Title:

M1 dynamics share similar inputs for initiating and correcting movement

Authors:

Peter J. Malonis¹, Nicholas G. Hatsopoulos^{2,†}, Jason N. MacLean^{3,*,†}, Matthew T. Kaufman^{2,4,*,†}

Affiliations:

¹ Committee on Computational Neuroscience

² Dept. of Organismal Biology and Anatomy

³ Dept. of Neurobiology

⁴ Neuroscience Institute, The University of Chicago, Chicago, IL 60637 USA

* Equal contributions

† Corresponding authors

Author contributions: MK and PM conceived the project. PM performed the analyses. JM and MK supervised the project with assistance from NGH. NGH provided the data. PM, JM and MK wrote the paper, and all authors edited the manuscript.

Competing interests: The authors declare no conflicts of interest.

Acknowledgements

The authors thank Josh Coles, Zach Haga, Jignesh Joshi, Dawn Paulsen, Jacob Reimer, Aaron Suminski, and Dennis Tkach for originally collecting the data, and Chethan Pandarinath for assistance with LFADS. This work was funded by R01 NS111982 (NGH), R01NS045853 (NGH) R01 EY022338 (JM) NSF CAREER Grant 0952686 (JM), The University of Chicago (MK), and the Neuroscience Institute at the University of Chicago (MK).

Abstract

Motor cortex is integral to generating voluntary movement commands. However, as a dynamical system, it is unclear how motor cortical movement commands are informed by either new or sensory-driven corrective instructions. Here, we examine population activity in the primary motor cortex of macaques during a continuous, sequential arm movement task in which the movement instruction is updated several times over the course of a trial. We use Latent Factor Analysis via Dynamical Systems (LFADS) to decompose population activity into a portion explainable via dynamics, and a stream of inferred inputs required to instruct that dynamical system. The time series of inferred inputs had several surprising properties. First, input timing was more strongly locked to target appearance than to movement onset, suggesting that variable reaction times may be a function of how inputs interact with ongoing dynamics rather than variability in instruction timing. Second, inferred inputs were tuned nearly identically for both initial and corrective movements, suggesting a commonality in the structure of inputs across visually-instructed and corrective movements that was previously obscured by the complexity of the dynamical system that is M1.

Introduction

Animal behavior requires the motor system to produce reliable movements with precise temporal structure. While some aspects of motor pattern generation occur in the brainstem and spinal cord (Marder & Calabrese, 1996), motor cortex exhibits activity patterns with complex temporal structure (Churchland, & Shenoy, 2007; Sergio et al., 2005) that relates to many aspects of movement as well (e.g. Evarts, 1968; Georgopoulos et al., 1982, 1986; Hatsopoulos et al., 2007; Kalaska, 2009). However, it remains unclear to what extent activity reflects intrinsically arising pattern generation (Brown, 1914; Yuste et al., 2005) versus external sensory and top-down inputs (Churchland et al., 2010; Omrani et al., 2017; Scott, 2008).

One approach to understanding the pattern generation aspects of this activity has been through the lens of dynamical systems analysis (Churchland et al., 2012; Seely et al., 2016; Perich et al., 2020; Rouse, 2018; Sauerbrei et al., 2020; Vyas et al., 2020) particularly in reaching primates. In the context of brief and largely open-loop reaching, dynamical systems have been used to show that motor cortex can be approximated well as having rotational dynamics (Churchland et al., 2012) or a variant thereof (Hennequin et al., 2014; Russo et al., 2018; Sabatini & Kaufman, 2021), and is initialized by movement-specific inputs (Churchland et al., 2012, 2010) and then triggered by a large, external, condition-invariant signal (Kaufman et al., 2016; Michaels et al., 2015). This approximation of motor cortex as having autonomous dynamics has worked well in cases where movements are brief. Yet, just as inputs external to motor cortex must initialize the system, successful motor behavior requires continual updates of these instructions. The dependence on ‘external’ inputs is especially clear at the time of movement initiation (Dacre et al., 2021; Sauerbrei et al., 2020) and during errors, when input is necessary to correct movement (Perich et al., 2020). Thus, we expect the motor system to differ from an autonomous dynamical system whenever the brain sets new goals or corrects mistakes dependent on feedback (Elliott et al., 2010). This difference should be particularly salient during continuous motor behavior, when task goals are repeatedly updated.

A recently developed method, Latent Factor Analysis via Dynamical Systems (LFADS) (Pandarathna et al., 2018), addresses this distinction by incorporating both autonomous dynamics and a separate stream of inputs in fitting neural activity with a deep neural network architecture. That is, this method attempts to learn the dynamics of the neural activity to the extent that they can be parsimoniously modeled as a dynamical system, and concurrently learns what inputs must be provided at what time points to this “generator” network to explain the full pattern of activity observed.

Here, we apply LFADS to neural activity recorded during a Random Target Pursuit task (Hatsopoulos et al., 2004), in which a monkey is presented with a new target whenever it contacts the current one, and in which subjects make frequent errors and execute corrections. We then analyzed the inferred inputs to the dynamical system and observed distinct, large input transients when new targets appeared. These transients were locked to the target presentation more strongly than to the movement itself. Further, we found that inferred inputs predicted the updated goal location in hand-centered coordinates better than shoulder-centered coordinates, consistent with prior work on target specification (Bremner & Andersen, 2012; Buneo et al., 2002; Pesaran et al., 2006). Finally, we found that input transients had highly similar tuning and

temporal structure for both initial, visually instructed movements and for error correction sub-movements. Thus, motor cortical activity corresponding to continuous movement is well-modeled as a dynamical system receiving brief external inputs which specify new goals and corrective movements similarly.

Results

We analyzed Utah array recordings from M1 while three monkeys performed a Random Target Pursuit (RTP) task (Fig. 1A). Movement was constrained to the horizontal plane by a two-link exoskeleton. In each trial of this task the monkey made continuous arm movements from target to target, each of which appeared at a random location on the screen when the previous one was hit with a cursor projected just above the hand (Fig. 1). The RTP task therefore requires continual updating of motor commands in response to stimuli throughout the trial, evoking continuous movement.

A dynamics-based fit to M1 activity required time-varying inputs

We used LFADS, a machine learning method based on artificial recurrent neural networks, to model the neural spiking data. This model treats neural data as arising from a dynamical system, initialized at the start of the trial with a seed that is specific to the upcoming activity, and subject to a time-varying set of inputs that arise from outside the dynamical system (Fig. 2A). This model is structured as an autoencoder: neural data are first transformed into the initial state for each trial, then run through a dynamical system to generate the state of the system over time (conceptually like the factors that would result from Principal Component Analysis), and finally these factors are related back to the same neural activity via a Generalized Linear Model. The non-dynamical inputs form a side loop (Fig. 2A, red path), where both the initial state and errors of the factors (failures to reconstruct the data) inform a stream of corrective external inputs to the dynamical generator. The output of LFADS is therefore a denoised set of firing rates over time on single trials for each neuron. This model has been shown to reliably distinguish between different conditions on single trials from the neural data alone; accurately predict external kinematic variables; and importantly for our purposes, correctly infer external inputs when a movement was perturbed on a given trial (Pandarinath et al., 2018). In the framework of the LFADS model, the presence of inputs at a given time in the trial means that the evolution of population activity cannot be parsimoniously explained by the autonomous dynamics of a recurrent neural network. Here, we were particularly interested in the nature of these inferred inputs, which reveal what inputs a dynamical system would require at each time point to recapitulate the neural activity recorded during continuous pursuit of random targets.

