Selective direct motor cortical influence during naturalistic climbing

Natalie Koh[1], Zhengyu Ma[1], Abhishek Sarup, Amy C. Kristl, Mark Agrios, Margaret Young[2], and Andrew Miri*

Department of Neurobiology, Northwestern University, Evanston, IL 60208, USA

[1]These authors contributed equally.
[2]Present address: University Medical Center Department of Ophthalmology, Bernstein Center for Computational Neuroscience, and International Max Planck Research School for Neurosciences, Göttingen, 37073, Germany

*Corresponding author (andrewmiri@northwestern.edu)

It remains poorly resolved when and how motor cortical output directly influences limb muscle activity through descending projections, which impedes mechanistic understanding of cortical movement control. Here we addressed this in mice performing an ethologically inspired all-limb climbing behavior. We quantified the direct influence of forelimb primary motor cortex (caudal forelimb area, CFA) on muscle activity comprehensively across the muscle activity states that occur during climbing. We found that CFA informs muscle activity pattern, mainly by selectively activating certain muscles while exerting much smaller, bidirectional effects on their antagonists. From Neuropixel recordings, we identified linear combinations (components) of motor cortical activity that covary with these effects, finding that these components differ from those that covary with muscle activity or kinematics. Collectively, our results reveal an instructive direct motor cortical influence on limb muscles that is selective within a motor behavior and reliant on a new type of neural activity subspace.

Introduction

Motor cortex and its descending projections have expanded in certain mammalian lineages, seemingly because of the fitness conferred by the motor performance they support, as the effects of motor cortical disturbance across species correlate with the degree of this expansion[1-5]. Without normal motor cortical output, certain types of movement cannot be executed[6-9], and many other types are slower, less agile, and less effective, especially when dexterity is challenged or movements must adapt during execution[7,10-15]. Yet it remains poorly resolved when and how motor cortical output directly influences muscle activity through its descending projections to mediate this influence. The involvement of motor cortex in motor learning[16-18] and movement preparation/initiation[19-21] complicates the interpretation of lesion and other inactivation results vis-à-vis movement execution, as disturbance to processes on which execution depends can impede execution itself. Moreover, recent results indicate that motor cortical influence on muscle activity at the shortest latencies (10-20 ms in mice) differs from its influence on even slightly longer timescales (~50 ms)[22]. The consequent ambiguity of direct motor cortical influence on muscles has stymied the development of more mechanistic models of descending motor control[23].

During tasks requiring motor cortex, existing results leave open several basic possibilities for the form this direct influence could take. First, motor cortex could drive the totality of limb muscle activity patterns, with substantial compensation provided by other motor system regions following motor cortical disturbance. This might result if the necessity for motor cortex to generate certain muscle activity patterns that cannot be achieved by the rest of the motor system[24] prompts motor cortex to shoulder all of the pattern-generating burden[25]. Second, motor cortex could participate in unison with the rest of the motor system in generating motor output, without playing a necessary role in determining its pattern. Here, loss of direct motor cortical influence on muscles would cause at most a non-specific, fractional attenuation of motor output. Third, motor cortex could selectively influence particular components of muscle activity, such that it informs ("instructs") ongoing muscle activity pattern and acts distinctly from the rest of the motor system. Loss of direct motor cortical influence would then cause changes in muscle activity that themselves vary as the state of muscle activity changes.

It also remains unclear whether, on balance, motor cortical output directly excites the activity of individual limb muscles, or has bidirectional effects that depend on context. Motor cortex is thought to drive online movement corrections and the adaptation of movements based on context[12,26-29]. Such a role could involve the activation and deactivation of individual muscles at different times to steer movement as context requires. If motor cortex instead was generating the totality of muscle activity patterns or participating in unison with the rest of the motor system in generating motor output, bidirectional control could be achieved by modulating the magnitude of descending motor cortical output. Consistent with this, brief focal stimulation in motor cortex excites the activity of muscle groups that act at particular limb joints (e.g. flexion or extension)[30-32]. However, this output may drive downstream circuits that alternately excite or inhibit the activity of individual muscles to achieve bidirectional control, especially if motor cortical output was instructing muscle activity pattern. Consistent with this, spiking in pyramidal tract neurons can be

both positively and negatively correlated with the activity of individual limb muscles[33-35] and substantial activation in motor cortex is observed in the absence of movement[36,37].

This ambiguity about when and how muscle activity is influenced by motor cortical output has in turn prevented characterization of the motor cortical activity patterns mediating this influence. Previous studies have focused on components of motor cortical activity that covary with limb muscle activity[38-40], or movement parameters like joint angles or reach direction (kinematics)[41-43]. Substantial variance in muscle activity or movement kinematics can be explained by these activity components, consistent with their functional relevance. However, if motor cortical output does not contribute to the totality of muscle activity patterns but instead selectively alters them, we might expect that the components of motor cortical output that drive muscle activity may not reflect muscle activity in total, but only some fraction of it. Moreover, motor cortical activity that covaries with muscle activity or kinematics in total may be a consequence of monitoring or predicting body state[44], perhaps to subserve aspects of motor control apart from directly driving muscle activation[45]. In line with this, muscle activity can be decoded from motor cortical activity during movements where this activity does not directly drive muscles[22]. Thus, the components of motor cortical activity responsible for its direct influence may differ from those to which functional significance has previously been imputed[41,46-50].

To address these basic questions regarding direct motor cortical influence on muscles, we developed a paradigm in which head-fixed mice perform a motor behavior that depends on motor cortex and at which mice are naturally adept: all-limb climbing. We combined electromyography (EMG), optogenetic silencing, and dimensionality reduction to comprehensively assess the direct influence of output from the caudal forelimb area (CFA), the mouse analog of forelimb primary motor cortex, on muscles during climbing. We found that CFA informs muscle activity pattern, acting mainly by selectively exciting activity in certain muscles while exerting much smaller and bidirectional effects on their antagonists. We then used large-scale multielectrode arrays (Neuropixels[51]) to record motor cortical firing patterns during climbing and identify components of these patterns that covary with influence on muscles. We found that such components partially overlap with those that covary with muscle activity, but are entirely different from those that covary with gross limb kinematics. Our results indicate that during a single motor behavior, motor cortex selectively influences muscles through a neural activity subspace distinct from those traditionally focused upon.

A naturalistic climbing paradigm

We first sought a behavioral paradigm that emphasizes performance of adaptive limb movements in response to unpredictable sensory information, as numerous previous studies indicate motor cortical involvement in such tasks[13,23,52]. Inspired by the natural movement repertoire of mice, we thus developed a paradigm in which head-fixed mice climb with all four limbs across a series of handholds that extend radially from a wheel, thereby rotating the wheel (Fig. 1a-c; Extended Data Fig. 1a; Supplementary Movie). After each handhold accessible to the right limbs rotates 180° past the mouse, a linear actuator embedded within the wheel moves the handhold to a new, randomly chosen mediolateral position, while the left handholds remain fixed (Fig. 1d; Extended Data Fig. 1b-e). This ensures that the sequence of right handholds the mouse climbs across is unpredictable (Fig. 1e), so sensory information must be used in real time to steer right limb movement.

In this paradigm, water-restricted mice climb intermittently in bouts throughout hour-long daily sessions, earning a volume of water reward when they stop climbing that depends on the distance of the previous bout. The likelihood of reward and volume dispensed are varied adaptively within sessions to encourage longer distance bouts and to maintain a relatively consistent reward rate across sessions. As mice climb, activity of four muscles in the right forelimb is measured with chronically implanted EMG electrodes[53,54] (Extended Data Fig. 1f), and limb orientation is charted using video-based tracking[55].

This climbing differs from motor behaviors in many other head-fixed paradigms where subjects learn new tasks and become increasingly skillful and stereotyped with repeated training[16,36,56]. First, we saw little indication of continued performance improvement across sessions. To assess this, we examined measures of bout length and climbing speed, as the reward scheme depends on them. If mice modified their climbing to improve reward attainment, it should show up in these measures. However, we found that after mice are acclimated to head fixation (2 sessions) and taught the pairing between climbing and reward (1-3 sessions), there was little change on average in the time spent climbing (Fig. 1f), the velocity of climbing (Fig. 1g), and the distance of climbing bouts (Fig. 1h). Thus, after beginning to climb for rewards, mice do not appear to progressively develop skill specific to our climbing paradigm.

Second, we found that the form of forelimb muscle activity does not show substantial progressive change across sessions. Much of the variance in matrices of muscle activity recorded during individual sessions (4 muscles by T timepoints) was captured by their first two principal components (PCs; 64-96%, median 76%, n = 9 mice). To assess the similarity of muscle activity across sessions, we computed the principal angles between 2D subspaces of muscle activity defined by the first two PCs for each session (2 PCs by T timepoints), comparing each of the first 20 daily sessions to the 20th session (Fig. 1i,j). The first principal angle was generally low, averaging less than 2°. Though adjacent sessions appeared more similar (see lower angles for session 19), there was little indication that increasingly distant sessions were increasingly more dissimilar, which would be expected for a progressive change in movement form.

Lastly, we found that the stereotypy in both muscle activity and limb kinematics did not show clear signs of increasing across sessions. We quantified stereotypy using sample entropy[57],

which measures regularity in time series. Sample entropy has previously been used to measure regularity in signals like ECG[58], EEG[59], and BOLD from fMRI[60]. Here we found no indication of increasing regularity in either muscle activity (Fig. 1k) or limb kinematics (Fig. 1l) across sessions.

Thus mice in our climbing paradigm do not appear to develop an increasingly stereotyped movement strategy. Instead, they perform a sophisticated locomotion that deviates from rhythmic pattern to negotiate unpredictable variation in the "terrain" presented by the handholds. The consistency of climbing form and performance across sessions, together with the natural climbing ability mice possess, obviate the need for extensive training. In fact, some mice exhibit long, continuous bouts of climbing during the first session during which they are taught the climbing-reward pairing. Our paradigm enables motor system interrogation during kinematically-diverse, nonstereotyped, and ethologically-relevant motor behavior in mice.

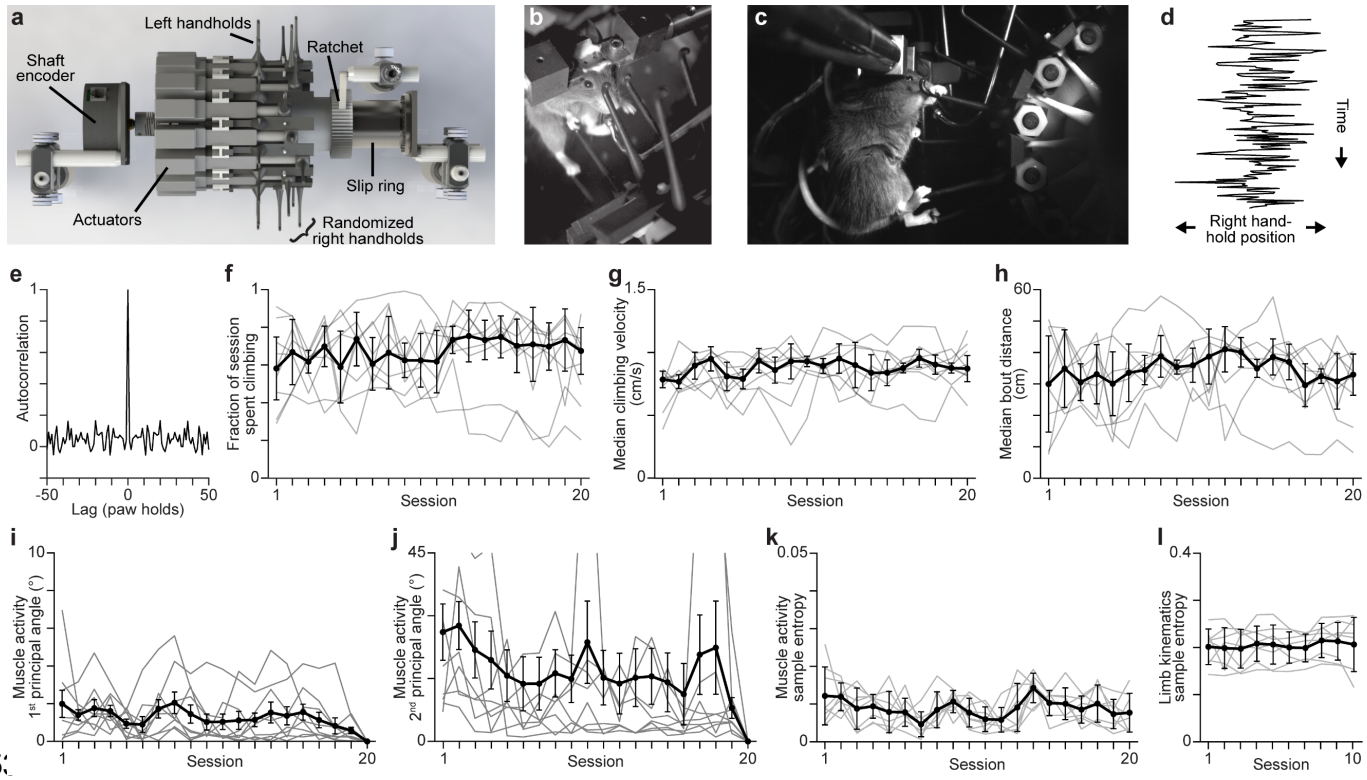

Figure 1 | Head-fixed climbing paradigm. a, Bird's-eye view of wheel apparatus for climbing. A shaft encoder measures the wheel's angular position. Actuators randomize the position of each right handhold when they reach a point 180° away from the mouse. A ratchet ensures the wheel only rotates in one direction. A slip ring commutes voltage signals to and from the actuators. **b**,**c**, A head-fixed mouse climbing in the paradigm. **d**, Example sequence of right handhold positions over time, illustrating randomization. **e**, Autocorrelation of right handhold positions. **f-h**, 1st, 2nd, and 3rd quartiles (n = 9 mice) for the fraction of time spent climbing (**f**), median climbing velocity (**g**), and median climbing bout distance (**h**) across sessions. Gray lines in **f-l** are for individual animals. Session 1 indicates the first session after mice had learned the pairing between climbing and reward, when reward dispensation switched from experimenter- to computer-controlled. **i**,**j**, 1st, 2nd, and 3rd quartiles for the 1st (**i**) and 2nd (**j**) principal angles for EMG time series collected during each of the first 20 climbing sessions versus the 20th climbing session. **k**,**l**, 1st, 2nd, and 3rd quartiles for the sample entropy of muscle activity (**k**) and limb kinematics (**l**) time series during each of the first 20 climbing sessions. For each session, we took the mean across sample entropy values for each muscle, or for the X and Y positions of each tracked limb point.

Quantifying direct motor cortical influence during climbing

We next sought to quantify direct motor cortical influence on contralateral forelimb muscles across the range of muscle activity states that occur during climbing. While mice (n = 8) were actively climbing, we sporadically and briefly inactivated the left caudal forelimb area (CFA, forelimb M1+S1) at random. We used transgenic mice that express channelrhodopsin2 in all cortical inhibitory interneurons, applying occasional 25 ms blue light pulses that covered the surface of CFA (10 mW/mm^2; Fig. 2a). Previous measurements indicate that this yields a ~50% activity reduction in deep cortical layers within 7 ms, which rises to 90-95% in < 20 ms[22,61]. Light pulses were always > 4 seconds apart to allow recovery of neural activity between events; on average, ~100-200 trials were collected during each daily session (11-37 sessions per animal). Twice as many equivalent events without blue light were notated in recordings to serve as control trials and enable comparisons between separate sets of control trials to aid statistical testing. Random trial timing ensured broad coverage of muscle activity states each mouse expressed during climbing.

Comparison between control and inactivation trial-averaged muscle activity indicated that CFA directly influences muscle activity during climbing. Inactivation and control trial averages diverged ~10 ms after light onset, which reflects the shortest latency at which CFA output influences muscles[22] (Fig. 2a,b, Extended Data Fig. 2a). We found that the mean across muscles of the absolute differences between control and inactivation averages was significantly greater than 0 at 25, 50 and 100 ms following light onset (Extended Data Fig. 2b; $p = 0.004$ in each case, Wilcoxon signed-rank test). Such direct influence has previously been observed in mice during a trained forelimb reaching task but is not seen during treadmill locomotion[22].

Our next goal was to assess CFA influence at different muscle activity states during climbing. To achieve adequate statistical power in these assessments despite the broad range of such states, we sought a means for collecting together inactivation and control trials that began at similar muscle activity states, and organizing them in a manner conducive to subsequent analysis. We thus used nonlinear dimensionality reduction (UMAP[62]) to generate a 2D map of muscle activity states in each mouse that well-preserves the proximity of similar states. By quantifying inactivation effects for trials beginning from states within local neighborhoods on these maps, we could then effectively measure CFA influence at different muscle activity states.

To represent muscle activity states robustly despite EMG signal variability and in a manner relevant to how CFA inactivation would affect muscles, we chose to use states spanning 50 ms epochs rather than individual time points. To ensure proximity on maps among states visited close together in time, we defined epochs that overlapped in time. Muscle activity traces surrounding each trial, together with their corresponding first derivatives, were subsampled in 5 ms increments and divided into overlapping 50 ms epochs that began every 10 ms (Fig. 2c). For each 50 ms epoch, the muscle activity and first derivative trace segments were concatenated into a single vector (8 segments x 10 time bins = 80 elements), and UMAP was applied to map vectors into two dimensions (Fig. 2d). On the resulting maps, embedded state vectors (points) from successive epochs form trajectories that reflect the sequence of states surrounding control and inactivation trials.

We then sought to quantify CFA influence on muscles across these muscle activity state maps by computing inactivation effects averaged over trials that began at states within local neighborhoods on the maps. To do this, we first defined a grid over each map (Fig. 2e). Separately for each muscle and each grid point, we computed the trial-averaged muscle activity for each muscle from -10 to +30 ms from trial/light onset, separately for inactivation and control trials. For these averages, we used all trials, but we weighted each by a Gaussian function of the Euclidean distance between the given grid point and the point for the epoch just prior to when an inactivation effect could begin on the given trial (-40 to +10 ms from trial/light onset, "weight epoch;" Fig. 2c,e). Trial weights are hence not influenced by inactivation effects. As a consequence, weight epoch states from control and inactivation trials are similarly distributed across maps (Fig. 2e). The standard deviation of the Gaussian function was chosen to be roughly 10% of the width of the map. This reflects a trade-off between differentiating effects across distinct map regions and the statistical power gained from combining trials. In effect, trial averages are heavily weighted toward trials beginning at states close to the given grid point. Only grid points close to a substantial number of weight epoch states were subsequently considered ("valid grid points;" see Methods).

Separately for each muscle, we measured the size of the inactivation effect at each valid grid point as the difference between the rate of change in inactivation and control trial averages from 0 to 20 ms after trial onset (Fig. 2f). This difference is proportional to the distance from the identity line of a point plotting the inactivation and control rates of change against one another (Fig. 2g). By focusing on the immediate effects of inactivation, our quantification captures the direct effect of motor cortical output on muscles and avoids longer latency effects that may be less specific. By using the original trials to measure effect sizes and not trial segments after embedding, these effect sizes are not themselves dimensionally reduced.

We then plot the resulting effect sizes at each valid grid point across the map, producing an "inactivation effect map" for each muscle recorded in each mouse (Fig. 2i; Extended Data Fig. 3a-c). Maps for different muscles in a given mouse show wide variation in the magnitude and sign of inactivation effects across grid points (Fig. 2j). For each grid point, we resampled from control trials to compute distributions of effect sizes expected by chance. We used these empirical null distributions to compute p-values for each grid point's effect size and assessed significance using a threshold that controls the rate of false discovery at 5% (Fig. 2k; Extended Data Fig. 3a-c).

UMAP does not guarantee that vectors embedded close to one another are necessarily similar; maps could for example contain rifts across which vectors are much more dissimilar than in other local map regions. We thus examined the similarity between nearby state vectors embedded in local neighborhoods across maps. For each grid point, we computed the average Euclidean distance (in the full 80D space) between all pairs of embedded vectors, weighting each by the same Gaussian function of distance from the given grid point as above. The distance used for weighting here was the mean for the two embedded vectors in the given pair. The only prominent variation we observed in the state vector similarity was a gradual increase in pairwise distance from the map center to the edges, indicating that vectors were locally more dissimilar around the edges (Fig. 2h; Extended Data Fig. 3d). Thus, although there is some variation in the

local similarity of states across maps, these changes are gradual at the resolution we use here for measuring inactivation effects.

There are a number of important caveats in evaluating inactivation effect maps. Collapsing 80D muscle activity state vectors onto a 2D map eliminates substantial information about those states. The smoothing we have applied in computing trial averages will also mask fine structure. UMAP organizes states according to similarity and so ignores other features potentially relevant to cortical influence. As a consequence, these maps do not reflect all the structure that may exist in the relation between muscle activity state and CFA influence.

However, our inactivation effect maps retain sufficient structure for the present purposes we introduce above. These maps reflect effects of inactivations beginning from distinct muscle activity states, and they reveal effects that vary across states and are statistically significant. We also found that map structure was not strongly dependent on the choice of key parameters: 20% increases and decreases in either Gaussian function standard deviation or the epoch duration did not substantially change results (Extended Data Figure 3e-h).

