

Mirror Neuron Populations Represent Sequences of Behavioral Epochs During Both Execution and Observation

Kevin A. Mazurek,^{1,2} Adam G. Rouse,^{2,3} and Marc H. Schieber^{1,2,3,4}

¹Department of Neurology, ²Del Monte Institute for Neuroscience, ³Department of Neuroscience, University of Rochester Medical Center, Rochester, New York 14642, and ⁴Department of Biomedical Engineering, University of Rochester Medical Center, Rochester, New York 14627

Mirror neurons (MNs) have the distinguishing characteristic of modulating during both execution and observation of an action. Although most studies of MNs have focused on various features of the observed movement, MNs also may monitor the behavioral circumstances in which the movement is embedded, including time periods preceding and following the observed movement. Here, we recorded multiple MNs simultaneously from implanted electrode arrays as two male monkeys executed and observed a reach, grasp, and manipulate task involving different target objects. MNs were recorded from premotor cortex (PM-MNs) and primary motor cortex (M1-MNs). During execution trials, hidden Markov models (HMMs) applied to the activity of either PM-MN or M1-MN populations most often detected sequences of four hidden states, which we named according to the behavioral epoch during which each state began: initial, reaction, movement, and final. The hidden states of MN populations thus reflected not only the movement, but also three behavioral epochs during which no movement occurred. HMMs trained on execution trials could decode similar sequences of hidden states in observation trials, with complete hidden state sequences decoded more frequently from PM-MN populations than from M1-MN populations. Moreover, population trajectories projected in a 2D plane defined by execution trials were preserved in observation trials more for PM-MN than for M1-MN populations. These results suggest that MN populations represent entire behavioral sequences, including both movement and non-movement. PM-MN populations showed greater similarity than M1-MN populations in their representation of behavioral sequences during execution versus observation.

Key words: grasping; hidden Markov models; manipulation; neural trajectory; principal component analysis; reaching

Significance Statement

Mirror neurons (MNs) are thought to provide a neural mechanism for understanding the actions of others. However, for an action to be understood, both the movement per se and the non-movement context before and after the movement need to be represented. We found that simultaneously recorded MN populations encoded sequential hidden neural states corresponding approximately to sequential behavioral epochs of a reach, grasp, and manipulate task. During observation trials, hidden state sequences were similar to those identified in execution trials. Hidden state similarity was stronger for MN populations in premotor cortex than for those in primary motor cortex. Execution/observation similarity of hidden state sequences may contribute to understanding the actions of others without actually performing the action oneself.

Introduction

Mirror neurons (MNs) discharge both as a movement is being performed by the subject (execution) and when the subject ob-

serves the same movement being performed by another individual (observation). MNs that discharged in relation to grasping actions of the hand were originally identified in area F5 of the macaque ventral premotor cortex (PMv) (di Pellegrino et al., 1992; Gallese et al., 1996; Rizzolatti et al., 1996). Neurons that discharge during both execution and observation also have been found in other cortical areas, including the dorsal premotor cortex, the primary motor cortex (M1), and the inferior parietal lobule (Cisek and Kalaska, 2004; Tkach et al., 2007; Rozzi et al., 2008; Kilner et al., 2009; Kraskov et al., 2009; Dushanova and

Received Dec. 8, 2017; revised March 26, 2018; accepted April 3, 2018.

Author contributions: K.A.M. wrote the first draft of the paper; K.A.M., A.G.R., and M.H.S. edited the paper; K.A.M. and M.H.S. designed research; K.A.M., A.G.R., and M.H.S. performed research; K.A.M. and M.H.S. analyzed data; K.A.M. wrote the paper.

This work was supported by the National Institute of Neurological Disorders and Stroke—National Institutes of Health (Grant F32NS093709 to K.A.M., Grant K99NS101127 to A.G.R., and Grant R01NS079664 to M.H.S.). We thank Marsha Hayles for editorial comments.

The authors declare no competing financial interests.

Correspondence should be addressed to Marc H. Schieber, Department of Neurology, University of Rochester Medical Center, 601 Elmwood Avenue, Box 673, Rochester, NY 14642. E-mail: mschiebe@ur.rochester.edu.

DOI:10.1523/JNEUROSCI.3481-17.2018
Copyright © 2018 the authors 0270-6474/18/384441-15\$15.00/0

Donoghue, 2010; Vigneswaran et al., 2013). MN activity traditionally has been interpreted as providing an abstracted representation of the action's goal (Gallese et al., 2004). However, MN activity can differ depending on a variety of circumstances under which the goal of the observed action remains unchanged. MN activity may vary depending on whether the observed movement occurs within versus beyond the reach of the subject's arm (Umiltà et al., 2001; Caggiano et al., 2009) or depending on the subject's point of view in observing the movement (Caggiano et al., 2011; Maranesi et al., 2017). MNs thus carry information about the context in which particular movements are made.

Although most studies of MNs have focused on discharge during the movement, contextual information is present before, during, and after the movement occurs. When having dinner with someone, you might expect to observe the other person pick up a piece of bread and put it in his/her mouth, but not to put it on his/her shoulder, although the movement would be very similar. MNs have been shown to discharge as monkeys both grasped-to-eat and grasped-to-place an object near the shoulder (Fogassi et al., 2005; Bonini et al., 2010). We hypothesized that MNs monitor not only observed movements per se, but entire sequences of behavioral epochs in which movements are embedded.

Although the firing rate of MNs before and after movements generally is relatively low, the simultaneous activity of MN populations nevertheless might carry information not only about the movements, but also about behavioral epochs preceding and following the movements. Hidden Markov models (HMMs) can be used to infer unobservable state sequences based on the statistics of observable neural events (Rabiner, 1989; Jones et al., 2007). Neural activity recorded from M1 and/or the PM of nonhuman primates has been shown to evolve through sequences of hidden states during trials of reaching and grasping movements (Abeles et al., 1995; Kemere et al., 2008; Kang et al., 2015). Such hidden states provide one way of representing high-dimensional neural population activity in a discretized lower dimensional state space.

We therefore applied HMMs to detect hidden state sequences in populations of MNs during both execution and observation of a reach, grasp, and manipulate (RGM) task. We recorded neural activity from both PM and M1 as monkeys performed the RGM task and again as the monkeys observed a human experimenter performing the task. We then selected neurons that were modulated significantly during both execution and observation, and applied HMMs to these MN populations in PM (PM-MNs) and in M1 (M1-MNs) separately to preserve any characteristics distinguishing the two cortical areas. Rather than showing one state during movement and another non-movement state, MN populations showed four distinct hidden states during individual RGM trials, which we named according to the behavioral epoch during which each state began: initial, reaction, movement, and final. HMMs trained on execution trials identified similar sequences of hidden states during observation trials more robustly for PM-MN than for M1-MN populations. Moreover, when mapped to a lower dimensional subspace, the four sequential hidden states corresponded to sequential regions of neural population trajectories.

Materials and Methods

Experimental preparation

All procedures for the care and use of nonhuman primates followed the *Guide for the Care and Use of Laboratory Animals* and were approved by the University Committee on Animal Resources at the University of Rochester. Two male rhesus monkeys (*Macaca mulatta*; Monkey L and Monkey X, weights 9 and 11 kg, respectively) were used in the present

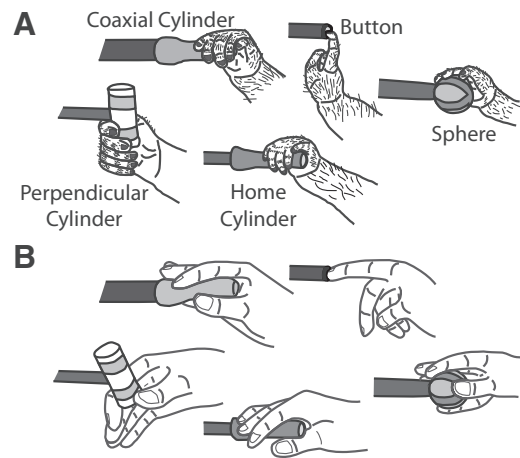


Figure 1. Grasp hand shapes. **A, B**, Line drawings made from video frames depict the hand shapes used by the monkey during execution trials (**A**) and those used by the human and observed by the monkey during observation trials (**B**). The one central object (home cylinder) and four peripheral target objects are labeled in **A** and illustrated here in positions reflecting their relative spatial arrangement as viewed by the monkey, although each object has been drawn in a side view as seen from a camera positioned to the monkey's left side.

experiments. As described previously (Mollazadeh et al., 2011), both monkeys had been implanted with floating microelectrode arrays (FMAs; MicroProbes) in the M1 and the PM of the left hemisphere. Each FMA had 16 recording electrodes of various lengths (1–9 mm). Recordings for the present study were obtained from four FMAs implanted in M1 and four in PM for Monkey L and from six FMAs in M1 and two in PM for Monkey X. For both monkeys, the PM electrodes were located in the posterior bank and lip of the arcuate sulcus, with the majority in the PMv.

Behavioral task

Execution. Each monkey was trained to sit in a primate chair and perform a RGM task with its right arm as described in detail previously (Mollazadeh et al., 2011). The center "home" object was a coaxial cylinder 20 mm in diameter and 60 mm in length. Surrounding the home cylinder were four target objects: a perpendicular cylinder, another coaxial cylinder, a button, and a sphere (Fig. 1). The target objects were located at 45° intervals on a circle of 13 cm radius centered on the home object. Each monkey used its right hand to perform the task while the left arm was restrained in the primate chair. The monkey initiated each trial by pulling the home object for a variable initial hold period (1.0–1.5 s). Then, a ring of blue LEDs illuminated around the base of one of the target objects, instructing the monkey to reach to, grasp, and manipulate that object, pulling the perpendicular or coaxial cylinder, pushing the button, or turning the sphere. When a given object was manipulated, a second ring of green LEDs at the base of that object was illuminated. The monkey then was required to maintain the manipulated position for a final hold of fixed duration (Monkey L: 0.9 s; Monkey X: 1.0 s). Upon successful completion of a trial, a liquid reward was delivered.

The target object for each trial was selected in a pseudorandom block design. Blocks consisted of one trial involving each of the four target objects presented in a random order. After successful completion of these four trials in a given block, the order of objects was randomly reshuffled for the next block. For trials ending in an error, the subsequent trial instructed the same target object until a trial was performed successfully. This prevented the monkey from avoiding trials of any particular object. The behavioral task was controlled by custom software written in TEMPO (Reflective Computing), which also sent 8-bit behavioral event marker codes into the recorded data stream. These codes marked the time at which certain behavioral events occurred, including the onset of the instruction, release of the center object indicating movement onset, start of the final hold period, and completion of the final hold period. Behavioral epochs were defined based on these behavioral event markers: the initial hold epoch was defined as the time from the start of a trial to

instruction onset; the reaction time epoch was defined from instruction onset to movement onset; the movement epoch was defined from movement onset to the start of the final hold; and the final hold epoch was defined from the start to the end of the final hold period.

Observation. During observation trials, with both hands restrained in the primate chair, the monkey observed an experimenter perform the same RGM task. Standing on the monkey's right side, an experimenter grasped the center object, received a blue LED instruction, reached to, grasped, and manipulated the target object, and then held the target object for the final hold duration. Upon successful completion of a trial performed by the experimenter, the monkey received a reward. In addition, to keep the monkey attentive, the experimenter occasionally made errors, after which the monkey received no reward. Incorrect trials were repeated until successfully completed. Because the monkey therefore could predict the object involved on any trial preceded by an error, only successful trials not preceded by an error trial were included in offline analyses of both execution and observation.

