The cortical dynamics orchestrating skilled prehension

Britton Sauerbrei*, Jian-Zhong Guo*, Jihong Zheng, Wendy Guo, Mayank Kabra, Nakul Verma, Kristin Branson, Adam Hantman****

Janelia Research Campus, Howard Hughes Medical Institute

***These authors contributed equally to this work.**
****Co-senior authors.**

Abstract
Skillful control of movement is central to our ability to sense and manipulate the world. Dexterous acts depend on cerebral cortex[1-10], and the activity of cortical neurons is correlated with movement[11-15]. By isolating the neural dynamics that command skilled movements from those that reflect other processes (such as planning and deciding to move), we were able to characterize and manipulate the motor commands underlying prehension. We showed that in mice trained to perform a reach / grab / supination / bring-to-mouth sequence (volitional prehension), multiple forms of optogenetic stimuli in sensorimotor cortex resulted in an involuntary, complete movement (opto-prehension). This result suggested that the trained brain could robustly transform a variety of aberrant stimuli into the dynamics sufficient for prehension. We measured the electrical activity of cortical populations and detailed limb kinematics during volitional prehension and opto-prehension. During volitional prehension, neurons fired before and during specific stages of the movement, and the population collectively tiled the entire behavioral sequence. During opto-prehension, most neurons recapitulated their volitional prehension activity patterns, but a physiologically distinct subset did not. On trials where the liminal optogenetic stimulus failed to produce these dynamics, movement did not occur, providing further evidence that a specific pattern of neural activity was causally coupled to prehension. Having identified these dynamics, we next tested their robustness to brief, closed-loop perturbation. Regardless of where along the reach we optogenetically halted cortical activity and the movement, relief of suppression resulted in cortical dynamics that immediately recapitulated all steps of the prehension program, and the animal completed the behavior. By combining electrophysiology and optogenetic perturbations, we have identified and characterized the cortical motor program driving a learned, dexterous movement sequence.

Main
Reaching, grasping, and object manipulation play a central role in the lives of mammals with prehensile forelimbs. The musculoskeletal complexity of the limb poses a challenging control problem for the central nervous system, which must orchestrate precisely-timed patterns of activity in many muscles to perform a wide diversity of tasks. Cortex is thought to be critical for producing these patterns, and single-unit recording studies have demonstrated that the activity of cortical neurons is correlated with muscle tension and limb kinematics. In order to dissect the role of sensorimotor cortex in dexterous movement, we developed a prehension task for head-fixed mice, in which animals learned over several weeks to reach for and grab a food pellet at a memorized position and deliver it to the mouth following an auditory cue (fig. 1a). Using high-speed video and computer vision techniques, we captured the animals' behavior and extracted the timing of 'waypoints' indexing the stages of the movement (lift hand from the perch,

open hand, grab food pellet, supinate hand, and bring hand to mouth), as well as three-dimensional hand position (fig. 1b, supplementary video 1). To study neural dynamics during the behavior, we used silicon probes to record spiking activity (mean 19 well-isolated units per recording) from a total of 584 neurons in sensorimotor cortex (fig. 1b, 1c, 1d), which showed strong fluctuations in firing rate around the movement. A majority of cells (453/584) were modulated before and during prehension, with a net rate increase for 243/584, and a decrease for 210/584 (fig. 1d; rank sum test on pre- and peri-lift spike counts with Benjamini-Hochberg correction, $q < .05$). While the responses of individual cells were highly consistent across trials (fig. 1c), we observed a wide diversity of patterns across neurons, including increases, decreases, and multi-phasic responses (fig. 1d). Single neurons can fire at selective points in the spatial trajectory of the hand (fig. 1e, supplementary video 2, extended data fig. 1c, rank sum test on the variance of the hand position at spike times versus the full trajectory, $p = 8e-26$), and the population tiled the entire extent of the behavior (extended data fig. 1d). Because movement is most likely generated by the coordinated activity of large ensembles, we examined the dynamics of the population activity. The population activity during reaching was highly precise: neural variability decreased sharply at movement onset (extended data fig. 1a, b)[16]. We estimated the trial-by-trial trajectory of the neural population using Gaussian Process Factor Analysis (GPFA)[17] and found that the neural trajectory was closely linked to movement (fig. 1f). Thus, mouse sensorimotor cortex generates a complex, time-varying signal tightly coupled to the reach-to-grasp-to-mouth movement sequence.

We wanted to confirm that this signal is sent out of the cortex to destinations that might lead to the production of movement. Cortex has two major output systems, the intratelencephalic (IT) system, which projects to the contralateral cortex and striatum, and the pyramidal tract (PT) system, which targets the pontine gray and the spinal cord. An adeno-associated virus (AAV) optimized for retrograde transmission (rAAV2-retro)[18] was used to induce expression of the genetically-encoded calcium indicator jRCaMP1b[19] so that the activity of each of these cortical output systems could be imaged. We found neurons of both projection classes that were robustly modulated during reaching (extended data fig. 1e, 1f), but the degree of modulation was higher in PT than IT neurons ($p = 8e-9$, rank sum test on mean post-pre $\Delta F/F$; extended data fig. 1g). Thus, cortical signals during reaching are sent to targets of both the IT and PT systems, but are more strongly represented in the pyramidal tract.

Although our analysis of spiking patterns establishes a clear correspondence between cortical dynamics and prehension, it is unclear whether this relationship is causal. We used optogenetic circuit perturbations in conjunction with electrophysiology to determine the necessity and sufficiency of neural dynamics for driving the movement. By optically stimulating inhibitory neurons expressing channelrhodopsin (ChR2), we were able to silence the spiking of putative excitatory neurons (fig. 2a) in VGAT-ChR2-EYFP mice. This inactivation produced two striking behavioral effects. First, inactivation blocked the initiation of reaching (fig. 2b, extended data fig. 2a, 2b). Second, at the offset of the laser stimulus, mice rapidly performed a reach (opto-prehension). These opto-prehensions were initiated with shorter reaction times than cued reaches (fig. 2b, median reaction time 94 ms vs 212 ms, $p = 9e-14$, rank sum test), suggesting that the inactivated region of cortex is downstream of areas involved in early stages of the volitional behavior, such as processing the cue and deciding to reach. Furthermore, it was possible to induce opto-prehensions even in the absence of a preceding cue (fig. 2b, magenta traces), and these also had shorter latencies than voluntary reaches (median reaction time 110 ms, p

= 3e-14, rank sum test). Thus, sensorimotor cortex appears to be a necessary and sufficient locus for triggering the movement sequence.

The ability to optogenetically initiate a skilled movement enables a direct comparison of the neural dynamics underlying voluntary and involuntary (opto-prehension) versions of the behavior. Activity patterns shared between the two conditions likely reflect the core motor program that animals learned to execute during training, while signals only present during voluntary reaches may represent volitional signals. Examination of the peri-lift firing rates (fig. 2c, left) showed that activity of many neurons before and during the execution of the movement was strikingly similar between voluntary and opto-prehensions. Similarly, we compared the position of the hand at spike times in both conditions, and found many were closely matched (fig 2d, supplementary videos 2 and 3, extended data fig. 2f). Taken together, these results show that the firing profiles of a subset of neurons were recapitulated between voluntary and opto-prehension. Other neurons had firing patterns that differed between the two conditions (fig 2c, extended data fig. 2g). Neurons that exhibited higher peri-lift rates during voluntary prehension were distinguished by higher baseline firing rates (q = 1.4e-5, rank sum test, fig. 2c, right). While firing rates have previously been used to classify cortical neurons, future work will be required to further characterize these cells.

We adopted two further approaches to characterize the similarities between the neural dynamics of voluntary and opto-prehensions. First, we used a neural decoding method to predict the times of movement waypoints from single neurons and from population spiking activity. After training classifiers for each waypoint (lift, hand open, etc.) on a subset of voluntary trials, we used these classifiers to predict waypoint times on the remaining voluntary trials and on opto-prehension trials. Population-based classifiers predicted each waypoint around its onset time on both voluntary and opto-prehensions (fig. 1b, 2e, extended data fig. 1i). Single-neuron-based classifiers also exhibited high prediction accuracy (extended data fig. 1h), and this accuracy was matched between voluntary and opto-prehensions (Spearman's rho = 0.69, p < 10e-20, extended data fig. 2c). Second, using GPFA, we computed the average population state at grab time on half the voluntary trials and then computed the distance of the neural population to this target state for the remaining voluntary and opto-prehension trials. When prehension occurred for either of these conditions, we found that this neural state distance decreased as the hand approached the pellet (fig. 2f, extended data fig. 2e, supplementary videos 4, 5). However, when the cortical network failed to generate the appropriate dynamics, the mouse did not reach to the target (fig. 2g, extended data fig. 2e). Post-laser firing rates were higher when opto-prehension was evoked than when no post-laser grab occurred, and these rates began to diverge around 20 ms after laser offset (extended data fig. 2d). Because this difference in firing rates emerged before the initiation of opto-prehension, it was likely driven by a difference in feedforward motor commands, rather than sensory feedback from the limb. These results show that a specific dynamical pattern starts and drives the complete movement program, regardless of whether this pattern is initiated voluntarily following a cue or evoked by optogenetic stimulation.

