Stable encoding of the natural variability of a dexterous motor skill in the cortex of freely behaving mice

Elizabeth A. de Laittre[1] and Jason N. MacLean[2]
[1] Committee on Computational Neuroscience, University of Chicago, Chicago, IL, USA
[2] Department of Neurobiology, University of Chicago, Chicago, IL, USA
Correspondence should be addressed to: jmaclean@uchicago.edu

Abstract: Skilled, goal-directed movements can exhibit trial-to-trial variability even in experts, particularly in response to dynamic environmental conditions or when perfect repetition is not required for success. Identifying where, to what extent, and how stably this variability is encoded in the nervous system will yield insight into how such learned movements are robustly maintained over time yet flexibly executed on each trial. We record calcium fluorescence activity in motor cortex—a key node in the multi-areal network responsible for movement control—in freely-moving mice performing a self-paced, precision reach-to-grasp task. High trial counts and rich single-trial variability enable rigorous statistical analysis of moment-to-moment movement encoding across matched behavioral sets over five days. We found that individual neurons in motor cortex stably encode details of paw, digit, and head movements, suggesting that reliable contributions from single cells support the consistent execution of skilled movements over time, even in complex, sensory-guided tasks like reach-to-grasp.

Even well-practiced behaviors exhibit variability with each attempt. In basketball, for instance, no two shots are identical, as each is shaped by factors such as starting posture, fatigue, and interactions with defenders and teammates. Yet, players consistently adapt and successfully perform this precise skill across these dynamic contexts. While generating goal-directed movements under such conditions remains a significant challenge in robotics (Billard and Kragic, 2019), biological nervous systems excel at this. To understand how motor skills are consistently executed across varying contexts, it is crucial to investigate how movement variability is represented in neural circuits and whether the representation of learned, complex motor behaviors remains stable over time. This insight can shed light on the circuit-level precision required to adapt to perturbations while maintaining reliable motor performance over extended periods of time.

Prehension, the act of reaching and grasping objects, is a vital yet complex movement requiring precise coordination of the trunk, shoulder, arm, wrist, and fingers, guided by rapid sensory feedback processing. The need to interact flexibly with diverse objects further enhances its complexity. Motor cortex plays a central role in the brain-wide circuit controlling precision reaching behaviors in mice (Guo et al., 2015; Wang et al., 2017; Miri et al., 2017; Galiñanes et al., 2018), which recent work suggests rely on online, active control rather than pre-planned, ballistic movements (Becker, Calame, et al., 2020).

Previously, mouse caudal forelimb area (CFA, or forelimb M1) has been shown to represent low-level features of forelimb movements, including muscle activation and kinematics, e.g. position and velocity, during forelimb-intensive behaviors performed by head-fixed mice (Omlor et al., 2019; Koh et al., 2024; Zhu et al., 2022; Grier et al., 2025). However, during a freely-moving reach-to-grasp behavior, which requires coordination of the whole body from gross movements of limbs and trunk to fine control of digits, the extent to which CFA encoding incorporates whole-body posture or digit movement details is unclear. While there is evidence for coordinated digit use in mice (Whishaw et al., 2017; Barrett et al., 2020; An et al., 2024) and posture-related activity in rats' sensorimotor cortices (Mimica et al., 2018; Disse et al., 2023), it is unknown

how these kinematic features are represented in mouse forelimb M1. Additionally, there has been little comparison of movement encoding properties in general across superficial layers of CFA, layer 2/3 (L2/3) and layer 5a (L5a), which have notable differences in connectivity and sit directly upstream of the better-studied deep projection neurons (Weiler et al., 2008; Oswald et al., 2013). Understanding how CFA encodes a range of movement features—such as proximal forelimb actions, whole-body posture, and fine digit control—is essential for clarifying its role in complex behaviors like reach-to-grasp. If M1 contributes to the online control of these movements as part of a broader sensorimotor circuit, it is likely to encode their moment-to-moment details.

Neural encoding in sensory and association brain areas often shows session-to-session variability, a phenomenon termed representational drift (reviewed in Micou and O'Leary, 2023; Driscoll et al., 2022; Chambers and Rumpel, 2017; Clopath et al., 2017; Lütcke et al., 2013). In contrast, motor areas exhibit more stable encoding over time (Jensen et al., 2022; Katlowitz et al., 2018; Dhawale et al., 2017; Gallego et al., 2020; Chestek et al., 2007; Fraser and Schwartz, 2012; Flint et al., 2016; but see Rokni et al., 2007 and Liberti et al., 2016). Most studies of the stability of motor encoding have relied on trial-averaged metrics, such as peri-event time histograms (PETHs), which can mask neural stability or instability by not accounting for behavioral variability within and across sessions. This raises the possibility that apparent neural instability could reflect changes in behavior rather than shifts in the mapping between neural activity and behavior (Liberti, Schmid, et al., 2022; Jensen et al., 2022). Additionally, the stability of moment-to-moment encoding for fine-grained movement features, such as joint position and velocity, remains underexplored. This is partly due to the challenge of studying behaviors that are variable enough to reveal distinct neural activity patterns for distinct movements while maintaining sufficient experimental control and sample sizes to ensure statistical rigor. The motor nervous system might maintain stable neural activity patterns for stereotyped learned movements but not for more flexible learned movements with higher trial-by-trial variance.

To address these gaps, we couple high speed video-based movement reconstruction with head-mounted one-photon calcium imaging in mice to examine motor cortical encoding of a whole-body, dexterous reach-to-grasp task that produces high trial counts with significant trial-to-trial variability within sessions. Despite this variability, individual mice consistently explore a similar kinematic movement space across days, enabling us to create matched datasets to directly compare the encoding of various movement variables—kinematics of the head, paw, and digits—across days. We find that the single trial variance of these movements is encoded in mouse forelimb motor cortex (CFA) and the nature of this encoding remains highly stable across days.

Results

A dexterous reach-to-grasp task elicits variable yet spatiotemporally coordinated forelimb and digit movements. To investigate the stability of motor cortical encoding of moment-to-moment forelimb and head movements, we trained unrestrained mice on a challenging reach-to-grasp task. Mice were trained to reach for a millet seed placed on a narrow pedestal positioned in one of two or three locations just outside an arena (Fig. 1A,B). Mice self-initiated delivery of a seed by moving a specific required distance away from the reach port (generally 1-2 body lengths away; see Methods for details of seed delivery systems). A critical aspect of the task is positioning the body and head after returning to the reach port, to ensure the pedestal is within a reachable distance and angle. In tandem, the mice must make micro-adjustments to forelimb and digit reaching movements based on body positioning. After returning to the reach port, mice performed a bout of reaches in quick succession during which they either successfully retrieved the seed or knocked it off. Throughout this paper, we treat each reach as a single trial. During a 40 minute daily session, a trained mouse

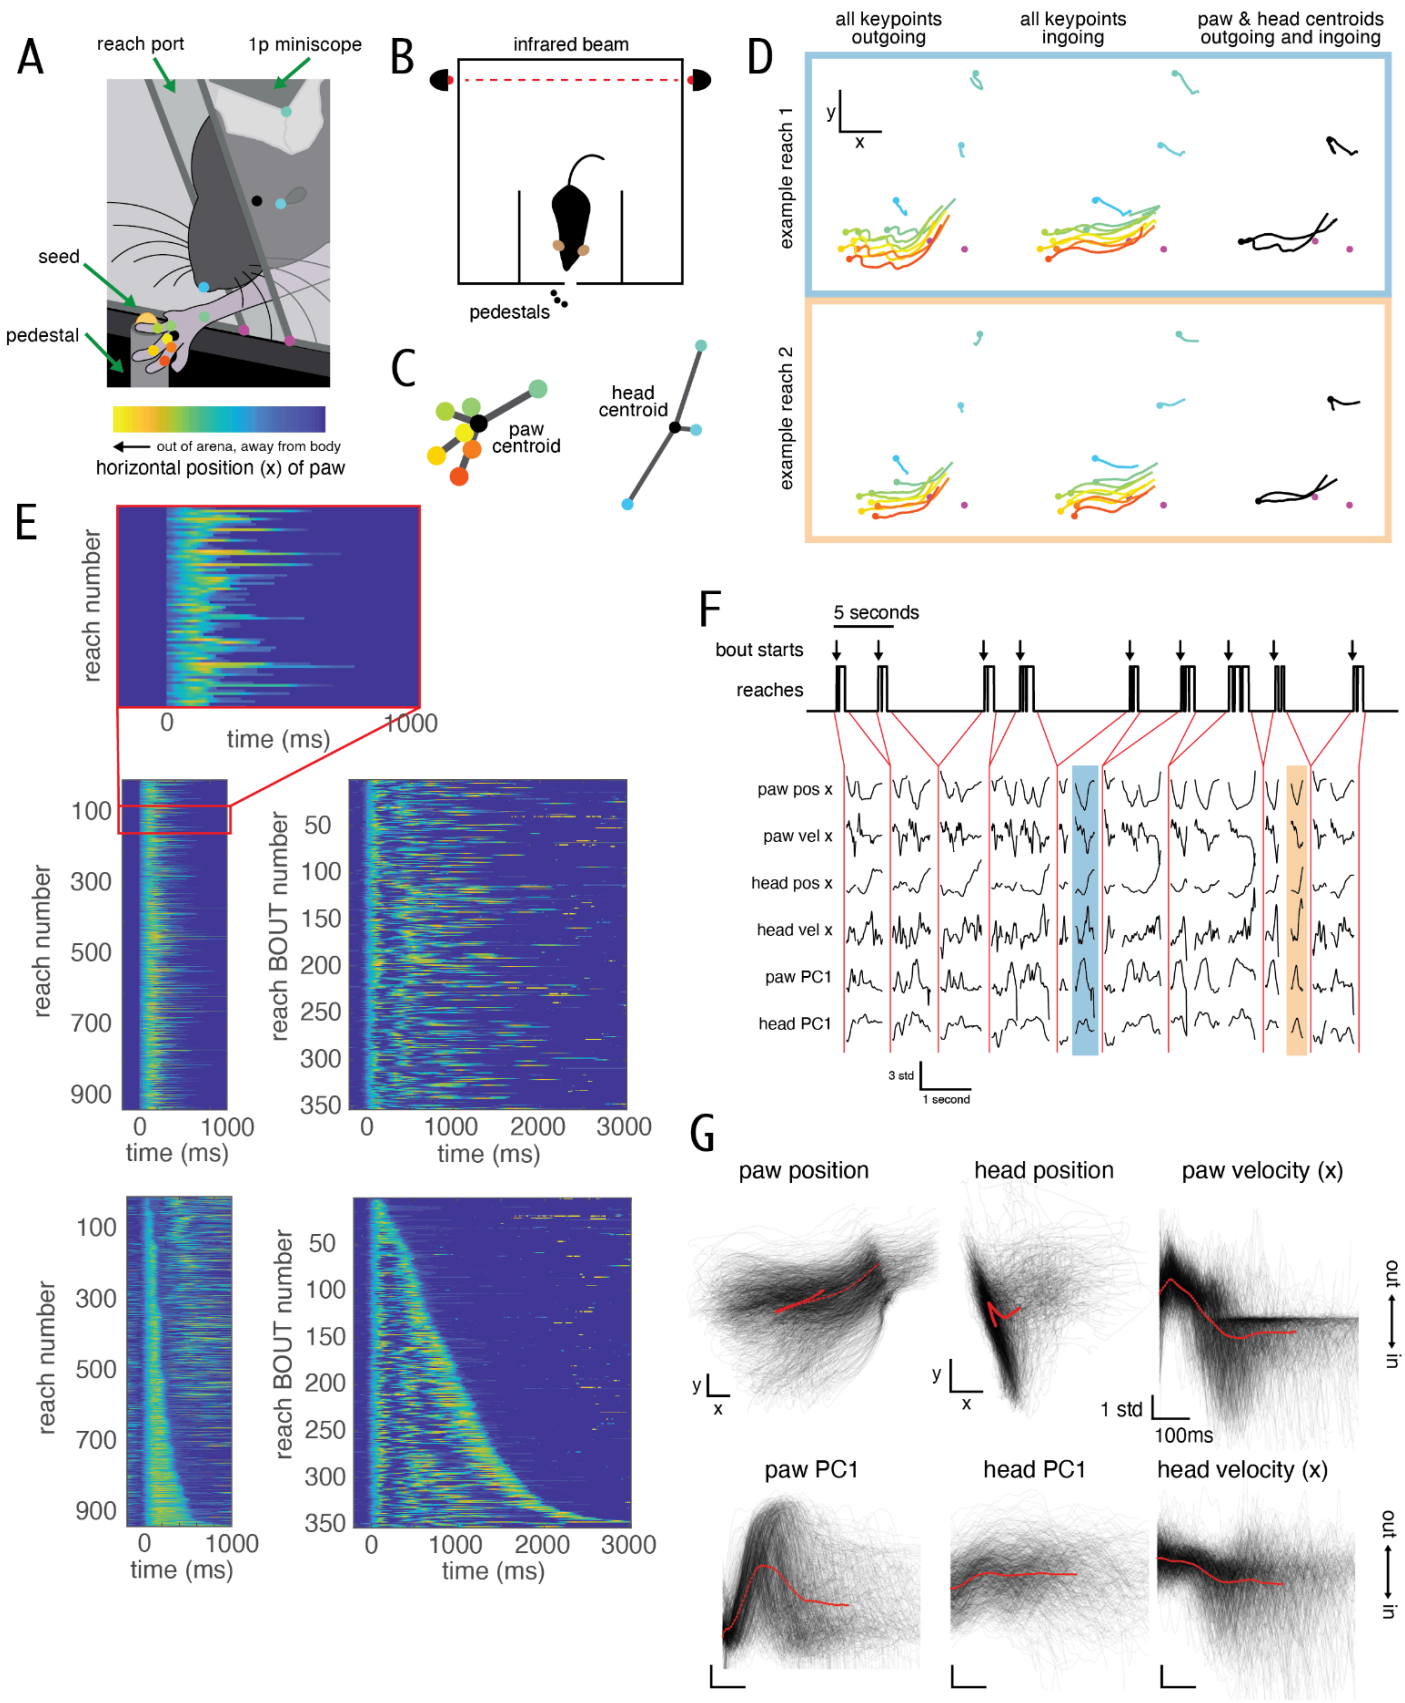
A
reach port
1p miniscope
seed
pedestal
out of arena, away from body
horizontal position (x) of paw
B
infrared beam
pedestals
C
paw centroid
head centroid
D
all keypoints outgoing
all keypoints ingoing
paw & head centroids outgoing and ingoing
example reach 1
example reach 2
y
x
E
reach number
time (ms)
reach BOUT number
F
5 seconds
bout starts
reaches
paw pos x
paw vel x
head pos x
head vel x
paw PC1
head PC1
3 std
1 second
G
paw position
head position
paw velocity (x)
paw PC1
head PC1
head velocity (x)
1 std
100ms
out
in

89

Figure 1: A dexterous reach-to-grasp task elicits variable yet spatiotemporally coordinated forelimb and digit movements.

