

Hierarchy in influence but not firing patterns among forelimb motor cortices

Akiko Saiki-Ishikawa†, Mark Agrios†, Sajishnu Savya†, Adam Forrest†, Hannah Sroussi,

Sarah Hsu, Diya Basrai, Feihong Xu, and Andrew Miri*

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

*Corresponding author. Email: andrewmiri@northwestern.edu

†These authors contributed equally to this work.

Abstract: Though hierarchy is commonly invoked in descriptions of neural system function, its presence and manifestation in firing patterns remain poorly resolved. Here we use rapid activity perturbations to demonstrate asymmetric influence between premotor and primary motor cortices on fast timescales during reaching in mice, confirming hierarchy among endogenous firing patterns. However, this hierarchy had a surprising manifestation: firing in each region was dominated by temporally-aligned patterns and was equally predictive of future firing in the other region. Network modeling suggests an explanation: functional hierarchy can be determined by the relative dependence of local circuits on across-region input, rather than asymmetry in that input. We propose that motor cortical firing patterns are primarily defined not by across-region input, but by the constraints of patterning motor output.

Main Text:

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 plan movements and primary motor cortex executes them (12, 13). A similar range of observations, in particular that of substantial premotor projections to brainstem and spinal cord (14-16), interconnection between premotor and primary motor cortices (17, 18), and activity related to movement planning in primary motor cortex and spinal cord (19, 20), inform an updated view of a partial hierarchy (21-23). Here, premotor and primary motor cortices interact and each drive downstream motor circuits, though primary motor cortex exerts comparatively more of its influence through its descending projections.

However, technical challenges have precluded demonstration of a hierarchy between premotor and primary motor cortices mediated by their naturally-occurring (endogenous) firing patterns. It is commonly assumed that these endogenous firing patterns will exhibit an asymmetric reciprocal influence, with activity in premotor cortex exerting a larger influence on primary motor cortex than vice versa, but existing observations have not yet resolved such an asymmetry. Notions of hierarchy are informed by observations like those of asymmetric effects on each region from electrical stimulation, lesion, or pharmacological inactivation of the other region (24, 25); of differing laminar targets of projections between regions (26, 27); and of differing degrees of activity related to movement preparation (28-30). Yet such observations do not necessarily imply asymmetric reciprocal influence of endogenous activity on the fast timescales of synaptic communication between regions (oligosynaptic timescales). Moreover, a number of observations

are consistent with a more symmetric influence between premotor and primary motor cortices (31-33). It has also been proposed that existing observations of asymmetry could instead reflect involvement of each region in somewhat different movement components (34, 35).

The ambiguity about hierarchy has in turn prevented resolution of how it manifests in firing patterns. Movement-related activity is frequently observed to begin in premotor cortices before it begins in primary motor cortex (9, 36, 37); this delay has been interpreted as reflecting feedforward influence. The application of time series prediction methods for discerning functional influence between neurons (38, 39) is grounded in the notion that feedforward influence leads to predictive relationships between firing patterns in upstream and downstream regions. We might expect that a (partial) hierarchy between regions would manifest as an asymmetry in firing pattern predictivity, where activity in the (more) upstream region is comparatively better at predicting subsequent activity in the (more) downstream region. 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 (40, 41). Thus, whether hierarchy manifests in firing pattern predictivity at the scale of brain regions within neural systems remains unresolved.

RFA and CFA activity during directional reaching

We sought to establish (1) the long-hypothesized hierarchy between premotor and primary motor cortex mediated by their endogenous firing patterns, and (2) how any such hierarchy would manifest in firing pattern predictivity. Building on recent work (42), we developed a head-fixed directional reaching paradigm in which mice learn to rest their hand on a rung and then reach to one of four spouts to grab a water droplet (Fig. 1A,B; Fig. S1A; 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. In trained VGAT-ChR2-EYFP mice (43), we briefly projected a small spot of blue light (50 ms, 1.5 mm diameter, 9 mW/mm²) onto the cortical surface unilaterally, triggered on reach onset, to silence either forelimb primary motor (caudal forelimb area, CFA) or premotor (rostral forelimb area, RFA) cortex. Forelimb muscle activity, recorded with chronically implanted electromyographic electrodes, rapidly diverged from controls (Fig. S1B-D), indicating that both regions exert oligosynaptic influence on muscles. Simultaneous large-scale multielectrode array (Neuropixel) recordings from both CFA and RFA (all layers; 21 sessions across six mice; Fig. S2A,B) confirmed that movement-related activity changes begin earlier in RFA (Fig. 1C,D), a feature of motor cortical activity assumed to reflect hierarchy.

We next showed that greater early activity change in RFA is a robust feature of firing patterns en masse. We used delayed latents across groups (DLAG; 44) 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). We decomposed each region’s activity into four across-region components (negative lags defined as RFA leading) and four within-region components. The distribution of resulting across-region lags that were significantly different from zero was biased in the direction of RFA leading (median = -18 ms; Fig. 1E; Fig. S2C), and remained so after weighting each count in the distribution by the fraction of neural activity variance it captured in the lagging region (Fig. 1F; Fig. S2D). The same was found when repeating this analysis but varying either the number of within-region (Fig. 1G; Fig. S2E) or across-region (Fig. 1H; Fig. S2F) components, while holding

the other fixed at four. Thus, a greater fraction of activity variance in CFA is captured by shared patterns in which RFA leads, compared to vice versa.

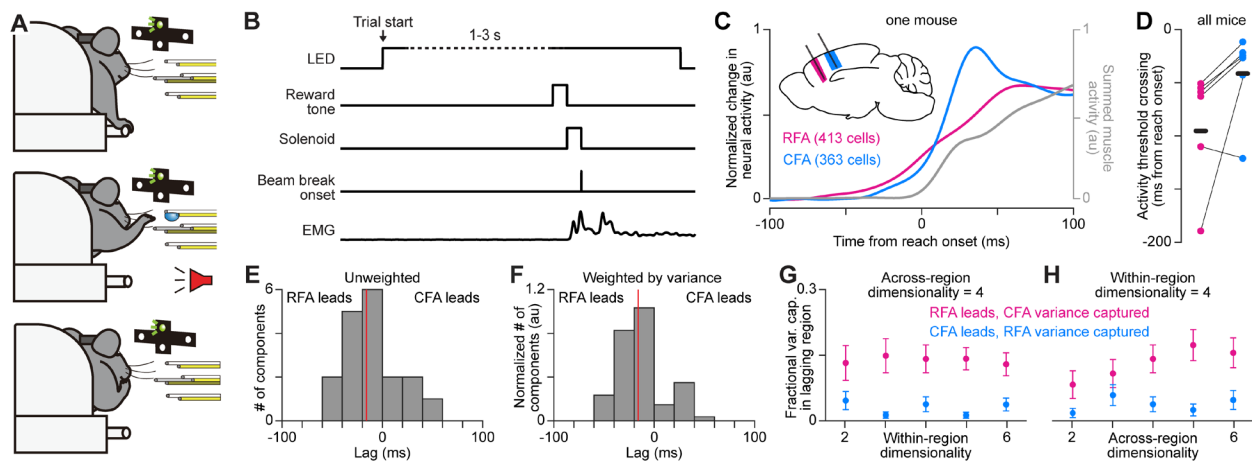


Fig. 1. RFA and CFA activity during directional reaching in mice. (A) Schematic depicting the directional reaching task. (B) Schematic depicting the time course of experimental control signals and muscle activity during the directional reaching task. (C) For one mouse, normalized activity change from baseline summed across the top three principal components (PCs) for all recorded RFA (magenta) or CFA (blue) neurons, and the top PC for muscles (gray). For calculations in (C) and (D), downward-going PC time courses were flipped vertically so changes from baseline were additive. (D) Time from reach onset at which the activity change from baseline summed across the top three PCs for RFA (magenta) or CFA (blue) neurons rose above a low threshold. (E),(F) Histogram of DLAG component lags that were significantly different from zero and between -100 and 100 ms, either unweighted (E) or weighted (F) by the variance each component captures in the lagging region, for all recording sessions together. (G),(H) Mean \pm SEM fractional activity variance ($n = 21$ recording sessions) captured in the lagging region by DLAG components when the lag was significantly different from zero and between -100 and 100 ms when varying the within-region (G) or the across-region (H) dimensionality.

