Title: Progressively shifting patterns of co-modulation among premotor cortex neurons carry dynamically similar signals during action execution and observation

Authors: Zhonghao Zhao[1] and Marc H. Schieber[1,2,3*]

Affiliations:
[1]Department of Biomedical Engineering, University of Rochester, Rochester, NY, 14627
[2]Department of Neurology, University of Rochester, Rochester, NY, 14642
[3]Department of Neuroscience, University of Rochester, Rochester, NY 14642

***Corresponding Author:**
Marc H. Schieber
E-mail: mschiebe@ur.rochester.edu

Number of:

Pages -	36
Figures -	15
Tables -	2
Multimedia -	0
3D models -	0

Words in:

Abstract -	306
Introduction -	1348
Discussion -	2255

Conflict of interest:

The authors declare no conflicts of interest.

Acknowledgements:

The authors thank John Housel and Jennifer Gardinier for technical assistance, and Gil Rivlis for custom task-control software. This work was supported by grant R01NS102343 (MHS) from the National Institute of Neurological Disorders and Stroke.

ABSTRACT

Many neurons in the premotor cortex show firing rate modulation whether the subject performs an action or observes another individual performing a similar action. Although such "mirror neurons" have been thought to have highly congruent discharge during execution and observation, many if not most show non-congruent activity. Studies of such neuronal populations have shown that the most prevalent patterns of co-modulation—captured as neural trajectories—pass through subspaces which are shared in part, but in part are visited exclusively during either execution or observation. These studies focused on reaching movements for which low-dimensional neural trajectories exhibit comparatively simple dynamical motifs. But the neural dynamics of hand movements are more complex. We developed a novel approach to examine prevalent patterns of co-modulation during execution and observation of a task that involved reaching, grasping, and manipulation. Rather than following neural trajectories in subspaces that contain their entire time course, we identified time series of instantaneous subspaces, calculated principal angles among them, sampled trajectory segments at the times of selected behavioral events, and projected those segments into the series of instantaneous subspaces. We found that instantaneous neural subspaces generally remained distinct during execution versus observation. Nevertheless, execution and observation could be partially aligned with canonical correlation, indicating some dynamical similarity of the neural representations of different movements relative to one another during execution and observation which may enable the nervous system to recognize corresponding actions performed by the subject or by another individual and/or may reflect social interaction between the two. During action execution, mirror neurons showed consistent patterns of co-modulation both within and between sessions, but other neurons that were modulated only during action execution and not during observation showed considerable variability of co-modulation. We speculate that during execution, mirror neurons carry a consistent forward model of the intended movement, while action-execution only neurons process more variable feedback.

INTRODUCTION

Although the premotor (PM) and primary motor cortex (M1) are generally thought to be involved in the planning and execution of movement, many neurons in these cortical motor areas have been found to discharge not only when the subject executes a movement, but also when the subject observes a similar movement being performed by another individual. Such neurons have been found in the ventral premotor cortex (PMv) (Bonini et al., 2014; Gallese et al., 1996), dorsal premotor cortex (PMd) (Cisek and Kalaska, 2004; Papadourakis and Raos, 2019; Pezzulo et al., 2022), and M1 (Dushanova & Donoghue, 2010; Kraskov et al., 2014a; Vigneswaran et al., 2013).

Early studies of these neurons emphasized those with congruent discharge during execution and observation conditions. Congruent neurons discharged during the same type of grasp (Gallese et al., 1996; Rizzolatti et al., 1996), or retained the same preferred direction (Dushanova & Donoghue, 2010; Kilner & Lemon, 2013) during both execution and observation. Emphasis on such congruent neurons led to the notion that they mediate understanding of observed actions by mirroring their own activity during execution (di Pellegrino et al., 1992; Rizzolatti and Craighero, 2004).

In addition to congruent neurons, however, even early studies also reported many other noncongruent neurons that also discharged during execution and during observation, but discharged differently in the two contexts (Gallese et al., 1996). For example, of the pyramidal tract neurons (PTNs) in PMv and M1 that showed modulation during both execution and observation, half showed substantially lower firing rates during observation than during execution (Kraskov et al., 2009; Vigneswaran et al., 2013; Kraskov et al., 2014). In many studies roughly half or more of the neurons modulated during both execution and observation were noncongruent (Dushanova and Donoghue, 2010; Kraskov et al., 2014; Mazurek et al., 2018; Jiang et al., 2020). Of PMv neurons modulated during both execution and observation, over the time course of behavioral trials only ~20% showed brief periods with strictly congruent firing rates (Pomper et al., 2023). And in both PMv (F5) and PMd, the proportion of congruent neurons may not be different from that expected by chance alone (Papadourakis and Raos, 2019). That so many neurons are active differentially during action execution versus observation calls into question the extent to which the representation of movements by these neuron populations actually matches in the two contexts. Although certain authors apply the term mirror neurons (MNs) strictly to highly congruent neurons, like some other authors, we will refer to all neurons modulated during both contexts—execution and observation—as MNs.

Addressing this issue at the population level is complex because of the wide variety of firing rate modulation found in individual neurons at different times in the course of behavioral trials, as well as among neurons in a recorded population. Behavior evolving in time may be represented more accurately by the temporal progression of co-modulation in populations of neurons than by the temporal pattern of firing rate in single neurons (Shenoy et al., 2013; Cunningham and Yu, 2014; Vyas et al., 2020). Patterns of co-modulation can be considered in a high-dimensional neural-state space where the firing rate of each neuron is a separate, orthogonal dimension. The instantaneous, simultaneous firing rates of all *N* neurons then is a point in this space, which traces out a trajectory over time. Neural population trajectories do not visit all regions of the *N*-dimensional state-space, however. Dimensionality reduction

109 techniques can be used to identify a small set of latent dimensions—a subspace—that captures
110 the most prevalent patterns of co-modulation among the population of *N* neurons.

111 Studies of neural trajectories underlying action execution that focused on reaching
112 movements made with the arm have revealed that rotational motifs in a low-dimensional
113 subspace capture much of the neural population's firing rate variance (Churchland et al., 2012).
114 While similar rotational dynamics were found in M1 during repetitive cycling movements, the
115 population trajectory of supplementary motor area neurons progressed as a helix through an
116 additional dimension during successive cycles (Russo et al., 2020). The M1 neural trajectories
117 underlying grasping movements (Suresh et al., 2020) or force production at the wrist (Dekleva
118 et al., 2024) are still more complex. The latent subspaces that capture the predominant
119 patterns of co-modulation among M1 neurons, for example, shift progressively over the time
120 course of behavioral trials involving reaching to, grasping, and manipulating (RGM) various
121 objects at various locations (Rouse and Schieber, 2018). Relatively few studies have examined
122 the trajectories of neural populations that are active during both execution and observation
123 (Mazurek et al., 2018; Jerjian et al., 2020; Jiang et al., 2020; Albertini et al., 2021; Pezzulo et al.,
124 2022).

125 A relevant but often overlooked aspect of such dynamics in neuron populations active
126 during both execution and observation has to do with the distinction between condition-
127 independent and condition-dependent variation in neuron activity (Rouse and Schieber, 2018).
128 The variance in neural activity averaged across all the conditions in a given task context is
129 condition-independent. For example, in an 8-direction, center-out reaching task, averaging a
130 unit's firing rate as a function of time across all 8 directions may show an initially low firing rate
131 that increases prior to movement onset, peaks during the movement, and then declines during
132 the final hold, irrespective of the movement direction. Subtracting this condition-independent
133 activity from the total activity during trials in each of the 8 directions gives the remaining
134 variance, and averaging separately across trials in each of the 8 directions then averages out
135 noise variance leaving the condition-dependent variance that represents the unit's modulation
136 among the 8 directions (conditions). Alternatively, condition-independent, condition-dependent,
137 and noise variance can be partitioned through demixed principal component analysis (Kobak et
138 al., 2016; Gallego et al., 2018). The condition-independent activity can be viewed as
139 representing the performance of a movement in some general class; the condition-dependent
140 activity as representing which particular movement in that class. The extent to which neural
141 dynamics occur in a subspace shared by execution and observation versus subspaces unique
142 to execution or observation may differ for the condition-independent versus condition-dependent
143 partitions of neural activity. Here, we tested the hypothesis that the condition-dependent activity
144 of PM mirror neuron populations progresses through distinct subspaces during execution versus
145 observation.

146 Because of the complexity of condition-dependent neural trajectories for movements
147 involving the hand, we developed a novel approach. Rather than examining trajectories over
148 the entire time course of behavioral trials, we first identified time series of instantaneous PM
149 mirror neurons subspaces covering the course of behavioral trials, identifying separate time
150 series for execution trials and for observation trials. Given that each subspace in these time
151 series is instantaneous (a snapshot in time), it captures condition-dependent variance in the
152 neural activity while minimizing condition-independent (time-dependent) variance. We then
153 tested the hypothesis that the condition-dependent subspace shifts progressively over the time

course of behavioral trials (Figure 1A) by calculating the principal angles between four selected instantaneous subspaces that occurred at times easily defined in each behavioral trial—instruction onset (I), go cue (G), movement onset (M), and the beginning of the final hold (H)—and every other instantaneous subspace in the time series. Because initial analyses showed that condition-dependent neural trajectories for different movements tended to separate increasingly over the course of behavioral trials, we additionally examined the combined effect of progressively shifting subspaces and increasing trajectory separation by decoding neural trajectory segments sampled for 100 msec after times I, G, M, and H and projected into the time series of instantaneous subspaces (Figure 1B). We then asked whether execution and observation trajectory segments could be aligned with canonical correlation (Figure 1C). Such alignment would indicate that the relationships among the trajectory segments in the execution subspace were similar to the relationships among the trajectory segments in the observation subspace, indicating a corresponding structure in the latent dynamic representations of execution and observation movements by the same PM MN population. And finally, because we previously have found that the activity of PM mirror neurons tends to lead that of neurons active only during action execution (AE) (Mazurek and Schieber, 2019), we performed parallel analyses of the instantaneous state space of PM AE neurons.

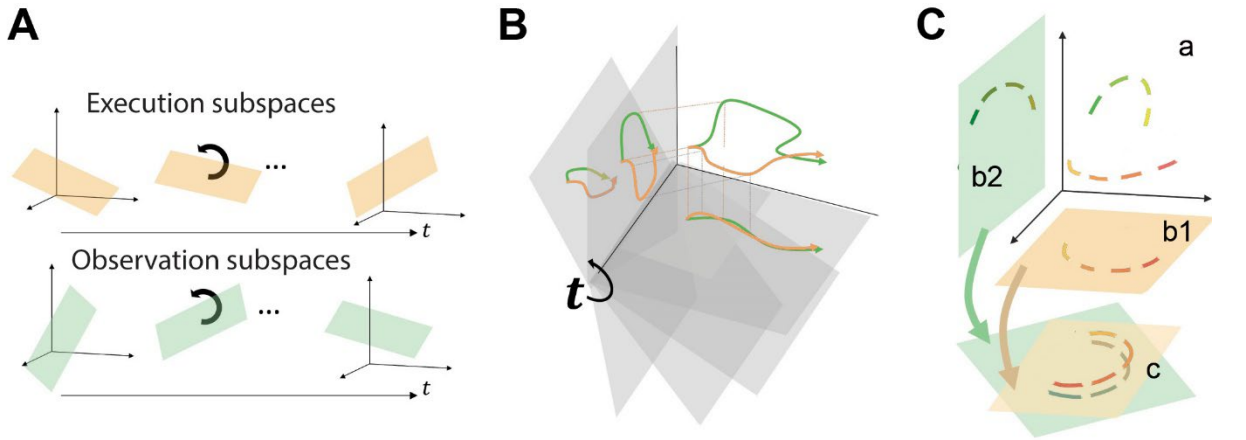

Figure 1. Conceptual approach. **A.** We hypothesized that the condition-dependent subspace of PM MN activity shifts progressively through the time course of behavioral trials both during execution (orange) and during observation (green). **B.** Neural trajectory segments (orange, green) were projected into time series of instantaneous subspaces (gray). **C.** Neural trajectory segments (latent dynamics) from the four RGM movements during execution (orange) and during observation (green) originate in the same high-dimensional space **(a),** but project into distinct low-dimensional subspaces **(b1, b2)**; nevertheless they can be aligned with CCA **(c)**, demonstrating similar latent dynamic relationships among the four movements during execution and observation.

RESULTS

We recorded spiking activity as each of three monkeys executed the RGM task, and then as each monkey observed the same RGM task being performed by an experimenter (Figure 2A). Because we chose to study relatively naturalistic movements, the reach, grasp, and manipulation components were not performed separately, but rather in a continuous fluid motion during the movement epoch of the task sequence (Figure 2B). In previous studies of a version of this task without separate instruction or delay epochs, we have shown that joint kinematics, EMG activity, and neuron activity in the primary motor cortex, all vary throughout the movement

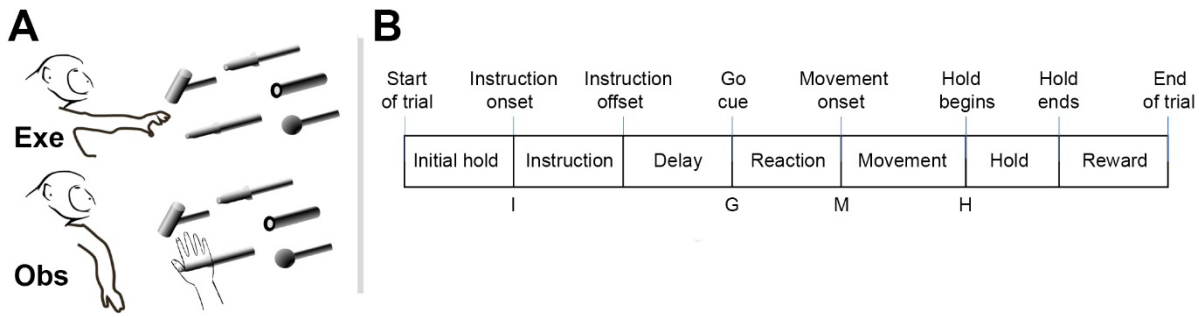

Figure 2. The reach-grasp-manipulate (RGM) task. **A.** In separate blocks of trials monkeys reached to, grasped, and manipulated four different objects themselves (Exe), and then observed a human perform the same task (Obs). **B.** The times of eight behavioral events from Start-of-trial to End-of-trial divided each trial into seven epochs from Initial hold to Reward. For analyses the data were aligned separately on, and trajectories were sampled for 100 msec following, the times of four selected events—Instruction onset (I), Go cue (G), Movement onset (M), and Hold (H).

epoch in relation to both reach location and object, with location predominating early in the movement epoch and object predominating later (Rouse and Schieber, 2015, 2016b, a). Our task thus did not dissociate the reach, the hand shape used to grasp the object, and the manipulation performed on the object. Additional details of the behavioral task are described in the Methods. Three sessions were recorded from each of the three monkeys, R, F, and T (a 6 kg female, 10 kg male, and 10 kg male, respectively). The numbers of successful execution trials (Exe) and observation trials (Obs) involving each of the four objects—sphere, button, coaxial cylinder, and perpendicular cylinder—are given in Table 1.