A) Closeup view of mouse reaching, with 10 keypoints on the paw and head marked, from the viewing angle of the tracking camera. Colorbar at bottom indicates paw position in x relative to the view shown here; relevant for the heatmaps in E.
B) Schematic of mouse in arena viewed from above, indicating location of pedestals, reachport, and infrared beam.
C) Keypoints on the paw and head, and their distances to the paw and head centroids (black dots). Colors match those in A.
D) Two example reconstructed reaches. Colors match those in A and C. For all panels, dots indicate the endpoint (point of the reach that is furthest from the arena). Left column, all paw and head keypoints during the outgoing part of the two reaches. Middle column, same for the ingoing part. Right column, the paw and head centroids only for both the outgoing and ingoing parts of the reaches. Magenta dots indicate the base corners of the reach port, as in A.
E) Paw position in the x dimension plotted for all reaches and bouts in one daily session for an example mouse (mouse 5). Colormap is under the schematic in A; yellow colors indicate further outside the arena, blue colors indicate closer to or inside the arena. Top, reaches sorted in the order in which they occurred in the session (left), and the same reaches resorted into bouts, again in the order in which they occurred (right). Bottom, the same reach and bout data sorted according to reach and bout duration.
F) A subset of kinematic variables for some example reaches. Top, binary variable indicating when the animal is reaching (paw is outside the arena), with arrows at the start of bouts. Bottom, time courses of 6 of the 14 kinematic variables extracted via keypoint tracking, during the bouts plotted above. Blue and orange shading indicate the two example reaches plotted in D.
G) All reaches in one daily session for an example mouse (mouse 5) for the same 6 variables as in F. Individual reaches are plotted in transparent gray, the average (time-locked to reach onset) is plotted on top in red. Scale bars for paw and head position correspond to 30 pixels in x and y; the view is the same as is schematized in A and for example reaches in D. Scale bar for paw velocity, head velocity, paw PC1, and head PC1 are 100ms and 1 standard deviation of each variable.

will perform hundreds to thousands of individual reaches (1936 reaches per session +/- 542 [789, 2831]; mean +/- st.dev. [range] across all sessions for all mice, n=25 sessions) and successfully retrieve and consume dozens of seeds (70 seeds +/- 34 [27, 139]; mean +/- st.dev. [range]).

The pedestal has a diameter similar to that of the seed, making it easy for the seed to be knocked off (Fig. 1B). Because of this and the speed at which the paw approaches the pedestal, the mouse has only one chance to contact the seed—resulting either in a successful grasp or, more likely, knocking the seed off and making it unreachable. Consequently, the task demands precise paw placement and highly coordinated, precisely timed digit control, as only a small range of movements will successfully secure the seed. The narrow pedestal also prevents the mouse from repositioning the seed within its digits after contact. Mice must then maintain their grip while transporting the seed back into the arena, adding to the task's complexity.

To quantitatively describe digit, paw, and head movements during reaching, we recorded high-speed video on all days and employed markerless keypoint tracking to reconstruct the kinematics of each reach (Fig. 1A,C,D; DeepLabCut; Mathis et al., 2018; Nath et al., 2019). From this keypoint tracking, we extracted the position and velocity of the paw and head centroids and used principal component analysis to estimate proxies for the shape and orientation of the paw and the orientation of the head (Fig. 1C; Methods). The position and velocity of the paw centroid reflect activation of proximal forelimb muscles controlling the elbow and shoulder, while principal components (PCs) of the paw markers reflect activation of distal muscles controlling the wrist and digits. Similarly, the position, velocity, and orientation of the head reflect whole-body posture and potentially activation of shoulder and neck muscles.

Within a 40 minute daily session, trained mice exhibit high reach-to-reach variability in the speed, shape, and duration of digit, paw, and head movements during reaching (Fig. 1E,F,G). This trial-to-trial variability persists even though the mice demonstrate stable success rates during the period of study (Supplementary Fig. 1). We consider this within-session variability advantageous, as it allows us to examine the moment-to-moment neural encoding of these movements.

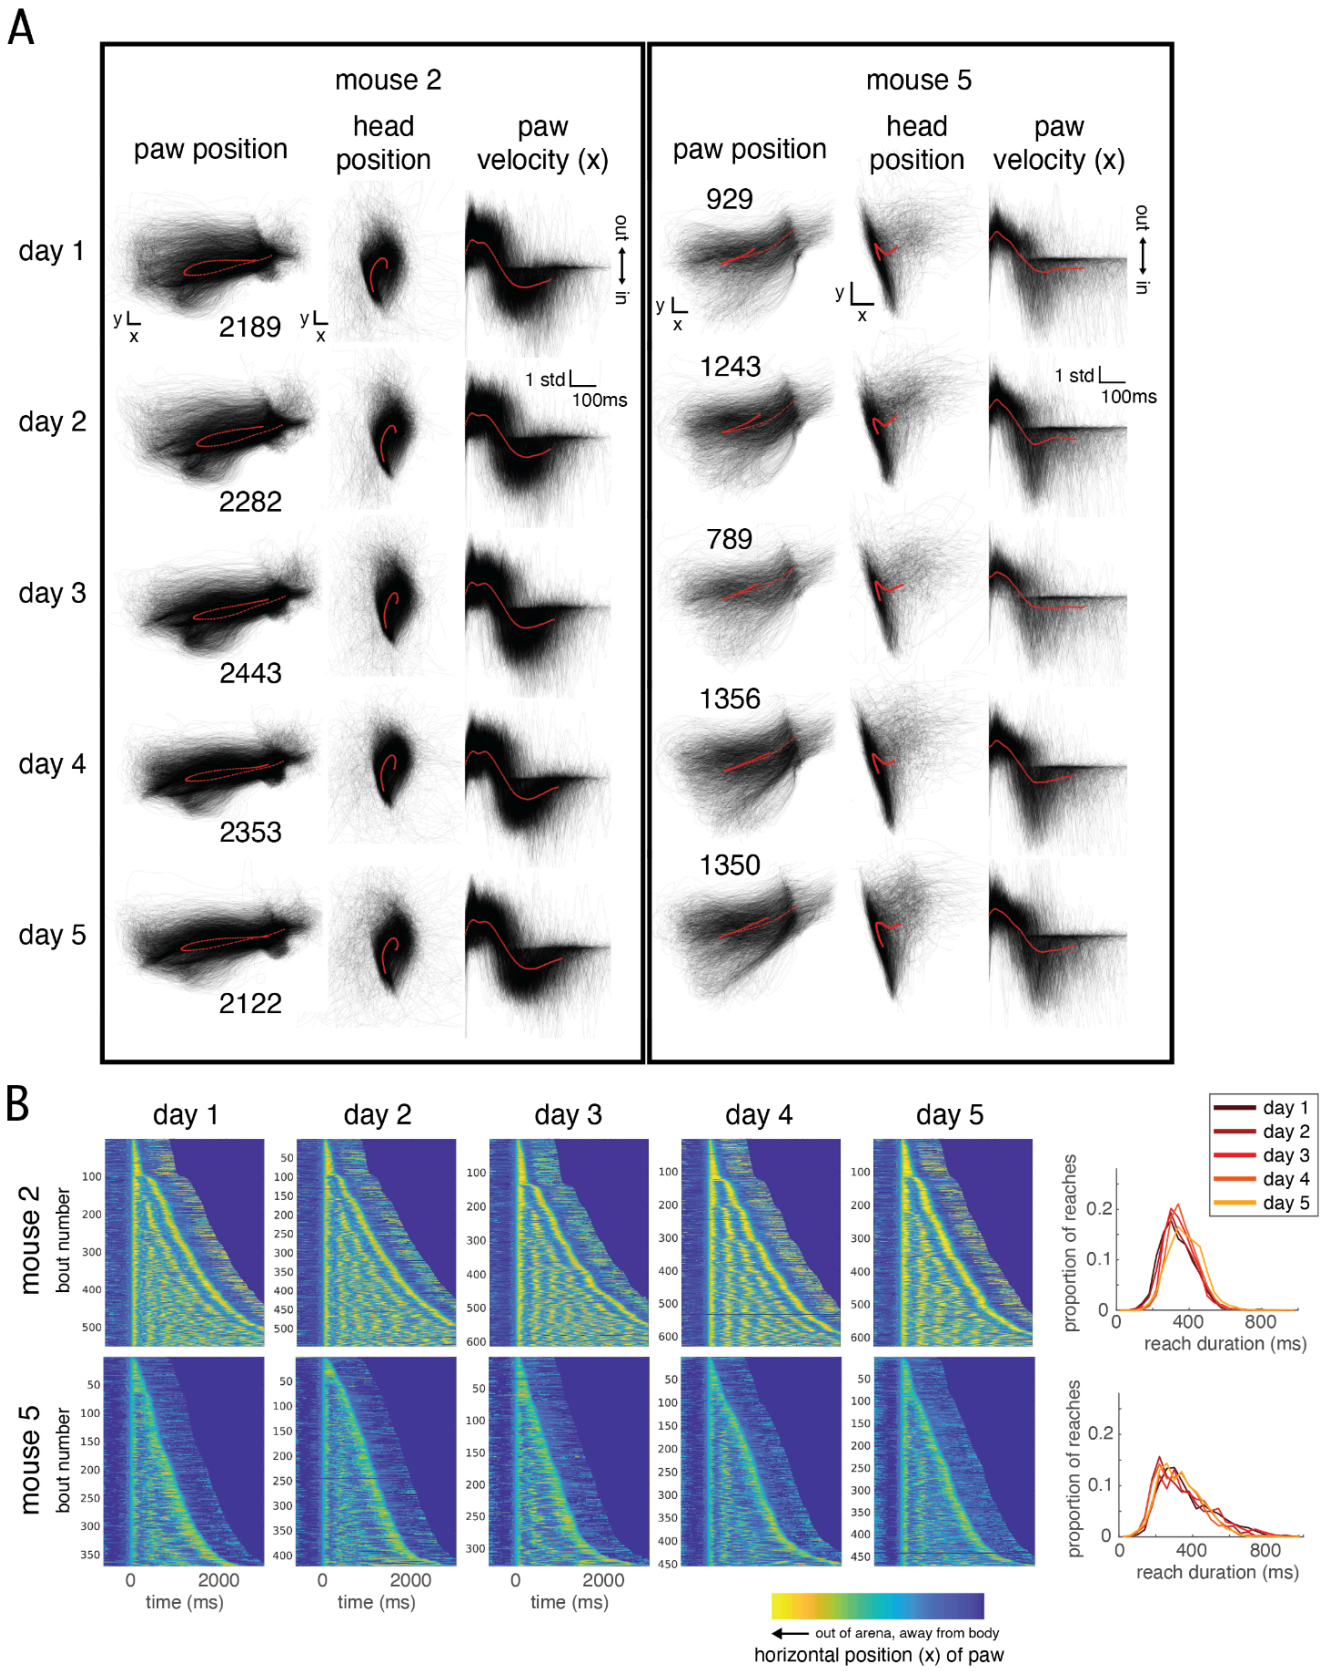
A
mouse 2
paw position
head position
paw velocity (x)
mouse 5
paw position
head position
paw velocity (x)
day 1
day 2
day 3
day 4
day 5
2189
2282
2443
2353
2122
929
1243
789
1356
1350
out
in
1 std
100ms
y
x
B
day 1
day 2
day 3
day 4
day 5
mouse 2
mouse 5
bout number
time (ms)
proportion of reaches
reach duration (ms)
out of arena, away from body
horizontal position (x) of paw

Figure 2: Each animal develops its own unique reaching pattern, yet consistently uses a similar range of movements across days.

A) All reaches in each of five daily sessions for two example mice, for 3 of the 14 kinematic variables extracted via keypoint tracking. As in Fig. 1G, individual reaches are plotted in transparent gray, the average (time-locked to reach onset) is plotted on top in red. Numbers next to the paw position plots indicate the number of individual reaches the animal made in that session.

B) For the same two example mice (mouse 2 and 5), for all five days: Left, paw position in the x dimension plotted the same way as in Fig. 1E. Right, probability density distributions of time duration of individual reaches, separated by session

Individual mice each develop a unique reaching pattern but consistently use a similar range of movements across days. To investigate whether neural coding for reach-to-grasp behavior changes or remains consistent over time, we analyzed data from five consecutive days late in training (starting between day 20 and 22 of training), a period during which mice demonstrate stable success rates (Supplementary Fig. 1). The within-session variability observed on a given day persists across the five days of study, with variability centered around a similar mean on each day (Fig. 2A). The set of specific digit, paw, and head movements employed differ across mice, as each mouse developed its own strategy for retrieving the seed (Fig. 2A,B, Supplementary Fig. 2). For each mouse, however, the distributions of all measured kinematic properties remain consistent across all five days. This was confirmed by a covariate balancing procedure used in later sections: after subsetting timepoints to match the joint distributions of kinematic variables across days, the majority of reaching timepoints were retained, indicating that the original distributions were already highly similar (Supplementary Fig. 3; Methods). This consistency in movement across days enables use of this dataset to compare neural encoding of reach, grasp, and carry behaviors over multiple days.

Neurons in CFA M1 motor cortex modulate activity during reaching. We recorded the activity of populations of individual neurons in caudal forelimb area of mouse motor cortex (CFA, or forelimb M1) using an Inscopix miniscope to perform one-photon calcium imaging in freely moving (non-headfixed) mice. We combined the miniscope with a prism to achieve an imaging plane perpendicular to the cortical surface (Andermann et al., 2013; Resendez et al., 2016). This setup enabled simultaneous imaging of putative cortical layers 2/3 and 5a, capturing hundreds of cells in a single field of view (244 cells +/- 50 [138, 347]; mean +/- st.dev. [range], n=25 datasets). By consistently returning to the same field of view each day, we were able to longitudinally track the activity of many of the same cells over multiple days (Fig. 3A,B).

The vast majority of neurons exhibited significant modulation during reaching relative to all other timepoints: 5544 out of 6105 cell-days (90.8%), where each cell-day represents a single day's observation for a given neuron, which constitutes 90.8% of each dataset (90.8% +/- 4.0% [82.1%, 96.5%]; mean +/- st.dev. [range]; ANOVA with $p = 0.05$, with Bonferroni correction for the number of cells in that dataset (field of view on that day); Methods). Across the recorded population, neurons displayed diverse temporal response profiles, with activity observed before, during, and/or after reaching (Fig. 3C). Additionally, cells were active during the intervals between reach bouts, which includes other behaviors such as consuming seeds, grooming, and locomotion to trigger the next seed delivery or to position the body at the reach port. In all analysis, each reach was treated as a trial. Neural activity showed notable variability across individual trials (Fig. 3C,E). Not all neurons were active during every reach, and even across reaches when a neuron was active, the magnitude of its activity varied. Despite this variability, neurons demonstrated stable overall patterns of activity across days (Fig. 3E,F), both in the pattern of single-trial activity relative to reach duration and in trial-averaged responses.