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. We inactivated RFA as described above while recording activity across layers in CFA (Fig. 2A-C, $n = 3$ mice), or similarly inactivated CFA while recording in RFA (Fig. 2D-F, $n = 3$ mice). This approach silences the vast majority of endogenous activity in projection neurons across all of RFA or most of CFA at least down through layer 5 (45, 46), while affecting similarly-sized cortical regions. Analysis of spike latencies after light onset found no evidence of direct light responses by ChR2-expressing neurons in the recorded region (Fig. S3A-D).

Results indicated an asymmetry in the oligosynaptic functional influence between CFA and RFA. Control and inactivation trial averages showed a larger and longer lasting reduction in CFA firing upon RFA inactivation compared to the effect on RFA firing upon CFA inactivation (Fig. 2A-F). To quantify the effect on firing in individual neurons, we computed the difference of control and inactivation trial averages 50 ms after light/trial onset using z-scored firing rates for each neuron.

Distributions of these relative firing rate changes showed more cells with a larger reduction in CFA compared to those in RFA, both for all neurons (Fig. 2G) and for wide-waveform, putative pyramidal neurons (Fig. 2H). These distributions did show more RFA neurons responding with an increase in firing; however, the average absolute change in firing was also substantially larger for CFA neurons than RFA neurons, again both for all neurons (Fig. 2I) and the wide-waveform subset (Fig. 2J). Moreover, the average absolute difference was significantly larger whether calculated over the 25, 50 or 100 ms following light/trial onset (Fig. 2K,L). Inactivation effects were similar for both wide-waveform neurons and narrow-waveform, putative inhibitory interneurons (Fig. S3E-P). To our knowledge, this is the first demonstration that the endogenous activity in premotor and primary motor cortices exert asymmetric effects on one another. We interpret this as a direct indication of functional hierarchy on oligosynaptic timescales.

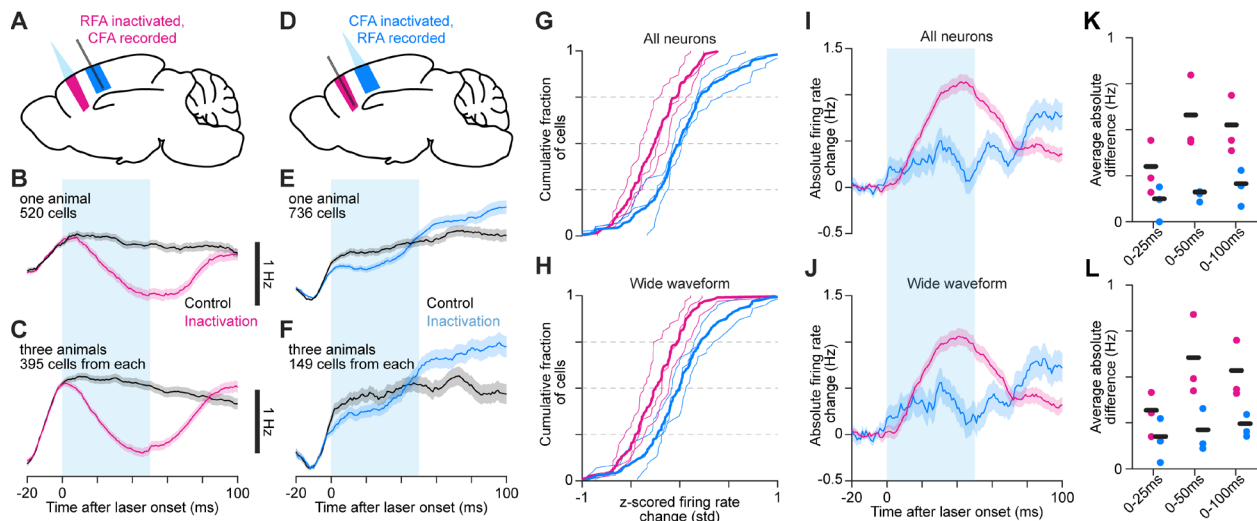


Fig. 2. Asymmetric reciprocal influence of endogenous RFA and CFA activity. (A-F) Schematics depicting inactivation and Neuropixel recording (A,D) and mean \pm SEM firing rate time series for neurons from one animal (B,E) or all three animals (C,F) from inactivating RFA and recording CFA (A-C), or vice versa (D-F). The cyan bar indicates light on. The same number of cells were used from each animal in (C) and (F). (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 trial/light onset, for the top 50 highest firing rate neurons (G) or wide waveform neurons (H) from all animals (thick lines) or individual animals (thin lines). (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. Baseline subtraction enables negative values. (K),(L) Average absolute firing rate difference from 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).

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. We first assessed the overall similarity of firing patterns in the two regions using 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; 47, 48), which maximizes the correlation between the time courses of component pairs, and partial least squares (PLS; 49), which maximizes their covariance. Both methods were applied to matrices comprising the trial averaged firing rates for all neurons from a given animal, concatenating separate averages aligned on reach onset and on grasp for reaches to each of the four spouts (Figs. 3A and S4A).

Both methods produced a similar basic result: firing patterns in RFA and CFA en masse are remarkably similar. 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 (Fig. 3B-D). PLS identified components that successively capture nearly the same amount of activity variance as corresponding principal components (i.e. the maximum of any possible components) while still maximizing the covariance of component time series (Fig. 3E,F).

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 relative to the other. This was not the case for RFA and CFA activity during reaching. Shifting CFA activity anywhere from -30 to 30 ms relative to RFA activity gave nearly identical alignment quality, both with CCA (Fig. 3G) and PLS (Fig. S4B). However, control analyses in which we aligned CFA activity to itself revealed clear maxima in alignment quality at the expected lags (Fig. S4C,D). That CFA-RFA alignment does not peak at certain lags could be explained by the across-region components DLAG found, with either RFA or CFA activity leading. More CFA variance was captured by components in which RFA activity led, compared to vice versa, but components in which CFA activity led were also identified. Though unequal in variance capture, these components can accommodate similar correlation or covariance when activity is aligned at a lag, since paired aligned components can differ in variance capture. Thus, collectively, these results indicate similar activity patterns exist in both regions, including those that align at a lag, though the variance they capture may vary between regions.

