavioral metric from k leading principal components of the neural firing vectors. The number of principal components k for each was optimized independently for each session and target.

More explicitly, for each k (between unity and the number of units recorded), we projected the firing rate vectors onto their k leading principal components. With 60% of the trials (randomized, but the same for all k), we computed optimal linear estimator coefficients by regressing a given movement metric onto the k of principal components of the firing rates. Half of the remaining trials (20%) were used as a cross-validation

set: we used the performance of the estimator on these trials to determine the k for which the corresponding optimal linear estimator was the best predictor of the behavior. Once the optimal k was found, the performance of the optimal linear estimator was measured on the remaining 20% of the trials and reported as a coefficient of determination.

Finally, we note that this aggressive optimization and cross-validation scheme is intended to be conservative. Because we will use this approach to show the lack of predictive power between some neural and behavioral time series, we wanted to be sure that we optimized the estimation process.

Multistate learning model. We considered a linear dynamic learning model similar to that used by Cheng and Sabes (2006) but with a two-dimensional state for each scalar observable, i.e., a second-order system to account for multiple timescales of learning (Smith et al., 2006). The observable y_t (behavioral metric or neural firing rate on trial t) is related to the internal state x_t by a linear transformation, parameterized by the “emission” vector \vec{E} :

$$y_t = \vec{E}^T \vec{x}_t + r_t, \quad (1)$$

where the emissions noise r_t is independent and normally distributed with 0 mean and variance R . The state update, including error-corrective learning driven by feedback of the observable, can be written as

$$\vec{x}_{t+1} = \mathbf{M}\vec{x}_t + \vec{B}y_t + \vec{w}_t, \quad (2)$$

where the parameters B and M describe the learning and retention rates of the system, respectively. w_t represents noise in the learning process and is independent and normally distributed with 0 mean and covariance $\mathbf{W} \equiv \langle \vec{w}_t \vec{w}_t^T \rangle$, where angle brackets are used to denote expectation value, averaged over trials.

In the steady-state regimen of this system—where the statistical description of the system, i.e., quantities such as $\langle \vec{x}_t \vec{x}_t^T \rangle$ and $\langle y_t y_t^T \rangle$, no longer evolves with t —it is straightforward to show that, for $\tau \geq 1$,

$$\begin{aligned} \langle y_t y_{t+\tau} \rangle &= \vec{E}^T \langle \vec{x}_t \vec{x}_{t+\tau}^T \rangle \vec{E} + \langle r_t r_{t+\tau} \rangle \vec{E} \\ &= \vec{E}^T [\mathbf{X}(\mathbf{K}^T)^\tau] \vec{E} + [R \vec{B}^T (\mathbf{K}^T)^{\tau-1}] \vec{E} \\ &= [\vec{E}^T \mathbf{X} \mathbf{K}^T + R \vec{B}^T] (\mathbf{K}^T)^{\tau-1} \vec{E}, \end{aligned} \quad (3)$$

with the shorthand $\mathbf{K} \equiv \mathbf{M} + \vec{B} \vec{E}^T$ and $\mathbf{X} \equiv \langle \vec{x}_t \vec{x}_t^T \rangle$. We define $A(\tau) \equiv \langle y_t y_{t+\tau}^T \rangle / \langle y_t y_t \rangle$, the autocorrelation of y_t in the steady state (for $\tau \geq 1$). Note first that $A(\tau)$ is proportional to the quantity in Equation 3. Furthermore, the dependence of this autocorrelation on the trial lag τ enters only through the exponentiation of the matrix \mathbf{K} , which is defined by the parameters of emission, learning, and retention (Eqs. 1–3). In particular, from the form of Equation 3, we can write $A(\tau) = \alpha_1 \lambda_1^\tau + \alpha_2 \lambda_2^\tau$, where the λ_j are simply the eigenvalues of \mathbf{K} , and the α_j are coefficients that can be derived from the other parameters of the system. We fit this autocorrelation function (and a simpler, first-order version; Eq. 5a) to the empirical autocorrelations using nonlinear least-squares regression.

Results

Neural and behavioral variability

In the May 8, 2011 Los Angeles Dodgers baseball game, Clayton Kershaw threw 74 four-seam fastballs. The variability across these repeated pitching movements can be measured in the release point, where the baseball leaves the pitcher’s hand (Fig. 1a), or the speed of the ball at that point. We chose to focus on these metrics, because they represent the endpoint of the pitching movement that is under the direct control of the pitcher. We analyzed release-point variability for all 160 instances in the 2011 MLB regular season in which a pitcher threw at least 70 four-seam fastballs in a game. On average, the SD of the release point and the pitch speed are 4.4 cm and 0.59 m/s, respectively. We are interested in why such variability remains, even in such highly trained movements.

To answer this question, we studied the neural basis of movement variation in the context of well practiced reaching movements in nonhuman primate. Two rhesus macaque monkeys were trained to make center-out reaching movements in the horizontal plane toward visual targets, with an instructed delay period. As seen in sample movement trajectories (Fig. 1b), animals displayed appreciable movement variability on repeated trials to the same target, despite receiving graded reward based on reach accuracy. To analyze this movement variability, we defined two primary movement metrics: (1) reach speed, defined as the peak tangential speed in the movement plane; and (2) initial direction, defined as the angle from initial hand position to the trajectory point where the instantaneous speed of the reach first exceeds 40% of its peak speed. We chose these metrics, because they are less affected by online corrections during the movement, yet results for reach endpoint are qualitatively similar. Variability in both initial reach direction (average SD of 14.3°) and reach speed (average SD of 0.06 m/s, or 12% of the average speed) are consistent with previous reports (Georgopoulos et al., 1981; Riehle and Requin, 1993; Churchland et al., 2006).

Neural activity predicts behavioral variability

During the experimental sessions, we also recorded extracellular activity from a large population of neurons (on average, 66 and 55 neurons for monkeys D and E, respectively) using chronically implanted 96-microelectrode arrays (Blackrock Microsystems; Maynard et al., 1997), one array in PMd of monkey D and one each in PMd and M1 of monkey E. (Here, we will not distinguish between PMd and M1 activity, because M1 neurons analyzed separately yield qualitatively similar results.) Except when otherwise noted, we describe neural activity related to movement preparation, quantified as the average firing rate for each neuron during the delay period.

