

Single neuron firing rate statistics in motor cortex during execution and observation of movement

Xiyuan Jiang, *IEEE Student Member*, Stephen I. Ryu, *IEEE Member*, Krishna V. Shenoy, *IEEE Senior Member*, and Jonathan C. Kao, *IEEE Member*

Abstract— Mirror neurons, which fire during both the execution and observation of movement, are believed to play an important role in motor processing and learning. However, much work still remains to understand the similarities and differences in how these neurons compute in the motor cortex during movement execution and observation. Here, we performed experiments where a monkey both executes and observes a center-out-and-back task within the same experimental session. By recording from putatively the same neural population, we were able to analyze and compare single neuron statistics between movement execution and observation. We found that a majority of neurons in the primary motor cortex (M1) and dorsal premotor cortex (PMd) have statistically different firing rate statistics between movement execution and observation. As a result of this difference, we then wondered if neurons during movement observation exhibited a similar characteristic to those during movement execution: changing of preferred directions as a function of movement speed. Interestingly, we found that while observed movement speed is encoded in the neural population, it only alters a small proportion of the neuron's firing rate statistics. These results suggest that neural populations in M1 and PMd process information related to movement differently between execution and observation.

I. INTRODUCTION

The observation of movement is believed to play an important role in motor processing and learning. For example, observing movement is a key component of action understanding (e.g., [1], [2]) and imitation [3]. Further, evidence suggests that motor learning can be aided by mentally rehearsing a movement [4]-[7]. Several studies have suggested that “mirror neurons” found in the ventral premotor cortex (PMv) [6], [8], dorsal premotor cortex (PMd) [9] and primary motor cortex (M1) [10] of both

monkeys [8] and humans [11], play an important role in motor processing and internalization during observation [12]. These mirror neurons are active during the observation and execution of goal-directed movements [13] and are thought to reflect motor processing and sensory consequences of actions that we perform or that are performed by others [13], [14]. Further, the activity of mirror neurons has been reported to be largely “congruent” between observation and execution in M1 and PMd. These congruent mirror neurons have similar statistics irrespective of whether the monkey is passively observing a movement or executing it [9]. In a prior study, where preferred direction was the key statistic of interest, congruent neurons were reported to constitute approximately 70% of M1 neurons and 60% of PMd neurons. Amongst these neurons, the mean preferred direction differences were small, averaging less than 15 degrees. We note that in this study, reaches were executed with a KINARM [9].

We wondered if mirror neuron congruence is also characteristic in a center-out-and-back task where the monkey moves his native arm freely to acquire targets in a virtual environment, as opposed to with a KINARM [15]. Further, we were interested in assessing if the preferred directions of neurons exhibited similar behavior when reach statistics changed. In particular, Churchland and colleagues reported that movement speed may change preparatory and reach preferred directions [16]. If neurons during movement observation behave congruently to neurons during movement execution, we would expect to see changes in preferred direction when movement observation speed changes. However, if this is not the case, it suggests that motor cortex may perform computation related to reach execution and observation differently.

To address these questions, we recorded the same neural population from M1 and PMd within the same experimental session while a monkey was instructed to execute and observe movement. We performed single neuron analyses to assess consistencies and differences between neural activity during movement execution and observation. We found that, for the center-out-and-back task, most neurons have different single neuron firing rate statistics between movement execution and observation. In addition, we found that at a single neuron level, the speed of the observed movement does not substantially alter neurons' firing rate statistics, in contrast to results in movement execution [16].

This work was supported by an NSF Graduate Research Fellowship (JCK), as well as the NIH and DARPA (KVS).

X. Jiang is with the Electrical and Computer Engineering Graduate Program, University of California, Los Angeles, CA, 90095 (e-mail: xiyuanjiang@ucla.edu).

S. I. Ryu is with the Department of Neurosurgery at the Palo Alto Medical Foundation, Palo Alto, CA 94301 and with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 (e-mail: seoulmanmd@gmail.com).

