

Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control

Highlights

- In early training, large uncorrelated neural variance produces variable cursor control
- Trial-to-trial uncorrelated neural and cursor variability decrease with training
- Task-relevant neural covariance increases and consolidates over training
- Consistent neural trajectories with task-relevant covariance produce skilled control

Authors

Vivek R. Athalye, Karunesh Ganguly, Rui M. Costa, Jose M. Carmena

Correspondence

ruicosta@neuro.fchampalimaud.org (R.M.C.),
jcarmena@berkeley.edu (J.M.C.)

In Brief

Athalye et al. ask how task-relevant neural populations coordinate to acquire activity patterns underlying novel skills. As subjects learn to control a brain-machine interface de novo, each neuron initially explores independently, and then coordinated activity patterns are consolidated, driving skillful performance.



Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control

Vivek R. Athalye,^{1,6} Karunesh Ganguly,^{4,5} Rui M. Costa,^{6,7,*} and Jose M. Carmena^{1,2,3,8,*}

¹Department of Electrical Engineering and Computer Sciences

²Helen Wills Neuroscience Institute

³UC Berkeley/UCSF Joint Graduate Program in Bioengineering
University of California, Berkeley, Berkeley, CA 94720, USA

⁴Neurology and Rehabilitation Services, San Francisco VA Medical Center, San Francisco, CA 94121, USA

⁵Department of Neurology, University of California, San Francisco, San Francisco, CA 94143, USA

⁶Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown, Avenida de Brasília, Doca de Pedrouços, Lisbon 1400-038, Portugal

⁷Department of Neuroscience, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY 10032, USA

⁸Lead Contact

*Correspondence: rui.costa@neuro.fchampalimaud.org (R.M.C.), jcarmena@berkeley.edu (J.M.C.)

<http://dx.doi.org/10.1016/j.neuron.2017.01.016>

SUMMARY

During motor learning, movements and underlying neural activity initially exhibit large trial-to-trial variability that decreases over learning. However, it is unclear how task-relevant neural populations coordinate to explore and consolidate activity patterns. Exploration and consolidation could happen for each neuron independently, across the population jointly, or both. We disambiguated among these possibilities by investigating how subjects learned *de novo* to control a brain-machine interface using neurons from motor cortex. We decomposed population activity into the sum of private and shared signals, which produce uncorrelated and correlated neural variance, respectively, and examined how these signals' evolution causally shapes behavior. We found that initially large trial-to-trial movement and private neural variability reduce over learning. Concomitantly, task-relevant shared variance increases, consolidating a manifold containing consistent neural trajectories that generate refined control. These results suggest that motor cortex acquires skillful control by leveraging both independent and coordinated variance to explore and consolidate neural patterns.

INTRODUCTION

When we begin learning a novel skill, we have no experience about how to move to achieve a particular goal. Early in training, movements exhibit variability from attempt to attempt, allowing the motor system to explore actions and select them based on consequences (Sutton and Barto, 1998; Tumer and Brainard, 2007; Wu et al., 2014). Gradually, movement variability de-

creases as the motor system consolidates particular movements that lead to success (Cohen and Sternad, 2009; Shmuelof et al., 2012). This suggests that a task-relevant neural population explores neural activity patterns to generate novel motor commands, and selects and consolidates those that achieve the desired goal. Indeed, neurophysiological motor learning studies have found that movement-related neural activity in various species follows a similar trend, exhibiting high variability in early training that reduces as particular ensembles and activity patterns are consolidated in late training (Arduin et al., 2013; Barnes et al., 2005; Cao et al., 2015; Costa et al., 2004; Kao et al., 2005; Mandelblat-Cerf et al., 2009; Peters et al., 2014; Santos et al., 2015; Zacksenhouse et al., 2007). However, it is unclear how exploration of particular patterns in a task-relevant neural population is accomplished, and also how goal-achieving patterns are refined and consolidated to improve performance.

The motor system is distributed and redundant, with parallel pathways and many more neurons than muscles, making it difficult to investigate the causal relationship between observed neural activity patterns and motor output (Briggman and Kristan, 2008). We therefore took advantage of a paradigm in which we could identify the output neurons that control behavior as well as the explicit transformation between output neuron activity and behavior. We used an operant learning brain-machine interface (BMI) in which stable recordings from ensembles of primary motor cortex (M1) neurons in macaque monkeys are input to a fixed mathematical transform ("decoder") outputting prosthetic two-dimensional cursor movements. The BMI provided a closed-loop feedback system operating within the natural motor system (Figure 1A). Because we defined the decoder to take input from observed cells ("direct cells"), this neuroprosthetic skill learning paradigm (Ganguly and Carmena, 2009, 2010) is uniquely advantageous to investigate how a task-relevant neural population coordinates to acquire skillful control.

Given that the decoder is novel, subjects must initially explore neural population activity space. This is a non-trivial search because population activity space is high dimensional, with

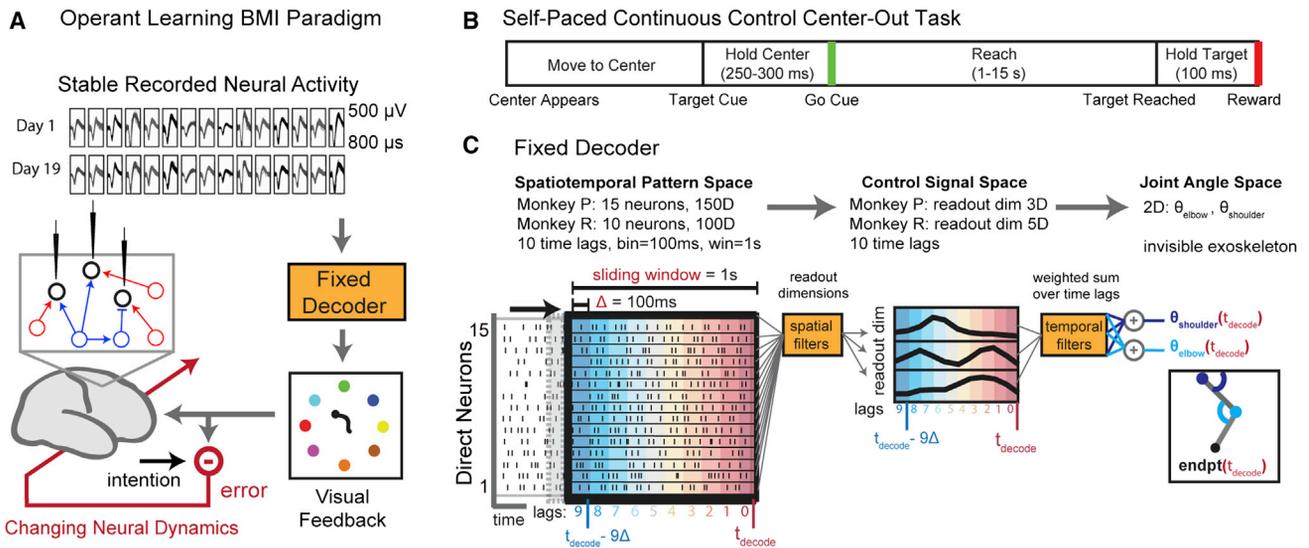


Figure 1. Experimental Paradigm

(A) The operant learning BMI paradigm. Stably recorded units are selected (direct cells) to provide input to a decoder with fixed parameters over training. (B) Center-out task timeline. See STAR Methods for details. (C) The task uses a position decoder, which maps a neural spatiotemporal pattern into a low-dimensional control signal that produces the cursor endpoint.

each neuron's activity represented as one dimension. Over training, subjects must use the behavioral consequences of explored activity patterns to select and refine goal-achieving patterns. This consolidation process is not well understood, especially as there are many neural activity patterns that can elicit the same behavioral output in systems mapping many neurons to fewer outputs, such as a BMI or our natural motor system. How might this learning happen? We consider three possibilities.

In “independent neuron learning,” each neuron explores independently, as has been proposed in computational models (Hélot et al., 2010; Legenstein et al., 2010), and each neuron tunes its independent patterns over learning (Figure 3A). Because independent neuron learning explores every dimension of population activity space, this learning can acquire any firing rate configuration and thus control any decoder. However, this learning cannot explain how subjects rapidly acquire (within 1 day) decoders that leverage neural covariance (Sadtler et al., 2014).

In contrast, “constrained-network learning” posits that the neural population possesses network connectivity constraints and thus preferentially produces covariation patterns within a fixed, lower-dimensional subspace of the full activity space. Indeed, both BMI and motor control studies have found that motor cortical neurons exhibit task-relevant, low-dimensional covariance during execution of well-learned behavior (Churchland et al., 2012; Kao et al., 2015; Sadtler et al., 2014). Under constrained-network learning, the neural population explores and changes covariation patterns within the fixed subspace (Sadtler et al., 2014) (Figure 3B). One limitation of the constrained exploration is that this learning can only acquire decoders reading out activity within the population's existing covariance subspace.

Finally, we consider “flexible-network learning,” which both permits the learning of many decoders and can explain the emer-

gence of new covariance. In this model, the neural population explores and changes both independent and covariation patterns. Critically, independent exploration permits subjects to initially explore a higher dimensional space of activity patterns. Upon discovering patterns with rewarding consequences, subjects may shape their network to reliably reproduce rewarding patterns, constraining patterns to lie within a low-dimensional subspace. This mechanism thus enables consolidation of covariance in a different subspace from initial covariance (Figure 3C).

To determine which of these three possibilities best explains how the task-relevant neural population coordinates to acquire neuroprosthetic skill de novo, we sought to disentangle how independent and coordinating sources of neural variance changed and contributed to BMI control over the course of learning. We used Factor Analysis (FA) (Everitt, 1984) to decompose population activity into the sum of private and shared signals, which estimate the effects of both private inputs to each cell and coordinating inputs that drive multiple cells simultaneously (Churchland et al., 2010; Yu et al., 2009). Using FA, we arbitrated between our three learning models in addressing how a task-relevant neural population explores, consolidates, and optimizes activity to acquire a neuroprosthetic skill de novo.

First, we asked how private and shared sources influence the exploration of different activity patterns from trial to trial. To isolate trial-to-trial variability from within-trial variance relevant for control, we focused on the variability of the population firing rate at coarse timescale (time window ~ 1 s) for a fixed cursor state (the center of the workspace) with a fixed movement intention (the target). Second, we asked how private and shared sources contribute to within-trial control over the course of learning. To assess control-relevant neural activity, we focused on spike counts at the decoder timescale of 100 ms. This decoder-timescale neural variance can be interpreted as signal strength for

BMI control. Finally, we asked how the spatial and temporal structure of late training neural activity supported skillful control. In particular, we asked whether neural activity preferentially occupied the control-relevant dimensions of the decoder, and whether a particular temporal sequence was consolidated that generated skillful control. We focused on fine-timescale trajectories at 50 ms resolution, faster even than the decoder timescale.

Exploration and consolidation via independent neuron learning predict that private trial-to-trial variability decreases and private decoder-timescale variance is used for control. Constrained-network learning predicts the opposite; shared trial-to-trial variability decreases and shared decoder-timescale variance is used for control. Combining these changes under flexible-network learning, we would expect that in early learning, private trial-to-trial variability contributes to activity exploration, and private decoder-timescale variance produces variable movements, while we would expect that in late learning, shared decoder-timescale variance consolidates and contributes to skillful control. Indeed, our analyses revealed that private and shared variance changed differentially to sculpt neural spatio-temporal patterns and improve movements, providing insight into how motor cortex coordinates to explore, consolidate, and optimize population activity underlying novel skills.

RESULTS

Neural Control Refinement Increases Movement Success Rate, Directness, and Consistency

We analyzed data from operant learning BMI experiments (Ganguly and Carmena, 2009) in which two rhesus macaques chronically implanted with microelectrode arrays in M1 learned to perform a two-dimensional, self-paced, continuous control center-out BMI reaching task (Figures 1A and 1B). Each trial, subjects drove the cursor under neural control to an instructed target (randomly selected from eight possible targets) for juice reward and controlled the cursor back to center to initiate the next trial.

The decoder was optimized to predict upper limb movement using input from a fixed subset of recorded neurons (“direct cells”) chosen because of their stability over days (see STAR Methods) (Ganguly and Carmena, 2009). Because both the direct cells and the decoder parameters were held fixed for the duration of training (Monkey P, 19 days; Monkey R, 7 days), performance improvement could only be driven by adaptive changes in population activity (Figure 1A).

The decoder transformed a neural activity sequence into the cursor’s x-y position (Figure 1C). The direct cells’ activity in the last 1 s was binned into 10 time lags of 100 ms bins, and linear combinations of the neuron-lag activity produced the cursor’s position. Because the number of neurons exceeded the two cursor dimensions (Monkey P, 15 neurons; Monkey R, 10 neurons), (infinitely) many activity patterns could produce the same behavior, as is the case in our motor system, which possesses more neurons than muscles. Changes in neural activity did not translate into changes in cursor position unless they occurred within the decoder’s control-relevant neural dimensions, the “readout space,” defined by the decoder weights.

Under these conditions, subjects gained proficient control, improving both accuracy and success rate for all targets (Figures 2A and 2B). Changes over learning were analyzed over “training epochs,” where each epoch contains a constant number of trials. This division of trials ensured that analysis results were not influenced by unequal numbers of trials performed on each day. Observed trends were consistent with analyses performed across days (Figures S2 and S3, available online). Note that analyses of behavioral and neural changes were performed for each target separately. All trends are plotted with error bars showing the SEM over targets.

We asked how control was refined on a trial-to-trial basis by analyzing the positions the cursor entered in each trial. We computed the cursor’s trial-to-trial workspace occupancy for each individual target, i.e., the probability that the cursor entered a given position within a trial (Figure 2C). We found that the fraction of workspace entered per trial decreased as the cursor movements became increasingly direct (Figure 2D), and that the subjects’ workspace occupancy became consistent over late training epochs (Figure 2E) (see STAR Methods). These two results show that the cursor’s trial-to-trial exploration of the workspace became refined and stable with training, extending previous reports on the straightening and stabilization of trial-averaged trajectories (Ganguly and Carmena, 2009). Finally, we found a decrease in the trial-to-trial variability of the entire set of positions entered within a trial (Figure 2F) (see STAR Methods). These results show that neural control was refined and consolidated over long-term BMI training, increasing movement success rate, directness, and consistency.

