

A distinct neural activity subspace for direct motor cortical influence on muscles

Natalie Koh¹, Zhengyu Ma¹, Abhishek Sarup, Amy C. Kristl, Mark Agrios, Margaret Young², and Andrew Miri*

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

¹These authors contributed equally.

²Present 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)

Among the patterns of motor cortical activity, those that directly drive muscles remain unresolved. Lesion studies have led to the hypothesis that motor cortex functions primarily to improve movement efficacy by enabling patterns of muscle activity that the rest of the motor system cannot achieve¹⁻³. Yet such studies weakly constrain when motor cortical output influences muscle activity in the unperturbed state. Analysis of motor cortical activity has consistently found, and imputed functional significance upon, signals that correlate with limb muscle activity⁴⁻⁶ or kinematics⁷⁻⁹. But a selective role in driving certain muscle activity patterns might rely on signals related only to those patterns and not others. Here we quantified the direct influence of forelimb motor cortex on muscle activity throughout a naturalistic climbing behavior, finding that this influence is selective for, and highly dependent upon, muscle activity states. We used multielectrode array recordings to identify linear combinations (components) of motor cortical activity patterns that covary with this influence. We find that these components differ substantially from those that covary with muscle activity or kinematics. Our results reveal a direct motor cortical influence on muscles that is selective within a motor behavior and reliant on a previously undescribed neural activity subspace.

One primary barrier to resolving how motor cortex drives movement is a persistent ambiguity regarding the components of muscle activity directly influenced by motor cortical output. Lesion studies have revealed a circumscribed set of motor behaviors that require motor cortex for their execution, like individuated finger movements^{1,10-12}. Yet at least in non-human mammals, broad swaths of the motor behavioral repertoire persist without normal motor cortical influence, including walking, jumping, climbing, goal-oriented manual tasks, and even highly precise, learned limb movement sequences¹³⁻¹⁸. In some cases, though, especially when dexterity is challenged, these movements are slower, less agile, and less effective^{1,15,19,20}. This has led to the hypothesis that motor cortex confers an added level of muscle control that improves movement efficacy and task learning by activating combinations of muscles in temporal patterns that cannot be achieved by other motor system structures^{2,3,21}. However, it remains unclear how this influence manifests during movement²²⁻²⁵. The involvement of motor cortex in motor learning^{18,26,27} and

movement preparation/initiation²⁸⁻³⁰ complicates lesion interpretation vis-à-vis movement execution, as disturbance to processes on which execution depends can impede execution itself. Recent results indicate that direct motor cortical influence can be selective for the performance of certain tasks³¹, but it remains unclear whether during such tasks, motor cortical influence is responsible for the totality of muscle activation or is selective to particular muscle activity patterns, and whether it informs muscle activity pattern, or is merely permissive³².

This fundamental ambiguity about when muscle activity is influenced by motor cortical output has in turn stymied characterization of the output activity components mediating this influence. Studies assessing the effect of motor cortical activity on muscles have focused on components of motor cortical activity that correlate with or explain the variance of muscle activity⁴⁻⁶, or movement parameters like joint angles or reach direction (kinematics)⁷⁻⁹. These correlations and explained variance can be substantial, consistent with the functional relevance of the activity components. However, if motor cortical output is primarily responsible for select muscle activity patterns, which may reflect only a fraction of muscle activity during movement, we might expect that the components of motor cortical output that drive muscle activity may only reflect those select patterns and not muscle activity in total. Moreover, motor cortical activity that covaries with muscle activity or kinematics in total may be a consequence of monitoring or predicting body state³³, perhaps to subserve aspects of motor control apart from directly driving muscle activation³⁴. In line with this, muscle activity can be decoded from motor cortical activity during movements where this activity does not directly drive muscles³¹. Thus, the components of motor cortical activity responsible for its influence may differ from those to which functional significance has previously been imputed (Fig. 1a)^{8,35-39}.

We sought to better resolve the components of motor cortical firing that drive muscles and test how similar they are to those that covary with muscle activity or movement kinematics. To address this in an ethologically relevant context in which a diversity of limb movements is performed, we developed a paradigm in which head-fixed mice perform a naturalistic climbing behavior. We combined electromyography, optogenetic silencing, and dimensionality reduction to assess the influence of motor cortical output across the limb movements mice express during climbing. Contrary to classical approaches, rapid optogenetic silencing of motor cortical output enabled us to quantify the direct effects of this output on muscle activity. We found that this influence (a) is selective for specific muscle activity states during climbing, (b) when present, has a magnitude highly dependent upon muscle activity state, and (c) exhibits selectivity and muscle state-dependence that themselves vary markedly across muscles. We then used multielectrode arrays (Neuropixels⁴⁰) to record motor cortical firing patterns during climbing and identify components that covary with influence on muscles. We found a substantial difference between these components and those that covary with muscle activity and movement kinematics. Our results indicate a selective motor cortical influence on muscles that is mediated by a neural activity subspace distinct from those traditionally focused upon.

An ethological movement paradigm

In our naturalistic climbing paradigm, head-fixed mice climb with all four limbs across a series of handholds that extend radially from a wheel, thereby rotating the wheel (Fig. 1b,c; Extended Data Fig. 1a,b; 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. 1c-f). 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. Thus, mice do not rely solely on a stereotyped movement strategy, performing adaptive forelimb movements expected to involve motor cortical influence⁴¹. Water-restricted mice climb intermittently in bouts throughout hour-long daily sessions, earning water rewards when they stop climbing based on the distance of the previous bout. The likelihood of reward and volume dispensed is varied adaptively within sessions to encourage longer distance bouts and 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^{42,43} (Extended Data Fig. 1g), and limb orientation is charted using video-based tracking⁴⁴ (Extended Data Fig. 2).

Despite the reward scheme, climbing in this paradigm does not show signs of learning or continued improvement across sessions. After mice are acclimated to head fixation (2 sessions) and taught the pairing between climbing and reward (1-3 sessions), we observe little change in climbing performance, including the time spent climbing (Fig. 1f), the velocity of climbing (Fig. 1g), and the distance of climbing bouts (Fig. 1h). Moreover, the form of muscle activity changes only minimally across sessions. We computed principal angles for the muscle activity matrices (4 muscles by T timepoints), comparing each of the first 20 daily sessions to the 20th session (Fig. 1i,j). The first principal angle for this comparison was generally low, averaging less than 2°, and there was little indication of gradual change across sessions. Thus, perhaps because climbing is a critical element of the mouse's natural movement repertoire, this paradigm obviates 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. This paradigm can facilitate motor system interrogation during a kinematically-diverse, nonstereotyped, and ethologically-relevant motor behavior. Though motor cortical function has been studied extensively in the context of learned motor tasks, our paradigm allows us to assay function during a task at which mice are naturally good.

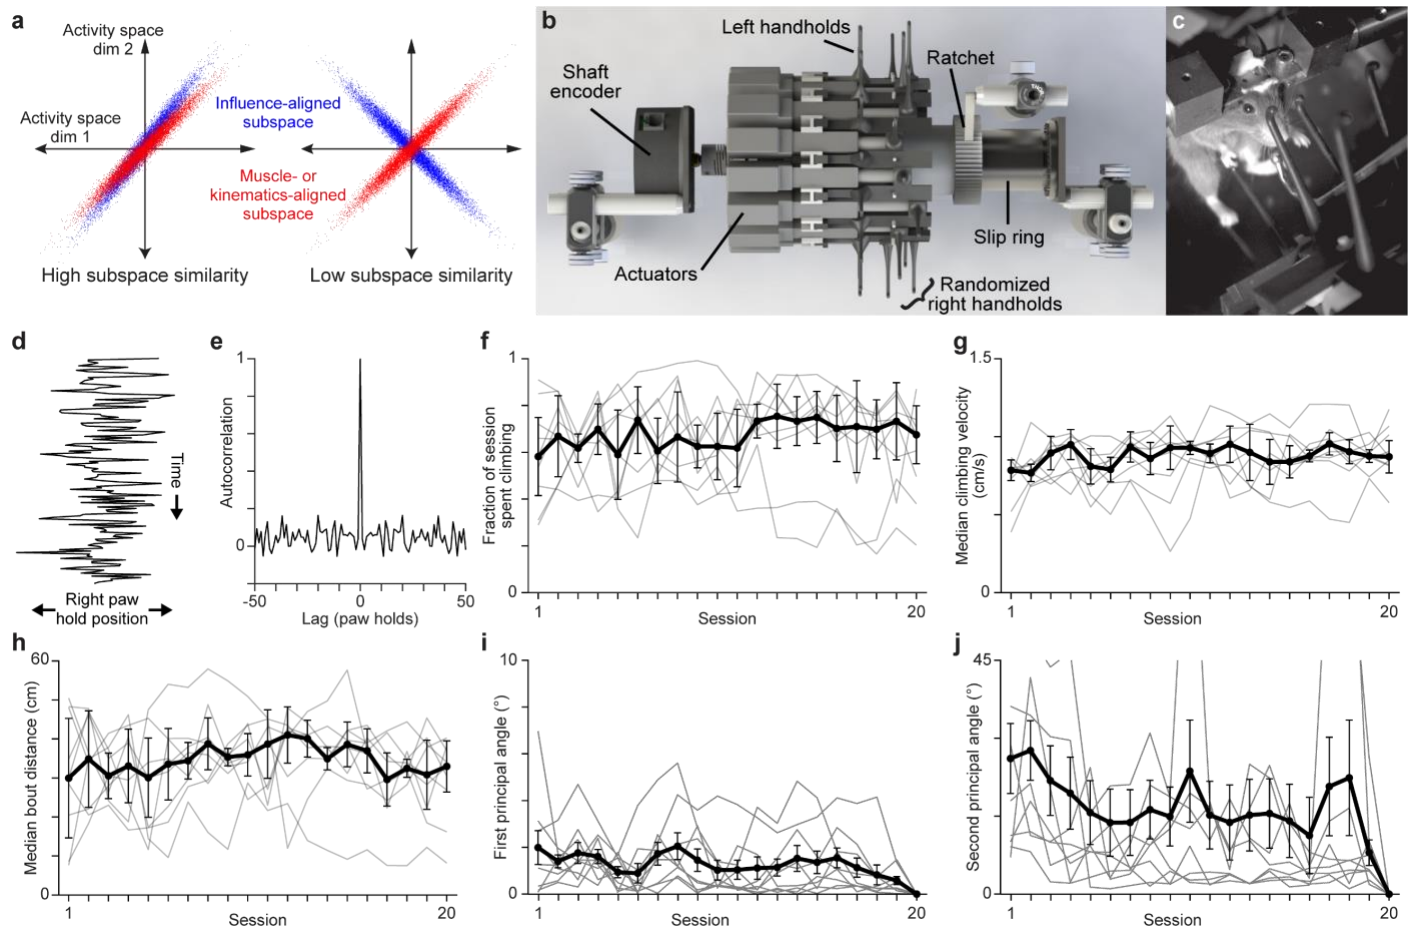


Figure 1 | Head-fixed climbing paradigm. **a**, Schematic of different scenarios for the overlap between neural activity subspaces discussed herein. **b**, Bird's-eye view of wheel apparatus for climbing. A shaft encoder measures the wheel's angular position. Actuators randomize the position of each 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. **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-j** are for individual animals. Session 1 indicates the first session after mice had learned the pairing between climbing and reward (1-3 sessions) and 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.