K.V. Shenoy is with the Departments of Electrical Engineering, Bioengineering and Neurobiology, Bio-X and the Stanford Neuroscience Institute, Stanford University, Stanford, CA 94305. He is also with the Howard Hughes Medical Institute at Stanford University (e-mail: shenoy@stanford.edu).

J. C. Kao is with the Department of Electrical and Computer Engineering, University of California, Los Angeles, CA, 90095 (e-mail: kao@seas.ucla.edu).

II. METHODS

A. Neuron recording

All procedures and experiments were approved by the Stanford University Institutional Animal Care and Use Committee (IACUC). Experiments were conducted with an adult male rhesus macaque (J), implanted with two 96 electrode Utah arrays (Blackrock Microsystems Inc., Salt Lake City, UT) using standard neurosurgical techniques. Monkey J was implanted 56 months prior to the experiments. One electrode array was implanted in M1 and the other in PMd, as estimated visually from local anatomical landmarks.

The monkey was trained to make point-to-point reaches to targets in a 2D plane with a virtual cursor controlled by the contralateral arm or by a neural decoder [17]. The virtual cursor and targets were presented in a 3D environment (MSMS, MDDF, USC, Los Angeles, CA) described in [18]. Hand position data were measured with an infrared reflective bead tracking system (Polaris, Northern Digital, Ontario, Canada). Behavioral control and neural decode were run on separate PCs using the Simulink/xPC platform (Mathworks, Natick, MA). Eye positions were recorded via ISCAN (ISCAN, Inc., MA) Neural data were initially processed by the Cerebus recording system (Blackrock Microsystems Inc., Salt Lake City, UT) according to specifications described in [18]. Spike events were detected by setting a threshold value to -4.5 times the RMS voltage of the channel.

After thresholding the activity, we spike sorted the data via MKsort (<https://github.com/ripple-neuro/mksort>). We incorporated a criterion for determining single neurons based on interspike interval (ISI) violations, since sorting spike waveforms incorporates a subjective component. Our criterion was based on prior studies [19], [20] that only counted neurons that had ISI violations (< 3ms) more than 3% of the time. We also performed all analyses with less stringent ISI requirements (< 2ms more than 3% of the time, or no ISI requirements), and found no change in any conclusions.

B. Trial selection during movement observation

It is possible that during movement observation, the monkey ceases to pay attention to the moving cursor. We addressed this in two ways. First, the experimenter observed the monkey's eye position throughout the duration of the experiment, and stopped the task if the monkey disengaged. Second, we only chose trials that met the following criteria: (a) the angle between the monkey's eye position vector and cursor movement vector was less than 45°, (b) the length of monkey's eye position vector was larger than half of the length of cursor movement vector. Together, these selection criteria resulted in trials selected where the monkey's eye movements were consistent with cursor movement.

C. Bootstrap hypothesis test

To test for significant differences in single neuron firing rate statistics, we performed a bootstrap hypothesis test. We first resampled the firing rate for a given direction, then fit a tuning curve using across all reach directions [21]. Suppose

the two resampled data vectors are x and y , and we are interested in whether they have the same mean (null hypothesis, $h = 0$). We achieve this with the following method [22]:

1. Mean subtraction: $\hat{x} = x - \bar{x}$, $\hat{y} = y - \bar{y}$
2. Concatenation: $c = [\hat{x}, \hat{y}]$
3. Sample vectors mix_x and mix_y from c with replacement; each vector comprises 1000 samples.
4. Compute mean-centered bootstrap statistic:
 $\Delta_{ctrl} = mix_x - mix_y$, $\Delta_c = \Delta_{ctrl} - \bar{\Delta}_{ctrl}$
5. Sample vectors $real_x$ and $real_y$ from x and y respectively with replacement; each vector comprises 1000 samples.
6. Compute statistic: $\Delta_{real} = real_x - real_y$
7. Calculate p-value: $p = \frac{\#\{\Delta_{real} < \Delta_c\}}{1000}$
8. If $p < 0.05$, then reject the null hypothesis ($h = 1$) arguing for x and y having different means, else do not reject the null hypothesis ($h = 0$).

