Cell Reports

Article

Structure in Neural Activity during Observed and Executed Movements Is Shared at the Neural Population Level, Not in Single Neurons

Graphical Abstract

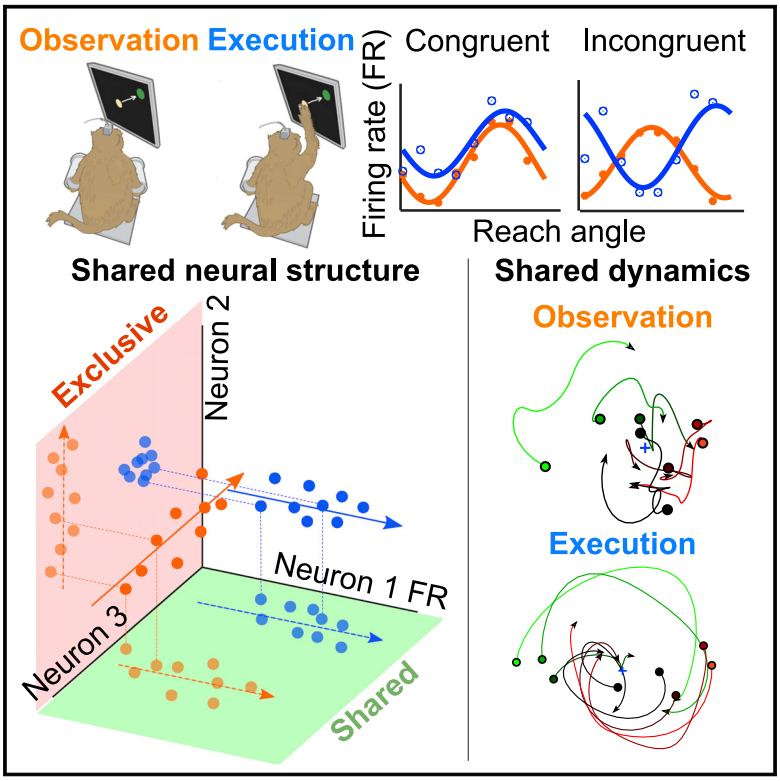

Authors

Xiyuan Jiang, Hemant Saggar, Stephen I. Ryu, Krishna V. Shenoy, Jonathan C. Kao

Correspondence

kao@seas.ucla.edu

In Brief

We can understand actions and learn movements by observing them. These processes are thought to involve neurons in motor cortex that are active during both movement observation and execution. Jiang et al. find that, when observing and executing movements, motor cortical neurons exhibit shared structure through consistent covariation and dynamics.

Highlights

- Neural activity for observed and executed movements occupies a shared subspace
- The shared subspace is composed of both congruent and incongruent neurons
- Neural dynamics are more similar within the shared subspace than outside it
- Observation and execution also have a context-specific exclusive subspace

Jiang et al., 2020, Cell Reports *32*, 108006
August 11, 2020
https://doi.org/10.1016/j.celrep.2020.108006

CellPress

Cell Reports

CellPress
OPEN ACCESS

Article

Structure in Neural Activity during Observed and Executed Movements Is Shared at the Neural Population Level, Not in Single Neurons

Xiyuan Jiang,[1] Hemant Saggar,[1] Stephen I. Ryu,[3,8] Krishna V. Shenoy,[3,4,5,6,7] and Jonathan C. Kao[1,2,9,*]
[1]Department of Electrical and Computer Engineering, University of California, Los Angeles, Los Angeles, CA 90024, USA
[2]Neurosciences Program, University of California, Los Angeles, Los Angeles, CA 90024, USA
[3]Department of Electrical Engineering, Stanford University, Stanford, CA 94305, USA
[4]Department of Neurobiology, Stanford University, Stanford, CA 94305, USA
[5]Department of Bioengineering, Stanford University, Stanford, CA 94305, USA
[6]Wu Tsai Neurosciences Institute and Bio-X Institute, Stanford University, Stanford, CA 94305, USA
[7]Howard Hughes Medical Institute, Stanford University, Stanford, CA 94305, USA
[8]Department of Neurosurgery, Palo Alto Medical Foundation, Palo Alto, CA 94301, USA
[9]Lead Contact
*Correspondence: kao@seas.ucla.edu
https://doi.org/10.1016/j.celrep.2020.108006

SUMMARY

In multiple cortical areas, including the motor cortex, neurons have similar firing rate statistics whether we observe or execute movements. These "congruent" neurons are hypothesized to support action understanding by participating in a neural circuit consistently activated in both observed and executed movements. We examined this hypothesis by analyzing neural population structure and dynamics between observed and executed movements. We find that observed and executed movements exhibit similar neural population covariation in a shared subspace capturing significant neural variance. Further, neural dynamics are more similar between observed and executed movements within the shared subspace than outside it. Finally, we find that this shared subspace has a heterogeneous composition of congruent and incongruent neurons. Together, these results argue that similar neural covariation and dynamics between observed and executed movements do not occur via activation of a subpopulation of congruent single neurons, but through consistent temporal activation of a heterogeneous neural population.

INTRODUCTION

Even when we don't move, motor cortex is involved in processing observed movements. Intriguingly, motor cortex does not use separate groups of neurons to process observed and executed movements: neurons in the ventral premotor cortex (PMv) (Gallese et al., 1996; Rizzolatti and Craighero, 2004; Bonini et al., 2014), dorsal premotor cortex (PMd) (Tkach et al., 2007; Cisek and Kalaska, 2004; Papadourakis and Raos, 2019), and primary motor cortex (M1) (Dushanova and Donoghue, 2010; Tkach et al., 2007; Vigneswaran et al., 2013; Mazurek et al., 2018) can fire during both executed and observed movements. This has led researchers to hypothesize that motor cortical neurons, in addition to generating movements, support an understanding of observed motor actions (Rizzolatti et al., 2001; Fogassi and Gallese, 2002), predict and interpret sensory consequences of our actions (Tkach et al., 2008), and reflect mental rehearsal (Cisek and Kalaska, 2004).

But how does motor cortex use the same neural circuit to both process observed movements and generate executed movements? Prior work suggests that this may occur through a subpopulation of "congruent" neurons, the activity of which is similar across observed and executed movement contexts (Gallese et al., 1996; Rizzolatti et al., 2001; Rizzolatti and Craighero, 2004; Tkach et al., 2007, 2008; Dushanova and Donoghue, 2010; Mazurek et al., 2018), for example, in preferred direction (PD) (Tkach et al., 2007, 2008; Dushanova and Donoghue, 2010) or in action encoding (Gallese et al., 1996; Mazurek et al., 2018). This has led researchers to hypothesize that congruent neurons support an understanding of observed movements by transforming an observed action into a motor representation, effectively activating (or "resonating" [Rizzolatti et al., 2001]) the same neural subcircuits (Gallese et al., 1996; Rizzolatti et al., 2001; Fogassi and Gallese, 2002; Tkach et al., 2007). More specifically, congruent neurons may reflect mental rehearsal of the action (Cisek and Kalaska, 2004) or internal generation of motor commands that are gated by downstream circuits, but otherwise would have produced the observed movement (Tkach et al., 2007, 2008). We call this the "congruent subpopulation hypothesis." On the other hand, there are also unique attributes of observed and executed actions. For example, motor cortex must uniquely generate movement in

executed, but not observed, actions, or motor cortex may represent the movements of external agents in observed, but not executed, actions. These unique computations may be carried out by a separate population of "incongruent" neurons. Dushanova and Donoghue (2010) found 62% of recorded M1 neurons were incongruent, having statistically significant changes in PD between observed and executed movements. Congruent and incongruent subpopulations may therefore play distinct roles in the motor circuit. For example, while congruent neurons may support an understanding of action (whether observed or executed), incongruent neurons may support the execution of the movement.

An alternative hypothesis is that shared computation between observed and executed movements occurs at the level of the neural population, rather than in congruent subpopulations (Papadourakis and Raos, 2019). This hypothesis is based on the framework that motor cortex is a dynamical system where the key feature is how neurons coordinate their activity through time. This coordinated activity can be analyzed in a low-dimensional neural subspace that reflects neural population covariation (Yu et al., 2009; Shenoy et al., 2013; Elsayed et al., 2016; Gallego et al., 2017). These subspaces have led to insight into several aspects of movement, including robust neural populations dynamics during movement generation (Churchland et al., 2012; Russo et al., 2018; Pandarinath et al., 2018a), how preparatory activity in PMd and M1 does not cause movement (Kaufman et al., 2014), neural constraints on learning (Sadtler et al., 2014), and shared computation between multiple motor tasks (Gallego et al., 2018). Under this alternative hypothesis, there is a low-dimensional *shared* neural subspace between observed and executed movements, irrespective of whether neurons are congruent or incongruent. Attributes shared between observed and executed movements may be neurally represented in this shared subspace, while attributes unique to observed and executed movements may be represented outside the shared subspace (an "exclusive" subspace). Importantly, the shared subspace may reflect coordinated activity patterns across both congruent and incongruent neurons, rather than just the congruent subpopulation.

To date, the relationship between observation and execution neural population activity is still opaque. While it may be intuitive to hypothesize that observation and execution activity must share neural population structure due to the presence of congruent neurons, this is not necessarily true. First, single neuron statistics used to define congruence, such as PD (used in this work), do not capture how multiple neurons coordinate their activity through time (reviewed in Cunningham and Yu, 2014). For example, in planned and executed movements, some PMd and M1 neurons exhibit a degree of congruence, having correlated PDs (Crammond and Kalaska, 2000); however, neural subspaces supporting computation for planned and executed movements are orthogonal, not overlapping (Elsayed et al., 2016). Second, neural subspaces involved in executed movements are hypothesized to comprise "output-potent" dimensions, where activity leads to generated movements (Kaufman et al., 2014; Elsayed et al., 2016; Kaufman et al., 2016; Gallego et al., 2018). Prior work established that even though neural activity during motor planning (without overt movement) can be decoded to predict executed movements, plan activity explicitly avoids these output-potent dimensions (Kaufman et al., 2014) and occupies neural subspaces orthogonal to execution activity (Elsayed et al., 2016). While these studies analyzed preparatory versus execution activity, these principles apply also to observation versus execution activity. In this study, there are no overt movements during observation, and therefore activity must reside in a putative output-null subspace that avoids driving the muscles. It is therefore possible that observation subspaces, like planning subspaces, are orthogonal to execution subspaces.

To elucidate the relationship between neural population activity between observed and executed movements, we addressed three categories of hypotheses, as illustrated in Figure 1. First, we investigated whether neural population structure between observed and executed movements is orthogonal (exclusive subspaces) or overlapping (shared subspace), as illustrated in Figure 1A. Second, we assessed if shared (or exclusive) structure between observed and executed movements was due primarily to congruent (or incongruent) neurons versus a heterogeneous population of congruent and incongruent neurons, as illustrated in Figure 1B. If congruent neurons support action understanding through consistent activation in observed and executed movements (Cisek and Kalaska, 2004; Tkach et al., 2007, 2008), they should primarily contribute to shared structure, while incongruent neurons should contribute to exclusive structure. Alternatively, a heterogeneous composition would imply that congruent and incongruent neurons participate in both shared and exclusive structure. Finally, we quantified whether neural dynamics were more similar in the shared subspace or congruent subpopulation via "tangling." Tangling quantifies how dissimilar future neural activity patterns can be from a similar starting point (i.e., present neural activity) (Russo et al., 2018). If the shared subspace has similar dynamics, they ought to have low tangling relative to other subspaces, as illustrated in Figure 1C.

In summary, we find that observation and execution activity share a neural subspace, composed of a heterogeneous population of congruent and incongruent neurons. This rejects the congruent subpopulation hypothesis. Further, the shared subspace has similar dynamics during observed and executed movements. Our results support the hypothesis that shared computation in observed and executed movements, reflected through the similar coordination of neural population activity through time, is carried out by a heterogeneous neural population.

RESULTS

We trained two rhesus macaques (Monkeys J and L) to perform a center-out-and-back task (see Method Details). In the executed movement context (Figure 2A), the monkey controlled the cursor with his hand position. In the observed movement context (Figure 2B), the monkey observed the cursor moving from target to target while both arms were restrained. To ensure the monkey was engaged in movement observation, we monitored his eye position during observation

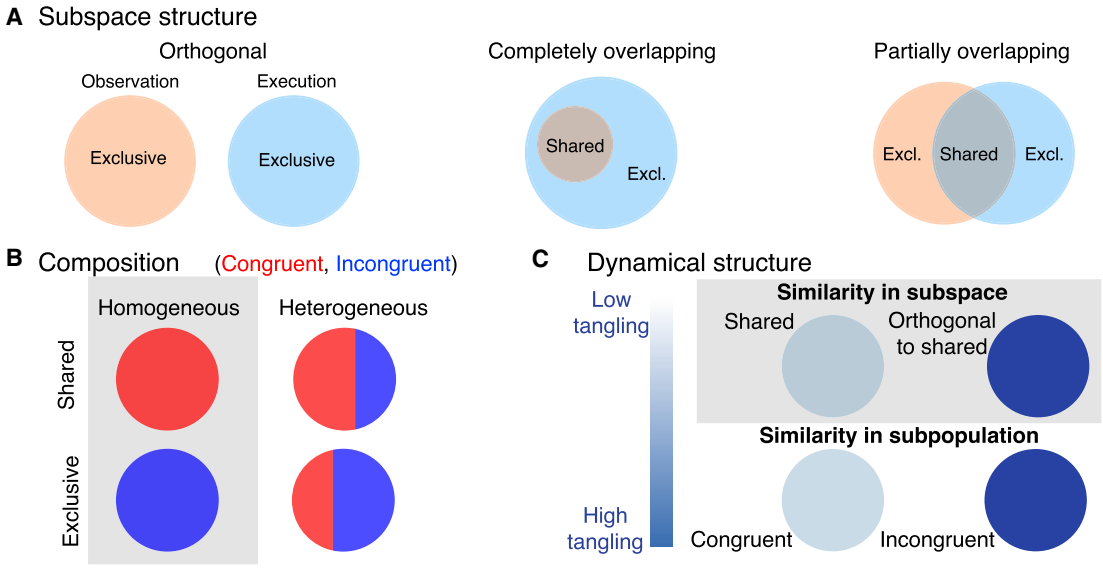

Figure 1. Hypothesis Categories

(A) Neural population subspaces between observation and execution may be orthogonal, completely overlapping, or partially overlapping with shared structure. Exclusive subspaces only capture neural variance in one context. Shared subspaces capture variance in both contexts.

(B) The shared (or exclusive) subspace may have a homogeneous composition of congruent (or incongruent) neurons, or a heterogeneous composition. Red (blue) denotes congruent (incongruent) neurons.

