

Hierarchy between forelimb premotor and primary motor cortices and its manifestation in their firing patterns

Akiko Saiki-Ishikawa, Mark Agrios, Sajishnu Savya, Adam Forrest, Hannah Sroussi, Sarah Hsu, Diya Basrai, Feihong Xu, Andrew Miri 

Department of Neurobiology, Northwestern University, Evanston, United States

 https://en.wikipedia.org/wiki/Open_access

 Copyright information

Reviewed Preprint

v2 • May 14, 2025

Revised by authors

Reviewed Preprint

v1 • November 11, 2024

eLife Assessment

This study provides **important** insights as to how interacting brain areas produce movement during the execution of a skilled multi-directional reaching task. Using a combination of single neuron and neural population analysis, optogenetic stimulation, and computational models, the authors provide **convincing** evidence of an asymmetrical influence between mouse premotor and motor cortex during the execution of a well-practiced behaviour. This asymmetry can only be captured by some but not all population analysis methods, which is a key lesson to the field in and of itself. Analyzing how activity that is shared and private to these areas relates to different aspects of movements, and why different methods provide different outcomes regarding the nature of inter-area interactions would further strengthen this work.

<https://doi.org/10.7554/eLife.103069.2.sa4>

Abstract

Though hierarchy is commonly invoked in descriptions of motor cortical function, its presence and manifestation in firing patterns remain poorly resolved. Here we use optogenetic inactivation to demonstrate that short-latency influence between forelimb premotor and primary motor cortices is asymmetric during reaching in mice, demonstrating a partial hierarchy between the endogenous activity in each region. Multi-region recordings revealed that some activity is captured by similar but delayed patterns where either region's activity leads, with premotor activity leading more. Yet firing in each region is dominated by patterns shared between regions and is equally predictive of firing in the other region at the single-neuron level. In dual-region network models fit to data, regions differed in their dependence on across-region input, rather than the amount of such input they received. Our results indicate that motor cortical hierarchy, while present, may not be exposed when inferring interactions between populations from firing patterns alone.

Introduction

A hierarchical organization characterized by feedforward influence between neuronal populations is commonly imputed to neural systems^{1–3}. In the motor system, a range of observations spanning anatomy^{4,5}, lesion^{6,7}, activity^{8,9}, and activity perturbation^{10,11} have been interpreted as reflective of a hierarchical organization in which premotor regions (PM) plan movements and exert their influence through primary motor cortex (M1), which executes them^{12,13}. Observation of substantial PM projections to brainstem and spinal cord^{14–16}, bidirectional connection between PM and M1^{17,18}, and activity related to movement planning in M1 and spinal cord^{19,20} together inform an updated view of a partial hierarchy^{21–23}. In this view, PM and M1 interact and each drive downstream motor circuits, though M1 exerts stronger influence on movement execution. Though “hierarchy” has divergent meanings in neuroscience discourse²⁴, here we refer specifically to asymmetric reciprocal influence between two neuronal populations, with activity in one exerting a larger functional influence on activity in the other, in line with certain contemporary usages^{23,25,26}. The empirical foundation for the view of a PM-M1 hierarchy remains incomplete. Such a hierarchy is realized when naturally occurring (endogenous) firing patterns in PM exert a larger influence on those in M1 than vice versa. However, existing observations have not yet resolved this asymmetric reciprocal influence between endogenous activity in PM and M1. Notions of PM-M1 hierarchy are instead based on observations like those of asymmetric effects on each region from electrical stimulation, lesion, or pharmacological inactivation of the other region^{27,28}; of differing laminar targets of projections between regions^{29,30}, and of differing degrees of activity in each region related to movement preparation^{31–33}. For example, in anesthetized monkeys³⁴ and rodents³⁵, stimulation in PM alters the effects of stimulation in M1 on muscles. Other work has used analysis of simultaneously measured activity in PM and M1 to support claims of feedforward PM-to-M1 influence^{9,36,37}. However, these sorts of observations do not necessarily imply asymmetric reciprocal influence of endogenous activity in PM and M1; such influence is not resolved by lesion or pharmacological inactivation, and need not match the influence of artificially induced activity or agree with activity correlations or anatomical connectivity. Moreover, other observations from electrical stimulation¹⁸ and simultaneous activity measurements in PM and M1^{38,39} are consistent with relatively symmetric reciprocal influence between PM and M1. Recent modeling of such activity measurements also suggests that PM-M1 interactions could change across task phases⁴⁰. It has also been proposed that differences in activity related to movement preparation could reflect involvement of each region in somewhat different aspects of movement^{41,42}, rather than hierarchy.

Because it is unclear if and when asymmetric reciprocal influence between PM and M1 exists, it also remains unclear how firing patterns reflect such influence. It is frequently assumed that feedforward signal flow from one region to another can manifest as the appearance of an activity pattern in one region and later in the other at a latency reflecting axonal conduction and synaptic transmission^{43–46}. Yet movement-related activity patterns in PM and M1 are diverse⁴⁷, and it remains to be seen how well activity in these regions can be captured by similar but temporally offset (delayed) activity patterns. Feedforward influence between regions is also thought to manifest in the ability to predict firing patterns in one region from those in the other. The application of time series prediction methods for discerning functional influence between neurons^{48,49} is grounded in the notion that feedforward influence leads to predictive relationships between firing patterns. We might expect that a partial hierarchy between PM and M1 would manifest as an asymmetry in firing pattern predictivity, where PM activity patterns share a stronger predictive relationship with subsequent M1 activity patterns than vice versa. However, local circuit dynamics can also be expected to play a substantial role in determining firing patterns. Recent theoretical results suggest that local circuit dynamics can be the primary

determinant of firing patterns, even in the presence of substantial across-region input like that between motor cortical regions^{50,51}. Thus, whether functional hierarchy manifests in firing pattern predictivity at the scale of brain regions within neural systems remains unresolved.

We sought to establish (1) hierarchy between PM and M1 mediated by their endogenous firing patterns, and (2) how any such hierarchy manifests in the firing patterns in each region. To address the former, during reaching in mice we examined interactions between the forelimb PM (rostral forelimb area, RFA) and M1 (caudal forelimb area, CFA). To assess interactions between endogenous activity patterns on the fast timescales of synaptic communication between regions (oligosynaptic timescales), we used rapid optogenetic inactivation of either cortical region while simultaneously recording with large-scale multielectrode arrays (Neuropixels⁵²) in the other region. We found that influence between regions is asymmetric, indicating a partial functional hierarchy. To address (2), we then recorded simultaneously with Neuropixels in both regions and performed a range of analyses on measured firing patterns. Consistent with partial hierarchy, we found that some activity can be captured by similar but delayed patterns in which either region's activity led, with RFA activity leading more often. However, we also found a high degree of similarity between firing patterns in both regions and that firing patterns in each region had similar predictive relationships with subsequent firing in the other region at the single-neuron level. In dual-region network models fit to our activity measurements with and without optogenetic inactivation, regions differed in their dependence on across-region input, rather than the amount of such input they received. These results suggest that functional hierarchy, while present among endogenous activity patterns in RFA and CFA during reaching, is not reflected in firing patterns as often assumed. This has important implications for contemporary attempts to infer interactions between populations from activity patterns alone.

Results

RFA and CFA activity during directional reaching

We first trained mice to perform a voluntary limb movement task that we expect to depend on activity in both RFA and CFA. Building on recent work⁵³, we developed a head-fixed directional reaching paradigm in which mice learn to rest their right hand on a rung and then reach to one of four spouts to grab a water droplet (**Figure 1A,B**; Movies S1,2). Rung touch illuminates a visual cue indicating the spout where the droplet will be dispensed after a subsequent (1-3 s later) auditory 'Go' cue. If mice moved their hand from the rung before the Go cue, the trial was aborted. The Go cue marked the start of a one-second response period, after which suction removed any uncollected water from the spout. Mice were acclimated to head fixation and trained to perform the task using methods similar to those previously described, a process which generally took between 10 and 20 daily sessions. To measure motor output during reaching, electromyographic (EMG) electrodes were chronically implanted into four muscle groups in the right forelimb: elbow flexors and extensors, and wrist flexors and extensors⁵⁴ (**Figure 1C**).

We next aimed to show that RFA and CFA output influence muscle activity at oligosynaptic latency during this directional water reaching. In trained VGAT-ChR2-EYFP mice⁵⁵ performing the reach task, we briefly projected a small spot of blue light (50 ms, 1.5 mm diameter, 9 mW/mm²) onto the surface of left RFA or CFA, triggered on reach onset. This approach silences the vast majority of endogenous activity in projection neurons at least down through layer 5^{54,56}, including those projecting to the brainstem and spinal cord; we show below that direct optogenetic effects were only seen in the targeted forelimb area and not the other. Activity was simultaneously recorded from right forelimb muscles to measure changes in motor output. For inactivations of either region, the absolute difference in activity summed across all muscles deviated from 0 very quickly after light onset (**Figure 1D,E**; **Figure S1**). This absolute difference was significantly different from zero when averaged over the first 25, 50, and 100 ms after light onset (**Figure 1F**; for RFA,




Figure 1.

RFA and CFA influence on forelimb muscles during directional reaching in mice.

(A) Schematic depicting the directional reaching task.

(B) Schematic depicting the time course of experimental control signals and muscle activity during the directional reaching task.

(C) Trial-averaged activity of four recorded muscles (black) $\pm$ SEM (gray, $n = 176$ trials) for one mouse during directional reaching towards each of four spouts. Separate averages are aligned on reach onset (Reach) or spout contact (Grasp). Vertical scale bars in (C) and (D) reflect standard deviation of z-scored muscle activity.

(D) For one example mouse, mean $\pm$ SEM muscle activity for trials without (black) or with inactivation (50 ms, cyan bar) of RFA (top, magenta) or CFA (bottom, blue) triggered on reach onset. Left image shows the position of the light stimulus on RFA and CFA.

(E) Mean $\pm$ SEM absolute difference between inactivation and control trial averages across all recorded muscles ($n = 12$ from 3 animals) for inactivation (50 ms, cyan bar) of RFA or CFA. For baseline subtraction, control trials were resampled to estimate the baseline difference expected by chance. This subtraction leads to values below zero.

(F) Absolute difference between inactivation and control trials averaged over three epochs after light/trial onset, for individual animals (circles) and the mean across animals (black bars).

$p = 9.6 \times 10^{-8}$ for 25 ms, $p = 1.3 \times 10^{-10}$ for 50 ms, and $p < 10^{-10}$ for 100 ms; for CFA, $p = 3.0 \times 10^{-5}$ for 25 ms, $p < 10^{-10}$ for 50 ms, and $p < 10^{-10}$ for 100 ms; one-tailed t-tests). This indicates that both RFA and CFA exert oligosynaptic influence on muscles during reaching.

We next examined whether neuronal firing patterns in the premotor RFA and primary motor CFA showed a commonly observed feature in primates and rodents alike^{9,19,36,57}, earlier activity preceding movement in the PM region. We acquired simultaneous large-scale multielectrode array (Neuropixel) recordings across cortical layers in both CFA and RFA in mice as they performed our reaching task (21 sessions across six mice, 74–293 qualifying single units per session in CFA, median 150; 17–194 per session in RFA, median 105; **Figure 2A–C**). As expected, the majority of recorded neurons in each region exhibited an elevated average firing rate during movement as compared to periods when forelimb muscles were quiescent, both for narrow-waveform, putative interneurons and wide-waveform, putative pyramidal neurons (**Figure 2D,E**; **Figure S2A,B**). The firing rate time series of a large fraction of neurons in each region were significantly correlated with the activity time series of at least one forelimb muscle (**Figure 2F**). For all subsequent analysis of simultaneous RFA and CFA recordings reported here, to ensure reach behavior was similar across analyzed trials, we only included trials where mice successfully reached to the correct spout. We also excluded outlying trials in which the reach duration, reaction time, or baseline muscle activity was more than three standard deviations above the mean for successful trials, as well as trials in which the initial rate of change in muscle activity was more than three standard deviations below the mean for successful trials (see Methods).

We then quantified the activity change around reaching onset in each region. We used the absolute change in trial-averaged neural activity state from a pre-reach baseline summed over top principal components (PCs) that together captured >95% of the variance (2–3 PCs; **Figure S2C,D**). We measured the time at which the activity state deviated from baseline preceding reach onset, which occurred earlier in RFA than in CFA (**Figure 2G,H**; $p = 0.031$, $n = 6$ mice, one-tailed signed rank test); the mean time for RFA was 53.9 ms earlier than in CFA. To ensure differences in onset were not the result of differences in baseline variance, we repeated the analysis using the same absolute threshold for each region. Here we used the higher of the two thresholds from each region. We found that RFA onset still preceded CFA onset ($p = 0.031$, $n = 6$ mice), but the timing difference was reduced, with RFA activation preceding CFA activation by 28.8 ms on average. We also performed a similar analysis on the trial-averaged firing rates of individual neurons. To better estimate onset times, we focused on neurons with mean firing rates above the 90th percentile value across all sessions within each animal. We found that the mean activity onset time for RFA neurons preceded that for CFA neurons (**Figure S2E,F**; $p = 0.031$, $n = 6$ mice). Finally, we found that the fraction of activity variance preceding movement onset was significantly higher for RFA than CFA (**Figure 2I**, $p = 0.016$, $n = 6$ mice). Thus, in our reaching paradigm, RFA exhibits earlier pre-movement activity compared to CFA.

Asymmetric reciprocal influence of endogenous RFA and CFA activity

We then sought to establish that the endogenous activity in RFA exerts a larger influence on CFA activity than vice versa. In trained VGAT-ChR2-EYFP mice performing the reach task, we inactivated RFA as described above while recording activity across layers with Neuropixels in CFA (**Figure 3A–C**, 13 sessions across three mice), or similarly inactivated CFA while recording in RFA (**Figure 3D–F**, 10 sessions across three mice). Blue light pulses were again triggered on reach onset. This will inactivate similarly-sized cortical regions that encompass all of RFA or most of CFA. Thus our approach provides a relatively comprehensive inactivation, enabling measurement of the influence each region has on neuronal firing in the other region.




Figure 2.

RFA and CFA activity during directional reaching in mice.

- (A) Mouse brain schematic depicting recordings in RFA and CFA.
- (B) Histograms showing the depth below pia of single unit isolations (neurons) recorded with Neuropixels and extracted with Kilosort. Dotted vertical lines depict where the cutoff for the bottom of each cortical region was defined.
- (C) Example of spike rasters (top) and muscle activity (bottom) recorded in one mouse during reach performance. Arrowheads indicate each onset (R) or spout contact (G).
- (D),(E) Scatter plots of the firing rates for neurons recorded in CFA (D) and RFA (E) during epochs when mice are activating their forelimb muscles, versus periods when all muscles are quiescent, separated for neurons that have wide and narrow waveforms.
- (F) The fractions of neurons recorded in both CFA and RFA whose firing rate time series was significantly correlated with that of at least one muscle (p -value threshold = 0.05), for individual animals (circles) and the mean across animals (black bars, $n = 6$ mice).
- (G) For one mouse, normalized absolute activity change from baseline summed across the top three principal components (PCs) for all recorded RFA or CFA neurons, and the top PC for muscle activity. Circles indicate time of detected activity onset. The pre-reach baseline epoch was from 150 ms to 100 ms before reach onset. Neural activity onset was detected as the first time at which activity rose 11 standard deviations above baseline.
- (H) Time from reach onset at which the activity change from baseline summed across the top PCs for RFA (magenta) or CFA (blue) neurons rose above a low threshold, for individual animals (circles) and the mean across animals (black bars, $n = 6$ mice).
- (I) The activity variance in the 150 ms before muscle activity onset, defined as a fraction of the total activity variance from 150 ms before to 150 ms after muscle activity onset, for each animal (circles) and the mean across animals (black bars, $n = 6$ mice).




Figure 3.

Asymmetric reciprocal influence of endogenous RFA and CFA activity.

(A-F) Schematics depicting inactivation and Neuropixel recording (A,D) and mean $\pm$ SEM firing rate time series for neurons from one animal (B,E) or all three animals (C,F) from inactivating RFA and recording CFA (A-C), or vice versa (D-F). The cyan bar indicates light on. The same number of cells were used from each animal in (C) and (F). The minimum firing rates for the inactivation trial averages are as follows: (B) 0.72 Hz, (C) 1.18 Hz, (E) 0.42 Hz, (F) 0.89 Hz.

(G),(H) Cumulative histograms of the difference between averaged z-scored firing rates for control and inactivation trials averaged from 45 to 55 ms after light/trial onset, for the top 50 highest firing rate neurons (G) or wide-waveform neurons (H) from each animal, combined.

(I),(J) Mean absolute firing rate difference $\pm$ SEM between control and inactivation trial averages for all (I) and wide-waveform (J) neurons recorded in the other area during RFA and CFA inactivation. The same number of cells were used from each animal. Baseline subtraction enables negative values.

(K),(L) Average absolute firing rate difference from 10 ms after light onset to 25, 50, and 100 ms after for all (K) and wide-waveform (L) neurons, for individual animals (circles) and the mean across animals (black bars).

We sought to verify that ChR2 activation in VGAT⁺ neurons was negligible in the region being recorded but not directly inactivated. We modified the Stimulus-Associated spike Latency Test (SALT⁵⁸) to test for significant effects of light on spike latency in the first 5 ms following light onset, specifically in narrow-waveform, putative inhibitory interneurons in the recorded region (**Figure S3A,B**). For each neuron, this yielded a p-value reflecting the likelihood that spike latency differences were due to chance (**Figure S3C-F**). The distribution of these p-values across all recording sessions was not significantly different from uniform, both for recordings in RFA and in CFA during inactivation of the other region (RFA: $p = 0.22$, $n = 328$ neurons; CFA: $p = 0.13$, $n = 207$ neurons; K-S test). We also observed that the effects on narrow-waveform neuron firing rates were similar in magnitude and time course to effects on wide-waveform, predominantly pyramidal neuron firing rates (**Figure S3G-R**). Thus we were not able to detect direct ChR2 activation effects on narrow-waveform neurons in the recorded region. We also verified that inactivation effects on firing in the downstream region were stable both within and across sessions

Results indicated an asymmetry in the oligosynaptic functional influence between CFA and RFA. Control and inactivation trial averages showed a larger reduction in CFA firing upon RFA inactivation compared to the effect on RFA firing upon CFA inactivation (**Figure 3A-F**). To quantify the effect on firing in individual neurons and to account for the possibility of both increases or decreases in their firing rates, we computed the difference of control and inactivation trial averages 50 ms after light/trial onset using z-scored firing rates for each neuron. To focus on the neurons with the best firing rate estimates, we used the neurons with the 50 highest average firing rates in each animal, either among all neurons, or all wide-waveform, putative pyramidal neurons. Distributions of the relative firing rate changes for these neurons showed that the vast majority of cells showed decreases in firing rate, and increases were rare, both for all neurons (**Figure 3G**) and for wide-waveform neurons (**Figure 3H**). These relative firing rate changes showed the substantial magnitude of reductions for many neurons; substantial fractions of neurons showed reductions that were more than half the standard deviation of their firing rate over the recording session. We note here that because we have focused on high firing rate neurons, we cannot distinguish whether the strong bias toward decreases in firing rate is unique to such neurons, though we see no reason to expect it to be.

Examining the average absolute change in firing over time, we again found a substantially larger effect for CFA neurons than RFA neurons, both for all neurons (**Figure 3I**) and the wide-waveform subset (**Figure 3J**). The average absolute change in firing was significantly larger whether calculated over the first 50 or 100 ms following light/trial onset, though it did not reach significance over the first 25 ms (**Figure 3K,L**; all neurons: 25 ms $p = 0.077$, 50 ms $p = 0.012$, 100 ms $p = 0.021$; wide-waveform: 25 ms $p = 0.101$, 50 ms $p = 0.027$, 100 ms $p = 0.022$; one-tailed t-test). Thus it is not the case that the smaller effect of CFA inactivation on RFA activity is solely due to offsetting increases and decreases in the firing rates of different neurons. These results demonstrate that the endogenous activity in RFA and that in CFA exert asymmetric effects on one another, with RFA activity exerting a comparatively larger effect on CFA activity. We interpret this as a direct indication of functional hierarchy on oligosynaptic timescales.