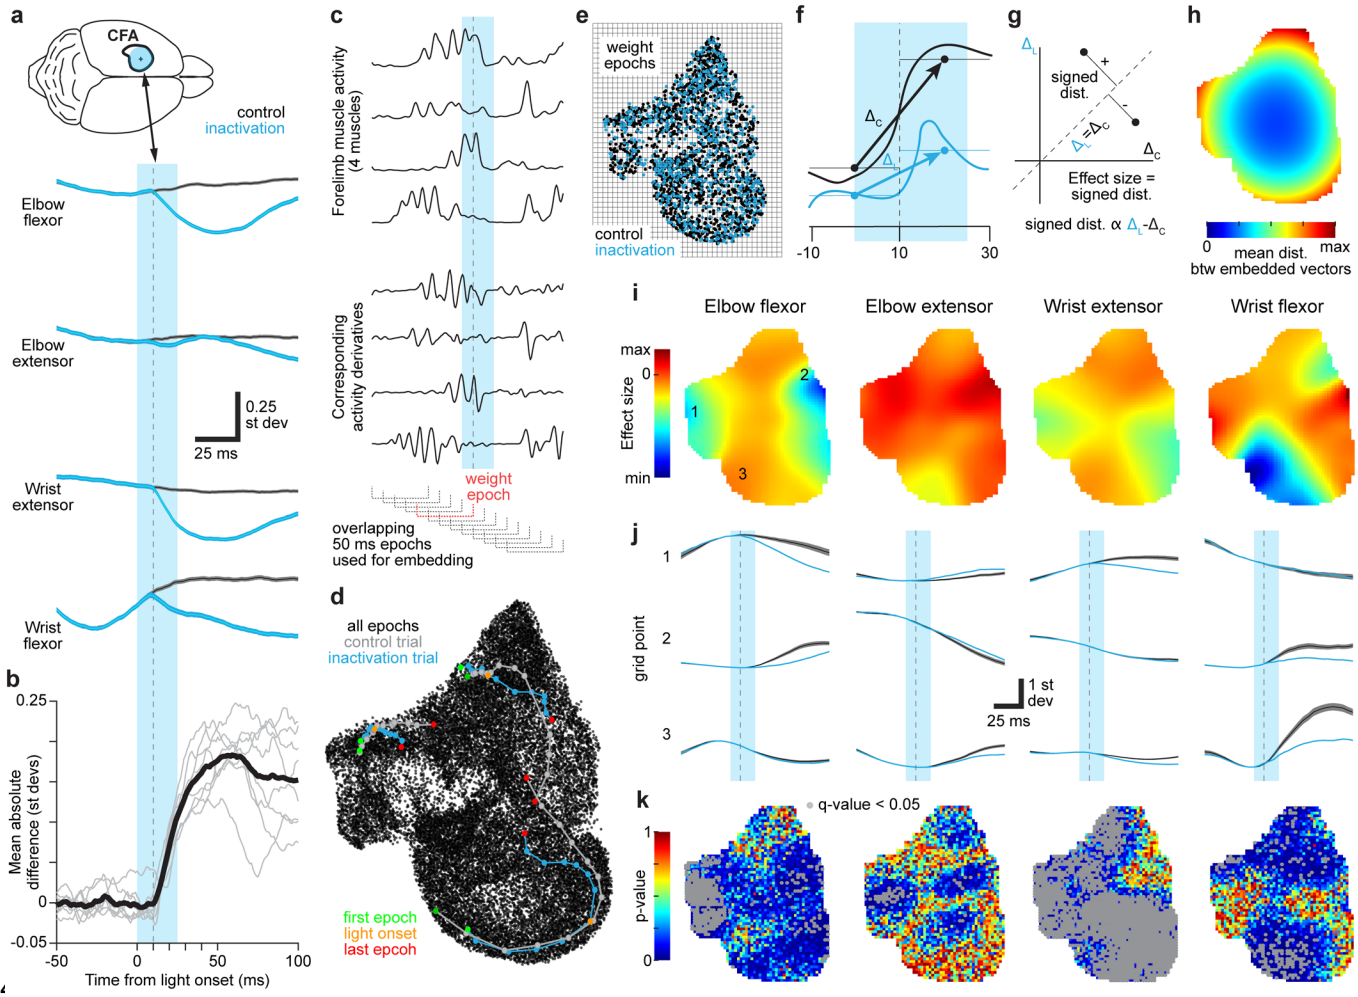

Figure 2 | Comprehensive assessment of CFA influence across muscle activity states. a, Control (gray, 18,397 trials) and inactivation (cyan, n = 9,029 trials) trial averages (mean ± sem) for four muscles across 8 mice. Vertical cyan bars in **a-c,f,j** indicate the 25 ms epoch of blue light applied to CFA and gray dotted lines are 10 ms after light onset. **b**, Mean absolute difference between inactivation and control trial averages across all four muscles. Light gray lines are individual animals, and the solid black line is the mean across animals. For baseline subtraction, control trials were resampled to estimate the baseline difference expected by chance. **c**, Example of muscle activity and its corresponding first derivative surrounding trials that were used for creating muscle activity state maps. The weight epoch immediately precedes when effects begin. **d**, Example muscle activity state map from one animal. Larger, connected dots show examples of states for sequential overlapping epochs from individual trials. Pairs were chosen based on their similarity during the weight epoch. **e**, Grid overlaying a map including only points from the weight epochs used for weighting trials in grid point trial averages. **f**, Schematic of the calculation of inactivation effect at each grid point from the control (black) and inactivation (cyan) trial-averaged muscle activity. Δ_C and Δ_L reflect the slopes of lines connecting the average activity just before to just after the inactivation effect begins. **g**, Schematic illustration of the effect size on a plot of Δ_L versus Δ_C. **h**, Map in which each grid point is colored by the mean distance, in the full 80D space,

between all pairs of embedded state vectors, with each individual distance weighted by a Gaussian function of the pair's mean distance from the grid point on the 2D map. **i**, Inactivation effect maps for the four recorded muscles. Panels **i**-**k** show representative results from one mouse. **j**, Grid point-averaged muscle activity from control (gray, mean ± sem) and inactivation (cyan, mean) trials, for three example grid points from the maps in **i**. **k**, Maps of p-values computed for inactivation effects at each grid point. Q-values (gray overlay) reflect the expected rate of false discovery below the corresponding p-value[63].

CFA acts primarily by selectively exciting physiological flexors

We then analyzed inactivation effect maps to characterize direct CFA influence on contralateral forelimb muscles and test existing ideas about motor cortical influence. We first confirmed that maps reproducibly show significant influence for at least some muscle activity states. We generated histograms of p-values computed for all valid grid points for each muscle. These histograms consistently showed a skew toward zero, reflective of a substantial fraction of grid points where the null hypothesis of no effect was false (Fig. 3a,b). From these distributions, we estimated the fraction of grid points showing effects (Fig. 3c). The mean estimated fractions were 0.62, 0.22, 0.73, and 0.37 for elbow flexor, elbow extensor, wrist extensor, and wrist flexor muscles, respectively. These estimates were significantly above zero for all four muscles ($p = 0.008$ or $p = 0.016$, Wilcoxon signed-rank test). Control maps generated from comparisons between separate sets of control trials showed uniform distributions, as expected under the null hypothesis (Extended Data Fig. 4a). These results show that direct CFA influence on muscles is specific to a subset of muscle activity states.

We next established that CFA influence varies in magnitude across muscle activity states. If this were true, then the 2D autocorrelation of inactivation effect maps should be significantly above what is expected by chance. We computed the 2D autocorrelation of both inactivation effect maps and control maps, observing substantially heightened autocorrelation in the former (Fig. 3d). To assess whether these differences were significant, we computed the mean difference between inactivation effect map autocorrelation and that of control maps, averaged over spatial lags up to 20 grid points (Fig. 3e). These differences were significantly above zero for all four muscles ($p = 0.008$ for each muscle, Wilcoxon signed-rank test).

We then demonstrated that the magnitude of CFA influence on different muscles was significantly different. Here we compared inactivation effect sizes across muscles for a select subset (~15-20) of grid points that were all at least three standard deviations of the aforementioned Gaussian weight function apart, so that effect sizes relied effectively on separate trials. A Kruskal-Wallis test found that effect sizes were indeed significantly different across muscles in six out of eight mice ($p < 0.004$, $p = 0.19$ and 0.83 in two other mice).

We next established that CFA influence magnitude across muscle activity states was not simply proportional to the size of muscle activity. This would be expected if CFA was providing a constant fraction of the drive to muscles during climbing, consistent with observations that the loss of motor cortical output leads to hypometric limb movements in some settings[12,29,64-66]. Plots of effect size versus muscle activity at trial onset were poorly captured by linear functions (Fig. 3f,g; Extended Data Fig. 4b,c). Residuals from linear fits were significantly different from uniform (K-S test, $p < 10^{-10}$ in all cases). The same held when fitting only the effect sizes that were themselves significantly above zero (false discovery rate < 0.05; $p < 0.004$ in all cases; Extended Data Fig. 4d,e).

Given these differences, we then examined how direct CFA influence varied across muscles. A number of previous observations indirectly suggest that primary motor cortex may preferentially control certain muscle groups more so than their antagonists[67-70]. We found that CFA

output preferentially influences elbow flexors and wrist extensors, which can be grouped together as *physiological flexors* because of their coactivation during locomotion and during the flexion reflex[71]. The distributions of effect sizes across all grid points and mice showed larger deviations compared to those of control maps for the elbow flexor and wrist extensor (Fig. 3h). The estimated fraction of grid points showing effects (false null hypotheses, Fig. 3c) was significantly greater for the elbow flexor ($p = 0.007$, Wilcoxon rank-sum test) and wrist extensor ($p = 0.015$), compared with their respective antagonists. The mean estimated fractions across mice were 61% higher for the elbow flexor and 43% higher for the wrist extensor, compared with their respective antagonists. Thus, CFA output directly influences physiological flexor muscles preferentially.

Finally, we found evidence that this preferential influence on physiological flexors is exclusively excitatory, in contrast to effects on their respective antagonists. The effect sizes significantly different from zero for the elbow flexor and the wrist extensor were exclusively negative, whereas the same was not true for their antagonists (Fig. 3i,j; Extended Data Fig. 4f). While the elbow extensor exhibited CFA influence fairly symmetric around zero, the wrist flexor did show a bias toward excitatory influence. We also observed that inactivation effect maps for the physiological flexors were more highly correlated compared to those for all other pairs of muscles (Fig. 3k), suggesting a greater degree of coordinated control of these muscles. Collectively, these results indicate that direct CFA influence instructs muscle activity pattern, exciting physiological flexor muscles to varying degrees and only at some muscle activity states (i.e., influence is selective), while having less frequent, weaker, and bidirectional effects on their respective antagonists.

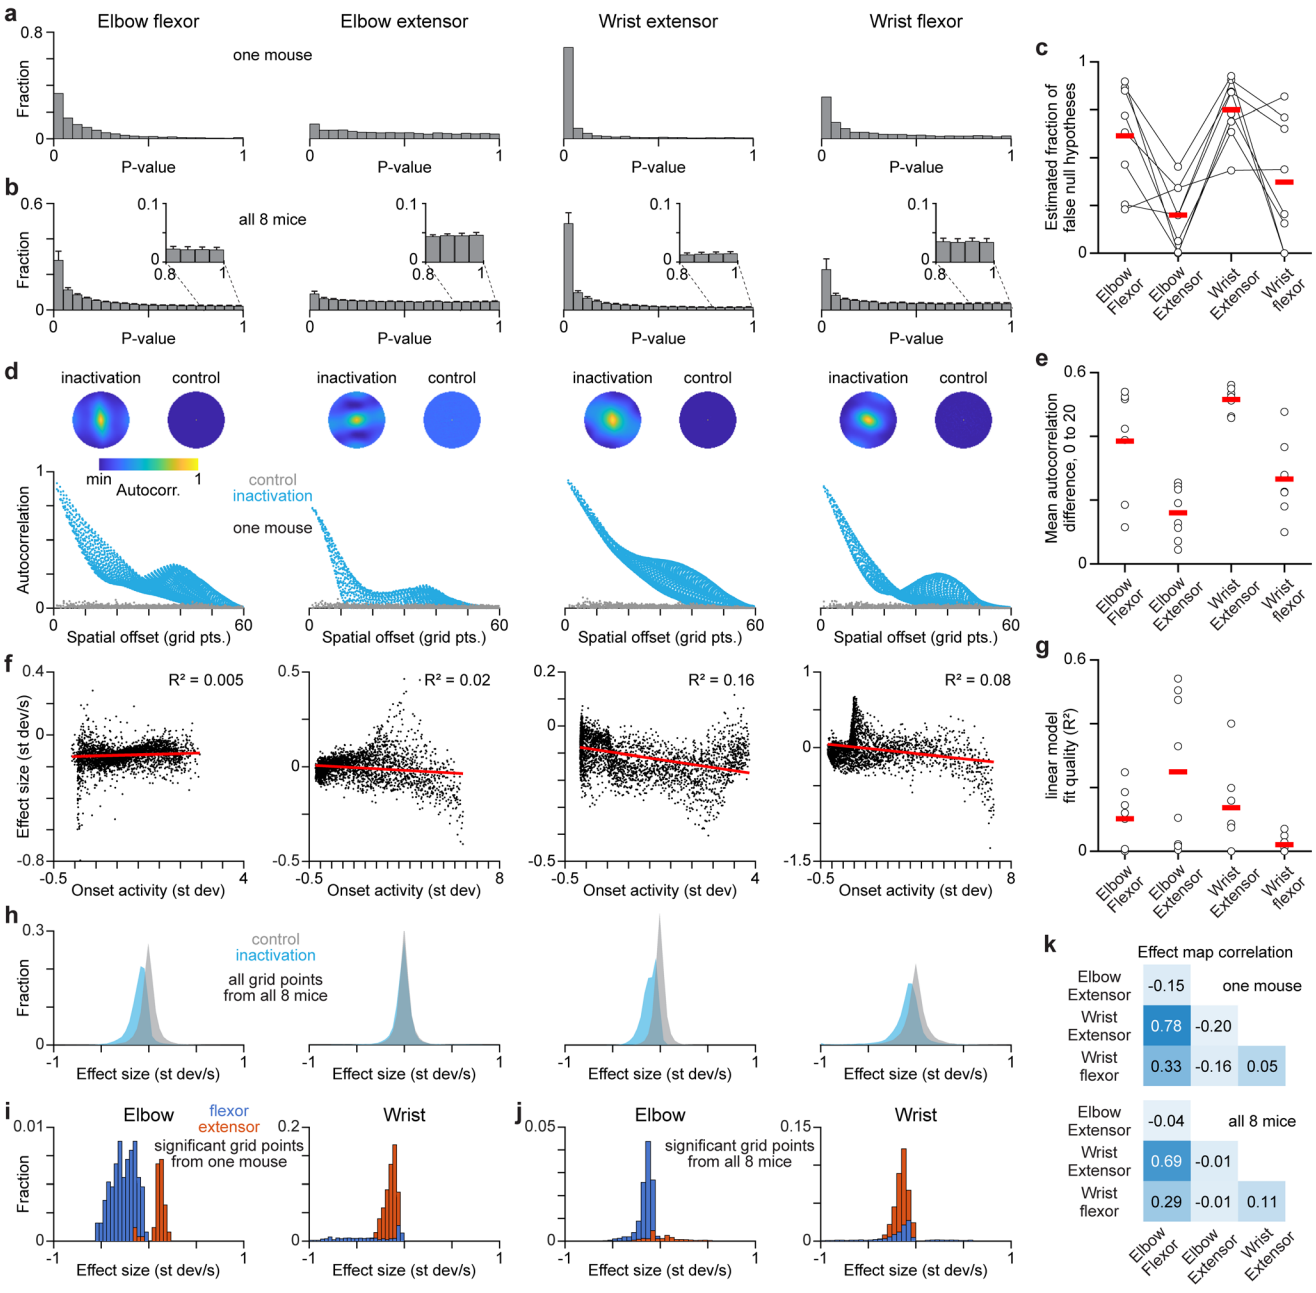

Figure 3 | CFA selectively excites physiological flexors. a,b, P-value distributions for inactivation effects on each muscle for all grid points in one mouse (**a**) and across all 8 mice (**b**). **c**, Estimated fraction of grid points at which the null hypothesis of no effect is false, calculated from p-value distributions, for individual animals (black circles) and the mean across animals (red bars). **d**, For one mouse, the 2D autocorrelation for inactivation effect maps and control maps generated with only control trials (top) and scatterplots of correlation values versus their spatial offset (lag from zero offset). **e**, Difference between inactivation effect maps and control maps in their mean autocorrelation over spatial offsets from 0 to 20 grid points for individual mice (black circles), and the mean across animals (red bars). **f**, Scatterplots of inactivation effect size versus muscle activity at trial onset for one animal. Each point reflects a different grid point. R^2 is for a

linear fit (red). **g**, R^2 for linear fits to scatterplots of inactivation effect size versus muscle activity at trial onset for individual mice (black circles), and the mean across animals (red bars). **h**, Effect size distributions for all grid points across all 8 mice, separately for inactivation effect maps and control maps. **i**,**j** Effect size distributions for all significant grid points from one mouse (**i**) and all 8 mice (**j**). **k**, 2D correlation between inactivation effect maps for different muscles for one mouse (top) and the means across all 8 mice (bottom).

Weak covariation between CFA influence and gross kinematic state of the forelimb

A number of previous observations suggest that motor cortex may, to some extent, control the limb via commands that dictate its kinematics rather than muscle activation[72-74], although this remains controversial[47,75]. If CFA were dictating contralateral forelimb kinematics, we reasoned that direct CFA influence on muscles should covary with the orientation of the contralateral forelimb. We therefore assessed how CFA influence correlated with the orientation of the contralateral forelimb.

Here we mimicked the approach we took to assess how CFA influence covaries with muscle activity state. We computed new state maps with UMAP using not muscle activity state vectors, but vectors composed of the horizontal and vertical positions of sites on the right forelimb tracked from video (Fig. 4a,b). Nearby points on these maps thus reflect 50 ms epochs of limb kinematics that are similar. We note that because video was captured at 100 Hz, time series segments used here had 10 ms spacing between points instead of 5 ms. The resulting 2D maps separated states that correspond to different limb orientations into different map regions (Fig. 4c), with the cyclic changes in limb orientation during iterative climbing ordered around the map.

Using these maps, we quantified the effects of CFA inactivation on each forelimb muscle at grid points covering the maps as above. We again computed p-values for effects at each grid point using comparisons between resampled control trials to generate empirical null distributions. Histograms of the inactivation effect sizes across all grid points did show deviations from controls (Fig. 4d). However, p-value distributions for effects on each muscle showed very limited skew toward zero, indicating discernible effects on only a small fraction of grid points (Fig. 4e). Thus, CFA influence on muscles does not covary with forelimb orientation nearly as well as it does with muscle activity state.

If CFA output dictates forelimb kinematics, we might expect that CFA inactivation would perturb limb kinematics themselves. We therefore quantified the effect of inactivation on the position of the sites tracked at the shoulder, elbow, wrist, and finger. Here again inactivation effects were computed as above, but using trial-averaged site position in place of trial-averaged muscle activity. Again, histograms of inactivation effect sizes did show deviation from controls (Fig. 4f). However, p-value distributions for effects on each site position showed very limited skew toward zero, indicating discernible effects on only a small fraction of grid points (Fig. 4g).

Given an expectation that muscle activity and limb kinematics should covary, these results may seem surprising. However, we note that during adaptive, nonstereotyped motor behaviors, like the climbing mice perform in our paradigm, linear covariation between muscle activity and limb orientation will not necessarily be strong, given the complex relationship between muscle activity and limb kinematics. To assess how well coupled muscle activity and limb kinematics were in our paradigm, we made state maps using either data type, here using the same time point spacing (10 ms) and equivalent amounts of data for both map types. We found that the proximity of muscle activity states on maps only weakly predicted the proximity of limb orientation states, and vice versa (Fig. 4h-j). For example, subsets of states close together on one map type corresponded in time to states that were widely distributed on the other map type. This decoupling between muscle activity and gross limb kinematics helped reveal here that CFA influence is not

well-organized by limb orientation. This finding suggests that CFA output directly specifies muscle activity and not limb orientation.