Direct motor cortical influence during climbing

To assess the direct influence of motor cortical output on muscle activity in this paradigm, and how it might vary with muscle activity state, we sporadically and briefly inactivated the left caudal forelimb area (CFA, forelimb M1+S1) at random as mice climbed. 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/mm²; 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^{31,45}. 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). Equivalent events without blue light were also notated in recordings to serve as control trials. Random trial timing ensured broad coverage of muscle activity states each mouse expressed during climbing. Inactivation and control trial averages diverged ~10 ms after light onset, which reflects the shortest latency at which CFA influences muscles (Fig. 2a, Extended Data Fig. 3). Such direct influence has previously been observed in mice during a trained forelimb reaching task but is not seen during treadmill locomotion³¹.

We then quantified the direct influence of left CFA on right forelimb muscles in each mouse (n = 8) across the range of muscle activity states they expressed during climbing. For statistical power in assessing inactivation effects, we used nonlinear dimensionality reduction (UMAP⁴⁶) to cluster trials occurring at similar muscle activity states in a readily parsable two dimensional map. To represent muscle activity states, muscle activity traces surrounding each trial, together with their corresponding first derivatives, were divided into overlapping 50 ms epochs that began every 10 ms (Fig. 2b). For each epoch, the muscle activity and derivative trace segments were concatenated into a single vector. We applied UMAP to map these vectors into two dimensions while preserving the proximity of nearby vectors (Fig. 2c). On the resulting maps, points from successive epochs form trajectories that reflect the sequence of states during control and inactivation trials. Control and inactivation trajectories that are similar before trial/light onset often subsequently diverge (Fig. 2c). Epochs from both control and inactivation trials were broadly distributed across maps, indicating that large clusters of activity states arising from only one trial type do not occur (Fig. 2d). We did not observe separate map regions clearly corresponding to specific phases of climbing, like reaching or pulling, perhaps because much recorded muscle activity related not to the basic cyclic motion of the limb, but to adaptive adjustments superimposed upon it.

To quantify CFA influence on muscles across these muscle activity state maps, we first defined grids over them (Fig. 2d). Separately for each muscle, we then measured the inactivation effect locally around each grid point. For each grid point, we first 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 weighted the muscle activity from each trial by a Gaussian function of the Euclidean distance between the given grid point and the point reflecting the trial epoch just prior to when an inactivation effect could be observed (-40 to +10 ms from trial/light onset, “weight epoch”; Fig. 2d). The standard deviation of the Gaussian function was chosen so

that averages reflect muscle activity states in a local neighborhood around each grid point. For each muscle, we quantified inactivation effect size for each grid point using the difference between the rate of change in inactivation and control trial averages from 0 to 20 ms after trial/light onset (Fig. 2e). 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. 2f). By plotting these effect sizes at each grid point across the map, we produced an “inactivation effect map” for each muscle recorded in each individual mouse (Fig. 2g; Extended Data Fig. 4a).

These inactivation effect maps reveal that CFA’s direct influence on muscles varies during climbing. Maps for different muscles in a given mouse show wide variation in the magnitude and direction of inactivation effects across grid points (Fig. 2g; Extended Data Fig. 4a). We resampled from control trials to compute empirical null distributions for effect sizes at each grid point. Distributions of the effect sizes aggregated over grid points and mice, when compared to their empirical null counterparts, showed that the elbow flexor and wrist extensor were nearly always deactivated when affected, while effects on the elbow extensor were roughly symmetric around zero and generally smaller in magnitude (Fig. 2h). Effects on the wrist flexor lie in between, showing a bias toward deactivation. We also use empirical null distributions to compute p-values for each grid point’s effect size on each map (Fig. 2i; Extended Data Fig. 4a). Histograms of the resulting p-values aggregated over grid points and mice were skewed toward 0 to different extents across muscles (Fig. 2j), indicating that the fraction of muscle activity states influenced by CFA varied across muscles. We estimated from these p-value distributions the fraction of grid points at which CFA influenced muscles⁴⁷, finding that estimates differed widely and significantly across muscles ($p = 2 \times 10^{-4}$, Kruskal-Wallis test; Extended Data Fig. 4b). 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 results indicate that CFA’s direct influence is only present at certain muscle activity states, varies in magnitude when present, and this variation across states itself differs markedly across muscles.

Importantly, the relationship between effect size and the level of muscle activity at trial/light onset was complicated and highly nonlinear (median Pearson correlation across muscles and mice = 0.22; Fig. 2k; Extended Data Fig. 4c). This indicates that CFA influence is not proportional to muscle activity, so that motor cortical activity components that resemble muscle activity over time would not strongly covary with CFA influence. Thus during climbing, direct CFA influence is neither pervasive nor absent, but is instead selective and highly state-dependent. This complex state dependence also rules out a merely permissive, or noninformative, influence on muscles.

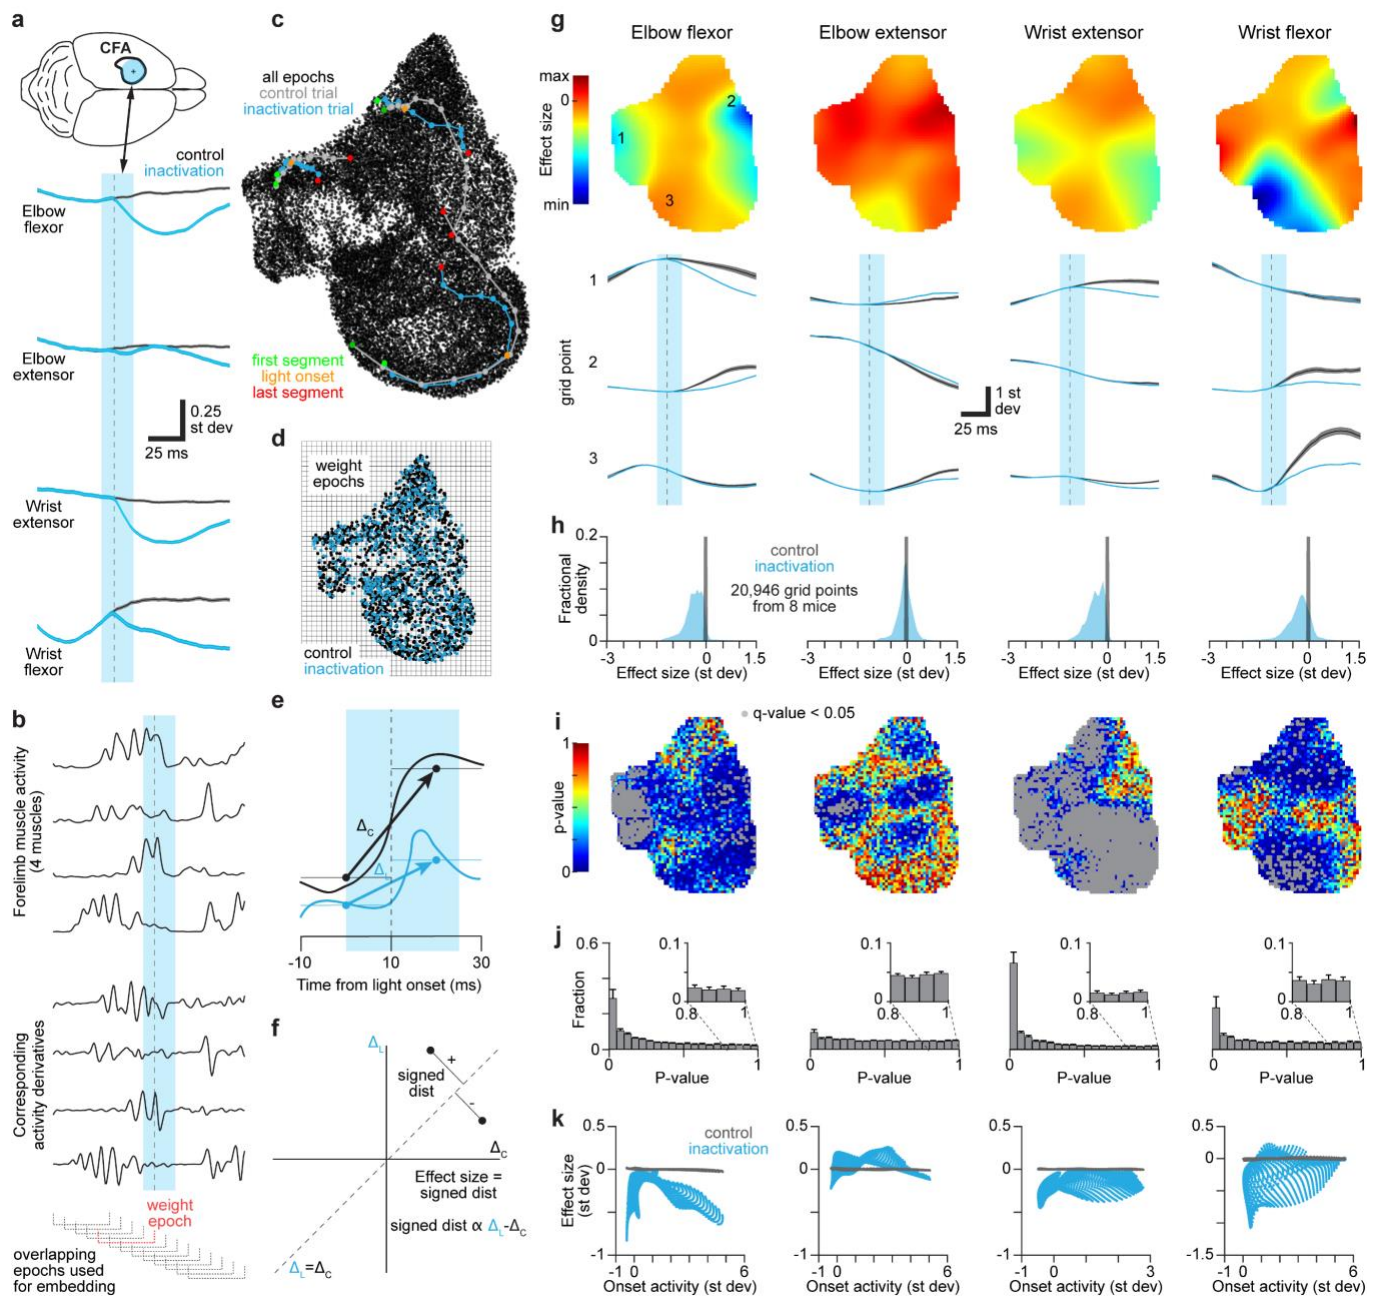


Figure 2 | Direct CFA influence across muscle activity states. **a**, Control (gray, 18,397 trials) and inactivation (cyan, n = 9,029 trials) trial averages (mean \pm sem) for four muscles across 8 mice. Vertical cyan bars in **a** and **b** indicate the 25 ms epoch of blue light applied to CFA and gray dotted lines are 10 ms after light onset. **b**, Example of muscle activity and its corresponding first derivative used for creating muscle activity state maps. **c**, Example map of muscle activity state vectors from one animal. Larger, connected dots show examples of state vectors for sequential overlapping epochs from individual trials. **d**, Grid overlaying a map including only points from the epochs used for weighting trials in grid point trial averages. **e**, 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. **f**, Schematic illustration of the effect size on a plot of Δ_L versus Δ_C . **g**, Inactivation effect maps for individual muscles (top) and grid point-averaged muscle activity from control (gray, mean \pm sem) and inactivation (cyan, mean) trials, for three example grid points. Panels **g,i,k** show representative results from one mouse. **h**, Aggregate effect size distributions for all grid points across all 8 mice. Control values in **h,k** are computed by replacing inactivation trials with a separate set of control trials without light. **i**, Maps of p-values computed for inactivation effects for individual muscles. **j**, Aggregate p-value distributions for all grid points across all 8 mice. **k**, Scatterplots of inactivation effect size versus muscle activity at light onset. Each point reflects a different grid point.

