Engaged decision-makers align spontaneous movements to stereotyped task demands

Authors: Chaoqun Yin*, Maxwell D. Melin*, Gabriel Rojas-Bowe, Xiaonan Richard Sun, Steven Gluf, João Couto, Alex Kostiuk, Simon Musall, Anne K. Churchland

* These authors contributed equally.

Abstract

Neural activity during sensory-guided decision-making is strongly modulated by animal movements. Although the impact of movements on neural activity is now well-documented, the relationship between these movements and behavioral performance remains unclear. To understand this relationship, we first tested whether the magnitude of animal movements (assessed with posture analysis of 28 individual body parts) was correlated with performance on a perceptual decision-making task. No strong relationship was present, suggesting that task performance is not affected by the magnitude of movements. We then tested if performance instead depends on movement timing and trajectory. We partitioned the movements into two groups: task-aligned movements that were well predicted by task events (such as the onset of the sensory stimulus or choice) and task independent movement (TIM) that occurred independently of task events. TIM had a reliable, inverse correlation with performance in head-restrained mice and freely moving rats. This argues that certain movements, defined by their timing and trajectories relative to task events, might indicate periods of engagement or disengagement in the task. To confirm this, we compared TIM to the latent behavioral states recovered by a hidden Markov model with Bernoulli generalized linear model observations (GLM-HMM) and found these, again, to be inversely correlated. Finally, we examined the impact of these behavioral states on neural activity measured with widefield calcium imaging. The engaged state was associated with widespread increased activity, particularly during the delay period. However, a linear encoding model could account for more overall variance in neural activity in the disengaged state. Our analyses demonstrate that this is likely because uninstructed movements had a greater impact on neural activity during disengagement. Taken together, these findings suggest that TIM is informative about the internal state of engagement, and that movements and state together have a major impact on neural activity.

Introduction

Animals make spontaneous movements almost constantly. Even in uncluttered environments, rodents will continuously crouch, rear and turn[1]. Both rodents and humans engage in extensive non-exercise “fidget-like” movements that consume considerable energy[2]. These “uninstructed movements” are distinct from the well-studied trained movements that an experimenter “instructs” a subject to make (e.g., movements a human or animal learns to make for reporting a decision).

Uninstructed movements have a major impact on neural activity. Modulation of neural activity in primary visual cortex by running sharpens the orientation tuning of excitatory neurons in passively-viewing mice[3,4]. This change is likely driven by a change in the operating point of neurons along a spiking nonlinearity[5] due to the recruitment of VIP interneurons which disinhibit cortical activity[6]. In addition to the coordinated limb and body movements required for running, mice make many smaller, idiosyncratic movements that can be captured using a camera positioned on the animal's face[7]. Facial movements strongly modulate neural activity, and are mostly orthogonal to sensory signals, suggesting that their modulation might not interfere with ongoing sensory processing[8]. Importantly, these conclusions stem from experiments conducted in passively-viewing animals. A natural conclusion from the work above is that these uninstructed movements have a large impact on neural activity particularly when animals are untrained: if an animal is simply passively viewing, neural activity might be dominated by uncontrolled state fluctuations. However, uninstructed movements likewise have a major impact on neural activity in trained animals, both in primary sensory cortices and elsewhere[9–11].

Although the impact of uninstructed movements on neural activity is now well-documented, the relationship of these movements to decision-making performance remains poorly understood. While including video data in encoding models designed to predict neural activity from behavior is a powerful approach to quantify the impact of movements on neural activity, the exact nature of these movements and how they interact with decision-related computations during decision-making remains unanswered. One hint about the relationship between uninstructed movements and decision-making performance comes from examining these movements over learning. Interestingly, as animals transition from novices to experts, the timing of their movements changes[9]. For example, whisks in novice animals are equally likely to occur at each moment in a decision. In contrast, expert animals make well-timed whisks that occur in a stereotyped manner relative to task events, such as the stimulus onset and reward delivery. In another study, higher levels of task-relevant motor variability were associated with faster learning on a reward based learning experiment[12].

The preponderance and timing of movements might fluctuate not only over learning, but also within an experimental session in an expert animal. Casual observation of behavioral videos reveals that periods of high movement tend to wax and wane, even in experts. Recent work examining neural states suggests that decision-making performance likewise fluctuates over an experimental session[13]. An intriguing possibility is therefore that these model-defined states might be related to movement-defined internal states. However, it is unclear how to quantify movement data to uncover this potential relationship.

