



OPEN

Selective modulation of cortical population dynamics during neuroprosthetic skill learning

Ellen L. Zippi^{1,5}, Albert K. You^{2,5}, Karunesh Ganguly^{3,4} & Jose M. Carmena^{1,2}✉

Brain-machine interfaces (BMIs) provide a framework for studying how cortical population dynamics evolve over learning in a task in which the mapping between neural activity and behavior is precisely defined. Learning to control a BMI is associated with the emergence of coordinated neural dynamics in populations of neurons whose activity serves as direct input to the BMI decoder (direct subpopulation). While previous work shows differential modification of firing rate modulation in this population relative to a population whose activity was not directly input to the BMI decoder (indirect subpopulation), little is known about how learning-related changes in cortical population dynamics within these groups compare. To investigate this, we monitored both direct and indirect subpopulations as two macaque monkeys learned to control a BMI. We found that while the combined population increased coordinated neural dynamics, this increase in coordination was primarily driven by changes in the direct subpopulation. These findings suggest that motor cortex refines cortical dynamics by increasing neural variance throughout the entire population during learning, with a more pronounced coordination of firing activity in subpopulations that are causally linked to behavior.

Learned behaviors are reinforced through mechanisms involving both cortical and subcortical structures^{1–4}. Just as behavioral actions are reinforced, so is the cortical population activity required to efficiently produce these actions^{5–7}. Studying mechanisms of cortical reinforcement underlying behavioral reinforcement can be challenging as the exact neural population controlling the desired behavior is unknown. Early studies of bio-feedback demonstrated that activity in motor cortex can be reinforced and volitionally controlled using reward and sensory feedback of the firing rate^{8,9}. Later, initial research on brain-machine interfaces (BMIs) showed that subjects could learn to control external devices (e.g. computer cursors or robotic arms) by learning to modulate the activity of a population of neurons and that the neural encoding of these prosthetic movements changed over time and decreased in variability^{10–14}. These BMIs allow for precisely defined mappings between recorded neural activity and behavior^{10,13,15}. Studies leveraging BMIs to study learning-related changes in cortical activity have demonstrated that neuroprosthetic skill learning can require the production of novel cortical dynamics to obtain skillful control^{11,16,17}.

While classical approaches examine individual neurons to understand fundamentally how motor cortical activity is reinforced, more recent methods looking at population-level activity have uncovered how dynamic processes may govern movement planning and execution^{18–29} as well as learning^{28,30–38}. Population-level activity is often characterized by low-dimensional dynamics that capture patterns of co-activation across neurons within a population³⁹. These population-level dynamics arise from input connectivity and within-population connectivity. Two parallel mechanisms have been proposed to reinforce specific cortical population dynamics; fast reinforcement of dynamics that naturally produce a desired behavior and slower reinforcement that refines them to result in more reliable production of neural activity patterns^{40,41}.

Previous studies have shown that neural populations are constrained to generate activity patterns within a pre-existing covariance structure within short timescales^{30,32,37,38}, suggesting that it is faster to learn to control and repurpose pre-existing cortical population dynamics than it is to modify them. When decoder perturbations that

¹Helen Wills Neuroscience Institute, University of California Berkeley, Berkeley, CA 94720, USA. ²Department of Electrical Engineering and Computer Sciences, University of California Berkeley, Berkeley, CA 94720, USA. ³Neurology and Rehabilitation Service, San Francisco VA Medical Center, San Francisco, CA 94121, USA. ⁴Department of Neurology, University of California, San Francisco, CA 94143, USA. ⁵These authors contributed equally: Ellen L. Zippi and Albert K. You. ✉email: jcarmena@berkeley.edu




Figure 1. Experimental setup and behavioral performance. **(a)** Experimental setup from Ganguly and Carmena, 2009. Activity recorded from direct neurons (blue) in M1/PMd was input into a fixed linear decoder and used to drive a computer cursor to perform a center-out task (see Methods for details). Activity from indirect neurons (red) was simultaneously recorded, but was not input into the decoder. **(b)** Performance improves over the first 15 training epochs for Monkey P and Monkey R. Each training epoch consists of 150 initiated trials. For some analyses, epochs were divided into groups of early (orange) and late (purple). Fraction of initiated trials that were successful increased over training epochs. **(c)** The time to reach a target decreased over training epochs. **(d)** Representative examples of single-trial cursor trajectories during the early and late learning.

change the behavioral output associated with specific neural activity were introduced after subjects had already achieved proficient control using a BMI there was an immediate deficit in performance. However, over training, subjects were able to recover performance of cursor control and furthermore, experienced a washout when the perturbation was removed, but only when the perturbation did not require alteration of the natural covariance pattern among the recorded neurons³². Other work has demonstrated that animals can learn to control BMIs that require neural patterns outside of the pre-existing covariance structure over the course of multiple days^{5,31}. This eventual modification of cortical dynamics suggests that learning novel skills requires the production of new underlying population activity that develops over longer timescales.

With motor cortical BMIs, a small subset of all possible neurons in motor cortex is selected to use as input to the decoder (direct neurons). These neurons exist within a large network of other motor cortical neurons (indirect neurons). While the selection of direct neurons from all recorded units in our experiments was arbitrary and there was initially no functional difference between the those selected to be direct and those that are not, previous work has shown that differences in the neural activity of these two groups emerge with learning^{16,42–47}. For example, it has been demonstrated that the task-relevant modulation of indirect neurons gradually reduces relative to direct neurons over learning¹⁶. Additionally, it has been shown in rodents that coherence develops between dorsal striatum and direct neurons, but not indirect neurons^{44–46}. In a study using 2-photon calcium imaging to record neural activity, mice initially modulated activity of both direct and indirect neurons, but predominantly modulated direct activity after learning⁴². Thus, it is likely that the initial cortical dynamics that produce desirable outcomes involve both direct neurons and the surrounding cortical network. Over time, as these cortical dynamics are refined, they may adapt to exclude neurons that do not directly drive behavior.

If this hypothesis is true, we expect differences in how cortical population-level dynamics within direct and indirect subpopulations change over time as well. As cortical dynamics are modified for more efficient control, the direct subpopulation would be expected to undergo further modification than the indirect subpopulation as additional modifications to indirect activity would not directly result in desirable outcomes. Here, we investigate this idea by studying recorded ensembles of motor cortical neurons while only a subset was assigned to have a causal role during BMI control and characterize the differential changes in coordinated neural dynamics between direct and indirect subpopulations.

Results

Two rhesus macaques (P and R) were chronically implanted with bilateral microelectrode arrays in primary motor and dorsal premotor cortices, with electrodes from a single hemisphere used for BMI control and subsequent analyses (see Methods). The monkeys learned to perform a two-dimensional, self-initiated, center-out BMI task, in which they drove a cursor under neural control to one of eight randomly instructed peripheral targets for a juice reward (Fig. 1a). The next trial was initiated by driving the cursor back to the center target. Trials from all days of the experiment were concatenated then separated into 150-trial epochs since the number of successful

trials was lower in early days of learning. Both animals increased the fraction of successful trials (Fig. 1b) and decreased the time to move the cursor from the center target to the peripheral target (Fig. 1c) over the course of the first 15 epochs. Example cursor trajectories from early and late learning are shown in Fig. 1d. To capture the correlates of learning before performance saturated, only the first 15 epochs were used for each animal.