Premotor and primary motor cortical activity are highly similar during reaching

We then examined how the asymmetric reciprocal influence between CFA and RFA manifests in the relationship between their firing patterns. Previous measurements with calcium sensor imaging⁵⁹ or non-simultaneous array recordings⁶⁰ have shown substantial similarity between CFA and RFA activity during forelimb tasks. We therefore sought to quantitatively assess the degree of similarity at high temporal resolution from our simultaneous recordings in both regions during reaching. We used methods that decompose two sets of variables into pairs of components (linear combinations of the original variables) that are highly similar. We used both canonical correlation analysis (CCA^{61,62}), which maximizes the correlation between the time courses for

pairs of components, and partial least squares (PLS⁶³), which maximizes the covariance of these time courses. Both methods were applied to matrices comprising the trial-averaged firing rates for all neurons from a given animal (**Figure 4A** and S4A). In these matrices, the separate averages aligned on reach onset and on grasp for reaches to each of the four spouts were concatenated. This analysis was performed separately for each of six mice using all qualifying single units aggregated across sessions.

CCA revealed that the vast majority of the trial-averaged activity variance in RFA and CFA is captured by components with nearly identical time series. On average, over 80% of activity in both regions is captured by components whose times series had Pearson correlations > 0.94 (**Figure 4B-D**). We compared the cumulative activity variance captured by successive canonical variables with that captured by equivalent numbers of principal components (i.e. the maximum of any possible components). On average, the first ten canonical variables captured $92.2 \pm 2.4\%$ (mean $\pm$ SEM) as much as the first ten principal components for CFA, and $92.6 \pm 1.5\%$ as much for RFA. PLS identified components that successively capture nearly the same amount of activity variance as corresponding principal components while still maximizing the covariance of component time series (**Figure 4E,F**). The first ten PLS components captured $98.7 \pm 0.3\%$ as much as the first ten principal components for CFA, and $98.2 \pm 0.5\%$ for RFA. To establish a baseline level of similarity expected by chance for both methods, we repeated calculations after replacing the firing rate time series segments used to generate trials averages with separate, randomly chosen segments for one of the two matrices. Similarity was markedly reduced in both cases (**Figure 4B-F**). Thus both methods produced a similar basic result: much of the activity variance in both regions is parsimoniously explained by components that are very similar. This implies that trial-averaged firing patterns in RFA and CFA en masse are highly aligned. This similarity in firing patterns, coupled with the presence of descending projection neurons in both regions⁶⁴, suggests redundant and collective influence on motor behavior. Because our results here involve the vast majority of trial-averaged activity variance, we expect that they encompass both components of activity that vary for different movement conditions (condition-dependent), and those that do not (condition-invariant)⁶⁵.

In a functional hierarchy where one region's activity rises earlier than another's, we might expect the alignment between sets of firing patterns to be maximal when one set is lagged in time relative to the other. We thus assessed whether there was a time lag greater than 0 at which RFA and CFA activity patterns during reaching would best align. We found this was not the case. We shifted CFA activity in time relative to RFA activity by each integer value from -30 to 30 ms and recomputed CCA and PLS as above. In all cases, alignment quality was nearly identical: the average Pearson correlation for canonical variables was nearly constant for all lags tested (**Figure 4G**), as was the summed covariance of PLS component pairs (**Figure S4B**). However, control analyses in which we aligned CFA activity to itself, or to itself shifted by 10 ms in time, revealed clear maxima in alignment quality at the expected lags (**Figure S4C,D**). Thus, the asymmetric reciprocal influence between RFA and CFA and the earlier activity change in RFA do not imply a maximal alignment at a lag.

An imbalance in activity patterns shared at a lag

In a functional hierarchy, we might expect there to be activity patterns in the region exerting the larger across-region influence that then appear later in the second region because of synaptic time delays. We thus used a newly-developed dimensionality reduction method, Delayed Latents Across Groups (DLAG⁴³) that is designed to identify such patterns. We used DLAG to decompose simultaneously recorded CFA and RFA firing patterns into latent variables (components) that are either unique to activity in each region (within-region) or shared between regions at an arbitrary temporal lag (across-region). For sessions of simultaneous CFA and RFA recording during reaching, we decomposed each region's activity into four across-region components (negative lags defined as RFA leading) and four within-region components.




Figure 4.

RFA and CFA firing patterns during reaching are highly similar.

(A) From one animal, four canonical variables aligned at reach onset or spout contact (grasp) for reaches to each spout. Bottom row shows the corresponding time series for the first principal component of muscle activity.

(B),(C) Mean $\pm$ SEM cumulative variance capture ($n = 21$ sessions across 6 mice) for canonical variables (CVs, color), principal components (PCs, black), and CVs using shifted firing rate time series segments for one region as a control (gray), for RFA (B) and CFA (C) activity. Red annotations facilitate comparisons across B-D. Results in B-F all reflect $n = 21$ sessions across 6 mice.

(D) Mean $\pm$ SEM Pearson correlation for canonical variable pairs.

(E),(F) Mean $\pm$ SEM cumulative variance capture for PLS components (color), principal components (PCs, black), and PLS components using shifted firing rate time series segments for one region as a control (gray), for RFA (E) and CFA (F) activity.

(G) Mean $\pm$ SEM Pearson correlation of CFA and RFA canonical variable pairs computed when shifting CFA activity relative to RFA activity, averaged over all canonical variable pairs either without (purple) or with (black) weighting each correlation value by the average variance captured by its corresponding pair.

Figure 5A [↗](#) shows representative results from one mouse, plotting the lag and variance capture in each region for all across-region components identified for three sessions. Some across-region components had a lag that was not significantly different from zero according to a permutation-based statistical test⁴³ [↗](#); others had a lag that failed to converge between the boundary values of ± 200 ms. Both of these component types were ignored in subsequent analysis. Here we focused analysis on 15 sessions across the six mice, excluding sessions with a lower number of trials, which hampered DLAG calculations. We also focused on across-region components with lags between -60 and 60 ms, on par with the difference between the pre-movement activity rise time in RFA and CFA (see above, 53.9 ms). We found that activity variance in each region could be captured by across-region components in which either CFA or RFA activity led with lags in this range (**Figure 5B** [↗](#)). This agrees with the breadth of results cited above that suggest bidirectional interaction between PM and M1.

We then examined whether there was any significant bias toward across-region components in which RFA activity led. For across-region components with lags between -60 and 60 ms, we found such a bias (**Figure 5C** [↗](#)). The median lag for all across-region components was -16 ms. The median of median lags for individual sessions was significantly less than 0 ($p = 0.021$, one-tailed Wilcoxon signed-rank test). The same bias toward negative lags remained when including all components that had lags between -200 and 200 ms (**Figure S5A** [↗](#)). We considered whether this bias remained when accounting for how much variance each component captured in either region. It did, as the bias remained prominent when we weighted components by their variance capture in CFA (**Figures 5D** [↗](#), S5B; mean variance-weighted lag = -10.8 ms) or RFA (**Figures 5E** [↗](#), S5C; mean variance-weighted lag = -13.4 ms). Thus, across-region activity components in which RFA activity led were more prominent than those in which CFA led.

We then examined specifically whether, consistent with a functional hierarchy, across-region components in which RFA activity led captured more activity variance in CFA than CFA-led components captured in RFA. Variance capture in the lagging region by RFA-led and CFA-led components indicated this was the case. The same held when repeating this analysis but varying either the number of within-region (**Figure 5F** [↗](#)) or across-region (**Figure 5G** [↗](#)) components, while holding the other fixed at four. The same again held when considering all components that had lags between -200 and 200 ms (**Figure S5D,E** [↗](#)). We note that the bias we observe here should not be due to CFA firing patterns being noisier, as firing rates were on average higher in CFA. Collectively, our results with DLAG reveal activity patterns shared at a lag between RFA and CFA, with a bias toward patterns in which RFA activity leads.

Symmetric firing pattern predictivity

We then asked whether asymmetric reciprocal influence between CFA and RFA manifests in an imbalance in the ability to predict firing patterns in one region from those in the other at the single neuron level. Given the lack of consensus on how to assess this predictivity, for comprehensiveness we applied three metrics that each make different assumptions about how one firing pattern enables or improves prediction of another: point-process Granger causality⁴⁹ [↗](#), transfer entropy⁴⁸ [↗](#), and convergent cross mapping⁶⁶ [↗](#). Though these methods can be used to probe for potential causal influence between neurons, we stress that we are not making causal claims here and are focusing only on what these metrics directly compute: how the firing of one neuron enables or improves prediction of firing in another.

We computed each metric for pairs of neurons recorded simultaneously during reaching, one in each region, with one defined as source and the other as target. To focus attention on the timescale over which we expect oligosynaptic influence to manifest, we considered models that use source activity at different lags ranging up to 30 ms preceding firing in the target neuron and chose the lag that maximized predictivity. To focus calculations on cell pairs for which statistical power was greatest, we took the product of the average firing rates for neurons in each pair and found the 10,000 pairs with the highest product. To avoid directional bias due to differing firing rates (CFA




Figure 5.

DLAG finds imbalance in variance capture by activity patterns shared at a lag

(A) A scatter plot of the lags of all across-region activity components detected by DLAG versus their fractional variance capture in each region, for one mouse (3 sessions). Open circles reflect components where the lag was not significantly different from zero. Lines connect variance capture for individual components. Values for one component are not shown because its lag failed to converge between -200 and 200 ms.

(B) Fractional variance capture for CFA activity (left) and RFA activity (right) by across-region components in which CFA activity leads (CFA>RFA) or RFA activity leads (RFA>CFA), and by within-region components. Connected dots reflect the average across sessions for individual mice (n = 6).

(C) Histogram of DLAG component lags that were significantly different from zero and between -60 and 60 ms, for all recording sessions (n = 15). Red line indicates the median component lag.

(D),(E) Histogram of DLAG component lags weighted by the variance each component captures in CFA (D) or RFA (E), for all recording sessions (n = 15). Red lines indicate the mean component lag after weighting by variance capture.

(F),(G) Mean ± SEM fractional activity variance captured in the lagging region by DLAG components when the lag was significantly different from zero and between -60 and 60 ms, when varying the within-region (F) or the across-region (G) dimensionality (n = 15 sessions).

firing rates were slightly higher on average), we algorithmically adjusted this set of pairs to match the overall firing rate distributions for neurons used from each area (**Figure S6A,B**). In calculations of point-process Granger causality, two models are fit to the firing of a target neuron: one using the past history of all neurons in the ensemble, and another that excludes the source neuron. The difference in fit quality quantifies the unique prediction improvement provided by the source neuron, and a p-value is assigned that indicates the likelihood of a similar improvement by chance. Here the ensemble involved all CFA and RFA neurons from the given recording session that fell within the 10,000 pair set after the algorithmic adjustment. Models were fit to activity from a set of reach trials each spanning 200 ms before reach onset to 800 ms after. The distributions of p-values for predictions in each direction (CFA source, RFA target, and vice versa) showed a skew toward zero and were significantly nonuniform (**Figure 6A**; KS tests, $p < 10^{-10}$). This shows that for some fraction of pairs, source neuron activity did improve predictions of target neuron activity. However, the degree of skew toward zero was similar for both distributions, indicating that a similar fraction of pairs exhibited prediction improvements. We estimated the fraction of pairs showing prediction improvements from the p-value distributions, finding a slightly larger fraction when the CFA neuron was the source (CFA source: 69% of pairs, RFA source: 60% of pairs). Furthermore, the degree of prediction improvement in each direction was similar, as distributions of model fit quality differences for pairs with p-values below a significance threshold were similar (**Figure 6B**). Interestingly, these prediction improvements were similar to those computed instead for pairs of neurons within each region (**Figure 6C**), though there were more large extreme values within regions, potentially reflecting strongly coupled cells. We found similar results with both transfer entropy and convergent cross mapping.

Transfer entropy measures how well the spiking of the target neuron can be predicted when considering its past history as well as the past history of the source neuron. The more the source neuron's past history improves the prediction of the target neuron's spiking activity compared to the target neuron's own past history alone, the higher the transfer entropy value. In this case, to increase our statistical power, rather than using the trial segments used for Granger causality calculations, we used spiking during all epochs of movement throughout a given recording session. To assign a p-value for each cell pair, we generated an empirical null distribution by recomputing transfer entropy values after many different circular permutations of one neuron's spike time series for the given session. Here again, p-value distributions for predictions in each direction showed a skew toward zero and were significantly nonuniform (**Figure 6D**; KS tests, $p < 10^{-10}$), evidencing prediction improvement for some pairs. The degree of skew toward zero here was very similar for both distributions. The estimated fraction of pairs showing prediction improvements was less than 1% different between distributions (CFA source: 29.4% of pairs, RFA source: 29.1% of pairs). The degree of prediction improvement was similar in each direction (**Figure 6E**) and was fairly similar to that computed for pairs within each region (**Figure 6F**).

Convergent cross mapping is designed to uncover casual interactions in a dynamical system with a weak to moderate deterministic component and where casual variables do not contain unique information. Here we used it to quantify how well the activity of the source neuron can be predicted by the historical dynamics of the target neuron. Cross mapping quality ("skill") was computed for 5,000 high firing rate neuron pairs using a set of reach trials each spanning 200 ms before reach onset to 800 ms after. P-value distributions for predictions in each direction again showed a skew toward zero and were significantly nonuniform (**Figure 6G**; KS tests, $p < 10^{-10}$), and the degree of skew toward zero was similar for both distributions. The estimated fraction of pairs showing prediction improvements was just 2.6% different between distributions (CFA source: 99.6% of pairs, RFA source: 97.0% of pairs). The degree of prediction improvement was similar in each direction (**Figure 6H**) and was fairly similar to that computed for pairs within each region (**Figure 6I**). Collectively, these results indicate that firing pattern predictivity computed for neuron pairs does not reflect an asymmetry like that seen in the actual influence of endogenous activity.




Figure 6.

RFA and CFA firing pattern predictivity.

(A,D,G) P-value distributions from calculations of Granger causality (A), transfer entropy (D), and convergent cross mapping (G) using firing patterns of across-region neuron pairs.

(B,E,H) For across-region neuron pairs, distributions of metric values reflecting the improved prediction of target neuron firing using source neuron firing (B and E), or the improved prediction of source neuron firing using target neuron firing (H). Here and in (C,F,I), only values for which the corresponding p-value fell below a threshold (set to ensure the false discovery rate = 0.10) are included, and box plots on top show the minimum, 1st, 2nd, and 3rd quartile, and maximum values.

(C,F,I) Distributions of metric values computed instead for within-region neuron pairs.

A network model recapitulates findings without asymmetric across-region input

Finally, we sought to develop intuition about how the sort of asymmetric reciprocal influence between CFA and RFA we observed could arise from two coupled neuronal populations. One possibility is that asymmetric influence could depend on asymmetric across-region input, with RFA sending stronger input to CFA than vice versa. On the other hand, asymmetric influence could instead arise from a difference in the dependence of activity in each region on that across-region input. In this scenario, differences between the local circuit dynamics in each region would cause a different degree of robustness to the loss of input from the other region.

We thus constructed a recurrent network model to observe how it would capture our results. We built a network of two populations (“RFA” and “CFA”), each 80% excitatory and 20% inhibitory, with local recurrent (within-region) connectivity and sparser connectivity between populations (across-region) to align with experimental observations (**Figure 7A** [↗](#)). In response to a ‘Go’ signal, the network generates a muscle activity output. We trained instances of the model to generate measured muscle activity (summed across all four muscles) and measured neural activity (summed across all neurons recorded in a given region), in three cases: recording from either region while inactivating the other, and paired recording of both regions without inactivation. For the two inactivation cases, an additional input was provided to inhibitory interneurons to mimic ChR2 stimulation. Measurements for each case were drawn from different animals and 30 model instances were computed from different random initial conditions. Both within-region and across-region connection weights were free to vary during training.

Analysis of fit results showed that networks do not capture the asymmetric influence we observed with greater input from “RFA” to “CFA.” Fit quality was very good, as measured neural and muscle activity were closely fit and individual model instances readily exhibited both the temporal offset in activity rise and the substantial asymmetry in inactivation effects (**Figure 7B** [↗](#)). Models collectively had slightly larger “RFA” to “CFA” across-region weights (**Figure 7C** [↗](#)). However, repeating fits while constraining across-region weights to be balanced on average (“symmetric”) yielded very good fits, with only a 14% average increase in fit error (**Figure 7D,E** [↗](#)).

Most strikingly, the total amount of across-region input from each region, measured as the product of across-region weights and the activity of corresponding “presynaptic” neurons, was not higher for “RFA” to “CFA” connections, both for unconstrained and symmetric models (**Figure 7F** [↗](#)). The across-region input was actually larger on average in the opposite direction, regardless of the time window over which the input was summed. We also observed that the “RFA” network, which was more robust to the silencing of across-region input, had a higher average strength of within-region inhibitory connections (**Figure 7G** [↗](#); unconstrained: $p = 1.9 \times 10^{-12}$, symmetric: $p = 8.2 \times 10^{-4}$, two-tailed t-test). These results indicate that asymmetric reciprocal influence can emerge in these networks from local recurrent connectivity and the activity dynamics they engender.

Discussion

Here we have explored the functional influence between the two main forelimb regions in mouse motor cortex, and its manifestations in firing patterns during reaching. Using relatively comprehensive, optogenetic inactivation of either CFA or RFA while recording neural activity in the other, we found an asymmetry in the functional influence between regions, one aligned with existing views of a partial hierarchy between forelimb regions. Analysis of simultaneously recorded activity patterns in both regions detected an asymmetry in the presence of activity patterns shared at a lag. However, we also found that activity patterns in each region were in general very similar, and equally predictive of one another at the single-neuron level. Network




Figure 7.

Network modeling of RFA and CFA activity.

- (A) Schematic of the dual network model fit to activity measurements.
- (B) Illustration of fit quality for one instance of the model with unconstrained weights of across-region synapses. In (B) and (E), cyan bars show the epochs of simulated inactivation.
- (C) Distribution of synaptic weights for all across-region connections from all instances of the unconstrained model. In (C),(D) and (F), box plots show the minimum, 1st, 2nd, and 3rd quartile, and maximum values.
- (D) Box plots for the distribution of error across 30 instances each of the unconstrained and constrained models.
- (E) Illustration of fit quality for one instance of the model with across-region weights constrained to be equal in each direction on average.
- (F) Distributions of summed across-region input (weights x activity) in each direction over different simulation epochs, for all instances of both the unconstrained and symmetric model types.
- (G) Distributions of within-region excitatory (Excit) and inhibitory (Inhib) synaptic weights, for all instances of both the unconstrained and symmetric model types.

models that capture population-level firing features and the asymmetric reciprocal influence between regions did so without an asymmetry in the amount of across-region input. Our results suggest that motor cortical hierarchy, while present in mice during reaching, is not reflected in motor cortical firing patterns as often assumed, and that firing patterns alone may not always clearly reflect the functional influence between regions.

Functional hierarchy between forelimb regions

To properly interpret our activity perturbation results, it is important to consider how the effects of inactivation on connected regions may arise. Both CFA and RFA receive input from a broad range of other cortical regions, and each project to a similarly broad range of regions^{29,68,69}. We expect that the most numerous inputs to neurons in each region are locally derived. Local recurrent circuits can give rise to activity patterns that evolve over timescales equivalent to many synaptic delays. When we silence either region, activity is thus disrupted in diverse synaptic pathways leading to neurons in the other region, including pathways through intervening regions. These disruptions will influence activity at different latencies, depending on conduction delays and the number of intervening synapses. By silencing either region, we can measure something approaching the full influence of the given region on the other, albeit one mediated through diverse synaptic paths. When assessing interactions between regions, we believe it can be important to account for this diversity of synaptic paths. Here we observed a significantly larger influence of RFA on CFA than vice versa on all timescales examined, indicating that an asymmetry in reciprocal influence is present across timescales.