D. Generative model to predict the reaching speed

To decode reach speed from movement observation activity, we performed PCA on the high-dimensional neural firing rates $y = [y_1, y_2, \dots, y_n]$ to avoid dimensions with low variability. Such dimensions may bias the classifier without appropriate regularization. The low dimensional data (which we refer to as the PC scores) is $\mathbf{s} = [s_1, s_2, \dots, s_m]$, with $m < n$, and the generative model is

$$\begin{aligned} k_{res} &= \max_k Pr(c = k | \mathbf{s}) \\ &= \max_k p(\mathbf{s} | c = k) Pr(c = k) \quad (1) \end{aligned}$$

In this model, c is a random variable denoting speed. The term $p(\mathbf{s} | c = k)$ is the density of the PC scores given a reach speed, and is assumed to obey a multivariate Gaussian distribution. Its parameters (mean and covariance) are estimated by maximizing the likelihood of the observed PC scores, which corresponds to the empirical mean and covariance respectively. We assume the covariance is diagonal and shared between classes.

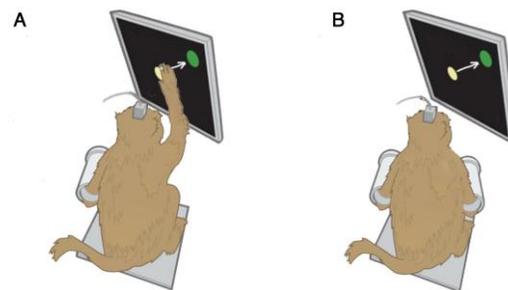


Figure 1. Behavioral task. (A) In execution of movement, the monkey controls the cursor with his hand. (B) In observation of movement, the monkey observes the cursor moving to each target while both arms are restrained.

III. RESULTS

We carried out experiments where a rhesus macaque performed a center-out-and-back task, as shown in Fig. 1. There were two conditions: either the monkey controls the cursor with his hand (execution) or observes the cursor moving while both of his arms are restrained (observation). To ensure the monkey was engaged in the observation task, we monitored his eye position during observation trials and only used trials where the monkey's eye positions were consistent with the cursor's start and end points. During execution, the monkey performed reaches in the space in front of him with the cursor position tied to his hand position; the targets were perceived in the space in front of him as they were presented via a Wheatstone stereograph configuration, previously described (e.g., [17], [18]). We recorded single neuron spiking activity from two Utah electrode arrays, implanted in M1 and PMd respectively. We recorded both observation and execution conditions in the same experimental session so that the recorded units were putatively the same between movement execution and observation. After spike sorting, we recorded anywhere from 70 to 82 single units between experimental sessions. The following sections describe the single unit analyses and statistics of the neurons between observation and execution.

A. A majority of neurons have different firing rate statistics between movement execution and observation

As we recorded from the same putative population between movement execution and observation, we first analyzed the degree of consistency in the neurons' firing rate statistics between conditions. In particular, we were interested in the extent to which the cosine tuning statistics were consistent in M1 and PMd activity during these tasks. It is worth noting that Tkach and colleagues studied a similar question in a random target pursuit task using a KINARM exoskeleton [15] and reported largely congruent activity during execution and observation [9]. Our task was different in that the monkey's arm is free to move in the space in front of him.

For each neuron recorded during movement execution and observation, we fit a cosine tuning model to the

recorded neural activity during the center-out-and-back task [21]. We analyzed three statistics from the cosine tuning model: the neuron's (1) mean firing rate, (2) modulation depth, defined as the maximum possible range of firing rates from the tuning curve model and (3) preferred direction, defined as the direction at which the tuning curve reaches its maximum firing rate. We then compared these statistics for the same neuron between movement execution and observation.

Interestingly, we found that a large proportion of neurons had significantly different preferred directions between the observation and execution conditions in M1 and PMd. Concretely, 49.5% of neurons in putative M1 and 60.0% of neurons in putative PMd had significantly different preferred directions ($p < 0.05$, bootstrap with 1000 resamples, see Methods). Histograms of the preferred direction changes are shown in Fig. 2A. The average magnitude change in preferred direction was 105.1 degrees amongst these neurons, indicating that statistically significant changes in preferred direction could also be very large in magnitude. We did not find that PMd had substantially more congruent mirror neurons than M1, as may be expected because PMd is anatomically closer to PMv Area F5, an area reported to contain a substantial proportion of mirror neurons [23]. These results show that the proportion of mirror neurons with congruent activity is approximately half of the population. Our observations are less than what was previously reported in M1 and PMd, although this may be related to task differences [9].