Because the inference of inputs was the focus here, we took particular care to fit the hyperparameters of the model affecting these external inputs. In general, the same data can be fit as arising from either a more complex dynamical system (higher dimensional or with ‘more nonlinear’ dynamics) receiving less exogenous input, or a simpler dynamical system receiving

stronger and more frequent inputs (Chicone, 2006). For LFADS, whether the system fits a more complex dynamical system or more complex inputs is informed by hyperparameters that are set by the user. If the “input penalty” hyperparameter is chosen to be too low, too much of the activity will be attributed to the inputs, not enough of the autonomous dynamics will be learned, and the model will not be able to leverage the dynamics to denoise the data optimally. If the model penalizes external input too heavily, then the model will not capture updates to the motor plan due to the changing state of the task.

To first evaluate the importance of external inputs to model performance, we measured how well the ongoing movement kinematics could be linearly decoded from the denoised single-trial firing rates (Keshtkaran et al., 2021; Pandarinath et al., 2018) inferred by models with and without external inputs (Fig. 2B). For the model with external inputs, we observed a mean cross-validated decoding r^2 of 0.75 ± 0.08 (mean $\pm$ s.d. across monkeys), compared to an r^2 of 0.39 ± 0.08 without external inputs. This result indicates that external inputs improved the model fit of firing rates for the purpose of accurately decoding kinematics. The denoised rates produced by LFADS also outperformed direct decoding from the recorded firing rates ($r^2 = 0.46$), as found previously with LFADS applied to motor cortex data from other tasks (Pandarinath et al., 2018). Consequently, for the remainder of the study we limited our study to LFADS containing external inputs.

Having established the importance of external inputs in the LFADS model to infer single-trial firing rates, we next evaluated how much external input in the model was optimal. To do so we iteratively considered multiple values of the input penalty hyperparameter. For each monkey, we found that the decoding performance peaked at an intermediate hyperparameter (Fig. 2C). For subsequent analyses, we set the controller penalty to a single value for all monkeys (2.0) which was at or near the peak decoding performance for each animal and yielded smooth input time series. We also considered the dimensionality of the inferred inputs, testing values of 0, 1, 2, 3, and 4 (Fig. 2E). Performance of decoding kinematics from the denoised factors was equivalent using any model with 1 or more inferred input dimensions. However, in decoding the position of the targets, inferred inputs of dimensionality 2 for monkeys RS and MK were optimal, while dimensionality 3 gave the best performance for monkey RJ (Fig. 2F). Therefore, we analyzed the models which used these optimal values.

Inferred inputs were sparse and locked to target presentation

In principle, external input could be continuous or pulsatile, and changes in the inputs could occur sparsely or frequently. LFADS penalizes the input magnitude at all time points, and thus will tend to remove sustained features in the inputs – which are easy for the autonomous generator to incorporate – much like a high-pass filter. We therefore focused our analysis on their sparsity, which allowed us to evaluate how frequently inputs changed.

We computed the Gini coefficient, a measure of sparsity (Hurley & Rickard, 2009), of each input as we varied the input penalty hyperparameter for model fitting. Raising the input penalty might have had two different effects on the resulting model. If the penalty had a

subtractive effect on inferred inputs, then the Gini coefficient would increase monotonically with input penalty as fewer peaks remained. Alternatively, if the penalty had a divisive effect on inferred inputs and simply rescaled the input trace, then the Gini coefficient would be unchanged. We found that the Gini coefficient first increased then decreased as the input penalty increased (Fig. 2D). The Gini coefficient peaked at values at or close to the ones that produced the best kinematic decoding performance (Fig. 2C,D). In other words, the best models of the dynamical structure of M1 activity were those for which external input was most concentrated at particular times, as opposed to being distributed more uniformly throughout the trial.

With this optimized model, inferred inputs maintained a relatively low and stable baseline with sparse, large, brief excursions that tended to occur immediately following the presentation of a new target (Fig. 3). Averaging the magnitude of inferred inputs across target presentations revealed a peak at ~50 ms after target appearance for all three monkeys (Fig. 4A). The timing of the inferred input peak coincides with previous reports of the latency from visual movement cues to the start of changes in M1 activity (Dickey et al., 2013; Lamarre et al., 1983). To examine how reliable the peak timing was on individual trials, we computed the latency from the target to the first inferred input above a threshold (Fig. 4B). As the threshold was increased, the peaks became more sharply peaked around 50 ms following target presentation, indicating that the largest transients were the most consistently timed.

To better evaluate the specific timing of the input transients, we examined whether the timing of the transients was more consistently related to the timing of target presentation or to the timing of movement initiation towards the target. The start of the initial movement was identified by examining the speed of the hand following target presentation and selecting the appropriate speed local minimum (see Methods). We then examined the latencies to each input peak as defined above to the target presentation and the initial movement (Fig. 4C). The distribution of latencies relative to the target had both a smaller variance and smaller width at half-height compared to the latencies relative to the initial movement ($p < 0.001$ for both, permutation test, see Methods). This result suggests that the post-target input transients are more consistently timed to stimulus appearance rather than movement onset. This in turn suggests that controller inputs are related to visual inputs, and therefore that the variable latency of responding to a new target may result from the interaction of these inputs with the dynamical system of M1.

We then verified that inferred inputs in the post-target presentation period were larger than the inputs during other parts of the trial. We calculated the total input magnitude in 100 ms sliding windows throughout the entirety of each trial. Using the magnitude of the input in each window we determined whether windows near target appearances could be distinguished from all other windows by iteratively calculating the area under the receiver operator characteristic (ROC) for each window. For each monkey, the window ranging from 0 ms to 100 ms following the target presentation was most distinguishable from all other times, with peak areas under the ROC curve of 0.85, 0.76, and 0.56 for monkeys RS, MK, and RJ, respectively (Fig. 4D). This establishes that large input transients are overrepresented in the period 150 ms after the target appearance.

Inferred input transients relate to target direction in hand-centered encoding

Given the strong correspondence between the timing of target appearance and of the inferred inputs, we evaluated the extent to which the values of the inferred inputs predicted the location of the target. To do so we made a sort of tuning curve by binning targets according to the direction from hand to target, then averaging the inputs for each bin to get a trace over time for each direction (Fig. 5A). Despite the continuous nature of the movement this approach revealed highly distinct inputs for different upcoming movements.

We then decoded target location to verify that inputs were meaningful, particularly around the time of target presentation. To do so, we compared the cross-validated decoding performance using either inferred inputs or M1 firing rates in different windows of time relative to the target presentation, with window sizes equal to 250 ms (Fig. 5B). M1 firing rates yielded a stable prediction of the target position throughout the full duration of the trial, consistent with M1's known correlations with movement across time (Hatsopoulos et al., 2007; Wu et al., 2006). In contrast, using inferred inputs, decoding performance rapidly decreased when the start of the window of input used was >50 ms after the target presentation. This contrast suggests that inputs in the model only briefly encode the movement, but they persistently alter the state of the dynamical system, such that the encoding is sustained in the evolution of M1 activity for the remainder of the movement.