Table 1. Numbers of trials in each session

	Monkey R		Monkey F		Monkey T	
	Exe	Obs	Exe	Obs	Exe	Obs
Session 1	(22,8,25,26)	(32,31,30,31)	(58,59,62,63)	(71,72,71,72)	(57,54,57,55)	(60,61,59,57)
Session 2	(34,26,34,38)	(40,41,40,37)	(59,58,60,56)	(73,72,75,74)	(47,53,52,43)	(57,53,58,58)
Session 3	(42,41,49,45)	(49,50,51,49)	(63,58,58,58)	(72,75,74,74)	(43,41,38,42)	(50,48,48,50)

Table 1. Numbers of trials in each session. For each of the three sessions from each of the three monkeys, numbers of trials involving each of the four objects (sphere, button, coaxial cylinder, perpendicular cylinder) are given in parentheses separately for execution and for observation.

Using object and task time period as factors, we performed two-way repeated measures analysis of variance (rmANOVA) on the firing rate of each sorted unit (see Methods). Because unit firing rates frequently differed during execution and observation, we performed such rmANOVAs separately on execution trials and observation trials. Table 2 gives the numbers of PM (PMv+PMd) units identified in each session as being modulated significantly during both execution and observation, which we refer to as mirror neurons (MN), along with the numbers of units modulated significantly during execution but not observation (AE), during observation but not execution (AO), or with no significant modulation during either execution or observation (NS). The numbers of AO and NS units were consistently small across monkeys and sessions. The present analyses therefore focus on MNs and, for comparison, AE neurons.

Table 2. Numbers of units in each session

Monkey	Session	MN	AE	AO	NS
R	1	48(19,29)	35(20,15)	3(1,2)	5(2,3)
	2	47(21,26)	25(16,9)	5(1,4)	11(4,7)
	3	37(19,18)	49(20,29)	1(1,0)	8(7,1)
T	1	79(37,42)	15(5,10)	2(0,2)	7(1,6)
	2	91(48,43)	22(6,16)	3(1,2)	7(1,6)
	3	100(48,52)	18(7,11)	0(0,0)	6(2,4)
F	1	44(24,20)	7(5,2)	1(1,0)	8(8,0)
	2	47(32,15)	10(9,1)	5(1,4)	3(3,0)
	3	42(28,14)	9(7,2)	3(1,2)	3(3,0)

Table 2. Numbers of PM units in each session. For each of the three sessions from each of the three monkeys (R, T, and F), numbers of PM units are given for each of four classes in the format of Total (PMv, PMd). MN – units modulated significantly during action execution and during action observation. AE – units modulated during execution but not during observation. AO – units modulated during observation but not execution. NS – units not modulated significantly during either execution or observation.

Condition-dependent versus condition-independent neural activity in PM MNs

Whereas a large fraction of condition-dependent neural variance during reaching movements without grasping can be captured in a two-dimensional subspace (Churchland et al., 2012; Ames et al., 2014), condition-dependent activity in movements that involve grasping is more complex (Suresh et al., 2020). In part, this may reflect the greater complexity of controlling the 24 degrees of freedom in the hand and wrist as compared to the 4 degrees of freedom in the elbow and shoulder (Sobinov and Bensmaia, 2021). Figure 3 illustrates similar complexity in a PM MN population during the present RGM movements. Here, PCA was performed on the activity of the PM MN population across the entire time course of execution trials involving all four objects. The colored traces in Figure 3A show the neural trajectories averaged separately across trials involving each of the four objects and then projected into the PC1 vs PC2 plane of the total neural space. Most of the variance in these four trajectories is comprised of a shared

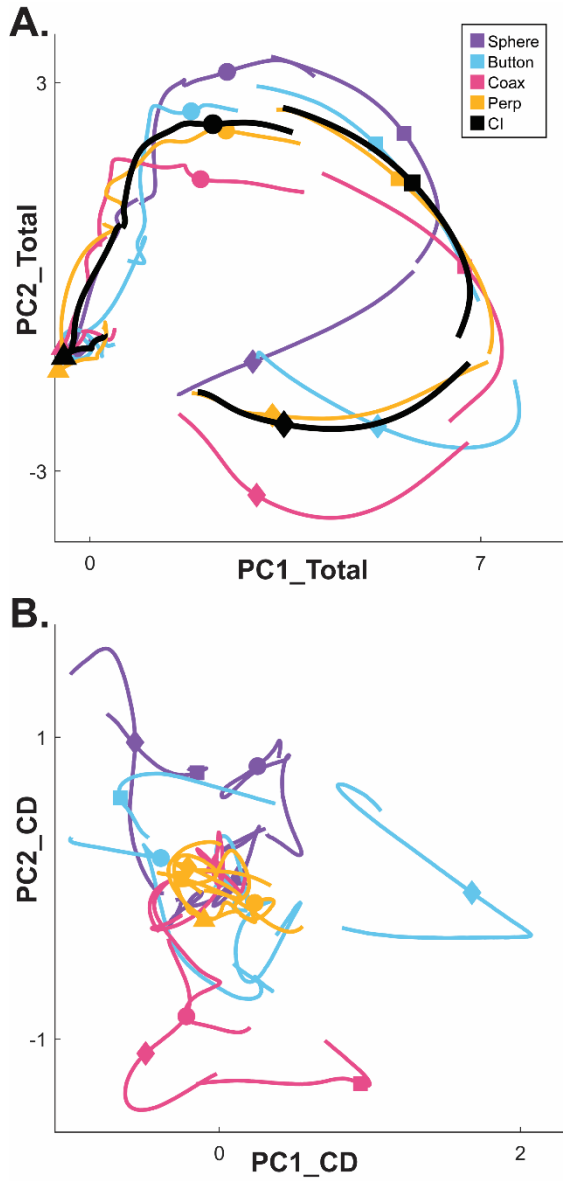

Figure 3. Neural trajectories of condition-independent versus condition-dependent activity. **A**. Neural trajectories of PM MN firing rates averaged across multiple execution trials involving each of the four objects (Sphere – purple, Button – cyan, Coaxial cylinder [Coax]– magenta, Perpendicular cylinder [Perp]– yellow) have been projected into the PC1 vs PC2 plane of the Total neural activity. Averaging these four trajectories gives their common, condition-independent (CI) trajectory (black). Time proceeds clockwise from left, with data separately aligned at four selected times: triangle – instruction onset (I); circle – go cue (G); square – movement onset (M); diamond – beginning of final hold (H). **B**. Condition-dependent trajectories obtained by subtracting the CI trajectory (black) from each of the four single-object trajectories (colors) in **A**, and then projected into the PC1 vs PC2 plane of their common subspace across the entire time course of trials. Data from monkey R, session 2.

rotational component. The black trajectory, obtained by averaging trajectories from trials involving all four objects together, represents this condition-independent (i.e. independent of the object involved) activity. The condition-dependent (i.e. dependent on which object was involved) variation in activity is reflected by the variation in the colored trajectories around the black trajectory. The condition-dependent portions can be isolated by subtracting the black trajectory from each of the colored trajectories. The resulting four condition-dependent trajectories have been projected into the PC1 vs PC2 plane of their own common subspace in Figure 3B. Rather than exhibiting a simple rotational motif, these trajectories appear knotted. To better understand how these complex, condition-dependent trajectories progress over the time course of RGM trials, we chose to examine their instantaneous subspaces.

Instantaneous subspaces shift progressively during both execution and observation

We identified an instantaneous subspace at each millisecond time step of RGM trials. At each time step, we applied PCA to the 4 trial-averaged neural states (i.e. the 4 points on the average neural trajectories representing trials involving the 4 different objects) yielding a 3-dimensional subspace at that time (see Methods). Note that because these 3-dimensional subspaces are essentially instantaneous, they capture the condition-dependent variation in neural states, but not the time-varying, condition-independent variation. To examine the temporal progression of these instantaneous subspaces, we then calculated the principal angles between each instantaneous subspace and the subspaces at four behavioral time points readily defined across trials, sessions, and monkeys: the onsets of the instruction (I), the go cue (G),

the movement (M), and the hold (H). As illustrated in the Methods (Figure 14A), in each session all three principal angles, each of which can range from 0° to 90°, tended to follow a similar time course, and in the Results we therefore illustrate only the first (i.e. smallest) principal angle. The closer the principal angles are to 0°, the closer two subspaces are to being identical; the closer to 90°, the closer the two subspaces are to being mutually orthogonal.

The left half of Figure 4 (columns A-D) illustrates the temporal progression of the first principal angle in the three sessions (red, green, and blue) from each monkey during execution trials. By definition the instantaneous subspace becomes identical to the subspace (principal angle of 0°) at each of the four selected times—I, G, M, and H—with the sharp decrease and then increase reflecting the 50 ms smoothing of neural firing rates. Of greater interest are the slower changes in the principal angle between these four time points. Figure 4A shows that after instruction onset (I) the instantaneous subspace shifted quickly away from the subspace at time I, indicated by a rapid increase in principal angle. Throughout the remainder of the instruction and delay epochs (from I to G), however, Figure 4B and C show that the instantaneous subspace shifted concurrently, not sequentially, toward the subspaces that would be reached at the end of the delay period (G) and then at the onset of movement (M), indicated by the progressive decreases in principal angle. In contrast, Figure 4D shows that shifting toward the H subspace did not begin until the movement onset (time M, monkeys R and T) or go cue (time G, monkey F). Considerable similarity was evident across sessions within each monkey (red, green, and blue traces) and across monkeys as well. These changes in principal angles indicate that after shifting briefly toward the subspace present at time the instruction appeared (I), the instantaneous subspace shifted progressively throughout the delay epoch toward the orientation that would be reached at the time of the go cue (G), then further toward that at the time of movement onset (M), and only thereafter toward that present at the time of the hold (H).

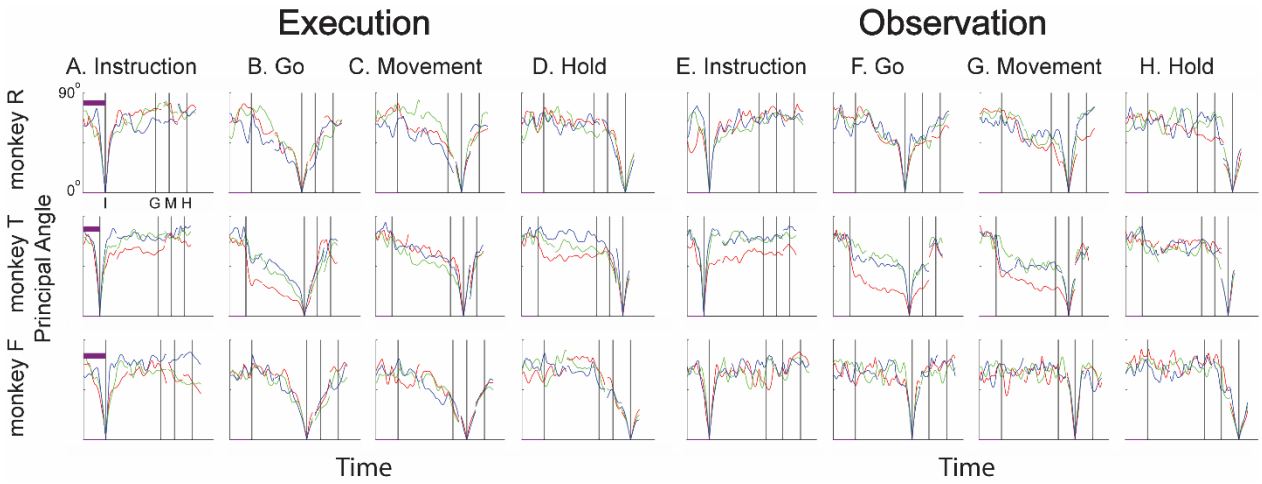

Figure 4. Time course of the first principal angle of MN instantaneous subspaces. Results from Execution trials are shown on the left; from Observation trials on the right. Each frame shows the time course of the first principal angle between each instantaneous subspace and that present at one of four selected times—**A, E.** instruction onset; **B, F.** go cue; **C, G.** movement onset; or **D, H.** the beginning of the final hold. The four vertical lines in each frame mark the times of I, G, M and H, as indicated in the upper left frame. Red, green, and blue traces represent sessions 1, 2, and 3, respectively, from each monkey. Purple bars in the left column (**A**) indicate 500 ms, which applies to the entire row for that monkey.

In addition to this progressive shifting of the instantaneous subspace common to all three monkeys, the principal angle time course showed another feature that was prominent only in monkey T, particularly in session 1 (red). Promptly after instruction onset (I) the principal angle abruptly decreased in relation to the subspace that would be present at time G (Figure 4B) and to a lesser degree in relation to the subspaces at times M and H (Figure 4C,D). These abrupt decreases in principal angle in this monkey indicate that during execution trials, after the instruction onset, the instantaneous subspace shifted with both a gradual and an abrupt component toward the subspaces that would be present at times G, M, and H, with the gradual component being more consistent across sessions and monkeys than the abrupt.

The right half of Figure 4 shows the progression of the first principal angle during observation trials. Overall, the temporal progression of the MN instantaneous subspace during observation was similar to that found during execution. Both monkey R and monkey T showed comparatively consistent gradual decreases in principal angle (along with the more variable abrupt decrease in monkey T) during observation that roughly paralleled the decreases found during execution. Monkey F was an exception, however, in that little if any decrease in principal angle occurred during the delay epoch, though a gradual decrease did occur approaching time H (Figure 4H).