Our results highlight the dual determinants of inter-regional influence among brain regions: the balance of across-region input and the relative dependence of regions on that input. When we inactivate a particular region, the effects we observe in a downstream region will depend both on how much input it received from synaptic paths originating in the inactivated region and how robust firing patterns in the downstream region are to the loss of that input. Though classical descriptions emphasized the feedforward synaptic influence between regions in hierarchically arranged neural systems, observations of the relative prominence of local synaptic input^{70–72} and contemporary recognition of the activity dynamics it can engender^{73–75} indicate a relevance of robustness to input loss. Our knowledge of the relative robustness of networks in different cortical regions to input loss remains limited, but there are indications of variations in local recurrence that may induce variation in local network robustness^{76–78}. Decoupling of functional influence between regions and their firing patterns like that we observed may be expected if local circuit dynamics are the primary determinant of inter-regional influence among brain regions.

These dual determinants of inter-regional influence should be kept in mind in interpreting the earlier pre-movement activity change we observed in RFA. Our findings thus call into question previous interpretations of earlier rising activity in PM as reflecting a hierarchy mediated by feedforward influence^{9,19,32,36}. Though our results show RFA has a stronger direct influence on CFA than vice versa, this may not depend on the earlier rising activity in RFA, but instead on a differing robustness to input loss. As has been proposed previously, the earlier rising activity in PM regions could reflect involvement in distinct aspects of movement, ones that require more pre-movement preparation²².

Manifestations of hierarchy in firing patterns

Our results indicate that asymmetries in the reciprocal influence between regions may not manifest at the level of firing patterns in ways we might have expected. In particular, we found a substantial imbalance in reciprocal influence but also that firing in each region was dominated by patterns shared between regions and was equally predictive of future firing in the other region at the single-neuron level. These results raise questions about contemporary attempts to infer interactions between regions through network modeling of activity patterns alone. To serve as

inputs that induce observed activity patterns, network models may exploit similar activity patterns seen in different regions. In a measurement regime where a very small fraction of relevant neurons are sampled, models could attribute influence that is actually local, but from unsampled neurons, to neurons with similar activity patterns observed in other regions. Models may then fail to reflect functional relationships between neurons that are revealed by activity perturbations. One way to at least partially address this may be to enforce in models a balance of across-region and within-region inputs that agrees with observations.

Using DLAG, we detected that some observed activity can be captured by similar but delayed patterns with either region's activity leading, and RFA's activity led to a greater extent. This generally agrees with the notion of a partial motor cortical hierarchy, with greater influence from the premotor RFA to the primary motor CFA than vice versa. It may then be tempting to interpret these activity components as mediating the asymmetric reciprocal influence we observed from optogenetic perturbations, but this remains supposition. We should also consider why other analyses we performed did not detect asymmetries like those we observed with DLAG. The failure to detect asymmetry with CCA and PLS applied after shifting one region's activity relative to the other's may indicate that the across-region components detected by DLAG, though differing in activity variance capture, can still accommodate similar correlation or covariance when activity is aligned at a lag. In addition, the firing pattern predictivity calculations we performed for pairs of neurons may belie features of firing that emerge at the population level and are detected by methods like DLAG. Importantly, the predictivity measures we used here do not account for the fraction of population-level activity variance each neuron's firing pattern reflects. Consider a neuron whose time-varying firing rate is twice that of another: its firing pattern could be roughly equally predictive of any other pattern (any effect of the increased number of spikes could be relatively small) but it would reflect twice as much population-level activity variance.

It is commonly assumed that feedforward influence between regions will create predictive relationships between the activity patterns in the upstream and downstream regions. However, we have found that this does not hold on the scale of neuronal populations in the motor cortex during reaching in mice. This helps explain previous findings from firing pattern analysis that have seemed at odds with notions of even a partial motor cortical hierarchy. For example, analyses of activity in PM and M1 have described highly similar firing patterns^{39,79,80}, similarities in preparatory activity⁸¹, and symmetry of modeled reciprocal input³⁸. Our results show that these observations do not contradict an actual imbalance in reciprocal functional influence between regions. In general, our results from firing pattern analysis suggest that care should be taken when interpreting the results from any single analysis in relation to system function. The accurate inference of interactions between populations may require accounting for more than just simultaneous activity patterns.

We note here that our firing pattern analysis has involved all recorded activity; it is possible that results may differ if we focused specifically on a subset of activity components that correlate strongly with muscle activity. However, if the symmetry we observed for firing pattern predictivity or CCA/PLS was replaced by an asymmetry for this subset of activity components, our results would imply an opposite asymmetry for the remaining activity components, which may be surprising.

Models of PM-M1 interaction

From many different random initial conditions, our dual-population network fitting converged on solutions that well-fit mean firing patterns across our three dataset types – recording from either region while inactivating the other, and paired recording of both regions without inactivation – and exhibited a difference in the dependence of activity in each region on across-region input. Here we were not able to directly fit the observed activity of individual neurons, since our three dataset types were collected in different mice. However, our results provide a proof of principle that population-level features we observed, like asymmetric reciprocal influence and earlier

activity change in one population, can emerge from local population dynamics, rather than asymmetric across-region input. This suggests that suppositions about how functional influence between populations may be reflected in firing patterns may not survive empirical testing, as the dominance of local population dynamics can decouple firing patterns from the functional influence between populations. These results also suggest caution in interpreting the relative strength of inputs between populations in network models, whether input weights or weights times activity, in relation to system function.

Our findings comport with recent theoretical results suggesting that local circuit dynamics can be the primary determinant of firing patterns in the presence of substantial across-region input^{50,51}. Across-region connections may instead function to coordinate firing patterns across regions or modulate higher-order firing pattern features, like their autocorrelation⁷⁷. Firing patterns may be defined primarily by other constraints, such as the need to generate appropriate motor output patterns. The propagation of local activity in the service of this systemic pattern generation could be another constraint, as a number of recent observations suggest that local circuit dynamics govern motor cortical firing^{82,83} to some extent⁸⁴. Our results underscore the recognized need⁵¹ for developing theory about across-region interactions.

Methods

Experimental Animals

All experiments and procedures were performed according to NIH guidelines and approved by the Institutional Animal Care and Use Committee of Northwestern University. A total of 26 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: 21 VGAT-ChR2-EYFP line 8 mice (B6.Cg-Tg(Slc32a1-COP4*H134R/EYFP) 8Gfng/J; Jackson Laboratories stock #014548); and 5 C57BL/6J mice (Jackson Laboratories stock #000664). All mice used in experiments were individually housed under a 12 hr light/dark cycle. At the time of the measurements reported, animals were 17–22 weeks old. Animals weighed approximately 23–28 g. All animals were being used in scientific experiments for the first time. This includes no previous exposure to pharmacological substances or altered diets.

Directional Reaching Task

We modified a recently published directional reaching task⁵³. Head-fixed male mice were trained to reach to one of four spouts to grab a water reward they could then bring to their mouth and ingest. Mice initiate trials by placing their right hand on a rung (during training) or relaxing their forelimb muscles (during neural recording) for a period randomly chosen for each trial between one and three seconds (rest period). One of four LED lights in front of the mouse illuminates at trial onset, the location of which corresponds to the randomly-selected spout location where water will be dispensed on the given trial (**Figure 1A**). If the forelimb remains at rest for the duration of the rest period, a 4 kHz tone (‘Go’ cue) sounds for 100 ms, a water droplet is dispensed, and the mouse is free to reach out and grasp it. If the mice move their forelimb before the rest period ends, a 400 Hz buzzer sounds for 50 ms, the LED turns off, and a 200 ms delay must pass before another trial can be initiated. If the water droplet is not retrieved within 1 second (response period), it is removed by a suction tube immediately beneath the dispensation spout and the buzzer sounds. If mice reach to the correct spout, a 2 s consumption period is imposed before a subsequent trial can be initiated. If the mice reach to the incorrect spout first, the buzzer sounds for 50 ms.

Apparatus

The training apparatus was housed inside a sound attenuating chamber (H10-24TA, Coulbourn). Head-fixed mice were positioned within a 3D printed enclosure with sections removed to allow fixation of the mouse's headplate to a headplate holder and to allow the right hand access to the rung and spouts. Enclosures also had a second rung for the left (non-reaching) hand and a divider below the mouse's chest that extended in front of the mouse to prevent the left hand from gaining access to the spouts. The waterspouts (blunted 21G needles) were positioned in front of the mouse on its right side in a diamond configuration. The waterspouts were 6 mm apart vertically and horizontally. Suction tubes (blunted 21G needles) that were 1.5 mm shorter were attached below each spout. The spouts were secured by a 3D printed holder and their position was adjusted using a three-axis manual micromanipulator (UMM-3C, Narishige). The spouts and suction tubes were connected to solenoid valves (161T012, Neptune Research) through flexible tubing (EW-06422-01, Cole-Parmer).

Waterspout and rung touches were detected with capacitive touch sensors (AT42QT1011, SparkFun) during training. Capacitive touch sensors could not be used during neural recording since they cause artifacts in recorded voltages. Thus during neural recording, spout touches were instead detected with infrared beam sensors positioned to detect a hand just in front of each spout (FT-KS40 and FX-502, Panasonic), and right forelimb muscle activity measured with EMG electrodes was used to initiate trials. Mice initiated trials by reducing muscle activity below a threshold set so that any limb movement during the rest period before the Go cue would cause a threshold crossing and abort the trial. The buzzer was also removed during neural recording.

Experimental control was performed using the Arduino Due. Four speakers and an Arduino Uno were used to play the reward tone. Four green LEDs (1.8 mm diameter) were secured by an LED holder in a diamond configuration at 10 mm apart vertically and 20 mm apart horizontally. The LEDs were positioned approximately 10 mm away from the front of the mouse's eyes and at the same height.

Training

Under anesthesia induced with isoflurane (1%–3%; Covetrus), mice were outfitted with 3D printed head plates (24 x 20 x 5mm) 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, or the skull was covered with clear cement to maintain the visibility of bregma.

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 and head-fixation using a modification of established procedures⁸⁵. After a day of acclimation to handling, mice were acclimated to head-fixation and the reaching task over 1–4 daily sessions during which they were head-fixed in the reaching apparatus and provided water rewards.

During head-fixed acclimation, the water droplets were dispensed from all 4 spouts at random intervals (5.5–6 s). The four spouts were placed in front of the mouth and all spouts dispensed a 5 μ L water droplet signaled with the reward tone. The four LEDs all turned on while the water was being presented. The water droplets were automatically removed by suction at the end of the interval regardless of whether the mouse collected it or not. Mice freely licked the waterspouts and quickly learned the association between sound and reward. Once this was learned, the waterspouts were moved to a lower right position that allowed easy access for the right hand.

Mice spontaneously performed reach-to-grasp movements to collect and consume the water droplets by licking their hand. Almost all mice completely switched to the reaching behavior during the first or second session.

Following acclimation, mice underwent a daily 60 min training session to learn to initiate trials by touching the rung and to associate the location of the illuminated LED with the waterspout location where reward would be dispensed. To encourage mice to reach to all four spout locations, the probability of each spout being selected for water dispensation on a given trial was computed based on the percentage of successful reaches to the given spout for past trials during the given session. Probabilities were inversely proportional to success rates. In addition, two task parameters were adaptively changed during training sessions to shape mouse performance: the rest period was gradually lengthened (starting from 0.1 s), and the response period during which the water was available was shortened (starting from 10 s). Over 7 to 18 daily training sessions, mice learned to associate the illuminated LED location with the spout location, to wait during increasingly long rest periods, and to reach within increasingly short response periods. The rest period gradually became longer until mice were able to wait for more than 2 s. The response period adaptively changed until they were able to reach within 1 s.

For some mice, upon reaching these behavioral thresholds, neural recordings were performed during subsequent training sessions ($n = 3$ mice, total of 10 recording sessions). For another cohort of mice, once they met these conditions, recordings were performed during subsequent sessions using an experimental control script that did not adaptively update task parameters. In this script, the probability of each spout being selected for water dispensation did not vary, the rest period randomly changed between 1 and 3 s in 0.2 s increments, and the response period was fixed at 1 s ($n = 3$ mice, total of 11 recording sessions). These data were analyzed separately but eventually combined because there was no substantive difference in the results.

In order to be considered “trained” and used for experiments involving optogenetic inactivation and/or Neuropixel recording, mice had to demonstrate the capacity to reach successfully to multiple spouts. Successful reaching to a given spout was defined as the mouse reaching to the spout much more often than expected by chance for trials rewarding the spout if a reach was attempted (trials where no reach was attempted were ignored). Chance was defined as 25% of trials with an attempted reach, since there were four spouts. The typical mouse reached this criteria for three of the four spouts.

EMG Recording

EMG electrodes were fabricated for forelimb muscle recording using established procedures^{54,86}. Briefly, each set consisted of four pairs of electrodes, each consisting 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 (33AC2364, 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 upper forelimb muscles and 4.5 cm for lower forelimb muscles. The ends of wires with bared regions had their tips stripped of insulation and then were twisted together and crimped inside of a 27-gauge needle that facilitated insertion into muscle.

Mice were chronically implanted with EMG electrodes during the surgery in which headplates were attached as described previously^{54,87}. Insertions targeted the biceps (elbow flexor), triceps (elbow extensor), extensor carpi radialis (wrist extensor) and palmaris longus (wrist flexor). As discussed previously, while our methods produce isolated recordings from antagonist muscle pairs, we cannot exclude the possibility that EMG recordings are influenced by the activity of nearby synergist muscles, since our methods do not readily allow for simultaneous recordings from synergist muscles in the mouse forelimb.

Recordings were amplified and bandpass filtered (1-75,000 Hz) using a differential amplifier (C3313, Intan technologies). Data was digitized and acquired at 30 kHz using the RHD2000 USB interface board and RHD USB interface GUI software (Intan technologies). Suprathreshold activity of any muscle was detected in this software to indicate forelimb movement. Typical thresholds were 150-500 μ V.

Optogenetic Inactivation

After VGAT-ChR2-EYFP mice reached proficiency after several days of training on the directional reaching task, dental cement above the skull was removed and a 2-2.5 mm diameter craniotomy was made above the left CFA or RFA. A thin layer of Kwik-Sil (WPI) was applied over the dura and a 4 mm diameter #1 thickness cover glass (64-0724, 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 small 0.5 mm diameter craniotomy was then opened over the other region for recording.

During subsequent behavioral sessions, a 400 μ m core, 0.39 NA optical patch cable (FT400EMT, Thorlabs) terminating in a 2.5 mm ceramic ferrule was attached to a micromanipulator (SM-25A, Narishige). We set the cable at a certain distance above the surface of the brain for each session using a micromanipulator to ensure that the cone of light emanating from the cable would project a spot of light 1.5 mm in diameter onto the surface of the brain. A Neuropixel probe was inserted into the open craniotomy and recordings were performed as described in the next section, except the agarose and paraffin were omitted to avoid covering the window above the other region. To attenuate firing throughout motor cortical layers, we used a 450 nm laser (MDL-III-450-200mW, Opto Engine LLC) to apply 50 ms pulses of light at an intensity of 9 mW/mm² to the brain surface. To inactivate the motor cortex near the outset of reaching, the light pulse was triggered when either the biceps or triceps EMG signal reached a threshold set to reflect activation above baseline inactivity (usually 100-500 μ V). Light was applied during a random 50% of the trials on which the stimulation conditions were met. Unstimulated trials were used as controls.

Neural Recording

Acute recordings using two Neuropixel 1.0 probes (Imec) were completed as mice performed the reaching task. To expose the recording area, dental cement above the skull was removed and a 2 mm diameter craniotomy was made over CFA and a 1 mm diameter was made over RFA. The exposed brain tissue was sealed with silicone elastomer (DentSilicone-V, Shofu, or Kwik-Cast, World Precision Instruments). A stainless-steel screw (U-1415-01, Hirotsugi-Keiki) was implanted above the contralateral cortex as a ground. The large tip electrode on the shank of Neuropixels was used as a reference. Before recording, the animal was head-fixed, the silicone elastomer was removed, and the Neuropixels were slowly inserted at an inclination of 30° in the coronal plane and 15° in the parasagittal plane using fine micromanipulators (SM-25A, Narishige). For CFA, probes were inserted between 1 to 2 mm lateral and 0.5 rostral to 0.5 mm caudal of bregma. For RFA, probes were inserted 0.35 to 1.35 mm lateral and 1.75 to 2.75 mm rostral of bregma. Once the probes were in place, the brain surface was covered by agarose gel (2% agarose-HGT, Nacalai Tesque) and a mixture of liquid and solid paraffin, to minimize the vibration of the brain. Data was acquired at 30 kHz using the PXIe Acquisition Module (Imec) and SpikeGLX software (Janelia Research Campus).

Quantification And Statistical Analysis

All analysis was completed in MATLAB versions R2021b or later (MathWorks).

EMG Processing and Analysis

With certain exceptions discussed below, EMG measurements were downsampled to 1 kHz, high-pass filtered at 250 Hz, rectified, and convolved with a Gaussian having a 10 ms standard deviation.

EMG During Optogenetic Inactivation

Before EMG trial averages were analyzed, outliers were removed. The total Euclidean distance between muscle time series segments from -50 ms to 0 ms before light/trial onset, summed across muscles, was computed (MATLAB function 'pdist2'). Control and inactivation trials were combined for each muscle. The distances were then averaged across trials. A threshold was set one standard deviation above the mean. Trials above this threshold were excluded and were not used in subsequent EMG analyses.

The absolute difference between control and inactivation trials was calculated using the difference between the means for EMG time series for each muscle. The fractional changes in the time series were corrected for the difference expected by chance due to the use of separate sets of trials, estimated as follows. On 1000 different iterations, we divided the control trials into random halves and similarly calculated an absolute difference time series using the two halves. The mean absolute difference time series across these 1000 iterations was computed, and this mean was subtracted from the absolute difference time series computed with the actual data. Lastly, we subtracted the mean value across the 20 ms preceding light/trial onset time from the resulting time series.

Identification of Reach Initiation using EMG

Muscle activity measured from EMG recordings was used to determine the time of reach onset. For this analysis, EMG across all recorded muscles was summed to obtain a single trace representing overall muscle activity. Reach epochs were first identified using the reward tone and port beam break sensor traces to find trials where the animal successfully obtained the water droplet. As the animal's reaching strategy could trigger an additional beam break at an incorrect port, successful reaches were defined as instances when the beam break at the correct port occurred within 50 ms of the first beam break at any port. As the rest period prior to the reward tone is at least 1 s in duration, the EMG baseline for each reaching trial was found by computing the average summed EMG between -500 ms and -100 ms relative to the reward tone. A global standard deviation from baseline during quiescence (no muscle activation) was identified by calculating the standard deviation of an arbitrarily selected quiescent period identified in the recording. A lower threshold for each trial is set to the EMG baseline plus 7 times the global standard deviation. The upper threshold is set to the 30th percentile EMG value across epochs from reward tone to beam break for all successful trials. The reach onset time was defined as the lower threshold crossing immediately preceding the first upper threshold crossing after the tone. The double threshold algorithm was used because on some trials, EMG would first go above and then return below the lower threshold before the reach, likely from a twitch or other involuntary reaction to the reward tone. Trials in which the EMG never crossed the upper threshold were excluded from further analysis.