(C) Dynamics in observation and execution may be more similar (low tangling) in the shared subspace (top, similarity in subspace), or in the congruent subpopulation (bottom, similarity in subpopulation).

trials and only used trials where the monkey's eye position was consistent with the cursor's start and end position (see Method Details). To further validate that observation activity was modulated for the task, we also trained a brain-machine interface (BMI) from observation trials on each experimental session. In closed-loop experiments, we found that Monkey J (L) was able to control the BMI to perform a center-out-and-back task at 100% (95.8% to 99.4% across datasets) success rate, consistent with prior observation-based BMIs (Gilja et al., 2012; Stavisky et al., 2017). This indicates that recorded neural population activity during observation was modulated for cursor movement during the task.

Simultaneous to behavior, we recorded spiking activity from Utah electrode arrays. Monkey J had two Utah arrays, one implanted in PMd and the other implanted in M1, as visually estimated from anatomical landmarks. Monkey L had one Utah array, implanted at the PMd/M1 border. All analyses directly comparing observation and execution activity were done with neurons recorded in the same experimental session, so that the recorded units were putatively the same in both observation and execution contexts. We initially recorded threshold crossings, detecting multiunit spikes whenever the voltage crossed a threshold at $-4.5 \times$ the root-mean-square voltage recorded on that electrode. We manually spike sorted this activity and implemented an interspike interval (ISI) threshold (see Method Details) to determine putative single units. We recorded anywhere from 41 to 83 putative single units in each experimental session. For these analyses, we intentionally did not discard single units based on their activity or correlation to movement, as single neurons may bear no obvious relationship to movement commands, but may contribute meaningfully to a population response (Shenoy et al., 2013; Cunningham and Yu, 2014; Gallego et al., 2017). Nevertheless, when possible, we confirmed our analyses on a smaller population where neurons with a cosine tuning curve (Georgopoulos et al., 1982) ($R^2 < 0.3$) were discarded (Table S1). We also confirmed our analyses on threshold-crossing activity, as it is possible to accurately estimate neural population dynamics without subjective spike sorting (Trautmann et al., 2019). Results from these analyses are summarized in Table S1, and uphold the conclusions from analysis of the single-unit population. Finally, the same conclusions also hold when considering only the subpopulation of PMd or M1 neurons (Table S5).

We first characterized the proportion of recorded congruent neurons in our recordings. We fit a cosine tuning model for each neuron recorded during observation and execution based on the average activity (in a window 200 to 500 ms after target onset) during center-out movements (Georgopoulos et al., 1982). We defined a neuron as being incongruent if its PD had a statistically significant change between observation and execution contexts (e.g., Figure 2D) by bootstrapping (Fan et al., 2014) (see Method Details); otherwise, it was congruent (e.g., Figure 2C). PD has been previously used to determine congruence (Tkach et al., 2007, 2008; Dushanova and Donoghue, 2010), and is reasonable because it is an important summary feature of single-neuron tuning (Georgopoulos et al., 1982), even comprising the basis for several BMI decoder algorithms (Taylor et al., 2002; Velliste et al., 2008; Kao et al., 2014).

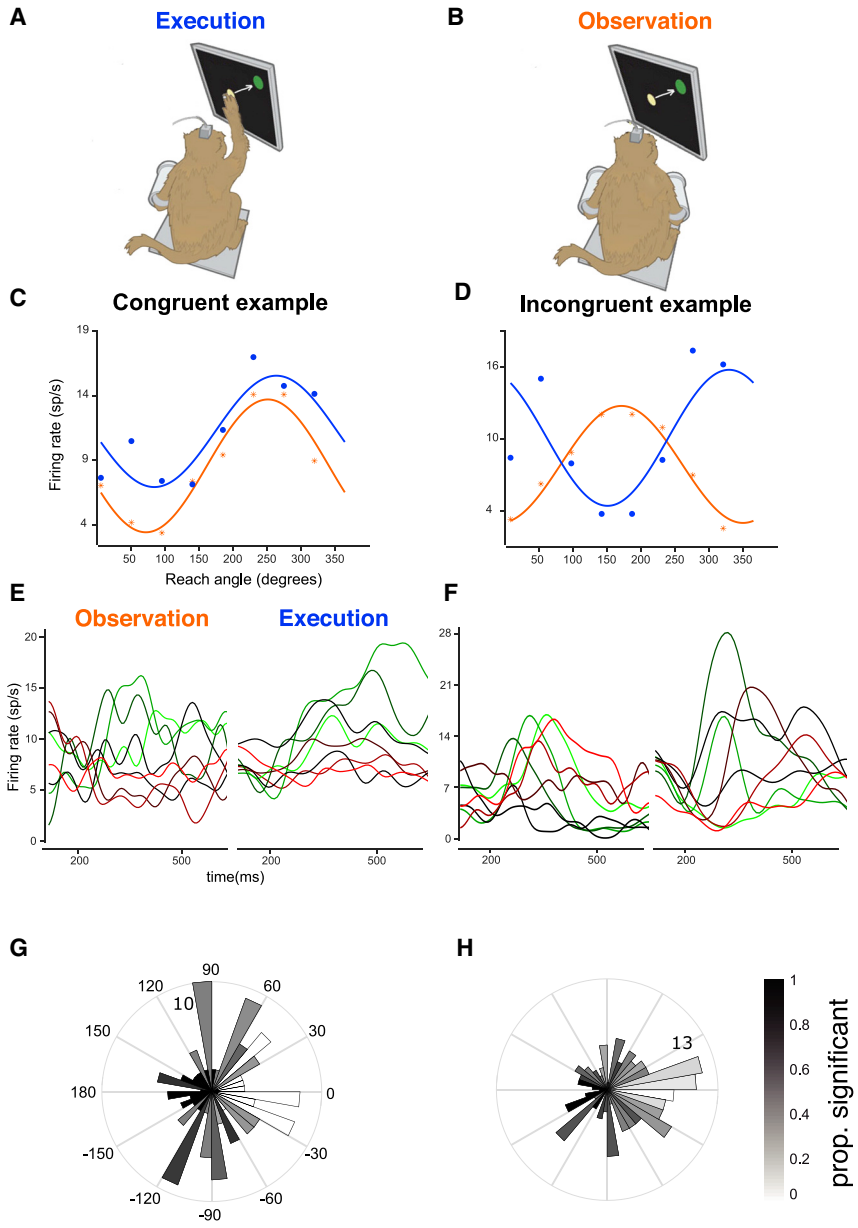

Figure 2. Behavior and Neuron Congruence

(A and B) Illustrations of (A) execution and (B) observation.

(C and D) Cosine tuning curves for a single neuron during execution (blue) and observation (orange) demonstrating (C) congruence and (D) incongruence.

(E and F) PSTHs for the (E) congruent and (F) incongruent neurons of (C) and (D) during execution and observation, with color denoting different target conditions.

(G and H) Histogram of PD differences between observation and execution for (G) Monkey J and (H) Monkey L. Grayscale indicates the proportion of neurons with a significant change in PD.

In Monkey J (L), we found that 47.4% (53.2%) of neurons in M1 and 52.7% (43.7%) of neurons in PMd were incongruent ($p < 0.05$, bootstrap with 1,000 resamples). Circular histograms of PD changes are shown in Figures 2G and 2H. The differences in PD change distributions between monkeys may be due to anatomical sampling, as Monkey J had two arrays in PMd and M1, while Monkey L had one array at the PMd/M1 border. We found more congruent neurons than the study by Dushanova and Donoghue (2010), where 38% of recorded neurons were congruent in M1. Finally, we also found that modulation depth and baseline firing rate were smaller on average in observation than in execution, as shown in Figure S1. These results show that our recorded neural population has a mixed degree of congruent and incongruent neurons.

Neural Population Subspaces for Observation and Execution Activity Are Not Orthogonal

We first examined the relationship between the neural population activity irrespective of whether neurons were congruent or incongruent. Neural population activity can be described in a high-dimensional neural space where each dimension reflects the activity of one recorded neuron. However, empirical results (Yu et al., 2009; Shenoy et al., 2013; Cunningham and Yu, 2014; Gallego et al., 2017) and theoretical studies (Gao et al., 2017) show that neural activity does not explore the entire space, but rather resides in a low-dimensional subspace. This subspace is described by a series of basis vectors or neural modes that capture how recorded neurons covary during behavior (Gallego et al., 2017, 2018). Are the neural modes of observation and execution activity orthogonal or overlapping? If observation and execution subspaces overlap, then neurons exhibit similar covariance in both contexts. If they are orthogonal, then neurons covary differently between contexts.

We first tested the hypothesis that observation and execution neural subspaces are orthogonal. This hypothesis is motivated by prior work in motor preparation that found that output-null activity during motor preparation is orthogonal to an execution subspace (Kaufman et al., 2014; Elsayed et al., 2016). This orthogonality allows motor cortex to have fundamentally different activity and computation between preparation and movement (Elsayed et al., 2016). In our study, observation activity is "putative output-null"; while we did not record electromyography (EMG) activity, analysis of videos from the experiment and BMI control suggest that the monkey was not contracting his muscles during the experiment (see Discussion and Figure S6). As a putative output-null space, observation activity, like preparatory activity, may be orthogonal to execution

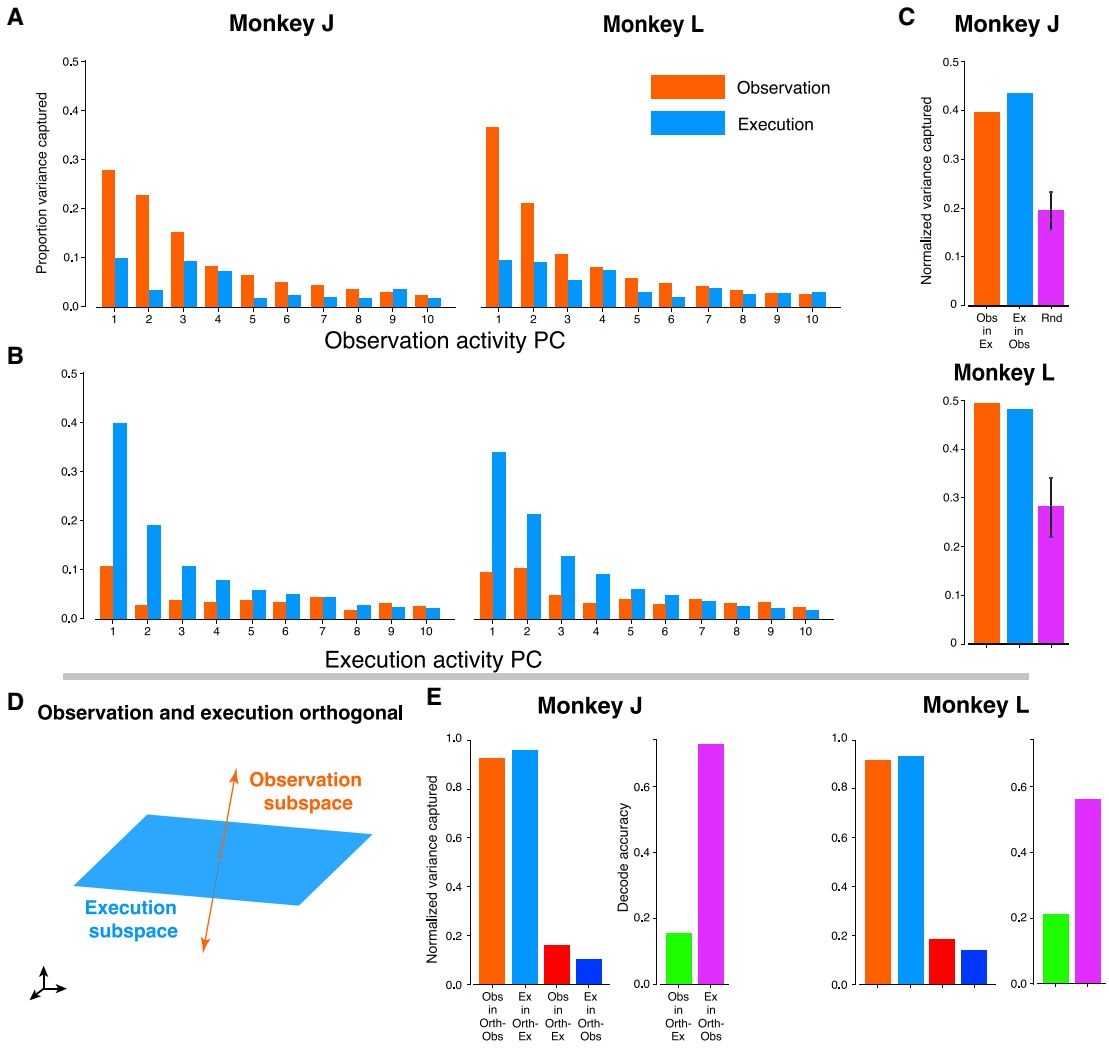

Figure 3. Observation and Execution Subspaces Are Not Orthogonal

(A and B) Observation and execution variance explained in (A) observation PCs and (B) execution PCs.

(C) Variance captured by cross-projections in the top 10 PCs (orange: Obs variance in Ex PCs; blue: Ex variance in Obs PCs), and variance captured by random observation and execution 10D subspaces (purple; error bars: 95% confidence interval).

(D) Hypothesis illustration.

(E) Orth-Obs and Orth-Ex subspace optimization. Variance captured within a context's own subspace (orange: Obs variance in Orth-Obs; blue: Ex variance in Orth-Ex) and across context "cross-projections" (red: observation variance in Orth-Ex; dark blue: execution variance in Orth-Obs). We also decoded Obs variance cross-projected in Orth-Ex (green) and Ex variance cross-projected in Orth-Obs (purple) (chance: 12.5%).

subspaces. We tested this by performing principal component analysis (PCA) separately on observation and execution trial-averaged peristimulus time histograms (PSTHs) (see Method Details). In observation, the top 10 PCs captured 74.1% (80.0%) of the total variance in Monkey J (L), while in execution they captured 83.9% (87.6%). We then projected observation activity on the execution PCs and vice versa to quantify the variance captured in cross-projections, as shown in Figures 3A and 3B. We found that 40.0% (49.5%) of cross-projected observation variance was captured in the top 10 execution PCs, and 44.0% (48.2%) of cross-projected execution variance was captured in the top 10 observation PCs. This cross-projection was significantly larger than expected by chance (Figure 3C, $p < 0.05$, comparison to random projections biased toward the original subspace; see Method Details). Because statistically significant observation variance is present in the execution PCs and vice versa, this analysis suggests that observation and execution subspaces are not orthogonal.