CFA activity components align with inactivation effects

In order to determine how CFA firing patterns covary with CFA influence on muscles, we used Neuropixels to measure the firing of CFA neurons in a subset of animals for which inactivation effect maps were computed ($n = 3$; Extended Data Fig. 5a). After completing collection of inactivation data for mapping influence in these mice, the next 3-4 daily behavioral sessions were devoted to neural recording in CFA with acutely-implanted Neuropixels. Template-based spike sorting yielded 366-684 units meeting our selection criteria per animal. Unit waveform centroid locations spanned all cortical layers, with the highest density of units found at depths corresponding to layer 5 (Extended Data Fig. 5b). 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^{41,45} (Extended Data Fig. 5c).

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. 3a). 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 state vectors were mapped to two dimensions 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 local to each grid point. As above, firing rate segments were weighted in averages by each corresponding muscle activity state vector's distance from the given grid point. Registering a neuron's firing to inactivation effect maps in this manner ensures that inactivation effects are matched with the neuron's expected activity immediately prior to the onset of those effects. The resulting maps of neural firing over muscle activity states (neural activity maps) showed a wide array of muscle-state dependent firing patterns (Fig. 3b). Statistical testing through permuting trial identities revealed that 72-81% of neurons with mean firing rates above 1 Hz had firing rates that varied with muscle state in each animal.

We used these neural activity maps to identify components of CFA firing that align across muscle activity states with CFA influence on muscles. We did so by combining singular value decomposition (SVD) and canonical correlation analysis (CCA) to align the neural activity maps with the inactivation effect maps computed for the same animals⁴⁸⁻⁵⁰. Neural activity maps for all wide-waveform neurons with overall mean firing rates above 0.1 Hz from a given mouse were converted into a grid points by neurons matrix, which was then replaced with a dimensionally-reduced, grid points by neural singular vectors version computed via SVD (Fig. 3c). These dimensionally-reduced matrices were then aligned through CCA with grid points by muscles matrices formed from inactivation effect maps for the four recorded muscles (Fig. 3c-f). The resulting canonical variables reflect components of CFA firing patterns that maximally correlate with components of inactivation effects but are mutually uncorrelated with one another. For all three animals, neural and inactivation effect canonical variables were highly correlated, and the

inactivation effect variables captured substantial fractions of inactivation effect variance (for canonical variables 1 through 4, mean correlation = 0.99, 0.97, 0.92 and 0.86; mean effect variance capture = 0.29, 0.26, 0.18 and 0.11; Fig. 3g). Correlation values were significantly higher than that of control values computed after permuting the map locations of trials from the neural recordings (for canonical variables 1 through 4, mean control correlation = 0.74, 0.72, 0.69, 0.65; $p = 0.001$ in all four cases). Plots of the cumulative variance captured across orthonormalized canonical vectors indicated that each inactivation effect variable captured a substantial amount of additional inactivation effect variance (Fig. 3h). Repeating CCA using the inactivation effect map of just one muscle found a CFA activity component highly correlated with the inactivation effects for the given muscle; this held for all muscles in each animal (median correlation = 0.97, range = 0.89 to 0.99, 12 total muscles in 3 animals). These results were robust to the particular trials used for computing inactivation effects (Extended Data Fig. 6a,b) and to the number of neuronal singular vectors used (Extended Data Fig. 6c-h). Thus, neural canonical vectors span a neural activity subspace where activity aligns closely with inactivation effects (influence subspace).

We examined the contribution of individual neuronal firing patterns to these influence subspaces. We used the weights of each neuron's firing in the singular vectors, and the weights of these singular vectors in the neural canonical vectors, to compute the relative contribution of each neuron's firing to each canonical vector. We then computed the norm of a vector composed of the four resulting weights for each neuron. The distributions of these vector norms across neurons were somewhat skewed (Fig. 3i), but the largest norm was never more than 5 times the median, and the middle half of neurons contributed almost as much aggregate weight as the top quarter of neurons (79.8-92.8% as much). Distributions of weights for individual muscles were somewhat more skewed, but still feature substantial contributions from a large fraction of neurons (Extended Data Fig. 6i). Thus activity components that align with inactivation effects involve substantial contributions from many CFA neurons.

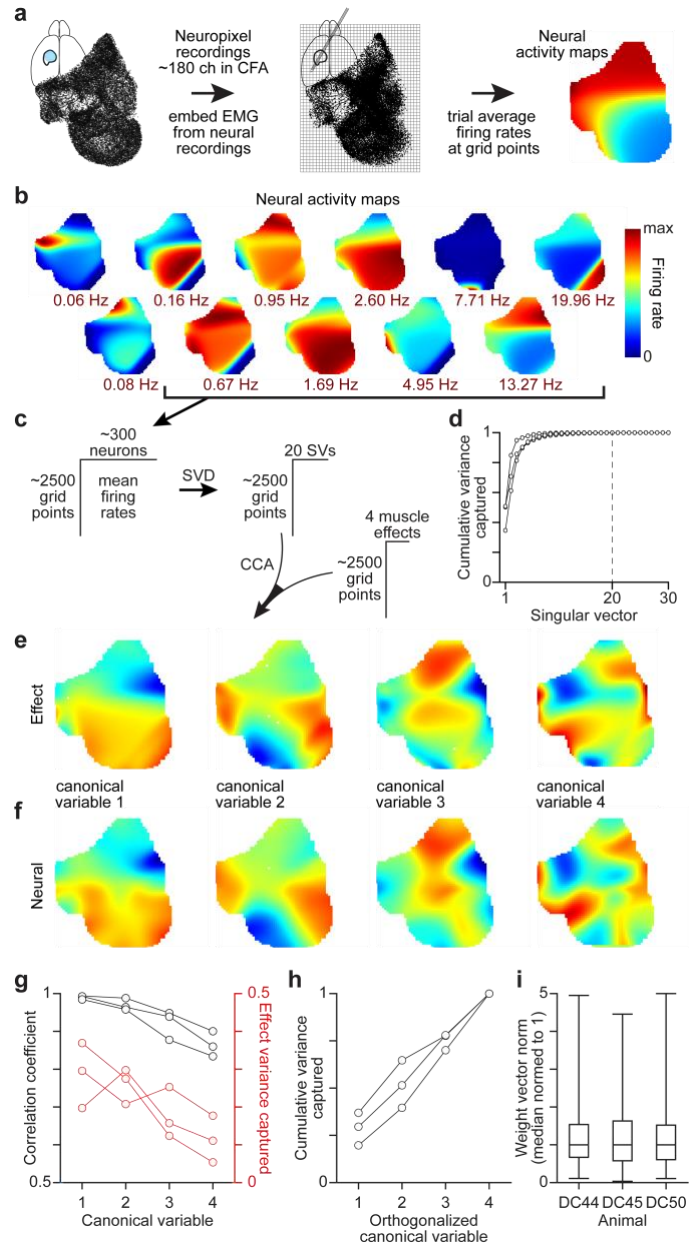


Figure 3 | A CFA activity subspace that aligns with CFA influence. **a**, Schematic of the calculation of an activity map for each recorded neuron that is registered to inactivation effect maps. **b**, Example neural activity maps from one mouse. Red text is the maximum firing rate. **c**, Schematic of computing the influence subspace. **d**, Cumulative variance captured versus singular vectors ordered by corresponding singular value, for three mice. Each connected set of dots is a separate mouse. **e, f**, Canonical variables for the inactivation effect (**e**) and neural activity (**f**) for one mouse. **g**, Correlation coefficient (black) and effect variance captured (red) for canonical variables. Each set of connected dots in **g, h** is from one animal. **h**, Cumulative variance captured by canonical variables after orthogonalizing their corresponding vectors. **i**, Distributions of the norms of vectors comprising the weights for the activity of individual neurons in the four canonical vectors (center line, median; box limits, 1st and 3rd quartiles; whiskers, minimum and maximum).

Influence subspace differs from muscle activity and limb kinematic equivalents

We then measured the overlap between the influence subspace and analogous subspaces that align to muscle activity and limb kinematics. To align neural and 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 maps were converted to a grid points by muscles matrix and aligned via CCA with the matrix of dimensionally-reduced neural activity (Fig. 4a,b; Extended Data Fig. 7a,b). To align neural activity and limb kinematics, new state maps were defined by UMAP not with muscle activity, but with the horizontal and vertical positions of sites on the right forelimb tracked as mice climbed during neural recordings. Nearby points on these maps reflect 50 ms epochs of limb kinematics that are similar (Extended Data Fig. 7c,d). We then assembled maps of the average horizontal and vertical site positions for grid points defined across these kinematic state maps (Extended Data Fig. 7e). We similarly assembled neural activity maps for each neuron using corresponding segments of their firing rates (Extended Data Fig. 7f). These two sets of maps were aligned via SVD and CCA (Fig. 4e,f; Extended Data Fig. 7g). For both muscle activity and limb kinematics, substantial fractions of variance were captured by canonical variables that were highly correlated with their corresponding neural canonical variables (Fig. 4c,d,g,h).

There is more overlap than would be expected by chance between influence and muscle activity subspaces, but the same is not true for their respective overlaps with limb kinematics subspaces. On a scale from 0 (no) to 1 (complete) overlap, subspace overlap between influence and muscle activity subspaces ranged from 0.343 to 0.497 across the three mice, while the overlap between influence and limb kinematics subspaces ranged from 0.017 to 0.025 and the overlap between muscle activity and limb kinematics subspaces ranged from 0.014 to 0.044 (Fig. 4i). Both of these ranges of overlap with limb kinematics subspaces were on par with those seen between subspaces defined at random while preserving the covariance captured by neural canonical variables. However, the overlap between influence and muscle activity subspaces was substantially above that of randomly-defined subspaces (Fig. 4i). The same relationships between subspaces were reflected in their principal angles as well (Fig. 4j-l).