To assess whether neurons at different cortical depths exhibit differences in movement encoding during reach-to-grasp, we split each field of view into a superficial and deep population, putatively corresponding to

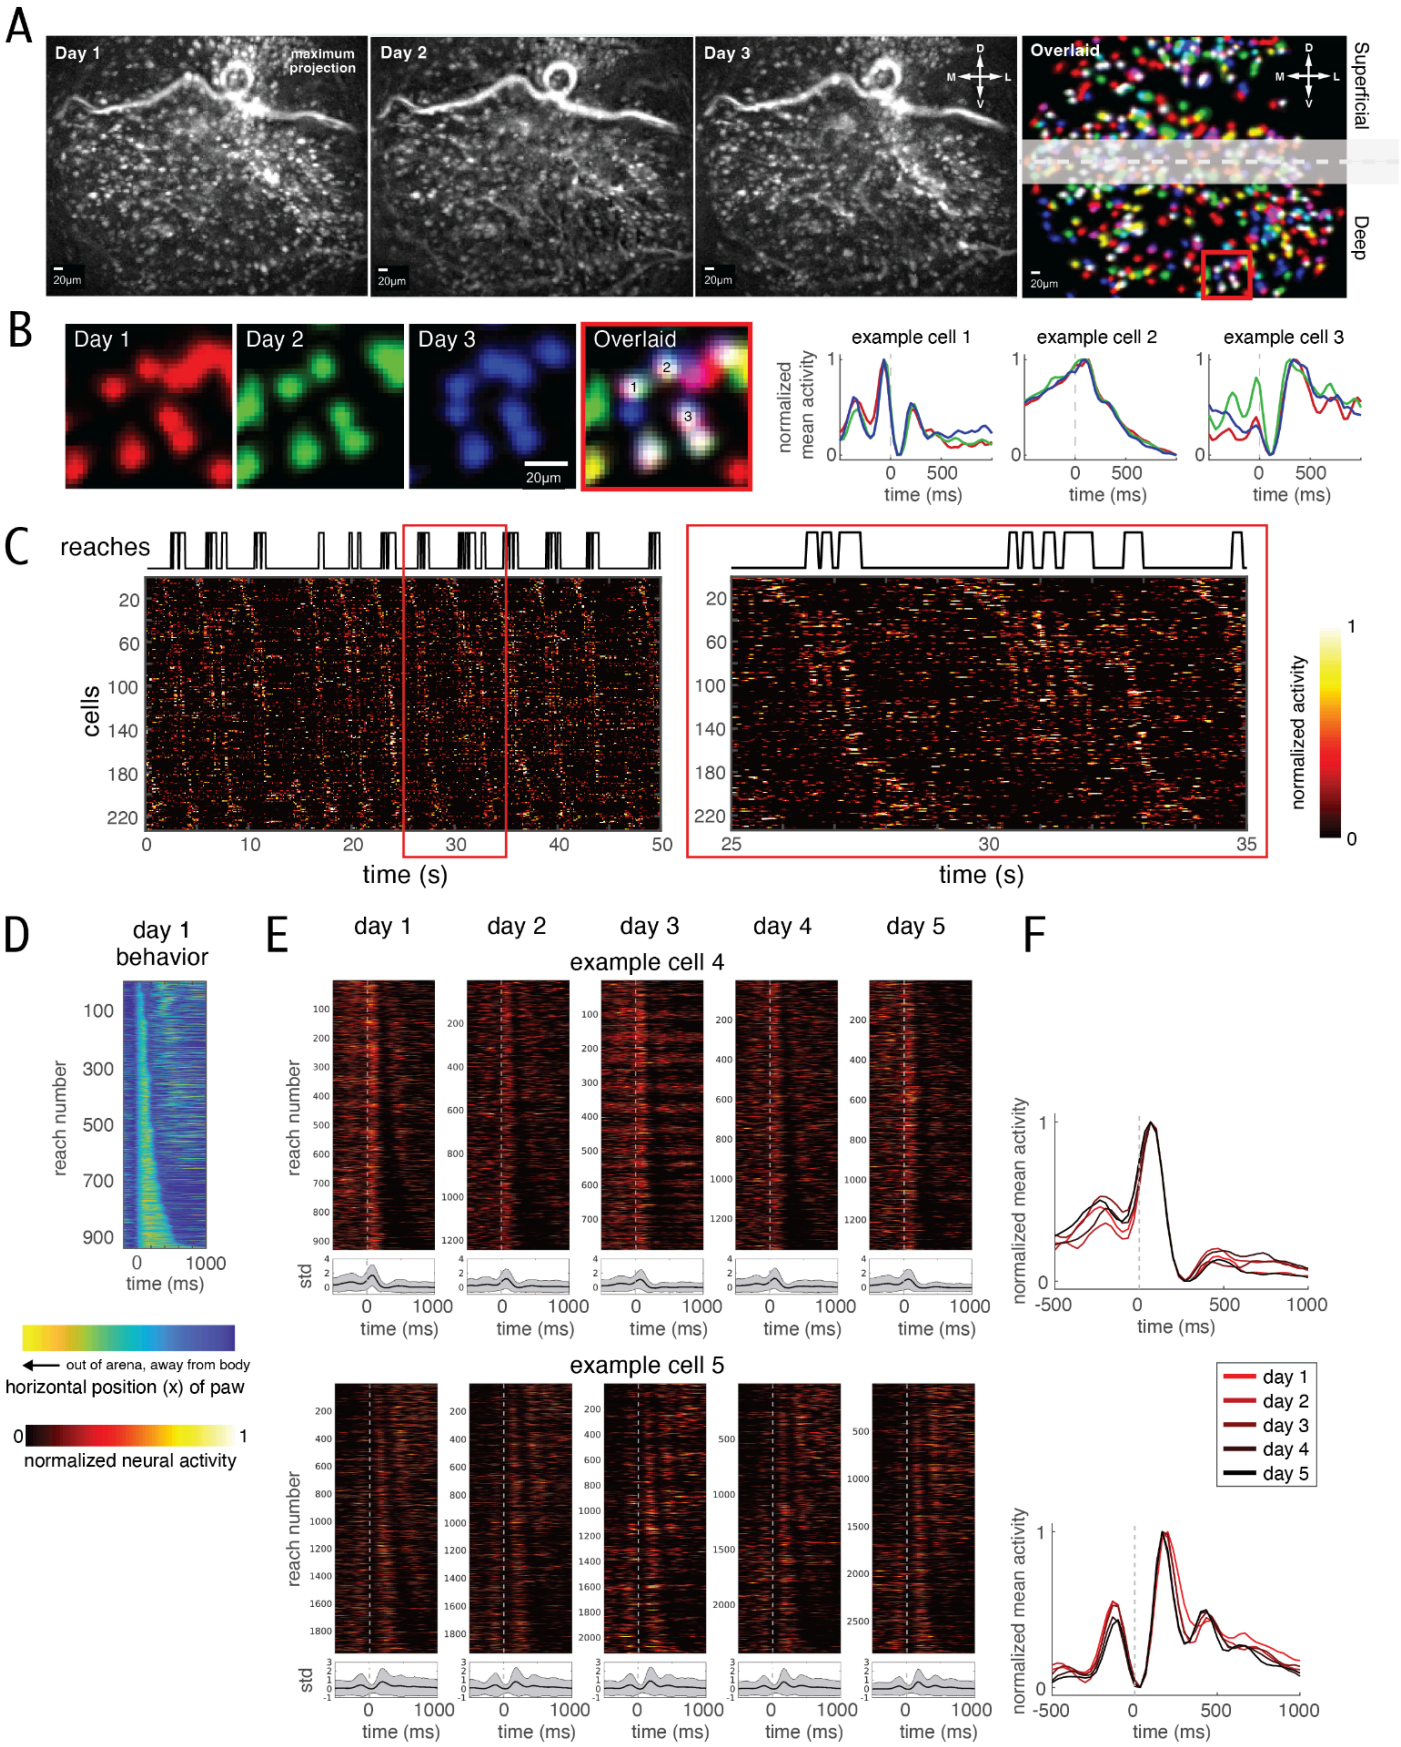
A
Day 1
maximum projection
Day 2
Day 3
Overlaid
Superficial
Deep
20μm
B
Day 1
Day 2
Day 3
Overlaid
example cell 1
example cell 2
example cell 3
normalized mean activity
time (ms)
C
reaches
cells
time (s)
normalized activity
D
day 1 behavior
reach number
time (ms)
out of arena, away from body
horizontal position (x) of paw
normalized neural activity
E
day 1
day 2
day 3
day 4
day 5
example cell 4
example cell 5
reach number
std
time (ms)
F
normalized mean activity
time (ms)
day 1
day 2
day 3
day 4
day 5

Figure 3: Neurons in CFA M1 motor cortex modulate during reaching

A) Far left, middle left, middle right: Maximum projections of the motion-corrected calcium imaging movie of the same field of view (mouse 4) on three consecutive days. Bright spots indicate cells; blood vessel visible at top. Compass indicates dorsal-ventral and medial-lateral. Far right: overlay of footprints of all individual cells identified on all three days, with day 1 in the red channel, day 2 green, and day 3 blue. White footprints indicate cells found on all three days.

B) Left: Zoomed in view of red square outline in A far right, separated by days and overlaid. Right: peri-event time histograms (PETHs) locked to the onset of reach (see text) for the three days shown, for three example cells from the zoomed view. PETHs are normalized for each day separately so that minimum value is 0, maximum value is 1. Colors indicate day.

C) Top: Binary variable indicating when the mouse was reaching. Bottom: Heatmap of inferred spiking activity of all recorded cells from one field of view for 50 seconds of recording. Right inset: Temporal zoom to 10 seconds of recording, corresponding to the red box in the full view. Each cell's activity has been normalized separately within this window: minimum value = 0, mean + 5*standard deviation = 1.

D) Kinematic summary for comparing with single trial neural activity in E and F, copied from Fig. 1E bottom left. Paw position in the x dimension plotted for all reaches in one daily session for an example mouse (mouse 5). Colormap is under the schematic in A; yellow colors indicate further outside the arena, blue colors indicate closer to or inside the arena. Reaches sorted according to reach duration. t = 0 indicates reach onset.

E) Single-trial inferred spiking activity of two example cells found on all five days, aligned to the start of reach. For each day, top panel shows a heatmap of the single trial activity. Bottom panel shows the average response across all trials on that day (PETH). Heatmaps are normalized within each day as 0 = mean activity, 1 = 10*standard deviation above the mean. Trials are sorted according to the duration of reaches, as for the kinematics in D: short reaches at the top, long reaches at the bottom. PETHs are plotted as mean +/- st dev, units refer to standard deviations above or below the mean.

F) PETHs for example cells in E, from all five days. PETHs are normalized for each day separately so that minimum value is 0, maximum value is 1. Colors indicate day. t = 0 indicates reach onset.

L2/3 and L5a (Methods) and compared the results of all encoding analyses described in later sections for these two populations. Although the differences observed were often consistent across analyses within a given mouse, they were not consistent across different mice (Supplementary Fig. 6).

Reach-locked average neuronal activity patterns (PETHs) are stable across days. We quantified the average response of each neuron during reaching using peri-event time histograms (PETHs) time-locked to the start of the reach. The diverse response timecourses seen in individual trials (Fig. 3C) were also present in PETHs (Fig. 4A). We found that, across the population, PETHs were highly similar across days. To quantify this similarity, we computed the correlation between each longitudinally tracked neuron's PETH on different days (Methods). For each mouse, we examined the distribution of these correlation values as a function of the interval between days and compared it to a control distribution of correlations between PETHs for non-same cells (Fig. 4B; Methods). We observed that the distribution of correlation values for each interval was clustered near one and was far from the median of the control distribution, which was close to zero (Fig. 4B,C). In other words, each neuron's activity was more similar to its own activity across days than to that of other neurons. This indicates that the time course of the average responses during and around reaching remained stable in this population over a five-day period.

Trial-by-trial variance in forelimb, digit, and head movements is captured in single-trial neuronal activity in CFA M1. Paw and head movements during reaching exhibit variability across multiple kinematic dimensions including speed, duration, timing of wrist pronation, and timing of digits opening and closing (Fig. 1 & 2). While PETHs quantify a neuron's response for the average movement, they do not reveal whether a neuron encodes the significant trial-by-trial variance in this behavior. To assess the neural encoding of the moment-to-moment details of reach-to-grasp kinematics, we constructed families of linear models to predict the activity of individual neurons based on the extensive set of kinematic variables we extracted using

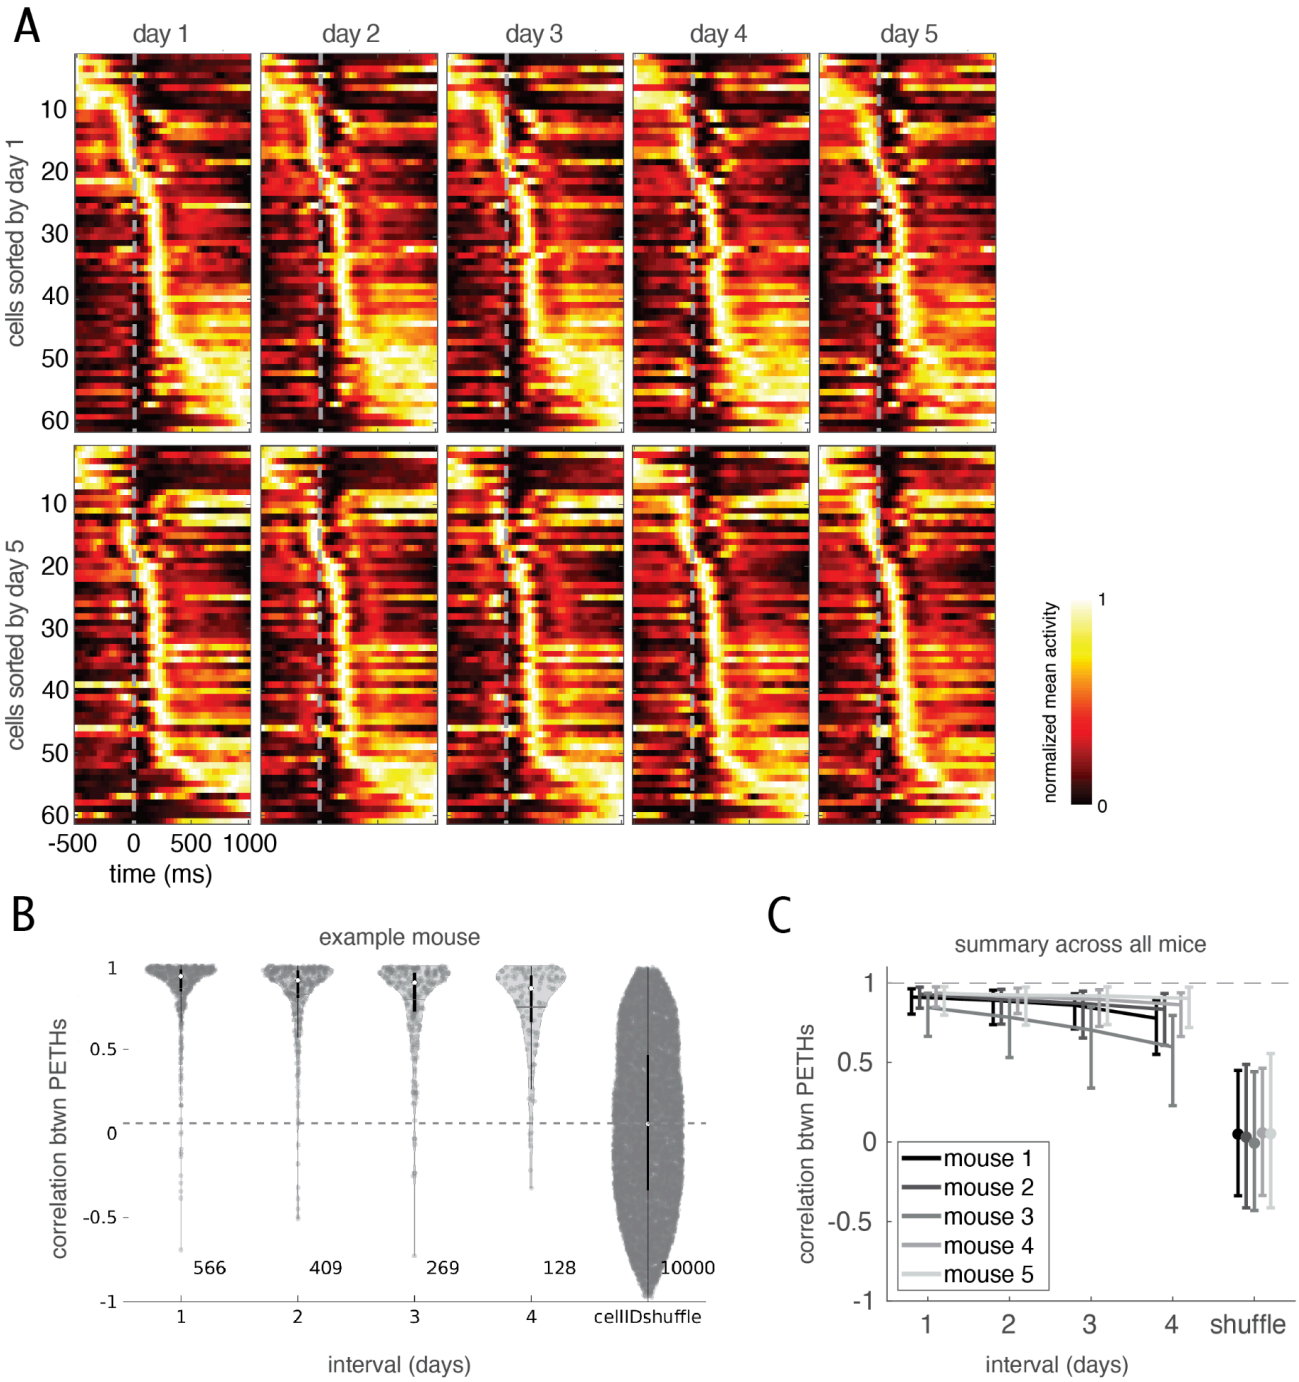

Figure 4: Reach-locked average neuronal activity patterns (PETHs) are stable across days

A) Normalized peri-event time histograms (PETHs) of neurons found on all five days for an example mouse (mouse 4). Cells are sorted according to the time of their peak activity on either day 1 (top row) or day 5 (bottom row). Each cell's PETH has been normalized separately within each day within the plotted window: minimum value = 0, maximum value = 1. t = 0 indicates reach onset.