Areas upstream of M1 in the fronto-parietal reach network have been shown to more strongly represent the difference vector from the hand to the target than the absolute position of the target relative to the torso (Bremner & Andersen, 2012; Pesaran et al., 2006). We therefore hypothesized that similar encoding might be present in the inferred inputs. We evaluated the performance of a nonlinear decoder predicting target position depending on whether the hand or shoulder was used as the origin for each target. Consistent with previous reports, the encoding of the targets by the inferred inputs was stronger relative to the hand than relative to the shoulder (Fig. 5C *left*; $r^2 = 0.70$ vs. 0.53 for Monkey RS, 0.63 vs. 0.39 for Monkey MK, and 0.46 vs. 0.21 for Monkey RJ; $p < 0.001$ for each monkey, corrected t-test, see Methods).

We also trained a decoder that fit the direction and magnitude of the vector pointing from the point of reference (hand or shoulder) to the target. The performance and difference between the hand-centric and shoulder-centric models was similar for the direction decoder (Fig. 5C, *center*). The magnitude decoder exhibited weaker performance in general (Fig. 5C, *right*), with no consistent differences by hand vs. shoulder origin. This indicates that the difference in decoding performance between the hand- and shoulder- centric decoders is largely due to the more accurate representation of target direction by the inferred inputs.

Corrective movements were also preceded by transient inputs

Many large inputs were time-locked to target presentation, but large inferred input transients were also observed outside of these time windows (Fig. 3). We hypothesized that

these transients might correspond to corrective submovements, which are common in the RTP task (Fig. 6A) and which have been suggested to coarsely share M1 activity patterns with initial movements (Rouse, 2018).

In example trials, we found that bends in the hand trajectory were associated with transients in the inferred inputs (Fig. 6B). To quantify this, we identified submovements by appropriate local minima in the speed of the hand (see Methods). We defined corrective submovements as all submovements that occurred after the initial submovement that followed the presentation of a new target. We then asked whether the time windows around corrective submovement onsets were associated with transients in the inferred inputs. For all three monkeys, we found a peak in the inferred input magnitude at the start of corrective submovements, 180 ms to 300 ms before the speed minimum (RS: 180 ms; RJ: 300 ms; MK: 210 ms) (Fig. 6C). As we found for initial movements, the transients in the 200 ms period prior to corrective submovements were statistically different from random time windows assessed via the area under the ROC curves (Fig. 6D, *bottom*). This corresponds to the period relative to the initial movement where the inputs are particularly active (Fig. 6D, *top*).

Inferred inputs similarly encode both initial movement and corrective submovements

To understand the information content of the inferred inputs for corrective submovements, we trained a decoder (using either support vector regression or random forest regression, Methods) to predict the position of the targets from the inferred inputs around corrective movements. This decoder explained the 23-28% of the variance in the hand-relative target position (Fig. 6E), and 31-41% of the target direction in held-out data (Fig. 6F). As with the post-target-presentation inputs, the prediction was much better for target direction relative to hand position than for target position relative to the shoulder ($r^2 = 0.31$ vs. 0.05 for Monkey RS, 0.41 vs. 0.11 for Monkey MK, and 0.36 vs. 0.13 for Monkey RJ, $p < 0.001$ for each monkey, corrected t-test), and the magnitude decoder had weaker performance and did not exhibit differences between the hand- and shoulder- centric direction decoders (Fig. 6G).

Given that the inferred inputs to M1 exhibited transients both following initial target presentation and around the time of corrective submovements, and that the direction of the upcoming movement could be decoded in both cases, we measured the similarity of these two representations. As in Figure 5A, we binned the target directions and computed the average input for each direction (Fig. 7A). Averaging the inferred input for each direction in a window from 350 to 50 ms before to the start of movement, we produced tuning curves for inferred inputs at target presentation and at corrective submovement onset. The two representations were highly similar (Fig. 7B), with an overall correlation of 0.89 (95% CI [0.83, 0.93]), and for each monkey individually (monkey RS: 0.93, 95% CI [0.84, 0.97]; MK: 0.97, [0.93, 0.98]; RJ: 0.82 [0.67, 0.9]).

This result was intriguing because in general the target representations in M1 are complex, evolving over the course of a trial (Churchland, & Shenoy, 2007; Fu et al., 1995). Accordingly, we found that the similarity in input encoding was not simply due to similarity in

firing rates. We repeated the comparison of tuning curves using the first 5 principal components of the M1 firing rates instead of inferred inputs (Fig. 7C) and found that the correlation in tuning between the initial movement period and the corrective submovements was 0.46, 95% CI [0.34, 0.57], (Monkey RS: 0.53, [0.32, 0.69]; MK: 0.26, [0.01, 0.48]; RJ: 0.57 [0.37, 0.72]). This argues that the tuning of inputs is more consistent between initial and corrective movements than the tuning of the firing rates themselves.

We also tested whether the consistency in inferred input between initial movements and corrections could be explained by consistent tuning throughout the entire trial. We tested the tuning correlation between inferred inputs in a 300 ms window centered on the peak speed of the initial movement and in a window from 350 to 50 ms before the start of the initial movement (Fig. 7D). This correlation was 0.58, compared to 0.89 for the tuning correlation between initial movement and corrective movement ($p < 0.001$, Z-test with Fisher's Z-transformation; Monkey RS, 0.56 vs 0.93, $p < 0.001$; MK: 0.81 vs 0.97, $p = 0.002$; RJ: 0.44 vs 0.82, $p = 0.005$). Therefore, the window prior to corrections had a stronger tuning correlation with the initial movement than the window centered on the speed maximum, despite the fact that, for each target, the initial movement was closer in time to the speed maximum than to the corrective movements. This again argues that the tuning consistency observed in the inferred inputs is surprising and is not simply a consequence of temporal proximity.

Discussion

We used LFADS to identify dynamical structure and infer exogenous inputs in M1 activity recorded during a continuous behavioral task with regular updates of movement goals. Most of the inferred input occurred immediately following target presentation and immediately preceding corrective movements. The post-target transients were more strongly time-locked to target appearance as compared to the start of movement. The details of these input transients coded for the direction of the target relative to the hand, and were highly similar for the initial movement and the corrective movements. Together these results demonstrate that motor cortex is well modeled as a dynamical system that is intermittently perturbed by external input following the appearance of a new goal or the realization that a correction is necessary to reach the desired goal.

Though the inputs are inferred and not observed directly, several aspects of their structure support the idea that the inferred inputs may strongly resemble real inputs to M1. First, the inferred inputs occur when known updates to the task goals occur: target presentation and corrections. While target-locked responses predominate in PMd (Hatsopoulos et al., 2004), M1 neurons are roughly evenly split between target-locked and movement-locked responses (Rao & Donoghue, 2014). Second, the inferred inputs resembled a hand-centered “reference frame” more than a shoulder-centered frame. Finally, the similarity of inputs for initial target presentation and corrections is surprising: firing rates are tuned somewhat differently for delayed and non-delayed movements (Ames et al., 2014; Crammond & Kalaska, 2000), and generally exhibit inconsistent tuning at different points in the trial (Churchland, & Shenoy, 2007; Hatsopoulos et al., 2007).

Previous work on corrective movements has proposed at least two distinct processes that produce corrections during arm movements (Elliott et al., 2010, 2017): online modifications to the initial impulse, and subsequent discrete position adjustments that bring the end effector to the target. Our findings argue that these processes may be more unified than previous evidence has suggested. First, the similarity of input tuning for initial and corrective movements suggests that the rest of the brain may not have to make a strong distinction between instructions and corrections. Second, the lack of inputs around the time of max speed does not support an initial impulse correction. These results may therefore argue for a greater compartmentalization of motor function than suspected, with M1 acting as a movement generator that can take movement context into account instead of requiring the rest of the brain to reflect current movement details in commands to M1 (Andrew Pruszynski et al., 2014).