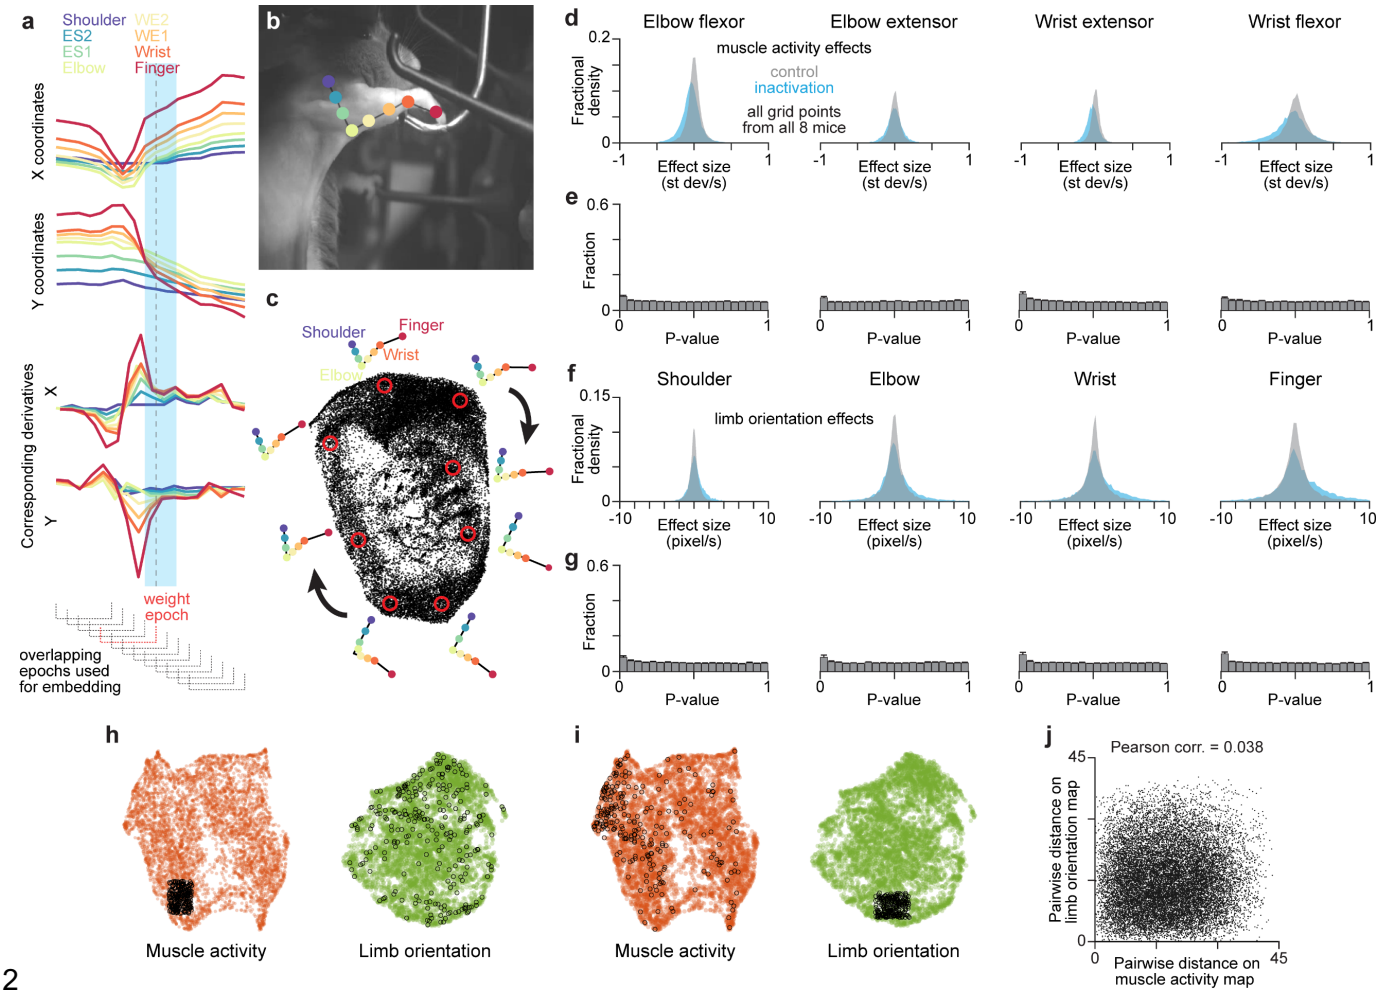

Figure 4 | Gross forelimb kinematics capture CFA influence worse than muscle activity. a, Time series of tracked forelimb sites and their corresponding first derivatives surrounding inactivation and control trials are windowed into overlapping segments. Segments are used to create a 2D embedding via UMAP. **b**, Image showing the locations of the eight sites tracked on the forelimb, according to the color code in **a** (ES1,2: 1st and 2nd sites between elbow and shoulder joints, WE1,2: 1st and 2nd sites between wrist and elbow joints). **c**, Example map of forelimb orientation states from one mouse, along with the trial-averaged positions of the forelimb sites at selected grid points within the map (red circles). **d-g**, Distributions of the sizes of inactivation effects on muscles (**d**), p-value distributions for inactivation effects on muscles (**e**), distributions of the sizes of inactivation effects on four main forelimb sites (**f**), and p-value distributions for inactivation effects on four main forelimb sites (**g**), calculated using forelimb orientation maps, across all grid points and all 8 mice. **h**,**i**, Maps of muscle activity states (left, orange) and limb orientation states (right, green) from one animal. Black circles in each panel mark corresponding sets of embedded vectors on the two maps (i.e., those coming from the same set of epochs). **j**, For the embedded vectors from 200 randomly selected epochs, the Euclidean distance between all possible pairs of those vectors on the muscle activity map in **h**,**i** plotted against the distance between their corresponding vectors on the limb orientation map.

CFA firing patterns during climbing

In order to determine how CFA firing patterns covary with CFA influence on muscles, we used Neuropixels to measure the firing of CFA neurons across cortical layers in a subset of animals for which inactivation effect maps were computed (n = 3; Fig. 5a). After completing collection of inactivation responses for mapping influence in these mice, we recorded neural activity in CFA with acutely-implanted Neuropixels during the next 3-4 daily behavioral sessions. Template-based spike sorting yielded 366-684 single units meeting our selection criteria per animal. Waveform-based grouping yielded 279-495 wide-waveform, putative pyramidal neurons and 87-189 narrow-waveform, putative inhibitory interneurons per animal, in line with previously observed proportions[23,61] (Fig. 5b). For the populations recorded in each animal, activity along the top principal components was modulated around climbing onset (Fig. 5c). The majority of both wide- and narrow-waveform neurons had higher firing rates during forelimb muscle activity as compared with periods when all recorded muscles were quiescent (Fig. 5d). Over 80% of recorded neurons had firing rate time series that were significantly correlated with the activity of at least one muscle. The same held for subsets of neurons with waveform centroids localized to different cortical layers (Fig. 5e).

To enable alignment with CFA influence, we measured variation in the firing of neurons across the muscle activity state maps used for quantifying inactivation effects (Fig. 5f). Segments of muscle activity (50 ms) collected as mice climbed during Neuropixel recordings, together with their first derivatives, were again assembled into vectors reflecting muscle activity state. These vectors were embedded in state maps using the transformation previously computed with UMAP from inactivation and control trials for the given animal. This positioned vectors on the same maps used for quantifying inactivation effects. For each recorded neuron, average firing rates during the 50 ms epochs reflected in muscle activity state vectors were used to compute an average firing rate at each grid point. The firing rate in each epoch was weighted in these averages by the corresponding muscle activity state vector's distance from the given grid point, using the same Gaussian function as above. Registering a neuron's firing to inactivation effect maps in this manner ensures that inactivation effects are matched with estimates of the neuron's average activity immediately prior to the onset of those effects.

The resulting maps of average firing rate across muscle activity states (neural activity maps) showed a wide array of muscle state-dependent firing patterns. In particular, we found many neurons with firing that was somewhat sparse across maps – that is, firing was heavily concentrated in subregions of the maps, and mostly or completely absent in others (Fig. 5g). To quantify this sparsity, we used an index that was originally developed to measure the place-dependent firing of hippocampal neurons[76]. Ordering neurons by this sparsity index revealed that even the neurons with a median level of sparsity had firing that was strongly biased towards a specific map subregion (Fig. 5h; Extended Data Fig. 5). Similar broad distributions of firing sparsity were seen in all three mice (Fig. 5i). Sparsity was only weakly dependent on mean firing rate during climbing (Fig. 5j). Neural activity maps did not vary substantially across cortical layers, as the average 2D correlation between maps for neurons assigned to different layers was similar to that expected under a null

470 hypothesis of no variation across layers (Fig. 5k). Thus, a substantial fraction of CFA neurons
471 across cortical layers each fire primarily at a very limited range of muscle activity states during
472 climbing. This observation was made possible by including the full range of muscle activity states
473 occurring during climbing in our analysis.
474 The finding that CFA neuron firing is highly structured across muscle activity state maps
475 illustrates their sufficiency for our present purposes laid out above. Although we can in no way
476 claim that these maps organize states in a manner that best reflects cortical control, results here
477 indicate that CFA neuron firing encodes substantial information about muscle activity states
478 organized by their similarity. We also note that the existence of neurons whose activity is strongly
479 biased toward a small fraction of muscle activity states is not consistent with the idea that cortical
480 control relies entirely on a low-dimensional neural activity subspace, and adds to recent
481 demonstrations of reproducible high-dimensional activity in primary motor cortex[77].
482
483

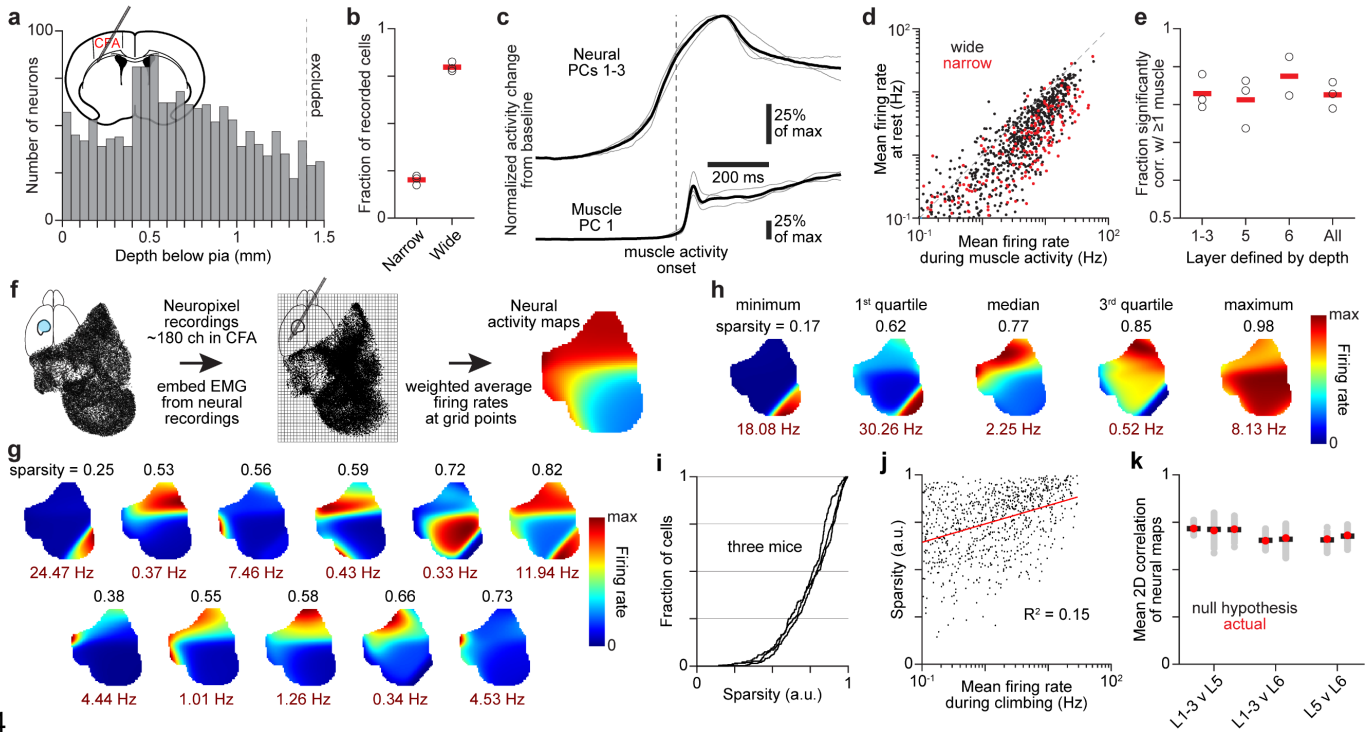

Figure 5 | CFA firing patterns vary in their sparsity across muscle activity states. a, Histogram of the depth below pia of recorded neurons, as assigned by Kilosort, across all three mice. Inset is a schematic coronal section showing placement of the Neuropixel in CFA. **b**, Fractions of neurons assigned to narrow- and wide-waveform subsets in each of three mice (black circles), and the mean across animals (red bars). **c**, Normalized absolute activity change from baseline summed across the top three principal components (PCs) for recorded CFA neurons, and the top PC for muscle activity, for individual mice (thin gray) and the mean across mice (black). **d**, Scatterplot of the mean firing rate for neurons during periods of no muscle activity (rest), and periods of muscle activity, across all three mice. **e**, Fraction of neurons in different layers whose firing rate time series were significantly correlated with the activity time series for at least one forelimb muscle. One mouse had very few neurons assigned to layer 6, preventing a reliable calculation in that case. **f**, Schematic of the calculation of an activity map for each recorded neuron that is registered to inactivation effect maps. **g**, Example neural activity maps exhibiting somewhat sparse firing (black values) for 11 neurons from one mouse, along with their maximum firing rates (red). **h**, Neural maps whose sparsity values reflect quartiles of the sparsity distribution for the given mouse. **i,** Cumulative histograms of sparsity values for all neurons in each of three mice. **j**, Scatterplot of sparsity versus mean firing rate during climbing for neurons in all three mice. R^2 is for a linear fit (red). **k**, Mean 2D correlation of activity maps for neurons assigned to different layers (red), compared to a null distribution computed by repeating the calculation 100 times after randomly permuting neuronal layer labels (gray), and the means thereof (black bars).

506 **CFA activity components that align with inactivation effects**
507 We then used neural activity maps to identify components of CFA firing that align across muscle
508 activity states with CFA influence on muscles. We did so by combining singular value
509 decomposition (SVD) and canonical correlation analysis (CCA) to align neural activity maps with
510 the inactivation effect maps computed for the same animals[78-80]. The neural activity map for each
511 wide-waveform neuron with overall mean firing rate above 0.1 Hz was converted to a vector by
512 concatenating its columns. For each mouse, these vectors were combined into a grid points by
513 neurons matrix, which was then replaced with a dimensionally-reduced grid points by 20 neural
514 singular vectors matrix computed with SVD (Fig. 6a,b). These dimensionally-reduced matrices
515 were then aligned through CCA with grid points by muscles matrices formed similarly by
516 concatenating the columns of inactivation effect maps for the four recorded muscles (Fig. 6a-d).
517 The resulting canonical variables reflect components of CFA firing patterns that maximally
518 correlate with components of inactivation effects but are mutually uncorrelated with one another.
519 For all three animals, neural and inactivation effect canonical variables were highly correlated, and
520 the inactivation effect variables captured substantial fractions of inactivation effect variance (Fig.
521 6e,f; mean correlation = 0.99, 0.97, 0.92 and 0.86 for canonical variables 1 through 4; mean effect
522 variance capture = 0.29, 0.26, 0.18 and 0.11). Plots of the cumulative variance captured across
523 orthonormalized canonical vectors indicated that each inactivation effect variable captured a
524 substantial amount of additional inactivation effect variance (Fig. 6g). Although CCA attempts to
525 maximize correlation between canonical variables, it is not fated that each canonical variable will
526 account for a great deal of the variance in effect maps, as they do here. Repeating CCA using the
527 inactivation effect map of each individual muscle found a CFA activity component highly
528 correlated with the inactivation effects for the given muscle in all cases (median correlation = 0.97,
529 range = 0.89 to 0.99, 12 total muscles in 3 animals). This indicates that CFA activity aligns well
530 with the effects on each individual muscle.
531 Here we have quantified alignment using in-sample fit (i.e., the fit to the same data) rather
532 than out-of-sample generalization. Thus to validate our results, we repeated CCA calculations after
533 randomly permuting the map locations of muscle activity states from neural recordings. The
534 resulting correlation between canonical variables was significantly lower than for the original data
535 (for canonical variables 1 through 4, mean control correlation = 0.74, 0.72, 0.69, 0.65; $p = 0.001$
536 in all 4 cases). We also repeated CCA calculations 300 times using separate, randomly-chosen
537 halves of trials for computing the inactivation effect maps. The canonical variables identified with
538 each half of trials were highly similar (Extended Data Fig. 6a), as were the 4D neural activity
539 subspaces spanned by the neural canonical vectors (Extended Data Fig. 6b). These results appear
540 to reflect the fact that maps are generally similar for separate sets of trials, which we illustrate in
541 the next section. In addition, we found that CCA results were not highly dependent on the number
542 of singular vectors used for neural dimensionality reduction (Extended Data Fig. 6c-f), nor on the
543 choice of key inactivation map parameters (Extended Data Fig. 6g). Thus, neural canonical vectors
544 span a neural activity subspace where activity aligns closely, and nontrivially, with inactivation
545 effects. Below we refer to this as the "influence subspace."

We examined the contribution of individual neurons to these influence subspaces, probing for evidence that they were localized to subsets of neurons. We used the weights of each neuron in the singular vectors, and the weights of these singular vectors in the neural canonical vectors, to compute the relative contribution of each neuron to each canonical vector. Here we focused on canonical vectors defined using the inactivation effect maps of individual muscles. We found a broad distribution of neuronal weights for each muscle in each animal (Extended Data Fig. 7a). To summarize across muscles, we computed the norm of a vector composed of the four resulting weights for each neuron, finding a similarly broad distribution of weight vector norms in each animal (Extended Data Fig. 7b). We tested these distributions for multimodality (i.e. multiple components), which would indicate distinct neuronal subsets, but results were not significant in all cases (p-values ranged from 0.34 to 0.99, Hartigan's dip test). We found similar results after scaling weights and weight vector norms by each given neuron's mean firing rate (Extended Data Fig. 7c,d; p ranged from 0.90 to 0.99).

Because descending commands from motor cortex emanate from layer 5, we checked to see if contribution to influence subspaces varied across layers. Though the distributions of weight vector norms and norms scaled by mean firing rate were substantially overlapping for neurons localized to different layers (Extended Data Fig. 7e,f), norms both with and without scaling were significantly higher for neurons localized to layers 5 (unscaled $p = 3 \times 10^{-5}$, scaled $p = 5 \times 10^{-14}$, Wilcoxon's rank sum test) and 6 (unscaled $p = 7 \times 10^{-4}$, scaled $p = 3 \times 10^{-7}$) compared to those localized to superficial layers. Thus, activity components that align with inactivation effects involve substantial contributions from many CFA neurons, but contributions from deep layer neurons are significantly larger.

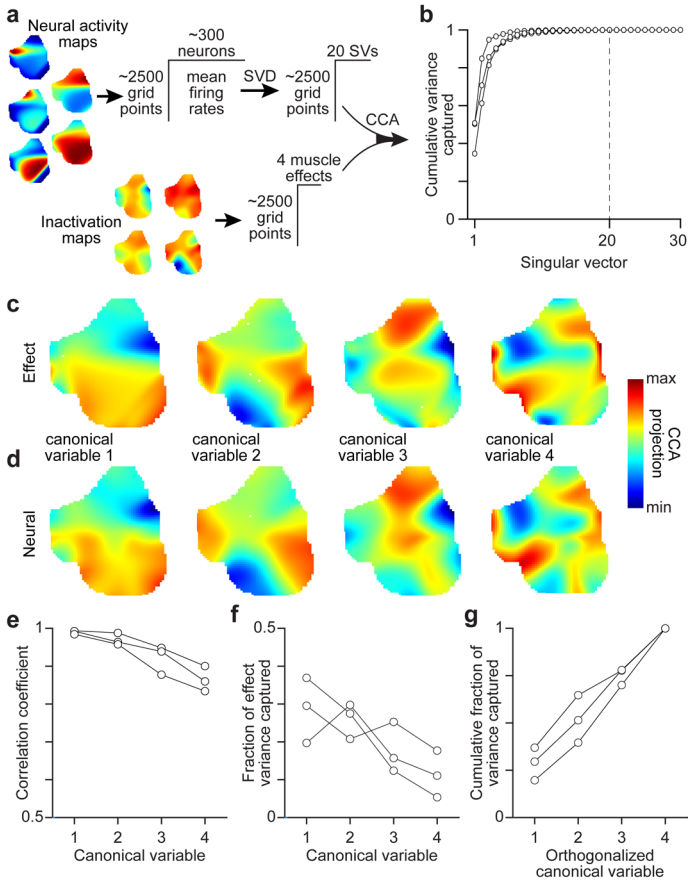

Figure 6 | A CFA activity subspace that aligns with CFA influence. a, Schematic of computing the influence subspace. **b**, Cumulative neural activity variance captured for singular vectors ordered by their singular values, for each of three mice. Each connected set of dots in **b** and **e**-**g** is from a separate mouse. **c**,**d**, Canonical variables for the inactivation effects (**c**) and neural activity (**d**) for one mouse. **e**, Correlation coefficient of canonical variables for each mouse. **f**, Fractional inactivation effect variance captured by canonical variables for each mouse. **g**, Cumulative fraction of inactivation effect variance captured by canonical variables after orthogonalizing their corresponding vectors, for each mouse.