We know, from data collected during our eight-target tuning trial blocks (see Materials and Methods) and from previous studies (Georgopoulos et al., 1986) that preparatory neural activity can be used to predict movement direction. However, we are specifically interested in how the variability in the neural response—i.e., variability across presentations of the same stimulus—is manifest in the variability of the behavioral response. Thus, we focus on trials from the three-target trial blocks, which include a large number (~ 300) of repeated reaches to each of three given targets in a Cartesian quadrant. To eliminate dependence on an explicit neuron tuning model (see Materials and Methods), we computed, separately for each target, the deviations from the mean neural and behavioral responses and analyzed the relationship between them. Specifically, for each target and experimental session, we constructed the optimal linear estimator relating the vector of delay-period firing rates to the movement speed and, separately, to the initial direction. These optimal linear estimators were computed using the leading principal components of the firing rates (typically 10–15), with the number of components determined by cross-validation; all reported R^2 values (coefficient of determination) are cross-validated (see Materials and Methods). With this approach, we found that PMd/M1 neural activity predicts a substantial fraction of the variation in movement speed and direction: on average, across sessions and targets, R^2 was 24 and 12% for reach speed and 15 and 28% for reach direction for monkeys D and E, respectively. These results appear to confirm a previous report that the precision of reach speed is limited by central variability in the movement plan (Churchland et al., 2006) and to extend that result to reach direction.

Drift in motor variability and cortical neural activity

The results of the previous section suggest that behavior is noisy in large part because the cortical neurons that drive them are noisy. However, this simple interpretation is complicated by the observation that there is nontrivial time structure to the behavioral variability. This structure can be seen in the slow drifts in speed and initial direction in the sequence of reaches to one target in a sample session (Fig. 2*a,c*), which is consistent with previous reports (Cheng and Sabes, 2007; Chestek et al., 2007). To quantify these drifts, we decomposed peak speed or initial direction x as a function of trial number t into three components:

$$x_t = \bar{x} + \Delta x_t^d + \Delta x_t^r, \quad (4a)$$

$$\Delta x_t^d = c_0 + c_1 t + c_2 t^2 + c_3 t^3 + c_4 t^4,$$

(4b)

the session mean \bar{x} , the “drift” Δx_t^d , and the trial-to-trial “fluctuations” Δx_t^r that remain after removing \bar{x} and Δx_t^d . Naively characterizing the drift, Δx_t^d , as a fourth-order polynomial in t (Eq. 4b), we found that almost all session targets exhibit significant drift: of 78 total session targets (for both monkeys), 75 and 72 show statistically significant drift in reach speed and direction ($p < 0.05$, permutation test). Drift accounts, on average, for 41 and 24% of the variance in reach speed (Fig. 2*b*) and 18 and 24% of the variance in reach direction (Fig. 2*d*) in monkeys D and E, respectively. The polynomial characterization was chosen for simplicity; the polynomial is meant only as an empirical description of the drift and not its underlying mechanism, which is addressed below, and the quantitative results are not sensitive to this choice.

To provide a qualitative impression, the behavioral drift curves for all experimental sessions are shown in Figure 3. Each curve shows the polynomial drift curve (Fig. 2*a,c*, black curves) for a single behavioral measure, target, and session. There is a variety of drift shapes and sizes. In particular, note that the drifts are generally not monotonic and that there is no obvious relationship between the concurrent drifts in the two behavioral parameters.

Given that behavioral variability consists of two distinct components, it is natural to ask whether neural activity in macaque PMd/M1 has an analogous structure across trials. As with behavior, we observed a slow drift in the neural firing rates across the course of many reach trials, illustrated in Figure 2*e*. Using the same characterization of drift as for the movement metrics (Eq. 4), we found that, on average, 9.0 and 7.6% of the variance across trials in individual neuronal firing rates is attributable to drift in monkey D and E, respectively (Fig. 2*f*). Although this is smaller than the fraction of the behavioral variability that is attributable to drift, the total variability in the response of a neuron is expected to be higher, e.g., as a result of intrinsic neural spiking noise, and that, for a particular target, many neurons will exhibit low mean firing rates, making drift difficult to measure.

Slow drifts in performance appear to be a hallmark of highly trained behaviors in expert humans as well (Fig. 2*g,h*): with the

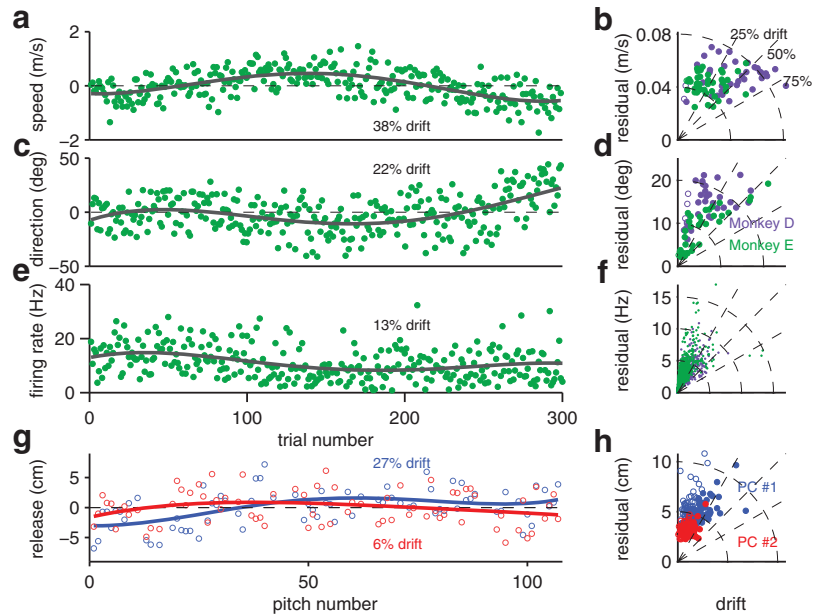


Figure 2. Quantification of behavioral and neural drift. *a*, Mean-subtracted speed for reaches to a single target during one experimental session (monkey E); the black trace shows the best polynomial fit of drift (Eq. 4). *b*, Breakdown of macaque reach speed variability for each session target into drift and residual components, expressed as SDs for both monkeys. Solid points indicate session targets with significant ($p < 0.05$, permutation test) drift. *c–h*, Analogous to *a* and *b* for initial reach direction (*c*, *d*), neural firing rate (*e*, *f*; neural firing rates were not mean subtracted, and in *f*, no distinction is made for cells with significant vs insignificant drift) and principal components (PC) of pitcher release point for 2011 MLB data (*g*, *h*).