Because the BMI decoder used in the experiment was novel to the subjects, they had to initially explore the neural population activity space. Over time, the subjects learn from the behavioral consequences of explored activity patterns and select target-achieving, rewarded patterns of activity. To characterize the neural dynamics associated with this neuroprosthetic learning, we examined recorded ensembles of motor cortical neurons from which only a subset was assigned a causal role during control. We define these neurons, whose activity was used as a direct input to the BMI decoder, as “direct neurons” (Monkey P, $N = 15$; Monkey R, $N = 10$). The remaining recorded motor cortical neurons, recorded using the same two 4×4 mm 64-channel microelectrode arrays (interelectrode distance 500 μm), whose activity was not used as direct input to the BMI, we define as “indirect neurons” (Monkey P, $N = 29$ – 69 ; Monkey R, $N = 87$ – 187). Spiking activity and waveforms for a representative direct and indirect unit are shown in Figure S1. For some analyses in which it is important to consider the same population across epochs we refer to indirect neurons that were stably recorded across all 15 epochs as “stable indirect neurons” (Monkey P, $N = 17$; Monkey R, $N = 14$). Stability of the indirect neurons was assessed using the methods described in Fraser and Schwartz⁴⁸. This method uses pairwise cross-correlograms, the autocorrelogram, waveform shape, and mean firing rate to classify neurons and has previously been used on recordings obtained from chronically implanted microelectrode arrays to assess the stability of neurons across days^{49–52}. Example waveforms from representative stable indirect units for each animal on the first day, middle day, and last day of recording are shown in Figure S2.

First, we examined how the neural firing rate variance changed in each subpopulation over learning. Changes in neural variance are often used as a proxy for neural exploration, as increasing the variance in firing rate allows for neurons to form different coordinated patterns of firing^{5,36,53–56}. Past work has shown neural activity fires in more coordinated patterns as behavior stabilizes, thereby decreasing the dimensionality in neural space over learning^{5,130–32}. We commonly refer to these low-dimensional spaces as manifolds or neural subspaces. In order to observe changes in these neural subspaces, epochs were separated between early and late for each animal (Epochs 1–7 and Epochs 9–15, respectively) to track differences as behavioral performance improved. The firing rate for each neuron was binned in 100 ms intervals. The binned firing rate variance for each neuron was then averaged across all neurons for each epoch. In early learning, the mean firing rate variance was significantly higher in the direct subpopulation than the indirect subpopulation for one subject and there was no significant difference for the other (Fig. S3). We normalized the firing rate variance for each subpopulation based on the mean variance in early learning to assess relative change in variance from early to late learning within each subpopulation. In both animals, we observed an increase in relative unit variance between early and late learning for both direct and indirect subpopulations (Fig. 2a). This increase in variance suggests a concerted effort of neural exploration that exists in a broader network that includes both direct and indirect neurons. To ensure that this increase in variance was not due to a change in the distribution of time spent at each target, we repeated this analysis within each target (Fig. 2b). Trials to each target were evenly divided across 15 epochs and relative unit variance during early (first 7 epochs) and late learning (last 7 epochs) was considered for each target individually. The relative unit variance increased within both direct and indirect subpopulations for Monkey P and for six of the eight targets for Monkey R (Table 1), indicating that this result was not due to a change in the distribution of time spent at each target. Because the change in relative unit variance was consistent across targets, subsequent analyses included trials to all targets to increase the statistical power associated with more trials. Furthermore, because the task required two-dimensional control to achieve success at all targets, grouping trials across all targets provided better insight into how learning occurs in a generalized two-dimensional space rather than for target-specific activity.

Previous work has shown that neurons fire in increasingly coordinated patterns as performance improves⁵. We consider these changes in coordinated firing as a proxy for consolidation of neural population dynamics since the neural variance is stabilizing onto low-dimensional subspaces. We use factor analysis (FA) to separate the neural variance in the population into two components—private and shared variances⁵⁷. The shared variance is the variance between neurons in the population and can be thought of as the underlying correlated firing pattern in the recorded population^{19,58}. Conversely, the private variance denotes the amount of variance each neuron has that is independent from the rest of the population. Past studies have explored the roles of these components in the direct neuron population, showing private variance as a proxy for exploration while an increase in shared variance is correlated to skill consolidation^{5,59}. Previous work has quantified the amount of coordinated neural activity as a measure of the balance between shared and private variance: the ratio of shared variance over total variance (SOT)^{5,58–61}. Here, we compared the proportion of the total neural variance that is captured in shared spaces for the combined direct and indirect population in each epoch as an estimate of coordination within the recorded population across learning.

Since the indirect population consisted of different units each epoch, we normalized the SOT to the mean SOT for early epochs (Fig. 3a). In both animals, we found that the relative SOT increased between early and late learning for the entire recorded population including both direct and indirect neurons. Together, with the increase in variance over learning, our results indicated a high level of increased coordination that occurs within the entire recorded population driven by increased exploration as BMI performance improves. To assess that the effect was not due to day-to-day differences in the population, we conducted the same analysis on neurons that were stably recorded across all 15 epochs, which yielded consistent results (Fig. 3b). Along with an increase in SOT, previous work has shown learning-related decreases in dimensionality of the shared neural subspace for the direct subpopulation⁵. We found a similar decrease in dimensionality for the entire stably recorded population (Fig. 3c) and furthermore found that the dimensionality of the neural subspace in each epoch was significantly correlated with the SOT of each epoch (Fig. 3d).




Figure 2. Neural variance increases with learning Variance was calculated for each neuron and then averaged across neurons. Relative variance was calculated by normalizing to the mean early variance within subpopulation. **(a)** Both direct and indirect subpopulations increased relative neural variance from early to late learning (Unpaired t-test; Monkey P: direct $p = 8.63e-5$, indirect $p = 0.002$; Monkey R: direct $p = 0.007$, indirect $p = 0.003$). **(b)** Analysis was repeated within target. Both direct and indirect subpopulations increased relative neural variance from early to late learning within target (Unpaired t-tests results reported in Table 1)

| Target | Monkey P | | Monkey R | |
|--------|----------|----------|--------------|----------|
| | Direct | Indirect | Direct | Indirect |
| 1 | 1.74e-04 | 0.003 | 0.006 | 8.64e-05 |
| 2 | 9.01e-04 | 1.72e-04 | 0.010 | 7.72e-04 |
| 3 | 0.004 | 0.001 | 0.015 | 0.002 |
| 4 | 0.006 | 1.35e-04 | 0.117 (n.s.) | 0.018 |
| 5 | 7.24e-05 | 0.012 | 0.001 | 2.12e-06 |
| 6 | 3.78e-04 | 1.25e-04 | 5.71e-04 | 4.51e-05 |
| 7 | 3.17e-04 | 5.25e-04 | 0.005 | 9.75e-04 |
| 8 | 6.07e-04 | 0.006 | 0.093 (n.s.) | 0.037 |

Table 1. Relative variance within target. P-values from unpaired t-test comparing relative neural variance in early and late learning calculated within each target separately.