Influence subspace differs from muscle activity and limb kinematic equivalents

Finally, we sought to test the hypothesis that CFA activity subspaces in which activity aligns with muscle activity or limb kinematics in total are mediating CFA's direct influence on muscles. If this were true, we would expect that such subspaces would be substantially similar to the influence subspace (Fig. 7a). To find a subspace in which CFA activity aligns with muscle activity, we assembled maps of the average activity at each grid point for each muscle, in the same manner as for the neural activity maps described above. For each mouse, these muscle activity maps were converted to vectors, assembled into a grid points by muscles matrix, and aligned via CCA with the matrix of dimensionally-reduced neural activity (Fig. 7b,c; Extended Data Fig. 8a,b). To find a subspace in which CFA activity aligns with limb kinematics, we embedded into the limb kinematic state maps described above vectors composed of limb site positions and their corresponding first derivatives from neural recording sessions. We then made maps of the average horizontal and vertical positions of each site at grid points defined on the limb kinematic state maps (Extended Data Fig. 8c). We assembled neural activity maps for each neuron using segments of their firing rate time series that correspond to embedded limb state vectors (Extended Data Fig. 8d). These site position and neural maps were aligned via SVD and CCA (Fig. 7d,e; Extended Data Fig. 8e). For both muscle activity and limb orientation, substantial fractions of variance were captured by canonical variables that were highly correlated with their corresponding neural canonical variables (Fig. 7f-i). The resulting neural canonical vectors span neural activity subspaces where activity aligns with muscle activity or limb kinematics. Activity in these subspaces is akin to the motor cortical activity patterns that correlate with muscle activity or limb kinematics in total that have been described in previous studies.

To evaluate the similarity between influence subspaces and the resulting muscle activity or limb kinematics subspaces, we compared them in two ways. First, we measured the overlap between pairs of subspaces and compared it to the overlap when one was replaced by a subspace defined at random while preserving the neural covariance captured by canonical variables (see Methods). On a scale from 0 (no) to 1 (complete) overlap, the overlap of the influence subspace with the muscle activity subspace ranged from 0.423 to 0.740 across the three mice, substantially above the overlap with 100 randomly defined subspaces (Fig. 7j). However, the overlap between influence and limb kinematics subspaces ranged only from 0.018 to 0.030 and the overlap between muscle activity and limb kinematics subspaces ranged only from 0.025 to 0.060 Both of these overlap ranges were on par with those seen when one subspace was randomly defined. The same relationships between subspaces were reflected in their principal angles as well (Fig. 7k-m). Overlap values were not strongly sensitive to the precise choice of key map parameters (Extended Data Fig. 8f). Thus, there is substantially more overlap between influence and muscle activity subspaces than would be expected by chance, but the same is not true for their respective overlaps with limb kinematics subspaces.

Second, we compared subspace overlap to another benchmark – the degree of overlap expected for the same subspace type but computed with separate sets of time series segments from the neural recordings. Here we computed each of the three subspace types twice, using separate

randomly selected halves of time series segments each time. We then computed overlaps for subspaces of the same type and of different types, both computed with separate segments. We repeated this for 300 different random parcellations of segments and computed the mean overlap across parcellations (Fig. 7n). The overlap between subspaces of different types was much lower here than that of the same types. Thus the overlap was less than expected by chance for two separate measurements of the same subspace. Collectively, these results suggest that subspaces where neural activity aligns with muscle activity are partially, though incompletely, related to CFA's direct influence on muscles, but those where neural activity aligns with gross forelimb kinematics are not.

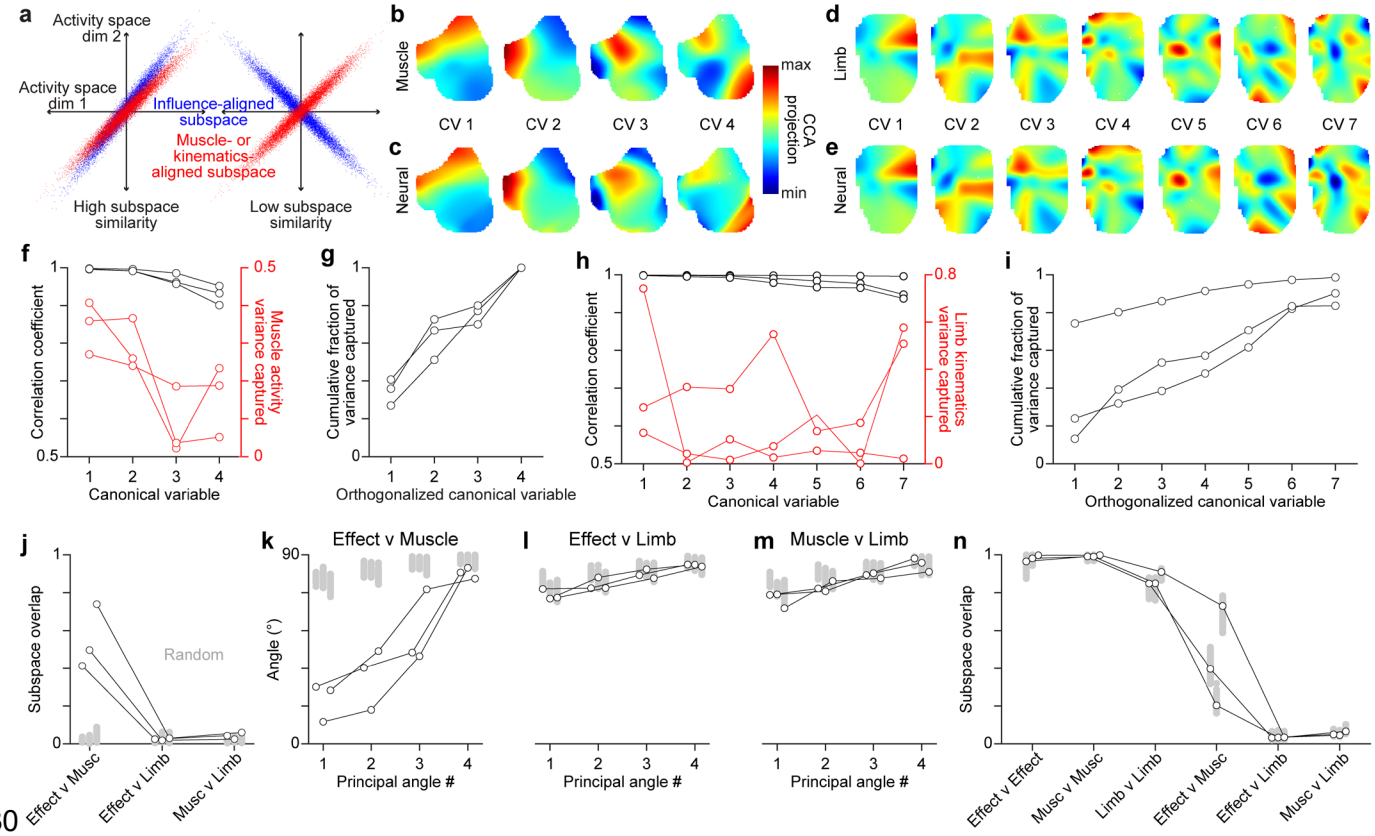

Figure 7 | Divergence between neural activity subspaces aligned with inactivation effects, muscle activity, and limb kinematics. a, Schematic of different scenarios for the overlap between neural activity subspaces. **b**,**c**, Canonical variables for muscle activity (**b**) and neural activity (**c**) for one mouse. **d**,**e**, Canonical variables for limb kinematics (**d**) and neural activity (**e**) for the mouse used in **b**,**c**. **f**, Correlation coefficient (black) and fractional muscle activity variance captured (red) for canonical variables. Each set of connected dots in **f-n** are from one mouse. **g**, Cumulative fraction of muscle activity variance captured by canonical variables after orthogonalizing their corresponding vectors. **h**, Correlation coefficient (black) and fractional limb kinematics variance captured (red) for canonical variables. **i**, Cumulative fraction of limb kinematics variance captured by canonical variables after orthogonalizing their corresponding vectors. **j**, Overlap of different activity subspaces (black circles) compared to 300 estimates of the overlap expected by chance for each animal (gray dots). **k-m**, Principal angles for different activity subspaces (black circles) compared to 300 estimates of the principal angles expected by chance for each animal (gray dots). **n**, Mean overlap (black circles) between subspaces defined from two maps made with separate halves of time series segments, for 300 different segment parcellations (gray dots).

Discussion

Here we have assessed the direct influence of motor cortical output on muscle activity during naturalistic climbing in mice. By quantifying this influence comprehensively – though again not at the finest likely extant resolution – across the muscle activity states that occur during climbing, we have shown that CFA instructs motor output pattern distinctly from the rest of the motor system. CFA acts by exciting physiological flexor muscles, while having less frequent, weaker, and bidirectional effects on their respective antagonists. We have also shown that the firing of individual CFA neurons exhibits a varying degree of specificity for subsets of similar muscle activity states. Finally, our comprehensive approach enabled us to distinguish a neural activity subspace aligned with CFA influence from those aligned with muscle activity or limb kinematics in total. These results suggest that during an ethologically relevant motor behavior, mouse motor cortex appears to selectively direct muscle activity through a neural activity subspace distinct from those to which functional significance has previously been imputed[41,46-50].

It has previously been proposed that motor cortex generates neural activity patterns that directly dictate muscle activity through descending pathways[25,81]. Motor cortex has activity components that are substantially predictive of muscle activity and movement kinematics[38-43], making plausible a primary role in generating the totality of muscle commands during motor cortically dependent behaviors. At the same time, disturbance to motor cortical output often causes hypometric limb movements[1,12,29,64-66], suggesting that motor cortex contributes in some way to driving muscle activity but does not actually instruct motor output pattern in many contexts. Here we have shown that in mice, forelimb primary motor cortex is not involved in generating all forelimb muscle commands, either independently or in unison with the rest of the motor system. The instructive nature of the influence we demonstrate is highlighted by the following: the muscle activity perturbations following optogenetic silencing during climbing match the latency of the fastest possible influence of motor cortex on forelimb muscles in mice[22], are reproducible across animals, are not observed during treadmill locomotion, and are only present during certain muscle activity states during climbing.

Our observation of highly correlated CFA influence on physiological flexors, which are coactivated during locomotion, could reflect an ethological adaptation in which specific cortical regions are adapted to drive specific motor behavior types[82,83]. Several results are consistent with CFA playing a more prominent role in driving rhythmic or iterated locomotor movements[84-87], while other, more rostral regions serve as primary drivers of reaching or eating[88,89]. Our results also align with those from other mammals showing broader extents of motor cortex where electrical stimulation evokes responses in physiological flexors, compared to their antagonists[90]. Ethological organization in motor cortex would seemingly imply organizational differences between species with very different behavioral repertoires. In line with this, humans, who depend less on their forelimbs for climbing and other locomotion, do not exhibit the bias toward responses in physiological flexor muscles with motor cortical stimulation seen in other mammals[68].

It has also been proposed that motor cortex confers an added level of muscle control that improves movement efficacy and task learning by generating muscle activity patterns that cannot

be achieved by other motor system structures[24,91,92]. Our results suggest that this role is mediated by a selective augmentation of ongoing motor output, rather than participation in generating the totality of that output. Moreover, this selective influence differs categorically between muscles of different functional types. We were surprised to observe that direct motor cortical influence on physiological flexors was uniformly excitatory, and that influence on their respective antagonists was distinct not just in magnitude, but also sign, showing a mix of excitation and inhibition. Given the non-stereotyped and adaptive nature of the limb movements mice perform to negotiate unpredictable handhold positions in our paradigm, we had expected a bidirectional influence on all muscles. For example, if motor cortex were monitoring limb state in relation to a target movement and introducing corrections to reduce the deviation from that target, as predicted by current theory[93-95], we would expect bidirectional influence if errors are symmetric around the target. However, we note that this aspect of direct motor cortical influence could differ in other behavioral contexts. In other words, the form of direct motor cortical influence itself could be adapted to behavioral context.

Our observations as well have important implications for the functional interpretation of kinematics-related neural activity. Activity in motor cortex is known to be substantially predictive of both forelimb muscle activity and kinematics[38,39,73,96], and controversy remains regarding how motor cortical activity encodes commands[97,98]. By focusing on a context in which muscle activity and gross limb kinematics are substantially uncoupled, we found that while there exists a neural activity subspace where activity aligns with gross forelimb kinematics, these kinematics do not well predict direct CFA influence, and the kinematics subspace is essentially orthogonal to the other two we identified. However, we cannot rule out that there are other, perhaps finer, kinematic features that more substantially covary with direct CFA influence. For example, CFA could be specifically commanding features relevant to the adaptive aspects of the climbing mice perform, like mediolateral hand position and hand posture during handhold grasping. Despite this caveat, our observations remain a significant revelation given the contemporary prevalence of video-based kinematic tracking for measuring nervous system output. Prominent kinematic features can correlate substantially with activity in a neuronal population, but correlate negligibly with the functional influence of that activity.

In addition, our observations lend significant support to the prominent hypothesis that direct motor cortical influence on muscles comes from a subspace of the overall motor cortical activity space, the "output potent" subspace[50]. We silenced motor cortical activity at a broad diversity of activity states, driving it towards a state of generally low activity in which a subset of inhibitory interneurons are highly active[22,61]. These activity perturbations are thus diverse in their directions through neural activity space. The observation that perturbations induce muscle effects of widely varying magnitudes, including essentially no effect at all, is consistent with the idea that only some activity space dimensions drive muscle activity[50,99]. However, the influence subspace we identified here differs meaningfully from previous definitions of an output potent subspace, and is distinct from subspaces that align with muscle activity or limb kinematics in total. The divergence between subspaces we observed suggests that the ability to predict muscle activity and

727 limb kinematics from motor cortical activity reflects processes beyond the direct control of
728 muscles, like monitoring or predicting body state for computing movement errors[22,44,45].

729

EXPERIMENTAL MODEL AND SUBJECT DETAILS

All experiments and procedures were performed according to NIH guidelines and approved by the Institutional Animal Care and Use Committee of Northwestern University.

Experimental Animals

A total of 50 adult male mice were used, including those in early experimental stages to establish methodology. Strain details and number of animals in each group are as follows: 44 VGAT-ChR2-EYFP line 8 mice (B6.Cg-Tg(Slc32a1-COP4*H134R/EYFP) 8Gfng/J; Jackson Laboratories stock #014548); and 6 C57BL/6J mice (Jackson Laboratories stock #000664).

All mice used in experiments were individually housed under a 12-hour light/dark cycle in a temperature- and humidity-controlled room with *ad libitum* access to food and water, except during experiments. At the time of the measurements reported, animals were 12-18 weeks old. Animals weighed 24–30 g. All animals were being used in scientific experiments for the first time. This included no previous exposure to pharmacological substances or altered diets.

METHOD DETAILS

Climbing apparatus

The climbing apparatus (Extended Data Fig. 1) was housed inside a sound attenuating chamber (H10-24A, Coulbourn). Experimental control was performed using the Matlab Data Acquisition Toolbox, the NI PCI-e-6323 DAQ, and two Arduino Duos. The climbing apparatus itself consisted of a 3D-printed cylindrical wheel with alternating handholds positioned 12 degrees apart from each other. The right handholds were affixed to linear actuators (L-12-30-50-12-I, Actuonix) while the left handholds were statically positioned. A ratchet mechanism was used to ensure that the climbing wheel could only rotate downwards from the mouse. One end of the wheel was supported by a shaft angular encoder (A2-A-B-D-M-D, U.S. Digital). Angular position signals were sent to the Arduinos to track the location of each handhold. When each right handhold reached a position 180° away from the mouse, the linear actuator moved the handhold to a new, randomly chosen mediolateral position. The other end of the wheel was supported by a slip ring (SR20M-LT, Michigan Scientific) that commuted voltage signals to and from the actuators embedded in the wheel. Water rewards were dispensed with a solenoid valve (161T012, NResearch) attached to a lick tube (01-290-12, Fisher), and this dispensation was controlled by Matlab through the NI PCI-e-6323 DAQ. A speaker was used to play a 5 kHz tone for 200 ms whenever rewards were dispensed.

Training

Under anesthesia induced with isoflurane (1–3%), mice were outfitted with titanium or plastic head plates affixed to the skull using dental cement (Metabond, Parkell). Headplates had an open center that enabled subsequent access to the skull, which was covered with dental cement. During headplate implantation, the position of bregma relative to marks on either side of the headplate was measured to facilitate the positioning of craniotomies during later surgeries. After recovery

from headplate implantation surgery, mice were placed on a water schedule in which they received 1 ml of water per day.

At least 4 days after the start of the water schedule, mice were acclimated to handling by the experimenter following established procedures[100]. After acclimation to handling, mice were acclimated to head-fixation over two daily sessions during which they were placed in a 3D printed hutch positioned directly in front of the climbing wheel apparatus and provided water rewards (3 µl per reward) at regular intervals.

Following acclimation, mice underwent daily hour-long training sessions on the wheel apparatus. Training involved an initial stage (1-3 sessions) aimed at training mice to grab for and pull at the handholds in order to rotate the wheel downward and receive water reward. Mice were head-fixed in an upright position facing the front of the wheel so that all four limbs could easily grab onto the handholds, and the right handholds remained fixed. Rewards were triggered by an experimenter's key press whenever a mouse performed any slight rotation of the wheel downward towards his body, and longer or faster bouts were rewarded with additional rewards. Over the course of these sessions, the mice generally learned to associate rotating the wheel with a water reward and began iteratively rotating the wheel.

During the next stage of training (4-10 sessions, median 5), right handholds were kept fixed and mice were encouraged to rotate the wheel for increasingly long bouts. Here, rewards were dispensed for continuous climbing bouts above a threshold distance after the bout ended, using an automated experimental control script written in Matlab. The first ten times during a training session that the threshold distance was met, mice automatically received a water reward. Subsequently, the total distance traveled was compared to those from the previous 10 bouts. If the time was above the 25th percentile value for those 10 bouts, the mouse received one water reward. If it was above the 60th percentile value, the mouse received two water rewards. And if it was above the 90th percentile value, the mouse received four water rewards. Otherwise, the mouse received no water reward. The threshold distance was adaptively adjusted to maintain the reward rate such that the mouse received approximately 1 mL of water over each training session. As a result, if the recent reward rate was too low, the threshold distance was lowered, and if the recent reward rate was too high, the threshold distance was raised. During all subsequent training sessions, the right handhold positions were randomly repositioned along the horizontal axis after rotating past the mouse, though the same reward scheme was used.

Electromyographic recording

Electromyographic (EMG) electrode sets were fabricated for forelimb muscle recording using established procedures[22,53,54]. Briefly, each set consisted of four pairs of electrodes. Each pair was comprised of two 0.001″ braided steel wires (793200, A-M Systems) knotted together. On one wire of each pair, insulation was removed from 1 to 1.5 mm away from the knot; on the other, insulation was removed from 2 to 2.5 mm away from the knot. The ends of the wires on the opposite side of the knot were soldered to an 8-pin miniature connector (11P3828, Newark). Different lengths of wire were left between the knot and the connector depending on the muscle a

given pair of electrodes would be implanted within: 3.5 cm for biceps and triceps, 4.5 cm for extensor carpi radialis and palmaris longus. The ends of wires with bared regions had their tips stripped of insulation, then were twisted together and crimped inside of a 27-gauge needle that facilitated insertion into muscle.

Mice were implanted with EMG electrodes during the surgery in which headplates were attached. The neck and right forelimb of the mouse were shaved, and incisions were made above the muscle to be implanted. Electrode pairs were led under the skin from the incision on the scalp to the incision at the site of implantation. Using the needle, electrodes were inserted through the muscle, and the distal portion of the electrodes was knotted to secure both 0.5 mm bared wire regions within the muscle. The needle and excess wire were then cut away. Incisions were sutured and the connector was affixed with dental cement to the posterior edge of the headplate.

EMG recordings were amplified and digitized using a 16-channel bipolar amplifying headstage (C3313, Intan Technologies). Data was acquired at 4 kHz using the RHD2000 USB interface board (Intan Technologies).

Optogenetic inactivation

After a VGAT-ChR2-EYFP mouse completed a few climbing sessions with randomly positioned handholds, dental cement above the skull was removed and a 2-mm-diameter craniotomy was made above the left caudal forelimb area centered at 1.5 mm lateral and 0.25 mm rostral of bregma. A thin layer of Kwik-Sil (World Precision Instruments) was applied over the dura and a 3-mm-diameter #1 thickness cover glass (Warner Instruments) was placed on the Kwik-Sil before it cured. The gap between the skull and the cover glass was then sealed with dental cement around the circumference of the glass. A custom stainless steel ferrule guide (Ziggy's Tubes and Wires) was then cemented to the headplate a certain distance above the surface of the brain. This distance was set to ensure that the cone of light emanating from a 400 μm core, 0.50 NA optical patch cable terminating in a 2.5 mm ceramic ferrule (M128L01, Thorlabs) would project a spot of light 2 mm in diameter onto the surface of the brain (Fig. 2a). The ferrule guide enabled quick and reliable positioning of the ferrule above the brain surface so that a large expanse of cortex could be illuminated. In previous experiments using this method for inactivation, control experiments in wildtype mice showed no discernible muscle activity perturbation in response to light[22]. Moreover, the short latency at which we measure effects here would preclude visually-driven responses.