Modeling Neural Variance Changes over Learning with FA

Having quantified control refinement over training, we sought to model how the direct cells coordinated to explore and change neural activity patterns underlying control. Under what conditions would independent patterns (described by private variance; Figure 3A) or covariation patterns (described by shared variance; Figure 3B) be preferred for generating movement variability? When a subject has no idea which neural readout dimensions the decoder uses, private variance is beneficial because this high-dimensional exploration would generate variability in any readout space (Figure S1A). Low-dimensional shared variance might be misaligned with the readout space and thus produce little movement variability (Figure S1B). On the other hand, when the subject has experience with the readout space, low-dimensional shared variance aligned with the readout space is more efficient in generating movements because it concentrates more variance in fewer neural dimensions (Figure S1C) and leads to faster learning (Sadtler et al., 2014). Thus, flexible-network learning combines these properties, allowing the population to expansively explore using private variance and then consolidate control-specific shared variance. Importantly, exploring the behavioral consequences of activity outside the initial shared space could guide the shared space to rotate or change dimensionality (Figure 3C).

Seeking evidence for these neural activity changes, we used FA to decompose population firing rate activity into the sum of (1) a mean rate; (2) private signals, which drive each neuron

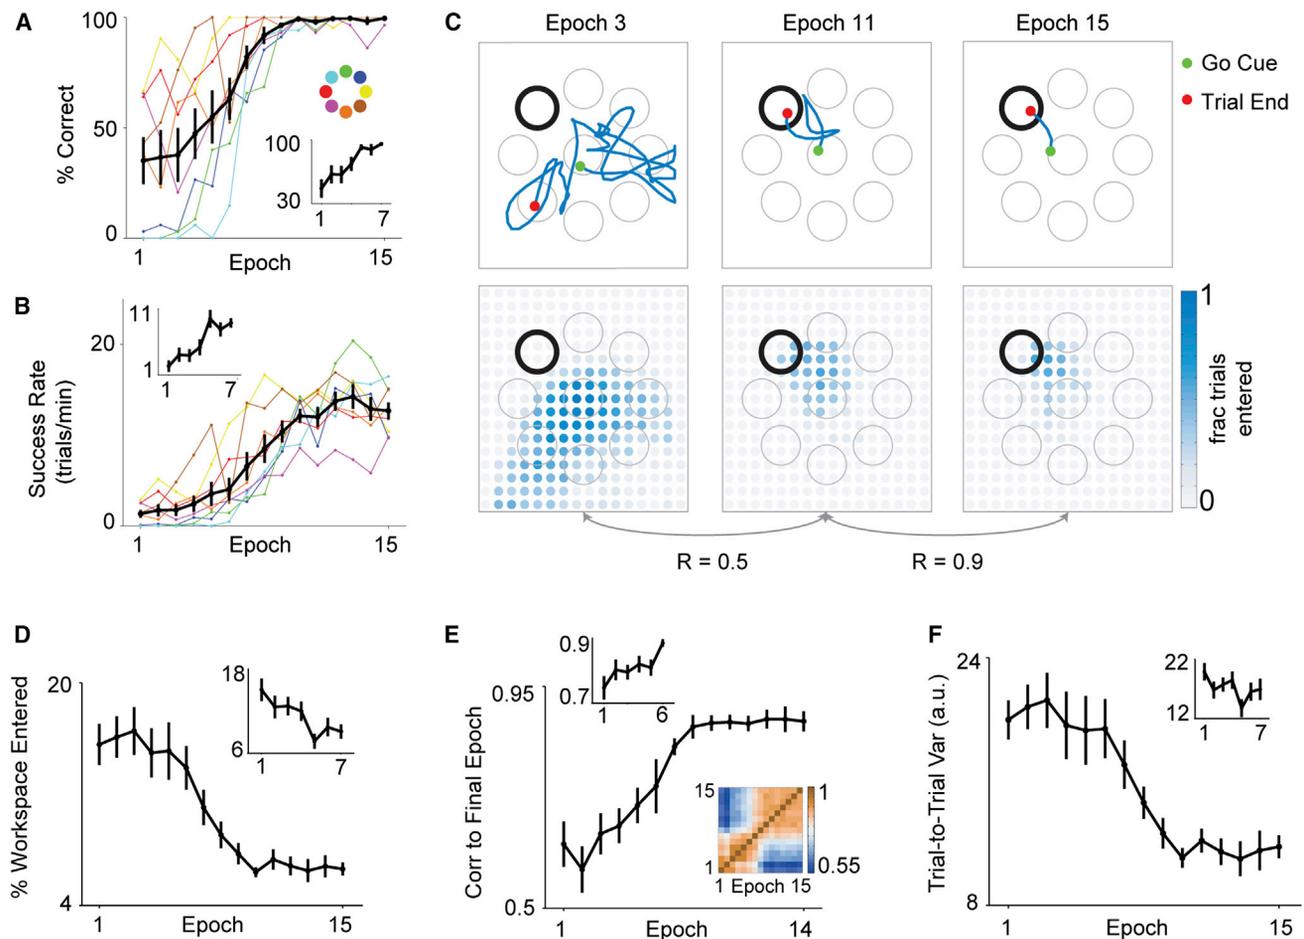


Figure 2. Cursor Control Is Refined

(A and B) Performance improved over training for Monkey P (main) and Monkey R (inset). (A) Percent correct increased (Monkey P, $r = 0.93$, $p = 5.8 \times 10^{-7}$, early < late $p = 3.3 \times 10^{-10}$; Monkey R, $r = 0.96$, $p = 4.7 \times 10^{-4}$, early < late $p = 1.5 \times 10^{-9}$). (B) Success rate in trials per minute increased (Monkey P, $r = 0.96$, $p = 1.6 \times 10^{-8}$, early < late $p = 1.4 \times 10^{-6}$; Monkey R, $r = 0.91$, $p = 4 \times 10^{-3}$, early < late $p = 4.0 \times 10^{-6}$). Targets (indicated by color) showed different time courses of improvement. Mean and SEM error bars are shown over targets ($n = 8$ for both subjects).

(C) (Top) Single-trial example trajectories over three epochs, for one example target. (Bottom) Workspace occupancy maps for the example target over three epochs. Intensity indicates the fraction of trials a position was occupied.

(D) Percent of workspace entered per trial decreased over training (Monkey P, $r = -0.94$, $p = 3.2 \times 10^{-7}$, early > late $p = 4.1 \times 10^{-9}$; Monkey R, $r = -0.87$, $p = 1 \times 10^{-2}$, early > late $p = 2.0 \times 10^{-4}$).

(E) Similarity of workspace occupancy to the final workspace map increased and stabilized with training (Monkey P, $r = 0.92$, $p = 2.3 \times 10^{-6}$, early < late $p = 5.9 \times 10^{-9}$; Monkey R, $r = 0.88$, $p = 2 \times 10^{-2}$, early < late $p = 1.2 \times 10^{-2}$). (Bottom inset) Heatmap shows the similarity between each epoch pair's occupancy maps. A heatmap was calculated for each target separately; the target average is shown here. The lower left red box indicates cursor scatter during gross exploration while the upper red box indicates late training stability of refined control.

(F) Trial-to-trial variability of workspace occupancy decreased (Monkey P, $r = -0.93$, $p = 4.6 \times 10^{-7}$, early > late $p = 1.7 \times 10^{-8}$; Monkey R, $r = -0.54$, $p = 2 \times 10^{-1}$, early > late $p = 2.4 \times 10^{-2}$).

without correlation (Figure 3E); and (3) low-dimensional shared signals, which drive multiple neurons simultaneously, produce correlations, and constrain activity to the shared space containing all possible shared signals (Figure 3F). Thus, FA models the population's total covariance matrix as the sum of the diagonal covariance matrix due to uncorrelated private signals and the low-rank covariance matrix due to shared signals, permitting variance only within the shared space (Figures 3D and 3G). Our use of private variance and shared variance specifically refers to their respective covariance matrices. We emphasize that FA

is fit for each target separately on the full simultaneously recorded direct cell population, not just neuron pairs as illustrated for geometric intuition (Figures 3E–3J). Model fitting assessed the shared space dimensionality (number of shared signals) needed to best describe the firing rates (see STAR Methods).

We used FA to analyze how the neural population's private and shared variances change over learning (Figure 3D). First, the magnitude of total variance can change (Figure 3H). Second, the balance of shared and private variance can change, a measure of coordination that we quantified with the shared-to-total

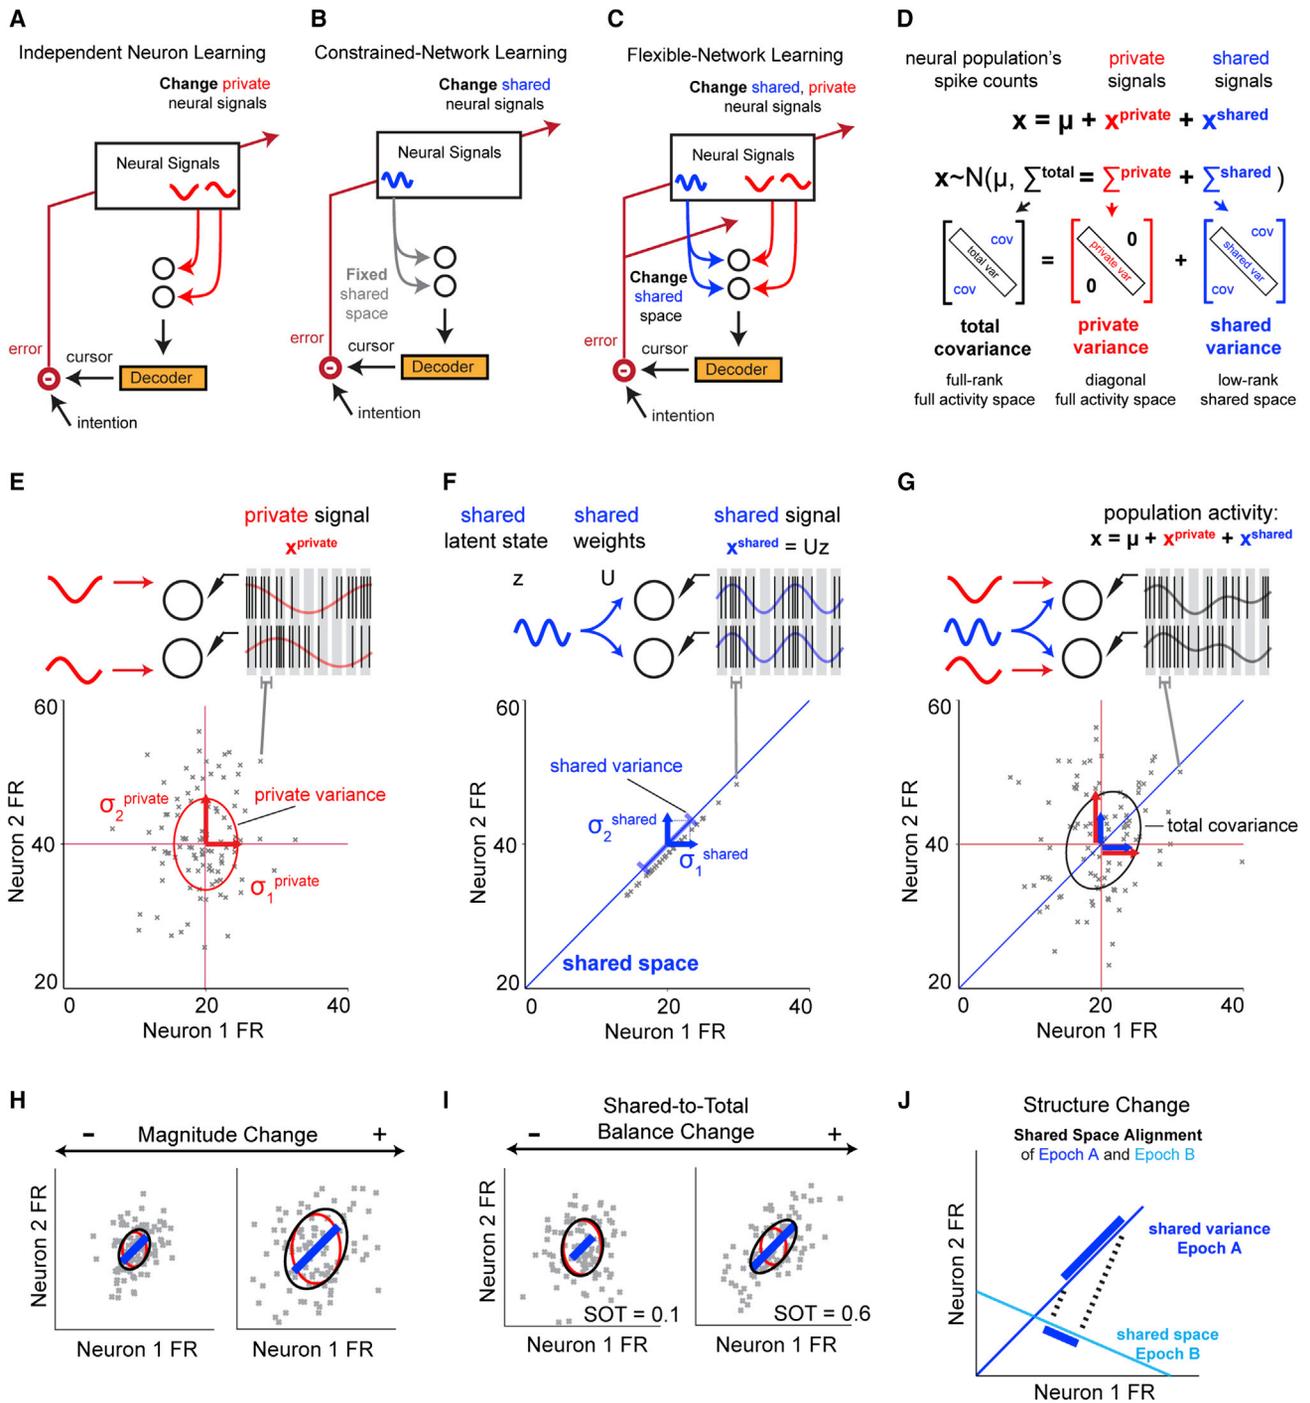


Figure 3. To Model Neural Changes over BMI Learning, FA Decomposes Population Spike Counts into the Sum of Private and Shared Sources of Variance

(A) Independent neuron learning posits that each neuron explores and acquires independent patterns to control the BMI.
 (B) Constrained-network learning posits that the neural population learns BMI control by exploring and changing covariation patterns in a fixed shared space, e.g., because of network connectivity constraints.
 (C) Flexible-network learning posits that the neural population learns BMI control by exploring and changing both independent and covariation patterns and by using its exploration to change the shared space of covariation patterns. For intuition of how independent and covariation patterns can contribute to movement variability, see Figure S1.
 (D) Factor Analysis (FA) decomposes the spike count covariance matrix into the sum of a diagonal private variance matrix and a low-rank shared variance matrix.