We also examined the temporal progression of the instantaneous subspace of AE neurons. The left half of Figure 5 shows that during execution the AE population in monkey R had a pattern of gradual decrease in principal angle similar to that found in the MN population. The gradual decrease was present though less evident in monkey T, and still less evident in monkey F, reflecting the smaller numbers of AE neurons recorded in these two monkeys which constrained the range of principal angles between three-dimensional subspaces (see Methods, Figure 14B). AE neurons were not modulated significantly during observation trials, and hence the right half of Figure 5 shows little if any progressive change in principal angle.

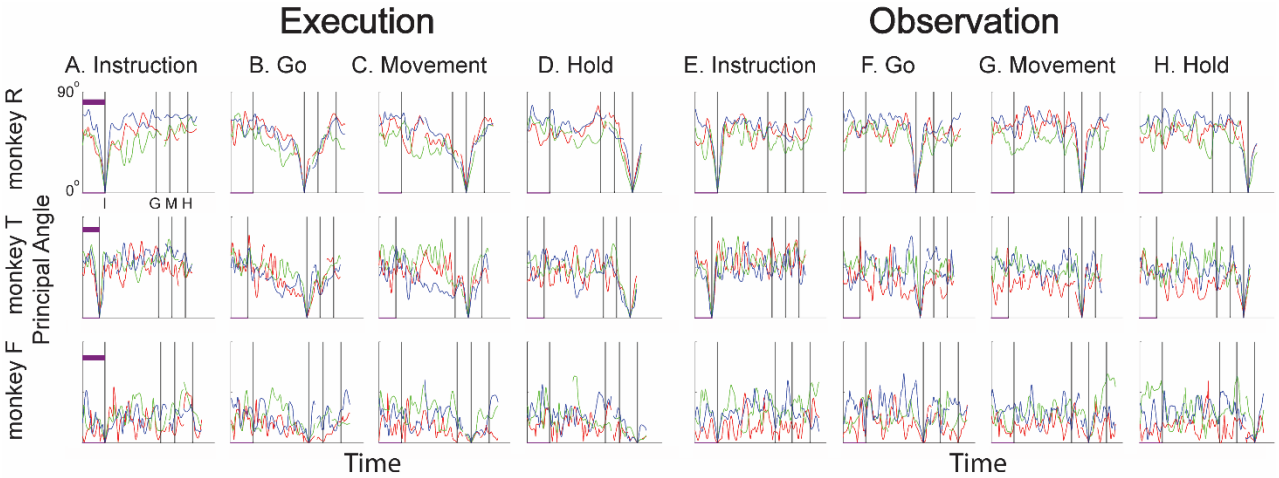

Figure 5. Time course of the first principal angle of AE neuron instantaneous subspaces. Results from Execution trials are shown on the left; from Observation trials on the right. Each frame shows the time course of the first principal angle between each instantaneous subspace and that present at one of four selected times—**A, E.** instruction onset; **B, F.** go cue; **C, G.** movement onset; or **D, H.** the beginning of the final hold. The four vertical lines in each frame mark the times of I, G, M and H, as indicated in the upper left frame. Red, green, and blue traces represent sessions 1, 2, and 3, respectively, from each monkey. Purple bars in the left column (**A**) indicate 500 ms, which applies to the entire row for that monkey.

291

292 **Neural trajectories separate progressively during both execution and observation**

293 The progressive changes in principal angles do not capture another important aspect of
294 condition-dependent neural activity. The neural trajectories during trials involving different
295 objects separated increasingly as trials progressed in time. To illustrate this increasing
296 separation, we clipped 100 ms segments of high-dimensional PM MN population trial-averaged
297 trajectories beginning at times I, G, M, and H, for trials involving each of the four objects. We
298 then projected the set of four object-specific trajectory segments clipped at each time into each
299 of the four instantaneous 3D subspaces at times I, G, M, and H. This process was repeated
300 separately for execution trials and for observation trials.

301 For visualization, we projected these trial-averaged MN trajectory segments from an
302 example session into the PC1 vs PC2 planes (which consistently captured > 70% of the
303 variance) of the I, G, M, or H instantaneous 3D subspaces. In Figure 6, the trajectory segments
304 for each of the four objects (sphere – purple, button – cyan, coaxial cylinder – magenta,
305 perpendicular cylinder – yellow) sampled at different times (rows) have been projected into each
306 of the four instantaneous subspaces (columns). Rather than appearing knotted as in Figure 3,
307 these short trajectory segments are distinct when projected into each instantaneous subspace.

308 Each set of trajectory segments is projected into its corresponding subspace along the
309 main diagonal, showing that during execution (Figure 6A) the trajectory segments for the four
310 objects are close together at the time of instruction onset (I), are more separated at the time of
311 the go cue (G), have separated further still at movement onset (M), and have become
312 somewhat less separated at the beginning of the final hold (H). During observation (Figure 6B)
313 a similar trend is evident along the main diagonal, although the separation is less, reflecting the
314 commonly described lower firing rates of MNs during observation than during execution (Ferroni
315 et al., 2021). In addition, during observation the separation of the four trajectories was
316 somewhat greater at the beginning of the hold (H) than at movement onset (M).

317 Off-diagonal frames along the rows (same trajectory segments, different instantaneous
318 subspaces) or along the columns (different trajectory segments, same instantaneous
319 subspaces) show less separation than along the main diagonal, both during execution and
320 during observation. These differences reflect the progressive shifting of the condition-dependent
321 instantaneous subspace of the PM MN population as trials progressed in time.

322 To summarize these changes in trajectory separation, we calculated the 3-dimensional
323 cumulative separation (C*S*, i.e. the summed pointwise Euclidean distance between all pairwise
324 combinations of the four object-specific trajectory segments, see Methods) for each set of four
325 segments projected into each of the four instantaneous subspaces. *CS* values, which we use
326 only to characterize the phenomenon of trajectory separation, are illustrated for execution from
327 this example session as the color matrix of Figure 6C; for observation, Figure 6D. For both
328 execution and observation, the highest *CS* values lie on the main diagonal, increasing in
329 temporal order from Instruction to Go to Movement to Hold, with the exception that for execution
330 *CS* for Hold was less than for Movement. Figure 6E and 6F show *CS* matrices averaged across
331 all three sessions from all three monkeys, demonstrating that the features seen in the example
332 session were relatively consistent across sessions. During both execution and observation, as
333 the instantaneous subspace of the PM MN population shifted progressively over the time course
334 of RGM trials, the separation of condition-dependent neural trajectories also increased.

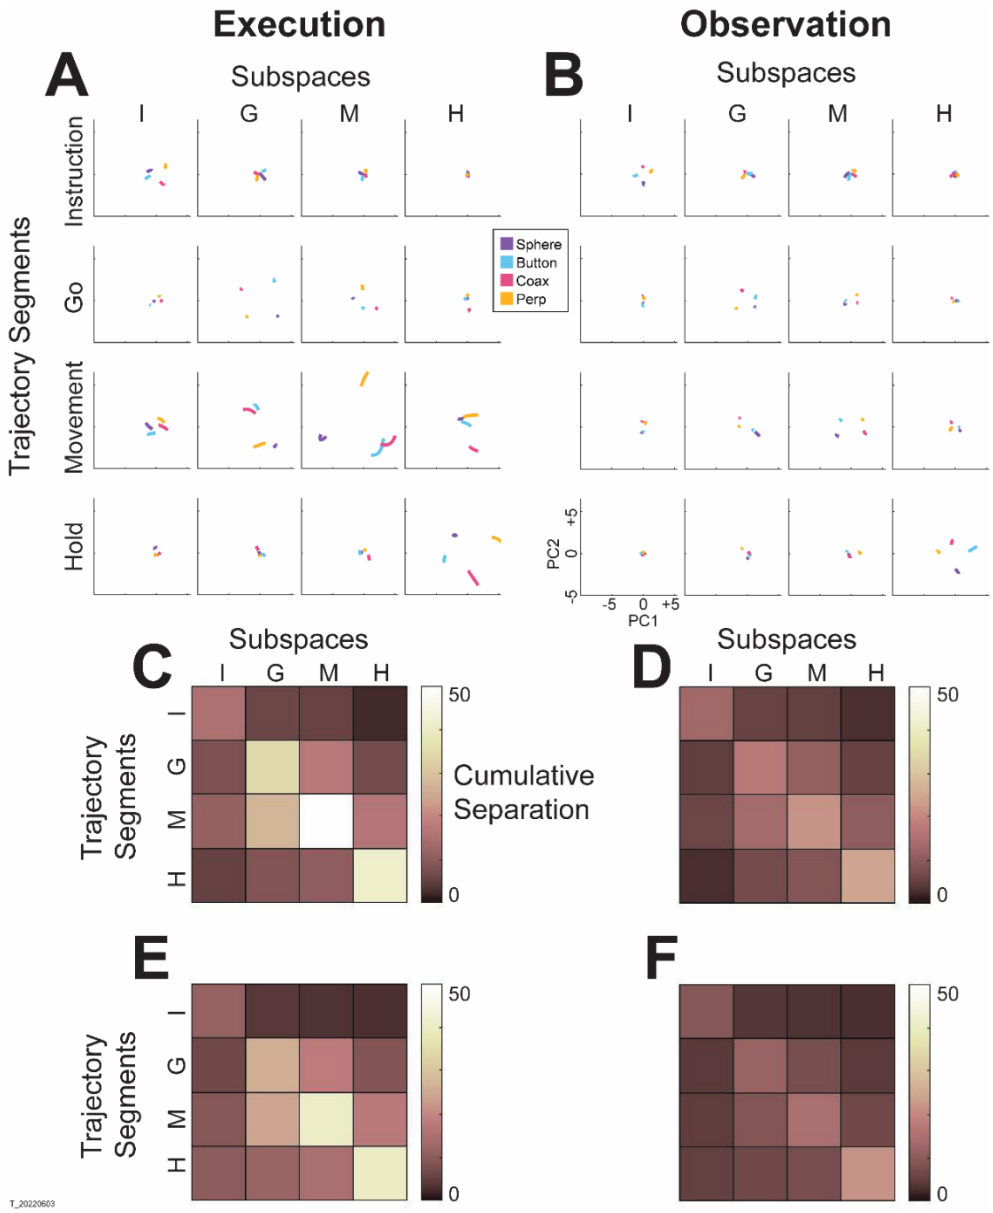

Figure 6. MN trajectory segments projected into instantaneous subspaces**. A.** Using execution data from an example session (monkey T, session 3), trajectory segments averaged across trials involving each of the four objects (sphere – purple, button – cyan, coaxial cylinder [coax] – magenta, perpendicular cylinder [perp] – yellow) were clipped for 100 ms immediately following each of four behavioral events (rows: Instruction onset, Go cue, Movement onset, Hold). Each set of these four segments then was projected into the PC1 vs PC2 plane of the instantaneous 3D subspace present at four different times (columns: I, G, M, H). **B.** The same process was performed using observation data from the same session. The PC1 vs PC2 scales at lower left apply to all frames in both **A** and **B**. **C and D.** Cumulative separation values (*CS*, see Methods) calculated in 3D for each of the frames projected in 2D in **A and B**, respectively, are shown as color matrices. **E** and **F** show *CS* values averaged across all 9 sessions from all 3 monkeys for execution and observation, respectively. Data in **A**, **B**, **C**, and **D** are from monkey T, session 3.

Decodable information changes progressively during both execution and observation

As RGM trials proceeded in time, the condition-dependent neural activity of the PM MN population thus changed in two ways. First, the instantaneous condition-dependent subspace shifted, indicating that the patterns of firing-rate co-modulation among neurons representing the four different RGM movements changed progressively, both during execution and during observation. Second, as firing rates generally increased, the neural trajectories representing the four RGM movements became progressively more separated, more so during execution than during observation.

To evaluate the net effect of these two progressive changes combined, we clipped 100 ms single-trial trajectory segments beginning at times I, G, M, or H, and projected these trajectory segments from individual trials into the instantaneous 3D subspaces at 50 ms time steps. At each of these time steps, we trained a separate LSTM decoder to classify individual trials according to which of the four objects was involved in that trial. We expected that the trajectory segments would be classified most accurately when projected into instantaneous subspaces near the time at which the trajectory segments were clipped. At other times we reasoned that classification accuracy would depend both on the similarity of the current instantaneous subspace to that found at the clip time, as evaluated by the principal angle (Figure 4), and on the separation of the four trajectories at the clip time (Figure 6).

Figure 7A-D (left half) shows the resulting classification accuracy as a function of trial time for the 100 ms Instruction, Go, Movement, or Hold execution trajectory segments, each projected into the same time series of instantaneous execution subspaces from the same session. Solid curves indicate classification accuracy averaged across 10-fold cross-validation (as described in the Methods); the surrounding shaded areas indicate ± 1 standard deviation from that average; different colors represent the three sessions from the same monkey, with black being their average. Horizontal lines indicate the range of classification accuracies that would have been obtained had the instantaneous subspaces been chosen randomly, which we estimated for each set of trajectory segments by bootstrapping—projecting the trajectory segments into a randomly selected 3D space, training an LSTM decoder, and classifying single trials, repeated 500 times (Natraj et al., 2022).

As might have been expected based both on principal angles and on trajectory separation, classification accuracy consistently peaked at a time point within or near the 100 ms duration of the corresponding trajectory segments (orange flags at the top of vertical lines). Classification accuracy decreased progressively at times preceding and following each of these peaks. In monkey R (top row), for example, mean classification of the Instruction trajectory segments (Figure 7A) initially was close to 0.25, rose to 0.49 around the time of the instruction onset, and then fell back to 0.25. Mean accuracy for the Go segments (Figure 7B) also began close to 0.25, rose gradually during the delay epoch to peak at 0.76 around the time of the Go cue, and decreased thereafter. For the Movement (Figure 7C) and Hold (Figure 7D) segments, classification accuracy started somewhat higher (reflecting greater trajectory segment separation at the time they were clipped, Figure 6) and peaked at 0.87 and 0.89 near the time of those events, respectively. Similar trends were seen for monkeys T (middle row) and F (bottom row). For each monkey, classification accuracy for each of the four sets of trajectory segments—Instruction, Go, Movement, and Hold—as a function of time was quite consistent.