B) Correlation between PETHs for all cells found on more than one day, as a function of the interval between the days, for an example mouse (mouse 4). Each datapoint is an observation of the same cell on two different days. Numbers to the bottom right of the violins indicate the number of observations included in that violin. For each violin, the white dot indicates the median of the distribution and the black box indicates interquartile range. Cell ID shuffle indicates the distribution of correlations between PETHs of two cells randomly chosen from any day (excluding any cell pairs that are registered as the same cell; see Methods). Gray dashed line indicates the median of this shuffle.

C) Correlation between PETHs for all mice. Datapoint is median, error bars are interquartile range across cells for that interval for that mouse. Grayscale indicates mouse. Shuffle as in B.

keypoint tracking (schematized in Fig. 5A; Methods).

Notably, the mean movement could be encoded even if trial-by-trial variability around the mean is not. This would suggest that motor cortex carries a general 'reach' signal—indicating whether the mouse is performing the behavior—without capturing the moment-to-moment details of the movements involved. While individual neurons exhibit significant trial-by-trial variance around their average responses (Fig. 3D), it is possible that that variance does not co-vary with the variance observed in movement. Additionally, even if movement details are indeed encoded in motor cortical activity on a given day, it is also possible that the nature of that encoding is not the same across days, indicating instability (drift).

Before fitting the models, we applied a covariate balancing subsampling procedure to the kinematic predictor sets for all five days, separately for each mouse, to ensure that we are comparing encoding of similar movements across days (Supplementary Fig. 3; Methods). This step is necessary for later analysis of encoding stability; without it, we would not be able to differentiate between true drift and changes in encoding caused by changes in behavior.

The majority of neurons on any given day have a modest but statistically significant amount of variance explained when predicted by the entire set of kinematic predictors, as compared to a timepoint-shuffled control (full single-trial model schematized in Fig. 5A). Specifically, 4866 out of 6105 cell-days (79.7%) showed significant variance explained by the full single-trial model, where each cell-day represents a single day's observation for a given neuron ($p = 0.05$, with Bonferroni correction for the number of cells in that dataset (field of view on that day); Methods). On average, this comprises 79.6% of neurons in each dataset (79.6% +/- 18.9% [38.8%, 96.0%]; mean +/- st.dev. [range], n=25 datasets; Supp Fig. 4B,C).

We have already observed that individual cells exhibit time-varying average responses (PETHs in Fig. 4), which can lead to significant predictions in our timepoint-based linear models even in the absence of coding for variance around the mean. In other words, the time-varying average movement could explain a cell's activity above and beyond the timepoint shuffle control, even if the cell's single-trial variance does not capture single-trial variance in the movements. To test for this possibility, we compare performance of the full single-trial model to that of a trial-average model, in which we predict the cell's single-trial activity from trial-averaged movements (schematized in Fig. 5B). In the predictor matrix of this trial-average model, we replace the timepoints of each reach with the reach-locked average, for each kinematic predictor separately (Methods). The majority of recorded neurons were significantly predicted by a trial-average model (4301 out of 6105 cell-days, or 70.5%; Methods), corresponding to 70.3% of each dataset (70.3% +/- 15.3% [30.2%, 93.2%]; mean +/- st.dev. [range]). Almost all of these were also significantly predicted by the full single-trial model (3924 out of 4301 cell-days, or 91.2%). Few neurons were significantly predicted by only the trial-average model, indicating they encode only a general reach signal (377 out of 6105 cell-days, or 6.2%), or by only the full single-trial model, indicating they encode only moment-to-moment details (942 out of 6105 cell-days, or 15.4%).

Of the neurons significantly predicted by the full single-trial model, we found that the vast majority were better modeled by the full single-trial model than by the trial-average model (Fig. 5C-E; 4310 out of 4866 (88.6%) cell-days, 88.3% +/- 5.0% of each dataset [74.1%, 96.4%]; mean +/- st.dev. [range]). This finding suggests that these neurons encode single-trial variance above and beyond what they encode about the average reach. Together, these results indicate that mouse motor cortex encodes not only a general "reaching" signal but also moment-to-moment details of movement. For the remainder of our analysis, we focus on the subset of neurons with significant full single-trial model predictions.

Next we consider the contribution of individual predictors. Due to the nature of the task and the physics of the mice's bodies, many kinematic variables are correlated (Supplementary Fig. 4A). Thus, the coefficients

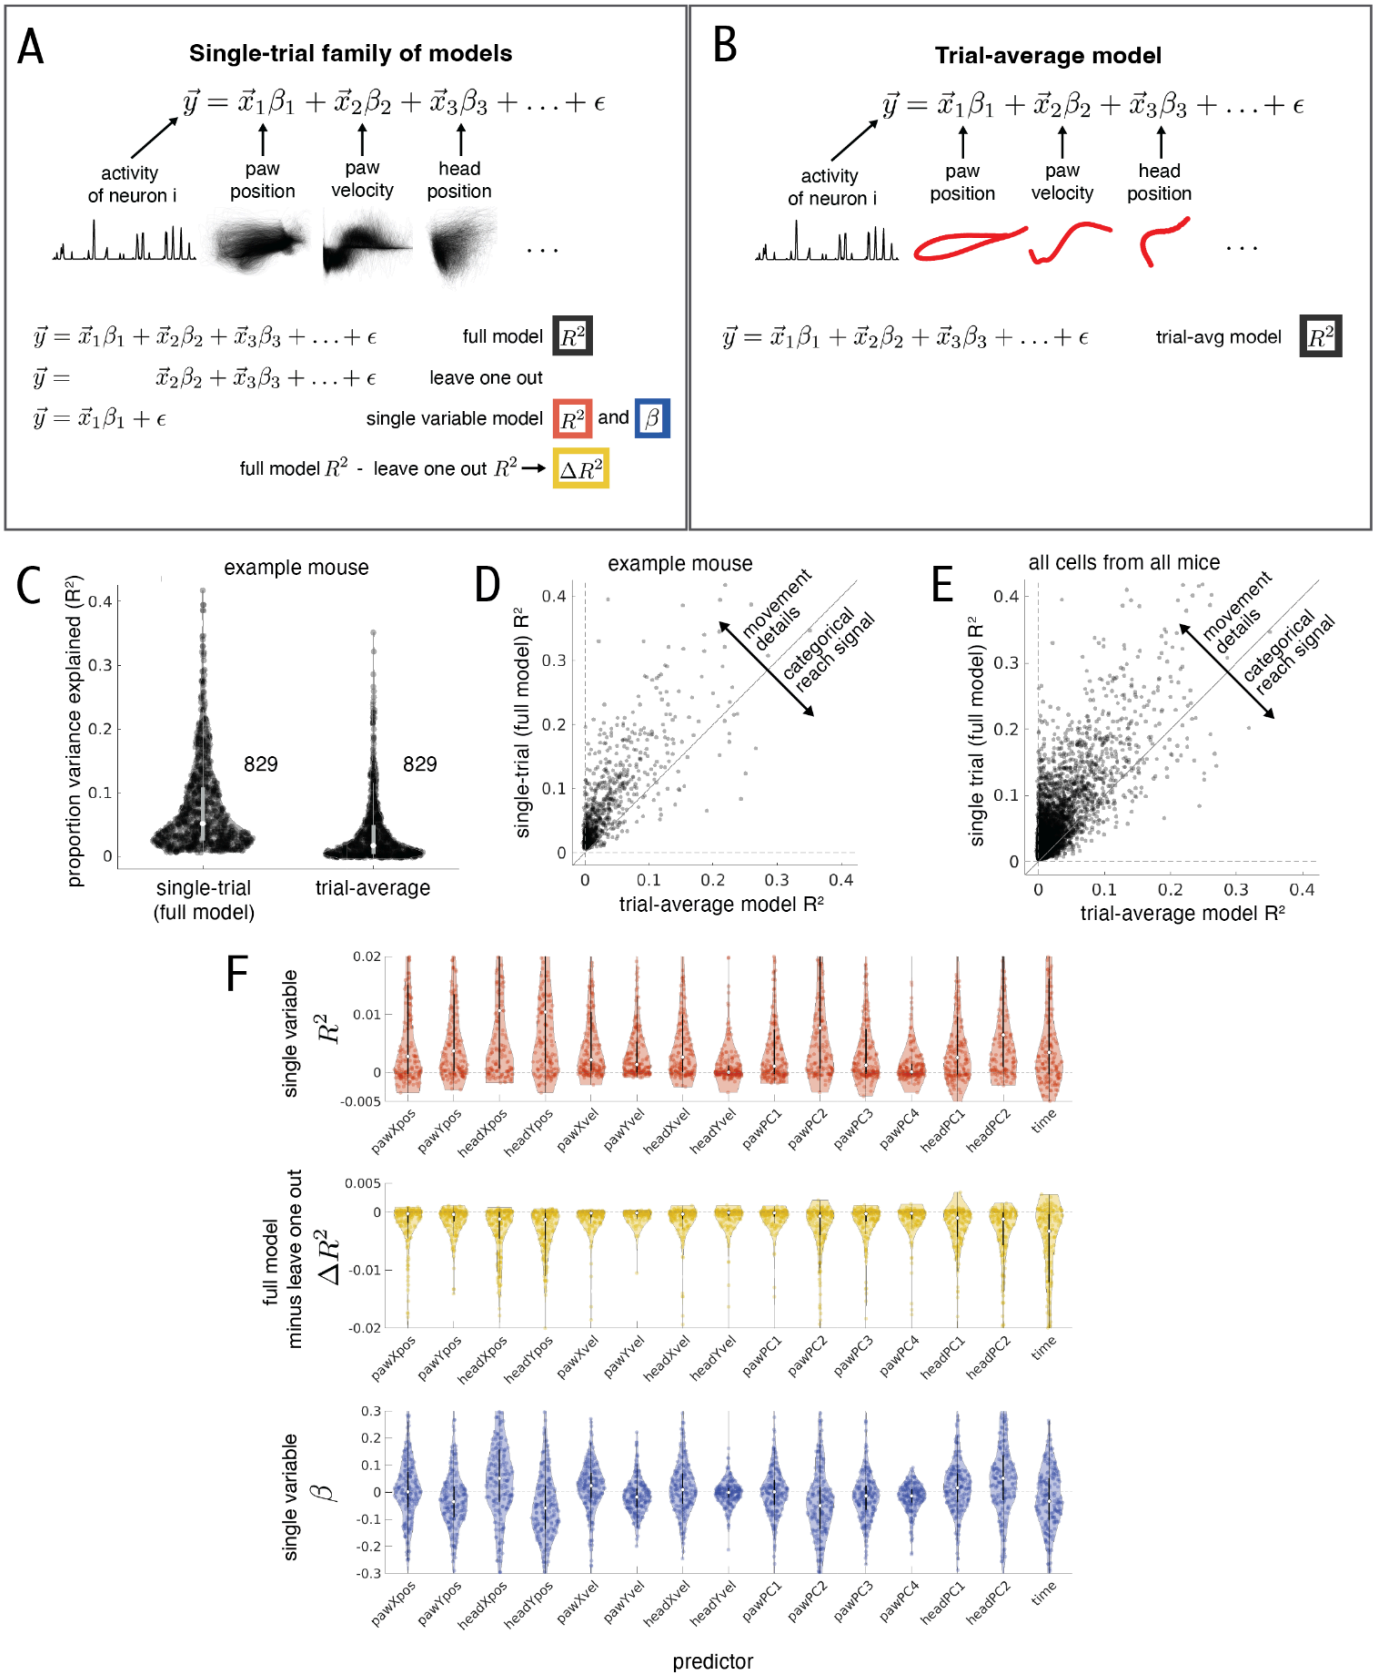
A
Single-trial family of models
$\vec{y} = \vec{x}_1\beta_1 + \vec{x}_2\beta_2 + \vec{x}_3\beta_3 + \ldots + \epsilon$
activity of neuron i
paw position
paw velocity
head position
$\vec{y} = \vec{x}_1\beta_1 + \vec{x}_2\beta_2 + \vec{x}_3\beta_3 + \ldots + \epsilon$ full model R^2
$\vec{y} = \vec{x}_2\beta_2 + \vec{x}_3\beta_3 + \ldots + \epsilon$ leave one out
$\vec{y} = \vec{x}_1\beta_1 + \epsilon$ single variable model R^2 and β
full model R^2 - leave one out R^2 → ΔR^2
B
Trial-average model
$\vec{y} = \vec{x}_1\beta_1 + \vec{x}_2\beta_2 + \vec{x}_3\beta_3 + \ldots + \epsilon$
activity of neuron i
paw position
paw velocity
head position
$\vec{y} = \vec{x}_1\beta_1 + \vec{x}_2\beta_2 + \vec{x}_3\beta_3 + \ldots + \epsilon$ trial-avg model R^2
C
example mouse
proportion variance explained (R²)
829
829
single-trial (full model)
trial-average
D
example mouse
single-trial (full model) R²
movement details
categorical reach signal
trial-average model R²
E
all cells from all mice
single trial (full model) R²
movement details
categorical reach signal
trial-average model R²
F
single variable R^2
full model minus leave one out ΔR^2
single variable β
pawXpos
pawYpos
headXpos
headYpos
pawXvel
pawYvel
headXvel
headYvel
pawPC1
pawPC2
pawPC3
pawPC4
headPC1
headPC2
time
predictor

Figure 5: Trial-by-trial variance in forelimb, digit, and head movements allows for accurate modeling of single-trial CFA M1 neuronal activity

A) Schematic of trial-to-trial linear model families built to predict the activity of individual cells from paw, digit, and head movements during reaching. The full single-trial model contains all kinematic variables and their observed values across trials. For each predictor, we build a leave-one-out model using all predictors except the one in question, and a single variable model, using only that predictor. The change in variance explained from the full model to the leave one out model indicates unique variance explained by that predictor.

B) Schematic of the trial-average linear model, where each trial of the kinematics is replaced by the average timecourse for that kinematic predictor, averaged across reaches, locked to reach onset (see Methods).

C) Proportion of variance explained (R-squared) by the full single-trial model or the trial average model for all cell-days with a significant full single-trial model prediction, for an example mouse (mouse 5). Numbers to the right of each violin indicate the number of cell-days in the distribution. Each datapoint corresponds to a cell on one day (cell-day). Cells registered across days have multiple dots if they have significant full single-trial model predictions on multiple days. For each violin, the white dot indicates the median of the distribution and the gray box indicates interquartile range.

D) Same data as in C, but with the full model and trial average model R-squared values plotted against each other for each cell-day.

E) Same as D, but for all cell-days with significant full model predictions from all mice.

F) Model features from the single-trial family of models for all cells from one example day from one mouse. Top: variance explained (R-squared) for single variable model fits. Each datapoint corresponds to a cell; each cell has one datapoint in each violin. Middle: same as top but for the Delta-R-squared. Bottom: same as top but for the coefficients from single variable model fits.

in our full single-trial model should not be interpreted as the maximum contribution of each individual kinematic predictor (Stevenson et al., 2018). To assess the contribution of individual predictors to explaining a neuron's activity, we compute two metrics: (schematized in Fig. 5A): 1) Single-variable R-squared values: For each predictor, we fit a model using only that predictor to estimate the total variance it explains. 2) Delta-R-squared values: Starting with the full single-trial model, we create a leave-one-out model for each predictor by excluding it and refitting the model. The resulting drop in R-squared indicates the unique variance that predictor contributes. Additionally, to capture how a neuron's activity is modulated relative to variations in each movement predictor, we also consider coefficients (betas) from the single variable models, which serve as a form of tuning curve.

