






Separating cognitive and motor processes in the behaving mouse

Received: 22 August 2023

Accepted: 21 November 2024

Published online: 04 February 2025

 Check for updates

Munib A. Hasnain ^{1,2,7}, Jaclyn E. Birnbaum ^{2,3,7}, Juan Luis Ugarte Nunez¹, Emma K. Hartman¹, Chandramouli Chandrasekaran ^{4,5,6} & Michael N. Economo ^{1,2,6} 

The cognitive processes supporting complex animal behavior are closely associated with movements responsible for critical processes, such as facial expressions or the active sampling of our environments. These movements are strongly related to neural activity across much of the brain and are often highly correlated with ongoing cognitive processes. A fundamental issue for understanding the neural signatures of cognition and movements is whether cognitive processes are separable from related movements or if they are driven by common neural mechanisms. Here we demonstrate how the separability of cognitive and motor processes can be assessed and, when separable, how the neural dynamics associated with each component can be isolated. We designed a behavioral task in mice that involves multiple cognitive processes, and we show that dynamics commonly taken to support cognitive processes are strongly contaminated by movements. When cognitive and motor components are isolated using a novel approach for subspace decomposition, we find that they exhibit distinct dynamical trajectories and are encoded by largely separate populations of cells. Accurately isolating dynamics associated with particular cognitive and motor processes will be essential for developing conceptual and computational models of neural circuit function.

Goal-directed behavior follows from an interplay between cognitive and motor processes in the mammalian brain. For example, while dribbling down the court, a basketball player must track the players around them, decide which play to run, plan their next move, execute fast and accurate movements and flexibly adapt those movements according to the actions of other players. The neural representations of these cognitive and motor processes are often distributed across many of the same brain areas^{1–3} and engaged concurrently^{4,5}. For these reasons, disentangling the neural signatures of specific cognitive and motor processes is challenging.

Many behavioral tasks are designed with this issue in mind, aiming to isolate constituent neural processes by separating them into discrete

temporal epochs. For example, in sensory-guided decision-making tasks, subjects are often trained to withhold responses while stimuli are presented so that perceptual decision-making is isolated in time from associated actions. Analogous paradigms have been developed to uncover the neural underpinnings of working memory^{6,7}, motor planning^{8,9}, contextual encoding¹⁰, reward prediction¹¹ and other cognitive processes¹².

These paradigms often rely on the assumption that experimental subjects withhold motor output during task epochs in which instructed responses are absent. However, uninstructed movements not required for task completion, such as changes in posture, facial expressions and gaze, are commonly observed in rodents, humans and

¹Department of Biomedical Engineering, Boston University, Boston, MA, USA. ²Center for Neurophotonics, Boston University, Boston, MA, USA. ³Graduate Program for Neuroscience, Boston University, Boston, MA, USA. ⁴Department of Psychological and Brain Sciences, Boston University, Boston, MA, USA.

⁵Department of Neurobiology & Anatomy, Boston University, Boston, MA, USA. ⁶Center for Systems Neuroscience, Boston University, Boston, MA, USA.

⁷These authors contributed equally: Munib A. Hasnain, Jaclyn E. Birnbaum. ✉ e-mail: mne@bu.edu

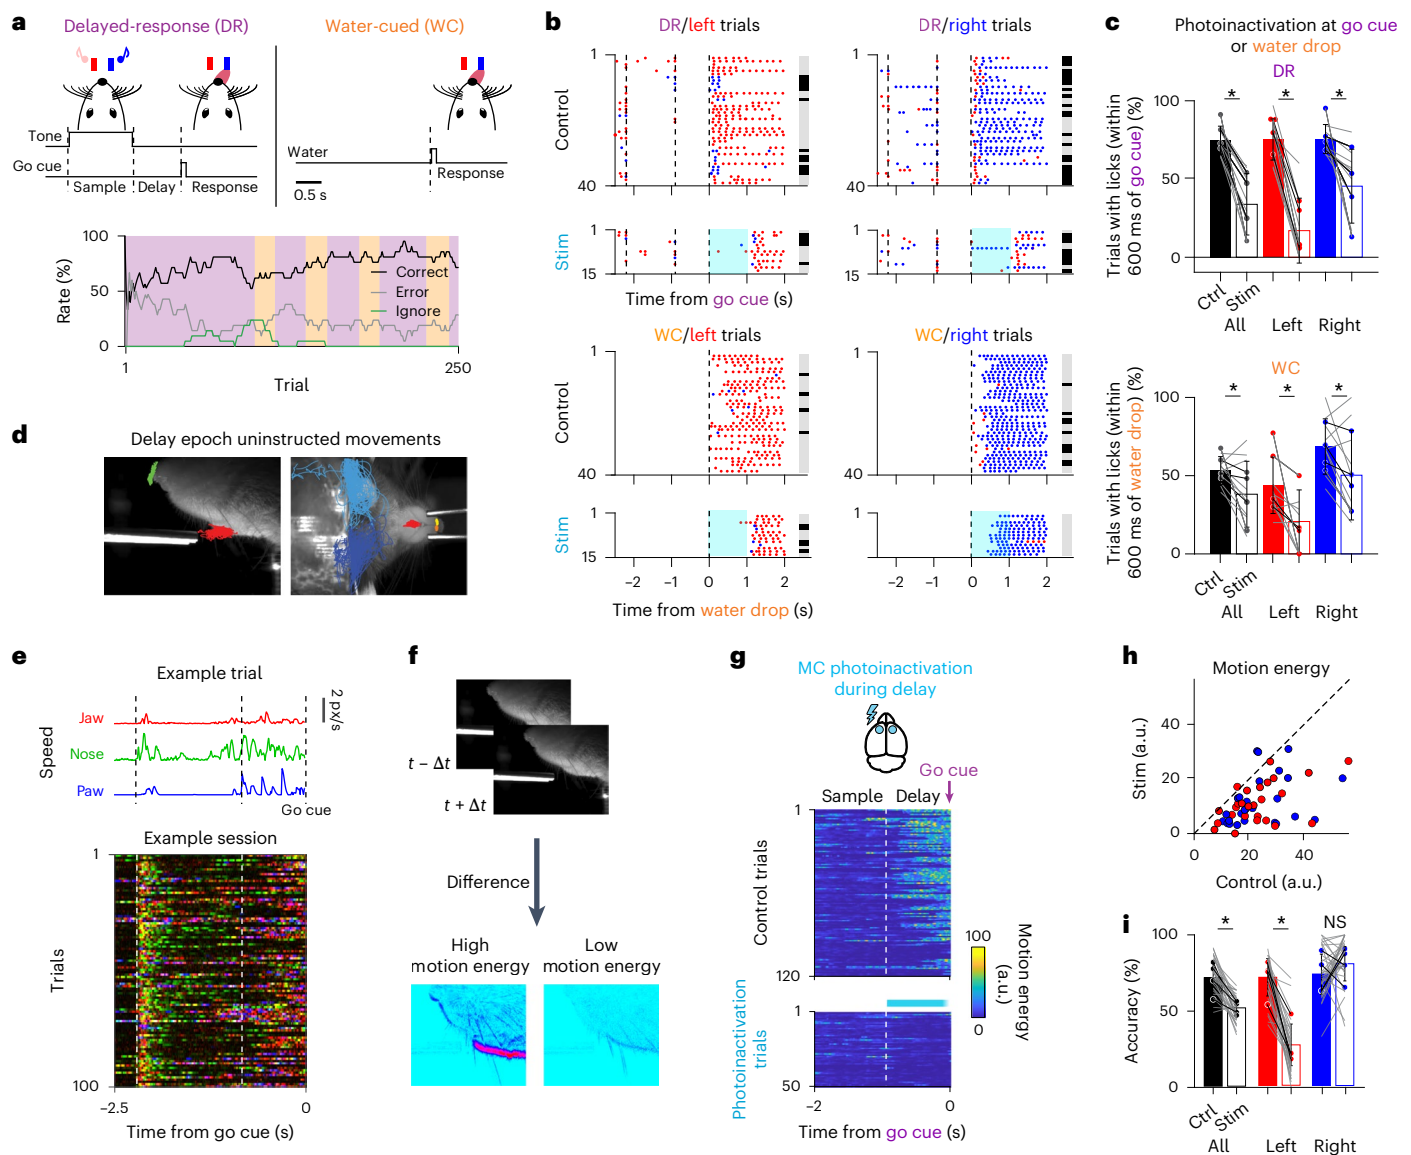


Fig. 1 | Cortical dependence and uninstructed movements during a two-context task. a, Top: schematic of the two-context task. In DR trials, an auditory stimulus was presented during the sample epoch (1.3 s) instructing mice to lick for reward to the right (white noise) or left (8-kHz tone). Mice were trained to withhold their response during a delay epoch (0.9 s) and initiated their response after an auditory go cue. In WC trials, mice were presented with a water reward at a random port at random inter-reward intervals. The WC task was introduced once mice became experts in the DR task. Bottom: 20-trial trailing average of correct, error and ignore rates for an example session. **b**, Performance for an example session separated by trial type (left versus right), control versus photoinactivation trials and by context. Photoinactivation (1 s) was initiated at the go cue or water drop. Blue dots indicate right licks; red dots indicate left licks. Gray bars, correct trials; black bars, error trials. Blue shading indicates the timing of photoinactivation. Vertical black dashed lines in the top panels indicate the onset, from left to right, of the sample period, the delay period and the go cue. In the bottom panels, vertical black dashed lines indicate the timing of the water drop. **c**, Percent of trials with correct licks within 600 ms of the go cue or water drop on control and photoinactivation trials. Colored points represent mean values for each animal ($n = 4$, individual animals connected by dark lines). Light gray lines denote individual sessions ($n = 10$). Bars are across-session means. Asterisks denote significant differences ($P < 0.05$) between control and photoinactivation trials (all DR trials: $P = 0.0001$; DR left trials: $P = 4.5 \times 10^{-6}$; DR right trials: $P = 0.0074$; all WC trials: $P = 0.0371$; WC left trials: $P = 0.0134$; WC right

trials: $P = 0.0316$, paired two-sided t -test). Error bars denote s.d. across sessions. **d**, Behavior was tracked with high-speed video from side (left) and bottom (right) views. Trajectories of delay epoch uninstructed movements from an example session overlaid. **e**, Uninstructed movements are highly variable across trials and across time. Top: jaw, nose and paw speed for an example trial. Bottom: feature overlay for a subset of trials in an example session. At each time bin t , an (r, g, b) color value was encoded as $(\text{jaw}_t, \text{nose}_t, \text{paw}_t)$. Vertical white dashed lines indicate, from left to right, the onset, from left to right, of the sample and delay periods. **f**, Schematic of motion energy calculation. Example frames depicting high and low motion energy are shown. **g–i**, Bilateral motor cortex photoinactivation during the delay epoch. **g**, Motion energy on control (top) and photoinactivation (bottom) trials for an example session. On photoinactivation trials, light was delivered for 0.8 s beginning at the start of the delay epoch. Vertical white dashed lines indicate the onset of the delay period. **h**, Average delay epoch motion energy on single sessions for control and photoinactivation trials ($n = 29$ sessions, four mice). Red points indicate left trials and blue points indicate right trials for a single session. **i**, Performance for control and delay epoch photoinactivation trials. Mice often defaulted to a right choice after delay epoch photoinactivation. Asterisks denote significant differences ($P < 0.05$, paired t -test) between control and photoinactivation trials ($n = 29$ sessions, four mice). Black lines and colored points indicate averages across sessions for individual animals, and light gray lines indicate sessions. Error bars denote s.d. across sessions. Stim, photoinactivation trials; px, pixels; NS, not significant.

non-human primates during learned behavioral tasks^{13–16}. Importantly, uninstructed movements explain much of the variance observed in brain-wide neural activity^{17–19} and can be strongly correlated with task variables^{7,20–24}.

These observations imply a serious challenge for experimental studies addressing a wide range of cognitive processes. If uninstructed movements are correlated with a latent cognitive variable of interest—a stimulus perceived, a decision made, a memory stored or a motor plan formed—then the neural dynamics leading to or resulting from uninstructed movements can be easily misconstrued as responsible for that cognitive process. One common approach to isolate putative cognitive dynamics is to track and ‘regress out’ movements^{14,17,25}, but this approach assumes that cognitive and motor processes are separable—that is, driven by independent neural dynamics. However, some cognitive processes may be inherently embodied such that their associated neural dynamics are linked to overt movement²⁶. Regressing out neural dynamics associated with embodied movements would lead to the inadvertent removal of the precise dynamics that one wants to study. Whether cognitive and motor dynamics are separable remains an open question and likely depends on the brain area, behavior and cognitive process of interest.

Here we address the separability of dynamics associated with cognitive processes and correlated movements. We developed a behavioral paradigm in which mice perform sensory-guided movements involving multiple cognitive signals and exhibit idiosyncratic, task-correlated uninstructed movements. We build upon work demonstrating how different neural processes can be multiplexed in a single brain region^{25,27,28} to develop a novel method for identifying separate subspaces of neural activity containing dynamics related to internal processes and dynamics related to movements. This approach is simple to adopt and does not require tracking or segmentation of body parts. It also does not require an explicit choice of models relating neural dynamics to movement, avoiding common assumptions of linearity. Decomposing activity into complementary internal and movement-related subspaces reveals new insights into the nature of dynamics supporting cognitive processes, their relationship to movements and the manner in which these representations are encoded at the level of single neurons.

Results

Task-switching behavioral paradigm

We designed a behavioral paradigm associated with multiple cognitive processes. In this paradigm, head-fixed mice performed two directional licking tasks that alternated block-wise within each behavioral session. These tasks varied in their cognitive demands but required the same instructed motor output.

This paradigm included an established delayed-response (DR) task in which motor planning is temporally dissociated from movement execution⁹ (Fig. 1a). In the DR task, an auditory stimulus indicated the location of a reward. After a brief delay epoch with no auditory stimuli, a separate auditory ‘go cue’ instructed the mouse to move. If the mouse performed a directional tongue movement to the correct target, a water reward was delivered. In separate blocks of trials, mice performed a water-cued (WC) task in which all auditory stimuli were omitted. Instead, a drop of water was presented at a random point in time at a randomly selected reward port. Animals consumed water upon its presentation. Mice received no explicit cues signaling DR and WC blocks. The DR task required a perceptual decision that guides the planning of a subsequent motor action. In the WC task, uncertainty in the timing and location of reward prevented formation of a motor plan before reward presentation, and, instead, mice detected reward availability using olfactory and/or vibrissal cues²⁹. This structure required animals to maintain an internal representation of the current block identity (context) to maintain high task performance.

We examined uninstructed movements in this behavioral paradigm using high-speed video. We tracked movement of the tongue,

jaw, nose and paws³⁰ and calculated instantaneous motion energy to quantify movement in a feature-agnostic manner¹⁸ (Fig. 1d–f). We found that mice performed uninstructed movements that varied in their identity and timing across both trials and contexts (Fig. 1e and Supplementary Movie 1).

Previous work established the necessity of the anterior lateral motor cortex (ALM) and tongue–jaw motor cortex (tjM1) for the planning and execution of tongue movements^{9,31–33}. To assess cortical involvement in instructed movements (directional licking) in both behavioral tasks, we used optogenetic photoinactivation to inhibit the ALM and tjM1 at the go cue (DR) or water presentation (WC) bilaterally in VGAT-ChR2-EYFP mice (10 sessions, four mice). Simultaneous photoinactivation of both regions at the go cue (DR task) or water drop (WC task) impaired the initiation of movement in both contexts (Fig. 1b,c; all DR trials: $41 \pm 20\%$ reduction, mean \pm s.d., $P = 1 \times 10^{-4}$; all WC trials: $15 \pm 20\%$, $P = 0.04$; and Extended Data Fig. 1c,d). To assess cortical involvement in the expression of uninstructed movements, and whether these movements might be controlled by ALM or tjM1 preferentially, we bilaterally silenced the ALM or tjM1 individually during the delay epoch of DR trials (ALM: 15 sessions, tjM1: nine sessions, four mice). Uninstructed movements were suppressed by photoinactivation of either area (Extended Data Fig. 1b; ALM: $53 \pm 31\%$ reduction, mean \pm s.d., $P = 3 \times 10^{-4}$; tjM1: $62 \pm 24\%$ reduction, $P = 2 \times 10^{-3}$, paired *t*-test) and more strongly suppressed by concurrent photoinactivation of both areas (Fig. 1g,h; 29 sessions, four mice; $73 \pm 35\%$ reduction, mean \pm s.d., $P = 7 \times 10^{-9}$, paired *t*-test), similar to previous observations²¹ (but see ref. 34). In the DR task, delay photoinactivation of both areas led to a significant decrease in overall performance (Fig. 1i; $22\% \pm 14\%$ reduction, mean \pm s.d., $P = 2 \times 10^{-6}$ paired *t*-test; Extended Data Fig. 1b) as expected from the disruption of choice/planning dynamics⁹, often leading to animals responding consistently to a preferred side rather than in a manner commensurate with the sensory cue. These results indicate that both the ALM and tjM1 are involved in the generation of uninstructed movements and the planning and execution of instructed directional tongue movements in the two-context paradigm.

Neural dynamics related to task variables

While movement-related and task-related information is encoded in both the ALM and tjM1 in the DR task^{31,35} and other directional licking tasks³³, task variables are more strongly represented in the ALM^{33,36}, and so we focused our subsequent analyses there. We recorded activity extracellularly in the ALM with high-density silicon probes during the two-context paradigm to track dynamics associated with planning, context and the execution of movements (Fig. 2; two-context paradigm: 12 sessions, six mice, 522 units, including 214 well-isolated single units; an additional three mice were trained only on the DR task: 25 sessions, nine mice, 1,651 units, including 483 well-isolated single units; see Extended Data Fig. 2a for per-session statistics). In the DR task, we found that choice selectivity (right versus left DR trials) was widespread across all task epochs as expected (sample: 36%; delay: 42%; response: 58% of 483 single units; 25 sessions, nine mice). Individual units showed a variety of activity patterns across trial types, including delay epoch preparatory activity, which has often been taken to be a neural signature of motor planning^{32,37} (Fig. 2b, left). Individual units were also selective for behavioral context during all task epochs, including during the inter-trial interval (ITI) (39% of single units, 12 sessions, six mice, 214 single units), suggesting persistent coding of context (Fig. 2b, right).

We next sought to understand how the selectivity patterns observed in individual cells were encoded at the population level. We examined neural activity in state space, in which the firing rate of each unit represents one dimension. We defined one-dimensional coding directions (CDs) within the high-dimensional state space that best encoded task variables^{32,38–40} (Fig. 2c). We first examined putative neural correlates of three cognitive processes: choice, urgency and contextual encoding.

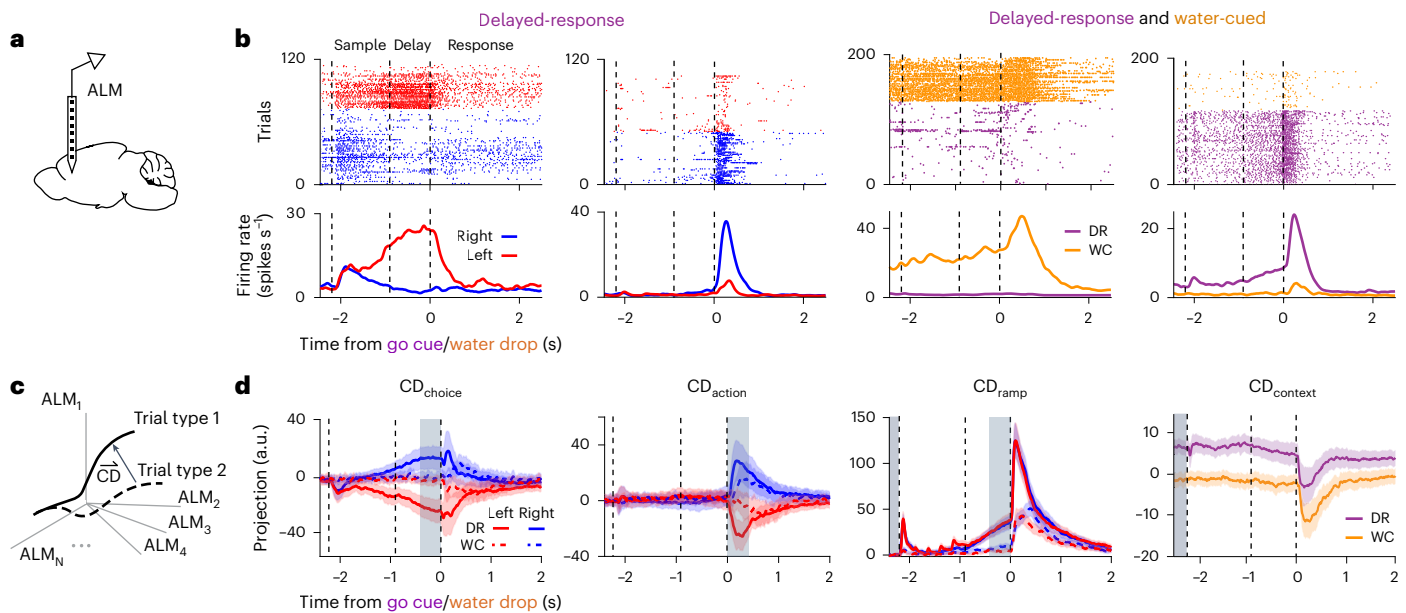


Fig. 2 | Single-cell activity and population dynamics encode task-relevant cognitive and motor processes. **a**, Schematic of silicon probe recordings in the ALM. **b**, Spike rasters (top) and peri-stimulus time histograms (PSTHs; bottom) for four example units in the ALM. Units show selectivity for trial type (left versus right) and context (DR versus WC) during all task epochs. PSTHs for DR and WC tasks shown on right are averages of equal numbers of left and right trials.

c, Schematic illustrating identification of CDs. CDs are defined as directions in state space that maximally separate activity between trajectories defined by trial types (CD_{choice} and CD_{action}), context ($CD_{context}$) or timepoints (CD_{ramp}). **d**, Projections along each of four defined CDs. Gray shaded regions indicate timepoints used for CD estimation. Mean and 95% confidence intervals of bootstrapped distributions are shown. spks, spikes.

We defined a CD for left–right choice, CD_{choice} , as the direction in state space that best differentiates activity on correct left versus right trials during the delay epoch of DR trials. Projections along CD_{choice} revealed that choice selectivity in the ALM, taken to represent a neural correlate of a motor plan^{32,37} or choice^{19,39}, emerges early in the sample epoch and continues to grow throughout the delay epoch before decaying at the onset of the response epoch (Fig. 2d). Animal's upcoming choice could be reliably decoded from projections along CD_{choice} across sessions (Extended Data Fig. 2c; area under the curve (AUC): 0.86 ± 0.11 , mean \pm s.d.). We next defined CD_{ramp} as the direction that captures trial-type-independent increases or decreases of activity between the onset of the stimulus and the go cue in DR trials^{32,38,39}. This non-selective ramping has been thought to represent an urgency or timing signal^{41,42}. Finally, we defined $CD_{context}$ as the direction that best differentiates activity during the ITI of DR and WC trials. Projections along $CD_{context}$ revealed that context (DR versus WC blocks) is encoded in the ALM throughout each trial (Fig. 2d) with strong modulation at the onset of the response epoch.

We also examined neural signals associated with the execution of movements. We defined an additional CD, CD_{action} , as the direction that captures trial-type-selective changes in neural activity that emerge after the go cue. Projections of neural activity along CD_{action} robustly encoded movement direction in the ALM, as expected. Together, these four CDs explained $53 \pm 11\%$ (mean \pm s.d.) of the variance in trial-averaged neural data (Extended Data Fig. 2b).

