Title: Cyclic, condition-independent activity in primary motor cortex predicts corrective movement behavior

Abbreviated Title: Cyclic neural activity in corrective movements

Author: Adam G. Rouse[1]*, Marc H. Schieber[2], Sridevi V. Sarma[3]

Affiliations:

[1]Department of Neurosurgery, Department of Molecular and Integrative Physiology, University of Kansas Medical Center, Kansas City, KS, 66160

[2]Department of Neurology, Department of Neuroscience, Del Monte Institute for Neuroscience, University of Rochester Medical Center, Rochester, NY, 14642

[3]Department of Biomedical Engineering, Institute for Computational Medicine, Johns Hopkins University, Baltimore, MD, 21218

Corresponding Author and Lead Contact: *Please address correspondence to:

Adam G. Rouse

Department of Neurosurgery

University of Kansas Medical Center

3901 Rainbow Blvd, Mailstop 3021

Kansas City, KS 66160

Figures: 8 figures and 2 tables in manuscript, 4 figures in extended data and figures

Word Counts: Abstract 123, Significance Statement 84, Introduction: 740 words, Discussion: 1,632 words

Acknowledgements: This work was supported by NIH NINDS K99/R00 NS101127.

Conflict of Interest: The authors declare no competing financial interests.

Abstract

Reaching movements have previously been observed to have large condition-independent neural activity and cyclic neural dynamics. A new precision center-out task was used to test whether condition-independent cyclic neural activity in the primary motor cortex (M1) occurred not only during initial reaching movements but also during subsequent corrective movements. Corrective movements were observed to be discrete with a time course and bell-shaped speed profiles similar to the initial movements. Condition-independent cyclic neural trajectories were similar for initial and corrective submovements. The phase of the cyclic condition-independent neural activity predicted when peak movement speeds occurred, even when the subject made multiple corrective movements. Rather than being controlled as continuations of the initial reach, a discrete cycle of motor cortex activity encodes each corrective submovement.

Significance Statement

During a precision center-out task, initial and subsequent corrective movements occur as discrete submovements with bell-shaped speed profiles. A cycle of condition-independent activity in primary motor cortex neuron populations corresponds to each submovement whether initial or corrective, such that the phase of this cyclic activity predicts the time of peak speeds—both initial and corrective. These submovements accompanied by cyclic neural activity offer important clues into how we successfully execute precise, corrective reaching movements and may have implications for optimizing control of brain-computer interfaces.

Introduction

Corrective movements with sensorimotor feedback are critical for elegant motor control. While a single, discrete movement like a pointing gesture may be mostly ballistic, more precise aiming movements typically require an error correction phase (Woodworth, 1899; Craik, 1947; Abrams et al., 1990; Sainburg et al., 1999; Elliott et al., 2010). Yet the way neurons in motor areas of the brain encode

and generate corrective movements to achieve movement precision is relatively unexplored. When examining populations of neurons in primary motor cortex during instructed movements, it has become clear that predictable dynamics occur with past neural spiking history being partially predictive of future spikes (Maynard et al., 1999; Jackson et al., 2003; Truccolo et al., 2005; Sarma et al., 2010). Yet behaving animals also respond to new sensorimotor information such as tasks that require precision and one might expect predictable neural dynamics to be reduced. Conversely, the predictable neural dynamics may explain limits on behavior.

Most experimental tasks to study individual neurons in primary motor cortex have used instructed movements with a clear onset and offset. These single movement tasks have identified a variety of movement features being encoded (Evarts, 1968; Thach, 1978; Georgopoulos et al., 1982; Kalaska et al., 1989; Kakei et al., 1999). In addition to encoded features, however, there is also a large condition-independent component in the firing rate of neurons in motor cortex (Kaufman et al., 2016; Rouse and Schieber, 2018). Condition-independent neural activity is the change in a neuron's firing rate from baseline over time that happens regardless of the instructed movement for any given trial within a given task. Techniques like demixed principal component analysis partition a neural population's activity into condition-independent modulation and the more classically described tuning to task conditions (Kobak et al., 2016).

In addition to being condition-independent or -dependent, changes in firing rate can also be synchronous or asynchronous/rotational across a population. Primary motor cortex neurons are not all synchronously time-locked for example, but rather have an asynchronous range of onset latencies from 200ms to 0ms before movement (Moran and Schwartz, 1999). Condition-independent activity is also not a synchronous increase and decrease in firing rate for each neuron but a more complex neural trajectory (Yu et al., 2007; Cunningham and Yu, 2014). The inherent cyclical structure for reaching movements is hypothesized to be related to intrinsic cortical dynamics (Churchland et al., 2012; Hall et al., 2014) with

the firing of some neurons reliably occurring earlier relative to other neurons with the generation of movement.

Cyclic, condition-independent neural dynamics presents a challenge in describing how corrective movements necessary for precision might be encoded using the traditional models typically ascribed to primary motor cortex (M1). For experimental tasks with single instructed reaches, it is straightforward to subtract away condition-independent firing rate and use the remaining firing rate to identify the encoding of movement direction. A linear "cosine-tuning" model fits this data well (Georgopoulos et al., 1986; Salinas and Abbott, 1994). Experiments varying instructed speeds have shown larger changes for faster movements suggesting some condition-independent neural activity is speed tuning (Moran and Schwartz, 1999). However, to decode small amplitude, precise movements with online correction, a more accurate model of condition-independent activity is likely needed to accurately identify the remaining activity that encodes the various movement conditions.

In this work, we focus on the condition-independent, rotational neural activity. These are the changes in firing rate that are consistent across movements but vary in timing across neurons. Better understanding these condition-independent cyclic dynamics will better define what constitutes a movement and how corrective movements are generated after an initial movement. It is unclear if neurons encode these smaller movements similarly with the same range of latencies. If there is an ordered set of latencies, this might better define a series of submovements related to behavior and also better identify and separate the condition-dependent encoding of movement features.

In the present analysis, we test whether neural dynamics in cortex are similar between initial and corrective movements, examining the condition-independent activity across a wide range of speeds. A precision center-out task that required moving to small targets was used to elicit visuomotor corrections. We examined whether corrective movements in our task are discrete submovements, behaviorally similar to initial movements. We then ask whether condition-independent activity follows a similar neural

trajectory that is cyclic for both initial and corrective submovements. Finally, we ask whether common neural dynamics improve our predictions of behavior for when initial and corrective movements occur.

Materials and Methods

Non-human primates

Two male rhesus monkeys, P and Q (weight 11 and 10 kg, ages 7 and 6 years old, respectively), were subjects in the present study. All procedures for the care and use of these nonhuman primates followed the Guide for the Care and Use of Laboratory Animals and were approved by the University Committee on Animal Resources at the [Author University].

Experimental Design

A precision center-out task was performed by the monkey using a joystick to control a cursor on a 24" LCD display. The monkey manipulated a custom, 18 cm long, cylindrical rod attached to the joystick housing. The joystick moved freely with minimal resistance as the spring mechanism for providing centering, restorative force was removed. The joystick could move approximately 15º from vertical which corresponds to a workspace of 9.3 cm in both the vertical and horizontal directions. The cursor directly represented measured joystick position and moved within a 1000 x1000 pixel workspace centered on the display. The cursor workspace limits were slightly within the physical limits of the joystick. Thus 110 pixels correspond to approximately 1 cm of movement at the end of the joystick. The cursor on the display was a small cross that represented a single point in the workspace. The custom software for task control sampled the joystick data, updated the scene, and stored the cursor position (equivalent to joystick position) and trial event times at 100 Hz.

The precision center-out task consisted of three sets of eight peripheral targets located equidistance and equally spaced in 45˚ intervals around a center, home target (Figure 1). The center target had a radius of 75 pixels. Each center-out target—defined in polar coordinates—was one of three

different sizes i) large targets spanning 45˚ of the workspace and covering 250-450 pixels from the center, ii) shallower targets also spanning 45˚ but covering a width only 325-375 pixels from the center, and iii) narrower targets only spanning 15˚ covering 250-450 pixels from the center. All 24 targets (3 sizes x 8 locations) were presented pseudo-randomly in equal amounts throughout a session.

For each trial, following the subject acquiring the home target and performing a required initial hold ranging from 300-500 ms, the instruction occurred with the given trial's correct target changing from black to green. Following instruction, the monkey could move the cursor immediately to contact the correct target. At contact, the outline of all targets changed colors from white to black providing visual feedback when the cursor was within the target boundaries. After contacting the desired target, the cursor was required to remain within the target for a variable hold time of 500-600 ms. If the cursor left the target during a final hold, the monkey was allowed to leave and contact the target again and complete a final hold. Once a successful final hold of 500-600 ms was completed, the animal received a liquid reward. Both the required initial and final hold times for each trial were randomly sampled from a uniform distribution.

A peak finding algorithm to identify local maxima (*findpeaks* function in Matlab (Mathworks, 2020)) was used for analysis of the timing of cursor speed peaks. Cursor speeds were identified as peaks if they met the following criteria: i) the peak speed was greater than 250 pixels/s and ii) the peak's prominence— the height difference between the peak and the larger of the two adjacent troughs (minimum speed before encountering a larger peak)—was at least 50% of the absolute height of the peak. All cursor speed peaks and surrounding time windows were identified as submovements within a trial. Initial peaks were identified as the first submovement that ended at least 150 pixels from the center (approximately halfway to the peripheral target). Any small movements before the initial speed peak—506 (4.6% of trials) for P and 616 (7.0% of trials) for Q—were discarded from further analysis. Additional speed peaks following the initial speed peak were defined as corrective submovements. To focus analysis on submovements made to successfully acquire the target, corrective submovements were

only included if some portion of the acceleration phase—time from preceding speed trough to speed peak—occurred outside the peripheral target.

Neural Recordings

Floating microelectrode arrays (MicroProbes for Life Science) were implanted in the anterior lip and bank of the central sulcus to record from primary motor cortex (M1) in each monkey, using methods described in detail previously [citation redacted]. For monkey P, recordings were collected from six 16-channel arrays implanted in M1. For monkey Q, two 32-channel arrays and one 16-channel array in M1 were used. The location of the implanted arrays were previously reported in (Fig. 2) of [citation redacted] and spanned the forelimb area of M1. Intracortical microstimulation on single electrodes with a current up to a maximum of 100 μA (0.2ms stimulation pulse width, 3ms interpulse interval) with the animal lightly anesthetized with ketamine showed a mixture of evoked forelimb movements. Of the 96 electrodes for monkey P, stimulation of 11 sites elicited proximal arm movements, 6 sites elicited wrist movements, and 21 sites elicited movement of the digits. Of the 80 electrodes for monkey Q, 34 sites were proximal, 9 sites were wrist, and 25 were digits. Channels with spiking activity were thresholded manually online and spike-snippets and spike times were collected with the Plexon MAP (Plexon, Inc.) and Cerebus (Blackrock Microsystems, LLC.) data acquisition systems. The spike snippets were sorted off-line with a custom, semi-automated algorithm. Both single- and multi-unit recordings were used in the present analysis.

Statistical Analyses

The firing rates for individual units were estimated by convolving spike times with a Gaussian filter (σ=30 ms). The firing rates were estimated with 10 ms bin resolution for 100 Hz signals to match the joystick data and events. The firing rates were then square-root transformed to equalize variance between high and low firing rates (Kihlberg et al., 1972; Snedecor and Cochran, 1980; Ashe and Georgopoulos, 1994). To make sure our results were robust to the square-root transform, identical

analysis was also performed without the square-root transform and the results for selected figures are presented in Extended Data and Figures. The overall results and conclusions were very similar in all analyses.

We wished to assess how firing rates changed across a large population of neurons as a function of time for both initial and corrective movements. To do this, we calculated the differences in firing rate between individual neurons and averaged across the population. Note, in the neural space this is equivalent to the L1 distance or city block distance. To aid in visualizing the average change in firing rate of all recorded spiking units across time, we chose to use classical multi-dimensional scaling (also known as Principal Coordinates Analysis, *cmdscale* in Matlab) on the calculated firing rate differences. Multidimensional scaling reduces the high-dimensional dataset to a chosen number of dimensions while preserving as well as possible the distance (differences in firing rates) between original datapoints.

Identifying condition-independent, rotational neural activity

We focused our analysis on the condition-independent, rotational activity in the recorded neural population. To explain the various features of neural activity, we'll consider an idealized example of four neurons, shown in Figure 1. First, the change in firing rate related to a task can be in part the same for all movement conditions (condition-independent) and/or in part vary depending on the particular movement condition (condition-dependent). The trial-averaged firing rates, including both condition-independent and –dependent firing rates (but excluding noise), for four idealized model neurons during a typical center-out reaching task are shown in Figure 1. In 1A, the condition-independent firing rates for each neuron are obtained by averaging the overall firing rates across all four target conditions, and condition-dependent firing rates then are obtained after subtracting away the condition-independent firing rates from the overall firing rate for each condition.