Although classification accuracy consistently peaked near the time of the behavioral event immediately after which each set of trajectory segments was sampled, the rise in accuracy before and the decline after the peak differed depending on the behavioral event. Peak classification accuracy for Instruction segments was modest, beginning to rise from mean chance levels ~100 ms before the instruction onset and quickly falling back thereafter. At times outside of this brief peak, however, the instantaneous subspace was no more similar to that at the time of instruction onset than could be expected from chance alone.

In contrast, classification accuracy of the Go trajectory segments was elevated above mean chance levels for more of the RGM trial duration. Though exceeding 3 standard deviations from mean chance only late in the delay epoch, Go classification accuracy rose

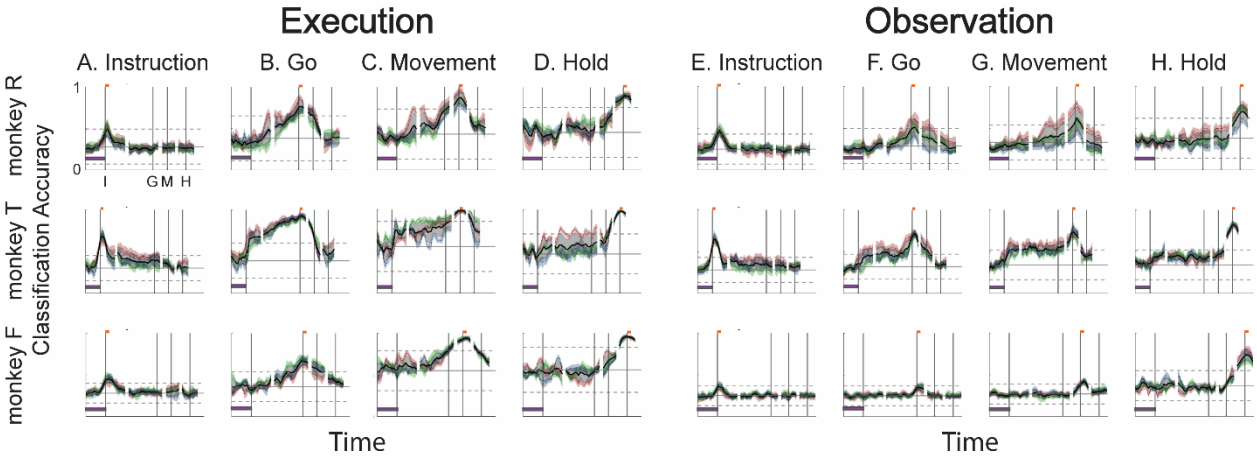

Figure 7. Decodable information as a function of time in PM MN populations. Classification accuracy for execution trajectory segments projected into instantaneous execution subspaces as a function of time is shown on the left, for observation trajectory segments projected into instantaneous observation subspaces on the right. **A,E.** Instruction trajectory segments. **B,F.** Go segments. **C,G.** Movement segments. **D,H.** Hold segments. In each frame, the short horizontal orange flag at the top of the vertical lines indicate the 100 ms during which each set of trajectory segments was clipped; the horizontal purple bar at lower left represents 500 ms. Results in 50 ms steps have been aligned separately at the times of the instruction onset (I), go cue (G), movement onset (M), and hold (H), each indicated by a vertical line as labeled in the frame at upper left. Solid curves indicate mean classification accuracy of 10-fold cross-validation as a function of time, with the shaded areas indicating 1 standard deviation. Colors red, green, and blue represent sessions 1, 2, and 3 from each monkey, with black being their average. Horizontal black lines indicate the mean (solid) ± 3 standard deviations (dashed) classification accuracy obtained by projecting each set of trajectory segments into 500 randomly selected 3D spaces.

steadily through the delay epoch, peaked near the go cue, then fell back to near mean chance levels during the reaction time (G to M) and movement (M to H) epochs. In monkeys R and F this rise began during the 500 ms instruction epoch, but only exceeded chance upper bound shortly before the Go cue. In monkey T, however, classification accuracy of the Go segments (and Movement segments) increased abruptly after the instruction onset, exceeding the chance upper bound, and then continuing to increase through the instruction and delay epochs. This abrupt increase reflects an abrupt shift of the instantaneous subspace (i.e. decrease in principal angle, Figure 4) toward that which eventually would be present at the time of the go cue (and time of movement onset) in monkey T. Such an abrupt shift was not present in monkeys R or F.

The rise in classification accuracy of the Movement trajectory segments (Figure 7C) began, not after the go cue, but earlier during the delay epoch (I-G). Although the classification accuracy of the Movement trajectory segments did not exceed chance until after the Go cue, by the time of the Go cue, Movement trajectory classification accuracy already was approaching its peak. Had the condition-dependent instantaneous subspaces during the delay epoch been orthogonal to those at the time of movement onset, the Movement trajectories would have had no projection in delay epoch subspaces and classification accuracy would have remained at baseline. The progressive increase in classification accuracy of Movement trajectory segments during the preparatory delay epoch indicates that rather than suddenly changing direction between the delay epoch and the reaction and movement epochs, as the delay epoch proceeded the condition-dependent neural trajectories of PM MNs thus gradually shifted toward where they would be at movement onset.

Classification accuracy of the Hold trajectory segments increased relatively late in execution trials. During the delay and reaction epochs the instantaneous subspaces were no more similar than chance to that at the beginning of the hold epoch. Classification accuracy of the Hold trajectory segments began to increase only after movement onset (M), rising through the movement epoch, peaking near the beginning of the hold epoch and decreasing thereafter.

We performed a similar classification accuracy analysis for observation trials. For Instruction trajectory segments (Figure 7E), the brief peak of classification accuracy occurring around the time of instruction onset (I) during observation trials was quite like that found during execution trials. For the Go and Movement segments (Figure 7F,G), classification accuracy tended to be lower, and the peaks near the times of the behavioral events (G, M) tended to be comparatively short lived. Though only brief, low peaks at times G and M were present in monkey F, a gradual rise through the delay epoch was present in monkey R, and the abrupt increase again was present in monkey T. In all three monkeys, classification accuracy of the Hold trajectory segments, during observation as during execution, began to increase only after movement onset, rising through the movement epoch, peaking near the beginning of the hold epoch and decreasing thereafter (Figure 7H).

We also performed a classification accuracy analysis of the AE neuron populations, the results of which are shown in Figure 8. During execution, classification accuracy for AE populations was lower overall than that for MN populations (Figure 7), most likely because of the smaller population sizes. In monkeys R and T, similar short-lived peaks were present around the time of instruction onset (Figure 8A), followed by gradual increases in classification accuracy of the Go and Movement trajectory segments (Figure 8B,C), and then an increase in classification accuracy of the Hold segment that began around movement onset (M). The additional abrupt increase found in monkey T's MN population was not present in the AE population, however. And for monkey F, classification accuracy remained at chance levels, presumably because of the small numbers of AE neurons recorded in this monkey (7, 10, and 9, in sessions 1, 2, and 3, respectively). During observation (Figure 8E-H), classification accuracy for AE populations again was lower overall than that for MN populations (Figure 7E-H). Only low-amplitude, short-lived peaks were present around times I, G, M, and H for the AE Instruction, Go, Movement, and Hold trajectory segments, respectively. Given that individual AE neurons showed no statistically significant modulation during observation trials, even these small peaks might not have been expected; however, previous studies have indicated that neurons not individually related to task events nevertheless may contribute to a population

response (Shenoy et al., 2013; Cunningham and Yu, 2014; Gallego et al., 2017; Jiang et al., 2020).

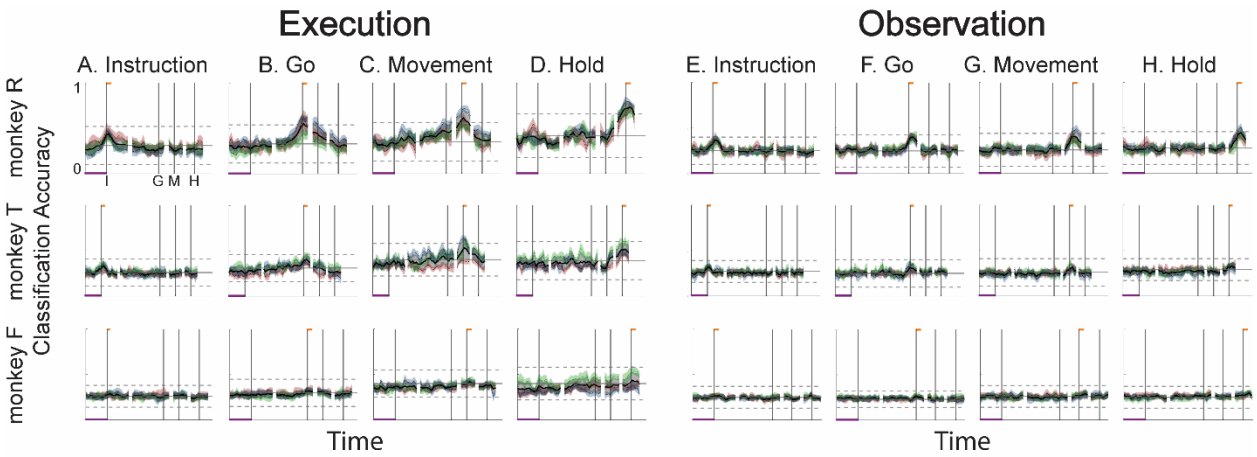

Figure 8. Decodable information as a function of time in PM AE populations. Classification accuracy for execution trajectory segments projected into instantaneous execution subspaces as a function of time is shown on the left, for observation trajectory segments projected into instantaneous observation subspaces on the right**. A,E.** Instruction trajectory segments. **B,F.** Go segments. **C,G.** Movement segments. **D,H.** Hold segments. In each frame, the short horizontal orange flag at the top of the vertical lines indicate the 100 ms during which each set of trajectory segments was clipped; the horizontal purple bar at lower left represents 500 ms. Results in 50 ms steps have been aligned separately at the times of the instruction onset (I), go cue (G), movement onset (M), and hold (H), each indicated by a vertical line as labeled in the frame at upper left. Solid curves indicate mean classification accuracy of 10-fold cross-validation as a function of time, with the shaded areas indicating 1 standard deviation. Colors red, green, and blue represent sessions 1, 2, and 3 from each monkey, with black being their average. Horizontal black lines indicate the mean (solid) ± 3 standard deviations (dashed) classification accuracy obtained by projecting each set of trajectory segments into 500 randomly selected 3D spaces.

Do PM mirror neurons progress through the same subspaces during execution and observation?

Having found that PM mirror neuron populations show similar progressive shifts in their instantaneous neural subspace during execution and observation of RGM trials, as well as similar changes in decodable information, we asked whether this progression passes through similar subspaces during execution and observation. To address this question, we first calculated the principal angles between the instantaneous execution subspace at times I, G, M, or H and the time series of instantaneous observation subspaces. Conversely, we calculated the principal angles between the instantaneous observation subspaces at times I, G, M, or H and the time series of instantaneous execution subspaces. As shown by relatively high and constant principal angles, in neither monkey R nor monkey F did the instantaneous observation subspaces shift toward the I, G, M, or H execution subspace (Figure 9A-D), nor did the instantaneous execution subspaces shift toward the I, G, M, or H observation subspace (Figure

9E-H). Throughout the time course of trials, Monkeys R and F thus showed little overlap of instantaneous execution and observation condition-dependent subspaces.

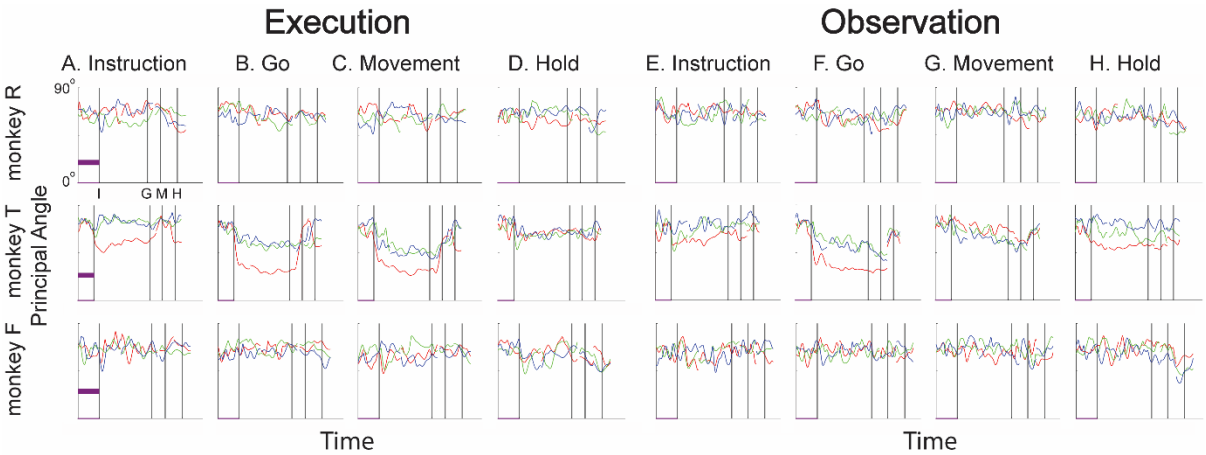

Figure 9. Principal angles cross-calculated between instantaneous execution and observation subspaces of PM MNs as a function of time. First principal angles between the time series of instantaneous observation subspaces and the I, G, M, or H instantaneous execution subspace (**A, B, C, D,** respectively) are shown on the left; between the time series of instantaneous execution subspaces and the I, G, M, or H instantaneous observation subspace (**E, F, G, H,** respectively) on the right. Red, green, and blue traces represent sessions 1, 2, and 3, respectively, from each monkey. Purple bars in the left column (**A**) indicate 500 ms, which applies to the entire row for that monkey.

In monkey T, however, an abrupt drop in principal angles indicated that the instantaneous observation subspace shifted abruptly after instruction onset (I) becoming similar to the I execution subspace (Figure 9A, first principal angle ~50° in session 1, red) and even more similar to the G and M execution subspaces (Figure 5 B,C, ~30°). Then around the time of the go cue or movement onset principal angles relatively abruptly returned to baseline. Conversely, at instruction onset the instantaneous execution subspace shifted abruptly, becoming most similar to the G observation subspace (Figure 9F, ~35°). The execution and observation subspaces of PM MNs thus overlapped considerably in monkey T from the time of instruction onset to movement onset, i.e. through the delay and reaction epochs.