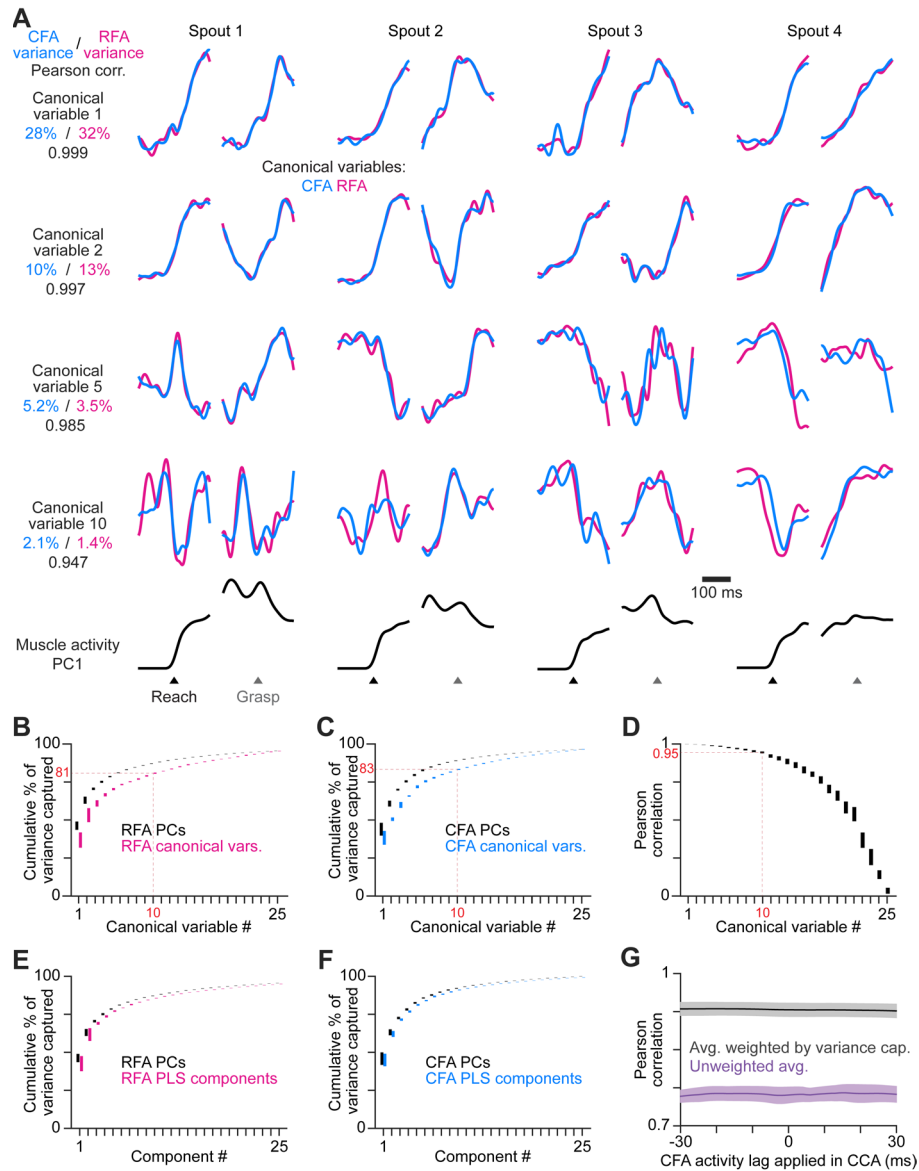


Fig. 3. RFA and CFA firing patterns are highly similar. (A) From one animal, four canonical variables aligned at reach onset or spout contact (grasp) for reaches to each spout. Bottom row shows the corresponding time series for the first principal component of muscle activity. (B),(C) Mean \pm SEM cumulative variance capture for canonical variables (color) and principal components (PCs, black) for RFA (B) and CFA (C) activity ($n = 21$ sessions). Red annotations facilitate comparisons across B-D. (D) Mean \pm SEM Pearson correlation for canonical variable pairs. (E),(F) Mean \pm SEM cumulative variance capture for PLS components (color) and principal components (PCs, black) 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, for the average over all pairs either without (purple) or with weighting each correlation value by the variance captured by the given pair (black).

Symmetric firing pattern predictivity

We then asked whether asymmetric reciprocal influence between CFA and RFA manifests in an imbalance in the ability to predict the activity in one region from that in the other. Though the methods we applied 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 subsequent firing in another. For comprehensiveness, we applied three metrics that each make different assumptions about how one firing pattern enables or improves prediction of another: transfer entropy (38), point-process Granger causality (39), and convergent cross mapping (50).

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 identified 10,000 pairs with the highest product of average firing rates. To avoid directional bias due to differing firing rates (CFA 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 (Fig. S5A,B). To assign a p-value to each metric computed for each cell pair, we generated an empirical null distribution by recomputing the metric after many different circular permutations of the spike time series.

Results revealed a striking symmetry in the degree of firing pattern predictivity in each direction. For all three metrics, p-value distributions for predictions in each direction were skewed toward zero (Fig. 4A-C), indicating that for some fraction of pairs, source activity did enable or improve predictions (51). Yet for all three metrics, p-value distributions for predictions in each direction were highly similar. This implies a similar fraction of pairs in each direction exhibit predictivity above chance. Furthermore, the degree of predictivity in each direction appeared similar; distributions of metric values were very similar for pairs with p-values below a significance threshold (false discovery rate < 0.10; Fig. 4D-F). Interestingly, metric value distributions computed instead for pairs of neurons within each region were also fairly similar, though there were more large extreme values, potentially reflecting strongly coupled cells (Fig. 4G-I). These results held when repeated using only spiking from the 300 ms before reach initiation (Fig. S5C-H). Thus, firing pattern predictivity does not reflect an asymmetry like that seen in the actual influence of endogenous activity.

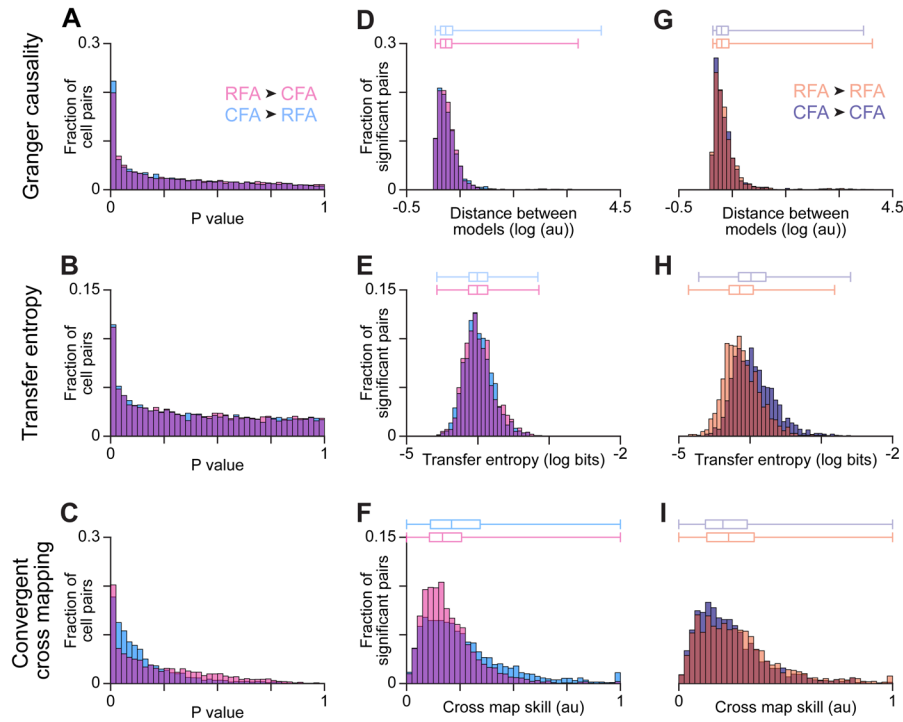


Fig. 4. RFA and CFA firing pattern predictivity. (A-C) P-value distributions from calculations of Granger causality, transfer entropy, and convergent cross mapping using firing patterns of across-region neuron pairs. (D-F) For across-region neuron pairs, distributions of metric values reflecting the improved prediction of target neuron firing using source neuron firing (D and E), or the improved prediction of source neuron firing using target neuron firing (F). In (D-I), only values for which the corresponding p-value fell below a false discovery-corrected threshold are included, and box plots on top show the minimum, 1st, 2nd, and 3rd quartile, and maximum values. (G-I) Distributions of metric values computed instead for within-region neuron pairs.

A network model recapitulates findings without asymmetric across-region input

How do we reconcile this symmetry in firing pattern predictivity with our observation of asymmetric reciprocal influence? Asymmetric influence might arise not from asymmetric across-region input, but instead from a difference in the relative dependence of activity in each region on that 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 questioned whether the asymmetric reciprocal influence we observed might naturally emerge from local circuit interactions, rather than asymmetric across-region input. To address this, we 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 (Fig. 5A). 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 indeed, 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 (Fig. 5B). Models collectively had slightly larger “RFA” to “CFA” across-region weights (Fig. 5C). 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 (Fig. 5D). 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 model instances (Fig. 5E). 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 (Fig. 5F). These results indicate that asymmetric reciprocal influence can emerge in these networks from local recurrent connectivity and the activity dynamics they engender.