Changes in firing rate can also be either synchronous or asynchronous. Synchronous, time-locked activity represents changes in firing rate that happen simultaneously across the neural population,

while asynchronous activity of varying time course in different neurons can lead to patterns of traveling waves or oscillations in the population with a predictable progression in time. As shown in the time-domain in Fig. 1B, the synchronous activity (top) with all neurons having the same time course can either be condition-independent (left) or depend on the target condition (right). Similarly, the asynchronous activity (bottom) with various neurons leading and lagging can either be condition-independent activity (left) or target dependent (right).

Alternatively, the firing rates of the neuron population can be visualized as neural trajectories in a Cartesian neural space where each neuron's firing rate is plotted on an orthogonal dimension. Synchronous activity (top) appears as a trajectory that moves along a single dimension away or towards the origin. Asynchronous activity (bottom), when firing rates change at different times among the various neurons, appears as trajectories that traverse across various neural dimensions. Thus, the overall neural trajectory is caused by changes in firing rate across the population of neurons and projecting the data into different neural dimensions, i.e. subspaces, results in identification of the different features of neural activity.

For a complex task with various submovements and less well-defined conditions such as our precision center-out task, the condition-independent activity provides a useful analysis to identify repeated activity that defines a submovement. While the synchronous rise and fall of global firing rate provides some information, the rotational activity has the potential to improve identification of the repeated patterns of neural activation related to movement execution. Like any model of a large number of neurons, our model relies on simplifying assumptions. The condition-independent averaging across all trials is only relevant for the experimental data under study and should not be interpreted as condition-independent across all behavior. Additionally, the neural activity is likely neither fully synchronous nor rotates perfectly as a circle but likely forms an intermediate ellipse. Finally, a high-dimensional neural dataset is likely to have more complex dynamics than can be modeled with only linear dynamics that are

limited to a combination of rotation about and growth/shrinkage from the origin. Nevertheless, identifying the condition-independent, rotational activity can be a useful tool to study brain activity.

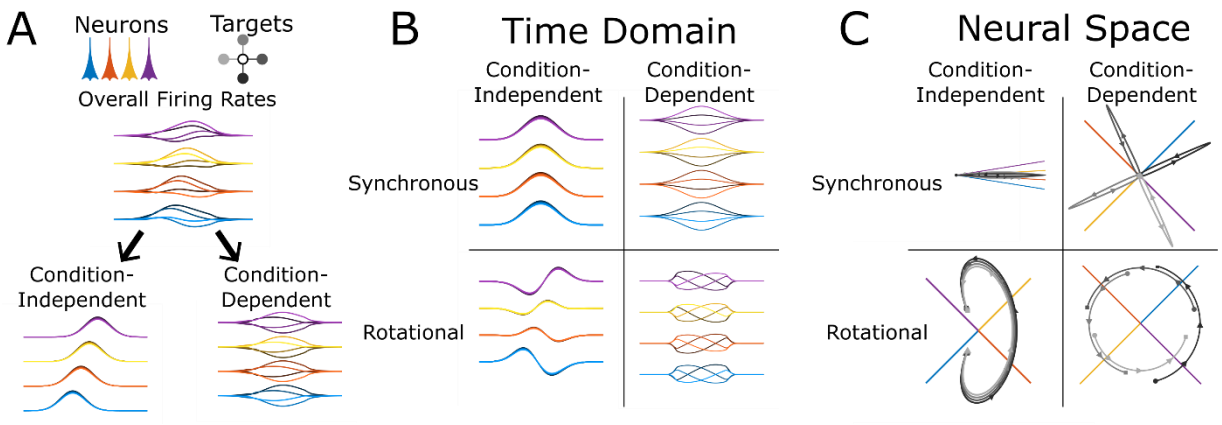

Figure 1. Idealized representation of both the synchronous and rotational components of condition-independent and -dependent changes in neuronal firing rate. A) The firing rates for four neurons (blue, orange, yellow, and purple) are shown for reaches to four target directions (light to dark grayscale). The overall firing rates are different firing rates for the four neurons for the four target condition. By time averaging across the four conditions, the condition-independent firing rates and the residual condition-dependent firing rates are both identified. B) Next, analysis of the population reveals that firing rates can be decomposed to either i) synchronous activity across all neurons at each time point or ii) rotational firing rate changes. When plotted in the time domain, the firing rates all modulate synchronously in time or switch between different neurons firing more at different times. C) The neural space visualizes the population activity by showing each neuron's firing rate as a point along an orthogonal dimension with time represented as a trajectory through this space. In this representation, the difference between synchronous and rotational activity is better appreciated. Synchronous activity is movement along a single neural dimension while rotational activity is movement between dimensions. Note, the dimensions defined by individual neurons are shown projected in a 2D plane. Only the given component (synchronous/rotational and condition-independent/-dependent) are shown for these four example neurons for visualization purposes. In a much higher dimensional space when recording from a large number of neurons, the possibility of finding dimensions with little overlap between components is much greater.

Dynamical Systems Model

Dynamical system modeling is a useful way to characterize the observed changes in firing rates. The simplest is a system of coupled, first-order ordinary differential equations where changes in firing rate depend only on a linear transform matrix multiplied by the current firing rates. In this model, the changes in firing rate can grow/shrink along a single dimension (synchronous) as well as rotate across dimensions (asynchronous). The eigen decomposition of the transform matrix yields eigenvalues with the real part representing growing or shrinking away from the origin while the imaginary part represents rotations.

To identify and analyze potential repeatable temporal dynamics of the neural population that correlated with movement, our neural data was modeled as a linear, time-invariant system. This model was built using only the condition-independent activity by averaging the firing rates for individual spiking units across all trials regardless of the movement condition (i.e. target location). The condition-independent activity was then submitted to jPCA (Churchland et al., 2012) to identify the two-dimensional neural plane with the most rotational/cyclic activity. Note, this utilization of the jPCA algorithm on only the condition-independent activity is different than the typical application of jPCA to data containing the condition-dependent activity. Once this jPC neural plane was identified, this work introduces a new analysis variable—condition-independent phase (CIφ)—which estimates at each time the instantaneous phase by using the Hilbert transform of the firing rates within the identified two-dimensional plane. More thorough descriptions of the dynamical systems model, jPCA dimensionality reduction, and CIφ estimation are given in the Extended data and figures.

The calculation of the jPC plane and the CIφ was performed using 5-fold cross-validation. Each recording session was divided into 5 testing sets of trials each containing 20% of the data. The jPC plane was calculated by training on the other 80% of the data and then tested on each test set. All presented

results for CIφ is using the test data projected into the jPC dimensions identified by the separate training set.

Firing Rate vs. Speed Model

As a comparison to our dynamical system model and phase analysis, the average firing rate across the neural population—the synchronous, condition-independent neural activity—was also used to identify when maximal firing rates occurred and could predict future speed peaks. Since some neurons increase their firing rate and others decrease their firing rate on average with movement, a linear prediction model was used to create a weighted sum of the neural firing rates to predict cursor speed peaks. The average firing rate of each unit from 300 ms before to 100 ms after each peak speed was used in a linear regression to predict the peak cursor speed for each submovement. The regression model thus generated weighting coefficients for each neuron's prediction of movement speed. These coefficients were then used to weight the instantaneous firing rates at any time point to generate a speed estimate. While this estimate predicts upcoming movement speed, we show the estimate from the firing rates with no lag added to report the time when the maximal firing rate occurred, not the peak speed of actual movement.

Neural Subspace Overlap

To measure the amount of overlap of condition-independent neural activity between initial and corrective submovements, we performed principal component analysis (PCA) on the average firing rates relative to peak speed for all neurons in the N-dimensional neural space. We identified the 6-PC space that accounted for the most condition-independent variance for either initial or corrective submovements. We then projected the condition-independent activity from the opposite submovement type in this 6 PC space. We defined the percentage aligned, A, using the same methods as Elsayed et al. (2016):

$$A\% = 100 \times \frac{TR(D_{Init}^{T} C_{Corr} D_{Init})}{\sum_{i=1}^{6} \sigma_{Corr}(i)}$$

D_{init} is the matrix that defines the first six PCs of the initial condition-independent neural activity. C_{Corr} is the covariance matrix of the corrective condition-independent neural activity and $\sigma_{Corr}(i)$ is the ith singular value of C_{Corr}. Tr(..) is the matrix trace. The values of the alignment index range from 0-100% and correspond to how much of the variance in the first 6 corrective PCs were contained in the neural dimensions of the first 6 initial PCs. The opposite projection of the initial PCs into the corrective PC space was also performed.

Circular Statistics

The CIφ values were tested for a significant relationship to the instantaneous cursor speed as well as the occurrence of speed peaks. All circular statistics—mean, variance, correlation, and Rayleigh test for non-uniformity—were calculated with CircStat, a Circular Statistics Toolbox for Matlab (Berens, 2009).

Software Availibility

Matlab code to perform the dynamical system modeling and to calculate the CIφ is available publicly on GitHub. Additional analysis code is available from the authors on request.

Results

Motor Behavior

Movement speed was analyzed throughout the trial from instruction until successful completion of the final target hold. The two monkeys successfully completed 10,963 (monkey P) and 8,737 (monkey Q) trials across 12 recording sessions each. There were 6478 and 3912 corrective submovements identified for monkeys P and Q, respectively. Across all trials, 68.3% (P) and 71.1% (Q) were completed in a single initial movement, 17.5% (P) and 20.3% (Q) of trials were completed with one additional corrective submovement, and 14.2% (P) and 8.6% (Q) of trials required two or more corrective submovements. The location of the identified speed peaks within example trials and the speed profiles

are shown in Figure 2A and 2B, respectively. The speed peaks tended to be distinctly separated with nearly zero velocity between most peaks. As shown in Figure 2C, 99.0% (P) and 97.7% (Q) of the minimum speed trough following the initial speed peak were less than 20% of the peak. Similarly, 82.4% (P) and 85.8% (Q) of the troughs were less than 20% of the preceding peak between corrective speed peaks. The mean peak speeds for initial submovements were 1533 (P) and 1182 (Q) pixels/s while corrective submovement peak speeds were 460 (P) and 400 (Q) pixels/s. Thus, the average peaks for corrective submovements were 30.0% and 33.8% of initial submovements and a low-speed trough almost always occurred between two speed peaks.

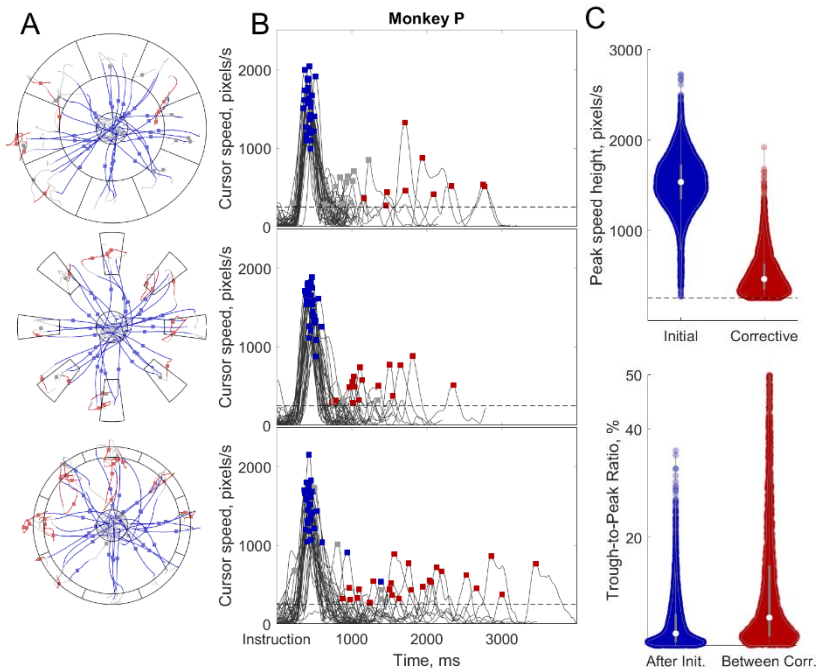

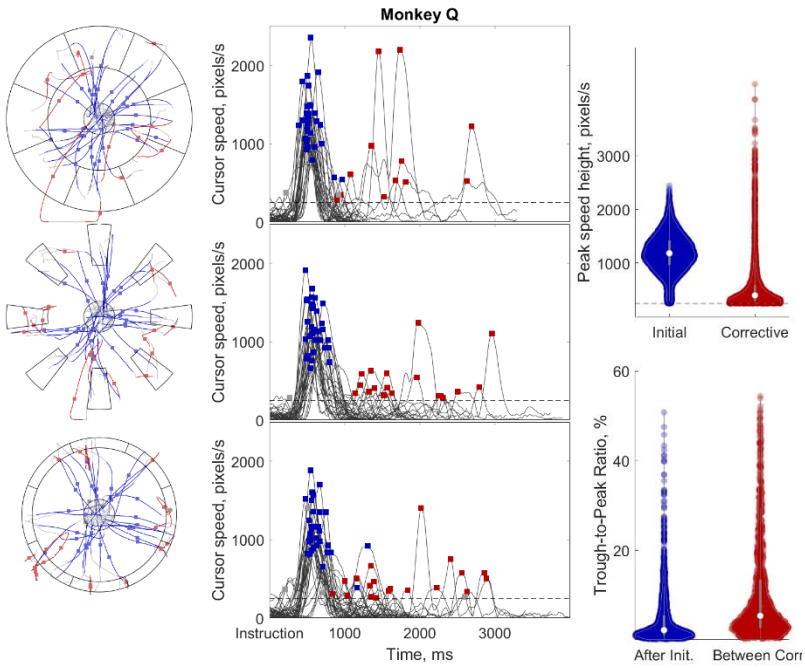

Figure 2. A) Cursor paths for four example trials to each target for the three target sizes: regular (top), narrow (middle), shallow (bottom). Initial submovements from 200ms before to 200ms after speed peaks are plotted in blue with the point when peak speed occurred shown with a blue dot. Corrective movements are similarly identified in red with a red dot. Grey lines connect the rest of a trial before, between, or after submovements with a speed peak. B) Cursor speed plotted versus time for a subset of