We also used classification accuracy to evaluate cross-projected trajectory segments. We projected the Instruction, Go, Movement, and Hold execution trajectory segments into the time series of instantaneous observation subspaces (Figure 10A-D) and projected the Instruction, Go, Movement, and Hold observation trajectory segments into the time series of execution subspaces (Figure 10E-H). In neither monkey R or monkey F did either of these cross-projections show gradual progression, abrupt changes, or peaks of classification accuracy. Nor did the classification accuracy in either cross-projection exceed that expected from chance alone. These results confirm that little if any overlap in instantaneous, condition-dependent execution and observation subspaces were present in monkey R or monkey F. In monkey T, however, abrupt changes with slight gradual progression in classification accuracy were present, at some times minimally exceeding chance, confirming overlap in monkey T.

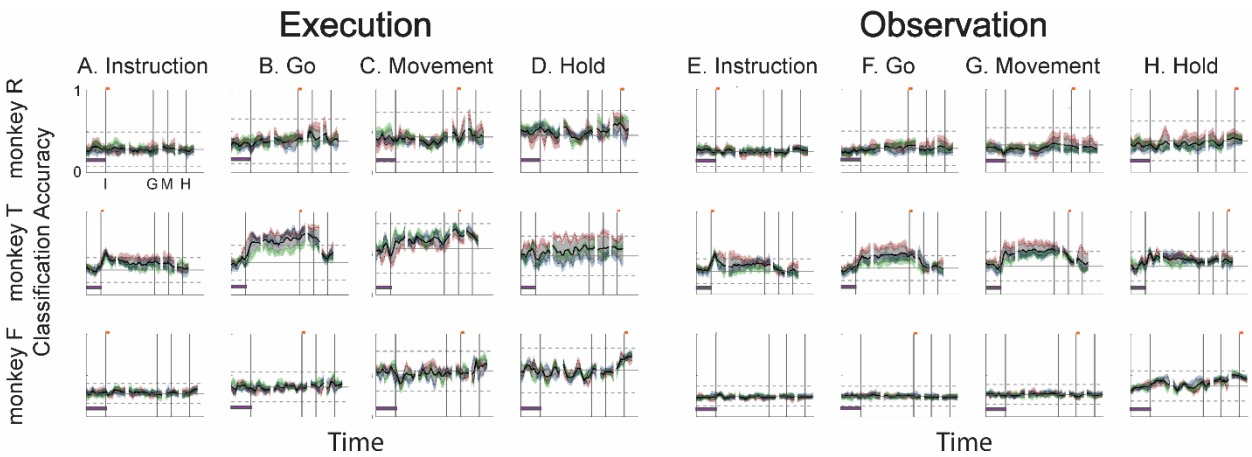

Figure 10. Classification accuracy of cross-projections between instantaneous execution and observation subspaces of PM MNs as a function of time. On the left, Instruction, Go, Movement, and Hold execution trajectory segments (**A, B, C, D,** respectively) from individual trials have projected into the time series of instantaneous observation subspaces and classified with a separate LSTM decoder at each time point; on the right, Instruction, Go, Movement, and Hold observation trajectory segments (**E, F, G, H,** respectively) have been projected into the time series instantaneous execution subspaces. Red, green, and blue traces represent sessions 1, 2, and 3, respectively, from each monkey, with black traces being the average across the three sessions for each monkey. Solid curves indicate mean classification accuracy of 10-fold cross-validation as a function of time, with the shaded areas indicating 1 standard deviation. Colors red, green, and blue represent sessions 1, 2, and 3 from each monkey, with black being their average. Horizontal black lines indicate the mean (solid) ± 3 standard deviations (dashed) classification accuracy obtained by projecting each set of trajectory segments into 500 randomly selected 3D spaces.

Alignment of latent dynamics

We next asked whether execution and observation trajectory segments, though in distinct subspaces, nevertheless could be aligned using canonical correlation. Such alignment would indicate that neural representations of trials involving the four objects bore a similar relationship to one another in neural space during execution and observation, even though they occurred in different subspaces. For example, the trajectories of PMd+M1 neuron populations recorded from two different monkeys during center-out reaching movements could be aligned well, indicating similar relationships among the neural representations of the eight movements even though the neural populations were recorded from two different monkeys (Safaie et al., 2023). In both brains the latent dynamic representation of movement to the target at 0° was closer to the representation of movement to the target at 45° than to the representation of movement to the target at 180°.

We therefore applied canonical correlation analysis (CCA, see Methods) to align the trajectories of execution trials with those of observation trials in the same recording session. An example of trial-averaged Hold trajectory segments in their original execution subspace and original observation subspace before alignment are shown in Figure 11A. The relationship among the execution trajectory segments appears substantially different than that among the observation trajectory segments. But when both sets of trajectory segments are projected into

another subspace identified with CCA, as in Figure 11B, a similar relationship among the neural representations of the four movements during execution and observation is revealed: in both behavioral contexts the neural representation of movements involving the sphere (purple) is closest to the representation of movements involving the coaxial cylinder (magenta) and farthest from that of movements involving the button (cyan).

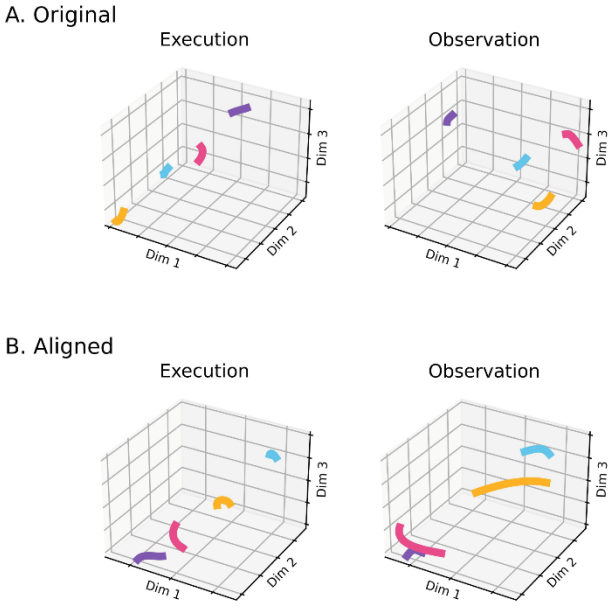

Figure 11. Alignment of execution and observation trajectory segments. **A.** Hold trajectory segments from execution trials have been projected into the original instantaneous execution subspace at time H (left), and from observation trials into the original instantaneous observation subspace also at time H (right). **B.** The same execution (left) and observation (right) trajectory segments have been projection into another subspace identified with canonical correlation. Colors indicate trajectory segments from trials involving the sphere – purple, coaxial cylinder (coax)– magenta, perpendicular cylinder (perp) – yellow, and button – cyan. Data from monkey F, session 2.

As a positive control, we first aligned MN execution trajectory segments from two different sessions in the same monkey. The 2 sessions in monkey R provided only 1 possible comparison, but the 3 sessions in monkeys T and F each provided 3 comparisons. For each of these 7 comparisons, we found the bootstrapped average CC1, average CC2, and average CC3. The mean ± standard deviation of these 7 averages has been has been plotted in Figure 12A (black) for Instruction, Go, Movement, and Hold trajectory segments. All three CCs increased progressively from Instruction, to Go, to Movement, and Hold, with the highest values being reached for the Movement trajectory segments ($\overline{CC1} = 0.89$, $\overline{CC2} = 0.77$, $\overline{CC3} = 0.61$). These values indicate consistent relationships among the Movement neural trajectory segments representing the four different RGM movements from session to session, as would have been expected from previous studies (Gallego et al., 2018; Gallego et al., 2020; Safaie et al., 2023).

Given that PM MN activity progressed largely through non-overlapping instantaneous subspaces during execution versus observation (particularly in monkeys R and F), we asked whether the relationship among the neural representations of the four RGM movements was similar during execution versus observation. To address this question, we aligned MN execution trajectory segments with MN observation trajectory segments from the same session (2 sessions from monkey R, 3 from monkey T, 3 from monkey F). The 3D mean ± standard deviation of these 8 means also has been plotted in Figure 12A (red). All three CCs again increased progressively from Instruction, to Go, to Movement, and Hold, with the highest values here being reached for the Hold trajectory segments ($\overline{CC1} = 0.73$, $\overline{CC2} = 0.54$, $\overline{CC3} = 0.39$).

Interestingly, even though considerable overlap of the execution and observation instantaneous subspaces was present in monkey T, coefficients for the Hold segments were only slightly higher when averaged across only the 3 sessions from this monkey ($\overline{CC1} = 0.79$, $\overline{CC2} = 0.65$, $\overline{CC3} = 0.50$) than across the 5 sessions from monkeys R and F ($\overline{CC1} = 0.70$, $\overline{CC2} = 0.48$, $\overline{CC3} = 0.32$). Execution/observation overlap was most marked at the time of the Go trajectory segments, however, and for Go segments coefficients averaged across the 3 sessions from monkey T ($\overline{CC1} = 0.68$, $\overline{CC2} = 0.49$, $\overline{CC3} = 0.32$) were substanially higher than

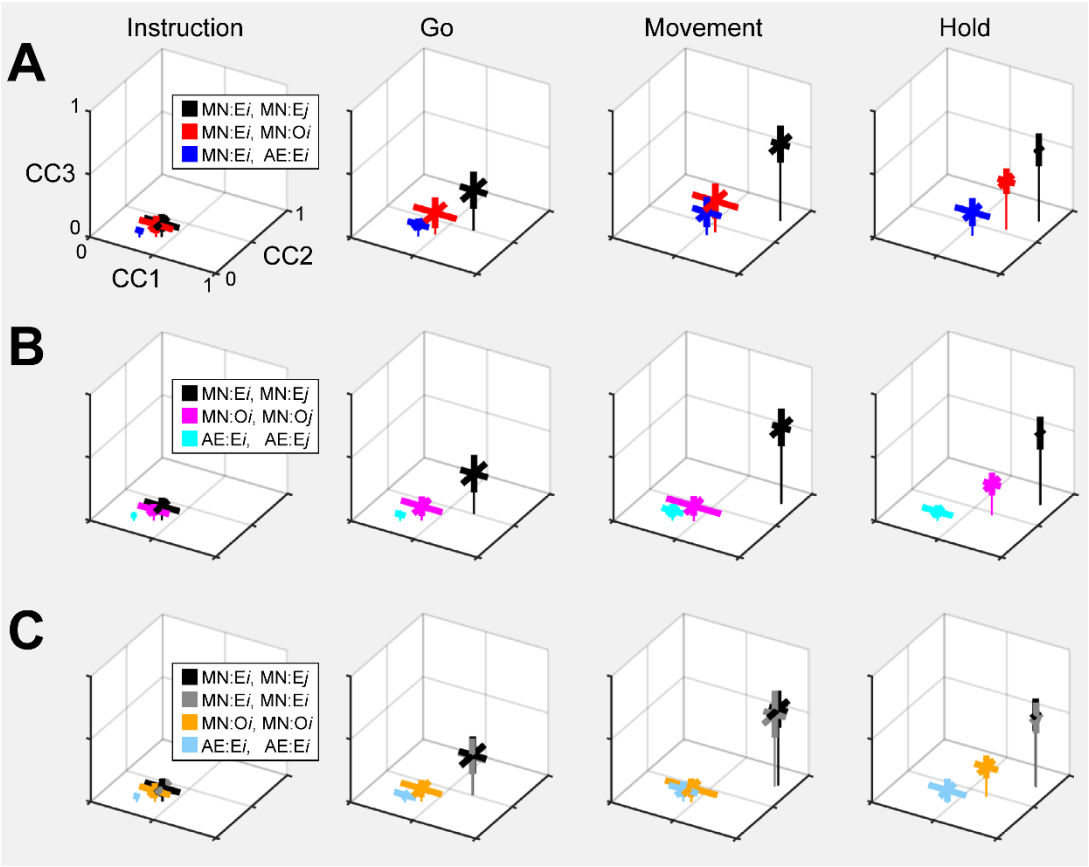

Figure 12. Alignment of trajectory segments. Canonical correlation analysis (CCA) was used to examine the alignment of Instruction, Go, Movement, and Hold trajectory segments. **A.** MN trajectory segments from execution trials recorded in two different sessions from the same monkey (black, MN:E*i*, MN:E*j*) is used as a point of reference with which to compare alignment of MN execution versus observation trials collected in the same session (red, MN:E*i*, MN:*Oi*), and MN versus AE neuron execution segments from the same session (blue, MN:E*i*, AE:E*i*). **B.** Alignment across sessions for MN observation trajectory segments (magenta, MN:O*i*, MN:*Oj*) and AE execution segments (cyan, AE:E*i*, AE:E*j*), again with MN execution segments for reference (black, MN:E*i*, MN:E*j*). **C.** Alignment within sessions for MN execution segments (gray, MN:E*i*, MN:E*i*),MN observation trajectory segments (orange, MN:O*i*, MN:*Oi*) and AE execution segments (light blue, AE:E*i*, AE:E*i*), again with MN execution segments aligned across sessions for reference (black, MN:E*i*, MN:E*j*). In each case, the three resulting correlation coefficients—CC1, CC2, and CC3—have been averaged across multiple comparisons. Thick bars representing the standard deviations of the three coefficients cross at their means, with a thin line dropped vertically from that point to the CC1 vs CC2 plane. See text for further description. (*i* and *j* indicate different sessions from the same monkey.)

when averaged across the 5 sessions from monkeys R and F ($\overline{CC1} = 0.36$, $\overline{CC2} = 0.22$, $\overline{CC3} = 0.09$). Though not as high as for execution/execution alignment, these findings indicate substantial alignment of MN trajectory segments during execution and observation. PM MNs thus showed some similarity in the relationships among their representations of the four RGM movements during execution and observation, particularly at the time of the hold.

During action execution, mirror neurons generally are considered to function in parallel with neurons active only during action execution. We therefore expected strong alignment between MN and AE neuron execution trajectory segments from the same session. Figure 12A (blue) shows the mean CCs between MN and AE execution trajectory segments across 8 comparisons (2 R, 3 T, 3 F). Though again increasing from Instruction to Movement and Hold, all three CCs remained comparatively low, reaching the highest values for the Hold segments ($\overline{CC1} = 0.57$, $\overline{CC2} = 0.35$, $\overline{CC3} = 0.19$). AE neurons thus show relationships among their representations of the four RGM movements less similar to those of MNs during execution than do MNs during execution vs. observation.