(legend continued on next page)

variance ratio (Figure 3I). Third, the shared space containing shared variance can change, which we quantified by calculating the alignment between two epochs' shared variances (Figure 3J; STAR Methods). The shared space alignment, ranging between 0 and 1, asks how well Epoch B's shared signals can approximate Epoch A's shared signals. If Epochs A and B have perfectly aligned shared spaces, then the alignment is 1, and if the shared spaces are orthogonal, the alignment is 0. Thus, if the shared space stays fixed over learning, as predicted by constrained-network learning, we expect alignment values to be high throughout training. In contrast, if the shared space consolidates with learning, as predicted by flexible-network learning, we expect high alignment values only in late learning. We used these three analyses to study how private and shared sources of variance contribute to the exploration and consolidation of activity patterns.

Over Learning, Private Trial-to-Trial Variability Decreases while the Shared Space Changes and Stabilizes

While previous studies have found trial-to-trial neural variability decreases for a particular intended action over motor learning, we used FA to ask how private and shared sources influenced trial-to-trial variability. We analyzed firing rate variability at the start of trials for each target separately. To avoid capturing within-trial variance used for control, we extracted one spike count vector from each trial in a large ~ 1 s window locked to Go Cue (Figure 4A, left). Subjects' trial times determined the window size (Monkey P, 0.9 s; Monkey R, 1.3 s; see STAR Methods for details), but the results were insensitive to window choice (Figures S4B–S4D). We assessed private trial-to-trial variability beyond that expected for a Poisson process by normalizing each neuron's private variance by its mean rate (see STAR Methods). Shared trial-to-trial variability was calculated with each neuron's un-normalized shared variance because in Poisson simulations, shared variance did not scale with the experimentally observed firing rates (Figure S5B). This analysis enabled us to disambiguate between independent neuron learning's private trial-to-trial variability decrease, constrained-network learning's shared trial-to-trial variability decrease, and flexible-network learning's private trial-to-trial variability decrease and shared space consolidation.

FA applied to a neuron pair in early and late training is shown for illustration (Figure 4A), showing a decrease in private trial-to-trial variability, an increase in shared-to-total variance ratio, and rotation of the shared space. At the direct cell population level, FA found only one shared dimension was needed for each target and epoch. Consistent with the example, we found initially large

private trial-to-trial variability, which decreased prominently over training (Figures 4B and S2A). In contrast, shared trial-to-trial variability showed no significant trend, resulting in an increase in the shared-to-total variance ratio (un-normalized by firing rate) over long-term training (Figures 4C and S2C). These results suggest that neural activity space exploration in early learning was primarily accomplished with private trial-to-trial variability. The decrease in private trial-to-trial variability was correlated with success rate over training (Figures 4D and S2B), indicating the subjects produced desired activity as private trial-to-trial variability reduced.

While the magnitude of shared trial-to-trial variability showed no change, its subspace consolidated over training. We assessed the shared space alignment between epoch pairs (Figure 4E, left; Figure S2D, left) and found the alignment with the final epoch's shared space increased with training (Figure 4E, right; Figure S2D, right), correlating with the subject's success rate (Figures 4F and S2E). Notably, the emergent shared space possessed a significantly different alignment from the initial shared space. Altogether, these results support flexible-network learning, showing that initial trial-to-trial variability is private to each neuron, and as private trial-to-trial variability decreased, the shared space evolved and consolidated.

Decoder-Timescale Shared Variance Strengthens and Consolidates to Coordinate Population Activity

Having analyzed coarse-timescale trial-to-trial variability, we next zoomed in temporal resolution and asked how private and shared sources contributed to control-relevant activity at the decoder timescale of 100 ms bins (decoder bin size) for each target separately. Decoder-timescale neural variance can be interpreted as signal strength for BMI control, in contrast to our previous analysis on trial-to-trial variability. We thus use FA to study if the independent neuron, constrained-network, or flexible-network model best describes how learning changes decoder-timescale population activity driving control.

We first observed population activity became lower dimensional, as the shared dimensionality describing the 100 ms spike counts within all trials to a single target decreased from ~ 4 to ~ 2 in late training (Figure 5B). Because the brain found a two-dimensional neural solution in late training, we compared shared variance over training in the two neural dimensions that captured the most shared variance, which we call the "main shared variance" (see STAR Methods). Note that high-dimensional shared variance did not constitute a significant fraction of total variance (Figure S5F).

FA applied to a neuron pair in early and late training is shown for illustration, in which shared variance increased and rotated

(E) Private signals modulating two neurons' firing rates. Each scattered gray "x" is a binned population spike count. Neuron 2 has a larger private variance than neuron 1.

(F) Shared signals modulating two neurons' firing rates within a one-dimensional shared space.

(G) FA describes the mixture of private and shared variance underlying population activity.

(H) FA detects changes in the total magnitude of variance.

(I) FA also detects changes in the shared-to-total variance ratio, quantifying how shared and private variance change in different proportions. The shared-to-total ratio plus the private-to-total ratio equals 1.

(J) FA detects changes in how the population co-varies using the shared space alignment between Epochs A and B, which is the fraction of Epoch A's shared variance captured in Epoch B's shared space.

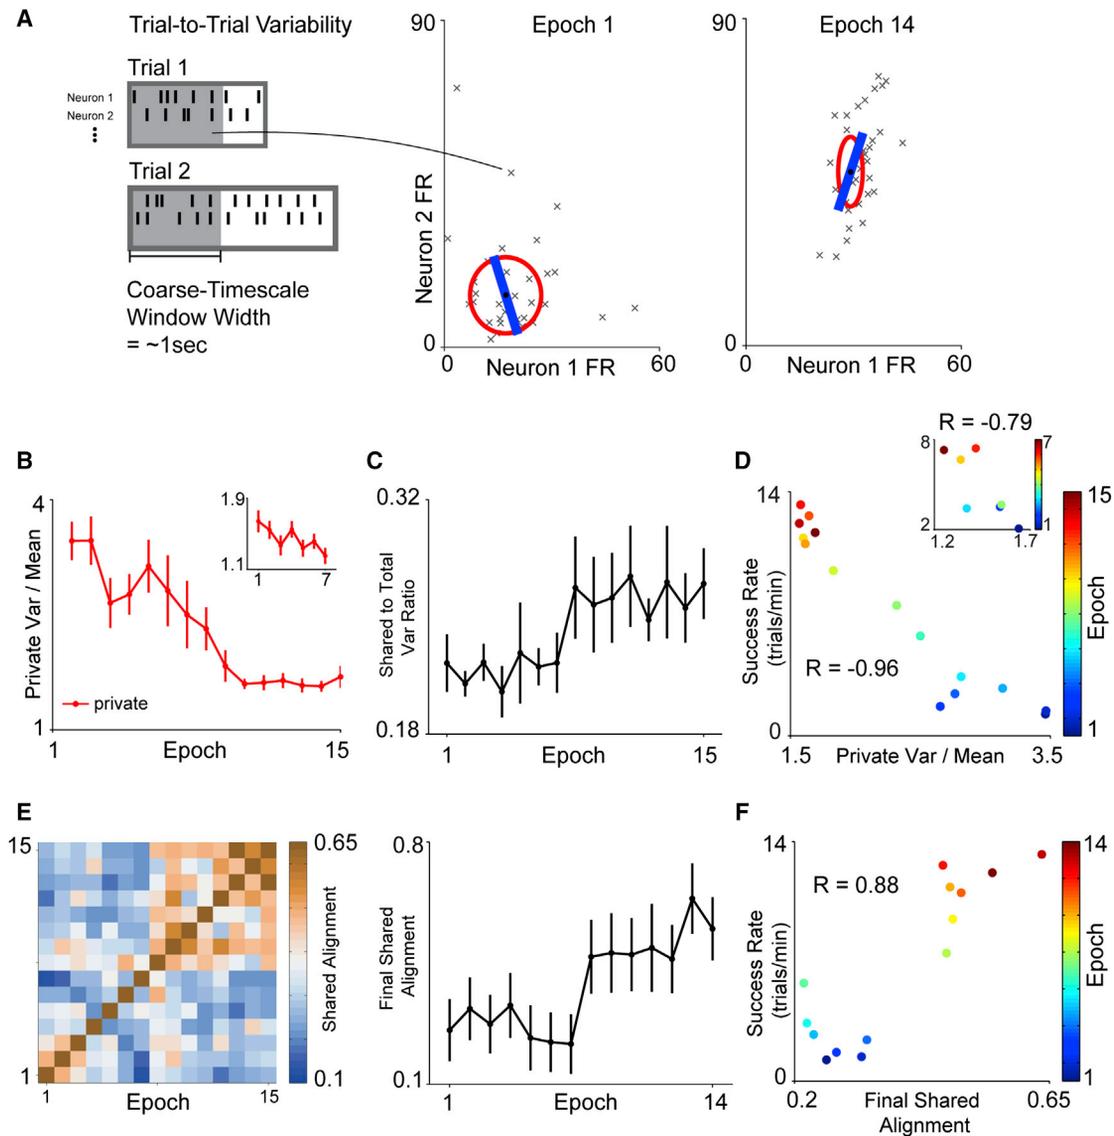


Figure 4. Private Trial-to-Trial Variability Decreases and the Shared Trial-to-Trial Variability Consolidates a Shared Space

(A) Trial-to-trial variability example: FA fit on two direct cells in early and late training for Monkey P on one example target.

(B) Private trial-to-trial variability for Monkey P (main) and Monkey R (inset) decreased (Monkey P, $r = -0.93$, $p = 6.7e-7$, early > late $p = 1.1e-8$; Monkey R, $r = -0.81$, $p = 2.7e-2$, early > late $p = 2.2e-2$).

(C) The shared-to-total variance ratio increased over long-term training (Monkey P, $r = 0.83$, $p = 1.1e-4$, early < late $p = 8.1e-3$), but not significantly over shorter training (Monkey R, $r = 0.27$, $p = 5.5e-1$, early < late $p = 8.2e-2$).

(D) The private trial-to-trial variability decrease correlated with success rate improvement (Monkey P, $r = -0.96$, $p = 1.9e-8$; Monkey R, $r = -0.79$, $p = 3.5e-2$).

(E) (Left) The shared space alignment map indicated consolidation. Each element is the alignment of the row epoch's shared variance with the column epoch's shared space. The upper right red block indicates that the shared space becomes relatively stable, while the preceding blue indicates the shared space is unstable in early training. (Right) The shared space alignment between the final epoch's shared space and the preceding epochs' shared variance increased with long-term training (Monkey P, $r = 0.83$, $p = 2.6e-4$, early < late $p = 5.8e-4$; Monkey R, $r = -0.12$, $p = 8.2e-1$, early > late $p = 4.9e-1$). The mean chance variance alignment is 0.07, and the 95th percentile chance alignment is 0.25 for Monkey P (see STAR Methods).

(F) The shared space alignment with the final epoch's shared space correlated with success rate over long-term training (Monkey P, $r = 0.88$, $p = 3.3e-5$; Monkey R, $r = 0.07$, $p = 9e-1$). Changes in mean firing rate do not explain these results (Figure S5). These results also held for analyses over training days rather than epochs (Figure S2) and different window lengths (Figure S4).

Error bars represent SEM over targets ($n = 8$ for each subject).