Uninstructed movements related to preparatory dynamics

We examined uninstructed movements in the DR task and their correlation with both task variables and preparatory dynamics (Fig. 3). We found that animals frequently performed uninstructed movements that were correlated with trial type and time within each trial (Fig. 3a). On single trials, upcoming choice could be decoded from uninstructed movements beginning immediately after the presentation of the sample tone, with accuracy increasing throughout the delay epoch (Fig. 3b).

Notably, uninstructed movements were also highly correlated with putative choice and urgency signals on a moment-by-moment

basis. The onset and magnitude of uninstructed movements often coincided with the onset and magnitude of choice selectivity and ramping (Fig. 3c,e). We used cross-validated multiple linear regression to predict projections along CD_{choice} and CD_{ramp} on a trial-by-trial basis from kinematic features of movement. Model predictions were often highly correlated with motor planning signals, indicating a close relationship between animals' uninstructed movements and preparatory dynamics (Fig. 3d,f; CD_{choice} : $R^2 = 0.41 \pm 0.23$, mean \pm s.d. across sessions, CD_{ramp} : $R^2 = 0.46 \pm 0.23$, $n = 25$ sessions), although this relationship was variable across animals, sessions (Fig. 3d,f, bottom row; CD_{choice} : minimum $R^2 = 0.002$, maximum $R^2 = 0.766$; CD_{ramp} : minimum $R^2 = 0.096$, maximum $R^2 = 0.853$) and in the kinematic features that were most predictive (Extended Data Fig. 3a–c). This variability is consistent with other studies⁴³ and highlights the need to assess the relationship between neural signals and related movements on a session-by-session basis.

We trained additional animals on a randomized delay task in which the timing of the go cue cannot be anticipated⁴⁴. Incorporating uncertainty into the length of the delay epoch leads to a qualitative change in population-level choice selectivity. Selectivity in CD_{choice} projections emerges earlier than in the standard fixed-delay DR task (Fig. 3h, top right; fixed: 0.61 ± 0.11 s from delay onset, $n = 25$ sessions; randomized: 0.16 ± 0.19 s, $P = 2.4 \times 10^{-12}$, $n = 19$ sessions, two-sided t -test; Fig. 3h, top left), suggesting that motor plans are prepared before the earliest possible go cue time⁴⁴. We found that the timing of uninstructed movements, too, shifted in an analogous fashion (Fig. 3h, bottom; fixed: 0.58 ± 0.43 s from delay onset; randomized: 0.19 ± 0.39 s, $P = 0.01$, two-sided t -test). Projections along CD_{choice} remained predictable from uninstructed movements on a single-trial level (Fig. 3i; average R^2 across sessions: 0.38 ± 0.20 , minimum $R^2 = 0.041$, maximum $R^2 = 0.734$, $n = 19$ sessions), providing further evidence of the tight link between preparatory dynamics and uninstructed movements.

Uninstructed movements related to behavioral context

We next compared uninstructed movements across task blocks in the two-context paradigm (Fig. 4). We found that animals perform qualitatively different uninstructed movements in each behavioral context,

despite both contexts requiring instructed movements to the same targets (Fig. 4a). Context could be decoded from kinematic features of movement across epochs, including during the ITI, in which contextual cues were absent (Fig. 4b, top). The timecourse of context decoding from neural population data and from movement kinematics was also similar (Fig. 4b, bottom).

We sought to understand whether trial-to-trial variability in the neural representation of context encoding could be predicted from uninstructed movements. We again trained a multiple linear regression model to predict single-trial projections along CD_{context} from kinematic features and found that, like putative choice and urgency dynamics, context-selective signals were associated with context-specific movements (Fig. 4c,d; $R^2 = 0.49 \pm 0.22$, mean \pm s.d across sessions, minimum $R^2 = 0.097$, maximum $R^2 = 0.751$, $n = 12$ sessions).

Taken together, these findings suggest that dynamics along CD_{choice} , CD_{ramp} and CD_{context} could be associated with uninstructed movements and/or the cognitive processes related to the anticipation and planning of upcoming instructed movements and encoding of task context.

Subspace decomposition of trial-averaged neural activity

We next examined whether neural signals encoding task variables could be isolated from movement-related signals, despite their correlation in time. We hypothesized that, if the neural dynamics associated with cognitive and motor processes are separable, their components in the ALM should be driven by distinct latent signals.

Previous theoretical work demonstrated a simple mechanism that explains how neural activity can vary dynamically while ensuring that specific variables remain stably encoded²⁸. In this formalism, a variable of interest may be decoded from a neural population, perhaps by a downstream circuit, through a linear transformation matrix, W , which could be taken to follow from a particular pattern of synaptic connectivity. If the dimensionality of neural activity is larger than the dimensionality of the output signals decoded (that is, if W is not square), then W must have a null space—dimensions in activity space along which neural activity can vary without changing the decoded output²⁸.

Applying this idea to recordings from the primate motor cortex demonstrated that motor planning dynamics are confined to the null space of a linear transformation between motor cortical activity and muscle activation²⁵, thereby allowing motor planning signals to evolve dynamically in a manner that is independent of motor output. The

orthogonal complement to this output-null subspace was termed the output-potent subspace; changes in activity along output-potent dimensions were then associated with changes in muscle activation²⁵.

Subsequent work identified analogous subspaces of neural activity in the primate motor cortex as the orthogonal subspaces that capture the most variance in neural activity during the delay and response epochs of a delayed-response task to capture preparatory and movement-related activity, respectively²⁷. This approach is effective when preparatory and movement-related dynamics are confined to distinct temporal epochs and when neural activity is exclusively related to motor preparation and motor execution, respectively, in those epochs. Conveniently, it avoids the need to measure muscle activation or any other descriptor of motor output. It also avoids the need to explicitly model the transformation between neural activity and motor output with a linear (or any other parametric) transformation. This transformation is generally unknown and can be highly nonlinear, particularly for movements directly controlled by central pattern generators, such as those that support much of the behavioral repertoire of rodents (for example, locomoting, breathing, whisking, licking, chewing, vocalizing and swallowing)^{45,46}.

We sought to determine if this computational framework could be used to isolate the neural dynamics associated with instructed and uninstructed movements from dynamics related to the encoding of choice, urgency and context in our paradigm. We identified subspaces that maximally capture variance in trial-averaged neural activity in the delay and response epochs of our DR trials²⁷ (Fig. 5a and Methods). We found that dynamics within ‘delay’ and ‘response’ subspaces were confined to the delay and response epochs, respectively, closely resembling results observed in primates²⁷ (Fig. 5b–d). This outcome is expected when trial-averaged dynamics in different epochs occupy subspaces that are largely orthogonal²⁷. However, activity within both subspaces remained highly correlated with uninstructed movements on single trials (Fig. 5b,e,f), suggesting that movement-related dynamics were not restricted to the ‘response’ subspace as intended. This result follows from the observation that planning-related and movement-related processes were not confined to distinct task epochs in our paradigm. Uninstructed movements were expressed across all task epochs (Figs. 1, 3 and 4).

Subspace decomposition of single-trial neural activity

To extend this formalism to the case where movements and cognitive processes are not strictly segregated into distinct task epochs, we

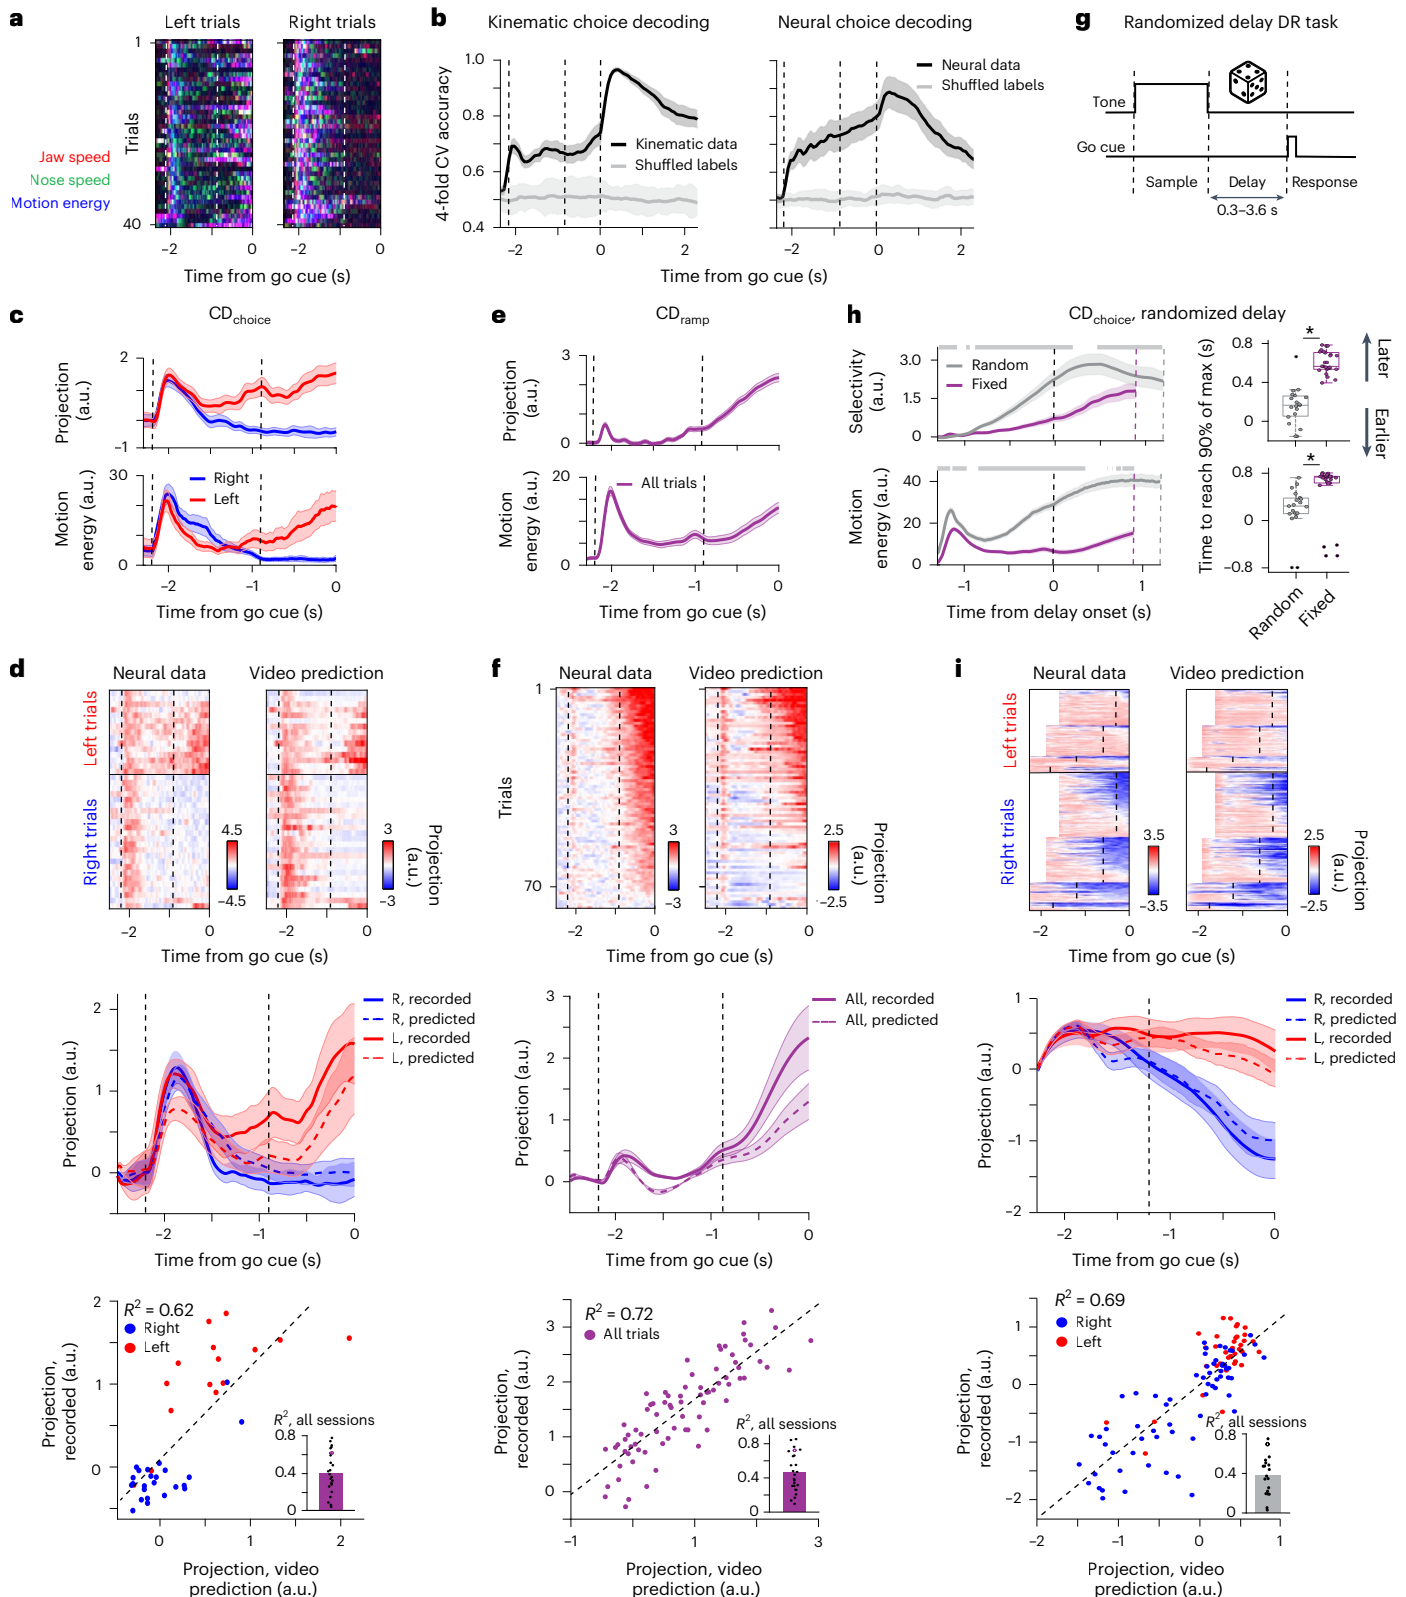
Fig. 3 | Uninstructed movements are tightly linked to putative planning dynamics.

a, Choice-dependent uninstructed movements. Overlaid trajectories from a random subset of DR trials in an example session. In this example, uninstructed movements are prominent on left, but not right, trials during the delay epoch. **b**, Choice (left versus right) decoding from kinematic features and neural population data. Lines with shaded regions depict mean and 95% confidence intervals of bootstrapped distributions. Gray lines denote shuffled choice labels. **c**, Top: average projection onto CD_{choice} on left (red) and right (blue) trials for the example session depicted in **a**. Bottom: average motion energy for all correct left and right trials. Lines with shaded regions depict mean and 95% confidence intervals across trials. **d**, Observed and video predictions of single-trial projections of neural activity onto CD_{choice} (Methods) for the same example session as in **a**. Top: heatmap of observed and predicted single-trial projections. Trials are sorted by the observed average projection magnitude in the late delay, with left and right trials sorted separately. Middle: trial-averaged CD_{choice} projections and predictions. Lines with shaded regions depict mean and 95% confidence intervals across trials. Bottom: scatter plot of the average delay epoch projection of neural data onto CD_{choice} versus corresponding video predictions. Dots denote single trials, and dashed line denotes linear fit. Inset: R^2 values for all sessions ($n = 25$ sessions). Open circle denotes the example session in **a**. **e**, Same as **c** but for CD_{ramp} (different example session). Left and right hit trials are grouped (purple). **f**, Same as **d** but for CD_{ramp} . Same session as **e**. **g**, Schematic of the randomized delay task. The delay epoch duration was

randomly selected from six values (Methods) with frequencies following an exponential distribution. **h**, Differences in choice selectivity and uninstructed movements between the randomized and fixed DR tasks. Top left: selectivity in CD_{choice} projections averaged across sessions for randomized delay (gray; $n = 19$ sessions) and fixed delay (purple; $n = 25$ sessions) tasks. Solid lines with shaded regions depict mean and 95% confidence intervals across sessions. Dashed vertical lines indicate time of delay onset (black), go cue for the fixed delay (purple) and/or go cue for the randomized delay (gray). Only trials with a delay duration of 1.2 s are shown for the randomized delay task for clarity. Gray bar at the top denotes timepoints where the slopes of the curves are significantly different ($P < 0.05$, paired two-sided t -test). Bottom left: same but for motion energy. Top right: time relative to delay onset that selectivity in CD_{choice} projections reaches 90% of its maximum value. Center line of the box denotes the median; bottom and top lines of the box denote the 25th and 75th percentiles, respectively; and whiskers denote the minimum and maximum data points not considered outliers. Colored dots, individual sessions; black dots, outlier sessions. Asterisk denotes significant difference ($P < 0.05$) between latencies in randomized versus fixed delay sessions ($P = 2.4 \times 10^{-12}$, two-sided t -test). Bottom right: same but for session-averaged motion energy ($P = 0.01$). **i**, Same as **d** for an example session of the randomized delay task. Top: trials with delay durations of 0.3 s, 0.6 s, 1.2 s and 1.8 s are shown. Middle and bottom: trials with delay duration of 1.2 s only. CV, cross-validation.

modified this analytical approach in two ways (see Supplementary Note 1 for a more detailed discussion). Because uninstructed movements vary considerably from trial to trial, we avoided analyses of trial-averaged data and focused on single-trial data. Second, rather than identifying subspaces capturing variance during the delay (when no movement is assumed) and response epochs, we instead annotated all timepoints across a session in which animals were moving and all timepoints at which animals were stationary. This could

be straightforwardly achieved using a threshold on motion energy (the magnitude of frame-by-frame changes in video). We found two orthogonal subspaces that best explained variance during all timepoints across a session at which animals were stationary and moving (Fig. 5h). We termed the resulting subspaces ‘movement-null’ and ‘movement-potent’ subspaces. The movement-potent subspace contains dynamics associated with movement. Dynamics that are observed in the absence of movements—those likely related to cognitive and



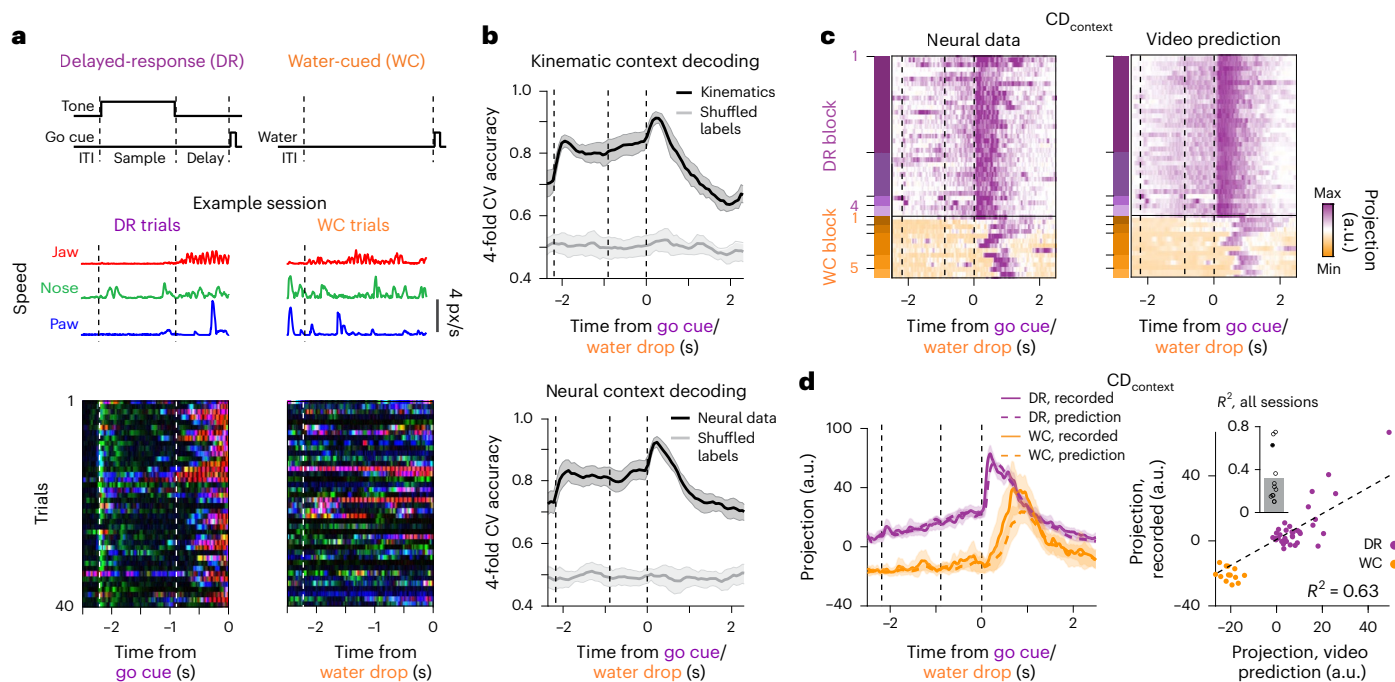


Fig. 4 | Uninstructed movements are closely related to the neural encoding of context. **a**, Schematic of tasks (top) and context-dependent uninstructed movements. Jaw/nose/paw movements on an example trial (middle) and overlaid trajectories of uninstructed movements across an example session (bottom) for DR task (left) and WC task (right). **b**, Context decoding from high-speed video (top) and neural data (bottom) as a function of time across all sessions. Gray lines denote shuffled context labels. **c**, Left: heatmap of single-trial projections of neural data onto CD_{context}. Right: CD_{context} projections predicted from video in an example session. The chronological DR or WC block within

the session is denoted by purple and orange rectangles, respectively, on the left of each plot. **d**, Left: trial-averaged projection of neural data onto CD_{context} and predictions from video (dashed lines) for the example session illustrated in **c**. Right: scatter plot of average projection of neural data onto CD_{context} during the ITI versus average video prediction. Dots denote single trials; dashed line denotes line of best fit. Inset: R^2 values for all sessions ($n = 12$ sessions). Filled circle denotes the example session in **c**. Lines with shaded regions depict mean and 95% confidence intervals of bootstrapped distributions throughout. CV, cross-validation.

other ‘internal’ processes—are contained within the movement-null subspace. Activity within the two subspaces explained 72% of the variance in single-trial neural activity (movement-null: $30 \pm 9.5\%$, mean \pm s.d.; movement-potent: $42 \pm 8.7\%$; $n = 25$ sessions). Neural dynamics in the movement-potent subspace were correlated with motion energy on single trials, as might be expected (Fig. 5i,l,m). Naively, movement-null subspace activity might be expected to be similarly anticorrelated with motion energy because that subspace is constructed to capture variance during periods of animal stationarity. This outcome was largely not observed. Rather, movement-null subspace activity did not display any consistent relationship with movement (Fig. 5i,l,m), suggesting that these ‘internal’ dynamics were not only prominent in the absence of movement but also persisted during movement, in line with the supposition that cognitive processes can persist during movement.

Uninstructed movements were commonly observed during the delay epoch of DR trials and were highly correlated with putative choice and urgency signals (Fig. 3c–h). To examine whether the movement-potent subspace might inadvertently capture choice or urgency dynamics that are correlated in time with movements, we estimated subspaces using trials from the WC context only. Here, the movement-potent subspace is determined in a context in which choice and urgency signals are absent¹³ (Fig. 2d and Extended Data Fig. 4c), precluding the possibility that dimensions of neural activity related to choice and urgency are inadvertently assigned to the movement-potent subspace due to their correlation with movement. This approach yielded subspaces that were highly similar (Extended Data Fig. 4d–f). In another control, examining movement-null and movement-potent subspaces estimated only using data recorded during the response epoch of both tasks, when planning and urgency

dynamics should be minimal, again yielded similar results (Extended Data Fig. 4g–i). Finally, we estimated the movement-null subspace as the subspace defined by the top principal components (PCs) of activity recorded during periods of stationarity and then determined the movement-potent subspace as the top PCs of dynamics not already captured in the movement-null subspace. This latter procedure, which is highly conservative in avoiding the spurious assignment of internal dynamics to the movement-potent subspace, again yielded similar results (Extended Data Fig. 4j–l and Methods). We also confirmed that the choice of subspace dimensionality had minimal impact on the dynamics contained within each (Extended Data Fig. 6). Thus, the movement-potent subspace captures movement-related signals and not simply dynamics related to internal processes that are correlated in time with movements.