Given that MNs showed relatively strong alignment of their execution trajectories between sessions, we asked whether i) MNs also showed strong alignment of their observation trajectories between sessions, and ii) whether AE neurons showed strong alignment of their execution trajectories between sessions (1 R, 3 T, and 3 F comparisons). Neither was the case, however. Between session alignment of MN observation trajectories (Figure 12B, magenta, Hold: $\overline{CC1} = 0.66$, $\overline{CC2} = 0.46$, $\overline{CC3} = 0.25$) was somewhat lower than alignment of MN execution and observation trajectories in the same session (Figure 12A, red). And between session alignment of AE execution trajectories (Figure 12B, cyan, Hold: $\overline{CC1} = 0.37$, $\overline{CC2} = 0.21$, $\overline{CC3} = 0.08$) was substantially lower than alignment of AE and MN execution trajectories in the same session (Figure 12A, blue).

The bootstrapping approach we used for CCA (see Methods) enabled us to evaluate the consistency of relationships among trajectory segments across repeated samplings of trials recorded from the same population in the same session. We therefore performed 500 iterations of CCA between two different random samples of MN execution, MN observation, or AE execution trajectory segments from a given session (2 R, 3 T, 3 F). Within-session alignment of MN execution trajectory segments (Figure 12C, gray, Movement: $\overline{CC1} = 0.87$, $\overline{CC2} = 0.75$, $\overline{CC3} = 0.58$), was as strong as between session alignment (Figure 12C, black). But within-session alignment of MN observation trajectory segments (Figure 12C, orange, Hold: $\overline{CC1} = 0.65$, $\overline{CC2} = 0.46$, $\overline{CC3} = 0.24$) and AE execution trajectory segments (Figure 12C, light blue, Hold: $\overline{CC1} = 0.46$, $\overline{CC2} = 0.25$, $\overline{CC3} = 0.10$) remained low. These results suggest that whereas MN execution trajectories are quite consistent both between and within sessions, MN observation trajectories and AE execution trajectories are substantially less so.

DISCUSSION

As neurophysiological studies have advanced from examination of single neurons to neuron populations, analytic approaches have advanced from analyses of single neuron firing rates to analyses of co-modulation patterns among neuron populations. The co-modulation in a neuronal population can be expressed as the trajectory of the simultaneous firing rates of the *N* neurons through their *N*-dimensional state space, and the predominant patterns of co-modulation can be extracted by projecting this high-dimensional trajectory into a low-dimensional space that captures a large proportion of the population's firing-rate variance. Compared to reaching movements, however, the low-dimensional trajectories of neuronal activity controlling hand movements are relatively complex (Rouse and Schieber, 2018; Suresh et al., 2020). To approach this problem, rather than examining neural trajectories in subspaces that capture their entire time course, we identified time series of instantaneous, condition-dependent subspaces that spanned the time course of behavioral trials.

Using this approach, we found that the instantaneous, condition-dependent subspace of PM MN populations shifts progressively during both execution and observation of RGM trials. The instantaneous subspace of AE neuron populations likewise shifts progressively during action execution. This progressive shifting of the instantaneous subspace resembles that found previously using fractional overlap of condition-dependent variance in M1 neuron populations performing similar movements in an RGM task without a delay epoch (Rouse and Schieber, 2018). Although the progressive shifting described here is rotation in the mathematical sense, it is not necessarily a smooth rotation in a few dimensions. We therefore have used the word "shift" to contrast with the smooth rotation of neural trajectories in a low-dimensional subspace described in other studies, particularly those using jPCA (Churchland et al., 2012; Russo et al., 2020; Rouse et al., 2022).

Features of the instantaneous subspace

Short bursts of "signal" related discharge are known to occur in a substantial fraction of PMd neurons beginning at latencies of ~60 ms following an instructional stimulus (Weinrich et al., 1984; Cisek and Kalaska, 2004). Here we found that the instantaneous subspace shifted briefly toward that present at the time of instruction onset (I), similarly during execution and observation. This brief trough in principal angle (Figure 4A) and the corresponding peak in classification accuracy (Figure 7A) in part may reflect smoothing of firing rates with a 50 ms Gaussian kernel. We speculate, however, that the early rise of this peak at the time of instruction onset also reflects the anticipatory activity often seen in PMd neurons in expectation of an instruction, which may not be entirely non-specific, but rather may position the neural population to receive one of the limited set of potential instructions (Mauritz and Wise, 1986). The relatively low amplitude of peak classification accuracy for Instruction trajectory segments we attribute to the likely possibility that only the last 40 ms of our 100 ms Instruction segments captured signal related discharge.

The firing rates of MNs in both PMv and PMd have been shown previously to modulate during preparatory delay periods (Cisek and Kalaska, 2004; Maranesi et al., 2014). During execution of a reaching task, condition-dependent subspaces during the preparatory delay are orthogonal to those found during the subsequent movement epochs (Kaufman et al., 2014; Elsayed et al., 2016). Studies that have identified such orthogonal subspaces specifically

optimized preparatory and movement subspaces to be orthogonal to one another, however, whereas the present approach did not. Here, we found that during the preparatory delay of the present RGM task, the condition-dependent, instantaneous subspace did not remain orthogonal to that which would be present at movement onset or during the movement epoch. Rather, as the preparatory delay proceeded, the instantaneous subspace shifted concurrently toward both the subspace that would be present at the time of the go cue ending the preparatory delay (G) and that which would be present at movement onset (M). By time G, the instantaneous subspace already had shifted approximately halfway toward the time M subspace. This difference in the orthogonality of preparatory versus movement subspaces may reflect differences in reaching without grasping, which involves coordinated motion in 4 degrees of freedom (DOFs) at the shoulder and elbow, versus the present RGM movements, which involve simultaneous, fluidly coordinated motion in at least 22 DOFs of the shoulder, elbow, wrist, and digits (Rouse and Schieber, 2015). Finally, we note that the progressive shift toward the subspace present at the onset of the final hold (H) did begin only after the delay period had ended (G) and around the time of movement onset (M).

PM MN populations during execution versus observation.

In general, instantaneous execution subspaces were distinct from instantaneous observation subspaces, indicated by large principal angles between them (Figure 9) and by low classification accuracy when execution trajectories were cross-projected into observation subspaces and vice versa (Figure 10). This was the case not only during corresponding time points in execution and observation trials, but throughout their entire time course. Moreover, in all three monkeys, progressive shifting of the instantaneous, condition-dependent subspace was absent both in the principal angles between execution and observation subspaces and in the decoding of execution trajectory segments cross-projected into observation subspaces (and vice versa). These findings indicate that the predominant modes of co-modulation among PM MNs that gradually shift as behavioral trials proceed are largely distinct during execution and observation.

Although mirror neurons originally were thought to provide highly congruent neural representations of action execution and action observation (Gallese et al., 1996; Rizzolatti et al., 1996), the present findings are consistent with recent studies that have emphasized the considerable fraction of neurons with non-congruent activity, as well as differences in neural population activity during action execution versus action observation (Jiang et al., 2020; Pomper et al., 2023). As more situations have been investigated, the number of conditions needed to define a “true” mirror neuron in the strict sense of being entirely congruent has grown, making the duration of such congruence brief and/or its likelihood comparable to chance (Papadourakis and Raos, 2019; Pomper et al., 2023).

In one of our monkeys (monkey T), however, abrupt drops in principal angles and abrupt increases in cross-projected classification accuracy indicated overlap between execution and observation instantaneous subspaces that began at instruction onset, persisted through the delay epoch, and then over the reaction and movement epochs returned to a lack of overlap during the hold. This additional mode of co-modulation among PM-MNs did not shift progressively. Instead, it maintained a comparatively fixed orientation that partially overlapped with the instantaneous subspaces present at both times G and M. This overlapping subspace that appeared and ended abruptly may correspond to the shared subspace found in other

661 studies that identified orthogonal execution and observation subspaces as well (Jiang et al.,
662 2020). Of course, we do not know the source of this inter-individual difference between monkey
663 T and monkeys R and F. It might reflect either a sampling difference in the particular neurons
664 recorded (even though similar microelectrode arrays were placed in similar cortical locations) or
665 a covert difference in monkey T's behavior.

666 **Alignment of trajectory segments with canonical correlation**

667 Rather than aligning entire neural trajectories, we applied canonical correlation to
668 trajectory segments clipped for 100 ms following four well-defined behavioral events: Instruction
669 onset, Go cue, Movement onset, and the beginning of the final Hold. In all cases, alignment
670 was poorest for Instruction segments, somewhat higher for Go segments, and strongest for
671 Movement and Hold segments. Although the reasons for this sequential increase in the
672 strength of alignment are not entirely clear, we speculate that it reflects the progressive increase
673 in the separation of the condition-dependent neural trajectory segments for trials involving the
674 four different objects (Figure 6) relative to the trial-by-trial noise variance in the trajectory
675 segments.

676 Corresponding neural representations of action execution and observation during task
677 epochs with higher neural firing rates have been described previously in PMd MNs and in PMv
678 MNs using representational similarity analysis RSA (Papadourakis and Raos, 2019). During
679 force production in eight different directions, neural trajectories of PMd neurons draw similar
680 "clocks" during execution, cooperative execution, and passive observation (Pezzulo et al.,
681 2022). Likewise in the present study, despite the difference between instantaneous execution
682 and observation subspaces, in all three monkeys execution and observation trajectory
683 segments showed some degree of alignment, particularly the Movement and Hold segments
684 (Figure 12A), indicating similar relationships among the latent dynamic representations of the
685 four RGM movements during execution and observation. Alignment between trajectory
686 segments of the same PM MN population during execution versus observation in the same
687 session, however, was substantially less than that found between MN execution segments from
688 two different sessions in the same monkey. In part, this may reflect the lower firing rates of PM
689 MNs typically found during observation as compared to execution trials (Ferroni et al., 2021).
690 Alternatively, the lower alignment may reflect a certain degree of trial-by-trial variability in MN
691 observation segments, as indicated by the limited alignment of MN observation trajectory
692 segments with themselves both between and within sessions (Figure 12B,C).

693 Based on the assumption that AE neurons and MNs function as a homogenous neuron
694 population during action execution, we had expected AE and MN execution trajectory segments
695 to align closely. During execution trials, the progression of instantaneous condition-dependent
696 subspaces and of classification accuracy in AE populations was quite similar to that in MN
697 populations. We were surprised to find, therefore, that alignment between execution trajectory
698 segment from AE populations and the simultaneously recorded MN populations was even lower
699 than alignment between MN execution and observation segments (Figure 12A). Moreover,
700 whereas alignment of MN execution trajectory segments both between and within sessions was
701 high, alignment of AE neuron execution trajectory segments both between and within sessions
702 was comparatively low (Figure 12B,C). These findings indicate that whereas the predominant
703 patterns of co-modulation among MNs during execution are quite consistent both between and
704 within sessions, the patterns of co-modulation among AE neurons are more variable. Together

705 with our previous finding that modulation of MNs leads that of non-mirror neurons in time, both
706 at the single neuron level and at the population level (Mazurek and Schieber, 2019), this
707 difference in consistency versus variability leads us to speculate that during action execution,
708 while MNs carry a consistent forward model of the intended movement, AE neurons carry more
709 variable feedback information.

710 **The role of mirror neuron populations**

711 Although we did not track extraocular movements, video monitoring demonstrated that
712 our monkeys remained attentive throughout the blocks of observation trials, actively scanning
713 the visual environment. Though perhaps not following the experimenter's movements closely
714 with eye movements, or even with covert visual attention, the present results in and of
715 themselves demonstrate that during observation trials the PM MN population was processing
716 information on the sequential epochs of the behavioral task (Mazurek et al., 2018) as well as the
717 object to which the experimenter's actions were directed on each trial. These findings are
718 consistent with the notion that the PM MN population predictively represents the sequence of
719 behavioral events during observation trials (Kilner et al., 2007; Maranesi et al., 2014; Ferroni et
720 al., 2021). Our finding that MN observation trajectory segments across repeated samplings
721 within sessions showed limited alignment compared to MN execution segments (Figure 12C),
722 however, indicates more trial-by-trial variability of MN co-modulation during observation than
723 during execution. In addition to any consistent, predictive, forward model of the observed
724 experimenter's expected performance, MNs may also receive visual input that incorporates
725 more variable, trial-by-trial deviation from the predicted performance being observed.

726 One classic interpretation of similar latent dynamics in the PM MN population during
727 execution and observation would be that this similarity provides a means for the brain to
728 recognize similar movements performed by the monkey during execution and by the
729 experimenter during observation. Through some as yet unknown process, perhaps akin to a
730 communication subspace (Semedo et al., 2019) and possibly involving brain regions beyond
731 PM, the brain is able to recognize the correspondence between the latent dynamics of the
732 executed and observed actions.

733 Alternatively, given that observation of another individual can be considered a form of
734 social interaction, PM MN population activity during action observation, rather than representing
735 movements made by another individual similar to one's own, instead may represent other
736 movements one might execute in response to those made by another individual (Ninomiya et
737 al., 2020; Bonini et al., 2022; Ferrucci et al., 2022; Pomper et al., 2023). Though neurons active
738 only during observation of others (AO units) have been hypothesized to drive observation
739 activity in MNs, the present AO populations were too small to analyze with the approaches we
740 applied here. Nevertheless, the similar relative organization of the execution and observation
741 population activity in PM MNs revealed by alignment of their latent dynamics through CCA could
742 constitute a mapping of particular movements that might be made by the subject in response to
743 particular movements made by the other individual, responses which would not necessarily be
744 motorically similar to the movements observed.

745 The present analyses as well as others have focused on the condition-dependent
746 variance in MN population activity (Jiang et al., 2020). Other studies that have not separated
747 the condition-dependent versus condition-independent variance in neural activity have
748 described perhaps even more similar latent dynamics during execution and observation

(Mazurek et al., 2018; Jerjian et al., 2020; Pezzulo et al., 2022). We speculate that while condition-dependent activity may represent a particular type of movement (e.g. sphere, button, coaxial cylinder, or perpendicular cylinder) within a class of action (e.g. grasping) in a manner that differs depending on the actor, condition-independent activity may provide a neural representation of a class of action (e.g. grasping, dropping, shaking, or throwing) independent of the actor.