(Figure 5A). In concordance, in the entire direct cell population, we found main shared variance increased prominently while private variance decreased mildly over long-term training (Figures

5C and S3A), increasing the main shared-to-total variance ratio by a factor of ~ 2 (Figures 5D and S3B). These results are not specific to our bin width selection, as they held over a range of bin

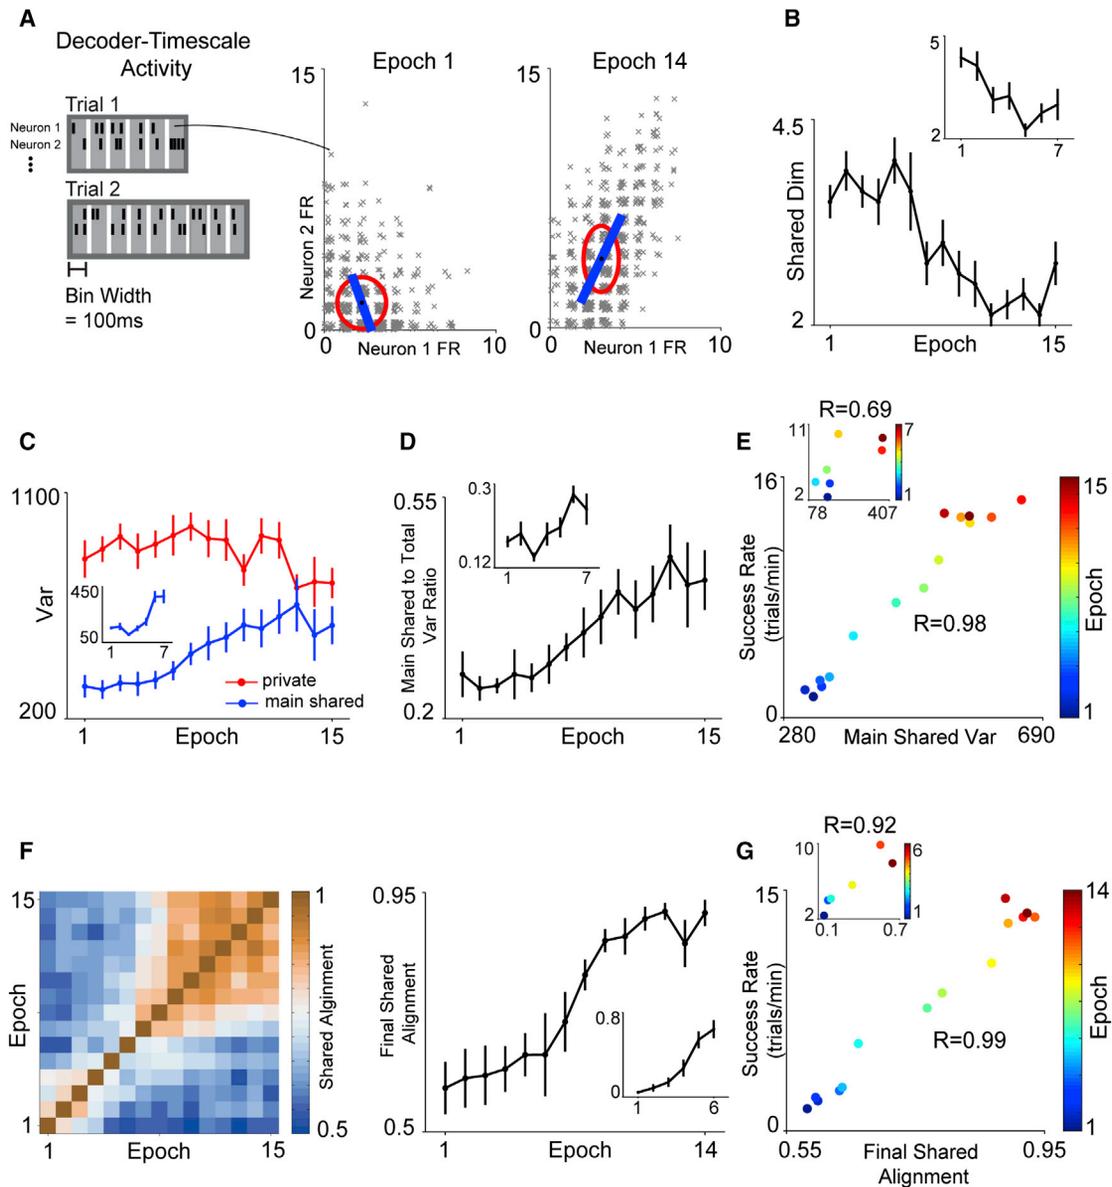


Figure 5. Decoder-Timescale Shared Variance Increases and Consolidates a Shared Space

(A) Decoder-timescale variance example: FA fit on two direct cells in early and late training for Monkey P on one example target.

(B) The shared dimensionality decreased from ~ 4 to ~ 2 (Monkey P, $r = -0.85$, $p = 5.7 \times 10^{-5}$, early > late $p = 1.4 \times 10^{-7}$; Monkey R, $r = -0.80$, $p = 3.1 \times 10^{-2}$, early > late $p = 9.8 \times 10^{-5}$). We used a main shared dimensionality of 2 (see STAR Methods).

(C) Main shared variance increased with training (Monkey P, $r = 0.92$, $p = 8.1 \times 10^{-7}$, early < late $p = 1.9 \times 10^{-5}$; Monkey R, $r = 0.84$, $p = 2.0 \times 10^{-2}$, early < late $p = 1.2 \times 10^{-5}$).

(D) The ratio of main shared variance to total variance increased with training (Monkey P, $r = 0.95$, $p = 7.1 \times 10^{-8}$, early < late $p = 2.2 \times 10^{-6}$; Monkey R, $r = 0.76$, $p = 4.6 \times 10^{-2}$, early < late $p = 1.2 \times 10^{-4}$).

(E) Main shared variance correlated with success rate (Monkey P, $r = 0.98$, $p = 1.9 \times 10^{-10}$; Monkey R, $r = 0.69$, $p = 8.6 \times 10^{-2}$).

(F) (Left) Shared space alignment map. The upper right red block indicates the shared space becomes stable, while the preceding blue indicates the shared space is changing in early training. (Right) Alignment with the final epoch's shared space increased with training (Monkey P, $r = 0.95$, $p = 1.8 \times 10^{-7}$, early < late $p = 9.6 \times 10^{-9}$; Monkey R, $r = 0.96$, $p = 2.2 \times 10^{-3}$, early < late $p = 2.2 \times 10^{-6}$). The mean chance alignment is 0.13 (Monkey P) and 0.20 (Monkey R) and 95th percentile chance alignment is 0.28 (Monkey P) and 0.40 (Monkey R).

(G) Alignment with the final epoch's shared space correlated with success rate improvement (Monkey P, $r = 0.99$, $p = 7.7 \times 10^{-11}$; Monkey R, $r = 0.92$, $p = 8.3 \times 10^{-3}$). Changes in mean firing rate do not explain these results (Figure S5). These results also held for analyses over training days rather than epochs (Figure S3) and different bin widths (Figure S4).

Error bars represent SEM over targets ($n = 8$ for each subject).

sizes (Figure S4E). Further, main shared variance correlated with success rate (Figures 5E and S3C), suggesting that it was used for control.

We next tested whether a shared space was consolidated over training for decoder-timescale variance. By analyzing the main shared space alignment, we found a period of relative instability in early learning followed by a period of stability in late learning (Figure 5F, left). Notably, the stable main shared space of late training was significantly different from the main shared spaces explored in early training (Figure 5F, right). Further, the alignment of main shared variance to the final epoch correlated with success rate (Figure 5G), suggesting main shared variance consolidation supported learning. These results support the flexible-network learning mechanism, suggesting that learning proceeds by increasing shared variance within an emergent shared space that coordinates activity used for BMI control.

Decoder-Timescale Private Variance Contributes Noisy yet Successful Control, while Consolidated Shared Variance Supports Skillful Control

Decoder-timescale private and shared variance changed differently over training, so we leveraged the decoder to ask how they contributed to cursor control. Mathematically, both sources are feasible strategies for generating a desired control signal (Figure 6A, left). Private variance represents the limit of high-dimensional activity, such that the population exhibits no correlations, and still can possess the temporal structure to generate desired control signals in the readout space (Figure 6A, middle). Shared variance captures low-dimensional covariation patterns and also can generate desired control signals in the readout space (Figure 6A, right). For each target and epoch, we fit an FA model on decoder-timescale spike counts and used it to calculate the contribution of main shared and private variance to each individual trial of activity (see STAR Methods). We then used the main shared and private variance contributions as separate inputs to the decoder and compared the resulting control signals (Figure 6B). Best trials produced by each source are shown, illustrating that both private and main shared variance could produce good control.

Under independent neuron learning, the population uses private variance for BMI control, while for constrained-network learning, the population uses shared variance. Because flexible-network learning permits both sources to contribute to control, we might expect private variance to dominate exploratory control in the beginning, while shared variance might contribute refined control in the end. Interestingly, both private and main shared variance improved how close they drove the cursor to the target (Figure 6C). However, when we compared late training trials for which both sources achieved the target (Monkey P, 832 trials; Monkey R, 385 trials; Figure 6D), we observed that main shared movements were faster (average time difference, Monkey P, -0.80 s; Monkey R, -0.66 s) and more direct (average distance difference, Monkey P, -1.6 ; Monkey R, -4.0 , in units of distance from center to target) (Figures 6E and 6F). Thus, in late training, main shared variance contributed more skillful control. Indeed, offline analysis suggests main shared variance can even produce more skillful control than the total activity (Figures S6B–S6D).

Main shared variance was not always better than private variance, however. In early training, private variance achieved positions closer to target (Figure 6C), and total activity-driven movements explored more of the workspace than main shared movements (Figure 6G). This suggests that initial main shared variance did not vary within the readout space, and that private variance served an exploratory role in early training by contributing behavioral variability. In support of flexible-network learning, private and main shared variance provided different contributions to performance improvement, with main shared variance contributing more skillful control in late learning.

Control-Relevant Shared Variance Increases because of Shared Variance Growth and Re-alignment

How does decoder-timescale shared variance support skillful control (Figure 6D)? We asked how main shared variance occupied the decoder's readout space to produce control signals (Figure 7A), finding that main shared variance in the readout space increased over training (Figure 7B). Two independent mechanisms exist that increase readout space shared variance. First, shared variance itself can increase (magnitude growth; Figure 7C, left), as we already found (Figure 6A). Second, the shared space can rotate to align with the readout space (re-alignment), thus making shared variance more efficient in producing control signals (Figure 7C, right). We quantified the relative contributions of these two mechanisms by computing the factor by which main shared variance grew and the factor by which main shared space alignment with the readout space grew. The product of these two yields the factor by which readout space main shared variance grew. Interestingly, both mechanisms significantly increased over training (Figure 7D). This provides one explanation for the changes in the main shared space; it can align with the readout space, making main shared variance more effective in producing control signals. These results provide a glimpse of how motor cortex can strengthen and optimize population coordination for neuroprosthetic control.

Fine-Timescale Shared Neural Trajectories Are Consolidated, Which Produces Accurate Movement

Having analyzed how the consolidated main shared variance relates to the readout space, we asked whether the covariation patterns consolidated a temporal sequence that supported skillful control, since shared variance can arise with or without consistent neural trajectories (Figure 8A). To investigate temporal structure, we used finer bins (50 ms) than the decoder timescale (100 ms), and for each individual target and epoch, we fit FA and calculated the main shared variance contribution to each trial's activity and trial averaged to form the main shared trajectory (see STAR Methods). Early epochs showed no pronounced pattern, while later epochs showed a large and consistent trajectory (Figures 8A and 8B). We found that the magnitude of each target's main shared trajectory increased (Figure 8D) and that the similarity of each target's main shared trajectory to the final epoch increased (Figure 8E), demonstrating neural trajectory consolidation.

Was all the observed main shared variance due to these emergent neural trajectories? We took each trial's main shared