Animals may think and move simultaneously. Thus, it follows that internal dynamics contained in the movement-null subspace may be present during the response epoch. Additionally, uninstructed movements occur during both the sample and delay epochs, and so we should observe dynamics within the movement-potent subspace during both task epochs. As expected, activity within movement-null and movement-potent subspaces captured variance across epochs (Fig. 5k; see also Supplementary Note 2). Interestingly, the proportion of units with activity largely confined to the movement-potent subspace was substantially larger than expected by chance (Fig. 5n and Extended Data Fig. 5b; $P = 6 \times 10^{-19}$, $n = 483$ single units; Methods) and larger than when subspaces were identified from trial-averaged data here (Fig. 5g) and in non-human primates²⁵. This observation suggests the identification of neuronal populations engaged solely in motor processes that could not be identified without properly accounting for uninstructed movements.

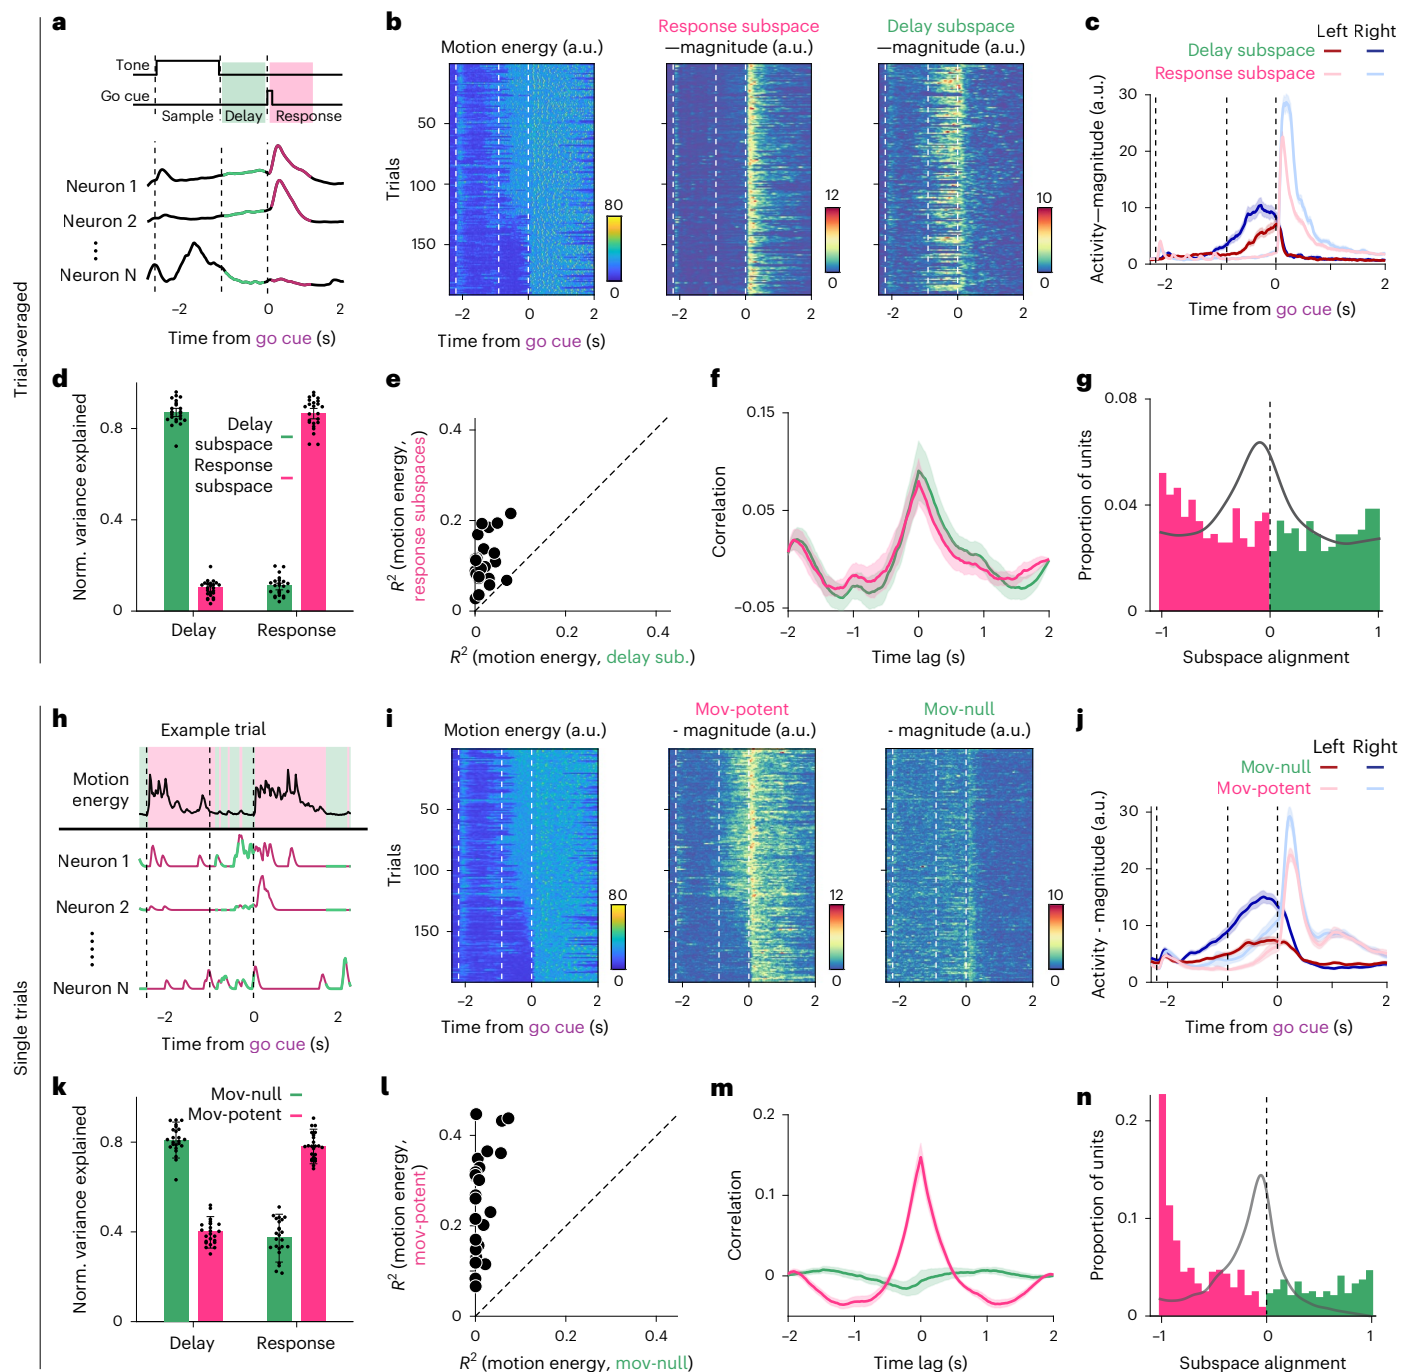


Fig. 5 | Subspace decomposition of neural activity using trial-averaged and single-trial data. **a**, Schematic of the approach for estimating delay (green) and response (pink) subspaces from trial-averaged neural data. Delay and response subspaces are determined as the orthogonal subspaces that maximally contain the variance of neural activity in the delay and response epochs, respectively. Correct trials from the DR context were used to identify subspaces. **b**, Activity during DR trials within each subspace for an example session. Left: motion energy across trials. Middle: sum-squared magnitude of activity in the response subspace. Right: sum-squared magnitude of activity in the delay subspace. Trials are sorted by average delay epoch motion energy. **c**, Sum-squared magnitude of activity in delay and response subspaces during DR lick-left and lick-right trials for an example session. Mean and s.e.m. across trials are shown. **d**, Normalized (Norm.) variance explained of the neural data during DR trials by the activity in the delay and response subspaces during the delay and response epochs. Delay and response subspaces selectively capture activity in the delay and response epochs, as expected. Points indicate sessions; bar height indicates the mean across sessions; and error bars indicate s.d. across session ($n = 25$ sessions). **e**, Variance explained (R^2) of motion energy by the sum-squared magnitude of

activity in the delay and response subspaces using single-trial DR and WC data. Each point represents the mean across trials for a session ($n = 25$ sessions). **f**, Cross-correlation between motion energy and activity in the delay and response subspaces using single-trial DR and WC data. Lines indicate mean across sessions, and shaded region represents s.e.m. across sessions ($n = 25$ sessions). **g**, Subspace alignment for single units across all sessions ($n = 483$ single units; Methods). Values closer to 1 indicate that more variance of a unit's activity is contained within the preparatory subspace. Chance distribution is shown as a gray line. **h**, Schematic of the approach for estimating movement-null and movement-potent subspaces from single-trial data. Motion energy is used to annotate when an animal is moving (pink) or stationary (green). These labels are then applied to single-trial neural data. Movement-null and movement-potent subspaces are determined as the orthogonal subspaces that maximally capture the variance of neural activity during periods of quiescence and movement, respectively. Trials from both DR and WC contexts were used to identify subspaces. **i–n**, Same as **b–f** but using the single-trial approach for estimating movement-null and movement-potent subspaces. Norm., normalized; Mov, movement.

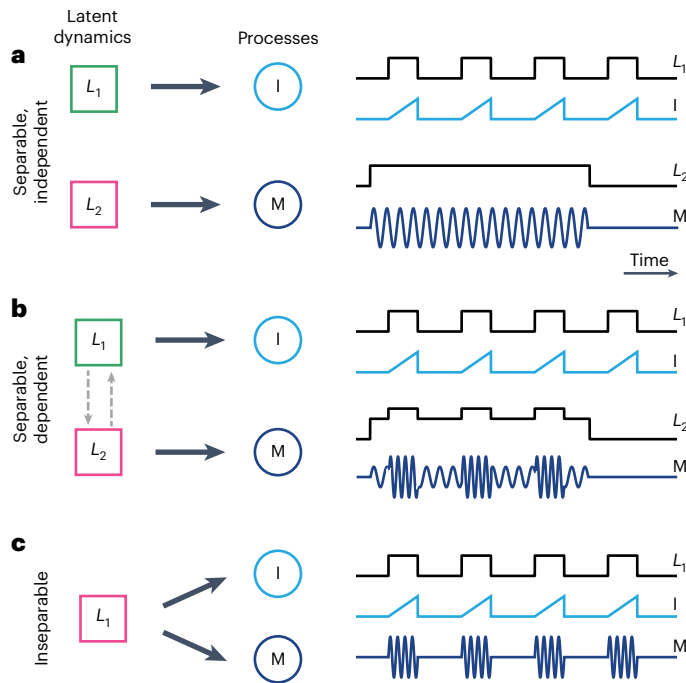


Fig. 6 | Schematic of potential relationships between internal and movement-related dynamics. **a**, Internal and movement-related dynamics that are independent and separable. Left: schematic depicting an internal process (I) and a motor process (M) that are each governed by separate latent dynamics, L_1 and L_2 , respectively. L_1 and L_2 evolve within the movement-null (green) and movement-potent (pink) subspaces, respectively. Right: cartoon time series of separable, independent latent dynamics (L_1 and L_2) and related processes (I and M). **b**, Same as **a** in the case of latent dynamics that are separable but dependent. I and M may be loosely correlated in time. **c**, Same as **a** and **b** in the case of inseparable processes governed by a single set of latent dynamics (L_1).

Internal and movement-related dynamics during motor planning

The movement-potent subspace may contain at least three classes of movement-related neural dynamics: (1) motor commands (and efference copies); (2) sensory feedback related to movements; and (3) internal dynamics that are present exclusively during periods of movement. That is, if a particular internal process is ‘embodied’, in the sense that it is mediated by the same latent dynamics responsible for associated movements and, thus, always observed in the presence of those movements, then we would expect to find those latent dynamics wholly contained within the movement-potent subspace (Fig. 6c). Dynamics supporting internal processes independent of movements should be contained within the movement-null subspace (Fig. 6a). Between these extremes, we expect to observe dynamics within both subspaces when a particular internal process (with associated latent dynamics in the movement-null subspace) biases the expression of movements (with associated latent dynamics in the movement-potent subspace) (Fig. 6b). In that latter case, internal processes may appear correlated with movements despite being largely mediated by distinct latent dynamics.

To evaluate these possibilities, we examined whether putative cognitive signals (Fig. 2) were contained within the movement-potent and/or movement-null subspaces. Prominent choice selectivity was contained in both subspaces. The existence of choice dynamics in the movement-null subspace indicates an internal representation of choice that is separable from dynamics related to the execution of choice-related uninstructed movements (Fig. 6b). The average time-courses of choice dynamics in these two subspaces displayed subtle yet notable differences, as did their expression on error trials. Selectivity emerged quickly after stimulus presentation in the movement-null subspace and did not change significantly through the delay epoch (Fig. 7a;

$P > 0.05$, paired t -test comparing selectivity during the last 100 ms of sample and delay epochs, $n = 25$ sessions; Extended Data Fig. 7c), consistent with an internal representation driven by sensory input. On error trials, selectivity in the movement-null subspace initially followed the same trajectory as correct trials but decayed with stimulus offset. In the movement-potent subspace, in contrast, selectivity increased slowly during the sample epoch and continued to increase monotonically during the delay epoch ($P = 3 \times 10^{-13}$, paired t -test comparing selectivity during the last 100 ms of sample and delay epochs, $n = 25$ sessions; $P = 1 \times 10^{-5}$, comparing delay epoch selectivity across subspaces; Extended Data Fig. 7b), and no significant movement-potent subspace selectivity emerged on error trials. These observations suggest that sensory stimuli initially drive appropriate dynamics within the movement-null subspace on error trials but that choice is not encoded stably and does not engage movement-potent dynamics related to uninstructed movements.

Although the trajectories of choice dynamics were broadly similar across subspaces in trial-averaged data (Fig. 7a), the existence of choice dynamics in both subspaces implies that they must differ markedly on single trials. The onset and magnitude of activity along the component of CD_{choice} within the movement-potent subspace tracked trial-type-selective motion energy on a moment-by-moment basis (Fig. 7b,c), as expected, whereas there was no consistent relationship between motion energy and activity along the component of CD_{choice} in the movement-null subspace. After the go cue, transient responses accompanying movement initiation were absent from the movement-null subspace (Extended Data Fig. 7a), further suggesting the existence of choice dynamics that are separable from ongoing movement dynamics. Notably, we found that many single units contribute to either the movement-null or the movement-potent subspace representations of CD_{choice} , but not both (Fig. 7d and Extended Data Fig. 6d), providing additional evidence of dynamics that are not only separable but also encoded by distinct populations of neurons (Fig. 7d).

In contrast, ramping dynamics were mostly confined to the movement-potent subspace (Fig. 7e)—perhaps surprising given that ramping activity has been interpreted as an internal urgency or timing signal^{41,42}. This observation remained consistent when estimating movement-null and movement-potent subspaces using only WC trials or using only the response epoch of WC and DR trials (Extended Data Fig. 8a), when ramping dynamics were absent, confirming that ramping dynamics were not observed in the movement-potent space because of the mis-assignment of internal ramping dimensions to the movement-potent subspace because of their correlation in time with movement. We further searched for the existence of movement-null subspace ramping dynamics by explicitly identifying the dimension within the movement-null and movement-potent subspaces that maximized ramping dynamics—which could be different than the dimension that maximized ramping within the full activity space—and again failed to identify a prominent ramping signal in the movement-null subspace (Extended Data Fig. 8b). The relative paucity of ramping dynamics in the movement-null subspace suggests that they are not readily separable from associated movements in the ALM in our behavioral paradigm (Discussion).

Persistent movement-null subspace representation of context

Next, we examined projections of activity along CD_{context} in each subspace (Fig. 8a and Extended Data Fig. 9). We found robust contextual selectivity in both subspaces, consistent with the interpretation that the ALM contains both an internal representation of context and activity related to context-dependent movements. The transient, response epoch dynamics observed along CD_{context} (Fig. 2d, far right) were contained within the movement-potent subspace (Fig. 8a), likely indicative of subtle context-specific differences in instructed movements. We then compared context selectivity (the difference in projections onto CD_{context} on DR versus WC trials) on trials with high motion energy during the ITI (‘High move’ trials) and trials with

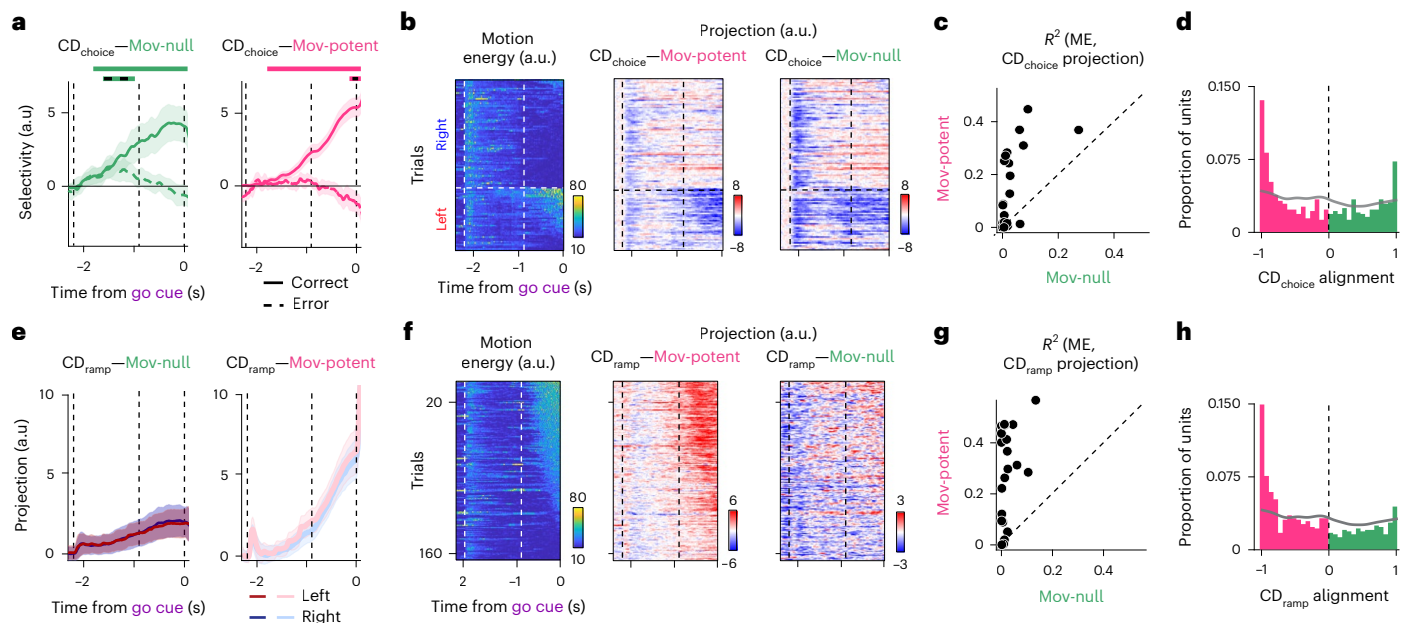


Fig. 7 | Subspace decomposition allows for the re-examination of population measures of motor planning. **a**, Selectivity (neural projections on lick-right trials minus lick-left trials) of movement-null (left) and movement-potent (right) subspace activity when projected along CD_{choice}. Mean and 5–95% confidence interval of the bootstrap distribution for correct (solid) and error (dashed) trials are shown. Horizontal bars (solid lines, correct trials; solid lines with black dashes, incorrect trials) above plot indicate when the selectivity trace is significantly different from zero ($P < 0.05$, one-sided test, bootstrap). **b**, Projections along CD_{choice} in the movement-potent subspace (middle), but not the movement-null subspace (right), follow the timecourse and magnitude of motion energy (left) on single lick-left and lick-right trials (example session). **c**, Variance explained (R^2) of motion energy by projections along CD_{choice} in the movement-null or movement-potent subspace on single trials. Each point is the average across all trials in a session. **d**, Distribution of

single unit alignment with CD_{choice} ($n = 483$ single units; Methods). Distribution of movement-potent tuned units (alignment ≤ -0.8) was significantly different than expected by chance ($P = 1 \times 10^{-13}$; Methods). Similarly, distribution of movement-null tuned units (alignment ≥ 0.8) was significantly different than expected by chance ($P = 7 \times 10^{-12}$). Chance distribution is shown as a gray line. **e**, Projection of movement-null (left) and movement-potent (right) subspace activity along CD_{ramp} on lick-left and lick-right trials. Mean and 95% confidence interval of the bootstrap distribution for correct trials are shown. **f–h**, Same as **b–d** for CD_{ramp}. **h**, Distribution of single unit alignment with CD_{ramp} ($n = 483$ single units). Distribution of movement-potent tuned units (alignment ≤ -0.8) was significantly different than expected by chance ($P = 6 \times 10^{-14}$). Similarly, distribution of movement-null tuned units (alignment ≥ 0.8) was significantly different than expected by chance ($P = 1 \times 10^{-12}$; Methods). ME, motion energy.

little or no ITI motion energy ('Low move' trials). Context selectivity within the movement-null subspace was indistinguishable in 'High move' and 'Low move' trials. In contrast, context selectivity within the movement-potent subspace was reduced by 72% on 'Low move' trials (Fig. 8b,c; full population: selectivity reduced by 0.101 ± 0.092 , mean \pm s.d., $P = 0.003$; movement-null: 0.026 ± 0.050 , $P = 0.095$; movement-potent: 0.098 ± 0.082 , $P = 1.6 \times 10^{-3}$, paired t -test; Δ selectivity between 'High move' and 'Low move' trials in movement-null versus movement-potent: $P = 9.5 \times 10^{-3}$, paired t -test). The magnitude and timing of the reduction in movement-potent subspace context selectivity mirrored the reduction in motion energy on 'Low move' trials (Fig. 8b, bottom). Together, these observations demonstrate that the movement-null, but not the movement-potent, subspace contains a stable representation of context that is unchanged in the presence of both instructed and uninstructed movements. Furthermore, we found that largely distinct populations of single neurons contribute preferentially to the movement-null and movement-potent representations of context (Fig. 8d). Thus, the analysis of context-dependent dynamics without subspace decomposition would have spuriously conflated separate latent dynamics, encoded by different populations of neurons, likely responsible for internal representations of context and related context-dependent movements.

Tongue kinematics encoded in the movement-potent subspace Finally, we asked whether trial-type-selective activity after the go cue relates to movement, according to previous suppositions^{35,47} (Extended Data Fig. 10). Left–right selectivity along CD_{action} existed within both the

movement-null and the movement-potent subspaces (Extended Data Fig. 10a), although the magnitude of selectivity in the movement-null subspace was substantially smaller in magnitude. Interestingly, on error trials, activity along CD_{action} flipped to resemble that of the other trial type in the movement-null subspace but flipped asymmetrically within the movement-potent subspace. This asymmetry closely corresponded to an asymmetry in tongue angle, with incorrect movements directed to less extreme angles when directed to the left, on average (Extended Data Fig. 10b,c). The component of CD_{action} within the movement-potent subspace also better tracked moment-to-moment changes in tongue angle during the response epoch (Extended Data Fig. 10d,e; $P = 2 \times 10^{-8}$, paired t -test between movement-null and movement-potent variance explained of tongue angle, $n = 25$ sessions). The similarity between tongue angle, a kinematic feature, and dynamics only within the movement-potent representation of CD_{action} is notable, as our analytical approach to the identification of subspaces does not incorporate any kinematic information. Action dynamics within the movement-null subspace could relate to an internal representation of the motor plan, or intention of the animal, to respond to one reward port or the other.