METHODS

Three Rhesus monkeys, R, T, and F (a 6 kg female, a 10 kg male, and an 11 kg male, *Macaca mulatta*) were used in the present study. All procedures for the care and use of these non-human 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, Rochester, New York.

Execution trials

Each monkey was trained to perform a delayed-response Reach-Grasp-Manipulate (RGM) task (Figure 2). Prior to each trial a ring of blue LEDs was illuminated around the pole supporting a center object and a 4 kHz tone began, both signaling the end of an inter-trial interval and the opportunity to begin a new trial. The monkey initiated the following sequence by pulling the center object for an initial hold epoch of variable duration (500-1000 ms). A ring of blue LEDs around the pole supporting one of four peripheral objects then was illuminated instructing the monkey as to the target object for the current trial. After 500 ms these instruction LEDs were extinguished, and the monkey was required to wait for a preparatory delay epoch lasting 500-2000 ms. At the end of this preparatory delay epoch, the blue LEDs for the center object were extinguished and the 4 kHz tone ceased, providing a go cue. The monkey then reached to, grasped, and manipulated the remembered target object: turning a sphere, pushing a button, pulling a coaxial cylinder (coax), or pulling a perpendicular cylinder (perp). The reach, grasp, manipulate sequence was performed as a single, uninterrupted, fluid movement of the entire upper extremity (Rouse and Schieber, 2015, 2016b, a). Once the instructed object had been manipulated, a ring of green LEDs around the object illuminated (indicating successful manipulation of the object) and the ring of blue LEDs for that object also illuminated (indicating correct object). The monkey then was required to hold the instructed object in its manipulated position for a final hold epoch of 1000 ms, after which the blue LEDs were extinguished. (The green LEDs extinguished whenever the monkey released the object.) After a 300 ms delay, the monkey received a liquid reward on each successful trial.

The selection and sequence of target objects in successive trials was controlled by custom software (Unified Task Control System, Gil Rivlis), which also 1) generated behavioral event marker codes (Figure 2B), and 2) arranged trials involving the four different objects in a pseudorandom block design. The behavioral event marker codes indicated the times at which specific behavioral events occurred: Start of trial, Instruction onset, Instruction offset, Go cue (delay epoch ended), Movement onset, Hold began, Hold ended, End of trial. One trial involving each of the four different objects was presented sequentially in a block. Once a block had been completed, the sequence of the four objects was shuffled randomly for the next block. To prevent the monkey from skipping more difficult objects, if the monkey failed to complete a trial successfully the same target was repeated until the monkey succeeded.

Observation trials

In a separate block of trials, the monkey observed an experimenter performing the same delayed-response RGM task. The experimenter occasionally made errors intentionally. The monkey received a reward each time the experimenter performed a successful trial, but not when the experimenter made an error, which kept the monkey attentive to the experimenter's performance. Although extraocular movements were not recorded or controlled, video

monitoring verified that the monkey remained alert and attentive throughout blocks of observation trials.

Neuron Recording

Although many studies have focused on neurons from either PMv or PMd, we combined units from these two cortical areas because neurons in each area have been shown to be modulated during both reaching and grasping (Stark et al., 2007) and during both execution and observation (Papadourakis and Raos, 2019). The three monkeys each were implanted with Floating Microelectrode Arrays (FMAs, Microprobes for Life Sciences), in the ventral premotor cortex (PMv) and in the dorsal premotor cortex (PMd). In monkeys R and T, 16-channel FMAs were implanted; in monkey F, 32-channel FMAs were used. Monkeys R and F each had a total of 64 recording electrodes implanted in PMd and 64 in PMv, whereas monkey T had 64 in PMd, but only 48 in PMv. Broadband signals were recorded simultaneously from all 128 electrodes using a Nomad/Trellis data acquisition system (Ripple, Salt Lake City, UT), which also recorded the behavioral event marker codes generated by the behavioral control system. In each recording session, data were collected during similar numbers of successful trials involving each target object during execution and then during observation, as summarized in Table 2. Off-line, spike waveforms were extracted and sorted using custom software. Sorted units were classified as definite single units, probably single units, multi-units, or noise based on their signal-to-noise ratio and estimated false-positive fraction using previously published criteria (Rouse and Schieber, 2016).

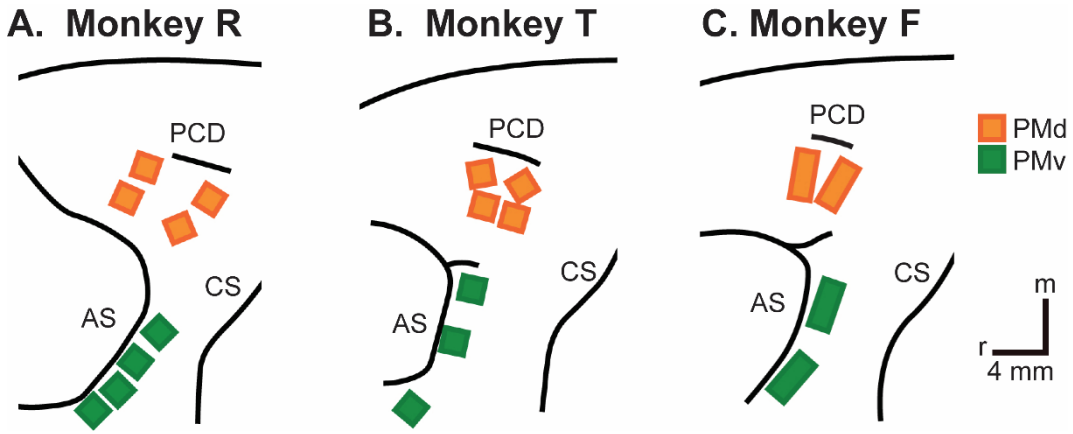

Figure 13. Recording array locations in **A.** Monkey R. **B.** Monkey T. **C.** Monkey F. PCD – precentral dimple; AS – arcuate sulcus; CS – central sulcus; r – rostral; m – medial. Scale bars apply to all three monkeys.

Mirror Neuron Identification

Each definite single unit, probable single unit or multi-unit was tested for task-related modulation. Because a given neuron's firing rates during execution and observation trials almost always differed (Ferroni et al., 2021; Pomper et al., 2023), we tested each unit for modulation using data from these two contexts separately. Spike counts from each successful behavioral trial were extracted during eleven 200 ms periods: i) before instruction onset, ii) after instruction onset, iii) before instruction offset, iv) after instruction offset (delay began), v) before delay ended, vi) after delay ended (reaction time began), vii) before movement onset, viii) after

movement onset (movement time began), ix) before movement ended, x) after movement ended (final hold began), xi) before hold ended. We then conducted two-way rmANOVA on these spike counts using object and time period as factors. We considered a unit task-related if it showed a significant main effect of either i) object or ii) time period, or a significant iii) interaction effect. Any unit modulated significantly both during execution and during observation was considered to be a mirror neuron (MN). Because each unit thus had six opportunities to show significance, we used a corrected significance criterion of $p<0.0083$ ($<0.05/6$). Any unit modulated during execution but not during observation was considered an action execution (AE) neuron. Any unit modulated during action observation but not during execution was considered an action observation neuron (AO). Units unmodulated during both execution and observation were considered not significantly (NS) related to the task.

Data analysis

Spike times for each neuron were binned (bin width = 1 ms), smoothed with a Gaussian kernel (σ = 50 ms) and square-root transformed to render variance similar from low to high firing rates (Kihlberg et al., 1972; Snedecor and Cochran, 1980). The activity of each neuron was time-aligned to four behavioral events and truncated before and after using the median delay, reaction, and movement times per object and per session as follows: i) instruction onset (I)—500 ms before, 500 ms after; ii) go cue (G)—median delay duration before, half the median reaction time after; iii) movement onset (M)—half the median reaction time before, 200 ms after; and iv) start of final hold (H)—200 ms before, 200 ms after. These four snippets of neural activity were concatenated for each trial. Neural activity then was stored as a three-dimensional tensor (N x K x T, where N is number of neurons, K the number of trials, and T the number of time points) for each of the four target objects.

Instantaneous subspace identification

Instantaneous neural subspaces were identified at 1 ms intervals. At each 1 ms time step, the N-dimensional neural firing rates from trials involving the four different objects—sphere, button, coaxial cylinder, and perpendicular cylinder—were averaged separately, providing four points in the N-dimensional space representing the average neural activity for trials involving the different objects at that time step. PCA then was performed on these four points. Because three dimensions capture all the variance of four points, three principal component dimensions fully defined each instantaneous subspace. Each instantaneous 3D subspace can be considered a filter described by a matrix, W, that can project high-dimensional neural activity into a low-dimensional subspace, with the time series of instantaneous subspaces, W_i, forming a progressively shifting series of filters (Figure 1B).

Trajectory visualization and separation

We projected 100 ms segments of neural activity into each instantaneous subspace by multiplying the neural activity, $X(t)$, by the transforming matrix for the ith subspace, W_i ,which yielded low dimensional trajectories, $L(t) = X(t)W_i \ (t \in T)$. This process was repeated for each instantaneous subspace in the time domain of interest ($t \in T$). To quantify the separation between the four trial-averaged trajectories from trials involving different objects in a given instantaneous subspace, we then calculated their cumulative separation (CS) as:

$$CS = \frac{1}{T}\sum_{t\in T} D(t) = \frac{1}{T}\sum_{t\in T}\ \sum_{i\neq j} d_{ij}(t)$$

where $d_{ij}(t)$ is the 3-dimensional Euclidean distance between the ith and j^{th} trajectories at time point t. We summed the 6 pairwise distances between the 4 trajectory segments across time points and normalized by the number of time points, $T = 100$. The larger the CS, the greater the separation of the trajectory segments.

Subspace Comparisons—Principal Angles

To assess the progressive shift of instantaneous subspaces, we computed the principal angles (Bjorck and Golub, 1973; Gallego et al., 2018) between the instantaneous subspace at each of four selected time points—onset of the instruction (I), go cue (G), onset of movement (M), and beginning of the hold (H)—and each of the other instantaneous subspaces in a time series. For example, given the 3-dimensional instantaneous subspace at the time of the of movement onset, W_M, and at any other time, W_i, we calculated their 3x3 inner product matrix and performed singular value decomposition to obtain:

$$W_M^T W_i = P_M C P_i^T$$

where 3x3 matrices P_M and P_i define new manifold directions which successively minimize the 3 principal angles specific to the two subspaces being compared. The elements of diagonal matrix C then are the ranked cosines of the principal angles, θ_i , ordered from smallest to largest:

$$C = diag(\cos(\theta_1), \cos(\theta_2), \cos(\theta_3)) .$$

We then plotted the three principal angles as a function of time, as illustrated in Figure 14A. Note that at time when $W_i = W_M$, all three principal angles are zero by definition, and the sharp decline before time M and the sharp rise afterward reflect the Gaussian kernel (σ = 50 ms) used to smooth unit firing rates. These sharp troughs thus are trivial, but both the gradual decline before and the gradual rise following the sharp troughs are not.

In this example from monkey R, session 1, the first principal angle never reached 90°. To determine whether this reflected a lack of orthogonality or a limitation of population size, we computed the first principle angle between a fixed 3-dimensional subspace, and 1000 3-dimensional subspaces randomly chosen from *N*-dimensional spaces, for *N* varying from 5 to 500. Figure 14B shows that for large *N*, principal angles between a fixed subspace and other randomly chosen subspaces are likely to be close to 90°. But as *N* decreases, these random principal angles are less likely to approach 90°, without necessarily indicating non-random overlap of the subspaces. In the Results, rather than using principal angles as a measure of subspace overlap, we therefore focus on trends in the time course of principal angles between instantaneous subspaces. In addition, given that the set of three principal angles typically followed similar time courses (Figure 14A), in the Results we illustrate only the first principal angle, θ_1 .

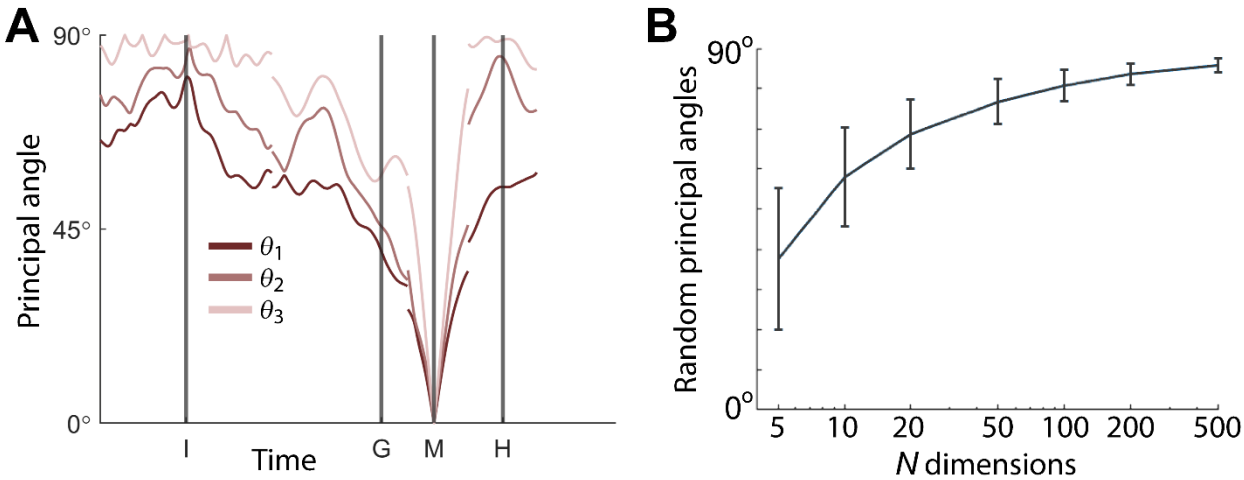

Figure 14. A. First, second, and third principal angles as a function of time. An example in which the three principal angles, $\theta_1, \theta_2, \theta_3$, between the instantaneous subspace at time M (movement onset) and the entire time series of instantaneous subspaces have been plotted as a function of time for PM MNs. (Data from monkey R, session 1.) **B.** First principal angles between a fixed 3D subspace and 1000 other 3D subspaces randomly chosen from subspaces of dimensionality varying from 5 to 500. Error bars indicate 1 standard deviation from the mean.