trials. Initial (blue) and corrective (red) submovement speed peaks are identified with squares. Gray squares identify speed peaks that were thrown out because they i) were small initial movements that did not move outside the center or ii) occurred entirely within the peripheral target. C) Top) Distribution of peak speeds for initial (blue) and corrective (red) submovements. Bottom) Distribution of the trough-to-peak ratio for the troughs following an initial submovement before a corrective submovements and following a corrective submovement before another corrective submovement.

The speed profiles were time aligned to peak speed to better examine the identified submovements (Figure 3A). Almost all submovements show a clear bell-shaped profile for both the initial and corrective movements. The time points when speed was one-half of the maximum speed were identified and the relationship between peak speed and submovement duration was examined. As shown in Figure 3B, the movement duration to half maximum speed was similar and close to symmetric for both initial and corrective submovements with initial submovements being slightly longer on average. Because the target was always the same distance from the center and the animals' initial reaches tended to get close to the target, there was an inverse relationship between the time duration of initial reaches and the maximum speed observed (Figure 2C). For corrective submovements, there was no clear relationship between movement duration and speed. Overall, all submovement durations as measured by the full width at half maximum occurred within a similar range, with 96.7%/88.0% (P/Q) of all initial and 96.3%/93.0% (P/Q) corrective submovements between 100-350 ms. The time between speed peaks—either initial to first corrective submovements or between subsequent corrective submovements—is plotted in Figure 2D. The median time between peaks was 570 ms for monkey P and 700 ms for monkey Q with the most common time between peaks being 450ms (P) and 550ms (Q). Only 3.2% (P) and 0.1% (Q) of speed peaks had a time between peaks less than 200ms and 6.1% (P) and 10.0% (Q) of speed peak pairs had times greater than 1200 ms. These observations suggest the movement behavior could be divided into submovements with similar bell-shaped velocity profiles and similar time durations.

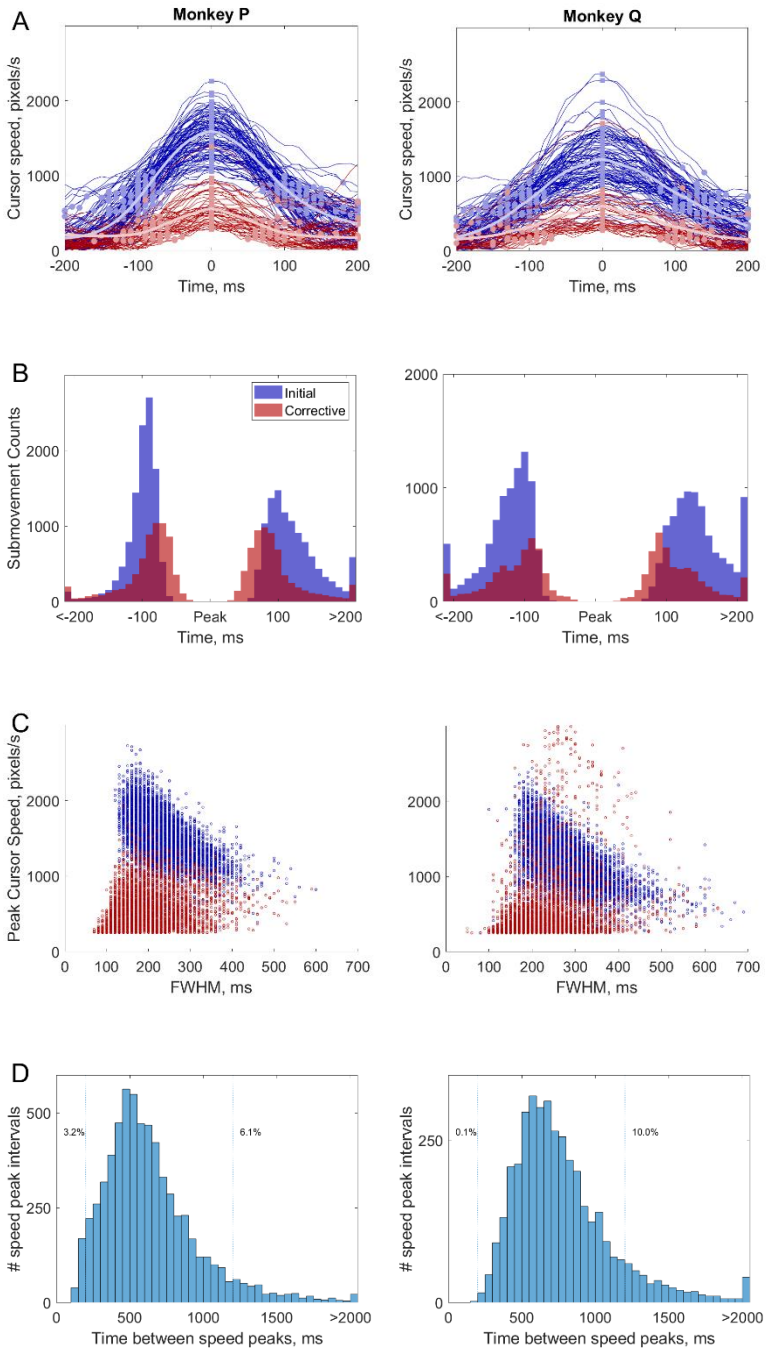

Figure 3. Time course of submovements. A) The cursor speeds are plotted aligned to speed peaks for initial (blue) and corrective (red) submovements. Squares identify speed peaks while circles are located at half the maximum of the speed peak. N.B. The cursor speeds shown are before the bandpass filter used for identifying peaks displayed in Figure 1B. Thus, the maximum of each trace may not align exactly with the plotted peak speed. B) Histogram of the timing of half-maximum speed of all submovements before and after peak speed. C) Peak cursor speed plotted versus the duration of the submovement as

measured by the full time width between half-maximum speed peaks (FWHM). D) The time duration between speed peaks including the times from initial submovement to first corrective submovement as well as between any consecutive pairs of corrective submovements.

Neural Population Firing Rates

The neural firing rates across the recorded population were also analyzed by time aligning to the movement speed peaks and examining the firing rates from 500 ms before until 300 ms after the peak speed. The average number of spikes (binned in 10 ms time bins without square root transform) for all analyzed units aligned to the peak speed for initial and corrective submovements are shown in Figure 4A. A clear peak occurs before the peak speed for both initial and corrective submovements in both monkeys. Monkey P's peak firing rates occurred 180 ms and 110 ms before initial and corrective submovements, respectively, while monkey Q's occurred 180 and 160 ms before. Thus, firing rates increased globally for corrective submovements in addition to the initial reach.

If all neurons had the same time lag that mimics the subsequent peak in movement speed, there would be a synchronized increase and decrease of all condition-independent firing rates simultaneously. However, when examining average firing rates from 10 example neurons from one recording session from monkey P, all aligned to peak speed, we see heterogenous timing of the rise and fall in firing rates relative to the peak speed (Figure 4B). This relationship tended to be conserved across initial and corrective movements with the purple spiking units being earlier firing for initial and corrective submovements while the orange tended to be late. This suggests that the condition-independent neural activity across the neurons might form a repeatable temporal structure—a neural trajectory—that is more complex than a simultaneous rise and fall in firing rate across the population.

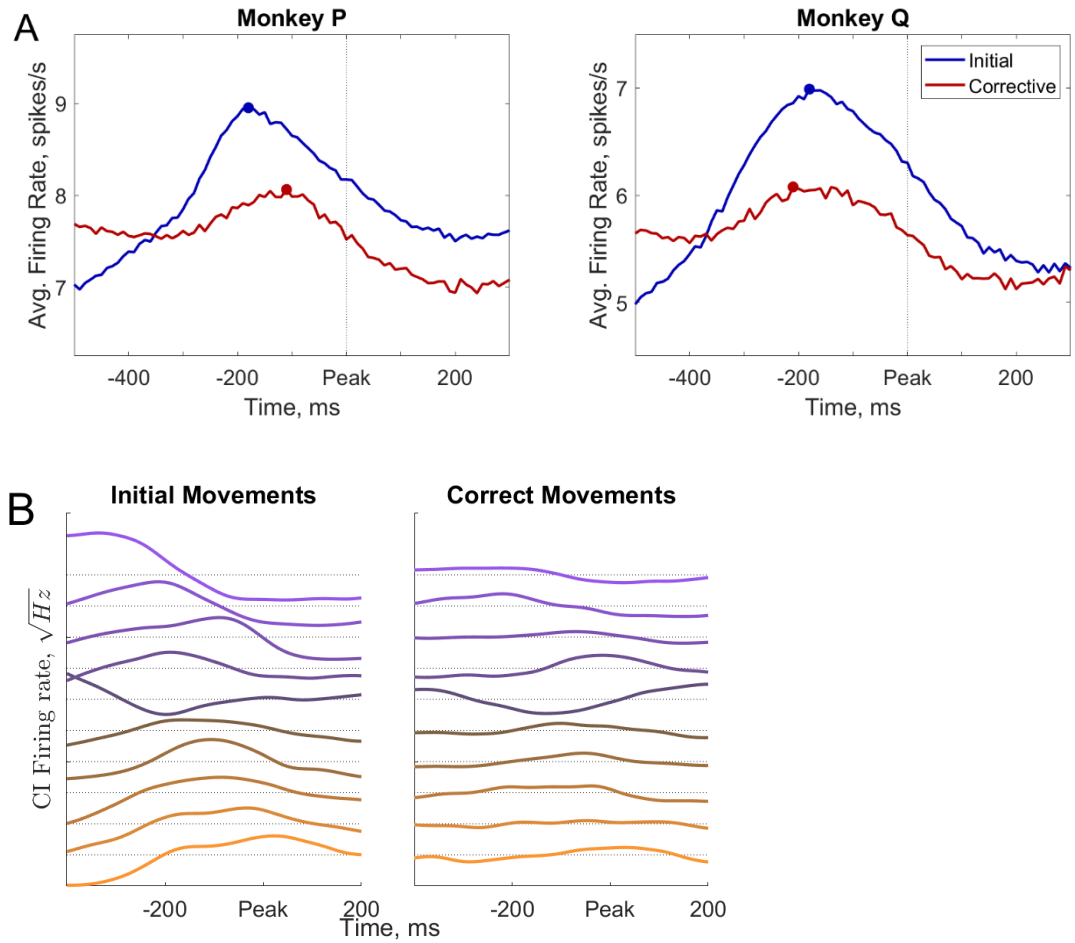

Figure 4. Mean firing rate relative to peak cursor speed. A) The firing rate for all spiking units was averaged for all initial (blue) and corrective (red) submovements. Circles indicate the time of peak firing rate for each condition. B) Average firing rates for 10 example spiking units recorded simultaneously from monkey P time-aligned relative to peak speed for all initial (left) and corrective (right) submovements. Firing rates are shown relative to the average firing rate within the given time window (initial or corrective) for each spiking unit. Units are colored based on the initial movement by whether their firing rates were greater early (purple) or late (orange).