Apart from preferred directions, we also observed that most neurons had different modulation depths (Δ) and baseline firing rates (r_0). We observed that modulation depth during observation was smaller than during execution ($p < 0.05$, bootstrap). Across all neurons, the average modulation depth was 3.79 spikes/s less for movement observation than movement execution in putative M1, and 3.54 spikes/s in putative PMd. These results are shown in Fig 2B. We found that these changes were statistically significant in 62.8% of neurons in putative M1 and 65.4% of neurons in putative PMd. We similarly found that 88.6% of neurons in putative M1 and 76.3% of neurons in putative PMd had statistically significant changes in baseline firing rate, with the average difference in baseline firing rate being 4.86 spikes/s less for observation than execution across all

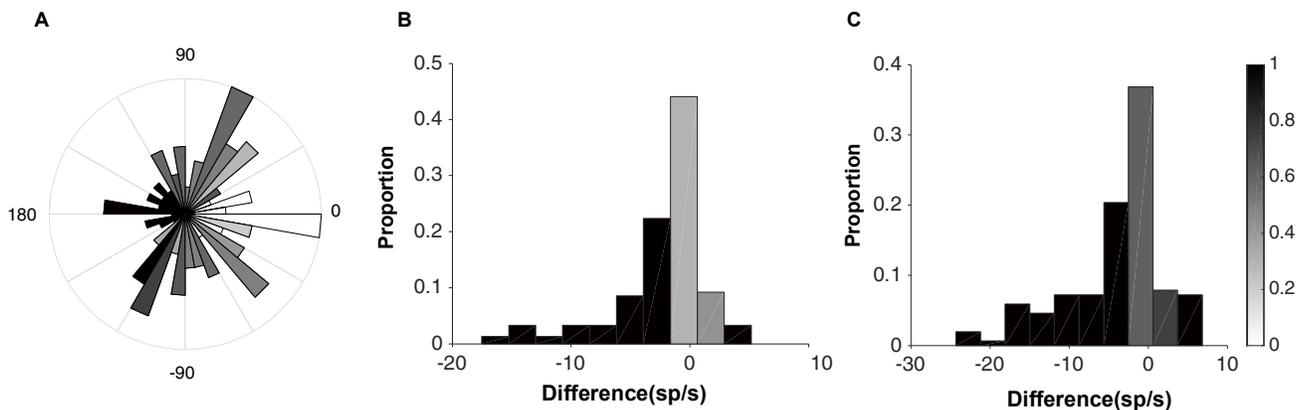


Figure 2. (A) A polar histogram of the difference in preferred direction (PD) between movement execution and observation. The grayscale indicates the proportion of neurons with a significant change in PD, with the color scale on the right. (B) Change in modulation depth. (C) Change in base firing rate.

neurons, as shown in Fig. 2C. Similar trends also held for the differences in variance and Fano factor between conditions, where 69.0% in M1 and 67.2% in PMd had significant changes in firing rate variance, and 43.2% in M1 and 49.0% in PMd had significant change in Fano factor. Across all neurons, the firing rate variance was smaller and the Fano factor was larger for observation condition.

B. Firing rate statistics during observation do not change substantially for different cursor speeds

Churchland and colleagues found that preparatory preferred directions were modulated by movement speed in a task where the monkey reached at two distinct speed ranges [16]. In the preparatory stage, around 32% of neurons in M1 and PMd had statistically significant changes in PD, with a mean of 39 degrees. Along these lines, we wondered if the neuron tuning statistics changed as a function of cursor speed during movement observation. If the mirror neuron system mimics the behavior of neurons during reaching, we may expect preferred directions to change across different observed movement speeds. Thus, we performed an additional experiment where the cursor speed was varied across five different levels, with speeds ranging from 4.62 cm/s to 6.53 cm/s. The different speeds were presented in separate experimental blocks, each comprising approximately 150 total center-out trials.