same characterization of drift (Eq. 4), we found that, of the 160 MLB pitcher games described above, 99 and 86 exhibit statistically significant ($p < 0.05$, permutation test) drift in release point and pitch speed, respectively. On average, the drift accounts for 9.5% of variance in fastball release points for this dataset and 12.6% of variance in pitch speed; the scale of these drifts is generally larger than the < 1.3 cm PITCHf/x release-point measurement error. (The smaller fraction of pitcher games with significant drift compared with monkey sessions is likely attributable in part to statistical limitations, e.g., the smaller number of movement repetitions.) We verified that these drifts are not attributable to systemic changes in PITCHf/x measurements by verifying that the release points of opposing pitchers performing during the same game (concurrently) do not drift in the same direction.

Central neural activity only weakly predicts trial-by-trial fluctuations in behavior

Having identified these two statistically distinct components of both neural and behavioral variability, we then asked whether either, or both, of the two distinct components of behavioral and neural variability are related. We first assessed the relationship between the trial-by-trial fluctuations. We isolate the fluctuations by explicitly modeling the drift, e.g., by using Equation 4 to subtract the drift from each sequence.

First, for individual neurons, we simply computed the correlation coefficient between the drift-removed variation in firing rate and the drift-removed behavioral metrics. We performed this computation for the three-target blocks in each session, separately for each neuron and reach target. The results of this analysis are shown in Figure 4, *a* and *b*: single cells predict only a small fraction of the variance in both reach speed and direction. Furthermore, the distributions of R^2 values mostly overlap with the distributions calculated in the control case, in which neural re-

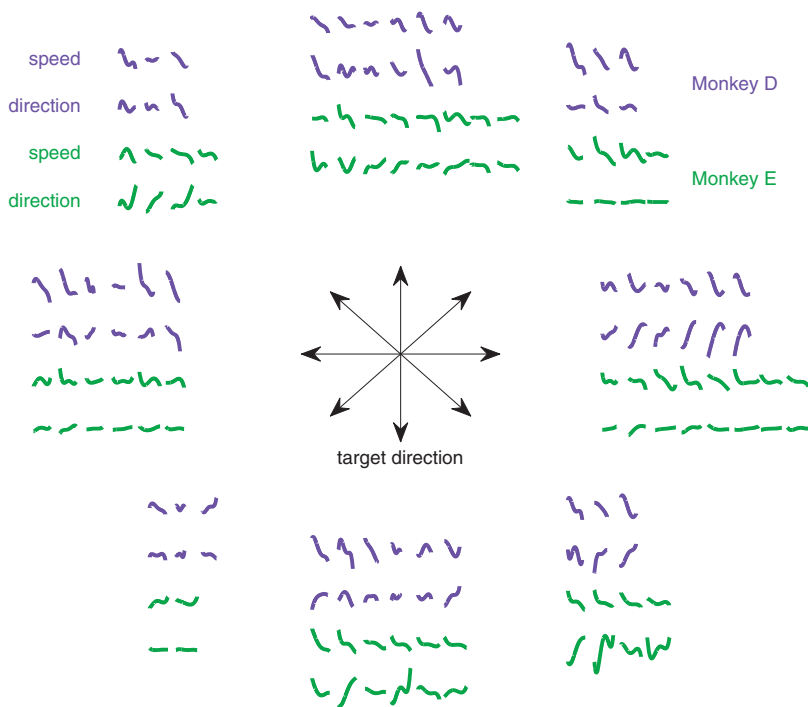


Figure 3. Time course for all behavioral drifts. Each trace corresponds to the fourth-order polynomial fit (Eq. 4) of the drift for a given behavioral metric (speed or direction) for a given target in a single session. Vertical pairs correspond to the speed and direction drifts measured in the same session. Each of the eight panels contains data from the corresponding target direction, with color indicating the animal's identity.

cordings from one experimental session are correlated with movement metrics from another experimental session (with identical target conditions). Thus, the firing rates of individual neurons are poor predictors of trial-by-trial behavioral fluctuations.

Next, we related behavioral and neural fluctuations at the population level. We computed optimal linear estimators (relating each movement metric to the vector of firing rates across cells) exactly as before, except that both the neural activity and behavior metrics were drift removed. The results of this analysis, again expressed as R^2 , are shown in Figure 4, *c* and *d*. Even with a large number of PMd/M1 cells that are apparently tuned for reach movements (Fig. 2), these neural populations have little predictive power for the trial-to-trial fluctuations in behavior, only a few percent better than the control case (correlating neural recordings with behavior from different experimental sessions, as above).

The weak correlations we observed between the fluctuations in neural firing rates and in behavior might have been attributable to our having averaged the firing rates over such a long preparatory period (800–1300 ms). However, we repeated these analyses, trying to predict behavioral fluctuations using the neural activity recorded during each quartile of the delay period (200–325 ms each), and obtained similarly small R^2 values (Fig. 4*e,f*, “motor prep. by quartile”). We also repeated these analyses using neural firing rates from other epochs of the trial timeline, including the period of movement onset, and again found similarly small R^2 values (Fig. 4*e,f*). We could find no point in the trial timeline during which the neural firing rates were good predictors of fluctuations in behavioral variability.

Finally, to ensure that the results of Figure 4 were not dependent on the use of Equation 4 (fourth-order polynomial fit of drift) to estimate the trial fluctuations, we repeated these analyses

with fluctuations isolated in an alternate, model-free way: by analyzing not the speed, initial direction, and firing rates themselves but rather the differences between these values on each consecutive pair of trials. Because the drift is slow, this also effectively removed the drift from these quantities; the resulting sequences had near 0 autocorrelation. The results were qualitatively the same as those shown in Figure 4.