Cessation of VGAT-ChR2-EYFP stimulation leads to the dynamics sufficient for evoking prehension; are these dynamics generated only by this particular perturbation, or can other stimuli also evoke the behavior? To address this question, we drove ChR2 expression selectively either in putative IT neurons or putative PT neurons using cell-type selective mouse lines (Tlx3(PL56)-Cre and Sim1(KJ18)-Cre, respectively)[20].

Activation of putative IT neurons blocked the movement and induced rapid opto-prehensions at laser offset (fig. 3a, extended data fig. 3a), similar to activation of VGAT+ neurons. Activation of putative PT neurons, by contrast, only partially blocked the movement, and less robust opto-prehension was induced at laser offset (fig. 3d, extended data fig. 3b). Following both perturbations, if opto-prehension was evoked, cortical dynamics moved towards the grab state (fig. 3b, 3e), and it was possible to decode lift times from single-neuron firing patterns after the offset of the laser (fig. 3c, f). Thus, direct activation of excitatory projection cell types can also disrupt and then extract the program of a skilled movement.

The combination of optogenetics and electrophysiology revealed the necessary and sufficient cortical dynamics that initiate the reach-to-grab sequence, but does cortex provide only a “go” signal, or does it orchestrate the full sequence? If cortical dynamics are required throughout the entire movement, then briefly perturbing these dynamics during an ongoing reach should disrupt or halt the movement. We tested this hypothesis by using real-time detection of movement onset to trigger a brief activation of the VGAT-ChR2-EYFP cortical neurons. This perturbation interrupted the reach, but following laser offset, the animals were able to quickly recover and successfully execute the full movement sequence (fig. 4a). Does cortex re-generate the normal dynamics driving movement after the perturbation, and if so, where do the dynamics resume? The spatial and neural distance to the target decreased rapidly after the laser offset (fig. 4b, 4c), and the spatial distribution of spikes on interruption trials largely recapitulated the voluntary pattern (fig. 4d, extended data fig. 4a). In order to address the question of where the dynamics restart, we applied the neural decoder to population activity after termination of the laser. Following the end of the perturbation, cortex reproduced the dynamics driving lift, and proceeded to generate the entire neural sequence (fig. 4e, extended data fig. 4b). Taken together, these results demonstrate that sensorimotor cortex does not merely initiate the behavior, but instead provides a continuous signal that controls the full movement sequence from lift to at mouth.

Discussion

The mammalian forelimb is a sophisticated effector having many degrees of freedom, highly complex mechanics, and the ability to sense and manipulate the environment. Our results demonstrate that cortical dynamics contain the temporally rich signals required to control skilled forelimb movements. While recent studies have argued that rodent sensorimotor cortex is involved in the online control of only dexterous movements of the hand[21, 22], we find that cortex is necessary and sufficient for coordinating the spatiotemporal pattern of muscle activity across the entire limb and movement sequence, not merely at the distal joints during object manipulation.

By simultaneously recording neural dynamics and leveraging the optogenetic toolbox available in the mouse, we show that multiple, aberrant optogenetic stimuli can be funneled into the structured pattern of cortical activity driving prehension. This raises the possibility that during normal prehension, upstream brain regions might be able to trigger the behavior by generating an unstructured “go” signal, which the cortex then transforms into the time-varying control signal. This highly stereotyped neural trajectory must result from changes in cortex and associated connections that constitute the engram of a complex motor skill acquired over weeks of learning[23-27]. An animal’s full behavioral repertoire likely requires a collection of engrams for different movements; how these compete during action selection and execution will require future investigation. The combination of optogenetic control and recently-developed techniques for marking,

monitoring, and manipulating neural ensembles at cellular resolution should facilitate study of how motor skills are stored and retrieved[28-30].

Methods

Behavioral task and video analysis

Mice were fitted with head posts, food restricted, and trained to reach for food pellets, as described previously[8]. Behavior was recorded at 500 Hz (IORodeo, BIAS: Basic Image Acquisition Software) using two high-speed cameras (PointGrey, Flea3), which were calibrated to allow 3D triangulation of hand position (Caltech Camera Calibration Toolbox for Matlab). Two types of information were extracted from video: ethograms labeling the frames in which lift, hand open, grab, supination, hand at mouth, and chew occurred, obtained using the Janelia Automatic Animal Behavior Annotator, and the position of the hand in space, obtained using the Animal Part Tracker. All procedures were approved by the Institutional Animal Care and Use Committee at Janelia Research Campus (protocol 13-99).

Automatic behavior classification

Using an adaptation of the Janelia Automatic Animal Behavior Annotator (JAABA)[31], we trained automatic behavior classifiers which input information from the video frames and output predictions of the behavior category -- *lift, hand-open, grab, supination, at-mouth,* and *chew.* We adapted JAABA to use Histogram of Oriented Gradient (HOG)[32] and Histogram of Optical Flow (HOF)[33] features derived directly from the video frames, instead of features derived from animal trajectories. The automatic behavior predictions were post-processed as described previously [8] to find the first lift-hand-open-grab and supination-at-mouth-chew sequences.

Electrophysiological recordings

Neural recordings were performed using the Whisper acquisition system (Janelia Applied Physics and Instrumentation Group) and 64-channel silicon probes (NeuroNexus A4x16-Poly2-5mm-23s-200-177-A64 or Janelia “Hires” 4x16 probes). These probes consisted of four shanks with 16 contacts at the tip of each, over a depth of 345µm (NeuroNexus) or 320µm (Janelia probes). On the day before the experiment, a small craniotomy was made over sensorimotor cortex contralateral to the limb, and a stainless steel reference wire was implanted in visual cortex. During the recording session, the probe tips were positioned at bregma +0.5mm, 1.7mm lateral, and slowly lowered to a depth of ~900µm from the cortical surface, and a silicone elastomer (Kwik-Sil, World Precision Instruments) was applied to seal the craniotomy. At the end of the session, the probe was removed, and the craniotomy was re-sealed with silicone to allow a subsequent session on the following day. Signals were amplified with a gain of 200 and digitized to 16 bits at 25-50 kHz, and spike sorting was performed with JRClust[34].

Optogenetic manipulations

Cell-type specific expression of ChR2 was achieved by either using VGAT-ChR2-EYFP mice expressing ChR2 in inhibitory neurons (Slc32a1-COP4*H134R/EYFP, The Jackson Laboratory), or by crossing a Cre driver line to a Cre-dependent ChR2 reporter mouse, Ai32 (Rosa-CAG-LSL-ChR2(H134R)-EYFP-WPRE, The Jackson Laboratory). Experiments were performed in VGAT-ChR2-EYFP (n = 9), Tg(Tlx3-Cre)PL56Gsat X Ai32 (n = 3), Tg(Sim1-Cre)KJ18Gsat X Ai32 (n = 4), or Tg(Rbp4-Cre)KL100Gsat X Ai32

(n = 2) mice. An optical fiber (200 µm or 400 µm, NA 0.39, Thorlabs) was coupled to a 473 nm laser (LuxX 473-80, Omikron Laserage) and positioned 2-4 mm over sensorimotor cortex in the head fixation apparatus, as described previously. A blue light emitting diode array was directed at the animal's eyes throughout the session in order to mask the laser stimulus. Three trial types were used: control trials, in which the cue was presented with no laser stimulation, laser + cue trials, in which both were presented, and laser-only trials, in which the laser was turned on without a cue or food administration. A two-second laser stimulus (40 Hz sine wave) was initiated synchronously with the cue for VGAT-ChR2-EYFP mice, or 200 ms before cue onset for Tlx3-Cre X Ai32 and Sim1-Cre X Ai32 mice. Laser power was calibrated to the minimum level necessary to block reaching in probe experiments in the final days of training; this ranged from 10-50 mW at the fiber tip for VGAT animals, and 0.5-6 mW for Tlx3 and Sim1 animals. In interruption experiments, a region of the video frame between the average lift and hand open locations was identified using BIAS software, and a contrast change in this region was used to open the laser shutter for 50-100 ms.

Two-photon imaging

Expression of the red-shifted calcium indicator jRCaMP1b was targeted selectively to pyramidal tract (PT) or intratelencephalic (IT) neurons by injecting rAAV2-retro-SYN-Cre into the pontine gray and AAV-2/1-SYN-FLEX-RCaMP1b into sensorimotor cortex (n = 3 mice), or by injecting rAAV2-retro-CAG-RCaMP1b into the contralateral striatum (n = 2) or ipsilateral pontine nuclei (n = 1). Pontine injections were targeted to lambda -3.9 mm, lateral 0.4 mm, at depths of 5.2, 5.4, 5.6, and 5.8 mm in each mouse, as described previously[35]. Striatum injections were targeted to bregma +0.4 mm, lateral 2.5, at depths of 2.2, 2.4, and 2.6 mm. Five injections were targeted to sensorimotor cortex at depths of 0.6 mm. Injection volumes were 50 µl at each site. Animals were fitted with head posts and cranial windows over sensorimotor cortex (2.2 mm window, Potomac Photonics, at bregma +0.6mm, lateral 1.6), food deprived, and trained to reach. Imaging was performed using a Fidelity-2 1070 nm laser (Coherent) on two microscopes, using either a 16X Nikon CFI LWD Plan Fluorite Objective (0.80 NA, 3.0 mm WD) or a custom objective (Jenoptik Inc), NA 1.0. Images were motion-corrected (cross-correlation method, Thunder Library), and regions of interest were defined manually (ImageJ).