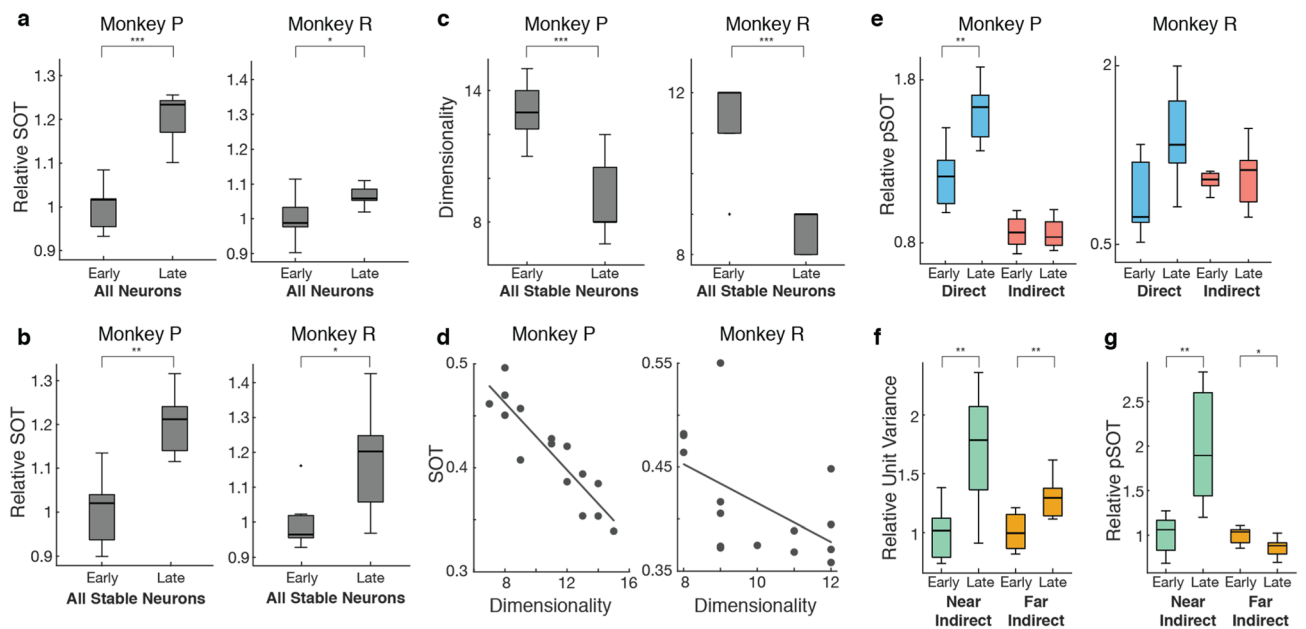


Figure 3. Increases in coordinated neural activity over learning are primarily driven by direct neurons. **(a)** The relative shared-over-total variance (SOT) ratio was calculated with respect to the mean SOT across the entire recorded population. Relative SOT increased between early and late learning, indicating an overall increase in coordination of neural activity in the entire recorded population (Unpaired t-test; Monkey P, $p = 1.31 \times 10^{-5}$; Monkey R, $p = 0.033$). **(b)** Relative SOT also increased between early and late learning for a stably recorded population consisting of the same units each epoch (Paired t-test; Monkey P, $p = 0.002$; Monkey R, $p = 0.027$). **(c)** Dimensionality of the neural subspace for the stably recorded population decreased from early to late learning (Unpaired t-test; Monkey P: $p = 4.10 \times 10^{-4}$; Monkey R: $p = 8.24 \times 10^{-5}$). **(d)** The dimensionality of the neural subspace is correlated with SOT (Linear regression; Monkey P: $R^2 = 0.813$, $p = 4.41 \times 10^{-6}$; Monkey R: $R^2 = 0.277$, $p = 0.044$). **(e)** Respective contributions of each sub-population to the SOT ratio (pSOT, see Methods for details) were calculated in early and late learning relative to contributions in early learning. Only direct pSOT relative to early learning increased from early to late learning (Paired t-test; Monkey P, direct $p = 0.003$, indirect $p = 0.805$; Monkey R, direct $p = 0.084$, indirect $p = 0.675$). To test that the change in direct was not due to a chance grouping of neurons, we repeated the t-test 500 times while shuffling direct and indirect labels and compared the true direct t-statistic to a distribution of t-statistics from the shuffled populations (Permutation test; Monkey P, direct $p = 0.014$; Monkey R, direct $p = 0.078$). **(f)** Both near and far indirect neurons exhibited significant increases in neural variance (Unpaired t-test; near $p = 0.006$, far $p = 0.007$). **(g)** Only near indirect neurons exhibited a significant increase in pSOT relative to early learning (Unpaired t-test; $p = 0.003$). Far indirect neurons exhibited a significant decrease in pSOT (Unpaired t-test; $p = 0.022$). To test that changes in pSOT are not due to a chance grouping of neurons, we repeated the t-tests 500 times while shuffling near and far labels (Permutation test; near, $p = 0.008$; far, $p = 0.002$).

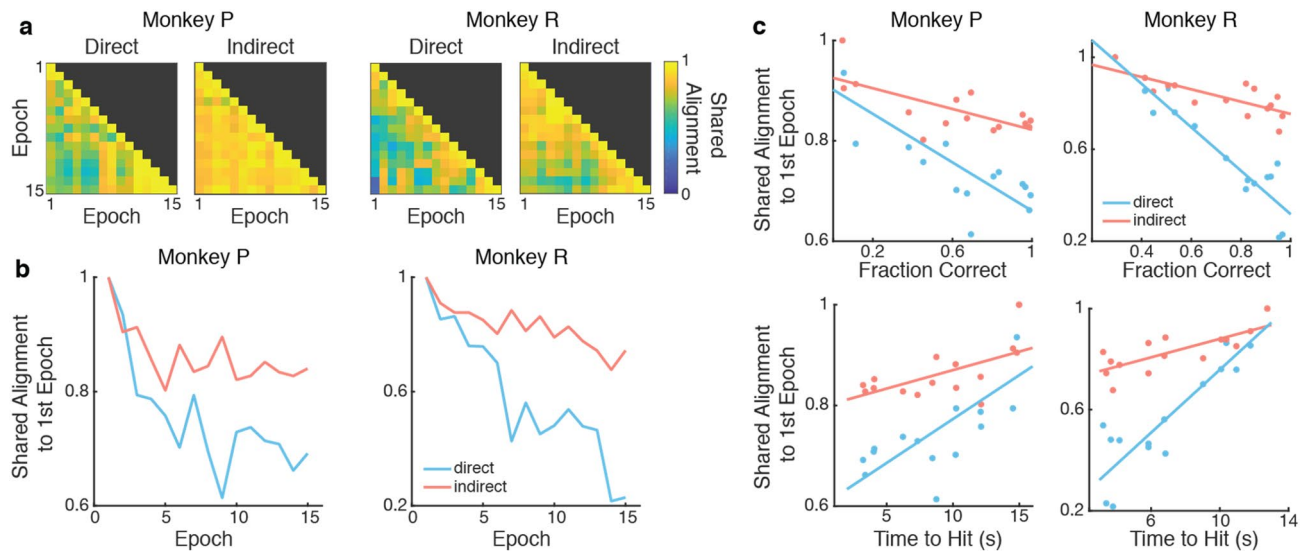


Figure 4. Rotation of low-dimensional neural subspace. **(a)** Shared alignment was calculated pairwise between epochs for each subpopulation. **(b)** Alignment diverges from the first epoch in both subpopulations (Linear regression; Monkey P: direct, $R^2 = 0.530$, $p = 0.001$, indirect, $R^2 = 0.352$, $p = 0.012$; Monkey R: direct $R^2 = 0.857$, $p = 4.63 \times 10^{-7}$, indirect $R^2 = 0.760$, $p = 1.41 \times 10^{-5}$). The slopes for direct and indirect shared alignment across epochs were significantly different (One-way ANCOVA; Monkey P: $p = 0.024$; Monkey R: $p = 5.84 \times 10^{-6}$). **(c)** Alignment is correlated with fraction correct (Top, linear regression; Monkey P: direct, $R^2 = 0.657$, $p = 2.0 \times 10^{-4}$, indirect $R^2 = 0.481$, $p = 4.2 \times 10^{-4}$; Monkey R: direct $R^2 = 0.870$, $p < 1.0 \times 10^{-6}$, indirect $R^2 = 0.589$, $p = 8.0 \times 10^{-4}$) and time to hit (Bottom, linear regression; Monkey P: direct, $R^2 = 0.535$, $p = 0.002$, indirect $R^2 = 0.380$, $p = 0.014$; Monkey R: direct $R^2 = 0.815$, $p < 1.0 \times 10^{-6}$, indirect $R^2 = 0.592$, $p = 8.0 \times 10^{-4}$).