To show this more conclusively, we designed an additional analysis to consider two important factors unaccounted for by this PCA. First, Elsayed et al. (2016) demonstrated that explicit joint optimization over two contexts is a stronger test of orthogonality. Our prior PC analysis was performed separately (not jointly) on observation and execution activity, which may

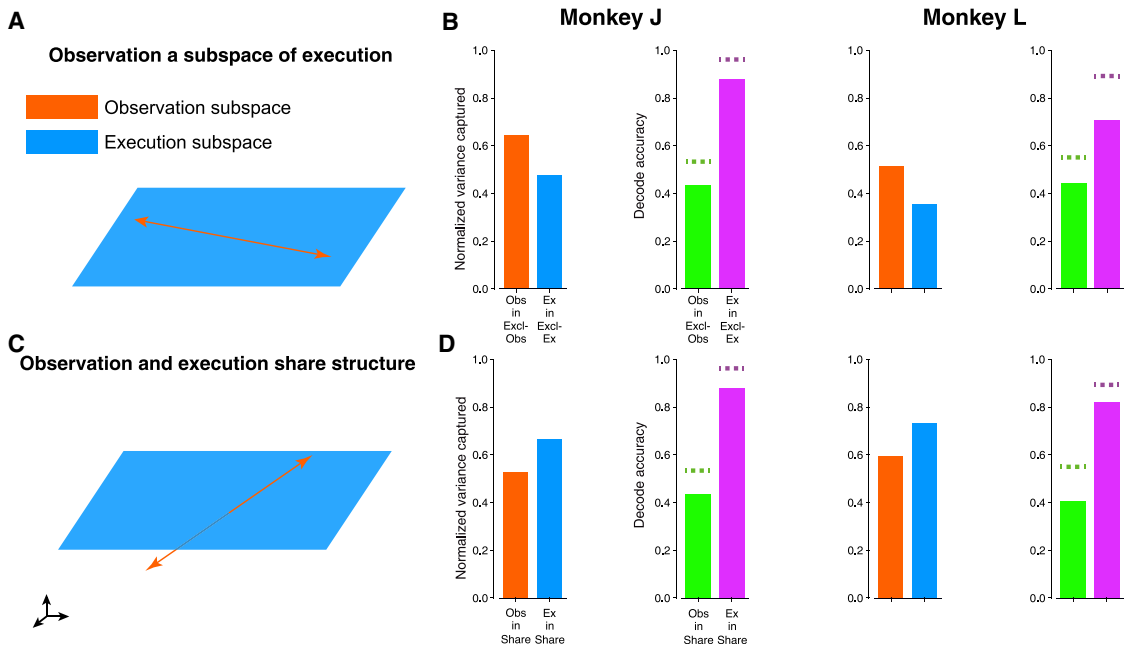

Figure 4. Exclusive and Shared Subspace

(A) Illustration of the hypothesis that the observation subspace is a subset of the execution subspace.

(B) Variance captured in the Excl-Obs (orange) and Excl-Ex (blue) subspaces. By design, Excl subspaces capture less than 1% cross-projected variance. Observation variance is high in Excl-Obs, and can be decoded at significantly above chance levels (green). The dotted line indicates maximum accuracy achieved using all neurons. The same results hold for execution activity.

(C) Illustration of the hypothesis that observation and execution activity share neural population structure.

(D) The shared subspace captures significant observation and execution variance, and was decodable above chance levels for observation (green) and execution (purple) variance.

have had spurious overlap due to noise. In contrast, a more rigorous approach is to ask *if it is possible* to find orthogonal observation and execution subspaces jointly. Second, cross-projected variance may be noise, rather than behaviorally informative. We addressed these concerns by performing an additional dimensionality reduction analysis and decoding analysis. We used manifold optimization to jointly identify mutually orthogonal subspaces, one for observation ("Orth-Obs") and one for execution ("Orth-Ex") (Koep and Weichwald, 2016). If observation and execution activity are orthogonal, then the Orth-Obs and Orth-Ex subspaces will have nearly zero cross-projected variance; otherwise, they will have significant cross-projected variance. To test if any cross-projected variance is signal or noise, we decoded target identity from the cross-projected variance. We report results for four-dimensional (4D) Orth-Obs and Orth-Ex subspaces, but our conclusions hold for different dimensionalities (Table S2). We use a metric called "normalized variance captured" (see Method Details), which is the variance captured in the 4D Orth subspace divided by the variance captured by the top four PCs (i.e., the maximum possible variance captured by four dimensions). All variance captured quantities in this study are normalized in this way, unless otherwise noted.

Our optimization identified an Orth-Obs and Orth-Ex subspace, each capturing over 90% of variance in their contexts across both monkeys, as shown in Figure 3E. Critically, we observed that there was significant cross-projected variance. For Monkey J (L), there was 15.9% (18.0%) cross-projected observation variance in the Orth-Ex, and 10.3% (13.9%) cross-projected execution variance in Orth-Obs. To test whether this cross-projected variance was noise, we decoded cross-projected activity to decode target identity using a linear discriminator (see Methods). We found that execution variance in Orth-Obs was task relevant, achieving 72.7% (55.7%) decode accuracy (Figure 3E, purple bars), significantly above chance levels of 12.5% ($p < 0.05$, bootstrap with 1,000 resamples). In contrast, observation variance in Orth-Ex was not always task relevant, achieving 15.3% (21.3%) decode accuracy (Figure 3C, green bars). These modest accuracies were significantly above chance in four of six datasets ($p < 0.05$, bootstrap with 1,000 resamples; $p = 0.330$ and 0.238 in remaining datasets).

The presence of task-relevant execution variance in the Orth-Obs subspace rejects the hypothesis that execution and observation activity are orthogonal. But subtly, observation activity has little or no task-relevant variance in Orth-Ex. As execution activity has variance in Orth-Obs, but observation activity has limited task-related variance in Orth-Ex, this leads to an alternative hypothesis: observation activity is a subset of execution activity. In this hypothesis, observation and execution share neural modes for processing actions.

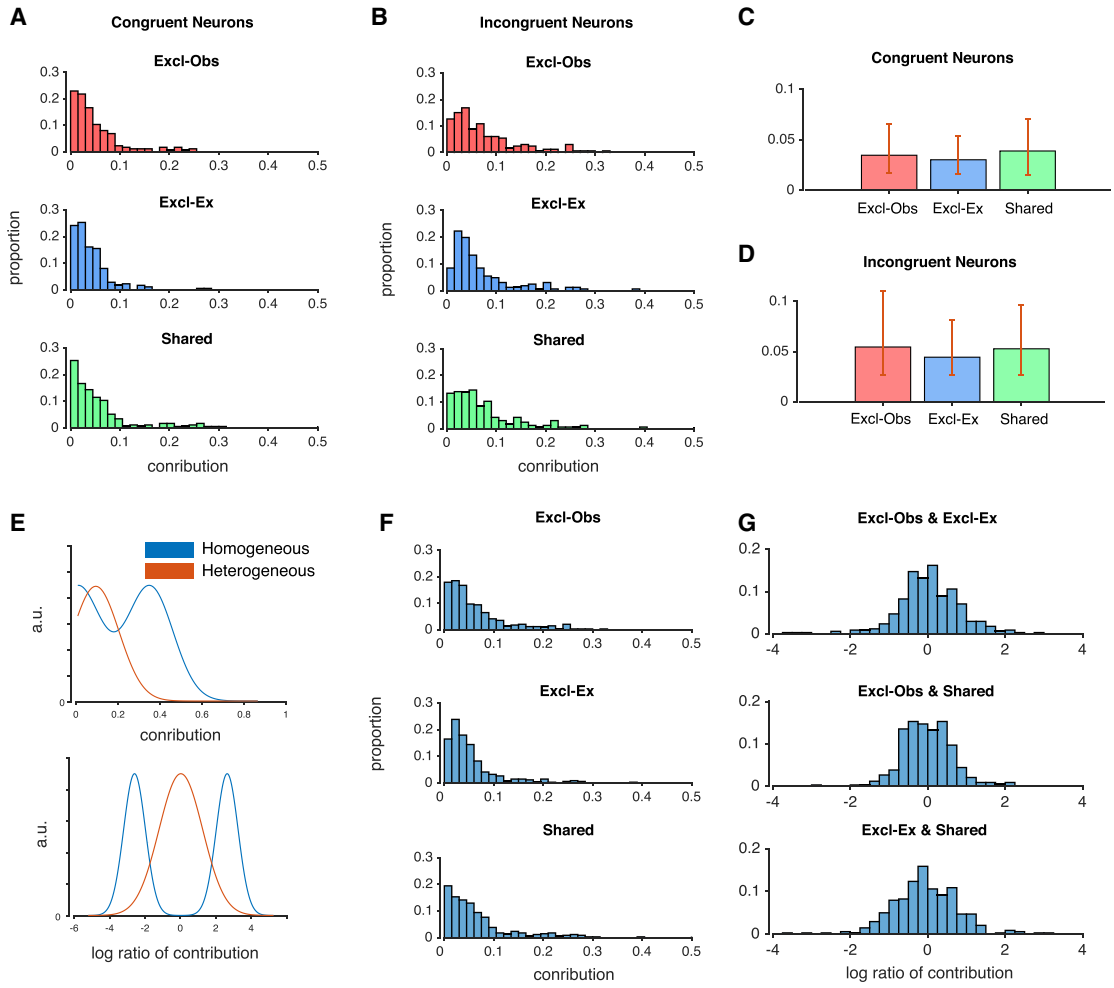

Figure 5. Congruent and Incongruent Composition of the Exclusive and Shared Subspaces

(A and B) Distribution of neuron subspace contributions in the Excl-Obs, Excl-Ex, and shared subspaces for congruent neurons (A) and incongruent neurons (B).

(C and D) Median (error bars are 25th and 75th percentiles) neuron subspace contributions in Excl-Obs, Excl-Ex, and shared subspaces for congruent neurons (C) and incongruent neurons (D).

(E) Hypothesis figures. If the homogeneous hypothesis is true, then the distribution of neuron subspace contributions should be bimodal across all subspaces. One subpopulation should have nearly zero contributions for the exclusive (or shared) subspaces and significant contributions for the shared (or exclusive) subspaces. Further, the log ratio of the subspace contributions between the shared and exclusive subspaces should be bimodal.

(F and G) The empirical distribution of (F) neuron subspace contributions and (G) log ratio of neuron subspace contributions in Excl-Obs, Excl-Ex, and shared subspaces is more consistent with the heterogeneous hypothesis.

Execution activity, however, exists in a larger subspace, comprising additional dimensions that lead to overt movement. These dimensions are explicitly avoided during observation.

Observation Activity Is Not a Subset of Execution Activity

Is the observation subspace a subset of the execution subspace, where shared task-related processing occurs, but is otherwise orthogonal to other execution dimensions involved in movement generation? This hypothesis is conceptually diagrammed in Figure 4A. Under this hypothesis, there would be no subspace that captures task-relevant observation variance and essentially zero execution variance. We designed a new constrained subspace optimization, which defined an "exclusive" observation (Excl-Obs) subspace that maximizes observation variance while capturing less than 1% execution variance (see Method Details). If the Excl-Obs subspace has task-relevant observation variance, then observation activity explores dimensions avoided during execution, rejecting the hypothesis.

We found, for Monkey J (L), that the Excl-Obs subspace captured 64.1% (51.3%) of normalized observation variance, as shown in Figure 4B. This Excl-Obs variance was task relevant, achieving a decode accuracy of 43.6% (44.4%), close to the maximum performance, with all neurons denoted by dotted lines in Figure 4B ($p < 0.05$, bootstrap with 1,000 resamples). In addition, when we decoded the 1% variance execution activity in the Excl-Obs subspace, decode accuracy was significantly

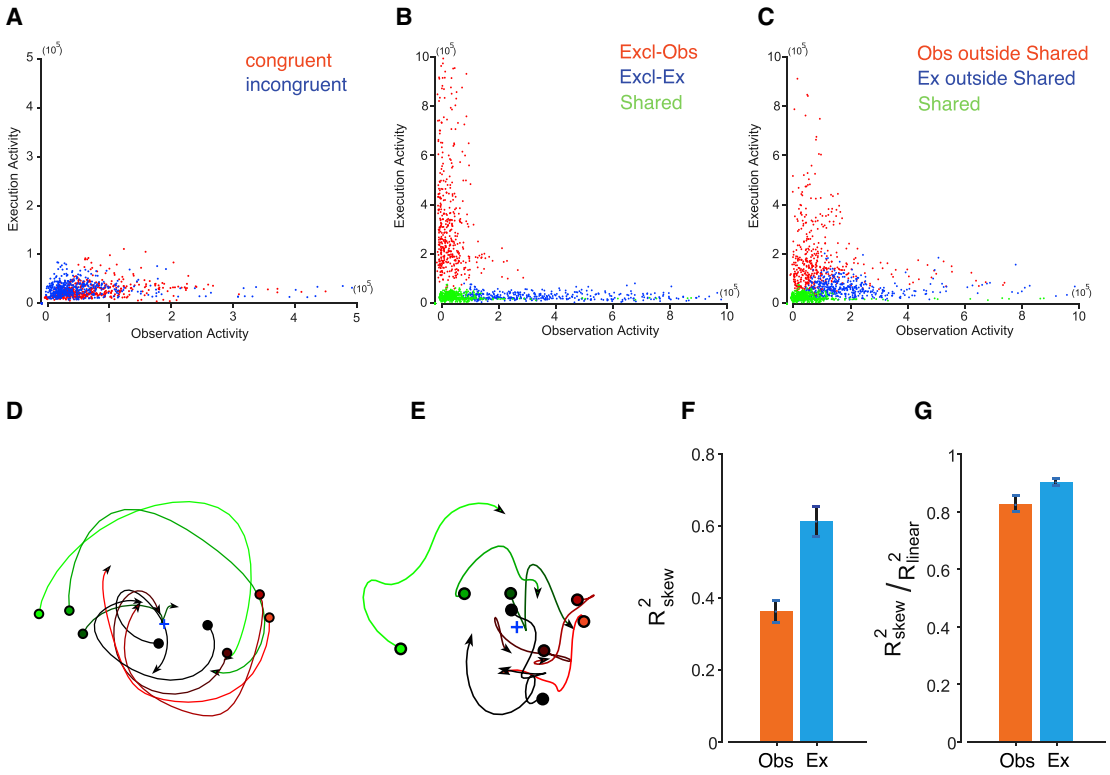

Figure 6. Dynamics in Observation and Execution

(A–C) Tangling results for Monkey J. Monkey L's results are consistent and not shown.

(A) Tangling in congruent and incongruent subpopulations during observation and execution. Each data point corresponds to a different time.

(B) Tangling in the Excl-Obs, Excl-Ex, and Shared subspace.