Statistical analysis

Peri-lift firing rates (Fig 1d, 2c). For each neuron, lift-centered spike trains were smoothed with a Gaussian kernel (σ = 50 ms) and averaged across trials. Lift modulation was assessed using a rank sum test comparing the raw spike counts in a 500 ms window centered at lift +200 ms with counts in a 500 ms window centered at lift -750 ms. Multiple comparisons were corrected using the false discovery rate framework ($q < 0.05$). For comparison of firing rates on voluntary and opto-prehensions (fig. 2c), a rank sum test was performed for each neuron and peri-lift time bin.

Spike-triggered hand position (Fig. 1e, 2d, 4d, E1d, E2f, E4a, supplementary videos 2, 3). Hand trajectories were extracted from lift -100 ms to supination +100 ms, and trajectories were aligned across datasets by applying a procrustes transformation to the within-dataset mean trajectories. The hand position was extracted at the time of each spike. A kernel density estimator (σ = 1.5mm) was used to obtain the spatial density of spikes, as well as the spatial density of the hand over the entire trajectory. The full-trajectory density was subtracted from the spike-triggered position density, and iso-density-difference contours were extracted at the following ratios of the peak density difference: +/- 0.02, 0.4, 0.6, 0.8, 0.95, 1. Three-dimensional surface plots of these

positive and negative contours (Fig. 1e, 2d, 4d, E1d) were generated for voluntary and interruption trials, and trials were compared by overlaying two-dimensional projections of the positive contours (Fig. 2d). In order to quantify the similarity between firing patterns on voluntary and opto-prehensions, we analyzed the pairwise distances between spike-triggered hand positions. If the spatial firing pattern on opto-prehensions recapitulates the pattern on normal reaches, then a randomly-chosen spike from an opto-prehension should lie closer to a randomly-chosen spike from the same neuron on a normal reach than to a randomly-chosen spike from a randomly-chosen neuron. Thus, for each neuron, we computed two distributions by bootstrapping with 50000 resamples. The first was the distribution of distances from the hand position at the time of a randomly-selected spike on an opto-prehension to the position at the time of a spike from that neuron on a normal reach. The second was the distribution of distances from a randomly-selected spike on a opto-prehension to a spike from a pool consisting of 50 spikes from each neuron in the full dataset. The medians of these distributions for each neuron were then compared (Fig. E2f, E4a).

Gaussian Process Factor Analysis (GPFA) and neural distance (Fig. 1f, 2f, 2g, 3b, 3e, 4b, 4c, E2e, supplementary videos 4, 5). Neural population activity was reduced to a five-dimensional latent variable space using GPFA[17] (bin size 20 ms). The target spatial and neural states were defined using the three-dimensional position of the hand, and the five-dimensional latent variable representation of sensorimotor cortex obtained using GPFA, respectively. In both cases, the states were sampled at grab times, and the target state was defined to be the central location of the grab-triggered states, computed using convex hull peeling[36]. Only the even control trials were used to calculate the target states. The Euclidean distance from the target was then computed for each trial and time point, and the resulting distance curves were centered either on grab time (Fig. 2f, 4b, E2e [upper]) or laser offset (Fig. 2g, 4c, E2e [lower]).

Neural decoder (Fig. 1b, 2e, 3c, 3f, 4e, E1h, E2c, E4b)
We trained a classifier to predict the mouse's behavior (*lift, hand-open, grab, at-mouth, supination, chew, or none-of-the-above*) from its neural activity. The behavior category label used for training this neural decoder was itself the output of an automatic classifier of behavior category from video of the animal behaving, trained using JAABA. We labeled time points within 10 ms of the video-based lift, hand-open, and grab of the first lift-hand-open-grab sequence as lifts, hand-opens, and grabs, respectively, and time points within 10 ms of the video-based supination, at-mouth, and chew of the first supination-at-mouth-chew sequence were labeled as supination, at-mouth, and chew, respectively. Time points more than 30 ms from these video-based behavior labels were labeled as none-of-the-above. We trained both population-level neural decoders and single-neuron decoders. For population-level neural decoders, neural activity from all neurons between 0 and 100 ms before a given time point were used to predict behavior category at that time point. Neural activity in the 100 ms time window was represented as a 110n-dimensional vector, where n is the number of neurons, encoding spike rates in sub-windows corresponding to Gaussians of varying bandwidths (2.5, 5, 10, 12.5, 15, 20, 25, 35, 50, 75, 100 ms) and offsets (-18, -16, …, -2, 0 ms) truncated to within the 100 ms window. We trained a multiclass, one-vs-all linear Support Vector Machine (SVM) to predict behavior category on voluntary reaches, with equal class priors. While the classifier was only trained on voluntary reaches, we applied it to laser-evoked reaches as well, to investigate the relationship between the neural activity sequence during voluntary and laser-evoked reaches. Single-neuron decoders were trained as above, but using 110-dimensional input corresponding to a single neuron.

Fano factor (Fig. E1a, E1b). For each neuron, lift-centered spike counts were computed in 50 ms bins. In each bin, the Fano factor was defined to be the slope of the regression line through the variance-vs-mean plot for all neurons[16].

Supplementary video captions

Video 1: Head-fixed prehension behavior and hand tracking. The video shows raw images from two cameras capturing the movement sequence, along with the triangulated three-dimensional location of the paw.

Video 2: Spatial position of the hand at spike times during voluntary prehension. The video shows the three-dimensional position of the hand during each voluntary prehension in an experimental session (green), along with the hand position at spike times for example neurons 1 (left) and 4 (right). Spikes occurring in the current trial are magenta, while spikes in earlier trials are yellow. The relative density of spikes (iso-density-difference contours between the hand position at spike times and the hand position over the full trajectory) is overlaid; red represents regions with many spikes, and blue represents regions with few spikes. Scale bars are 5 mm.

Video 3: Spatial position of the hand at spike times during opto-prehension. As in video 2, but for opto-prehensions in the same dataset. Hand position at spike times is indicated in cyan.

Video 4: Neural population activity and hand position during voluntary and opto-prehension, centered on cue. Neural state, computed using GPFA, is shown in the left panel, and hand position is shown in the right panel. Each point corresponds to a single trial, with yellow indicating voluntary prehension and blue indicating opto-prehension. Lift and grab times are green and magenta, respectively.

Video 5: Neural population activity and hand position during voluntary and opto-prehension, centered on grab. As in video 4, but trajectories are aligned to the grab time.

Contributions

B.S., J.G., and A.H. designed the experiments. B.S. and J.G. performed electrophysiological recordings. J.G. performed behavioral experiments. B.S. analyzed electrophysiology and behavior data. J.G. and W.G. analyzed behavior data. J.Z. performed imaging experiments and analyzed imaging data. N.V. and K.B. developed neural decoding analyses. M.K. and K.B. developed computer vision algorithms and software. B.S., A.H., and K.B. wrote the paper with input from all authors. A.H. and K.B. supervised the project.

Acknowledgements

We thank Byron Yu and the Yu Lab, Steve Edgley, Joshua Dudman, Albert Lee, Brett Mensh, Misha Ahrens, Matteo Mischiati, and Jeremy Cohen for discussions and comments on the manuscript; Allen Lee for tracking software; Tim Harris, Brian Barbarits, Bill Karsh, Steve Sawtelle, Peter Polidoro, and Dan Flickinger for instrumentation development and support; James Jun for spike sorting software; and Julia Kuhl for the mouse drawing.