To further quantify the difference between influence subspace and those that align with muscle activity or limb kinematics, we compared the overlap between subspaces with that expected for a given subspace type if computed using separate sets of trials. When two subspaces of each type are computed using separate random halves of trials, the two subspaces are highly similar. For the three mice, the mean fractional subspace overlap for 300 random trial partitions ranged from 0.965 to 0.995 for influence, 0.997 to 1.00 for muscle activity, and 0.849 to 0.895 for limb kinematics. However, when the same comparison using separate trials was done for different subspace types, the overlap was much lower. The mean overlap between influence and muscle activity subspaces ranged from 0.203 to 0.397 across animals, and the mean overlap between influence and kinematics subspaces ranged from 0.016 to 0.059 (Fig. 4m).

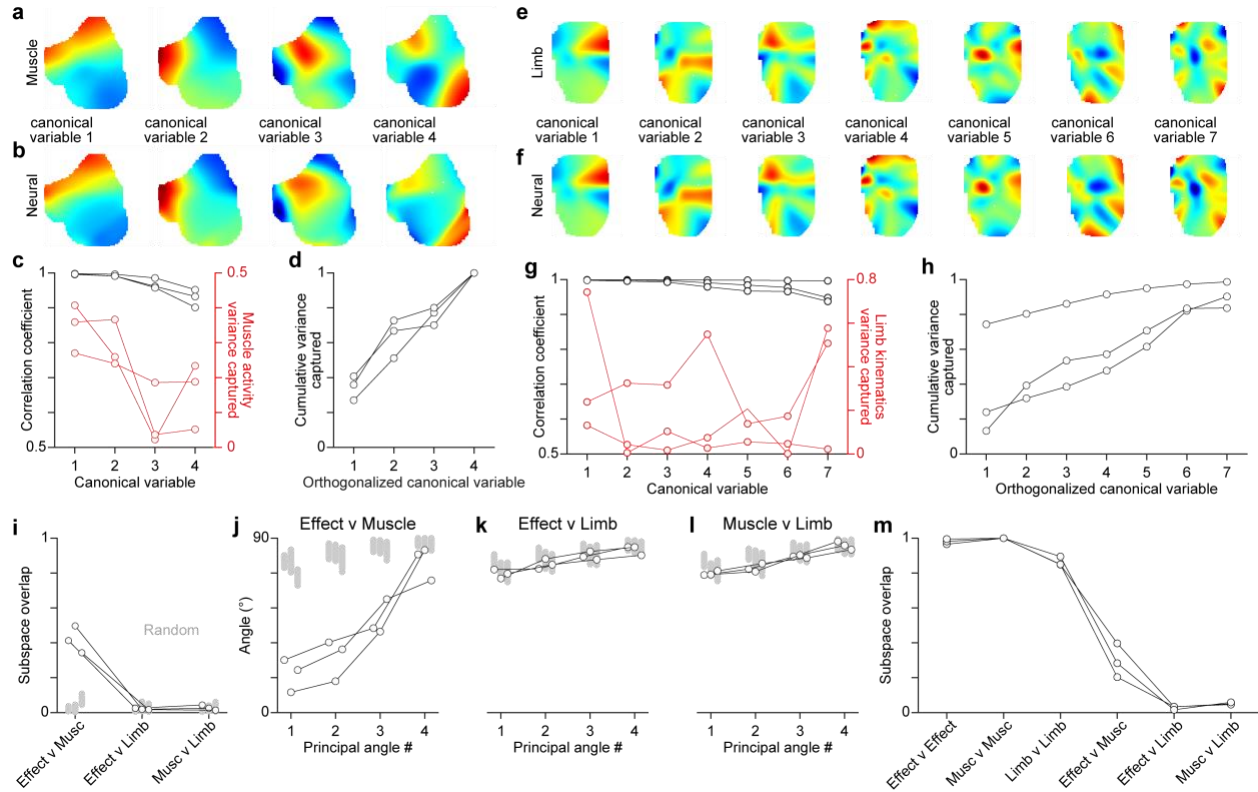


Figure 4 | Divergence between neural activity subspaces aligned with inactivation effects, muscle activity, and limb kinematics. **a,b**, Canonical variables for muscle activity (**a**) and neural activity (**b**) for one mouse. **c**, Correlation coefficient (black) and muscle activity variance captured (red) for canonical variables. Each set of connected dots in **c,d** are from one animal. **d**, Cumulative variance captured by canonical variables after orthogonalizing their corresponding vectors. **e,f**, Canonical variables for limb kinematics (**e**) and neural activity (**f**) for the mouse used in **a-d**. **g**, Correlation coefficient (black) and limb kinematics variance captured (red) for canonical variables. Each set of connected dots in **g-m** are from one animal. **h**, Cumulative variance captured by canonical variables after orthogonalizing their corresponding vectors. **i**, Overlap of different activity subspaces (black circles) compared to 300 estimates of the overlap expected by chance for each animal (gray dots). **j-l**, Principal angles for different activity subspaces (black circles) compared to 300 estimates of the principal angles expected by chance for each animal (gray dots). **m**, Overlap for subspaces defined from maps made with separate trials.

Discussion

Our findings reconcile a prominent dichotomy in the literature between the selective apparent involvement of motor cortex during most mammalian motor behaviors^{1,10-20} and the prevalent focus on motor cortical firing that correlates with motor outputs like muscle activity or limb kinematics en masse⁴⁻⁹. During a naturalistic motor behavior, mouse motor cortex appears to exert a selective, state-dependent influence through a neural activity subspace distinct from those to which functional significance has previously been imputed^{8,35-39}. Our observations lend significant support to two prominent hypotheses: first, that direct motor cortical influence on muscles is specific to certain muscle activity patterns^{2,3,21}, and second, that this influence comes from a subspace of the overall motor cortical activity space, the “output potent” subspace³⁹.

Our first key result came from quantifying variation in the direct influence of motor cortical output on muscle activity during naturalistic climbing behavior. Until now, it had remained unclear whether in behavioral contexts such as this motor cortical output directly influenced muscle activity pervasively, selectively, or not all, since large motor cortical lesions do not abolish climbing and other basic motor behaviors in nonhuman primates, and the observed degradation of movement following lesion could stem from motor cortical involvement in motor learning and movement preparation/initiation, rather than movement execution. We found that CFA influence during climbing is selective to certain muscle activity states, varies in magnitude and direction across states, and exhibits a state dependence that itself varies across muscles. Moreover, this influence bears a complex dependence on muscle activity state, indicating that CFA informs muscle activity patterns. While previous results indicated that direct motor cortical influence could be selective for certain motor behaviors³¹, our results here reveal a finer granularity of influence, one that varies markedly within a particular motor behavior. Our results agree with the long-standing hypothesis that during much of motor behavior, motor cortical output selectively drives specific muscle activity to broaden task capacity². It remains to be seen whether this influence is better viewed as selectively adding muscle activity patterns or steering muscle activity state.

Our second key result came from identifying a subspace where CFA activity closely covaries with its influence on muscles. 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^{31,45}. 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^{39,51}. However, the influence subspace we identified here from aligning CFA activity with inactivation effects differs from previous definitions of an “output potent” subspace, and this influence subspace is distinct from those that align with muscle activity or limb kinematics. This indicates that motor commands emanating from motor cortex during a dexterous ethological behavior like climbing do not match muscle activity or limb kinematics in total, in agreement with a selective and limited motor cortical influence across muscle activity states.

The divergence between subspaces that align with influence, muscle activity, or limb kinematics suggests that the ability to predict muscle activity and limb kinematics from motor cortical activity reflects processes other than the direct driving of muscles, like monitoring or predicting body state^{33,34} for computing motor errors. The orthogonality of the kinematics subspace to the other two subspaces has important implications for the functional interpretation of kinematics-related neural activity. This is a significant revelation given the contemporary prevalence of video-based kinematic tracking for measuring nervous system output.

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 unidirectional turning of the wheel. 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 carried 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 four days after the start of the water schedule, mice were acclimated to handling by the experimenter following established procedures⁵². 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 μ l 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 the handholds in order to rotate the wheel downward and receive water rewards. 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 longer bouts. Here, rewards were dispensed for continuous climbing bouts exceeding a threshold distance, after the bout ended. 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 distance was above the 25th percentile value for the previous 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^{31,42,43}. 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 into muscle, and the distal portion of the electrodes was knotted. 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 light³¹. 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 a 25 ms light pulse at an intensity of 10 mW/mm² to the brain surface. During each session that involved this inactivation, trials were initiated after the current bout distance exceeded a random threshold between 0° and 20°, provided the mouse was still actively climbing. This ensured that trials were broadly distributed across the muscle activity states that occur during climbing. Light pulses were applied during a random third of trials, while the remaining trials served as controls (no light). This ratio of control trials to inactivation trials afforded equivalent statistical power in comparing light effects (inactivation vs. control) to effects expected by chance (control vs. control). Inactivation trials were never triggered less than 5

seconds apart to allow ample time for recovery. Counting both inactivation and control, a total of 2292-5115 trials (median = 2715) were collected in each mouse, 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 (IMEC) was subsequently inserted to a depth of 1.5 mm into the brain (Extended Data Fig. 5a) at a rate of 3-5 $\mu\text{m/s}$ 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 2.0⁵⁴ (<https://github.com/MouseLand/Kilosort2>).

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 trial/light onset. Annotation of behavior was accomplished using DeepLabCut⁴⁴. To enable better markerless tracking, the right forelimbs of recorded mice were shaved and tattooed (Black tattoo paste, Ketchum Mfg., Lake Luzerne, New York) at several sites along the right arm. All videos were also adjusted with ffmpeg (ffmpeg.org) to improve brightness and contrast. In DeepLabCut, we provided manually labeled locations of the eight forelimb sites on ~4000 randomly sampled video frames across mice and sessions for training: shoulder, two sites between the shoulder and elbow, elbow, two sites between the elbow and wrist, wrist, and last digit (Extended Data Fig. 2a). DeepLabCut's ResNet-50 neural network with an Adam optimizer was trained on the annotated images for 1,030,000 iterations. 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. 2b-c). First, sites in each video time series that were assigned a likelihood lower than 0.75 by DeepLabCut (i.e., its confidence that a site was correctly labeled) were replaced with an interpolated value using the median of the ten previous and ten following values (Matlab function `fillmissing`). Outlier sites in the time series were also identified using the median absolute deviation (MAD), and these were replaced with an

interpolated value using a moving median of window length 10. Shoulder coordinates were constrained to lie within 1.5 scaled 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. 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 extend 180 degrees, and the distances between each of these joints were within 2 MAD of their medians (e.g., the length of a line connecting the shoulder to the elbow should stay relatively constant throughout the videos recorded from a single session).

EMG preprocessing

For both optogenetic inactivation and Neuropixel recording sessions, EMG measurements were high-pass filtered at 250 Hz, rectified, and convolved with a Gaussian kernel having a 10 ms standard deviation but amplitudes for times < 0 – its “backwards-in-time” side – set uniformly to zero. This latter feature ensured causal filtering to enable precise assessment of perturbation latencies. EMG traces were then z-scored across time and downsampled to 1 kHz.