(C) Tangling in the highest variance subspace orthogonal to the Shared subspace for observation and execution.

(D and E) Projection of (D) execution and (E) observation activity onto the top jPC plane. Colors denote target conditions.

(F) Average R^2_{skew} in observation and execution activity across six datasets (error bars: SEM).

(G) Average ratio of R^2_{skew}/R^2_{linear}, (error bars: SEM).

smaller than decoding random projections of execution activity ($p < 0.05$, bootstrap with 1,000 resamples). These results show that observation activity has task-relevant variance in unique dimensions avoided by execution activity. This rejects the hypothesis that observation activity is a subset of execution activity.

Results from prior literature (Kaufman et al., 2014), as well as Figure 3, indicate that an exclusive execution (Excl-Ex) subspace must exist: execution activity uniquely explores dimensions for movement generation (output-potent) that are avoided by observation activity (putative output-null). We similarly applied the exclusive subspace optimization to quantify the amount of Excl-Ex variance during executed movements, which partially reflects neural variance dedicated to generating movement. We found that the Excl-Ex subspace captured 47.2% (35.1%) of normalized execution variance. This Excl-Ex variance was task relevant, achieving a decode accuracy of 87.9% (70.7%), significantly higher than chance levels ($p < 0.05$, bootstrap with 1,000 resamples).

In addition to these exclusive subspaces, observation and execution activity must explore a common shared subspace because they are not orthogonal. We therefore designed an optimization to identify this shared subspace. We defined the shared subspace to be orthogonal to Excl-Obs and Excl-Ex, while jointly maximizing observation and execution variance (see Method Details). We found that the shared subspace captured 52.8% (59.2%) of normalized observation variance in Monkey J (L) and 66.3% (73.1%) of normalized execution variance, as shown in Figure 4D. This variance was task relevant for both contexts ($p < 0.05$, bootstrap with 1,000 resamples): observation activity could be decoded at 43.5% (40.7%) accuracy and execution activity at 88.1% (81.8%) accuracy. These results identify similar patterns of neural population covariation during observation and execution occurring in a shared subspace, and dissimilar covariation occurs in exclusive subspaces. We next asked: how does neuron congruence and incongruence relate to the shared and exclusive subspaces?

The Shared Subspace Has a Heterogeneous Composition of Congruent and Incongruent Neurons

The congruent subpopulation hypothesis poses the concept that congruent neurons mediate an understanding of a motor action through congruent activation of a motor cortical circuit. If congruent neurons reflect shared computation between

observed and executed movements, they should primarily contribute to the shared subspace. Likewise, if incongruent neurons reflect unique computation, they should primarily contribute to the exclusive subspaces. We term this the "homogeneous" composition hypothesis, diagrammed in Figure 1B. In contrast, the "heterogeneous" composition hypothesis posits that both congruent and incongruent neurons contribute to the shared and exclusive subspaces.

To test if the shared (or exclusive) subspace is primarily composed of congruent (or incongruent) neurons, we analyzed the contribution of incongruent and congruent neurons in the Excl-Obs, Excl-Ex, and shared subspaces. We defined a neuron's "subspace contribution" as the square root of the variance of the projected neuron's activity in the subspace divided by the neuron's total variance, times the mean firing rate of the neuron (in both contexts). This simplifies to the magnitude of the weights times the mean firing rate, a metric previously used to quantify channel contributions to subspaces (see Method Details) (Nuyujukian et al., 2014). As shown in Figures 5A and 5B, we observed that congruent and incongruent neurons had similar subspace contribution distributions for both the shared and exclusive subspaces. To further quantify this, we compared the PD changes of congruent neurons with a subspace contribution of at least 0.01 to see if congruent neurons contributing to the shared subspace may have smaller PD shifts than those contributing to the exclusive subspaces. As shown in Figures S4A–S4C, there was no statistically significant difference in PD shifts between congruent neurons in the shared versus exclusive subspaces. We also repeated the optimizations to find Excl-Obs, Excl-Ex, and shared subspaces using only subpopulations of either congruent or incongruent neurons. Even with homogeneous subpopulations, congruent neurons had significant exclusive subspaces and incongruent neurons had significant shared subspaces (Figure S2).

These results reject the homogeneous composition hypothesis, and suggest that shared neural population structure is achieved through a heterogeneous population, rather than through a congruent subpopulation of neurons. However, are there any other homogeneous subpopulations (rather than congruent and incongruent categories) that define neuron classes primarily contributing to shared or exclusive subspaces? If the shared and exclusive subspaces were primarily driven by any homogeneous subpopulation, we would expect (1) the distribution of subspace contributions to be bimodal for both shared and exclusive subspaces, and (2) the log ratio of neuron subspace contributions to the shared and exclusive subspaces to be bimodal. This is because, if such subpopulations exist, then a neuron having a large (small) weight for the shared subspace should also have a small (large) weight for the exclusive subspace. These hypotheses are diagrammed in Figure 5E. The empirical contributions and log ratios, shown in Figures 5F and 5G, were consistent with the heterogeneous hypothesis. This suggests that there do not appear to be any subpopulations of neurons that contribute only to the shared or exclusive subspaces. As an additional control, we assessed if there was a bimodal population of neurons that prefer observation versus execution, which could explain the exclusive subspace composition. To test this, we calculated each neuron's preference index for neuron i as $B(i) = B_{obs}(i)/\overline{B}_{obs} - B_{ex}(i)/\overline{B}_{ex}$, where B is tuning strength and $\overline{B}$ is average tuning strength across all neurons (Elsayed et al., 2016). We then performed Hartigan's dip test on these preference indices to test if there was a bimodal distribution of observation versus execution-preferring neurons. We did not find that these neuron preference indices were bimodal ($p = 0.974$ [0.978] in Monkey J [L]). Together, these results support the concept that a heterogeneous population is responsible for shared structure between observed and executed movements, rather than a subpopulation of neurons.

Neural Dynamics Are Similar in Shared Subspaces, Dissimilar in Exclusive Subspaces

Neural population structure describes how neurons covary across all time. Another key feature of neural population activity is its dynamics: how neural population activity evolves predictably through time (Churchland et al., 2012; Russo et al., 2018; Pandarinath et al., 2018a). If similar computations occur during observation and execution, then neural population activity ought to evolve through time in similar ways. Specifically, shared subspaces should exhibit more similar dynamics between observation and execution than exclusive subspaces.

We quantified the similarity of dynamics between observation and execution in the shared and exclusive subspaces, as well as in congruent and incongruent subpopulations. We computed similarity of dynamics through tangling (Russo et al., 2018). Neural dynamics can be thought of as a flow field in the subspace, where the neural population activity is propelled according to the direction and strength of the flow field. Tangling is a measure of how inconsistent the flow field is in the subspace. Dynamics are tangled if, at the same location in the subspace, the flow field is inconsistently oriented. When the flow field is consistent, dynamics are untangled.

We first compared tangling for congruent and incongruent neuron subpopulations. If the congruent subpopulation hypothesis is true, then the congruent subpopulation should exhibit less tangling in observation and execution. In observation, the median congruent neuron tangling for Monkey J (L) was 6.36×10^4 (8.18×10^4), while incongruent neuron tangling was 4.00×10^4 (6.48×10^4). In execution, the median congruent neuron tangling for Monkey J (L) was 2.21×10^4 (3.21×10^4), while incongruent neuron tangling was 2.65×10^4 (3.17×10^4). Although this difference is statistically significant in all pairs ($p < 0.05$, Wilcoxon rank sum test), the effect is relatively small, and the tangling distributions of congruent and incongruent subpopulations are highly overlapping (Figure 6A). Congruent and incongruent subpopulations therefore have similar levels of tangling in observation and execution.

We next compared tangling for the shared (green), Excl-Obs (orange), and Excl-Ex (blue) subspaces (Figure 6B). We observed that, in the Excl-Obs subspace, observation activity was less tangled than execution activity. Likewise, in the Excl-Ex subspace, execution activity was less tangled than observation activity. These results are not surprising, given the negligible execution (observation) variance in the Excl-Obs (Excl-Ex) subspace. This nevertheless supports the hypothesis that the exclusive subspaces support context-specific activity, as execution

activity in Excl-Obs and observation activity in Excl-Ex have inconsistent dynamics. In contrast, we found that the shared subspace had relatively low tangling for both observation and execution activity, as shown by the green dots in Figure 6B. We also found that tangling for execution (observation) activity in the Excl-Ex (Excl-Obs) subspace was, on average, smaller than tangling in the shared subspace. This is not entirely surprising: the Excl-Ex and Excl-Obs subspaces are optimized to capture significant task-related execution and observation variance, and therefore should exhibit context-specific dynamics. The shared subspace, however, is unique in that, unlike Excl-Ex and Excl-Obs, it has relatively consistent dynamics for both contexts (less tangling than in across-context exclusive subspaces, $p < 0.05$, Wilcoxon rank sum test). We also evaluated tangling through time over the course of the trial in each subspace, but did not observe any consistent evidence that tangling was higher during particular points in the trial. Finally, we computed the speed of neural population trajectories through time, which have previously been reported to decrease over the course of a reach (Afshar et al., 2011). We observed this general decrease in trajectory speed in the shared and within-context exclusive subspaces, but not in the opposite context exclusive subspace, further supporting the idea that relevant dynamics occur in the shared and within-context exclusive subspaces (Figure S5D).

The Excl-Ex and Excl-Obs subspaces are designed to capture little cross-projected variance and, as such, have high tangling in observation and execution contexts, respectively. We performed an additional control to demonstrate that the shared subspace uniquely captures low tangling in both observation and execution. After computing the shared subspace, we computed the next best 4D orthogonal subspaces that maximally captured the remaining execution (blue) and observation (orange) variance. In these orthogonal subspaces, we found that there was higher tangling relative to the shared subspace, as shown in Figure 6C. These results suggest that the shared subspace captures shared neural population dynamics between observation and execution, potentially reflecting similar computations in both contexts.

A consequence of these results is that dynamical motifs present in executed movements may also be present during observed movements. We therefore assessed if observation activity exhibits rotational dynamics, a salient feature in execution activity (Churchland et al., 2012; Kaufman et al., 2016; Pandarinath et al., 2015). We used jPCA (Churchland et al., 2012) to calculate R^2_{skew} and R^2_{linear}, which are coefficients of determination for how well future neural states can be predicted using strictly rotational and strictly linear dynamics, respectively (see Method Details). Consistent with prior results, we found that purely rotational dynamics performed reasonably well at predicting future neural states during execution ($R^2_{skew} = 0.61$, $R^2_{linear} = 0.69$, averaged across both monkeys). Interestingly, we found that rotational dynamics could also predict future neural states in observation activity reasonably well ($R^2_{skew} = 0.36$, $R^2_{linear} = 0.45$), albeit worse than in execution. Example jPCA trajectories are shown for execution (Figure 6D) and observation activity (Figure 6E). Across all datasets, we found that observation activity could be described by a rotational dynamical system explaining, on average, 40% of future neural state variance (Figure 6F). The relative performance of using solely rotational dynamics versus general linear dynamics (which allow for contraction or expansion) was comparable between observation and execution (Figure 6G), suggesting that rotational dynamical motifs can be prominent in observation activity. We found that the shared subspace had relatively strong rotational dynamics in both observation and execution contexts, while exclusive subspaces with high tangling (i.e., opposing context) did not have strongly rotational dynamics (Figures S5B and S5C). We did not observe a significant difference in rotational dynamics between incongruent and congruent subpopulations in both execution and observation (Figure S5A).

DISCUSSION

The congruent subpopulation hypothesis posits that mirror neurons support understanding of an observed motor action through activation (or "resonance") of a group of congruent neurons that would have fired during executed movements. However, recent literature has suggested that shared computation in observed and executed movements may occur not in certain cell subpopulations, but across a heterogeneous neural population (Papadourakis and Raos, 2019; Mazurek et al., 2018). Our results provide new insights into neural population structure between observed and executed movements. First, observation and execution activity exhibit both shared and exclusive structure in their neural population activity (Figure 1A). Second, we did not identify any neural subpopulations in the shared and exclusive subspaces. In particular, congruent and incongruent neurons contributed to both subspaces (Figure 1B, "heterogeneous"). Third, neural dynamics are relatively untangled in the shared subspace compared with other subspaces, but the subpopulation of congruent neurons has similar tangling to incongruent neurons (Figure 1C, "similarity in subspace"). These results argue that shared computation—referring to similar neural activity covariation and temporal evolution across observation and execution—occurs across a heterogeneous neural population, and not in a congruent subpopulation of neurons.

These results also suggest that there is context-specific processing that occurs in exclusive subspaces. Our results of an Excl-Ex subspace agree with prior literature. The analyses of Figure 3 show it was possible to find an Excl-Ex subspace where projected observation activity could not be decoded above chance. A consistent interpretation of these results is that the Excl-Ex, which contains minimal observation variance, is partially aligned with output-potent dimensions (Kaufman et al., 2014). If observation activity is accompanied by EMG silence (i.e., no muscle activation), then observation activity would explicitly avoid these output-potent dimensions. Therefore, it is likely that the Excl-Ex subspace aligns with output-potent dimensions involved in executed movements.

Although we did not measure EMG activity, it is unlikely that the monkey was co-contracting his muscles during observation trials for three reasons. (1) During observation experiments, we observed no visible overt movement of the hands. To quantify this, we used DeepLabCut (Mathis et al., 2018) to track Monkey J's fist and Monkey L's digits in a video recording of a representative session. Digit movement was, on average, less than 1 pixel/s for Monkey J (L) across all observation trials. For

comparison, Monkey J's (L's) hand in video had an area of approximately 195 pixels^2 (158 pixels^2). This small, non-zero velocity is consistent with artifactual micromotion due to tracking precision in DeepLabCut, which we observed occasionally had errors of a few pixels (Figure S6). (2) Observation data were used to train a biomimetic BMI decoder that achieved very high performance from BMI onset without learning. BMI trials required moving the cursor to the target followed by a 500-ms hold time, for eight unique radial targets. Monkeys often had to "dial in" to the target, requiring precise fine control (Gilja et al., 2012). Achieving this level of control at BMI onset would require (a) distinct co-contraction patterns to move to each target, (b) a distinct co-contraction pattern for holding the target, (c) a co-contraction pattern that enables accurate fine control, and (d) executing this strategy at peak performance at the onset of BMI control. These systematic and quickly changing co-contractions would likely be visible on video; we did not observe this. (3) Monkeys J and L were able to sustain BMI control for hours. Sustaining such a precise, yet latent, co-contraction strategy for this long is unlikely, given the effort and energy required. Because the BMI is biomimetic, immediately achieving peak performance without learning, the lack of a co-contraction strategy in BMI control implies a lack of co-contraction strategy during observation. Together, these conclusions support the idea that the monkey was likely not co-contracting his muscles, so that observation activity would avoid output-potent dimensions.