We conclude that the representation of trial-to-trial behavioral fluctuations in PMd/M1 activity is weak.

Relating neural and behavioral drifts: correlation of drift magnitude

Having observed that the trial-by-trial fluctuations in neural activity and behavior are only weakly correlated, we now ask whether their drifts are related. It might seem natural to perform the analogous analysis to that used for the fluctuations, i.e., to compute optimal linear estimators relating the neural and behavioral drifts, and then to compute the R^2 between the actual and predicted behavioral drifts. However, these drifts, by construction, have very strong trial-by-trial correlations, so the space of possible drifts is low-dimensional, despite a large number of trials. In our case, the space of drifts has only four degrees of freedom, because we fit the drift with fourth-order polynomials. This means that the resultant R^2 is not meaningful: combining ~ 60 neural drifts into a prediction of behavioral drift will necessarily lead to a perfect fit. Indeed, when we repeated the analyses in Figure 3, *c* and *d*, using the neural drifts to predict the behavior drifts, we found that, in each case, we could “predict” the behavior drift with an R^2 value of unity. Therefore, overfitting precludes the use of a simple regression approach to compare these low-dimensional objects. However, two alternative analyses demonstrate the relation between the neural and behavioral drifts: we show that the magnitudes of the drifts are correlated and that they also have similar shapes.

First, we show the magnitudes of behavioral and neural drifts are correlated across sessions. For each target and session, we quantified the magnitude of the behavioral drift as its SD across trials (Fig. 2*b,d*, abscissas). We quantified the magnitude of the neural drift as a population average: we started with the SDs (Fig. 2*f*, abscissas) and then normalized by the mean activity for each unit, yielding the coefficient of variation, and finally averaged this value across units, separately for each target and session. Because target-dependent differences in behavioral or neural drift could artificially inflate the correlation between these measures, we took the conservative step of separately subtracting from both quantities the means for session targets with matched experimental conditions. The resulting values are positive for each animal and each movement metric (Fig. 5*a,c*): $r = 0.58$ and 0.23 for reach speed and $r = 0.54$ and 0.27 for reach direction, for monkeys D and E, respectively. Only the correlations for monkey D are significant ($p < 0.05$, *t* test) when analyzed separately for each animal and movement metric. However, the results for both metrics are significant when data from the two animals are combined,

and the results for both animals are significant when data from speed and direction are combined. Thus, the magnitude of drift in neural activity is indeed predictive, across targets and sessions, of the magnitude of drift that will be present in the behavior.

For comparison, the magnitudes of the trial-to-trial fluctuations in neural activity and behavior are shown in Figure 5, *b* and *d*. In all four cases, the magnitudes are anticorrelated ($r = -0.17$ and -0.18 for reach speed and $r = -0.07$ and -0.18 for reach direction for monkeys D and E, respectively), although these correlations are not significantly negative. Furthermore, the correlation between neural and behavioral drift magnitudes are significantly stronger than the trial-to-trial fluctuation magnitudes ($p < 0.05$, *t* test), even when assessed for each animal and each movement metric separately. These results provide additional evidence that the neural and behavioral fluctuations are not strongly linked but suggest that the neural and behavioral drifts are linked.

Relating neural and behavioral drifts: similarity of drift shape

Having established that the drifts in neural activity and behavior have correlated magnitudes, we proceeded to determine whether they exhibit similar shapes. Specifically, we analyzed autocorrelation functions $A(\tau)$, the correlation of the variable of interest with itself τ trials delayed, of the neural firing rates and behavioral metrics. As shown in Figure 6*a–c*, the autocorrelation functions, averaged over session targets (and neurons), all have similar shapes. We first asked whether these autocorrelation functions had an exponential form,

$$A(\tau) = \alpha\lambda^\tau, \quad (5a)$$

which would be consistent with a random-walk drift in the mean behavior or neuronal activity. However, the best-fit exponentials (Fig. 6*a–c*, gray lines) do not capture the shape of the autocorrelations: the empirical data show a more rapid drop in autocorrelation at short time differences and a slower decay at long time difference. This observation suggested the presence of more than one timescale in the drift process, and in fact we find (Fig. 6*a–c*, colored lines) that the autocorrelations are very well fit by a double exponential:

$$A(\tau) = \alpha_1\lambda_1^\tau + \alpha_2\lambda_2^\tau, \quad (5b)$$

where $\lambda_1 < \lambda_2$ and where the timescales of the fast and slow drifts processes correspond to λ_1 and λ_2 , respectively.

We will provide additional motivation for the double-exponential autocorrelation below. First, however, we argue that a common mechanism underlies the behavior and neural autocorrelations in Figure 6*a–c*. The first piece of evidence is shown in