Trajectories in Neural Space

To examine the shape of the condition-independent neural trajectory for initial and corrective submovements, the average difference in firing rate across all neurons between various time points was calculated. If there was simultaneous global modulation of firing rate, all the change in condition-independent firing rates would occur along a single neural dimension and the neural trajectory would form a straight line away from and back to baseline. However, a more complex or cyclic trajectory would indicate repeatable dynamics within the population of neurons. The distance between the condition-independent firing rates is shown in Figure 5. The distances are shown using multidimensional scaling to most closely represent the distances between the high-dimensional neural firing rates in a two-dimensional plot.

The time points of -450, -300, -150, 0, and +150ms relative to peak speed were examined. These times were selected because the highest global firing rate occurred at approximately -150 ms before peak speed, -450 and +150 ms were times of low firing rates, and -300 and 0 ms were times of intermediate firing rates as shown in Figure 3. Yet we observe in Figure 5 that the average change in firing rates between the intermediate times (-300 and 0 ms), was only slightly less than the firing rate change between the early (-450 and -150 ms) or late (-150 and +150 ms) time intervals. Thus, changes in firing rate for individual spiking units did not all synchronously align with the global change in firing rate. This was true for both initial and corrective submovements. This almost circular neural trajectory—a width nearly equal to height—demonstrates that the condition-independent activity does not occur as a single rise and fall in global firing rate, but rather as a neural trajectory with certain neuron's firing rates modulating earlier relative to movement while others modulate later. The neural trajectories are also not perfectly circular but appear to spiral. This is likely due to the changes of firing rate being largest at the onset of movement for the initial movements and then large again at the end of corrective movements. As shown in Figure 4A, this corresponds to rising from a lower baseline at the initial movements and falling back to

a similar low baseline at the end of corrective movements but only reaching an intermediate firing rate at the time between initial and corrective movements.

Additionally, the average change in firing rate per neuron that was observed for each trial was calculated. This change in firing rate includes not just the condition-independent change but also the change dependent on the movement condition as well as any trial-by-trial noise. For initial submovements, the average total change in firing rate across all three time points was 93.2% (monkey P) and 94.0% (Q) with the condition-independent changes being 40.0% (P) and 50.1% (Q). Thus, surprisingly, the average change in firing rate per neuron for each trial was nearly as large for corrective as initial submovements. The smaller amplitude submovements did not result in less change in firing rate within a given trial but only the condition-independent firing rate change when averaged across all submovements.

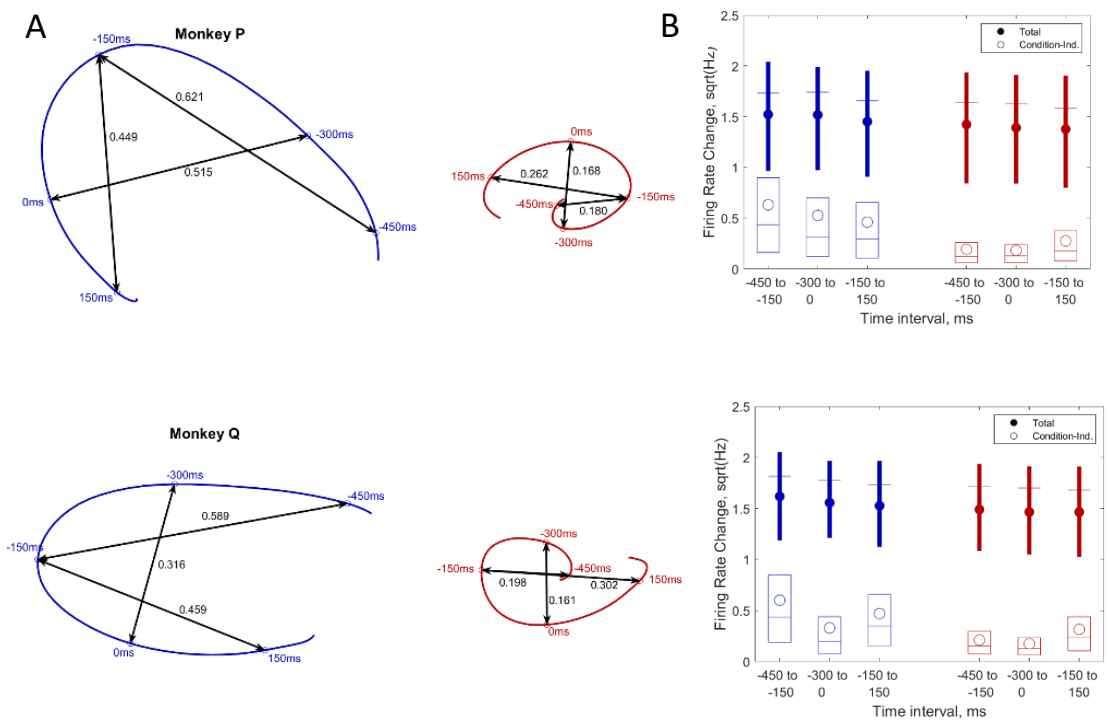

Figure 5. Neural distance between condition-independent firing rates. A) The differences between average firing rates for initial (blue) and corrective (red) for times before (-) and after (+) peak speed (0ms). Distances are shown for three timepoint pairs: -450ms & -150ms, -300 & 0ms, and -150ms & 150ms. The trajectories and distances are shown as a two-dimensional approximation of the high-dimensional neural space using multidimensional scaling. The dimensions for each panel were derived independently with no relationship across panels. B) Changes in firing rate for all analyzed neurons. The change in firing rate when averaged across all trials (Condition-Independent--open circles) and the total trial-by-trial change in firing rate (Total including Condition-Independent, Condition-Dependent, and Noise--closed circles). The circles represent the mean across all neurons while the 25th, 50th, and 75th percentile neuron are shown with the horizontal lines and boxes/whiskers. Changes in firing rate were measured between three time increments relative to peak speed: -450 to -150ms, -300 to 0ms, and -150 to 150 ms. For the corrective movements (red), there was a significant drop in Condition-Independent firing rates compared to initial submovements (blue). However, the overall change in firing rates per trial remained similar. Extended data Figure 5-1 displays similar results but with no square-root transform.

We next wished to examine the firing rate variance and dimensionality of the trial-averaged condition-independent neural activity. The variance of the condition-independent firing rates from a baseline average for corrective submovements was 37.0% (P) and 48.8% (Q) of the condition-independent variance for initial submovements. Principal component analysis was used to assess the dimensionality of the condition-independent activity identified by time-averaging firing rates relative to peak speed. Greater than 97% of all condition-independent variance was observed to occur in 6 dimensions for both animals in both initial and corrective submovements (Figure 6A). In all cases, a single principal component had a large fraction (55-70%) of the variance but the second dimension also had considerable variance (~30% for initial and ~20% for corrective) with the remaining dimensions being relatively small. When comparing the amount of overlap in condition-independent activity between the initial and corrective submovements, there was considerable but not complete overlap i.e. neural trajectories were neither completely parallel or orthogonal to each other. Approximately 80-85% of the corrective condition-independent neural activity occurred in the same 6 dimensions as the initial subspace. Conversely, only 65-70% of the initial condition-independent activity occurred in the corrective subspace with the rest lying in other neural dimensions.

Consistent Neural Dynamics

Despite the smaller magnitude of the condition-independent neural activity during corrective movements, the repeated oscillations in speed and cyclic neural trajectories that overlap in the neural space suggested a portion of neural activity was repeatable and common for both initial and corrective submovements. To examine this possibility further, we built a simple linear system model using the neural firing rates from both initial and corrective submovements to characterize any common temporal dynamics that might be present. The neural firing rates were again averaged across all movement directions to examine only the common condition-independent activity. The plane in the neural space with the most rotational neural activity was identified using jPCA. We labeled the two neural dimensions with the most rotational condition-independent activity as CIx and CIy. To compare consistently across

recording sessions as well as monkeys, the pair of dimensions from each session was rotated such that CIx was aligned most closely with the time course of the global average firing rate across the population (shown in Fig 4A).

The average firing rates for all initial and corrective submovements in our identified CI plane are shown in Figure 6. The neural trajectory in the 2-dimensional CIx/CIy plane are shown in Figure 6B, while the CIx and CIy dimensions are plotted as a function of time in Figure 6C. Considering the initial versus corrective submovements, the peaks and troughs of both CIx and CIy, while differing in magnitude, were similar in their timing. The peak in the CIx dimension—defined as the dimension in the plane that best correlated with the global average firing rate of the population—occurred approximately 150-200 ms before peak speed for both initial and corrective submovements, whereas the peak in the CIy dimension occurred near the time of peak speed for both. The initial and corrective neural trajectories are very similar in their shape and direction of rotation within the plane, with the trajectories for corrective submovements resembling a smaller, scaled version of the larger trajectories for initial submovements.

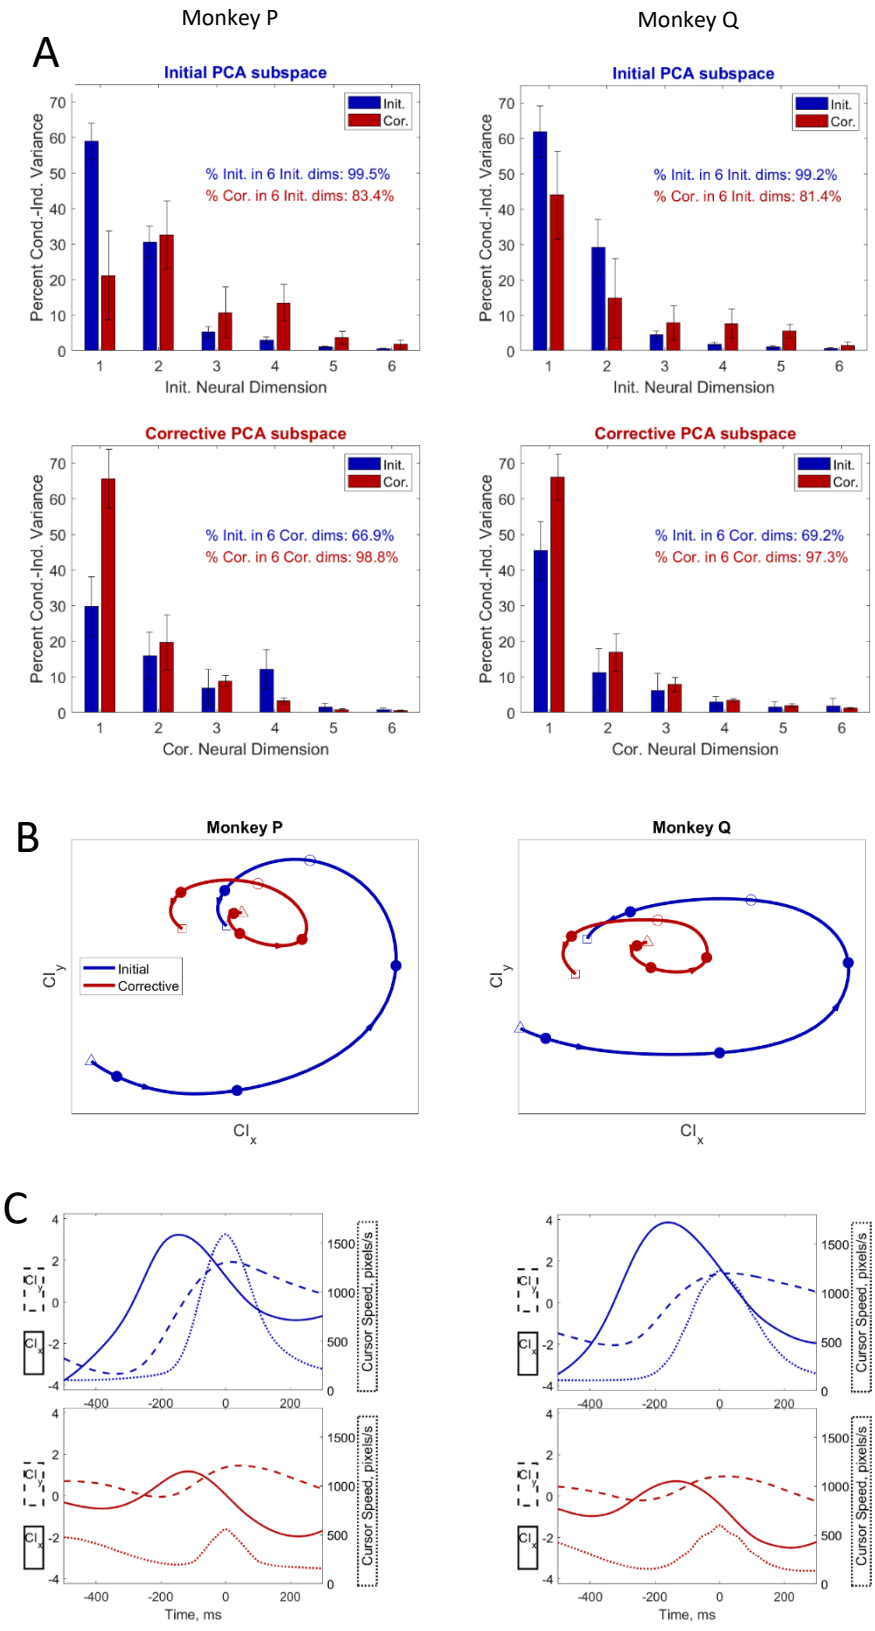
Monkey P
Monkey Q
A
Initial PCA subspace
Corrective PCA subspace
Percent Cond.-Ind. Variance
Init. Neural Dimension
Cor. Neural Dimension
Init.
Cor.
% Init. in 6 Init. dims: 99.5%
% Cor. in 6 Init. dims: 83.4%
% Init. in 6 Init. dims: 99.2%
% Cor. in 6 Init. dims: 81.4%
% Init. in 6 Cor. dims: 66.9%
% Cor. in 6 Cor. dims: 98.8%
% Init. in 6 Cor. dims: 69.2%
% Cor. in 6 Cor. dims: 97.3%
B
Monkey P
Monkey Q
Initial
Corrective
CI_x
CI_y
C
Cursor Speed, pixels/s
Time, ms