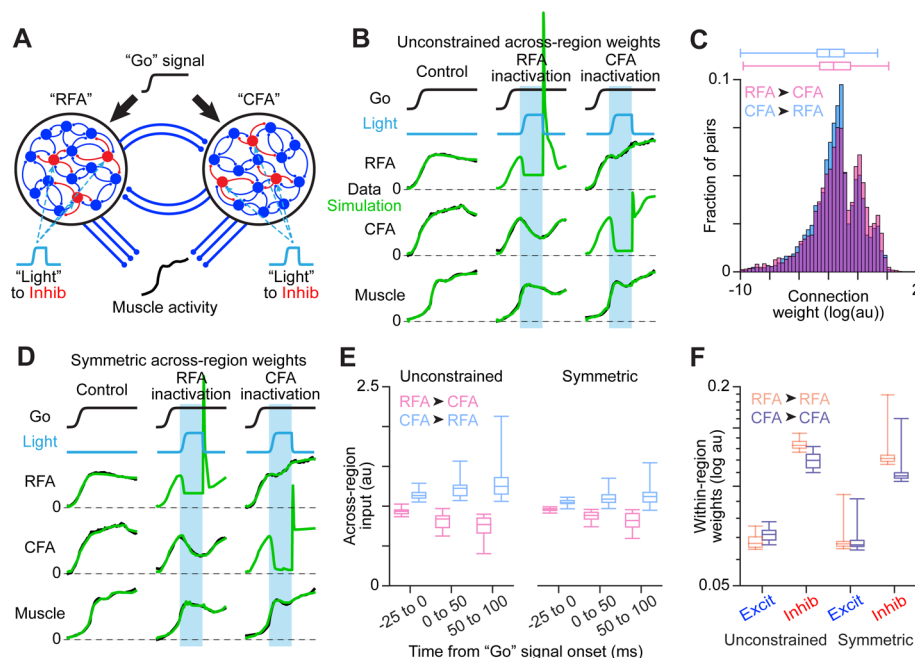


Fig. 5. Network modeling of RFA and CFA activity. (A) Model schematic. (B) Illustration of fit quality for one instance of the model with unconstrained weights of across-region synapses. In (B) and (D), cyan bars show the epochs of simulated inactivation. (C) Distribution of synaptic weights for all across-region connections from all instances of the unconstrained model. In (C), (E), and (F), box plots show the minimum, 1st, 2nd, and 3rd quartile, and maximum values. (D) Illustration of fit quality for one instance of the model with across-region weights constrained to

be equal in each direction on average. (E) 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. (F) Distributions of within-region excitatory (Excit) and inhibitory (Inhib) synaptic weights, for all instances of both the unconstrained and symmetric model types.

Discussion

Our network modeling suggests an explanation for the decoupling between functional influence and firing pattern we observe: functional hierarchy can be determined by the robustness of local circuits to loss of input, rather than asymmetric across-region input. These results highlight the dual determinants of hierarchy among brain regions: the balance of across-region input and the relative dependence of regions on that input. Though classical descriptions emphasized the former, contemporary recognition of the prominence of local recurrence and of the activity dynamics it engenders underscore the latter. Imagine a scenario in which inactivation of either region had symmetric effects on the other, indicating a lack of functional hierarchy. This could result if across-region input was imbalanced, but the dependence of each region on that input was also imbalanced, such that the effects offset. Here local circuit dynamics would substantially determine hierarchy. Our findings thus call into question previous interpretations of earlier rising activity in premotor cortex as reflecting a hierarchy mediated by feedforward influence (9, 19, 29, 37). 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 differing robustness to input loss.

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. Certain analyses of activity in premotor and primary motor cortices have described highly similar firing patterns (32, 52, 53), similarities in preparatory activity (54), and symmetry of modeled reciprocal input (31). Our results show that these observations do not contradict notions of motor cortical hierarchy. Differences in firing patterns between premotor and primary motor cortices could reflect not hierarchy, but differences in functional roles (22). For example, premotor cortex may exhibit more and earlier preparatory activity because the components of movement it drives demand more and earlier preparation.

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 (40, 41). Across-region connections may instead function to coordinate firing patterns across regions or modulate higher-order firing pattern features, like their autocorrelation (55). Firing patterns themselves 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 (56, 57) to some extent (58). Our results underscore the recognized need (40) for developing theory about across-region interactions.

Materials and 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 task (42). 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 (Fig. 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 consume 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 electrical artifacts. Thus during neural recording, spout touches were detected with an infrared beam sensor (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 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 procedures (59). 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 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 uL 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.

EMG Recording

EMG electrodes were fabricated for forelimb muscle recording using established procedures (45, 60). 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 (45, 61). 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/mm^2 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 uV). 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 performed 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, HiroSugi-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 Kilosort3 (62) 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 having 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 threshold (45). 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.

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 the n highest firing rate neurons from the other 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 SALT (63), 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 (Fig. S3A-D) 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 x 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 method (44). 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 proportions of variance captured by across-region latents with statistically significant delays. We grouped the latents by the sign of their delays (i.e. which region led), and each across-region latent's variance was defined as the variance captured in the lagging region (i.e. the downstream region).

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 N_{CFA} neurons and 1600 time points \times N_{RFA} neurons for each recording session, where N_{CFA} is the number of included neurons from CFA and N_{RFA} is the number of included neurons from RFA. We denote these matrices X_{CFA} and X_{RFA} 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 variant (49) was used. PLS-SVD is done by performing singular value decomposition on the cross-covariance matrix $X^T Y$ 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 (Fig. S4).

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.

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.

Transfer Entropy

Transfer entropy analysis was performed using a publicly available algorithm (38). 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 Ito 2011. 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 (GC; 39), 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.

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 (Fig. S6A). 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 (Fig. S6B-C).

The GC algorithm additionally computes the model difference test statistic for all within-region pairs. Using a subset of these pairs, the predictive power of the CFA and RFA intraregional neuronal population was also assessed. GC calculations for both across-region and within-region pairs were also repeated using spike trains containing only the 300 ms immediately prior to reach initiation.

Convergent Cross Mapping

Convergent Cross Mapping (CCM) was implemented using a modified version of MATLAB function ‘Sugi’ (64). 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 Transfer Entropy (TE) and Granger Causality (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” (50, 65).

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 activity (66). 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 ($E+1$). τ , which indicates how many timesteps each shadow manifold dimension is lagged in time $[t, t-1*\tau \text{ ms}, t-2*\tau \text{ ms}...]$, was set to 1.

Network Modeling

A recurrent neural network was trained to replicate experimental data. The package PsychRNN (67) 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:

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

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→CFA and CFA→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.

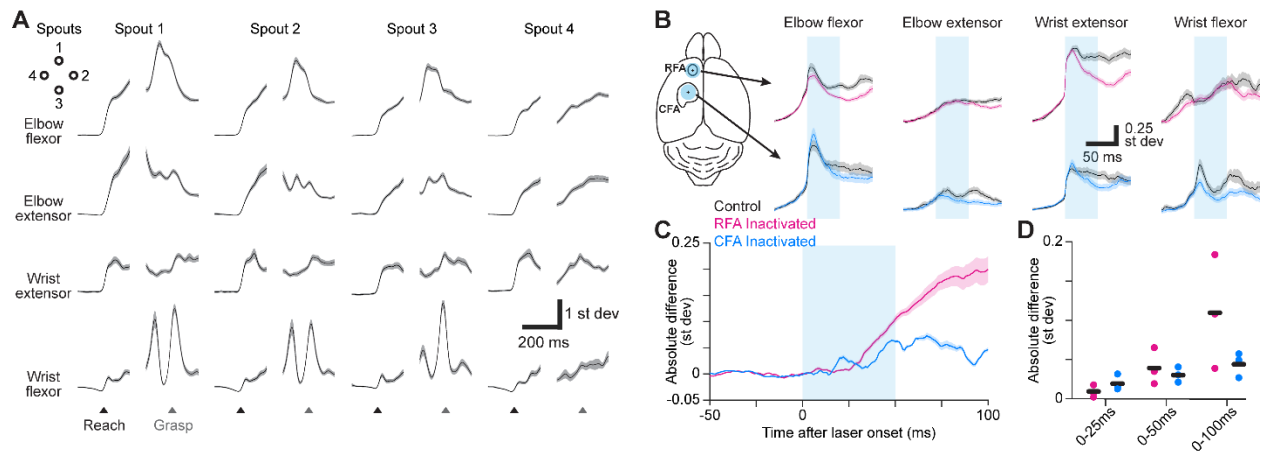


Fig. S1. Forelimb muscle activity and cortical inactivation. (A) Trial-averaged activity of four recorded muscles (black) \pm SEM (gray, $n=176$ trials) for one mouse during directional reaching towards each of four spouts. Separate averages are aligned on reach onset or spout contact (grasp). Vertical scale bar in (A) and (B) reflects standard deviation of Z-scored muscle activity. (B) Mean muscle activity \pm SEM 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. (C) 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. (D) Absolute difference between inactivation and control trials averaged over three epochs after trial/light onset, for individual animals (circles) and the mean across animals (black bars).