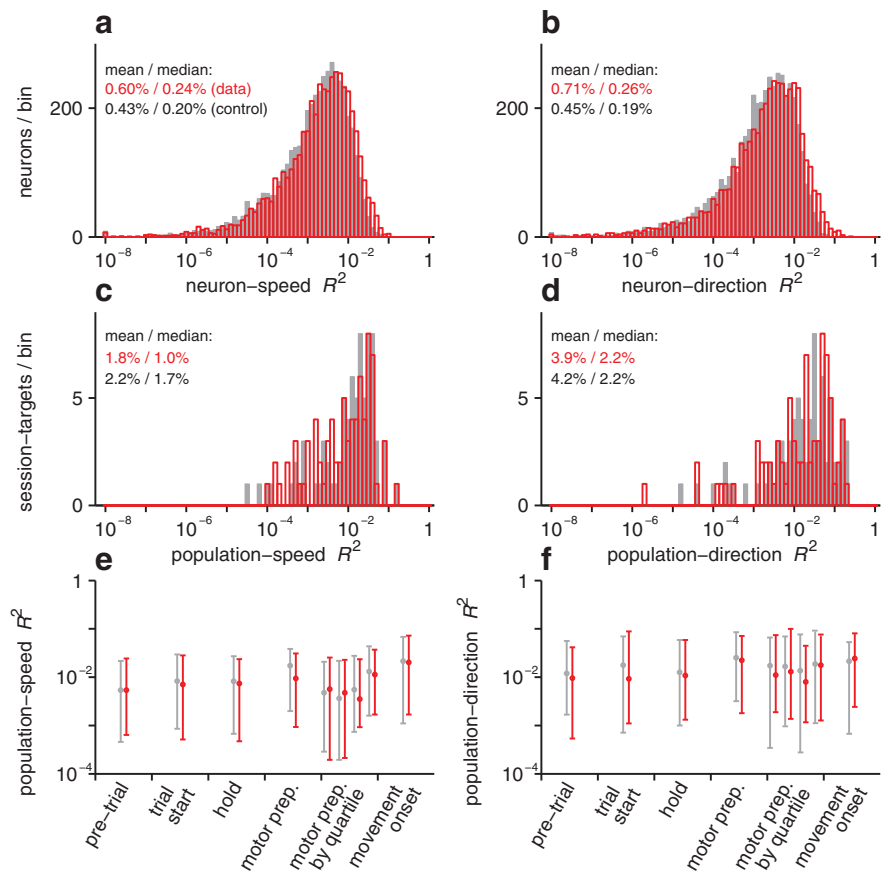


Figure 4. Relating trial-by-trial fluctuations in neural activity and behavior. *a, b*, Histogram (with bins logarithmically spaced) of R^2 between trial-to-trial fluctuations in behavior (reach speed, *a*; initial direction, *b*) and trial-to-trial responses of individual neurons, with drift removed from both. Control histogram for null hypothesis of no correlation (gray) was determined by correlating neural activity with behavioral data from different sessions with matched experimental conditions. Data are combined from both monkeys, but results hold for each. *c, d*, Analogous to *a* and *b*, but correlating behavior to fluctuations in PMd/M1 population response via optimal linear estimator (see Materials and Methods). *e, f*, Summary of population behavior R^2 analysis (*c, d*) but correlating behavior to neural firing rates observed in different trial epochs. Each data point represents distribution median, and error bars denote the smallest interval containing 68% of data points; colors are analogous to *a–d*. “pre-trial” is a 1 s window before the trial begins. “trial start” begins with the appearance of a central start target; “hold” begins when the monkey has moved to that target. “motor prep.” begins with the appearance of the visual reach target and ends with the go cue. We also repeated the analysis using the neural firing rates measured during each of the four quartiles of the preparatory period. (For more details on the trial timeline, see Materials and Methods.)

Figure 6*d*, which shows the time constants for the fits of Equation 5*b*, plotted on a log scale as $1 - \lambda_1$ and $1 - \lambda_2$, i.e., the fraction of the correlation that decays with each trial, for the fast and slow processes, respectively. This plot shows that there is consistency in the best-fit time constants across variables—reach speed, reach direction, and neural activity—and across animals. Next, if these behavioral and neural time constants reflect a common underlying mechanism, then they should also correlate with each other across experimental sessions and across targets. Testing this prediction is complicated by the fact that there was not sufficient data from a single target session to reliably fit the two-time-constant model of Equation 5*b*. Instead, we fit these reduced datasets with the single-time-constant model of Equation 5*a* and were able to obtain reliable fits. The single time constant λ primarily reflects the slower time constant of Equation 5*b*, although it is also influenced by the fast time constant [the log of $1 - \lambda$ and $1 - \lambda_2$ correlated significantly ($p < 0.05$) for all three variables, but the log of $1 - \lambda$ and $1 - \lambda_1$ only correlated for reach speed]. Figure 7 shows the correlations across target sessions in the log of $1 - \lambda$, pairwise for the two behavioral measures and for the

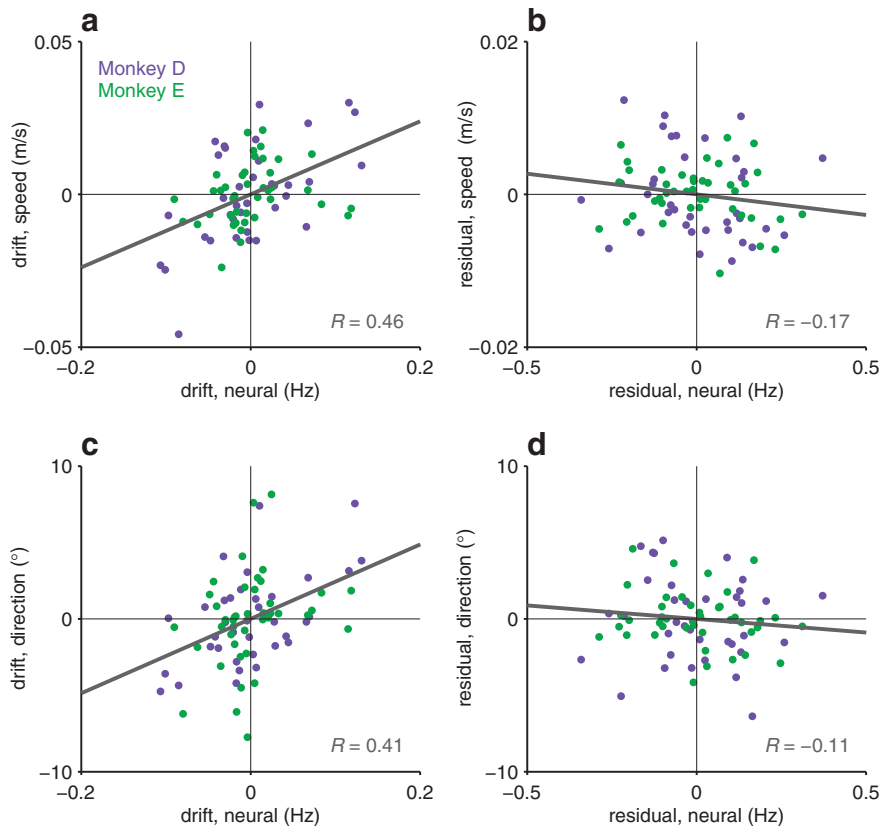


Figure 5. Relating neural and behavioral drift magnitudes. **a**, Correlation across sessions and targets between the amounts of variability attributable to drift in reach speed and in preparatory neural activity. **b**, Analogous to **a** but comparing residual trial-to-trial variability instead of variability attributable to drift. **c**, **d**, Analogous to **a** and **b** but with reach direction as the behavioral variable. In all panels, the behavioral variability is expressed as SD; neural variability is expressed as coefficient of variation (SD/mean), averaged across all neurons. All values are mean subtracted for matched experimental conditions (see text). The gray trace shows the regression slope for both monkeys combined. For both behavioral metrics, the magnitude of behavioral drift correlates positively with the magnitude of neural drift for each monkey (**a**, **c**) and is significant with both monkeys combined ($p < 0.05$, t test).

neural firing rates. The correlations are significant for each pairwise correlation and each animal, with the exception of the speed-rate correlation for monkey E. Together, these results suggest that the behavioral and neural autocorrelations arise from the same underlying process.