1. Lawrence, D.G. and H.G. Kuypers, *The functional organization of the motor system in the monkey. I. The effects of bilateral pyramidal lesions.* Brain, 1968. **91**(1): p. 1-14.
2. Whishaw, I.Q., et al., *Proximal and distal impairments in rat forelimb use in reaching follow unilateral pyramidal tract lesions.* Behav Brain Res, 1993. **56**(1): p. 59-76.
3. Castro, A.J., *The effects of cortical ablations on digital usage in the rat.* Brain Res, 1972. **37**(2): p. 173-85.
4. Passingham, R.E., V.H. Perry, and F. Wilkinson, *The long-term effects of removal of sensorimotor cortex in infant and adult rhesus monkeys.* Brain, 1983. **106 (Pt 3)**: p. 675-705.
5. Martin, J.H. and C. Ghez, *Impairments in reaching during reversible inactivation of the distal forelimb representation of the motor cortex in the cat.* Neurosci Lett, 1991. **133**(1): p. 61-4.
6. Penfield, W., *Mechanisms of voluntary movement.* Brain, 1954. **77**(1): p. 1-17.
7. Graziano, M.S., C.S. Taylor, and T. Moore, *Complex movements evoked by microstimulation of precentral cortex.* Neuron, 2002. **34**(5): p. 841-51.
8. Guo, J.Z., et al., *Cortex commands the performance of skilled movement.* Elife, 2015. **4**: p. e10774.
9. Harrison, T.C., O.G. Ayling, and T.H. Murphy, *Distinct cortical circuit mechanisms for complex forelimb movement and motor map topography.* Neuron, 2012. **74**(2): p. 397-409.
10. Porter, R.L., R, *Corticospinal Function and Voluntary Movement*. 1993: Oxford University Press.
11. Evarts, E.V., *Pyramidal tract activity associated with a conditioned hand movement in the monkey.* J Neurophysiol, 1966. **29**(6): p. 1011-27.
12. Kakei, S., D.S. Hoffman, and P.L. Strick, *Muscle and movement representations in the primary motor cortex.* Science, 1999. **285**(5436): p. 2136-9.
13. Scott, S.H., *The role of primary motor cortex in goal-directed movements: insights from neurophysiological studies on non-human primates.* Curr Opin Neurobiol, 2003. **13**(6): p. 671-7.
14. Churchland, M.M., et al., *Neural population dynamics during reaching.* Nature, 2012. **487**(7405): p. 51-6.
15. Schaffelhofer, S. and H. Scherberger, *Object vision to hand action in macaque parietal, premotor, and motor cortices.* Elife, 2016. **5**.
16. Churchland, M.M., et al., *Stimulus onset quenches neural variability: a widespread cortical phenomenon.* Nat Neurosci, 2010. **13**(3): p. 369-78.
17. Yu, B.M., et al., *Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity.* J Neurophysiol, 2009. **102**(1): p. 614-35.
18. Tervo, D.G., et al., *A Designer AAV Variant Permits Efficient Retrograde Access to Projection Neurons.* Neuron, 2016. **92**(2): p. 372-382.
19. Dana, H., et al., *Sensitive red protein calcium indicators for imaging neural activity.* Elife, 2016. **5**.

20. Gerfen, C.R., R. Paletzki, and N. Heintz, *GENSAT BAC cre-recombinase driver lines to study the functional organization of cerebral cortical and basal ganglia circuits.* Neuron, 2013. **80**(6): p. 1368-83.
21. Kawai, R., et al., *Motor cortex is required for learning but not for executing a motor skill.* Neuron, 2015. **86**(3): p. 800-12.
22. Otchy, T.M., et al., *Acute off-target effects of neural circuit manipulations.* Nature, 2015. **528**(7582): p. 358-63.
23. Xu, T., et al., *Rapid formation and selective stabilization of synapses for enduring motor memories.* Nature, 2009. **462**(7275): p. 915-9.
24. Fu, M., et al., *Repetitive motor learning induces coordinated formation of clustered dendritic spines in vivo.* Nature, 2012. **483**(7387): p. 92-5.
25. Peters, A.J., S.X. Chen, and T. Komiyama, *Emergence of reproducible spatiotemporal activity during motor learning.* Nature, 2014. **510**(7504): p. 263-7.
26. Whishaw, I.Q., *Loss of the innate cortical engram for action patterns used in skilled reaching and the development of behavioral compensation following motor cortex lesions in the rat.* Neuropharmacology, 2000. **39**(5): p. 788-805.
27. Ramanathan, D., J.M. Conner, and M.H. Tuszynski, *A form of motor cortical plasticity that correlates with recovery of function after brain injury.* Proc Natl Acad Sci U S A, 2006. **103**(30): p. 11370-5.
28. Tonegawa, S., et al., *Memory Engram Cells Have Come of Age.* Neuron, 2015. **87**(5): p. 918-31.
29. Packer, A.M., et al., *Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo.* Nat Methods, 2015. **12**(2): p. 140-6.
30. Shemesh, O.e.a., *Temporally precise single-cell-resolution optogenetics.* Nature Neuroscience, 2017. **20**(12): p. 1796-1806.
31. Kabra, M., et al., *JAABA: interactive machine learning for automatic annotation of animal behavior.* Nat Methods, 2013. **10**(1): p. 64-7.
32. Dalal, N. and B. Triggs, *Histograms of oriented gradients for human detection.* 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Vol 1, Proceedings, 2005: p. 886-893.
33. Laptev, I., et al., *Learning realistic human actions from movies.* 2008 Ieee Conference on Computer Vision and Pattern Recognition, Vols 1-12, 2008: p. 3222-+.
34. Jun, J.J.e.a., *Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction.* bioRxiv, 2017. **101030**.
35. Huang, C.C., et al., *Convergence of pontine and proprioceptive streams onto multimodal cerebellar granule cells.* Elife, 2013. **2**: p. e00400.
36. Liu, R.Y., Parelius, J. M., & Singh, K, *Multivariate analysis by data depth: descriptive statistics, graphics and inference.* The annals of statistics, 1999. **27**(3): p. 783-858.

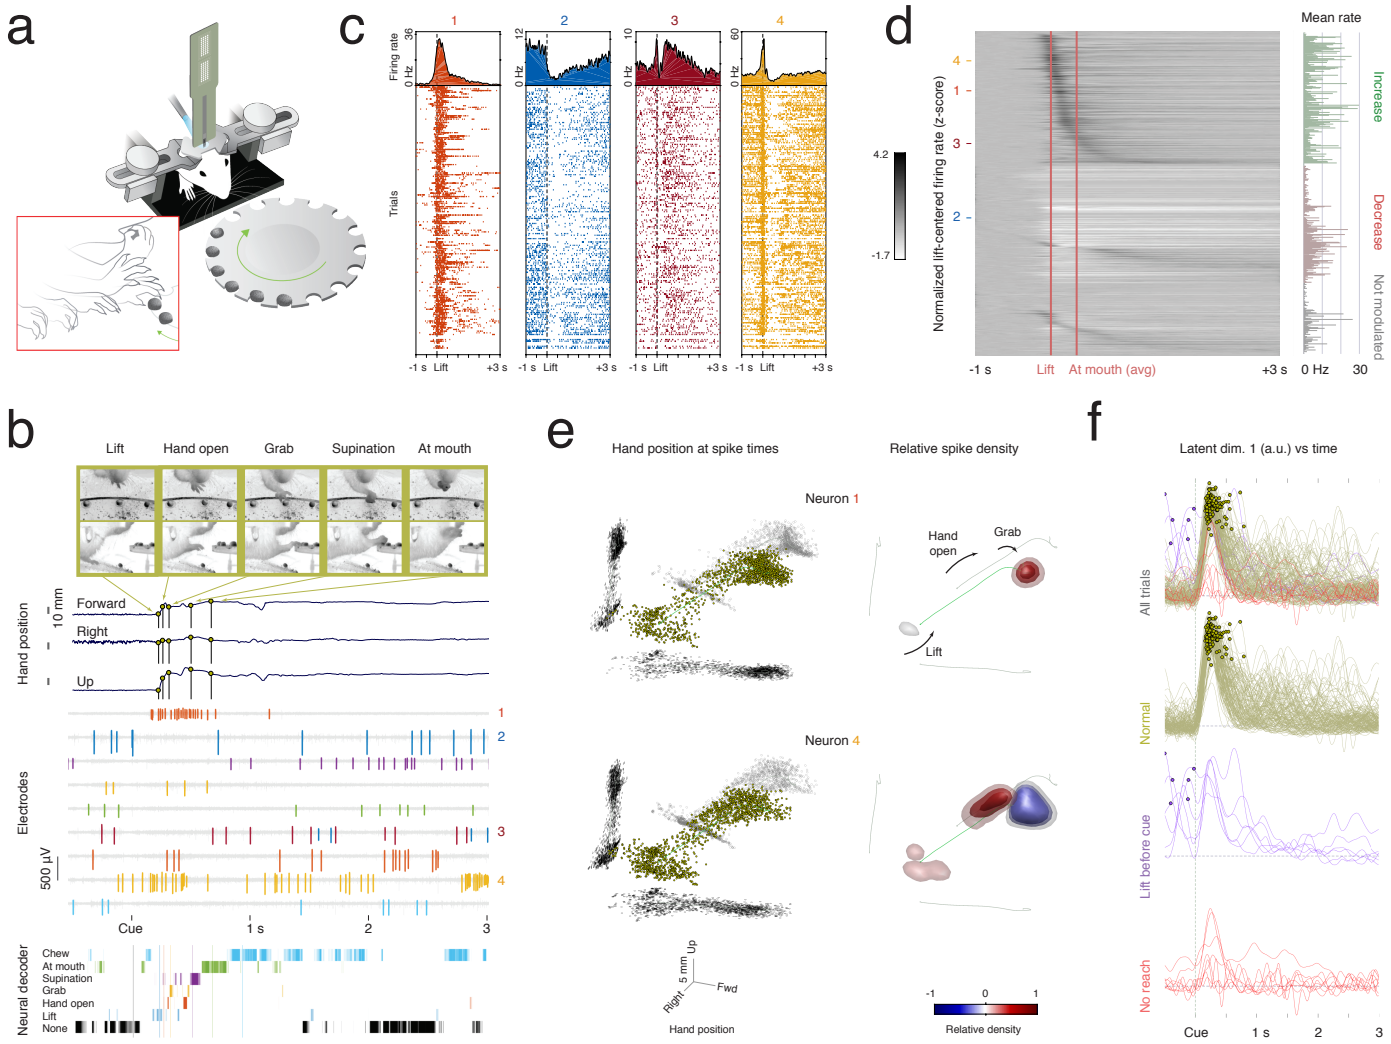

Figure 1: Neural dynamics in sensorimotor cortex during a prehension task. **a**, Experimental setup. Head-fixed mice reached for a pellet of food following an acoustic cue during recording and optogenetic perturbation of cortical activity. **b**, Raw video, electrophysiological recordings, mouse behavior, and neural decoding on a single trial. Three-dimensional hand trajectories and the timing of each waypoint in the behavioral sequence were extracted from video using computer vision methods. **c**, Spike raster plots and peri-event time histograms for four example neurons, centered on lift. **d**, Average z-scored firing rates and mean firing rates for all 584 motor cortical neurons. During prehension, most neurons exhibited increases (42%) or decreases (36%) in spike counts around lift (rank sum test with Benjamini-Hochberg correction, $q < 0.05$). **e**, Left: spatial locations of the hand at spike times for two example neurons. Right: difference between the spatial density of the spike-triggered hand position and the density of the hand position over the full trajectory. Shaded regions represent iso-density-difference contours; many spikes occur in the red regions, and few spikes occur in the blue regions. **f**, Neural population states for single trials, estimated using Gaussian Process Factor Analysis (GPFA). Circles correspond to lift times. Colors indicate normally timed reaches (yellow), early reaches (lift before cue, purple), and trials without a reach (red).