Discussion

We aimed to understand whether the neural correlates of three cognitive variables commonly examined in sensorimotor decision-making tasks—choice, urgency and context—were separable from the neural encoding of related uninstructed movements in the mouse ALM. We addressed this question by adapting an analytical formalism in which neural activity is decomposed into orthogonal subspaces—here, a

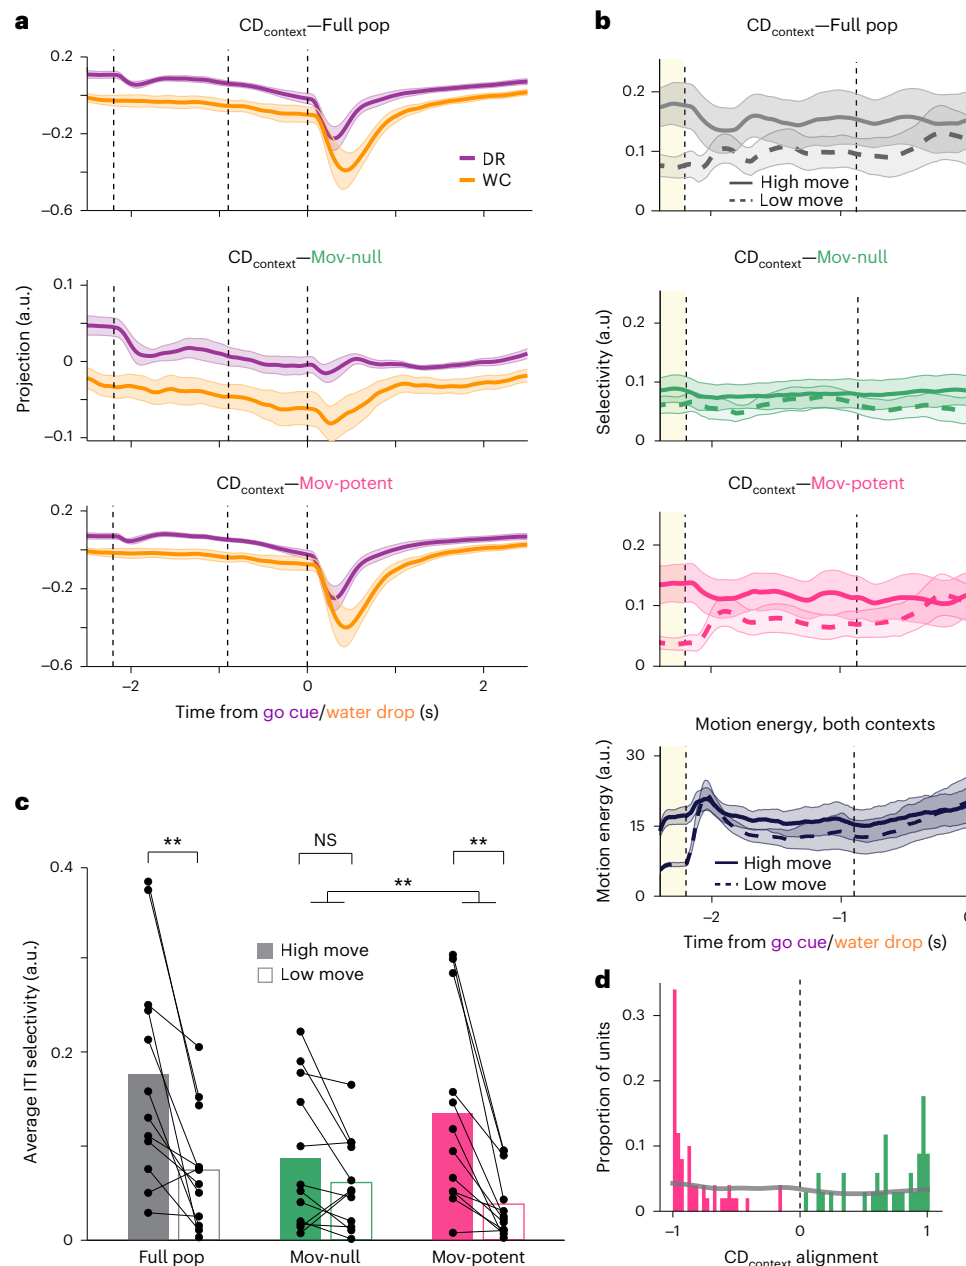


Fig. 8 | A persistent, cognitive representation of context in the movement-null subspace. a, Projections along $CD_{context}$ on DR hit trials (purple) and WC hit trials (orange) within the full population (top), movement-null subspace (middle) or the movement-potent subspace (bottom). Solid lines and shaded area denote the mean projection and 95% confidence interval across sessions ($n = 12$ sessions). Note that transient response epoch dynamics are absent in the movement-null subspace. **b**, Top three panels: selectivity in $CD_{context}$ projections (DR–WC trials) on ‘High move’ (solid lines) or ‘Low move’ (dashed) trials. Lines denote the mean selectivity across sessions ($n = 12$) in the full space of neural activity (top), activity within the movement-null subspace (second) and movement-potent subspace (third). Bottom panel: average motion energy across sessions on ‘High move’ or ‘Low move’ trials. DR and WC trials are grouped together for each session. Shaded area, 95% confidence interval across sessions. Yellow region denotes ITI used to define ‘High

move’ and ‘Low move’ trials. **c**, Average $CD_{context}$ selectivity during the ITI (yellow shaded region in **b**) for ‘High move’ (filled bars) versus ‘Low move’ (open bars) trials. Dots denote single sessions; bars denote mean across sessions ($n = 12$ sessions). Asterisks denote significant differences ($**P < 0.01$) in $CD_{context}$ selectivity (P values from left to right comparing ‘High move’ versus ‘Low move’ trials, full population: $P = 0.003$, movement-null: $P = 0.095$, movement-potent: $P = 1.6 \times 10^{-3}$; Δ selectivity between ‘High move’ and ‘Low move’ trials in movement-null versus movement-potent: $P = 9.5 \times 10^{-3}$, paired two-sided t -test). **d**, Distribution of single unit alignment with $CD_{context}$ (Methods). Distribution of movement-potent tuned units (alignment ≤ -0.8) was significantly different than expected by chance ($P = 2 \times 10^{-8}$; Methods). Similarly, distribution of movement-null tuned units (alignment ≥ 0.8) was significantly different than expected by chance ($P = 9 \times 10^{-6}$, $n = 84$ single units). Chance distribution is shown as a gray line. NS, not significant; pop, population.

movement-potent subspace containing dynamics related to the execution of movements and a complementary movement-null subspace containing dynamics related to cognitive and other internal processes^{25,27,28}. Extending upon this framework allowed us to consider single-trial neural data recorded in the presence of uninstructed movements that can vary dramatically across trials in their timing and identity (Fig. 5).

Separability of cognitive dynamics from movement-dynamics Using this approach, we demonstrated that the ALM contains neural dynamics related to encoding of choice and context that could be readily separated from movements, despite both cognitive variables being strongly correlated with movements in time (Figs. 7 and 8). Choice-selective signals were present in both movement-null and

movement-potent subspaces (Fig. 7a), consistent with the interpretation that a cognitive representation of choice within the movement-null subspace biases the probability and identity of uninstructed movements in a choice-selective manner. Context was stably encoded within the movement-null subspace during all trial epochs and the ITI (Fig. 8a–c). Although context was strongly represented in the movement-potent subspace as well, these dynamics were significantly reduced in the absence of movement and were also modulated by animals' instructed movements. These results suggest a persistent internal representation of context in the movement-null subspace in addition to distinct movement-potent-subspace dynamics related to context-dependent movements, again indicative of related, but separable, latent dynamics (Fig. 6b).

Surprisingly, ramping dynamics proposed to underlie urgency (or timing)^{41,42} were principally represented in the movement-potent subspace (Fig. 7e), indicating an absence of an internal representation of urgency that is separable from movements in the ALM. In contrast to these cognitive variables, encoding of kinematic features of movements, which were not used to determine subspaces, were confined to the movement-potent subspace (Extended Data Fig. 10).

That many individual units contribute solely to one subspace (Figs. 5n, 7d,h and 8d) suggests that the complementary subspaces that we identified could map to distinct cell types within the underlying cortical circuit^{40,48}. The suggestion that individual neurons appear to code for only either internal or movement-related variables when uninstructed movements are accounted for (Figs. 5n, 7d,h and 8d) underscores the importance of properly considering movements in future work focused on the roles of functionally, anatomically and/or transcriptomically defined cell types in neural computation.

Subspace decomposition

The approach to subspace decomposition presented here represents a means for assessing the issue of separability and for isolating separable dynamics into movement-related and internal components for further analyses across a wide range of experimental preparations. This method only requires annotation of when animals are moving and uses a robust analytical approach that does not require fine-tuning of parameters.

Alternative methods have been proposed to evaluate neural dynamics associated with cognitive processes in the presence of movement (Supplementary Note 1). The most common approach, which assumes separability of cognitive and motor dynamics, is to parameterize movements as fully as possible and 'regress them out'. However, parameterization of movement, particularly the orofacial and postural movements that have been associated with strong, brain-wide neural dynamics^{3,17,18}, can be challenging. Furthermore, this approach typically assumes a linear relationship between neural activity and kinematic (or electromyographic) variables—an assumption unlikely to hold for common movements mediated by central pattern generators and other highly nonlinear mechanisms^{45,46}.

Determination of movement-null and movement-potent subspaces as the orthogonal subspaces containing the most variance in neural activity during periods of stationarity and movement is conservative in the assignment of dynamics to the movement-null subspace. It is more likely that dynamics associated with internal processes are mis-assigned to the movement-potent subspace than vice versa. For example, dynamics associated with 'embodied' cognitive processes that always occur during periods of movement will be represented only in the movement-potent subspace (Fig. 6c). Thus, the dynamics within the movement-null subspace are highly likely to indicate signals related to internal processes. Straightforward variations of this approach can be used to determine subspaces in a manner that is more conservative in assigning dynamics to the movement-potent subspace—for example, estimating subspaces from response epoch data ('instructed' movements), from WC trials (where uninstructed

movements occur in the absence of choice and urgency signals) or simply through PC decomposition of activity recorded during periods of stationarity (Extended Data Fig. 4j–l). Nevertheless, the similarity of results obtained using these analytical variations argue against the possibility that cognitive dynamics were inadvertently assigned to the movement-potent subspace here.

Despite its robustness, our approach to subspace decomposition makes several simplifying assumptions. First, movements of the posterior torso, hindlimbs or tail of the animal were not routinely captured by videography. It remains possible, albeit unlikely, that animals routinely perform movements confined to these portions of the body without concomitant movement of the forepaws, head, neck, face or whiskers¹⁵. Periods of stationarity were no doubt algorithmically classified imperfectly and could have contained subtle movements below our threshold for detection. Second, we did not consider potential time lags between motion energy captured by videography and the dynamics associated with motor and/or sensory signals⁴³. However, these time lags, on the order of a few tens of milliseconds, should be much shorter than the timescale of dynamics explored in this study. Third, the set of signals related to movement and/or cognitive processes may also be better described as occupying nonlinear manifolds rather than linear subspaces⁴⁹. Subspaces (or manifolds) could also shift dynamically during behavior after changes in the activation of upstream or downstream brain regions^{50,51} as a result of changes in context⁵². Considering this additional complexity may improve estimates of the dynamics associated with specific neural processes.

Variable relationships between cognition and movement

The dynamics associated with some cognitive processes and movements may be largely independent (Fig. 6a)—perhaps in the case of the locomotor patterns of an individual walking down the street while thinking about what they will cook for dinner. On the other hand, some cognitive processes may be embodied, in the sense that they are tightly linked to externally observable changes in the body (Fig. 6c), such as the relationship between arousal and changes in pupil diameter⁵³. Between these extremes may be movements that are correlated with, but mediated by neural dynamics that are separable from, those supporting cognitive and other internal processes (Fig. 6b). Regardless of their relationship, understanding the separability of cognitive processes and associated movements in a wide variety of experimental settings will be essential for the interpretation of neural dynamics observed during behavior. In particular, methods for properly interpreting cognitive signals in the presence of related movements will be vital for increasingly common experimental paradigms examining freely moving animals performing complex, naturalistic behaviors^{16,54–57}.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41593-024-01859-1>.

References

- Alexander, G. E. & Crutcher, M. D. Neural representations of the target (goal) of visually guided arm movements in three motor areas of the monkey. *J. Neurophysiol.* **64**, 164–178 (1990).
- Cisek, P. & Kalaska, J. F. Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* **45**, 801–814 (2005).
- Chen, S. et al. Brain-wide neural activity underlying memory-guided movement. *Cell* **187**, 676–691 (2024).
- Zimnik, A. J. & Churchland, M. M. Independent generation of sequence elements by motor cortex. *Nat. Neurosci.* **24**, 412–424 (2021).

5. Ames, K. C., Ryu, S. I. & Shenoy, K. V. Simultaneous motor preparation and execution in a last-moment reach correction task. *Nat. Commun.* **10**, 2718 (2019).
6. Fuster, J. M. & Alexander, G. E. Neuron activity related to short-term memory. *Science* **173**, 652–654 (1971).
7. Erlich, J. C., Bialek, M. & Brody, C. D. A cortical substrate for memory-guided orienting in the rat. *Neuron* **72**, 330–343 (2011).
8. Tanji, J. & Evarts, E. V. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *J. Neurophysiol.* **39**, 1062–1068 (1976).
9. Guo, Z. V. et al. Flow of cortical activity underlying a tactile decision in mice. *Neuron* **81**, 179–194 (2014).
10. Wallis, J. D., Anderson, K. C. & Miller, E. K. Single neurons in prefrontal cortex encode abstract rules. *Nature* **411**, 953–956 (2001).
11. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
12. Freedman, D. J., Riesenhuber, M., Poggio, T. & Miller, E. K. Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* **291**, 312–316 (2001).
13. Zagha, E. The importance of accounting for movement when relating neuronal activity to sensory and cognitive processes. *J. Neurosci.* **42**, 1375–1382 (2022).
14. Drew, P. J., Winder, A. T. & Zhang, Q. Twitches, blinks, and fidgets: important generators of ongoing neural activity. *Neuroscientist* **25**, 298–313 (2019).
15. Hulsey, D., Zumwalt, K., Mazzucato, L., McCormick, D. A. & Jaramillo, S. Decision-making dynamics are predicted by arousal and uninstructed movements. *Cell Rep.* **43**, 113709 (2024).
16. Tremblay, S., Testard, C., DiTullio, R. W., Inchauspé, J. & Petrides, M. Neural cognitive signals during spontaneous movements in the macaque. *Nat. Neurosci.* **26**, 295–305 (2023).
17. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677–1686 (2019).
18. Stringer, C. Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
19. Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
20. Mangin, E. N., Chen, J., Lin, J. & Li, N. Behavioral measurements of motor readiness in mice. *Curr. Biol.* **33**, 3610–3624 (2023).
21. Terada, S.-I., Kobayashi, K. & Matsuzaki, M. Transition of distinct context-dependent ensembles from secondary to primary motor cortex in skilled motor performance. *Cell Rep.* **41**, 111494 (2022).
22. Lowet, E. et al. Enhanced neural processing by covert attention only during microsaccades directed toward the attended stimulus. *Neuron* **99**, 207–214 (2018).
23. Popescu, S. T. & Wexler, M. Spontaneous body movements in spatial cognition. *Front. Psychol.* **3**, 136 (2012).
24. Hajnal, J. V. et al. Artifacts due to stimulus correlated motion in functional imaging of the brain. *Magn. Reson. Med.* **31**, 283–291 (1994).
25. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* **17**, 440–448 (2014).
26. Cisek, P. & Kalaska, J. F. Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **33**, 269–298 (2010).
27. Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nat. Commun.* **7**, 13239 (2016).
28. Druckmann, S. & Chklovskii, D. B. Neuronal circuits underlying persistent representations despite time varying activity. *Curr. Biol.* **22**, 2095–2103 (2012).
29. Galiñanes, G. L., Bonardi, C. & Huber, D. Directional reaching for water as a cortex-dependent behavioral framework for mice. *Cell Rep.* **22**, 2767–2783 (2018).
30. Mathis, A. et al. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* **21**, 1281–1289 (2018).
31. Mayrhofer, J. M. et al. Distinct contributions of whisker sensory cortex and tongue-jaw motor cortex in a goal-directed sensorimotor transformation. *Neuron* **103**, 1034–1043 (2019).
32. Inagaki, H. K. et al. Neural algorithms and circuits for motor planning. *Annu. Rev. Neurosci.* **45**, 249–271 (2022).
33. Xu, D. et al. Cortical processing of flexible and context-dependent sensorimotor sequences. *Nature* **603**, 464–469 (2022).
34. Yin, X., Wang, Y., Li, J. & Guo, Z. V. Lateralization of short-term memory in the frontal cortex. *Cell Rep.* **40**, 111190 (2022).
35. Li, N., Chen, T.-W., Guo, Z. V., Gerfen, C. R. & Svoboda, K. A motor cortex circuit for motor planning and movement. *Nature* **519**, 51–56 (2015).
36. Chen, T.-W., Li, N., Daie, K. & Svoboda, K. A map of anticipatory activity in mouse motor cortex. *Neuron* **94**, 866–879 (2017).
37. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: a dynamical systems perspective. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
38. Li, N., Daie, K., Svoboda, K. & Druckmann, S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* **532**, 459–464 (2016).
39. Yang, W., Tipparaju, S. L., Chen, G. & Li, N. Thalamus-driven functional populations in frontal cortex support decision-making. *Nat. Neurosci.* **25**, 1339–1352 (2022).
40. Economo, M. N. et al. Distinct descending motor cortex pathways and their roles in movement. *Nature* **563**, 79–84 (2018).
41. Narayanan, N. S. Ramping activity is a cortical mechanism of temporal control of action. *Curr. Opin. Behav. Sci.* **8**, 226–230 (2016).
42. Cisek, P., Puskas, G. A. & El-Murr, S. Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* **29**, 11560–11571 (2009).
43. Wang, Z. A. et al. Not everything, not everywhere, not all at once: a study of brain-wide encoding of movement. Preprint at *bioRxiv* <https://doi.org/10.1101/2023.06.08.544257> (2023).
44. Inagaki, H. K., Fontolan, L., Romani, S. & Svoboda, K. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* **566**, 212–217 (2019).
45. Moore, J. D., Kleinfeld, D. & Wang, F. How the brainstem controls orofacial behaviors comprised of rhythmic actions. *Trends Neurosci.* **37**, 370–380 (2014).
46. Arber, S. & Costa, R. M. Networking brainstem and basal ganglia circuits for movement. *Nat. Rev. Neurosci.* **23**, 342–360 (2022).
47. Svoboda, K. & Li, N. Neural mechanisms of movement planning: motor cortex and beyond. *Curr. Opin. Neurobiol.* **49**, 33–41 (2018).
48. Christensen, A. J., Ott, T. & Kepecs, A. Cognition and the single neuron: how cell types construct the dynamic computations of frontal cortex. *Curr. Opin. Neurobiol.* **77**, 102630 (2022).
49. Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A. Neural manifolds for the control of movement. *Neuron* **94**, 978–984 (2017).
50. Dahmen, D. et al. Strong and localized recurrence controls dimensionality of neural activity across brain areas. Preprint at *bioRxiv* <https://doi.org/10.1101/2020.11.02.365072> (2023).
51. Morales-Gregorio, A. et al. Neural manifolds in V1 change with top-down signals from V4 targeting the foveal region. *Cell Rep.* **43**, 114371 (2024).
52. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).

53. Joshi, S. & Gold, J. I. Pupil size as a window on neural substrates of cognition. *Trends Cogn. Sci.* **24**, 466–480 (2020).
54. Pereira, T. D., Shaevitz, J. W. & Murthy, M. Quantifying behavior to understand the brain. *Nat. Neurosci.* **23**, 1537–1549 (2020).
55. Miller, C. T. et al. Natural behavior is the language of the brain. *Curr. Biol.* **32**, R482–R493 (2022).
56. Datta, S. R., Anderson, D. J., Branson, K., Perona, P. & Leifer, A. Computational neuroethology: a call to action. *Neuron* **104**, 11–24 (2019).
57. Dennis, E. J. et al. Systems neuroscience of natural behaviors in rodents. *J. Neurosci.* **41**, 911–919 (2021).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature America, Inc. 2025

Methods

Animals

This study used data collected from 17 mice; both male and female animals between 8 weeks and 11 weeks of age were used. Six animals were used for the two-context task and were either C57BL/6J (The Jackson Laboratory (JAX), 000664) or VGAT-ChR2-EYFP (–/–) (bred in-house; breeders: JAX, 014548). An additional three animals, either C57BL/6J or VGAT-ChR2-EYFP (–/–), were trained only on the DR task. A separate cohort of four animals was used in the randomized delay task (VGAT-ChR2-EYFP (+/–) or VGAT-ChR2-EYFP (–/–)). Finally, four VGAT-ChR2-EYFP (+/–) animals were used in optogenetics experiments. Mice were housed in a 12-h reverse dark/light cycle room (ambient temperature of 68–79 °F (20–26 °C); humidity of 30–70%) with ad libitum access to food. Access to water was restricted during behavioral and electrophysiology experiments (see ‘Mouse behavior’ subsection). Sample sizes were not determined using any statistical tests.

Surgical procedures

All surgical procedures were performed in accordance with protocols approved by the Boston University Institutional Animal Care and Use Committee. For postoperative analgesia, mice were given ketoprofen (0.1 ml of 1 mg ml^{–1} solution) and buprenorphine (0.06 ml of 0.03 mg ml^{–1} solution) before the start of all surgical procedures. Mice were anesthetized with 1–2% isoflurane and placed on a heating pad in a stereotaxic apparatus. Artificial tears (Akorn, sodium chloride 5% ophthalmic ointment) were applied to their eyes, and a local anesthetic was injected under the skin (bupivacaine; 0.1 ml of 5 mg ml^{–1} solution) above the skull. The skin overlying the skull was removed to expose the ALM (AP: +2.5 mm, ML: +1.5 mm), bregma and lambda. The periosteum was removed and the skin was secured to the skull with cyanoacrylate (Krazy Glue) around the margins. For electrophysiology experiments, a headbar was implanted just posterior to bregma and secured with superglue and dental cement (Jet). Wells to hold cortex buffer (NaCl 125 mM, KCl 5 mM, glucose 10 mM, HEPES 10 mM, CaCl₂ 2 mM, MgSO₄ 2 mM, pH 7.4) during electrophysiology recordings were sculpted using dental cement, and a thin layer of superglue was applied over any exposed bone.

For optogenetics experiments, after headbar implantation, bone overlying frontal cortex was thinned bilaterally with a dental drill and removed, exposing the frontal cortex. A glass window was then implanted over each hemisphere and secured to the skull with cyanoacrylate⁵⁸. Before performing photoinactivation experiments, any bone regrowth was removed.

Mouse behavior

After surgery, mice were allowed to recover for approximately 1 week and then water restricted, receiving 1.0 ml of water per day. Behavioral training commenced after animals had been water restricted for 3–4 d. If animals received less than 1.0 ml during training, they were supplemented with additional water.