Our findings also have implications for the understanding of reaction times. It has previously been argued that reaction times reflect variable timing of the movement commands themselves (Cisek, 2007; Roitman & Shadlen, 2002), possibly due to sensory variability (Osborne et al., 2005), and commensurate with M1 firing rates being strongly locked to movement onset (Rao & Donoghue, 2014). However, we found that inferred input was most tightly locked to the target presentation itself, not the movement onset. This suggests that the variable reaction time of the monkey is likely due to an interaction of the input with the current state of the dynamical system, and that depending on this state the system may be faster or slower to initiate. However, it also is possible that this is only true in a continuous-movement task like RTP.

Interestingly, inferred inputs were best described as encoding target position in hand-centric coordinates rather than body-centric, consistent with representations in cortical areas upstream of M1 in the fronto-parietal reach network such as posterior parietal cortex (Bremner & Andersen, 2012; Buneo et al., 2002, 2008), and dorsal premotor cortex (Pesaran et al., 2006). Reference frames found across the parieto-frontal reach network are highly heterogeneous (Wu & Hatsopoulos, 2006) and this type of description leaves out important aspects of motor responses (Shenoy et al., 2013; Scott, 2008; Omrani et al., 2017; Fetz, 1992). Nonetheless, to the extent that reference frames capture encoding, our work indicates that inputs to M1 are closer to hand-centric and suggests that this coordinate system informs the production of motor commands in M1. We did not measure eye position in these experiments, so eye-centric reference frames could not be considered here.

Using LFADS to infer inputs has some limitations. In trying to separate activity into a dynamical portion and non-dynamical inputs, LFADS fits a dynamical model capturing regular statistical structure in population recordings. Inputs that occur predictably on each trial cannot be distinguished from autonomous dynamics of the system (Perich et al., 2020), and so LFADS is incapable of inferring some possible kinds of external inputs to the recorded population. For instance, LFADS would be incapable of inferring sustained inputs that occurred throughout a trial. Similarly, feedback inputs that are predictably related to the current state of the system will generally be subsumed into the inferred dynamics (Kalidindi et al., 2020). Therefore, LFADS likely misses some kinds of inputs to M1. However, dynamical models such as LFADS are capable of inferring unpredictable inputs reflecting new external information and updated internal goals. For instance, it has been shown that LFADS-inferred inputs encode an unexpected jump in a target to one of two locations (Pandarinath et al., 2018). Here, using a

more continuous task without physical perturbations, we found that exogenous input is stronger in the period between the target presentation and the first movement, as opposed to during the movement itself.

Relatedly, the hyperparameter space for LFADS is high-dimensional, making it impractical to perform a global optimization of fits. In fitting LFADS, we selected a set of hyperparameters that led to high-performance decoding of movement kinematics, far better than using the firing rates alone (Fig. 2). However, the only hyperparameters for which we performed a full search were those that related to the level of inferred input. It is possible that exploring a larger set of LFADS hyperparameters could help model M1 dynamics during continuous movement at an even higher level of precision (Keshtkaran et al., 2021).

Future experiments could enable further understanding of the role of inputs to M1 in the generation of movement. For example, cueing with a non-visual input, such as a directional sound cue, could help determine whether the shared features of target-evoked and corrective movements is a commonality due to vision. Or, using a passive movement condition could serve to identify the effect of isolated proprioceptive inputs in the absence of efference copy.

The current generation of dynamical systems models enable powerful inference about cortical population activity. Moreover, this inference permits an examination of the nature of inputs, including about their effects on movement variability and initiation under different conditions. While at present the source of these inputs is not clear, future pairing of dynamical systems techniques with greater biological access will enable even deeper understanding of the distribution of computation in the motor system.

Methods

Behavioral task

Three adult male rhesus monkeys (*Macaca mulatta*) were operantly trained to control a cursor in a two-dimensional workspace using a two-link robotic exoskeleton (Scott, 1999). The animals sat in a primate chair with their dominant arm in the exoskeleton. Their shoulder joint was abducted 90° and supported by the manipulandum such that all movements were made within the horizontal plane. Direct vision of the limb was precluded by a horizontal projection screen above the monkey's arm. Visual feedback was available via a visual cursor projected onto the screen. Cartesian coordinates of the visual cursor were determined by digitizing the shoulder and elbow angle along with angular velocity at 500 Hz and transforming these variables into a visual cursor position (in centimeters) using the forward kinematic equations for the exoskeleton.

In this configuration, the monkeys performed a random target pursuit (RTP) task. This task required the monkeys to move a cursor (6.7 mm in diameter for monkey RS; MK: 6.7; RJ: 10) to a series of 7 square targets (10 mm on a side for monkey RS; MK: 13 mm; RJ: 20 mm). When the cursor reached the target, the target immediately disappeared and a new one appeared at a random location in the workspace. In Monkey RJ, it was necessary to exclude some trials where the monkey was disengaged in the task for part of the trial. Trials were excluded if the monkey spent more than 500 ms total with a cursor speed less than 1 mm/s.

The raw cursor position was filtered both forwards and backwards with a 3rd order low-pass Butterworth filter with a cutoff frequency of 20 Hz. After computing the velocity via numerical differentiation, the same filter was applied again. Because the algorithm requires trials of equal length, trials were truncated to a fixed length, and trials that were shorter than this length were excluded. The trial length was selected for each data set so as to maximize the total number of time bins across all of the included trials (the time bins per trial times number of trials retained).

Neural recordings

The neural recordings were collected using a microelectrode Utah array composed of 100 silicon electrodes (1.0 mm electrode length; 400 µm inter-electrode separation). The arrays were implanted in the arm area of the primary motor cortex (M1) of each monkey. During a recording session, signals from up to 96 electrodes were amplified (gain, 5000), bandpass filtered between 0.3 Hz and 7.5 kHz, and recorded digitally (14-bit) at 30 kHz per channel using a Cerebus acquisition system (BlackRock Microsystems, Salt Lake City, UT). Only waveforms (1.6 ms duration) that crossed a threshold were stored and spike-sorted into single units using Offline Sorter (Plexon, Inc., Dallas, TX). For each monkey, a single recording session was included, with 100 neurons, 49 neurons, and 51 neurons from monkeys RS, MK, and RJ, respectively.

Identification of submovements

Submovements were identified using the speed of the cursor over time. Candidates for the start of submovements were identified as the local minima in the smoothed speed profile. The smoothing was performed using a Savitsky-Golay filter with a window of 101 ms and polynomial order of 2. These candidates had to satisfy two additional criteria to be accepted as the start of a submovement. First, we required that the peak speed of the ensuing movement was at least 25 mm/s greater than the speed at the local minimum. This excluded noise and small movements during periods of hesitancy. Second, we required that the minimum time interval between candidates was 200 ms. This limited the extent to which time windows around submovements overlapped. For candidates within 200 ms, we applied a pairwise selection rule to eliminate one candidate until all candidates were a minimum of 200 ms apart. This selection rule was as follows. For each candidate we computed the speed increase: the next peak speed minus the speed at the minimum. In most cases we selected the candidate with the larger speed increase. However, there were rare cases in which a pair of candidates represented, respectively, the beginning of a submovement and a brief, shallow deceleration interrupting the acceleration phase of the submovement. In these cases, we aimed to define the first candidate as the beginning of the submovement and overrode the “larger increase” rule. These cases were identified if two conditions were both met. First, the dip was shallow: the difference between the preceding speed maximum and the second candidate was <10 mm/s. Second, the speed increases for the two candidates were similar, defined as neither being more than 5 times greater than the other.