To address these questions, we developed a new, model-based method to partition movement variance. We related this metric both to decision-making performance and to cortex-wide neural activity. We defined a new metric, Task Independent Movements (TIM), that quantifies movements that are independent of task events. Using this new metric, we discovered that the magnitude of TIM is inversely related to decision-making performance, both in head-fixed mice and freely moving rats. Next, we related the magnitude of TIM to model-defined states of engagement and found that TIM tended to be highest when animals were disengaged. Finally, we compared neural activity during

engaged and disengaged states, and observed multiple signatures of engagement that are apparent both in the raw fluorescence signal and also in the outputs of an encoding model. Taken together, these observations suggest that uninstructed movements not only modulate neural activity in a state-dependent fashion, they are also an indicator of an animal's latent state of engagement. Importantly, it is not the magnitude of movements that is informative but their timing and trajectories, as animals often make frequent movements when highly engaged.

Results

To understand how an animal's movements relate to decision-making performance, we trained animals on a 2-choice sensory decision task (Fig. 1a). We first examined the decisions and movements in head-fixed mice. Mice were trained to judge whether the number of auditory clicks was greater on the left versus the right, and to report their decision by licking one of two reward spouts[14]. Neural and behavioral data from some subjects has been reported previously[15]. Movements were tracked via two video cameras and quantified using DeepLabCut (DLC)[16] (Fig. 1B). We then considered the positions of these DLC labels on each trial. Interestingly, the time-varying positions of these DLC labels were informative about many aspects of the animals' decisions. For instance, a cross-validated choice decoder could predict animal choices with high reliability, even before the animal started licking (Fig. 1C, blue line). Decoder performance clearly exceeded chance levels (50%) during the stimulus period and also increased during the delay period when mice could form their decisions based on the stimulus. Surprisingly, a decoder that predicted whether mice obtained a reward in a given trial ('Outcome', Fig. 1C, red line) was consistently above chance levels, even during the baseline period before the animal could observe the stimulus. This suggested to us that some aspect of the animal's movement was related to its overall decision-making performance, allowing a decoder to make a reasonable guess at the trial's outcome based only on these movements. We wondered whether we might gain insight into such movements by separating the animal's average task dependent movements (Fig. 1D, thick black trace schematizes movements that differ for left vs. right trials) from single-trial idiosyncratic deviations from this mean (Fig. 1D, thin gray traces).

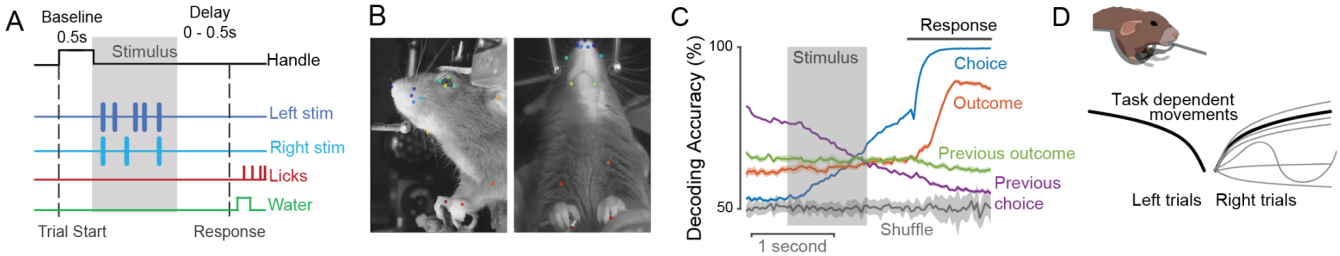

Figure 1. Movements during a decision-making task are informative about choices and current/past outcomes. A. Schematic of an auditory decision-making task in which animals report their judgment about a series of auditory events. B. DeepLabCut (DLC) was used to label 28 body parts in mice. C. A support vector machine (SVM) decoder could predict many aspects of the animals' decisions based on the movements of the DLC labels. D. Animals make uninstructed movements, for instance, reaching to the spout (top). These movements might have a task dependent component (thick black lines) such as differing on average between left and right trials. Movements on individual trials might deviate from these averages (thin gray lines).

To test this hypothesis, we developed a new linear regression model. The goal of the model was to explain the trial-to-trial variance of the DLC label positions using the variables that are part of our decision-making task. These variables include the animal's current/previous choice, the current/previous reward outcome, stimulus rate, and the interaction terms of choice and outcome (Methods). The model learned weights that best mapped these variables to predict the position of each DLC label. As expected from our previous analysis (Fig. 1C), our model could predict as much as half of the movement of the DLC labels (Supp. Fig 1A). However, for each session, some variance in DLC labels was left unexplained by the model. We term this residual variance, *Task Independent Movement* (TIM), as it encompasses movements that are not predicted by task events.