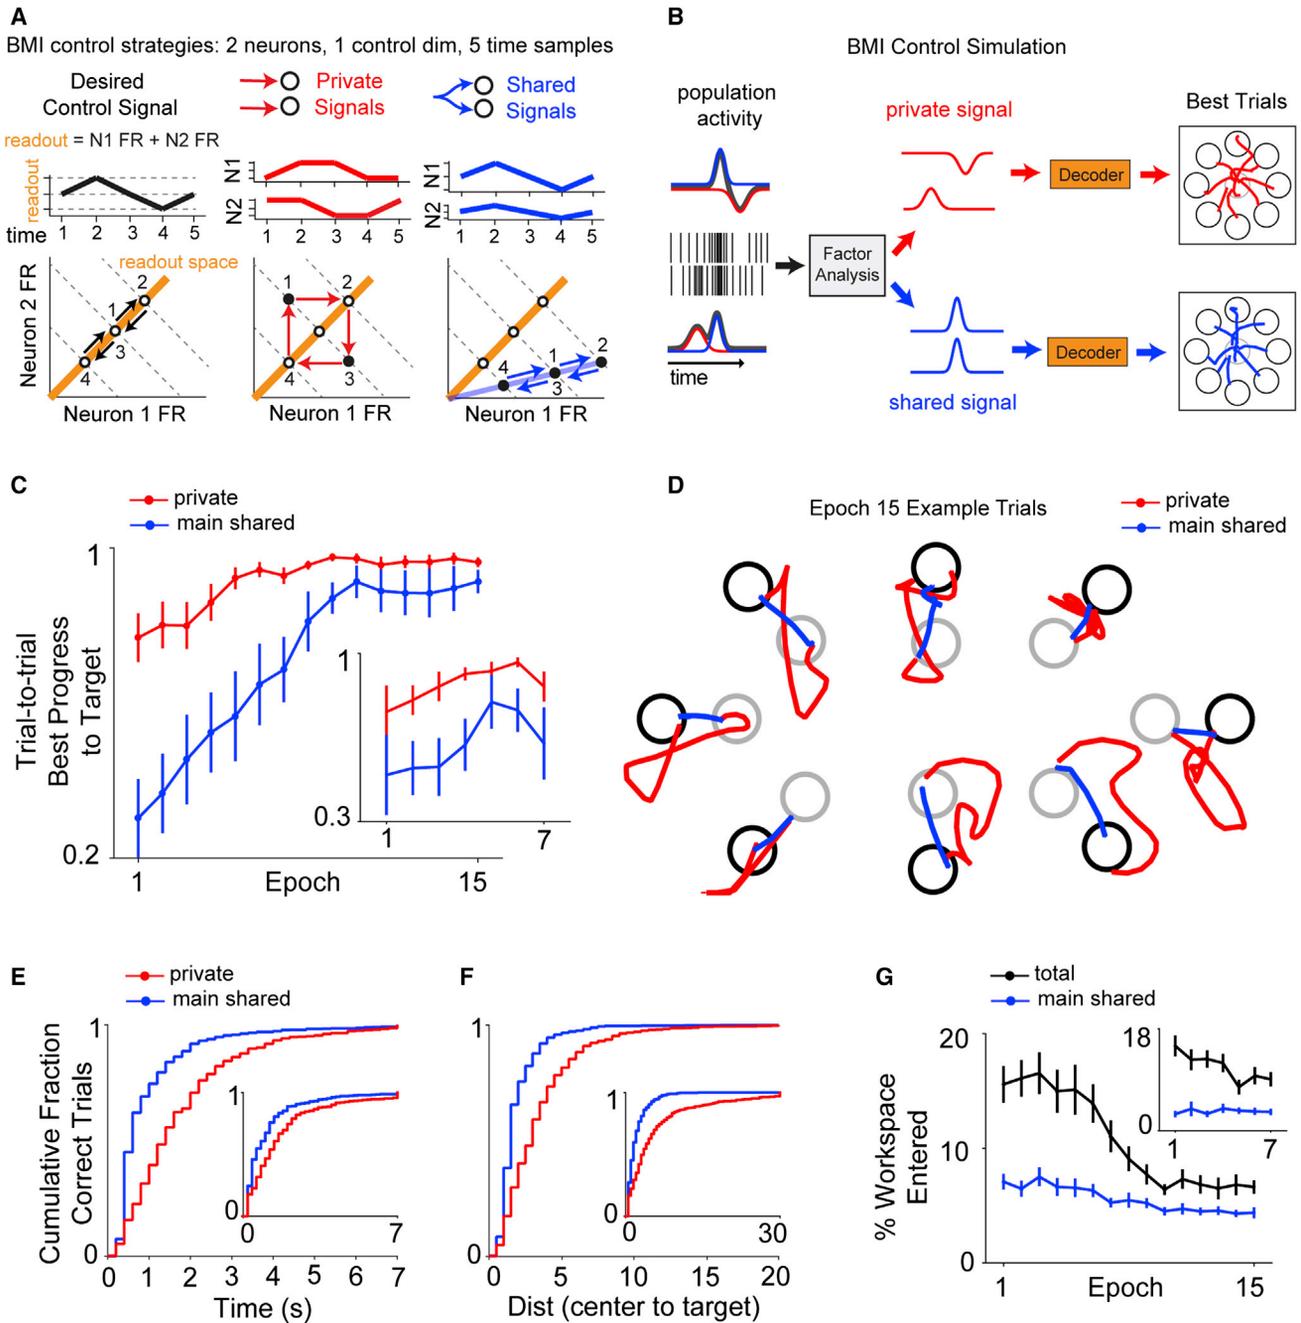


Figure 6. Decoder-Timescale Private Variance Produces Noisy yet Successful Movements, While Shared Variance Generate Skillful Movements

(A) Both private and shared signals can mathematically produce the same desired control signal. (Left) The desired control signal is shown over time, and its values are plotted on neuron firing rate axes. The readout value is the sum of the two neurons' firing rates. The dashed lines are contours that capture firing rates yielding the same sum. (Middle) The neurons' activity is uncorrelated and produces the desired control signal. (Right) The neurons' activity is perfectly correlated and produces the desired control signal.

(B) The contribution of private and shared variance to each trial's activity can be estimated and simulated through the decoder separately to determine how they independently contribute to movement.

(C) Both main shared and private variance increased best progress to target over training. Private variance; Monkey P, $r = 0.85$, $p = 5.7e-5$, early < late $p = 2.7e-6$; Monkey R, $r = 0.75$, $p = 5.3e-2$, early < late $p = 3.0e-2$. Main shared variance; Monkey P, $r = 0.94$, $p = 2.0e-7$, early < late $p = 3.1e-8$; Monkey R, $r = 0.74$, $p = 5.8e-2$, early < late $p = 9.6e-3$.

(D) Example trials driven by main shared and private variance from the final epoch for each target.

(legend continued on next page)

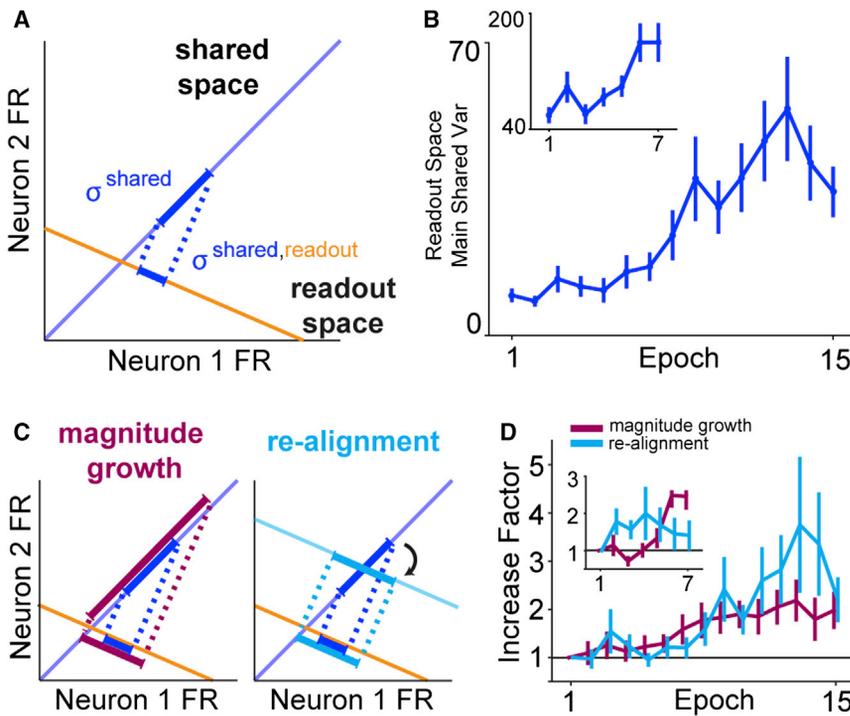


Figure 7. Task-Relevant Shared Variance Increases because of Shared Variance Growth and Re-alignment

(A) Shared variance can drive the decoder by producing variance in the decoder's readout space.

(B) Main shared variance in the readout space increased (Monkey P, $r = 0.90$, $p = 6.6e-6$, early < late $p = 8.5e-9$; Monkey R, $r = 0.85$, $p = 1.5e-2$, early < late $p = 8.4e-5$).

(C) Shared variance can increase in the readout space by two mechanisms. (Left) An increase in shared variance magnitude increases the readout space shared variance, as long as the shared space and readout space initially align. (Right) Re-alignment of the shared space with the readout space increases the readout space shared variance. Readout space shared variance growth is the product of the magnitude growth and re-alignment.

(D) Magnitude growth and re-alignment both increased for main shared variance, plotted as the ratio of each epoch's value to the first epoch's value. Magnitude growth factor; Monkey P, $r = 0.93$, $p = 7.3e-7$, early < late $p = 3.1e-4$; Monkey R, $r = 0.82$, $p = 2.4e-2$, early < late $p = 1.6e-5$. Re-alignment factor; Monkey P, $r = 0.84$, $p = 9.9e-5$, early < late $p = 8.2e-6$; Monkey R, $r = 0.17$, $p = 7.2e-2$, early < late $p = 6.2e-1$. While shared

space alignment with the readout space increased, note that a significant fraction of shared variance remained outside the readout space in late learning (Figure S6A).

Error bars represent SEM over targets ($n = 8$ for each subject).

variance contributions and subtracted the trial-averaged main shared trajectory, leaving behind main shared variation beyond that due to the average trajectory, and calculated how much main shared variance remained. The large fraction of remaining main shared variance shows that the observed coordination of population activity is not simply explained by the activity's average time course. Interestingly, initial main shared variance showed no trial-averaged neural trajectory. The fraction of main shared variance due to the average neural trajectory increased with training (Figure 8F), showing that main shared variance became increasingly temporally structured over learning. Finally, we tested whether the main shared trajectories were sufficient for generating neuroprosthetic control. We applied them through the decoder and found that they evolved over learning to generate straight and accurate movements in late training (Figures 8G and 8H). These results show that emergent shared temporal structure produces control signals for skillful movement.

DISCUSSION

There are substantial challenges to investigating the neural basis of classical motor skill learning, as experimenters cannot observe all cells driving movement, cannot track the precise neural-movement transformation, and cannot analyze how a complete task-relevant population coordinates over learning. In this work, we leverage the reduced preparation of the BMI paradigm to study how the direct cells driving the decoder explored activity patterns from trial to trial, and how they consolidated control signals underlying skillful movement. Using FA, we uncovered that private and shared signals both evolve and contribute to control differently in a task-relevant population over training.

Private Signals Contribute Initial Neural and Prosthetic Variability

Previous studies have shown that neural variability reduction in motor brain areas accompanies motor adaptation (Mandelblat-Cerf

(E) In late training, successful main shared variance-driven movements were significantly faster than private variance-driven movements (main shared minus private time to target; Monkey P, mean = -0.80 s, main shared < private $p = 1.8e-56$; Monkey R, mean = -0.66 s, main shared < private: $p = 4.3e-11$).

(F) In late training, successful main shared variance-driven movements were significantly more direct than private variance-driven movements (main shared minus private distance traveled; Monkey P, mean = -1.6 , main shared < private $p = 3.2e-52$; Monkey R, mean = -4.0 , main shared < private $p = 2.49e-22$). Interestingly, main shared variance achieved higher performance than even total activity (Figure S6).

For (E) and (F), late training trials were analyzed for which both private and main shared variances were successful (Monkey P, epochs 8–15, 832 trials; Monkey R, epochs 4–7, 385 trials).

(G) Total activity-driven movements explored the workspace more than main shared variance-driven movements, especially in early training (main shared < total, Monkey P, $p = 2.5e-9$; Monkey R, $p = 3.2e-9$).

Error bars represent SEM over targets ($n = 8$ for each subject).

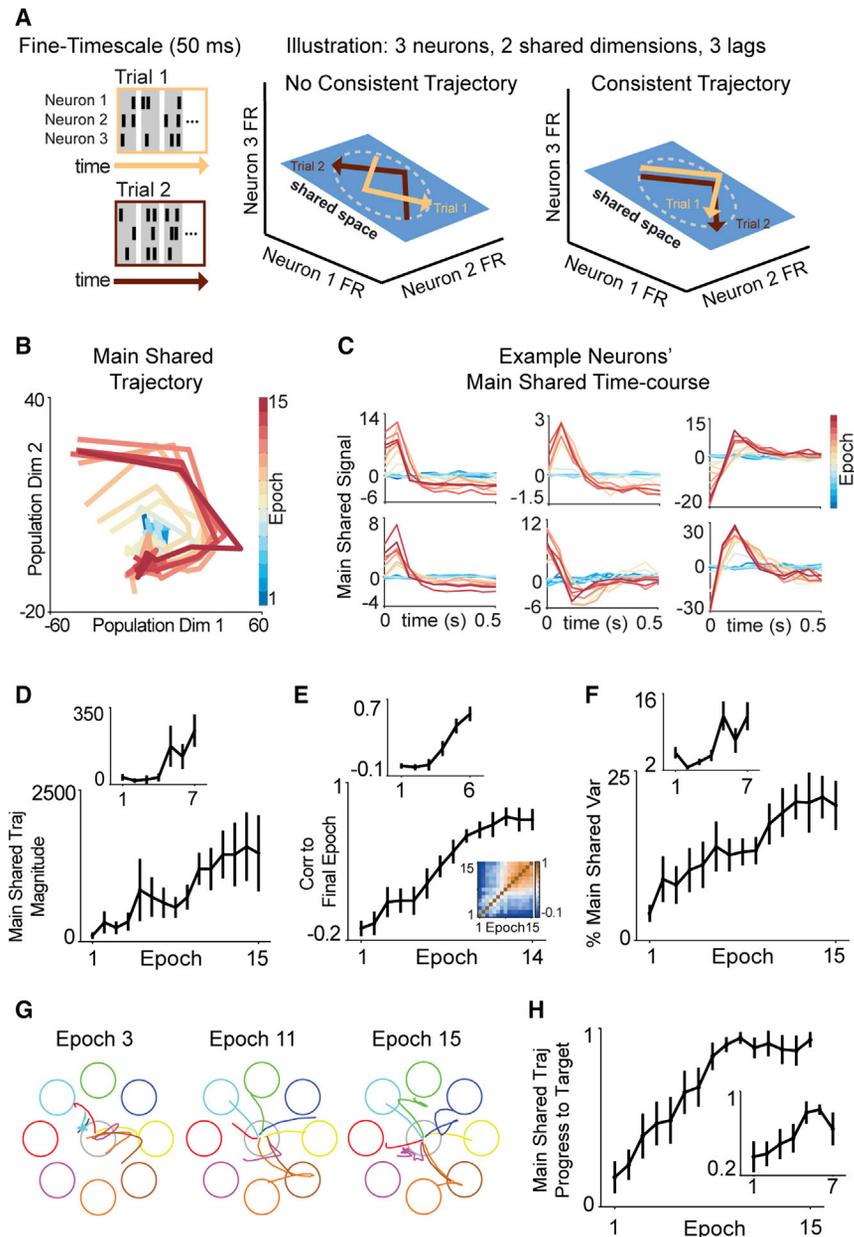


Figure 8. Shared Neural Trajectories Are Consolidated at Fine Timescale, Which Produces Accurate Movement

(A) Shared variance can arise without a consistent neural trajectory (left) or with a consistent neural trajectory (right).

(B) For one example target, every epoch's trial-averaged main shared trajectory is plotted in the same informative two-dimensional plane of the high-dimensional neural space (see STAR Methods).

(C) Six example cells' main shared time course over training epochs for one target.

(D) The magnitude of the main shared trajectory increased, as quantified by the squared norm (Monkey P, $r = 0.95$, $p = 8.8e-8$, early < late $p = 9.1e-7$; Monkey R, $r = 0.86$, $p = 1.3e-2$, early < late $p = 1.0e-6$).

(E) The correlation between each epoch's main shared trajectory and the final epoch increased and stabilized (Monkey P, $r = 0.98$, $p = 3.0e-9$, early < late $p = 2.2e-9$; Monkey R, $r = 0.93$, $p = 6.3e-3$, early < late $p = 5.5e-8$).

(F) The percent of fine-timescale main shared variance due to the trial-averaged main shared trajectory increased (Monkey P, $r = 0.96$, $p = 9.7e-9$, early < late $p = 4.4e-6$; Monkey R, $r = 0.77$, $p = 4.4e-2$, early < late $p = 7.1e-5$).