To attenuate firing throughout motor cortical layers, we used a 450 nm laser (MDL-III-450, Opto Engine) to sporadically apply 25 ms light pulses at an intensity of 10 mW/mm^2 to the brain surface. During each session that involved inactivation of the motor cortex, trials were initiated after the distance climbed on the current bout exceeded a random threshold between 0° and 20°, and a mouse was actively climbing. Light was applied on a random third of trials. This ensured that inactivation and control (no light) trials were broadly distributed across the muscle activity states that occur during climbing. Inactivation trials were never triggered less than 5 seconds apart. A total of 2292-5115 trials (median = 2715) were collected in each mouse (roughly

1/3 inactivation and 2/3 control), spanning 11-37 (median = 18.5) climbing sessions.

Neural recording

For some mice, optogenetic inactivation sessions were followed by 3-4 daily neural recording sessions that typically lasted an hour each. One day prior to the first neural recording session, the cover glass and Kwik-Sil were removed, a small durotomy was made in the craniotomy center, and a Pt/Ir reference wire was implanted to a depth of 1.5 mm in left posterior parietal cortex. Opaque silicone elastomer (Kwik-Cast, World Precision Instruments) was used to cover the craniotomy after surgery and between recording sessions. At the time of recording, the exposed brain surface was covered with agarose and silicone oil or liquid paraffin oil. A Neuropixel 1.0 (IMEC) was subsequently inserted to a depth of 1.5 mm into the brain (Fig. 5a) at a rate of 3-5 μm/s using a motorized micromanipulator (MP-225A, Sutter Instrument). Electrode voltages were acquired at 30 kHz and bandpass filtered at 0.3 to 10 kHz using SpikeGLX (Bill Karsh, https://github.com/billkarsh/SpikeGLX), and then sorted with Kilosort 2.0[101] (https://github.com/MouseLand/Kilosort2). Sorted units were assigned to separate cortical layers based on the depth below pia assigned to the waveform centroid in Kilosort and the laminar depths reported in[102]. To be conservative, units near laminar boundaries were ignored in analyzing neural responses by layer.

EMG preprocessing

For both optogenetic inactivation and Neuropixel recording sessions, EMG measurements were downsampled to 1 kHz, high-pass filtered at 250 Hz, rectified, and convolved with a modified Gaussian filter kernel. We used causal filtering to enable precise assessment of perturbation latencies, and hence used a Gaussian filter kernel that was initially defined with a 10 ms standard deviation, but then had amplitudes for times < 0 – its "backwards in time" side – set uniformly to zero. EMG traces were then z-scored across all time points during the given session.

Video recording and analysis

A high-speed, high-resolution monochrome camera (Blackfly S USB 3.0, 1.6 MP, 226 FPS, Sony IMX273 Mono; Teledyne FLIR, Wilsonville, OR) with a 12 mm fixed focal length lens (C-Mount, Edmund Optics) was positioned to the right of the head-fixed mouse during inactivation and neural recording sessions, and videos were acquired under a near-infrared light source at 100 frames per second with a resolution of 400 by 400 pixels. During optogenetic inactivation sessions, the camera was triggered to start recording using StreamPix software (NorPix, Inc; Montreal, Quebec, Canada) 20 ms before each inactivation or control trial. Each recording lasted 500 ms beyond the 25 ms light or control command pulse that marked the trial. Annotation of behavior was accomplished using DeepLabCut[55]. To enable better markerless tracking, the right forelimbs of mice were shaved and tattooed (Black tattoo paste, Ketchum Mfg., Lake Luzerne, New York) at 8 different sites along the right arm. All videos were also adjusted with ffmpeg (ffmpeg.org) to improve brightness and contrast. DeepLabCut using the ResNet-50 neural network with an Adam

optimizer was trained on the annotated images for 1,030,000 iterations; on ~4000 randomly sampled video frames across mice and sessions, we provided manually labeled locations of the 8 forelimb sites for training: shoulder, two sites between the shoulder and elbow, elbow, two sites between the elbow and wrist, wrist, and tip of the last digit (Fig. 4b). The training set comprised 80% of the labeled frames.

All DeepLabCut-tracked forelimb site trajectories were then exported to Matlab for further post-processing to remove outliers (Extended Data Fig. 9). First, sites in each video time series that were assigned by DeepLabCut a likelihood (i.e. its confidence that a site was correctly labeled) lower than 0.75 were replaced with an interpolated value using the median of the ten previous and ten following values (Matlab function fillmissing). Next, outliers in site position time series were identified using the median absolute deviation (MAD): shoulder coordinates were constrained to lie within 1.5 MAD from their median, digit coordinates to be within 3 MAD from their median, and all other joints to be within 2 MAD from their respective medians. Outliers were replaced with an interpolated value using a moving median of window length 10. Lastly, limits were imposed on the pairwise distances and angles between neighboring joints (the shoulder-elbow, the elbow-wrist, and the wrist-digit tip) such that the angle between shoulder-elbow and elbow-wrist could not exceed 180 degrees, and the distances between each of these joints were within 2 MAD of their medians. Site positions not meeting these criteria was also replaced with an interpolated value using a moving median of window length 10.

Behavioral analysis

For each animal, we computed the principal angles between muscle activity time series for each of the first 20 climbing sessions that used the automated control script, compared to the 20th session. PCA was first performed on the muscle activity time series for the four recorded muscles. The top 2 PCs were used for the subsequent calculations, since they captured over 90% of activity variance. Principal angles were then computed using a standard approach: The cross-covariance matrix was computed for the two PC time series from the given session and those from the 20th session, and SVD was applied to the result. The principal angles were then computed as the real part of the inverse cosine of the resulting singular values.

Sample entropy[57] was computed for the muscle activity time series for the four recorded muscles using the 'sampen' function downloaded from the Matlab file exchange, with the embedding dimension set to 10.

Muscle activity state maps

To obtain 2D muscle activity state maps on which nearby states are highly similar, UMAP was applied to segments extracted from the EMG time series collected during the optogenetic inactivation sessions, and their corresponding first derivatives. For each control or inactivation trial, we defined 13 overlapping 50 ms epochs centered every 10 ms from -55 to +75 ms from command pulse onset (Fig. 2c). Then, for each 50 ms epoch, we averaged the EMG traces for each muscle over 5 ms bins, and concatenated the resulting values for the four muscles and their

corresponding first derivatives, yielding one 80 by 1 vector. Thus in these vectors, the first 40 values reflect the EMG signals from 4 muscles, and the last 40 values reflect their first derivatives (Fig. 2c). UMAP (Matlab function run_umap, from the Matlab File Exchange, with the following parameters: n_neighbors = 30, n_components = 2, min_dist = 0.3, metric = Euclidean) was applied to embed all resulting vectors from both inactivation and control trials, generating the 2D muscle activity state maps (Fig. 2d). Using two dimensions here simplified subsequent quantification of inactivation effects across muscle activity states. Using overlapping epochs ensured that embedded state vectors from individual trials formed continuous trajectories across the resulting maps (Fig. 2d).

Muscle inactivation effect maps

To quantify the influence of CFA inactivation on muscles across the muscle activity state maps, we first excluded outlying embedded state vectors. For each embedded vector on a given map, we computed its mean Euclidean distance on the map from all other vectors. A vector was deemed an outlier if its mean distance was more than three standard deviations above the mean for all other vectors. Some mice had no embedded vectors classified as outliers, and the percentage classified as such was overall less than 1%.

We then defined uniformly spaced grid points across each map at which we would calculate the effects of CFA inactivation. Since the number of grid points of a fixed spacing across a map would depend on the scale of the map, which could vary between maps, we rescaled the coordinates of each 2D map $\boldsymbol{rd}$ to $\boldsymbol{rd_{rescaled}} = \frac{\boldsymbol{rd} - min(\boldsymbol{rd})}{std(\boldsymbol{PD})} * 10 + weightPara * 2$, where PD is the set of all Euclidean distances between embedded vectors, and $weightPara$ is the standard deviation of the Gaussian function used to weight trials in computing the inactivation effect size at each grid point (see below; "spatial filter width" in Extended Data Figs. 3, 6 and 8). For all maps, we used $weightPara = 5$, which yielded ample statistical power for testing the significance of inactivation effects.

We then sought to compute the inactivation effect at each grid point, based on trials that began from muscle activity states close to each given grid point. Our approach was to compute effects for each grid point from inactivation and control trial averages made from all trials of a given type, but where trials were weighted by the distance on the 2D map from the given grid point to the weight epoch state vector for the given trial. The weight epoch state vector was defined as the embedded vector from the epoch spanning -40 to +10 ms from command pulse onset on the given trial (i.e., the epoch just before any direct effect could begin; Fig. 2c,e). However, the density of these vectors varied across maps, meaning that the summed weight of trials contributing to trial averages would vary across grid points as well. Given this, to ignore grid points around which the weight epoch state vectors were too sparse for reliable trial averaging, a grid point was designated as "valid" only if $\sum_i W_{g,i} >$ a threshold. $W_{g,i}$ was in essence a Gaussian function, defined as $W_{g,i} = exp\left(-\frac{(Loc_g - Loc_i)^2}{2 * weightPara^2}\right)$, where i is the weight epoch state index (one per trial), g is the grid point index, and $Loc_g - Loc_i$ is the Euclidean distance between the weight epoch state and the

grid point on the 2D map. We used a threshold of 10, which led nearly all grid points falling within a convex hull surrounding embedded vectors to be classified as valid, while nearly all grid points falling outside this hull were not. Grid points not designated as valid were ignored in subsequent calculations. All embedded muscle activity state vectors (those not deemed outliers) were close to valid grid points, and so this criteria did not prevent an appreciable contribution from any vectors. Together with the very small fraction of embedded vectors ignored as outliers, this implies that our analysis involves practically all the muscle activity states that occur during climbing.

Next, we calculated the trial-averaged activity for each muscle at each valid grid point, separately for the inactivation and control trials. For each muscle and trial type, we extracted segments of their activity time series from -10 to +30 ms relative to command pulse onset. We then took a weighted average of these segments across each trial type, where each segment was weighted by $W_{g,i}$. At each grid point, this produced control and inactivation trial averages for each muscle, where each trial is weighted by the distance from the grid point to the muscle activity state just prior to any direct inactivation effect during the trial.

Finally, we quantified the size of the CFA inactivation effect at each valid grid point by comparing the rates of change (slopes) in inactivation and control trial averages 0 to +20 ms from command pulse onset (Fig. 2f). Muscle activity at time 0 ms was defined as the mean from -10 to +10 ms, and likewise activity at +20 ms was defined as the mean from +10 to +30 ms. Quantifying the inactivation effect by taking the ratio or difference between the control and inactivation slopes returned qualitatively similar results. We used the difference because it was more easily interpretable (Fig. 2g).

We also sought to assess the similarity between nearby state vectors embedded in local neighborhoods across maps, and thus the similarity of states from which trials contributing strongly to given grid point trial averages began. To do this, we computed the average Euclidean distance in the full 80D space between state vectors embedded near each grid point on the 2D maps. We first calculated the Euclidean distances between all possible pairs of embedded vectors in the full 80D space. For each grid point, we then computed a locally weighted average of these distances. In these averages, each distance d was weighted by a Gaussian function of the distance on the 2D map between the grid point and the midpoint between the corresponding two embedded vectors (i.e., those separated by d), using the same Gaussian function as above ($weightPara$ = 5).

Analysis of inactivation effect maps

To determine which muscle activity states were significantly influenced by CFA inactivation, we performed a two-tailed nonparametric permutation test at each grid point by computing the probability of obtaining the observed inactivation effect size by chance. For each animal, and each grid point, 300 permutations were performed by first randomly splitting control trials into two groups, each with a number of trials equal to the number of inactivation trials. The number of trials in both the control ($N_{control}$) or inactivation ($N_{inactivation}$) trials groups was such that if $N_{total_control}/2 > N_{total_inactivations}$, we would set $N_{control} = N_{total_inactivations}$; otherwise, $N_{inactivation} = N_{total_control}/2$. Since our experiments were designed to collect twice

as many control trials as inactivation trials, control trials could be sampled without replacement during the splitting process. For each grid point, and for each permutation, we calculated the inactivation effect size using the control trial average, computed as above, for one randomly selected group of control trials and the inactivation trial average, also computed as above. For each permutation, and at each grid point, we also calculated the effect size expected by chance ("null") using trial averages for the two control trial groups. Then, at each valid grid point, we randomly picked one inactivation effect size from the 300 permutations and compared that with all 300 null effect sizes. To calculate the p-value, we compared the 300 null effects with the randomly chosen inactivation effect, calculated the fractions of the null effects where the null effects were smaller than or greater than the inactivation effect, and multiplied the smaller fraction by 2. To correct for multiple comparisons, the Benjamini-Hochberg method was used to control the false discovery rate (Matlab function fdr_bh). The effect size at each grid point was considered to be statistically significant if the FDR-corrected p-value was less than 0.05 (i.e., the likelihood of false discovery was 5% or less). We also used p-values across all grid points to estimate the fraction of null hypotheses that were false for each mouse (Matlab function mafdr), and used this as the fraction of grid points exhibiting inactivation effects.

The 2D autocorrelation for inactivation effect maps was computed using the Matlab function xcorr2, and normalized by the maximum value for the given map. For controls, maps were generated using null effect sizes at each grid point from one of the 300 comparisons of effect sizes computed with two separate sets of control trials described above. To test that inactivation effect maps differ across muscles for each animal using the Kruskal-Wallis test, we first removed inactivation effect sizes from the maps until the minimum distance between the grid points for remaining effect sizes was greater than 3x the standard deviation of the Gaussian function used for computing grid point trial averages. This ensured that remaining effect sizes were effectively computed from separate trials. Linear models were fit to scatter plots using the Matlab function fitlm. The 2D correlation between inactivation effect maps was computed using the Matlab function corr.

Neural activity maps

To compare firing patterns in CFA to inactivation effects, we generated maps of average firing rate across grid points for each recorded CFA neuron. Since neural recordings and optogenetic inactivations were carried out in separate sessions, the first step was to align muscle activity states from the neural recording sessions with those from the optogenetic inactivation sessions by embedding the former set of states on the same 2D map used to make inactivation effect maps. To do this, we identified a large number of 50 ms epochs during active climbing from the neural recording sessions. To identify these epochs, peaks in the activity summed across all four muscles were first detected using the Matlab function findpeaks, where a peak was identified if a sample in the time series surpassed a threshold equal to the mean + ½*std. Based on visual inspection, this identified many peaks that always fell during active climbing. Next, two time points ("onset points") were randomly selected on either side of the peak, such that the expected interval from

the peak to each point was 50 ms. Next, any onset point was eliminated if it was less than 50 ms from either the previous or subsequent onset point. This ensured that 50 ms epochs defined around onset points would not overlap. For each remaining onset point, EMG time series segments from the epoch spanning -45 to +5 ms relative to the onset point were extracted. Each recording session typically yielded upwards of 30,000 epochs. Of these, 10,000 were randomly selected for each session to limit compute time for subsequent calculations. State vectors (80 x 1) were assembled for each epoch as above.

We then embedded the resulting 10,000 state vectors from each neural recording session onto the muscle activity state maps previously defined via UMAP to make inactivation effect maps for the given mouse. We next extracted the segments of each neuron's firing rate time series corresponding to each state vector. For each neuron and at each valid grid point, the firing rate segments were averaged as above, weighting each segment by a Gaussian function of the distance on the map from their corresponding state vector to the grid point ($weightPara = 5$). The resulting segment-averaged firing rates were then averaged across time, yielding a single scalar firing rate value for each neuron at each grid point (Fig. 5f, right, and g).

For analysis described below, only putative pyramidal neurons were included. To identify these neurons, we pooled the spike widths of all sorted single units, obtaining a bimodal distribution with separate modes reflecting wide-waveform, putative pyramidal neurons and narrow-waveform, putative interneurons[22,103]. Neurons with spike widths > 450 µs were classified as putative pyramidal neurons.

To estimate the number of neurons that showed behaviorally-dependent firing across muscle activity states, we first generated an empirical null distribution for the degree of variation across neural activity maps separately for each neuron. To do so, we reassigned each state vector from neural recording sessions to the location of a different, randomly selected vector on the map and recomputed the neural activity maps. We repeated this 500 times to yield 500 permuted maps for each neuron. To assess behaviorally-dependent variation, the skewness of the original neural activity map values and those on the 500 permuted neural activity maps was calculated (Matlab function kurtosis) as $k = \frac{E(x-\mu)^4}{\sigma^4}$, where x is the set of firing rates at all grid points, μ is the mean of x, σ is the standard deviation of x, and $E(\cdot)$ represents the expected value. Significance was assessed using a p-value, defined as the fraction of null distribution skewness values greater than the original. From the p-values for all neurons in a given mouse, we estimated the fraction of false null hypotheses (Matlab function mafdr) and used this as the fraction of cells with behaviorally-dependent firing.

The sparsity for neural activity maps was computed using the formula given in [76]. Here the sparsity is the mean of the firing rates at each grid point on the map, squared and divided by the mean of the squares of the firing rates at each grid point. Note that as in the original reference, this minimizes the dependence of sparsity values on the frequency that each state on the map is visited. The 2D correlation between neural maps was computed using the Matlab function corr.

Neural subspace for inactivation effects

Separately for each animal (n = 3 mice), neural subspaces were identified using singular vector canonical correlation analysis (SVCCA[80]) applied to the inactivation effect and neural activity maps. A singular value decomposition (SVD)-based approach was taken here because the number of recorded neurons was much larger than the number of recorded muscles. Neurons with mean firing rates lower than 0.1 Hz over the 10,000 epochs were excluded. Two matrices were generated for alignment with SVCCA. One matrix had a column for each neuron generated by vertically concatenating the successive columns of that neuron's neural activity map. The resulting matrix had dimensions of N_{grid} by N_{neuron} , where N_{neuron} is the number of neurons recorded in given mouse across recording sessions (putative pyramidal, mean firing rate over 0.1 Hz), while N_{grid} is the number of valid grid points. The second matrix was made in similar fashion by vertically concatenating the successive columns of the inactivation effect map for each muscle. This matrix had dimensions of N_{grid} by 4 since there were four muscles recorded.

SVCCA was conducted in two steps. Neural activities were first soft normalized[104] using a soft normalization constant of 5 Hz. SVD was performed using numpy.linalg.svd in python to decompose the neural data into left singular vectors, a diagonal matrix containing singular values and right singular vectors. Next, the diagonal matrix, truncated to just the top 20 singular values, was multiplied with the corresponding top 20 right singular vectors, resulting in 20 neural activity components. CCA was applied to the N_{grid} by 20 matrix of these components and the N_{grid} by 4 matrix of inactivation effects. Twenty dimensions were retained here because the amount of variance captured and CCA alignment quality (canonical variable correlation) saturated at around 20 dimensions. CCA was also repeated with individual columns of the matrix of inactivation effect sizes in place of the full 4D matrix to show that there were CFA activity components correlated with the inactivation effects for each individual muscle.

To compute the additional variance captured by each successive canonical vector, canonical vectors were orthogonalized with the Gram-Schmidt process. In addition, we randomly shuffled the neuronal firing rate segments relative to the embedded vectors on the maps to generate shuffled neural activity maps, and performed SVCCA again as a negative control. The highest correlations between the canonical neural and effect vectors for all animals were less than 0.75. To further verify the effectiveness of SVCCA, separately for each mouse we split the inactivation and control trials from inactivation experiments randomly into two groups, calculated separate inactivation effect maps for each group, used SVCCA to find subspaces where CFA activity aligns with each of them, and calculated the principal angles between the two resulting neural activity subspaces. This procedure was repeated 300 times (Extended Data Fig. 6b).

Using the CCA results for individual columns of the matrix of inactivation effect sizes (i.e., results for effects on individual muscles), we computed the effective weight of each neuron's activity in each canonical variable by matrix multiplying the neuron to singular vector coefficients and the singular vector to canonical vector coefficients. The four individual muscle effect size vectors yielded four effective weights for each neuron. To compare across animals in which we had recorded different numbers of neurons, we normalized these weights to have a median of 1 for each animal (Extended Data Fig. 7). To measure the relative contribution of neurons across all

four muscles, we computed the norm of the 4-element vector composed of the weights for each muscle for a given neuron, and again normalized these so the median for each animal was 1.

Overlap between subspaces

To find a neural activity subspace that aligned with muscle activity itself, we again performed SVCCA using the N_{grid} by N_{neuron} matrix of mean firing rates. However, in place of the matrix of inactivation effect sizes, an N_{grid} by 4 matrix was used where each column reflected the average activity at each grid point for one of the four muscles, computed just as was done for the neural activity maps, except muscle activity from inactivation sessions was used. Similar methods were used to find a neural activity subspace that aligned with limb kinematics. Here, instead of 4 columns, we began with 16 – the horizontal and vertical coordinates for the 8 positions tracked along the right forelimb. Since the 16 kinematic variables were highly correlated, SVD was also used on this matrix to reduce its dimensionality to 7 before performing CCA.