While the increase in SOT indicates more coordination within the entire population, it does not explain whether these changes are driven by a specific subpopulation. To answer this question, we considered the partial shared-over-total variance (pSOT) ratio of the stably recorded population to see how the same population of neurons change coordinated firing activity over learning (see Methods). Intuitively, the pSOT ratio asks how much of the overall change in coordinated activity was driven by one subpopulation versus the other. We see that, relative to early learning, there was an increase in $\text{pSOT}_{\text{direct}}$ in late learning but not in $\text{pSOT}_{\text{indirect}}$ for both animals (Fig. 3e), indicating that the increase in coordination of population activity seen across the stably recorded population was driven by the increase in coordination of population activity within the direct subpopulation. While this result was only statistically significant in Monkey P, Monkey R (who had fewer direct neurons) exhibited the same trend. The larger increase in $\text{pSOT}_{\text{direct}}$ suggests that the increased neural exploration in the network was primarily a consequence of changes in coordinated patterns specific to the direct neurons. To further characterize these changes in the indirect neurons, we separated all of the indirect neurons in Monkey P into “far” and “near” indirect neurons. “Far” indirect neurons ($N = 29\text{--}69$) were those recorded on electrodes not containing direct neurons. In contrast, “near” indirect neurons ($N = 7\text{--}14$) were indirect neurons that existed on the same electrode shanks as direct neurons. Monkey R was excluded from these analyses due to recording too few near indirect neurons during several epochs ($N = 0\text{--}10$). We found that neural variance increased for both far and near indirect neurons between early and late learning (Fig. 3f). However, the pSOT only increased for the near indirect subpopulation and significantly decreased for the far indirect subpopulation (Fig. 3g). Together, these results suggest that while neural exploration exists in broader networks consisting of both direct and indirect neurons, activity in neurons closer in proximity to direct neurons becomes more coordinated than activity in neurons farther away from direct neurons. While the probability of synaptic connections among neurons does not depend strongly on interneuron distance, differences in near and far indirect neurons could be due to common synaptic inputs between the direct neurons and near indirect neurons from other brain regions into M1⁶².

To characterize how neural exploration modified the direct and indirect neural subspaces differently, we quantified these changes by calculating the shared alignment pairwise between each training epoch’s shared covariance matrix for each subpopulation according to the methods described in Athalye et al., 2017 (Fig. 4a). The shared alignment measures the similarity of covariance planes to compare how much of the shared space of one epoch projects onto the shared space of another epoch. Intuitively, given both two-dimensional shared subspaces, the shared alignment compares the angle between the two planes. Orthogonal planes, or subspaces, would result in a shared alignment of 0 and perfectly aligned planes would result in a shared alignment of 1. If the shared subspace consolidates with learning, as has been shown in direct subpopulations⁵, we would expect the shared subspace to rotate away from the initial subspace over learning. If the shared subspace remains fixed over learning, we would predict that the alignment between the first epoch and later epochs remains high, indicating little change in the coordinated activity of the population. Since we are interested in how the subspaces pertaining to specific populations change over time, we analyzed only the neurons that were stable across learning. We found that the shared alignment decreased from the first epoch for both subpopulations (Fig. 4b). This indicates that both




Figure 5. Neural variance modulates concomitantly between subpopulations. (a) Monkey P learned to perform BMI with a new decoder following proficient control with the old decoder. 8 experimental blocks were performed, alternating between a new decoder (cross) and the previously learned decoder (circle). Fraction of initiated trials that were successful increased over training blocks (left) and the time to reach a target decreased over training blocks (right). (b) Both subpopulations increase their neural variance over blocks (Linear regression; Direct $R^2=0.623$, $p=0.020$, Indirect $R^2=0.655$, $p=0.015$). The relative variances across blocks are correlated between subpopulations (Pearson's $r=0.856$, $p=0.007$). (c) Each point represents the change in relative variance between two consecutive blocks. The changes in relative variance within the direct and indirect subpopulations are correlated (Pearson's $r=0.820$, $p=0.024$).

subpopulations rotated their low-dimensional subspaces, suggesting that neurons may adapt on a network level that includes both direct and indirect neurons. Furthermore, this rotation of the low-dimensional subspaces is correlated with behavior (Fig. 4c). As the shared subspace diverges from where it began in the first epoch, the fraction of correct trials significantly increases and the time it takes for the cursor to reach the target significantly decreases. While this is true for both the direct and indirect subpopulations, the extent of rotation as measured by the shared alignment with the first epoch and the proportion of the variation in the shared alignment that is predictable from the behavior were higher for the direct subpopulation than the indirect subpopulation.

To further explore how indirect and direct neural activity may adapt together, we analyzed data from a second experiment in which Monkey P learned to perform the same BMI task with a new decoder following proficient control with the original learned decoder (Fig. 5a). The new decoder used the same direct neurons as the original decoder but the decoder weight assigned to each direct neuron was changed so that the same activity patterns result in different cursor movements when using the different decoders. Eight experimental blocks were performed over the course of four days, alternating between control with the new and previously learned decoder each day. The neural variance for direct and indirect neurons was calculated within each of these eight blocks. Note that only stable indirect neurons were used for this analysis since we wanted to explicitly track how the variance changed as a function of block number. We found that both subpopulations increased and decreased neural variance together over blocks, with similar changes in variance between blocks occurring in both direct and indirect neurons (Fig. 5b). Thus, increases in neural variability with changing decoders over shorter time-scales involved increased exploration not only by the direct neurons, but also by the supporting indirect neurons. Because there was an increase in neural variability in both subpopulations with each decoder swap, we were also able to assess whether the amount change in neural variance was similar between the two subpopulations. We found that the changes in the firing rate variance of the direct subpopulation were correlated to changes in the indirect subpopulation (Fig. 5c).

Discussion

In this study, we explored how changes in cortical population dynamics underlie skill learning. Specifically, we examined how cortical dynamics of subpopulations of neurons change over learning when the mapping of neural activity to behavior is precisely defined using a BMI. That is, how does the adaptation of the subpopulations used as inputs into a BMI decoder (direct neurons) compare to that of a subpopulation not used for decoding (indirect neurons)?

Our results revealed that learning-related neural state space exploration included neurons from both the direct and indirect subpopulations. Both subpopulations increased variance in their firing rates (Fig. 2) and there was an increase in coordination of neural firing patterns across the combined population that was correlated with a decrease in dimensionality of the neural space (Fig. 3). Previous literature has suggested that volitional modulation of neural activity in animals may be related to movements, cognitive imagery, or shifts in attention^{63,64}. Thus, it is possible that these learning-related changes in neural variance are the result of an underlying behavioral strategy through which the animals learn to modulate their neural activity. While we could not directly compare these changes in neural variance and coordinated activity to undetected movements made by the animals, both animals were observed to have minimal movements during BMI control¹¹.

An increase in neural variability with learning, as observed in our results, has also been seen in previous studies^{5,36}. One explanation for this increase in variability is that it allows the brain to explore new activity patterns that may improve behavior⁵. Another study finding the same increase in neural variability in early learning proposed that this increase may be the result of changing internal states or increased neural engagement³⁶. Thus,

this increase in variance of the firing rates of both direct and indirect neurons may be a result of an increased exploratory drive or a change in the animals' internal state, which may be indicative of its arousal or uncertainty about its environment.