We calculated the magnitude of TIM for each trial in the stimulus and delay windows (1.5-2 s). TIM fluctuated both within and across experimental sessions (Fig. 2A, blue trace). When comparing the magnitude of TIM to the smoothed correct rate on the decision-making task (performance, yellow trace), we observed that these metrics were anti-correlated (Pearson's correlation coefficient; r=-0.51). This negative correlation was consistent across mice (Fig. 2B, left bar). To determine whether this anti-correlation reflects that animals simply make more frequent movements when decision accuracy is low, we calculated the correlation between correct rate and the motion energy of DLC label positions (Fig. 2A, thin gray trace, r=-0.203). This correlation was inconsistent across mice (Fig. 2B, right), arguing that, on average, mice make movements with equal frequency during periods of high vs. low performance. What differs instead for periods of high vs low performance is the trial-to-trial variance of movements (as quantified by TIM); which is consistently smaller during periods of high performance. To better understand the relationship between TIM and decision-making performance, we computed their cross-correlation with shifts of 0 up to 200 trials (Fig. 2C). The negative correlation tended to be smallest when there was close to no lag, suggesting that changes in TIM and decision-making performance change together.

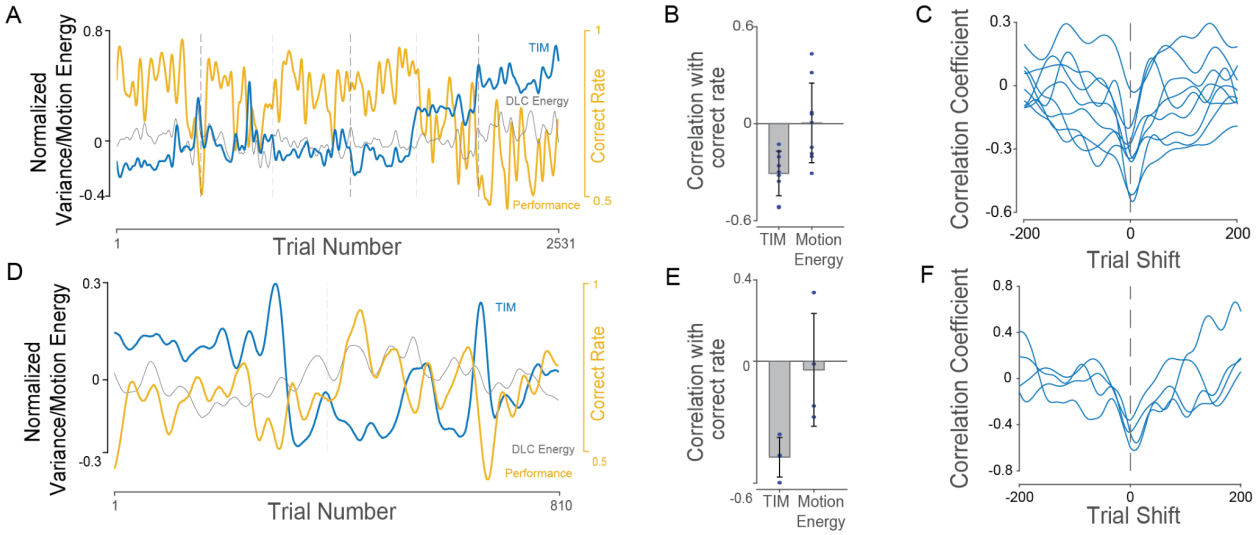

Figure 2. In rats and mice, decision-making performance and task invariant movements (TIM) are inversely related. A. Data from a single mouse over 4 sessions (separated by dashed lines). B. The correlation between correct rate and TIM (left) or motion (right) energy for 7 mice. C. The correlation between TIM and correct rate is most strongly negative when there is a zero lag. D. Same as A but for two sessions from a freely moving rat. E,F. Same as B,C but for 4 freely moving rats.

A possible explanation for the relationship between TIM and decision-making performance is that both metrics are governed by temporal patterns that are stereotyped across sessions (e.g., low performance at the beginning and end, with high performance in the middle). To address this, we measured the cross-session correlation of both TIM and decision-making performance (Supp. Fig. 1B). The cross-session correlation coefficients would be high if these metrics were stereotyped across sessions[17]. The average cross-session correlation for the magnitude of TIM was 0.35, indicating that the dynamics of TIM have some similarity across sessions (specifically, TIM tended to be higher towards the end of each session). However, the average cross-session correlation for decision-making performance averaged was only 0.02, indicating that decision-making performance fluctuates largely independently on each session. Therefore, the relation between TIM and task performance cannot be explained by them having similar patterns that are conserved across sessions.

One possibility is that the relationship between decision-making performance and TIM emerged because the animals in our study were head restrained. Indeed, TIM could simply reflect that animals struggle against restraint when they are uncomfortable and that these moments of discomfort are (unsurprisingly) accompanied by poor decisions. To test this, we repeated this analysis in 4 freely moving rats judging the repetition rate of a series of auditory or visual clicks[18] (Methods; Supp. Fig. 1C). The same inverse relationship between TIM and task performance was present (Fig. 2D, r = -0.46) and was consistent across animals (Fig. 2E,F). Again we observed no consistent relationship between performance and motion energy (Fig. 2E, r=0.337, Fig. 2F). This extends the generality of our observation beyond the restricted context of head-restrained mice and demonstrates that TIM is a strong indicator of an internal state change that is reflected in task performance across different experimental conditions and species.