Once the submovements were defined, they were identified as initial movements or corrective submovements. The initial movements were defined as the first movement following the target presentation. Corrective submovements were defined as submovements occurring after the initial submovement but before the presentation of the next target.

LFADS fitting and model evaluation

The spiking data recorded during each trial were modeled using Latent Factor Analysis via Dynamical Systems (LFADS), described in detail in Pandarinath et al. (2018). This model uses a type of variational autoencoder (Kingma & Welling, 2014; Doersch, 2016) to learn a dynamical representation of each trial. The central structure of this network consists of two artificial recurrent neural networks (RNNs), an “encoder” and a “generator.” The generator is trained to produce autonomous dynamics that can be mapped onto the neural activity of a trial. The encoder learns to map the activity of a trial onto a distribution of initial states based on how likely they were to produce the activity on that trial. In addition, a “controller” RNN can be added that learns to compensate for the dynamics that cannot be learned by the autonomous generator. The output of the controller is a distribution of inputs to the generator that are passed on each time step. To obtain the inferred inputs to the neural population for each trial, LFADS samples from the distribution of controller outputs and averages the samples.

Most of the LFADS hyperparameters were adapted from those used for similar recordings in Pandarinath et al. (2018), with the number of factors outputted from the generator set to 50. The full set of hyperparameters is listed in Table 1. Because we were most interested

in the inputs from the controller to the generator in the LFADS model, we most thoroughly explored the hyperparameters directly responsible for determining the structure of the controller outputs (the inferred inputs). These were 1) the penalty on the inferred inputs and 2) the dimensionality of the inferred inputs, ($|u_t|$ in Pandarinath et al. (2018)), including values of zero.

Models were evaluated by training a linear decoder of movement kinematics from the LFADS factors and evaluating the decoder's performance (see below). We evaluated a range of inferred input penalties, ranging from values low enough to produce overfitting to values large enough to yield underfitting.

After performing this search we selected an inferred input penalty (= 2.0) that was at or near optimal decoder performance for all monkeys. We fixed the inferred input penalty to this value for the search of dimensionality of the inferred input. In the search for the dimensionality of the inputs, the decoder performance plateaued at 1 input and larger. We therefore selected the input dimensionality for each monkey using the decoding performance for the target (see Fig. 2).

Decoding of movement kinematics

To evaluate the LFADS models, we decoded movement kinematics from the LFADS factors. An ordinary least-squares linear model was used to map the factors onto the position and velocity of the endpoint (hand), with a 100 ms lag between the LFADS factors and the kinematics. To determine the decoder performance, we performed 5-fold cross-validation, splitting the time bins into random test sets. Performance of the model was evaluated using the held-out variance explained. Performance across the four kinematic variables was averaged to obtain a single measure of performance for each training/test set. For the controls in which kinematics were decoded directly from neural activity, we obtained continuous firing rates by smoothing the raw spike trains with a Gaussian filter with s.d. equal to 25 ms.

Analysis of input timing

Analyses were performed on the time-varying inferred inputs for each trial. We used all trials in the analysis of the inferred inputs, including both those that were used for training and validation in training the LFADS network. For analysis of inferred input timing, and decoding from inferred inputs below, we excluded the response to the first target on each trial in order to eliminate artifacts resulting from the variable behavior of the monkey before the start of the trial. To determine whether the timing of the input transients was more closely related to the target presentation or to the initial movement (Fig. 4C), we examined the timing of the maximum inferred input magnitude in the period between the presentation of the target and the initial movement. We included any target presentation in this analysis, regardless of how large the maximum input magnitude was. Since the timing of the transients of the individual inputs is more exact when the peaks reach a larger threshold (Fig. 4B), and for some directions there is on average only a small transient (Fig. 5A) which may not be detected, taking the timing of the maximum is conservative for detecting systematic timing precision. But by not setting a

threshold for input transients to use, we included all relevant target presentations and thus captured the structure of timing in the model in the most general way.

To compute statistics for Figure 4C, we used a permutation test to evaluate which distribution of latencies had the larger spread. Spread was defined as the full width at half-height (FWHH) of the peak of the distribution. First we subtracted the modal value from each set of latencies to center the peak of both latency distributions on 0. Then we combined the two sets of latencies into a single population that we split in half randomly 1000 times. For each of these random partitions we computed the FWHH, and took the absolute difference between the FWHH of the two random subgroups. These absolute differences constituted the null distribution for the permutation test. We compared the difference in FWHH between the actual samples of the distribution and the null distribution of spreads to obtain a p-value for the hypothesis that the distribution of latencies to the different events had the same spread. Because by definition the maxima we examined were before the initial movement, there was a sharp cutoff on the right side of the distribution of latencies to the initial movement (Fig. 4C). However, this simply means that the FWHH measure is possibly conservative in detecting the greater spread of the initial movement distribution.

We also considered the possibility that the inputs were better aligned to the acceleration phase of the movement. We performed the same analysis, but defining the initial movement time by the time at which the movement reaches half of its peak change in speed. Defining the initial movement in this way did not affect the results or significance levels.

To determine whether windows near target appearances could be distinguished from all other windows, we calculated the area under the receiver operator characteristic (ROC) for each window. To do so, we computed the distribution of input magnitude for a 100 ms long window of interest at a fixed latency to the target. We then compute a “background” distribution of input magnitude in all other 100 ms windows outside this period. We then computed the area under the ROC, which is equal to the probability that the input magnitude in a randomly drawn window at the latency of interest is greater than a random window in the background distribution.

Decoding from inferred inputs

We decoded target direction from all the samples of inferred input in 250 ms windows around either target presentation or corrective movement. We first split the targets into a training and test set, with each containing 50% of the data. We optimized the decoder for each monkey on the training set. As part of this optimization, we chose between two methods of decoding, support vector regression and random forest regression, using the sklearn implementation of both methods (Pedregosa et al., 2011). The optimization procedure included fitting what are generally the most important hyperparameters for these two methods (Probst et al., 2019), as well preprocessing parameters. For the support vector regression, we used a radial basis function and fit the L2 regularization hyperparameter from a range of 0.1 to 1.5. All other parameters were set to the sklearn defaults. For the random forest regressor, we fit the minimal samples for each leaf (from the values 1, 5, 10, 15), and the maximum number of features to consider at each split. The number of features (optimized) was either the full number of features, the square root of the number of features, or the base 2 logarithm of the number of features. Two preprocessing parameters were optimized: (1) the time window relative to the

reference event to use, and (2) the lag between the event and the hand position when computing the hand-centered direction to the target. The decoder hyperparameters were optimized using 5-fold cross validation on the training set. We then evaluated the selected decoder by training on the full training set and predicting target directions on the test set. For comparison, we also decoded from firing rates using the same procedure. Firing rates were computed as described above in *Decoding of movement kinematics*.

To compute performance of the position decoder, we predicted the x- and y-position in the relevant reference frame, then computed the general coefficient of determination (r^2) between the true and predicted values on each held-out set. To obtain a single performance score, we averaged the r^2 for the x- and y-coordinate, weighted by the variance in each dimension. For the direction and magnitude decoders, we first normalized each target position vector to the unit circle. We then jointly predicted the x- and y- coordinate of the normalized vector along with the magnitude of the unnormalized vector. The performance of the direction decoder was defined as the average of the r^2 between the predicted and actual x- and y-coordinates. Cartesian coordinates were chosen instead of angles to eliminate wraparound issues for decoding.