Neural recordings

During recording sessions, each monkey performed RGM trials itself (execution) and then observed the experimenter performing RGM trials (observation) in separate blocks of trials. Before beginning data collection, thresholds were set for spiking activity from each of 128 recording electrodes using Plexon's Sort Client. During both execution and observation trials, neuron spike times and action potential waveforms were captured using a MAP System (Plexon MAP Software, RRID:SCR_003170), along with the behavioral event marker codes generated by the TEMPO task-control system. Offline, neuron spikes were sorted further using Plexon's Offline Sorter followed by a custom sorting algorithm.

Each sorted unit was analyzed for task-related modulation during execution and observation trials using repeated-measures ANOVA (RM-ANOVA). For each successful trial involving any of the four target objects, the spike rate of the unit was binned in 100 ms nonoverlapping windows aligned on instruction onset. Using the MATLAB functions `rmfit` and `ranova` (The MathWorks, RRID:SCR_001622), a repeated-measures model was fit to the data (`rmfit`) and a two-way ANOVA (time \times object) was performed to detect significant effects (RM-ANOVA). The unit was considered to modulate significantly during execution or observation if a significant effect of time or of the interaction between time and object was detected (Bonferroni corrected for the number of effects being examined, $p < 0.05/2 = 0.025$). Any unit that modulated significantly for both behavioral conditions (execution and observation) was considered to be a MN.

Each MN was classified as a definite single-unit, probable single-unit, or multiunit based on the signal-to-noise ratio (SNR) of its captured waveforms and the fraction of false-positive spikes (f_F) that may have originated from another neuron or from noise (Meunier et al., 2003; Hill et al., 2011; Rouse and Schieber, 2016). The value of f_F was estimated from the number of interspike interval (ISI) violations (< 1 ms), ν_{ISI} , using the following equation (Hill et al., 2011):

$$f_F = 1 - \left(\sqrt{\frac{1}{4} - \frac{\nu_{ISI} \times T}{2(t_{max} - t_{min})N^2}} + \frac{1}{2} \right) \quad (1)$$

where N is the total number of spikes, T is the duration of the recorded session, t_{max} is the refractory period (chosen at 1 ms), and t_{min} is the waveform sort width after threshold crossing, during which no other spikes could have been detected (0.675 ms). The SNR was calculated as follows: $SNR = A/2\sigma_{noise}$, where A is the peak-to-peak amplitude of the mean spike waveform and σ_{noise} is the SD of the residual noise after subtracting the mean waveform from each of 32 sampled time points. Each MN was classified as follows: (1) MNs with an $SNR \geq 3$ and no ISI violations ($\nu_{ISI} = 0, f_F = 0$) were considered "definite" single units; (2) MNs with $SNR \geq 2.5$ and an estimated 90% or more true spikes ($f_F \leq 0.1$) were considered to be "probable" single units; and (3) MNs with $2.5 > SNR \geq 1.5$ and/or $f_F > 0.1$ were considered "multiunits." Any units with $SNR < 1.5$ were discarded from further analysis.

For each MN, we calculated two metrics to quantify the MN's selectivity across the four objects and its execution/observation similarity,

both focused on the movement epoch during which MNs typically have been studied. First, we calculated the change in firing rate between baseline (averaged over the 500 ms immediately preceding the instruction onset) and the movement epoch (averaged from movement onset to the start of the final hold) in each trial. We averaged this change in firing rate across all trials separately for each object (\bar{R}_{obj}). We then considered these four average firing rates as a vector, and normalized them by the length of that vector (L2 norm, $\|\bar{R}\|_2$) as follows:

$$r_{obj} = \frac{\bar{R}_{obj}}{\|\bar{R}\|_2}$$

thereby capturing the MN's relative changes in firing rate associated with each of the four objects as a 4D unit vector. If a MN was ideally selective for one object, then the normalized firing rates, r_{obj} , would be 1 for one of the objects and 0 for the others (e.g., $r_1 = 1, r_2 = 0, \dots, r_4 = 0$). If a MN was ideally unselective, then the average firing rate would be equal for all four objects ($\bar{R}_{obj} = \bar{R}_{unsel}$), yielding four normalized firing rates, r_{unsel} , all of equal magnitude as follows:

$$r_{unsel} = \frac{\bar{R}_{unsel}}{\|\bar{R}_{unsel}\|_2} = \frac{1}{\sqrt{N_{obj}}} = \pm \frac{1}{2}$$

To calculate a selectivity index, SEL, the normalized firing rates for each MN were rectified to treat increases and decreases equivalently. The unit vectors for any given MN, for an ideally unselective MN, and for an ideally selective MN then are as follows:

$$\mathbf{r}_{abs} = \begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ r_4 \end{bmatrix}, \mathbf{r}_{unsel} = \begin{bmatrix} r_{unsel} \\ r_{unsel} \\ r_{unsel} \\ r_{unsel} \end{bmatrix}, \mathbf{r}_{sel} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

The arc length distance (d_{max}) between the ideally selective (\mathbf{r}_{sel}) and the ideally unselective (\mathbf{r}_{unsel}) MN is then:

$$d_{max} = \cos^{-1}(\mathbf{r}_{unsel} \cdot \mathbf{r}_{sel})$$

and the arc length between the unit vector representing any given MN and the ideally unselective MN is:

$$d = \cos^{-1}(\mathbf{r}_{unsel} \cdot \mathbf{r}_{abs}).$$

We then computed our selectivity index (SEL) as follows:

$$SEL = \frac{d}{d_{max}}$$

SEL can vary between 0 (ideally unselective) and 1 (ideally selective for one object, regardless of which object). For each MN, a SEL was calculated separately for execution and for observation.

To quantify the similarity of MN activity during execution and observation, for each MN, we considered the four r_{obj} (without rectification) from execution trials as one vector (\mathbf{r}_E) and the four r_{obj} from observation trials as another vector (\mathbf{r}_O):

$$\mathbf{r}_E = \begin{pmatrix} r_{E1} \\ \cdot \\ \cdot \\ \cdot \\ r_{E4} \end{pmatrix}, \mathbf{r}_O = \begin{pmatrix} r_{O1} \\ \cdot \\ \cdot \\ \cdot \\ r_{O4} \end{pmatrix}.$$

We then calculated our execution/observation similarity index (SIM) as follows:

$$SIM = \mathbf{r}_E \cdot \mathbf{r}_O$$

If the MN's movement epoch activity during execution and observation was ideally similar, then the two vectors would align and the SIM would be 1. If the MN's activity was completely dissimilar during execution versus observation, then the two vectors would be orthogonal and the SIM would be 0. If activity was modulated in the opposite directions

during execution versus observation but with the same relative magnitude for each of the four objects, then the SIM would be -1 .

HMMs

For both monkeys, HMMs were trained to detect four states in successful execution trials separately for each target object and for each MN population (PM or M1) in each session. Training on PM-MN and M1-MN populations separately maintained any distinguishing population characteristics during the task. Successful trials included for analysis required that the trial had been performed correctly (i.e., a reward had been delivered) and that the preceding trial also had been performed correctly (so that the monkey could not anticipate the instructed object).

For each trial, spike times of all MNs in the simultaneously recorded population were used to create a single input observation sequence for training the HMMs. The input observation sequence for each trial spanned the time from 500 ms before the instruction onset until 500 ms after completion of the final hold. Because the monkey's reaction and movement times varied from trial to trial, so did the duration of the input observation sequences. Spike times were binned in 2 ms windows. A unique MN identifier (ID) ranging from 1 to N , the total number of MNs in the population, was assigned to any bin in which a spike from a given MN had occurred. A zero was assigned to any bin in which no MN spike occurred. In the event that multiple MNs discharged during the same 2 ms bin, the ID of one of the discharging MNs was selected randomly and assigned to that bin. The resulting 1D sequence of MN IDs and zeroes then described the population activity in each trial. These sequences were used to train HMMs. The Baum–Welch algorithm (Baum et al., 1970) recursively estimated the following HMM parameters: a transition matrix (the probabilities of a transition from the current state to any of the four states) and an emission matrix (the probabilities that a given unit will fire in a given state). Estimates of the model parameters were updated either for 500 iterations or until the log-likelihood converged to less than a tolerance factor ($<10^{-6}$). The Baum–Welch algorithm guaranteed approaching a local maximum log-likelihood, though not necessarily the global maximum. Therefore, each model was trained starting at 10 different initial conditions to improve the chances of approximating the global maximum log-likelihood. From these 10 trained models, the resulting model with the greatest log-likelihood was selected for subsequent analyses (Jones et al., 2007). Note that the training of HMMs incorporated no information about the timing of the behavioral events.

Hidden state analyses

Once an HMM had been trained for a given object and MN population, we quantified the extent to which the sequence of hidden states was consistent from trial to trial. A state was considered active if its probability exceeded an arbitrarily chosen threshold of 0.6. An ordinal sequence of active states then was determined for each trial. For execution trials, the state sequence consistency was quantified as the percentage of all trials for a given target object that had the same order of hidden states. Because the monkeys' reaction times and movement times varied from trial to trial, hidden state sequences were analyzed from 500 ms before until 1500 ms after instruction onset, sufficient to capture the early portion of the final hold epoch across all trials.

The HMMs trained on the execution trials then were used to decode hidden states during the observation trials for the same MN population. Similar to the execution trials, the state sequence consistency across observation trials was quantified as the fraction of all trials involving a given target object that had the same order of active hidden states from 500 ms before until 1500 ms after instruction onset. During 273 of the 1114 observation trials (24.5%), only one state, and hence no sequence of states, was detected for the duration of the trial. Such trials were excluded from analyses identifying the most common state sequence.

Data shuffling

To examine the factors in the original data that enabled HMMs to identify consistent state sequences, shuffled datasets were used to train additional HMMs. For all types of shuffling, the original spike times of a given unit within a given trial were conserved and these original spike time sequences were shuffled without replacement. First, we shuffled the data among the trials involving a given object ("trial-shuffled"). For example,

trial-shuffled "sphere trial 1" might consist of the recording from unit 1 during original sphere trial 7, the recording from unit 2 during original sphere trial 3, the recording from unit 3 during original sphere trial 9, etc. Trial shuffling thus preserved each unit's firing rate histogram within and across objects, but destroyed any precise temporal synchrony between units. Second, we shuffled each unit's spike times across trials involving various objects ("object-shuffled"), which necessarily meant also shuffling data from different trials. For example, object-shuffled "sphere trial 1" might consist of the recording from unit 1 during original button trial 5, the recording from unit 2 during original sphere trial 13, the recording from unit 3 during original coaxial cylinder trial 4, etc. Shuffling across objects preserved each neuron's firing rate histogram across all objects but not within each object. Third, we shuffled the ID of each neuron within each original trial ("neuron-shuffled"). This neuron shuffling preserved the precise timing and true simultaneity of the original recordings in each trial, but shuffled the recordings from different units among the IDs used by the HMM. For example, neuron-shuffled "sphere trial 1" consisted entirely of unit recordings from the original sphere trial 1, but the recording from unit 1 might be assigned to ID = 4, the recording from unit 2 to ID = 9, the recording from unit 3 to ID = 2, etc. Therefore, the recording for ID = 1 in shuffled sphere trials 1, 2, 3, etc., might be the recording from unit 8 in original sphere trial 1, from unit 3 in original sphere trial 2, from unit 7 in original sphere trial 3, etc. Shuffling the identity of the neurons preserved the precise timing simultaneity of the original recordings in each trial, but not the mean firing rate histograms for each neuron ID.

For each of the shuffled datasets, HMMs were trained in the same fashion as the original, unshuffled dataset (described above). Spike times from each trial were binned in 2 ms windows, 10 different HMMs were estimated at different initial conditions, and the HMM with the greatest log-likelihood was then selected as the model with which to continue analysis. This process was repeated 15 times to obtain several HMMs resulting from different iterations of shuffling. The resulting 15 HMMs then were analyzed for consistent state sequences (described above) to compare with the original unshuffled dataset.