To compute principal angles between two neural activity subspaces, we orthonormalized the neural canonical vectors defining each subspace, computed the cross-covariance matrix for the two sets of vectors, computed the singular value decomposition of the matrix, and calculated the inverse cosine of the singular values in degrees[105]. We also measured the degree of overlap between neural activity subspaces using the metric $OL = \frac{\sum var(C_{s2} * M_{reproj})}{\sum var(M_{proj})}$, where M is the original matrix of neural activity, C_{s1} and C_{s2} are matrices comprising coefficient vectors that that define the two subspaces, $M_{proj} = C_{s1} * M$, and $M_{reproj} = C_{s1}^{T} * M_{proj}$.

To verify the significance of the overlap, we permuted the order of coefficients in the columns of C_{s2} to get $\underline{C_{s2}}$. We only used $\underline{C_{s2}}$ when the total variance of $\underline{C_{s2}} * M$ was in between 0.8 - 1.2 times the total variance of $C_{s2} * M$. We then repeated subspace overlap calculations for 100 different $\underline{C_{s2}}$ meeting this criteria. This ensured that the resulting null distribution of overlap values did not differ from the actual value simply because the $\underline{C_{s2}}$ captured much less variance than C_{s2}. We also calculated subspace overlap for subspaces that were each computed using just one half of the total epochs from the neural recording, repeating this for 300 different random parcellations of the epochs.

Acknowledgments
The authors would like to thank J. Glaser for helpful advice and D. Dombeck, M. Elbaz, and D. Xing for comments on the manuscript. This research was supported by a Searle Scholar Award, a Sloan Research Fellowship, a Simons Collaboration on the Global Brain Pilot Award, a Whitehall Research Grant Award, The Chicago Biomedical Consortium with support from the Searle Funds at The Chicago Community Trust, and NIH grant DP2 NS120847 awarded to A.M.

Author Contributions
The climbing wheel apparatus was designed and built by A.S. with guidance from A.M. Data were collected by N.K. and Z.M., with help from A.S. and A.K. Data analysis methods were developed and applied by Z.M., N.K., and A.M., with early assistance from M.A. and M.Y.

Competing interests
The authors declare no competing interests.

Materials & Correspondence
Correspondence and request for materials should be addressed to A.M. (andrewmiri@northwestern.edu).

Data availability
The data that support the findings of this study are available from the corresponding author upon reasonable request.

Code availability
All Matlab code for data analyses and CAD files for 3D printed wheel components will be made available on GitHub upon publication.

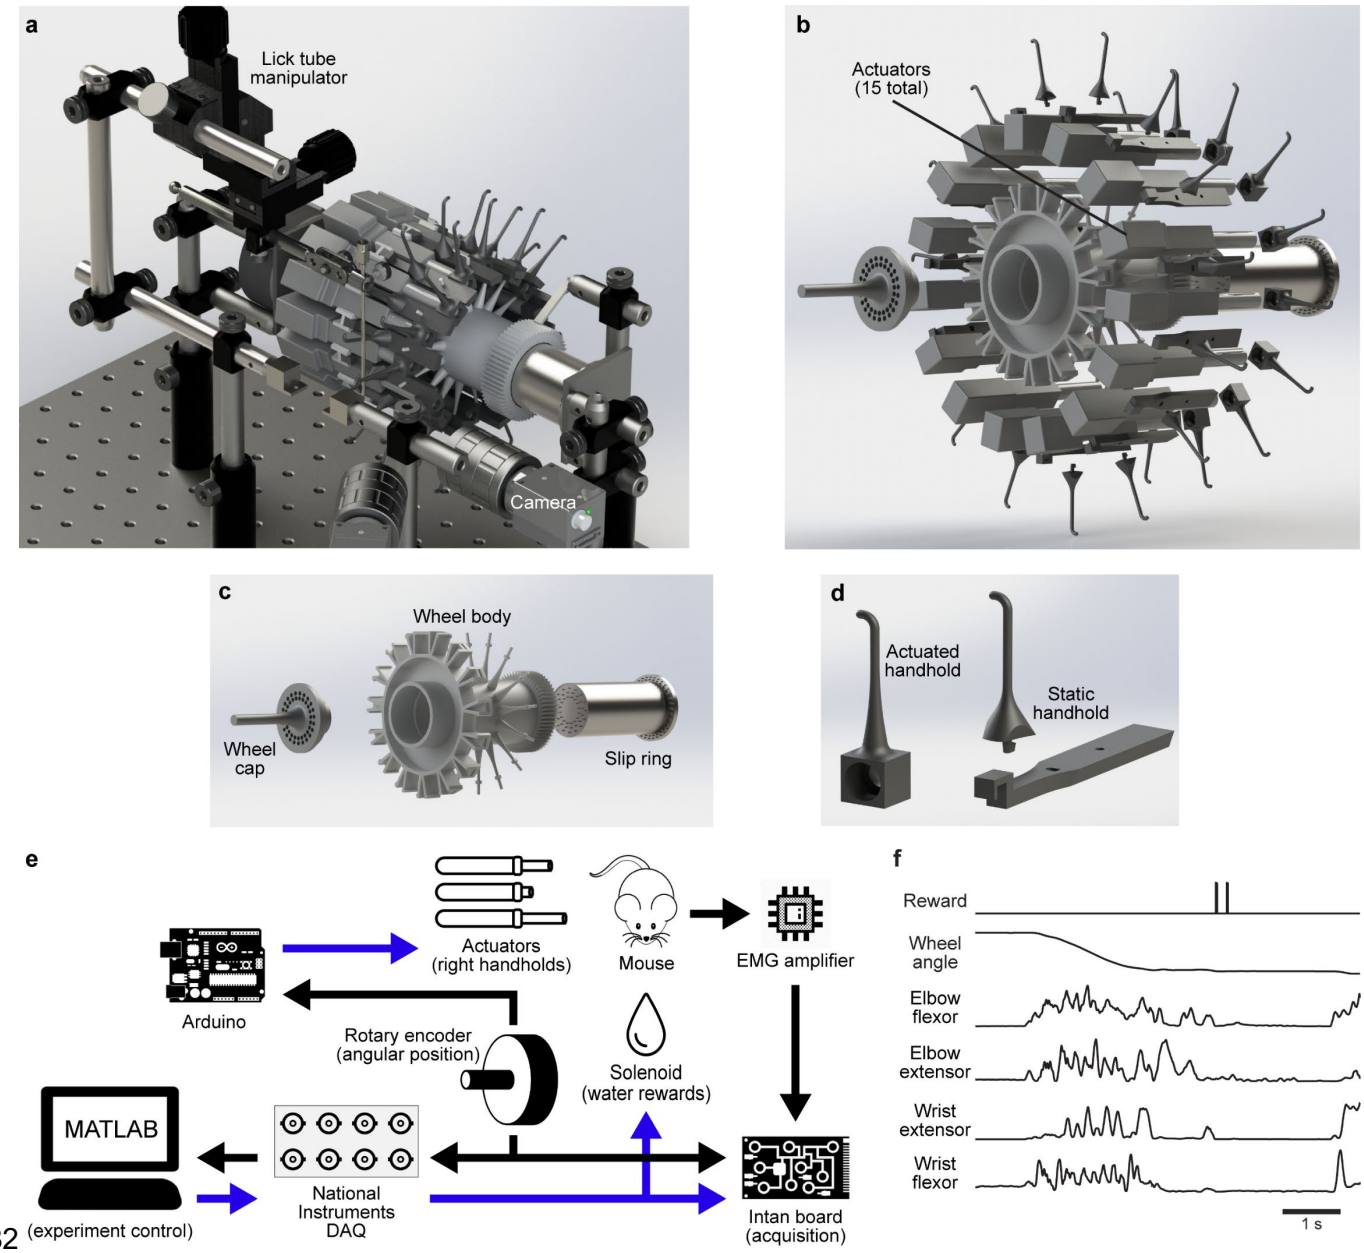

Extended Data Figure 1 | Head-fixed climbing paradigm. a, 3D illustration of the climbing apparatus, including the lick tube manipulator with lick tube, mounts for head fixation, and cameras. **b**, 3D exploded illustration of the wheel with the separate components visible, showing the slip ring, actuators, and 3D printed parts. **c**, Same as **b**, but without actuators and handholds. **d**, The two types of 3D printed handholds. Actuated handholds slide onto actuators. Static handholds attach to the wheel body between actuators. **e**, Flow chart illustrating data acquisition and experimental control. Blue arrows indicate command output signals. Black arrows indicate measured signals. **f**, Example time series during climbing: solenoid command for dispensing reward, wheel angle from optical encoder, and four channels of EMG.

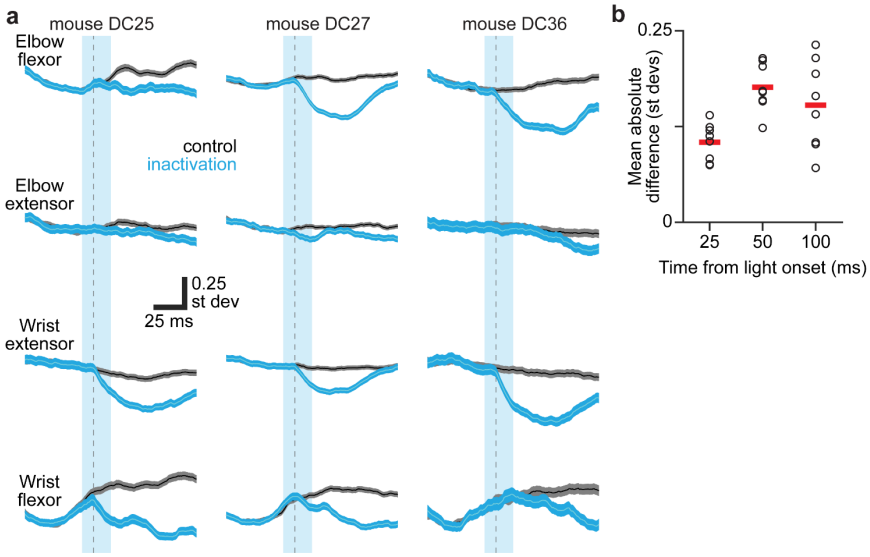

Extended Data Figure 2 | Effect of CFA inactivation on muscle activity. a, Mean ± sem muscle activity for control (gray) and inactivation (cyan) trials for the recorded forelimb muscles in 3 individual mice. Vertical cyan bars indicate the 25 ms epoch of blue light applied to CFA. Gray dotted lines are 10 ms after light onset, the shortest latency at which effects can be observed. **b**, Absolute difference between inactivation and control trials at 25, 50, and 100 ms after trial onset for individual animals (black circles) and the mean across animals (red bars).

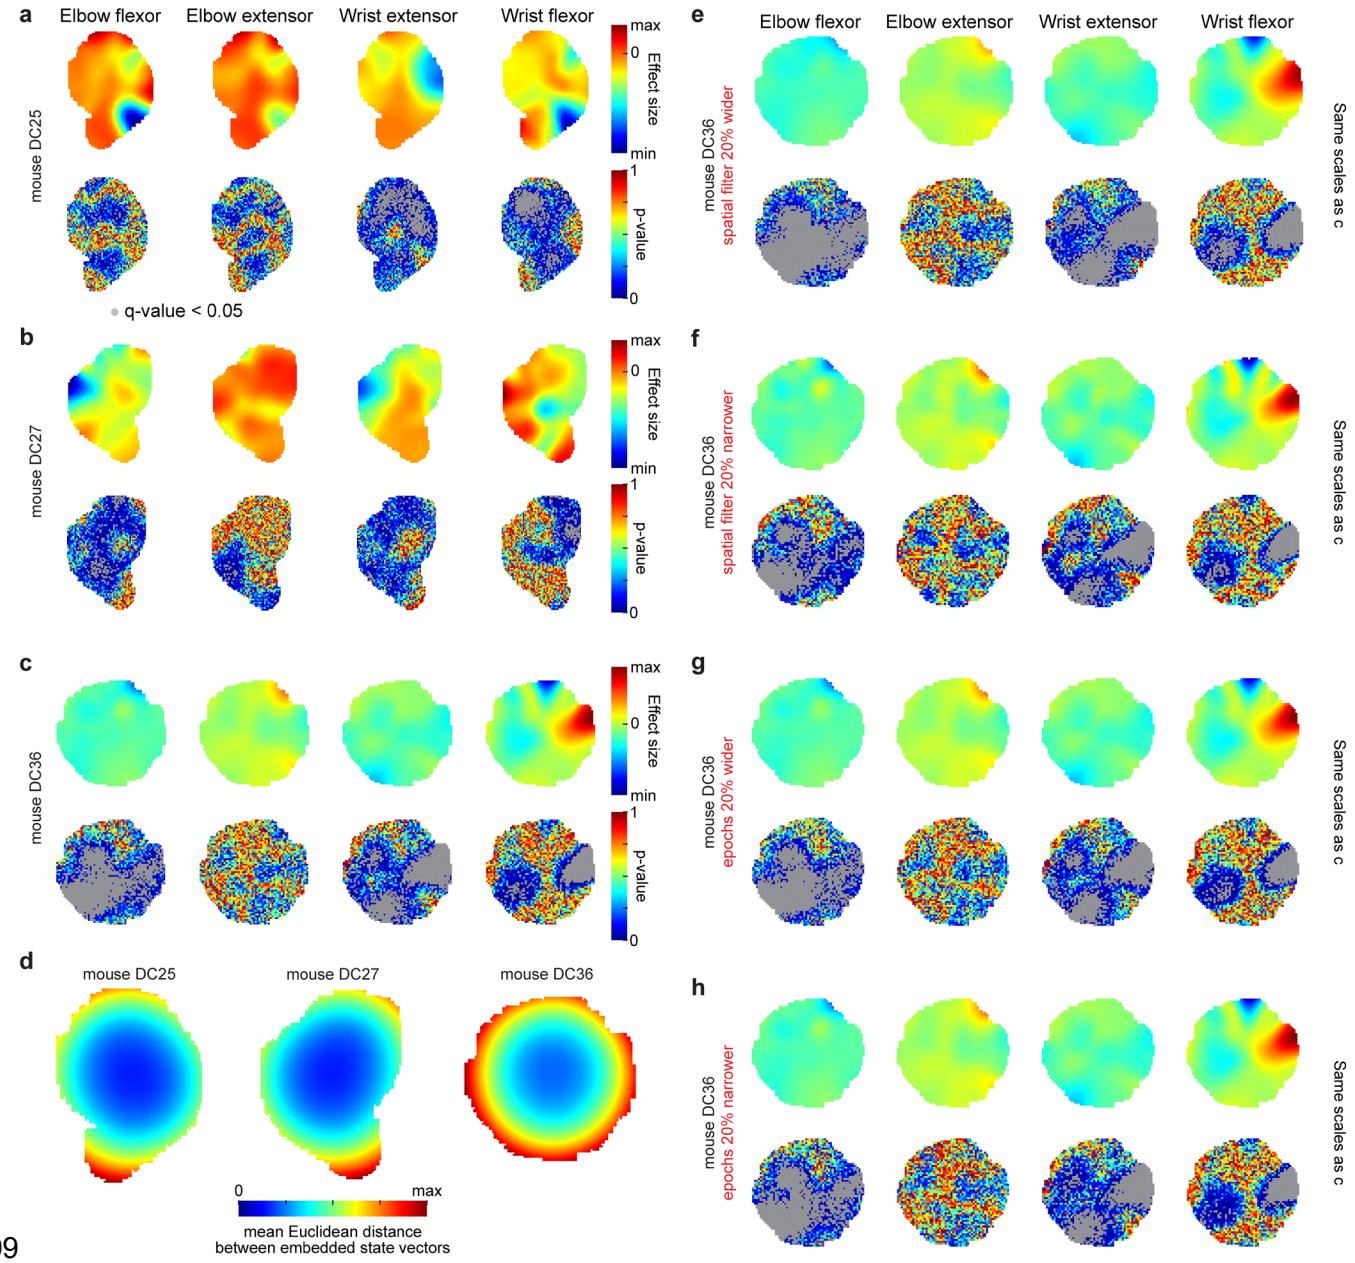

Extended Data Figure 3 | CFA inactivation effects across muscle activity state space. a-c, Inactivation effect maps and corresponding p-value maps for each of the four muscles in three additional mice. Q-values (gray overlay) reflect the expected rate of false discovery below the corresponding p-value[63]. **d**, For three additional mice, maps in which each grid point is colored by the mean distance (in 80D) between all pairs of embedded states, with each distance weighted by a Gaussian function of the pair's mean distance from the grid point on the 2D map. **e-h**, Same as **c**, but following calculation with slightly different parameters (red text).

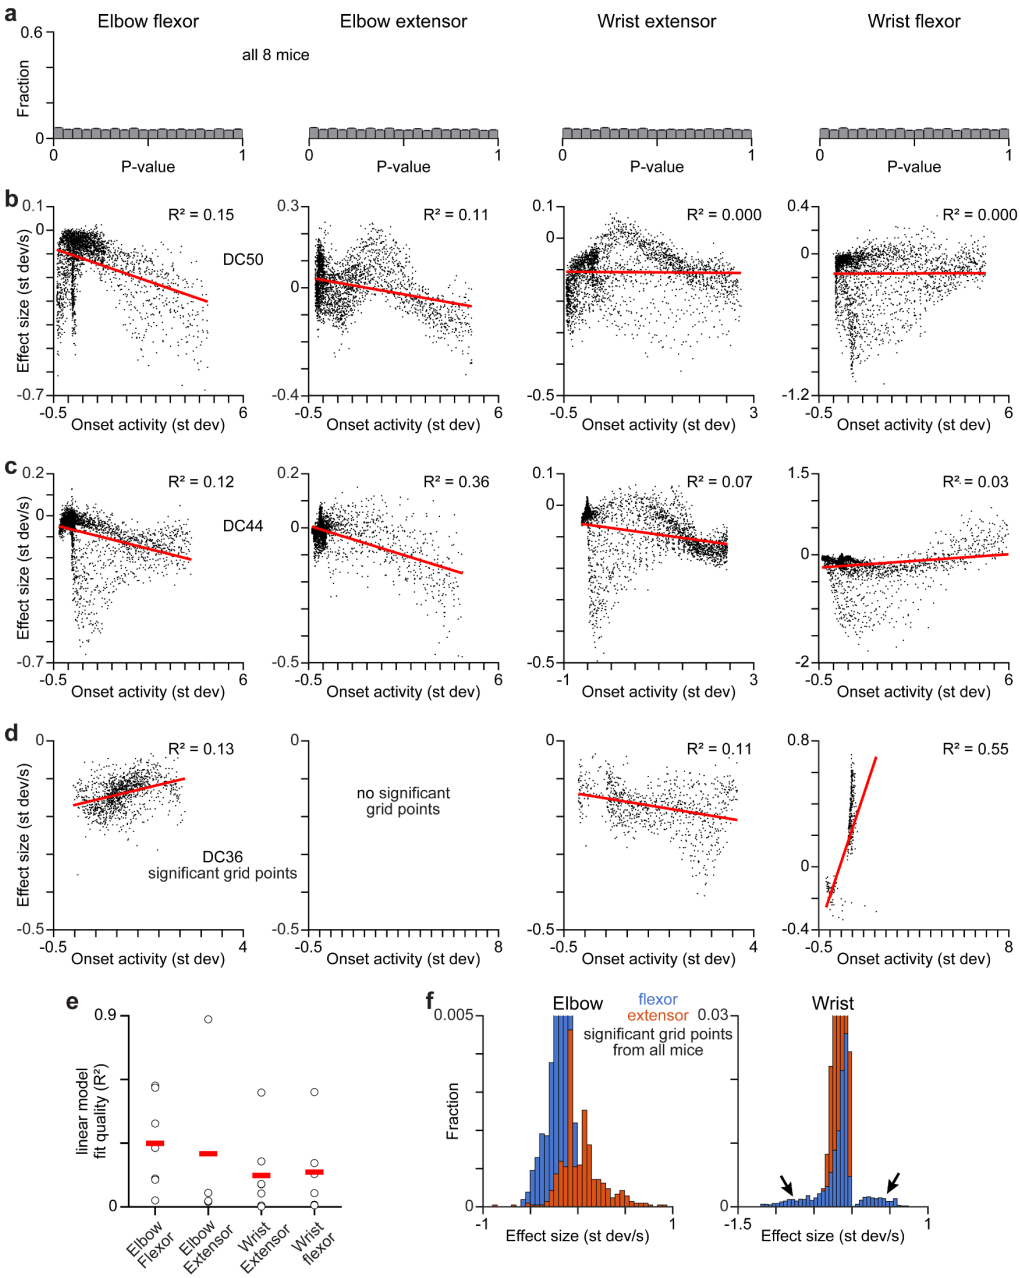

Extended Data Figure 4 | CFA selectively excites physiological flexors. a, P-value distributions for inactivation effects on each muscle when calculated by comparing control trials to a separate set of control trials, for all grid points across all 8 mice. **b**,**c**, Scatterplots of inactivation effect size versus muscle activity at trial onset for two additional mice. Each point reflects a different grid point. R^2 is for a linear fit (red). **d**,**e** Same as Fig. 3f (**d**) and 3g (**e**), except only including grid points where effect size was significantly different from zero. Note that the y-axis scale in **e** differs from that in Fig. 3g. **f**, Same as Fig. 3j, but zoomed in 10-fold to clarify rare larger effects.