(G and H) The trial-averaged main shared trajectory simulated through the decoder increased movement accuracy (Monkey P, $r = 0.92$, $p = 1.4e-6$, early < late $p = 1.6e-8$; Monkey R, $r = 0.84$, $p = 1.8e-3$, early < late $p = 5.9e-4$).

Error bars represent SEM over targets ($n = 8$ for each subject).

possesses flexibility in finding goal-achieving patterns throughout population activity space.

How did initial decoder-timescale private variance contribute to cursor variability? We supplied shared and private variance separately as simulated inputs to the decoder and found that private variance contributed large cursor workspace exploration in early training (Figure 6G), which achieved positions closer to the

et al., 2009), procedural learning (Barnes et al., 2005), skill learning (Costa et al., 2004; Kao et al., 2005; Peters et al., 2014; Santos et al., 2015), and neuroprosthetic learning (Arduin et al., 2013; Zacksenhouse et al., 2007), suggesting that the brain searches for goal-achieving patterns by modulating neural variability. We asked how a task-relevant neural population coordinates to search for goal-achieving activity patterns. By dissecting total trial-to-trial variability with FA, we found that initially large trial-to-trial variability that is private to each neuron decreases over training (Figure 4B). This result extends Zacksenhouse et al. (2007), who found that early neuroprosthetic training was accompanied by increased total firing rate variability that was unexplained by BMI kinematics. Thus, a task-relevant population can explore activity space by increasingly varying each neuron independently, suggesting that the brain

target, while shared variance initially had little impact on control (Figure 6C). Thus, our findings suggest that each neuron in the task-relevant population uses independent patterns to explore behavioral states in early learning.

Role of Private Signals in Solving the Credit Assignment Problem

Neuroprosthetic learning studies provide intriguing evidence that the brain can solve the credit assignment problem by specifically adapting the neurons that contribute to the global error signal provided by prosthetic cursor feedback (Ganguly et al., 2011; Gulati et al., 2014; Jarosiewicz et al., 2008; Koralek et al., 2013). To explain changes specific to the neurons driving errors, neuroprosthetic learning models require private neural

noise to provide a beneficial exploratory signal for learning (Héliot et al., 2010; Legenstein et al., 2010). These models posited a fixed level of exploratory neural noise, and simulations indicated that this noise can hurt final cursor performance. Thus, our observation of initially large and later reduced private trial-to-trial variability is consistent with the view that motor cortex varies neurons independently to learn their contribution to control and then reduces this variability to generate refined control.

Shared Variance Consolidates for Skillful Neuroprosthetic Control

It is known that particular neural ensembles (Cao et al., 2015; Peters et al., 2014) and spatiotemporal patterns are consolidated during natural motor learning (Costa et al., 2004; Peters et al., 2014; Santos et al., 2015) and neuroprosthetic learning (Carmena et al., 2003; Ganguly and Carmena, 2009; So et al., 2012a), but less is known about whether and how the population coordinates over learning. We asked the following: for skilled control, does each neuron use independent patterns, or does the population utilize covariation patterns? For both trial-to-trial variability and decoder-timescale activity, we found that the shared-to-total variance ratio increased (Figures 4C and 5D) and that population activity is coordinated in a stable, low-dimensional shared space (Figures 4E and 5F). Using decoder simulations, we observed that emergent decoder-timescale shared variance contributed more skillful control than both private variance and total activity (Figures 6E, 6F, S6C, and S6D). We hypothesize that the population finds this solution by selecting particular shared inputs that produce goal-achieving activity within a characteristic manifold.

Flexibility in Acquiring Neural Patterns

Recent studies have asked how flexible the brain is in acquiring neural patterns for neuroprosthetic control (Hwang et al., 2013; Sadtler et al., 2014). Sadtler et al. (2014) found that the structure of shared neural variance plays a causal role in neuroprosthetic learning; given one training session, it was easier for a subject to produce control signals within an existing BMI manifold of shared variance than outside the BMI manifold. Another study (Hwang et al., 2013) found that subjects can operate a BMI by re-purposing activity patterns associated with their physical movements. Thus, given a familiar task context, subjects more readily select neural patterns within their pre-existing repertoire, facilitating rapid control.

In our study, subjects possess no starting BMI control strategy as they learn a neuroprosthetic skill *de novo* over days of training. Using a decoder that permitted decoder-timescale variance from both private and shared sources as input, the neural population's initial shared variance turned out to be small (Figure 5C), misaligned with the decoder (Figures 7B and S6A), and produced neither significant cursor variability (Figure 6G) nor accurate movements (Figure 6C). Instead, the population initially generated private variance unconstrained to a manifold (Figure 5C), and over training, shared variance increased and a BMI manifold emerged containing consolidated neural trajectories for skillful control (Figures 6E, 6F, and 8H). Because private

variance permits subjects to observe the behavioral consequence of activity patterns outside the initial shared space, it may enable the consolidation of BMI manifolds and increase the space of patterns and thus decoders that subjects can learn. We hypothesize this flexible neural exploration and consolidation mechanism is employed for long-term novel skill learning, distinct from the more constrained neural exploration used for faster learning.

Refined Control through Matching of Neural Activity to the Effector

To control a novel effector, subjects must discover motor control neural activity that is matched to the effector. A recent neuroprosthetic learning study found that changes in neurons' mean firing rates across movement directions correlate with the decoder's properties (Orsborn et al., 2014). In this work, we asked how neural dynamics for individual movement directions drive the decoder (Figure 7A), finding that shared variance in the decoder's readout space increased over training (Figure 7B). This occurs both because shared variance increases overall, and because the shared space aligns with the readout space (Figure 7D). Thus, neuroprosthetic skill emerges as subjects consolidate shared variance matched to the decoder. Given that subjects could also generalize control for untrained reaches (Ganguly and Carmena, 2009), these neural changes may reflect the development of an internal model for BMI control (Golub et al., 2015; Héliot et al., 2010).

How might our findings extend to acquiring novel motor skill, which requires generating novel control signals to the body? The brain's task in BMI approximates current models of the brain's function in natural motor learning. Our decoder is a linear readout of neural activity into control signals, just as a common model posits muscles perform a linear readout of motor cortical activity (Churchland et al., 2012; Kaufman et al., 2014; Lillicrap and Scott, 2013). In analogy to motor cortical control of the decoder, motor control studies suggest that motor cortex generates neural control signals for our biomechanical plant, showing that motor cortical activity contains information about muscle-level dynamics as well as kinematics (Sergio et al., 2005) and possesses statistics that reflect the biomechanics it must control (Lillicrap and Scott, 2013). Thus, our findings on private and shared neural variance in motor cortex could apply to exploring and consolidating physical movements.

Refined Control through Consolidation of Coordinated Neural Trajectories

Because we observed decoder-timescale shared variance could produce skilled cursor movements (Figures 6D–6F), we asked whether covariation patterns acquired a temporal sequence at a fine timescale that supported learning. Interestingly, initial covariation patterns do not traverse a consistent trajectory (Figure 8D). Over training, as a shared space consolidates, neural trajectories consolidate within the shared space (Figures 8E and 8F) that are sufficient to drive straight and accurate cursor movements (Figures 8G and 8H). Future work will need to analyze whether these emergent temporal patterns obey dynamical rules such as those underlying natural motor control (Churchland et al., 2012; Kao et al., 2015).

One interesting observation is that a large fraction of shared variance exists outside the readout space and thus does not directly drive the cursor (Figure S6A), possibly reflecting internal neural dynamics necessary to generate the control signal to the decoder. This is consistent with motor control hypotheses that large neural populations exploit their high dimensionality such that neural activity can obey lawful dynamics while a linear readout of the activity can produce flexible movements (Churchland et al., 2012) and can prepare movements without executing them (Kaufman et al., 2014). Interpreting neural data will be aided by modeling how neural networks might learn dynamics for BMI control as is being done for motor control (Hennequin et al., 2014; Sussillo et al., 2015).

Biological Sources of Private and Shared Variance: Noise and Supporting Networks

We note that private variance in the direct cells can arise from multiple processes (Faisal et al., 2008): (1) “noisy” intrinsic spiking variance at a fixed firing rate, such as in a Poisson Process with fixed firing rate; (2) the action of private input activity changing the cell’s underlying firing rate; and (3) input activity structure that FA does not model, for example, time-varying shared variance. We used Poisson simulations to ensure changing mean firing rate did not explain our observed trial-to-trial variability changes (Figure S5). In the latter two cases, some of the variance that is private in the population of direct cells may be shared with unobserved cells in connected networks. Thus, private variance need not be purely corruptive noise, a simplifying assumption sometimes made (Santhanam et al., 2009).

Cells in M1 likely possess tens of thousands of synapses (Cragg, 1967); the input activity space is very high dimensional. Some of these inputs may be independent for each recorded cell while some may be coordinated across cells. These two classes of inputs would drive distinct statistical population variance, which FA attempts to describe with shared and private variance. There are many potential biological implementations for the observed increase in consistency and coordination in the direct cell population. Private trial-to-trial variability can decrease because variable private inputs become more consistent or their synaptic weights to direct cells weaken. Decoder-timescale shared variance can increase because anatomically shared inputs strengthen in signal or synaptic weight or because non-anatomically shared inputs become coordinated. Our FA results suggest new hypotheses about the neural substrate of skill learning in the direct cells and the supporting network.

Indeed, a growing body of research indicates that during BMI control, the brain can enter a novel, wide-scale state producing new representations in indirect cells (i.e., not connected to the BMI) in M1 (Clancy et al., 2014; Ganguly et al., 2011; Gulati et al., 2014; So et al., 2012b), in primary sensory cortex (Clancy et al., 2014), across cortical areas (Wander et al., 2013), and in striatum (Koralek et al., 2013, 2012). These functional changes are subserved by ensemble re-activation during slow wave sleep (Gulati et al., 2014) and corticostriatal plasticity (Koralek et al., 2012), and suggest that emergent task-relevant coordination is driven by the formation of large-scale circuits including corticostriatal loops.

Implications for Neural Prostheses

Advances in neural engineering have produced algorithms to facilitate high performance even when decoders are re-trained frequently. How might decoder and neural adaptation synergize? Recent work shows closed-loop decoder adaptation (CLDA), which uses information about the task goal to fit the decoder to the subject’s explored activity patterns (Dangi et al., 2013; Gilja et al., 2012; Shanechi et al., 2016), can improve performance in conjunction with neural learning (Orsborn et al., 2014). Future decoders might benefit from more detailed models of neural population dynamics and how they change with learning. Indeed, a recent algorithmic approach yielded significant performance improvement by modeling neural population dynamics underlying natural movements to decode the subject’s intent while moving freely (Kao et al., 2015). Perhaps neural learning can help to generalize this approach to immobile patients, as we found coordinated neural dynamics can be consolidated over training in the absence of overt movement. Given our findings that main shared variance achieves better performance than total activity in simulations (Figures S6B–S6D), a performance-motivated extension would be to design a decoder that is able to denoise neural observations based on learned neural dynamics (Shenoy and Carmena, 2014).

In conclusion, this study uncovers the emergence and consolidation of coordinated dynamics for BMI control, highlighting the potential for neural prostheses to become more natural and automatic to the user, as well as an experimental tool for future investigations on the neural basis of skill learning.

STAR★METHODS

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

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
 - Electrophysiology
 - Experimental Setup and Behavioral Training
 - BMI Decoder
 - Quantifying Learning Over Training Epochs
 - Behavioral Metrics
 - Factor Analysis
 - Trial-to-trial Variability of Shared and Private Signals
 - Decoder-timescale Private and Shared Variance
 - Sample-by-sample Contribution of Private and Shared Variance
 - Trial-to-trial Decoder Simulations
 - Relationship of Shared Variance to the Readout Space
 - Shared Neural Trajectories at Fine-Timescale
- QUANTIFICATION AND STATISTICAL ANALYSES

SUPPLEMENTAL INFORMATION

Supplemental Information includes six figures and one table and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2017.01.016>.

AUTHOR CONTRIBUTIONS

K.G. and J.M.C. conceived, designed, and performed the original experiments. V.R.A., R.M.C., and J.M.C. conceived and designed this study. V.R.A. analyzed the data. V.R.A., R.M.C., and J.M.C. contributed materials and analysis tools. V.R.A., R.M.C., and J.M.C. wrote the paper.

ACKNOWLEDGMENTS

We thank A. Motiwala, D. Kobaks, J. Semedo, P. Khanna, and S. Gowda for helpful discussions, and R. Neely and S. Santacruz for manuscript comments. We also thank the reviewers for helpful comments and constructive suggestions. This work was supported by the National Science Foundation Graduate Research Fellowship to V.R.A.; grants from ERA-NET, European Research Council (COG 617142), and Howard Hughes Medical Institute (IEC 55007415) to R.M.C.; and grants from the National Science Foundation (CBET-0954243 and EFRI-M3C 1137267) and Office of Naval Research (N00014-15-1-2312) to J.M.C.