Neural trajectory analyses

Patterns of covariation in a simultaneously recorded neural population can be analyzed as trajectories through a neural space in which each neuron's firing rate constitutes an orthogonal dimension (Churchland et al., 2012; Ames et al., 2014; Cunningham and Yu, 2014). The major components of these high-dimensional trajectories then can be examined in a low-dimensional subspace. We projected the neural trajectories of MN populations into low-dimensional subspaces using principal component analysis (PCA). To convert the binned inputs used for the HMMs to firing rate estimates, a 50 ms Gaussian filter was applied to the spike times of each MN in each trial. The firing rates of each MN were averaged across execution trials involving each object and PCA was performed on the population of MNs for each object separately. The joint neural trajectories during individual trials were mapped into the resulting subspace defined by the first two PCs and averaged across execution and observation trials separately.

Results

Neural activity was recorded and analyzed initially from execution trials (263 trials for Monkey L; 294 for Monkey X) and from observation trials (239 for Monkey L; 278 for Monkey X) of the RGM task from three sessions in each monkey. A summary of the number of execution trials, reaction times, and movement times for each monkey and for the human experimenter is given in Table 1. Reaction times for Monkey L, Monkey X, and the experimenter were significantly different from one another (pairwise Wilcoxon rank-sum tests with Bonferroni correction for three comparisons, $p < 0.05/3 = 0.0167$; $p_{L,X} = 5e-70$, $p_{L,E} = 3e-109$, $p_{X,E} = 7e-39$). Movement times for both monkeys differed significantly from those of the experimenter ($p_{L,E} = 4e-104$, $p_{X,E} = 2e-90$), but not from one another ($p_{L,X} = 0.055$). The experimenter took longer to react and to move than either monkey.

Using RM-ANOVA, we identified 272 single units and multiunits from Monkey L and 258 from Monkey X that had significant task-related modulation during execution trials. Table 2 summarizes the total numbers of significantly modulated units from each monkey, along with a breakdown for each session. Of the units with significant modulation during execution, 101 (37.1%) units from Monkey L also demonstrated task-related modulation during observation, as did 139 (53.9%) from Monkey X. Although our criteria may differ from those used by others, we refer to units that modulated significantly during both execution and observation as MNs. We also divided MNs into subpopulations recorded from PM or M1. Of the 101 MNs recorded in the three sessions from Monkey L, 75 (74.3%) were recorded in PM and 26 (25.7%) in M1; for Monkey X, 23 (16.6%) were recorded in PM and 116 (83.5%) in M1. Across all experimental sessions from both monkeys, 49 MNs were classified as definite single units, 51 as probable single units, and 140 as multiunits. Because our analyses are based on simultaneously recorded pop-

ulations, we made no attempt to determine which of these units were recorded in more than one session.

Although MNs often are thought to discharge almost identically during execution and observation of a particular movement and not discharge during other movements, previous studies have shown that many MNs discharge during both execution and observation of multiple different grasps (Rochat et al., 2010; Bonini et al., 2014b) and that the discharge of a given MN during execution versus observation of the same grasp may vary from “congruent” to “noncongruent” (Gallese et al., 1996; Vigneswaran et al., 2013). We likewise observed that individual MNs showed different degrees of similarity in the timing and amplitude of their activity during execution versus observation (Fig. 2). Some MNs had similar discharge patterns during both execution and observation for each target object. The example MN in Figure 2A began to discharge shortly after movement onset (circle in each raster line) and increased its firing rate until the start of the final hold (upward triangle). Although this MN discharged at similar times in the trials involving all four objects, its peak firing rate varied among objects, being highest for the button during both execution and observation. For other MNs, however, the modulation patterns during execution and observation were less similar. The MN shown in Figure 2B discharged a burst between the appearance of the instruction and movement onset, but thereafter, its firing rate decreased below its tonic baseline during execution trials, whereas during observation trials the burst was less intense and firing remained above baseline except for trials involving the perpendicular cylinder. Moreover, whereas this neuron achieved similar peak firing rates during execution trials involving the button and the coaxial cylinder in execution trials, in observation trials, its firing rates were substantially lower for the coaxial cylinder. Note that, although the discharge of these MNs had differing degrees of similarity during execution and observation trials, the overall pattern of discharge tended to be conserved in each MN. The former (Fig. 2A) had no baseline firing and discharged a burst during movement, whereas the latter (Fig. 2B) had a tonic firing rate that increased after the instruction appeared and then decreased.

We developed two metrics to quantify how selective each MN was for a particular object (a selectivity index, SEL) and how

Table 1. Reaction and movement times during execution and observation

	Monkey L (execution)	Monkey X (execution)	Experimenter (observation)
No. of sessions	3	3	6
No. of trials	263	294	517
Reaction times (ms) (median [25% 75%])	278 [253 300]	379 [336 427]	489 [405 591]
Movement times (ms) (median [25% 75%])	298 [244 344]	292 [250 390]	544 [488 634]

Table 2. Summary of the recorded number of neurons

	Monkey L				Monkey X			
	1	2	3	All	1	2	3	All
Session no.	1	2	3	All	1	2	3	All
Execution	90	87	95	272	71	83	104	258
MN	38	30	33	101	44	35	60	139
Non-MN	52	57	62	171	27	48	44	119
PM-MN	30	22	23	75	3	7	13	23
M1-MN	8	8	10	26	41	28	47	116
PM non-MN	22	21	20	63	5	11	7	23
M1 non-MN	30	36	42	108	22	37	37	96

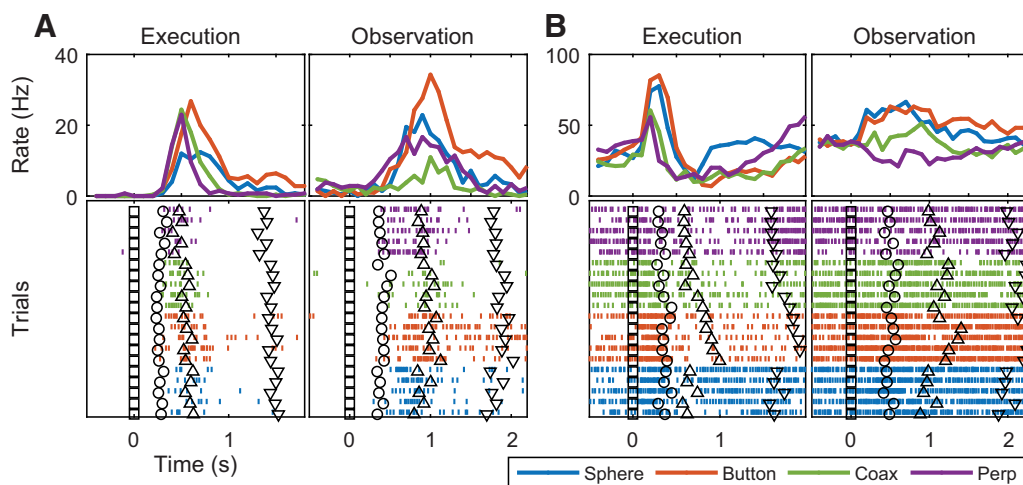


Figure 2. Example MNs. **A**, Definite single-unit MN recorded from PM with similar discharge during execution and observation. **B**, Multiunit MN recorded from M1, which modulated during both execution and observation but with less similarity than the MN shown in **A**. The perievent time histograms shown above have overlapping traces averaging ~20 trials for each of the four objects, whereas the spike rasters below depict only five trials per object. Behavioral events in each rastered trial are represented by black markers (instruction onset: square, movement onset: circle, start of the final hold period: upright triangle, and end of the final hold period: inverted triangle). All data have been aligned on instruction onset (squares).

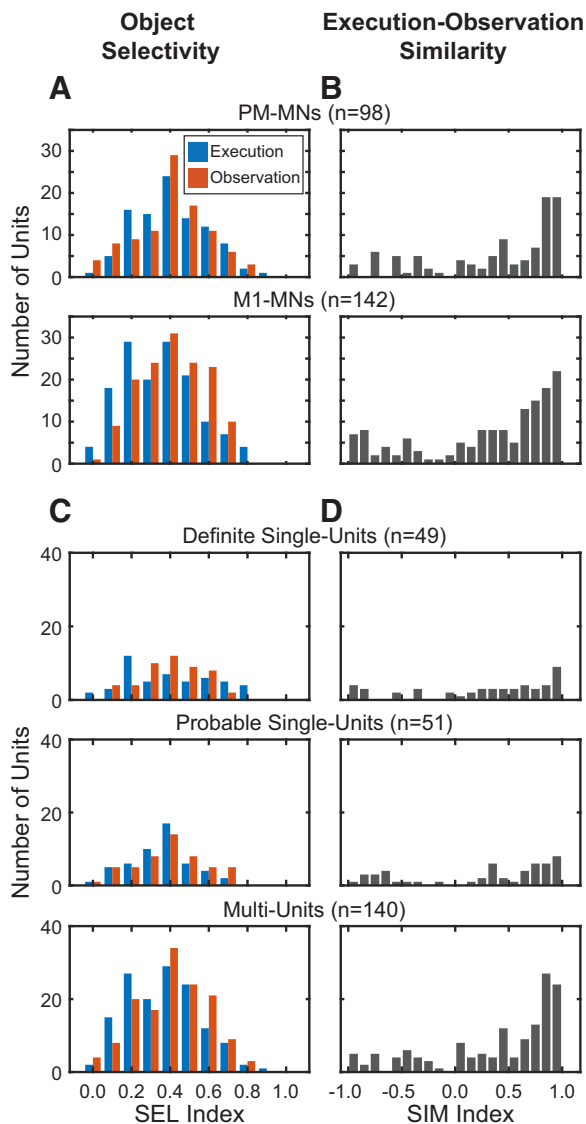


Figure 3. Distributions of selectivity and similarity indexes for individual MNs. The object selectivity index (SEL, left column) ranged from 0 to 1, where 0 indicates the unit was unselective, modulating equivalently in trials involving any of the four objects, and 1 indicates the unit fired selectively for only one object. For each MN, a SEL was computed separately for execution trials and observation trials, indicated by blue and red bars, respectively. The execution–observation similarity index (SIM, right column) compared the execution and observation activity of each MN across the four objects. The SIM ranged from -1 to 1 , where 0 indicates that activity during execution versus observation was orthogonal (dissimilar), 1 indicates that activity was ideally similar across all four objects during execution and observation, and -1 indicates that activity was ideally similar in magnitude during execution and observation but opposite in direction; that is, increasing versus decreasing firing rate. **A** and **B** show the SEL and SIM distributions, respectively, for PM-MNs above and M1-MNs below. **C** and **D** show distributions of the same data pooled across PM-MNs and M1-MNs and sorted instead by unit type: definite single units, top row; probable single units, middle row; and multiunits, bottom row.

similar each MN's discharge was during the movement epoch of execution versus observation trials (a similarity index, SIM). The SEL compared the absolute change in firing rate from the initial hold to the movement epoch averaged across trials involving each object. An SEL of 1 indicates that the MN discharged for only one object, whereas an SEL of 0 indicates that the MN discharged equivalently for all four objects.

For each MN, we calculated an SEL separately for execution and for observation trials. Figure 3A depicts the SEL distributions

for PM-MNs versus M1-MNs, each pooled across sessions and monkeys. For PM-MNs, SEL values for execution trials (0.46 [0.32–0.58], median [25th, 75th percentiles]) and for observation trials (0.47 [0.34–0.56]) were not significantly different ($p = 0.52$, Wilcoxon signed-rank test). For M1-MNs, however, SEL values for execution trials (0.40 [0.23–0.54]) were lower than for observation trials (0.45 [0.33–0.59]) ($p = 0.009$, Wilcoxon signed-rank test). The selectivity of PM-MNs and M1-MNs differed for execution trials ($p = 0.01$, Kruskal–Wallis test), but not for observation trials ($p = 0.92$, Kruskal–Wallis test). M1-MNs thus tended to be slightly less selective during execution trials.