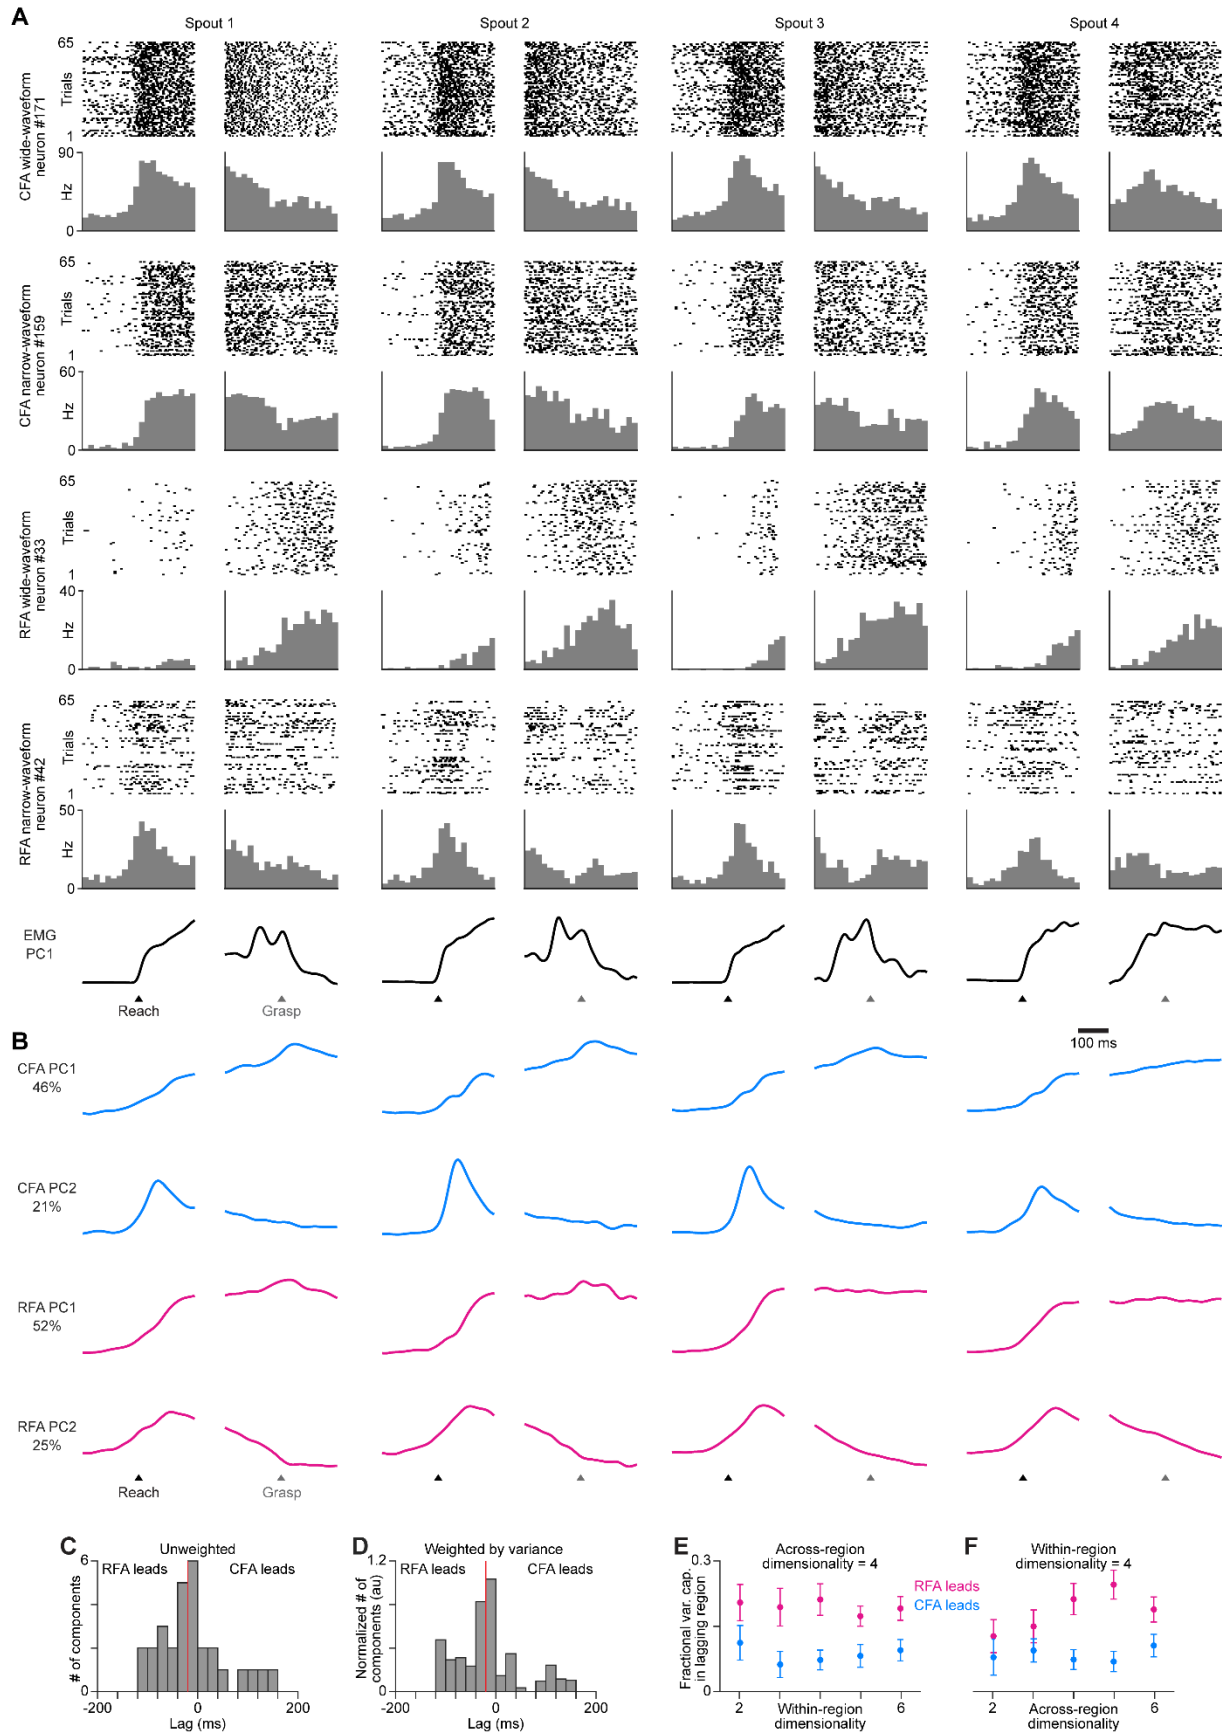


Fig. S2. RFA and CFA activity during directional reaching in mice. (A) For four example neurons, spike rasters (top) and binned firing rate (bottom) across all successful reach trials to each spout, aligned on reach onset (black arrowheads) and spout contact (Grasp, gray arrowheads). Bottom row shows the corresponding time series for the first principal component of muscle activity. (B) CFA (blue) and RFA (magenta) activity projected onto their respective first two principal components, from the recording involving the neurons shown in (A). (C)-(F) are similar to Fig. 1E-H, but include all components identified by DLAG having lags between -200 and 200 ms, not just those having lags between -100 and 100 ms. (C),(D) Histogram of DLAG component lags that were significantly different from zero, either unweighted (C) or weighted (D) by the variance each component captures in the lagging region, for all recording sessions together. (E),(F) Mean \pm SEM fractional activity variance ($n = 21$ recording sessions) captured in the lagging region by DLAG components when the lag was significantly different from zero when varying the within-region (E) or the across-region (F) dimensionality.