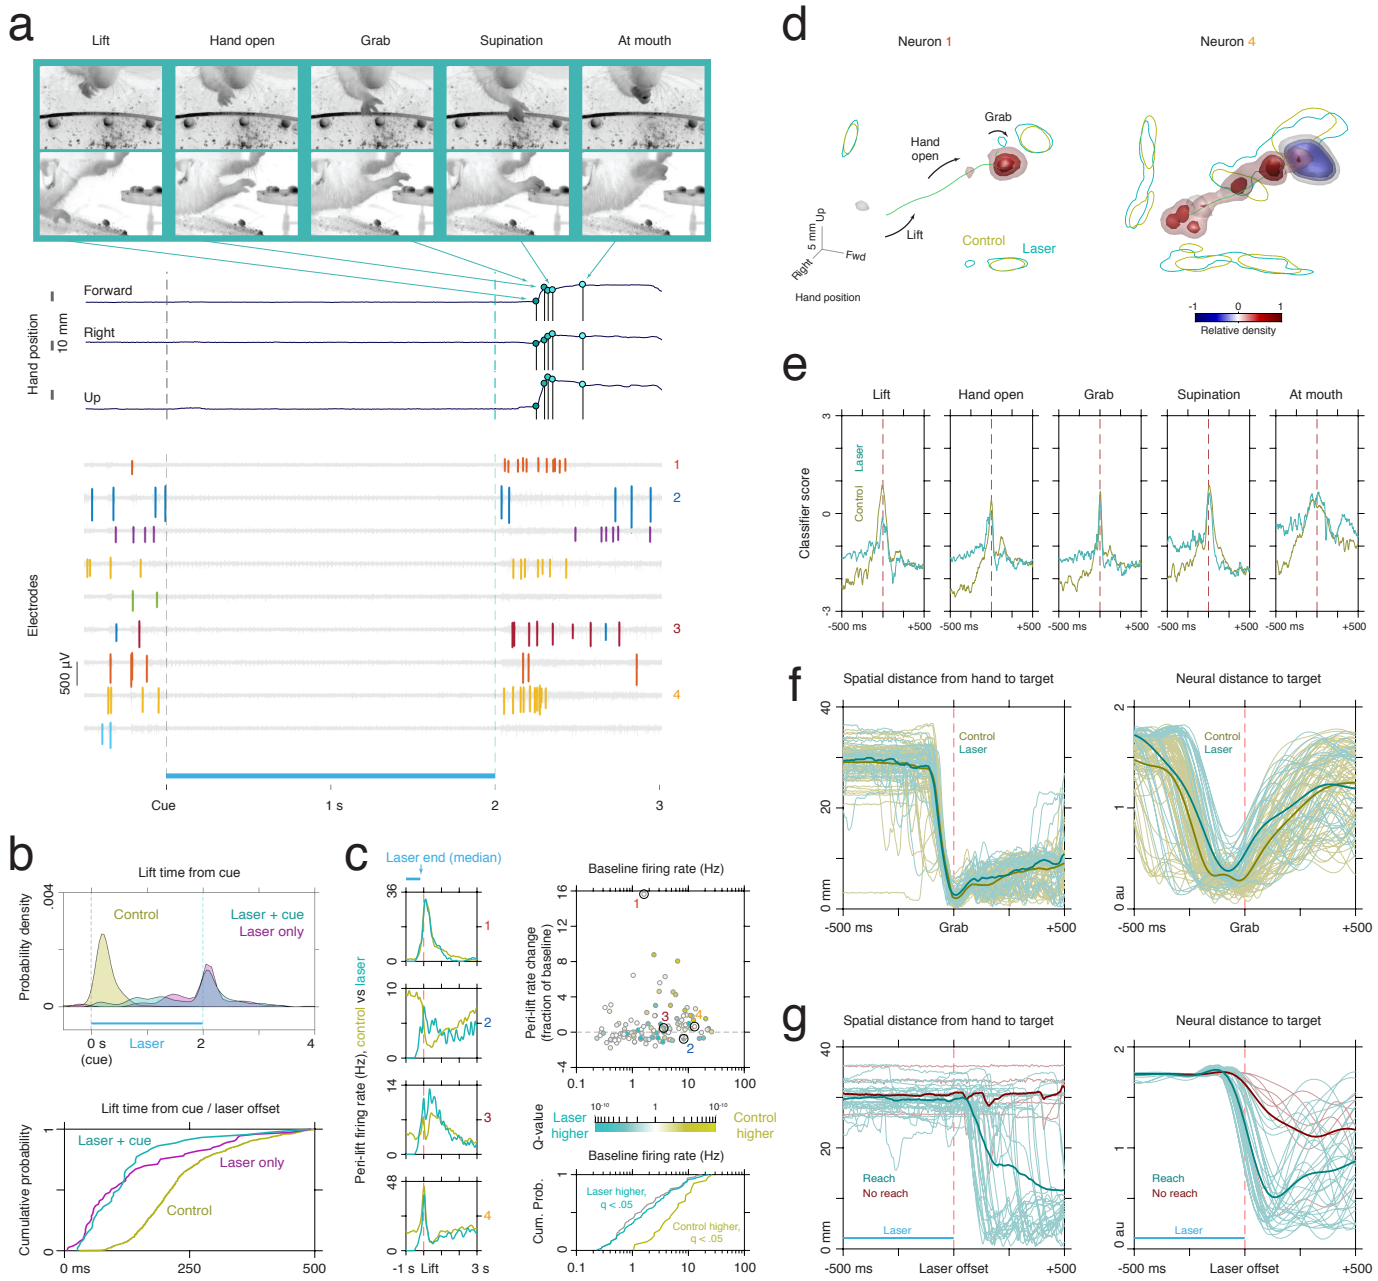

Figure 2: Neural dynamics during prehension under optogenetic silencing of motor cortex in VGAT-ChR2-EYFP mice. **a**, Electrophysiology and behavior during a laser trial. The four example neurons are the same ones displayed in figure 1. **b**, Upper: distribution of lift times on control (yellow), laser + cue (blue), and laser-only (magenta) trials. Curves indicate probability density obtained from a Gaussian kernel density estimate with σ = 100 ms. Lower: distribution of cue-to-lift times (yellow) and laser-offset-to-lift times for laser + cue trials (blue) and laser-only trials (magenta). **c**, Left: peri-lift firing rates for the four example neurons on control (yellow) and laser (blue) trials. Right, upper: baseline firing rate and lift modulation, color-coded by average q-value from lift to lift +500 ms. Right, lower: distribution of baseline firing rates for neurons with higher peri-lift rates on control trials ($q < 0.05$, yellow) or laser trials ($q < 0.05$, blue). **d**, Relative spatial density of spikes during laser and control trials. Shaded regions represent iso-density-difference contours for the spike-triggered position and the full trajectory on laser trials. Yellow and blue contours represent projection of control and laser data, respectively. **e**, Trial-averaged mean classifier scores for each behavior. The classifier was trained to detect each behavioral waypoint using neural population activity. **f**, Left: spatial distance from hand to target on control and laser trials, centered on grab. Right: neural distance from population state to target state, estimated using GPFA. **g**, Spatial (left) and neural (right) distance to target, centered on laser offset, for trials with (blue) and without (red) post-laser rebound reaches.