The best-performing hyperparameters on the hyperparameter training set were used to evaluate the model using the test set, and to compare hand-centric and shoulder-centric decoding using cross-validation. Generalization performance was evaluated using randomly shuffled 10-fold cross validation using the test set. This was repeated 10 times, for a total of 100 splits of the data. Both the hand-centric model and the shoulder-centric model were fit and tested using each of the 100 splits to get an estimate of the generalization performance of the models. To test the significance of the difference between the hand-centric and shoulder-centric model, we applied a paired-sample t-test with corrected variance for cross-validation (Nadeau & Bengio, 2003; Bouckaert & Frank, 2004).

Comparing coding of initial and corrective submovements

We binned the angle of the target direction relative to the hand into 12 equal-sized bins. For each bin, we computed the mean value of inferred input in a window from 350 to 50 ms before the initial movement. We performed the same procedure for corrective submovements, except using the same window relative to the corrective submovement.

This resulted in 24 direction-means each for the initial submovements and the corrective submovements (12 direction bins $\times$ 2 inferred inputs). We computed the Pearson correlation coefficient between the direction-means for the initial submovements and the corrective submovements. To combine across the two inferred inputs when computing this correlation, we first z-scored the direction-means in each movement condition before computing the correlation coefficient.

For Figure 7C, we performed this same procedure, except that instead of computing the means of inferred input for each direction bin, we computed the mean of the first 5 principal components of the smoothed firing rates (see above). This was done to denoise the tuning curve given that there were fewer corrective submovements than initial movements.

Figures

Figure 1

Figure 1. The Random Target Pursuit task elicits continuous movement in response to changing movement goals. **(a)** The monkey performed the random target pursuit task in the horizontal plane using a two-link exoskeletal robot. The monkey moves from a starting position (*gray square*) to the target (*blue square*). Once the target is acquired, a new target immediately appears (*green square*). **(b)** An example hand trajectory from one trial from Monkey RS. Starting position is indicated by the red dot, and the target locations as blue squares. The trajectory for each target is drawn with a separate color. **(c)** The speed profile (*green*) and population firing rate (*blue*) in recorded units for an example trial in Monkey MK. Multiple local minima were present in the speed profile for some targets, indicating submovements / corrections.

Figure 2

Figure 2. Modeling M1 population activity with Latent Factor Analysis via Dynamical Systems (LFADS). **(a)** Schematic of LFADS approach. For each trial, the spatiotemporal pattern of population activity is reduced to a vector representing an initial state of a dynamical system (*blue*), and a low dimensional time series that represents exogenous input not explained by the autonomous dynamics of that system (*red*). **(b)** Kinematic decoding from LFADS factors for models with and without a controller. Decoding performance using firing rates is shown for reference. Error bars represent standard deviation of performance across cross-validation splits. **(c)** Kinematic decoding performance as a function of LFADS input penalty. **(d)** Gini coefficient of the inferred inputs for models with a range of input penalties. **(e)** Kinematic decoding performance as a function of the dimensionality of inferred inputs used in the LFADS model. **(f)** Target position decoding performances, using the LFADS inferred inputs, as a function of the dimensionality of the inputs used in the model.

Figure 3

Figure 3. Example inferred input trace for one trial from monkeys (a) RS, (b) MK, and (c) RJ.

Figure 4

Figure 4. LFADS inputs exhibit transients that are locked to the appearance of new targets and prior to corrective movements. **(a)** Average controller magnitude relative to target presentation. SEMs are largely hidden by the traces. **(b)** Distribution of the timing of peaks in the inferred input relative to the appearance of a new target. Distributions are shown for four different threshold values above which peaks are included. Distributions for monkey RS. **(c)** Distribution of the timing of maximum inferred input magnitude aligned to target presentation (orange) and movement (dark blue). **(d)** Area under the receiver operating characteristic (ROC) curve showing how reliably particular windows in time can be distinguished from the rest of the trial using controller magnitude. Values for a range of 100 ms windows (see Methods). Shaded error is 95% bootstrapped confidence interval (the error region at some points may be narrower than the line thickness).

Figure 4 child figure 1

Figure 4 child figure 1. Same as Figure 4b for monkeys MK and RJ.

Figure 5

Figure 5. LFADS controller values predict the location of the next target. **(a)** Average of the inferred input for different directions of the target relative to the hand. Shading represents $\pm$ SEM. Monkey RS. **(b)** Cross-validated decoding performance of target position prediction using different time windows. X-axis values are centers of 250 ms windows used to predict target direction. **(c)** Decoding performance on held-out data of target position (*left*), target direction (*center*) and target distance (*right*) relative to both the hand and shoulder. Error bars represent standard deviation of performance across cross-validation splits. Shld, shoulder.

Figure 6

Figure 6. LFADS inferred inputs also predict corrective movements. **(a)** Speed of the hand across an example trial for monkey RJ. Target presentation times, gray lines; initial movements (*purple dots*) and corrections (*green dots*) identified from local speed minima. **(b)** Movement trajectory and inferred input value 1 for the same trial as in **(a)**. Black squares, target location or initial position of hand at start of trial. Arrows, sequence that targets were shown. Dots, position of hand at one LFADS sample (100 Hz). The color of the pale dot trail represents value of inferred input 1 at that time point. Initial movements and corrections shown as in **(a)**. **(c)** Average controller magnitude relative to corrective submovement onset. Shaded areas are SEMs. **(d)** Area under ROC curve across different windows as in Figure 4d, but relative to movement onset for initial and corrective movements rather than target appearance. Shaded error is 95% bootstrapped confidence interval (the error region at some points may be narrower than the line thickness). **(e)** Cross-validated decoding of target position, **(f)** direction, and **(g)** magnitude for corrective movements.

Figure 7

Figure 7. Comparison of controller representation of initial movements and corrective movements. **(a)** Average of inferred inputs for monkey RS for different directions for initial movements (*left*) and corrections (*right*). Input 1 shown on top, input 2 on bottom. **(b)** Correlation between mean controller value in window from 350 to 50 ms before the movement for initial movements vs. corrective movements. Each point represents the mean for one direction of movement for one input, z-scored using the mean and s.d. for that input. **(c)** Same as **(b)** but using 5 principal components of firing rates rather than controller values. **(d)** The correlation of direction-averaged inferred input around the initial movement to direction-averaged input around the maximum speed.

Table 1

Parameter Name	Value
_clip_value	80
batch_size	32
cell_clip_value	5
cell_weight_scale	1
ci_enc_dim	128
co_mean_corr_scale	0
co_prior_var_scale	0.1
con_dim	128
controller_input_lag	1
do_causal_controller	FALSE
do_feed_factors_to_controller	TRUE
do_reset_learning_rate	FALSE
do_train_encoder_only	FALSE
do_train_io_only	FALSE
do_train_prior_ar_atau	TRUE
do_train_prior_ar_nvar	TRUE
do_train_readin	TRUE
ext_input_dim	0
factors_dim	50
feedback_factors_or_rates	factors
gen_cell_input_weight_scale	1

gen_cell_rec_weight_scale	1
gen_dim	100
ic_dim	200
ic_enc_dim	64
Parameter Name (cont.)	Value
ic_post_var_min	0
ic_prior_var_max	0.1
ic_prior_var_min	0.1
ic_prior_var_scale	0.1
inject_ext_input_to_gen	FALSE
keep_prob	0.98
kl_ic_weight	1
kl_increase_steps	900
kl_start_step	0
l2_con_scale	500
l2_gen_scale	500
l2_increase_steps	900
l2_start_step	0
learning_rate_decay_factor	0.98
learning_rate_init	0.01
learning_rate_n_to_compare	6
learning_rate_stop	0.01
max_ckpt_to_keep	5
max_ckpt_to_keep_lve	5

max_grad_norm	200
ndatasets	1
output_dist	poisson
prior_ar_atau	10
prior_ar_nvar	0.1
temporal_spike_jitter_width	0

Table 1. Hyperparameter values for fitting LFADS. Parameter names are as in the code package.