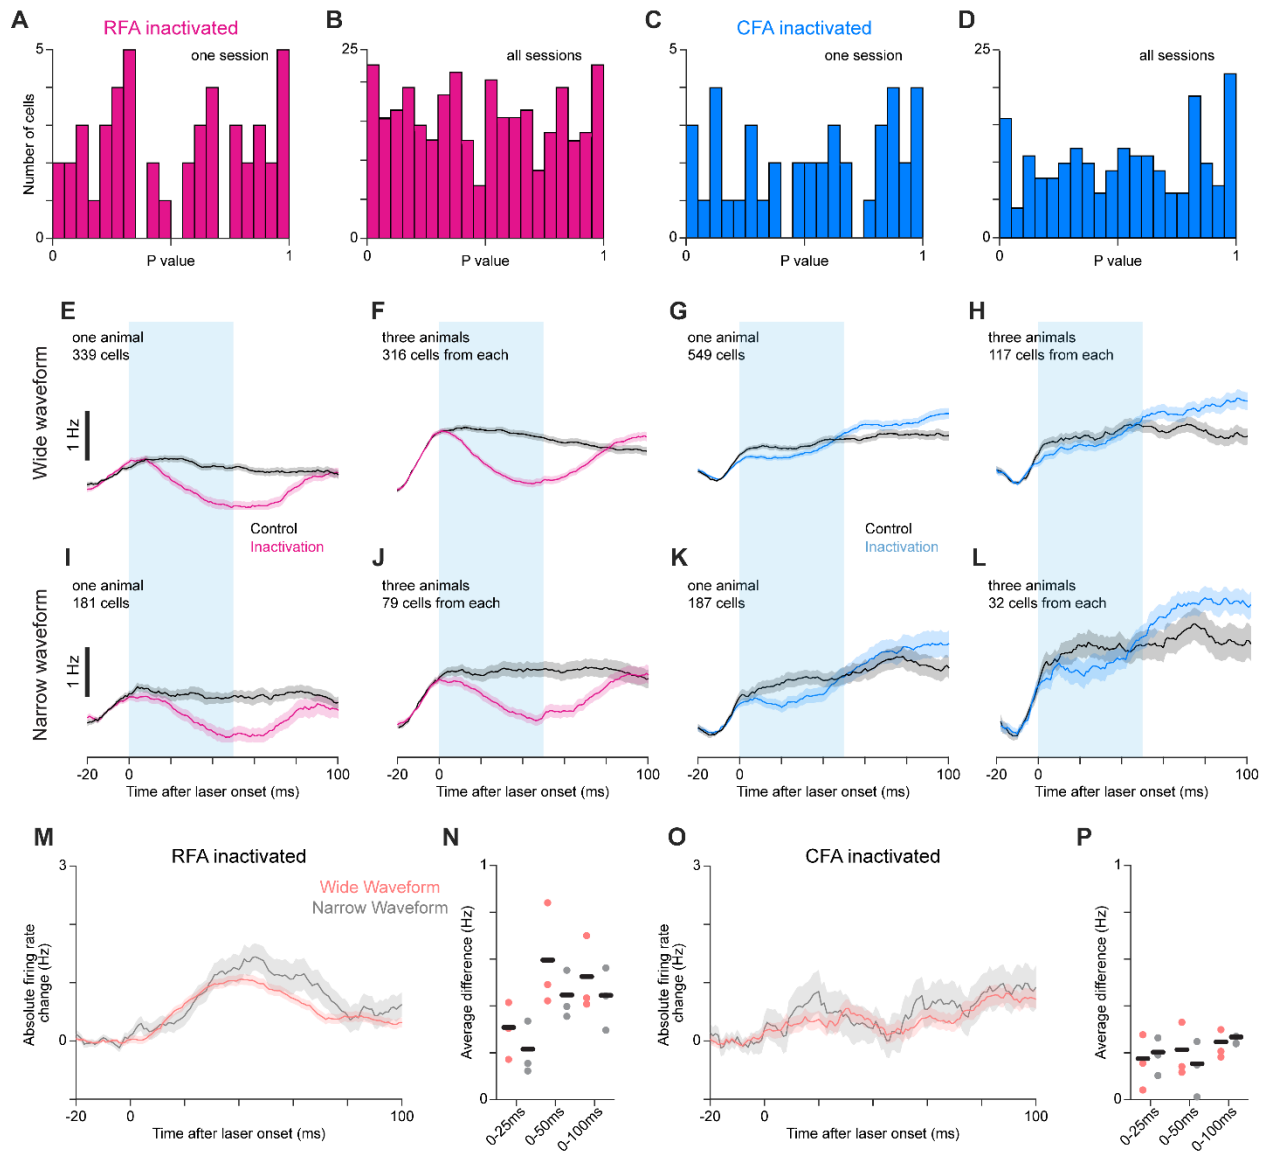


Fig. S3. Effects of CFA and RFA inactivation on firing rates in the other region. (A-D) Histograms of p-values from our modified version of SALT for narrow-waveform neurons recorded in one session (A,C) or all sessions (B,D) in either CFA (A,B) or RFA (C,D) while inactivating the other region. The uniformness of these distributions indicates an absence of appreciable violation of the null hypotheses that neurons are not directly activated by light. (E-L) Mean firing rate \pm SEM for wide-waveform (E-H) or narrow-waveform (I-L) neurons for one animal (E,G,I,K) or three animals (F,H,J,L) recorded in CFA (E,F,I,J) or RFA (G,H,K,L) while inactivating the other region. Averages combining cells from multiple animals (F,J,H,L) used the same number of cells from each animal. (F). The cyan rectangle indicates when the light was applied. (M-P) Mean absolute firing rate change \pm SEM between control and inactivation trials (M,O) and mean absolute firing rate difference between control and inactivation trials averaged from light/trial onset to 25, 50, or 100 ms afterwards (N,P) for wide- and narrow-waveform

neurons recorded in CFA (**M,N**) or RFA (**O,P**) during inactivation of the other region. Black bars show mean across animals.

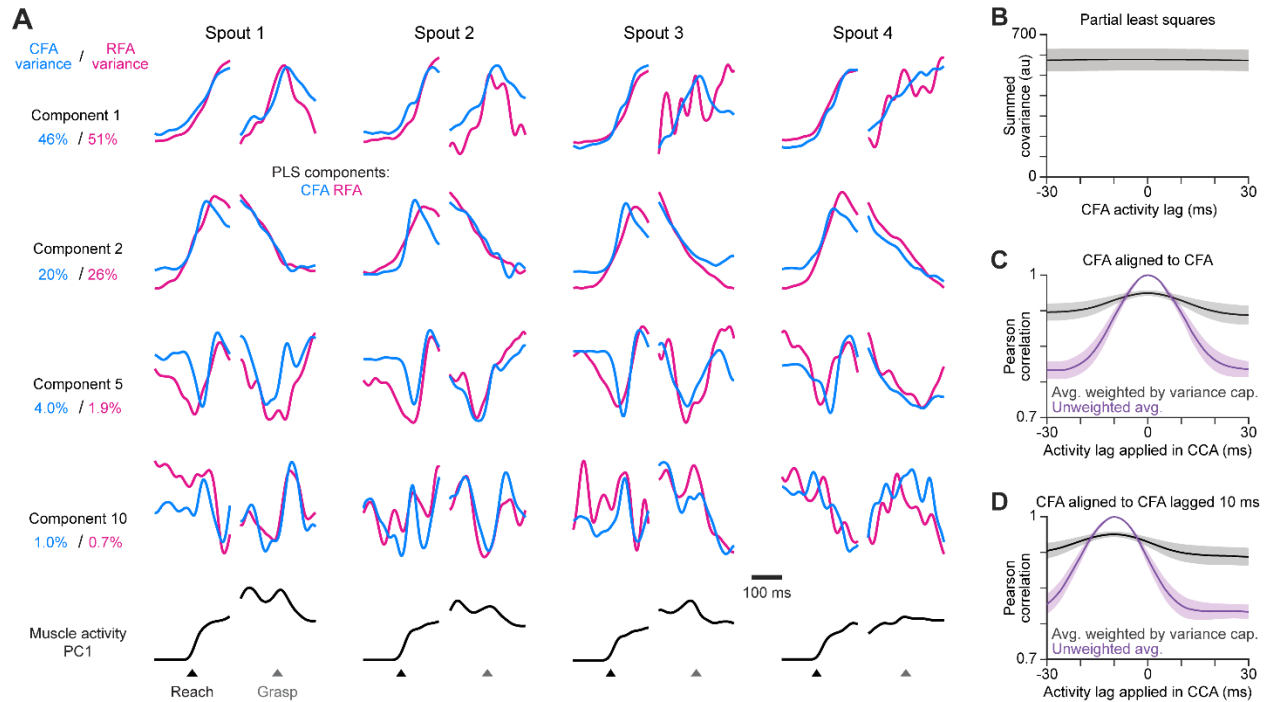


Fig. S4. PLS alignment of RFA and CFA activity. (A) For the same animal used in Fig. 3A, four PLS components aligned on reach onset or spout contact (grasp) for reaches to each spout. Bottom row shows the corresponding time series for the first principal component of muscle activity. (B) Mean \pm SEM 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.