Many predictors exhibit non-zero single-variable R-squared values and delta-R-squared values (Fig. 5F top, middle), indicating that these movement variables are encoded by single cells in motor cortex. Additionally, distributions of single variable model coefficients are dispersed away from zero (Fig. 5F bottom). Notably, despite partial correlation with paw-related predictors (Supplementary Fig. 4A), several head-related predictors have non-zero delta R-squared values for many neurons, indicating that they contribute unique explanatory power. Thus, the moment-to-moment details of the reaching movements encoded in caudal forelimb area of motor cortex include not only the position of the paw but its shape and orientation (posture of the digits and wrist) and the position and orientation of the head.

Linear model encoding profiles of single cells are stable across days. To assess whether the encoding of trial-by-trial variability is consistent over time, we tested whether the linear models generalized to days other than the ones they were trained on (Fig. 6A; Methods). Specifically, we applied single-trial full models forward and backwards, restricting our analysis to cells that showed statistically significant movement encoding—that is, a significant single-trial full model prediction. For cells with significant fits on some days but not others, we only compared days with significant fits, which constitutes 79.7% of all cell-days (see previous section). Model performance declined only slightly as the interval between training and testing days increased (Fig. 6B,C). Of the 4,969 cell-day pairs examined, 4,586 (92.3%) showed significantly similar

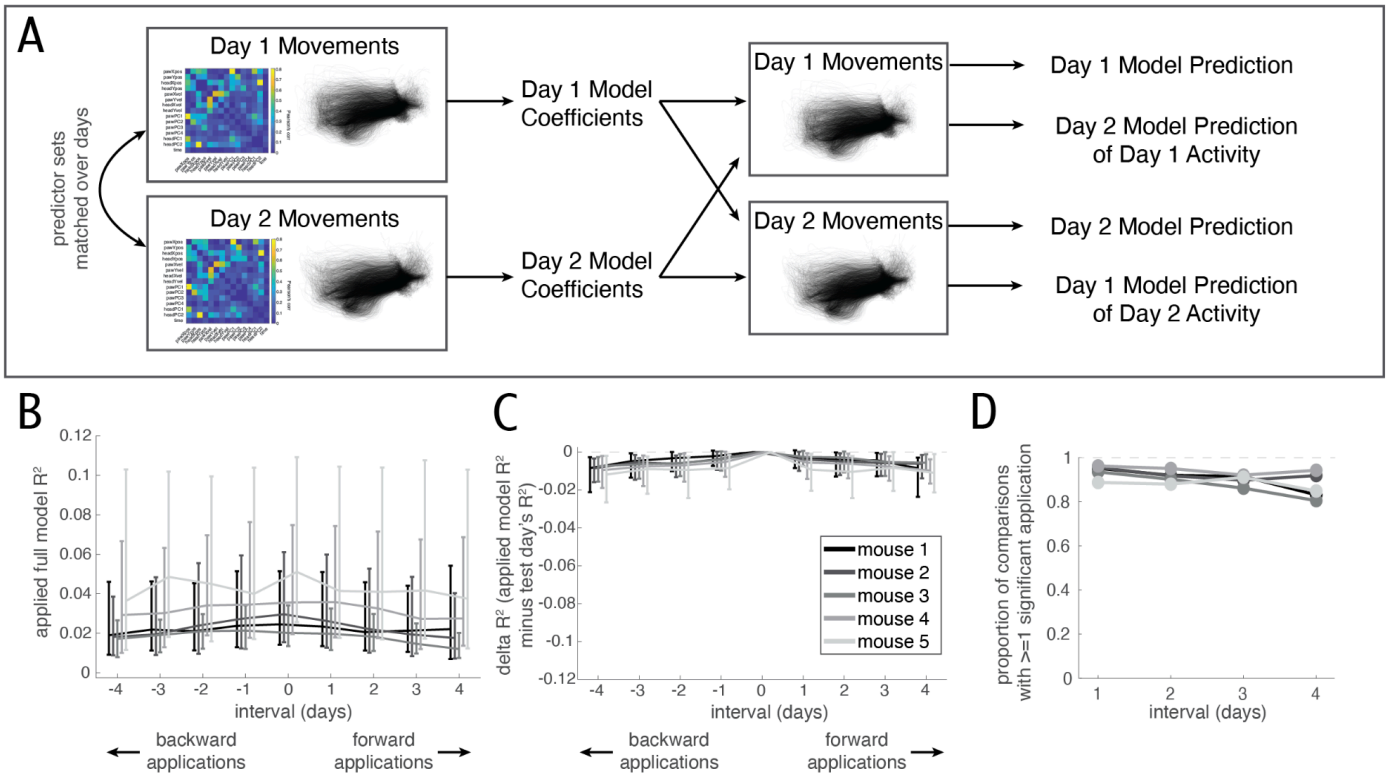

Figure 6: Linear encoding models predict single neuron activity across days

A) Schematic of applying linear models across days: The full single-trial linear model built for a given cell on a given day is applied to other days for that cell. Kinematic variable predictor sets are matched before fitting models (see text and Methods). Only cell-days with significant single-trial full model predictions are included in across-day linear model analysis.

B) Resulting R-squared value when a model fitted on one day is applied to another day, for all intervals between train and test days. Datapoint is median, error bars are interquartile range across train-test day pairs for that interval for that mouse. Grayscale indicates mice, as in C. Forward applications indicate the train day comes before the test day; backward applications indicate the train day comes after the test day.

C) Performance of applied models relative to a test day's own model. Difference between R-squared from an applied model and the R-squared of the test day's own fitted model. Datapoint is median, error bars are interquartile range across train-test day pairs for that interval for that mouse. Forward and backward applications as in B.

D) Significance of applied model predictions. Proportion of cell-day pairs with significant model application in at least one direction, as a function of interval between the cell-days. Grayscale indicates mice, as in C. Applied models are compared to test day's own timepoint shuffle control to determine significance (see Methods).

encoding, defined as at least one significant forward or backward model application, with significance assessed relative to a time-shuffled control for the test day (Methods). The vast majority of these—4,258 pairs (92.8% of significant pairs, or 85.7% of all pairs)—exhibited significant encoding in both directions. These findings demonstrate that the encoding of moment-to-moment movement details in motor cortical activity remains largely stable at the single-neuron level across multiple days.

Discussion

Here, we investigated motor cortical encoding for movements of the paw, digits, and head during a self-paced, skilled reach-to-grasp task and assessed how stable this encoding is over days. By recording motor cortex activity during the performance of a challenging motor skill characterized by many repetitions within a day

and high reach-to-reach variability—features that make the skill more reflective of real-world motor behaviors—we find that neurons in the mouse motor cortex encode both general reaching signals and the detailed movements of the head, forelimb, and digits during reach-to-grasp behaviors. This encoding remains remarkably stable over days, providing a reliable neural foundation for producing flexible skilled movements consistently over time.

Quantifying neural encoding during naturalistic behaviors with high trial-to-trial variability poses significant challenges, which we overcame through a series of carefully designed approaches. First, we used encoding models that do not assume stereotypy in movements across trials. This analytical flexibility allowed us to account for the inherent variability of naturalistic behaviors. Second, we designed the task to balance variability with sufficient trial counts. Key elements of the task, such as requiring the mice to reposition themselves after each bout of reaches, randomizing seed positions across multiple locations, and omitting penalties for failed attempts, encouraged trial-to-trial variability while discouraging the development of stereotyped, automatized movement patterns. The task's inherent difficulty likely further contributed to the diversity in movements. Third, we captured these movements in fine-grained detail using high-speed, high-resolution cameras, allowing us to precisely quantify trial-to-trial variability above the noise threshold. High-resolution behavioral data is crucial to accurately identify variations in movement encoding. Finally, to ensure that our stability analysis was not confounded by behavioral shifts over days, we implemented a covariate balancing procedure before fitting linear models. This approach allowed us to quantify encoding within a consistent region of the high-dimensional kinematic parameter set each day, effectively removing neural changes resulting from behavioral variability. Together, these strategies enabled us to rigorously estimate and compare neural encoding across days in a statistically robust manner. Our findings not only highlight the stability of motor cortical encoding during complex, variable behaviors but also underscore the value of incorporating detailed behavioral analyses into neural encoding studies (Sadeh and Clopath, 2022; Musall, Kaufman et al., 2019).

We used single-trial encoding models to investigate whether the significant trial-to-trial variability in movements is encoded in motor cortex and whether this encoding remains stable over days. While more complex models could likely explain additional variance, they often come at the cost of interpretability. In contrast, our linear approach allowed us to directly link the activity of individual neurons to a carefully curated set of 14 kinematic variables, capturing the movements of specific body parts (e.g., position and velocity of the paw and head, as well as conformation of the digits and orientation of the paw and head). We found that moment-to-moment movement details during reach-to-grasp explained a statistically significant fraction of variance in most recorded neurons. Furthermore, these relationships remained stable across multiple days, representing the longest duration over which such detailed single-cell encoding of movement has been examined. Collectively, the stability of PETHs and linear encoding models strongly suggest that stable behaviors arise from the consistent contributions of individual neurons within the motor circuit.

The cell populations studied here include the upper layers of CFA, putatively including L2/3 and L5a, which have notable differences in connectivity that could potentially lead to differences in encoding stability (Oswald et al., 2013; Anderson et al., 2010; Weiler et al., 2008; Yamawaki et al., 2021; Hira et al., 2013; Shepherd, 2009). However, our comparison of superficial and deep populations, as we defined them here, is inconclusive. It is possible that the variability across mice obscured any differences, or that the strong interconnectedness of these layers causes them to share many encoding properties.

The finding that motor cortex encodes single-trial variability in reach-to-grasp movements suggests a role in the monitoring and possibly online control of this sensory-driven behavior, likely operating as part of broader sensorimotor loops involving other cortical and subcortical regions such as somatosensory cortex,

thalamus, brainstem, and cerebellum (Yamawaki et al., 2021; Sauerbrei et al., 2020; Wagner et al., 2019; Becker and Person, 2019). Further work is needed to determine whether this encoding reflects movement control signals for the ongoing reach, as observed in cerebellum (Becker and Person, 2019) and/or corollary discharge, useful for longer-term processes such as reach-to-reach adjustment (Becker, Calame, et al., 2020; Dhawale et al., 2019) or continual learning as task demands change (Dhawale et al., 2019; Kawai et al., 2015; Hwang et al., 2019). The linear models showed that principal components of paw markers, which reflect wrist and digit states, contributed explanatory power comparable to paw position, which reflects the state of proximal joints such as shoulder and elbow, implicating CFA not just in precision reaching but also in grasping (Wang et al., 2017). Interestingly, position and orientation of the head also contributed explanatory power comparable to paw- and digit-related predictors. This could reflect two mutually compatible explanations: first, head-related variables may be correlated with omitted kinematic variables that are more directly linked to neural activity (Stevenson et al., 2018). Second, head position and orientation may themselves be directly relevant to computations in motor cortex, providing crucial information about body positioning relative to the reach port and pedestal. This latter explanation implies a role for egocentric planning signals in guiding the paw to the pedestal, as observed in encoding of whole body posture in rats (Mimica et al., 2018). These findings highlight the importance of examining how multiple body parts are encoded in a given brain area, rather than limiting the focus to movements traditionally associated with that area. A comprehensive understanding of whole-body movement generation requires this broader perspective.

These findings extend our understanding of how the motor nervous system maintains the ability to produce learned, goal-directed movements. While previous studies have demonstrated that the brain can maintain stable neural activity patterns associated with stereotyped behaviors, i.e. a motor tape (Jensen et al., 2022; Katlowitz et al., 2018), here we demonstrate that it also maintains stable representations of a more flexible behavior. To allow for adaptive control, both online (e.g. within a reach) and offline (e.g. adjustments from one reach to the next), the nervous system must process moment-to-moment details of the constituent movements of motor skills. Our findings indicate that the motor nervous system supports flexible motor skills via stable representations of these details, allowing for robust but adaptive performance in dynamic contexts over extended periods of time.

To uncover how the brain produces the flexible, context-dependent, and coordinated movements essential for daily human activities, it is crucial to study comparable behaviors in animal models. The freely-moving reach-to-grasp behavior analyzed here exemplifies such complexity, combining whole-body posture with precise forelimb and digit control. This task closely parallels primate reaching and grasping behaviors (Sacrey et al., 2009; Whishaw et al., 1992; Guo et al., 2015; Becker, Calame, et al., 2020; Whishaw et al., 2017) and serves as a powerful model for investigating the neural mechanisms underlying these intricate actions.

References

An, X., Matho, K., Li, Y., Mohan, H., Xu, X. H., Whishaw, I. Q., Kepecs, A., & Huang, Z. J. (2024). A cortical circuit for orchestrating oromanual food manipulation. bioRxiv. https://doi.org/10.1101/2022.12.03.518964

Andermann, M. L., Gilfoy, N. B., Goldey, G. J., Sachdev, R. N. S., Wölfel, M., McCormick, D. A., Reid, R. C., & Levene, M. J. (2013). Chronic Cellular Imaging of Entire Cortical Columns in Awake Mice Using Microprisms. *Neuron*, *80*(4), 900–913. https://doi.org/10.1016/j.neuron.2013.07.052

Anderson, C. T., Sheets, P. L., Kiritani, T., & Shepherd, G. M. G. (2010). Sublayer-specific microcircuits of corticospinal and corticostriatal neurons in motor cortex. *Nature Neuroscience*, *13*(6), 739–744. https://doi.org/10.1038/nn.2538

Barrett, J. M., Raineri Tapies, M. G., & Shepherd, G. M. G. (2020). Manual dexterity of mice during food-handling involves the thumb and a set of fast basic movements. *PLOS ONE*, *15*(1), e0226774. https://doi.org/10.1371/journal.pone.0226774

Becker, M. I., Calame, D. J., Wrobel, J., & Person, A. L. (2020). Online control of reach accuracy in mice. *Journal of Neurophysiology*, *124*(6), 1637–1655. https://doi.org/10.1152/jn.00324.2020

Becker, M. I., & Person, A. L. (2019). Cerebellar Control of Reach Kinematics for Endpoint Precision. *Neuron*, *103*(2), 335-348.e5. https://doi.org/10.1016/j.neuron.2019.05.007

Billard, A., & Kragic, D. (2019). Trends and challenges in robot manipulation. *Science*, *364*(6446), eaat8414. https://doi.org/10.1126/science.aat8414

Chambers, A. R., & Rumpel, S. (2017). A stable brain from unstable components: Emerging concepts and implications for neural computation. *Neuroscience*, *357*, 172–184. https://doi.org/10.1016/j.neuroscience.2017.06.005

Chestek, C. A., Batista, A. P., Santhanam, G., Yu, B. M., Afshar, A., Cunningham, J. P., Gilja, V., Ryu, S. I., Churchland, M. M., & Shenoy, K. V. (2007). Single-Neuron Stability during Repeated Reaching in Macaque Premotor Cortex. *Journal of Neuroscience*, *27*(40), 10742–10750. https://doi.org/10.1523/JNEUROSCI.0959-07.2007

Clopath, C., Bonhoeffer, T., Hübener, M., & Rose, T. (2017). Variance and invariance of neuronal long-term representations. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1715), 20160161. https://doi.org/10.1098/rstb.2016.0161

Dhawale, A. K., Miyamoto, Y. R., Smith, M. A., & Ölveczky, B. P. (2019). Adaptive Regulation of Motor Variability. *Current Biology*. https://doi.org/10.1016/j.cub.2019.08.052

Dhawale, A. K., Poddar, R., Wolff, S. B., Normand, V. A., Kopelowitz, E., & Ölveczky, B. P. (2017). Automated long-term recording and analysis of neural activity in behaving animals. *ELife*, *6*, e27702. https://doi.org/10.7554/eLife.27702