Previous work observing increases in coordinated population activity suggested that these changes are characteristic of more stereotyped behavior over learning and this increased neural covariance has been associated with subjects making straighter, more direct paths to the targets⁵. Furthermore, changes in covariance structure have been shown to relate to synaptic connectivity⁶⁵. Thus, when considering the entire recorded population, it appears as though the whole population adapts together to facilitate learning. However, past work has shown differences in adaptation between direct and indirect neurons^{16,43,45}. When we considered the relative contribution of each subpopulation to the overall increase in coordination, we found that the indirect subpopulation contributed very little to the increase in coordinated patterns (Fig. 3e). That is, while we witnessed an increase in coordination in the entire recorded population, there was less within-group coordination in indirect neurons compared to direct neurons. Our metrics of coordination (e.g., SOT) rely on averaging the amount of correlated activity between pairs of neurons. Coordination of firing activity occurring more heavily in one subpopulation would nevertheless increase the SOT in the entire population. Altogether, while both subpopulations exhibit similar levels of exploration over learning as demonstrated by an increase in neural variance, the exploration by the indirect subpopulation results in less of an increase in coordination of neural firing patterns than that of the direct subpopulation.

Changes in coordinated neural activity resulted in rotations of the neural space over learning in both direct and indirect subpopulations (Fig. 4). The neural space can be intuitively thought of as the lower-dimensional space in which co-activations of neurons exist and rotating this neural space corresponds to adjusting which neurons are more active given the state of other neurons in the population. Following this interpretation, both direct and indirect subpopulations similarly adapted their coordinated firing patterns within subpopulation over learning. This adaptation in coordinated firing patterns was also correlated with behavioral improvements in the task for both the direct and indirect subpopulations, however, the relationship was stronger for the direct subpopulation than the indirect subpopulation for both subjects. This result, along with the increase in pSOT observed in the direct subpopulation, suggests that as the direct subpopulation increases coordination of neural activity its covariance exhibits greater changes than that of the indirect subpopulation over learning.

We also examined how the neural firing rate variance of both the direct and indirect subpopulations changed in an experiment where two decoders were swapped each day, requiring changes in neural activity over short timescales. We found that changes in neural variance occurred with each decoder swap and were proportional between the two subpopulations (Fig. 5). This suggests that increases in neural variability occur over both short timescales, as seen in this experiment, and over longer timescales, as seen in the initial 15 epochs, involve both neurons within the supporting cortical network and direct neurons. Furthermore, the changes in neural variance between the two subpopulations were correlated. Consequently, both the direct and indirect subpopulations may be adapting via the same mechanism but to different extents. However, this experiment was limited to only four days of switching between decoders. It is unclear whether or not these parallel changes in neural variance between the direct and indirect subpopulations would continue if the animal was given more extensive practice.

Our findings that both direct and indirect subpopulations increase neural variability similarly but exhibit differential changes in coordination could be explained by existing hypotheses on how the brain learns to refine coordinated neural dynamics⁴⁰. Specifically, small networks of cortical neurons may be driven by upstream cortical and subcortical inputs. We found that indirect neurons adapted in similar ways as direct neurons, suggesting that some indirect neurons may in fact be adapting with or alongside direct neurons. When we investigated how distance from direct neurons influenced these results, we found that indirect neurons in closer spatial proximity to direct neurons increased coordination more than indirect neurons that were farther away (Fig. 3g). While the upstream projections from cortical and subcortical structures are not necessarily spatially organized, our results are consistent with the hypothesis that upstream structures may be driving changes in smaller groups of inter-connected neurons³⁵. Furthermore, it has previously been shown that when disparities are present between the control space and neural space (i.e. how well the decoders aligned to the natural firing patterns of the neurons), neurons with larger disparities adapt more over learning compared to neurons with smaller disparities^{5,17,66,67}. A recent study has shown that task-related neurons, consisting of direct neurons as well as task-modulated indirect neurons, increase coherency to slow-wave activity (SWA) during sleep which has been linked with consolidation⁴³. This suggests that in addition to online task practice, neural reactivations during sleep can aid in exploring the contributions of direct and indirect neural populations relative to successful outcomes and reward. Notably, indirect neurons that were closely tied to reward were preserved and resembled direct neurons; this might explain why some indirect neurons were modified during neuroprosthetic skill acquisition. This also provides further evidence that mechanisms of reinforcement learning may underlie our observed phenomena. Thus, it is quite plausible that adaptation of neural activity over BMI learning is attributed to finding the clusters of neurons with a direct effect on behavior, which may include both direct and indirect neurons, depending on their specific network connectivity and temporal association with successful outcomes.

In this study, we used factor analysis (FA) to find correlations in the neural activity of the recorded population. Underlying this model are latent factors, which are variables that coarsely group neurons together based on coordinated activity patterns. Importantly, the activity of a single neuron can be associated with multiple latent factors. While we were agnostic to what the latent factors in FA may correspond to in this study, they may be physiologically analogous to upstream connections from cortical or subcortical structures that drive changes in small clusters of neurons containing direct neurons. The idea that neural reinforcement is dependent on cortico-cortical and cortico-striatal circuits, similar to behavioral reinforcement, has been previously supported by studies using BMIs. For example, as rodents learned to produce specific patterns of cortical activity, coherence between these neurons and dorsal striatum emerged and neurons in dorsal striatum developed target-predictive

modulation of firing activity^{44–46}. Furthermore, mice without functional NMDA receptors in striatal projection neurons could not learn to re-enter a cortical pattern that led to reward. Thus, cortico-striatal plasticity is necessary for learning to efficiently produce the cortical activity patterns required to obtain rewards. These findings along with the results from our study further support the hypothesis that smaller clusters of neurons, which may include both direct and indirect neurons, are adapted over learning more than clusters of neurons that do not drive behavior.

Overall, our results demonstrate that the brain learns to modify cortical population dynamics in subpopulations relevant for behavioral control. When using a BMI, we find that neurons with direct input to the decoder as well as neurons in the surrounding cortical network increase exploration and consolidate their firing activity onto a low-dimensional neural space. The degree of coordination among the population is dependent on the relationship of the neural activity to the behavioral output. Thus, the brain may not be reinforcing the activity of single neurons, but rather reinforcing cortical population dynamics that are relevant to producing a desired behavior. These findings indicate that the brain learns to control a BMI by refining cortical population-level dynamics, suggesting that BMI decoders extracting information based on population-level statistics, such as the covariance structure of the population, may be more effective compared to traditional decoding methods based on the statistics of individual neurons. Dimensionality reduction techniques such as FA allow us to pull out the correlated activity in a population of neurons. That is, the shared variance obtained using FA represents the concerted activity of a population and if we assume that the uncorrelated activity is largely noise, then using these population-level statistics effectively increases the signal-to-noise ratio of the neural activity. Building decoders based on these smoothed neural signals may translate to smoother output signals (e.g. cursor movements).

In conclusion, this study demonstrates an emergence of coordinated population dynamics within both populations of neurons whose activity is directly used as input for BMI control as well as within the surrounding network. The extent to which these subpopulations modify their coordinated activity varies, with the direct subpopulation exhibiting larger changes. Understanding the role of modifications of adjacent indirect activity in obtaining precise control of a BMI may help us understand the neural adaptation that is required for achieving long-term, stable control of a BMI.

Methods:

Experimental design. *Animal subjects.* 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.

Two adult male rhesus monkeys (*Macaca mulatta*) were chronically implanted in the brain with arrays of 64 microelectrodes (Innovative Neurophysiology, Durham NC)¹¹. 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 included in the direct ensembles (Monkey P: right M1; Monkey R: left M1) and only activity from the same hemisphere was included in the indirect ensemble. 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 brain⁶⁸.