Received: January 13, 2016

Revised: September 7, 2016

Accepted: January 19, 2017

Published: February 9, 2017

REFERENCES

- Arduin, P.J., Frégnac, Y., Shulz, D.E., and Ego-Stengel, V. (2013). "Master" neurons induced by operant conditioning in rat motor cortex during a brain-machine interface task. *J. Neurosci.* *33*, 8308–8320.
- Barnes, T.D., Kubota, Y., Hu, D., Jin, D.Z., and Graybiel, A.M. (2005). Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* *437*, 1158–1161.
- Briggman, K.L., and Kristan, W.B. (2008). Multifunctional pattern-generating circuits. *Annu. Rev. Neurosci.* *31*, 271–294.
- Cao, V.Y., Ye, Y., Mastwal, S., Ren, M., Coon, M., Liu, Q., Costa, R.M., and Wang, K.H. (2015). Motor learning consolidates Arc-expressing neuronal ensembles in secondary motor cortex. *Neuron* *86*, 1385–1392.
- Carmena, J.M., Lebedev, M.A., Crist, R.E., O'Doherty, J.E., Santucci, D.M., Dimitrov, D.F., Patil, P.G., Henriquez, C.S., and Nicolelis, M.A.L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol.* *1*, E42.
- Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* *13*, 369–378.
- 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.
- Clancy, K.B., Koralek, A.C., Costa, R.M., Feldman, D.E., and Carmena, J.M. (2014). Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat. Neurosci.* *17*, 807–809.
- Cohen, R.G., and Sternad, D. (2009). Variability in motor learning: relocating, channeling and reducing noise. *Exp. Brain Res.* *193*, 69–83.
- Costa, R.M., Cohen, D., and Nicolelis, M.A.L. (2004). Differential corticostriatal plasticity during fast and slow motor skill learning in mice. *Curr. Biol.* *14*, 1124–1134.
- Cragg, B.G. (1967). The density of synapses and neurones in the motor and visual areas of the cerebral cortex. *J. Anat.* *101*, 639–654.
- Dangi, S., Orsborn, A.L., Moorman, H.G., and Carmena, J.M. (2013). Design and analysis of closed-loop decoder adaptation algorithms for brain-machine interfaces. *Neural Comput.* *25*, 1693–1731.
- Dempster, A.P., Laird, N.M., and Rubin, D.B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. B* *39*, 1–38.
- Everitt, B.S. (1984). *An Introduction to Latent Variable Models* (London: Chapman & Hall).
- Faisal, A.A., Selen, L.P.J., and Wolpert, D.M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.* *9*, 292–303.
- Ganguly, K., and Carmena, J.M. (2009). Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol.* *7*, e1000153.
- Ganguly, K., and Carmena, J.M. (2010). Neural correlates of skill acquisition with a cortical brain-machine interface. *J. Mot. Behav.* *42*, 355–360.
- Ganguly, K., Dimitrov, D.F., Wallis, J.D., and Carmena, J.M. (2011). Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nat. Neurosci.* *14*, 662–667.
- 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.
- Golub, M.D., Yu, B.M., and Chase, S.M. (2015). Internal models for interpreting neural population activity during sensorimotor control. *eLife* *4*, e10015.
- Gulati, T., Ramanathan, D.S., Wong, C.C., and Ganguly, K. (2014). Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. *Nat. Neurosci.* *17*, 1107–1113.
- Héliot, R., Ganguly, K., Jimenez, J., and Carmena, J.M. (2010). Learning in closed-loop brain-machine interfaces: modeling and experimental validation. *IEEE Trans. Syst. Man Cybern. B Cybern.* *40*, 1387–1397.
- Hennequin, G., Vogels, T.P., and Gerstner, W. (2014). Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* *82*, 1394–1406.
- Hwang, E.J., Bailey, P.M., and Andersen, R.A. (2013). Volitional control of neural activity relies on the natural motor repertoire. *Curr. Biol.* *23*, 353–361.
- Jarosiewicz, B., Chase, S.M., Fraser, G.W., Velliste, M., Kass, R.E., and Schwartz, A.B. (2008). Functional network reorganization during learning in a brain-computer interface paradigm. *Proc. Natl. Acad. Sci. USA* *105*, 19486–19491.
- Kao, M.H., Doupe, A.J., and Brainard, M.S. (2005). Contributions of an avian basal ganglia-forebrain circuit to real-time modulation of song. *Nature* *433*, 638–643.
- 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.
- 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.
- Koralek, A.C., Jin, X., Long, J.D., 2nd, Costa, R.M., and Carmena, J.M. (2012). Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* *483*, 331–335.
- Koralek, A.C., Costa, R.M., and Carmena, J.M. (2013). Temporally precise cell-specific coherence develops in corticostriatal networks during learning. *Neuron* *79*, 865–872.
- Legenstein, R., Chase, S.M., Schwartz, A.B., and Maass, W. (2010). A reward-modulated hebbian rule can explain experimentally observed network reorganization in a brain control task. *J. Neurosci.* *30*, 8400–8410.
- Lillicrap, T.P., and Scott, S.H. (2013). Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. *Neuron* *77*, 168–179.
- Mandelblat-Cerf, Y., Paz, R., and Vaadia, E. (2009). Trial-to-trial variability of single cells in motor cortices is dynamically modified during visuomotor adaptation. *J. Neurosci.* *29*, 15053–15062.
- Orsborn, A.L., Moorman, H.G., Overduin, S.A., Shanechi, M.M., Dimitrov, D.F., and Carmena, J.M. (2014). Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron* *82*, 1380–1393.
- Paxinos, G., Huang, X.-F., and Toga, A. (2000). *The Rhesus Monkey Brain in Stereotaxic Coordinates* (San Diego: Academic Press).

- Peters, A.J., Chen, S.X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature* 510, 263–267.
- 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.
- Santhanam, G., Yu, B.M., Gilja, V., Ryu, S.I., Afshar, A., Sahani, M., and Shenoy, K.V. (2009). Factor-analysis methods for higher-performance neural prostheses. *J. Neurophysiol.* 102, 1315–1330.
- Santos, F.J., Oliveira, R.F., Jin, X., and Costa, R.M. (2015). Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning. *eLife* 4, e09423.
- 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.
- Shanechi, M.M., Orsborn, A.L., and Carmena, J.M. (2016). Robust brain-machine interface design using optimal feedback control modeling and adaptive point process filtering. *PLoS Comput. Biol.* 12, e1004730.
- Shenoy, K.V., and Carmena, J.M. (2014). Combining decoder design and neural adaptation in brain-machine interfaces. *Neuron* 84, 665–680.
- Shmuelof, L., Krakauer, J.W., and Mazzoni, P. (2012). How is a motor skill learned? Change and invariance at the levels of task success and trajectory control. *J. Neurophysiol.* 108, 578–594.
- So, K., Koralek, A.C., Ganguly, K., Gastpar, M.C., and Carmena, J.M. (2012a). Assessing functional connectivity of neural ensembles using directed information. *J. Neural Eng.* 9, 026004.
- So, K., Ganguly, K., Jimenez, J., Gastpar, M.C., and Carmena, J.M. (2012b). Redundant information encoding in primary motor cortex during natural and prosthetic motor control. *J. Comput. Neurosci.* 32, 555–561.
- Sussillo, D., Churchland, M.M., Kaufman, M.T., and Shenoy, K.V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* 18, 1025–1033.
- Sutton, R.S., and Barto, A.G. (1998). *Reinforcement Learning: An Introduction* (Cambridge: MIT Press).
- Tumer, E.C., and Brainard, M.S. (2007). Performance variability enables adaptive plasticity of 'crystallized' adult birdsong. *Nature* 450, 1240–1244.
- Wander, J.D., Blakely, T., Miller, K.J., Weaver, K.E., Johnson, L.A., Olson, J.D., Fetz, E.E., Rao, R.P.N., and Ojemann, J.G. (2013). Distributed cortical adaptation during learning of a brain-computer interface task. *Proc. Natl. Acad. Sci. USA* 110, 10818–10823.
- Wu, H.G., Miyamoto, Y.R., Gonzalez Castro, L.N., Ölveczky, B.P., and Smith, M.A. (2014). Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat. Neurosci.* 17, 312–321.
- 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.
- Zacksenhouse, M., Lebedev, M.A., Carmena, J.M., O'Doherty, J.E., Henriquez, C., and Nicolelis, M.A.L. (2007). Cortical modulations increase in early sessions with brain-machine interface. *PLoS ONE* 2, e619.

STAR★METHODS

KEY RESOURCES TABLE

| REAGENT or RESOURCE | SOURCE | IDENTIFIER |
|---|----------------------------|---|
| Experimental Models: Organisms/Strains | | |
| Rhesus macaque (<i>Mucacca mulatta</i>) | UC Davis | N/A |
| Software and Algorithms | | |
| MATLAB | MathWorks | https://www.mathworks.com/products/matlab.html |
| Plexon Sort Client and Wavetracker | Plexon | http://www.plexon.com/products/software |
| Other | | |
| Plexon MAP System | Plexon | http://www.plexon.com/products/multichannel-acquisition-processor-map-data-acquisition-system |
| Microwire arrays | Innovative Neurophysiology | http://www.inphysiology.com/fixe-d-arrays/ |
| NHP Kinarm exoskeleton | BKIN Technologies | http://www.bkintechnologies.com/bkin-products/nhp-kinarm-exoskeleton-lab/ |

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Dr. Jose M. Carmena (jcarmena@berkeley.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Two adult male rhesus monkeys (*Macaca mulatta*) (Monkey P, age: 7 years 1 month, weight: 15 kg; Monkey R, age 6 years, weight: 10.7 kg) were chronically implanted in the brain with arrays of 64 Teflon-coated tungsten microelectrodes (35 μm in diameter, 500 μm pitch; Innovative Neurophysiology, Durham NC) in an 8 \times 8 array configuration (Ganguly and Carmena, 2009). Monkey P was implanted in the left hemisphere in the arm area of both primary motor cortex (M1) and dorsal premotor cortex (PMd), and in the right hemisphere in the arm area of M1, with a total of 192 microwires across three implants. Monkey R was implanted bilaterally in the arm area of M1 and PMd (256 microwires across four implants). Only activity from M1 was used in these BMI experiments (Monkey P: right M1; Monkey R: left M1). Array implants were targeted for pyramidal tract neurons in layer 5. Localization of target areas was performed using stereotactic coordinates from a neuroanatomical atlas of the rhesus brain (Paxinos et al., 2000).

All procedures were conducted in compliance with the NIH Guide for the Care and Use of Laboratory Animals and were approved by the University of California at Berkeley Institutional Animal Care and Use Committee.

METHOD DETAILS

Electrophysiology

Neural activity was recorded using the MAP system (Plexon). Stable units were selected based on waveform shape, amplitude, relationship to other units on the same channel, inter-spike interval distribution, and the presence of an absolute refractory period. Only units from primary motor cortex were used which had a clearly identified waveform with signal-to-noise ratio of at least 4:1. Activity was sorted prior to recording sessions using an online spike-sorting application (Sort Client; Plexon). Stability of waveforms was confirmed by analyzing the stability of PCA projections over days (Wavetracker; Plexon).

Experimental Setup and Behavioral Training

BMI Task

Subjects performed a center-out reaching task to eight targets (Figure 1A). The cursor was continuously controlled by neural activity. Subjects self-initiated trials by moving to the center target. Peripheral targets were presented in pseudorandom order. A successful trial required a short hold at the center, moving to the peripheral target within 15 s, and a brief hold at the target (Figure 1B). Successful trials resulted in a liquid reward; failed trials were repeated.

During BMI control, both arms were removed from the workspace, lightly restrained, and thus unable to perform reaches. During selected sessions, video and surface electromyogram (EMG) recordings from proximal muscle groups were performed. Neither animal moved their upper extremity during brain control.

Manual Control Training Before BMI

Before starting the BMI learning experiments, subjects were overtrained on the task performed with arm movements using a Kinarm (BKIN Technologies) exoskeleton which restricted shoulder and elbow to move in the horizontal plane.

BMI Decoder

Decoding Algorithm

In this study, the decoder used neural activity to control the joint angles of a virtual, invisible two-link arm. A Wiener Filter linearly mapped the binned spike counts of N neurons $x(t) \in \mathbb{R}^N$ in the last one second into virtual shoulder and elbow joint angles $y(t) \in \mathbb{R}^2$ via

$$y(t) = \mathbf{b} + \sum_{k=0}^M \mathbf{a}(k\Delta)x(t - k\Delta).$$

The fixed decoder parameters $\mathbf{a}(k\Delta) \in \mathbb{R}^{2 \times N}$ for $k = 0, \dots, M$ and $\mathbf{b} \in \mathbb{R}^2$ used in this study were fit via least-squares linear regression to predict the subjects' actual shoulder and elbow joint angles in a calibration manual control session (Ganguly and Carmena, 2009). In this experiment, $\text{num_lags} = 10 = M + 1$, and the spike count bin size was $\Delta = 100$ ms. The joint kinematics forward model mapped shoulder and elbow angles $y(t)$ to endpoint position, the only feedback signal the subjects received. The joint positions $y(t)$ and endpoint feedback signal were updated every $\Delta = 100$ ms.

Decoder Readout Space

We analyzed the decoder's weights and determined that applying neural activity through the decoder is equivalent to first projecting neural activity into a subspace (the decoder readout space) and then applying it through the decoder. The decoder produced 2-dimensional output by summing $\text{num_lags} = 10$ linear combinations of neural activity, so we analyzed how many neural dimensions the decoder primarily leveraged to produce control. At maximum, it could be 10, e.g., if each lag used a different individual neuron's activity. We found the readout space dimensionality by evaluating the expression

$$\mathbf{a}(k\Delta) \approx \mathbf{a}(k\Delta)\mathbf{P},$$

for $k = 0, \dots, M$, where $\mathbf{a}(k\Delta)$ are the decoder parameters, $\mathbf{P} \in \mathbb{R}^{N \times N}$ is a projection matrix of rank L , and $\text{num_lags} = 10 = M + 1$. (Note that for $L = N$ we have perfect equality.) \mathbf{P} was found by performing PCA on the decoder parameters to compress neural dimensionality (as opposed to temporal dimensionality). $L = 3$ was sufficient for $\mathbf{a}(k\Delta)\mathbf{P}$ to capture $\sim 97\%$ of $\mathbf{a}(k\Delta)$ parameter variance for Monkey P, and $L = 5$ captured $\sim 99\%$ of parameter variance for Monkey R.

The decoder "readout space" is the column space of \mathbf{P} , as we can reproduce BMI cursor movements by first projecting neural activity into this space. Thus, we note that for our decoder, the readout space dimensionality is not the same as the cursor dimensionality.

Quantifying Learning Over Training Epochs

We analyzed the behavioral and neural learning process for each target separately over "training epoch" blocks of a constant number of trials. We chose the "training epoch" rather than the more standard "training day" because we wished to eliminate the effect of variable trial numbers per training day on analysis results. Note that because subjects performed a different number of trials to each target overall, the number of trials in a training epoch differs for each target. For Monkey P (Monkey R), the smallest epoch size is 23 trials (31 trials) and the largest is 36 trials (56 trials).