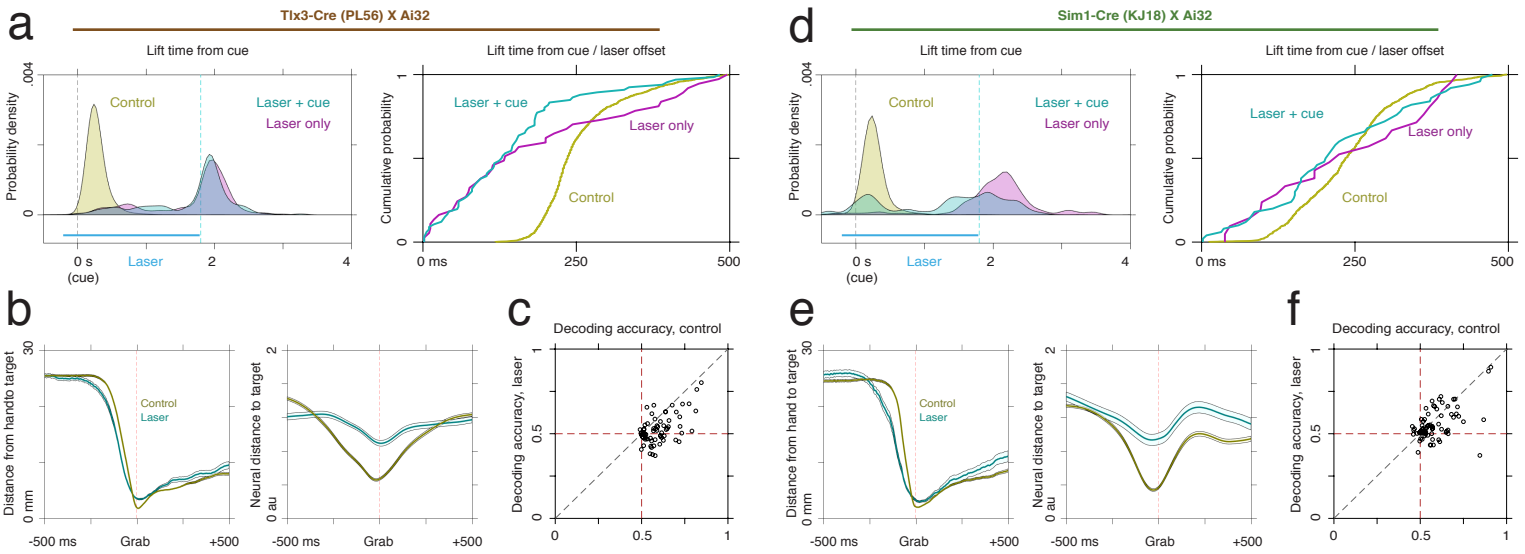

Figure 3: Influence of cell-type-specific perturbation of excitatory neurons on neural dynamics and behavior. **a**, Lift times for control (yellow), laser + cue (blue), and laser only (magenta) trials under optogenetic perturbation in Tlx3-Cre (PL56) X Ai32 mice. **b**, Spatial (left) and neural (right) distance to target on control (yellow) and laser (blue) trials for Tlx3-Cre X Ai32 mice. **c**, Single-neuron lift decoding accuracy on laser and control trials for Tlx3-Cre X Ai32 mice. **d**, Lift times under optogenetic perturbation in Sim1-Cre (KJ18) X Ai32 mice. **e**, Spatial and neural distance to target for Sim1-Cre X Ai32 mice. **f**, Single-neuron lift decoding accuracy for Sim1-Cre X Ai32 mice.

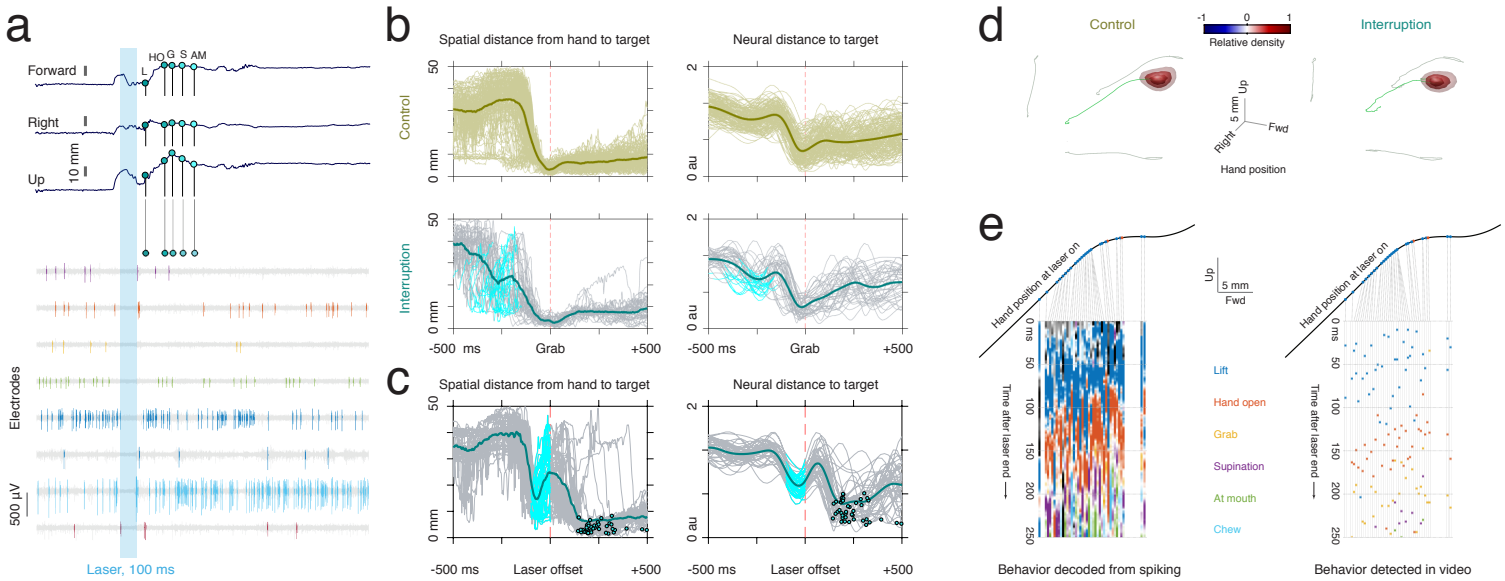

Figure 4: Real-time interruption of neural dynamics and behavior. **a**, Hand trajectory, behavior times, and neural recording during an interrupted reach. **b**, Spatial (left) and neural (right) distance to target on control (upper) and interruption (lower) trials, centered on grab. **c**, Spatial (left) and neural (right) distance to grab on interruption trials, centered on laser offset. Circles indicate grab times, and the blue regions indicate laser-on periods. **d**, Relative density of spike location on control (left) and interruption (right) trials for and example neuron. Shaded regions represent iso-density-difference contours for the spike-triggered position and the full trajectory. **e**, Decoding of behavioral waypoints from neural population activity after interruption of neural activity and movement. Upper traces: hand position at the time of laser onset for interruption trials, projected onto the mean hand trajectory. Left inset: waypoints decoded from spiking following the end of the laser. Right inset: behavioral waypoints detected in the video following the end of the laser.