To more deeply explore the relationship between movements and internal state, we took advantage of a recently developed, generalized linear model coupled to a hidden Markov model[13] (GLM-HMM). This model relies purely on behavioral data to reveal when subjects switch between different behavioral strategies. Previous work using this model showed that mice alternate between distinct behavioral strategies in bouts lasting tens to hundreds of trials[13]. These results suggest that decision-making performance is better described by multiple “internal states”, each of which defines a specific relationship between behavioral variables (e.g. stimulus evidence and bias) and the animals’ choice. An example of these hypothetical internal states is schematized in Figure 3A, reflecting a subject that alternates between bouts of engaged decision making, right-biased behavior, and left-biased behavior. In an “engaged” state, a unit change in stimulus strength leads to a large change in the proportion of a particular choice (Fig. 3A, red curve). In “disengaged” states, by contrast, the relationship is weaker and strong side biases are present. We fit this model to our behavioral data (Methods) and, in keeping with previous results, observed that the mice in our study occupy multiple states in a single experimental session: engaged, left biased, and right biased. Decisions in the left- and right-biased states, which we will refer to collectively as “disengaged” states, rely only weakly on sensory stimuli (Fig. 3B, blue/orange lines). Decisions in the engaged state, by contrast, rely strongly on the stimulus (Fig. 3B, black line). After fitting, we used the forward-backward algorithm to recover posterior state probabilities. This is the probability that a mouse was in a particular state on a given trial (Fig 3C, one example session). We split trials by their most likely state and fit psychometric functions for each. Behavioral data and psychometric curves match what was expected from the GLM-HMM weights (Fig

3D). The biased states show a corresponding shift in the psychometric curves and decreased dependence on stimulus. When we compared these states to TIM, we found a strong relationship: high TIM, shown above to be inversely related to decision-making performance, was likewise inversely related to the probability of being in the engaged state (Fig. 3E-G, r = -0.5119). These results confirm that TIM is closely related to the behavioral states defined by the GLM-HMM.