Behavioral training was conducted using custom scripts written for Bpod r0.5 (<https://www.sanworks.io/shop/products.php?productFamily=bpod>). All mice used in experiments were first trained on the DR task with a fixed delay epoch until they reached at least 70% accuracy. At the beginning of each trial, one of two auditory tones lasting 1.3 s was played; the tone indicating a ‘right’ trial was white noise, and the tone indicating a ‘left’ trial was a pure tone (8-kHz pulses). The delay epoch (0.9 s for 12 mice, 0.7 s for one mouse and linearly time warped to 0.9 s) started after the completion of the sample tone. After the delay epoch, an auditory go cue was played (swept-frequency cosine (chirp), 10 ms), after which animals were rewarded with approximately 3 µl of water for contacting the correct lickport. Lickport contacts before the response epoch resulted in the current task epoch restarting, ensuring that they could not progress to the next task epoch until they had stopped licking for the specified

length of the current epoch. If the animal did not respond within 3 s of the go cue, this was considered an ‘ignore’ trial, although responses were typically registered within 300 ms of the go cue.

Mice used for the two-context task (Figs. 1, 2, 4 and 8 and Extended Data Figs. 1–3 and 9) were introduced to the WC context after they were fully trained on the DR context and at least 2 d before the first electrophysiology recording session. A behavioral session began with approximately 100 DR trials and was then followed by alternating blocks of WC and DR trials. Each interleaved block was 10–25 trials, with no fixed block duration. All sessions started with DR trials. Pilot sessions beginning with WC trials had to be terminated early due to a high ‘ignore’ rate, likely due to the mouse becoming sated early in the session. In a WC trial, all auditory cues were omitted, and an approximately 3-µl water reward was presented randomly from either lickport. Trials in which the animal contacted the lickport before the reward (‘early lick’) were omitted from analyses. ITIs were randomly drawn from an exponential distribution with mean 1.5 s. There were no explicit cues during the ITI to indicate to the animal which context block it was in.

For the randomized delay task (Fig. 3g,h,i), mice were first fully trained on the fixed delay DR task (delay duration of 0.9 s). The delay epoch length was then randomized; the duration for each trial was randomly selected from six possible values (0.3 s, 0.6 s, 1.2 s, 1.8 s, 2.4 s and 3.6 s). The probability of each delay duration being selected was assigned such that it resembled a probability density function of an exponential distribution with $\tau = 0.9$ s, as in ref. 44. Animals were trained on this version of the task until they became experts (>70% accuracy and <20% early lick rate).

Videography analysis

High-speed video was captured (400-Hz frame rate) from two cameras (FLIR, Blackfly monochrome camera, BFS-U3-16S2M-CS). One provided a side view of the mouse, and the other provided a bottom view. We tracked the movements of specific body features using DeepLabCut³⁰ (Figs. 1, 3 and 4). The tongue, jaw and nose were tracked using both cameras, whereas the paws were tracked using only the bottom view. Position and velocity of each tracked feature was calculated from each camera. The *x* and *y* position of each kinematic feature was extracted from the output of DeepLabCut. Missing values were filled in with the nearest available value for all features, except for the tongue. The velocity of each feature was then calculated as the first-order derivative of the position vector. Tongue angle and length were found using the bottom camera. Tongue angle was defined as the angle between the vector pointing from the jaw to the tip of the tongue and the vector defining the direction the mouse was facing. Tongue length was calculated as the Euclidean distance from the jaw to the tip of the tongue.

Plots of kinematic feature overlays (Figs. 1e, 3a and 4a and Extended Data Fig. 3) were generated by plotting an (*r*, *g*, *b*) value for each timepoint, *t*, where the values were specified by (KinFeature1_{*t*}, KinFeature2_{*t*}, KinFeature3_{*t*}). All kinematic features (speed or motion energy) were first standardized by taking the 99th percentile across time and trials and normalizing to this value.

Motion energy

The motion energy for a given frame and pixel was defined as the absolute value of the difference between the median value of the pixel across the next five frames (12.5 ms) and the median value of the pixel across the previous five frames (12.5 ms). Motion energy for each frame was then converted to a single value by taking the 99th percentile (~700 pixels) of motion energy values across the frame. Motion energy calculated in this manner was high during overt movements over small regions of pixel space, such as during whisking, while remaining relatively low during passive respiration that corresponded to subtle motion across the animal’s body. A threshold above which an animal was classed as moving was defined on a per-session basis manually. Motion energy distributions, per session, were bimodal, showing a large peak with

little variance at low motion energy values and a second, smaller peak with high variance at large values. The threshold was set as the motion energy value separating these two modes. We found that this method of setting the threshold captured both short-duration and long-duration movements. Alternative methods, such as defining a baseline period of no movement against which to compare, were less successful due to the variability in timing of the movements across trials.

Photoinactivation experiments

Optogenetic photoinactivation was deployed on approximately 30% of trials selected at random. A ‘masking flash’ (470-nm LEDs; LUXEON Rebel LED) controlled by an Arduino Teensy microcontroller (100-ms pulses at 10 Hz) was delivered constantly for the duration of the session to prevent mice from differentiating control and photoinactivation trials. A 488-nm laser (Coherent, OBIS 488-nm LX 150-mW laser system) was used for all photoinactivation experiments. ChR2-assisted photoinactivation (Fig. 1b,c,g and Extended Data Fig. 1b–d) was performed through a cranial window⁵⁸ (see ‘Surgical procedures’ subsection), which replaced bone over the frontal cortex. Light was delivered either at the start of the delay epoch or at the start of the response epoch (only one epoch tested per session). A power density of 1.5 mW mm^{−2} was used for all photoinactivation experiments.

For delay epoch photoinactivation (Fig. 1g–i and Extended Data Fig. 1b), light was delivered at the onset of the delay epoch and lasted for 0.6 s, followed by a 0.2-s linear ramp down. We targeted one of three regions per session: bilateral motor cortex (ALM and tJM1), bilateral ALM or bilateral tJM1. A scanning galvo system (Thorlabs, GPS011) was used to simultaneously target both hemispheres by scanning at 40 Hz. The beam (2-mm diameter) was centered around the following locations: ALM (AP 2.5 mm, ML 1.5 mm), tJM1 (AP 1.75 mm, ML 2.25 mm) and motor cortex (that is, ALM and tJM1; AP 2.25 mm, ML 1.85 mm). Due to their proximity, when specifically targeting ALM or tJM1, Kwik-Cast (World Precision Instruments) was applied to the surrounding regions to prevent photoinactivation of other regions. For photoinactivation at the go cue or water drop (Fig. 1b,c and Extended Data Fig. 1c,d), light was delivered to the motor cortex for 0.8 s, followed by a 0.2-s ramp down beginning at the go cue/water drop presentation.

Electrophysiology recordings

Extracellular recordings were performed in the ALM using one of two types of silicon probes: H2 probes, which have two shanks, each with 32 channels with 25-μm spacing (Cambridge Neurotech), or Neuropixels 1.0 (ref. 59), which have one shank and allow recording from 384 channels arranged in a checkerboard pattern (IMEC). For recordings with H2 probes, the 64-channel voltage signals were multiplexed using a Whisper recording system (Neural Circuits, LLC), recorded on a PXI-8133 board (National Instruments) and digitized at 14 bits. The signals were demultiplexed into 64 (for H2 probes) or 384 (for Neuropixels 1.0) voltage traces sampled at 25 kHz and stored for offline analysis.

At least 6 h before recording, a small craniotomy (1–1.5-mm diameter) was made over the ALM (AP 2.5 mm, ML 1.5 mm). Then, 2–5 recordings were performed on a single craniotomy on consecutive days. After inserting the probes between 900 μm and 1,100 μm (New Scale Technologies, MPM System), brain tissue was allowed to settle for at least 5 min before starting recordings. All recordings were made using SpikeGLX (<https://github.com/billkarsh/SpikeGLX>, phase3a release).

Behavioral analysis

Behavioral data were collected using Bpod r0.5 (<https://www.sanworks.io/shop/products.php?productFamily=bpod>). All sessions used for behavioral analysis had at least 40 correct DR trials for each direction (left or right) and 20 correct WC trials for each direction, excluding early lick and ignore trials, which were omitted from all analyses. To assess the impact of ALM photoinactivation on movement initiation, we first calculated the percent of trials with a correct lick within 600 ms

of the go cue/water drop (Fig. 1c). To find the fraction of time with the tongue visible during the photoinactivation period (Extended Data Fig. 1d), for each trial we found the number of timepoints within the 1 s after the go cue/water drop where the tongue was labeled as visible by DeepLabCut and divided that by the total number of timepoints within the photoinactivation period. This was then averaged for all control or photoinactivation trials for a given session.

To assess the impact of bilateral motor cortex/ALM/tJM1 photoinactivation during the delay epoch, the average motion energy during the delay epoch was calculated separately for right and left control versus photoinactivation trials (Fig. 1h and Extended Data Fig. 1b).

Electrophysiology recording analysis

JRCLUST⁶⁰ (<https://github.com/JaneliaSciComp/JRCLUST>, version 4.1.0) and/or Kilosort 3 (ref. 61) (<https://github.com/MouseLand/Kilosort>) with manual curation using Phy 2 (<https://github.com/cortex-lab/phy>) were used for spike sorting. A unit was considered a well-isolated single unit based on manual inspection of its inter-spike interval (ISI) histogram, separation from other units and its stationarity across the session⁶². Units that passed manual curation but had a higher ISI violation rate were called multiunits. Recording sessions were included for analysis only if they had at least 10 units (see Extended Data Fig. 2 for a distribution of neuron counts across sessions). For the DR task, we recorded 1,651 units (483 single units) in ALM from 25 sessions using nine mice. In 12 sessions from six mice, animals performed the two-context task. In total, 522 units (214 well-isolated single units) were recorded in these sessions. Finally, for the randomized delay task, we recorded 845 units (288 well-isolated single units) in ALM from 19 sessions using four mice. For subspace alignment (Figs. 5g,n, 7d,h and 8d) and single-unit selectivity analyses (Fig. 3g and Extended Data Fig. 4c), only well-isolated single units with firing rates exceeding 1 Hz were included. All units with firing rates exceeding 1 Hz were included in all other analyses.

To find choice-selective neurons, 40 trials were subsampled for both right and left correct trials, and a rank-sum test was used to compare spike counts for each unit during the sample, delay or response epochs. Selective units were those with significantly different spike counts ($P < 0.05$) during a given epoch. Context-selective units were defined in a two-step process to ensure that we were not conflating context selectivity with slow changes in animal engagement/motivation over each session. First, 40 trials each were subsampled for DR and WC trials, and spike counts were compared ($P < 0.05$, rank-sum test) during the ITI (the 300 ms preceding the sample tone onset) when the animal received no external cues to indicate which context it was in. Because sessions always began with a DR block and often ended with a WC block, it is possible that differences in firing rates across contexts defined in this way could be representing time within the session. To account for this, for each unit identified in the first step, we calculated the difference in spike rate for each pair of DR and WC blocks (for example, a session with five blocks, DR–WC–DR–WC–DR, would have $n_{\text{Pairs}} = 4$ adjacent, overlapping block pairs) and included units as context selective only if their *preferred* context (context with a higher spike rate) was the same for at least $n_{\text{Pairs}} - 1$ of the block pairs.

Selectivity in the neural population (Extended Data Fig. 4c) was calculated as the difference in spike rate on *preferred* – *non-preferred* trials in choice-selective units, with *preferred* trials referring to the trial type with a higher spike rate for each unit.

CD analysis

CDs were defined as directions in neural activity state space, defined by firing rates, that maximally separated trajectories of two conditions.

CD_{choice} and CD_{action} were calculated as:

$$\mathbf{v} = \frac{\bar{\mathbf{x}}_{\text{lickright}} - \bar{\mathbf{x}}_{\text{lickleft}}}{\sqrt{\text{var}(\mathbf{x}_{\text{lickright}}) + \text{var}(\mathbf{x}_{\text{lickleft}})}}$$

$$CD = \frac{\mathbf{v}}{\sum |\mathbf{v}|}$$

For each unit, the mean spike rate difference between right lick trials, $\bar{x}_{\text{lickright}}$, and left lick trials, $\bar{x}_{\text{lickleft}}$, was calculated in a 400-ms time interval. $\bar{x}_{\text{lickright}}$ and $\bar{x}_{\text{lickleft}}$ were then individually baseline subtracted and divided by baseline s.d. before calculating CDs (baseline: -2.4 s to -2.2 s relative to the go cue, during the ITI). CD_{choice} was calculated in the last 400 ms of the delay epoch, and CD_{action} was calculated in the first 400 ms of the response epoch. The vector \mathbf{v} was then obtained by normalizing by the square root of the sum of the variances across conditions. Finally, \mathbf{v} was normalized by its L_1 norm to ensure that projections do not scale with the length of \mathbf{v} , the number of units simultaneously recorded.

CD_{ramp} was calculated as:

$$\mathbf{v} = \frac{\bar{x}_{\text{delay}} - \bar{x}_{\text{ITI}}}{\sqrt{\text{var}(x_{\text{delay}}) + \text{var}(x_{\text{ITI}})}}$$

$$CD = \frac{\mathbf{v}}{\sum |\mathbf{v}|}$$

This calculation was similar to the calculations of CD_{choice} and CD_{action} but incorporated data from all correct DR trials during the last 400 ms of the delay epoch, \bar{x}_{delay} , and during the ITI (300 ms preceding the sample tone onset), \bar{x}_{ITI} . \bar{x}_{delay} and \bar{x}_{ITI} were both standardized using the baseline statistics from all correct DR trials.

To find CD_{context} , we first calculated a CD_{context}^p for each pair, p , of DR and WC blocks in a session. If the session ended with a WC block, that block was excluded. CD_{context}^p was calculated as:

$$\mathbf{v}_p = \frac{\bar{x}_{\text{DR}^p} - \bar{x}_{\text{WC}^p}}{\sqrt{\text{var}(x_{\text{DR}^p}) + \text{var}(x_{\text{WC}^p})}}$$

$$CD_p = \frac{\mathbf{v}_p}{\sum |\mathbf{v}_p|}$$

CD_{context} was then defined as the average over all CD_{context}^p in a session. The calculation of CD_{context}^p was similar to the calculation of other CDs but incorporated data from the ITI of correct DR, \bar{x}_{DR^p} , and correct WC, \bar{x}_{WC^p} , trials in a given pair of blocks. \bar{x}_{DR^p} and \bar{x}_{WC^p} were standardized using the baseline statistics from all correct DR and WC trials.

CD_{action} was orthogonalized to CD_{choice} to exclusively capture response epoch selectivity that emerges after the go cue. CD_{ramp} was orthogonalized to CD_{action} and CD_{choice} to remove selectivity that may emerge during the response epoch that is also captured in CD_{action} . Orthogonalization was performed using the Gram–Schmidt process.

Projections of the population activity, $\mathbf{x} \in \mathbb{R}^{(B \times K) \times N}$, along the CD were calculated as:

$$\text{proj} = \mathbf{CD}^T \times \mathbf{x}$$

where B is the number of time bins; K is the number of trials; N is the number of neurons; and T is the transpose operator.

For the randomized delay task, trials with non-1.2-s delay lengths were used to calculate CD_{choice} (fit trials), always using the last 600 ms of the delay epoch as the time interval for calculation⁴⁴. Population activity from trials with 1.2-s delays (test trials) was then projected along the CD_{choice} (Fig. 3h, top left, and Fig. 3i, middle). For visualization purposes, trials with all delay lengths (including fit trials) were projected along CD_{choice} in Fig. 3i, top.

Reliability of choice decoding from CD_{choice}

A receiver operating curve (ROC)-AUC analysis was performed to estimate the reliability of choice decoding from projections along CD_{choice} on a single-session basis (Extended Data Fig. 2c). For each session, equal numbers of correct left and right lick trials were split into a training set (70%) and a testing set (30%). CD_{choice} was calculated from the activity of the neural population using training trials as described above. Activity from testing trials was projected along CD_{choice} and provided as input to a logistic regression model (fitglm() in MATLAB with distribution = 'Binomial' and link = 'logit'). The model was optimized to predict the animal's choice on a particular trial from the delay epoch CD_{choice} activity. The model output was then passed into MATLAB's perfcurve() function to obtain an ROC. Reliability of decoding was measured as the area under the ROC (AUC), shown in the inset of Extended Data Fig. 2c.

Choice and context decoding from neural population or kinematic features

We trained logistic regression models to predict either animal choice (Fig. 3b) or behavioral context (Fig. 4b) from either kinematic features or single-trial neural activity. The kinematic regressors were made up of the x and y positions and velocities of DeepLabCut-tracked features (tongue, jaw and nose) as well as the tongue length, angle and motion energy. The neural regressors were the firing rates of the units simultaneously recorded within each session. Equal numbers of correct left and right lick trials were used for this analysis. A separate model was trained at each time bin for both neural and kinematic decoding. Models were trained with ridge regularization and four-fold cross-validation with 30% of trials held out for testing. Chance accuracy was obtained by shuffling choice/context labels across trials. Accuracy of the prediction was defined as:

$$\text{Accuracy} = \frac{\text{Number of correct predictions}}{\text{Total number of predictions}}$$

Emergence of CD_{choice} selectivity and motion energy in randomized versus fixed delay tasks

To find selectivity within the CD_{choice} projection (Fig. 3h, top left), we found the difference between the trial-averaged projection on right and left trials for each session. We then found the maximum selectivity value before the go cue and identified the first timepoint, relative to delay onset, for the selectivity trace to exceed 90% of the maximum value (Fig. 3h, top right). The same metric was found using session-averaged motion energy in Fig. 3h, bottom panels.

Predicting projections along CDs from kinematic features

We trained support vector machines (SVMs) to predict projections along CD_{choice} , CD_{ramp} and CD_{context} . Correct DR trials were used for predicting CD_{choice} and CD_{ramp} projections; correct DR and WC trials were used for predicting CD_{context} . For predicting CD_{context} , the regressors were made up of the x and y positions and velocities of DeepLabCut-tracked features (multiple points on the tongue, jaw and nose taken from two camera angles) as well as the tongue length, angle and motion energy (totaling 60 regressors). Only the sample and delay epoch were considered when predicting CD_{choice} and CD_{ramp} projections; because of this, any tongue-related metrics were not included as regressors in these analyses (totaling 55 regressors). Projections and kinematic features were first mean centered and scaled to unit variance before input to the regression model. The previous B time bins of kinematic data were used to predict the current bin's neural data ($B = 6$; each bin is 5 ms). Models were trained with ridge regularization and four-fold cross-validation. The models were tested on held-out testing data (40% of trials for CD_{choice} and CD_{ramp} ; tested on 30% of trials for CD_{context}). To assess the prediction quality between projections and model predictions, we calculated the variance explained as:

$$R^2 = 1 - \frac{\sum_t (y_t - \hat{y}_t)^2}{\sum_t (y_t - \bar{y})^2}$$

where y_t is the value of the projection at time t ; \hat{y}_t is the prediction; and \bar{y} is the mean of the projection across all time.

When predicting projections along CD_{choice} during the randomized delay task, unregularized models were used (Fig. 3h). This was due to the small numbers of trials for each delay length, which precluded the partitioning of data into three separate sets (training, validation and testing trials). Instead, models were trained with linear regression and four-fold cross-validation and tested on held-out data from each cross-validation fold. For assessing prediction quality between projections and model predictions, trials with 1.2-s delay epochs were used.

To understand which groups of kinematic features were most informative for predicting projections along CD_{choice} or CD_{context} (Extended Data Fig. 3), the absolute value of the beta coefficients for each kinematic regressor in the trained model was first obtained. For each feature group (jaw, nose and motion energy), the average beta coefficient across all regressors belonging to this kinematic feature group ($8 \times B$ for jaw, $6 \times B$ for nose and $1 \times B$ for motion energy; for example, the x and y positions and velocities for one point on the jaw on each camera totals eight jaw regressors for each time bin, B) was calculated. The average beta coefficient for each feature group was then expressed as a fraction of the total of all beta coefficients.

Estimating dimensionality of neural data

The dimensionality of neural data (Extended Data Fig. 6a) was estimated using Parallel Analysis^{63,64}. Null distributions were created for each eigenvalue by shuffling the time bins separately for each column of a matrix $X \in \mathbb{R}^{B \times N}$, where B is the number of time bins, and N is the number of units. The remaining covariance structure is then due to chance. We repeated this procedure 1,000 times to generate a null distribution for each eigenvalue. The number of eigenvalues of the unshuffled data that exceed the 95th percentile of the null distribution was used as an upper bound estimate of dimensionality.

Subspace identification (trial-averaged data)

Delay and response subspaces were identified following the general approach described in ref. 27. For identifying subspaces using trial-averaged data (Fig. 5a), neural activity was first trial averaged separately for correct left and right trials. Activity across trial types was then concatenated, to produce a matrix $X \in \mathbb{R}^{(B \times C) \times N}$, where B is the number of time bins; C is the number of trial types; and N is the number of units. Trial-averaged neural activity was soft normalized (normalization factor = 5) as described in ref. 27 and subsequently baseline subtracted (baseline: -2.4 s to -2.2 s before go cue, during the ITI). The matrix, X , was further divided into two matrices, X_{delay} and X_{response} . X_{delay} contained activity from [-1, 0] seconds relative to the go cue, and X_{response} contained activity from [0, 1] seconds relative to the go cue. Delay and response subspaces were then identified by solving the following optimization problem:

$$[\hat{Q}_{\text{delay}}, \hat{Q}_{\text{response}}] = \underset{[Q_{\text{delay}}, Q_{\text{response}}]}{\operatorname{argmax}} \left[\frac{1}{2} \frac{\operatorname{Tr}(Q_{\text{delay}}^T C_{\text{delay}} Q_{\text{delay}})}{\sum_{i=1}^{d_{\text{delay}}} \sigma_{\text{delay}}(i)} + \frac{1}{2} \frac{\operatorname{Tr}(Q_{\text{response}}^T C_{\text{response}} Q_{\text{response}})}{\sum_{i=1}^{d_{\text{response}}} \sigma_{\text{response}}(i)} \right]$$

$$\text{subject to } Q_{\text{delay}}^T Q_{\text{response}} = 0, \quad Q_{\text{delay}}^T Q_{\text{delay}} = I, \quad Q_{\text{response}}^T Q_{\text{response}} = I$$

where C_{delay} and C_{response} are the covariances of X_{delay} and X_{response} ; σ_{delay} and σ_{response} are the singular values of C_{delay} and C_{response} ; and d_{delay} and d_{response} are the dimensionality of the subspaces. This optimization problem jointly identifies the subspaces that maximally contain the

variance in neural activity during the delay and response epochs. The dimensionality of each space was chosen to be half the number of simultaneously recorded neurons or 20, whichever was smaller. Therefore, the dimensionality of the full population was equal to the dimensionality of the delay or response subspaces for sessions containing 40 or fewer neurons. Optimization was performed using the Manopt⁶⁵ toolbox for MATLAB.

Subspace identification (single-trial data)

For identifying subspaces using single-trial data (Fig. 5h), single-trial neural activity was first binned in 5-ms intervals and smoothed with a causal Gaussian kernel with a half width of 35 ms. Each trial's activity was subtracted by the average baseline activity across trials and scaled by the s.d. across trials (baseline: -2.4 s to -2.2 s before go cue, during the ITI). $X_{\text{moving}} \in \mathbb{R}^{B_p \times N}$ and $X_{\text{stationary}} \in \mathbb{R}^{B_s \times N}$ were defined using the single-trial neural activity when the animal was moving or stationary (see 'Motion energy' subsection), respectively. B_s is the number of time bins during which the animal was annotated as stationary, and B_p is the number of timepoints annotated as moving. All correct and error trials from DR and WC contexts were used unless specified otherwise. Once the data were partitioned into moving and stationary timepoints, subspace identification was carried out as described for trial-averaged data in the preceding paragraph:

$$[\hat{Q}_{\text{mov-null}}, \hat{Q}_{\text{mov-pot}}] = \underset{[Q_{\text{mov-null}}, Q_{\text{mov-pot}}]}{\operatorname{argmax}} \left[\frac{1}{2} \frac{\operatorname{Tr}(Q_{\text{mov-null}}^T C_{\text{stationary}} Q_{\text{mov-null}})}{\sum_{i=1}^{d_{\text{mov-null}}} \sigma_{\text{stationary}}(i)} + \frac{1}{2} \frac{\operatorname{Tr}(Q_{\text{mov-pot}}^T C_{\text{moving}} Q_{\text{mov-pot}})}{\sum_{i=1}^{d_{\text{mov-pot}}} \sigma_{\text{moving}}(i)} \right]$$

$$\text{subject to } Q_{\text{mov-null}}^T Q_{\text{mov-pot}} = 0,$$

$$Q_{\text{mov-null}}^T Q_{\text{mov-null}} = I, \quad Q_{\text{mov-pot}}^T Q_{\text{mov-pot}} = I$$

We quantified the normalized variance explained in a subspace as in refs. 27,66 (Fig. 5d,k):

$$\text{norm VE} = \frac{\operatorname{Tr}(Q^T C Q)}{\sum_{i=1}^d \sigma(i)}$$

where Q is the subspace; C is the covariance of neural activity; and σ are the singular values of C . This normalization provides the maximum variance that can be captured by d dimensions.