More interestingly, we found that there was an Excl-Obs subspace. This evidence rejects the hypothesis that observation activity is merely a subspace (i.e., a putative "output-null" component) of execution activity. The motor cortex, rather, has a subspace for observation activity that is orthogonal to execution subspaces. This Excl-Obs subspace may be involved in cognitive and motor processes unique to observed movements, such as representing the agent (or lack thereof) carrying out the observed movement. It may also be linked to execution activity, much like orthogonal preparatory subspaces are linked to execution activity (Elsayed et al., 2016).

Observation and execution activity also occupy a shared subspace with similar dynamics. We note that we also found there was a shared subspace in both PMd and M1 subpopulations, as shown in Table S5. Differences from Mazurek et al. (2018), who found that premotor cortex neurons had more similar trajectories between observation and execution than did M1 neurons, may be due to task differences, location of recording (we recorded in PMd, while Mazurek and colleagues recorded primarily in PMv), or that, in our observation context, the monkey did not observe an experimenter performing the task. Nevertheless, this shared subspace reflects that, whether observing or executing movement, neurons covary in similar ways. Even more, this coordinated activity is structured through time, exhibiting relatively consistent dynamics (low tangling). This means that, in the shared subspace, a certain neural activity pattern will lead to similar future neural activity patterns, whether observing or executing movement. We found tangling outside of this shared subspace to be high, meaning that the temporal evolution of execution and observation activity outside of the shared subspace is context dependent.

Because the shared subspace has more consistent (1) covariation among neurons, and (2) dynamics through time, we hypothesize that the shared subspace is related to shared computation between observed and executed movements. For example, behavioral studies demonstrate that humans can improve their performance on sports and other motor tasks through mental rehearsal (Sheahan et al., 2018; Feltz and Landers, 1983; Frank et al., 2014; Clark, 1960; Dechent et al., 2004; Mokienko et al., 2013) and action observation (Zhang et al., 2011; Mattar and Gribble, 2005). Hence, motor learning and skill transfer can occur without overt movement. The shared subspace, which exhibits similar neural dynamics across both observed and executed movements, but comprises congruent and incongruent neurons, may mediate this motor learning and skill transfer (Jeannerod and Decety, 1995; Sirigu et al., 1996; Cisek and Kalaska, 2004; Zhang et al., 2011; Vyas et al., 2018; Sheahan et al., 2018). Further experiments should assess, in the context of task learning, if shared subspace activity reflects neural correlates of learning.

As neural population dynamics are shared between execution and observation, we expect to see similar dynamical motifs in both contexts. A prominent dynamical motif in execution activity is the rotational dynamics (Churchland et al., 2012; Pandarinath et al., 2015). In particular, Kaufman et al. (2016) found that rotational dynamics occurred during the transition from preparing to executing moving, suggesting a role for these dynamics in generating movements. In addition to movement generation, rotational dynamics have also been observed during covert mental rehearsal (Vyas et al., 2018). Our results add to this literature by demonstrating that rotational dynamics are also present during observed movements, albeit to a smaller degree than in executed movements. These results suggest that a component of the rotational dynamical motif may be related to motor processing that is not explicitly involved in generating movement.

We used PD to quantify whether neurons fired congruently between execution and observation, a metric commonly used to describe congruent activity for a center-out task (Tkach et al., 2007, 2008; Dushanova and Donoghue, 2010). In other tasks, neural responses across task conditions can be used to calculate a similarity metric, analogous to congruence (Mazurek et al., 2018). It is worth noting, however, that PDs in motor cortex can change in different contexts, including whether a prepared reach is slow or fast (Churchland et al., 2006), across time in a trial (Churchland and Shenoy, 2007), between preparation and movement (Churchland et al., 2010), under different arm orientations (Scott et al., 1997; Scott and Kalaska, 1997), between isometric or isotonic tasks (Sergio et al., 2005), and based on location of the hand during an isometric-force task (Sergio and Kalaska, 2003). It is therefore not surprising that PDs change between observation and execution, as reported by Dushanova and Donoghue (2010) and our results. In particular, we found that around half of the cells changed their PD, leading to sizeable congruent and incongruent subpopulations. It is worth noting that, even though PDs may change between contexts, this does not lead to a prediction on the composition of congruent or incongruent neurons in the shared and exclusive subspaces.

In this study, we focused on the role of congruent and incongruent neurons in the shared and exclusive subspaces. We

note that there were also neurons that were primarily active during only the execution or observation context. As such, they have little variance during the execution or observation contexts. Because the subspace optimization techniques that we used maximize the variance of the neuron firing rates across time and conditions, neurons that are relatively quiescent during a context (e.g., observation) have less variance across time, and should therefore have less contribution to the shared subspace and that context's exclusive subspace (e.g., Excl-Obs). These neurons may still contribute if they have across-condition variance. We show the subspace contribution of neurons active in both contexts and one context in Figure S4. Neurons active in both contexts in general had larger subspace contributions. Among neurons active in both contexts, we found that 44.6% of neurons were congruent and 55.4% of neurons were incongruent. It is worth re-emphasizing that, in prior analyses, we did not distinguish whether neurons were active or not, but whether they were congruent or incongruent across contexts. Our analyses therefore include all neurons, including those inactive in one context.

We note that we did not impose mandatory eye fixation. Does the shared subspace therefore reflect neural variance related to similar gaze trajectories between observation and execution? We argue that prior literature militates against this alternative interpretation. First, while PMd neurons may encode eye position (relative to hand and target) (Pesaran et al., 2006), gaze-related variables explain a smaller percentage of the variance (i.e., <20%) than that captured by the shared subspace (Cisek and Kalaska, 2002). We found that the total (unnormalized) variance captured by the shared subspace was 28% and 35% in Monkeys J and L, respectively, which is more than gaze-related variance. These results are also supported by prior analyses using partial correlations to argue that apparent correlations between motor cortical activity and eye movement during observed movements is primarily due to correlation between cursor movement and eye position (Tkach et al., 2007 [their Figure 6]). These analyses suggest that the shared subspace is not solely explained by similar eye movements. Another alternative hypothesis is that the shared subspace may primarily reflect visual feedback. If this was the case, then neural activity in the future should be most correlated with kinematics at present. We performed a decoding analysis to predict the cursor's position from the neural activity at different time lags. We found that the optimal decoding accuracy occurred when the neural data instead preceded cursor position, as shown in Figure S3. Decoding accuracy was worse when the cursor position trailed the neural data, arguing that the recorded activity is not primarily representing visual feedback.

These results have implications for BMIs, which can be designed by having users observe movements or imagine executing movements (Shenoy and Carmena, 2014; Hochberg et al., 2012; Nuyujukian et al., 2011). While there are shared components of neural population activity between observed and executed movements, the activity across both contexts remains distinct. Decoding imagined movements, which can have an output-potent component, may result in better performance than observing movement, especially considering that execution activity has distinctly different structure than observation activity. Further, because execution activity exhibits more linear dynamical motifs, such activity may be better suited for BMIs that learn and incorporate neural dynamics (Kao et al., 2015, 2017; Sussillo et al., 2016; Pandarinath et al., 2014, 2018a, 2018b).

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
  - Lead Contact
  - Materials Availability
  - Data and Code Availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Experimental design and data preprocessing
  - Trial selection during movement observation
  - Single neuron firing rate statistics
  - Classifying reach direction from neural data
  - Neural data smoothing and soft-normalization
  - Neural population matrix
  - Normalized variance captured in a subspace
  - Computing chance alignment indices
  - Subspace optimization
  - *Optimization 1: orthogonal subspaces*
  - *Optimization 2: Exclusive subspaces*
  - *Optimization 3: Shared subspaces*
  - Neuron subspace contribution
  - Variance explained for neuron subclasses
  - Quantifying dynamics
- QUANTIFICATION AND STATISTICAL ANALYSIS

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at https://doi.org/10.1016/j.celrep.2020.108006.

ACKNOWLEDGMENTS

We thank Mackenzie Risch, Michelle Wechsler, Liana Yates, Rosie Steinbach, and Shannon Smith for surgical assistance and veterinary care, as well as Evelyn Castaneda and Beverly Davis for administrative support. We thank W.L. Gore Inc. for donating preclude used in our implant procedure. This work was supported by the National Science Foundation Graduate Research Fellowship (J.C.K.); Christopher and Dana Reeve Paralysis Foundation (S.I.R. and K.V.S.); and the following to K.V.S.: NIH National Institute of Neurological Disorders and Stroke Transformative Research Award R01NS076460, NIH National Institute of Mental Health Transformative Research Award R01MH09964703, NIH Director's Pioneer Award 8DP1HD075623, Defense Advanced Research Projects Agency (DARPA) Biological Technology Office (BTO) "REPAIR" Award N66001-10-C-2010, DARPA BTO "NeuroFAST" Award W911NF-14-2-0013, Simons Foundation Collaboration on the Global Brain awards 325380 and 543045, the Office of Naval Research grant N000141812158, and the Howard Hughes Medical Institute.

AUTHOR CONTRIBUTIONS

X.J. and H.S. were involved in data analysis and manuscript writeup; S.I.R. was responsible for surgical implantation and assisted in manuscript review; K.V.S. was involved in experimental design and manuscript review; J.C.K. designed

and conducted the experiments and was involved in data review, data analysis, and manuscript writeup.

DECLARATION OF INTERESTS

K.V.S. is a consultant for Neuralink and is on the scientific advisory boards of CTRL-Labs, MIND-X, Inscopix, and Heal. All other authors have no competing interests.

Received: March 15, 2020
Revised: May 24, 2020
Accepted: July 16, 2020
Published: August 11, 2020; corrected online August 21, 2020

SUPPORTING CITATIONS

The following reference appears in the Supplemental Information: Berens and Velasco (2009).

REFERENCES

Afshar, A., Santhanam, G., Yu, B.M., Ryu, S.I., Sahani, M., and Shenoy, K.V. (2011). Single-trial neural correlates of arm movement preparation. Neuron *71*, 555–564.

Berens, P., and Velasco, M.J. (2009). The circular statistics toolbox for matlab.

Bonini, L., Maranesi, M., Livi, A., Fogassi, L., and Rizzolatti, G. (2014). Space-dependent representation of objects and other's action in monkey ventral premotor grasping neurons. J. Neurosci. *34*, 4108–4119.

Boumal, N., Mishra, B., Absil, P.-A., and Sepulchre, R. (2014). Manopt, a matlab toolbox for optimization on manifolds. J. Mach. Learn. Res. *15*, 1455–1459.

Boyd, S., and Vandenberghe, L. (2004). Convex Optimization (Cambridge University Press).

Chandrasekaran, C., Peixoto, D., Newsome, W.T., and Shenoy, K.V. (2017). Laminar differences in decision-related neural activity in dorsal premotor cortex. Nat. Commun. *8*, 614.

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

Churchland, M.M., Santhanam, G., and Shenoy, K.V. (2006). Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. J. Neurophysiol. *96*, 3130–3146.

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I., and Shenoy, K.V. (2010). Cortical preparatory activity: representation of movement or first cog in a dynamical machine? Neuron *68*, 387–400.

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. Nature *487*, 51–56.

Cisek, P., and Kalaska, J.F. (2002). Modest gaze-related discharge modulation in monkey dorsal premotor cortex during a reaching task performed with free fixation. J. Neurophysiol. *88*, 1064–1072.

Cisek, P., and Kalaska, J.F. (2004). Neural correlates of mental rehearsal in dorsal premotor cortex. Nature *431*, 993–996.

Clark, L.V. (1960). Effect of mental practice on the development of a certain motor skill. Research Quarterly. American Association for Health. Physical Education and Recreation *31*, 560–569.

Crammond, D.J., and Kalaska, J.F. (2000). Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. J. Neurophysiol. *84*, 986–1005.

Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. Nat. Neurosci. *17*, 1500–1509.

Cunningham, J.P., Nuyujukian, P., Gilja, V., Chestek, C.A., Ryu, S.I., and Shenoy, K.V. (2011). A closed-loop human simulator for investigating the role of feedback control in brain-machine interfaces. J. Neurophysiol. *105*, 1932–1949.

Dechent, P., Merboldt, K.-D., and Frahm, J. (2004). Is the human primary motor cortex involved in motor imagery? Brain Res. Cogn. Brain Res. *19*, 138–144.

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

Elsayed, G.F., Lara, A.H., Kaufman, M.T., Churchland, M.M., and Cunningham, J.P. (2016). Reorganization between preparatory and movement population responses in motor cortex. Nat. Commun. *7*, 13239.

Fan, J.M., Nuyujukian, P., Kao, J.C., Chestek, C.A., Ryu, S.I., and Shenoy, K.V. (2014). Intention estimation in brain-machine interfaces. J. Neural Eng. *11*, 016004.

Feltz, D.L., and Landers, D.M. (1983). The effects of mental practice on motor skill learning and performance: a meta-analysis. JSEP *5*, 25–57.

Fogassi, L., and Gallese, V. (2002). The neural correlates of action understanding in non-human primates. In Advances in Consciousness Research, *Vol. 42*, M.I. Stamenov and V. Gallese, eds. (John Benjamins Publishing Company), pp. 13–35.

Frank, C., Land, W.M., Popp, C., and Schack, T. (2014). Mental representation and mental practice: experimental investigation on the functional links between motor memory and motor imagery. PLoS One *9*, e95175.

Gallego, J.A., Perich, M.G., Miller, L.E., and Solla, S.A. (2017). Neural manifolds for the control of movement. Neuron *94*, 978–984.

Gallego, J.A., Perich, M.G., Naufel, S.N., Ethier, C., Solla, S.A., and Miller, L.E. (2018). Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. Nat. Commun. *9*, 4233.

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

Gao, P., Trautmann, E., Yu, B.M., Santhanam, G., Ryu, S., Shenoy, K., and Ganguli, S. (2017). A theory of multineuronal dimensionality, dynamics and measurement. bioRxiv. https://doi.org/10.1101/214262.

Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., and Massey, J.T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J. Neurosci. *2*, 1527–1537.

Gilja, V., Nuyujukian, P., Chestek, C.A., Cunningham, J.P., Yu, B.M., Fan, J.M., Churchland, M.M., Kaufman, M.T., Kao, J.C., Ryu, S.I., and Shenoy, K.V. (2012). A high-performance neural prosthesis enabled by control algorithm design. Nat. Neurosci. *15*, 1752–1757.

Hochberg, L.R., Bacher, D., Jarosiewicz, B., Masse, N.Y., Simeral, J.D., Vogel, J., Haddadin, S., Liu, J., Cash, S.S., van der Smagt, P., and Donoghue, J.P. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. Nature *485*, 372–375.