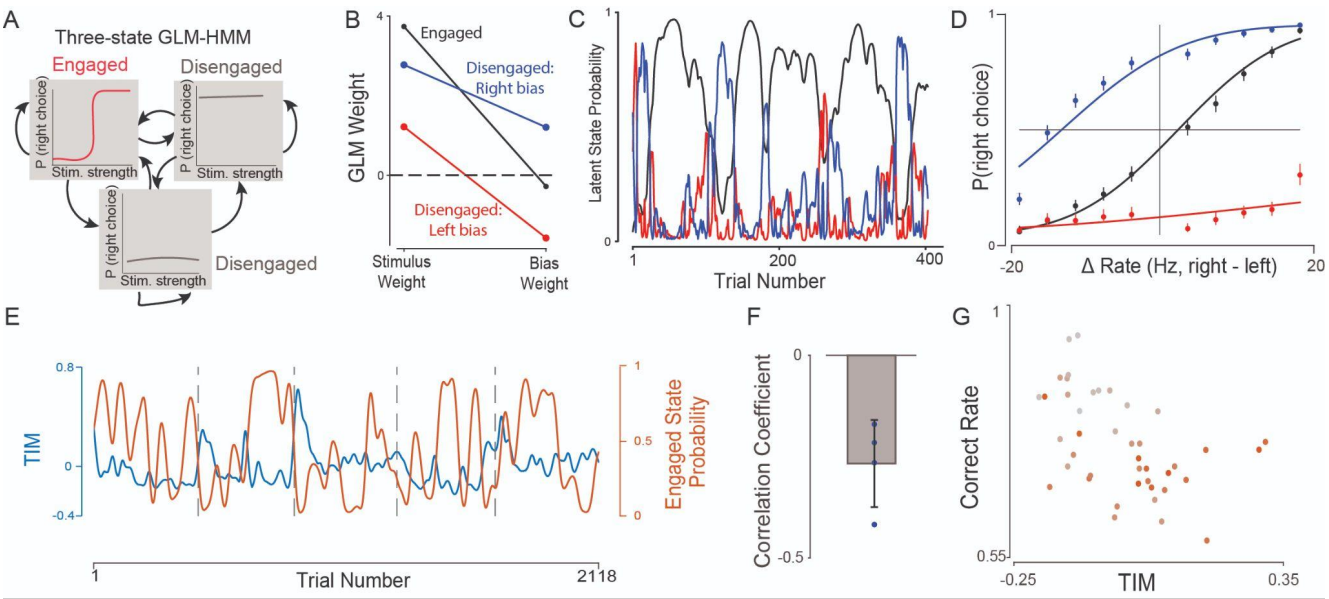

Figure 3. The states generated by a GLM-HMM are related to TIM. A. A hidden Markov model (HMM) was combined with a generalized linear model (GLM) to allow a different function to govern the impact of sensory stimuli depending on the animal's state. Transition probabilities (arrows) describe how likely the animal is to transition from each state into each other state. B-D. Example data from the GLM-HMM (4 mice, 19,578 trials). B. Weight assigned by the GLM to the sensory stimulus and the side bias; each line is a distinct state discovered by the HMM. C. Posterior state probabilities for one session. D. Psychometric functions from each of the three states (color convention same as in B). Error bars: Wilson binomial confidence intervals. Horizontal line: chance performance. Vertical line: category boundary. E. TIM (blue) and the probability of being in the engaged state (orange) for a single mouse. Dashed lines: session breaks. Pearson's correlation coefficient: -0.405. F. A negative correlation between TIM and state was observed in all 4 mice with trained GLM-HMM. G. Scatter plot of TIM and correct rate for the same animal as in (E). Each point reflects the value of TIM for one trial (horizontal axis) and the correct rate for the same moment (calculated over a 50-trial window). Values of TIM were downsampled to match the precision of the correct rate. Points are colored according to whether the animal was in the engaged state (orange) or a disengaged state (gray). $r_{\text{TIM-correct rate}} = -0.4828$, $r_{\text{TIM-engaged state}} = -0.5119$.

The relationship between TIM, decision-making performance, and state suggests that TIM might be informative about the animal's overall engagement level. Pupil diameter is a widely used metric to estimate arousal[19–21], which might be related to the task engagement we estimated with our model[22]. We first analyzed pupil measurements in the same mice reported above (Fig. 4A). The relationship between pupil diameter and correct rate was highly variable, and showed a weak negative correlation (Fig. 4B, r=-0.136). This inconsistent relationship was surprising given prior work[22,23] (see Discussion) and was in contrast to TIM, for which a much stronger and more consistent relationship was present (Fig. 2B,E, left bars). To ensure that the inconsistent relationship between pupil diameter and performance was not due to the video quality of the original dataset, we analyzed 5 additional mice for

which pupil diameter was measured with a much higher quality camera positioned on the animal's eye. These animals were likewise head-restrained and engaged in a visual version of the same task (Methods). Again, we observed a strong negative relationship between correct rate and TIM (Fig. 4C, left, r=-0.267) and a weak, inconsistent relationship between correct rate and standardized pupil diameter (Fig. 4C, right, r=0.023). This difference was also clear when all data from all sessions of a given animal were pooled (Fig. 4D).

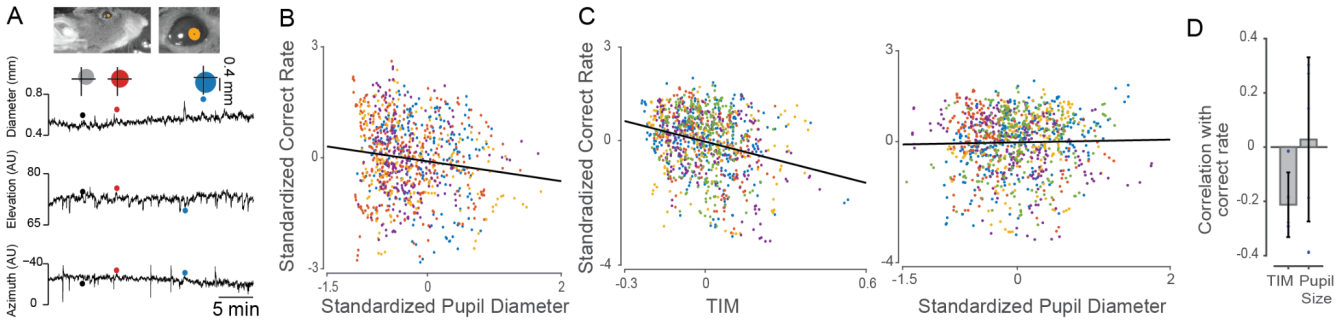

Figure 4. The relationship between TIM and decision-making performance is stronger than the relationship between pupil diameter and decision-making performance. A. Top: Frames from a video camera were used to track pupil diameter. Pupil diameter was estimated using custom written software[24]: the size of the eye was used to estimate the pixel to mm conversion and the area extracted via thresholding after morphological and filtering operations. Bottom: Filtered pupil diameter (area shown as circles) and eye position (shown in the examples as deviation from crosshairs) traces. B. Scatter plot from the same animals in Figs. 2,3. Points of a given color are from the same animal. C. Scatter plot from 5 additional animals showing correlation between correct rate and TIM (left) and pupil diameter (right). Correct rates were standardized for each animal to remove the baseline difference in correct rate. D. Correlation with performance of TIM and Pupil size for each animal; correlations computed over all data from each animal.