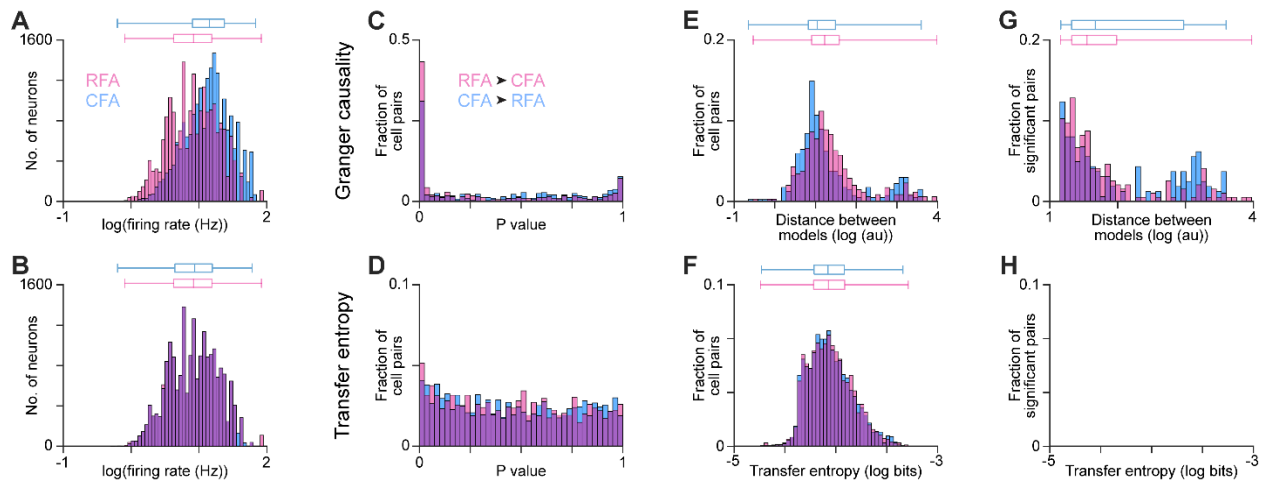


Fig. S5. Additional plots from firing pattern predictivity calculations. (A),(B) Firing rate distribution for CFA (blue) and RFA (magenta) neurons belonging to the 10,000 pairs with the highest firing rate product across all 21 recording sessions, before (A) and after (B) applying our algorithm for equalizing the firing rate distributions to avoid imparting a directional bias in firing pattern prediction calculations. In (A),(B), and (E-H), box plots on top show the minimum, 1st, 2nd, and 3rd quartiles, and maximum values. (C),(D) P-value distributions from calculations of Granger causality and transfer entropy using firing patterns of across-region neuron pairs, but only from the 300 ms prior to reach onset. (E-H) Distributions of metric values reflecting the improved prediction of target neuron firing using source neuron firing from the 300 ms prior to reach onset for across-region neuron pairs, including all pairs (E,F), and those for which the corresponding p-value fell below a false discovery-corrected threshold (G,H). H is empty because no pairs fell below the threshold.

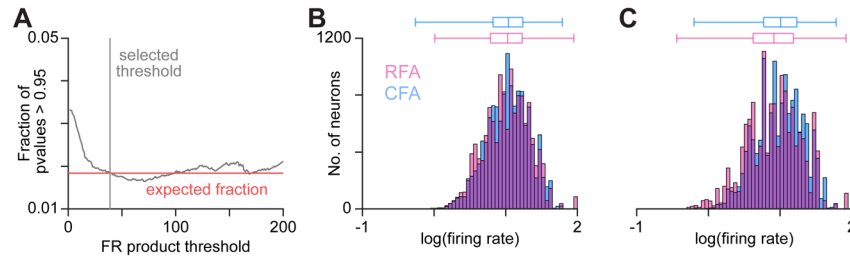


Fig. S6. Further neuron pair exclusion to avoid calculation anomalies. (A) Illustrates the calculation of a firing rate product threshold to eliminate an abnormal number of p values near 1. We determined that these anomalous values resulted from pairs with lower firing rate products, for which statistical 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 FR product threshold exclusion. In (B),(C), box plots on top show the minimum, 1st, 2nd, and 3rd quartiles, and maximum values.

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.

Acknowledgments: The authors would like to thank J. Glaser and L. Pinto for helpful conversations, and J. Glaser and M. Elbaz for comments on the manuscript.

Funding:

Japan Society for the Promotion of Science Research Fellow Grant-in-aid JP18J01678 (ASI)
Uehara Memorial Foundation Postdoctoral fellowship (ASI)
Searle Scholar Award (AM)
Sloan Research Fellowship (AM)
Simons Collaboration on the Global Brain Pilot Award 875841SPI (AM)
Whitehall Research Grant Award 2018-12-108 (AM)
Chicago Biomedical Consortium with support from the Searle Funds at The Chicago Community Trust C-099 (AM)
National Institutes of Health grant DP2 NS120847 (AM)

Author contributions:

Conceptualization: AI, AM
Methodology: AI, MA, SS, AF, AM
Investigation: AI, MA, SS, AF, HS, SH, DB, FX
Funding acquisition: AM
Supervision: AM
Writing – original draft: all authors
Writing – review & editing: all authors

Competing interests: Authors declare that they have no competing interests.

Data and materials availability: The data that support the findings of this study will be available at the lab GitHub page upon publication.

References:

1. J. Vezoli *et al.*, Cortical hierarchy, dual counterstream architecture and the importance of top-down generative networks. *Neuroimage* **225**, 117479 (2021).
2. D. J. Felleman, D. C. V. Essen, Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cereb Cortex* **1**, 1--47 (1991).
3. E. Bullmore, O. Sporns, Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci* **10**, 186-198 (2009).
4. A. W. Campbell, Histological Studies on the Localisation of Cerebral Function. *J Ment Sci* **50**, 651-662 (1904).
5. Y. Ueta, T. Otsuka, M. Morishima, M. Ushimaru, Y. Kawaguchi, Multiple layer 5 pyramidal cell subtypes relay cortical feedback from secondary to primary motor areas in rats. *Cereb Cortex* **24**, 2362-2376 (2014).
6. R. E. Passingham, Premotor cortex: sensory cues and movement. *Behav Brain Res* **18**, 175-185 (1985).
7. C. M. Gremel, R. M. Costa, Premotor cortex is critical for goal-directed actions. *Front Comput Neurosc* **7**, (2013).
8. P. E. Roland, B. Larsen, N. A. Lassen, E. Skinhoj, Supplementary motor area and other cortical areas in organization of voluntary movements in man. *J Neurophysiol* **43**, 118-136 (1980).
9. T. L. Veuthey, K. Derosier, S. Kondapavulur, K. Ganguly, Single-trial cross-area neural population dynamics during long-term skill learning. *Nat Commun* **11**, 4057 (2020).
10. E. Schmidlin, T. Brochier, M. A. Maier, P. A. Kirkwood, R. N. Lemon, 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 (2008).
11. K. C. Elliott, J. A. Borrell, S. Barbay, R. J. Nudo, Functional Suppression of Premotor Activity in a Transient Model of Motor Cortex Injury. *bioRxiv*, 2020.2006.2012.148403 (2020).
12. I. Fried, P. Haggard, B. J. He, A. Schurger, Volition and Action in the Human Brain: Processes, Pathologies, and Reasons. *J Neurosci* **37**, 10842-10847 (2017).
13. S. H. Scott, Population vectors and motor cortex: neural coding or epiphenomenon? *Nat Neurosci* **3**, 307-308 (2000).
14. R. P. Dum, P. L. Strick, The origin of corticospinal projections from the premotor areas in the frontal lobe. *J Neurosci* **11**, 667-689 (1991).
15. F. S. Hausmann, J. M. Barrett, M. E. Martin, H. Zhan, G. M. G. Shepherd, Axonal Barcode Analysis of Pyramidal Tract Projections from Mouse Forelimb M1 and M2. *J Neurosci* **42**, 7733-7743 (2022).
16. K. M. Fisher, B. Zaaimi, S. A. Edgley, S. N. Baker, Extensive Cortical Convergence to Primate Reticulospinal Pathways. *J Neurosci* **41**, 1005-1018 (2021).
17. R. P. Dum, P. L. Strick, Frontal lobe inputs to the digit representations of the motor areas on the lateral surface of the hemisphere. *J Neurosci* **25**, 1375-1386 (2005).
18. D. T. Bundy *et al.*, Stimulation-Evoked Effective Connectivity (SEEC): An in-vivo approach for defining mesoscale corticocortical connectivity. *J Neurosci Methods* **384**, 109767 (2023).
19. M. Weinrich, S. P. Wise, K. H. Mauritz, A neurophysiological study of the premotor cortex in the rhesus monkey. *Brain* **107 (Pt 2)**, 385-414 (1984).