Disse, G. D., Nandakumar, B., Pauzin, F. P., Blumenthal, G. H., Kong, Z., Ditterich, J., & Moxon, K. A. (2023). Neural ensemble dynamics in trunk and hindlimb sensorimotor cortex encode for the control of postural stability. *Cell Reports*, *42*(4), 112347. https://doi.org/10.1016/j.celrep.2023.112347

Driscoll, L. N., Duncker, L., & Harvey, C. D. (2022). Representational drift: Emerging theories for continual learning and experimental future directions. *Current Opinion in Neurobiology*, *76*, 102609. https://doi.org/10.1016/j.conb.2022.102609

Flint, R. D., Scheid, M. R., Wright, Z. A., Solla, S. A., & Slutzky, M. W. (2016). Long-Term Stability of Motor Cortical Activity: Implications for Brain Machine Interfaces and Optimal Feedback Control. *Journal of Neuroscience*, *36*(12), 3623–3632. https://doi.org/10.1523/JNEUROSCI.2339-15.2016

Fraser, G. W., & Schwartz, A. B. (2012). Recording from the same neurons chronically in motor cortex. *Journal of Neurophysiology*, *107*(7), 1970–1978. https://doi.org/10.1152/jn.01012.2010

Galiñanes, G. L., Bonardi, C., & Huber, D. (2018). Directional Reaching for Water as a Cortex-Dependent Behavioral Framework for Mice. *Cell Reports*, *22*(10), 2767–2783. https://doi.org/10.1016/j.celrep.2018.02.042

Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A., & Miller, L. E. (2020). Long-term stability of cortical population dynamics underlying consistent behavior. *Nature Neuroscience*. https://doi.org/10.1038/s41593-019-0555-4

Grier, H. A., Salimian, S., & Kaufman, M. T. (2025). Mouse sensorimotor cortex reflects complex kinematic details during reaching and grasping. bioRxiv. https://doi.org/10.1101/2025.02.19.639049

Guo, J.-Z., Graves, A. R., Guo, W. W., Zheng, J., Lee, A., Rodríguez-González, J., Li, N., Macklin, J. J., Phillips, J. W., Mensh, B. D., Branson, K., & Hantman, A. W. (2015). Cortex commands the performance of skilled movement. *ELife*, *4*. https://doi.org/10.7554/eLife.10774

Hira, R., Ohkubo, F., Tanaka, Y. R., Masamizu, Y., Augustine, G. J., Kasai, H., & Matsuzaki, M. (2013). In vivo optogenetic tracing of functional corticocortical connections between motor forelimb areas. *Frontiers in Neural Circuits*, *7*, 55. https://doi.org/10.3389/fncir.2013.00055

Hwang, E. J., Dahlen, J. E., Hu, Y. Y., Aguilar, K., Yu, B., Mukundan, M., Mitani, A., & Komiyama, T. (2019). Disengagement of motor cortex from movement control during long-term learning. *Science Advances*, *5*(10), eaay0001. https://doi.org/10.1126/sciadv.aay0001

Jensen, K. T., Kadmon Harpaz, N., Dhawale, A. K., Wolff, S. B. E., & Ölveczky, B. P. (2022). Long-term stability of single neuron activity in the motor system. *Nature Neuroscience*, *25*(12), 1664–1674. https://doi.org/10.1038/s41593-022-01194-3

Katlowitz, K. A., Picardo, M. A., & Long, M. A. (2018). Stable Sequential Activity Underlying the Maintenance of a Precisely Executed Skilled Behavior. *Neuron*, *98*(6), 1133-1140.e3. https://doi.org/10.1016/j.neuron.2018.05.017

Kawai, R., Markman, T., Poddar, R., Ko, R., Fantana, A. L., Dhawale, A. K., Kampff, A. R., & Ölveczky, B. P. (2015). Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron*, *86*(3), 800–812. https://doi.org/10.1016/j.neuron.2015.03.024

Koh, N., Ma, Z., Sarup, A., Kristl, A. C., Agrios, M., Young, M., & Miri, A. (2024). Selective direct motor cortical influence during naturalistic climbing. bioRxiv. https://doi.org/10.1101/2023.06.18.545509

Liberti, W. A., Markowitz, J. E., Perkins, L. N., Liberti, D. C., Leman, D. P., Guitchounts, G., Velho, T., Kotton, D. N., Lois, C., & Gardner, T. J. (2016). Unstable neurons underlie a stable learned behavior. *Nature Neuroscience*, *19*(12), 1665–1671. https://doi.org/10.1038/nn.4405

Liberti, W. A., Schmid, T. A., Forli, A., Snyder, M., & Yartsev, M. M. (2022). A stable hippocampal code in freely flying bats. *Nature*, *604*(7904), 98–103. https://doi.org/10.1038/s41586-022-04560-0

Lütcke, H., Margolis, D. J., & Helmchen, F. (2013). Steady or changing? Long-term monitoring of neuronal population activity. *Trends in Neurosciences*, *36*(7), 375–384. https://doi.org/10.1016/j.tins.2013.03.008

Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., & Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, *21*(9), 1281–1289. https://doi.org/10.1038/s41593-018-0209-y

Micou, C., & O'Leary, T. (2023). Representational drift as a window into neural and behavioural plasticity. *Current Opinion in Neurobiology*, *81*, 102746. https://doi.org/10.1016/j.conb.2023.102746

Mimica, B., Dunn, B. A., Tombaz, T., Bojja, V. P. T. N. C. S., & Whitlock, J. R. (2018). Efficient cortical coding of 3D posture in freely behaving rats. *Science*, *362*(6414), 584–589. https://doi.org/10.1126/science.aau2013

Miri, A., Warriner, C. L., Seely, J. S., Elsayed, G. F., Cunningham, J. P., Churchland, M. M., & Jessell, T. M. (2017). Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex. *Neuron*, *95*(3), 683-696.e11. https://doi.org/10.1016/j.neuron.2017.06.042

Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S., & Churchland, A. K. (2019). Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience*, *22*(10), 1677–1686. https://doi.org/10.1038/s41593-019-0502-4

Nath, T., Mathis, A., Chen, A. C., Patel, A., Bethge, M., & Mathis, M. W. (2019). Using DeepLabCut for 3D markerless pose estimation across species and behaviors. *Nature Protocols*, *14*(7), 2152–2176. https://doi.org/10.1038/s41596-019-0176-0

Omlor, W., Wahl, A.-S., Sipilä, P., Lütcke, H., Laurenczy, B., Chen, I.-W., Sumanovski, L. T., van 't Hoff, M., Bethge, P., Voigt, F. F., Schwab, M. E., & Helmchen, F. (2019). Context-dependent limb movement encoding in neuronal populations of motor cortex. *Nature Communications*, *10*(1), 4812. https://doi.org/10.1038/s41467-019-12670-z

Oswald, M. J., Tantirigama, M. L. S., Sonntag, I., Hughes, S. M., & Empson, R. M. (2013). Diversity of layer 5 projection neurons in the mouse motor cortex. *Frontiers in Cellular Neuroscience*, *7*. https://doi.org/10.3389/fncel.2013.00174

Resendez, S. L., Jennings, J. H., Ung, R. L., Namboodiri, V. M. K., Zhou, Z. C., Otis, J. M., Nomura, H., McHenry, J. A., Kosyk, O., & Stuber, G. D. (2016). Visualization of cortical, subcortical and deep brain neural circuit dynamics during naturalistic mammalian behavior with head-mounted microscopes and chronically implanted lenses. *Nature Protocols*, *11*(3), 566–597. https://doi.org/10.1038/nprot.2016.021

Rokni, U., Richardson, A. G., Bizzi, E., & Seung, H. S. (2007). Motor Learning with Unstable Neural Representations. *Neuron*, *54*(4), 653–666. https://doi.org/10.1016/j.neuron.2007.04.030

Sacrey, L.-A. R., Alaverdashvili, M., & Whishaw, I. Q. (2009). Similar hand shaping in reaching-for-food (skilled reaching) in rats and humans provides evidence of homology in release, collection, and manipulation movements. *Behavioural Brain Research*, *204*(1), 153–161. https://doi.org/10.1016/j.bbr.2009.05.035

Sadeh, S., & Clopath, C. (2022). Contribution of behavioural variability to representational drift. *ELife*, *11*, e77907. https://doi.org/10.7554/eLife.77907

Sauerbrei, B. A., Guo, J.-Z., Cohen, J. D., Mischiati, M., Guo, W., Kabra, M., Verma, N., Mensh, B., Branson, K., & Hantman, A. W. (2020). Cortical pattern generation during dexterous movement is input-driven. *Nature*, *577*(7790), 386–391. https://doi.org/10.1038/s41586-019-1869-9

Shepherd, G. M. G. (2009). Intracortical cartography in an agranular area. *Frontiers in Neuroscience*, *3*(3), 337–343. https://doi.org/10.3389/neuro.01.030.2009

Stevenson, I. H. (2018). Omitted Variable Bias in GLMs of Neural Spiking Activity. *Neural Computation*, *30*(12), 3227–3258. https://doi.org/10.1162/neco_a_01138

Wagner, M. J., Kim, T. H., Kadmon, J., Nguyen, N. D., Ganguli, S., Schnitzer, M. J., & Luo, L. (2019). Shared Cortex-Cerebellum Dynamics in the Execution and Learning of a Motor Task. *Cell*, *177*(3), 669-682.e24. https://doi.org/10.1016/j.cell.2019.02.019

Wang, X., Liu, Y., Li, X., Zhang, Z., Yang, H., Zhang, Y., Williams, P. R., Alwahab, N. S. A., Kapur, K., Yu, B., Zhang, Y., Chen, M., Ding, H., Gerfen, C. R., Wang, K. H., & He, Z. (2017). Deconstruction of Corticospinal Circuits for Goal-Directed Motor Skills. *Cell*, *171*(2), 440-455.e14. https://doi.org/10.1016/j.cell.2017.08.014

Weiler, N., Wood, L., Yu, J., Solla, S. A., & Shepherd, G. M. G. (2008). Top-down laminar organization of the excitatory network in motor cortex. *Nature Neuroscience*, *11*(3), 360–366. https://doi.org/10.1038/nn2049

Whishaw, I. Q., Faraji, J., Kuntz, J., Mirza Agha, B., Patel, M., Metz, G. A. S., & Mohajerani, M. H. (2017). Organization of the reach and grasp in head-fixed vs freely-moving mice provides support for multiple motor channel theory of neocortical organization. *Experimental Brain Research*, *235*(6), 1919–1932. https://doi.org/10.1007/s00221-017-4925-4

Whishaw, I. Q., Pellis, S. M., & Gorny, B. P. (1992). Skilled reaching in rats and humans: evidence for parallel development or homology. *Behavioural Brain Research*, *47*(1), 59–70. https://doi.org/10.1016/S0166-4328(05)80252-9

Yamawaki, N., Raineri Tapies, M. G., Stults, A., Smith, G. A., & Shepherd, G. M. (2021). Circuit organization of the excitatory sensorimotor loop through hand/forelimb S1 and M1. *ELife*, *10*, e66836. https://doi.org/10.7554/eLife.66836

Zhu, F., Grier, H. A., Tandon, R., Cai, C., Agarwal, A., Giovannucci, A., Kaufman, M. T., & Pandarinath, C. (2022). A deep learning framework for inference of single-trial neural population dynamics from calcium imaging with subframe temporal resolution. *Nature Neuroscience*, *25*(12), 1724–1734. https://doi.org/10.1038/s41593-022-01189-0

Acknowledgments

The authors would like to thank Gabriella Wheeler Fox, Zach Morrissey, Siwei Wang, Harold Rockwell, Sunnie Hong, Steven A. Redford, Stephanie E. Palmer, Nicholas G. Hatsopoulos, Matthew T. Kaufman, William A. Liberti, and David L. McLean for useful discussions and feedback on the manuscript. The authors would also like to thank Dalton D. Moore, Chadbourne M. B. Smith, Gabriella Wheeler Fox, and Steven A. Redford for technical assistance. This work was supported by National Science Foundation Graduate Research Fellowship Program (NSF GRFP) fellowship DGE-1746045 (E.A.d.); and National Institutes of Health grants R01-NS094184 (J.N.M). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Methods

Mice

All experimental procedures were approved by the University of Chicago Institutional Animal Care and Use Committee and were performed in compliance with the Guide for the Care and Use of Laboratory Animals. Data were collected from B6.Cg-Igs7tm148.1 (tetO-GCaMP6f, CAG-tTA2) Hze/J mice (Jax strain 030328, also known as Ai148D) of either sex (n = 2 female, n = 3 male). The line was originally obtained from Jackson Labs and bred in house. Mice were between 10 and 12 weeks old at the start of the shaping and training protocol.

Surgical procedures

Animals underwent three surgical procedures: AAV-Cre injections, prism implantation, and baseplate implantation. For all surgeries, anesthesia was induced with isoflurane (induction at 4%, maintenance at 1–1.5%). During AAV injection, two small burr-hole craniotomies were drilled to allow insertion of a micropipette and injection of virus. Injections were made at three depths (650, 400, and 200um below the surface) at two locations targeting caudal forelimb area (0.25mm anterior and 1.25mm lateral to bregma, and again at 0.25mm anterior and 1.75mm lateral to bregma). After injections, the burr holes were sealed with dental cement and the skin sutured. Animals recovered for a minimum of seven days before prism implantation. During prism implantation, a 2mm diameter craniotomy was drilled, centered at 0.25mm anterior and 1.5mm lateral to bregma. To ease insertion of the prism into the cortical tissue, a stereotaxically mounted scalpel was used to make a 1.25mm wide incision running along the medial lateral axis (parallel to the sagittal plane), centered at the same coordinates as the craniotomy; the incision ran to a depth of 800um. The prism was then inserted into the tissue along this incision to the same depth, with the angled part of the prism directed posteriorly. The resulting field of view is thus anterior to the site of insertion. The remaining surface of brain tissue was covered with Kwil-sil (World Precision Instruments) and the prism was cemented in place. Animals recovered for a minimum of seven days before baseplate implantation. During baseplate implantation, a plastic baseplate was cemented above the prism and grin lens, placed to allow the best alignment of the microscope with the focal plane of the prism. This final procedure is non-invasive, requiring only the application of cement to previously implanted hardware; there is no contact with soft tissue. Animals recovered for a minimum of one full day before the start of food restriction and habituation in the arena.

Behavioral training

The entire behavioral experience consisted of three parts: habituation, shaping, and training. During all phases, the mice were briefly (~1min) and lightly anesthetized with isoflurane so that the miniscope could be attached to their headplates. Mice were then placed in the arena and allowed to recover completely (defined by showing unencumbered, normal ambulatory movement, at least 5 minutes) before the behavioral protocol for that day began. The arena was 15cm x 15cm with a reach port (6mm wide rectangular opening that ran from the floor up ~8cm) located in the center of one of the walls (see schematic in Fig. 1B). The wall with the reach port was slanted away from the arena at a 30 degree angle relative to vertical to enable the mice to maneuver better at the reach port while wearing the miniscopes on their heads. Removable walls were placed on either side of the reach port to limit possible body positions around the reach port to make the pedestals reachable with only the left paw during the early phases of shaping and training (see schematic in Fig. 1B). For some animals these walls were present throughout training, including the days used in this study; for other animals, these walls were removed before the days included in this study. There was no notable impact on the overall set of body positions these animals used in a given session before or after removal of the walls.

Food restriction, Habituation and Shaping

At least one full day after baseplate implantation, mice began food restriction, in which they received a limited amount of pellet food, titrated to maintain their weight between 85% and 90% of their baseline weight. Mice were food restricted for at least one week before habituation began, sometimes longer in order to achieve a weight between 85% and 90% of baseline. Mice were singly housed throughout food restriction, habituation, and training.

Mice were habituated to the training arena and the miniscope for at least seven days before shaping began. During habituation, mice were placed in the behavioral arena for twenty to thirty minutes, once per day. At the start of habituation, twenty millet seeds were placed on the ground inside the arena next to the reach port, to encourage an association between the reach port and food.