Figure 6. A) The percentage of condition-independent variance by dimension using principal component analysis. The top plots show the dimensions in descending variance for the initial submovements while the bottom plots show descending dimensions for the corrective submovements. B) The average firing rate for initial and corrective submovements are plotted in the CIx/CIy plane identified with jPCA. The trajectories start at the triangles and end at the squares. Each filled circle is a 150 ms time step. Extended data Figure 6-1 displays similar results but with no square-root transform. C). Average CIx (solid lines) and CIy (dashed lines) plotted as a function of time relative to average cursor speed (dotted lines).

We then examined the instantaneous phases of CIx and CIy activity, calculated using the Hilbert transform. The average phase of this neural activity is shown in Figure 7A. The phase of CIx (solid lines) and that of CIy (dashed lines) each were similar for both initial and corrective submovements, with the zero phase of CIx occurring about 150 ms before the peak speed while CIy lagged by approximate $\pi/2$ with the zero crossing around peak speed. Histograms of the phase of CIx and of CIy at peak speed on individual trials are shown in Figure 7B. Mean values of CIx were 0.34π (initial) and 0.43π (corrective) for monkey P and 0.31π (initial) and 0.45π (corrective) for monkey Q. Mean values of CIy were -0.14π (initial) and -0.15π (corrective) for monkey P and -0.15π (initial) and -0.06π (corrective) for monkey Q. Thus, there was a clear relationship between peak speed and the phase of condition-independent activity. CIx leads peak speed by $\sim \pi/2$ whereas CIy peaks close to peak speed. Similarly, the time when the phase of CIx = 0 and CIy = $-\pi/2$ was observed to occur 150-200 ms before the peak speed for most initial and corrective submovements.

Since the phase in the CI plane appeared to define the neural dynamics and predict upcoming speed peaks, we created a metric we call the condition-independent phase (CIφ) by averaging the phase of CIx and phase of CIy + $\pi/2$ to calculate the current phase in the CI plane. We then examined the percentage of the total neural variance explained when the condition-independent activity was defined by CIφ compared to trial averaging based on time-alignment. Note, this percentage of condition-independent activity is relative to the total change in firing rates which also includes condition-dependent changes as well as trial-by-trial variability. As shown in Figure 7C, more neural variance was captured by binning

the data based on CIφ than by binning the data based on time relative to speed peaks. While the difference was modest for initial submovements, a larger, statistically significant increase occurred for corrective submovements with the variance explained going from 2.9% for time-aligned to 5.9% for CIφ-aligned for monkey P and 2.2% to 4.2% for monkey Q (Table 1). The amount of the CIφ explained variance that occurred only in the CIx and CIy dimensions was calculated next. For the two monkeys and the two conditions—initial and corrective—the variance in the two dimensions, CIx and CIy, accounted for between 54% and 76% of the condition-independent variance in the four instances. The CIx dimension accounted for more in all instances with the CIy being nearly equal for monkey P initial submovements while the biggest disparity occurred for monkey Q corrective submovements with the CIy variance being 37% of the CIx variance. Altogether, these observations suggest that compared to aligning condition-independent neural activity across submovements based on time relative to peak speed, more consistent alignment is achieved using CIφ.

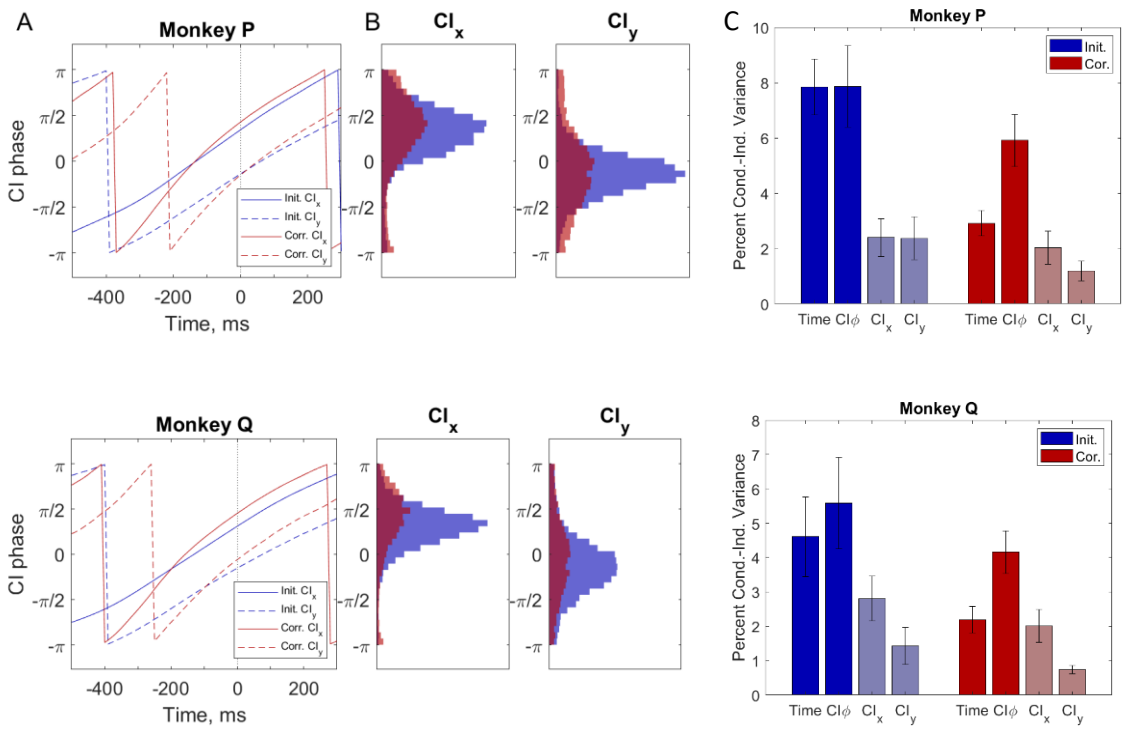

Figure 7. Phase of CIx and CIy relative to peak cursor speed. A) Average phase of CIx (solid lines) and CIy (dashed lines) time-aligned to peak speed (Time = 0) for all initial (blue) and corrective (red) submovements. B) Histogram of the phase of CIx and CIy at the time of peak speed for initial (blue) and corrective (red) submovements. All distibutions of CIx and CIy were non-uniformly distributed at peak speed (Rayleigh test, all $p < 0.001$). C) Percentage of total variance explained by condition-independent neural activity when the condition-independent activity is calculated by i) time averaging (Time) or ii) averaging by condition-independent phase (CIφ). The condition-independent variance explained was greater for statistical greater for CIφ alignment in all conditions except initial submovements for monkey P (Table 1). The percentage of the CIφ variance that lies in the two dimensions CIx and CIy identified by the jPCA algorithm is also shown. Error bars represent the standard deviation across the 12 sessions.

Table 1. Percent of total neural variance that is condition-independent when using Time-alignment vs. CIφ alignment.

	Time-alignment Median	CIφ - alignment Median	Paired Median Difference	95% Confid. interval	Data Comparison	Statistical Test
Initial					Paired median difference across 12 sessions	Two-sided permutation t-test
Monkey P	7.84	7.86	0.34	[-0.21, 0.49]		
Monkey Q	4.61	5.59	0.85	[0.62, 1.19]		
Corrective						
Monkey P	2.91	5.91	3.03	[2.53, 3.44]		
Monkey Q	2.19	4.16	2.06	[1.73, 2.25]		

We therefore examined the relationship between CIφ and hand movement. Figure 8A illustrates the temporal distribution of when CIφ =0 relative to the time of peak hand speed. These distributions consistently peaked 100-150 ms before the speed peak for both initial and corrective submovements. A relatively consistent relationship between neural activity in this CIx/CIy plane and peak speed thus was present for both initial and corrective submovements across all trials regardless of target size or reach direction. We next compared these distributions with distributions of the time of peak firing rate across the population relative to the time of peak speed. For the comparison we built a simple linear regression model to estimate peak speed with a weighted sum of the instantaneous firing rate of all spiking units. The weights were positive for units that increased their firing rates with movement and negative for units that decreased their firing rates with movement. Using this firing rate model, we calculated when the peak firing rate across the neural population occurred. Figure 8B then shows the temporal distributions these peak firing rates relative to peak speed for both initial and corrective submovements. Like the distributions using the dynamical model above (Figure 8A), those for the firing rate model peaked 150 to 100 ms before peak speed. The peaks were broader by 10-20 ms, however, as characterized by the corresponding standard deviations (σ) given for each distribution. Furthermore, although ≥ 84% of

submovements were included in each of these distributions (percentages are given in Fig. 8), a small fraction of submovements could not be aligned, lacking a CIφ=0 in the dynamical systems model and/or a peak in the firing rate model within the -300 to 100 ms time window examined, and the percentage of these unaligned trials was consistently smaller for the dynamical systems model. Compared to using only the synchronous firing rates, using the temporal dynamics of the neural population thus improved the accuracy and consistency with which the time of peak speed could be predicted.

Finally, we examined directly the relationship between cursor speed and the phase of cyclic neural activity in the identified rotational CIx/CIy plane. In figure 8C, we have plotted the cursor speed as a function of CIφ. While the CIφ is an angle that ranges between +/-π radians when calculated, for purposes of display here we incremented CIφ in steps of 2π to show how successive cycles of neural activity (abscissa) were related to movement speed (ordinate) as individual trials progressed through both initial and subsequent corrective submovements. The speed of movement clearly oscillated with the cyclic neural activity, the cursor speeds for most trials rising and falling in 2π cycles of CIφ. The circular correlation between speed and CIφ was 0.44 and 0.42 ($p<0.001$) with the largest speeds occurring at CIφ = 0.32π and 0.31π (+2kπ) for monkeys P and Q, respectively. Both the speed averaged across trials (white) and the non-uniform occurrence of peak speeds in individual trials (black) demonstrate that movement speed was consistently correlated with the cycles of condition-independent neural activity.