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 EMG time series and their corresponding first derivatives. For each control or inactivation trial, we extracted 13 overlapping 50 ms segments centered every 10 ms from -55 ms to 75 ms from trial/light onset. For each segment, we averaged the EMG traces for each muscle over 5 ms bins and concatenated the EMG segments for the four muscles and their corresponding first derivatives, yielding one 80 by 1 vector. In these vectors, the first 40 values reflect the EMG signals from four muscles, and the last 40 values reflect their first derivatives (Fig. 2b). UMAP (Matlab function `run_umap`, from the Matlab File Exchange, `n_neighbors = 30`, `n_components = 2`, `min_dist = 0.3`, `metric = Euclidean`) was applied to all vectors, including those from both inactivation and control trials, to get the muscle activity state maps (Figs. 2c-d). Using two dimensions here simplified subsequent quantification of inactivation effects across muscle activity states. Using overlapping segments ensured that segments from individual trials formed continuous trajectories across the resulting maps (Fig. 2c).

Muscle inactivation effects

To quantify the influence of CFA inactivation on muscles across the muscle activity state map, we first excluded outlier segments. A segment was considered an outlier if its median/mean distance from other segments on the map was more than 3 times the standard deviation of all the pairwise distances between all pairs of segments. We then defined uniformly spaced grid points across each map and proceeded to quantify the effect of inactivation at each grid point. 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 map rd to $rd_{rescaled} = \frac{rd - \text{mean}(rd)}{\text{std}(rd)} * 10 + \text{weightPara} * 2$, with $\text{weightPara} = 5$. weightPara controls the standard deviation of a Gaussian function for defining the weights used to weight trials in computing the inactivation

effect size at each grid point. For each grid point, we calculated the L2 distance between each segment on the map and the grid point, and then took the sum of a Gaussian function of each of these distances to obtain $\sum_i W_{g,i}$, where $W_{g,i} = \exp\left(-\frac{L2(\text{coordinates}_g - \text{coordinates}_i)^2}{2 * \text{weightPara}^2}\right)$, i is the segment index, and g is the grid point index. A grid point was designated as “valid” if $\sum_i W_{g,i} > 10$; grid points not designated as valid were ignored in subsequent calculations as they had few segments nearby.

Next, we calculated the trial averaged EMG for each muscle and for each grid point separately for the inactivation and control trials. For each muscle, we extracted a segment of the EMG time series associated with each trial, from -10 ms to 30 ms relative to trial/light onset. We then took a weighted average of these segments around each grid point, where each was weighted by $W_{g,i}$, with i set to the index of the embedded segment from the given trial that spanned from -40 to 10 ms relative to pulse onset. This results in control and inactivation trial averages for each muscle and for each grid point where each trial is effectively weighted by the distance from the grid point to its muscle activity state just prior to any inactivation effect. 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 across a 20 ms window spanning from 0 ms to 20 ms from pulse onset (Fig 2e). 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. 2e,f).

Testing significance of inactivation effects

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 or 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 for one randomly selected group of control trials and the inactivation trial average, both 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 (Fig. 2e-i; Extended Data Fig. 4). 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 as the p-value (Fig. 2j). To correct

for multiple comparisons, p-values for each valid grid point were subject to the Benjamini-Hochberg method to control the false discovery rate (Matlab function `fdr_bh`; $q < 0.05$ as the criteria). The effect size at each grid point was considered to be statistically significant if the FDR-corrected p-value was less than 0.05.

Neural activity maps

To compare firing in CFA to inactivation effects, we generated maps of the average firing rate across muscle activity states (grid points) for each recorded neuron. Since Neuropixel recordings and optogenetic inactivations were carried out in separate sessions, the first step for revealing the association was to align muscle activity states obtained from the neural recording sessions and the optogenetic inactivation sessions by embedding them on the same 2D map. Thus we extracted 10,000 50 ms segments of the EMG time series during climbing bouts from the Neuropixel recording sessions. First, peaks of the summed EMG traces across muscles were detected (using the Matlab function `findpeaks`, where a peak is identified if a sample in the time series surpasses a threshold of the mean + $\frac{1}{2}$ *std). Second, two time points were randomly selected around each peak, with two mean distances of ± 50 ms respectively, as “onset points”. Next, any onset point was eliminated if its distances to its two neighbors were smaller than 50 ms. This step ensured that the onset points were sufficiently spaced apart from each other. For each remaining onset point, the EMG segment spanning from -45 ms to +5 ms relative to the onset point was selected. Each recording session typically yielded upwards of 30,000 segments. Of these, 10,000 were randomly selected to limit compute time for subsequent calculations.

We embedded the 10,000 segments from each neural recording session onto the muscle activity state maps defined using the UMAP template obtained from the inactivation sessions (Fig. 3a, middle; see above). We then extracted the corresponding segments of neuronal firing rate time series corresponding to these EMG segments. For each neuron and at each valid grid point, the firing rate segments were averaged with the weights defined by the Gaussian function of their distance on the map to the grid point with *weightPara* = 5. The resulting trial-averaged firing rates were then averaged across time, yielding a single scalar firing rate value for each neuron at each grid point (Fig. 3a, right, and b).

For analysis described below, only putative pyramidal neurons were included here. To select for these putative pyramidal neurons, we pooled the spike widths of all neurons and obtained a bimodal distribution, which separated putative pyramidal neurons and interneurons. Neurons with narrower spike widths were considered to be interneurons and were excluded from further analyses (Extended Data Fig. 5b).

To estimate the number of neural activity maps that showed significant variation across muscle activity states, we generated an empirical null distribution for the degree of variation across neural activity maps separately for each neuron. To do so, we reassigned each EMG segment from neural recording sessions to the location of a different, randomly selected segment on the map, doing so 500 times to yield 500 permuted maps. For each one, we recomputed the neural activity map as

described above. The skewness of the original neural activity map and the 500 permuted neural activity maps were calculated (Matlab function kurtosis) as $k = \frac{E(x-\mu)^4}{\sigma^4}$, where x is the set of map values, 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 $\frac{N_{k_{surrogate} > k_{empirical}}}{N_{k_{surrogate}}}$. With p-values calculated, we calculated the fraction of false null hypotheses (Matlab function mafdr) and used this as the fraction of cells with behaviorally-dependent firing.

Neural activity subspace for inactivation effects

Separately for each animal ($n = 3$ mice), neural activity subspaces were identified using singular vector canonical correlation analysis (SVCCA⁵⁰) with the averaged neural firing rates and the muscle inactivation effects across a time window (0 ms to 20 ms) for all the valid grid points. 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. Pyramidal neurons with mean firing rates lower than 0.1 Hz over the 10,000 embedded segments were excluded. This firing rate threshold was selected to lie within a range where the number of singular vectors required to capture 95% variances was relatively stable. 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) and 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 4 muscles recorded.

SVCCA was conducted in two steps (Fig. 3c-d). Neural activity values were first soft normalized⁵³ using a soft normalization constant of 5 Hz. SVD was performed using `numpy.linalg.svd` in Python to decompose the neural activity into left singular vectors, a diagonal matrix containing singular values, and right singular vectors. Then a truncated diagonal matrix with the top 20 values was multiplied with the corresponding top 20 right singular vectors, resulting in a reduced, 20-dimensional matrix of neural activities. CCA was applied to this matrix 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. This process was also repeated with individual columns of the matrix of inactivation effect sizes in place of the full matrix to show that there were CFA activity components highly correlated with the inactivation effects for each individual muscle.

To visualize the goodness of fit, neural canonical vectors were projected back onto the embedding map. To compute additional variance capture by sequential canonical vectors, canonical vectors were orthogonalized with the Gram-Schmidt process. In addition, we randomly shuffled the neural firing rate segments relative to the EMG segments on the maps to generate shuffled neural activity

maps, and performed SVCCA again as a negative control. The highest correlations between these 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 the separate inactivation effect maps for each group, used SVCCA to find neural activity 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. 3c).

Using the CCA results obtained with individual columns of the matrix of inactivation effect sizes, 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 sets of effective weights for each neuron (Extended Data Fig. 5e, h).

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. In place of the matrix of inactivation effect sizes, however, an N_{grid} by 4 matrix was used where each entry reflected the average muscle activity at each grid point for one of the four muscles, computed in the same way as described above for the neural activity maps. Similar methods were used to find a neural activity subspace that aligned with limb kinematics. Here instead of four dimensions, we began with 16 – horizontal and vertical coordinates for the eight 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 (Fig. 4e-h).

To compute principal angles between neural activity subspaces, we orthonormalized each subspace, built a matrix of inner products of the basis vectors from the two subspaces, performed SVD on this matrix, and calculated the inverse cosine of the singular values in degrees⁵⁴. We also measured the degree of overlap between two neural activity subspaces using the metric OL given by $OL = \frac{\sum_{i=1}^R var(M_{reproj_{s2}}^i)}{\sum_{i=1}^R var(M_{proj}^i)}$, where M is the original data matrix, M^i is the i^{th} row of M , R is the number of rows in M , C_{s1} and C_{s2} are matrices comprising orthonormal coefficient vectors that define the two subspaces, $M_{proj} = C_{s1} * M$ (i.e., the projection onto the first subspace), $M_{reproj} = C_{s1}^T * M_{proj}$ (i.e., the pseudo-full rank matrix projected back from the first subspace), and $M_{reproj_{s2}} = C_{s2} * M_{reproj}$ (i.e., the projection of M_{reproj} onto the second subspace).

To compute the overlap of randomly chosen neural activity subspaces that preserved the covariance captured by neural canonical variables, we recomputed the overlap 100 times after permuting the order of coefficients in the vector C_{s2} to get \underline{C}_{s2} , and only using \underline{C}_{s2} for which the total variance of $\underline{C}_{s2} * M$ was between 0.8 - 1.2 times the total variance of $C_{s2} * M$. This ensured

that the variance captured by these sets of permuted vectors was similar to that of the original subspace. We also calculated subspace overlap for subspaces that were each computed using just one half of the total segments from the neural recording for 300 random parcellations of the segments.