Reaching Trial Exclusion

Due to the presence of successful reaching trials with outlying muscle activity traces, certain exclusionary criteria were implemented to curate a collection of reaching trials to be utilized for subsequent analysis. Reach duration was calculated by subtracting the reach onset time from the beam break time. Successful trials in which the reach duration was longer than the mean duration of all successful trials in the given session plus 3 times the standard deviation were flagged for exclusion. Reaction time was calculated by subtracting the reward time (solenoid command pulse

onset) from the reach onset time. Successful trials in which the reach onset was less than -50 ms were flagged. Negative reaction times were possible in sessions where a fixed rest period was used because the animal was able to predict the reward tone. Ramp time was calculated by subtracting the reach onset time from the second threshold crossing time. Successful trials in which the ramp time was longer than the mean ramp time for all successful trials in that session plus 3 times the standard deviation were flagged. Finally, successful trials in which the standard deviation of the trial's individual EMG baseline was greater than the mean baseline standard deviation for all successful trials in that session plus 3 times the standard deviation were flagged. The final set of qualifying successful trials for each session was obtained by removing any flagged trials.

Spike Sorting and Unit Curation

Putative spikes were detected and sorted with Kilosort3⁸⁸ using default parameters. Sorting was improved (i.e. single unit yield was increased) by excluding pathological channels before sorting.

These channels were identified using abnormalities in high frequency components. The fast Fourier transform (MATLAB function 'fft') was computed for each channel and the weights over the interval from 300 to 1000 Hz were summed. We observed that the magnitude of this sum varied smoothly across channels, with the exception of certain channels which were characterized by much larger, outlier values. After median subtracting over a window of 10 channels, pathological channels were identified using a threshold of the median plus or minus 3 times the standard deviation. Channels identified in this manner for exclusion were consistent across recordings for the same probe. We excluded between 5 and 10 channels per session.

Single units identified by Kilosort3 were further curated by removing those with abnormal refractory violations. Inter-spike intervals (ISIs) were calculated by taking the autocorrelation of the spike trains binned at 0.1 ms. A test statistic was generated by summing ISI frequencies from 0.3 to 1.0 ms and normalizing by the summed ISI frequencies between 10 and 50 ms to account for overall firing rate. Units with a test statistic greater than 0.18 were excluded from any further analysis. This threshold is based on the frequency of violations expected in a sorted unit resulting from two neurons that have Poisson spiking and with one of the units accounting for 90% of the spikes.

Depths of sorted units were analyzed to identify cortical cells. CFA has a thickness of approximately ~1500 μm , and RFA has a thickness of ~1800 μm . For both CFA and RFA recordings, the depth of the most superior unit was used as an approximation for the cortical surface. Units within 1500 μm of this superior unit were considered cortical cells and subsequently analyzed. Firing rates were estimated at each ms during recordings by summing Gaussians with a 10 ms standard deviation centered on each spike time.

Putative pyramidal cells were identified based on waveform width. Widths for each unit were calculated by finding the trough-to-peak duration of the assigned waveform template. A histogram of waveform widths collected across all dual-probe recording sessions exhibited a bimodal distribution that could be well approximated by the sum of two gaussians. Using this model, we defined a width threshold that permits a 5% misclassification rate where tails of the fitted gaussians fall above or below threshold⁵⁴. A threshold of 0.417 ms was defined as the difference between wide-waveform and narrow-waveform cells.

Units were also classified based on whether their peak or trough had a greater magnitude. This was determined by analyzing the assigned waveform template and determining whether the maximum absolute value was originally positive or negative.

Identification of Neurons Significantly Correlated with muscles

In order to identify the neurons that significantly correlated with muscle activation, instantaneous firing rates were calculated from spike trains by smoothing each train with 10 ms standard deviation Gaussian. This allowed the correlation to be measured between a neuron's instantaneous firing rate time series with each muscle's recorded EMG activity. A bootstrapping approach was implemented to determine which of these correlations were significant and not by chance. Spike trains were circularly shifted at least ten seconds and no more than ten seconds less than the duration of the experiment (to ensure that the circular shift adequately removed any temporal structure between the neurons and the muscles) and correlations were calculated again. This process was repeated 300 times producing a null distribution for each neuron's correlation 'by chance' to each muscle. The observed correlations for each muscle was measured against the null distributions using a two-tailed test (to find both correlated and anti-correlated neurons) and if any p value corresponding to a muscle was significant ($p < 0.05$) the neuron was labeled as significantly correlated to muscle activity.

Optogenetic Inactivation Effects on Firing Patterns

For analysis of optogenetic inactivation effects, outlier trial exclusion was done differently since the activity perturbation may affect trial outcome. This exclusion was done by combining both inactivation and control trials for each neuron. We determined the distance between firing rates over the 20 ms before laser/trial onset between all the trials (MATLAB function 'pdist2') under the assumption that before reach initiation, the firing rate series for each neuron should look similar.

These distances were then averaged across neurons for each trial. A threshold was set at the mean of these distances plus a quarter of the standard deviation. Trials falling above this threshold were excluded.

Trial-averaged neuronal firing rates for control and inactivation trials were baseline subtracted by subtracting the mean firing rate from -20 to 0 ms before light/trial onset from the entire time series for each neuron. Across-animal trial averages were calculated using the same number of neurons from each animal to prevent animals with more recorded neurons from dominating the averages. We first found which mouse of the three similarly inactivated had the lowest number of recorded neurons, n , and used n randomly chosen neurons from the other two mice.

The absolute difference between inactivated and control trials was calculated for individual neurons and averaged across similarly-inactivated animals for comparison. We subtracted the control trial average from the inactivation trial average for each neuron, and then took the absolute value. The averages again used the same number of neurons from each of the three mice based on whichever had the least amount recorded. The resulting average difference time series for RFA- and CFA-inactivated mice were baseline corrected by subtracting the mean change from -20 to 0 ms before laser/trial onset from the entire time series.

To verify that Channelrhodopsin2 (ChR2) stimulation of inhibitory interneurons only occurred in the target region and did not spread to the recorded region, we modified SALT⁵⁸, a method developed to detect light-responsive neurons in a statistically-based, unsupervised manner. A window size of 5 ms was used for baseline and test epochs and the analysis used a time resolution of 1 ms. In order to ensure that the resulting p-values followed a uniform distribution from 0 to 1 under the null hypothesis of no direct light effect, a random sample from the Jansen-Shannon divergence values obtained from the post-stimulus firing patterns was used for comparison to an empirical null distribution based on firing during non-stimulus epochs. This step replaced the step in the original algorithm where the median of the Jansen-Shannon divergence values was taken, which did not yield a uniform p-value distribution under the null hypothesis. For each recording session, a p-value was computed for the firing of each narrow- waveform neuron following light

onset, where a low p-value indicates a neuron more likely to directly respond to light. For our purposes, the approximately uniform distributions of p-values we observed (**Figure S3C-F**) demonstrate general conformance to the null hypothesis of no direct ChR2 activation in the recorded region.

Delayed Latents Across Groups (DLAG)

The application of DLAG used firing patterns during qualifying successful trials (see above). To avoid covariance matrix rank deficiency in fitting DLAG models, we excluded low-firing neurons by setting a threshold such that any neuron that spiked fewer times than half the number of trials was excluded. For remaining neurons simultaneously recorded in a given session, firing rate matrices were assembled for each region, where each element was a given neuron's firing rate average over a 20 ms window. Trials spanned from 200 ms before to 800 ms after reach initiation. Matrices were thus N neurons $\times$ 50 time bins. The DLAG model was then fit using these matrices for each trial. Positive delay values indicated latent variables where CFA led RFA, and a negative delay value indicated latent variables where RFA led CFA.

To probe the robustness of observed results, we varied the numbers of across-region and within-region latent variable dimensionalities (i.e. the numbers of latent variables). Models were fit to each session's data with every combination of across-region dimensionalities of 2 to 6 and within-region dimensionalities of 2 to 6 in each area; therefore, sessions were eligible only if each area had at least 12 remaining neurons ($n = 15$ sessions). Any across-region latent variable with a delay that failed to converge or reached the boundary values of ± 200 ms in 10,000 iterations was removed from the analysis.

To test if identified delays were significantly different from zero, bootstrapping was performed by sampling 100 trials from the set of existing trials with replacement, as in the originally published method⁴³. A delay was considered statistically significant if less than 5% of the bootstrapped samples performed just as well with a model with a 0 ms delay as they did with the model with the original calculated delay. Across-region latent variables with non-significant delays were removed from the analysis. For each session, we calculated the proportion of variance captured by across-region latents with statistically significant delays, grouping latents by the sign of their delays (i.e. which region led).

Canonical Correlation Analysis (CCA) and Partial Least Squares (PLS)

CCA (MATLAB function 'canoncorr') and PLS were performed on matrices comprising the trial-averaged firing rates for all neurons from an animal, concatenating averages aligned on both reach onset and spout contact for reaches to each of the four spouts (i.e. eight trial averages are concatenated). Averages spanned from 99 ms before to 100 ms after the alignment point, resulting in matrices of size 1600 time points $\times$ NCFA neurons and 1600 time points $\times$ NRFA neurons for each recording session, where NCFA is the number of included neurons from CFA and NRFA is the number of included neurons from RFA. We denote these matrices XCFA and XRFA in what follows. PCA was performed on these matrices and on average, 25 principal components were able to capture 95% of the neural activity variance, and so 25 was chosen as the number of PLS components and canonical variables for the subsequent analyses, enabling comparison to PCA results.

To mitigate matrix rank deficiency for CCA, principal component analysis was performed on the neural activity matrices and the first 25 principal components were then used for the CCA. Since the resulting 25 canonical vectors (CVs) for each data matrix are not necessarily orthogonal, they were orthogonalized in order to compute the additional neural activity captured by successive CVs, enabling comparisons with principal components. Weighted averages of Pearson correlations for all pairs of CVs were computed by weighting the Pearson correlation of each CV pair by the

average of the additional neural activity variance captured by each corresponding orthogonalized CV from the pair. To account for synaptic delay between RFA and CFA, lags were introduced in the same way as stated below for the PLS analysis.

For PLS, the PLS-SVD variant⁶³ was used. PLS-SVD is done by performing singular value decomposition on the cross-covariance matrix X^{TY} which in our case is $X_{RFA}^T X_{CFA}$, yielding

$$USV^T = X_{RFA}^T \cdot X_{CFA}$$

where the columns of U define axes in RFA activity space and the columns of V define axes in CFA activity space. The matrix S is a diagonal matrix where the n^{th} diagonal element is the covariance of RFA activity projected onto the n^{th} column of U and CFA activity projected onto the n^{th} column of V . The trace of S can be interpreted as how much total covariation between RFA and CFA activity is captured by the PLS components. A necessary additional step is to divide this number by the total variance in neural activity in these two matrices in order to normalize by how active these brain regions are during the reach and grasp epochs that are being analyzed.

$$c = \text{trace}(S) / (v_{RFA} + v_{CFA})$$

where v_{RFA} and v_{CFA} are the sum of the variances of the rows of X_{RFA} and X_{CFA} respectively. The original activity matrices can then be projected onto their respective axes (by $X_{RFA} * U$ for RFA, $X_{CFA} * V$ for CFA) to produce time courses in neural activity space such that corresponding pairs will maximally covary, i.e. the projection of X_{RFA} onto the k^{th} column of U and X_{CFA} onto the k^{th} column of V will have covariance equal to the k^{th} diagonal element of S .

PLS-SVD was similarly performed after shifting one activity matrix in time relative to the other. This was done by keeping the same X_{RFA} matrix for each animal, but constructing a new X_{CFA} matrix. The reach and grasp onset times were shifted by an amount between -30 and 30 ms and the trial epochs -99 to 100 ms around each shifted time point were extracted, trial averaged, and concatenated just as before. PLS-SVD was then performed and c was recomputed for each lag from -30 to 30 for each animal.

An additional control analysis was performed to test whether the similarity between CFA and RFA activity found with CCA and PLS was beyond what could be expected by chance for activity measured and processed according to our approach. To do this, we repeated either alignment method, keeping X_{CFA} the same, but shifting each reach and grasp onset used for aligning trials for X_{RFA} by a random number of milliseconds between 5000 and 10000 into the future. Checks were implemented to make sure that each shifted trial did not overlap with another (original or new) trial. This removes the behavior-related activity correlations between the original matrices used for alignment.

Firing Pattern Predictivity Measurements

These measurements were performed on a subset of simultaneously recorded neuron pairs from dual probe recordings. Only wide-waveform, putative pyramidal neurons having a maximum waveform deflection in the negative direction were used. Average firing rate was calculated for each neuron across all qualifying successful trials over the window from 200 ms before to 800 ms after reach initiation. The ten thousand neuron pairs with the highest average firing rate products across all sessions were identified. Note that pairs only included neurons that belonged to the same session, as simultaneous recording was required for the trial-by-trial predictivity analysis performed here.

After calculating the firing rate distribution of the CFA and RFA cells belonging to the top ten thousand pairs, CFA cells had higher firing than RFA cells on average. Firing pattern prediction can be heavily influenced by differences in firing rate, since statistical power can be related to the

number of spikes observed. Thus, we developed an algorithm to match the distributions of firing rates by excluding certain pairs. We first generated histograms of the log firing rate for the CFA and RFA cells belonging to the top ten thousand pairs binned at 0.05 log spikes per second. Note that the multiplicity of a given neuron in these distributions matched the number of pairs the given neuron appeared in. To start the algorithm, a firing rate bin with more CFA cells than RFA cells was chosen. We then randomly selected a cell pair that contained a CFA cell from this bin and an RFA cell from a firing rate bin which had more RFA cells than CFA cells. To better match the firing rate distributions, the high firing rate CFA cell was then replaced with a different CFA cell from the same session and firing rate bin as the selected RFA cell. The selection of new CFA cells was allowed to occur with replacement. This selection process was repeated until all firing rate bins contained the same number of CFA and RFA cells, or when no other cell pairs could be created to match the binned firing rate distributions.

From distributions of p-values calculated as described below for each metric, we estimated the fraction of false null hypotheses as one minus an estimate of the *a priori* fraction of true null hypotheses⁶⁷ (MATLAB function ‘mafdr’).

Transfer Entropy

Transfer entropy analysis was performed using a publicly available algorithm⁴⁸. This algorithm takes the spike trains of two neurons (one source neuron, one target neuron), and compares how well the spiking of the target neuron can be predicted when considering its past history as well as the past history of the source neuron. The better the source neuron’s past history improves the prediction of the target neuron’s spiking activity compared to the source neuron’s own past history alone, the higher the transfer entropy (TE) value. In this case, to increase our statistical power, rather than using the trial segments used for Granger causality described below, we extracted and used spiking during epochs of movement in the following way. First, rectified and filtered EMG recordings from the four muscles were summed. Then, a period of approximately one second of muscle quiescence was manually identified in this summed time series. A threshold was set at the average value over this period plus 7 times the standard deviation over this period, and movement epochs were defined as when summed muscle activity surpassed this threshold. To include brief reductions in muscle activity that happen during ongoing movement, any epoch of 100 ms or less below this threshold was reclassified as movement. Additionally, periods of movement less than 10 ms were reclassified as non-movement since this was unlikely to reflect meaningful movement. The neural activity during these movement epochs was then extracted and concatenated. However, when concatenating epochs, the discontinuous skips in time from the end of one epoch to the beginning of the next could, although very sparse, cause inaccurate calculations of the transfer entropy by counting spikes from one time point as causing spikes at a time point much farther in the future. To circumvent this issue, pads of zeros were inserted between each movement epoch so that no spikes in one epoch could be calculated as causing spikes in the next movement epoch. To calculate transfer entropy, we used $TE_{J \rightarrow I}(d)$ from eq. 5 in ref. 43. We calculated $TE_{J \rightarrow I}(d)$ for $d = 0, 1, 2, \dots, 30$ ms and took the maximum value as our measurement for the transfer entropy between two neurons, the observed TE value.

In order to determine the significance of the observed TE value for each neuron pair, we obtain a corresponding p value by creating empirical null distributions, i.e. TE values expected in the absence of coupled firing between the neurons. For each neuron pair, the concatenated spike train of the source neuron is circularly permuted (MATLAB function ‘circshift’) by values no less than 3 seconds and no more than T-3 seconds (where T is the duration of the concatenated spike train) so that any temporal relationship between the two spike trains is extinguished while preserving the spiking statistics and firing rate patterns of each individual spike train. After this permutation, the TE was measured as before. This process was done 300 times to form a null distribution of TE values for each neuron pair and a p value was calculated by taking the fraction of null TE values that were higher than the observed TE value.

Granger Causality

Point-process Granger causality with exogenous temporal modulations (GC⁴⁹), was utilized to assess directional bias in the predictivity of neuronal firing between CFA and RFA. Using GC, predictive models are generated to fit spiking data of a target neuron using the spiking data from the entire neuronal ensemble, both with and without the source neuron. In our conditions, this neural ensemble contains neurons from both CFA and RFA. The method yields a test statistic for each source-target neuron pair by computing the difference in accuracy of the prediction from these two models. The larger difference, the more unique information the source neuron contains about the spiking of the target neuron. This test statistic follows a chi-squared distribution, yielding p-values for each source-target pair.

Binary spike trains for each neuron belonging to the top ten thousand firing rate CFA/RFA neuronal pairs were generated. Only spiking data between 200 ms prior to reach initiation and 800 ms after reach initiation were used. GC was conducted on a per-session basis using a random selection of 40 qualifying successful trials (the same for all pairs) to limit computation time. Global regression of exogenous temporal modulations were optimized to bins of either 1 ms, 5 ms, 8 ms, 10 ms, 20 ms, or 25 ms. The duration of spiking history incorporated in the model was fixed at 30 ms to force all computed test statistics to follow the same chi-square distribution.

For certain source-target pairs, the model did not converge on the target spiking data within tolerance. We removed these pairs from the analysis. Additionally, for some pairs, the model would be overfit and predict the target data perfectly. We also removed these pairs from further analysis. As a consequence, 35% of the 10,000 pairs with the highest product of average firing rates were excluded from the GC analysis for these reasons. However, this should not affect our results in a meaningful way. While it could affect the degree of skew toward 0 in p-value distributions, here we are not ascribing significance to the degree of skew, just the similarity between the skew for the CFA source and RFA source conditions. We also do not draw distinctions between the results from different prediction methods that would be affected by only using a subset of neurons for GC.

Following these exclusions, we observed a preponderance of pairs with p-values greater than 0.95. We found that the firing rate product of these high p-values pairs were heavily skewed towards low values (**Figure S7A**). These low firing rate pairs likely did not meet the assumptions of GC and were removed from further analysis. Following this exclusion, the firing rate distributions of RFA and CFA neurons belonging to this final set of neuronal pairs were reanalyzed and found not to be appreciably different (**Figure S7B,C**).

The GC algorithm additionally computes the model difference test statistic for all within- region pairs. Using these values, the predictive power of the CFA and RFA intraregional neuronal population was also assessed.

Convergent Cross Mapping

Convergent Cross Mapping (CCM) was implemented using a modified version of MATLAB function ‘Sugi’⁸⁹. Here we used the algorithm to quantify how well the activity of the source (casual) neuron can be predicted by the historical dynamics of a target (effector) neuron. While TE and GC assume the observations underlie a purely stochastic system, CCM can uncover casual interactions in a dynamical system with a weak to moderate deterministic component and where casual variables do not contain unique information.

For each pair of source and target neurons, the historical dynamics of the target neuron can be represented by a "shadow manifold," where every point on the shadow manifold is constructed from time-lagged observations. Intuitively, we can think of nearby points on the shadow manifold

as having similar dynamics. Thus, if the source neuron has casual interactions with the target neuron, then activity of the source neuron at time t can be predicted from a cluster of nearby points at time t in the shadow manifold. CCM quantifies these casual interactions by finding the correlation coefficient between the predicted source activity calculated using a cluster of nearby points in the target manifold at time t and the actual source activity at the same time, a metric termed “cross-map skill”^{66,90}.