We first assessed if cursor speed was encoded in the firing rates of M1 and PMd during observation. To decode, we used a Gaussian generative model on the principal components of the neural activity (see Methods). We found that it was possible to decode the cursor speed far above chance from the neuron firing rates (Fig. 3A). In particular, we decoded the correct speed 34.6% of the time (chance 20%). These results indicate that the recorded neurons in M1 and PMd encode the speed of the observed cursor.

We next asked if this information was present in the single neuron firing rate statistics. Interestingly, when varying the cursor movement speed, we found that most neurons do not change their preferred directions (Fig. 3C). In particular, we found that 10.4% of neurons in M1 and 8.6% of neurons in PMd had significant changes in preferred direction for different cursor movement speeds. Among these neurons, the maximal average change in preferred directions (across any two speeds) was 39.7 degrees. These trends were similar for the modulation depth, where 10.4% of neurons in M1 and 7.1% of neurons in PMd had significantly different modulation depth. The average change in modulation depth was 0.72 spikes/s.

This indicates that while observation speed information is indeed encoded in the neural population, it is not readily apparent at the single neuron level. As an additional control to ensure that the approximately 10% of neurons changing their preferred directions were not solely responsible for the speed encoding, we decoded movement observation speed after removing the neurons that significantly changed their preferred directions (no matter during which speeds pairs) and observed that we could still reliably decode movement speed. In particular, when only considering neurons that do

not changed their preferred direction, we decoded the correct speed 29.5% of the time (change 20%), as shown in Fig. 3B.

This decoded speed is relatively modest due to a small proportion of neurons. However, population decoding results are often robust even when recording multi-units. Although we still used sorted units, we also performed this decoding analysis using units with less stringent ISI violation requirements. Concretely, in this population, the proportion of neurons with ISI violations less than 2ms more than 3% of the time is approximately 2% of the neurons. With this population, we found that the decoded speed was with these sorted units was 43.5% and when considering neurons that did not change their preferred direction, was 38.2%. These results indicate that speed is still strongly encoded, even amongst units that do not significantly change their firing rate statistics between movement observation and execution.

Together, these results demonstrate that the behavior of single neurons during movement observation is qualitatively different than during execution. Single neurons have smaller changes in tuning curve statistics than what may be expected from their behavior during execution. However, information about cursor speed is still encoded in the population. These results argue that the behavior of the mirror neurons differs qualitatively between movement execution and observation.

IV. DISCUSSION

Our results show that while neurons have a degree of congruent activity between movement execution and observation as previously reported, the activity of motor cortical neurons during the execution and observation of movement may differ in important ways. First, we found that a large proportion of neurons have statistically significant

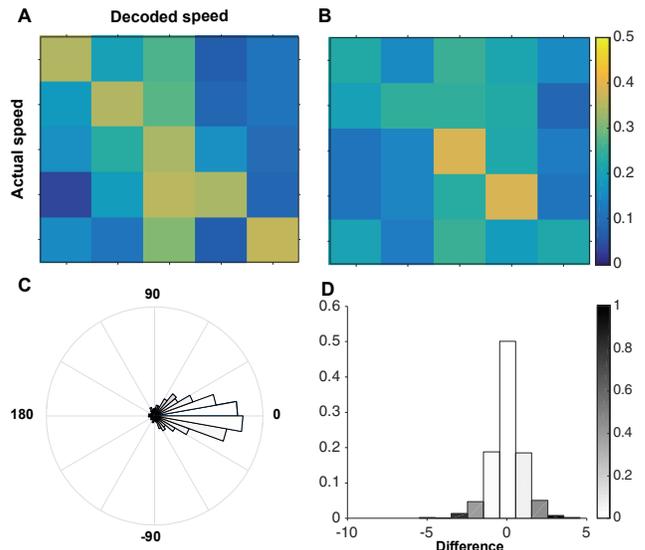


Figure 3. (A) Classification of observation speeds, using all neurons. Results are shown as a confusion matrix. (B) Classification using neurons without a significant change in preferred direction. (C) The polar histogram for the difference in preferred direction between different observation speeds. (D) The histogram for the difference in modulation depth(spikes/s) between different observation speeds.