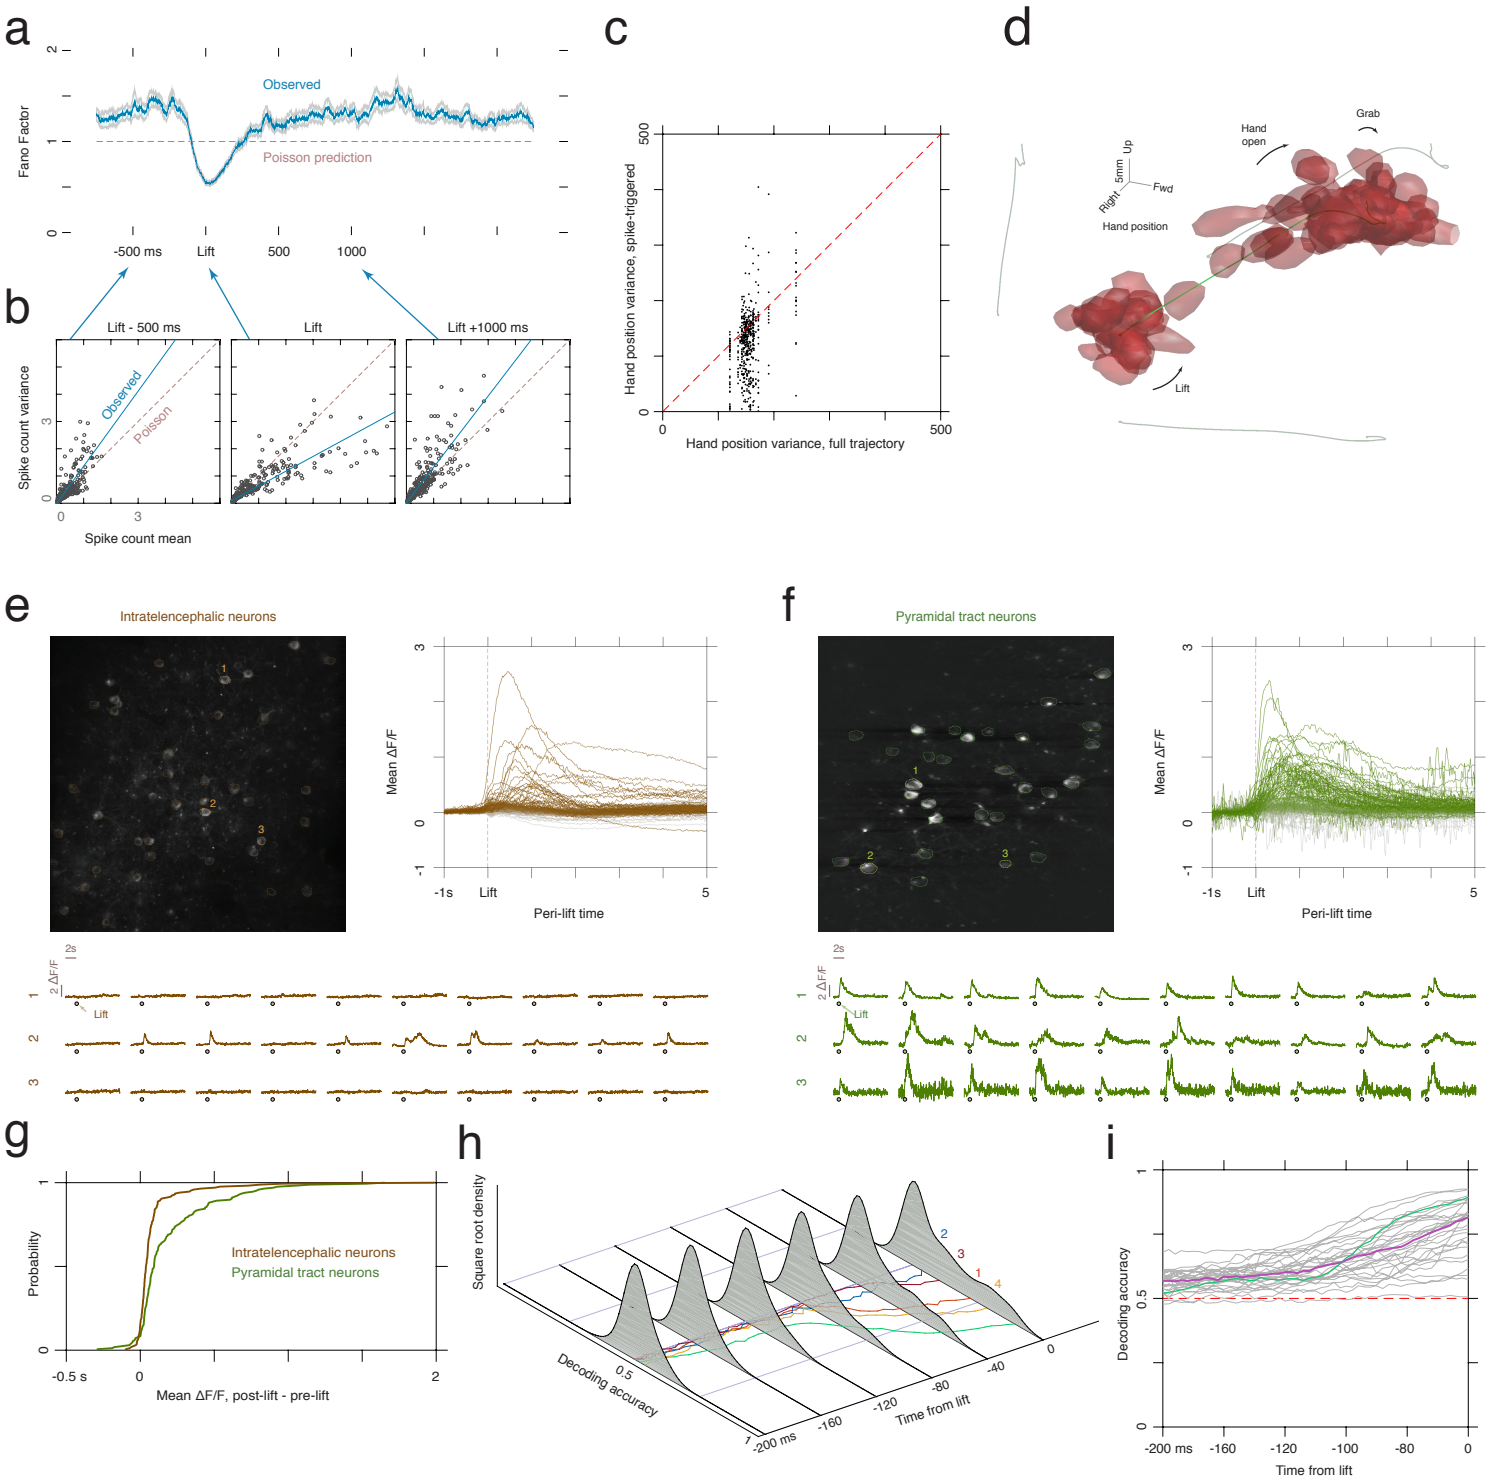

Extended data figure 1: Cortical activity during normal reaching. **a**, Variance-to-mean ratio (Fano factor) for lift-centered spike counts. Poisson spiking should result in a constant value of one. **b**, Scatterplots of spike count variance versus mean for all recorded neurons at different temporal offsets from lift. **c**, Variance of hand position sampled at spike times vs variance of hand position over the full trajectory. **d**, Spatial tuning of neurons. Contours indicate iso-density-difference regions where the spike-triggered density is larger than the full trajectory density (0.75 of maximum density difference). Neurons plotted have at least 100 spikes during reaching, and have a ratio of full-trajectory position variance to spike-triggered position variance greater than the median of 1.12 (n = 134). **e**, Upper left: two-photon calcium imaging from intratelencephalic (IT) neurons (injection of rAAV2-retro carrying jRCaMP1b into contralateral striatum). Lower: ΔF/F traces from three IT neurons on ten consecutive trials. Upper right: lift-centered mean ΔF/F for all neurons. **f**, Two-photon calcium imaging during the task, as in **d**, from identified PT neurons (injection of rAAV2-retro with jRCaMP1b into the ipsilateral pontine nuclei). **g**, Distribution of mean difference between pre- and post-lift ΔF/F (1 s post-lift vs 1 s pre-lift) across IT and PT neurons. **h**, Distribution of single-neuron lift decoding accuracy at offsets ranging from lift -100ms to lift time. The example neurons from figure 1 are displayed as line plots, and the full-sample decoding accuracy for the dataset is displayed in the green line. **i**, Accuracy of decoding using all simultaneously-recorded neurons for each dataset. The example dataset is in green, and the median is magenta (median accuracy 0.82 at lift time).