The preprocessing and trial extraction followed the same steps used for Granger causality calculations. Five thousand RFA-CFA neuron pairs were curated by taking the top ten thousand pairs described above and excluding the bottom 50% when sorted by lower firing rate of the pair. This avoided anomalous results from CCM calculations that seemed to depend on epochs with very few or no spikes. Spike trains were then binned at 1 ms and smoothed using a gaussian kernel with a sigma of 10 ms. For every neuron pair, 30 trials were randomly sampled from all qualifying successful trials to reduce computation time, and shadow manifolds were constructed from the previous 500 ms of activity in the target neuron within the same trial. CCM was then applied to these 30 trials to calculate a cross map skill value, quantifying the correlation between the time series of predicted source activity and the time series of actual source activity. To account for non-instantaneous interactions between neurons, this process was repeated at time delays up to 30 ms, incrementing every 3 ms (i.e. 0, 3ms..., 30 ms), keeping the maximal cross map skill value and its associated delay.

In order to determine whether these observed cross map skill values were significant, an empirical null distribution was generated by disrupting the temporal relationship between pairs of neurons. For every pair, we repeated the same circular permutation method used with transfer entropy (300 permutations), and applied CCM to generate cross-map skill values, making sure to use the same 30 trials and optimal time delay used to obtain the actual cross map skill for the given pair. P-values for each pair were obtained by counting the fraction of null cross map skill values that were greater than the observed cross map skill.

Applying CCM involved choosing several hyperparameters. E , the embedding dimension of the shadow manifold, can be thought of as the number of time lags that optimally captures the historical dynamics of the target neuron; if E is set to 4, then each dimension on the shadow manifold represents the neurons activity at times $[t, t-1 \text{ ms}, t-2 \text{ ms}, t-3 \text{ ms}]$. To determine the optimal value for E , we applied a simplex projection to each individual neuron in CFA and RFA to determine how many lagged dimensions best forecast a neuron’s own future activity⁹¹. This optimal embedding dimension for both CFA and RFA neurons was found to be 4. The number of nearby neighbors, K , searched for in the target shadow manifold was kept at 5 ($E+1$). $\tau\tau$, which indicates how many timesteps each shadow manifold dimension is lagged in time $[t, t-1*\tau\tau \text{ ms}, t-2*\tau\tau \text{ ms}...]$, was set to 1.

Network Modeling

A recurrent neural network was trained to replicate experimental data. The package PsychRNN⁹² was used to configure the network architecture, inputs, outputs, training, and simulation in the following way. Two neural populations of 500 units each, one representing RFA and the other representing CFA, were created with 80% excitatory units and 20% inhibitory units. The connection probability between neurons within each region was 5% and additional sparse connections were created from excitatory units in each region to neurons in the other region with a connection probability of 0.1%; no inhibitory units projected to the opposite region. A subset of 100 excitatory units from each region were chosen to be the units whose activity was trained to reproduce the observed muscle activity traces; these neurons were also stipulated not to have any projections to the opposite brain region. The simulated muscle activity output was the weighted sum of the activity of these 200 neurons, with weights allowed to vary during training. A second subset of 100 excitatory units from each region were chosen to be the units whose activity was trained to reproduce the experimental neural activity traces.

Experimental data used for training all models consisted of averaged neural or muscle activity spanning from 50 before to 100 ms after reach onset from successful reaches. Averages from one animal were used for each of three conditions (CFA inactivation, RFA inactivation, dual recording/no inactivation). However, since animals were used for only one type of experiment, training data reflected data from three different animals.

The simulated summed neural activity for each region was the weighted sum of each set of 100 neurons, with weights held constant to replicate how experimentally observed averages were calculated. To simulate activation of inhibitory interneurons expressing ChR2, all inhibitory units in the “inactivated” region also received a “light” signal as input. We modeled the light signal by concatenating two sigmoid curves, each with a maximum value of 1, using the following equation:

$$s(t) = \pm \frac{1}{(1+e^{-k \cdot t})}$$

The $\pm$ denotes that the sign of the numerator of the two sigmoids varies. The first sigmoid had a positive numerator and ramps up from 0 to 1, while the second had a negative numerator so that it starts at 1 and decays to 0. The first sigmoid had k set to 0.5 so that the curve started to ramp up (surpassing a value of 0.01) at reach onset and reached 0.95 15 ms later. The signal sustained its maximum value of 1 until 50 ms after reach onset, and the second sigmoid had $k = 1$ so that the signal had a value < 0.01 5 ms later. All units also received a ‘Go’ signal as input in each of the three conditions. The ‘Go’ signal similarly was a sigmoid function that started at a value of 0 and had $k = 0.3$, such that the curve started to ramp up (surpassing a value of 0.01) 40 ms before reach onset, reached 0.95 25 ms later, and then sustained its maximum value of 1 for the remainder of the trial.

Training consisted of up to 100,000 trials where each was randomly selected to be one of the three conditions. For each trial, the output of the muscle output units was measured along with the neural activity output, either from the non-inactivated brain region or both regions if the trial was a no-inactivation trial. The mean squared error between the model outputs and experimentally observed traces was computed and used as input to the cost function, which was minimized as model weights were updated. Weights for connections between units within and across regions were allowed to update but certain constraints were applied. Positive weights remained positive and negative weights remained negative in order to keep the number of excitatory and inhibitory units constant. Weights between units that were initialized to be 0 remained 0 throughout training to keep the number of connections within and across regions constant. In the symmetric case, an additional term was added to the cost function: the absolute value of the difference between the sums of the weights of the across-region connections in each direction (RFA $\rightarrow$ CFA and CFA $\rightarrow$ RFA). As mentioned previously, output weights from the muscle output units were allowed to fluctuate throughout training but output weights from the neural activity output units remained fixed. Biases on all units also remained fixed during training.

Supplementary figures




Figure S1.

Extended time series from Figure F1D.

For one example mouse, mean $\pm$ SEM muscle activity for trials without (black) or with inactivation (50 ms, cyan bar) of RFA (top, magenta) or CFA (bottom, blue) triggered on reach onset. Left image shows the position of the light stimulus on RFA and CFA.




Figure S2.

RFA and CFA activity during directional reaching in mice.

(A),(B) Histograms of waveform widths for recorded neurons in CFA (A) and RFA (B), showing the bimodal distribution of narrow and wide waveforms. Dotted line shows the threshold above which neurons were considered wide waveform.

(C) For four example neurons, spike rasters (top) and binned firing rate (bottom) across all successful reach trials to each spout, aligned on reach onset (black arrowheads) and spout contact (Grasp, gray arrowheads). Bottom row shows the corresponding time series for the first principal component of muscle activity.

(D) CFA (blue) and RFA (magenta) activity projected onto their respective first two principal components, from the recording involving the neurons shown in (A).

(E) For one mouse, a histogram of the time around reach onset when the firing rate of individual neurons rises above a pre-reach baseline, for RFA or CFA neurons. Vertical lines indicate the mean across neurons for each region. To avoid neurons with noisy mean firing rates, the analysis depicted in (E),(F) included only the neurons with mean firing rates above the 90th percentile value for each given animal. Onset time for each neuron was detected by finding the time at which its trial-averaged firing rate first rose above a threshold defined as 7 standard deviations above the mean for a pre-reach baseline epoch (150 ms to 100 ms before reach onset).

(F) Mean time from reach onset at which the firing rate of individual neurons rises above a pre-reach baseline for RFA or CFA neurons, for individual animals (circles) and the mean across animals (black bars, $n = 6$ mice).




Figure S3.

Effects of CFA and RFA inactivation on firing rates in the other region.

(A),(B) Histograms of waveform widths for recorded neurons in CFA (A) and RFA (B), showing the bimodal distribution of narrow and wide waveforms. Dotted line shows the threshold above which neurons were considered wide waveform. (C)-(F) Histograms of p-values from our modified version of SALT for narrow-waveform neurons recorded in one session (C,E) or all sessions (D,F) in either CFA (C,D) or RFA (E,F) while inactivating the other region. The uniformness of these distributions indicates an absence of appreciable violation of the null hypotheses that neurons are not directly activated by light. (G)-(N) Mean firing rate $\pm$ SEM for wide-waveform (G-J) or narrow-waveform (K-N) neurons for one animal (G,I,K,M) or three animals (H,J,L,N) recorded in CFA (G,H,K,L) or RFA (I,J,M,N) while inactivating the other region. Averages combining cells from multiple animals used the same number of cells from each animal. The cyan rectangle indicates when the light was applied. (O)-(R) Mean absolute firing rate change $\pm$ SEM between control and inactivation trials (O,Q) and mean absolute firing rate difference between control and inactivation trials averaged from 10 ms after light/trial onset to 25, 50, or 100 ms afterwards (P,R) for wide- and narrow-waveform neurons recorded in CFA (O,P) or RFA (Q,R) during inactivation of the other region. Black bars show mean across animals. (T)-(AA) Mean firing rate $\pm$ SEM for all three animals recorded in CFA (T,U,X,Y) or RFA (V,W,Z,AA) while inactivating the other region, either for two separate sessions (T-W) or the first and second half of trials from all sessions (X-AA). The cyan rectangle indicates when the light was applied. Average inactivation effects show remarkable consistency, both within and across sessions.

Figure S4.

PLS alignment of RFA and CFA activity.

(A) For the same animal used in **Figure 4A**, four PLS components aligned on reach onset or spout contact (Grasp) for reaches to each spout. Bottom row shows the corresponding time series for the first principal component of muscle activity. (B) Mean $\pm$ SEM summed covariance of CFA and RFA PLS components computed when lagging CFA activity relative to RFA. We found no lag where components exhibit an appreciably greater total covariance. (C) Mean $\pm$ SEM Pearson correlation of canonical variable pairs computed by aligning CFA activity to itself, but lagging one copy relative to the other, for the average over all pairs either without (purple) or with weighting each correlation value by the variance captured by the given pair (black). (D) Same as (C), but when initially shifting one copy 10 ms relative to the other. In this case, we expect the alignment to be maximal at a lag of -10 ms.




Figure S5.

Additional plots from DLAG calculations.

Results in this figure are similar to **Figure 5C-G**, but include all components identified by DLAG that have lags between -200 and 200 ms

(A) Histogram of DLAG component lags that were significantly different from zero and between -200 and 200 ms, for all recording sessions ($n = 15$). Red line indicates the median component lag. (B),(C) Histogram of DLAG component lags weighted by the variance each component captures in CFA (B) or RFA (C), for all recording sessions. Red lines indicate the mean component lag after weighting by variance capture. (D),(E) Mean $\pm$ SEM fractional activity variance captured in the lagging region by DLAG components when the lag was significantly different from zero and between -200 and 200 ms when varying the within-region (D) or the across-region (E) dimensionality ($n = 15$ sessions).




Figure S6.

Additional plots from firing pattern predictivity calculations.

(A),(B) Firing rate distribution for CFA (blue) and RFA (magenta) neurons belonging to the 10,000 pairs with the highest firing rate product across all 21 recording sessions, before (A) and after (B) applying our algorithm for equalizing the firing rate distributions to avoid imparting a directional bias in firing pattern prediction calculations. Box plots on top show the minimum, 1st, 2nd, and 3rd quartiles, and maximum values.





Figure S7.

Further neuron pair exclusion to avoid calculation anomalies.

(A) Illustrates the calculation of a firing rate product threshold to eliminate an abnormal number of p values near 1. We determined that these anomalous values resulted from pairs with lower firing rate products, for which statistical assumptions of calculations likely were not met. Trace shows the fraction of p values greater than 0.95 at different firing rate product thresholds on the pairs used for Granger causality calculations. Red line shows the expected fraction greater than 0.95 based on the p value distribution.

(B),(C) Firing rate distribution for CFA and RFA neurons included in pairs analyzed using Granger causality (B) or convergent cross mapping (C) after firing rate product threshold exclusion. In (B),(C), box plots on top show the minimum, 1st, 2nd, and 3rd quartiles, and maximum values.



Acknowledgements

We are grateful to J. Glaser and L. Pinto for helpful conversations, and J. Glaser and M. Elbaz for comments on the manuscript. A.S.I. was supported by the Japan Society for the Promotion of Science and the Uehara Memorial Foundation. A.M. was supported by a Searle Scholar Award, a Sloan Research Fellowship, a Whitehall Research Grant Award, The Chicago Biomedical Consortium with support from the Searle Funds at The Chicago Community Trust, the Simons Foundation, and NIH grant DP2 NS120847.

Additional information

Author contributions

A.S.I. and A.M. devised the project and designed experiments. A.S.I., M.A., and S.S. performed the experiments. A.F., A.S.I., M.A., S.S., H.S., S.H., D.B., and F.X. analyzed data. A.M. wrote the manuscript, with all other authors contributing text for the Methods.

Additional files

Movie S1. *Water reaching, side view. A side view of a mouse completing a trial of the water reaching task. LEDs are hidden behind the water ports from this view angle.* [↗](#)

Movie S2. *Water reaching, rear view. A rear view of a mouse completing three trials of the water reaching task. LEDs can be seen in the middle of the screen. Graphics indicate when the LED, Go cue tone ('Cue'), and water dispensation ('Reward') occur. Rewards are achieved when the mouse maintains its hand on the rung for the duration of the rest period.* [↗](#)

References

1. Felleman D.J., Essen D.C.V (1991) **Distributed Hierarchical Processing in the Primate Cerebral Cortex** *Cereb Cortex* **1**:1–47
2. Bullmore E., Sporns O (2009) **Complex brain networks: graph theoretical analysis of structural and functional systems** *Nat Rev Neurosci* **10**:186–198
3. Vezoli J., Magrou L., Goebel R., Wang X.J., Knoblauch K., Vinck M., Kennedy H (2021) **Cortical hierarchy, dual counterstream architecture and the importance of top-down generative networks** *Neuroimage* **225**:117479
4. Campbell A.W (1904) **Histological Studies on the Localisation of Cerebral Function** *J Ment Sci* **50**:651–662
5. Ueta Y., Otsuka T., Morishima M., Ushimaru M., Kawaguchi Y (2014) **Multiple layer 5 pyramidal cell subtypes relay cortical feedback from secondary to primary motor areas in rats** *Cereb Cortex* **24**:2362–2376
6. Passingham R.E (1985) **Premotor cortex: sensory cues and movement** *Behav Brain Res* **18**:175–185
7. Gremel C.M., Costa R.M (2013) **Premotor cortex is critical for goal-directed actions** *Front Comput Neurosci* **7**
8. Roland P.E., Larsen B., Lassen N.A., Skinhoj E (1980) **Supplementary motor area and other cortical areas in organization of voluntary movements in man** *J Neurophysiol* **43**:118–136
9. Veuthey T.L., Derosier K., Kondapavulur S., Ganguly K (2020) **Single-trial cross-area neural population dynamics during long-term skill learning** *Nat Commun* **11**:4057
10. Schmidlin E., Brochier T., Maier M.A., Kirkwood P.A., Lemon R.N (2008) **Pronounced reduction of digit motor responses evoked from macaque ventral premotor cortex after reversible inactivation of the primary motor cortex hand area** *J Neurosci* **28**:5772–5783
11. Elliott K.C., Borrell J.A., Barbay S., Nudo R.J. (2020) **Functional Suppression of Premotor Activity in a Transient Model of Motor Cortex Injury** *bioRxiv* <https://doi.org/10.1101/2020.06.12.148403>
12. Scott S.H (2000) **Population vectors and motor cortex: neural coding or epiphenomenon?** *Nat Neurosci* **3**:307–308
13. Fried I., Haggard P., He B.J., Schurger A (2017) **Volition and Action in the Human Brain: Processes, Pathologies, and Reasons** *J Neurosci* **37**:10842–10847
14. Dum R.P., Strick P.L (1991) **The origin of corticospinal projections from the premotor areas in the frontal lobe** *J Neurosci* **11**:667–689
15. Fisher K.M., Zaimi B., Edgley S.A., Baker S.N (2021) **Extensive Cortical Convergence to Primate Reticulospinal Pathways** *J Neurosci* **41**:1005–1018

16. Hausmann F.S., Barrett J.M., Martin M.E., Zhan H., Shepherd G.M.G (2022) **Axonal Barcode Analysis of Pyramidal Tract Projections from Mouse Forelimb M1 and M2** *J Neurosci* **42**:7733–7743
17. Dum R.P., Strick P.L (2005) **Frontal lobe inputs to the digit representations of the motor areas on the lateral surface of the hemisphere** *J Neurosci* **25**:1375–1386
18. Bundy D.T., Barbay S., Hudson H.M., Frost S.B., Nudo R.J., Guggenmos D.J (2023) **Stimulation-Evoked Effective Connectivity (SEEC): An in-vivo approach for defining mesoscale corticocortical connectivity** *J Neurosci Methods* **384**:109767
19. Weinrich M., Wise S.P., Mauritz K.H (1984) **A neurophysiological study of the premotor cortex in the rhesus monkey** *Brain* **107**:385–414
20. Prut Y., Perlmutter S.I., Fetz E.E (2001) **Distributed processing in the motor system: spinal cord perspective** *Prog Brain Res* **130**:267–278
21. Fulton J.F (1935) **A note on the definition of the "motor" and "premotor" areas** *Brain* **58**:311–316
22. Graziano M.S.A. (2009) **The intelligent movement machine : an ethological perspective on the primate motor system** Oxford University Press
23. Morandell K., Huber D (2017) **The role of forelimb motor cortex areas in goal directed action in mice** *Sci Rep* **7**:15759
24. Hilgetag C.C., Goulas A (2020) **'Hierarchy' in the organization of brain networks** *Philos Trans R Soc Lond B Biol Sci* **375**:20190319
25. Siegle J.H., Jia X.X., Durand S., Gale S., Bennett C., Graddis N., Heller G., Ramirez T.K., Choi H., Luviano J.A., et al. (2021) **Survey of spiking in the mouse visual system reveals functional hierarchy** *Nature* **592**:86–92
26. Merel J., Botvinick M., Wayne G (2019) **Hierarchical motor control in mammals and machines** *Nat Commun* **10**:5489
27. Bucy P.C (1933) **Electrical excitability and cyto-architecture of the premotor cortex in monkeys** *Arch Neuro Psychiatr* **30**:1205–1225
28. Stepniewska I., Gharbawie O.A., Burish M.J., Kaas J.H (2014) **Effects of muscimol inactivations of functional domains in motor, premotor, and posterior parietal cortex on complex movements evoked by electrical stimulation** *J Neurophysiol* **111**:1100–1119
29. Rouiller E.M., Moret V., Liang F (1993) **Comparison of the connectional properties of the two forelimb areas of the rat sensorimotor cortex: support for the presence of a premotor or supplementary motor cortical area** *Somatosens Mot Res* **10**:269–289
30. Hira R., Ohkubo F., Tanaka Y.R., Masamizu Y., Augustine G.J., Kasai H., Matsuzaki M (2013) **In vivo optogenetic tracing of functional corticocortical connections between motor forelimb areas** *Front Neural Circuit* **7**
31. Mushiake H., Inase M., Tanji J (1991) **Neuronal activity in the primate premotor, supplementary, and precentral motor cortex during visually guided and internally determined sequential movements** *J Neurophysiol* **66**:705–718