20. Y. Prut, S. I. Perlmutter, E. E. Fetz, Distributed processing in the motor system: spinal cord perspective. *Prog Brain Res* **130**, 267-278 (2001).
21. J. F. Fulton, A note on the definition of the "motor" and "premotor" areas. *Brain* **58**, 311-316 (1935).
22. M. S. A. Graziano, *The intelligent movement machine : an ethological perspective on the primate motor system*. (Oxford University Press, Oxford ; New York, 2009), pp. ix, 224 pages.
23. K. Morandell, D. Huber, The role of forelimb motor cortex areas in goal directed action in mice. *Sci Rep* **7**, 15759 (2017).
24. P. C. Bucy, Electrical excitability and cyto-architecture of the premotor cortex in monkeys. *Arch Neuro Psychiatr* **30**, 1205-1225 (1933).
25. I. Stepniewska, O. A. Gharbawie, M. J. Burish, J. H. Kaas, 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 (2014).
26. R. Hira *et al.*, In vivo optogenetic tracing of functional corticocortical connections between motor forelimb areas. *Front Neural Circuit* **7**, (2013).
27. E. M. Rouiller, V. Moret, F. Liang, 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 (1993).
28. H. Mushiake, M. Inase, J. Tanji, Neuronal activity in the primate premotor, supplementary, and precentral motor cortex during visually guided and internally determined sequential movements. *J Neurophysiol* **66**, 705-718 (1991).
29. M. A. Umiltà, T. Brochier, R. L. Spinks, R. N. Lemon, Simultaneous recording of macaque premotor and primary motor cortex neuronal populations reveals different functional contributions to visuomotor grasp. *J Neurophysiol* **98**, 488-501 (2007).
30. T. C. Dixon, C. M. Merrick, J. D. Wallis, R. B. Ivry, J. M. Carmena, Hybrid dedicated and distributed coding in PMd/M1 provides separation and interaction of bilateral arm signals. *PLoS Comput Biol* **17**, e1009615 (2021).
31. W. Truccolo, L. R. Hochberg, J. P. Donoghue, Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat Neurosci* **13**, 105-111 (2010).
32. R. Kimura, A. Saiki, Y. Fujiwara-Tsukamoto, Y. Sakai, Y. Isomura, Large-scale analysis reveals populational contributions of cortical spike rate and synchrony to behavioural functions. *J Physiol* **595**, 385-413 (2017).
33. R. D'Aleo, A. G. Rouse, M. H. Schieber, S. V. Sarma, Cortico-cortical drive in a coupled premotor-primary motor cortex dynamical system. *Cell Rep* **41**, 111849 (2022).
34. M. Graziano, The organization of behavioral repertoire in motor cortex. *Annu Rev Neurosci* **29**, 105-134 (2006).
35. M. Wiesendanger *et al.*, Input and Output Organization of the Supplementary Motor Area. *Ciba F Symp* **132**, 40-62 (1987).
36. J. Tanji, K. Kurata, Comparison of movement-related activity in two cortical motor areas of primates. *J Neurophysiol* **48**, 633-653 (1982).
37. H. Makino *et al.*, Transformation of Cortex-wide Emergent Properties during Motor Learning. *Neuron* **94**, 880-890 e888 (2017).
38. S. Ito *et al.*, Extending transfer entropy improves identification of effective connectivity in a spiking cortical network model. *PLoS One* **6**, e27431 (2011).

39. A. Casile, R. T. Faghieh, E. N. Brown, 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 (2021).
40. O. Gozel, B. Doiron, Between-area communication through the lens of within-area neuronal dynamics. *bioRxiv*, 2022.2004.2011.487906 (2023).
41. L. Bachschmid-Romano, N. G. Hatsopoulos, N. Brunel, J. Diedrichsen, Interplay between external inputs and recurrent dynamics during movement preparation and execution in a network model of motor cortex. *Elife* **12**, (2023).
42. G. L. Galinanes, C. Bonardi, D. Huber, Directional Reaching for Water as a Cortex-Dependent Behavioral Framework for Mice. *Cell Rep* **22**, 2767-2783 (2018).
43. S. Zhao *et al.*, Cell type-specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function. *Nat Methods* **8**, 745-752 (2011).
44. E. Gokcen *et al.*, Disentangling the flow of signals between populations of neurons. *Nat Comput Sci* **2**, 512-+ (2022).
45. A. Miri *et al.*, Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex. *Neuron* **95**, 683-696 e611 (2017).
46. Z. V. Guo *et al.*, Flow of cortical activity underlying a tactile decision in mice. *Neuron* **81**, 179-194 (2014).
47. H. Hotelling, Relations between two sets of variates. *Biometrika* **28**, 321-377 (1936).
48. D. Sussillo, M. M. Churchland, M. T. Kaufman, K. V. Shenoy, A neural network that finds a naturalistic solution for the production of muscle activity. *Nat Neurosci* **18**, 1025-1033 (2015).
49. E. Le Floch *et al.*, 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 (2012).
50. G. Sugihara *et al.*, Detecting causality in complex ecosystems. *Science* **338**, 496-500 (2012).
51. J. D. Storey, A direct approach to false discovery rates. *Journal of the Royal Statistical Society Series B-Statistical Methodology* **64**, 479-498 (2002).
52. B. Hyland, Neural activity related to reaching and grasping in rostral and caudal regions of rat motor cortex. *Behav Brain Res* **94**, 255-269 (1998).
53. A. Riehle, J. Requin, 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 (1993).
54. G. F. Elsayed, A. H. Lara, M. T. Kaufman, M. M. Churchland, J. P. Cunningham, Reorganization between preparatory and movement population responses in motor cortex. *Nat Commun* **7**, 13239 (2016).
55. R. Chaudhuri, K. Knoblauch, M. A. Gariel, H. Kennedy, X. J. Wang, A Large-Scale Circuit Mechanism for Hierarchical Dynamical Processing in the Primate Cortex. *Neuron* **88**, 419-431 (2015).
56. K. V. Shenoy, M. Sahani, M. M. Churchland, Cortical control of arm movements: a dynamical systems perspective. *Annu Rev Neurosci* **36**, 337-359 (2013).
57. J. S. Seely *et al.*, Tensor Analysis Reveals Distinct Population Structure that Parallels the Different Computational Roles of Areas M1 and V1. *PLoS Comput Biol* **12**, e1005164 (2016).
58. B. A. Sauerbrei *et al.*, Cortical pattern generation during dexterous movement is input-driven. *Nature* **577**, 386-391 (2020).

59. Z. V. Guo *et al.*, Procedures for behavioral experiments in head-fixed mice. *PLoS One* **9**, e88678 (2014).
60. T. Akay, H. J. Acharya, K. Fouad, K. G. Pearson, Behavioral and electromyographic characterization of mice lacking EphA4 receptors. *Journal of Neurophysiology* **96**, 642-651 (2006).
61. C. L. Warriner, S. Fageiry, S. Saxena, R. M. Costa, A. Miri, Motor cortical influence relies on task-specific activity covariation. *Cell Rep* **40**, 111427 (2022).
62. M. Pachitariu, N. Steinmetz, S. Kadir, M. Carandini, H. Kenneth D., Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv*, (2016).
63. D. Kvitsiani *et al.*, Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. *Nature* **498**, 363-366 (2013).
64. A. Krakovská, J. Jakubík, H. Budáčová, M. Holeciová, Causality studied in reconstructed state space. Examples of uni-directionally connected chaotic systems. *arXiv preprint arXiv:1511.00505*, (2015).
65. H. Ye, E. R. Deyle, L. J. Gilarranz, G. Sugihara, Distinguishing time-delayed causal interactions using convergent cross mapping. *Sci Rep* **5**, 14750 (2015).
66. A. T. Clark *et al.*, Spatial convergent cross mapping to detect causal relationships from short time series. *Ecology* **96**, 1174-1181 (2015).
67. D. B. Ehrlich, J. T. Stone, D. Brandfonbrener, A. Atanasov, J. D. Murray, PsychRNN: An Accessible and Flexible Python Package for Training Recurrent Neural Network Models on Cognitive Tasks. *Eneuro* **8**, (2021).