During shaping, mice were again placed in the behavioral arena for twenty to thirty minutes, once per day. Instead of being presented inside the arena, seeds were presented in two places outside the arena: approximately ten seeds were presented on the ground outside the arena immediately in front of the reach port. These were reachable by the animal's tongue. Seeds were also

presented in a low square trough outside the arena near the reach port. The trough was positioned a few millimeters laterally to the center of the reach port, such that the animal would only be able to successfully reach the trough with its contralateral forelimb. All mice were shaped and trained to use their left paws. The goal of shaping was to train the animal to associate reaching its forelimb outside the arena with getting food. Functionally, during shaping, the mice would scoop or knock seeds out of the trough and onto the ground, then lick them up with their tongues. Any large magnitude forelimb movement through the reach port was considered a reach. In order to move from shaping to training, mice needed to make at least 50 reaches with their left forelimb during a twenty minute period within the shaping session, and at least 80% of their total reaches in that period needed to be with their left forelimb. This shaping procedure is based on that described in Chen et al., 2014. Shaping was performed daily until mice achieved these criteria within a single shaping session. Often during the last shaping session before training, mice far exceeded these criteria, reaching 75-100 times with their left forelimbs. Mice began the training protocol the next day after achieving the shaping criteria. If, on the first training day, the animal did not show interest in reaching for the seeds on the pedestal, or reached very few times, the training session was converted back to a shaping session, to reinforce the association between using their forelimbs and getting seeds. The animal was held to the same requirements as before for returning to training.

We found that on the first day of shaping there is a critical moment requiring manual involvement. Sometimes at the beginning of shaping, the first reaches that mice make out of the port do not make contact with the seed trough or do not make strong enough contact to knock any seeds out. We observed that many mice give up reaching entirely after only a few tries if their initial reaches do not lead to reward. We found that if we reward the first one or two left-paw reaches by manually knocking seeds out of the trough immediately after the reach, the animals continue to reach and quickly learn the desired association between left-paw reaches and food reward.

Training

During training, seeds were only presented to the mouse on the seed delivery pedestals. Pedestals were positioned relative to the reach port such that they were within reaching distance with the paw, but out of reach of the tongue; this positioning was adjusted for each mouse because of size differences across mice, particularly across gender. The pedestals were narrow such that they could only hold one seed each and the seed was easy to knock off and out of reach. On a given training day, the mice were in the arena for between 75 and 90 minutes. Periods of time where seeds were presented (and mice subsequently reached) were interleaved with break periods during which no seeds were presented. Seed presentation periods lasted 3-5 minutes, followed by 2-4 minutes of break. There were 8-12 seed presentation periods per session, totaling 40 minutes of seed presentation (and subsequent reaching behavior) per day. Training was performed at approximately the same time each day for a given mouse. Imaging took place every training day. For each mouse, five sessions from late in training were used in this study, starting between day 20 and 22.

Seed delivery

There were two different kinds of seed delivery methods used for the mice included in this study. Mice only ever experienced one seed delivery method throughout their training.

For two of the mice (mice 1 and 2), there were three permanent seed delivery pedestals positioned in an arc outside the reach port approximately equidistant from the close edge of the reach port and referred to as medial, lateral, and far lateral (schematized in Fig 1B). Their positioning was offset to the mouse's right such that the pedestals were not reachable with the right forelimb. Seeds were delivered manually by the experimenter. Mice were required to turn away from the seed port and walk towards the back of the arena in order for a new seed to be delivered, thus initiating a new "trial". Only one seed was placed per trial. The order of seed placement was random but counterbalanced such that there was approximately the same number of deliveries per pedestal in a given session.

For one of the mice (mouse 3), the same manual delivery method was used with the three permanent pedestals, but seeds were only placed on the center pedestal, referred to as the lateral pedestal.

For the final two mice (mice 4 and 5), there was a single mobile seed delivery pedestal which was moved laterally between two different positions throughout the session by a linear actuator. Their positioning was offset to the mouse's right such that the pedestals were not reachable with the right forelimb. The two pedestal positions are a subset of the three positions used for manual delivery (schematized in Fig 1B), using only the two pedestals closest to the mouse's midline, referred to as the medial and lateral pedestals. For these mice, seeds were delivered automatically: the pedestal was lowered by a linear actuator into a large trough of seeds and then lifted again, having caught a seed. The design of our automated seed delivery system is very similar to that described in Ellens et al., 2016, using the pellet hopper and delivery arm, but without a reaching shelf. The trough of seeds was multiple centimeters below the floor level of the arena and thus out of reach. The schematic in Fig 1A is based on our automated seed delivery system. The order of the two delivery positions was random in a given session, determined by a Bernoulli random variable in custom arduino code. To trigger the delivery of another seed, mice had to break an infrared beam at the back of the arena.

For the manual seed delivery, if mice turned away to trigger another seed delivery but the previous seed was still present on the pedestal, the seed was left in position for up to two additional trials before that seed was removed and a new seed was placed. For the automated seed delivery, there was no timeout for unsuccessful reaches and no requirement that the seed be removed from the pedestal before triggering another seed delivery.

For mice that had seeds delivered to multiple pedestal locations, they were trained on all pedestal locations from the beginning of training.

Behavioral video recording and kinematic reconstruction

Videography and keypoint extraction

We took high speed, color video of the mouse during reaching at a frame rate of 200Hz using a Flir 16S2C USB3 Blackfly color camera (BFS-U3-16S2C-CS). All kinematic data used in this paper is captured in 2D, with a camera that is at an approximately 30 degree angle from head on, offset to the mouse's left to best view the left forelimb and paw. Additionally, the camera view is angled approximately 20 degrees from above. The schematic in Fig 1A represents the camera viewing angle. We used DeepLabCut to extract a variety of kinematic features from this video data (Mathis et al., 2018; Nath et al., 2019). On the paw, we tracked the joints of the first three digits and the wrist; on the head, we tracked the nose tip, corner of the eye, and a prominent corner of the surgical implant (Fig. 1A). After extracting keypoints using DeepLabCut, kinematic data was further processed with custom MATLAB software. Points that were identified as poorly labeled (because of large distances traveled from the previous timepoint, or because markers that are anatomically close together were labeled as too far apart) were linearly interpolated over. Data was then smoothed: missing frames were linearly interpolated over, we smoothed using a butterworth filter (4th order, low-pass filter with 50Hz cutoff), then removed the interpolated (originally missing) frames.

Alignment across days

The position of multiple non-moving features of the arena (referred to as static markers) were extracted from the videos and used to align kinematic data across days. We adjusted for translation only: the x and y position of the most reliably tracked static marker was subtracted from all the kinematic variables. Generally this translation was minimal because the cameras were fixed in place relative to the arena. We also extracted the position of the pedestal and the presence of a seed for each reach automatically from the video data.

Calculating kinematic variables of interest

We calculated the centroid of the paw as the average position of seven markers on the paw (wrist, digits 1-3 bases and middles). The centroid of the head was calculated as the average of all three markers on the head (nose, front corner of the eye, and corner of the microscope). See Fig 1A and C for visualization of these markers. To calculate a low dimensional estimate of paw shape and orientation, we subtracted the paw centroid position from the seven paw markers, in x and y separately, then performed pca on these 14 dimensions (seven markers, in x and y). We projected data onto the first 4 dimensions, then included it in the linear model. We performed the same procedure for the three markers on the head to capture an estimate of head orientation, keeping 2 principal components of the 6 total to include in the linear model. We pooled aligned data across all five days before performing PCA, such that the PCA values were meaningful relative to each other across days. For one mouse (mouse 3), PCs of the paw markers were not included in analysis due to poor digit tracking.

Note that our kinematic data is two dimensional and extracted from a single video, such that the head and paw PCs capture mixtures of rotation and translation as well as shape. The position of the paw and head in 2D are subtracted out from paw and head markers, respectively, before performing PCA to remove the effects of translation in x and y on our estimates of paw and head rotation and shape. However, the translation in z (depth, towards and away from the camera) cannot be subtracted out and thus is captured in our PCA. To minimize the impact of translation in z on our ability to capture paw orientation and shape from a 2D image, we aligned the camera such that the major axis of paw movement is parallel to the imaging plane. Hence, we refer to the paw PCs as capturing orientation and shape of the paw.

We calculate velocity from the position data using the following approach: we first calculate the difference between position in consecutive timepoints and divide that by the amount of time between timepoints (generally 5ms). We then treat that as the velocity at the halfway point between the observed timepoints and use linear interpolation to infer the velocity at the observed timepoints.

Interpolation

For the linear encoding model analysis, kinematic data (200Hz) was interpolated to the timepoints of the neural data (30Hz) using the MATLAB function interp1.

Segmenting reaches

Reaches are defined using the paw centroid position relative to a static marker on the reach port (right magenta dot in Fig 1A): the paw was considered outside the arena when it was to the left of the static marker, and inside the arena when it was to the right of it. The onset of a reach was identified as the first video frame where the paw centroid was outside the arena and similarly the end of the reach was identified as the last video frame when the paw centroid was outside the arena before it re-entered. By this definition, reaches often consist of one out and back motion, but could include multiple out and back motions if the paw centroid did not pass to the right of the static marker in between.

Reaches were clustered into multiple reaches in quick succession, followed by a longer break during which the animal left the reach port to trigger the delivery of another seed (see Fig 1F top for illustration). We refer to such a cluster as a bout. Bout starts were identified using the inter-reach interval: an inter-reach interval of greater than 1 second separated two bouts. For some sessions for some mice, we used a shorter threshold between bouts (as low as 700ms), depending on the distribution of inter-reach intervals and visual inspection of the behavior.

All reaches to all pedestals were included in the neural coding analyses, since all reaches constituted movements that can be used to study the motor cortical encoding of different kinematic features.

Performance on the task

A successful reach is one where the animal grasped the seed and brought it into the arena. Success rate is defined as the proportion of reaches that are successful out of all reaches made with a seed present on the pedestal in the medial or lateral position. Reaches where a seed is not present at the beginning of the reach are not included in the total. Success was scored manually by watching the video recording of each reach.

For all animals, only reaches when the seed was on the medial and lateral pedestals were included to calculate the success rates. Of the two animals trained with seeds presented on the far lateral pedestal, neither ever successfully retrieved a seed from that pedestal, and it was deemed too far for them to reach from the port. For these animals, reaches while a seed was presented at that position were not included in success rate calculations but were included in neural encoding analyses.

Inclusion of reaching data in analysis

Success versus fail: Due to the nature of the task, successful reaches naturally exhibit lower variance across multiple kinematic features compared to failed reaches. For example, the paw must pass within a certain distance of the pedestal during the reach for the animal to grasp the seed. In all neural coding analysis, we include both successful and failed reaches, as both sets of movements can be used to study the motor cortical encoding of different kinematic features.

Timepoints during reaching: We had the least occlusion and the most reliable tracking of the paw and digit markers during reaching (when the paw was outside of the arena) so we focus our analysis of neural coding on the animal's movements there. Hence, only timepoints during reaching were included in the linear model analyses.

One-photon calcium imaging

Image acquisition

One-photon calcium imaging was performed using an Inscopix nVista 3.0 miniscope (excitation centered at 475 +/- 7 nm, collection centered at 535 +/- 25 nm) controlled by Inscopix Data Acquisition Software (IDAS). A baseplate was chronically attached to the animal's skull and centered over the chronically implanted grin lens and prism (see *Surgical procedures*). The lightweight (2.2g) miniscope was attached to the baseplate during recording sessions. Data was collected from a single imaging plane at 30Hz, referred to as the field of view (FOV) for this mouse. The field of view corresponds to approximately 900 x 650 micrometers, captured in 1280 x 800 pixels. Use of the 45-degree prism means that our field of view is perpendicular to the surface of cortical tissue and putatively spans cortical layers 2/3 and 5a based on the size of the FOV, with the cortical surface apparent at the top of the FOV. Each recording session lasted 40 minutes, broken up into 3 to 5 minute recording chunks, each separated by a break of anywhere from 2 to 4 minutes. Imaging was continuous during a recording chunk, but paused during breaks. When analyzing imaging data from a given session, data from all chunks in that session were concatenated and analyzed together.

Locating the same field of view over days

The same field of view was imaged across all days of training. Because the prism is chronically implanted into the neural tissue, the only degree of freedom in identifying the field of view is depth (because of the 90 degree rotation from the prism and its orientation, more "shallow" corresponds to more posterior, whereas deeper into the tissue corresponds to more anterior). On the first day of training, the depth with the most active cells was selected while the animal was moving around the arena but not performing the

task. Then, some small blood vessels were identified and used as fiduciary markers to find the same plane on subsequent days. Imaged FOVs are estimated between 50 and 100 micrometers deep in the tissue from the surface of the prism.

Pre-processing of calcium imaging data

Data was pre-processed using the Inscopix Data Processing Software (IDPS) MATLAB API and included 4x downsampling by averaging (a 4x4 square of pixels is averaged to one pixel), spatial filtering by frame (bandpass with low cutoff 0.005Hz and high cutoff 0.5Hz, where Hz refers to units of 1/pixel), and across-frame motion correction using TurboReg (Thevanaz et al., 1998). We exported the resulting stacks of images for cell identification using Caiman CNMFE (Zhou et al., 2018; Giovannucci et al., 2019) to extract spatial footprints and single-cell fluorescence time series. Due to computational limitations, each 40 minute session was processed in 3 batches. Data from a session was broken up into 3 batches of approximately equal size by recording chunk (see previous methods section *Image acquisition*), distributing the recording chunks such that each batch contained neural data from the beginning, middle, and end of a session (e.g. batch 1 contained recording chunks 1, 4, 7, etc.). After passing each batch of imaging data through the Caiman CNMFE algorithm, we matched cells across batches using CellReg (Sheintuch et al., 2017), with the requirement that weighted spatial footprints had to have a spatial correlation of 0.9 across batches to be considered the same cell. Only cells found in all batches were included in the final cell set for a given session. The spatial footprint for a given cell for a given session was computed by averaging across its spatial footprints from each batch. Time series variables (Caiman variables C, C_raw, and S) for each cell were z-scored within batch before de-interleaving and concatenating across batches. Then, for each cell separately, we convolved its inferred spike data (S output from Caiman) with a gaussian of 30ms standard deviation. This is the neural activity signal used for all analysis.

Registering cells across days

We registered cells across days using CellReg (Sheintuch et al., 2017), with a strict, fixed requirement that across two consecutive days, weighted spatial footprints had to have a spatial correlation of 0.9 to be considered the same cell.

Separating cells into superficial and deep sets

For each mouse, we computed an all cells-all days map of centroids, which included all cells found on any day. For cells found only on one day, their centroids on this map was the centroid from the one day. For cells found on more than one day, their centroid on this map was computed by averaging across their centroids on all days they were found. After computing the all cells-all days centroid map, we separated cells into two groups according to their vertical position in the field of view (corresponding to cortical depth). We found the median vertical position across cells, and assigned cells above that to the superficial set and below that to the deep set. We removed cells whose centroids were within 10 pixels (vertically) of the identified median position, to better disambiguate between superficial and deep sets. Note that due to the depth of cortical tissue that we were able to access with our prism lens, our imaging FOV includes only the upper layers of CFA. Our superficial and deep sets thus putatively map onto L2/3 and L5a, respectively.

Reach modulation via ANOVA

Reach modulation was determined for each cell individually via an ANOVA comparing its activity during reaching timepoints with its activity during all other timepoints, which include other behaviors such as locomotion, grooming, food manipulation, and food consumption. Significance of reach modulation was assessed with a p-value of 0.05 with Bonferroni correction for the number of cells in the dataset (that day's recording session for that mouse).

Peri-event time histograms (PETHs)

We calculated PETHs by averaging neural activity in each time bin across trials. Each reach is one trial. Neural data from all reaches were included in PETH calculation, regardless of success or seed position. Here, t = 0 indicates the start of the reach, also referred to as reach onset. For visualization, PETHs were normalized such that the lowest value = 0 and highest value = 1 within the shown time window, [-500 1000]ms relative to reach onset. For comparing PETHs over days, PETHs were normalized in the same way but within a narrower window, [-100 500]ms relative to reach onset. This shorter window was used when calculating the correlation between PETHs for the same cell found on different days.