changes in tuning curve parameters, including preferred direction and modulation depth. Given the Utah array's random sampling of neurons, as well as its 400-micron electrode separation, this suggests congruent mirror neurons do not make up a majority of cells in layer 3 / 4 motor cortical circuits. Rather, many neurons behave incongruently between conditions. Another qualitative difference in activity between movement execution and observation is that cursor movement speeds during movement observation do not substantially alter firing rate statistics. This contrasts with prior observations during movement execution, where preferred directions change in greater proportion.

These results have implications on the design of brain-machine interfaces (BMIs) in how training sets are collected. In pilot clinical trials, human participants are instructed to imagine making reaching movements (e.g. [24]-[26]). It is possible that the motor commands during imagination may have produced movement in the absence of motor injury. However, animal models used to develop algorithms for BMIs typically use able-bodied monkeys (e.g. [27]-[31]). In some instances, decoders are trained via movement execution and in others, via movement observation [32]. It is worth noting that our results contribute to a growing literature that decoders trained under these two paradigms may differ. While a limitation of animal models using movement execution incorporate proprioceptive feedback not present in a typical human BMI participant [33], activity during movement observation may differ substantially from imagined movements. For example, imagined movements may have produced muscle activity (as in movement execution) and have been shown to exhibit similar dynamics to macaque motor cortex during reaching [34].

To this end, future work should analyze neural population dynamics underlying movement execution and movement observation. For example, how do the dynamics of neural populations in M1 and PMd differ between movement execution and observation? Several studies have assessed the dynamics of neural populations, which describe how the neural population modulates itself through time in lawful ways through its recurrent connectivity. Recent work has resulted in concrete models of these dynamics [35]-[39]. In the motor system, these dynamics are predictive of movement generation [40]-[44]. For example, the dynamics describe how quickly one may make a reach [40] as well as if the reach was prepared or not [41].