We also compared the SEL distributions for definite single units, probable single units, and multiunits, pooling across PM-MNs and M1-MNs from all three sessions in both monkeys (Fig. 3C). For definite single units, the median SEL for execution trials was 0.42 [0.26, 0.64] and for observation trials was 0.44 [0.35, 0.56]. For probable single units, the median SEL for execution trials was 0.42 [0.32, 0.50] and for observation trials was 0.46 [0.32, 0.58]. For multiunits, the median SEL for execution trials was 0.44 [0.25, 0.55] and for observation trials was 0.47 [0.33, 0.59]. In none of these three MN classes did SEL differ between execution and observation (Wilcoxon signed-rank tests, $p = 0.95$ for definite single units, $p = 0.32$ for probable single units, and $p = 0.12$ for multiunits). In addition, SEL did not differ among the three classes for either execution or observation trials (Kruskal–Wallis tests, $p = 0.69$ for execution trials, $p = 0.95$ for observation trials).

The similarity index, SIM, compared each MN's firing rates for the four movements during execution versus observation (Fig. 3B). A SIM of 1 indicates that the ratios of the normalized firing rates during the four movements were identical during execution and observation. A SIM of 0 indicates the normalized firing rates during the four movements were completely different (orthogonal) during execution versus observation. A SIM of -1 indicates that the ratios of the normalized firing rates for the four movements were all of the same relative magnitude, but were opposite in direction during execution versus observation. The median SIM was 0.60 [0.02–0.88] for PM-MNs and 0.51 [-0.03 –0.82] for M1-MNs, consistent with the common observation that most MNs tend to modulate similarly during execution and observation. The SIM distributions of PM-MNs and M1-MNs were not significantly different ($p = 0.11$, two-tail two-sample Komogorov–Smirnov test).

We also compared the SIM distributions of the three unit classes, again combining PM-MNs and M1-MNs from all sessions in both monkeys (Fig. 3D). For definite single units, the median SIM was 0.43 [-0.15 –0.84]; for probable single units, 0.49 [-0.39 –0.86]; and for multiunits, 0.64 [0.05–0.86]. The distributions of the SIMs were not significantly different between any of the MN classes (two-tail two-sample Komogorov–Smirnov tests, definite vs probable, $p = 0.98$; definite versus multiunit, $p = 0.54$; probable versus multiunit, $p = 0.41$). Definite single units, probable single units, and multiunits thus all had similar distributions of both SEL and SIM. We therefore combined MNs in all three of these classes for subsequent analyses.

Hidden state sequences during action execution

Figure 4A illustrates the sequence of hidden states detected during three individual trials involving each of the four objects, all from the same session. Each frame shows a raster of spike times for the 23 simultaneously recorded PM-MNs. Colored lines and the shaded areas beneath show the time course of the probability

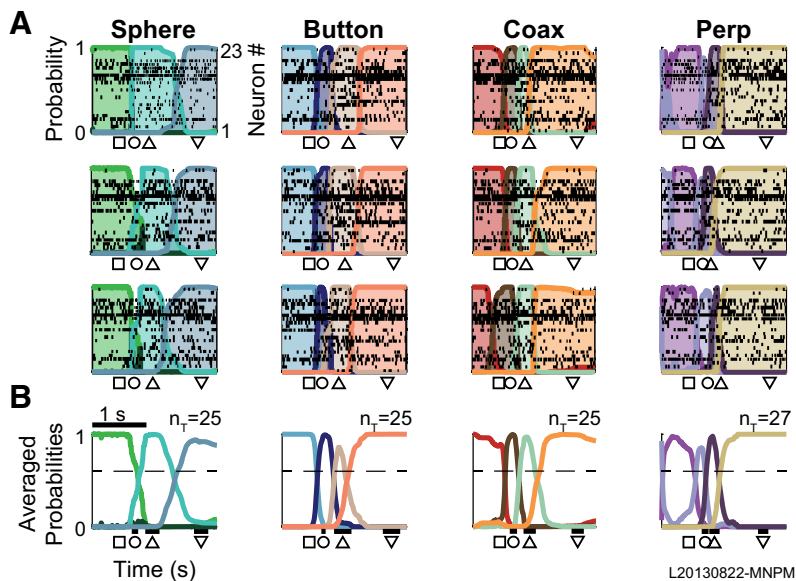


Figure 4. HMMs trained on execution trials for PM-MN population recorded from Monkey L in one experimental session. Each column represents trials involving a different target object: sphere, button, coaxial cylinder (Coax), and perpendicular cylinder (Perp). **A**, Three rows of individual trials illustrating the degree to which the state sequences were consistent across trials. Background colors indicate the ordinal sequence of hidden states. Different color schemes were used for each object. Spike rasters show the simultaneous activity of individual MNs in the sampled population. The times of behavioral events are represented by black markers beneath each plot: instruction onset (\square), movement onset (\circ), start of the final hold (\triangle), and completion of the final hold (∇). **B**, Averaged hidden state probabilities for each of the four objects. Dashed lines indicate the 0.6 probability level and the number of trials averaged for each object (n_T) is displayed above each subplot. All data have been aligned at the time of instruction onset (\square). Subsequent behavioral event markers have been plotted at their average time after instruction onset, with horizontal black bars above the markers indicating ± 2 SDs for the time of each marker across all analyzed trials. The 1.0 s time scale bar above the bottom left plot applies to all plots.

of each hidden state. Markers beneath the frame indicate the times of four behavioral events in that trial: (1) instruction onset (square), (2) movement onset (circle), (3) start of the final hold period (upright triangle), and (4) the end of the final hold period (inverted triangle). An initial state typically was present during the initial hold epoch at the beginning of the trial and a different, final state appeared after completion of the movement and continued beyond the end of the final hold epoch. Between these initial and final states, two sequential states most often were detected, with the probability of each state rising above 0.6, as illustrated by the trials involving the button, coaxial cylinder, and perpendicular cylinder in Figure 4. When both of these states were present, the former appeared during the reaction time epoch (between instruction onset and movement onset) and the other appeared during the movement epoch (between movement onset and start of the final hold). We will refer to these four sequential states based on the behavioral epoch in which each typically became active (i.e., achieved probability >0.6), namely as the initial, reaction, movement, and final states. In some trials, however, only one state became active between the initial and final states, as seen in the three sphere trials shown in Figure 4A. Overall for this 23 PM-MN population, during execution trials, 4 states were active sequentially in 87% (89/102) of the individual trials, 3 states in 12% (12/102), 2 states in 1% (1/102), and 1 state in 0% (0/102). Figure 4B shows the state probabilities aligned at instruction onset and averaged across all individual trials involving each object from this session. That the average probability of each state rose above 0.6 for trials involving each of the four objects despite the trial-to-trial variability in reaction and movement epoch durations, indicates that, for this population of PM-MNs, the HMMs detected each state with high consistency.

Figure 5 shows such averaged hidden state probability time courses for PM-MN and M1-MN populations for each session from each monkey, all for execution trials involving the coaxial cylinder. Four states were detected consistently in 9 of these 12 MN populations, with the average probability of each of the 4 states rising above 0.6. In contrast, only 3 states were detected in 3 of the 12 sessions: Monkey L's M1-MNs from session 3 and Monkey X's PM-MNs from sessions 2 and 3. Although the HMMs were trained to find four states, in these populations, the average probability of the third, "movement" state (green) never rose substantially above 0.0 and only the second state (brown) became active (probability >0.6) between the initial and final states.

For each of the 12 combinations of cortical area, monkey, and session, Figure 6 depicts the percentage of individual trials (including all four objects) in which 4, 3, 2, or 1 hidden states were detected as becoming active (probability >0.6). Overall, four states were detected most often, followed by three states. Indeed, four states were detected in $>50\%$ of individual trials with all MN populations except M1-MNs from the third session from Monkey L. Two states were detected in only a small percentage of trials and one state alone rarely if ever. White numerals in each column give the number of MNs in each population. Note that 1, 2, and even 3 states were detected primarily when fewer than 20 MNs were available. When >20 MNs were available, four hidden states were detected consistently in $>85\%$ of individual execution trials.

Timing of hidden state transitions during action execution

We examined the timing of hidden state transitions relative to the behavioral events, focusing only on trials in which all four hidden states (initial, reaction, movement, and final) became active. In each trial, we identified the times at which the probability of each hidden state rose above 0.6 and/or fell below 0.6 and then computed the time difference between these hidden state transitions and each of three behavioral events: instruction onset, movement onset, and start of the final hold. Combining the results from all 12 MN populations across trials involving all four objects, Figure 7, A and B, shows the distributions of these time differences.

These distributions reveal that each transition on average occurred closest to one of the three behavioral events (instruction onset, movement onset, or start of the final hold) and during one of the behavioral epochs (reaction epoch, movement epoch, or final hold epoch). The fall of the initial state and rise of the reaction state occurred, on average, 136 and 156 ms, respectively, after instruction onset during the reaction time epoch. The fall of the reaction state and the rise of the movement state occurred, on average, 63 and 78 ms, respectively, after movement onset during the movement epoch. The fall of the movement state and rise of the final state occurred, on average, 64 and 89 ms, respectively, after the start of the final hold during the final hold epoch. Therefore, the transitions from one hidden state to the next occurred most often after the nearest behavioral event. The SD of the times

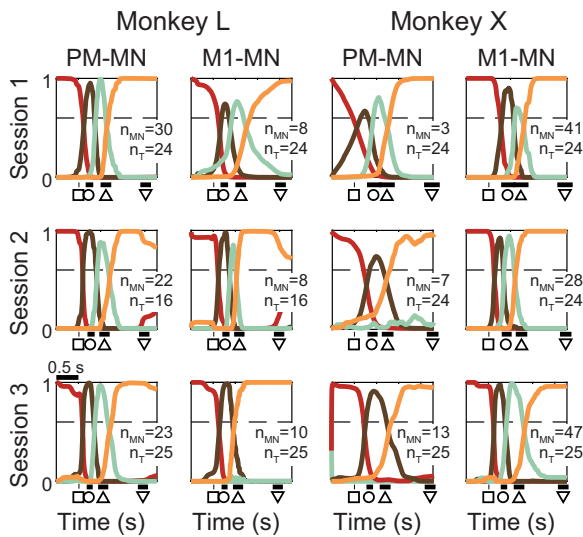


Figure 5. Trial-averaged hidden state probabilities for coaxial cylinder trials. Average state probabilities are displayed for the MN population recorded from each cortical area (PM or M1) in each monkey (L or X) in each session (1, 2, or 3). Conventions are similar to Figure 4, with dashed lines indicating the 0.6 probability level and the number of MNs (n_{MN}) in the population and the number of trials averaged (n_T) displayed in each subplot. All data have been aligned at the time of instruction onset (\square). Subsequent behavioral event markers have been plotted at their average time after instruction onset, with horizontal black bars above the markers indicating ± 2 SDs across all analyzed trials. The 0.5 s time scale bar above the bottom left plot applies to all plots. Conventions are same as for Figure 4.




Figure 6. Number of hidden states detected in execution trials. Bar segments represent the percentage of trials (combined across the four objects) in which HMMs trained on the MN population from each monkey, cortical area, and session detected different numbers of states: 4 = blue, 3 = red, 2 = yellow, and 1 = purple. With one exception, four states were detected in the majority of trials using each of the 12 MN populations. The number of MNs in each population is given in white on the bar.

between the closest behavioral event and transitions, however, ranged from 121 to 161 ms for falls and from 119 to 165 ms for rises, values almost as large if not larger than the means. Therefore, although occurring most often after the closest behavioral event marker, each state transition often occurred before that marker as well.