32. Umiltà M.A., Brochier T., Spinks R.L., Lemon R.N (2007) **Simultaneous recording of macaque premotor and primary motor cortex neuronal populations reveals different functional contributions to visuomotor grasp** *J Neurophysiol* **98**:488–501
33. Dixon T.C., Merrick C.M., Wallis J.D., Ivry R.B., Carmena J.M (2021) **Hybrid dedicated and distributed coding in PMd/M1 provides separation and interaction of bilateral arm signals** *PLoS Comput Biol* **17**:e1009615
34. Cerri G., Shimazu H., Maier M.A., Lemon R.N (2003) **Facilitation from ventral premotor cortex of primary motor cortex outputs to macaque hand muscles** *J Neurophysiol* **90**:832–842
35. Deffeyes J.E., Touvykine B., Quessy S., Dancause N (2015) **Interactions between rostral and caudal cortical motor areas in the rat** *J Neurophysiol* **113**:3893–3904
36. Makino H., Ren C., Liu H., Kim A.N., Kondapaneni N., Liu X., Kuzum D., Komiyama T (2017) **Transformation of Cortex-wide Emergent Properties during Motor Learning** *Neuron* **94**:880–890
37. Terada S.I., Kobayashi K., Matsuzaki M (2022) **Transition of distinct context-dependent ensembles from secondary to primary motor cortex in skilled motor performance** *Cell Rep* **41**:111494
38. Truccolo W., Hochberg L.R., Donoghue J.P (2010) **Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes** *Nat Neurosci* **13**:105–111
39. Kimura R., Saiki A., Fujiwara-Tsukamoto Y., Sakai Y., Isomura Y (2017) **Large-scale analysis reveals populational contributions of cortical spike rate and synchrony to behavioural functions** *J Physiol* **595**:385–413
40. D'Aleo R., Rouse A.G., Schieber M.H., Sarma S.V (2022) **Cortico-cortical drive in a coupled premotor-primary motor cortex dynamical system** *Cell Rep* **41**:111849
41. Wiesendanger M., Hummelsheim H., Bianchetti M., Chen D.F., Hyland B., Maier V., Wiesendanger R (1987) **Input and Output Organization of the Supplementary Motor Area** *Ciba F Symp* **132**:40–62
42. Graziano M (2006) **The organization of behavioral repertoire in motor cortex** *Annu Rev Neurosci* **29**:105–134
43. Gokcen E., Jasper A.I., Semedo J.D., Zandvakili A., Kohn A., Machens C.K., Yu B.M (2022) **Disentangling the flow of signals between populations of neurons** *Nat Comput Sci* **2**:512
44. Issa E.B., Cadieu C.F., DiCarlo J.J (2018) **Neural dynamics at successive stages of the ventral visual stream are consistent with hierarchical error signals** *eLife* **7**
45. Siegel M., Buschman T.J., Miller E.K (2015) **Cortical information flow during flexible sensorimotor decisions** *Science* **348**:1352–1355
46. Schwiedrzik C.M., Freiwald W.A (2017) **High-Level Prediction Signals in a Low- Level Area of the Macaque Face-Processing Hierarchy** *Neuron* **96**:89–97
47. Churchland M.M., Shenoy K.V (2007) **Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex** *J Neurophysiol* **97**:4235–4257

48. Ito S., Hansen M.E., Heiland R., Lumsdaine A., Litke A.M., Beggs J.M (2011) **Extending transfer entropy improves identification of effective connectivity in a spiking cortical network model** *PLoS One* **6**:e27431
49. Casile A., Faghih R.T., Brown E.N (2021) **Robust point-process Granger causality analysis in presence of exogenous temporal modulations and trial-by-trial variability in spike trains** *PLoS Comput Biol* **17**:e1007675
50. Bachschmid-Romano L., Hatsopoulos N.G., Brunel N., Diedrichsen J (2023) **Interplay between external inputs and recurrent dynamics during movement preparation and execution in a network model of motor cortex** *eLife* **12**
51. Gozel O., Doiron B. (2023) **Between-area communication through the lens of within- area neuronal dynamics** *bioRxiv*
52. Jun J.J., Steinmetz N.A., Siegle J.H., Denman D.J., Bauza M., Barbarits B., Lee A.K., Anastassiou C.A., Andrei A., Aydın Ç., et al. (2017) **Fully integrated silicon probes for high-density recording of neural activity** *Nature* **551**:232–236
53. Galinanes G.L., Bonardi C., Huber D (2018) **Directional Reaching for Water as a Cortex-Dependent Behavioral Framework for Mice** *Cell Rep* **22**:2767–2783
54. Miri A., Warriner C.L., Seely J.S., Elsayed G.F., Cunningham J.P., Churchland M.M., Jessell T.M (2017) **Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex** *Neuron* **95**:683–696
55. Zhao S., Ting J.T., Atallah H.E., Qiu L., Tan J., Gloss B., Augustine G.J., Deisseroth K., Luo M., Graybiel A.M., Feng G (2011) **Cell type-specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function** *Nat Methods* **8**:745–752
56. Guo Z.V., Li N., Huber D., Ophir E., Gutnisky D., Ting J.T., Feng G., Svoboda K (2014) **Flow of cortical activity underlying a tactile decision in mice** *Neuron* **81**:179–194
57. Tanji J., Kurata K (1982) **Comparison of movement-related activity in two cortical motor areas of primates** *J Neurophysiol* **48**:633–653
58. Kvitsiani D., Ranade S., Hangya B., Taniguchi H., Huang J.Z., Kepecs A (2013) **Distinct behavioural and network correlates of two interneuron types in prefrontal cortex** *Nature* **498**:363–366
59. Terada S.I., Kobayashi K., Ohkura M., Nakai J., Matsuzaki M (2018) **Super-wide- field two-photon imaging with a micro-optical device moving in post-objective space** *Nat Commun* **9**:3550
60. Saiki A., Kimura R., Samura T., Fujiwara-Tsukamoto Y., Sakai Y., Isomura Y (2014) **Different modulation of common motor information in rat primary and secondary motor cortices** *PLoS One* **9**:e98662
61. Hotelling H (1936) **Relations between two sets of variates** *Biometrika* **28**:321–377
62. Sussillo D., Churchland M.M., Kaufman M.T., Shenoy K.V (2015) **A neural network that finds a naturalistic solution for the production of muscle activity** *Nat Neurosci* **18**:1025–1033

63. Le Floch E., Guillemot V., Frouin V., Pinel P., Lalanne C., Trinchera L., Tenenhaus A., Moreno A., Zilbovicius M., Bourgeron T., et al. (2012) **Significant correlation between a set of genetic polymorphisms and a functional brain network revealed by feature selection and sparse Partial Least Squares** *Neuroimage* **63**:11–24
64. Wang Z., Maunze B., Wang Y., Tsoulfas P., Blackmore M.G (2018) **Global Connectivity and Function of Descending Spinal Input Revealed by 3D Microscopy and Retrograde Transduction** *J Neurosci* **38**:10566–10581
65. Kaufman M.T., Seely J.S., Sussillo D., Ryu S.I., Shenoy K.V., Churchland M.M (2016) **The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type** *Eneuro* **3**
66. Sugihara G., May R., Ye H., Hsieh C.H., Deyle E., Fogarty M., Munch S (2012) **Detecting causality in complex ecosystems** *Science* **338**:496–500
67. Storey J.D (2002) **A direct approach to false discovery rates** *Journal of the Royal Statistical Society Series B-Statistical Methodology* **64**:479–498
68. Oh S.W., Harris J.A., Ng L., Winslow B., Cain N., Mihalas S., Wang Q., Lau C., Kuan L., Henry A.M., et al. (2014) **A mesoscale connectome of the mouse brain** *Nature* **508**:207–214
69. Urban Iii E.T., Hudson H.M., Li Y., Nishibe M., Barbay S., Guggenmos D.J., and Nudo R.J. (2024) **Corticocortical connections of the rostral forelimb area in rats: a quantitative tract-tracing study** *Cereb Cortex* **34**
70. Rossi L.F., Harris K.D., Carandini M (2020) **Spatial connectivity matches direction selectivity in visual cortex** *Nature* **588**:648–652
71. Thomson A.M., Lamy C (2007) **Functional maps of neocortical local circuitry** *Front Neurosci* **1**:19–42
72. Braitenberg V., Sch*uz A. (2014) **Cortex: Statistics and Geometry of Neuronal Connectivity, 2nd Edition** Berlin: Springer
73. Ahmadian Y., Miller K.D (2021) **What is the dynamical regime of cerebral cortex?** *Neuron* **109**:3373–3391
74. Sadeh S., Clopath C (2021) **Inhibitory stabilization and cortical computation** *Nat Rev Neurosci* **22**:21–37
75. Seung H.S (1996) **How the brain keeps the eyes still** *Proc. Natl. Acad. Sci. USA* **93**:13339–13344
76. Wang P., Kong R., Kong X., Liegeois R., Orban C., Deco G., van den Heuvel M.P., Thomas Yeo B.T. (2019) **Inversion of a large-scale circuit model reveals a cortical hierarchy in the dynamic resting human brain** *Sci Adv* **5**:eaat7854
77. Chaudhuri R., Knoblauch K., Gariel M.A., Kennedy H., Wang X.J (2015) **A Large- Scale Circuit Mechanism for Hierarchical Dynamical Processing in the Primate Cortex** *Neuron* **88**:419–431
78. Sanzeni A., Akitake B., Goldbach H.C., Leedy C.E., Brunel N., Histed M.H (2020) **Inhibition stabilization is a widespread property of cortical networks** *eLife* **9**

79. Riehle A., Requin J (1993) **The predictive value for performance speed of preparatory changes in neuronal activity of the monkey motor and premotor cortex** *Behav Brain Res* **53**:35–49
80. Hyland B (1998) **Neural activity related to reaching and grasping in rostral and caudal regions of rat motor cortex** *Behav Brain Res* **94**:255–269
81. Elsayed G.F., Lara A.H., Kaufman M.T., Churchland M.M., Cunningham J.P (2016) **Reorganization between preparatory and movement population responses in motor cortex** *Nat Commun* **7**:13239
82. Shenoy K.V., Sahani M., Churchland M.M (2013) **Cortical control of arm movements: a dynamical systems perspective** *Annu Rev Neurosci* **36**:337–359
83. Seely J.S., Kaufman M.T., Ryu S.I., Shenoy K.V., Cunningham J.P., Churchland M.M (2016) **Tensor Analysis Reveals Distinct Population Structure that Parallels the Different Computational Roles of Areas M1 and V1** *PLoS Comput Biol* **12**:e1005164
84. Sauerbrei B.A., Guo J.Z., Cohen J.D., Mischiati M., Guo W., Kabra M., Verma N., Mensh B., Branson K., Hantman A.W (2020) **Cortical pattern generation during dexterous movement is input-driven** *Nature* **577**:386–391
85. Guo Z.V., Hires S.A., Li N., O'Connor D.H., Komiyama T., Ophir E., Huber D., Bonardi C., Morandell K., Gutnisky D., et al. (2014) **Procedures for behavioral experiments in head-fixed mice** *PLoS One* **9**:e88678
86. Akay T., Acharya H.J., Fouad K., Pearson K.G (2006) **Behavioral and electromyographic characterization of mice lacking EphA4 receptors** *Journal of Neurophysiology* **96**:642–651
87. Warriner C.L., Fageiry S., Saxena S., Costa R.M., Miri A (2022) **Motor cortical influence relies on task-specific activity covariation** *Cell Rep* **40**:111427
88. Pachitariu M., Steinmetz N., Kadir S., Carandini M., Kenneth D., H. (2016) **Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels** *bioRxiv*
89. Krakovská A., Jakubík J., Budáčová H., Holeciová M. (2015) **Causality studied in reconstructed state space. Examples of uni-directionally connected chaotic systems** *arXiv* <https://doi.org/10.48550/arXiv.1511.00505>
90. Ye H., Deyle E.R., Gilarranz L.J., Sugihara G (2015) **Distinguishing time-delayed causal interactions using convergent cross mapping** *Sci Rep* **5**:14750
91. Clark A.T., Ye H., Isbell F., Deyle E.R., Cowles J., Tilman G.D., Sugihara G (2015) **Spatial convergent cross mapping to detect causal relationships from short time series** *Ecology* **96**:1174–1181
92. Ehrlich D.B., Stone J.T., Brandfonbrener D., Atanasov A., Murray J.D (2021) **PsychRNN: An Accessible and Flexible Python Package for Training Recurrent Neural Network Models on Cognitive Tasks** *Eneuro* **8**

Author information

Akiko Saiki-Ishikawa[†]

Department of Neurobiology, Northwestern University, Evanston, United States

[†]These authors contributed equally to this work.

Mark Agrios[†]

Department of Neurobiology, Northwestern University, Evanston, United States

[†]These authors contributed equally to this work.

Sajishnu Savya[†]

Department of Neurobiology, Northwestern University, Evanston, United States

[†]These authors contributed equally to this work.

Adam Forrest[†]

Department of Neurobiology, Northwestern University, Evanston, United States

[†]These authors contributed equally to this work.

Hannah Sroussi

Department of Neurobiology, Northwestern University, Evanston, United States

Sarah Hsu

Department of Neurobiology, Northwestern University, Evanston, United States

Diya Basrai

Department of Neurobiology, Northwestern University, Evanston, United States

Feihong Xu

Department of Neurobiology, Northwestern University, Evanston, United States

Andrew Miri[#]

Department of Neurobiology, Northwestern University, Evanston, United States

For correspondence: andrewmiri@northwestern.edu

[#]lead contact

Editors

Reviewing Editor

Juan Alvaro Gallego

Imperial College London, London, United Kingdom

Senior Editor

Tamar Makin

University of Cambridge, Cambridge, United Kingdom

Reviewer #1 (Public review):

This study examined the interaction between two key cortical regions in the mouse brain involved in goal-directed movements, the rostral forelimb area (RFA) - considered a premotor region involved in movement planning, and the caudal forelimb area (CFA) - considered a primary motor region that more directly influences movement execution. The authors ask whether there exists a hierarchical interaction between these regions, as previously hypothesized, and focus on a specific definition of hierarchy - examining whether the neural activity in the premotor region exerts a larger functional influence on the activity in the primary motor area, than vice versa. They examine this question using advanced experimental and analytical methods, including localized optogenetic manipulation of neural activity in either region while measuring both the neural activity in the other region and EMG signals from several muscles involved in the reaching movement, as well as simultaneous electrophysiology recordings from both regions in a separate cohort of animals.

The findings presented show that localized optogenetic manipulation of neural activity in either RFA or CFA resulted in similarly short-latency changes of the muscle output and in firing rate changes in the other region. However, perturbation of RFA led to a larger absolute change in the neural activity of CFA neurons. The authors interpret these findings as evidence for reciprocal, but asymmetrical, influence between the regions, suggesting some degree of hierarchy in which RFA has a greater effect on the neural activity in CFA. They go on to examine whether this asymmetry can also be observed in simultaneously recorded neural activity patterns from both regions. They use multiple advanced analysis methods that either identify latent components in the population level or measure the predictability of firing rates of single neurons in one region using firing rates of single neurons in the other region. Interestingly, the main finding across these analyses seems to be that both regions share highly similar components that capture a high degree of the variability of the neural activity patterns in each region. Single units' activity from either region could be predicted to a similar degree from the activity of single units in the other region, without a clear division into a leading area and a lagging area, as one might expect to find in a simple hierarchical interaction. However, the authors find some evidence showing a slight bias towards leading activity in RFA. Using a two-region neural network model that is fit to the summed neural activity recorded in the different experiments and to the summed muscle output, the authors show that a network with constrained (balanced) weights between the regions can still output the observed measured activities and the observed asymmetrical effects of the optogenetic manipulations, by having different within-region local weights. These results emphasize the challenges in studying interactions between brain regions with reciprocal interactions, multiple external inputs, and recurrent within-region connections.

Strengths:

The experiments and analyses performed in this study are comprehensive and provide a detailed examination and comparison of neural activity recorded simultaneously using dense electrophysiology probes from two main motor regions that have been the focus of studies examining goal-directed movements. The findings showing reciprocal effects from each region to the other, similar short-latency modulation of muscle output by both regions, and similarity of neural activity patterns, are convincing and add to the growing body of evidence that highlight the complexity of the interactions between multiple regions in the motor system and go against a simple feedforward-like hierarchy.

The neural network model complements these findings and adds an important demonstration that the observed asymmetry can, in theory, also arise from differences in

local recurrent connections and not necessarily from different input projections from one region to the other. This sheds an important light on the multiple factors that should be considered when studying the interaction between any two brain regions, with a specific emphasis on the role of local recurrent connections, that should be of interest to the general neuroscience community.

Weaknesses:

While the reciprocal interaction and similarity in neural activity across RFA and CFA is an important observation that is supported by the authors' findings, the evidence for a hierarchical interaction between the two regions appears to be weaker. The primary evidence for a hierarchical interaction comes from a causal optogenetic manipulation, carried out at the onset of the reaching movement and conducted with $n = 3$ in each experimental group, which shows an effect in both regions, yet the effect is greater when silencing the activity in RFA and examining the resulting change in CFA, than vice versa. Analysis of the simultaneously recorded neural activity, on the other hand, reveals mostly no clear hierarchy with leading or lagging dynamics between the regions. The findings of the optogenetic manipulation might be more compelling if similar effects were observed when the same manipulation was applied at different stages of movement preparation and execution, indicating a consistent interaction that is independent from the movement phase.

The methods used to investigate hierarchical interactions through analysis of simultaneously recorded activity yielded inconsistent results. For instance, CCA and PLS showed no clear lead-lag relationship, while DLAG provided some evidence suggesting RFA leads CFA. Overall, these methods largely failed to demonstrate a clear hierarchical interaction. Assuming a partial hierarchy exists, this inconsistency may indicate that the hierarchy is not reflected in the activity patterns or that these analytical methods are inadequate for detecting such interactions within complex neural networks that are influenced by multiple external inputs, reciprocal inter-regional connections, and dominant intra-regional recurrent activity.

As is also argued by the authors, these inconsistent findings underscore the need for caution when interpreting results from similar analyses used to infer inter-regional interactions from neural activity patterns alone. However, the study lacks sufficient explanation for why different methods yielded different results and more elaborate clarification is needed for the findings presented. For example, in the population-level analyses using CCA and PLS, the authors show that both techniques reveal components that are highly similar across regions and explain a substantial portion of each region's variance. Yet, shifting the activity of one region relative to the other to explore potential lead-lag relationships does not alter the results of these analyses. If the regions' activities were better aligned at some unknown true lead-lag time (or aligned at zero), one would expect a peak in alignment within the tested range, as is observed when these same analyses are applied to activity within a single region. It is thus unclear why shifting one region's activity relative to the other does not change the outcome. The interpretation of these results therefore, remains ambiguous and would benefit from further clarification.

<https://doi.org/10.7554/eLife.103069.2.sa3>

Reviewer #2 (Public review):

Summary:

While technical advances have enabled large-scale, multi-site neural recordings, characterizing inter-regional communication and its behavioral relevance remains challenging due to intrinsic properties of the brain such as shared inputs, network complexity, and external noise. This work by Saiki-Ishikawa et al. examines the functional

hierarchy between premotor (PM) and primary motor (M1) cortices in mice during a directional reaching task. The authors find some evidence consistent with an asymmetric reciprocal influence between the regions, but overall, activity patterns were highly similar and equally predictive of one another. These results suggest that motor cortical hierarchy, though present, is not fully reflected in firing patterns alone.

Strengths:

Inferring functional hierarchies between brain regions, given the complexity of reciprocal and local connectivity, dynamic interactions, and the influence of both shared and independent external inputs, is a challenging task. It requires careful analysis of simultaneous recording data, combined with cross-validation across multiple metrics, to accurately assess the functional relationships between regions. The authors have generated a valuable dataset simultaneously recording from both regions at scale from mice performing a cortex-dependent directional reaching task.

Using electrophysiological and silencing data, the authors found evidence supporting the traditionally assumed asymmetric influence from PM to M1. While earlier studies inferred a functional hierarchy based on partial temporal relationships in firing patterns, the authors applied a series of complementary analyses to rigorously test this hierarchy at both individual neuron and population levels, with robust statistical validation of significance.

In addition, recording combined with brief optogenetic silencing of the other region allowed authors to infer the asymmetric functional influence in a more causal manner. This experiment is well designed to focus on the effect of inactivation manifesting through oligosynaptic connections to support the existence of a premotor to primary motor functional hierarchy.

Subsequent analyses revealed a more complex picture. CCA, PLS, and three measures of predictivity (Granger causality, transfer entropy, and convergent cross mapping) emphasized similarities in firing patterns and cross-region predictability. However, DLAG suggested an imbalance, with RFA capturing CFA variance at a negative time lag, indicating that RFA 'leads' CFA. Taken together these results provide useful insights for current studies of functional hierarchy about potential limitations in inferring hierarchy solely based on firing rates.