V. REFERENCES

- [1] G. Rizzolatti, L. Fogassi, and V. Gallese, "Neurophysiological mechanisms underlying the understanding and imitation of action," *Nat. Rev. Neurosci.*, vol. 2, no. 9, pp. 661–670, Sep. 2001.
- [2] G. Rizzolatti and G. Luppino, "The cortical motor system," *Neuron*, vol. 31, no. 6, pp. 889–901, Sep. 2001.
- [3] M. Jeannerod, "The representing brain: Neural correlates of motor intention and imagery," *Behav. Brain Sci.*, vol. 17, no. 2, pp. 187–202, Jun. 1994.
- [4] M. Jeannerod and J. Decety, "Mental motor imagery: a window into the representational stages of action," *Curr. Opin. Neurobiol.*, vol. 5, no. 6, pp. 727–732, Dec. 1995.
- [5] A. Sirigu, J. R. Duhamel, L. Cohen, B. Pillon, B. Dubois, and Y. Agid, "The mental representation of hand movements after parietal cortex damage," *Science*, vol. 273, no. 5281, pp. 1564–1568, Sep. 1996.
- [6] P. Cisek and J. F. Kalaska, "Neural correlates of mental rehearsal in dorsal premotor cortex," *Nature*, vol. 431, no. 7011, pp. 993–996, Oct. 2004.
- [7] Vyas S, Even-Chen N, Stavisky SD, Ryu SI, Nuyujukian P, Shenoy KV "Neural population dynamics underlying motor learning transfer". *Neuron*. In press, 2018.
- [8] G. Rizzolatti and L. Craighero, "The mirror-neuron system," *Annu. Rev. Neurosci.*, vol. 27, pp. 169–192, 2004.
- [9] D. Tkach, J. Reimer, and N. G. Hatsopoulos, "Congruent activity during action and action observation in motor cortex," *J. Neurosci.*, vol. 27, no. 48, pp. 13241–13250, Nov. 2007.
- [10] J. Järveläinen, M. Schürmann, and R. Hari, "Activation of the human primary motor cortex during observation of tool use," *Neuroimage*, vol. 23, no. 1, pp. 187–192, Sep. 2004.
- [11] L. M. Oberman, J. A. Pineda, and V. S. Ramachandran, "The human mirror neuron system: a link between action observation and social skills," *Soc. Cogn. Affect. Neurosci.*, vol. 2, no. 1, pp. 62–66, Mar. 2007.
- [12] M. Iacoboni and M. Dapretto, "The mirror neuron system and the consequences of its dysfunction," *Nat. Rev. Neurosci.*, vol. 7, no. 12, pp. 942–951, Dec. 2006.
- [13] S. Acharya and S. Shukla, "Mirror neurons: Enigma of the metaphysical modular brain," *J. Nat. Sci. Biol. Med.*, vol. 3, no. 2, pp. 118–124, Jul. 2012.
- [14] D. Tkach, J. Reimer, and N. G. Hatsopoulos, "Observation-based learning for brain-machine interfaces," *Curr. Opin. Neurobiol.*, vol. 18, no. 6, pp. 589–594, 2008.
- [15] S. H. Scott, "Apparatus for measuring and perturbing shoulder and elbow joint positions and torques during reaching," *J. Neurosci. Methods*, vol. 89, no. 2, pp. 119–127, Jul. 1999.
- [16] M. M. Churchland, G. Santhanam, and K. V. Shenoy, "Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach," *J. Neurophysiol.*, vol. 96, pp. 3130–3146, 2006.
- [17] J. P. Cunningham, P. Nuyujukian, V. Gilja, C. A. Chestek, S. I. Ryu, and K. V. Shenoy, "A closed-loop human simulator for investigating the role of feedback control in brain-machine interfaces," *J. Neurophysiol.*, vol. 105, pp. 1932–1949, 2011.
- [18] V. Gilja et al., "A high-performance neural prosthesis enabled by control algorithm design," *Nat. Neurosci.*, vol. 15, no. 12, pp. 1752–1757, Nov. 2012.
- [19] C. Chandrasekaran, D. Peixoto, W. T. Newsome, and K. V. Shenoy, "Laminar differences in decision-related neural activity in dorsal premotor cortex," *Nat. Commun.*, vol. 8, no. 1, p. 614, Sep. 2017.
- [20] U. Rutishauser, E. M. Schuman, and A. N. Mamelak, "Online detection and sorting of extracellularly recorded action potentials in human medial temporal lobe recordings, in vivo," *J. Neurosci. Methods*, vol. 154, no. 1–2, pp. 204–224, Jun. 2006.
- [21] A. P. Georgopoulos, J. F. Kalaska, R. Caminiti, and J. T. Massey, "On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex," *J. Neurosci.*, vol. 2, no. 11, pp. 1527–1537, Nov. 1982.
- [22] J. M. Fan, P. Nuyujukian, J. C. Kao, C. A. Chestek, S. I. Ryu, and K. V. Shenoy, "Intention estimation in brain-machine interfaces," *J. Neural Eng.*, vol. 11, no. 1, p. 016004, Jan. 2014.
- [23] J. M. Kilner and R. N. Lemon, "What we know currently about mirror neurons," *Curr. Biol.*, vol. 23, no. 23, pp. R1057–R1062, 2013.
- [24] L. R. Hochberg et al., "Reach and grasp by people with tetraplegia using a neurally controlled robotic arm," *Nature*, vol. 485, no. 7398, pp. 372–375, May 2012.
- [25] J. L. Collinger et al., "High-performance neuroprosthetic control by an individual with tetraplegia," *Lancet*, vol. 381, no. 9866, pp. 557–564, Feb. 2013.
- [26] V. Gilja et al., "Clinical translation of a high-performance neural prosthesis," *Nat. Med.*, vol. 21, no. 10, pp. 1142–1145, 2015.
- [27] P. Nuyujukian, J. M. Fan, V. Gilja, P. S. Kalanithi, C. A. Chestek, and K. V. Shenoy, "Monkey models for brain-machine interfaces: the need for maintaining diversity," in *Proceedings of the 33rd Annual Conference of the IEEE EMBS*, 2011, vol. 2011, pp. 1301–1305.
- [28] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz, "Cortical control of a prosthetic arm for self-feeding," *Nature*, vol. 453, no. 7198, pp. 1098–1101, Jun. 2008.