Behavioral Metrics

Task-performance Metrics

We quantified task performance using percent correct (fraction of initiated trials completed successfully) and success rate (successful trials per minute). Success rate combines accuracy and speed and is given by $(\text{num_successful_trials} / \sum_{\text{all_trials}} \text{time_to_initiate_trial} + \text{time_to_complete_trial})$.

Cursor Variability Analyses

To analyze cursor variability, we analyzed the cursor trial-to-trial Spatial Occupancy Map (SOM) by discretizing the workspace (each bin was 0.34% of the workspace) and computing the fraction of trials the cursor would enter each position bin. This was computed for each target and training epoch as: $(\text{SOM}(\text{position}_{ij}) = \text{num_trials_cursor_enters_position}_{ij} / \text{num_trials})$. We performed three analyses. First, we assessed control refinement via the fraction of workspace entered per trial: $(1 / \text{num_trials}) (\sum_{\text{trial}_i=1}^{\text{num_trials}} (\text{num_bins_entered}(\text{trial}_i) / \text{num_bins}))$. Second, we assessed the stability of the learned cursor control strategy via the correlation between the final epoch SOM and previous epochs: $\text{corr}(\text{SOM}_{\text{epoch}_i}, \text{SOM}_{\text{epoch}_{\text{final}}})$, for $\text{epoch}_i = 1, \dots, \text{epoch}_{\text{final}-1}$, where $\text{corr}(\mathbf{a}, \mathbf{b}) = (\sum_{i=1}^N \mathbf{a}_i \mathbf{b}_i / (\|\mathbf{a}\| \|\mathbf{b}\|))$ is the standard vector correlation. Third and finally, we assessed the cursor's trial-to-trial movement variability by considering each trial as a vector of 0's and 1's indicating the cursor's spatial occupancy. We computed the covariance matrix using this parameterization of each trial and calculated the variability via the trace. We note that a decrease in trial-to-trial variability is not a mathematical consequence of reduction in workspace explored, as it is possible for subjects to enter

the same large fraction of the workspace consistently as well as to enter different small fractions of the workspace from trial-to-trial. The trends found by these analyses are not sensitive to the discretization resolution; the same results were achieved with coarser resolution, e.g., bins 4 times larger (data not shown).

Factor Analysis

Model

We used FA to model the joint distribution of N neurons' spike counts $x \in \mathbb{R}^N$ as the sum of 1) a mean rate $\mu \in \mathbb{R}^N$, 2) private signals with diagonal covariance $\Psi \in \mathbb{R}^{N \times N}$, and 3) shared signals due to a low-dimensional latent variable $z \in \mathbb{R}^k$, $k < N$:

$$z \sim N(0, I)$$

$$x | z \sim N(\mu + Uz, \Psi)$$

$$x \sim N(\mu, UU^T + \Psi).$$

$U \in \mathbb{R}^{N \times k}$ are the weights (or factors) capturing how z drives x . We make explicit FA's decomposition of population activity into shared signals $x^{\text{shared}} = Uz \in \mathbb{R}^N$ and private signals $x^{\text{private}} \in \mathbb{R}^N$:

$$x^{\text{private}} \sim N(0, \Sigma^{\text{private}})$$

$$x^{\text{shared}} \sim N(0, \Sigma^{\text{shared}})$$

$$x^{\text{private}}, x^{\text{shared}} \text{ independent}$$

$$x = \mu + x^{\text{shared}} + x^{\text{private}},$$

$$\text{where } \Sigma^{\text{shared}} = UU^T, \Sigma^{\text{private}} = \Psi, \text{ and } \Sigma^{\text{total}} = \Sigma^{\text{shared}} + \Sigma^{\text{private}}.$$

For each target separately, we fit the model's parameters by maximizing the log-likelihood of the data with the EM algorithm for a chosen shared dimensionality $k < N$. The best-fitting shared dimensionality k is estimated by using cross-validated log-likelihood to determine which k best describes held-out data (Dempster et al., 1977). To ensure we didn't choose too high of shared dimensionality, we analyzed the eigenvalues of the shared variance found using the dimensionality which maximized cross-validated log-likelihood and evaluated how many dimensions were needed to capture 90% of shared variance. We used this estimate of shared dimensionality throughout the work.

Comparison to Pairwise Correlation and PCA

To quantify correlated population activity, FA has advantages over more familiar measures such as average pairwise correlation and PCA. First, pairwise correlation addresses how well one cell's activity can predict another's, while the cell's shared-to-total variance ratio indicates how well that cell's activity can be predicted by the rest of the population together. Second, average pairwise correlation does not tell us the data's dimensionality. While PCA provides information about dimensionality, FA provides additional features. FA is a probabilistic model which posits total covariance as the sum of shared and private variance, permitting analysis of the two sources separately, while PCA analyzes dimensionality of total covariance.

Shared and Private Variance and the Shared-to-total Variance Ratio

Neuron i 's shared variance is $\Sigma_{ii}^{\text{shared}}$, private variance is $\Sigma_{ii}^{\text{private}}$, and total variance is $\Sigma_{ii}^{\text{total}} = \Sigma_{ii}^{\text{shared}} + \Sigma_{ii}^{\text{private}}$. The population-average variance is $(1/N)\text{trace}(\Sigma) = (1/N)\sum_{i=1}^N \Sigma_{ii}$. The population's shared variance to total variance ratio is $(\text{trace}(\Sigma^{\text{shared}})/\text{trace}(\Sigma^{\text{total}}))$.

Extracting Main Shared Variance

After fitting FA, we ordered shared signals by their variance by diagonalizing Σ^{shared} via the singular value decomposition (SVD) (Yu et al., 2009). Throughout the work we compared the decoder-timescale "main shared" variance, i.e., variance within the top p orthogonalized factors which provide the best rank- p approximation of Σ^{shared} . For both subjects, $p = 2$, as that was the shared dimensionality in late learning.

Shared Space Alignment

We used the "shared space alignment" to measure the similarity between the shared variance (or main shared variance) of Epoch A and Epoch B. The shared space alignment is the fraction of epoch A shared variance captured in epoch B's shared space and thus ranges from 0 to 1. For some geometric intuition, in the one-dimensional case (i.e., $\text{rank}(\Sigma^{A,\text{shared}}) = \text{rank}(\Sigma^{B,\text{shared}}) = 1$), the space alignment is equivalent to $\cos \theta$, where θ is the angle between epoch A and epoch B's one-dimensional shared space. We note that the shared space alignment is asymmetric when shared dimensionality is greater than 1, such that alignment of A with B need not be equal to the alignment of B with A.

Let $\Sigma^{A, \text{shared}}$ be Epoch A's shared variance, and Epoch B's shared space is $\text{col}(U^B)$, the column space of Epoch B's factor matrix U^B . We computed the shared space alignment in three steps:

1. Compute $P_{U^B} \in \mathbb{R}^{N \times N}$, the projection matrix into $\text{col}(U^B)$. This is $P_{U^B} = VV^T$, where V is an orthonormal basis for $\text{col}(U^B)$, e.g., calculated via the singular value decomposition $\text{SVD}(U^B U^{B^T}) = VSV^T$.
2. Project $\Sigma^{A, \text{shared}}$ onto $\text{col}(U^B)$, which is $P_{U^B} \Sigma^{A, \text{shared}} P_{U^B}^T$.
3. Compute the alignment via $(\text{trace}(P_{U^B} \Sigma^{A, \text{shared}} P_{U^B}^T) / \text{trace}(\Sigma^{A, \text{shared}}))$.

This computation is used in [Figures 4E, 5F, and 7D](#). In [Figure 7D](#), main shared variance is projected into the readout space.

We computed the chance alignment by generating 100,000 random pairs of 1) a projection subspace of high dimensional neural space and 2) identity covariance in a random subspace of high dimensional neural space, and calculating the mean and 95th percentile of alignment. For [Figure 4E](#), the projection subspace and covariance are both 1-dimensional in a 15-dimensional ambient space (yielding mean chance alignment of 0.07 and 95th percentile chance alignment of 0.25). For [Figure 5E](#), the projection subspace and covariance are both 2-dimensional in a 15-dimensional ambient space for Monkey P (yielding mean chance alignment of 0.13 and 95th percentile chance alignment of 0.28) and 10-dimensional ambient space for Monkey R (yielding mean chance alignment of 0.20 and 95th percentile chance alignment of 0.40).

Trial-to-trial Variability of Shared and Private Signals

We analyzed trial-to-trial variability by applying FA to a data matrix $\in \mathbb{R}^N \times \text{num_trials}$ in which each column is one trial's population spike counts in a window (~ 1 s) locked to Go Cue. For each monkey, we chose the window length to be the longest amount of time such that at least 70% of trials for every target and training epoch were longer (Monkey P 0.9 s; Monkey R 1.3 s) ([Figure S4A](#)). We removed neurons with an average rate below 1.5 Hz because we had too few trials of non-zero activity to accurately estimate their variance. Over targets and epochs, a median of 11 out of 15 neurons were used for Monkey P and 9 out of 10 neurons were used for Monkey R. For all targets and epochs, the best-fitting shared dimensionality was 1, thus shared space alignment over training was between shared spaces of equal dimensionality ([Figure 4E](#)).

We considered two components to spike count variability: fixed firing rate spiking variability (such as in a Poisson process) and underlying firing rate variability ([Churchland et al., 2010](#)). We analyzed whether neural variability changes differed from changes expected with Poisson spiking variability by 1) normalizing each neuron's private variance by its mean rate and 2) performing FA on Poisson-simulated neurons ([Figures S5B–S5E](#)). The ratio of shared variance to total variance was calculated without any normalization of private variance ([Figure 4C](#)).

Decoder-timescale Private and Shared Variance

We analyze decoder-timescale population variance for each target and epoch by binning trials in 100 ms bins and applying FA. Neurons with firing rate below 1.5 Hz were excluded, as in the trial-to-trial variability analysis.

Sample-by-sample Contribution of Private and Shared Variance

Given an observed spike count at time t : $x(t)$, we used FA to estimate the shared and private variance contribution:

$$\hat{x}^{\text{shared}}(t) = E[x^{\text{shared}}(t) | x(t)] = E[Uz(t) | x(t)] = UU^T (UU^T + \Psi)^{-1} (x(t) - \mu)$$

$$\hat{x}^{\text{private}}(t) = x(t) - \hat{x}^{\text{shared}}(t) - \mu.$$

The main shared variance contribution is the linear projection of $\hat{x}^{\text{shared}}(t)$ into the main shared subspace (the top eigenvectors of Σ^{shared}).

Trial-to-trial Decoder Simulations

To test the how decoder-timescale main shared variance and private variance independently contribute to control, we separately supplied each trial's main shared trajectory $\hat{x}^{\text{main shared}}(t) + \mu$ and private trajectory $\hat{x}^{\text{private}}(t) + \mu$ as input to the decoder for the length of the observed trial. We quantified simulated performance using best progress to target, time to target hit (seconds), and normalized distance traveled to target hit. Best progress to target was computed as $(\max(\text{closest_distance_to_target} - \text{target_radius}, 0) / (\text{distance_center_to_target} - \text{target_radius})) \in [0, 1]$, such that a value of 1 indicates the target was reached and 0 indicates the cursor got no closer than the starting center point. Distance traveled to target was normalized by $\text{distance_center_to_target}$. To compare late learning shared and private control, we used trials for which both were successful ([Figures 6E and 6F](#)).

Relationship of Shared Variance to the Readout Space

We quantified shared variance in the decoder's readout space via $\text{trace}(P_R \Sigma^{\text{shared}} P_R^T)$, where P_R is the projection matrix onto the readout space ([Figure 7B](#)). It is the product of the shared variance magnitude $\text{trace}(\Sigma^{\text{shared}})$ and shared space alignment with the readout space $(\text{trace}(P_R \Sigma^{\text{shared}} P_R^T) / \text{trace}(\Sigma^{\text{shared}}))$, i.e., the fraction of shared variance captured in the readout space ([Figures 7D](#)

and S6A). For each target and training epoch, we calculated the main shared variance magnitude and readout space alignment and normalized by training epoch 1 values to yield a “magnitude growth factor” and an “alignment growth factor” (Figure 7D).

Shared Neural Trajectories at Fine-Timescale

To quantify shared neural trajectories, we zoomed in bin resolution to fine-timescale with 50 ms bins, performed FA on each target and epoch separately, and calculated $E[x^{\text{main shared}}(t)|x(t)]$ on each spike count time sample. We then trial averaged each trial's main shared trajectory (Figure 8BC). We quantified the magnitude of the main shared trajectory by taking the squared Frobenius norm (Figure 8D) and the stability of the trajectory by taking the correlation between each epoch's main shared trajectory and the final main shared trajectory (Figure 8E). The fraction of main shared variance explained by the average main shared trajectory was calculated as follows. We took the difference between each trial's main shared trajectory and the trial-averaged main shared trajectory within a 2 s window and calculated how much main shared variance remained in these residuals. The main shared variance due to the trial-averaged main shared trajectory was $1 - (\text{residual main shared variance}/\text{main shared variance})$ (Figure 8F). Finally, we supplied the trial-averaged main shared trajectory as input to the decoder and quantified the progress to target (Figure 8H).

QUANTIFICATION AND STATISTICAL ANALYSES

All analyses were performed for each target separately, and each plotted trend shows the mean and SEM over targets ($n = 8$) for each training epoch, as indicated in the Results and figure legends. Trends were analyzed for significance with linear regression. Additionally, for both subjects we grouped the first 3 epochs into an early block and the final 3 epochs into a late block, resulting in $n = 24$ in each group, from 3 epochs of 8 targets. For each analysis, we then performed a non-parametric one-sided Wilcoxon rank sum test of the appropriate comparison between early and late median. $p < 0.05$ were considered statistically significant. All statistical analyses were performed with custom scripts in MATLAB.