One might not have expected hidden state transitions in PM-MN and M1-MN populations to occur, on average, after rather than before behavioral events. For comparison, we therefore trained HMMs on populations of neurons recorded simultaneously with the present MNs, but which modulated

significantly only during execution trials and not during observation trials (non-MNs, the number in each population is given in Table 2). Figure 7, C and D, shows the distributions of hidden state transition times relative to each behavioral event marker in trials during which all four hidden states became active combined across the 12 non-MN populations and four objects. As for MN populations, state transitions in non-MN populations also occurred on average after the closest behavioral event, with the exception of the first change of state. The fall of the initial state and rise of the reaction state for non-MNs also occurred during the reaction time, on average, 183 and 195 ms, respectively, after instruction onset, but only -126 ms and -114 ms, respectively, before movement onset. Compared with the first change of state in MN populations, however, these transitions in non-MN populations occurred on average later: whether aligned on instruction onset or movement onset, both the fall of the initial state and the rise of the reaction state occurred later in the non-MN populations than in the MN populations ($p < 10^{-6}$, two-tailed t tests, Bonferroni corrected for 4 tests). The subsequent fall of the reaction state and the rise of the movement state in non-MN populations occurred, on average, 93 and 103 ms, respectively, after movement onset. Although closest to movement onset like the equivalent transitions in the MN populations, these transitions in the non-MN populations again occurred later than those in the MN populations ($p < 10^{-3}$, two-tailed t tests, Bonferroni corrected for 2 tests). The fall of the movement state and rise of the final state, then, occurred, on average, 84 and 106 ms, respectively, after the start of the final hold, again slightly later than the equivalent transitions in the MN populations, although in this case the difference was not significant. Although equivalent sequences of hidden states thus occurred with similar timing during execution trials in both the non-MN and MN populations, hidden state transitions in the MN populations occurred, on average, before those in the non-MN populations.

Factors affecting consistency of state sequences

Figure 8A shows a scatterplot of state sequence consistency versus the number of MNs for each of the 12 combinations of cortical area (PM, M1), monkey (L, X), and session (1, 2, 3). Consistency was $>75\%$ for all 6 sessions in which >20 MNs were available. For three of the six sessions for which fewer than 20 MNs were available, consistency was $>70\%$, but for the other three, consistency was $<55\%$. Even though high consistency more often was associated with larger numbers of MNs, some degree of consistency could be obtained with as few as seven MNs.

HMMs rely on the simultaneity of the recorded neural activity. We examined the extent to which truly simultaneous recordings from each MN population contributed to state sequence consistency by shuffling the recordings from each MN of a given population among different trials involving the same object, re-training the HMMs, and then reassessing consistency. Trial-shuffled consistency for each session is shown in the scatterplot of Figure 8B. Overall, this shuffling did not reduce the consistency across the 12 combinations of area, monkey, and session, averaging $81 \pm 17\%$ compared with the unshuffled data, which averaged $78 \pm 21\%$ ($p = 0.91$, Kruskal–Wallis, Tukey’s *post hoc* test). The consistency of hidden state sequences thus did not rely on precise temporal synchrony across a MN population; rather, the firing modulation pattern of each unit during trials involving a particular object was sufficient.

We next evaluated the extent to which object tuning of MNs contributed to HMM consistency. We shuffled MN activity across trials involving different objects in the same session, re-

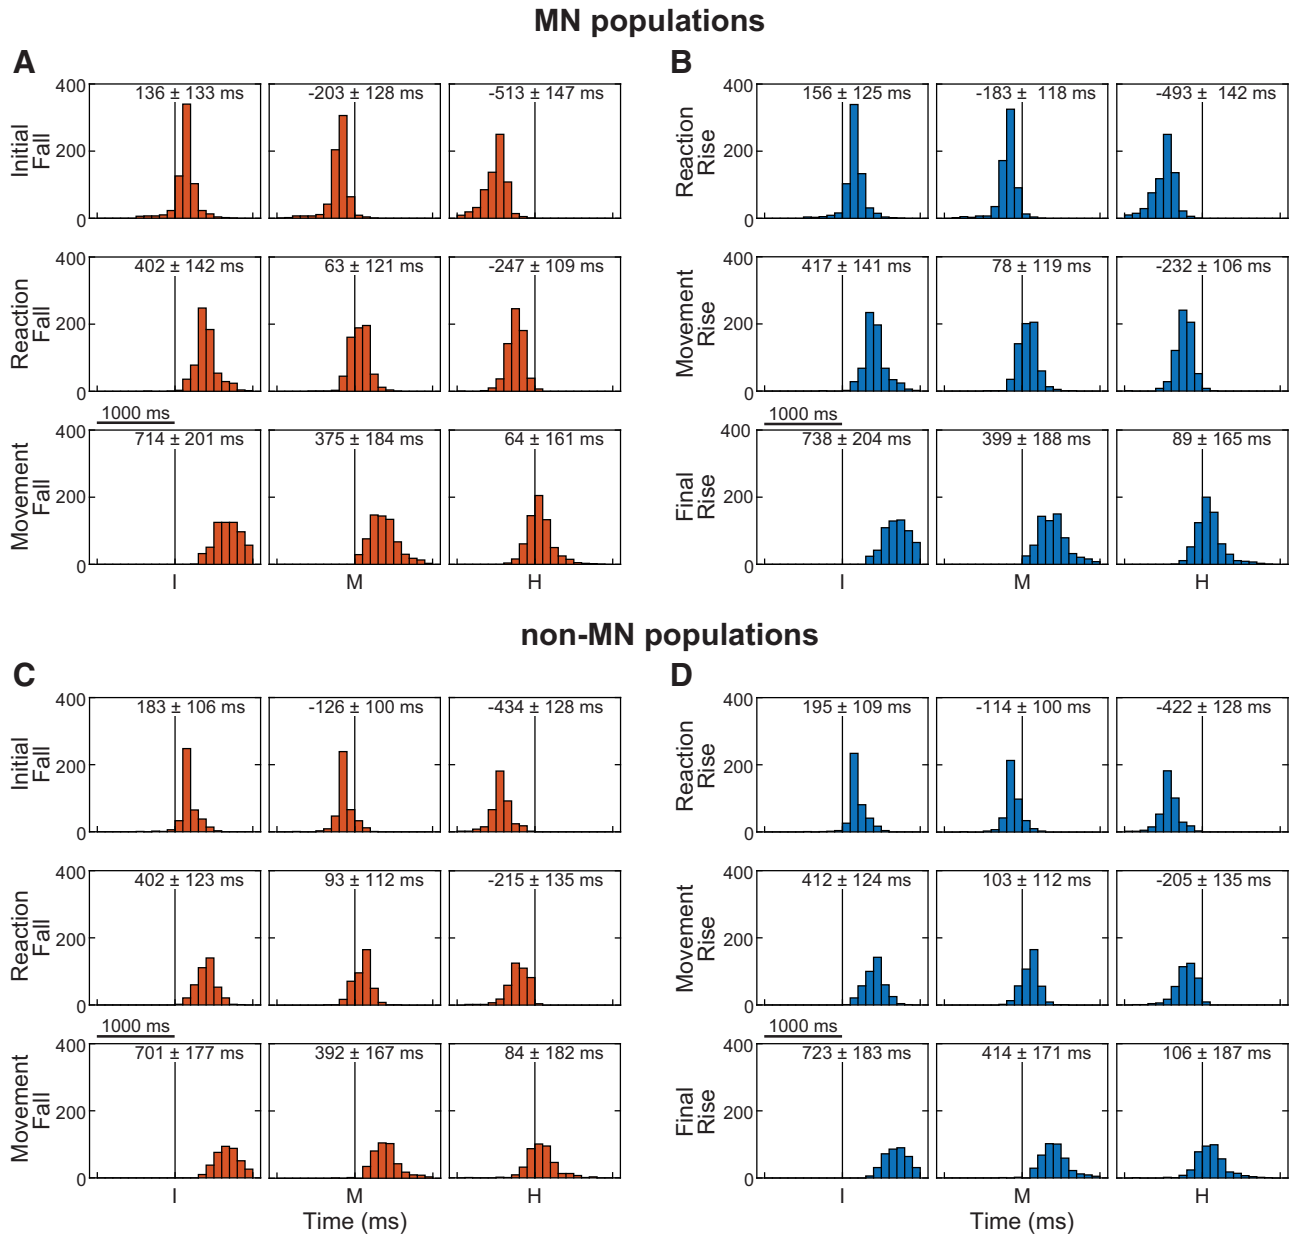


Figure 7. Hidden state transition times. Distributions of hidden state fall times (*A, C*), and rise times (*B, D*), relative to the onset of the instruction (I), the onset of movement (M), and the start of the final hold (H) are shown for the initial, reaction, movement, and final hidden states. Data were combined across objects, monkeys, and sessions separately for MN populations (*A, B*) and non-MN populations (*C, D*). Transition times were detected when the probability of a given state fell below (red) or rose above (blue) 0.6. Numbers at the top of each histogram give the mean and SD (in milliseconds) of the distribution relative to the behavioral event.

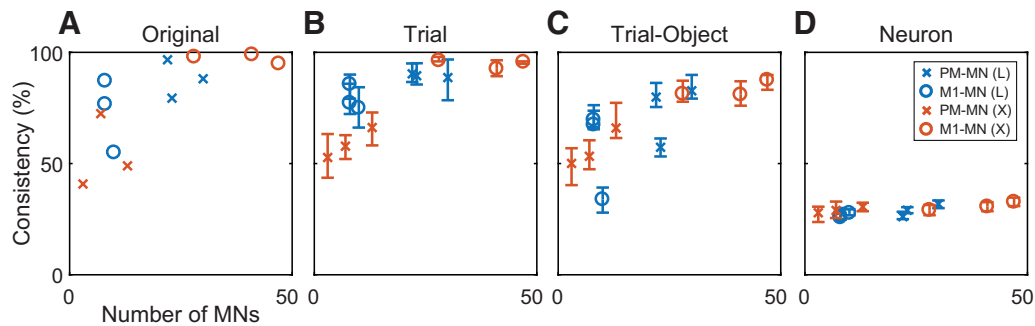


Figure 8. State sequence consistency of the 12 MN populations. *A*, Original dataset. *B*, Trial-shuffled datasets. *C*, Trial–object-shuffled datasets. *D*, Neuron-shuffled datasets. Blue symbols represent mean state consistency for Monkey L; red symbols for Monkey X. Error bars indicate SDs across 15 reshufflings. PM-MN populations are represented by “x” symbols; M1-MN populations by “o” symbols.

trained the HMMs, and reevaluated consistency (Fig. 8C). Shuffling across different objects produced only a slight decrease in HMM consistency averaging $68 \pm 18\%$, statistically different from the trial shuffled datasets ($p < 8e-7$, Kruskal–Wallis, Tukey’s *post hoc* test), but not the original dataset ($p = 0.08$, Kruskal–Wallis, Tukey’s *post hoc* test).

Last, we kept the same neurons in the same trials but shuffled the ID assignments of the MN recordings within each trial, again retrained the HMMs, and again reassessed consistency (Fig. 8D). In all 12 MN populations, shuffling neuron IDs reduced consistency to an average of $29 \pm 4\%$, a decrease of 49% relative to the original, unshuffled data and significantly different from each of the other shuffled datasets ($p < 4e-9$, Kruskal–Wallis, Tukey’s *post hoc* test)

Hidden state sequences during action observation

We also trained HMMs on trials in which each monkey observed the RGM task being performed by an experimenter. HMMs were trained in the same manner as the execution trials presented above. Figure 9 shows examples from a single session in which four hidden states were detected consistently in trials involving the sphere, button, or coaxial cylinder, but only three states in trials involving the perpendicular cylinder. Overall, the number of hidden states detected was both lower and more variable during observation than during execution trials. Across all 12 combinations of monkey, area, and session, 4, 3, 2, and 1 states were detected in 80.9%, 18.0%, 1.0%, and 0.1% of execution trials, but only in 44.1%, 32.6%, 17.7%, and 5.6% of observation trials, respectively.