Unit activity was reconstructed from movement-null and movement-potent subspaces according to the following equation:

$$X_{\text{subspace}} = XQ^T$$

$$X_{\text{recon}} = X_{\text{subspace}} Q$$

where X_{recon} is the reconstructed neural activity; X_{subspace} is the projected neural activity within the movement-null or movement-potent subspace; X is either single-trial or trial-averaged firing rates; and Q is either the movement-null or movement-potent subspace.

Bootstrapped distributions of CDs within the movement-null and movement-potent subspaces were obtained through two separate methods. For Fig. 7 and Extended Data Fig. 7a,c, we first estimated movement-null and movement-potent subspaces per session and reconstructed neural activity from each subspace. Then, we performed the hierarchical bootstrapping procedure as described in the 'Hierarchical bootstrapping' subsection. For each iteration, we used the original neural activity to estimate the CDs and then projected the reconstructed neural activity onto the CDs. For Extended Data Fig. 5b, for each bootstrap iteration, the reconstructed neural activity itself was used to estimate the CDs. Activity within each subspace was then projected along the respective CDs. For Fig. 7b,f, CDs were directly

identified from the reconstructed neural activity, X_{recon} , for the individual example sessions.

Subspace and CD alignment for an individual unit was calculated as:

$$A = \frac{VE_{\text{mov-null}} - VE_{\text{mov-pot}}}{VE_{\text{mov-null}} + VE_{\text{mov-pot}}}$$

where VE is the variance explained of each individual unit by the movement-null or movement-potent subspace or by the activity along a CD within the movement-null or movement-potent subspace. VE for a single unit was calculated as:

$$VE = 1 - \frac{\sum_t (x_t - \hat{x}_t)^2}{\sum_t (x_t - \bar{x})^2}$$

where x_t is the trial-averaged firing rate; \hat{x} is the trial-averaged firing rate reconstructed from either a subspace or CD; and t is the time bin.

Unit activity was reconstructed from movement-null and movement-potent CD components according to the following equation:

$$X_{\text{CD}} = X \times \mathbf{CD}^T$$

$$X_{\text{recon}} = X_{\text{CD}} \times \mathbf{CD}$$

where X_{recon} is the reconstructed neural activity; X_{CD} is the projected neural activity within the movement-null or movement-potent CD components; X is single-trial firing rates; and \mathbf{CD} is either the movement-null or movement-potent CD.

To ask if the number of single units observed to be aligned to either the movement-null or movement-potent subspace was different than expected by chance, we compared the distributions to those obtained from randomly sampled subspaces or CDs using a procedure similar to the one described in ref. 27. Each element comprising a subspace or CD was randomly sampled from a normal distribution with zero mean and unit variance but was biased by the covariance structure of the neural activity for a given session. Biasing the randomly sampled subspaces by the covariances controls for the unbalanced variance between stationary and moving timepoints (or between delay and response epochs). That is to say that the shuffled distributions take into account the relative amount of movement tuning across the neural population. The dimension for each random subspace on each iteration was randomly sampled uniformly from the dimensions of the movement-null and movement-potent subspaces across sessions. The dimension for each random CD was always 1.

We sampled random subspaces or CDs, $\mathbf{v}_{\text{mov-null}}$ and $\mathbf{v}_{\text{mov-pot}}$, as follows:

$$\mathbf{v}_{\text{mov-null}} = \text{orth} \left(\frac{\mathbf{U}_{\text{mov-null}} \sqrt{\mathbf{S}_{\text{mov-null}}} \mathbf{v}}{\|\mathbf{U}_{\text{mov-null}} \sqrt{\mathbf{S}_{\text{mov-null}}} \mathbf{v}\|_2} \right)$$

$$\mathbf{v}_{\text{mov-pot}} = \text{orth} \left(\frac{\mathbf{U}_{\text{mov-pot}} \sqrt{\mathbf{S}_{\text{mov-pot}}} \mathbf{v}}{\|\mathbf{U}_{\text{mov-pot}} \sqrt{\mathbf{S}_{\text{mov-pot}}} \mathbf{v}\|_2} \right)$$

where \mathbf{U} and \mathbf{S} are the left and right singular vectors of their associated covariance matrices, $\mathbf{C}_{\text{stationary}}$ and $\mathbf{C}_{\text{moving}}$, and \mathbf{v} is a matrix, or vector in the case of CDs, whose elements are independently drawn from a normal distribution with zero mean and unit variance. As described above, covariance matrices $\mathbf{C}_{\text{stationary}}$ and $\mathbf{C}_{\text{moving}}$ are defined by neural activity during stationary and movement timepoints for single-trial data. $\text{orth}(A)$ computes the orthonormal basis of a matrix A . Neural

activity from each session was projected along the randomly sampled subspaces, and alignment indices were calculated. This procedure was repeated 1,000 times to generate null distributions, which represents the alignment indices of our data with random subspaces (Extended Data Fig. 5). To generate P values associated with Figs. 5n, 7d, h and 8d, on each iteration, we computed the proportion of strongly tuned units (alignment ≥ 0.8 and alignment ≤ -0.8). This provided two distributions, one for alignment ≥ 0.8 and another for alignment ≤ -0.8 . These chance distributions were separately fit with Gaussian distributions using the `fitgmdist()` function in MATLAB. Then, P values were computed as the probability of observing the proportion of tuned units that we observed in the data from the fitted distributions.

Subspaces were identified using data from DR and WC correct and error trials (Figs. 5, 7 and 8). Control analyses were performed using separate sets of trials to assess if the movement-potent subspace erroneously contained movement-null-subspace dynamics. First, we estimated subspaces using WC trials only (Extended Data Figs. 4 and 8). Second, we estimated subspaces using DR and WC trials, but we restricted timepoints used to those in the response epoch only (Extended Data Figs. 4 and 8). Both controls allowed us to estimate movement-null and movement-potent subspaces in the presence of uninstructed movements but in the absence of planning dynamics. These controls thus allowed us to measure the degree to which the movement-potent subspace erroneously captured movement-null dynamics.

In a further attempt to validate that the movement-potent subspace is not inadvertently capturing movement-correlated internal dynamics, we identified movement-null and movement-potent subspaces using a two-stage principal component analysis (PCA) approach (Extended Data Fig. 4). The movement-null subspace was first identified as the dominant PCs (first five PCs) of the single-trial firing rates when mice are not moving. Second, the activity within the movement-null subspace was removed from single-trial firing rates. Then, the movement-potent subspace was calculated as the first five PCs of the single-trial residuals:

$$X_{\text{mov-null}} = X Q_{\text{mov-null}}^T$$

$$X_{\text{recon, mov-null}} = X_{\text{mov-null}} Q_{\text{mov-null}}$$

$$Q_{\text{mov-pot}} = \text{PCA}(X - X_{\text{recon, mov-null}})$$

$$X_{\text{mov-pot}} = X Q_{\text{mov-pot}}^T$$

$X \in \mathbb{R}^{(B \times K) \times N}$ is the single-trial firing rates where B is the number of timepoints; K is the number of trials; and N is the number of units. $Q_{\text{mov-null}}$ is the first five PCs of the single-trial firing rates when animals are stationary, and $X_{\text{mov-null}}$ is the projection along those PCs. $X_{\text{recon, mov-null}}$ is the reconstructed neural activity obtained from the multiplication of activity within the movement-null subspace, $X_{\text{mov-null}}$, and the movement-null subspace, $Q_{\text{mov-null}}$. $\text{PCA}(Z)$ indicates computing the PCs of the matrix Z . $Q_{\text{mov-pot}}$ is the first five PCs of the single-trial residuals obtained from subtracting $X_{\text{mov-null}}$ from the single-trial firing rates, X . Finally, $X_{\text{mov-pot}}$ is obtained from projecting single-trial firing rates onto $Q_{\text{mov-pot}}$. $X_{\text{recon, mov-null}}$ contains activity that is explainable by the first five PCs in the absence of movement, and, therefore, the residuals contain movement-related neural dynamics. Thus, performing PCA on these residuals provides a movement-potent subspace that is orthogonal to the movement-null subspace. This approach ensures that dynamics observed during stationarity are contained within the movement-null subspace before assigning any dynamics to the movement-potent subspace. Therefore, it is more conservative in avoiding the mis-assignment of movement-correlated internal dynamics to the movement-potent subspace.

CD_{context} selectivity on ‘High move’ and ‘Low move’ trials

For each session, DR and WC trials were first separated into ‘High move’ and ‘Low move’ trials. ‘High move’ trials were those where the average motion energy in the ITI was greater than the movement threshold for that session (see ‘Motion energy’ subsection). ‘Low move’ trials were the remaining trials. The last 40 trials in each session were excluded from analysis to account for the decrease in uninstructed movements that are observed toward the end of behavioral sessions as animals become satiated and disengaged. CD_{context} was calculated from the full population or from neural activity reconstructed from either subspace. Population activity from ‘High-move’ and ‘Low-move’ trials was then projected along CD_{context}. Selectivity was defined as the trial-averaged projection along CD_{context} on DR trials minus the projection along CD_{context} on WC trials (Fig. 8b,c).

Hierarchical bootstrapping

Projections along CDs were obtained via a hierarchical bootstrapping procedure^{67,68} (Figs. 2 and 7 and Extended Data Figs. 7, 8 and 10). Pseudopopulations were constructed by randomly sampling with replacement M mice, two sessions per sampled mouse, 50 correct trials of each type, 20 error trials of each type and 20 neurons. M is the number of mice in the original cohort. Bootstrapping was repeated for 1,000 iterations. In each iteration, data derived from some individual mice (and sessions, trials and neurons) will be overrepresented, and some will be omitted. Average effects driven by small subsets of animals, sessions, trials and/or units will be accompanied by large confidence intervals. For all results obtained through this bootstrapping procedure, mean and 95% confidence intervals (shaded area) of the bootstrap distribution are shown, except for selectivity (Fig. 7a and Extended Data Fig. 7a), where 5–95% confidence intervals are shown to indicate where projections significantly differ from zero ($P < 0.05$, one-sided test, bootstrap).

Statistics and reproducibility

No statistical methods were used to predetermine sample sizes. All *t*-tests were two-sided unless stated otherwise. All animals from which electrophysiology and/or behavior data were collected were included in the study. Trial type and optogenetic perturbations were randomly determined by a computer program during electrophysiology and behavioral data collection. Experimenters were blinded to trial type during all data collection and spike sorting. Experimenters were not blinded to group allocation during data analysis.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The data described in this study are available on Zenodo at <https://doi.org/10.5281/zenodo.13941414> (ref. 69).

Code availability

MATLAB and Python code for subspace identification is available at <https://github.com/economolab/subspaceID>. Custom MATLAB code used for analyses is available on Zenodo at <https://doi.org/10.5281/zenodo.13941414> (ref. 69).

References

58. Kim, T. H. et al. Long-term optical access to an estimated one million neurons in the live mouse cortex. *Cell Rep.* **17**, 3385–3394 (2016).
59. Jun, J. J. et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).

60. Jun, J. J. et al. Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction. Preprint at *bioRxiv* <https://doi.org/10.1101/101030> (2017).
61. Pachitariu, M., Sridhar, S., Pennington, J. & Stringer, C. Spike sorting with Kilosort4. *Nat. Methods* **21**, 914–921 (2024).
62. Vincent, J. P. & Economo, M. N. Assessing cross-contamination in spike-sorted electrophysiology data. *eNeuro* <https://doi.org/10.1523/ENEURO.0554-23.2024> (2024).
63. Altan, E., Solla, S. A., Miller, L. E. & Perreault, E. J. Estimating the dimensionality of the manifold underlying multi-electrode neural recordings. *PLoS Comput. Biol.* **17**, e1008591 (2021).
64. Horn, J. L. A rationale and test for the number of factors in factor analysis. *Psychometrika* **30**, 179–185 (1965).
65. Boumal, N., Mishra, B., Absil, P.-A. & Sepulchre, R. Manopt, a Matlab toolbox for optimization on manifolds. *J. Mach. Learn. Res.* **15**, 1455–1459 (2014).
66. Jiang, X., Saggari, H., Ryu, S. I., Shenoy, K. V. & Kao, J. C. Structure in neural activity during observed and executed movements is shared at the neural population level, not in single neurons. *Cell Rep.* **32**, 108006 (2020).
67. Aarts, E., Verhage, M., Veenivliet, J. V., Dolan, C. V. & van der Sluis, S. A solution to dependency: using multilevel analysis to accommodate nested data. *Nat. Neurosci.* **17**, 491–496 (2014).
68. Leeden, R. v., Meijer, E. & Busing, F. M. in *Handbook of Multilevel Analysis* (eds Leeuw, J. & Meijer, E.) 401–433 (Springer, 2008).
69. Hasnain, M., Birnbaum, J. & Economo, M. Data and code accompanying Hasnain, Birnbaum et al., Nature Neuroscience 2024 [data set]. *Zenodo* <https://doi.org/10.5281/zenodo.13941415> (2024).

Acknowledgements

We thank K. Svoboda, B. DePasquale, S. Druckmann and B. Scott for helpful discussions; T. Wang for helpful comments on the manuscript; and J. Jiang for help with crystal skull surgeries. This work was supported by the Whitehall Foundation, the Klingenstein Fund, the Simons Foundation and National Institutes of Health R01NS121409 and U19NS137920.

Author contributions

M.A.H., J.E.B. and M.N.E. conceived of the project. M.N.E. and C.C. supervised research. M.A.H., J.E.B. and M.N.E. designed experiments. M.A.H., J.E.B., J.L.U.N. and E.K.H. performed experiments. M.A.H., J.E.B. and M.N.E. analyzed data. M.A.H., J.E.B. and M.N.E. wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

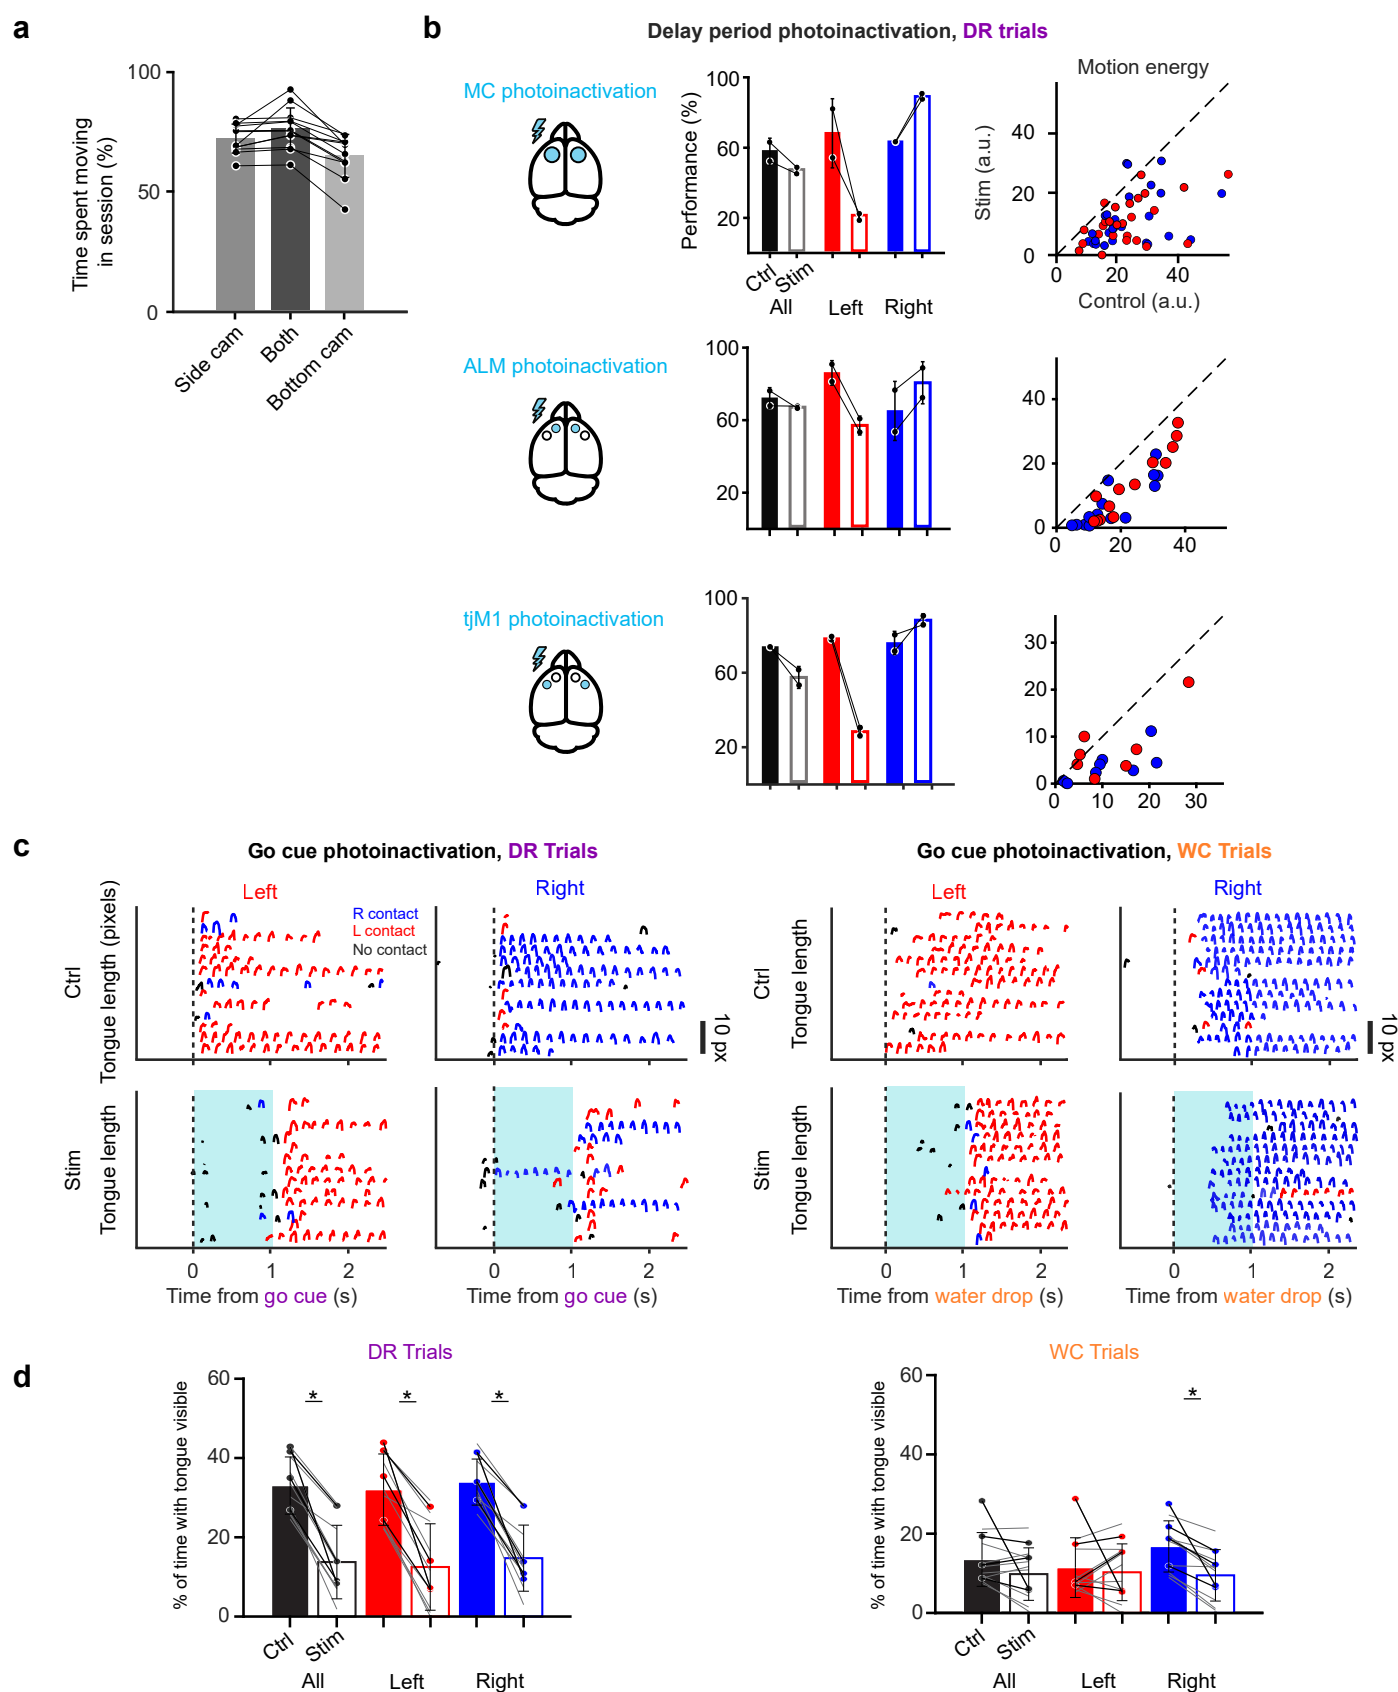
Extended data is available for this paper at <https://doi.org/10.1038/s41593-024-01859-1>.

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41593-024-01859-1>.

Correspondence and requests for materials should be addressed to Michael N. Economo.