Acknowledgments

The authors would like to thank 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

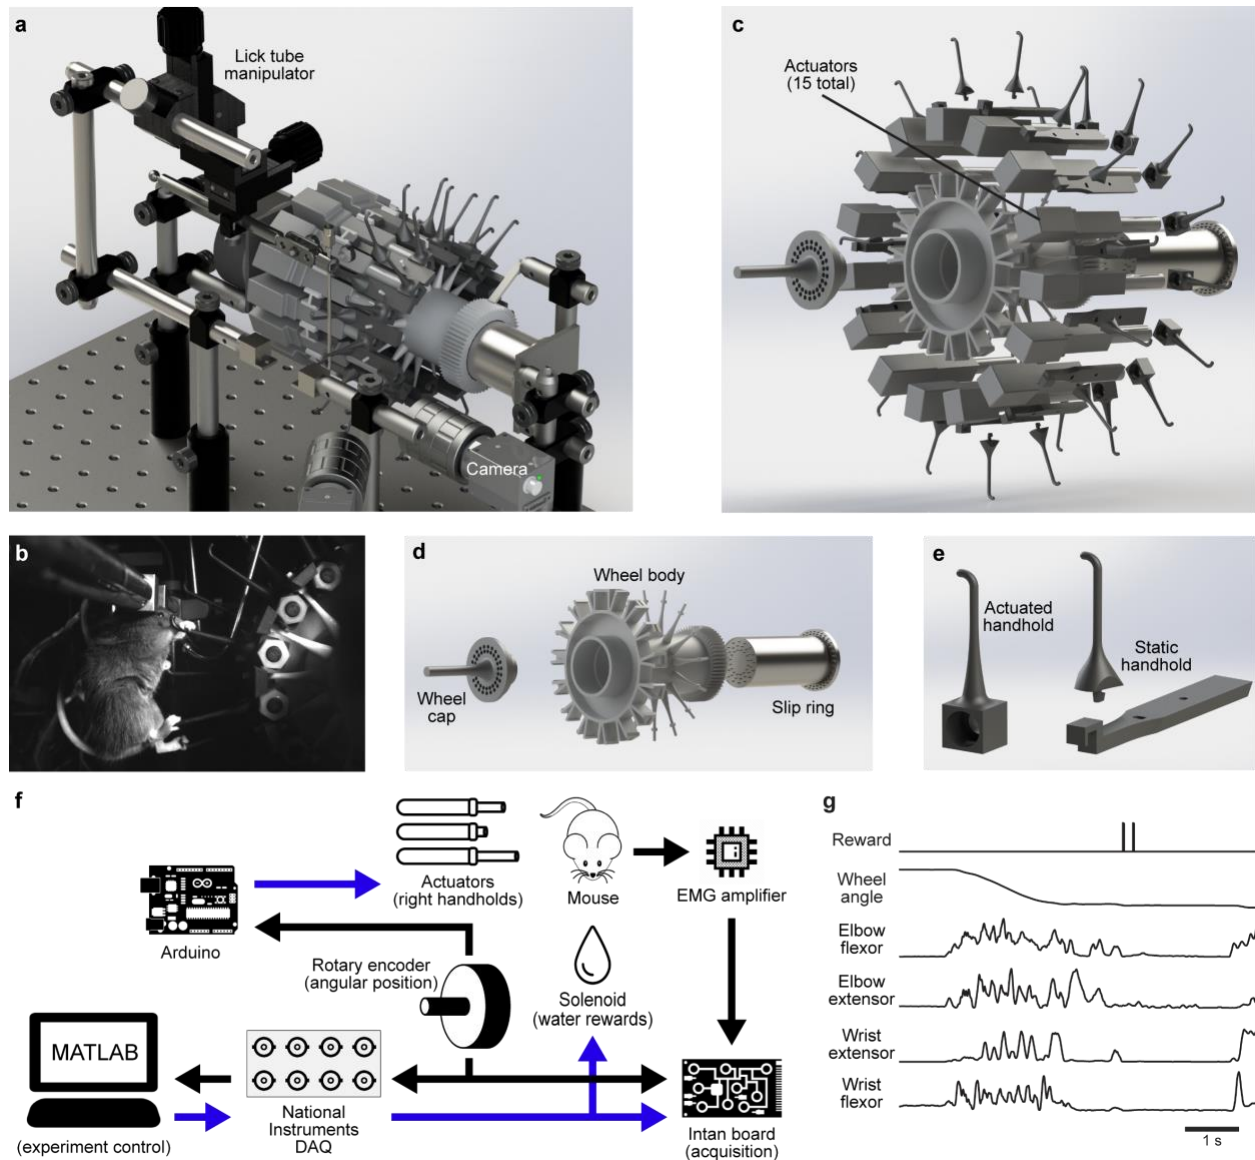
Correspondence and request for materials should be addressed to A.M. (andrewmiri@northwestern.edu).

Data availability

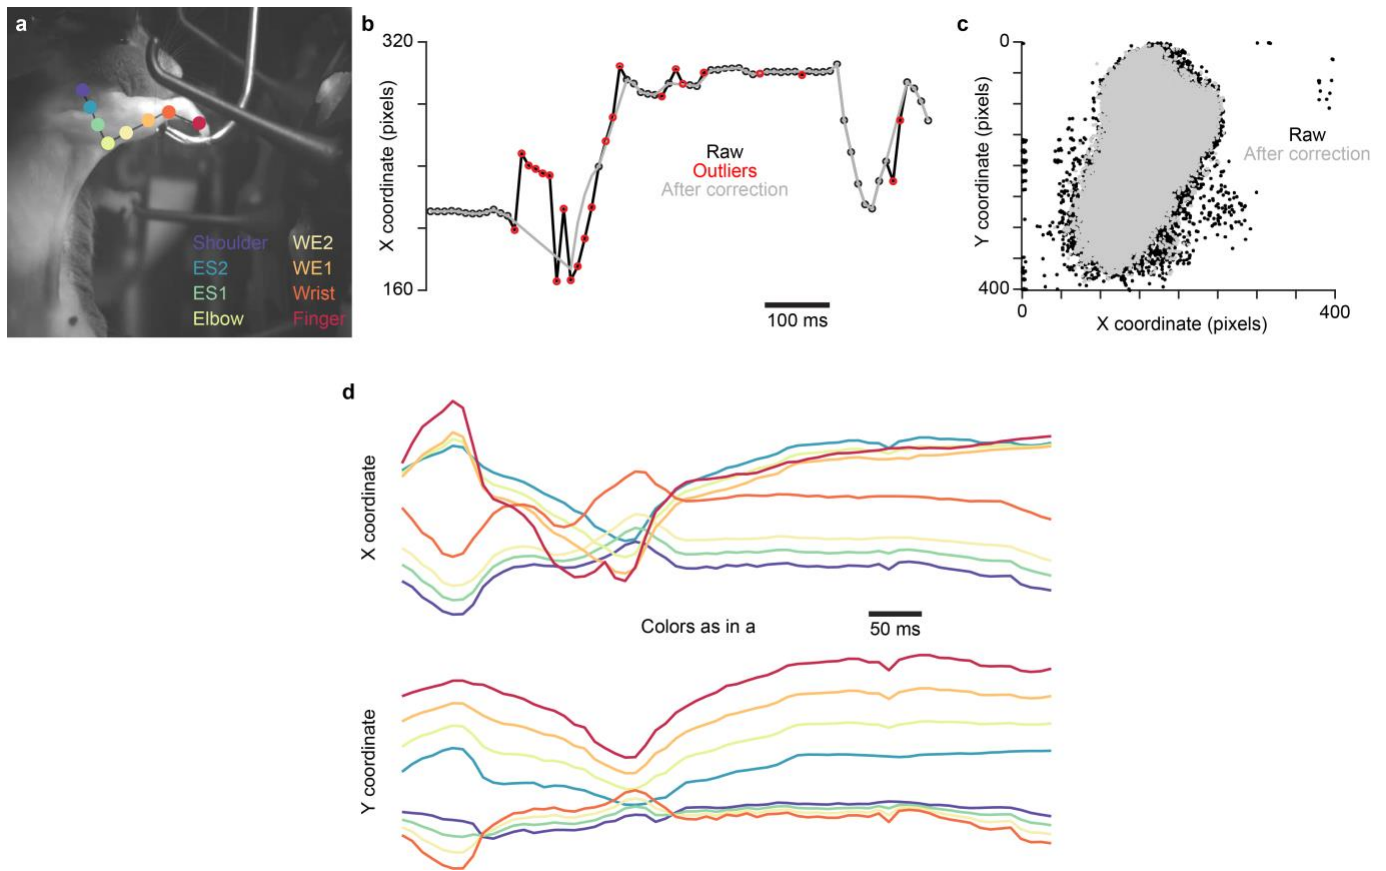
The data that support the findings of this study are available from the corresponding author upon reasonable request.

Code availability

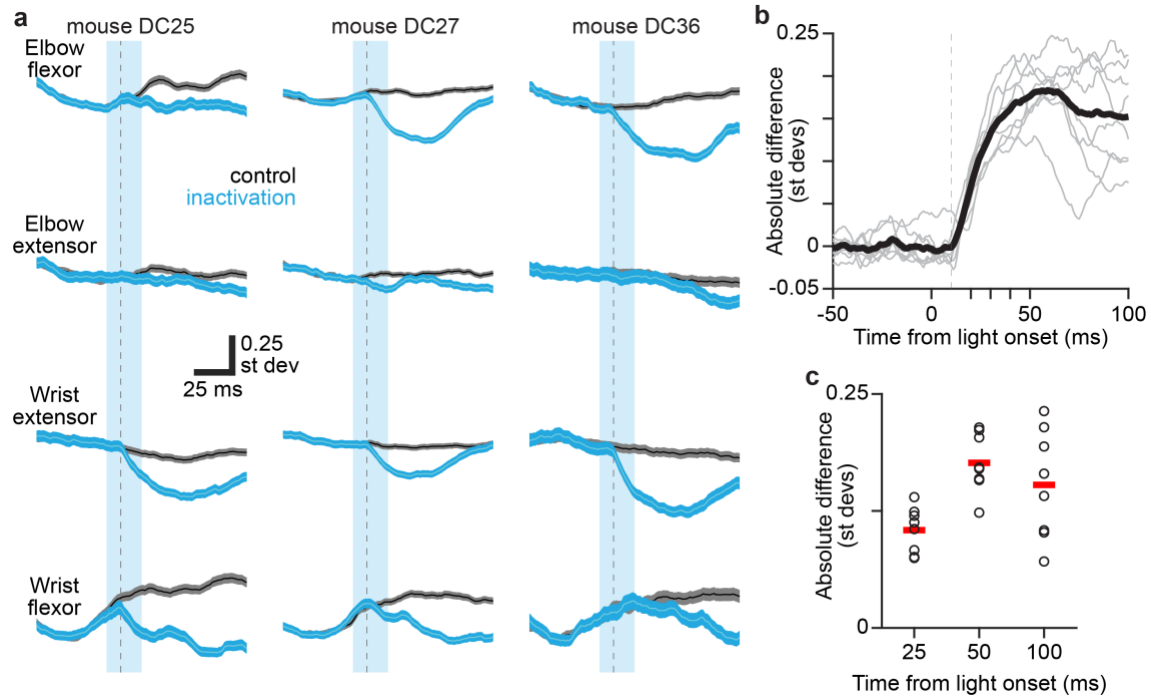
All Matlab code for data analyses and CAD files for 3D printed wheel components will be made available on GitHub upon publication.