Consistency also was lower for HMMs trained on observation trials. For execution trials across all sessions and objects, PM-MN populations from Monkey L had consistent state sequences in 86.7% [85.2%, 88.2%] (mean [25th and 75th percentiles]) of individual trials, and M1-MN populations, in 72.2% [70.3%, 74.1%]. PM-MN populations from Monkey X had consistent state sequences in 54.1% [52.0%, 56.1%] of individual trials, and M1-MN populations in 97.3% [96.6%, 98.0%]. However, in observation trials, PM-MN populations from Monkey L had consistent state sequences in 62.8% [60.7%, 64.9%], and M1-MN populations in 26.4% [24.3%, 28.5%]. PM-MN populations from Monkey X had consistent state sequences in 34.5% [32.7%, 36.3%] of individual trials, and M1-MN populations in 45.0% [42.8%, 47.1%].

Similar sequences of hidden states detected during execution and observation

Although MNs are commonly thought to represent movements per se, hidden state sequences in MN populations from both PM and M1 appeared to reflect the entire sequence of behavioral epochs in execution trials, including the initial hold, reaction time, and final hold epochs during which no motion occurred. Did the same sequence of hidden states occur in MN populations during observation trials? Using HMMs trained on execution trials, we decoded the observation trials recorded during the same session. Figure 10 illustrates HMMs trained on sphere execution

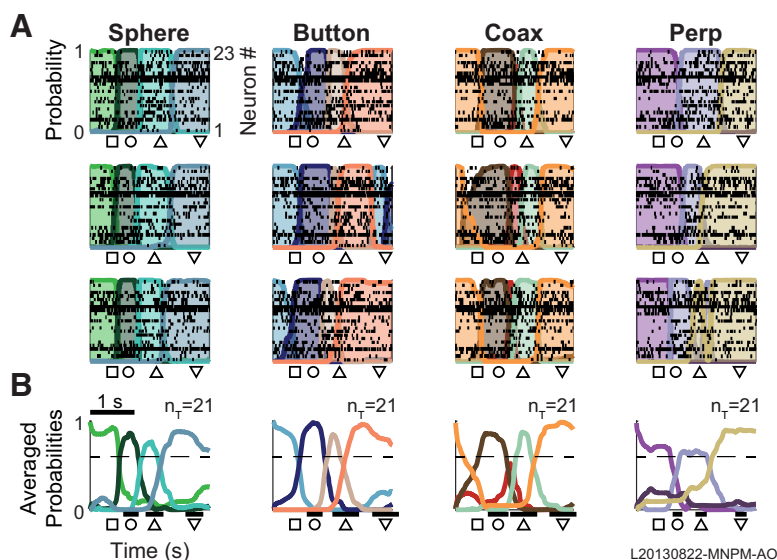


Figure 9. HMMs trained on observation trials for the PM-MN population recorded from Monkey L in a single experimental session. **A**, Similar state sequence transitions are detected for each object on a trial by trial basis. **B**, Hidden state probabilities averaged across trials for each object. Conventions the same as for Figure 4.

trials and then used to decode sphere observation trials separately for MN populations from PM (Fig. 10A–D) and from M1 (Fig. 10E–H), all recorded simultaneously in the same session from Monkey L. In the three single-trial examples, the PM-MN population consistently detected the four sequential states—initial, reaction, movement, and final—in the execution trials used for training (Fig. 10A) and decoded similar sequences in observation trials that had not been used to train the model (Fig. 10C). Because of trial-to-trial variability in the timing of state transitions, trial-averaged state probabilities for the PM-MN population did not necessarily reach 1.0 (Fig. 10B,D). Nevertheless, the hidden state consistency of this PM-MN population for execution trials was 100%; that is, the initial, reaction, movement, and final state each attained probability >0.6 in that order in every trial. During observation trials, the hidden state consistency of this PM-MN population decreased to 73.7%, which was not unexpected given that the monkey may not have paid close attention on every observation trial (Krakauer et al., 2017).

Substantially different results were obtained, however, on the same trials using the simultaneously recorded M1-MN population (Fig. 10E–H). In the single-trial examples, four states were detected during sphere execution trials (Fig. 10E), but when the HMM trained with M1-MNs on execution trials was used to decode observation trials, only the initial state was detected as active and remained so throughout each of the illustrated single trials (Fig. 10G). Similar results for this M1-MN population are evident in the trial-averaged probabilities (Fig. 10F,H). For this M1-MN population, state sequence consistency across all sphere trials was 100% for execution trials but fell to 5.3% for observation trials. For the sphere trials in this session, similar hidden state sequences were present during execution trials and most observation trials in the PM-MN population, but not in the M1-MN population.

For each MN population in each session, we trained an HMM separately for each object on execution trials and then used that HMM to decode hidden states during observation trials. We then calculated the hidden state consistency for the execution trials and for the observation trials separately. Figure 11 shows these hidden state consistencies for the PM-MN and M1-MN popula-

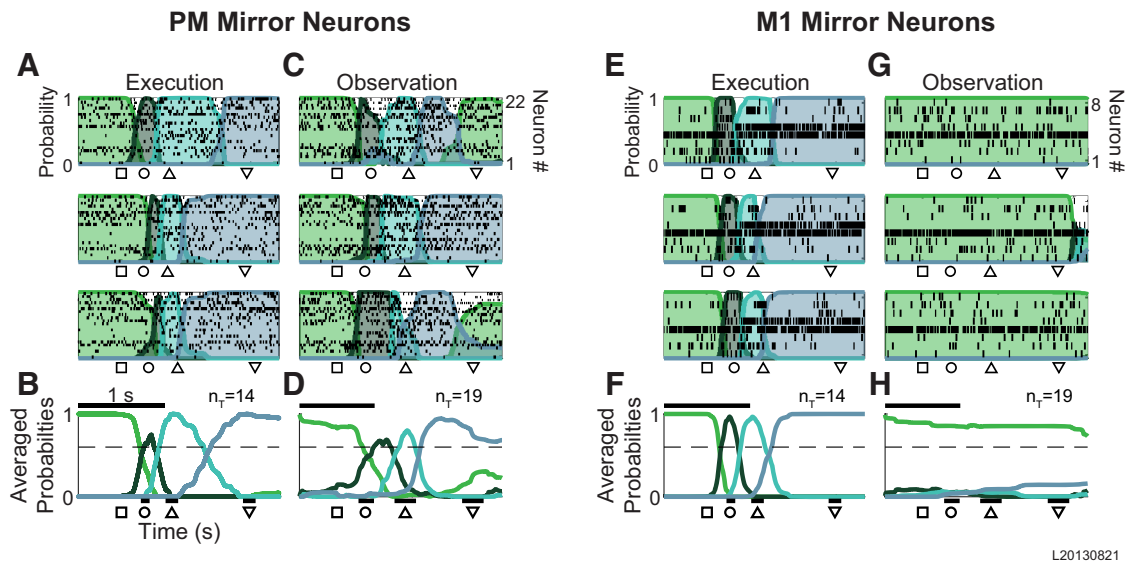


Figure 10. Sphere trials decoded with HMMs trained on execution trials. State probabilities are shown for three individual trials above (*A, C, E, G*) and averaged across all trials below (*B, D, F, H*) for PM-MNs (*A–D*) and M1-MNs (*E–H*) in Monkey L. In the execution trials used to train the HMMs, four states were consistently detected in sequential order—initial, reaction, movement, final—whether the MN population came from PM (*A, B*) or M1 (*E, F*). When those HMMs trained on execution trials were used to decode observation trials, however, the sequence of four hidden states was found in the PM-MN population (*C, D*), but not in the M1-MN population (*G, H*). Conventions are the same as in Figure 4, with black markers representing the different behavioral events and the time scale bars representing 1.0 s for both execution and observation trials.

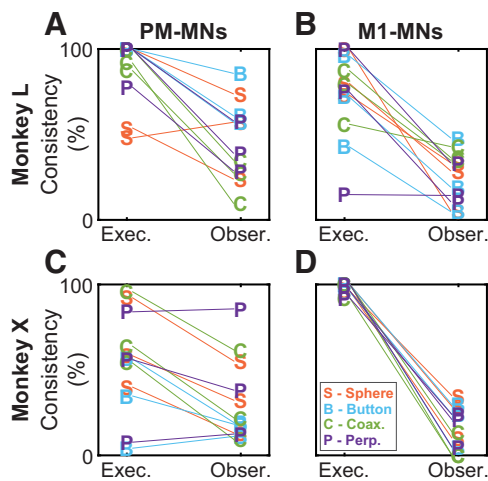


Figure 11. State consistency for models trained on execution trials and then used to decode observation trials. Each panel compares model consistency across trained execution trials (Exec.) with that across decoded observation trials (Obser.) involving each of the four objects in the three sessions for different MN populations: *A*, PM-MN populations from Monkey L; *B*, M1-MN populations from Monkey L; *C*, PM-MN populations from Monkey X; and *D*, M1-MN populations from Monkey X. Different objects represented by different symbols and colors given in the legend in *D* (sphere: orange “S”; button: blue “B”; coaxial cylinder: green “C”; and perpendicular cylinder: purple “P”). Consistency generally was greater for execution trials than for observation trials, but the ratio of observation to execution consistency was greater on average for PM-MNs than M1-MNs.

tions from each of the two monkeys. Not unexpectedly, for each type of MN population from each monkey, comparing across the 12 samples (4 objects × 3 sessions) showed lower consistency for the decoded observation trials than for the execution trials used to train the HMMs (Wilcoxon signed-rank tests, $p < 0.01$ after Bonferroni correction for 4 tests). To compare the degree of the drop in consistency from execution to observation, we calculated the ratio of the consistency during observation versus execution trials (observation consistency/execution consistency) for each of the 12 samples (4 objects × 3 sessions). In Monkey L, the mean

ratio across the 12 samples was 0.54 for PM-MN populations and 0.40 for M1-MN populations; in Monkey X, the mean ratios were 0.79 and 0.17, respectively. In each monkey, the ratios for M1-MN populations thus were lower than those for PM-MN populations, although this trend reached significance only in Monkey X (Wilcoxon rank-sum tests, $p = 0.2$ for Monkey L, $p = 0.0003$ for Monkey X). The hidden states found during execution thus tended to persist during observation trials more in PM-MN than in M1-MN populations.

Hidden states correspond to specific subspaces in population activity during execution and observation

Sequences of four hidden states thus became active consistently across the majority of execution trials in both PM-MN and M1-MN populations. The same sequences could be detected in many observation trials, particularly for PM-MN populations. These findings suggest temporal patterns of covariation in the activity of MN populations during execution trials that remain similar during observation trials. To examine such covariation in greater detail, we considered the firing rates of the N MNs in a given population as an N -dimensional space and performed PCA on the N firing rates for execution trials involving each object separately. Across all 12 combinations of monkey, session, and cortical area, the first two PCs accounted for 71.2–95.1% of the variance in MN firing rates during execution trials.

For each object, we then projected the N -dimensional firing rates for all execution trials onto the plane of the first two PCs for that object and averaged across trials. Figure 12 illustrates such averaged neural trajectories for the PM-MN and M1-MN populations (Fig. 12*A, B*, respectively) recorded simultaneously in session 2 from Monkey L. In both the PM-MN and M1-MN populations, these 2D trajectories progressed from the start of the trial, through instruction onset (□), movement onset (○), start of the final hold (Δ), and completion of the final hold (▽) during execution trials involving each of the four objects.