Electrophysiology. Neural activity was recorded using the MAP system (Plexon, Dallas TX). Stable units, to be part of the direct ensemble, were selected based on waveform shape, amplitude, relationship to other units on the same channel, interspike 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 (Wave-tracker; Plexon).

Direct units are defined as the units being used to control the BMI. Indirect units consisted of the remaining recorded units. For analyses including only stable units from the same hemisphere, stability in the indirect ensemble was assessed using pairwise cross-correlograms, autocorrelograms, waveform shapes, and mean firing rates⁴⁸.

Experimental setup and behavioral training. 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, Kingston ON) exoskeleton which restricted shoulder and elbow movements to the horizontal plane.

BMI tasks. Data from Ganguly and Carmena¹¹, in which subjects performed a self-initiated, eight-target, center-out reaching task, was analyzed. In these experiments, a cursor on a screen was continuously controlled by neural activity. Subjects self-initiated trials by moving the cursor to a center target. One of the eight peripheral targets was randomly selected each trial. Self-initiated trials consisted of those in which the animal moved the cursor to the center target and held for 250–300 ms. Successful trials required the animal to move the cursor to the peripheral target within 15 s of initiating the trial and hold the cursor at the target for 250–300 ms. Successful trials resulted in a juice reward; failed trials were repeated. During BMI control, both arms were lightly restrained to restrict arm movement during BMI control. During selected sessions, video and surface electromyogram (EMG) recordings from proximal muscle groups were performed to confirm minimal arm movements occurred during BMI control.

After the initial 19-days of performing the BMI task with a fixed-decoder, Monkey P learned a second decoder over the course of four days. Within each of these four days, Monkey P performed one training block of the new decoder, followed by one training block of the old decoder. Both of these decoders were fixed and used the same direct ensemble as input.

Preprocessing pipeline. For all analyses, neural data was binned into 100 ms bins to match the decoder time-scale. Additionally, learning was analyzed over “training epochs,” where each epoch consisted of 150 self-initiated trials. Learning took place over the first 15 training epochs (2250 self-initiated trials). We chose to analyze the data across training epochs, rather than days, to eliminate the effect of variable numbers of trials each day. Only the first 15 training epochs (2250 self-initiated trials) were analyzed; we defined early and late learning as the first and last seven of these training epochs, respectively. Monkey P initiated a total of 3589 trials. Monkey R initiated a total of 2357 trials.

Statistical analysis. Factor analysis. Shared-over-total variance ratio. Factor analysis (FA) was conducted on the neural population for each epoch to observe underlying correlated neural activity. FA decomposes population signals into correlated and uncorrelated components. For a given neuron i , correlated activity is represented by the shared variance Σ_{ii}^{shared} , and the degree to which the activity was correlated over learning was represented by the ratio of shared-over-total variance of the neural population (SOT). We calculated the SOT ratio according to the methods described in Althaye et al.⁵.

$$SOT = \frac{\text{trace}(\Sigma^{shared})}{\text{trace}(\Sigma^{total})}$$

Since FA decomposes the neural activity into correlated and uncorrelated components, we can reduce the dimensionality of the neural data by examining how much of the neural variance is captured by the shared (correlated) components. The number of dimensions was selected such that 90% (or more) of the shared variance was captured. A scree plot quantifying the variance captured by each factor for the combined population is shown in Figure S4.

Total variance. We also considered how the total variance changed between early and late learning. This was the sum of the private and shared variances $\Sigma_{ii}^{total} = \Sigma_{ii}^{shared} + \Sigma_{ii}^{private}$.

Partial shared-over-total variance ratio. We quantified the respective contributions of subpopulations to the SOT ratio using the partial shared-over-total variance (pSOT) ratio. Here, we compared the sum of the shared variance for each subpopulation over the total variance for the entire population. A relative measure was used to account for the fact that the direct and indirect ensembles were different sizes. That is,

$$pSOT_{direct} = \frac{\text{trace}(\Sigma_{direct}^{shared})}{\text{trace}(\Sigma_{all}^{private}) + \text{trace}(\Sigma_{all}^{shared})}$$

$$pSOT_{indirect} = \frac{\text{trace}(\Sigma_{indirect}^{shared})}{\text{trace}(\Sigma_{all}^{private}) + \text{trace}(\Sigma_{all}^{shared})},$$

where i denotes the covariance matrix for the (sub)population i .

Shared space alignment. We used the “shared space alignment” to measure the similarity between the shared variance of two different training epochs. The shared space alignment is the fraction of shared variance from one epoch captured in the shared space of a second epoch and thus ranges from 0 to 1. We calculated shared alignment according to the methods described in Athalye et al.⁵. Given two epochs, A and B, we first compute the projection matrix into Epoch B’s shared space, $col(U^B)$. We then project $\Sigma^{A,shared}$ onto B’s shared space, $P_{UB} \Sigma^{A,shared} P_{UB}^T$. Finally, the alignment is calculated,

$$Shared\ Alignment = \frac{\text{trace}(P_{UB} \Sigma^{A,shared} P_{UB}^T)}{\text{trace}(\Sigma^{A,shared})}$$

Quantification and statistical analyses. All analyses were performed on each training epoch separately. Trends were analyzed for significance with linear regressions. Epochs 1–7 and Epochs 9–15 were grouped into “early” and “late,” respectively. Epoch 8 was omitted so that there were equal numbers of epochs for both early and late learning. Groupings of early and late epochs were compared using an unpaired two-sample t-test when comparing unstable populations of neurons and a paired one-sample t-test when comparing stable populations. Significance was considered as $p < 0.05$. A single asterisk was used to denote significance of at least $p < 0.05$, a double asterisk was used to denote $p < 0.01$, and a triple asterisk was used to denote $p < 0.001$.

For t-tests resulting in a significant change in one group but not the other (e.g. direct but not indirect; near but not far), a random permutation test in which the group labels were shuffled 500 times was calculated to

ensure that the results were not a consequence of a small sample size. In all cases, the t-statistic for the group with the significant change was outside 95% of the distribution of t-statistics obtained from the shuffled labels tests indicating that the significant group was not a chance combination of neurons.

Ethical approval. All procedures were conducted in compliance with the NIH Guide for the Care and Use of Laboratory Animals and reporting in the manuscript follows the recommendations in the ARRIVE guidelines. All procedures were approved by the University of California at Berkeley Institutional Animal Care and Use Committee.

Data availability

The datasets analyzed during the current study will be publicly available on Figshare (https://figshare.com/articles/dataset/Selective_modulation_of_cortical_population_dynamics_during_neuroprosthetic_skill_learning/21111688) upon publication.