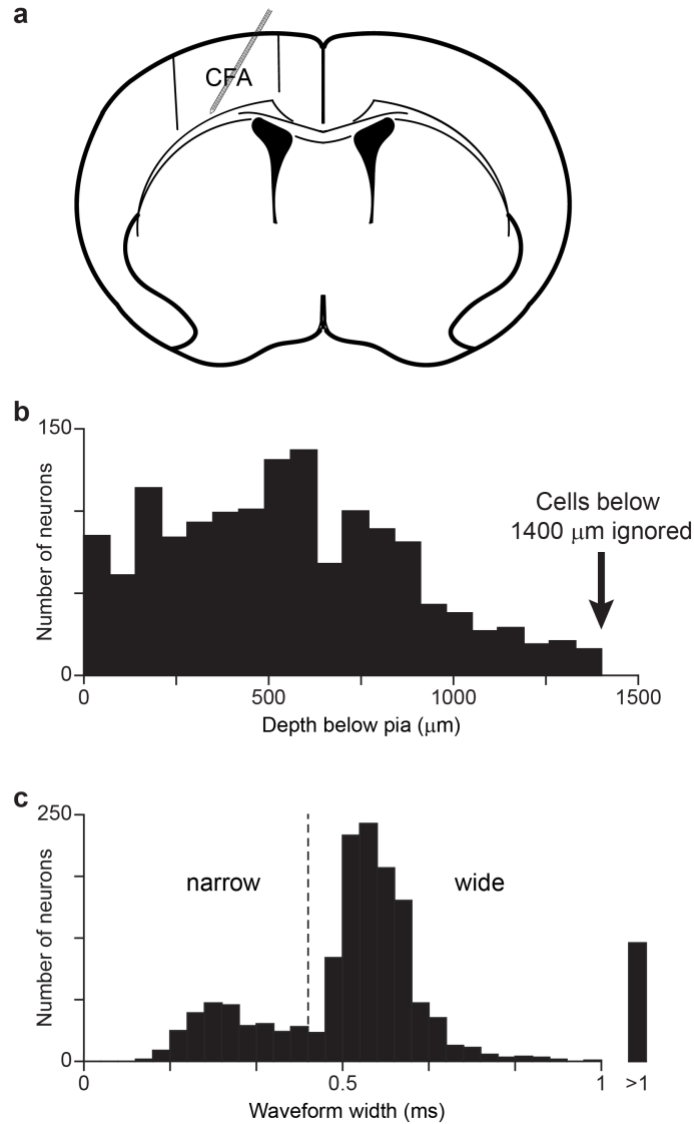
Extended Data Figure 1 | Head-fixed climbing paradigm. **a**, 3D illustration of the climbing apparatus, including the lick tube manipulator with lick tube, mounts for head fixation, and cameras. **b**, Side view photo of a head-fixed mouse during a climbing session. **c**, 3D exploded illustration of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



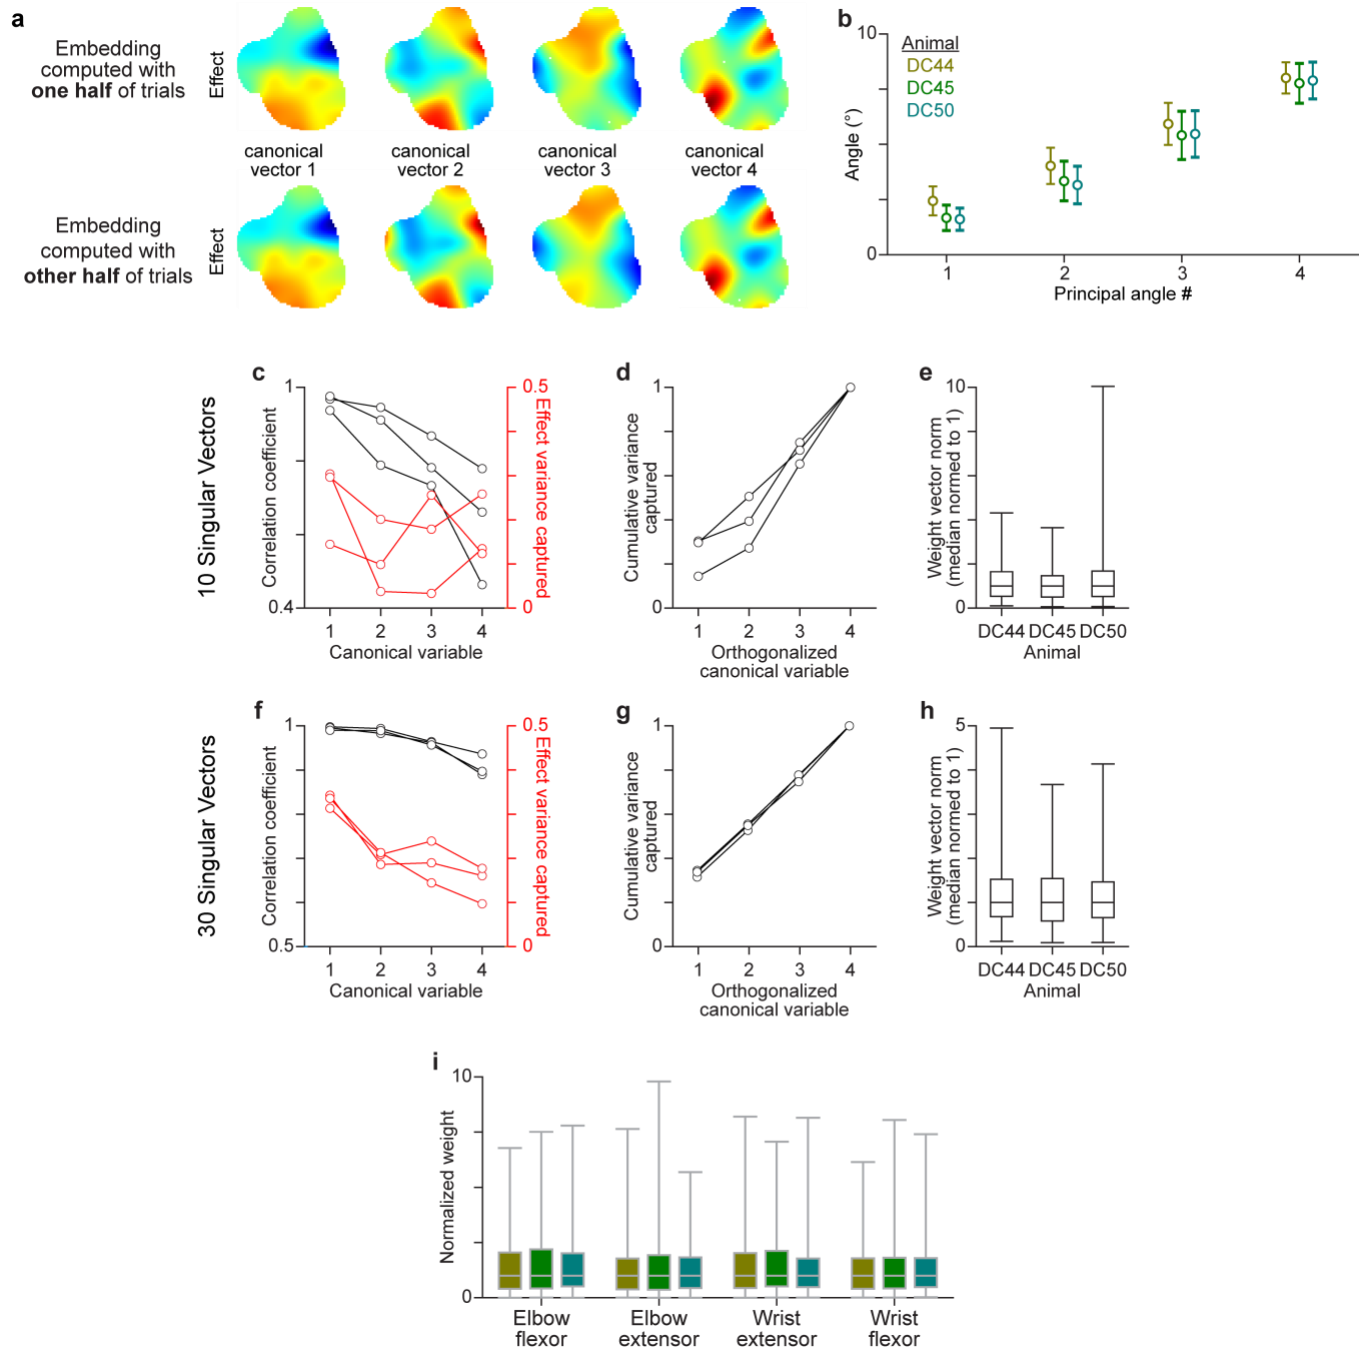
Extended Data Figure 2 | Video processing to extract forelimb kinematics. **a**, DeepLabCut-labeled video image, showing the eight sites along the right forelimb that were tracked (ES1,2: 1st and 2nd sites between elbow and shoulder joints, WE1,2: 1st and 2nd sites between wrist and elbow joints). **b**, Example time series for the finger site from post-hoc processing of DeepLabCut labels to correct outliers. Raw is the time series output by DeepLabCut. **c**, Scatter plot of X and Y coordinates of the finger site tracked in **b**, showing fewer outliers after correction. **d**, Example X and Y coordinate traces for site positions after post-hoc processing, with colors corresponding to the labels shown in **a**.



Extended Data Figure 3 | Effect of CFA inactivation on muscle activity. **a**, Mean \pm 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. **b**, Absolute difference between inactivation and control trial averages summed 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**, Absolute difference between inactivation and control trials at 25, 50, and 100 ms after trial/light onset for individual animals (black circles) and the mean across animals (red bars).