Peer review information *Nature Neuroscience* thanks Andrei Khilkevich, Thomas Mrsic-Flogel and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

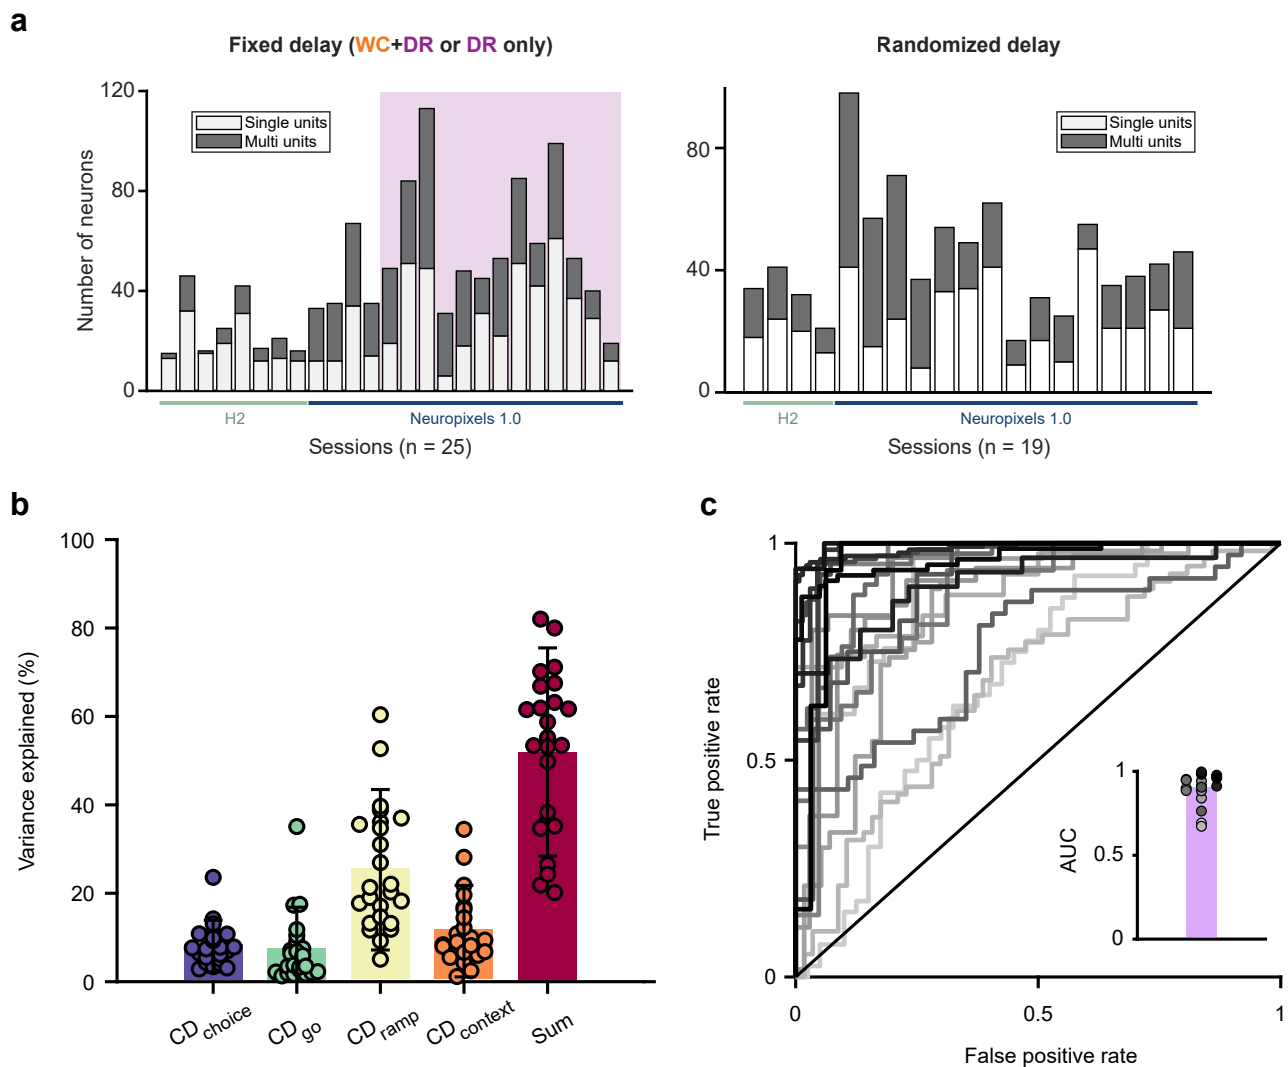
Reprints and permissions information is available at www.nature.com/reprints.



Extended Data Fig. 1 | Region-specific photoinactivation of ALM and tJM1.

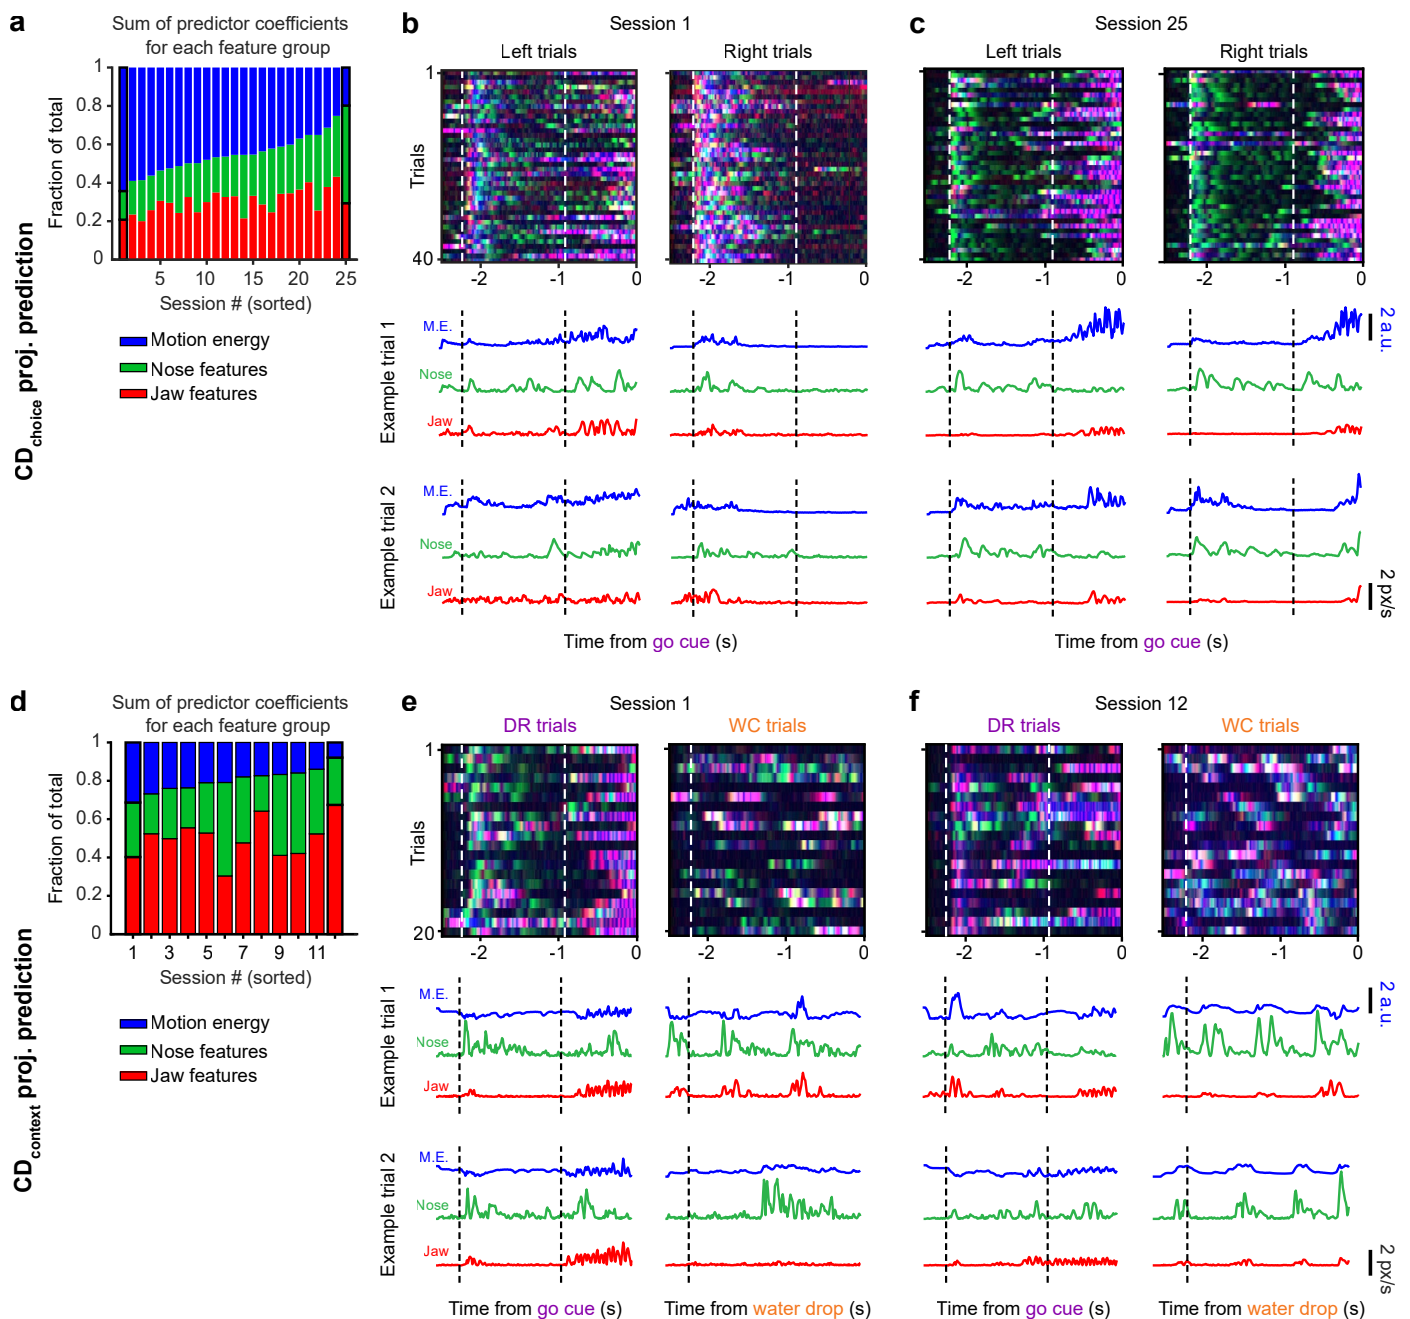
a. Percentage of time spent moving (see Methods) in a session determined by video recordings using the side view, bottom view, and both views. Points indicate each of the 12 randomly selected sessions used for this analysis. Error bars denote standard deviation across sessions ($n = 12$). **b.** Effect of delay epoch photoinactivation on behavioral performance (*middle column*) and uninstructed movements (*right column*) when photoinactivation was directed to the MC (ALM + tJM1; *top row*; same data as Fig. 1g–i $n = 14$ sessions, 2 mice), the ALM (*middle row*, $n = 15$ sessions, 2 mice), and the tJM1 (*bottom*, $n = 9$ sessions, 2 mice). Photoinactivation of ALM and tJM1 led to similar behavioral impairment and reduction in uninstructed movements, with larger effects observed with MC (ALM + tJM1) photoinactivation. **c.** Tongue length during control and go cue/water drop photoinactivation trials for delayed-response (*left*) and water-cued (*right*) contexts. Blue traces indicate right lickport contacts, red traces indicate

left contacts, and black traces indicate no contact. Vertical dashed line indicates go cue or water drop onset. Blue shaded region indicates photoinactivation period. **d.** Percentage of time with tongue visible during photoinactivation period for DR trials (*left*) and WC trials (*right*). Each colored point indicates mean value for an animal ($n = 4$ animals), individual animals are connected by black lines. Light gray lines denote individual sessions ($n = 10$ sessions). Bars are the mean across all sessions. Asterisks denote significant differences ($p < 0.05$) between control and photoinactivation trials (Percent reduction on all DR trials: $19 \pm 7\%$, mean \pm s.d., $p = 1.6 \times 10^{-5}$; DR left trials: $20 \pm 8\%$, $p = 3.0 \times 10^{-5}$; DR right trials: $20 \pm 7\%$, $p = 1.6 \times 10^{-5}$; All WC trials: $4 \pm 8\%$, $p = 0.154$; WC left trials: $1\% \pm 9\%$, $p = 0.702$; WC right trials: $7\% \pm 5\%$, $p = 0.002$; paired two-sided t -test, $n = 10$ sessions). Error bars indicate standard deviation across sessions. In WC trials, tongue protrusion was only significantly impaired on one trial type, while ability to successfully contact the lickport was impaired in all conditions (see Fig. 1c).



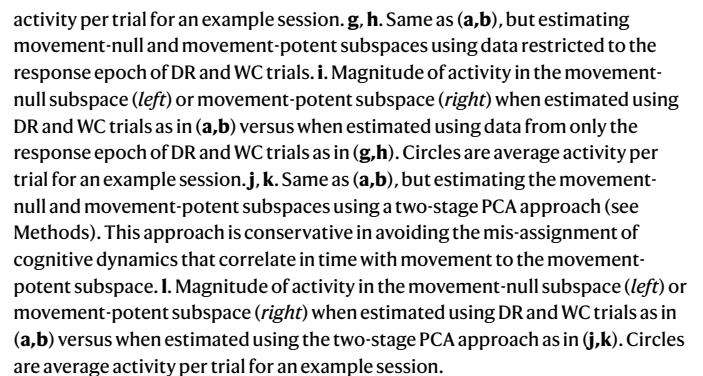
Extended Data Fig. 2 | Session-by-session statistics. a. Number of recorded single- and multi-units per session for the fixed delay task (*left*) or the randomized delay task (*right*). *Left*, Purple shaded region indicates sessions in which animals were only presented with DR trials. Green and blue bars underneath plots indicate the probe type used for a given session. **b.** Variance explained of trial-averaged neural activity by each coding direction. The coding


directions were calculated using neural activity from individual sessions ($n = 25$). Bar height represents the mean across sessions and error bars indicate standard deviation across sessions. **c.** Receiver operating curves (ROC) demonstrating choice decoding accuracy from delay epoch CD_{choice} projections across all individual sessions (see Methods). Inset: area under the ROC curve (AUC). Bar height represents the mean across sessions and points indicate sessions.



Extended Data Fig. 3 | Session-by-session variability in the relationship between kinematic features and putative cognitive dynamics. **a.** Regression weights for each group of kinematic predictors of CD_{choice} projections, as a fraction of all predictor coefficients (see Methods). Sessions are sorted in descending order by motion energy fraction. Outlined bars indicate example sessions shown in (b) and (c). **b.** Example session where motion energy made up the largest fraction of regression weights for predicting CD_{choice} projections.

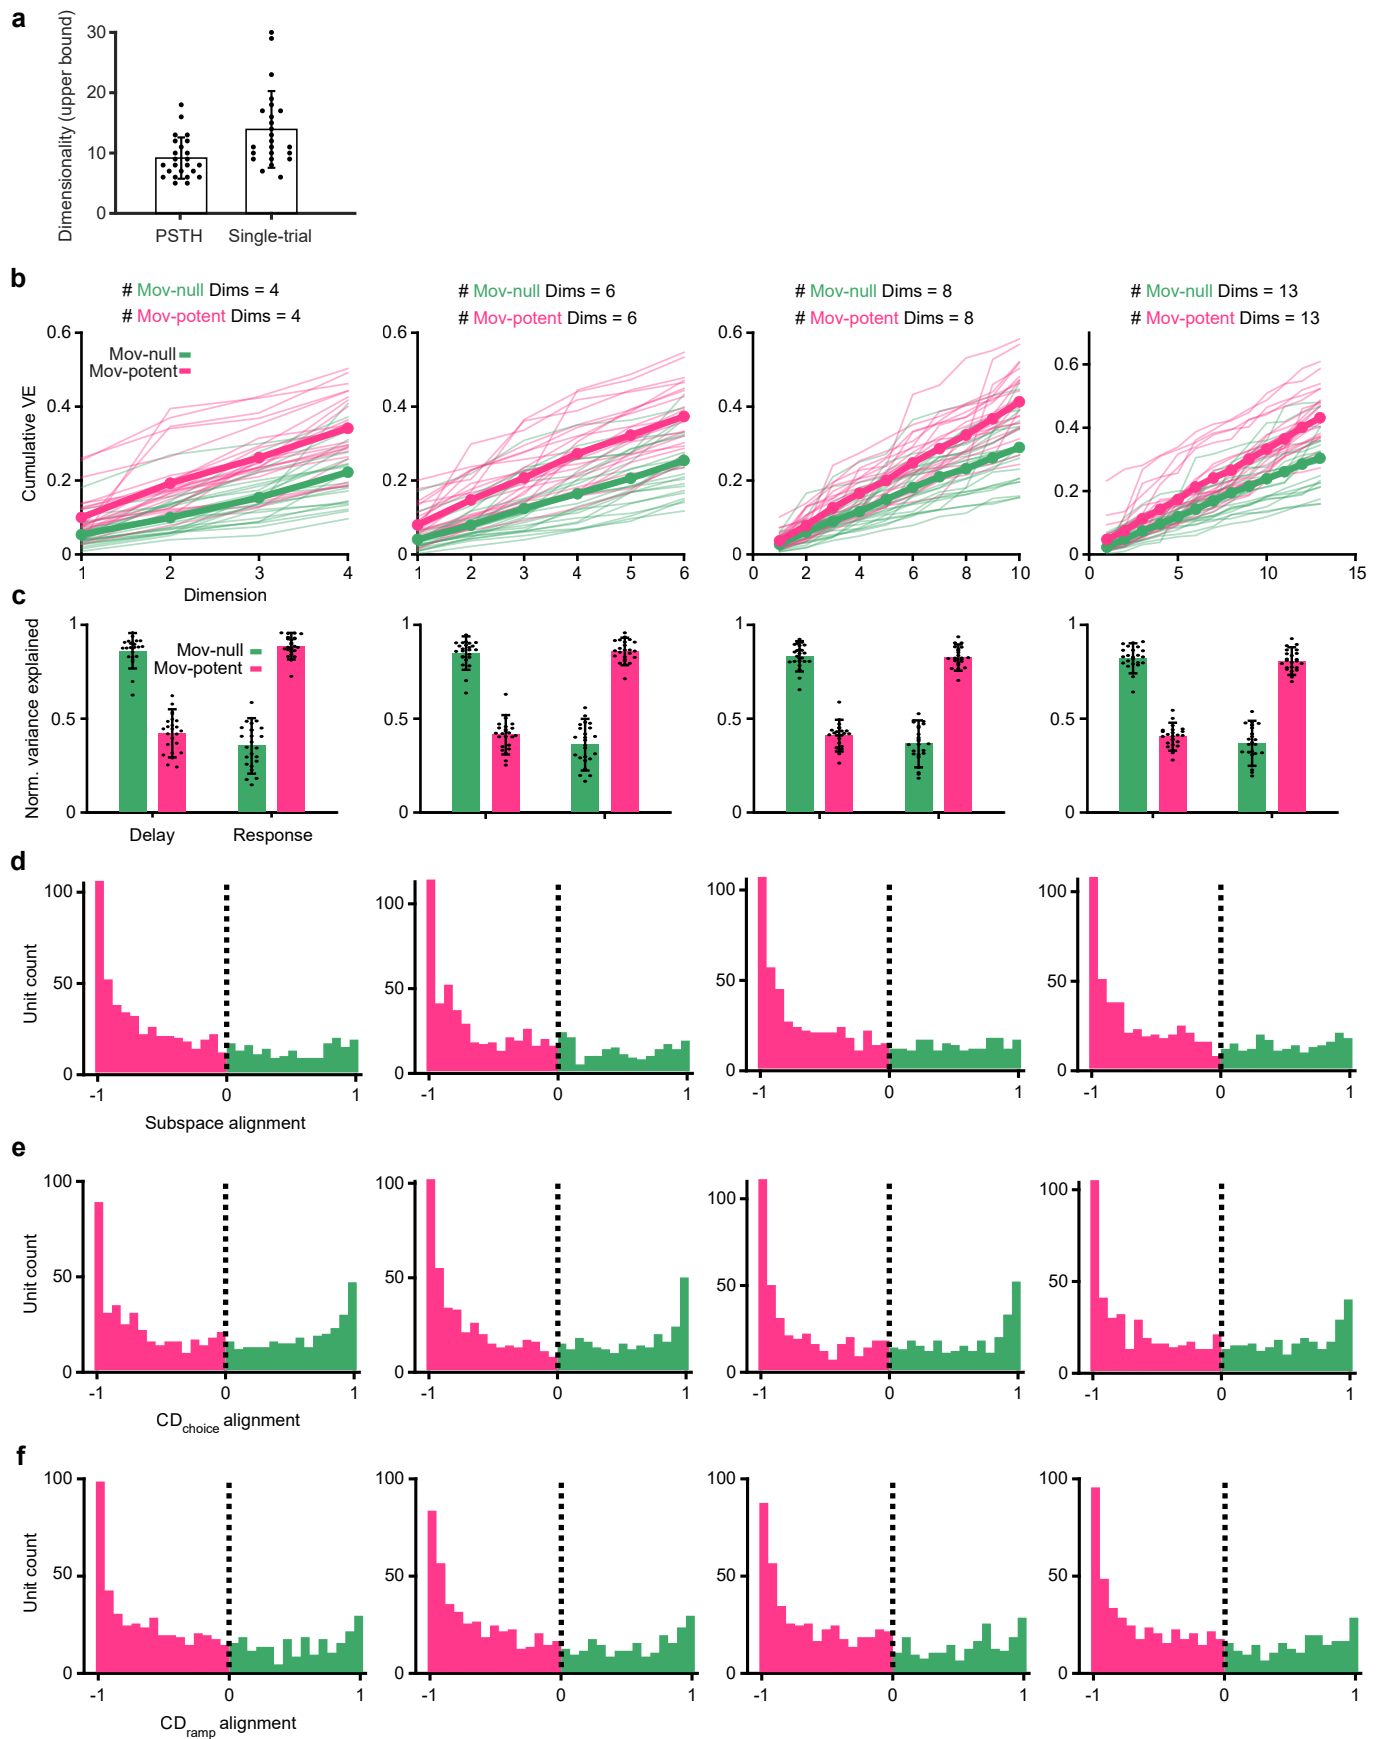
Top, overlaid jaw/nose speed or motion energy for a subset trials. *Bottom*, two example trials of kinematic feature trajectories. **c.** Same as (b) but for an example session where jaw and nose features made up a larger fraction of regression weights. **d.** Same as (a) but for predicting $CD_{context}$ projections. **e, f.** Same as (b) and (c) but for two example sessions with different regression weight fractions for predicting $CD_{context}$ projections.





Extended Data Fig. 5 | Alignment of single-units to random subspaces. Random subspaces were constructed by independently and identically drawing from a normal distribution with zero mean and unit variance. Each random subspace was then biased towards the covariance structure of the actual data (see Methods). **a.** Null distributions of alignment indices for trial-averaged data.