Jeannerod, M., and Decety, J. (1995). Mental motor imagery: a window into the representational stages of action. Curr. Opin. Neurobiol. *5*, 727–732.

Kao, J.C., Stavisky, S.D., Sussillo, D., Nuyujukian, P., and Shenoy, K.V. (2014). Information systems opportunities in brain-machine interface decoders. Proc. IEEE *102*, 666–682.

Kao, J.C., Nuyujukian, P., Ryu, S.I., Churchland, M.M., Cunningham, J.P., and Shenoy, K.V. (2015). Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. Nat. Commun. *6*, 7759.

Kao, J.C., Ryu, S.I., and Shenoy, K.V. (2017). Leveraging neural dynamics to extend functional lifetime of brain-machine interfaces. Sci. Rep. *7*, 7395.

Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. Nat. Neurosci. *17*, 440–448.

Kaufman, M.T., Seely, J.S., Sussillo, D., Ryu, S.I., Shenoy, K.V., and Churchland, M.M. (2016). The largest response component in the motor cortex reflects movement timing but not movement type. eNeuro *3*, ENEURO.0085-16.2016.

Koep, N., and Weichwald, S. (2016). Pymanopt: a python toolbox for optimization on manifolds using automatic differentiation. J. Mach. Learn. Res. *17*, 1–5.

Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. Nat. Neurosci. *21*, 1281–1289.

Mattar, A.A., and Gribble, P.L. (2005). Motor learning by observing. Neuron *46*, 153–160.

Mazurek, K.A., Rouse, A.G., and Schieber, M.H. (2018). Mirror neuron populations represent sequences of behavioral epochs during both execution and observation. J. Neurosci. *38*, 4441–4455.

Mokienko, O.A., Chervyakov, A.V., Kulikova, S.N., Bobrov, P.D., Chernikova, L.A., Frolov, A.A., and Piradov, M.A. (2013). Increased motor cortex excitability during motor imagery in brain-computer interface trained subjects. Front. Comput. Neurosci. *7*, 168.

Nuyujukian, P., Fan, J.M., Gilja, V., Kalanithi, P.S., Chestek, C.A., and Shenoy, K.V. (2011). Monkey models for brain-machine interfaces: the need for maintaining diversity. In Proceedings of the 33rd Annual Conference of the IEEE EMBS (Boston), pp. 1301–1305.

Nuyujukian, P., Kao, J.C., Fan, J.M., Stavisky, S.D., Ryu, S.I., and Shenoy, K.V. (2014). Performance sustaining intracortical neural prostheses. J. Neural Eng. *11*, 066003.

Pandarinath, C., Gilja, V., Blabe, C., Hochberg, L.R., Shenoy, K.V., and Henderson, J.M. (2014). Corresponding neural signatures of movement and imagined movement in human motor cortex. COSYNE *2014*, 207.

Pandarinath, C., Gilja, V., Blabe, C.H., Nuyujukian, P., Sarma, A.A., Sorice, B.L., Eskandar, E.N., Hochberg, L.R., Henderson, J.M., and Shenoy, K.V. (2015). Neural population dynamics in human motor cortex during movements in people with ALS. eLife *4*, e07436.

Pandarinath, C., Ames, K.C., Russo, A.A., Farshchian, A., Miller, L.E., Dyer, E.L., and Kao, J.C. (2018a). Latent factors and dynamics in motor cortex and their application to brain-machine interfaces. J. Neurosci. *38*, 9390–9401.

Pandarinath, C., O'Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D., Kao, J.C., Trautmann, E.M., Kaufman, M.T., Ryu, S.I., Hochberg, L.R., et al. (2018b). Inferring single-trial neural population dynamics using sequential auto-encoders. Nat. Methods *15*, 805–815.

Papadourakis, V., and Raos, V. (2019). Neurons in the macaque dorsal premotor cortex respond to execution and observation of actions. Cereb. Cortex *29*, 4223–4237.

Pesaran, B., Nelson, M.J., and Andersen, R.A. (2006). Dorsal premotor neurons encode the relative position of the hand, eye, and goal during reach planning. Neuron *51*, 125–134.

Rizzolatti, G., and Craighero, L. (2004). The mirror-neuron system. Annu. Rev. Neurosci. *27*, 169–192.

Rizzolatti, G., Fogassi, L., and Gallese, V. (2001). Neurophysiological mechanisms underlying the understanding and imitation of action. Nat. Rev. Neurosci. *2*, 661–670.

Russo, A.A., Bittner, S.R., Perkins, S.M., Seely, J.S., London, B.M., Lara, A.H., Miri, A., Marshall, N.J., Kohn, A., Jessell, T.M., et al. (2018). Motor cortex embeds muscle-like commands in an untangled population response. Neuron *97*, 953–966.e8.

Rutishauser, U., Schuman, E.M., and Mamelak, A.N. (2006). Online detection and sorting of extracellularly recorded action potentials in human medial temporal lobe recordings, in vivo. J. Neurosci. Methods *154*, 204–224.

Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu, B.M., and Batista, A.P. (2014). Neural constraints on learning. Nature *512*, 423–426.

Scott, S.H., and Kalaska, J.F. (1997). Reaching movements with similar hand paths but different arm orientations. I. Activity of individual cells in motor cortex. J. Neurophysiol. *77*, 826–852.

Scott, S.H., Sergio, L.E., and Kalaska, J.F. (1997). Reaching movements with similar hand paths but different arm orientations. II. Activity of individual cells in dorsal premotor cortex and parietal area 5. J. Neurophysiol. *78*, 2413–2426.

Sergio, L.E., and Kalaska, J.F. (2003). Systematic changes in motor cortex cell activity with arm posture during directional isometric force generation. J. Neurophysiol. *89*, 212–228.

Sergio, L.E., Hamel-Pâquet, C., and Kalaska, J.F. (2005). Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. J. Neurophysiol. *94*, 2353–2378.

Sheahan, H.R., Ingram, J.N., Žalalytė, G.M., and Wolpert, D.M. (2018). Imagery of movements immediately following performance allows learning of motor skills that interfere. Sci. Rep. *8*, 14330.

Shenoy, K.V., and Carmena, J.M. (2014). Combining decoder design and neural adaptation in brain-machine interfaces. Neuron *84*, 665–680.

Shenoy, K.V., Sahani, M., and Churchland, M.M. (2013). Cortical control of arm movements: a dynamical systems perspective. Annu. Rev. Neurosci. *36*, 337–359.

Sirigu, A., Duhamel, J.R., Cohen, L., Pillon, B., Dubois, B., and Agid, Y. (1996). The mental representation of hand movements after parietal cortex damage. Science *273*, 1564–1568.

Stavisky, S.D., Kao, J.C., Ryu, S.I., and Shenoy, K.V. (2017). Motor cortical visuomotor feedback activity is initially isolated from downstream targets in output-null neural state space dimensions. Neuron *95*, 195–208.e9.

Sussillo, D., Stavisky, S.D., Kao, J.C., Ryu, S.I., and Shenoy, K.V. (2016). Making brain-machine interfaces robust to future neural variability. Nat. Commun. *7*, 13749.

Taylor, D.M., Tillery, S.I.H., and Schwartz, A.B. (2002). Direct cortical control of 3D neuroprosthetic devices. Science *296*, 1829–1832.

Tkach, D., Reimer, J., and Hatsopoulos, N.G. (2007). Congruent activity during action and action observation in motor cortex. J. Neurosci. *27*, 13241–13250.

Tkach, D., Reimer, J., and Hatsopoulos, N.G. (2008). Observation-based learning for brain-machine interfaces. Curr. Opin. Neurobiol. *18*, 589–594.

Trautmann, E.M., Stavisky, S.D., Lahiri, S., Ames, K.C., Kaufman, M.T., O'Shea, D.J., Vyas, S., Sun, X., Ryu, S.I., Ganguli, S., and Shenoy, K.V. (2019). Accurate estimation of neural population dynamics without spike sorting. Neuron *103*, 292–308.e4.

Velliste, M., Perel, S., Spalding, M.C., Whitford, A.S., and Schwartz, A.B. (2008). Cortical control of a prosthetic arm for self-feeding. Nature *453*, 1098–1101.

Vigneswaran, G., Philipp, R., Lemon, R.N., and Kraskov, A. (2013). M1 corticospinal mirror neurons and their role in movement suppression during action observation. Curr. Biol. *23*, 236–243.

Vyas, S., Even-Chen, N., Stavisky, S.D., Ryu, S.I., Nuyujukian, P., and Shenoy, K.V. (2018). Neural population dynamics underlying motor learning transfer. Neuron *97*, 1177–1186.e3.

Yu, B.M., Cunningham, J.P., Santhanam, G., Ryu, S.I., Shenoy, K.V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. J. Neurophysiol. *102*, 614–635.

Zhang, X., de Beukelaar, T.T., Possel, J., Olaerts, M., Swinnen, S.P., Woolley, D.G., and Wenderoth, N. (2011). Movement observation improves early consolidation of motor memory. J. Neurosci. *31*, 11515–11520.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental Models: Organisms/Strains		
Rhesus macaque (Mucacca mulatta)	Wisconsin, Davis, and Yerkes Primate Centers	N/A
Software and Algorithms		
Single neuron firing rate statistics	Fan et al., 2014	https://doi.org/10.1088/1741-2560/11/1/016004
Manopt	Boumal et al., 2014	https://www.manopt.org/index.html
Tangling	Russo et al., 2018	https://doi.org/10.1016/j.neuron.2018.01.004
jPCA	Churchland et al., 2012	https://doi.org/10.1038/nature11129
MKsort	Ripple LLC	https://github.com/ripple-neuro/mksort
Subspace optimization	This paper	https://github.com/jcykao/subspace-opt
Other		
Cerebus System	Blackrock Microsystems	https://www.blackrockmicro.com/%20neuroscience-research-products/neural-data-acquisition-systems/cerebus-daq-system/
Utah Microelectrode Arrays	Blackrock Microsystems	https://www.blackrockmicro.com/electrode-types/utah-array/

RESOURCE AVAILABILITY

Lead Contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Jonathan C. Kao (kao@seas.ucla.edu).

Materials Availability

This study did not generate new unique reagents.

Data and Code Availability

Subspace optimization code is available at https://github.com/jcykao/subspace-opt. The dataset supporting the current study are available from the Lead Contact upon reasonable request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

All procedures and experiments were approved by the Stanford University Institutional Animal Care and Use Committee (IACUC). Experiments were conducted with two adult male rhesus macaques (Macaca mulatta, monkeys J and L), implanted with 96 electrode Utah arrays (Blackrock Microsystems Inc., Salt Lake City, UT) using standard neurosurgical techniques. At the time of experiments, Monkey J was 11 years old and weighed 15 kg, while Monkey L was 18 years old and weighed 9 kg. Monkey J (L) was implanted 56 (67) months prior to experiments. Monkey J had two Utah arrays, one implanted in PMd and the other implanted in M1, as visually estimated from anatomical landmarks. Monkey L had one Utah array, implanted at the PMd / M1 border.

METHOD DETAILS

Experimental design and data preprocessing

There were two contexts in our experiments: (1) observation (Obs) and (2) execution (Ex). In both contexts, there were 8 targets, equally spaced in 45° intervals at 0°, 45°, ..., 315° in a 2D plane. Each target was 12 cm from the center target. In the observation context, the monkey observed the cursor moving at a constant speed of 9.25 cm/s between the center and the 8 targets. In the execution context, the monkey made center-out-and-back reaches in this 2D plane with a virtual cursor controlled by the contralateral arm. In both cases, targets were selected after the cursor was held for 500 ms in an acceptance window of size 4 cm by 4 cm centered on the target. There was no delay period. There were no rest times between trials within an experimental block. Inter-trial intervals in both contexts were 40 ms. For data analysis, only reaches starting from the center were analyzed. The virtual cursor and

targets were presented in a 3D environment (MSMS, MDDF, USC, Los Angeles, CA) described in Cunningham et al. (2011). The monkey was unable able to see his hand during experiments, only the virtual cursor. We chose trial counts that aim to maximize the number of trials collected across all conditions within one experimental session. The number of trials for observation in each session ranges from 89 to 258, while for execution this number ranges from 251 to 266.

Hand position data were measured with an infrared reflective bead tracking system (Polaris, Northern Digital, Ontario, Canada). Behavioral control and neural decode were run on separate PCs using the Simulink/xPC platform (Mathworks, Natick, MA). Eye positions were recorded via ISCAN (ISCAN, Inc., MA). Neural data were initially processed by the Cerebus recording system (Blackrock Microsystems Inc., Salt Lake City, UT) according to specifications described in Cunningham et al. (2011). Spike events were detected by setting a threshold value to -4.5 times the RMS voltage of the channel.

After thresholding the activity, we spike sorted the data via MK sort, available at https://sites.google.com/site/antimatt/software, and found anywhere from 41 to 83 putative single units in each experimental session. We incorporated this more stringent and conservative criterion for determining single neurons, since sorting spike waveforms incorporates a subjective component. Our criterion was based on prior studies (Chandrasekaran et al., 2017; Rutishauser et al., 2006) that only counted neurons that had ISI violations (< 3 ms) less than 3% of the time. Our empirical dataset labels for this work are: J140507, J140508, L140829, L140911, L140912, L140916, the first letter indicates the monkey label (J or L), and the following numbers indicates the date of data acquisition. For example, J140507 means the data is collected from monkey J on May 07, 2014.

All presented analyses used all datasets, with the exception of the subspace composition analysis presented in Figure S2 where dataset L140829 was excluded. This was because this dataset did not have enough congruent or incongruent units to reliably determine an exclusive and shared subspace. We were not able to do this analysis with single units satisfying a tuning threshold of $R^2 > 0.3$ due to there being not enough units in all but one dataset.

Trial selection during movement observation

The monkey may cease to pay attention to the moving cursor during movement observation. We addressed this in two ways. First, we observed the monkey's eye position throughout the duration of the experiment, and stopped the task if the monkey disengaged. Second, we only analyzed trials that met the following criteria: (a) the angle between monkey's eye position vector (between the start and end of trial) and cursor movement vector was less than 45°, (b) the length of monkey's eye position vector was larger than half of the length of cursor movement vector. Together, these selection criteria resulted in trials selected where the monkey's eye movements were consistent with cursor movement. The percentage of observation trials excluded due to this criteria in each dataset was: 12.2%, 0.3%, 35.2%, 6.7%, 12.3%, 6.5%. It is also worth noting that, as described in the Results, these datasets were adequate training sets for BMIs achieving near-perfect performance on a center-out-and-back task.