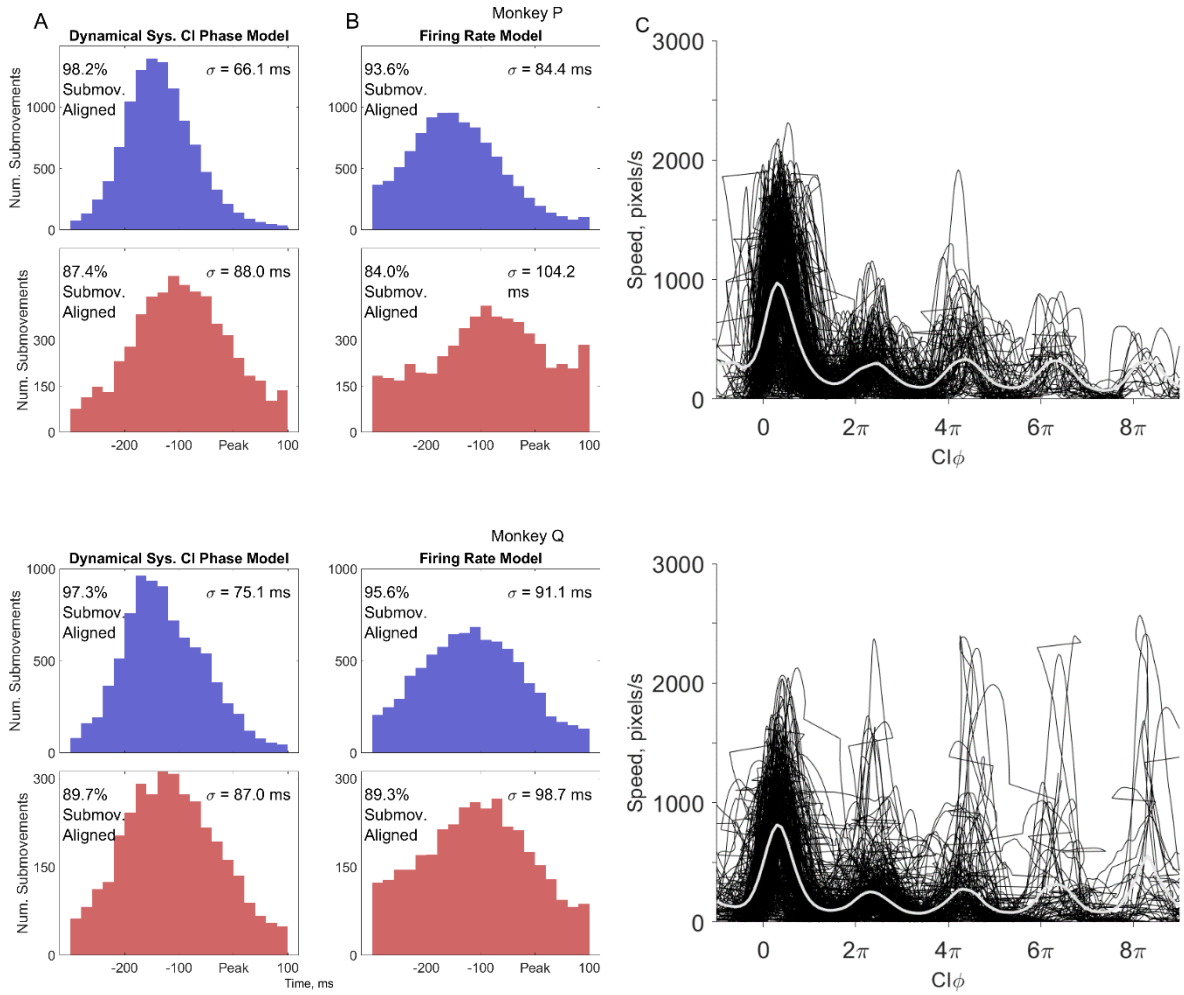

Figure 8. Relationship between CIφ and cursor speed. Identifying the times of peak speeds with a dynamical systems model (A) or with an instantaneous firing rates (B). The time point when CIφ = 0 (A) or peak firing rate (B) was used as a prediction of the upcoming submovement. Each histogram shows only those submovements for which the neural data aligned with the movement data, i.e. CIφ = 0 (A) or maximum firing rate (B) occurred within the time range examined (-300 to 100 ms relative to submovement peak speed). The percentage of total aligned trials is shown for each distribution as well as the standard deviation (σ) for the aligned trials. In all cases, the dynamical systems model predictions were more precise, with a narrower standard deviation (F-test, all $p<0.001$ for both initial and corrective submovements for both monkeys, see Table 1) and fewer unaligned trials. Extended data Figure 8-1 displays predictions from the CIφ model with no square-root transform with slightly worse predictions but still improved compared to the firing rate model. C) Cursor speed is plotted as a function of CIφ for 200 trials with at least one corrective submovement. The average speed of all trials as a function of CIφ is shown in white, illustrating the oscillation in cursor speed depending on the phase of neural activity.

The circular correlations between CIφ and cursor speed for all corrective trials were 0.43 and 0.43 for monkeys P and Q, respectively, $p<0.001$ in both cases. Note, the unwrapped CIφ is not always a monotonically increasing value as occasionally the neural activity could reverse and move clockwise rather than counter-clockwise in the neural plane shown in Figure 6B.

Table 2. Comparison of predication accuracy as measured with standard deviation in predictions using the dynamical system CIφ model vs. an instantaneous firing fate model.

	σ_1, CIφ, (ms)	σ_2, Firing Rate (ms)	F-stat, $\frac{\sigma_1^2}{\sigma_2^2}$	95% Confid. interval	Data Comparison	Statistical Test
Initial					All submovements with a prediction between -300:100 ms, assuming normal distribution	F-test
Monkey P	66.1	84.4	0.61	[0.59, 0.64]		
Monkey Q	75.1	91.1	0.69	[0.65, 0.71]		
Corrective						
Monkey P	88.0	104.2	0.71	[0.68, 0.75]		
Monkey Q	87.0	98.7	2.06	[0.73, 0.83]		

Discussion

Our precision center-out task utilized small targets to elicit many trials that were not completed in an initial reach and required one or more corrective submovements. We found a temporal relationship between both initial and corrective reaching movements with cyclic, condition-independent neural activity. The speed profiles of both the initial and corrective submovements aligned with cyclic neural activity in primary motor cortex providing a useful neural marker encoding a series of submovements. Discrete submovements defined by multiple speed peaks have previously been described in behavioral studies of reaching (Pratt et al., 1994; Lee et al., 1997; Hatsopoulos et al., 2007; Polyakov et al., 2009), turning a knob (Novak et al., 2000), isometric contractions (Massey et al., 1992; Hall et al., 2014), and object manipulation for tactile discrimination (Pruszynski et al., 2018). The experimental results and

analysis presented here provide new evidence of a relationship between previously reported neural dynamics (Churchland et al., 2012; Hall et al., 2014) and behaviorally observed submovements.

Distinct movement speed peaks with corrective movements

In our precision center-out task, the monkeys' movements showed consistent bell-shaped speed profiles. These speed profiles were evident for both the larger initial movements from the center toward the target as well as the subsequent corrective movements. A large majority of both initial and corrective submovements had durations of 100-350 ms with a low-speed trough between almost all submovements. For initial movements, the movement duration was longer for slower reach speed, which was likely due to the reach target being at the same distance for all trials. For corrective movements, there was no relationship between movement duration and speed with all corrective submovements having approximately the same duration regardless of speed.

Similar but smaller condition-independent activity for corrective movements

The orientation and rotation direction of the cyclic, condition-independent neural modulation was similar between initial and corrective submovements and occupied overlapping dimensions of the neural space. The magnitude of the condition-independent neural activity that occurred for corrective submovements was approximately one-third to one-half the magnitude of that for the initial submovements. This decreased change in average firing rate approximately matched the decrease in movement speed for the corrective compared to the initial submovements suggesting speed tuning of the condition-independent activity. This does not imply, however, that each individual trial and each individual neuron have smaller changes of firing rate during smaller amplitude movements. In fact, the overall total change in firing rates—the combined condition-independent activity, condition-dependent encoding of task variables, and unexplained noise—for corrective submovements was of nearly the same magnitude (90-95%, Figure 5) as that for initial submovements. Only when averaging firing rates—time aligned to the peak movement speed or the decoded condition-independent phase—were the differences

in firing rate modulation between initial and corrective movements apparent. On average, the encoding of movement speed is clearly present in primary motor cortex (Moran and Schwartz, 1999; Paninski et al., 2004). By identifying the neural dimensions with condition-independent activity, it is possible to identify this speed dependent encoding. Yet precisely identifying encoded speed on a trial-by-trial basis with the neural activity remains challenging as there are often large changes of firing rates for individual neurons that are more variable and idiosyncratic during any particular corrective submovement.

Condition-independent phase predictive of cursor speed

Since rotational neural dynamics were presented by Churchland et al. (2012), considerable debate has existed about their importance and origin. While originally presented as similar rotations but with different starting points and radial magnitudes for varying conditions, others have pointed out the how neurons with various but consistent temporal lags (Lebedev et al., 2019) may explain many of the results. More recently, Zimnik and Churchland (2021) demonstrated two repeated cycles of neural activity, each shortened in time, when a pair of movements were simultaneously instructed to be performed in rapid succession. These faster neural dynamics for a sequence of two movements appear similar to our observation that cycles of neural activity being quicker for the corrective than initial submovements (Fig. 6). Here, by focusing on condition-independent neural activity, we found that similar cycles of neural activity appear for not only initial planned reaches but also for the highly variable, corrective submovements that are made online with visual feedback. Our results highlight i) the various time lags between individual cortical neurons' firing and the upcoming movement and ii) the complex neural trajectories in a high-dimensional space, both accompany reaching movements whether they are large or small and whether instructed and planned or made online with feedback.

The cyclic CIφ values aligned with cursor speed, confirming that the observed cyclic neural dynamics were behaviorally relevant with the cursor speed peaks and troughs corresponding to particular CIφ values. Separate corrective submovements thus are encoded by sequential cycles of neural activity.

Rather than continuously moving the hand toward the target at all times, the subjects' movements appeared as pulsatile bursts of joystick movement and presumably muscle contraction. This work supports the notion of movement fragments (Hatsopoulos et al., 2007) and extends it to a task with a larger range of movement amplitudes including small, corrective movements. Additionally, our CIφ algorithm provides a useful way to identify the timing of each movement fragment by utilizing recent advances in applying dynamical system modeling to neural populations (Churchland et al., 2012). Large condition-independent activity has been suggested to make brain dynamics more robust to noise by increasing the differences in neural signals even when the muscle activation pattern at certain time points look very similar (Russo et al., 2018). In the context of precise, corrective movements, one might speculate cyclic brain dynamics organize movements into distinct submovements to create time-varying neural and musculoskeletal dynamics that are more reliable for motor control. Additionally, by having larger changes in firing rate even for corrective movements that occupy different dimensions of the neural space, the larger separation of the neural signals for different movements may be hypothesized to encode the desired movement more accurately, either large initial or small corrective.

It is difficult to determine how much of the cyclic brain activity and behavior observed is caused by intrinsic neural dynamics in M1 compared to sensorimotor feedback control and/or a cognitive strategy. When the behavior has sufficient time delay between each submovement defined by each speed peak as occurred in our task (Fig. 3D), the neural activity can fit both descriptions. Observations of additional submovements defined by second or third speed peaks do not necessarily require a feedback controller with discrete updates. A single, continuous optimal feedback controller with appropriate delays and signal dependent noise can also generate additional submovements with multiple speed peaks (Li et al., 2018). Results by Susilaradeya et al. (2019) argue that extrinsic effects of a task interact with the intrinsic dynamics of the brain in a manner consistent with an optimal feedback controller and may provide a framework for assessing these effects across a variety of tasks including our precision center-out task. Further work examining neural activity in various tasks and/or additional sensorimotor brain

areas is needed for clearer descriptions of the neural dynamics of the sensory processing, cognitive planning, and motor execution for precise, corrective movements.

It also is unclear how consistent our observed condition-independent cyclic neural dynamics are for other forelimb movements besides reaching confined to a single limb. For instance, Suresh et al. (2020) showed low-dimensional cyclic dynamics appear to occur much less for grasping than reaching with the authors arguing that the different muscle groups and amount of increased afferent input required for grasping may explain these differences. Additionally, the increased degrees-of-freedom and diversity of hand grasping shapes compared to planar reaching may prevent a simple low-dimensional neural representation. The neural signals in a given hemisphere for cyclic movements of the contra- and ipsilateral arms have also have been reported to be in orthogonal subspaces (Ames and Churchland, 2019) raising additional questions about how consistent neural dynamics might be across behaviors. Previous work has typically not focused on how close neural dynamics align with behavior and not analyzed only condition-independent neural dynamics. The different analytical methods for previous studies make it difficult to determine how well our algorithm might generalize and how well it identifies truly condition-independent activity on very different tasks. However, firing rate changes independent of movement features with various time lags seem likely to be present across various tasks. Thus, we hypothesize the portion of cyclic dynamics identified with our condition-independent focused algorithm might be more common and robust than other rotational neural dynamics and have a more observable direct link to behavior.