To test how internal states, and their associated movement patterns, impact neural activity, we analyzed calcium signals measured during decision-making. Neural activity was measured using widefield calcium imaging (Fig. 5A), allowing us to measure neural activity simultaneously across the entire dorsal cortex of four mice (Ai93; Emx-Cre; LSL-tTA; CaMK2α-tTA). Before analyzing neural activity, we used the GLM-HMM posterior probabilities to assign trials to either "engaged" or "disengaged" categories. We first compared neural activity for trials within these two categories, focusing on four periods of the trial: baseline, stimulus, delay, and response (Fig. 5B). Importantly, we counterbalanced correct vs error trials so that their ratio was equal within the engaged and disengaged trial groups. This ensures that any differences in neural activity were not driven by differences in reward rate. Neural activity was broadly comparable in the two states during the baseline and stimulus periods, but was higher in the engaged state during the delay period (Fig. 5B, bottom row). This elevated activity was most pronounced in sensory and association areas. The temporal dynamics of upper limb somatosensory (Fig. 5C, top) and primary visual cortices (Fig. 5C, bottom) confirmed that the differences in fluorescence were restricted to the delay and response periods. Differences in auditory cortex, which might have been expected on this auditory task, were not evident although the very lateral position of auditory cortex limits its visibility in widefield imaging.

The differences we observed in the magnitude of neural activity (Fig. 5B,C) may have been driven by differences in how the brain encodes either task variables, instructed movements, or uninstructed movements. To distinguish these possibilities, we fit linear encoding models to the neural data for each state. We first fit a model which includes task variables, instructed movements (trial initiation and licking responses), and uninstructed movements (e.g pupil, hindlimb movements, facial video) as regressors. This model accounted for considerable variance in each state (engaged