We also computed the autocorrelation functions for pitch release point and speed in the MLB data. These were similarly well described by a double exponential, with time constants notably similar to those obtained with the nonhuman primate reaching data (Fig. 6*d*). This similarity in the statistics of the pitching and reaching again suggests a common underlying mechanism.

Drift and error-corrective learning

Given that the drifts in neural activity and behavior appear to arise from a common mechanism, the natural question is: what might that mechanism be? For the reaching experiment, drift could arise from time-varying changes in the motivational or attentional state of the animal (Arieli et al., 1996; Xu-Wilson et al., 2009). For example, increasing fatigue or reward satiety would seem to predict simple monotonic patterns of drift over time. In fact, the behavioral drifts exhibit complex features. In particular, drifts in movement metrics are generally not monotonic, and we find no evidence that drifts in reach speed and reach direction are correlated in any straightforward way (Fig. 3). If the level of motivation, attention, or fatigue had been drifting non-

monotonically over the course of a session, we would expect these changes to be reflected in a parallel drift in the variance of the trial-to-trial fluctuations. As is illustrated in example data of Figure 3, *a*, *c*, and *e*, the variance of the trial-to-trial fluctuations was flat across experimental sessions. This last observation, along with the lack of predictable drifts in reach direction, also suggest that the drift is not attributable to the sharpening of a Bayesian prior over target locations in the three-target blocks (Verstynen and Sabes, 2011).

The fact that the neural and behavioral autocorrelations are so well described by a double exponential (Eq. 4*b*; Fig. 6) support the idea that these drifts are attributable to a process with well defined statistics. Here we consider the possibility that this process is error-corrective learning, a constant feature of skilled movements such as reaching, even when no external perturbations are experienced (Cheng and Sabes, 2007; van Beers, 2009). Error-corrective learning is well modeled as a linear dynamical system with variability arising in both the motor output and the learning itself, attributable to noisy error feedback, state maintenance, etc., and this process can drive drifts in movement metrics over time, with exponential autocorrelations (Cheng and Sabes, 2006, 2007). Smith et al. (2006) have argued that error-corrective learning is better modeled as a two-state linear dynamical system with two timescales of learning. In Materials and Methods, we show that this

two-state learning model necessarily leads to behavioral drifts that exhibit a double-exponential autocorrelation function (Eqs. 1–3), just like those we observe in the present study (Fig. 6*a–c*). To test whether these autocorrelations may indeed arise from a two-state error-corrective learning process, we used Equations 1–3 to infer the autocorrelation time constants from the two learning rates that Smith et al. (2006) measured directly from the behavioral corrections after dynamic perturbation. These time constants are comparable with those we obtain from fits of the autocorrelation functions in our neural and behavioral reaching data, as well as the pitching data (Fig. 6*d*). (For the latter, to reduce sensitivity to noise, we only included the 64, of 160, pitcher games in which measured drift is found to be significant at $p < 0.003$, by permutation test.) In summary, we find that the learning model suggested by Smith et al. (2006) quantitatively predicts the behavioral and neural autocorrelations that we find in macaque reaching and the autocorrelations observed in MLB pitching data. These results are suggestive that the drifts are attributable to a noisy error-corrective learning process.

Discussion

We have shown that, across repeated reaches to a target, the behavioral metrics of the movements and the neural activity in PMd/M1 exhibit two statistically distinct components of variability: (1) a slow drift in the mean neural/behavioral responses; and

(2) trial-to-trial fluctuations. We found that trial-to-trial fluctuations in central neural activity do not strongly predict trial-to-trial fluctuations in reach behavior. Conversely, drifts in neural activity are linked with the drifts in behavior and appear to arise from the same process, one that is well described as the accumulation of noise in a simple model of error-corrective learning.

Although our conclusions rely primarily on linear analyses, we did not assume linearity in the full mapping from neural activity to behavior (e.g., cosine tuning of velocity). Rather, we built local, first-order models relating fluctuations in neural activity to fluctuations in behavior on nominally identical trials. Thus, by focusing on many trials to the same target, we were able to link neural activity and behavior, without requiring precise knowledge of the tuning properties of the cells or the full mapping from neural activity to behavior. This approach is valid as long as the true causal link between neural activity and behavior depends on firing rate and is continuous and differentiable, both mild and reasonable assumptions.

The weak relationship between trial-to-trial fluctuations in behavior and in central neural activity—at any point in the trial timeline we considered—was unexpected. This implies that behavioral fluctuations, which constitute the majority of total behavioral variability (Fig. 5*b,d*), primarily arise either downstream of PMd/M1 or otherwise independently. The weak propagation of trial-to-trial fluctuations in central neural activity to behavior also suggests that the neural circuits downstream of PMd/M1 effectively average out the noise in this activity. This is especially surprising given that these neural populations do exhibit nontrivial neuron–neuron correlations, as has been reported in these and other analogous cortical areas (Lee et al., 1998; Huang and Lisberger, 2009).