While I would detail some questions and issues on specifics of data analyses and modeling below, I appreciate the authors' effort in training RNNs that match some behavioral and recorded neural activity patterns including the inactivation result. The authors point out two components that can determine the across-region influence - 1) the amount of inputs received and 2) the dependence on across-region input, i.e., relative importance of local dynamics, providing useful insights in inferring functional relationships across regions.

Weaknesses:

(1) Trial-averaging was applied in CCA and PLS analyses. While trial-averaging can be appropriate in certain cases, it leads to the loss of trial-to-trial variance, potentially inflating the perceived similarities between the activity in the two regions (Figure 4). Do authors observe comparable degrees of similarity, e.g., variance explained by canonical variables? Also, the authors report conflicting findings regarding the temporal relationship between RFA and CFA when using CCA/PLS versus DLAG. Could this discrepancy be due to the use of trial-averaging in former analyses but not in the latter?

(2) A key strength of the current study is the precise tracking of forelimb muscle activity during a complex motor task involving reaching for four different targets. This rich behavioral data is rarely collected in mice and offers a valuable opportunity to investigate the behavioral relevance of the PM-M1 functional interaction, yet little has been done to explore this aspect in depth. For example, single-trial time courses of inter-regional latent variables

acquired from DLAG analysis can be correlated with single-trial muscle activity and/or reach trajectories to examine the behavioral relevance of inter-regional dynamics. Namely, can trial-by-trial change in inter-regional dynamics explain behavioral variability across trials and/or targets? Does the inter-areal interaction change in error trials? Furthermore, the authors could quantify the relative contribution of across-area versus within area dynamics to behavioral variability. It would also be interesting to assess the degree to which across-area and within-area dynamics are correlated. Specifically, can across-area dynamics vary independently from within-area dynamics across trials, potentially operating through a distinct communication subspace?

(3) While network modeling of RFA and CFA activity captured some aspects of behavioral and neural data, I wonder if certain findings such as the connection weight distribution (Figure 7C), across-region input (Figure 7F), and the within-region weights (Figure 7G), primarily resulted from fitting the different overall firing rates between the two regions with CFA exhibiting higher average firing rates. Did the authors account for this firing rate disparity when training the RNNs?

(4) Another way to assess the functional hierarchy is by comparing the time courses of movement representation between the two regions. For example, a linear decoder could be used to compare the amount of information about muscle activity and/or target location as well as time courses thereof between the two regions. This approach is advantageous because it incorporates behavior rather than focusing solely on neural activity. Since one of the main claims of this study is the limitation of inferring functional hierarchy from firing rate data alone, the authors should use the behavior as a lens for examining inter-areal interactions.

Comments on revisions:

I appreciate the authors' thoughtful revisions in response to prior reviews, which I believe have substantially improved the manuscript. In particular, I found the addition of the new section "Manifestations of hierarchy in firing patterns" to be valuable, as it begins to address some of the more complex and potentially conflicting observations

<https://doi.org/10.7554/eLife.103069.2.sa2>

Reviewer #3 (Public review):

This study investigates how two cortical regions which are central to the study of rodent motor control (rostral forelimb area, RFA, and caudal forelimb area, CFA) interact during directional forelimb reaching in mice. The authors investigate this interaction using (1) optogenetic manipulations in one area while recording extracellularly from the other, (2) statistical analyses of simultaneous CFA/RFA extracellular recordings, and (3) network modeling. The authors provide solid evidence that asymmetry between RFA and CFA can be observed, although such asymmetry is only observed in certain experimental and analytical contexts.

The authors find asymmetry when applying optogenetic perturbations, reporting a greater impact of RFA inactivation on CFA activity than vice-versa. The authors then investigate asymmetry in endogenous activity during forelimb movements and find asymmetry with some analytical methods but not others. Asymmetry was observed in the onset timing of movement-related deviations of local latent components with RFA leading CFA (computed with PCA) and in a relatively higher proportion and importance of cross-area latent components with RFA leading than CFA leading (computed with DLAG). However, no asymmetry was observed using several other methods that compute cross-area latent dynamics, nor with methods computed on individual neuron pairs across regions. The authors follow up this experimental work by developing a two-area model with asymmetric

dependence on cross-area input. This model is used to show that differences in local connectivity can drive asymmetry between two areas with equal amounts of across-region input.

Overall, this work provides a useful demonstration that different cross-area analysis methods result in different conclusions regarding asymmetric interactions between brain areas and suggests careful consideration of methods when analyzing such networks is critical. A deeper examination of why different analytical methods result in observed asymmetry or no asymmetry, analyses that specifically examine neural dynamics informative about details of the movement, or a biological investigation of the hypothesis provided by the model would provide greater clarity regarding the interaction between RFA and CFA.

Strengths:

The authors are rigorous in their experimental and analytical methods, carefully monitoring the impact of their perturbations with simultaneous recordings and providing valid controls for their analytical methods. They cite relevant previous literature that largely agrees with the current work, highlighting the continued ambiguity regarding the extent to which there exists an asymmetry in endogenous activity between RFA and CFA.

A strength of the paper is the evidence for asymmetry provided by optogenetic manipulation. They show that RFA inactivation causes a greater absolute difference in muscle activity than CFA interaction (deviations begin 25-50 ms after laser onset, Figure 1) and that RFA inactivation causes a relatively larger decrease in CFA firing rate than CFA inactivation causes in RFA (deviations begin <25ms after laser onset, Figure 3). The timescales of these changes provide solid evidence for an asymmetry in impact of inactivating RFA/CFA on the other region that could not be driven by differences in feedback from disrupted movement (which would appear with a ~50ms delay).

The authors also utilize a range of different analytical methods, showing an interesting difference between some population-based methods (PCA, DLAG) that observe asymmetry, and single neuron pair methods (granger causality, transfer entropy, and convergent cross mapping) that do not. Moreover, the modeling work presents an interesting potential cause of "hierarchy" or "asymmetry" between brain areas: local connectivity that impacts dependence on across-region input, rather than the amount of across-region input actually present.

Weaknesses:

There is no attempt to examine neural dynamics that are specifically relevant/informative about the details of the ongoing forelimb movement (e.g., kinematics, reach direction). Thus, it may be preemptive to claim that firing patterns alone do not reflect functional influence between RFA/CFA. For example, given evidence that the largest component of motor cortical activity doesn't reflect details of ongoing movement (reach direction or path; Kaufman, et al. PMID: 27761519) and that the analytical tools the authors use likely include this component (PCA, CCA), it may not be surprising that CFA and RFA do not show asymmetry if such asymmetry is related to control of movement details. An asymmetry may still exist in the components of neural activity that encode information about movement details, and thus it may be necessary to isolate and examine the interaction of behaviorally-relevant dynamics (e.g., Sani, et al. PMID: 33169030).

The idea that local circuit dynamics play a central role in determining the asymmetry between RFA and CFA is not supported by experimental data in this paper. The plausibility of this hypothesis is supported by the model but is not explored in any analyses of the experimental data collected. Further experimental investigation is needed to separate this hypothesis from other possibilities.

Comments on revisions:

The authors have improved the manuscript by reviewing several aspects of the text and the addition of supplemental materials. I believe these revisions have clarified some important aspects of the results.

<https://doi.org/10.7554/eLife.103069.2.sa1>

Author response:

The following is the authors' response to the original reviews

Public Reviews:

Reviewer #1 (Public review):

This study examined the interaction between two key cortical regions in the mouse brain involved in goal-directed movements, the rostral forelimb area (RFA) - considered a premotor region involved in movement planning, and the caudal forelimb area (CFA) - considered a primary motor region that more directly influences movement execution. The authors ask whether there exists a hierarchical interaction between these regions, as previously hypothesized, and focus on a specific definition of hierarchy - examining whether the neural activity in the premotor region exerts a larger functional influence on the activity in the primary motor area than vice versa. They examine this question using advanced experimental and analytical methods, including localized optogenetic manipulation of neural activity in either region while measuring both the neural activity in the other region and EMG signals from several muscles involved in the reaching movement, as well as simultaneous electrophysiology recordings from both regions in a separate cohort of animals.

The findings presented show that localized optogenetic manipulation of neural activity in either RFA or CFA resulted in similarly short-latency changes in the muscle output and in firing rate changes in the other region. However, perturbation of RFA led to a larger absolute change in the neural activity of CFA neurons. The authors interpret these findings as evidence for reciprocal, but asymmetrical, influence between the regions, suggesting some degree of hierarchy in which RFA has a greater effect on the neural activity in CFA. They go on to examine whether this asymmetry can also be observed in simultaneously recorded neural activity patterns from both regions. They use multiple advanced analysis methods that either identify latent components at the population level or measure the predictability of firing rates of single neurons in one region using firing rates of single neurons in the other region. Interestingly, the main finding across these analyses seems to be that both regions share highly similar components that capture a high degree of variability of the neural activity patterns in each region. Single units' activity from either region could be predicted to a similar degree from the activity of single units in the other region, without a clear division into a leading area and a lagging area, as one might expect to find in a simple hierarchical interaction. However, the authors find some evidence showing a slight bias towards leading activity in RFA. Using a two-region neural network model that is fit to the summed neural activity recorded in the different experiments and to the summed muscle output, the authors show that a network with constrained (balanced) weights between the regions can still output the observed measured activities and the observed asymmetrical effects of the optogenetic manipulations, by having different within-region local weights. These results put into question whether previous and current findings that demonstrate asymmetry in the output of regions can be interpreted as evidence for asymmetrical (and thus hierarchical) inputs between regions, emphasizing the challenges in studying interactions between any brain regions.

Strengths:

The experiments and analyses performed in this study are comprehensive and provide a detailed examination and comparison of neural activity recorded simultaneously using dense electrophysiology probes from two main motor regions that have been the focus of studies examining goal-directed movements. The findings showing reciprocal effects from each region to the other, similar short-latency modulation of muscle output by both regions, and similarity of neural activity patterns without a clear lead/lag interaction, are convincing and add to the growing body of evidence that highlight the complexity of the interactions between multiple regions in the motor system and go against a simple feedforward-like network and dynamics. The neural network model complements these findings and adds an important demonstration that the observed asymmetry can, in theory, also arise from differences in local recurrent connections and not necessarily from different input projections from one region to the other. This sheds an important light on the multiple factors that should be considered when studying the interaction between any two brain regions, with a specific emphasis on the role of local recurrent connections, that should be of interest to the general neuroscience community.

Weaknesses:

While the similarity of the activity patterns across regions and lack of a clear leading/lagging interaction are interesting observations that are mostly supported by the findings presented (however, see comment below for lack of clarity in CCA/PLS analyses), the main question posed by the authors - whether there exists an endogenous hierarchical interaction between RFA and CFA - seems to be left largely open.

The authors note that there is currently no clear evidence of asymmetrical reciprocal influence between naturally occurring neural activity patterns of the two regions, as previous attempts have used non-natural electrical stimulation, lesions, or pharmacological inactivation. The use of acute optogenetic perturbations does not seem to be vastly different in that aspect, as it is a non-natural stimulation of inhibitory interneurons that abruptly perturbs the ongoing dynamics.

We do believe that our optogenetic inactivation identifies a causal interaction between the endogenous activity patterns in the excitatory projection neurons, which we have largely silenced, and the downstream endogenous activity that is perturbed. The effect in the downstream region results directly from the silencing of activity in the excitatory projection neurons that mediate each region's interaction with other regions. Here we have performed a causal intervention common in biology: a loss-of-function experiment. Such experiments generally reveal that a causal interaction of some sort is present, but often do not clarify much about the nature of the interaction, as is true in our case. By showing that a silencing of endogenous activity in one motor cortical region causes a significant change to the endogenous activity in another, we establish a causal relationship between these activity patterns. This is analogous to knocking out the gene for a transcription factor and observing causal effects on the expression of other genes that depend on it.

Moreover, our experiments are, to our knowledge, the first that localize a causal relationship to endogenous activity in motor cortex at a particular point during a motor behavior. Lesion and pharmacological or chemogenetic inactivation have long-lasting effects, and so their consequences on firing in other regions cannot be attributed to a short-latency influence of activity at a particular point during movement. Moreover, the involvement of motor cortex in motor learning and movement preparation/initiation complicates the interpretation of these consequences in relation to movement execution, as disturbance to processes on which execution depends can impede execution itself. Stimulation experiments generate spiking in excitatory projection neurons that is not endogenous.

That said, we would agree that the form of the causal interaction between RFA and CFA remains unaddressed by our results. These results do not expose how the silenced activity patterns affect activity in the downstream region, just as knocking out a transcription factor gene does not expose how the transcription factor influences the expression of other genes. To show evidence for a specific type of interaction dynamics between RFA and CFA, a different sort of experiment would be necessary. See Jazayeri and Afraz, *Neuron*, 2017 for more on this issue.

Furthermore, the main finding that supports a hierarchical interaction is a difference in the absolute change of firing rates as a result of the optogenetic perturbation, a finding that is based on a small number of animals ($N = 3$ in each experimental group), and one which may be difficult to interpret.

Though $N = 3$, we do show statistical significance. Moreover, using three replicates is not uncommon in biological experiments that require a large technical investment.

As the authors nicely demonstrate in their neural network model, the two regions may differ in the strength of local within-region inhibitory connections. Could this theoretically also lead to a difference in the effect of the artificial light stimulation of the inhibitory interneurons on the local population of excitatory projection neurons, driving an asymmetrical effect on the downstream region?

We (Miri et al., *Neuron*, 2017) and others (Guo et al., *Neuron*, 2014) have shown that the effect of this inactivation on excitatory neurons in CFA is a near-complete silencing (90-95% within 20 ms). There thus is not much room for the effects on projection neurons in RFA to be much larger. We have measured these local effects in RFA as part of other work (Kristl et al., *bioRxiv*, 2025), verifying that the effects on RFA projection neuron firing are not larger.

Moreover, the manipulation was performed upon the beginning of the reaching movement, while the premotor region is often hypothesized to exert its main control during movement preparation, and thus possibly show greater modulation during that movement epoch. It is not clear if the observed difference in absolute change is dependent on the chosen time of optogenetic stimulation and if this effect is a general effect that will hold if the stimulation is delivered during different movement epochs, such as during movement preparation.

We agree that the dependence of RFA-CFA interactions on movement phase would be interesting to address in subsequent experiments. While a strong interpretation of lesion results might lead to a hypothesis that premotor influence on primary motor cortex is local to, or stronger during, movement preparation as opposed to execution, at present there is to our knowledge no empirical support from interventional experiments for this hypothesis. Moreover, existing results from analysis of activity in these two regions have produced conflicting results on the strength of interaction between these regions during preparation. Compare for example BachschmidRomano et al., *eLife*, 2023 to Kaufman et al., *Nature Neuroscience*, 2014.

That said, this lesion interpretation would predict the same asymmetry we have observed from perturbations at the beginning of a reach - a larger effect of RFA on CFA than vice versa.

Another finding that is not clearly interpretable is in the analysis of the population activity using CCA and PLS. The authors show that shifting the activity of one region compared to the other, in an attempt to find the optimal leading/lagging interaction, does not affect the results of these analyses. Assuming the activities of both regions are better aligned at some unknown groundtruth lead/lag time, I would expect to see a peak

somewhere in the range examined, as is nicely shown when running the same analyses on a single region's activity. If the activities are indeed aligned at zero, without a clear leading/lagging interaction, but the results remain similar when shifting the activities of one region compared to the other, the interpretation of these analyses is not clear.

Our results in this case were definitely surprising. Many share the intuition that there should be a lag at which the correlations in activity between regions may be strongest. The similarity in alignment across lags we observed might be expected if communication between regions occurs over a range of latencies as a result of dependence on a broad diversity of synaptic paths that connect neurons. In the Discussion, we offer an explanation of how to reconcile these findings with the seemingly different picture presented by DLAG.

Reviewer #2 (Public review):

Summary:

While technical advances have enabled large-scale, multi-site neural recordings, characterizing inter-regional communication and its behavioral relevance remains challenging due to intrinsic properties of the brain such as shared inputs, network complexity, and external noise. This work by Saiki-Ishikawa et al. examines the functional hierarchy between premotor (PM) and primary motor (M1) cortices in mice during a directional reaching task. The authors find some evidence consistent with an asymmetric reciprocal influence between the regions, but overall, activity patterns were highly similar and equally predictive of one another. These results suggest that motor cortical hierarchy, though present, is not fully reflected in firing patterns alone.

Strengths:

Inferring functional hierarchies between brain regions, given the complexity of reciprocal and local connectivity, dynamic interactions, and the influence of both shared and independent external inputs, is a challenging task. It requires careful analysis of simultaneous recording data, combined with cross-validation across multiple metrics, to accurately assess the functional relationships between regions. The authors have generated a valuable dataset simultaneously recording from both regions at scale from mice performing a cortex-dependent directional reaching task.

Using electrophysiological and silencing data, the authors found evidence supporting the traditionally assumed asymmetric influence from PM to M1. While earlier studies inferred a functional hierarchy based on partial temporal relationships in firing patterns, the authors applied a series of complementary analyses to rigorously test this hierarchy at both individual neuron and population levels, with robust statistical validation of significance.

In addition, recording combined with brief optogenetic silencing of the other region allowed authors to infer the asymmetric functional influence in a more causal manner. This experiment is well designed to focus on the effect of inactivation manifesting through oligosynaptic connections to support the existence of a premotor to primary motor functional hierarchy.

Subsequent analyses revealed a more complex picture. CCA, PLS, and three measures of predictivity (Granger causality, transfer entropy, and convergent cross-mapping) emphasized similarities in firing patterns and cross-region predictability. However, DLAG suggested an imbalance, with RFA capturing CFA variance at a negative time lag, indicating that RFA 'leads' CFA. Taken together these results provide useful insights for current studies of functional hierarchy about potential limitations in inferring hierarchy solely based on firing rates.

While I would detail some questions and issues on specifics of data analyses and modeling below, I appreciate the authors' effort in training RNNs that match some behavioral and recorded neural activity patterns including the inactivation result. The authors point out two components that can determine the across-region influence - 1) the amount of inputs received and 2) the dependence on across-region input, i.e., the relative importance of local dynamics, providing useful insights in inferring functional relationships across regions.

Weaknesses:

(1) Trial-averaging was applied in CCA and PLS analyses. While trial-averaging can be appropriate in certain cases, it leads to the loss of trial-to-trial variance, potentially inflating the perceived similarities between the activity in the two regions (Figure 4). Do authors observe comparable degrees of similarity, e.g., variance explained by canonical variables? Also, the authors report conflicting findings regarding the temporal relationship between RFA and CFA when using CCA/PLS versus DLAG. Could this discrepancy be due to the use of trial-averaging in former analyses but not in the latter?

We certainly agree that the similarity in firing patterns is higher in trial averages than on single trials, given the variation in single-neuron firing patterns across trials. Here, we were trying to examine the similarity of activity variance that is clearly movement dependent, as trial averages are, and to use an approach aligned with those applied in the existing literature. We would also agree that there is more that can be learned about interactions from trial-by-trial analysis. It is possible that the activity components identified by DLAG as being asymmetric somehow are not reflected strongly in trial averages. In our Discussion we offer another potential explanation that is based on other differences in what is calculated by DLAG and CCA/PLS.

We also note here that all of the firing pattern predictivity analysis we report (Figure 6) was done on single trial data, and in all cases the predictivity was symmetric. Thus, our results in aggregate are not consistent with symmetry purely being an artifact of trial averaging.

(2) A key strength of the current study is the precise tracking of forelimb muscle activity during a complex motor task involving reaching for four different targets. This rich behavioral data is rarely collected in mice and offers a valuable opportunity to investigate the behavioral relevance of the PM-M1 functional interaction, yet little has been done to explore this aspect in depth. For example, single-trial time courses of inter-regional latent variables acquired from DLAG analysis can be correlated with single-trial muscle activity and/or reach trajectories to examine the behavioral relevance of inter-regional dynamics. Namely, can trial-by-trial change in inter-regional dynamics explain behavioral variability across trials and/or targets? Does the inter-area interaction change in error trials? Furthermore, the authors could quantify the relative contribution of across-area versus within-area dynamics to behavioral variability. It would also be interesting to assess the degree to which across-area and within-area dynamics are correlated. Specifically, can acrossarea dynamics vary independently from within-area dynamics across trials, potentially operating through a distinct communication subspace?