- [29] P. J. Ifft, S. Shokur, Z. Li, M. A. Lebedev, and M. A. L. Nicolelis, "A Brain-Machine Interface Enables Bimanual Arm Movements in Monkeys," *Sci. Transl. Med.*, vol. 5, no. 210ra154, Nov. 2013.
- [30] P. T. Sadtler et al., "Neural constraints on learning," *Nature*, vol. 512, no. 7515, pp. 423–426, Aug. 2014.
- [31] M. M. Shanechi et al., "Rapid control and feedback rates enhance neuroprosthetic control," *Nat. Commun.*, vol. 8, p. 13825, 2017.
- [32] K. V. Shenoy and J. M. Carmena, "Combining decoder design and neural adaptation in brain-machine interfaces," *Neuron*, vol. 84, no. 4, pp. 665–680, 2014.
- [33] A. J. Suminski, D. C. Tkach, A. H. Fagg, and N. G. Hatsopoulos, "Incorporating feedback from multiple sensory modalities enhances brain-machine interface control," *Journal of Neuroscience*, vol. 30, no. 50, pp. 16777–16787, 2010.
- [34] C. Pandarinath et al., "Neural population dynamics in human motor cortex during movements in people with ALS," *Elife*, vol. 4, pp. 1–9, 2015.
- [35] M. M. Churchland et al., "Neural population dynamics during reaching," *Nature*, vol. 487, no. 7405, pp. 51–56, Jul. 2012.
- [36] J. C. Kao, P. Nuyujukian, S. I. Ryu, M. M. Churchland, J. P. Cunningham, and K. V. Shenoy, "Single-trial dynamics of motor cortex and their applications to brain-machine interfaces," *Nat. Commun.*, vol. 6, p. 7759, Jul. 2015.
- [37] L. Buesing, J. Bill, B. Nessler, and W. Maass, "Neural Dynamics as Sampling: A Model for Stochastic Computation in Recurrent Networks of Spiking Neurons," *PLoS Comput. Biol.*, vol. 7, no. 11, p. e1002211, 2011.
- [38] J. H. Macke, L. Buesing, J. P. Cunningham, B. M. Yu, K. V. Shenoy, and M. Sahani, "Empirical models of spiking in neural populations," in *Advances in Neural Information Processing Systems 24*, J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2011, pp. 1350–1358.
- [39] K. L. Briggman and W. Denk, "Towards neural circuit reconstruction with volume electron microscopy techniques," *Curr. Opin. Neurobiol.*, vol. 16, no. 5, pp. 562–570, Oct. 2006.
- [40] A. Afshar, G. Santhanam, B. M. Yu, S. I. Ryu, M. Sahani, and K. V. Shenoy, "Single-trial neural correlates of arm movement preparation," *Neuron*, vol. 71, no. 3, pp. 555–564, Aug. 2011.
- [41] K. C. Ames, S. I. Ryu, and K. V. Shenoy, "Neural dynamics of reaching following incorrect or absent motor preparation," *Neuron*, vol. 81, no. 2, pp. 438–451, Jan. 2014.
- [42] M. T. Kaufman, M. M. Churchland, S. I. Ryu, and K. V. Shenoy, "Cortical activity in the null space: permitting preparation without movement," *Nat. Neurosci.*, vol. 17, no. 3, pp. 440–448, Mar. 2014.
- [43] M. T. Kaufman, M. M. Churchland, S. I. Ryu, and K. V. Shenoy, "Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex," *Elife*, vol. 4, p. e04677, May 2015.
- [44] M. T. Kaufman, J. S. Seely, D. Sussillo, S. I. Ryu, K. V. Shenoy, and M. M. Churchland, "The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type," *eNeuro*, vol. 3, no. 4, 2016.