Single neuron firing rate statistics

To test for significant differences in single neuron firing rate statistics, we performed a bootstrap hypothesis test (Fan et al., 2014). We first resampled the firing rate for a given direction. Then, we fit a tuning curve for the resampled firing rates. Our null hypothesis tests if a tuning curve statistic has the same mean between two different samples (i.e., observation and execution activity). For example, for the preferred direction of a single neuron, $\mathbf{x}$ and $\mathbf{y}$ are both n-dimensional column vectors representing the preferred directions for observation and execution contexts across many re-samples. We tested for statistically significant changes as follows:

1. Mean subtract statistics from two different samples: $\hat{\mathbf{x}} = \mathbf{x} - \bar{\mathbf{x}}$, $\hat{\mathbf{y}} = \mathbf{y} - \bar{\mathbf{y}}$, where $\bar{\mathbf{x}}$ denotes the mean of $\mathbf{x}$.
2. Concatenate all mean-subtracted statistics column-wise: $\mathbf{c} = [\hat{\mathbf{x}}, \hat{\mathbf{y}}]$, so $\mathbf{c}$ is a $2n$ by 1 vector.
3. Sample vectors $\mathbf{x}_{\text{mix}}$ and $\mathbf{y}_{\text{mix}}$ from $\mathbf{c}$ with replacement; each vector comprises 1000 samples.
4. Compute the mean-centered bootstrap statistic: $\Delta_{\text{ctrl}} = \mathbf{x}_{\text{mix}} - \mathbf{y}_{\text{mix}}$ and $\Delta_c = \Delta_{\text{ctrl}} - \overline{\Delta}_{\text{ctrl}}$.
5. Sample vectors $\mathbf{x}_{\text{real}}$ and $\mathbf{y}_{\text{real}}$ from $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ respectively with replacement; each vector comprises 1000 samples.
6. Compute the statistic: $\Delta_{\text{real}} = \mathbf{x}_{\text{real}} - \mathbf{y}_{\text{real}}$.
7. Calculate the p value: $p = \frac{\#(\left|\overline{\Delta}_{\text{real}}\right| < |\Delta_{\text{ctrl}}|)}{1000}$.

We rejected the null hypothesis if $p < 0.05$.

Classifying reach direction from neural data

We utilized a generative model with shared covariance matrices across classes (i.e., the 8 reach directions), thereby implementing a linear discriminant. The data was already projected into a low dimensional subspace of interest. We denote the low dimensional latent states as $\mathbf{s} = [s_1, s_2, \ldots, s_D]$, where D is number of dimensions. D is smaller than the number of total neurons. We modeled the latent states in a given class as being a multivariate Gaussian distribution, i.e.,

$$p(\mathbf{s}|C = k) \sim \mathcal{N}(\mu_k, \Sigma) \qquad \text{(Equation 1)}$$

where μ_k is class k's mean, and Σ is an across-class covariance. We constrain the across-class covariance to be diagonal. The maximum-likelihood mean is the sample mean of the low-dimensional states in class k, while the maximum-likelihood covariance

is the weighted sum of sample covariance matrices across all classes. The weights are the proportions of samples in each class. To decode, we calculate the class that maximizes the posterior probability:

$$k^* = \arg\max_k p(C = k|\mathbf{s}) \quad \text{(Equation 2)}$$

where k^* denotes the decoded class. In decoding reach direction, we found that the peak performance using all the data was approximately 60% (100%) for observation (execution). This relatively poor performance for observation indicates that activity is less modulated during observation; we note that the same decoder applied in a BMI context trained from observation achieved performance over 95%. We did not observe performance significantly increase with multi-class SVMs (data not shown).

Neural data smoothing and soft-normalization

We smoothed the firing rates with a Gaussian kernel having standard deviation 25 ms. We also soft-normalized the data before optimization (Churchland et al., 2012; Elsayed et al., 2016). This entailed normalizing the firing rate of a neuron, denoted by x in units of spikes/s, according to

$$\frac{x}{\lambda + \max(x) - \min(x)} \quad \text{(Equation 3)}$$

where $\max(x) - \min(x)$ is the range of one neuron's firing rate across all time. We set λ = 10 spikes/s based on the fact that $\max(x)$ - $\min(x)$ range from 1 to 42 spikes/s for observation, with mean at 12 spikes/s, and 1 to 82 spikes/s for execution, with mean 22 spikes/s. All conclusions of this study were upheld when varying the Gaussian kernel smoothing (25, 50, or 75 ms) and soft-normalization λ (5, 10, or 15 spikes/s).

Neural population matrix

In all neural population analysis, including PCA, subspace optimization analysis, and dynamics, we performed analysis on the trial-averaged PSTHs across neurons, conditions, and time. We formed a matrix $\mathbf{M} \in \mathbb{R}^{N \times TC}$, where N is the number of neurons, T is the number of time points, and C is the number of reach conditions (8 directions in this work). Each row of this matrix corresponds to the trial-averaged firing rates for a neuron over time, concatenated over all conditions. We formed this matrix for both observation and execution activity. We performed subspace optimization on these trial-averaged firing rates, rather than single trials, since our subspace optimizations focus on maximizing variance. Trial-averaging firing rates reduces single trial variance, so that the subspace optimizations more strongly reflect across-condition variance. This is a commonly used approach when performing dimensionality reduction on a neural population (Elsayed et al., 2016; Churchland et al., 2012).

Normalized variance captured in a subspace

We used the following equation to quantify the normalized variance captured of neural population data in a subspace, which is the metric used throughout this study. We note that Elsayed and colleagues have also called this quantity 'alignment index' (Elsayed et al., 2016). The normalized variance captured is:

$$A = \frac{\mathrm{Tr}(\mathbf{Q}^T \mathbf{C} \mathbf{Q})}{\sum_{i=1}^{d} \lambda_i},$$

where $\mathbf{Q} \in \mathbb{R}^{N \times d}$ is a d-dimensional basis over N neurons, $\mathbf{C}$ is the covariance matrix of the neural activity, and λ_i is the i^{th} largest eigenvalue for $\mathbf{C}$. The matrices $\mathbf{Q}$ and $\mathbf{C}$ can be computed for both observation and execution activity. The quantity $\sum_{i=1}^{d} \lambda_i$ is the maximum variance that can be captured in d dimensions. As such, A is a quantity between 0 and 1.

Computing chance alignment indices

We performed a random subspace selection method (reported by Elsayed et al. (2016), Supplementary Note 3) to quantify the chance level normalized variance captured (alignment index) between observation and execution subspaces. This random subspace selection is biased to the covariance structure of the data. It assumes that there is a fixed correlation between the neurons at all times, dictating a subspace where computation occurs, and randomly samples from this subspace.

For example, consider observation activity. The covariance matrix of observation activity is denoted as $\mathbf{C}_{\text{Obs}}$, and its eigendecomposition is

$$\mathbf{C}_{\text{Obs}} = \mathbf{Q}_{\text{Obs}} \Lambda_{\text{Obs}} \mathbf{Q}_{\text{Obs}}^{\mathsf{T}}, \quad \text{(Equation 4)}$$

where the columns of $\mathbf{Q}_{\text{Obs}}$ are eigenvectors of $\mathbf{C}_{\text{Obs}}$, and Λ_{Obs} is a diagonal matrix of the sorted eigenvalues of $\mathbf{C}_{\text{Obs}}$. A randomly sampled observation subspace, denoted as $\mathbf{Q}_{\text{Obs,rnd}}$, is obtained by computing:

$$\mathbf{Q}_{\text{Obs,rnd}} = \mathrm{orth}\left(\frac{\mathbf{Q}_{\text{Obs}} \Lambda_{\text{Obs}}^{1/2} \mathbf{v}}{\| \mathbf{Q}_{\text{Obs}} \Lambda_{\text{Obs}}^{1/2} \mathbf{v} \|_2} \right) \quad \text{(Equation 5)}$$

where $\mathbf{v} \in \mathbb{R}^{N \times d}$ is a matrix whose entries are independently drawn from a Gaussian distribution with zero mean and unit variance. The **orth**() operator computes the left singular values of its argument, resulting in an orthonormal basis. An analogous procedure is used to find $\mathbf{Q}_{\text{Ex,rnd}}$. To sample one random alignment index, A_{rnd}, we then compute:

$$A_{\text{rnd}} = \frac{\text{Tr}\left(\mathbf{Q}_{\text{Obs,rnd}}^T \mathbf{Q}_{\text{Ex,rnd}} \mathbf{Q}_{\text{Ex,rnd}}^T \mathbf{Q}_{\text{Obs,rnd}}\right)}{d}. \quad \text{(Equation 6)}$$

We chose $d = 10$ for this calculation. The empirical alignment index for execution activity projected onto observation PCs is

$$A_{\text{Ex-on-Obs}} = \frac{\text{Tr}\left(\mathbf{Q}_{\text{Obs}}^T \mathbf{C}_{\text{Ex}} \mathbf{Q}_{\text{Obs}}\right)}{\sum_{i=1}^{d} \lambda_i}, \quad \text{(Equation 7)}$$

where λ_i are the sorted eigenvalues of $\mathbf{C}_{\text{Ex}}$. We can also compute the analogous quantity,

$$A_{\text{Obs-on-Ex}} = \frac{\text{Tr}\left(\mathbf{Q}_{\text{Ex}}^T \mathbf{C}_{\text{Obs}} \mathbf{Q}_{\text{Ex}}\right)}{\sum_{i=1}^{d} \lambda_i}. \quad \text{(Equation 8)}$$

To perform a statistical test, we sampled 10,000 random A_{rnd} to build a chance distribution. We then compared the values of $A_{\text{Obs-on-Ex}}$ and $A_{\text{Ex-on-Obs}}$ (empirical alignments) to these 10,000 values. We observed that $A_{\text{Obs-on-Ex}}$ and $A_{\text{Ex-on-Obs}}$ were statistically significant larger than the average A_{rnd}. We computed a p value by calculating the number of A_{rnd} samples larger than the empirical alignment index.

Subspace optimization

We performed three different subspace optimization problems in this study. Here, we describe each one in detail. All subspace optimization was performed using the manopt toolbox for MATLAB (Boumal et al., 2014). When necessary, we discuss how we modified the optimization problem to be used with the toolbox.

Optimization 1: orthogonal subspaces

In our first optimization, we sought to find an orthogonal observation subspace $\mathbf{Q}_{\text{Orth-Obs}} \in \mathbb{R}^{N \times d_{\text{Obs}}}$ and an orthogonal execution subspace $\mathbf{Q}_{\text{Orth-Ex}} \in \mathbb{R}^{N \times d_{\text{Ex}}}$ that were mutually orthogonal, i.e. $\mathbf{Q}_{\text{Orth-Obs}} \perp \mathbf{Q}_{\text{Orth-Ex}}$ or equivalently, $\mathbf{Q}_{\text{Orth-Obs}}^T \mathbf{Q}_{\text{Orth-Ex}} = 0$. This is an optimization over the Stiefel manifold, which is a submanifold of real $n \times k$ matrices with orthogonal columns, i.e., $M_k(\mathbb{R}^n) = \{\mathbf{X} \in \mathbb{R}^{n \times k} : \mathbf{X}^T\mathbf{X} = \mathbf{I}_k\}$, where $\mathbf{I}_k$ is the $k \times k$ identity matrix. Our objective was to find $\mathbf{Q}_{\text{Orth-Obs}}$ and $\mathbf{Q}_{\text{Orth-Ex}}$ such that the observation and execution variance, respectively, were maximized. To solve this optimization problem, we let $\mathbf{C}_{\text{Obs}}$ and $\mathbf{C}_{\text{Ex}}$ denote the covariance matrices for observation and execution activity, respectively. We let λ_i^{Obs} and λ_i^{Ex} denote the i^{th} eigenvalue of $\mathbf{C}_{\text{Obs}}$ and $\mathbf{C}_{\text{Ex}}$ respectively, where the eigenvalues are in sorted order, i.e., $\lambda_i \geq \lambda_j$ for $i \leq j$. Thus, we solved the optimization problem:

$$\underset{\mathbf{Q} \in M_{d_{\text{Obs}}+d_{\text{Ex}}}(\mathbb{R}^N)}{\text{maximize}} \; \frac{\text{Tr}\left(\mathbf{Q}_{\text{Orth-Obs}}^T \mathbf{C}_{\text{Obs}} \mathbf{Q}_{\text{Orth-Obs}}\right)}{\sum_{i=1}^{d_{\text{Obs}}} \lambda_i^{\text{Obs}}} + \frac{\text{Tr}\left(\mathbf{Q}_{\text{Orth-Ex}}^T \mathbf{C}_{\text{Ex}} \mathbf{Q}_{\text{Orth-Ex}}\right)}{\sum_{i=1}^{d_{\text{Ex}}} \lambda_i^{\text{Ex}}} \quad \text{(Equation 9)}$$

where $\mathbf{Q} = [\mathbf{Q}_{\text{Orth-Obs}}, \mathbf{Q}_{\text{Orth-Ex}}] \in M_{d_{\text{Obs}}+d_{\text{Ex}}}(\mathbb{R}^N)$. This optimization therefore finds mutually orthogonal subspaces, $\mathbf{Q}_{\text{Orth-Obs}}$ and $\mathbf{Q}_{\text{Orth-Ex}}$, that maximize the proportion of variance captured in the observation and execution subspaces respectively. As in the text, we note that while $\mathbf{Q}_{\text{Orth-Obs}}$ and $\mathbf{Q}_{\text{Orth-Ex}}$ are constrained to be orthogonal, it need not be the case that $\mathbf{Q}_{\text{Orth-Obs}}^T \mathbf{C}_{\text{Ex}} \mathbf{Q}_{\text{Orth-Obs}} = 0$ and vice versa, since $\mathbf{C}_{\text{Ex}}$ has variance in dimensions beyond the top d_{Ex} dimensions. Rather, this quantity is a measure of the degree to which observation and execution activity are orthogonal. The four quantities reported (observation or execution) variance projected onto (observation or execution) subspaces were:

(a) Observation on observation : $\frac{\text{Tr}\left(\mathbf{Q}_{\text{Orth-Obs}}^T \mathbf{C}_{\text{Obs}} \mathbf{Q}_{\text{Orth-Obs}}\right)}{\sum_{i=1}^{d_{\text{Obs}}} \lambda_i^{\text{Obs}}}$

(b) Observation on execution : $\frac{\text{Tr}\left(\mathbf{Q}_{\text{Orth-Ex}}^T \mathbf{C}_{\text{Obs}} \mathbf{Q}_{\text{Orth-Ex}}\right)}{\sum_{i=1}^{d_{\text{Ex}}} \lambda_i^{\text{Obs}}}$