The cyclic dynamics of corrective movements have important implications for brain-computer interfaces (BCI). To date, most BCI decoders are time-invariant when individual submovements occur. Decoders are typically first constructed from observed or imagined movements that assume single, straight-line movements. When algorithms for updating BCI decoders consider the change in movement direction for corrective movements, it is assumed the intended path is updated continuously (Gilja et al., 2012; Shanechi et al., 2016). Experiments have suggested that BCI control can be improved with two

states: active control and rest (Kim et al., 2011; Williams et al., 2013, 2016; Sachs et al., 2016). The ability to predict the phase of a movement—like CIφ presented here—might further allow for much better prediction of when the subject intends to move, contact the target, and make a corrective movement. The cyclic dynamics of the condition-independent neural activity may also lead to better descriptions of the condition-dependent activity that encode task features, specifically the observation that direction encoding is often not fixed but shifts with different phases of movement (Sergio and Kalaska, 1998; Churchland and Shenoy, 2007; Suminski et al., 2015; Suway et al., 2017), enabling decoders of movement direction that change with time. All leading to a more robust description of the of neural encoding of precise and corrective movements.

Extended Data 1. Matlab code to calculate the CIφ is available on Github.

Since the trial data contains corrective movements in addition to the large initial movements that were not precisely time aligned to trial events for averaging condition-independent neural activity, we developed a novel algorithm to iteratively average the firing rates, calculate CIφ, then average the firing rates again based on the CIφ. This iterative process involves three steps: i) Each unit's firing rate is averaged across all trials to determine its condition-independent firing rate. ii) Dimensionality reduction is performed using PCA and jPCA on the condition-independent firing rates to identify the neural plane with the most rotational/cyclic condition-independent activity. iii) The instantaneous phase is calculated using the Hilbert transform on the first two jPC dimensions for all data points. Matlab code is available on GitHub. Further details are available in the Readme document attached to the code.

References

Abrams RA, Meyer DE, Kornblum S (1990) Eye-hand coordination: oculomotor control in rapid aimed limb movements. J Exp Psychol Hum Percept Perform 16:248–267.

Ames KC, Churchland MM (2019) Motor cortex signals for each arm are mixed across hemispheres and neurons yet partitioned within the population response. Elife 8:1–36.

Ashe J, Georgopoulos AP (1994) Movement parameters and neural activity in motor cortex and area 5. Cereb Cortex 6:590–600.

Berens P (2009) CircStat : A MATLAB Toolbox for Circular Statistics. J Stat Softw 31:1–21.

Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu S, Shenoy K V (2012) Neural population dynamics during reaching. Nature 487:51–56.

Churchland MM, Shenoy K V (2007) Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. J Neurophysiol 97:4235–4257.

Craik KJW (1947) Theory of the human operator in control systems. I. the operator as an engineering system. Br J Psychol 38:56–61.

Cunningham JP, Yu BM (2014) Dimensionality reduction for large-scale neural recordings. Nat Neurosci 17:1500–1509.

Elliott D, Hansen S, Grierson LE, Lyons J, Bennett SJ, Hayes SJ (2010) Goal-Directed Aiming: Two Components but Multiple Processes. Psychol Bull 136:1023.

Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP (2016) Reorganization between preparatory and movement population responses in motor cortex. Nat Commun 7:1–15.

Evarts E V (1968) Relation of pyramidal tract activity to force exerted during voluntary movement. J

Neurophysiol 31:14–27.

Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci 2:1527–1537.

Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal Population Coding of Movement Direction. Science (80-) 233:1416–1419.

Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Fan JM, Ryu S, Shenoy K V (2012) A brain machine interface control algorithm designed from a feedback control perspective. Conf Proc IEEE Eng Med Biol Soc 2012:1318–1322.

Hall TM, de Carvalho F, Jackson A (2014) A common structure underlies low-frequency cortical dynamics in movement, sleep, and sedation. Neuron 83:1185–1199.

Hatsopoulos NG, Xu Q, Amit Y (2007) Encoding of movement fragments in the motor cortex. J Neurosci 27:5105–5114.

Jackson A, Gee VJ, Baker SN, Lemon RN (2003) Synchrony between Neurons with Similar Muscle Fields in Monkey Motor Cortex. Neuron 38:115–125.

Kakei S, Hoffman DS, Strick PL (1999) Muscle and movement representations in the primary motor cortex. Science 285:2136–2139.

Kalaska JF, Cohen DAD, Hyde ML, Prud'homme M (1989) 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.

Kaufman MT, Seely JS, Sussillo D, Ryu SI, Shenoy K V, Churchland MM (2016) The largest response component in motor cortex reflects movement timing but not movement type. eNeuro

3:ENEURO.0085-16.2016.

Kihlberg JK, Herson JH, Schotz WE (1972) Square Root Transformation Revisited. J R Stat Soc 21:76–81.

Kim S-P, Simeral JD, Hochberg LR, Donoghue JP, Friehs GM, Black MJ (2011) Point-and-click cursor control with an intracortical neural interface system by humans with tetraplegia. IEEE Trans neural Syst Rehabil Eng 19:193–203.

Kobak D, Brendel W, Constantinidis C, Feierstein CE, Kepecs A, Mainen ZF, Qi XL, Romo R, Uchida N, Machens CK (2016) Demixed principal component analysis of neural population data. Elife 5:1–36.

Lebedev MA, Ossadtchi A, Mill NA, Urpí NA, Cervera MR, Nicolelis MAL (2019) Analysis of neuronal ensemble activity reveals the pitfalls and shortcomings of rotation dynamics. Sci Rep 9:1–14.

Lee D, Port NL, Georgopoulos AP (1997) Manual interception of moving targets. Exp brain Res 116:421–433.

Li Z, Mazzoni P, Song S, Qian N (2018) A Single, Continuously Applied Control Policy for Modeling Reaching Movements with and without Perturbation. Neural Comput 30:397–427.

Massey JT, Lurito JT, Pellizzer G, Georgopoulos AP (1992) Three-dimensional drawings in isometric conditions: relation between geometry and kinematics. Exp Brain Res 88:685–690.

Mathworks (2020) Natick, Massachusetts: The MathWorks Inc.

Maynard EM, Hatsopoulos NG, Ojakangas CL, Acuna BD, Sanes JN, Normann R a, Donoghue JP (1999) Neuronal interactions improve cortical population coding of movement direction. J Neurosci 19:8083–8093.

Moran DW, Schwartz AB (1999) Motor Cortical Representation of Speed and Direction During

Reaching. J Neurophysiol 82:2676–2692.

Novak KE, Miller LE, Houk JC (2000) Kinematic properties of rapid hand movements in a knob turning task. Exp Brain Res 132:419–433.

Paninski L, Fellows MR, Hatsopoulos NG, Donoghue JP (2004) Spatiotemporal tuning of motor cortical neurons for hand position and velocity. J Neurophysiol 91:515–532.

Polyakov F, Stark E, Drori R, Abeles M, Flash T (2009) Parabolic movement primitives and cortical states: Merging optimality with geometric invariance. Biol Cybern 100:159–184.

Pratt J, Chasteen AL, Abrams RA (1994) Rapid aimed limb movements: age differences and practice effects in component submovements. Psychol Aging 9:325–334.

Pruszynski JA, Flanagan JR, Johansson RS (2018) Fast and accurate edge orientation processing during object manipulation. Elife 7:e31200.

Rouse AG, Schieber MH (2018) Condition-dependent neural dimensions progressively shift during reach to grasp. Cell Rep 25:3158–3168.

Russo AA, Bittner SR, Perkins SM, Seely JS, London BM, Lara AH, Miri A, Marshall NJ, Kohn A, Jessell TM, Abbott LF, Cunningham JP, Churchland MM (2018) Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. Neuron 97:953-966.e8.

Sachs NA, Ruiz-Torres R, Perreault EJ, Miller LE (2016) Brain-state classification and a dual-state decoder dramatically improve the control of cursor movement through a brain-machine interface. J Neural Eng 13:016009.

Sainburg RL, Ghez C, Kalakanis D (1999) Intersegmental dynamics are controlled by sequential anticipatory, error correction, and postural mechanisms. J Neurophysiol 81:1045–1056.

Salinas E, Abbott LF (1994) Vector reconstruction from firing rates. J Comput Neurosci 1:89–107.

Sarma S V., Eden UT, Cheng ML, Williams ZM, Hu R, Eskandar E, Brown EN (2010) Using point process models to compare neural spiking activity in the subthalamic nucleus of parkinsons patients and a healthy primate. IEEE Trans Biomed Eng 57:1297–1305.

Sergio LE, Kalaska JF (1998) Changes in the Temporal Pattern of Primary Motor Cortex Activity in a Directional Isometric Force Versus Limb Movement Task. J Neurophysiol 80:1577–1583.

Shanechi MM, Orsborn AL, Carmena JM (2016) Robust Brain-Machine Interface Design Using Optimal Feedback Control Modeling and Adaptive Point Process Filtering. PLoS Comput Biol 12:1–29.

Snedecor G, Cochran W (1980) Statistical Methods, 7th ed. Ames, IA: Iowa State UP.

Suminski AJ, Mardoum P, Lillicrap TP, Hatsopoulos NG (2015) Temporal evolution of both premotor and motor cortical tuning properties reflect changes in limb biomechanics. J Neurophysiol 113:2812–2823.

Suresh AK, Goodman JM, Okorokova E V., Kaufman M, Hatsopoulos NG, Bensmaia SJ (2020) Neural population dynamics in motor cortex are different for reach and grasp. Elife 9:1–16.

Susilaradeya D, Xu W, Hall TM, Galán F, Alter K, Jackson A (2019) Extrinsic and intrinsic dynamics in movement intermittency. Elife 8:1–27.

Suway SB, Orellana J, McMorland AJC, Fraser GW, Liu Z, Velliste M, Chase SM, Kass RE, Schwartz AB (2017) Temporally Segmented Directionality in the Motor Cortex. Cereb Cortex 12:1–14.

Thach WT (1978) Correlation of Neural Discharge with Pattern and Force of Muscular Activity, Joint Position, and Direction of Intended Next Movement in Motor Cortex and Cerebellum. J Neurophysiol 41:654–676.

Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN (2005) A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate

Effects. J Neurophysiol 96:1074–1089.

Williams JJ, Rouse AG, Thongpang S, Williams JC, Moran DW (2013) Differentiating closed-loop cortical intention from rest: building an asynchronous electrocorticographic BCI. J Neural Eng 10:046001.

Williams JJ, Tien RN, Inoue Y, Schwartz AB (2016) Idle state classification using spiking activity and local field potentials in a brain computer interface. 2016 38th Annu Int Conf IEEE Eng Med Biol Soc:1572–1575.

Woodworth RS (1899) The accuracy of voluntary movement. Psychol Rev 3:1–119.

Yu BM, Kemere C, Santhanam G, Afshar A, Ryu S, Meng TH, Sahani M, Shenoy K V (2007) Mixture of trajectory models for neural decoding of goal-directed movements. J Neurophysiol 97:3763–3780.

Zimnik AJ, Churchland MM (2021) Independent generation of sequence elements by motor cortex. Nat Neurosci 24:412–424.

Calculation of jPCA plane and Condition-Independent Phase

Since the trial data contains corrective movements in addition to the large initial movements that were not precisely time aligned to trial events for averaging condition-independent neural activity, we developed a novel algorithm to iteratively average the firing rates, calculate CIφ, then average the firing rates again based on the CIφ. This iterative process involves three steps: i) Each unit's firing rate is averaged across all trials to determine its condition-independent firing rate. ii) Dimensionality reduction is performed using PCA and jPCA on the condition-independent firing rates to identify the neural plane with the most rotational/cyclic condition-independent activity. iii) The instantaneous phase is calculated using the Hilbert transform on the first two jPC dimensions for all data points. A schematic of the iterative algorithm is shown below.

i) Trial averaging to identify condition-independent activity

The condition-independent neural activity is the average firing rate for each recorded spiking unit for all experimental trials regardless of the movement condition (ie target location). For classic neurophysiology experiments, this can be calculated by averaging time-aligned data. However, since trials were of varying durations and many included corrective movements, simple time alignment of trials was difficult. Data was therefore aligned based on each time-point's calculated CIφ rather than absolute time.