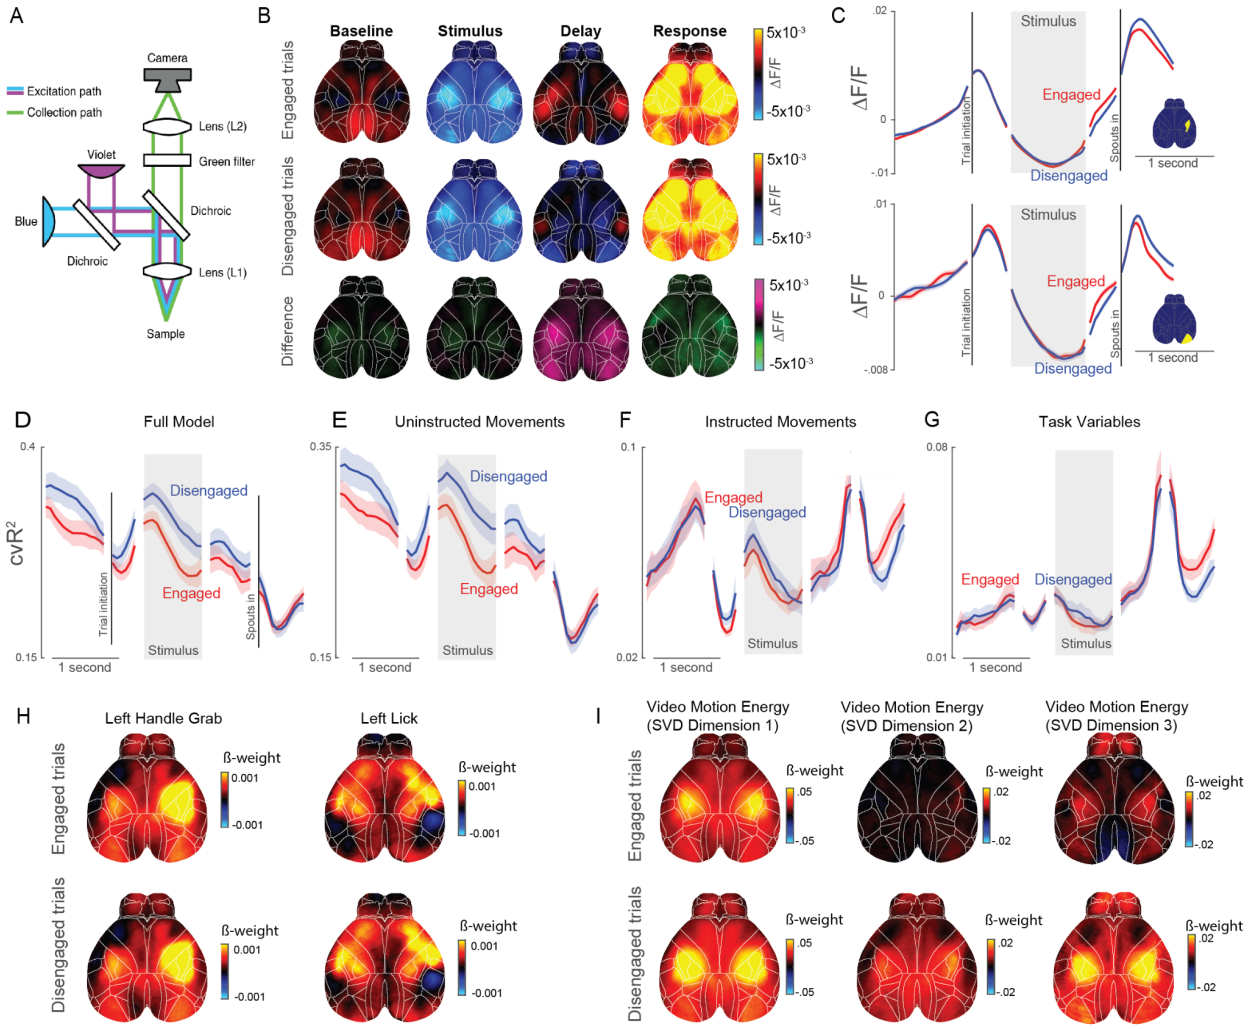

Figure 5. Neural activity differs in engaged vs. disengaged states. A. Schematic and photograph of widefield imaging macroscope (previously published[25]). B. Fluorescence activity (ΔF/F, colorbar) for 4 trial epochs, plotted separately for trials assigned to the engaged (top row) vs. disengaged (middle row) states. Bottom row: the difference in neural activity between the engaged and disengaged states. Correct vs error trials were counterbalanced within sessions so that the ratio of rewarded trials was the same in engaged and disengaged states. Both left and right trials were included. Heatmaps show average responses over 22 sessions from 4 mice. C. Whole-trial temporal dynamics for the engaged (red) and disengaged (blue) state. Data reflect neural activity in primary somatosensory cortex, upper limb area (top) and primary visual cortex (bottom). D-G. Cross-validated R^2 (cvR^2) from the linear encoding model; same color conventions as in (C). D. Full model. E. Uninstructed movements model. F. Instructed movements model. G. Task variables model. H. ß-weights for two instructed movements (left handle grab and a lick to the left spout, 133 ms after event occurrence) are similar for engaged trials (top) and disengaged trials (bottom). ß-weights for uninstructed movements (video motion energy) are

lower for engaged trials than for disengaged trials. The three columns show the beta weights in each state by the top three dimensions of video motion energy.

state cvR^2=0.426, disengaged state cvR^2=0.443; Supp Fig. 2). Interestingly, the full model explained more variance in the disengaged state for most of the trial (Fig. 5D, blue line above red line). We then fit individual models for three groups of regressors: task variables, instructed movements, and uninstructed movements. Uninstructed movements explained more neural variance during disengagement, especially during the stimulus and delay period (Fig. 5E). In contrast, instructed movements (Fig. 5F) and task variables (Fig. 5G) accounted for similar amounts of neural variance in each state during all periods of the task except for the response window.

The greater variance explained by uninstructed movements in the disengaged state (Fig. 5E) could arise either because individual movements impacted neural activity more during the disengaged state, or because the animal made more movements when disengaged. To test this, we compared the ß-weights learned by the model for each state. We verified that the ß-weights for instructed movements were similar for the engaged and disengaged states (Fig. 5H). This implies that a given instructed movement impacts neural activity similarly regardless of the animal's state. By contrast, the ß-weights for video camera motion energy were considerably higher during the disengaged state (Fig. 5I). This suggests that uninstructed movements have a larger impact on neural activity when animals are disengaged.

Discussion

Our observations define a link between decision-making performance, internal state, and animal movements. We demonstrate that the magnitude of TIM, a well-defined subset of animal movements, is inversely related to performance. These movements are defined by their independence to task events and have a much stronger relationship to internal state fluctuations compared to raw motion energy (Fig. 2B,E). Interestingly, during periods of high TIM, animals are likely to be in a state of disengagement, as defined by a separate model, the GLM-HMM (Fig. 3F,G). To understand how TIM and internal states are reflected in the brain, we compared cortex-wide neural activity from the engaged vs. disengaged state. We discovered that neural activity was elevated in the engaged state, particularly during a delay period between the stimulus and choice-response periods (Fig. 5B,C). This is consistent with the possibility that superior performance during the engaged state comes from enhanced working memory capacity that allows the decision to be preserved during the delay. Moreover, uninstructed movements accounted for more cortical activity in the disengaged versus the engaged state (Fig. 5D-G). These movements likely accounted for more variance in neural activity not because they were more frequent during disengagement, but because they had a stronger impact on neural activity during that time (Fig. 5I). Taken together, these observations demonstrate a concrete link between movements and internal state, and argue that the temporal alignment of stereotyped movement patterns, rather than their frequency, is predictive of task engagement.