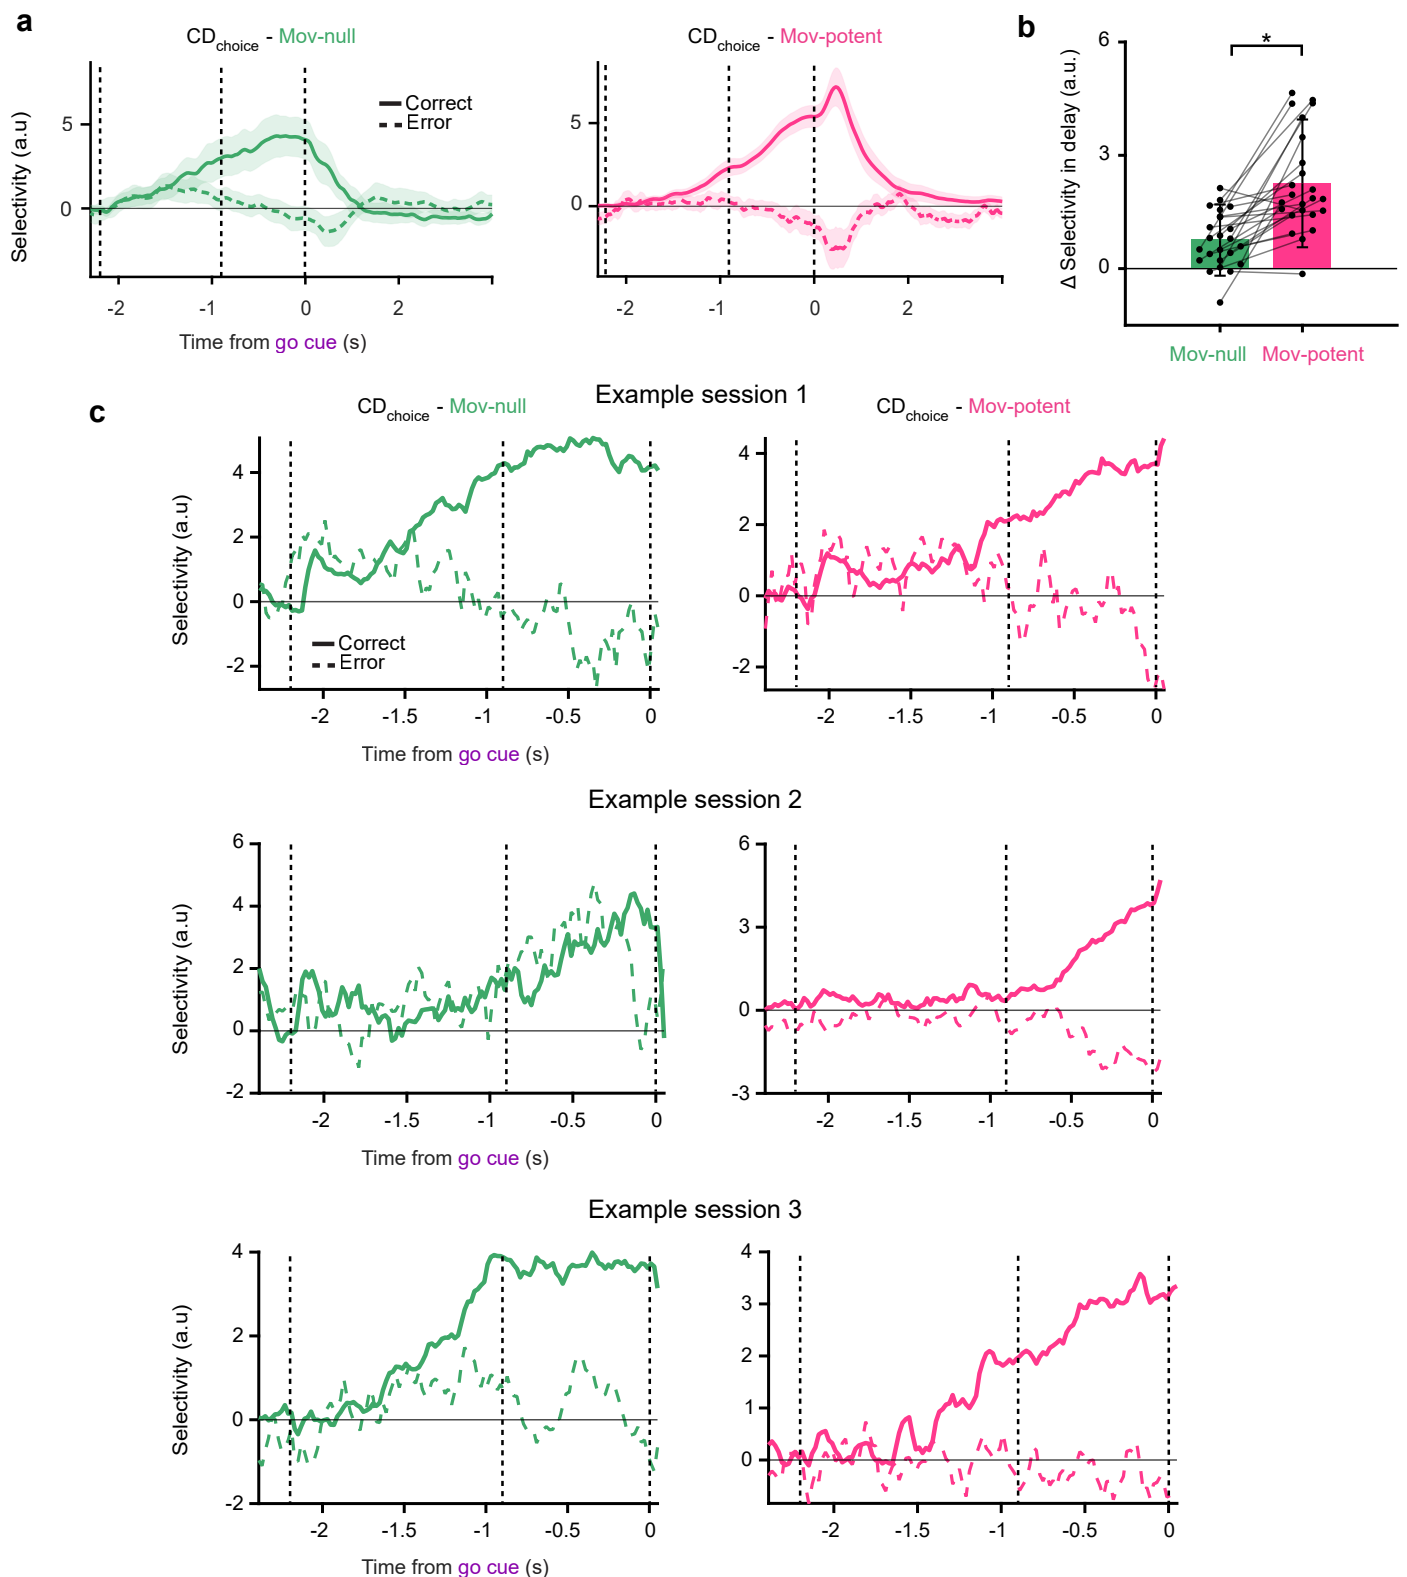
b. Null distributions of alignment indices for single-trial data. Null alignment distributions are skewed towards the movement-potent subspace due to the unbalanced variance between delay and response epochs (**a**) or between stationary and movement time points (**b**), reflecting the strong movement tuning of many units.



Extended Data Fig. 6 | See next page for caption.

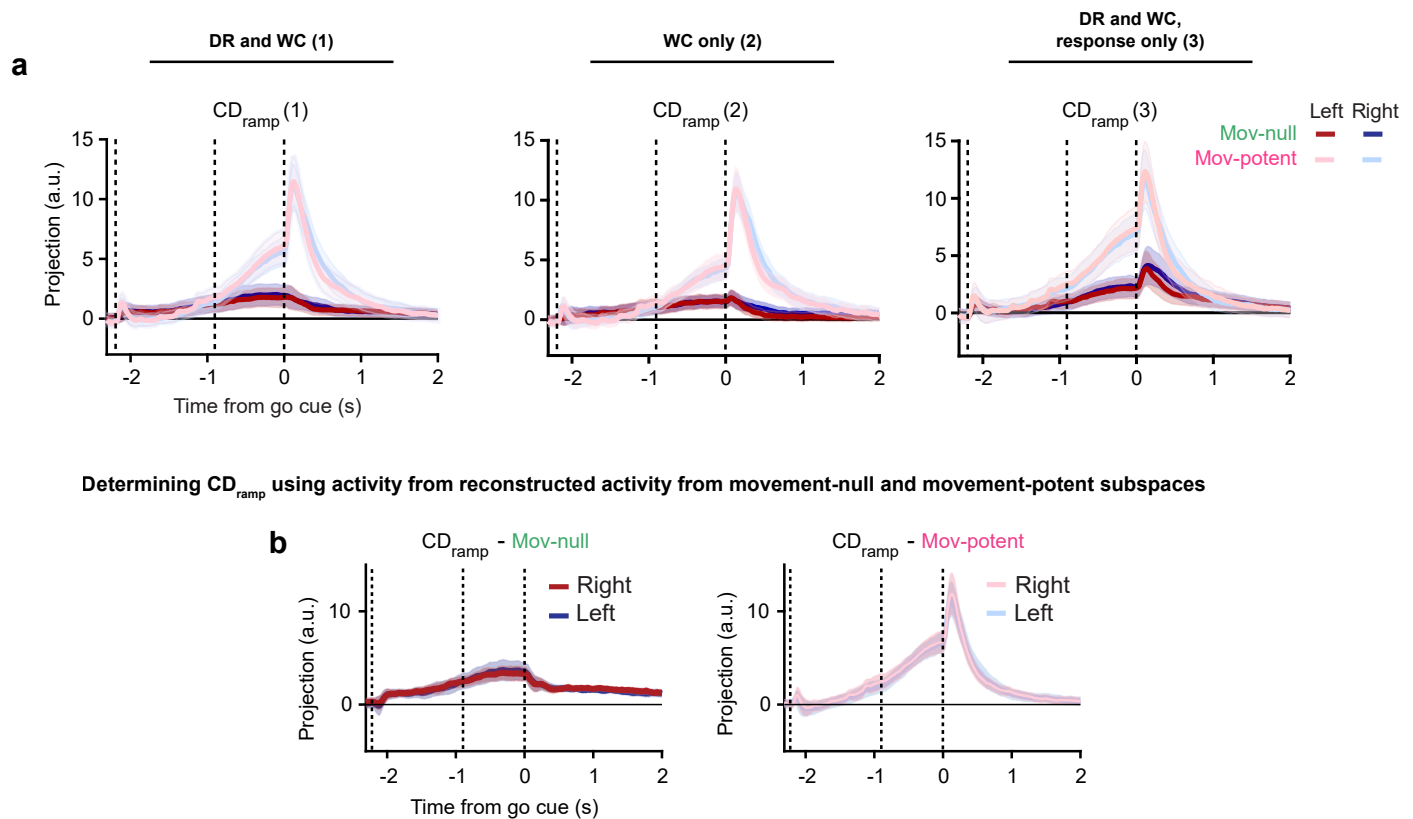
Extended Data Fig. 6 | Varying dimensionality of subspaces. Analyses were repeated while varying the dimensionality of movement-null and movement-potent subspaces. Each subspace was constrained to be 4 (*left*), 6 (*middle left*), 8 (*middle right*), or 13 dimensions (*right*). **a.** Upper bound estimate of dimensionality for trial-averaged (PSTH) data and single-trial data. Bar heights indicate mean across sessions, points indicate sessions, and error bars indicate standard deviation across sessions ($n = 25$ sessions). **b.** Cumulative variance explained of the neural activity by the activity in movement-null and movement-

potent subspaces. Bold lines and points indicate mean across sessions. Thin lines represent single sessions **c.** Normalized variance explained of neural activity during the delay or response epoch by the activity in movement-null and movement-potent subspaces. Points indicate sessions, bar height indicates mean across sessions, and error bars indicate standard deviation across sessions ($n = 25$ sessions). **d–f.** Subspace (**d**), CD_{choice} (**e**), and CD_{ramp} (**f**) alignment distributions when varying dimensionalities of each subspace.



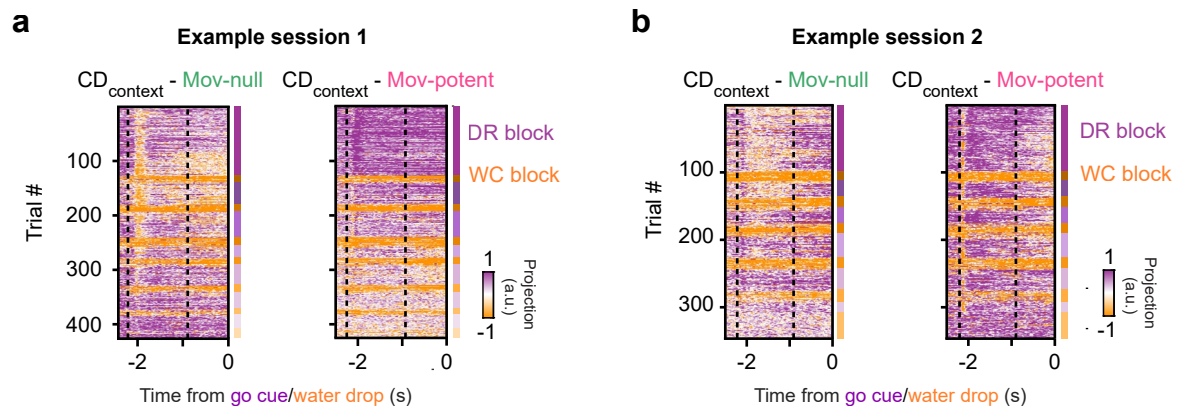
Extended Data Fig. 7 | Projections along movement-null and movement-potent components of CD_{choice}. **a.** Same data as in Fig. 7a except all time in trial shown to highlight activity during the response epoch. Selectivity (projections onto CD_{choice} on lick-right trials minus projections on lick-left trials) of movement-null (left) and movement-potent (right) subspace activity. Mean and 5–95% CI of the bootstrap distribution for correct (solid) and error (dashed) trials shown. **b.** Change in selectivity between the last 100 ms of the delay epoch and the last 100 ms of the sample epoch in movement-null and movement-potent components

of projections along CD_{choice} (Movement-potent: 2.25 ± 1.57 , mean \pm s.d., Movement-null: 0.76 ± 0.8 , $p = 1 \times 10^{-5}$, paired two-sided t -test, $n = 25$ sessions). Points indicate individual sessions, bar height indicates mean across sessions, and error bars indicate standard deviation across sessions. **c.** Three example sessions from three different mice depicting selectivity along CD_{choice} as in Fig. 7a. Solid lines denote the mean projection on correct trials and dashed lines denote the mean projection on error trials.



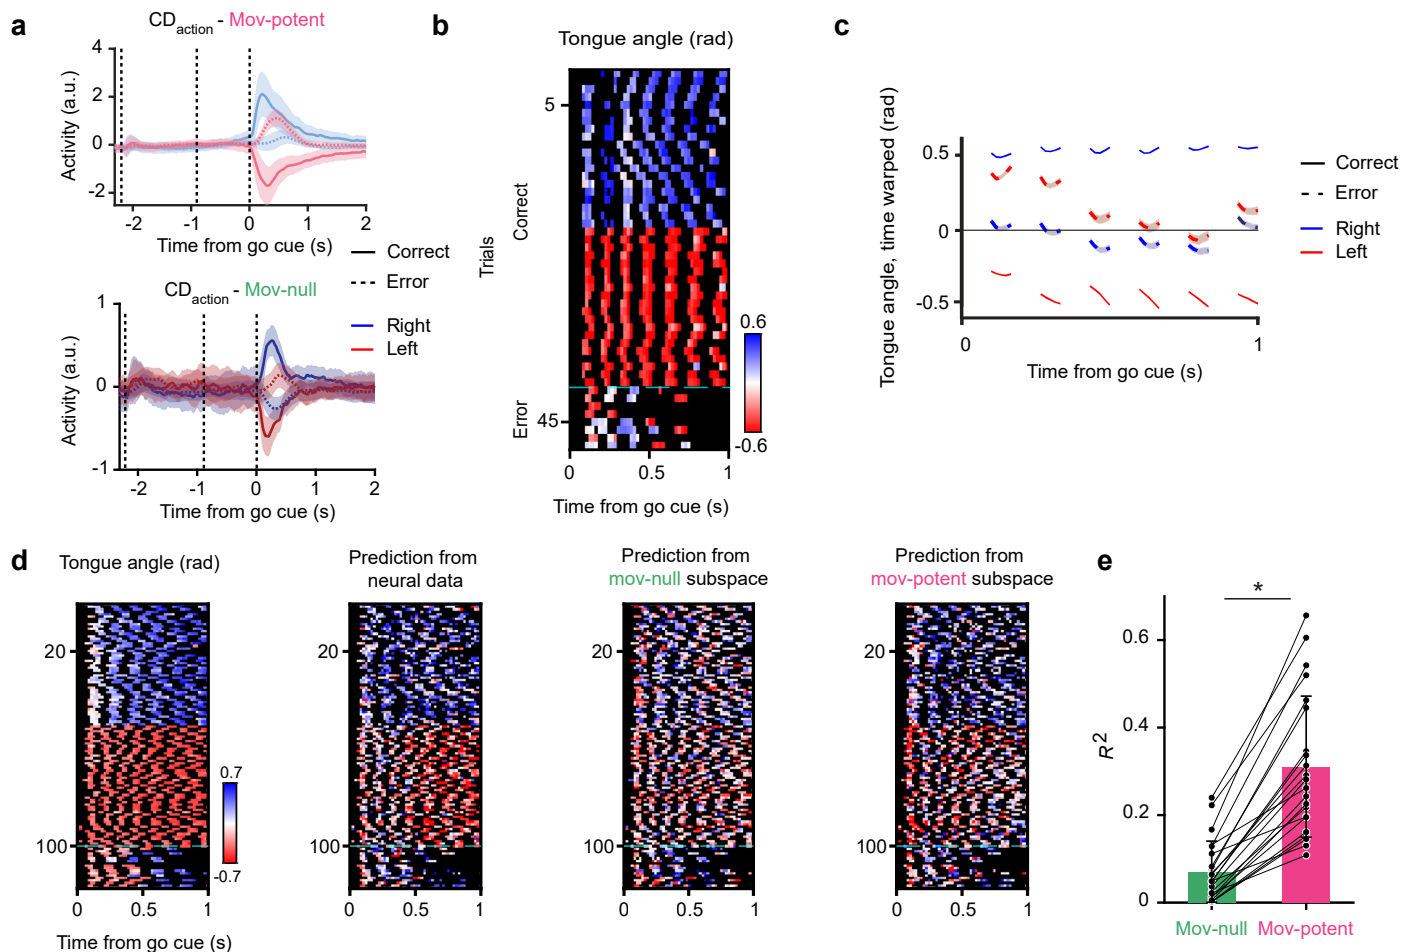
Extended Data Fig. 8 | Within-subspace CD projections using variations on procedure to determine subspaces. a. Projections of movement-null and movement-potent subspace activity along CD_{ramp} for each of three analytical variations. Movement-null and movement-potent subspaces were identified using both DR and WC trials (*left*), WC trials only (*middle*), and the response

epoch of DR and WC trials (*right*). Mean and 95% CI of bootstrap distribution shown. **b.** Projections along movement-null (*left*) and movement-potent (*right*) components of CD_{ramp} when determined from activity within each subspace individually, rather than from the full neural population. Mean and 95% CI of bootstrap distribution shown.



Extended Data Fig. 9 | Encoding of context in both the null and potent subspaces tracks block-wise task structure. a. Heatmap of single-trial projections of null and potent subspace activity along CD_{context} for an example

session. The chronological DR or WC block within the session is denoted by differently shaded purple and orange rectangles, respectively, on the right of each plot. **b.** Same as (a) but for another example session, from a different animal.



Extended Data Fig. 10 | Relationship between tongue angle and neural activity in the movement-null and movement-potent subspaces. **a.** Projections along movement-potent (top) and movement-null (bottom) components of CD_{action} . Correct trials shown in solid lines and error trials shown in dashed lines. Shaded region depicts 95% CI of bootstrap distribution. **b.** Tongue angle for an example session for correct and error trials. Black values indicate tongue not visible. **c.** Tongue angle on correct and error right and left trials. Tongue angle was linearly time warped to allow for averaging over trials and sessions. Mean

and s.e.m. across sessions shown. **d.** Tongue angle (left) and predictions from the full population neural activity (middle left), null subspace activity (middle right), and potent subspace activity (right) for an example session. **e.** Variance explained (R^2) of tongue angle by prediction from movement-null (green) and movement-potent (pink) subspaces. Asterisks denote significant differences between predictions from null and potent subspaces and error bars indicate standard deviation across sessions ($p = 2 \times 10^{-8}$, paired two-sided t-test, $n = 25$ sessions).

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed
<input type="checkbox"/>	<input checked="" type="checkbox"/> The exact sample size (<i>n</i>) for each experimental group/condition, given as a discrete number and unit of measurement
<input type="checkbox"/>	<input checked="" type="checkbox"/> A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
<input type="checkbox"/>	<input checked="" type="checkbox"/> The statistical test(s) used AND whether they are one- or two-sided <i>Only common tests should be described solely by name; describe more complex techniques in the Methods section.</i>
<input type="checkbox"/>	<input checked="" type="checkbox"/> A description of all covariates tested
<input type="checkbox"/>	<input checked="" type="checkbox"/> A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
<input type="checkbox"/>	<input checked="" type="checkbox"/> A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
<input type="checkbox"/>	<input checked="" type="checkbox"/> For null hypothesis testing, the test statistic (e.g. <i>F</i> , <i>t</i> , <i>r</i>) with confidence intervals, effect sizes, degrees of freedom and <i>P</i> value noted <i>Give P values as exact values whenever suitable.</i>
<input checked="" type="checkbox"/>	<input type="checkbox"/> For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
<input checked="" type="checkbox"/>	<input type="checkbox"/> For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
<input checked="" type="checkbox"/>	<input type="checkbox"/> Estimates of effect sizes (e.g. Cohen's <i>d</i> , Pearson's <i>r</i>), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection	Behavioral data was collected using Bpod r0.5 (https://www.sanworks.io/shop/products.php?productFamily=bpod). Electrophysiology data was collected with SpikeGLX (https://github.com/billkarsh/SpikeGLX , phase3A release). Spike sorting was done using JRClust v4.1.0 (https://github.com/JaneliaSciComp/JRCLUST), Kilosort3 (https://github.com/MouseLand/Kilosort), and Phy 2 (https://github.com/cortex-lab/phy) .
Data analysis	Data was analyzed using custom MATLAB (2022b) code. Subspace identification code is available at: https://github.com/economolab/subspaceID . Subspace optimization was performed using the manopt toolbox for MATLAB (https://github.com/NicolasBoumal/manopt). Analysis code is available at: https://doi.org/10.5281/zenodo.13941414 .

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

The data described in this study are available on Zenodo: <https://doi.org/10.5281/zenodo.13941414>.

Research involving human participants, their data, or biological material

Policy information about studies with [human participants or human data](#). See also policy information about [sex, gender \(identity/presentation\), and sexual orientation](#) and [race, ethnicity and racism](#).

Reporting on sex and gender

Reporting on race, ethnicity, or other socially relevant groupings

Population characteristics

Recruitment

Ethics oversight

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size

Data exclusions

Replication

Randomization

Blinding

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology and archaeology
<input type="checkbox"/>	<input checked="" type="checkbox"/> Animals and other organisms
<input checked="" type="checkbox"/>	<input type="checkbox"/> Clinical data
<input checked="" type="checkbox"/>	<input type="checkbox"/> Dual use research of concern
<input checked="" type="checkbox"/>	<input type="checkbox"/> Plants

Methods

n/a	Involved in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/> MRI-based neuroimaging

Animals and other research organisms

Policy information about [studies involving animals](#); [ARRIVE guidelines](#) recommended for reporting animal research, and [Sex and Gender in Research](#)

Laboratory animals	This study used data collected from 17 mice; both male and female animals between 8 and 11 weeks were used. Six animals were used for the two-context task and were either C57BL/6J (JAX 000664) or VGAT-ChR2-EYFP (-/-) (bred in-house; breeders: JAX 014548). An additional three animals, either C57BL/6J or VGAT-ChR2-EYFP (-/-) were trained only on the delayed response task. A separate cohort of four animals were used in the randomized delay task (VGAT-ChR2-EYFP (+/-) or VGAT-ChR2-EYFP (-/-)). Finally, four VGAT-ChR2-EYFP (+/-) animals were used in optogenetics experiments. Mice were housed in a 12-hour reverse dark/light cycle room (ambient temperature of 68-79° F (20-26° C); humidity of 30-70%) with ad libitum access to food. Access to water was restricted during behavioral and electrophysiology experiments.
Wild animals	This study did not involve wild animals.
Reporting on sex	This study used both male (n=7) and female (n=10) mice. The results of this study do not relate to sex.
Field-collected samples	No field samples were collected or used in this study.
Ethics oversight	All performed surgical procedures were in accordance with protocols approved by Boston University Institutional Animal Care and Use Committee.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Plants

Seed stocks	<i>Report on the source of all seed stocks or other plant material used. If applicable, state the seed stock centre and catalogue number. If plant specimens were collected from the field, describe the collection location, date and sampling procedures.</i>
Novel plant genotypes	<i>Describe the methods by which all novel plant genotypes were produced. This includes those generated by transgenic approaches, gene editing, chemical/radiation-based mutagenesis and hybridization. For transgenic lines, describe the transformation method, the number of independent lines analyzed and the generation upon which experiments were performed. For gene-edited lines, describe the editor used, the endogenous sequence targeted for editing, the targeting guide RNA sequence (if applicable) and how the editor was applied.</i>
Authentication	<i>Describe any authentication procedures for each seed stock used or novel genotype generated. Describe any experiments used to assess the effect of a mutation and, where applicable, how potential secondary effects (e.g. second site T-DNA insertions, mosaicism, off-target gene editing) were examined.</i>