Since CIφ depends on the averaged condition-independent activity and the condition-independent activity was averaged by CIφ alignment, an iterative approach was required. An initial estimate using simple time-alignment averaging of all initial and corrective submovements was used for the first iteration. For each subsequent iteration, the condition-independent firing rates for each unit were calculated by averaging all data points when the CIφ values were similar. The average firing rate was estimated using a sliding window of CIφ values with a step size of $\pi/50$ and a window size of $\pi/25$ to generate 100 equally spaced samples ranging from $-\pi$ to π.

ii) PCA and jPCA to identify rotational/cyclic neural activity

Next, the condition-independent firing rates were submitted to PCA and jPCA (Churchland et al., 2012) to identify the neural dimensions with the most cyclic activity. PCA was performed on the high-dimensional neural space to reduce the condition-independent firing rates to the six dimensions with the most variance. jPCA was then performed on this six-dimensional space. jPCA is a dimensionality reduction technique to identify the neural planes with the most rotational activity and is more fully described in Churchland et al (2012). Briefly, jPCA fits a first order dynamical system model to the neural activity : $\dot{X} = AX$

to predict change in firing rate $(\dot{X})$ based on the current firing rate (X). The transform matrix of this model (A) can be separated into a symmetric matrix representing pure scaling and a skew-symmetric matrix representing pure rotational dynamics. By taking the eigendecomposition of the skew-symmetric matrix, we obtain pairs of purely imaginary eigenvalues and corresponding eigenvectors that define planes of rotation in the neural space rank-ordered from greatest to least rotation. In the present analysis, only the first plane with the greatest condition-independent rotation was used and we defined the two dimensions of the plane as CIx and CIy. Additionally, to obtain a consistent CIφ across recording sessions, CIx was defined as the dimension in the CIx/CIy plane with the most variance. The positive CIx direction was defined as having more positive than negative coefficients in the neural space which corresponds to the direction where more units have an increased firing rate. Choosing this convention causes an increase in CIx to generally align with the onset of initial movement, since a majority of units increase firing rates at the onset of movement.

iii) Instantaneous phase estimate

Finally, the instantaneous phase were estimated by i) bidirectional bandpass filtering of the activity in both the CIx and CIy dimensions between 0.5-5 Hz with a 1st order Butterworth filter, ii) performing the Hilbert transform of both filtered signal (s_x and s_y) to generate a transformed signal

($\hat{s}_x$ and $\hat{s}_y$) that is a 90˚ phase shift of every Fourier component in the frequency domain to create an analytic representation of CIx and CIy, and iii) calculating the angle of the resulting analytic signal ($s_1 + i\hat{s}_1$) to estimate the instantaneous phase in each dimension, and iv) which we then sum (with a $\frac{\pi}{2}$ phase shift added to ϕy) to obtain a single instantaneous phase estimate for the neural activity within the plane which we call the condition-independent phase (*CIφ*).

i) $s_x(t) = filt_{0.5-5Hz}[CIx(t)]$;

ii) $\hat{s}_x(t) = \mathcal{H}[s_x(t)]$

iii) $\phi x\,(t) = \arg\{s_x(t) + i\hat{s}_x(t)\}$

iv) $CI\phi\,(t) = \phi x + (\phi y + \frac{\pi}{2})$

$Note -$ i, ii, and iii are performed for both x and y

The bandpass filtering reduces the low-frequency drift and high-frequency variability in CIx to generate a more consistent subsequent phase estimate. The condition-independent phase (CIφ) thus represents the instantaneous phase in the dimension of the neural space that has the most cyclic, condition-independent activity.

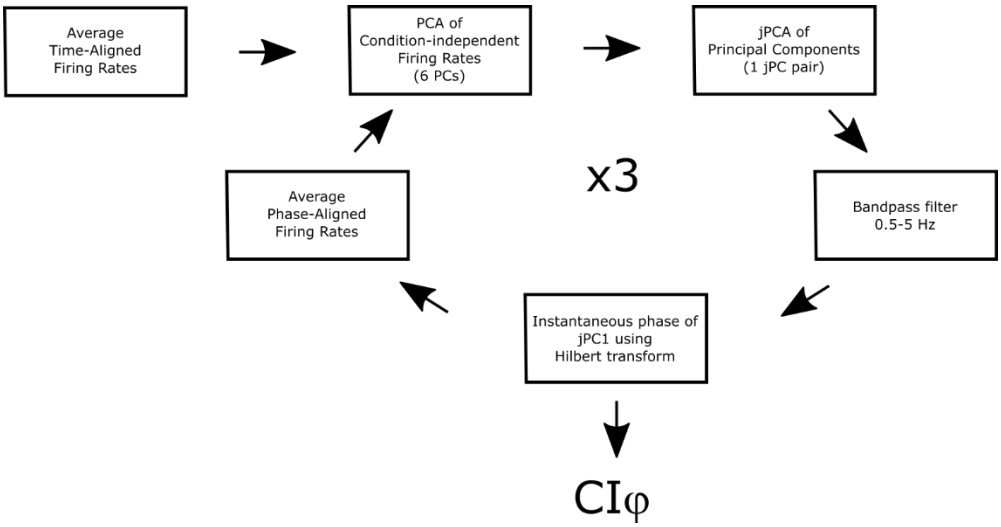

Schematic of CIϕ estimation. The cyclic condition-independent neural activity was estimated using an iterative technique. After initial time-alignment and averaging of the firing rates, PCA and jPCA was performed to estimate the dimensions neural activity with the greatest cyclic condition-activity were identified. Using the phase of this first jPC dimension, the firing rates were realigned for averaging. The PCA and jPCA were again performed a total of three times. The final CIϕ was then used for the presented analysis.

Matlab code is available on GitHub:

https://github.com/arouseKUMC/CIphase

Extended data figures

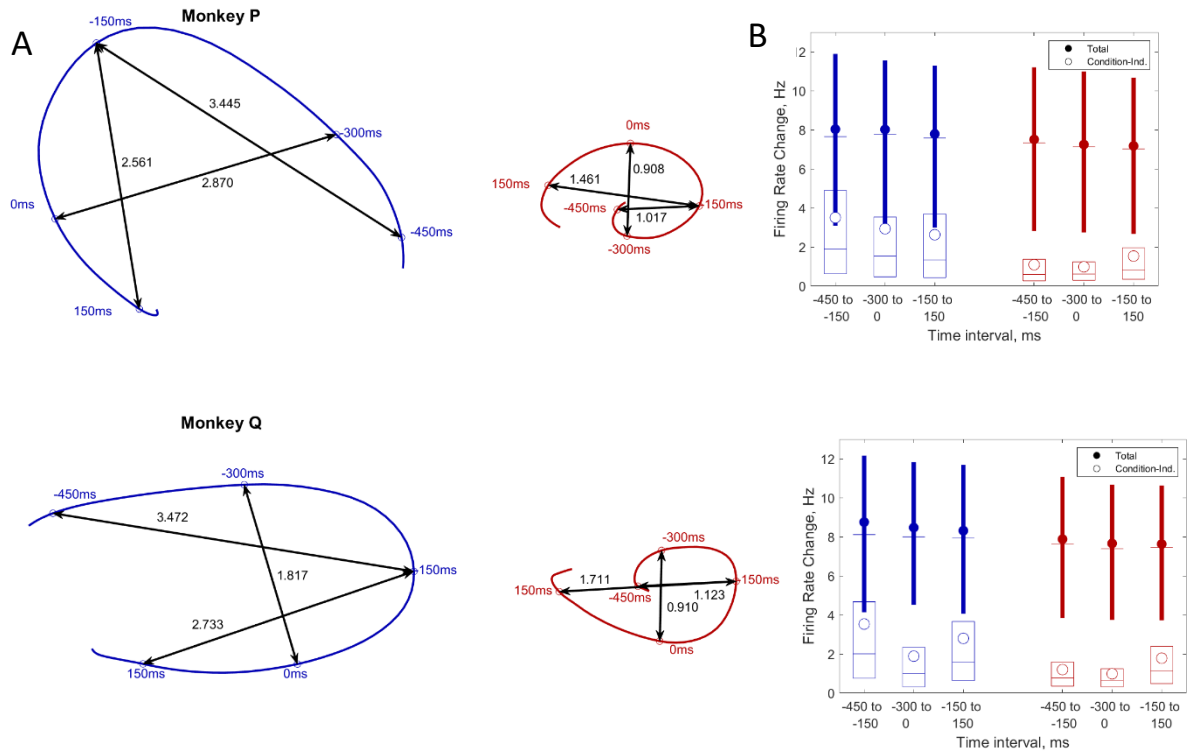

Figure 5-1. Same as Figure 5 but with no square root transform of firing rate. Neural distance between condition-independent firing rates. A) The differences between average firing rates (in arbitrary units) for initial (blue) and corrective (red) for times before (-) and after (+) peak speed (0ms). Distances are shown for three timepoint pairs: -450ms & -150ms, -300 & 0ms, and -150ms & 150ms. The trajectories and distances are shown as a two-dimensional approximation of the high-dimensional neural space using multidimensional scaling. The dimensions for each panel were derived independently with no relationship across panels. B) Changes in firing rate for all analyzed neurons. The change in firing rate when averaged across all trials (Condition-Independent--open circles) and the total trial-by-trial change in firing rate (Total including Condition-Independent, Condition-Dependent, and Noise--closed circles). The circles represent the mean across all neurons while the 25th, 50th, and 75th percentile neuron are shown with the horizontal lines and boxes/whiskers. Changes in firing rate were measured between three time increments relative to peak speed: -450 to -150ms, -300 to 0ms, and -150 to 150 ms. For the corrective movements (red), there was a significant drop in Condition-Independent firing rates compared to initial submovements (blue). However, the overall change in firing rates per trial remained similar.

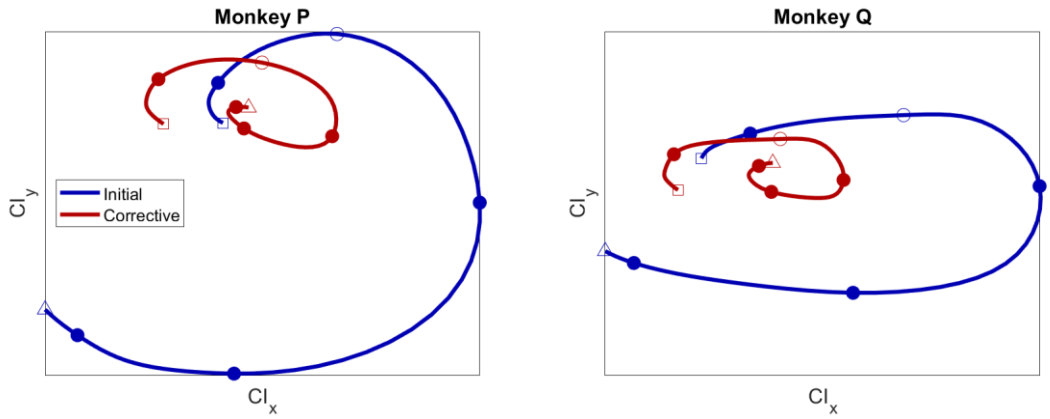

Figure 6-1. Same as Figure 6B. but without the square-root transform. The average firing rate for initial and corrective submovements are plotted in the CIx/CIy plane identified with jPCA. The trajectories start at the triangles and end at the squares. Each filled circle is a 150 ms time step.

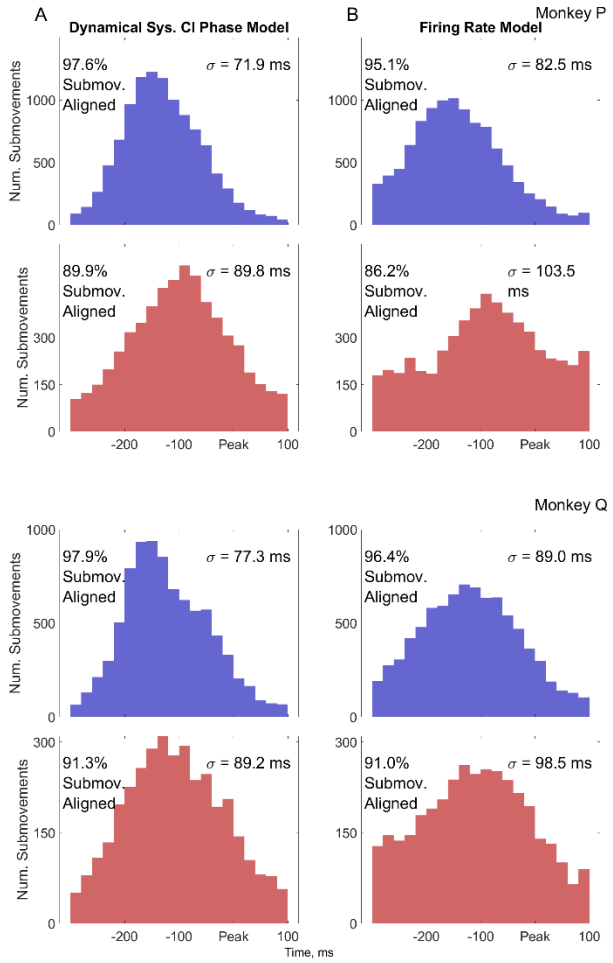

Figure 8-1. Same as Figure 8 A and B but CIϕ was calculated without the square-root transform. The standard deviation (σ) of the peak speed predictions for the CI phase model ranged from 2-6ms worse than the square-root transformed data but all were still significantly better than the firing rate model (F-test, all $p<0.001$ for both initial and corrective submovements for both monkeys).