We then projected the firing rates of the same MNs averaged across observation trials onto the same PC1 versus PC2 plane

defined by execution trials for each object. For the PM-MNs, these 2D neural trajectories during observation trials were similar to, albeit somewhat smaller than, those found during execution trials. For the M1-MNs, however, the 2D trajectories during observation trials collapsed, indicating that comodulation of these M1-MNs during observation trials had little to no projection in the plane defined by the activity of the same MNs during execution trials.

The area subtended by each unclosed average trajectory was estimated (using the “polyarea” function in MATLAB) from the average trajectory beginning at instruction onset and ending at the completion of the final hold. We then compared these areas for each type of MN population from each monkey across the 12 samples (4 objects × 3 sessions). For both the PM-MN and M1-MN populations in both monkeys, area was smaller for observation trials than for execution trials (Wilcoxon signed-rank tests, $p < 0.001$ after Bonferroni correction for 4 tests). To compare the degree of the decrease in area from execution to observation, we calculated the ratio of the area during observation versus execution trials (observation area/execution area) for each of the 12 samples. In Monkey L, the mean ratio across the 12 samples was 0.33 for PM-MN populations and 0.08 for M1-MN populations; in Monkey X, the mean ratios were 0.22 and 0.05, respectively. In each monkey the observation to execution area ratios for M1-MN populations thus were lower than those for PM-MN populations (Wilcoxon signed-rank tests, $p = 0.005$ for Monkey L, $p = 0.012$ for Monkey X). The temporal patterns of covariation in MN firing rates found during execution thus tended to persist during observation trials more in PM-MN than in M1-MN populations.

Discussion

Populations of MNs were found to encode sequences of hidden neural states as detected by HMMs. These hidden states corresponded approximately to behavioral epochs of the RGM task. By decoding neural activity during observation trials with models trained on execution trials, we found that populations of PM-MNs tended to show similar hidden state sequences and neural trajectories during execution and observation, whereas M1-MNs did not.

MN populations encode sequences of hidden states related to multiple behavioral epochs

Rather than only detecting a baseline state and a movement state, HMMs most often detected four hidden states in the activity of MNs during individual execution or observation trials. These four hidden states corresponded approximately to behavioral epochs in each trial. MN populations thus represented not only the movement per se, but also multiple non-movement behavioral

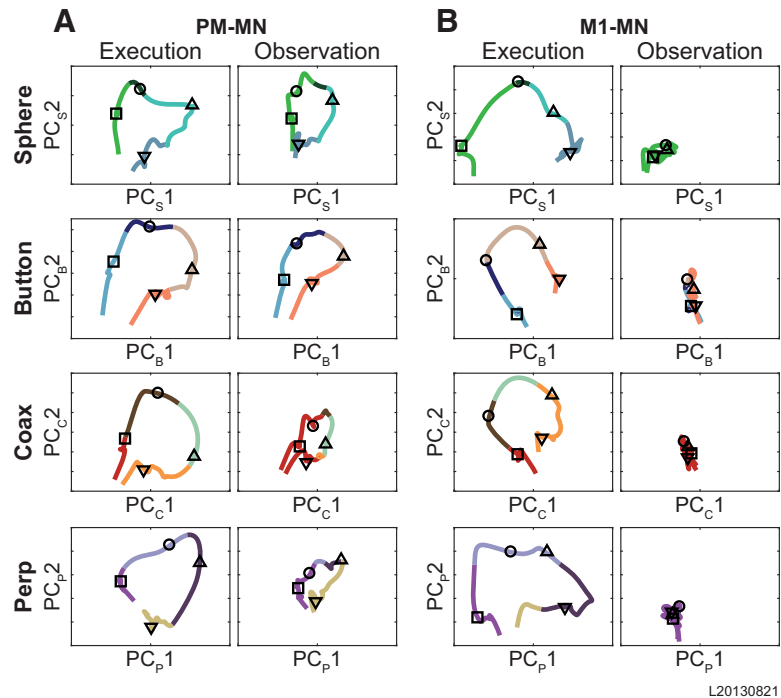


Figure 12. Neural population trajectories. PCA projections of the mean neural activity during execution and observation trials are shown for PM-MNs (**A**) and M1-MNs (**B**) in session 2 from Monkey L. The neural population trajectory averaged across multiple trials involving each object has been plotted in a 2D PC space computed based on PCA of the execution trial neural activity (PC_2 : sphere; PC_B : button; PC_C : coaxial cylinder; and PC_P : perpendicular cylinder). The most probable trial-averaged hidden state at each time point along the trajectory has been represented with a different color, maintaining the color scheme used in Figure 4. Observation trials were projected into the same PC space derived from execution trial data for each respective object and MN population. During execution trials, the neural trajectories of both MN populations progress through the average time of instruction onset (\square), movement onset (\circ), start of the final hold (\triangle), and completion of the final hold (∇). During observation trials, however, PM-MN trajectories (**A**) progressed similarly, with the hidden states occurring in similar PC-space locations, whereas the M1-MN trajectories (**B**) during observation did not resemble their execution counterparts.

epochs—initial hold, reaction time, and final hold—of the RGM trials.

The fall of one hidden state and the rise of the next typically occurred during one of the behavioral epochs, most often after the closest behavioral event marker (instruction onset, movement onset, or the start of the final hold). We therefore chose to label each hidden state based on the behavioral epoch (initial hold, reaction time, movement time, final hold) during which that state typically began. We found that, during execution trials, non-MN populations progressed through similar hidden state transitions with similar timing; therefore, the sequence of states and the timing of their transitions were not unique to MNs. Non-MNs, of course, were not modulated significantly during observation trials.

The initial state generally was present during the initial hold epoch and was replaced by the reaction state during the reaction time epoch, consistent with the appearance of a visual instruction cue resulting in neural activity that initiated movement in execution trials (Kaufman et al., 2016). In observation trials, however, the rise of the reaction state in the reaction time epoch could not have resulted from observation of the experimenter’s movement, which had not yet started. Rather, the rise of the reaction state during observation trials most likely resulted from the monkey seeing the instruction with the expectation that the experimenter was about to move, even though the monkey itself would not be generating an RGM movement in response to that visual cue.

The fall of the reaction state and rise of the movement state then occurred most often (but not always) after movement onset. We speculate that two factors may have contributed to this change of state. First, reafferent feedback, both visual and/or proprioceptive, may have contributed by arriving in PM and M1 after movement onset. PM and M1 both receive substantial somatosensory input (Lemon and Porter, 1976; Fetz et al., 1980; Lemon, 1981; Rizzolatti et al., 1981b) and visual information arriving in PM (Rizzolatti et al., 1981a; Fogassi et al., 1999; Graziano, 1999) could influence M1 via their interareal connections (Muakkassa and Strick, 1979; Shimazu et al., 2004). Whereas proprioceptive feedback would not be present during observation trials, visual input from the experimenter's movement would. Second, during RGM movements, the activity of most M1 neurons (not distinguishing MNs vs non-MNs) shows two sequential phases, the first phase encoding the location to which the reach is being made and the second encoding the object being grasped (Rouse and Schieber, 2016). The timing of this transition between location and object encoding occurs near movement onset but varies among individual M1 neurons—before movement onset in some and after movement onset in others—and therefore may also have contributed to the change from reaction to movement state.

The subsequent fall of the movement state and rise of the final state occurred, on average, after the start of the final hold. This change of state could reflect the transition in control from movement to posture (Venkadesan and Valero-Cuevas, 2008), the loss of dynamic feedback from the moving limb (Ghez and Sainburg, 1995; Sainburg et al., 1995), and/or the somatosensory input from contact with the object (Johansson and Flanagan, 2009).

In most reach-to-grasp tasks, the difference between the initial and final states might be attributable to the hand grasping an object during the final hold but not during the initial hold (Caggiano et al., 2009; Bonini et al., 2010). In the present RGM task, however, the hand grasped and manipulated an object during both the initial hold and the final hold. That the posture of the arm and hand were different during the final hold compared with the initial hold, however, may have contributed to the HMMs detecting distinct initial and final states (Caminiti et al., 1990; Sergio and Kalaska, 2003).

Previous studies also support the notion that MN populations not only respond to observed motion of another individual's arm and hand, but also monitor the preceding behavioral context. Some MNs discharge when the other individual withholds an expected movement in the NoGo trials of a Go/NoGo task (Bonini et al., 2014a). Some MNs show anticipatory discharge preceding the onset of observed movements (Cisek and Kalaska, 2004; Maranesi et al., 2014). Neither the NoGo discharge nor the anticipatory activity of such MNs can be attributed to observation of a movement and instead indicate that the context of expecting to observe a movement engages at least some MNs. We found that hidden state transitions in MN populations generally occurred slightly before, rather than after, the equivalent transitions in non-MN populations, consistent with the notion that even during execution MN populations incorporate some expectation of what is about to happen rather than simply responding to what has happened.

Although each RGM trial had the same behavioral structure—initial hold, reaction time, movement, and final hold—regardless of which object was instructed, we chose to train HMMs on each object separately. As illustrated in Figure 2, MNs often discharged differentially depending on the particular grasp being performed and this variance during the movement epoch might have degraded the performance of the HMMs had they been trained across objects. We therefore used random shuffling of the data

from execution trials to examine the extent to which the different objects affected hidden state consistency (Fig. 8). When we shuffled recordings of MN firing among trials from the same object (Fig. 8B), the consistency of the hidden states was similar to the original, unshuffled dataset (Fig. 8A). Although this shuffling to a certain extent disrupted the relative timing of both the behavioral events and the spikes discharged by MNs in different trials and therefore may have blurred state transition times, the retained consistency indicates that MNs generally progressed through a similar sequence of hidden states in different trials involving the same object. When we shuffled the assignments across objects, hidden state sequence consistency decreased somewhat (Fig. 8C), confirming that the four different grasps did affect HMM performance to some extent. Nevertheless, state sequence consistency by and large was maintained, indicating that the MN populations captured the sequence of the four behavioral epochs in the RGM task without relying entirely on object-specific information. This hidden state representation of the general structure of the RGM task was most substantially decreased, however, by shuffling neuron IDs within a trial (Fig. 8D). The hidden states of MN populations thus encoded sequential behavioral epochs based on differences in the general pattern of modulation of different MNs across objects. For example, the MN illustrated in Figure 2A had a very low baseline rate and discharged a burst during the movement for all four objects, whereas the MN illustrated in Figure 2B had a substantial baseline rate, fired a burst after instruction onset, and was suppressed after movement onset for all four objects. We infer that HMMs detected hidden states largely based on such overall patterns of modulation in the different MNs of a population distinct from each MN's individual variance in its own firing related to particular grasps/objects.

PM versus M1 MN populations

During observation, MNs are generally thought to monitor the movements of other individuals. Our findings reveal that MN populations, particularly those in PM, monitor not only the observed movement, but also the preceding and following behavioral epochs during which no movement occurs. What about MN activity during execution? MNs constitute a substantial fraction of the task-related population in both PM and M1 and therefore commonly have been assumed to participate in driving movements along with non-MNs. Nevertheless, our findings, like those of most other studies of MNs, do not address the question of whether MNs causally contribute to driving movements during execution or simply monitor the movement as it is executed. Addressing this issue directly awaits future investigations in which MNs versus non-MNs can be stimulated and/or inactivated selectively.

Two studies, however, have suggested that MNs with corticospinal axons (pyramidal tract neurons, PTNs) contribute to driving movements during execution while serving to prevent the subject's own movement during observation (Kraskov et al., 2009; Vigneswaran et al., 2013). The firing rate of PTN MNs in both PMv and M1 often was found to increase during execution, but to increase less or even to decrease during observation. Insufficient PTN excitation of the spinal cord could prevent the subject's own movement during observation.