These are all very interesting questions. Our study does not attempt to parse activity into components predictive of muscle activity and others that may reflect other functions. Distinct components of RFA and CFA activity may very well rely on distinct interactions between them.

(3) While network modeling of RFA and CFA activity captured some aspects of behavioral and neural data, I wonder if certain findings such as the connection weight distribution (Figure 7C), across-region input (Figure 7F), and the within-region weights (Figure 7G), primarily resulted from fitting the different overall firing rates between the two regions with CFA exhibiting higher average firing rates. Did the authors account for this firing rate disparity when training the RNNs?

The key comparison in Figure 7 is shown in 7F, where the firing rates are accounted for in calculating the across-region input strength. Equalizing the firing rates in RFA and CFA would effectively increase RFA rates. If the mean firing rates in each region were appreciably dependent on across-region inputs, we would then expect an off-setting change in the RFA → CFA weights, such that the RFA → CFA distributions in 7F would stay the same. We would also expect the CFA → RFA weights would increase, since RFA neurons would need more input. This would shift the CFA → RFA (blue) distributions up. Thus, if anything, the key difference in this panel would only get larger.

We also generally feel that it is a better approach to fit the actual firing rates, rather than normalizing, since normalizing the firing rates would take us further from the actual biology, not closer.

(4) Another way to assess the functional hierarchy is by comparing the time courses of movement representation between the two regions. For example, a linear decoder could be used to compare the amount of information about muscle activity and/or target location as well as time courses thereof between the two regions. This approach is advantageous because it incorporates behavior rather than focusing solely on neural activity. Since one of the main claims of this study is the limitation of inferring functional hierarchy from firing rate data alone, the authors should use the behavior as a lens for examining inter-areal interactions.

As we state above, we agree that examining interactions specific to movement-related activity components could reveal interesting structure in interregional interactions. Since it remains a challenge to rigorously identify a subset of neural activity patterns specifically related to driving muscle activity, any such analysis would involve an additional assumption. It remains unclear how well the activity that decoders use for predicting muscle activity matches the activity that actually drives muscle activity in situ.

To address this issue, which related to one raised by Reviewer #3 below, we have added an additional paragraph to the Discussion (see “Manifestations of hierarchy in firing patterns”).

Reviewer #3 (Public review):

This study investigates how two cortical regions that are central to the study of rodent motor control (rostral forelimb area, RFA, and caudal forelimb area, CFA) interact during directional forelimb reaching in mice. The authors investigate this interaction using

(1) optogenetic manipulations in one area while recording extracellularly from the other, (2) statistical analyses of simultaneous CFA/RFA extracellular recordings, and (3) network modeling.

The authors provide solid evidence that asymmetry between RFA and CFA can be observed, although such asymmetry is only observed in certain experimental and analytical contexts.

The authors find asymmetry when applying optogenetic perturbations, reporting a greater impact of RFA inactivation on CFA activity than vice-versa. The authors then

investigate asymmetry in endogenous activity during forelimb movements and find asymmetry with some analytical methods but not others. Asymmetry was observed in the onset timing of movement-related deviations of local latent components with RFA leading CFA (computed with PCA) and in a relatively higher proportion and importance of cross-area latent components with RFA leading than CFA leading (computed with DLAG). However, no asymmetry was observed using several other methods that compute cross-area latent dynamics, nor with methods computed on individual neuron pairs across regions. The authors follow up this experimental work by developing a twoarea model with asymmetric dependence on cross-area input. This model is used to show that differences in local connectivity can drive asymmetry between two areas with equal amounts of across-region input.

Overall, this work provides a useful demonstration that different cross-area analysis methods result in different conclusions regarding asymmetric interactions between brain areas and suggests careful consideration of methods when analyzing such networks is critical. A deeper examination of why different analytical methods result in observed asymmetry or no asymmetry, analyses that specifically examine neural dynamics informative about details of the movement, or a biological investigation of the hypothesis provided by the model would provide greater clarity regarding the interaction between RFA and CFA.

Strengths:

The authors are rigorous in their experimental and analytical methods, carefully monitoring the impact of their perturbations with simultaneous recordings, and providing valid controls for their analytical methods. They cite relevant previous literature that largely agrees with the current work, highlighting the continued ambiguity regarding the extent to which there exists an asymmetry in endogenous activity between RFA and CFA.

A strength of the paper is the evidence for asymmetry provided by optogenetic manipulation. They show that RFA inactivation causes a greater absolute difference in muscle activity than CFA interaction (deviations begin 25-50 ms after laser onset, Figure 1) and that RFA inactivation causes a relatively larger decrease in CFA firing rate than CFA inactivation causes in RFA (deviations begin <25ms after laser onset, Figure 3). The timescales of these changes provide solid evidence for an asymmetry in the impact of inactivating RFA/CFA on the other region that could not be driven by differences in feedback from disrupted movement (which would appear with a ~50ms delay).

The authors also utilize a range of different analytical methods, showing an interesting difference between some population-based methods (PCA, DLAG) that observe asymmetry, and single neuron pair methods (granger causality, transfer entropy, and convergent cross mapping) that do not. Moreover, the modeling work presents an interesting potential cause of "hierarchy" or "asymmetry" between brain areas: local connectivity that impacts dependence on across-region input, rather than the amount of across-region input actually present.

Weaknesses:

There is no attempt to examine neural dynamics that are specifically relevant/informative about the details of the ongoing forelimb movement (e.g., kinematics, reach direction). Thus, it may be preemptive to claim that firing patterns alone do not reflect functional influence between RFA/CFA. For example, given evidence that the largest component of motor cortical activity doesn't reflect details of ongoing movement (reach direction or path; Kaufman, et al. PMID: 27761519) and that the analytical tools the authors use likely isolate this component (PCA, CCA), it may not be

surprising that CFA and RFA do not show asymmetry if such asymmetry is related to the control of movement details.

An asymmetry may still exist in the components of neural activity that encode information about movement details, and thus it may be necessary to isolate and examine the interaction of behaviorally-relevant dynamics (e.g., Sani, et al. PMID: 33169030).

To clarify, we are not claiming that firing patterns in no way reflect the asymmetric functional influence that we demonstrate with optogenetic inactivation. Instead, we show that certain types of analysis that we might expect to reflect such influence, in fact, do not. Indeed, DLAG did exhibit asymmetries that matched those seen in functional influence (at least qualitatively), though other methods we applied did not.

As we state above, we do think that there is more that can be gleaned by looking at influence specifically in terms of activity related to movement. However, if we did find that movement-related activity exhibited an asymmetry following functional influence, our results imply that the remaining activity components would exhibit an opposite asymmetry, such that the overall balance is symmetric. This would itself be surprising. We also note that the components identified by CCA and PLS do show substantial variation across reach targets, indicating that they are not only reflecting condition-invariant components. These analyses were performed on components accounting for well over 90% of the total activity variance, suggesting that both condition-dependent and condition-invariant components should be included.

To address the concern about condition-dependent and condition-invariant components, we have added a sentence to the Results section reporting our CCA and PLS results: “Because our results here involve the vast majority of trial-averaged activity variance, we expect that they encompass both components of activity that vary for different movement conditions (condition-dependent), and those that do not (condition-invariant).” To address the general concerns about potential differences in activity components specifically related to muscle activity, we have also added an additional paragraph to the Discussion (see “Manifestations of hierarchy in firing patterns”).

The idea that local circuit dynamics play a central role in determining the asymmetry between RFA and CFA is not supported by experimental data in this paper. The plausibility of this hypothesis is supported by the model but is not explored in any analyses of the experimental data collected. Given the focus on this idea in the discussion, further experimental investigation is warranted.

While we do not provide experimental support for this hypothesis, the data we present also do not contradict this hypothesis. Here we used modeling as it is often used - to capture experimental results and generate hypotheses about potential explanation. We do feel that our Discussion makes clear where the hypothesis derives from and does not misrepresent the lack of experimental support. We expect readers will take our engagement with this hypothesis with the appropriate grain of salt. The imaginable experiments to support such a hypothesis would constitute another substantial study, requiring numerous controls - a whole other paper in itself.

Recommendations for the authors:

Reviewer #1 (Recommendations for the authors):

(1) There are a few small text/figure caption modifications that can be made for clarity of reading:

(2) Unclear sentence in the second paragraph of the introduction: "For example, stimulation applied in PM has been shown to alter the effects on muscles of stimulation in M1 under anesthesia, both in monkeys and rodents."

This sentence has been rephrased for clarity: "For example, in anesthetized monkeys³⁴ and rodents³⁵, stimulation in PM alters the effects of stimulation in M1 on muscles."

(3) The first section of the results presents the optogenetic manipulation. However, the critical control that tests whether this was strictly a local manipulation that did not affect cells in the other region is introduced only much later. It may be helpful to add a comment in this section noting that such a control was performed, even if it is explained in detail later when introducing the recordings.

We have added the following to the first Results section: "we show below that direct optogenetic effects were only seen in the targeted forelimb area and not the other."

(4) Figure 1D - I imagine these averages are from a single animal, but this is not stated in the figure caption.

"For one example mouse," has been added to the beginning of the Figure 1D legend.

(5) Figure 2F - N=6 is not stated in the panel's caption (though it can make it clearer), while it is stated in the caption of 2H.

"n = 6 mice" has been added to the Figure 2F legend.

(6) There's some inconsistency with the order of RFA/CFA in the figures, sometimes RFA is presented first (e.g., Figure 1D and 1F), and sometimes CFA is presented first (e.g., panels of Figure 2).

We do not foresee this leading to confusion.

(7) "As expected, the majority of recorded neurons in each region exhibited an elevated average firing rate during movement as compared to periods when forelimb muscles were quiescent (Figure 2D,E; Figure S1A,B)" - Figure S1A,B show histograms of narrow vs. wide waveforms, is this the relevant figure here?

We apologize for the cryptic reference. The waveform width histograms were referred to here because they enabled the separation of narrow- and wide-waveform cells shown in Figure 2D,E. We have added the following clause to the referenced sentence to make this explicit: ", both for narrow-waveform, putative interneurons and wide-waveform putative pyramidal neurons."

(8) Figure 2I caption - "The fraction of activity variance from 150 ms before reach onset to 150 ms after it that occurs before reach onset" - this sentence is not clear.

The Figure 2I legend has been updated to "The activity variance in the 150 ms before muscle activity onset, defined as a fraction of the total activity variance from 150 ms before to 150 ms after muscle activity onset, for each animal (circles) and the mean across animals (black bars, n = 6 mice)."

(9) Figure 4B-G - is this showing results across the 6 animals? Not stated clearly.

Yes - the 21 sessions we had referred to are drawn from all six mice. We have updated the legend here to make this explicit.

(10) DLAG analysis - is there any particular reasoning behind choosing four across-region and four within-region components?

In actuality, we completed this analysis for a broad range of component numbers and obtained similar results in all cases. Four fell in the center of our range, and so we focused the illustrations shown in the figure on this value. In general, the number of components is arbitrary. The original paper from Gokcen et al. describes a method for identifying a lower bound on the number of distinct components the method can identify. However, this method yields different results for each individual recording session. For the comparisons we performed, we needed to use the same range of values for each session.

(11) Figure 5A seems to show 11 across-session components, it's unclear from the caption but I imagine this should show 12 (4 components times 3 sessions?)

As we state in the Methods, any across-region latent variable with a lag that failed to converge between the boundary values of ± 200 ms was removed from the analysis. In the case illustrated in this panel, the lag for one of the components failed to converge and is not shown. We have now clarified this both in the relevant Results paragraph and in the figure legend.

(12) Figure 5B - is each marker here the average variance explained by all across/within components that were within the specified lag criteria across sessions per mouse? In other words, what does a single marker here stand for?

We apologize for the lack of clarity here. These values reflect the average across sessions for each mouse. We have updated the legend to make this explicit.

Reviewer #2 (Recommendations for the authors):

As I have addressed most of my major recommendations in the public review, I will use this section to include relatively minor points for the authors to consider.

(1) The EMG data in Figure 1C shows distinct patterns across spouts, both in the magnitude and complexity of muscle activations. It would be interesting to investigate whether these differences in muscle activity lead to behavioral variations (e.g., reaction time, reach duration) and how they relate to the relative involvement of the two areas.

We agree that it would be interesting to examine how the interactions between areas vary as behavior varies. While the differences between reaches here are limited, we have addressed this question for two substantially different motor behaviors (reaching and climbing) in a follow-up study that was recently preprinted (Kristl et al., biorxiv, 2025).

(2) How do the authors account for the lingering impact of RFA inactivation on muscle activity, which persists for tens of milliseconds after laser offset? Could this effect be due to compensatory motor activity following the perturbation? A further illustration of how the raw limb trajectories and/or muscle activity are perturbed and recovered would help readers better understand the impact of motor cortical inactivation.

To clarify the effects of inactivation on a longer timescale, we have added a new supplemental figure showing the plots from Figure 1D over a longer time window extending to 500 ms after trial onset (new Figure S1). Lingering effects do persist, at least in certain

cases. In general, we find it hard to ascertain the source of optogenetic effects on longer timescales like this. On the shortest timescales, effects will be mediated by relatively direct connections between regions. However, on these longer timescales, effects could be due to broader changes in brain and behavioral state that can influence muscle activity. For example, attempts to compensate for the initial disturbance to muscle activity could cause divergence from controls on these longer timescales. Muscle tissue itself is also known to have long timescale relaxation dynamics, and it would not be surprising if the relevant control circuits here also had long timescales dynamics, such that we would not expect an immediate return to control when the light pulse ends. Because of this ambiguity, we generally avoid interpretation of optogenetic effects on these longer timescales.

Reviewer #3 (Recommendations for the authors):

(1) Page 9: ". We measured the time at which the activity state deviated from baseline preceding reach onset," - I cannot find how this deviation was defined (neither the baseline nor the threshold).

We have added text to the Figure 2G legend that explicitly states how the baseline and activity onset time were defined.

(2) Given the shape of the curves in Figure 2G, the significance of this result seems susceptible to slight modifications of what defines a baseline or a deviation threshold. For example, it looks like the circle for CFA has a higher y-axis value, suggesting the baseline deviance is higher, but it is unclear why that would be from the plot. If the threshold for deviation in neural activity state were held uniform between CFA and RFA is the difference still significant across animals?

We have repeated the analysis using the same absolute threshold for each region. We used the higher of the two thresholds from each region. The difference remains significant. This is now described in the last paragraph of the Results section for Figure 2.

(3) Since summed deviation of the top 3 PCs is used to show a difference in activity onset between CFA/RFA, but only a small proportion of variance is explained pre-movement (<2% in most animals), it seems relevant to understand what percentage of CFA/RFA neuron activity actually is modulated and deviates from baseline prior to movement and to show the distribution of activity onsets at the single neuron level in CFA/RFA. Can an onset difference only be observed using PCA?

Because many neurons have low firing rates, estimating the time at which their firing rate begins to rise near reach onset is difficult to do reliably. It is also true that not all neurons show an increase around onset - some show a decrease and others show no discernible change. Using PCs to measure onset avoids both of these problems, since they capture both increases and decreases in individual neuron firing rates and are much less noisy than individual neuron firing rates.

However, based on this comment, we have repeated this analysis on a single-neuron level using only neurons with relatively high average firing rates. Specifically, we analyzed neurons with mean firing rates above the 90th percentile across all sessions within an animal. Neurons whose activity never crossed threshold were excluded. Results matched those using PCs, with RFA neurons showing an earlier average activity onset time. This is now described in the last paragraph of the Results section for Figure 2.

(4) It is stated that to study the impact of inactivation on CFA/RFA activity, only the 50 highest average firing rate neurons were used (and maybe elsewhere too, e.g., convergent cross mapping). It is unclear why this subselection is necessary. It is justified

by stating that higher firing rate neurons have better firing rate estimates. This may be supportable for very low firing rate units that spike sorting tools have a hard time tracking, but I don't think this is supported by data for most of the distribution of firing rates. It therefore seems like the results might be biased by a subselection of certain high firing rate neuron populations. It would be useful to also compute and mention if the results for all neurons/neuron pairs are the same. If there is worry about low-quality units being those with low firing rates, a threshold for firing rate as used elsewhere in the paper (at least 1 spike / 2 trials) seems justified.

The issue here is that as firing rates decrease and firing rate estimates get noisier, estimates of the change in firing rate get more variable. Here we are trying to estimate the fraction of neurons for which firing rates decreased upon inactivation of the other region. Variability in estimates of the firing rate change will bias this estimate toward 50%, since in the limit when the change estimates are entirely based on noise, we expect 50% to be decreases. As expected, when we use increasingly liberal thresholds for this analysis, the fraction of decreases trends closer to 50%.

As a consequence of this, we cannot easily distinguish whether higher firing rate neurons might for some reason have a greater tendency to exhibit decreases in firing compared to lower firing rate neurons. However, we see no positive reason to expect such a difference. We have added a sentence noting this caveat in interpreting our findings to the relevant paragraph of the Results.

The lack of min/max axis values in Figure 3B-F makes it hard to interpret - are these neurons almost silent when near the bottom of the plot or are they still firing a substantial # of spikes?

To aid interpretation of the relative magnitude of firing rate changes, we have added minimum firing rates for the averages depicted in Figure 3B,C,E and F to the legend. Our original thinking was that the plots in Figure 3G and H would provide an indication of the relative changes in firing.

It would be interesting to know if the impact of optogenetic stimulation changed with exposure to the manipulation. Are all results presented only from the first X number of sessions in each animal? Or is the effect robust over time and (within the same animal) you can get the same results of optogenetic inactivation over time? This information seems critical for reproducibility.

We have now performed brief optogenetic inactivations in several brain areas in several different behavioral paradigms, and have found that inactivation effects are stable both within and across sessions, almost surprisingly so. This includes cases where the inactivations were more frequent (every ~1.25 s on average) and more numerous (>15,000 trials per animal) than in the present manuscript. Thus we did not restrict our analysis here to the first X sessions or trials within a session. We have added additional plots as Figure S3T-AA showing the stability of optogenetic effects both within and across sessions.

Given that it can be difficult to record from interneurons (as the proportion of putative interneurons in Figure S1 attests), the SALT analyses would be more convincing if a few recordings had been performed in the same region as optogenetic stimulation to show a "positive control" of what direct interneuron stimulation looks like. Could also use this to validate the narrow/wide waveform classification.

We have verified that using SALT as we have in the present manuscript does detect vGAT+ interneurons directly responding to light. This is included in a recent preprint from the lab

(Kristl et al., biorxiv, 2025). We (Warriner et al., Cell Reports, 2022) and others (Guo et al., Neuron, 2014) have previously used direct ChR2 activation to validate waveform-based classification.

Simultaneous CFA/RFA recordings during optogenetic perturbation would also allow for time courses of inhibition to be compared in RFA/CFA. Does it take 25ms to inhibit locally, and the cross-area impact is fast, or does it inactivate very fast locally and takes ~25ms to impact the other region?

Latencies of this sort are difficult to precisely measure given the statistical limits of this sort of data, but there does appear to be some degree of delay between local and downstream effects. We do not have a statistical foundation as of yet for concluding that this is the case. It will be interesting to examine this issue more rigorously in the future.

Given the difference in the analytical methods, the authors should share data in a relatively unprocessed format (e.g., spike times from sorted units relative to video tracking + behavioral data), along with analysis code, to allow others to investigate these differences.

We plan to post the data and code to our lab's Github site once the Version of Record is online.

<https://doi.org/10.7554/eLife.103069.2.sa0>