Received: 2 February 2022; Accepted: 9 September 2022

Published online: 24 September 2022

References

- Donchin O. *et al.*, Cerebellar regions involved in adaptation to force field and visuomotor perturbation, *J. Neurophysiol.*, **14** (2012).
- Krakauer, J. W. *et al.* Differential cortical and subcortical activations in learning rotations and gains for reaching: A PET study. *J. Neurophysiol.* **91**, 10 (2004).
- Sing, G. C. & Smith, M. A. Reduction in learning rates associated with anterograde interference results from interactions between different timescales in motor adaptation. *PLoS Comput. Biol.* **6**(8), e1000893. <https://doi.org/10.1371/journal.pcbi.1000893> (2010).
- Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction* (The MIT Press, 1998).
- Athalye, V. R., Ganguly, K., Costa, R. M. & Carmena, J. M. Emergence of coordinated neural dynamics underlies neuroprosthetic learning and skillful control. *Neuron* **93**(4), 955–970.e5. <https://doi.org/10.1016/j.neuron.2017.01.016> (2017).
- Costa, R. M. A selectionist account of de novo action learning. *Curr. Opin. Neurobiol.* **21**(4), 579–586. <https://doi.org/10.1016/j.conb.2011.05.004> (2011).
- Tumer, E. C. & Brainard, M. S. Performance variability enables adaptive plasticity of ‘crystallized’ adult birdsong. *Nature* **450**(7173), 1240–1244. <https://doi.org/10.1038/nature06390> (2007).
- Fetz, E. E. Operant Conditioning of Cortical Unit Activity, **163**(3870), 955–958, <https://doi.org/10.1126/science.163.3870.955> (1969).
- Fetz, E. E. & Baker, M. A. Operantly conditioned patterns on precentral unit activity and correlated responses in adjacent cells and contralateral muscles. *J. Neurophysiol.* **36**(2), 179–204. <https://doi.org/10.1152/jn.1973.36.2.179> (1973).
- Carmena, J. M. *et al.* Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol.* **1**(2), e42. <https://doi.org/10.1371/journal.pbio.0000042> (2003).
- Ganguly, K. & Carmena, J. M. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol.* **7**(7), e1000153. <https://doi.org/10.1371/journal.pbio.1000153> (2009).
- Musallam, S., Corneil, B. D., Greger, B., Scherberger, H. & Andersen, R. A. Cognitive control signals for neural prosthetics. *Science* **305**(5681), 258–262. <https://doi.org/10.1126/science.1097938> (2004).
- Taylor, D. M. Direct cortical control of 3D neuroprosthetic devices. *Science* **296**(5574), 1829–1832. <https://doi.org/10.1126/science.1070291> (2002).
- Zacksenhouse, M. *et al.* Cortical Modulations Increase in Early Sessions with Brain-Machine Interface. *PLoS ONE* **2**(7), e619. <https://doi.org/10.1371/journal.pone.0000619> (2007).
- Shenoy, K. V. & Carmena, J. M. Combining decoder design and neural adaptation in brain-machine interfaces. *Neuron* **84**(4), 665–680. <https://doi.org/10.1016/j.neuron.2014.08.038> (2014).
- Ganguly, K., Dimitrov, D. E., Wallis, J. D. & Carmena, J. M. Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nat. Neurosci.* **14**(5), 662–667. <https://doi.org/10.1038/nn.2797> (2011).
- Orsborn, A. L. *et al.* Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron* **82**(6), 1380–1393. <https://doi.org/10.1016/j.neuron.2014.04.048> (2014).
- Ames, K. C. & Churchland, M. M. Motor cortex signals for each arm are mixed across hemispheres and neurons yet partitioned within the population response. *Elife* **8**, e46159. <https://doi.org/10.7554/eLife.46159> (2019).
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical preparatory activity: Representation of movement or first cog in a dynamical machine?. *Neuron* **68**(3), 387–400. <https://doi.org/10.1016/j.neuron.2010.09.015> (2010).
- Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**(7405), 51–56. <https://doi.org/10.1038/nature11129> (2012).
- Heming, E. A., Cross, K. P., Takei, T., Cook, D. J. & Scott, S. H. Independent representations of ipsilateral and contralateral limbs in primary motor cortex. *Elife* **8**, e48190. <https://doi.org/10.7554/eLife.48190> (2019).
- Kao, J. C. *et al.* Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nat. Commun.* **6**(1), 7759. <https://doi.org/10.1038/ncomms8759> (2015).
- Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* **17**(3), 440–448. <https://doi.org/10.1038/nn.3643> (2014).
- Pandarinath, C. *et al.* Latent factors and dynamics in motor cortex and their application to brain-machine interfaces. *J. Neurosci.* **38**(44), 9390–9401. <https://doi.org/10.1523/JNEUROSCI.1669-18.2018> (2018).
- Russo, A. A. *et al.* Motor cortex embeds muscle-like commands in an untangled population response. *Neuron* **97**(4), 953–966.e8. <https://doi.org/10.1016/j.neuron.2018.01.004> (2018).
- Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: A dynamical systems perspective. *Annu. Rev. Neurosci.* **36**(1), 337–359. <https://doi.org/10.1146/annurev-neuro-062111-150509> (2013).
- Suresh, A. K. *et al.* Neural population dynamics in motor cortex are different for reach and grasp. *Elife* **9**, e58848. <https://doi.org/10.7554/eLife.58848> (2020).
- Zhou, X., Tien, R. N., Ravikumar, S. & Chase, S. M. Distinct types of neural reorganization during long-term learning. *J. Neurophysiol.* **121**(4), 1329–1341. <https://doi.org/10.1152/jn.00466.2018> (2019).
- Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nat. Commun.* **7**(1), 13239. <https://doi.org/10.1038/ncomms13239> (2016).
- Golub, M. D. *et al.* Learning by neural reassociation. *Nat. Neurosci.* **21**(4), 607–616. <https://doi.org/10.1038/s41593-018-0095-3> (2018).