Decodable information—LSTM

As illustrated schematically in Figures 1B, the same segment of high-dimensional neural activity projected into different instantaneous subspaces can generate low-dimensional trajectories of varying separation. The degree of separation among the projected trajectory segments will depend, not only on their separation at the time when the segments were clipped, but also on the similarity of the subspaces into which the trajectory segments are projected. To quantify this separation, we projected high-dimensional neural trajectory segments (each including 100 points at 1 ms intervals) from successful trials involving each of the four different target objects into a time series of 3-dimensional instantaneous subspaces at 50 ms intervals. In each of these instantaneous subspaces, the neural trajectory segment from each trial thus became a 100 point x 3 dimensional matrix. For each instantaneous subspace in the time series, we then trained a separate long short-term memory (LSTM, (Hochreiter and Schmidhuber, 1997)) classifier to attribute each of the neural trajectories from individual trials to one of the four target object labels: sphere, button, coaxial cylinder, or perpendicular cylinder. Using MATLAB's Deep Learning Toolbox, each LSTM classifier had 3 inputs (instantaneous subspace dimensions), 20 hidden units in the bidirectional LSTM layer, and a softmax layer preceding the classification layer which had 4 output classes (target objects). The total number of successful trials available in each session for each object is given in Table 1. To avoid bias based on the total number of successful trials, we used the minimum number of successful trials across the four objects in each session, selecting that number from the total available randomly with replacement. Each LSTM classifier was trained with MATLAB's adaptive moment estimation (Adam) optimizer on 40% of the selected trials, and the remaining 60% were decoded by the trained classifier. The success of this decoding was used as an estimate of classification accuracy from 0 (no correct classifications) to 1 (100% correct classifications). This process was repeated 10 times and the mean ± standard deviation across the 10 folds was

reported as the classification accuracy at that time. Classification accuracy of trials projected into each instantaneous subspace at 50 ms intervals was plotted as a function of trial time.

Similarity of aligned latent dynamics

We used Canonical Correlation Alignment (CCA) to compare the similarity of latent dynamics in different subspaces (Gallego et al., 2020). In brief, given latent dynamics (trajectory segments) in two original subspaces, L_A and L_B, CCA finds a linear transformation of each original subspace such that, when projected into a common subspace, the aligned latent dynamics, $\tilde{L}_A$ and $\tilde{L}_B$, are maximumly correlated in each dimension of the common subspace. Larger canonical correlation coefficients (CCs) indicate a higher degree of similarity.

CCA was performed as follows: The original latent dynamics, L_A and L_B, first are transformed and decomposed as $L_A^T = Q_A R_A$ and $L_B^T = Q_B R_B$. The first *m = 3* column vectors of each Q_i provide an orthonormal basis for the column vectors of L_i^T (where $i = A, B$) . Singular value decomposition on the inner product matrix of Q_A and Q_B then gives $Q_A^T Q_B = USV^T$, and new manifold directions that maximize pairwise correlations are provided by $M_A = R_A^{-1}U$ and $M_B = R_B^{-1}V$. We then project the original latent dynamics into the new, common subspace: $\tilde{L}_A^T = L_A^T M_A$; $\tilde{L}_B^T = L_B^T M_B$. Pairwise correlation coefficients between the aligned latent dynamics sorted from largest to smallest then are given by the elements of the diagonal matrix $S = \tilde{L}_A \tilde{L}_B^T$.

We used a bootstrapping approach to CCA. From each of two data sets we randomly selected 20 trials involving each target object (totaling 80 trials) with replacement, clipped trajectory segments from each of those trials for 100 ms (100 points at 1 ms intervals) after the instruction onset, go cue, movement onset, or beginning of the final hold, and performed CCA as described above. With 500 iterations, we obtained a distribution of the correlation coefficients (CCs) between the two data sets in each of the three dimensions of the aligned subspace. Figure 15 illustrates such bootstrapped distributions of CCs for Instruction, Go, Movement, and Hold trajectory segments. In these 3D scatter plots, a point has been plotted at the coordinates (CC1, CC2, CC3) for only 50 (for visualization) of the 500 iterations and a green line has been drawn from the centroid (mean) of all 500 such points to the CC1 vs CC2 plane. This bootstrapping approach enabled us to assess the consistency of neural trajectories in a given neural population within a single session by drawing two separate random samples of 80 trials from the same population, which would not have been possible had we concatenated trajectory segments from all trials in the session (Gallego et al., 2020; Safaie et al., 2023). We

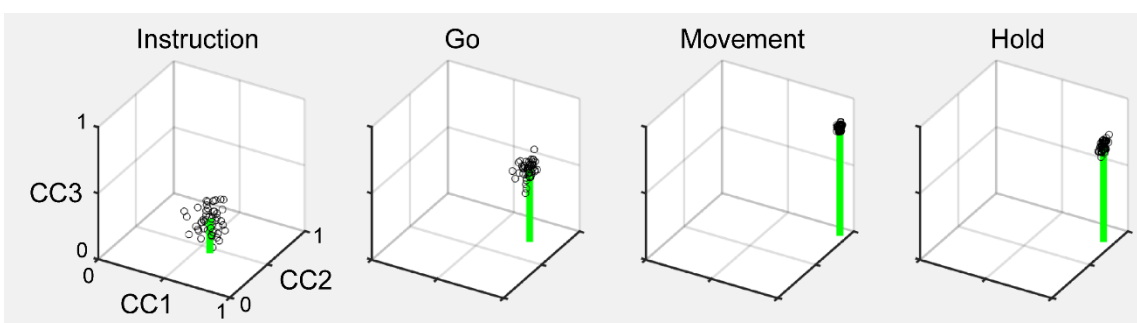

Figure 15. Bootstrapped CCA. Points representing 50 of 500 CCA iterations as described in the text have been plotted in three dimensions at the coordinates of their correlation coefficients, CC1, CC2, and CC3. A green line has been drawn from the mean of all 500 iterations to the CC1 vs CC2 plane. This example CCA aligns MN execution trajectory segments within session 3 from Monkey T.

then used the same approach to evaluate alignment of latent dynamics between different contexts (e.g. execution and observation), between different sessions (e.g. execution trials on two different days), and between different neural populations (e.g. MNs and AE neurons).

Code Availability

Code packages for all analyses performed in this work are available at: https://github.com/ShiftingSubspace/shiftsubs .

REFERNCES

Albertini D, Lanzilotto M, Maranesi M, Bonini L (2021) Largely shared neural codes for biological and nonbiological observed movements but not for executed actions in monkey premotor areas. J Neurophysiol 126:906-912.

Ames KC, Ryu SI, Shenoy KV (2014) Neural dynamics of reaching following incorrect or absent motor preparation. Neuron 81:438-451.

Bjorck A, Golub GH (1973) Numerical Methods for Computing Angles between Linear Subspaces. Math Comput 27:579-594.

Bonini L, Rotunno C, Arcuri E, Gallese V (2022) Mirror neurons 30 years later: implications and applications. Trends Cogn Sci 26:767-781.

Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV (2012) Neural population dynamics during reaching. Nature 487:51-56.

Cisek P, Kalaska JF (2004) Neural correlates of mental rehearsal in dorsal premotor cortex. Nature 431:993-996.

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

Dekleva BM, Chowdhury RH, Batista AP, Chase SM, Yu BM, Boninger ML, Collinger JL (2024) Motor cortex retains and reorients neural dynamics during motor imagery. Nat Hum Behav.

di Pellegrino G, Fadiga L, Fogassi L, Gallese V, Rizzolatti G (1992) Understanding motor events: a neurophysiological study. Experimental Brain Research 91:176-180.

Dushanova J, Donoghue J (2010) Neurons in primary motor cortex engaged during action observation. Eur J Neurosci 31:386-398.

Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP (2016) Reorganization between preparatory and movement population responses in motor cortex. Nat Commun 7:13239.

Ferroni CG, Albertini D, Lanzilotto M, Livi A, Maranesi M, Bonini L (2021) Local and system mechanisms for action execution and observation in parietal and premotor cortices. Curr Biol 31:2819-2830 e2814.

Ferrucci L, Nougaret S, Falcone R, Cirillo R, Ceccarelli F, Genovesio A (2022) Dedicated Representation of Others in the Macaque Frontal Cortex: From Action Monitoring and Prediction to Outcome Evaluation. Cereb Cortex 32:891-907.

Gallego JA, Perich MG, Miller LE, Solla SA (2017) Neural Manifolds for the Control of Movement. Neuron 94:978-984.

Gallego JA, Perich MG, Chowdhury RH, Solla SA, Miller LE (2020) Long-term stability of cortical population dynamics underlying consistent behavior. Nat Neurosci 23:260-270.

Gallego JA, Perich MG, Naufel SN, Ethier C, Solla SA, Miller LE (2018) Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. Nature Communications In Press:1-13.

Gallese V, Fadiga L, Fogassi L, Rizzolatti G (1996) Action recognition in the premotor cortex. Brain 119:593-609.

Hochreiter S, Schmidhuber J (1997) Long short-term memory. Neural Comput 9:1735-1780.

Jerjian SJ, Sahani M, Kraskov A (2020) Movement initiation and grasp representation in premotor and primary motor cortex mirror neurons. Elife 9.

Jiang X, Saggar H, Ryu SI, Shenoy KV, Kao JC (2020) Structure in Neural Activity during Observed and Executed Movements Is Shared at the Neural Population Level, Not in Single Neurons. Cell Rep 32:108006.

Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014) Cortical activity in the null space: permitting preparation without movement. Nat Neurosci 17:440-448.
Kihlberg JK, Herson JH, Schotz WE (1972) Square Root Transformation Revisited. J Roy Stat Soc C-App 21:76-&.
Kilner JM, Friston KJ, Frith CD (2007) Predictive coding: an account of the mirror neuron system. Cogn Process 8:159-166.
Kobak D, Brendel W, Constantinidis C, Feierstein CE, Kepecs A, Mainen ZF, Qi XL, Romo R, Uchida N, Machens CK (2016) Demixed principal component analysis of neural population data. Elife 5.
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.
Kraskov A, Philipp R, Waldert S, Vigneswaran G, Quallo MM, Lemon RN (2014) Corticospinal mirror neurons. Philos Trans R Soc Lond B Biol Sci 369:20130174.
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.
Mauritz KH, Wise SP (1986) Premotor cortex of the rhesus monkey: neuronal activity in anticipation of predictable environmental events. Experimental Brain Research 61:229-244.
Mazurek KA, Schieber MH (2019) Mirror neurons precede non-mirror neurons during action execution. J Neurophysiol 122:2630-2635.
Mazurek KA, Rouse AG, Schieber MH (2018) Mirror Neuron Populations Represent Sequences of Behavioral Epochs During Both Execution and Observation. J Neurosci 38:4441-4455.
Natraj N, Silversmith DB, Chang EF, Ganguly K (2022) Compartmentalized dynamics within a common multi-area mesoscale manifold represent a repertoire of human hand movements. Neuron 110:154-+.
Ninomiya T, Noritake A, Kobayashi K, Isoda M (2020) A causal role for frontal cortico-cortical coordination in social action monitoring. Nat Commun 11:5233.
Papadourakis V, Raos V (2019) Neurons in the Macaque Dorsal Premotor Cortex Respond to Execution and Observation of Actions. Cereb Cortex 29:4223-4237.
Pezzulo G, Donnarumma F, Ferrari-Toniolo S, Cisek P, Battaglia-Mayer A (2022) Shared population-level dynamics in monkey premotor cortex during solo action, joint action and action observation. Prog Neurobiol 210:102214.
Pomper JK, Shams M, Wen S, Bunjes F, Thier P (2023) Non-shared coding of observed and executed actions prevails in macaque ventral premotor mirror neurons. Elife 12.
Rizzolatti G, Craighero L (2004) The mirror-neuron system. Annual Review of Neuroscience 27:169-192.
Rizzolatti G, Fadiga L, Gallese V, Fogassi L (1996) Premotor cortex and the recognition of motor actions. Brain Research Cognitive Brain Rese:131-141.
Rouse AG, Schieber MH (2015) Spatiotemporal distribution of location and object effects in reach-to-grasp kinematics. J Neurophysiol 114:3268-3282.
Rouse AG, Schieber MH (2016a) Spatiotemporal Distribution of Location and Object Effects in Primary Motor Cortex Neurons during Reach-to-Grasp. J Neurosci 36:10640-10653.
Rouse AG, Schieber MH (2016b) Spatiotemporal distribution of location and object effects in the electromyographic activity of upper extremity muscles during reach-to-grasp. J Neurophysiol 115:3238-3248.
Rouse AG, Schieber MH (2018) Condition-Dependent Neural Dimensions Progressively Shift during Reach to Grasp. Cell Rep 25:3158-3168 e3153.
Rouse AG, Schieber MH, Sarma SV (2022) Cyclic, Condition-Independent Activity in Primary Motor Cortex Predicts Corrective Movement Behavior. eNeuro 9.

Russo AA, Khajeh R, Bittner SR, Perkins SM, Cunningham JP, Abbott LF, Churchland MM (2020) Neural Trajectories in the Supplementary Motor Area and Motor Cortex Exhibit Distinct Geometries, Compatible with Different Classes of Computation. Neuron 107:745-758 e746.

Safaie M, Chang JC, Park J, Miller LE, Dudman JT, Perich MG, Gallego JA (2023) Preserved neural dynamics across animals performing similar behaviour. Nature 623:765-771.

Semedo JD, Zandvakili A, Machens CK, Yu BM, Kohn A (2019) Cortical Areas Interact through a Communication Subspace. Neuron 102:249-259 e244.

Shenoy KV, Sahani M, Churchland MM (2013) Cortical control of arm movements: a dynamical systems perspective. Annu Rev Neurosci 36:337-359.

Snedecor GW, Cochran WG (1980) Statistical Methods, 7 Edition. Ames: Iowa State University Press.

Sobinov AR, Bensmaia SJ (2021) The neural mechanisms of manual dexterity. Nat Rev Neurosci 22:741-757.

Stark E, Asher I, Abeles M (2007) Encoding of reach and grasp by single neurons in premotor cortex is independent of recording site. J Neurophysiol 97:3351-3364.

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

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.

Vyas S, Golub MD, Sussillo D, Shenoy KV (2020) Computation Through Neural Population Dynamics. Annual Review of Neuroscience, Vol 43 43:249-275.

Weinrich M, Wise SP, Mauritz KH (1984) A neurophysiological study of the premotor cortex in the rhesus monkey. Brain 107:385-414.