Cell ID shuffle control

The cell ID shuffle distribution was generated by calculating the correlation between PETHs for non-same cells found on different days. For each pair of days (n = 5 days, yielding n = 10 day-pairs), we randomly sampled a cell from each day, ensured they were not registered as the same cell, and computed the correlation between their PETHs. We repeated this process 1000 times for each

day-pair, yielding a total of 10,000 samples in the shuffle distribution. Shuffle distributions were generated separately for each mouse.

Linear encoding models of single neurons

On each day separately, we fit a linear model where we predicted the activity of each cell based on measured behavioral and task variables. We used an instantaneous encoding model: neural activity at time t was predicted from behavior and task variables at time t. No time-shifted versions of any predictors were included. Behavioral variables were all interpolated from the original 200Hz to the 30Hz framerate of the neural data (see *Interpolation* section in *Behavioral video recording and kinematic reconstruction*). Here and throughout this work, dataset refers to the neural and behavioral data from one mouse from one day, e.g. there are 5 mice and 5 days for each mouse, yielding 25 datasets.

Data inclusion and predictor sets/design matrices

For each mouse, data from five sessions from late in training were used in this study, starting between day 20 and 22 (see *Training*). In total, 15 predictor variables were included: the position and velocity in x and y of the paw and the head (8 variables), the first four principal components (PCs) of the paw markers, the first two PCs of the head markers, and time within the session (not including breaks between recording chunks; time variable is equivalent to the index of the calcium imaging data). For one mouse (mouse 3), PCs of the paw markers were not included due to poor digit tracking. Only timepoints during reaching (when the paw was outside the arena) were used to fit the linear models; timepoints not during reaching were excluded (see next section and Supp Fig 4D for the number of timepoints per dataset broken down by mouse and day). As with the PETHs, all reaches were included in this analysis, regardless of success or seed position.

Covariate balancing across days to allow for direct comparison of linear fits

In order to directly compare properties of our linear model fits across days, we must ensure that each predictor has a similar distribution to itself over days, and also that the joint distribution of predictors is similar over days. Without this step, we cannot guarantee that any observed encoding instability is not a result of the animals exploring a slightly different part of kinematic space each day. We subsample from each day's collection of reach timepoints in the following way to create a similar distribution for each day within the high-dimensional space of all predictors.

Because we want to compare the encoding for the same absolute values of the predictors, we first z-score on all days together rather than on each day separately. For each mouse separately, we treated each predictor as follows: We concatenated all reach timepoints for that predictor across all five days, then z-scored. We then separated this data back into days and performed the covariate balancing procedure described below.

We use a greedy matching algorithm without replacement to balance the joint distributions of all behavioral predictors across days. In a given day's joint distribution, each sample is a reach timepoint; it corresponds to a vector of all kinematic predictor values at that timepoint (14D for mice 1, 2, 4, 5, and 10D for mouse 3). The time predictor is not included in this vector during covariate balancing. We use the day with the fewest reach timepoints as the reference day. For each timepoint in the reference day, we identify the timepoint in each other day that is most similar to the reference day timepoint. Similarity is calculated as the correlation between the vectors of kinematic predictors. Timepoints on the reference day are matched sequentially. After a timepoint on a non-reference day is matched to a reference day timepoint, it is no longer available for matching to later reference day timepoints; hence, the algorithm is greedy. After finding matches for all reference day timepoints, we remove matches in two phases. First, we remove non-reference day timepoints that have a correlation with their matched reference day timepoint that is below 0.9. There may be a different reference day timepoint for which this timepoint has a correlation greater than 0.9, but the algorithm guarantees that reference day timepoint already has a better match on the given non-reference day. Second, we remove reference day timepoints (and all their associated matched non-reference day timepoints) that have fewer than 2 remaining matches, meaning that reference day sample has only one other day where a sufficiently similar timepoint was found. Because the algorithm is greedy, we randomize the order that the reference timepoints are matched in; this means that the best matches are not biased towards the beginning of the session, but instead distributed throughout. If an animal's behavior is dissimilar over days, with less overlapping predictor sets, the covariate balancing procedure would yield a very small number of samples left to fit the models with. Even before subsampling, the probability density functions of our predictor sets are largely overlapping, indicating very similar behavior across days. Thus, covariate balancing leaves us with a large number of samples for each day to fit the models with (12889 +/- 4933 [7276 22175] samples per dataset (mean +/- st. dev. [range]), which constitutes 68.7% +/- 17.2% [31.5% 95.%] of timepoints in the original datasets (mean +/- st. dev. [range]); Supplementary Fig 4D). Datasets where lower percentages of timepoints were kept generally constitute days where the animal performed a higher number of reaches but with consistent distributions of behavioral features, such

that subsetting is matching the number of timepoints across days rather than to removing dissimilar behavior (for probability density functions for all predictors for all days for an example animal, see Supplementary Fig 3).
After covariate balancing, each predictor is again z-scored within day before fitting the models. All single day linear model results presented in this paper are fit with predictor sets that have been covariate balanced across all five days for that animal.

Fitting the single-trial family of models

Single-trial models refers to the standard instantaneous linear model fit where we predict neural activity at time t from behavioral and task variables observed at time t. Models were fit using ridge regression, using the MATLAB R2018a function ridge() with a ridge parameter equal to the number of timepoints (samples) in that day; see previous methods section *Covariate balancing across days to allow for direct comparison of linear fits* for statistics on the number of timepoints per day.

We fit a family of models for each cell on each day. Full models were fit using all 15 predictors specified in the methods section *Data inclusion and predictor sets/design matrices*. As described in the main text, we also fit a single variable model for each predictor that included only that predictor, as well as a leave-one-out model for each predictor that included all predictors except for the one in question. Delta-r-squared values are computed as the drop in r-squared from the full model to the leave one out model for the predictor in question, as schematized in Fig 5a.

We used a train-test split of 70% train, 30% test. We did the following procedure to ensure that each train-test split contained a balanced sampling of timepoints from across the session. For all the timepoints included in a day for a given mouse, we split the data into 100 chunks. We then grouped every 10th chunk together, resulting in 10 sets. For example, chunks 1, 11, 21, etc. went into set 1, while chunks 2, 12, 22, etc went into set 2. We then chose 7 random sets to use as the train set, and held out 3 chunks as the test set. We repeated this procedure for 10 cross-validations for all model fits. Within a day, the train-test splits for each cross-validation were the same for all cells.

All reported r-squared and delta-r-squared values are averaged across 10 cross-validations of the train-test split. Delta-r-squareds are calculated within each cross-validation, then averaged across cross-validations. Note that due to the train-test split, r-squared values can sometimes be negative. Similarly, the Delta R-squared (leave one out models) can sometimes be positive, indicating that removing the predictor in question improved the test R-squared values. This typically occurs for small R-squared values or poor fits, where the model is fitting to noise in the train set and thus cannot generalize to the test set.

Significance testing of full model r-squared values

To test whether a full model fit was better than expected by chance, we compare each cell's full model r-squared value to a timepoint shuffle null distribution. For each cell, we generate a null distribution by shuffling the timepoints of the neural activity (predicted variable) and refitting the full model for that day. We repeated this procedure 500 times per cross-validation, for a total of 5000 shuffles. If the observed, cross-validated full model r-squared value was greater than a threshold percentile of time-shuffled r-squared values, we deemed the full model fit for that cell on that day to be significant. We used Bonferroni correction on the threshold percentile, dividing the p-value by the number of cells in each day, such that the threshold percentile is calculated as 100 - 5/(number of cells for this mouse on this day). See Supplementary Fig 4B and C for counts and proportions of cells with significant predictions, broken down by mouse and day.

Fitting trial-average models

All linear models in this study treat each timepoint as an independent observation and regress across all observations in a session. This approach is standard for quantifying the neural encoding of movements, but it conflates encoding for the mean movement with encoding for the variability around the mean. Thus, we compared them to models that predict a neuron's single-trial activity from trial-averaged movements.

For each predictor separately, we calculated a time-varying average movement during reaching, locked to the onset of reach. This is very similar to calculating a PETH, but where we average across snippets of a kinematic variable, instead of snippets of neural activity. We built a new design matrix where for each predictor we replaced each reach with a copy of that predictor's average. Each reach was replaced with a corresponding number of timepoints of the average: for a reach that is 10 timepoints long, we replaced those 10 timepoints with the first 10 timepoints of the reach-locked average, whereas for a reach that is 20 timepoints long, we replaced those 20 timepoints with the first 20 timepoints of the reach-locked average. Only reach timepoints were used for the models, as for the single-trial models. The time predictor was treated in the same manner as the behavioral predictors, yielding a predictor that simply increases linearly over the course of a reach. Trial-average models were fit using all behavioral predictors and the time predictor. Since trial-average models serve as a within-session comparison, we do not perform across-day covariate balancing on trial-average design matrices. All trial average models were fit and tested for significance using the same cross-validated train-test procedure described for single-trial full models above.

Comparison of linear encoding models across days

As stated in the main text, for a given cell, when comparing its encoding of trial-by-trial variance across days, we only consider days where the cell has a significant single-trial full model fit. This represents the majority of cell-days recorded in this study (79.7% of all cell-days). Cell-days without significant single-trial full model fits do not show significant tuning for the movement predictors we consider here. Thus, their model fit components do not represent the activity of the cell in a statistically significant way and should not be used to quantify similarity or dissimilarity of encoding.

Applying models across days

To apply a model from one day (train day) to data from another day (test day), we weight the predictor vectors from the test day with the coefficients from the train day full model fit and add them together to generate a prediction for test day neural activity (schematized in Fig 6A), for which we calculate an r-squared value. As when fitting the models on their train days, when applying models to test day data, the predictors and neural activity are z-scored before applying, so no constant term is required. We used all cell-days with a significant full single-trial model fit as both train and test days: for a given cell, we applied all significant models to all other days with significant models. To test the applied model r-squared value for significance, we compare to the timepoint shuffle null distribution of the test day, the same distribution used to identify if the test day model itself provides a significant prediction. Forward application refers to when the train day occurs before the test day; backward application refers to when the test day occurs before the train day. Because of our requirement that both test and train days have significant full model predictions, all train and test day pairs have a forward and a backward application. For a given train and test day pair, we interpret the neural encoding as stable across these two days if the application in at least one direction is significant.

Notes on comparing linear model fits in the presence of correlated predictors

When predictors are sufficiently correlated, as in our study, there is an infinite set of possible coefficient combinations that can produce the same optimal fit in the full models. Even with appropriate regularization, model fits for a given day will be specific to that day's predictor correlations. If the correlations between predictors change from the training day to the test day, the coefficients derived from the training day may not accurately reflect the predictor relationships on the test day, potentially resulting in a drop in performance on the test day. This drop can occur even if the underlying relationships of interest, e.g. single neuron encoding of individual predictors, remain stable. If predictor correlations are not carefully accounted for, applying fitted models across days may mistakenly suggest that activity-behavior relationships for individual predictors are changing, when in fact, it is the relationships among the predictors that are shifting.

To mitigate this possibility and increase confidence in attributing observed instability to changes in neural encoding rather than changes in behavior, we take the following steps in our across-day analysis: 1) We perform covariate balancing on the movement predictor sets before fitting and comparing linear models. Because this procedure balances the joint distributions of predictors by considering all movement predictors together when matching timepoints, a by-product is partial enforcement of similarity (over days) in the relationships between pairs of predictors. 2) When applying models over days, we do not require model application to be significant in both directions for a pair of cell-days to qualify as having significantly similar encoding. In theory, if there is explanatory power in the predictor set that is maintained across days, this approach allows us to detect that stable encoding even if predictor correlations change meaningfully over days, as illustrated by the following example. On day 1, predictors x_1 and x_2 are correlated with each other and also with the predicted variable y in the same way, so model coefficients will give approximately equal weight to x_1 and x_2, perhaps with bias to one or the other dependent upon noise. On day 2, predictor x_1 is still correlated with predicted variable y in the same way, but x_2 is less correlated with both x_1 and y or perhaps correlated in a different way, so model coefficients for day 2 will give more weight to x_1 than x_2. Applying the model from day 1 to day 2 will not weight x_1 enough and will weight x_2 too highly, leading to a poor prediction. However, applying the model from day 2 to day 1 will yield a fit similar to day 1's own model, since on day 1 the contribution from x_1 or x_2 is approximately equivalent. Requiring model application to be significant in both directions requires predictor correlations to be precisely the same on each day in order to detect stable encoding. Not requiring symmetry allows us to detect stable encoding in the face of potentially changing predictor correlations. In practice, we observe that only 7.2% of pairs of cell-days deemed to have significantly stable encoding by this metric are significant in only one direction (328 out of 4586 pairs of cell-days), meaning that the vast majority (92.8%, or 4258 out of 4586 pairs of cell-days) are bidirectionally significant.

Finally, as with all similar studies, we are able to quantify encoding for the movement predictors we included but cannot speak to whether cells significantly encode other unrecorded features of this behavior. There are likely other features of the animals' behavior that are encoded by a portion of the neural population and are correlated with the variables measured here. As such, in the case of a cell with significant encoding of recorded movement predictors on some days but not all, we are unable to differentiate

between the following two situations: 1) that cell encodes the recorded movement features and has lost its relationship to them on some days or 2) that cell encodes different, unrecorded movement features that are correlated with the recorded features on some days but not all. We thus focus our across-day analyses on asking whether cells that maintain significant encoding of our collection of recorded movement variables maintain the nature of their encoding over days.

Methods References

Chen, C.-C., Gilmore, A., & Zuo, Y. (2014). Study Motor Skill Learning by Single-pellet Reaching Tasks in Mice. *Journal of Visualized Experiments : JoVE*, *85*. https://doi.org/10.3791/51238

Ellens, D. J., Gaidica, M., Toader, A., Peng, S., Shue, S., John, T., Bova, A., & Leventhal, D. K. (2016). An automated rat single pellet reaching system with high-speed video capture. *Journal of Neuroscience Methods*, *271*, 119–127. https://doi.org/10.1016/j.jneumeth.2016.07.009

Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B. L., Koay, S. A., Taxidis, J., Najafi, F., Gauthier, J. L., Zhou, P., Khakh, B. S., Tank, D. W., Chklovskii, D. B., & Pnevmatikakis, E. A. (2019). CaImAn an open source tool for scalable calcium imaging data analysis. *ELife*, *8*, e38173. https://doi.org/10.7554/eLife.38173

Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., & Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, *21*(9), 1281–1289. https://doi.org/10.1038/s41593-018-0209-y

Nath, T., Mathis, A., Chen, A. C., Patel, A., Bethge, M., & Mathis, M. W. (2019). Using DeepLabCut for 3D markerless pose estimation across species and behaviors. *Nature Protocols*, *14*(7), 2152–2176. https://doi.org/10.1038/s41596-019-0176-0

Sheintuch, L., Rubin, A., Brande-Eilat, N., Geva, N., Sadeh, N., Pinchasof, O., & Ziv, Y. (2017). Tracking the Same Neurons across Multiple Days in Ca2+ Imaging Data. *Cell Reports*, *21*(4), 1102–1115. https://doi.org/10.1016/j.celrep.2017.10.013

Thevenaz, P., Ruttimann, U. E., & Unser, M. (1998). A pyramid approach to subpixel registration based on intensity. *IEEE Transactions on Image Processing*, *7*(1), 27–41. https://doi.org/10.1109/83.650848

Zhou, P., Resendez, S. L., Rodriguez-Romaguera, J., Jimenez, J. C., Neufeld, S. Q., Giovannucci, A., Friedrich, J., Pnevmatikakis, E. A., Stuber, G. D., Hen, R., Kheirbek, M. A., Sabatini, B. L., Kass, R. E., & Paninski, L. (2018). Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *ELife*, *7*, e28728. https://doi.org/10.7554/eLife.28728