31. Oby, E. R. *et al.* New neural activity patterns emerge with long-term learning. *Proc. Natl. Acad. Sci.* **116**(30), 15210–15215. <https://doi.org/10.1073/pnas.1820296116> (2019).
32. Sadtler, P. T. *et al.* Neural constraints on learning. *Nature* **512**(7515), 423–426. <https://doi.org/10.1038/nature13665> (2014).
33. Vyas, S. *et al.* Neural population dynamics underlying motor learning transfer. *Neuron* **97**(5), 1177–1186.e3. <https://doi.org/10.1016/j.neuron.2018.01.040> (2018).
34. Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation through neural population dynamics. *Annu. Rev. Neurosci.* **43**(1), 249–275. <https://doi.org/10.1146/annurev-neuro-092619-094115> (2020).
35. Vyas, S., O’Shea, D. J., Ryu, S. I. & Shenoy, K. V. Causal role of motor preparation during error-driven learning. *Neuron* **106**(2), 329–339.e4. <https://doi.org/10.1016/j.neuron.2020.01.019> (2020).
36. Hennig, J. A. *et al.* Learning is shaped by abrupt changes in neural engagement. *Nat. Neurosci.* **24**(5), 727–736. <https://doi.org/10.1038/s41593-021-00822-8> (2021).
37. Perich, M. G., Gallego, J. A. & Miller, L. E. A neural population mechanism for rapid learning. *Neuron* **100**(4), 964–976.e7. <https://doi.org/10.1016/j.neuron.2018.09.030> (2018).
38. Sun, X. *et al.* Cortical preparatory activity indexes learned motor memories. *Nature* **602**(7896), 274–279. <https://doi.org/10.1038/s41586-021-04329-x> (2022).
39. Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A. Neural manifolds for the control of movement. *Neuron* **94**(5), 978–984. <https://doi.org/10.1016/j.neuron.2017.05.025> (2017).
40. Athalye, V. R., Carmena, J. M. & Costa, R. M. Neural reinforcement: re-entering and refining neural dynamics leading to desirable outcomes. *Curr. Opin. Neurobiol.* **60**, 145–154. <https://doi.org/10.1016/j.conb.2019.11.023> (2020).
41. Dayan, E. & Cohen, L. G. Neuroplasticity subserving motor skill learning. *Neuron* **72**(3), 443–454. <https://doi.org/10.1016/j.neuron.2011.10.008> (2011).
42. Clancy, K. B., Koralek, A. C., Costa, R. M., Feldman, D. E. & Carmena, J. M. Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat. Neurosci.* **17**(6), 807–809. <https://doi.org/10.1038/nn.3712> (2014).
43. Gulati, T., Ramanathan, D. S., Wong, C. C. & Ganguly, K. Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. *Nat. Neurosci.* **17**(8), 1107–1113. <https://doi.org/10.1038/nn.3759> (2014).
44. Koralek, A. C., Jin, X., Long, J. D. II., Costa, R. M. & Carmena, J. M. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* **483**(7389), 331–335. <https://doi.org/10.1038/nature10845> (2012).
45. Koralek, A. C., Costa, R. M. & Carmena, J. M. Temporally precise cell-specific coherence develops in corticostriatal networks during learning. *Neuron* **79**(5), 865–872. <https://doi.org/10.1016/j.neuron.2013.06.047> (2013).
46. Neely, R. M., Koralek, A. C., Athalye, V. R., Costa, R. M. & Carmena, J. M. Volitional modulation of primary visual cortex activity requires the Basal Ganglia. *Neuron* **97**(6), 1356–1368.e4. <https://doi.org/10.1016/j.neuron.2018.01.051> (2018).
47. So, K., Koralek, A. C., Ganguly, K., Gastpar, M. C. & Carmena, J. M. Assessing functional connectivity of neural ensembles using directed information. *J. Neural Eng.* **9**(2), 026004. <https://doi.org/10.1088/1741-2560/9/2/026004> (2012).
48. Fraser, G. W. & Schwartz, A. B. Recording from the same neurons chronically in motor cortex. *J. Neurophysiol.* **107**(7), 1970–1978. <https://doi.org/10.1152/jn.01012.2010> (2012).
49. Golub, M. D., Yu, B. M., Schwartz, A. B. & Chase, S. M. Motor cortical control of movement speed with implications for brain-machine interface control. *J. Neurophysiol.* **112**(2), 411–429. <https://doi.org/10.1152/jn.00391.2013> (2014).
50. Downey, J. E., Schwed, N., Chase, S. M., Schwartz, A. B. & Collinger, J. L. Intracortical recording stability in human brain–computer interface users. *J. Neural Eng.* **15**(4), 046016. <https://doi.org/10.1088/1741-2552/aab7a0> (2018).
51. Vasileva, L. N. & Bondar, I. V. Long-term stable recording of single-neuron spike activity in the amygdala in conscious rabbits. *Neurosci. Behav. Physiol.* **51**(3), 322–331. <https://doi.org/10.1007/s11055-021-01075-5> (2021).
52. Yoo, S. B. M., Tu, J. C. & Hayden, B. Y. Multicentric tracking of multiple agents by anterior cingulate cortex during pursuit and evasion. *Nat. Commun.* **12**(1), 1985. <https://doi.org/10.1038/s41467-021-22195-z> (2021).
53. Arduin, P.-J., Fregnac, Y., Shulz, D. E. & Ego-Stengel, V. ‘Master’ neurons induced by operant conditioning in rat motor cortex during a brain–machine interface task. *J. Neurosci.* **33**(19), 8308–8320. <https://doi.org/10.1523/JNEUROSCI.2744-12.2013> (2013).
54. Dhawale, A. K., Smith, M. A. & Ölveczky, B. P. The Role of Variability in Motor Learning. *Annu. Rev. Neurosci.* **40**(1), 479–498. <https://doi.org/10.1146/annurev-neuro-072116-031548> (2017).
55. Mandelblat-Cerf, Y., Paz, R. & Vaadia, E. Trial-to-trial variability of single cells in motor cortices is dynamically modified during visuomotor adaptation. *J. Neurosci.* **29**(48), 15053–15062. <https://doi.org/10.1523/JNEUROSCI.3011-09.2009> (2009).
56. Sternad, D. It’s not (only) the mean that matters: variability, noise and exploration in skill learning. *Curr. Opin. Behav. Sci.* **20**, 183–195. <https://doi.org/10.1016/j.cobeha.2018.01.004> (2018).
57. Everitt, B. S. *An Introduction to Latent Variable Models* (Springer, 1984). <https://doi.org/10.1007/978-94-009-5564-6>.
58. Yu, B. M. *et al.* Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* **102**(1), 614–635. <https://doi.org/10.1152/jn.90941.2008> (2009).
59. Athalye, V. R., Santos, F. J., Carmena, J. M. & Costa, R. M. Evidence for a neural law of effect. *Science* **359**(6379), 1024–1029. <https://doi.org/10.1126/science.aao6058> (2018).
60. Khanna, P. *et al.* Low-frequency stimulation enhances ensemble co-firing and dexterity after stroke. *Cell* **184**(4), 912–930.e20. <https://doi.org/10.1016/j.cell.2021.01.023> (2021).
61. A. You, A. Singhal, H. Moorman, S. Gowda, and J. M. Carmena, “Neural Correlates of Control of a Kinematically Redundant Brain-Machine Interface,” in *2019 9th International IEEE/EMBS Conference on Neural Engineering (NER)*, San Francisco, CA, USA, Mar. 2019, pp. 554–557. <https://doi.org/10.1109/NER.2019.8717010>.
62. Song, S., Sjöström, P. J., Reigl, M., Nelson, S. & Chklovskii, D. B. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol.* **3**(3), e68. <https://doi.org/10.1371/journal.pbio.0030068> (2005).
63. Hwang, E. J., Bailey, P. M. & Andersen, R. A. Volitional control of neural activity relies on the natural motor repertoire. *Curr. Biol.* **23**(5), 353–361. <https://doi.org/10.1016/j.cub.2013.01.027> (2013).
64. Fetz, E. E. Volitional control of neural activity: implications for brain–computer interfaces: Volitional control of neural activity. *J. Physiol.* **579**(3), 571–579. <https://doi.org/10.1113/jphysiol.2006.127142> (2007).
65. Okun, M. *et al.* Diverse coupling of neurons to populations in sensory cortex. *Nature* **521**(7553), 511–515. <https://doi.org/10.1038/nature14273> (2015).
66. Chase, S. M., Kass, R. E. & Schwartz, A. B. Behavioral and neural correlates of visuomotor adaptation observed through a brain-computer interface in primary motor cortex. *J. Neurophysiol.* **108**(2), 624–644. <https://doi.org/10.1152/jn.00371.2011> (2012).
67. Jarosiewicz, B. *et al.* Functional network reorganization during learning in a brain–computer interface paradigm. *Proc. Natl. Acad. Sci.* **105**(49), 19486–19491. <https://doi.org/10.1073/pnas.0808113105> (2008).
68. G. Paxinos, X.-F. Huang, and A. W. Toga, *The Rhesus Monkey Brain in Stereotaxic Coordinates*. (2000).

Acknowledgements

We thank N. Vendrell-Llopis for helpful comments on the manuscript and V. Athalye for valuable discussions. This work was supported by the National Science Foundation Graduate Research Fellowship (to E.L.Z. and to A.K.Y.), the Department of Veterans Affairs, Veterans Health Administration, Rehabilitation Research and Development, and the American Heart Association/American Stroke Association (to K.G.), the Alfred P. Sloan

Foundation, the Christopher and Dana Reeve Foundation, the National Science Foundation CAREER Award #0954243, the Defense Advanced Research Projects Agency contract N66001-10-C-2008, and the National Institute of Health Award R01NS106094 (to J.M.C.).

Author contributions

E.L.Z. and A.Y.K. analyzed the data. K.G. and J.M.C. designed and performed research. E.L.Z., A.Y.K., and J.M.C. wrote the manuscript text. E.L.Z. and A.Y.K. prepared the manuscript figures. K.G. edited the manuscript. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-20218-3>.

Correspondence and requests for materials should be addressed to J.M.C.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022