References

- Ames, K. C., Ryu, S. I., & Shenoy, K. V. (2014). Neural dynamics of reaching following incorrect or absent motor preparation. *Neuron*, 81(2), 438–451.
<https://doi.org/10.1016/j.neuron.2013.11.003>
- Andrew Pruszynski, J., Omrani, M., & Scott, S. H. (2014). Goal-dependent modulation of fast feedback responses in primary motor cortex. *Journal of Neuroscience*, 34(13), 4608–4617.
<https://doi.org/10.1523/JNEUROSCI.4520-13.2014>
- Batista, A. P., Buneo, C. A., Snyder, L. H., & Andersen, R. A. (1999). Reach plans in eye-centered coordinates. *Science*, 285(5425), 257–260.
<https://doi.org/10.1126/SCIENCE.285.5425.257>
- Bouckaert, R. R., & Frank, E. (2004). Evaluating the replicability of significance tests for comparing learning algorithms. *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 3–12.
- Bremner, L. R., & Andersen, R. A. (2012). Coding of the Reach Vector in Parietal Area 5d. *Neuron*, 75(2), 342–351. <https://doi.org/10.1016/j.neuron.2012.03.041>
- Brown, T. G. (1914). On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of rhythmic activity in progression, and a theory of the evolution of function in the nervous system. *The Journal of Physiology*, 48(1), 18–46.
<https://doi.org/10.1113/JPHYSIOL.1914.SP001646>

- Buneo, C. A., Batista, A. P., Jarvis, M. R., & Andersen, R. A. (2008). Time-invariant reference frames for parietal reach activity. *Experimental Brain Research*, 188(1), 77–89. <https://doi.org/10.1007/s00221-008-1340-x>
- Buneo, C. A., Jarvis, M. R., Batista, A. P., & Andersen, R. A. (2002). Direct visuomotor transformations for reaching. *Nature*, 416(6881), 632–636. <https://doi.org/10.1038/416632a>
- Chicone, C. (2006). *Ordinary Differential Equations with Applications* (2nd ed.). Springer-Verlag. <https://doi.org/10.1007/0-387-35794-7>
- Churchland, M., Cunningham, J., & Kaufman, M. (2012). *Neural population dynamics during reaching*. [http://www.stat.columbia.edu/\\$\sim\\$scunningham/pdf/ChurchlandNature2012.pdf](http://www.stat.columbia.edu/$\sim$scunningham/pdf/ChurchlandNature2012.pdf)
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., & Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, 487(7405), 51–56. <https://doi.org/10.1038/nature11129>
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., & Shenoy, K. V. (2010). Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron*, 68(3), 387–400. <https://doi.org/10.1016/j.neuron.2010.09.015>
- Churchland, M. M., & Shenoy, K. V. (2007). Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. *https://Doi.Org/10.1152/Jn.00095.2007*, 97(6), 4235–4257. <https://doi.org/10.1152/JN.00095.2007>
- Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485), 1585–1599.
- Crammond, D. J., & Kalaska, J. F. (2000). Prior Information in Motor and Premotor Cortex: Activity During the Delay Period and Effect on Pre-Movement Activity. *Journal of Neurophysiology*, 84(2), 986–1005. <https://doi.org/10.1152/jn.2000.84.2.986>
- Dacre, J., Colligan, M., Clarke, T., Ammer, J. J., Schiemann, J., Chamosa-Pino, V., Claudi, F., Harston, J. A., Eleftheriou, C., Pakan, J. M. P., Huang, C.-C., Hantman, A. W., Rochefort, N. L., & Duguid, I. (2021). A cerebellar-thalamocortical pathway drives behavioral context-dependent movement initiation. *Neuron*, 109(14), 2326–2338.e8. <https://doi.org/10.1016/j.neuron.2021.05.016>
- Dickey, A. S., Amit, Y., & Hatsopoulos, N. G. (2013). Heterogeneous neural coding of corrective movements in motor cortex. *Frontiers in Neural Circuits*, 0(MAR), 51. <https://doi.org/10.3389/FNCIR.2013.00051>
- Doersch, C. (2016). *Tutorial on Variational Autoencoders*. <http://arxiv.org/abs/1606.05908>

Elliott, D., Hansen, S., Grierson, L. E. M., Lyons, J., Bennett, S. J., & Hayes, S. J. (2010). Goal-Directed Aiming: Two Components but Multiple Processes. *Psychological Bulletin*, 136(6), 1023–1044. <https://doi.org/10.1037/a0020958>

Elliott, D., Lyons, J., Hayes, S. J., Burkitt, J. J., Roberts, J. W., Grierson, L. E. M., Hansen, S., & Bennett, S. J. (2017). The multiple process model of goal-directed reaching revisited. In *Neuroscience and Biobehavioral Reviews* (Vol. 72, pp. 95–110). Elsevier Ltd. <https://doi.org/10.1016/j.neubiorev.2016.11.016>

Evarts, E. V. (1968). Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of Neurophysiology*, 31(1), 14–27. <https://doi.org/10.1152/jn.1968.31.1.14>

Fetz, E. E. (1992). Are movement parameters recognizable coded in the activity of single neurons?.pdf. *Behavior and Brain Sciences*, 15(4), 679–690.

Fu, Q. G., Flament, D., Coltz, J. D., & Ebner, T. J. (1995). Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *Journal of Neurophysiology*, 73(2), 836–854. <https://doi.org/10.1152/jn.1995.73.2.836>

Georgopoulos, a P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J.Neurosci.*, 2(11)(11), 1527–1537. <https://doi.org/citeulike-article-id:444841>

Georgopoulos, a P., Schwartz, a B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science (New York, N.Y.)*, 233(4771), 1416–1419. <https://doi.org/10.1126/science.3749885>

Hatsopoulos, N. G., Joshi, J., & O’Leary, J. G. (2004). Decoding continuous and discrete motor behaviors using motor and premotor cortical ensembles. *Journal of Neurophysiology*, 92(2), 1165–1174. <https://doi.org/10.1152/JN.01245.2003>

Hatsopoulos, N. G., Xu, Q., & Amit, Y. (2007). Encoding of Movement Fragments in the Motor Cortex. *Journal of Neuroscience*, 27(19), 5105–5114. <https://doi.org/10.1523/JNEUROSCI.3570-06.2007>

Hennequin, G., Vogels, T. P., & Gerstner, W. (2014). Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron*, 82(6), 1394–1406. <https://doi.org/10.1016/j.neuron.2014.04.045>

Hurley, N., & Rickard, S. (2009). Comparing Measures of Sparsity. *IEEE Transactions on Information Theory*, 55(10), 4723–4741. <https://doi.org/10.1109/TIT.2009.2027527>