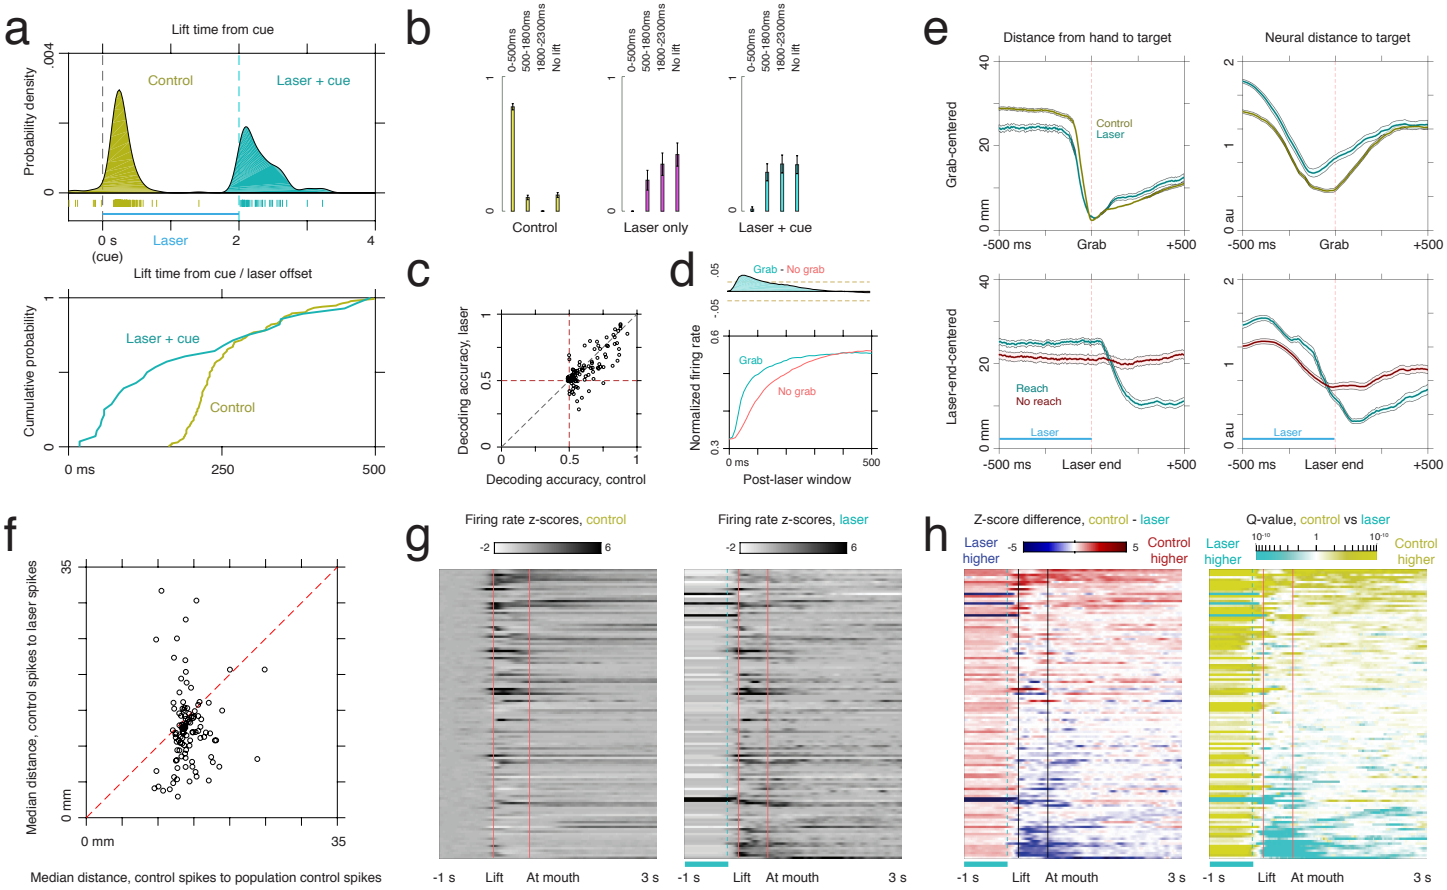

Extended data figure 2: Characterization of voluntary prehension and opto-prehension. **a**, Upper: lift times for the example dataset in figure 1. Yellow and blue curves represent densities for control and laser trials, respectively, and ticks represent single trials. Lower: distribution of reaction times from cue to lift (yellow, control trials) and from the end of the laser to lift (blue, laser + cue trials). **b**, Fraction of trials with lifts immediately following the cue (cue to cue +500 ms), during the laser epoch (cue +500 ms to laser end), after the laser (end of laser to end +500 ms), and trials with no lift in VGAT-ChR2-EYFP mice. **c**, Zero-lag decoding accuracy for single neurons on control and laser trials. **d**, Lower: normalized post-laser-end firing rates for trials with (blue) and without (red) opto-prehensions. For each post-laser window duration, the normalized firing rate for a given neuron is the rank of the spike counts in the window, divided by the duration of the window. The curves show the average values across the sample of neurons. Upper: difference between normalized firing rates on trials with and without post-laser grabs. **e**, Upper: grab-centered spatial (left) and neural (right) distance to target for control and laser trials, pooled over all datasets using VGAT-ChR2-EYFP mice. Lower: laser-end-centered distance to target on laser trials with and without rebound reaches. **f**, Similarity of spatial distribution of the hand at spike times during voluntary prehension and opto-prehension. X-axis: median spatial distance from a spike on a control trial for a given neuron to a spike on a control trial for a randomly-selected neuron. Y-axis: median spatial distance from a spike on a laser trial to a spike on a control trial for the same neuron. Rank sum test, p = 1.3e-6. **g**, Comparison of firing rates during voluntary and opto-prehension. Left: average lift-centered firing rate z-scores, sorted by the q-values for the comparison between firing rates on voluntary and opto-prehensions. Right: lift-centered z-scores for laser trials. **h**, Left: differences in firing rate z-scores between control and laser. Right: q-values for rank-sum test, color-coded by whether the firing rate is higher on control or laser trials.

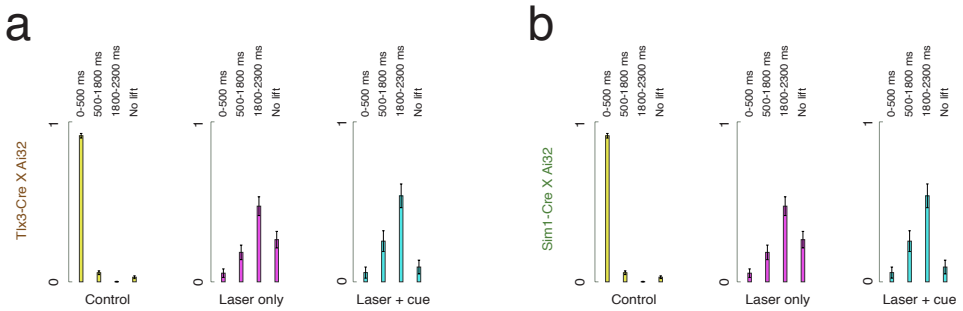

Extended data figure 3: Behavior summary for perturbation of excitatory subpopulations. **a**, Fraction of trials with lifts immediately following the cue (cue to cue + 500ms), during the laser epoch (cue + 500ms to laser offset), after the laser (offset of laser to offset + 500ms), and trials with no lift in Tlx3-Cre X Ai32 mice. **b**, Lift times, as in **a**, for Sim1-Cre X Ai32 mice.

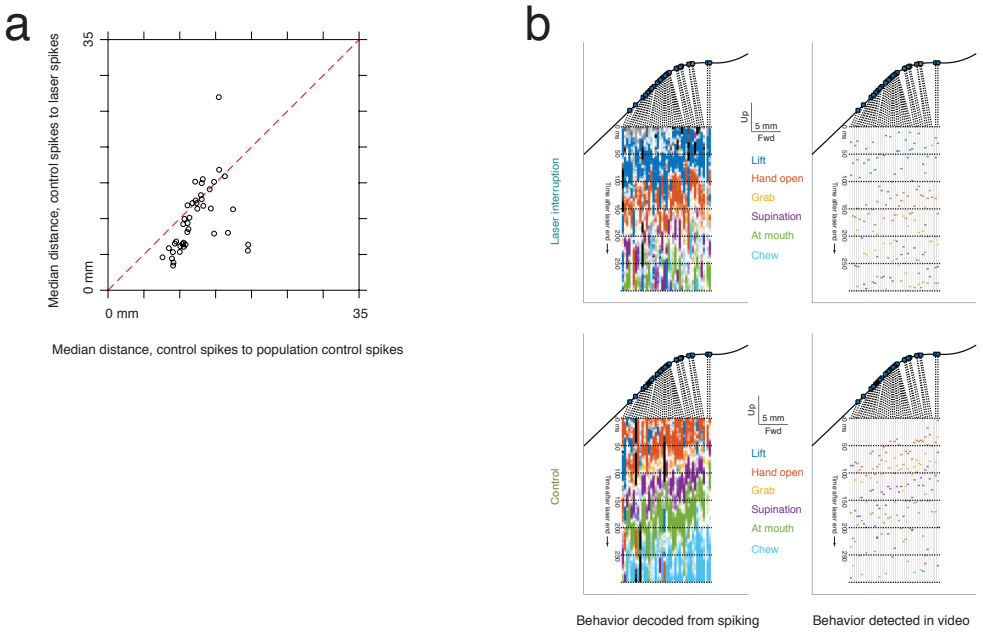

Extended data figure 4: Comparison of neural activity on control and interrupted reaches. **a**, X-axis: median spatial distance from a spike on a control trial for a given neuron to a spike on a control trial for a randomly-selected neuron. Y-axis: median spatial distance from a spike on a laser interruption trial to a spike on a control trial for the same neuron. **b**, Decoding of behavioral waypoints from neural data on laser interruption and control trials. The upper panels show the decoding results for laser interruption trials, as in figure 4e. The lower panels show the decoding for matched control trials in which the hand was in the same position as on a corresponding interruption trial at the beginning of the laser period.

Question	Approach	Result	Figure
What are the neural dynamics in mouse motor cortex during the reach-to-grab sequence?	Record neural ensemble activity during the movement.	Single neurons have highly reliable responses across trials, but the population exhibits a diverse range of response patterns tiling the entire movement sequence.	1b-f, E1a-d
Where in the brain does the cortex distribute reach-related signals?	Image activity of identified projection classes during the movement.	Both intratelencephalic (striatum-projecting) and pyramidal tract (pons-, medulla-, and spinal-projecting) neurons are active during the movement, but this activity is more pronounced in pyramidal tract neurons.	E1e-g
Is cortex necessary and sufficient for generating the movement sequence?	Activate inhibitory neurons during the task.	Silencing cortex blocks voluntary prehension (necessity), and termination of silencing induces a rapid opto-prehension (sufficiency).	2a-b, E2a-b
Do neural dynamics reflect motor commands?	Compare dynamics during voluntary and opto-prehension using (1) peri-lift firing rate statistics, (2) spatial analysis of hand position at spike times, (3) neural decoding of behavior waypoints from neural activity, and (4) distance-to-target analysis with Gaussian Process Factor Analysis.	Neural dynamics during opto-prehension largely recapitulate dynamics during voluntary prehension. However, a subset of neurons had different firing patterns in the two reach types, and cells with higher peri-lift firing rates during voluntary prehension were distinguished by higher baseline firing rates outside of the behavior.	2c-g, E2c-h
Are these dynamics robust to multiple types of perturbation?	Activate projection-class-specific subtypes of excitatory neurons while recording neural ensembles during the movement.	Activation of subsets of excitatory neurons blocked voluntary prehension and induced opto-prehension. Neural dynamics during opto-prehension recapitulated the dynamics during voluntary prehension.	3a-f, E3a-b
Are these dynamics required continously throughout the movement sequence?	Silence cortex in the middle of the movement while recording neural ensembles.	Mid-movement cortical inactivation halted behavior, and following this perturbation cortical dynamics restarted and drove completion of the full movement sequence.	4a-e, E4a-b

Extended data figure 5: Diagram of the relationship between the experiments and results.