(c) Execution on observation : $\frac{\text{Tr}\left(\mathbf{Q}_{\text{Orth-Obs}}^T \mathbf{C}_{\text{Ex}} \mathbf{Q}_{\text{Orth-Obs}}\right)}{\sum_{i=1}^{d_{\text{Obs}}} \lambda_i^{\text{Ex}}}$

(d) Execution on execution : $\frac{\text{Tr}\left(\mathbf{Q}_{\text{Orth-Ex}}^T \mathbf{C}_{\text{Ex}} \mathbf{Q}_{\text{Orth-Ex}}\right)}{\sum_{i=1}^{d_{\text{Ex}}} \lambda_i^{\text{Ex}}}$

Optimization 2: Exclusive subspaces

In our second optimization, we sought to find a subspace $\mathbf{Q}_{\text{Excl}-\text{Ex}} \in M_{d_{\text{Ex}}}(\mathbb{R}^N)$, that maximized the variance of one context subject to a constraint that the variance in the other context was less than a small number, v. By maximizing the variance of one context (e.g., observation) while constraining the other (e.g., execution), this quantifies an exclusive subspace describing variance in only one context. For example, an exclusive subspace for observation activity would capture observation variance but negligible execution variance. Thus, we sought to solve the optimization problem:

$$\underset{\mathbf{Q}_{\text{Excl}-\text{Ex}} \in M_{d_{\text{Ex}}}(\mathbb{R}^N)}{\text{maximize}} \quad \frac{\text{Tr}\left(\mathbf{Q}_{\text{Excl}-\text{Ex}}^T \mathbf{C}_{\text{Ex}} \mathbf{Q}_{\text{Excl}-\text{Ex}}\right)}{\sum_{i=1}^{d_{\text{Ex}}} \lambda_i^{\text{Ex}}}$$

$$\text{subject to} \quad \frac{\text{Tr}\left(\mathbf{Q}_{\text{Excl}-\text{Ex}}^T \mathbf{C}_{\text{Obs}} \mathbf{Q}_{\text{Excl}-\text{Ex}}\right)}{\sum_{i=1}^{d_{\text{Ex}}} \lambda_i^{\text{Obs}}} \leq v$$

In Figure 4B, we show results where $v = 0.01$. The objective function and inequality constraint are convex; therefore, the constraint can be incorporated as a log barrier and the overall objective remains a convex optimization problem (Boyd and Vandenberghe, 2004). Hence, we solved the following optimization problem:

$$\underset{\mathbf{Q}_{\text{Excl}-\text{Ex}} \in M_{d_{\text{Ex}}}(\mathbb{R}^N)}{\text{maximize}} \quad \frac{\text{Tr}\left(\mathbf{Q}_{\text{Excl}-\text{Ex}}^T \mathbf{C}_{\text{Ex}} \mathbf{Q}_{\text{Excl}-\text{Ex}}\right)}{\sum_{i=1}^{d_{\text{Ex}}} \lambda_i^{\text{Ex}}} + \frac{1}{t} \log\left(v - \frac{\text{Tr}\left(\mathbf{Q}_{\text{Excl}-\text{Ex}}^T \mathbf{C}_{\text{Obs}} \mathbf{Q}_{\text{Excl}-\text{Ex}}\right)}{\sum_{i=1}^{d_{\text{Ex}}} \lambda_i^{\text{Obs}}}\right) \qquad \text{(Equation 10)}$$

where $t > 0$ is a hyperparameter. Note that as $x \to 0^+$, $\log(x) \to -\infty$. That is, the log-barrier function will become infinite when the inequality constraint is not satisfied. This log function thus incorporates a "barrier" into the objective function. As t grows larger, the log-barrier will asymptotically have a value of 0 if the inequality constraint is satisfied, and $-\infty$ if it is violated. Therefore, as long as the optimization is initialized with a parameter setting in the feasible set, the output of the optimization will remain feasible. To initialize the optimization, we choose $\mathbf{Q}_{\text{Excl-Ex}}$ to be the d_{Ex} eigenvectors corresponding to the smallest eigenvalues of $\mathbf{C}_{\text{Obs}}$. This subspace $\mathbf{Q}_{\text{Excl-Ex}}$ minimizes the variance of data with covariance $\mathbf{C}_{\text{Obs}}$. If the initial $\mathbf{Q}_{\text{Excl-Ex}}$ lies in the feasible set, we proceed with the optimization. However, sometimes even this $\mathbf{Q}_{\text{Excl-Ex}}$ does not satisfy the constraint; in this case, the optimization problem has no solution.

We repeatedly annealed the value of t, so that $t \to (1 + \beta)t$ for $\beta \in (0, 1)$ at each step. We chose $\beta = 0.01$. We iteratively solved the optimization problem until convergence. We note that the variance captured will depend on the time length of the neural population activity used to find each subspace. In the main text, we found the observation and execution exclusive subspaces from neural activity from 100 ms to 1000 ms after target onset. We chose this interval, in part, because we wanted to quantify observation and execution subspaces for the entire duration of the trial. We obtain the same conclusions if we only analyze activity from 200 ms to 500 ms after target onset.

Optimization 3: Shared subspaces

We sought to find a shared subspace that was orthogonal to the exclusive subspace defined in optimization 2. Intuitively, this optimization thus asks: in dimensions orthogonal to the observation and execution exclusive subspaces, what subspace maximizes the variance of both observation and execution activity? This definition presumes that no shared computation occurs in exclusive subspaces. Hence, we sought to find a shared subspace $\mathbf{Q}_{\text{Shared}} \in M_{d_{\text{Shared}}}(\mathbb{R}^N)$ that maximized the sum of projected observation and execution variance, subject to the constraint that $\mathbf{Q}_{\text{Shared}}$ is orthogonal to the observation and execution exclusive subspaces.

$$\underset{\mathbf{Q}_{\text{Shared}} \in M_{d_{\text{Shared}}}(\mathbb{R}^N)}{\text{maximize}} \quad \frac{\text{Tr}\left(\mathbf{Q}_{\text{Shared}}^T \mathbf{C}_{\text{Ex}} \mathbf{Q}_{\text{Shared}}\right)}{\sum_{i=1}^{d_{\text{Shared}}} \lambda_i^{\text{Ex}}} + \frac{\text{Tr}\left(\mathbf{Q}_{\text{Shared}}^T \mathbf{C}_{\text{Obs}} \mathbf{Q}_{\text{Shared}}\right)}{\sum_{i=1}^{d_{\text{Shared}}} \lambda_i^{\text{Obs}}}$$

$$\text{subject to} \quad \mathbf{Q}_{\text{Shared}} \perp \mathbf{Q}_{\text{Excl}-\text{Obs}}$$
$$\mathbf{Q}_{\text{Shared}} \perp \mathbf{Q}_{\text{Excl}-\text{Ex}}$$

where $\mathbf{Q}_{\text{Shared}}$ is the shared subspace, while $\mathbf{Q}_{\text{Excl-Obs}}$ and $\mathbf{Q}_{\text{Excl-Ex}}$ are the exclusive subspaces for observation and execution, respectively. In order to fit this problem into the manopt framework, we created a new, related optimization problem. We first introduced a subspace $\mathbf{Q}_{\text{Excl0}} = \textbf{orth}([\mathbf{Q}_{\text{Excl-Obs}}, \mathbf{Q}_{\text{Excl-Ex}}])$, where **orth**() finds the orthonormal basis of a subspace. This $\mathbf{Q}_{\text{Excl0}}$ spans the subspace $[\mathbf{Q}_{\text{Excl-Obs}}, \mathbf{Q}_{\text{Excl-Ex}}]$ and has dimension $d_{\text{Excl0}} \leq d_{\text{Excl-Ex}} + d_{\text{Excl-Obs}}$, since there is no guarantee from optimization 2 that $\mathbf{Q}_{\text{Excl-Obs}}$ and $\mathbf{Q}_{\text{Excl-Ex}}$ are orthogonal. Our requirement is that $\mathbf{Q}_{\text{Shared}} \perp \mathbf{Q}_{\text{Excl0}}$. Thus, we solve the following optimization problem to find a $\mathbf{Q} = [\mathbf{Q}_{\text{Excl}}, \mathbf{Q}_{\text{Shared}}] \in M_{d_{\text{Excl0}} + d_{\text{Shared}}}(\mathbb{R}^N)$ instead, as follows:

$$\underset{\mathbf{Q} = [\mathbf{Q}_{\text{Excl}}, \mathbf{Q}_{\text{Shared}}] \in M_{d_{\text{Excl0}} + d_{\text{Shared}}}(\mathbb{R}^N)}{\text{maximize}} \quad \text{Tr}\left(\mathbf{Q}_{\text{Excl}}^T \mathbf{Q}_{\text{Excl0}}\right) + \frac{\text{Tr}\left(\mathbf{Q}_{\text{Shared}}^T \mathbf{C}_{\text{Ex}} \mathbf{Q}_{\text{Shared}}\right)}{\sum_{i=1}^{d_{\text{Shared}}} \lambda_i^{\text{Ex}}} + \frac{\text{Tr}\left(\mathbf{Q}_{\text{Shared}}^T \mathbf{C}_{\text{Obs}} \mathbf{Q}_{\text{Shared}}\right)}{\sum_{i=1}^{d_{\text{Shared}}} \lambda_i^{\text{Obs}}}. \qquad \text{(Equation 11)}$$

Since $\mathbf{Q}$ is optimized over the Stiefel manifold, by construction $\mathbf{Q}_{\text{Excl}} \perp \mathbf{Q}_{\text{Shared}}$ and the objective function forces $\mathbf{Q}_{\text{Excl}} \approx \mathbf{Q}_{\text{Excl0}}$.

Neuron subspace contribution

As in the previous section, we define a subspace $\mathbf{Q} \in M_d(\mathbb{R}^N)$, with its i^{th} row being a row vector $\mathbf{w}_i \in \mathbb{R}^{1\times d}$. We define the neuron activity matrix to be $\mathbf{X} \in R^{N\times TC}$. Further, we define $\mathbf{X}_i \in R^{N\times TC}$ as neuron i's activity, by setting every entry of $\mathbf{X}$ to zero except for i^{th} row. We define neuron i's "subspace contribution," c_i , to this subspace $\mathbf{Q}$ to be the square root of the variance of the projected neuron's activity in the subspace divided by the neuron's total variance, times the neuron's mean firing rate (in both observation and execution contexts), which we denote *FR*. This equivalent to taking the magnitude of the weights and multiplying it by the firing rate, which has previously been used to quantify neuron contributions to a projection (Nuyujukian et al., 2014).

$$c_i = \sqrt{\frac{\text{Tr}\left(\mathbf{Q}^T\mathbf{X}_i\mathbf{X}_i^T\mathbf{Q}\right)}{\text{Tr}\left(\mathbf{X}_i\mathbf{X}_i^T\right)}}\cdot FR = \sqrt{\frac{\text{Tr}\left(\mathbf{X}_i\mathbf{X}_i^T\mathbf{Q}\mathbf{Q}^T\right)}{\text{Tr}\left(\mathbf{X}_i\mathbf{X}_i^T\right)}}\cdot FR = \sqrt{\frac{var(x_i)\|\mathbf{w}_i\|^2}{var(x_i)}}\cdot FR = \|\mathbf{w}_i\|\cdot FR \quad \text{(Equation 12)}$$

where $var(x_i)$ is neuron i's total variance, and $\|\mathbf{w}_i\|^2$ is the squared sum of $\mathbf{Q}$'s i^{th} row. The reason for multiplying by FR is that neurons having small firing rate should have small contribution to a subspace, even though its activity might be aligned with this subspace, as described in Nuyujukian et al. (2014).

Variance explained for neuron subclasses

We performed the optimizations prior described with incongruent and congruent neuron populations. In these analyses, we restricted the subspace optimization to congruent or incongruent neurons. We note that because the number of congruent and incongruent neurons are not necessarily equalized, to bootstrap, we performed subspace optimization multiple times, each time sampling equal amounts of incongruent or congruent neurons with replacement. Concretely, we sampled 0.9 times the minimum number of congruent neurons and incongruent neurons. It is worth noting that due to a low number of neurons, we were sometimes unable to find a subspace that satisfied the exclusive subspace optimization constraint. If this was the case, we were unable to use this dataset for this analysis and therefore excluded it. In our analyses, dataset L140829 was excluded for this reason.

Quantifying dynamics

In this work, we quantified the degree of tangling and the degree of rotational dynamics in the neural population. All results were analyzed using a time window 200 ms to 500 ms after target onset, during which the monkey was either reliably executing or observing movement. Both analyses followed methods developed in prior work. We performed tangling analysis, as described by Russo et al. (2018). Tangling quantifies how consistent the velocity of the neural trajectory is at a location in state space.

$$Q(t) = \max_{t'} \frac{\|\dot{\mathbf{x}}_t - \dot{\mathbf{x}}_{t'}\|^2}{\|\mathbf{x}_t - \mathbf{x}_{t'}\|^2 + \varepsilon} \quad \text{(Equation 13)}$$

where $\mathbf{x}_t$ is the neural trajectory at time t, $\dot{\mathbf{x}}_t$ is its temporal derivative. $\|\cdot\|$ is the Euclidean norm, and ε is a small constant to prevent dividing by zero. If tangling is larger, then dynamics are less consistent. In this work, the subspaces for tangling analysis are all 4D, consistent with the dimensionality for subspace optimizations. When comparing tangling of the congruent versus incongruent subpopulations, we used PCA to reduce dimensionality to 4D.

Second, we performed jPCA analysis, as described by Churchland et al. (2012). jPCA finds a projection of the neural activity that captures rotational structure in the data. We first reduced the dimensionality of data to 6 using PCA. We reported coefficients of determination in predicting future neural activity using only the current neural activity via a dynamics matrix learned via jPCA

$$\dot{\mathbf{X}}_{\text{red}} = \mathbf{M}_{\text{skew}}\mathbf{X}_{\text{red}} \quad (14)$$

or linear regression,

$$\dot{\mathbf{X}}_{\text{red}} = \mathbf{M}\mathbf{X}_{\text{red}}. \quad (15)$$

For statistical tests comparing jPCA R^2 between congruent and incongruent neurons, we used the same bootstrap used to calculate PD difference, but on z-scored R^2_{skew}.

QUANTIFICATION AND STATISTICAL ANALYSIS

Data are presented as mean, mean ± 95% confidence interval, mean ± SD, or median ± 25% and 75% percentile, based on analysis and is indicated in the figure legends or main text. Single neuron statistical analysis are described in the "Single neuron firing rate statistics" Method section. Bootstrap tests were performed by performing a shuffle analysis 1000 times with replacement. Hartigan's dip test was used to perform a bimodality test. We used the Wilcoxon rank-sum test was used to compare difference in tangling. Statistically significant differences were tested at $p < 0.05$, unless otherwise noted.