There has been considerable interest in whether the tuning properties of sensorimotor neurons are stable over both time and repeated usage. The drift in neural activity that we observed is not inconsistent with previous studies suggesting that the tuning of target direction for individual neurons changes over time (Carmena et al., 2005; Rokni et al., 2007). However, there are methodological difficulties with those studies as a result of poor statistical power in measuring tuning parameters with a center-out experimental design (Stevenson et al., 2011), and other reports have shown that tuning properties of neurons exhibit stability over timescales at least as




Figure 6. Autocorrelation functions and the error-corrective learning model. *a–c*, Mean (and SE) autocorrelation functions of reach speed (*a*), direction (*b*), and neural activity (*c*), averaged across target and session for behavior and across targets, sessions, and neurons for neural activity. The data were fit to a linear dynamical learning model characterized by either one decay time constant (first-order model, Eq. 5a; gray lines) or two decay time constants (second-order model, Eq. 5b; colored lines). *d*, Comparison of fit decay time constants using the second-order model; fit errors are from the bootstrapping technique. These time constants are also compared with analogous time constants found in the 2011 MLB pitching data. (To reduce our sensitivity to noise, we only included the 64, of 160, pitcher games in which measured drift is found to be significant at $p < 0.003$, permutation test.) We also show motor learning time constants from Smith et al. (2006), who directly measured the time course of learning in response to external perturbations. (There, learning is described with a two-state learning model; we used Equation 3 to express their empirical parameters as autocorrelation time constants. We assumed that their stated measurement errors were independent but verified that this does not significantly affect the size of the errors on the time constants.)

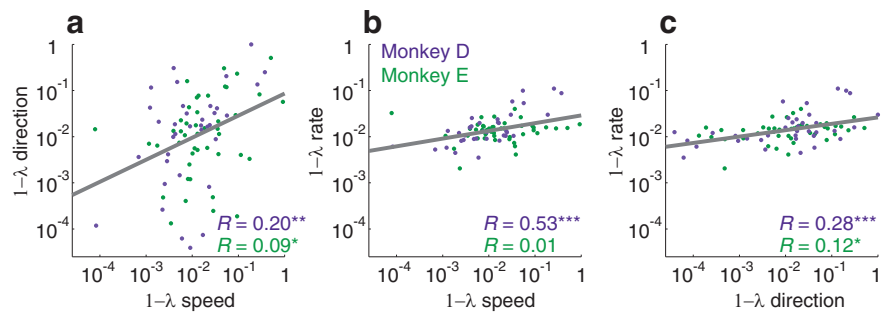


Figure 7. Relating neural and behavioral autocorrelation time constants. Correlation across sessions and targets between the autocorrelation time constants, fit using the first-order model (Eq. 5a). *a*, Correlation between time constants for movement direction and speed. *b*, Correlation between time constants for neural firing rate and movement direction. *c*, Correlation between time constants for neural firing rate and movement speed. For all plots, the gray line is the regression fit for data from both monkeys combined. Correlation coefficients, R , are shown for data from each monkey separately, with p values denoted as follows: * $p < 0.05$; ** $p < 0.005$; *** $p < 0.001$.

long as typical experimental sessions (Chestek et al., 2007; Stevenson et al., 2011; Fraser and Schwartz, 2012). Here, we avoid the statistical problems associated with fitting full tuning curves by focusing on whether and how neural activity and behavior

change over time for nominally identical trials. It is clear from our results and those of others (Chestek et al., 2007; van Beers et al., 2013) that both neural activity and behavior can drift over time. Here, we have shown that the neural and behavioral drifts are linked, suggesting that at least some of the changes in neural activity occur at or upstream of PMd and are behaviorally relevant.

We agree with the conclusion of Rokni et al. (2007) that drifts in the patterns of neural activity likely arise from a noisy learning process, i.e., a continual processes of adjusting the state of the system to reduce performance errors (Cheng and Sabes, 2006, 2007). Specifically, we showed that the statistical structure of both the neural and behavioral variability are well described by an error-corrective learning model with two autocorrelation time constants. These autocorrelation time constants are also consistent with those found in a previous study of motor learning with external force perturbations (Smith et al., 2006), suggesting that a single mechanism may account for the dynamics of learning with or without external perturbations. Thus, although Rokni et al. (2007) argued that drift only occurs along dimensions that do not affect behavior, we show, by analyzing the mapping from neural activity to behavior at one target over longer timescales, that these drifts in activity do appear to drive behavioral drift.

Given the sizable drifts that we observe in both neural activity and behavior, it is reasonable to ask why they have not been reported more frequently in the past. It is likely that that behavioral drift has been overlooked in previous studies when it was present. Most studies in which the value of movement metrics have been tracked across trials are studies of motor learning, which typically include an exogenous perturbation that drives large changes in behavior; the slow drifts observed here would be small by comparison and thus difficult to observe. These studies also typically focus on mean learning curves, which average out Brownian motion drift. Last, because of the slow time course of the drift we observe, it would be difficult to detect over shorter sequences of reach movements. Indeed, the drift analysis is data limited even in the present dataset, with ~ 300 trials per target, per session. Data collected over longer timeframes would permit a more complete characterization of the drift statistics (e.g., additional, longer timescales may be involved) and would also allow for a more direct correlation analysis between the neural and behavioral drifts.

Although we propose that the drift observed here is attributable to error-corrective learning, the degree to which learning is expected to drive drift depends on the details of the behavioral paradigm, in particular the error feedback and reward structure. Error-corrective learning is only one of several mechanisms for motor learning, and the relative importance of these depends on the details of the task, in particular the nature of the error feedback (Izawa and Shadmehr, 2011). When feedback is effectively binary (hit or miss), the motor system may be more likely to use an exploration-based reinforcement learning strategy (Shmuelof et al., 2012). In this case, the dynamics of the motor state may be dominated not by slow drift (as in our model, Eqs. 1–3) but rather by “intentional” trial-by-trial variability to explore the reward landscape (Thrun, 1992). Indeed, rapid changes in motor intention under such conditions could also explain why Churchland et al. (2006) observed short-timescale correlations between PMd activity and movement speed, whereas we do not. In that case, changes in the central representation of the motor plan would occur on the same timescale as trial-by-trial fluctuations (i.e., noise) and so are confounded statistically. When changes in the

plan are dominated by slow drift, as appears to be the case here, the behavioral drift represents a reasonable estimate (or at least lower bound) of the relative size of the central contributions to motor variability. Indeed, our measurements of 18–41% of behavioral variability attributable to drift are consistent with estimates of the central contribution arrived at through other psychophysical means (van Beers, 2009).