Kalaska, J. F. (2009). From Intention to Action: Motor Cortex and the Control of Reaching Movements. In D. Sternad (Ed.), *Progress in Motor Control: A Multidisciplinary Perspective* (pp. 139–178). Springer US. https://doi.org/10.1007/978-0-387-77064-2_8

- Kalidindi, H. T., Cross, K. P., Lillicrap, T. P., Omrani, M., Falotico, E., Sabes, P. N., & Scott, S. H. (2020). *Rotational dynamics in motor cortex are consistent with a feedback controller* (p. 2020.11.17.387043). <https://doi.org/10.1101/2020.11.17.387043>
- Kaufman, M. T., Seely, J. S., Sussillo, D., Ryu, S. I., Shenoy, K. V., & Churchland, M. M. (2016). The largest response component in motor cortex reflects movement timing but not movement type. *ENeuro*, ENEURO.0085-16.2016. <https://doi.org/10.1523/ENeuro.0085-16.2016>
- Keshtkaran, M. R., Sedler, A. R., Chowdhury, R. H., Tandon, R., Basrai, D., Nguyen, S. L., Sohn, H., Jazayeri, M., Miller, L. E., & Pandarinath, C. (2021). A large-scale neural network training framework for generalized estimation of single-trial population dynamics 1. *BioRxiv*, 2021.01.13.426570. <https://doi.org/10.1101/2021.01.13.426570>
- Kingma, D. P., & Welling, M. (2014, December). Auto-encoding variational bayes. *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings*. <https://arxiv.org/abs/1312.6114v10>
- Lamarre, Y., Busby, L., & Spidalieri, G. (1983). Fast ballistic arm movements triggered by visual, auditory, and somesthetic stimuli in the monkey. I. Activity of precentral cortical neurons. *Journal of Neurophysiology*, 50(6), 1343–1358. <https://doi.org/10.1152/jn.1983.50.6.1343>
- Marder, E., & Calabrese, R. L. (1996). Principles of rhythmic motor pattern generation. *Physiological Reviews*, 76(3), 687–717.
- Michaels, J. A., Dann, B., Intveld, R. W., & Scherberger, H. (2015). Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. *Journal of Neuroscience*, 35(32), 11415–11432. <https://doi.org/10.1523/JNEUROSCI.1714-15.2015>
- Nadeau, C., & Bengio, Y. (2003). Inference for the generalization error. *Machine Learning*, 52(3), 239–281.
- Omrani, M., Kaufman, M. T., Hatsopoulos, N. G., & Cheney, P. D. (2017). Perspectives on classical controversies about the motor cortex. *Journal of Neurophysiology*, 118(3), 1828–1848. <https://doi.org/10.1152/jn.00795.2016>
- Osborne, L. C., Lisberger, S. G., & Bialek, W. (2005). A sensory source for motor variation. *Nature*, 437(7057), 412–416. <https://doi.org/10.1038/nature03961>
- Pandarinath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., Henderson, J. M., Shenoy, K. V., Abbott, L. F., & Sussillo, D. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10), 805–815. <https://doi.org/10.1038/s41592-018-0109-9>

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., & Dubourg, V. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12(10).

Perich, M. G., Conti, S., Badi, M., Bogaard, A., Barra, B., Wurth, S., Bloch, J., Courtine, G., Micera, S., Capogrosso, M., & Milekovic, T. (2020). *Motor cortical dynamics are shaped by multiple distinct subspaces during naturalistic behavior* (p. 2020.07.30.228767). <https://doi.org/10.1101/2020.07.30.228767>

Pesaran, B., Nelson, M. J., & Andersen, R. A. (2006). Dorsal Premotor Neurons Encode the Relative Position of the Hand, Eye, and Goal during Reach Planning. *Neuron*, 51(1), 125–134. <https://doi.org/10.1016/j.neuron.2006.05.025>

Probst, P., Boulesteix, A.-L., & Bischl, B. (2019). Tunability: Importance of Hyperparameters of Machine Learning Algorithms. *Journal of Machine Learning Research*, 20(53), 1–32.

Rao, N. G., & Donoghue, J. P. (2014). Cue to action processing in motor cortex populations. <https://doi.org/10.1152/Jn.00274.2013>, 111(2), 441–453. <https://doi.org/10.1152/JN.00274.2013>

Roitman, J. D., & Shadlen, M. N. (2002). Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task. *Journal of Neuroscience*, 22(21), 9475–9489. <https://doi.org/10.1523/JNEUROSCI.22-21-09475.2002>

Rouse, A. (2018). Cyclic, condition-independent activity in primary motor cortex predicts corrective movement behavior. *BioRxiv*, 453746. <https://doi.org/10.1101/453746>

Russo, A. A., Bittner, S. R., Perkins, S. M., Seely, J. S., London, B. M., Lara, A. H., Miri, A., Marshall, N. J., Kohn, A., Jessell, T. M., Abbott, L. F., Cunningham, J. P., & Churchland, M. M. (2018). Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron*, 97(4), 953–966.e8. <https://doi.org/10.1016/J.NEURON.2018.01.004>

Sabatini, D. A., & Kaufman, M. T. (2021). *A curved manifold orients rotational dynamics in motor cortex* (p. 2021.09.09.459647). <https://doi.org/10.1101/2021.09.09.459647>

Sauerbrei, B. A., Guo, J. Z., Cohen, J. D., Mischiati, M., Guo, W., Kabra, M., Verma, N., Mensh, B., Branson, K., & Hantman, A. W. (2020). Cortical pattern generation during dexterous movement is input-driven. *Nature*, 577(7790), 386–391. <https://doi.org/10.1038/s41586-019-1869-9>

Scott, S. (2008). Inconvenient Truths about neural processing in primary motor cortex. *J Physiol*. <http://e.guigon.free.fr/rsc/article/ScottSH08.pdf>

Seely, J. S., Kaufman, M. T., Ryu, S. I., Shenoy, K. V., Cunningham, J. P., & Churchland, M. M. (2016). Tensor Analysis Reveals Distinct Population Structure that Parallels the Different

Computational Roles of Areas M1 and V1. *PLOS Computational Biology*, 12(11), e1005164. <https://doi.org/10.1371/journal.pcbi.1005164>

Sergio, L. E., Hamel-Pâquet, C., & Kalaska, J. F. (2005). Motor Cortex Neural Correlates of Output Kinematics and Kinetics During Isometric-Force and Arm-Reaching Tasks. *Journal of Neurophysiology*, 94(4), 2353–2378. <https://doi.org/10.1152/jn.00989.2004>

Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience*, 36(1), 337–359. <https://doi.org/10.1146/annurev-neuro-062111-150509>

Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation Through Neural Population Dynamics. *Annual Review of Neuroscience*, 43, 249–275. <https://doi.org/10.1146/annurev-neuro-092619-094115>

Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P., & Black, M. J. (2006). Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural Computation*, 18(1), 80–118. <https://doi.org/10.1162/089976606774841585>

Wu, W., & Hatsopoulos, N. (2006). Evidence against a single coordinate system representation in the motor cortex. *Experimental Brain Research* 2006 175:2, 175(2), 197–210. <https://doi.org/10.1007/S00221-006-0556-X>

Yuste, R., MacLean, J. N., Smith, J., & Lansner, A. (2005). The cortex as a central pattern generator. *Nature Reviews Neuroscience*, 6(6), 477–483. <https://doi.org/10.1038/nrn1686>