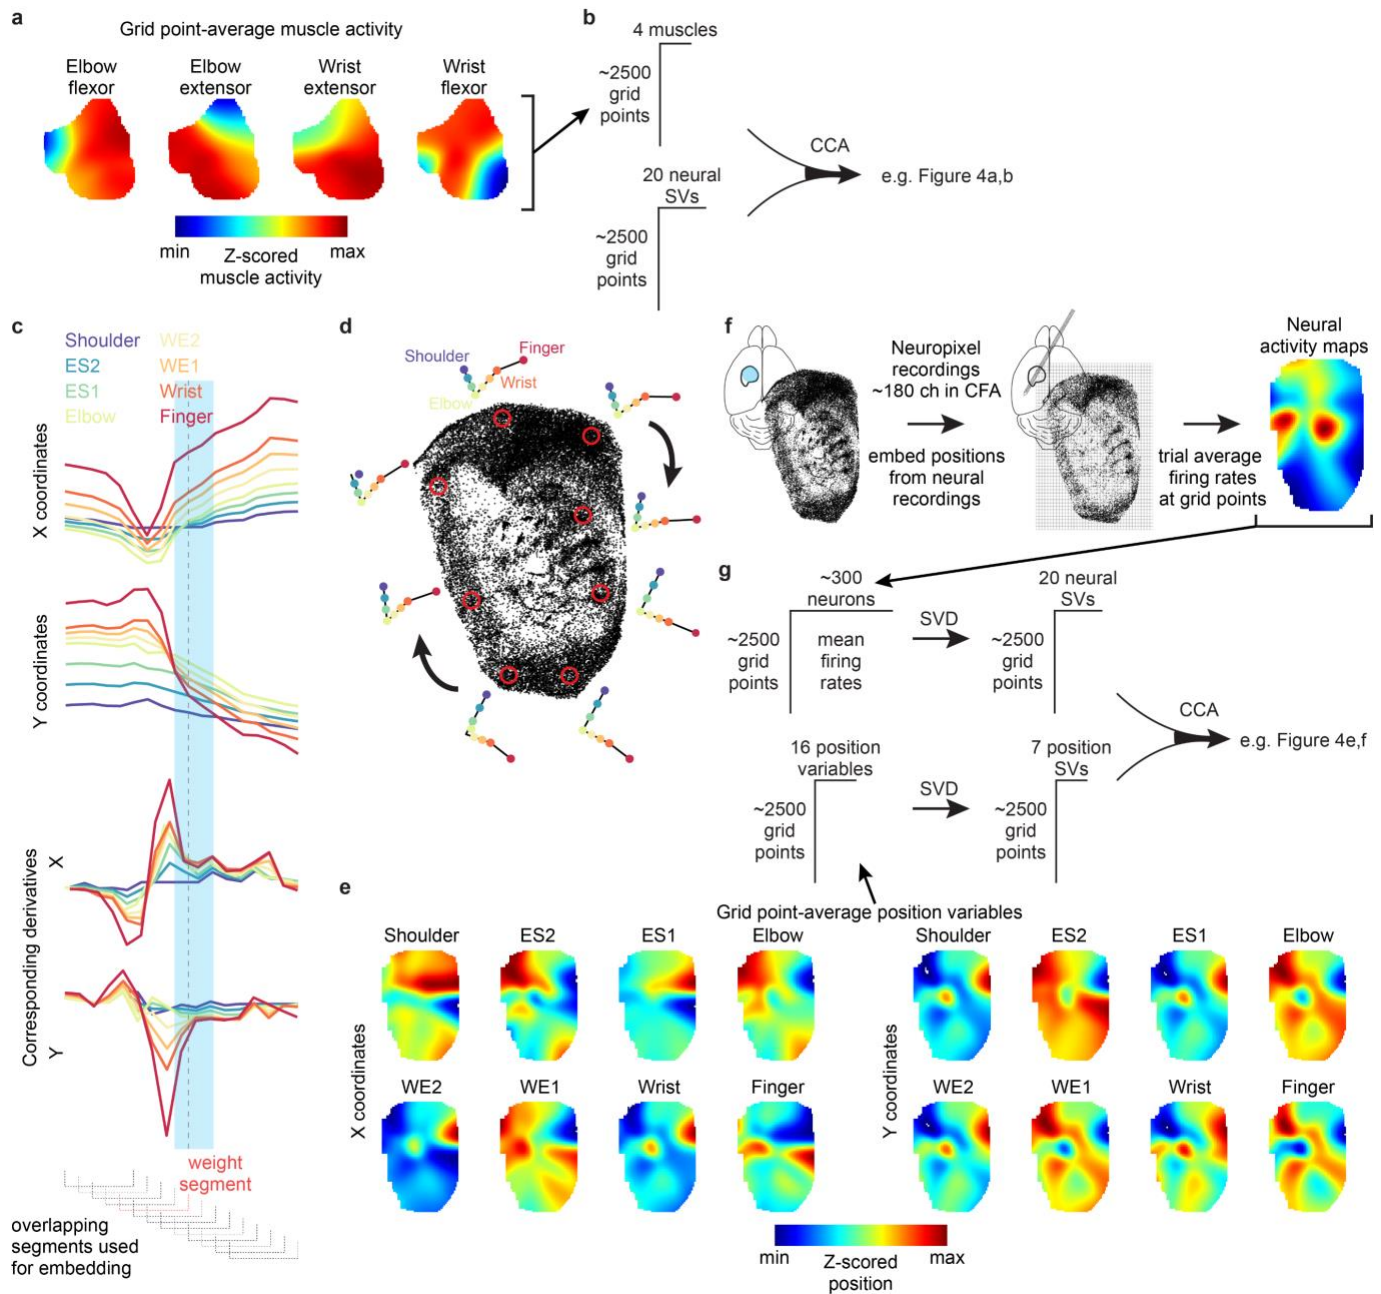


Extended Data Figure 5 | Recording neural activity in CFA. **a**, Schematic coronal section showing placement of the Neuropixel in CFA. The Neuropixel was inserted at an angle of 30° with respect to the midline. **b**, Histogram showing the depth below pia of recorded neurons, as assigned by Kilosort. **c**, Histogram of waveform widths for recorded neurons, showing the bimodal distribution of narrow and wide waveforms. Dotted line shows the threshold above which neurons were considered wide waveform.



Extended Data Figure 6 | CFA neural activity subspaces that align with CFA influence. a, Canonical variables for the inactivation effect maps resulting from CCA with neural activity maps, using effect maps computed with separate halves of trials. **b,** Mean \pm 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-e** show results when SVD is used to reduce neural activity dimensionality to 10. **c,** Correlation coefficient (black) and effect variance captured (red) for canonical variables. Each set of connected dots in **c,d** is from one animal. **d,** Cumulative variance captured by canonical variables after

orthogonalizing their corresponding vectors. **e**, Distributions of the norms of vectors comprising the weights for the activity of individual neurons in the four canonical vectors (center line, median; box limits, 1st and 3rd quartiles; whiskers, minimum and maximum). **f-h**, Same as **c-e**, but when SVD is used to reduce neural activity dimensionality to 30. **i**, Distributions of weights for the activity of individual neurons in canonical vectors defined by CCA using the inactivation effect maps of individual muscles (box plot elements as defined in **e**).



Extended Data Figure 7 | 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 4a,b. **c**, Time series of limb kinematics data and the corresponding first derivatives from CFA inactivation sessions are windowed into overlapping segments that are used to create a 2D embedding via UMAP. **d**, Example map of limb kinematics state vectors for an example animal, along with the trial-averaged positions of the forelimb sites at selected grid points within the map (red circles). **e**, Grid point-averaged position maps for the eight

tracked sites along the right forelimb for an example animal. **f**, Site positions tracked during neural recording sessions are embedded into the 2D limb kinematics state map, and the neural activity corresponding to the embedded segments is used to calculate grid point-averaged neural activity maps. **g**, The maps in **e** 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, together with the grid points by neural singular vectors matrix obtained from the neural activity maps, serve as input to CCA, which then identifies the canonical variables shown in Figure 4e,f.



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 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).
- 2 Shmuelof, L. & Krakauer, J. W. Are we ready for a natural history of motor learning? *Neuron* **72**, 469-476 (2011).
- 3 Schieber, M. H. How Might the Motor Cortex Individuate Movements. *Trends Neurosci* **13**, 440-445 (1990).
- 4 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).
- 5 Evarts, E. V. Relation of pyramidal tract activity to force exerted during voluntary movement. *J Neurophysiol* **31**, 14-27 (1968).
- 6 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).
- 7 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).
- 8 Kalaska, J. F. From intention to action: motor cortex and the control of reaching movements. *Adv Exp Med Biol* **629**, 139-178 (2009).
- 9 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).
- 10 Fulton, J. F. & Kennard, M. A. A Study of Flaccid and Spastic Paralysis Produced by Lesions of the Cerebral Cortex in Primates. *Res Publ Assoc Res N* **13**, 157-209 (1934).
- 11 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).
- 12 Leyton, A. S. F. & Sherrington, C. S. *Observations on the excitable cortex of the chimpanzee, orang-utan, and gorilla.* (Griffin, 1917).
- 13 Grunbaum, A. S. F. & Sherrington, C. S. "Observations on the physiology of the cerebral cortex of the anthropoid apes." *P R Soc London* **72**, 152-155 (1903).
- 14 Ogden, R. F., S.I. . On cerebral motor control: The recovery from experimentally produced hemiplegia. *Psychobiology* **1**, 33-47 (1917).
- 15 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).
- 16 Denny-Brown, D. & Botterell, E. H. The motor functions of the agranular frontal cortex. *Res Publ Assoc Res Nerv Ment Dis* **27 (1 vol.)**, 235-345 (1948).
- 17 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).
- 18 Kawai, R. *et al.* Motor cortex is required for learning but not for executing a motor skill. *Neuron* **86**, 800-812 (2015).
- 19 Lashley, K. S. Studies of cerebral function in learning V The retention of motor habits after destruction of the so-called motor areas in primates. *Arch Neuro Psychiatr* **12**, 249-276 (1924).

- 20 Travis, A. M. Neurological deficiencies after ablation of the precentral motor area in
Macaca mulatta. *Brain* **78**, 155-173 (1955).
- 21 Tower, S. S. Pyramidal lesion in the monkey. *Brain* **63**, 36-90 (1940).
- 22 Palop, J. J., Chin, J. & Mucke, L. A network dysfunction perspective on neurodegenerative
diseases. *Nature* **443**, 768-773 (2006).
- 23 Fetsch, C. R. *et al.* Focal optogenetic suppression in macaque area MT biases direction
discrimination and decision confidence, but only transiently. *Elife* **7** (2018).
- 24 Calaim, N., Dehmelt, F. A., Goncalves, P. J. & Machens, C. K. The geometry of robustness
in spiking neural networks. *Elife* **11** (2022).
- 25 Aerts, H., Fias, W., Caeyenberghs, K. & Marinazzo, D. Brain networks under attack:
robustness properties and the impact of lesions. *Brain* **139**, 3063-3083 (2016).
- 26 Hwang, E. J. *et al.* Disengagement of motor cortex from movement control during long-
term learning. *Science Advances* **5** (2019).
- 27 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).
- 28 Balasubramanian, K. *et al.* Propagating Motor Cortical Dynamics Facilitate Movement
Initiation. *Neuron* **106**, 526-536 e524 (2020).
- 29 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).
- 30 Guo, J. Z. *et al.* Cortex commands the performance of skilled movement. *Elife* **4**, e10774
(2015).
- 31 Miri, A. *et al.* Behaviorally Selective Engagement of Short-Latency Effector Pathways by
Motor Cortex. *Neuron* **95**, 683-696 e611 (2017).
- 32 Wolff, S. B. E. & Olveczky, B. P. The promise and perils of causal circuit manipulations.
Current Opinion in Neurobiology **49**, 84-94 (2018).
- 33 Hatsopoulos, N. G. & Suminski, A. J. Sensing with the motor cortex. *Neuron* **72**, 477-487
(2011).
- 34 Adams, R. A., Shipp, S. & Friston, K. J. Predictions not commands: active inference in the
motor system. *Brain Struct Funct* **218**, 611-643 (2013).
- 35 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).
- 36 Todorov, E. Direct cortical control of muscle activation in voluntary arm movements: a
model. *Nat Neurosci* **3**, 391-398 (2000).
- 37 Miri, A., Azim, E. & Jessell, T. M. Edging toward entelechy in motor control. *Neuron* **80**,
827-834 (2013).
- 38 Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A. Neural Manifolds for the Control
of Movement. *Neuron* **94**, 978-984 (2017).
- 39 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).
- 40 Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity.
Nature **551**, 232-236 (2017).
- 41 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).

- 42 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).
- 43 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).
- 44 Mathis, A. *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with
deep learning. *Nat Neurosci* **21**, 1281-1289 (2018).
- 45 Guo, Z. V. *et al.* Flow of cortical activity underlying a tactile decision in mice. *Neuron* **81**,
179-194 (2014).
- 46 Becht, E. *et al.* Dimensionality reduction for visualizing single-cell data using UMAP. *Nat*
Biotechnol **37**, 38–44 (2019).
- 47 Storey, J. D. A direct approach to false discovery rates. *Journal of the Royal Statistical*
Society Series B-Statistical Methodology **64**, 479-498 (2002).
- 48 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).
- 49 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).
- 50 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).
- 51 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).
- 52 Guo, Z. V. *et al.* Procedures for behavioral experiments in head-fixed mice. *PLoS One* **9**,
e88678 (2014).
- 53 Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51-56
(2012).
- 54 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).