The two-state learning model (Eqs. 1 and 2) shows that pure error-corrective learning can result in drift, but the magnitude of that drift depends on the learning parameters. van Beers (2009) has argued that minimizing total output variance requires parameters that eliminate drift. Here we see a regimen in which learning is not strictly variance-minimizing. This may be a result of the potential for instability with noisy feedback, attributable to an intrinsic cost of learning being factored into the optimization or the desire to preserve learning-related exploration. That we see drift with very similar time constants in both rhesus macaques and professional athletes suggests that the reward regimen of our experiments is, if not naturalistic, at least ethologically relevant.

In summary, because we observe weak trial-to-trial neuron–behavior correlations in the motor cortex, short-timescale movement variability, or motor noise cannot generally be of central origin. The neuron–behavior correlations we observe are primarily attributable to a slow drift in the mean cortical activity pattern and the mean behavior, whose statistics suggest a process of trial-by-trial learning. This central source of movement variability may thus paradoxically reflect a key neural mechanism for maintaining accurate control.

References

- Albert J (2010) Baseball data at season, play-by-play, and pitch-by-pitch levels. *J Stat Educ* 18:3.
- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871. [Medline](#)
- Carmena JM, Lebedev MA, Henriquez CS, Nicolelis MAL (2005) Stable ensemble performance with single-neuron variability during reaching movements in primates. *J Neurosci* 25:10712–10716. [CrossRef Medline](#)
- Cheng S, Sabes PN (2006) Modeling sensorimotor learning with dynamical systems. *Neural Comput* 18:760–793. [CrossRef Medline](#)
- Cheng S, Sabes PN (2007) Calibration of visually guided reaching is driven by error-corrective learning and internal dynamics. *J Neurophysiol* 97:3057–3069. [CrossRef Medline](#)
- Chestek CA, Batista AP, Santhanam G, Yu BM, Afshar A, Cunningham JP, Gilja V, Ryu SI, Churchland MM, Shenoy KV (2007) Single-neuron stability during repeated reaching in macaque premotor cortex. *J Neurosci* 27:10742–10750. [CrossRef Medline](#)
- Churchland MM, Afshar A, Shenoy KV (2006) A central source of movement variability. *Neuron* 52:1085–1096. [CrossRef Medline](#)
- Faisal AA, Selen LP, Wolpert DM (2008) Noise in the nervous system. *Nat Rev Neurosci* 9:292–303. [CrossRef Medline](#)
- Fraser GW, Schwartz AB (2012) Recording from the same neurons chronically in motor cortex. *J Neurophysiol* 107:1970–1978. [CrossRef Medline](#)
- Georgopoulos AP, Kalaska JF, Massey JT (1981) Spatial trajectories and reaction times of aimed movements: effects of practice, uncertainty and change in target location. *J Neurophysiol* 46:725–743. [Medline](#)
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 223:1416–1419. [CrossRef Medline](#)
- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394:780–784. [CrossRef Medline](#)
- Huang X, Lisberger SG (2009) Noise correlations in cortical area MT and their potential impact on trial-by-trial variation in the direction and speed of smooth-pursuit eye movements. *J Neurophysiol* 101:3012–3030. [CrossRef Medline](#)
- Izawa J, Shadmehr R (2011) Learning from sensory and reward prediction errors during motor adaptation. *PLoS Comput Biol* 7:e1002012. [CrossRef Medline](#)

- Lee D, Port NL, Kruse W, Georgopoulos AP (1998) Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci* 18:1161–1170. [Medline](#)
- Maynard EM, Nordhausen CT, Normann RA (1997) The Utah Intracortical Electrode Array: a recording structure for potential brain-computer interfaces. *Electroencephalogr Clin Neurophysiol* 102:228–239. [CrossRef Medline](#)
- McGuire LM, Sabes PN (2011) Heterogeneous representations in the superior parietal lobule are common across reaches to visual and proprioceptive targets. *J Neurosci* 31:6661–6673. [CrossRef Medline](#)
- Osborne LC, Lisberger SG, Bialek W (2005) A sensory source for motor variation. *Nature* 437:412–416. [CrossRef Medline](#)
- Riehle A, Requin J (1993) The predictive value for performance speed of preparatory changes in neuronal activity of the monkey motor and premotor cortex. *Behav Brain Res* 53:35–49. [CrossRef Medline](#)
- Rokni U, Richardson AG, Bizzi E, Seung HS (2007) Motor learning with unstable neural representations. *Neuron* 54:653–666. [CrossRef Medline](#)
- Shmuelof L, Huang VS, Haith AM, Delnicki RJ, Mazzoni P, Krakauer JW (2012) Overcoming motor “forgetting” through reinforcement of learned actions. *J Neurosci* 32:14617–14621. [CrossRef Medline](#)
- Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* 4:e179. [CrossRef Medline](#)
- Stevenson IH, Cherian A, London BM, Sachs NA, Lindberg E, Reimer J, Slutzky MW, Hatsopoulos NG, Miller LE, Kording KP (2011) Statistical assessment of the stability of neural movement representations. *J Neurophysiol* 106:764–774. [CrossRef Medline](#)
- Thrun, S (1992) The role of exploration in learning control. In: *Handbook for intelligent control: neural, fuzzy and adaptive approaches* (White DA, Sofge DA, eds). Florence, KY: Van Nostrand Reinhold.
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235. [CrossRef Medline](#)
- van Beers RJ (2009) Motor learning is optimally tuned to the properties of motor noise. *Neuron* 63:406–417. [CrossRef Medline](#)
- van Beers RJ, Brenner E, Smeets JB (2013) Random walk of motor planning in task-irrelevant dimensions. *J Neurophysiol* 109:969–977. [CrossRef Medline](#)
- Verstynen T, Sabes PN (2011) How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *J Neurosci* 31:10050–10059. [CrossRef Medline](#)
- Xu-Wilson M, Zee DS, Shadmehr R (2009) The intrinsic value of visual information affects saccade velocities. *Exp Brain Res* 196:475–481. [CrossRef Medline](#)