Although the present execution/observation similarity index (SIM) distributions were not significantly different between populations of PM-MNs and M1-MNs (in which we made no attempt to identify PTNs), the simultaneous MN activity in these two areas did show differences in hidden state sequences. HMMs trained on execution trials decoded full sequences of hidden states in observation trials more frequently for PM-MN than for

M1-MN populations. Furthermore, PM-MN population trajectories progressed similarly during execution and observation, whereas the trajectories of M1-MN populations collapsed during observation. M1-MN trajectories traveling through one neural subspace during execution, but an orthogonal subspace during observation may contribute to the absence of movement during observation (Kaufman et al., 2014). If such execution/observation differential activity of M1-MN populations, together with insufficient PTN activation of the spinal cord, prevents the subject's own movement during observation, then PM-MN populations would be free to continue discharging as if the subject were performing the movement, enabling the brain to monitor the entire sequence of behavioral epochs being observed.

References

- Abeles M, Bergman H, Gat I, Meilijson I, Seidemann E, Tishby N, Vaadia E (1995) Cortical activity flips among quasi-stationary states. *Proc Natl Acad Sci U S A* 92:8616–8620. [CrossRef Medline](#)
- Ames KC, Ryu SI, Shenoy KV (2014) Neural dynamics of reaching following incorrect or absent motor preparation. *Neuron* 81:438–451. [CrossRef Medline](#)
- Baum LE, Petrie T, Soules G, Weiss N (1970) A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics* 41:164–171.
- Bonini L, Rozzi S, Serventi FU, Simone L, Ferrari PF, Fogassi L (2010) Ventral premotor and inferior parietal cortices make distinct contribution to action organization and intention understanding. *Cereb Cortex* 20:1372–1385. [CrossRef Medline](#)
- Bonini L, Maranesi M, Livi A, Fogassi L, Rizzolatti G (2014a) Ventral premotor neurons encoding representations of action during self and others' inaction. *Curr Biol* 24:1611–1614. [CrossRef Medline](#)
- Bonini L, Maranesi M, Livi A, Fogassi L, Rizzolatti G (2014b) Space-dependent representation of objects and other's action in monkey ventral premotor grasping neurons. *J Neurosci* 34:4108–4119. [CrossRef Medline](#)
- Caggiano V, Fogassi L, Rizzolatti G, Thier P, Casile A (2009) Mirror neurons differentially encode the peripersonal and extrapersonal space of monkeys. *Science* 324:403–406. [CrossRef Medline](#)
- Caggiano V, Fogassi L, Rizzolatti G, Pomper JK, Thier P, Giese MA, Casile A (2011) View-based encoding of actions in mirror neurons of area f5 in macaque premotor cortex. *Curr Biol* 21:144–148. [CrossRef Medline](#)
- Caminiti R, Johnson PB, Urbano A (1990) Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. *J Neurosci* 10:2039–2058. [CrossRef Medline](#)
- Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV (2012) Neural population dynamics during reaching. *Nature* 487:51–56. [CrossRef Medline](#)
- Cisek P, Kalaska JF (2004) Neural correlates of mental rehearsal in dorsal premotor cortex. *Nature* 431:993–996. [CrossRef Medline](#)
- Cunningham JP, Yu BM (2014) Dimensionality reduction for large-scale neural recordings. *Nat Neurosci* 17:1500–1509. [CrossRef Medline](#)
- di Pellegrino G, Fadiga L, Fogassi L, Gallese V, Rizzolatti G (1992) Understanding motor events: a neurophysiological study. *Exp Brain Res* 91:176–180. [CrossRef Medline](#)
- Dushanova J, Donoghue J (2010) Neurons in primary motor cortex engaged during action observation. *Eur J Neurosci* 31:386–398. [CrossRef Medline](#)
- Fetz EE, Finocchio DV, Baker MA, Soso MJ (1980) Sensory and motor responses of precentral cortex cells during comparable passive and active joint movements. *J Neurophysiol* 43:1070–1089. [CrossRef Medline](#)
- Fogassi L, Raos V, Franchi G, Gallese V, Luppino G, Matelli M (1999) Visual responses in the dorsal premotor area F2 of the macaque monkey. *Exp Brain Res* 128:194–199. [CrossRef Medline](#)
- Fogassi L, Ferrari PF, Gesierich B, Rozzi S, Chersi F, Rizzolatti G (2005) Parietal lobe: from action organization to intention understanding. *Science* 308:662–667. [CrossRef Medline](#)
- Gallese V, Fadiga L, Fogassi L, Rizzolatti G (1996) Action recognition in the premotor cortex. *Brain* 119:593–609. [CrossRef Medline](#)
- Gallese V, Keysers C, Rizzolatti G (2004) A unifying view of the basis of social cognition. *Trends Cogn Sci* 8:396–403. [CrossRef Medline](#)
- Ghez C, Sainburg R (1995) Proprioceptive control of interjoint coordination. *Can J Physiol Pharmacol* 73:273–284. [CrossRef Medline](#)
- Graziano MS (1999) Where is my arm? the relative role of vision and proprioception in the neuronal representation of limb position. *Proc Natl Acad Sci U S A* 96:10418–10421. [CrossRef Medline](#)
- Hill DN, Mehta SB, Kleinfeld D (2011) Quality metrics to accompany spike sorting of extracellular signals. *J Neurosci* 31:8699–8705. [CrossRef Medline](#)
- Johansson RS, Flanagan JR (2009) Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nat Rev Neurosci* 10:345–359. [CrossRef Medline](#)
- Jones LM, Fontanini A, Sadacca BF, Miller P, Katz DB (2007) Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci U S A* 104:18772–18777. [CrossRef Medline](#)
- Kang X, Sarma SV, Santaniello S, Schieber M, Thakor N (2015) Task-Independent cognitive state transition detection from cortical neurons during 3D reach-to-grasp movements. *IEEE Trans Neural Syst Rehabil Eng*, 23, 676–682. [CrossRef Medline](#)
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014) Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci* 17:440–448. [CrossRef Medline](#)
- Kaufman MT, Seely JS, Sussillo D, Ryu SI, Shenoy KV, Churchland MM (2016) The largest response component in the motor cortex reflects movement timing but not movement type. *eNeuro* 3.
- Kemere C, Santhanam G, Yu BM, Afshar A, Ryu SI, Meng TH, Shenoy KV (2008) Detecting neural-state transitions using hidden Markov models for motor cortical prostheses. *J Neurophysiol* 100:2441–2452. [CrossRef Medline](#)
- Kilner JM, Neal A, Weiskopf N, Friston KJ, Frith CD (2009) Evidence of mirror neurons in human inferior frontal gyrus. *J Neurosci* 29:10153–10159. [CrossRef Medline](#)
- Krakauer JW, Ghazanfar AA, Gomez-Marín A, MacIver MA, Poeppel D (2017) Neuroscience needs behavior: correcting a reductionist bias. *Neuron* 93:480–490. [CrossRef Medline](#)
- Kraskov A, Dancause N, Quallo MM, Shepherd S, Lemon RN (2009) Corticospinal neurons in macaque ventral premotor cortex with mirror properties: a potential mechanism for action suppression? *Neuron* 64:922–930. [CrossRef Medline](#)
- Lemon RN (1981) Functional properties of monkey motor cortex neurones receiving afferent input from the hand and fingers. *J Physiol* 311:497–519. [CrossRef Medline](#)
- Lemon RN, Porter R (1976) Afferent input to movement-related precentral neurones in conscious monkeys. *Proc R Soc Lond B Biol Sci* 194:313–339. [CrossRef Medline](#)
- Maranesi M, Livi A, Fogassi L, Rizzolatti G, Bonini L (2014) Mirror neuron activation prior to action observation in a predictable context. *J Neurosci* 34:14827–14832. [CrossRef Medline](#)
- Maranesi M, Livi A, Bonini L (2017) Spatial and viewpoint selectivity for others' observed actions in monkey ventral premotor mirror neurons. *Sci Rep* 7:8231. [CrossRef Medline](#)
- Meunier N, Marion-Poll F, Lansky P, Rospars JP (2003) Estimation of the individual firing frequencies of two neurons recorded with a single electrode. *Chem Senses* 28:671–679. [CrossRef Medline](#)
- Mollazadeh M, Aggarwal V, Davidson AG, Law AJ, Thakor NV, Schieber MH (2011) Spatiotemporal variation of multiple neurophysiological signals in the primary motor cortex during dexterous reach-to-grasp movements. *J Neurosci* 31:15531–15543. [CrossRef Medline](#)
- Muakkassa KF, Strick PL (1979) Frontal lobe inputs to primate motor cortex: evidence for four somatotopically organized 'premotor' areas. *Brain Res* 177:176–182. [CrossRef Medline](#)
- Rabiner LR (1989) A tutorial on hidden Markov models and selected applications in speech recognition. *Proc IEEE* 77:257–286. [CrossRef](#)
- Rizzolatti G, Scandolara C, Matelli M, Gentilucci M (1981a) Afferent properties of periarculate neurons in macaque monkeys. II. visual responses. *Behav Brain Res* 2:147–163. [CrossRef Medline](#)
- Rizzolatti G, Scandolara C, Matelli M, Gentilucci M (1981b) Afferent properties of periarculate neurons in macaque monkeys. I. somatosensory responses. *Behav Brain Res* 2:125–146. [CrossRef Medline](#)
- Rizzolatti G, Fadiga L, Gallese V, Fogassi L (1996) Premotor cortex and the recognition of motor actions. *Brain Res Cogn Brain Res* 3:131–141. [CrossRef Medline](#)
- Rochat MJ, Caruana F, Jezzini A, Escola L, Intskirveli I, Grammont F, Gallese V, Rizzolatti G, Umiltà MA (2010) Responses of mirror neurons in area F5 to hand and tool grasping observation. *Exp Brain Res* 204:605–616. [CrossRef Medline](#)

- Rouse AG, Schieber MH (2016) Spatiotemporal distribution of location and object effects in primary motor cortex neurons during reach-to-grasp. *J Neurosci* 36:10640–10653. [CrossRef Medline](#)
- Rozzi S, Ferrari PF, Bonini L, Rizzolatti G, Fogassi L (2008) Functional organization of inferior parietal lobule convexity in the macaque monkey: electrophysiological characterization of motor, sensory and mirror responses and their correlation with cytoarchitectonic areas. *Eur J Neurosci* 28:1569–1588. [CrossRef Medline](#)
- Sainburg RL, Ghilardi MF, Poizner H, Ghez C (1995) Control of limb dynamics in normal subjects and patients without proprioception. *J Neurophysiol* 73:820–835. [CrossRef Medline](#)
- Sergio LE, Kalaska JF (2003) Systematic changes in motor cortex cell activity with arm posture during directional isometric force generation. *J Neurophysiol* 89:212–228. [CrossRef Medline](#)
- Shimazu H, Maier MA, Cerri G, Kirkwood PA, Lemon RN (2004) Macaque ventral premotor cortex exerts powerful facilitation of motor cortex outputs to upper limb motoneurons. *J Neurosci* 24:1200–1211. [CrossRef Medline](#)
- Tkach D, Reimer J, Hatsopoulos NG (2007) Congruent activity during action and action observation in motor cortex. *J Neurosci* 27:13241–13250. [CrossRef Medline](#)
- Umiltà MA, Kohler E, Gallese V, Fadiga L, Keysers C, Rizzolatti G (2001) I know what you are doing. a neurophysiological study. *Neuron* 31:155–165. [CrossRef Medline](#)
- Venkadesan M, Valero-Cuevas FJ (2008) Neural control of motion-to-force transitions with the fingertip. *J Neurosci* 28:1366–1373. [CrossRef Medline](#)
- Vigneswaran G, Philipp R, Lemon RN, Kraskov A (2013) M1 corticospinal mirror neurons and their role in movement suppression during action observation. *Curr Biol* 23:236–243. [CrossRef Medline](#)