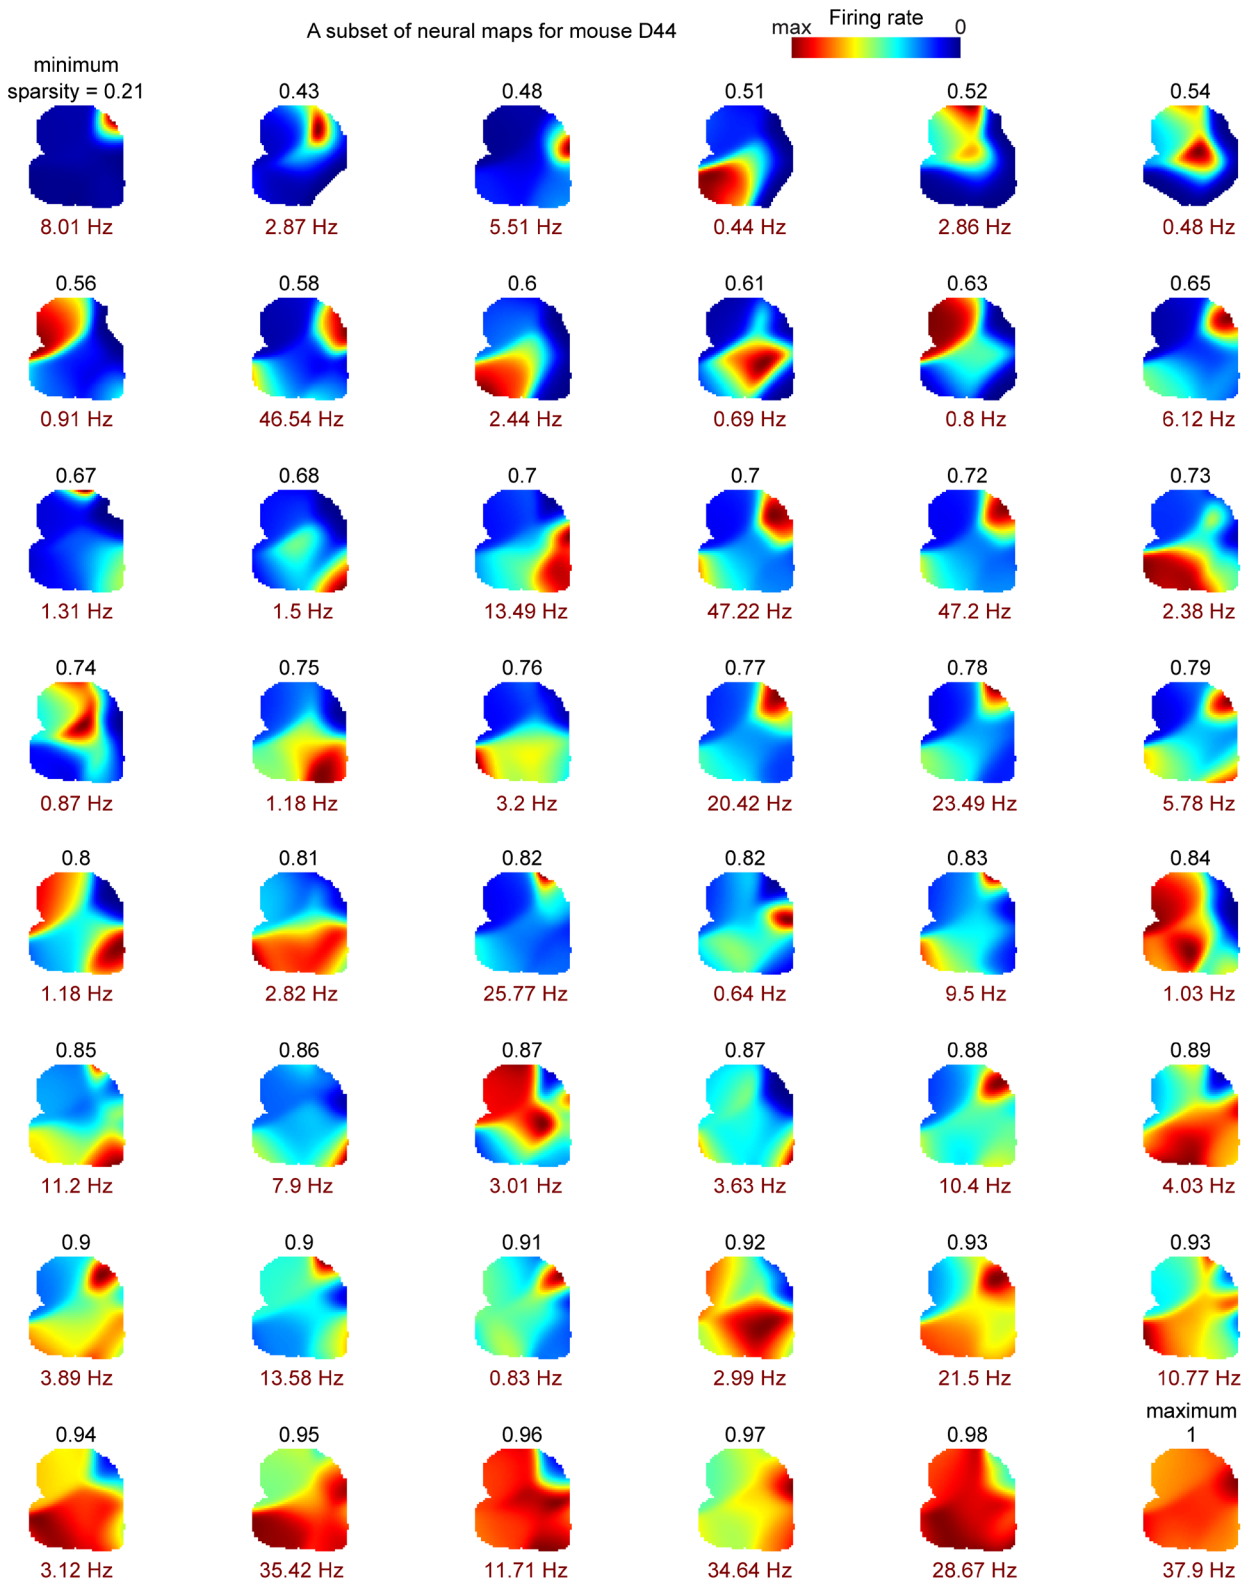

Extended Data Figure 5 | Neural activity maps for 48 neurons from one mouse. To illustrate varying sparsity, neurons shown have values equally spaced along the full distribution of sparsity values for the given mouse. Black text is sparsity, red text is each map's maximum firing rate.

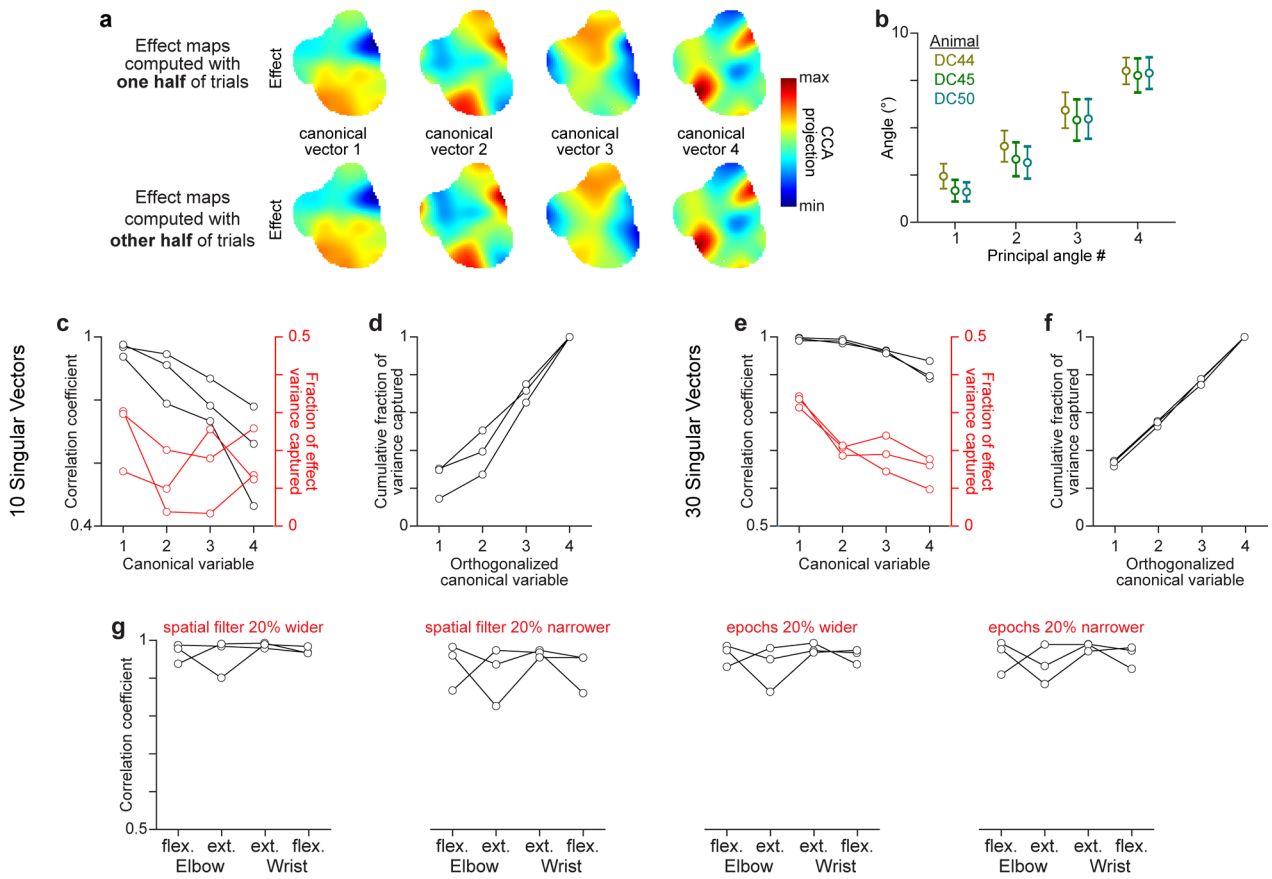

Extended Data Figure 6 | Finding a CFA activity subspace that aligns with CFA influence. **a**, Canonical variables for inactivation effect maps resulting from CCA with neural activity maps, using effect maps computed with separate halves of trials. **b**, Mean ± standard deviation for principal angles between subspaces spanned by the neural canonical vectors from CCA that used effect maps computed with separate, randomly-chosen halves of trials, over 300 iterations. **c**,**d** Correlation coefficient (black) and fractional inactivation effect variance captured (red) for canonical variables (**c**), and cumulative fraction of effect variance captured by canonical variables after orthogonalizing their corresponding vectors (**d**) when SVD is used to reduce neural activity dimensionality to 10 instead of 20 before applying CCA. Each set of connected dots in **c-g** is from one animal. **e**,**f**, Same as **c**,**d**, but when SVD is used to reduce neural activity dimensionality to 30. **g**, Correlation coefficients from CCA aligning neural activity maps and inactivation effect maps for individual muscles, calculated after changing key inactivation map parameters (red text).

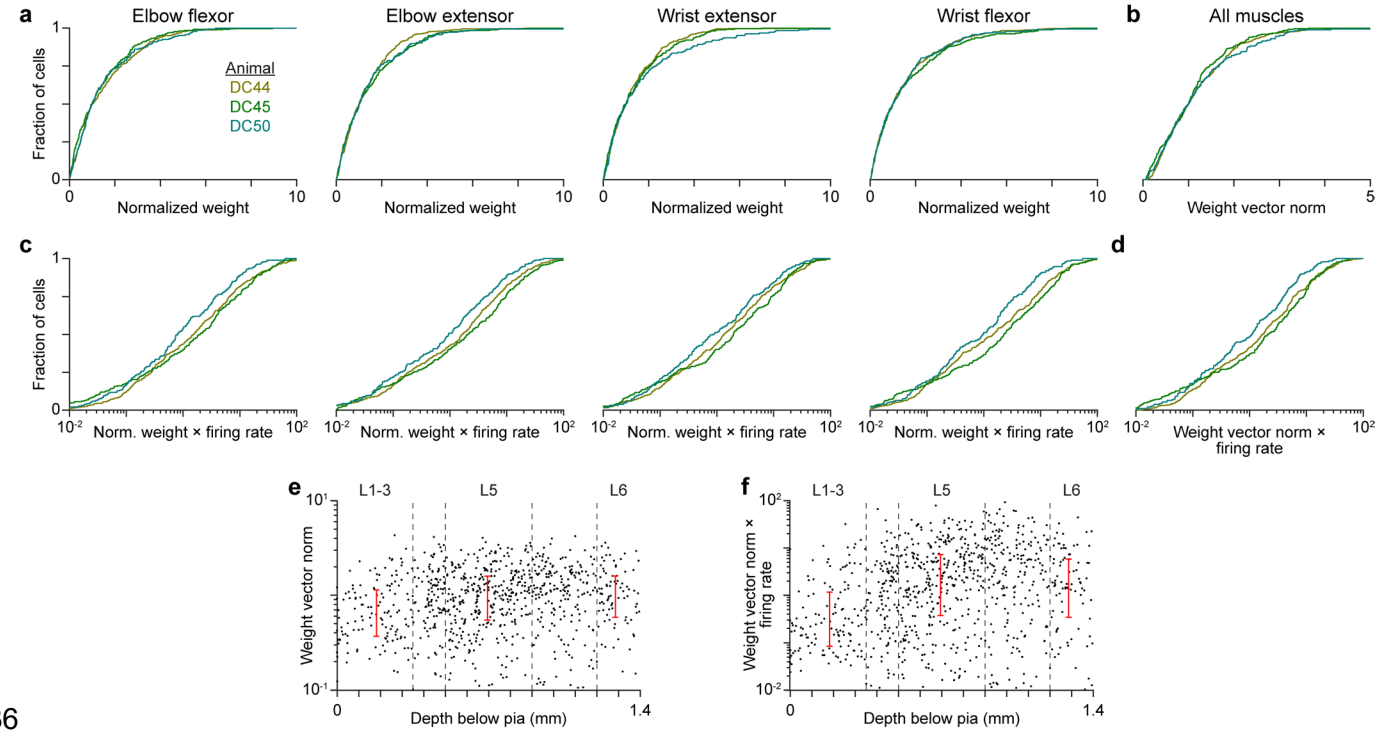

Extended Data Figure 7 | Influence subspace contributions across neurons. a, Cumulative distributions of weights for the activity of individual neurons in canonical vectors defined by CCA using the inactivation effect maps of individual muscles. Weights were normalized so the median value for each mouse is 1 to enable comparison across mice, as different numbers of neurons were recorded in each mouse. **b**, Cumulative distribution of norms of the four-element vectors composed of each given neuron's canonical vector weight from alignment with the inactivation effect maps of individual muscles. **c**,**d** Same as **a**,**b**, but weights (**c**) and weight vector norms (**d**) are scaled by each neuron's mean firing rate during climbing. **e**, Scatterplot of weight vector norms as in **b** versus the depth below pia assigned to the waveform centroid of each neuron. Red whisker plots span from the first to the third quartile of neurons assigned to each laminar group (those within dotted boundaries). Data from three mice are combined. **f**, Same as **e**, but weight vector norms are scaled by each neuron's mean firing rate during climbing.

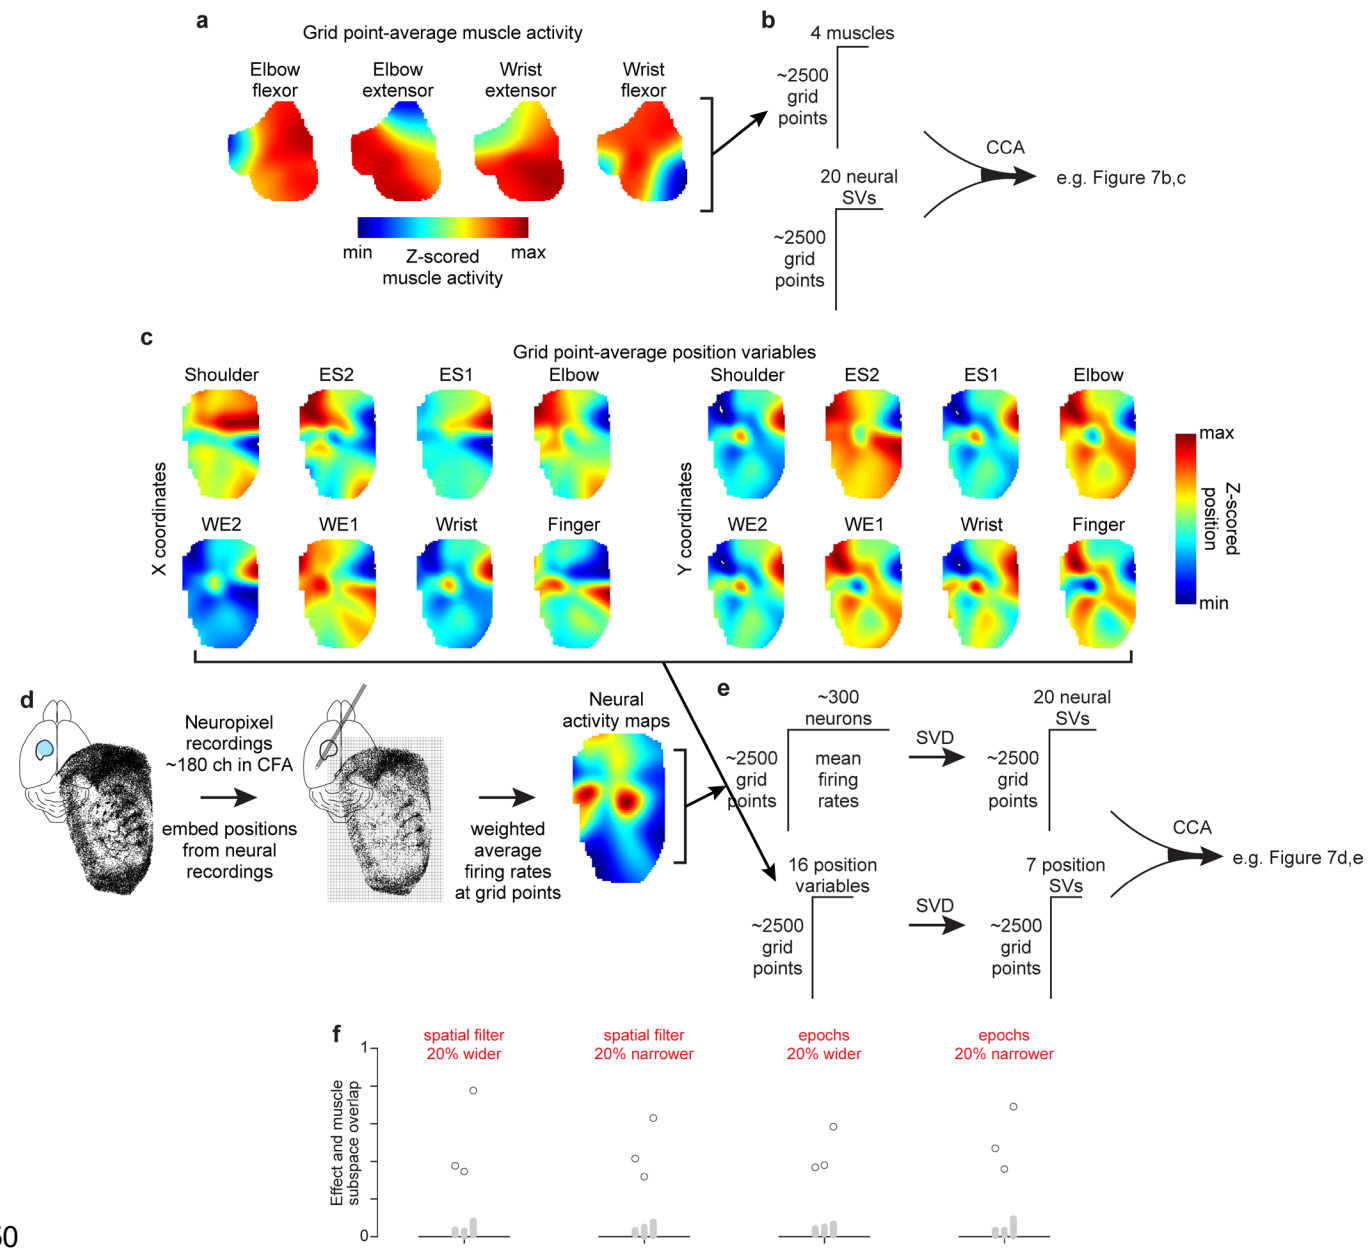

Extended Data Figure 8 | Extracting neural activity subspaces aligned with muscle activity and limb kinematics. a, Grid point-averaged muscle activity maps for an example animal. **b**, Grid point-averaged muscle activity maps are used to create a grid points by muscles matrix that, together with a grid points by 20 neural singular vectors matrix, serve as inputs to CCA. CCA then yields the canonical variables illustrated in Figure 7b,c. **c**, Grid point-averaged position maps for the eight tracked sites along the right forelimb for an example animal. **d**, 50 ms segments of the site positions tracked during neural recording sessions are embedded into the 2D limb orientation state map, and the neural activity corresponding to the embedded vectors is used to calculate grid point-averaged neural activity maps. **e**, The maps in **c** can be collapsed into a grid points by position variables matrix which is then dimensionally reduced to a grid points by 7 position singular vectors matrix using SVD. This matrix and the grid points by 20 neural singular vectors

matrix obtained from the neural activity maps serve as input to CCA. CCA then identifies the canonical variables shown in Figure 7d,e. **f**, Overlap between the influence and muscle activity subspaces (black circles) compared to 300 estimates of the overlap expected by chance (gray dots) for each animal calculated after slight changes to the parameters used for map calculations (red text).

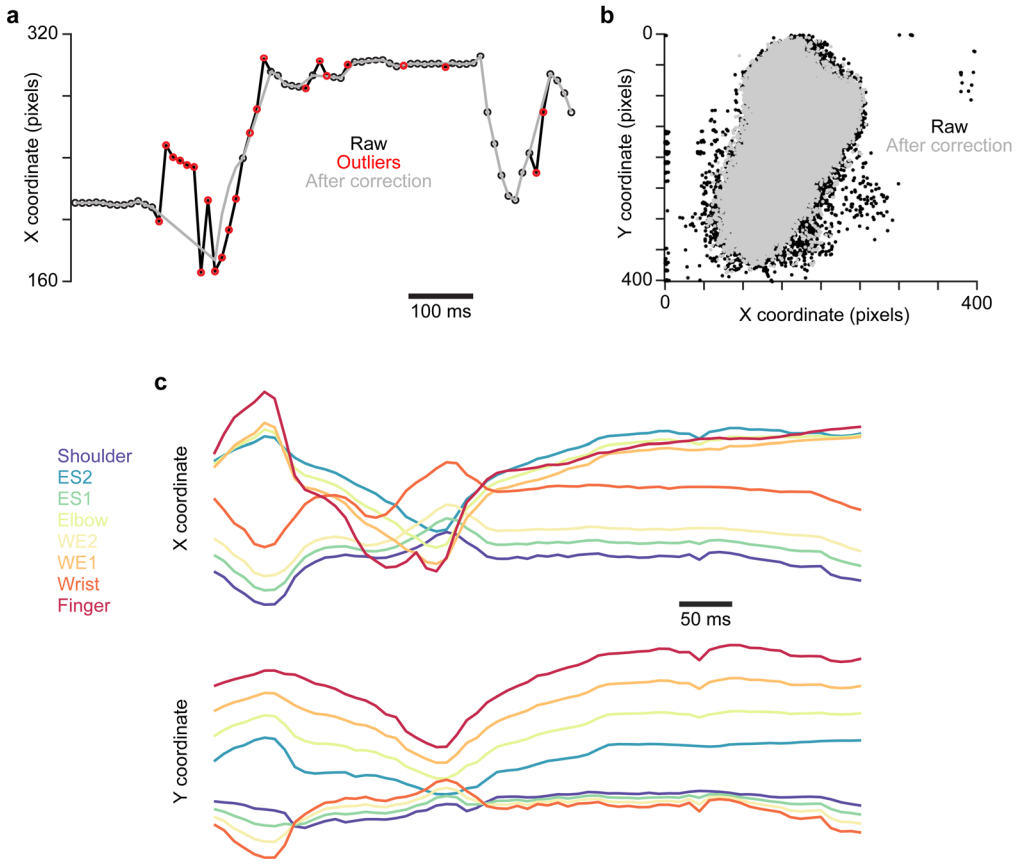

Extended Data Figure 9 | Post-processing of video-based limb tracking. a, Example time series for the point on the finger, illustrating post-hoc processing to correct outliers. Raw is the time series output by DeepLabCut before correction. **b**, Scatter plot of X and Y coordinates of the finger site tracked in **a**, showing fewer outliers after correction. **c**, Example X and Y coordinate traces for site positions (defined in Fig. 4) after post-hoc processing.

Supplementary Movie.mp4

Supplementary Movie | Head-fixed climbing. An example of a mouse climbing in our paradigm. The bird's-eye view towards the end makes apparent the stagger of the handholds on the mouse's right; this stagger itself is ever changing and thus unpredictable.

References

1 Darling, W. G., Pizzimenti, M. A. & Morecraft, R. J. Functional recovery following motor cortex lesions in non-human primates: experimental implications for human stroke patients. *J Integr Neurosci* **10**, 353-384 (2011).

2 Lemon, R. N. Descending pathways in motor control. *Annu Rev Neurosci* **31**, 195-218 (2008).

3 Kuypers, H. G. J. M. in *Comprehensive Physiology* 597-666 (1981).

4 Kaas, J. H. The evolution of the complex sensory and motor systems of the human brain. *Brain Res Bull* **75**, 384-390 (2008).

5 Krubitzer, L. The magnificent compromise: cortical field evolution in mammals. *Neuron* **56**, 201-208 (2007).

6 Fulton, J. F. & Kennard, M. A. A Study of Flaccid and Spastic Paralyses Produced by Lesions of the Cerebral Cortex in Primates. *Res Publ Assoc Res N* **13**, 157-209 (1934).

7 Lawrence, D. G. & Kuypers, H. G. The functional organization of the motor system in the monkey. I. The effects of bilateral pyramidal lesions. *Brain* **91**, 1-14 (1968).

8 Lawrence, D. G. & Hopkins, D. A. The development of motor control in the rhesus monkey: evidence concerning the role of corticomotoneuronal connections. *Brain* **99**, 235-254 (1976).

9 Leyton, A. S. F. & Sherrington, C. S. *Observations on the excitable cortex of the chimpanzee, orang-utan, and gorilla*. (Griffin, 1917).

10 Travis, A. M. Neurological deficiencies after ablation of the precentral motor area in Macaca mulatta. *Brain* **78**, 155-173 (1955).

11 Passingham, R. E., Perry, V. H. & Wilkinson, F. The long-term effects of removal of sensorimotor cortex in infant and adult rhesus monkeys. *Brain* **106 (Pt 3)**, 675-705 (1983).

12 Drew, T., Jiang, W., Kably, B. & Lavoie, S. Role of the motor cortex in the control of visually triggered gait modifications. *Can J Physiol Pharmacol* **74**, 426-442 (1996).

13 Warren, R. A. *et al.* A rapid whisker-based decision underlying skilled locomotion in mice. *Elife* **10** (2021).

14 Bollu, T. *et al.* Cortex-dependent corrections as the tongue reaches for and misses targets. *Nature* **594**, 82-87 (2021).

15 Farr, T. D. & Whishaw, I. Q. Quantitative and qualitative impairments in skilled reaching in the mouse (Mus musculus) after a focal motor cortex stroke. *Stroke* **33**, 1869-1875 (2002).

16 Kawai, R. *et al.* Motor cortex is required for learning but not for executing a motor skill. *Neuron* **86**, 800-812 (2015).

17 Hwang, E. J. *et al.* Disengagement of motor cortex from movement control during long-term learning. *Science Advances* **5** (2019).

18 Glees, P. & Cole, J. Recovery of Skilled Motor Functions after Small Repeated Lesions of Motor Cortex in Macaque. *Journal of Neurophysiology* **13**, 137-148 (1950).

19 Balasubramanian, K. *et al.* Propagating Motor Cortical Dynamics Facilitate Movement Initiation. *Neuron* **106**, 526-536 e524 (2020).

20 Hepp, R., Trouche, E. & Wiesendanger, M. Effects of unilateral and bilateral pyramidotomy on a conditioned rapid precision grip in monkeys (Macaca fascicularis). *Exp Brain Res* **21**, 519-527 (1974).

21 Guo, J. Z. *et al.* Cortex commands the performance of skilled movement. *Elife* **4**, e10774 (2015).

22 Miri, A. *et al.* Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex. *Neuron* **95**, 683-696 e611 (2017).
23 Warriner, C. L., Fageiry, S. K., Carmona, L. M. & Miri, A. Towards Cell and Subtype Resolved Functional Organization: Mouse as a Model for the Cortical Control of Movement. *Neuroscience* **450**, 151-160 (2020).
24 Shmuelof, L. & Krakauer, J. W. Are we ready for a natural history of motor learning? *Neuron* **72**, 469-476 (2011).
25 Yuste, R., MacLean, J. N., Smith, J. & Lansner, A. The cortex as a central pattern generator. *Nat Rev Neurosci* **6**, 477-483 (2005).
26 Archambault, P. S., Ferrari-Toniolo, S., Caminiti, R. & Battaglia-Mayer, A. Visually-guided correction of hand reaching movements: The neurophysiological bases in the cerebral cortex. *Vision Res* **110**, 244-256 (2015).
27 Evarts, E. V. & Tanji, J. Reflex and intended responses in motor cortex pyramidal tract neurons of monkey. *J Neurophysiol* **39**, 1069-1080 (1976).
28 Pruszynski, J. A. *et al.* Primary motor cortex underlies multi-joint integration for fast feedback control. *Nature* **478**, 387-390 (2011).
29 Bollu, T. *et al.* Motor cortical inactivation impairs corrective submovements in mice performing a hold-still center-out reach task. *J Neurophysiol* (2024).
30 Penfield, W. & Boldrey, E. Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain* **60**, 389-443 (1937).
31 Tennant, K. A. *et al.* The organization of the forelimb representation of the C57BL/6 mouse motor cortex as defined by intracortical microstimulation and cytoarchitecture. *Cereb Cortex* **21**, 865-876 (2011).
32 Armstrong, D. M. & Drew, T. Electromyographic responses evoked in muscles of the forelimb by intracortical stimulation in the cat. *J Physiol* **367**, 309-326 (1985).
33 Cheney, P. D., Fetz, E. E. & Palmer, S. S. Patterns of facilitation and suppression of antagonist forelimb muscles from motor cortex sites in the awake monkey. *J Neurophysiol* **53**, 805-820 (1985).
34 Park, M. C., Belhaj-Saif, A. & Cheney, P. D. Properties of primary motor cortex output to forelimb muscles in rhesus macaques. *J Neurophysiol* **92**, 2968-2984 (2004).
35 Griffin, D. M. & Strick, P. L. The motor cortex uses active suppression to sculpt movement. *Sci Adv* **6** (2020).
36 Peters, A. J., Chen, S. X. & Komiyama, T. Emergence of reproducible spatiotemporal activity during motor learning. *Nature* **510**, 263-267 (2014).
37 Ebbesen, C. L. & Brecht, M. Motor cortex - to act or not to act? *Nat Rev Neurosci* **18**, 694-705 (2017).
38 Oby, E. R., Ethier, C. & Miller, L. E. Movement representation in the primary motor cortex and its contribution to generalizable EMG predictions. *J Neurophysiol* **109**, 666-678 (2013).
39 Evarts, E. V. Relation of pyramidal tract activity to force exerted during voluntary movement. *J Neurophysiol* **31**, 14-27 (1968).
40 Armstrong, D. M. & Drew, T. Locomotor-related neuronal discharges in cat motor cortex compared with peripheral receptive fields and evoked movements. *J Physiol* **346**, 497-517 (1984).
41 Kalaska, J. F. From intention to action: motor cortex and the control of reaching movements. *Adv Exp Med Biol* **629**, 139-178 (2009).

42 Georgopoulos, A. P., Kalaska, J. F., Caminiti, R. & Massey, J. T. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* **2**, 1527--1537 (1982).
43 Paninski, L., Fellows, M. R., Hatsopoulos, N. G. & Donoghue, J. P. Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *Journal of Neurophysiology* **91**, 515-532 (2004).
44 Hatsopoulos, N. G. & Suminski, A. J. Sensing with the motor cortex. *Neuron* **72**, 477-487 (2011).
45 Adams, R. A., Shipp, S. & Friston, K. J. Predictions not commands: active inference in the motor system. *Brain Struct Funct* **218**, 611-643 (2013).
46 Lillicrap, T. P. & Scott, S. H. Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. *Neuron* **77**, 168-179 (2013).
47 Todorov, E. Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat Neurosci* **3**, 391-398 (2000).
48 Miri, A., Azim, E. & Jessell, T. M. Edging toward entelechy in motor control. *Neuron* **80**, 827-834 (2013).
49 Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A. Neural Manifolds for the Control of Movement. *Neuron* **94**, 978-984 (2017).
50 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**, 440-448 (2014).
51 Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232-236 (2017).
52 Drew, T., Prentice, S. & Schepens, B. Cortical and brainstem control of locomotion. *Prog Brain Res* **143**, 251-261 (2004).
53 Akay, T., Acharya, H. J., Fouad, K. & Pearson, K. G. Behavioral and electromyographic characterization of mice lacking EphA4 receptors. *Journal of Neurophysiology* **96**, 642-651 (2006).
54 Pearson, K. G., Acharya, H. & Fouad, K. A new electrode configuration for recording electromyographic activity in behaving mice. *J Neurosci Methods* **148**, 36-42 (2005).
55 Mathis, A. *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat Neurosci* **21**, 1281-1289 (2018).
56 Kargo, W. J. & Nitz, D. A. Improvements in the signal-to-noise ratio of motor cortex cells distinguish early versus late phases of motor skill learning. *J Neurosci* **24**, 5560-5569 (2004).
57 Richman, J. S. & Moorman, J. R. Physiological time-series analysis using approximate entropy and sample entropy. *Am J Physiol Heart Circ Physiol* **278**, H2039-2049 (2000).
58 Alcaraz, R. & Rieta, J. J. A novel application of sample entropy to the electrocardiogram of atrial fibrillation. *Nonlinear Anal-Real* **11**, 1026-1035 (2010).
59 Liang, Z. *et al.* EEG entropy measures in anesthesia. *Front Comput Neurosci* **9**, 16 (2015).
60 Nezafati, M., Temmar, H. & Keilholz, S. D. Functional MRI Signal Complexity Analysis Using Sample Entropy. *Front Neurosci* **14**, 700 (2020).
61 Guo, Z. V. *et al.* Flow of cortical activity underlying a tactile decision in mice. *Neuron* **81**, 179-194 (2014).
62 Becht, E. *et al.* Dimensionality reduction for visualizing single-cell data using UMAP. *Nat Biotechnol* **37**, 38–44 (2019).

63 Storey, J. D. A direct approach to false discovery rates. *Journal of the Royal Statistical Society Series B-Statistical Methodology* **64**, 479-498 (2002).
64 Asante, C. O. *et al.* Cortical control of adaptive locomotion in wild-type mice and mutant mice lacking the ephrin-Eph effector protein alpha2-chimaerin. *J Neurophysiol* **104**, 3189-3202 (2010).
65 Hoffman, D. S. & Strick, P. L. Effects of a primary motor cortex lesion on step-tracking movements of the wrist. *J Neurophysiol* **73**, 891-895 (1995).
66 Martin, J. H. & Ghez, C. Differential impairments in reaching and grasping produced by local inactivation within the forelimb representation of the motor cortex in the cat. *Exp Brain Res* **94**, 429-443 (1993).
67 Liu, G. *et al.* The Muscle Activation Differences in Post-Stroke Upper Limb Flexion Synergy Based on Spinal Cord Segments: A Preliminary Proof-of-Concept Study. *Front Neurol* **12**, 598554 (2021).
68 Roux, F. E., Niare, M., Charni, S., Giussani, C. & Durand, J. B. Functional architecture of the motor homunculus detected by electrostimulation. *J Physiol* **598**, 5487-5504 (2020).
69 Sanei, K. & Keir, P. J. Independence and control of the fingers depend on direction and contraction mode. *Hum Mov Sci* **32**, 457-471 (2013).
70 Capaday, C., Devanne, H., Bertrand, L. & Lavoie, B. A. Intracortical connections between motor cortical zones controlling antagonistic muscles in the cat: a combined anatomical and physiological study. *Exp Brain Res* **120**, 223-232 (1998).
71 Drew, T. Motor cortical activity during voluntary gait modifications in the cat. I. Cells related to the forelimbs. *J Neurophysiol* **70**, 179-199 (1993).
72 Crutcher, M. D. & Alexander, G. E. Movement-Related Neuronal-Activity Selectively Coding Either Direction or Muscle Pattern in 3 Motor Areas of the Monkey. *Journal of Neurophysiology* **64**, 151-163 (1990).
73 Kakei, S., Hoffman, D. S. & Strick, P. L. Muscle and movement representations in the primary motor cortex. *Science* **285**, 2136-2139 (1999).
74 Georgopoulos, A. P., Schwartz, A. B. & Kettner, R. E. Neuronal Population Coding of Movement Direction. *Science* **233**, 1416--1419 (1986).
75 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**, 387-400 (2010).
76 Jung, M. W., Wiener, S. I. & McNaughton, B. L. Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *J Neurosci* **14**, 7347-7356 (1994).
77 Marshall, N. J. *et al.* Flexible neural control of motor units. *Nat Neurosci* **25**, 1492-1504 (2022).
78 Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A. & Miller, L. E. Long-term stability of cortical population dynamics underlying consistent behavior. *Nat Neurosci* **23**, 260-270 (2020).
79 Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A neural network that finds a naturalistic solution for the production of muscle activity. *Nat Neurosci* **18**, 1025-1033 (2015).
80 Raghu, M., Gilmer, J., Yosinski, J. & Sohl-Dickstein, J. SVCCA: Singular Vector Canonical Correlation Analysis for Deep Learning Dynamics and Interpretability. *Adv Neur In* **30** (2017).

81 Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: a dynamical systems perspective. *Annu Rev Neurosci* **36**, 337-359 (2013).
82 Graziano, M. The organization of behavioral repertoire in motor cortex. *Annu Rev Neurosci* **29**, 105-134 (2006).
83 Krubitzer, L., Campi, K. L. & Cooke, D. F. All rodents are not the same: a modern synthesis of cortical organization. *Brain Behav Evol* **78**, 51-93 (2011).
84 Ramanathan, D., Conner, J. M. & Tuszynski, M. H. A form of motor cortical plasticity that correlates with recovery of function after brain injury. *P Natl Acad Sci USA* **103**, 11370-11375 (2006).
85 Brown, A. R. & Teskey, G. C. Motor cortex is functionally organized as a set of spatially distinct representations for complex movements. *J Neurosci* **34**, 13574-13585 (2014).
86 Harrison, T. C., Ayling, O. G. & Murphy, T. H. Distinct cortical circuit mechanisms for complex forelimb movement and motor map topography. *Neuron* **74**, 397-409 (2012).
87 Resta, F. *et al.* Large-scale all-optical dissection of motor cortex connectivity shows a segregated organization of mouse forelimb representations. *Cell Rep* **41**, 111627 (2022).
88 Yang, W. Z., Kanodia, H. & Arber, S. Structural and functional map for forelimb movement phases between cortex and medulla. *Cell* **186**, 162-+ (2023).
89 Wang, X. *et al.* Deconstruction of Corticospinal Circuits for Goal-Directed Motor Skills. *Cell* **171**, 440-455 e414 (2017).
90 Capaday, C., Ethier, C., Van Vreeswijk, C. & Darling, W. G. On the functional organization and operational principles of the motor cortex. *Front Neural Circuits* **7**, 66 (2013).
91 Tower, S. S. Pyramidal lesion in the monkey. *Brain* **63**, 36-90 (1940).
92 Schieber, M. H. How Might the Motor Cortex Individuate Movements. *Trends Neurosci* **13**, 440-445 (1990).
93 Mathis, M. W. & Schneider, S. Motor control: Neural correlates of optimal feedback control theory. *Curr Biol* **31**, R356-R358 (2021).
94 Scott, S. H. Optimal feedback control and the neural basis of volitional motor control. *Nat Rev Neurosci* **5**, 532-546 (2004).
95 Todorov, E. & Jordan, M. I. Optimal feedback control as a theory of motor coordination. *Nat Neurosci* **5**, 1226-1235 (2002).
96 Rastogi, A. *et al.* The Neural Representation of Force across Grasp Types in Motor Cortex of Humans with Tetraplegia. *Eneuro* **8** (2021).
97 Ebbesen, C. L. *et al.* More than Just a "Motor": Recent Surprises from the Frontal Cortex. *J Neurosci* **38**, 9402-9413 (2018).
98 Omrani, M., Kaufman, M. T., Hatsopoulos, N. G. & Cheney, P. D. Perspectives on classical controversies about the motor cortex. *J Neurophysiol* **118**, 1828-1848 (2017).
99 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**, 13239 (2016).
100 Guo, Z. V. *et al.* Procedures for behavioral experiments in head-fixed mice. *PLoS One* **9**, e88678 (2014).
101 Pachitariu, M., Sridhar, S. & Stringer, C. Solving the spike sorting problem with Kilosort. *bioRxiv*, 2023.2001.2007.523036 (2023).
102 Mao, T. *et al.* Long-range neuronal circuits underlying the interaction between sensory and motor cortex. *Neuron* **72**, 111-123 (2011).

103 Warriner, C. L., Fageiry, S., Saxena, S., Costa, R. M. & Miri, A. Motor cortical influence relies on task-specific activity covariation. *Cell Rep* **40**, 111427 (2022).
104 Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51-56 (2012).
105 Knyazev, A. V. & Argentati, M. E. Principal angles between subspaces in an A-based scalar product: Algorithms and perturbation estimates. *Siam J Sci Comput* **23**, 2008-2040 (2002).