Our work links together multiple previous observations about uninstructed movements, pupil diameter and neural activity. In previous work on decision-making experts, uninstructed movements account for a large fraction of the variance in neural activity. Movement modulation was present in mice making

auditory/visual decisions[9,10,26], and across multiple types of excitatory cortical neurons[15]. Movement modulation was also evident in area 8 of macaque monkeys making visual decisions[27] arguing that movements have a pronounced impact on neural activity across diverse species. However, how these movements relate to animal cognition was left unclear. The relationship of these movements to pupil diameter and arousal[20,28], as well as to internal states[13], was likewise unknown.

One recent paper examined the relationship between movements and internal state[22]. The authors calculated the probability of being in the optimal state and related this metric to both pupil diameter and a “movement index” that was reflected by face motion energy and running speed. Our work is broadly in keeping with these results: both our approach and theirs found a relationship between internal state and the animal’s movements. However, our study differs in that we found no consistent relationship between performance and the magnitude of video motion energy (Fig. 2B,E), nor between performance and pupil diameter (Fig. 4). A possible explanation for these differences is that our mice were placed on a platform instead of a running wheel, which may have limited the occurrence of high arousal states that are associated with locomotion. Nonetheless, the similar relationship between movements and decision-making performance that we observed in freely-moving rats (Fig. 2D-F) argues against the possibility that the platform in the mouse experiments drove the difference. Another possibility is that there were differences in the criteria used by the experimenters to end the session. Additional states of disengagement, with accompanying movement signatures might be present if animals are permitted to stay in the testing setup for longer durations even as their performance begins to decline. Lastly, a possible difference between our study and theirs is the task structure. Whereas our inter-trial delays are between 1-2 s, their inter-trial delays range from 3-7 s. Longer inter-trial delays would allow the pupil size to recover from reward-induced pupil dilations and for general arousal to decrease before the next trial starts. This is not the case in our task. Thus, one explanation for the inconsistent relationship between pupil size and performance observed here is that the pupil size in our animals was within a more restricted range compared to other studies.

More work is needed to better understand the relationship between movements, internal states and neural activity. First, a study of subcortical structures is needed. Our analysis focused on cortical structures because widefield imaging offered us access to many cortical areas simultaneously, a benefit since internal state could have impacted any of a number of areas. However, future work on subcortical structures could be informative; of particular interest is the posterior nucleus of the thalamus, which has previously shown to be implicated in state-dependent regulation of sensory processing[29]. Second, a manipulation experiment could establish a causal link between movements and internal state. Our approach cannot resolve whether TIM drives the animal to distraction or instead reflects a distracted state that was already initiated. Stimulating movements artificially could potentially resolve this, but the experiments are difficult: if a stimulated movement led to worse performance, it is hard to rule out the interpretation that the stimulation itself caused the change of state. Instead, targeted manipulation of state-modulating regions, such as the hypothalamus[30] or locus coeruleus, could be used to promote an engaged or disengaged state. We could then observe if TIM would reflect such artificial state changes accordingly. Cell-type specific measurements are also needed, and perhaps should extend to astrocytes as well as neurons[31].

Studies in other animal species are likewise needed. Although macaques show clear movement modulation in frontal structures[27], modulation of V1 neurons is more limited[32]. In marmosets, V1 appears at first similar to macaques: running modulates individual neurons only weakly. However, similar to mice, trial-to-trial fluctuations of shared gain modulations are present across marmoset V1, suggesting that internal state could modulate sensory cortex similarly in mice and marmosets[33]. If so, differences between mice and primates in terms of V1 modulation might have more to do with a different link between running and state of engagement in the two species.

Our observations about movements and performance, combined with existing studies on movements and neural activity, point to a closer relationship between movements and cognition than was previously appreciated. Some of our observations are in keeping with conventional wisdom: movements unrelated to the task at hand lend a great deal of insight into an individual's attentional state. This was known by first-grade teachers and is now formalized by our work. However, movements must not be considered monolithically in terms of their impact on cognition. Partitioning movements based on their relationship to task events is essential in determining when they are beneficial or rather harmful to cognition. What remains mysterious is *why* uninstructed movements are related to task performance, as well as the causal direction of this relationship. An intriguing possibility is that uninstructed movements are modified by training or engagement to help the brain more robustly encode task variables, and that these movements could be encouraged to enhance decision making and possible learning. This remains to be seen.

Acknowledgements

This work was supported by U19NS123716, a Multi-university Research Grant from the Office of Naval Research (AW911NF-16-1-0368) and an NSF-NCS collaborative award (2219946). We thank Anup Khanal for assistance with animal handling and training.

Methods

All surgical and behavioral procedures conformed to the guidelines established by the National Institutes of Health and were approved by the Institutional Animal Care and Use Committee of Cold Spring Harbor Laboratory and the University of California, Los Angeles David Geffen School of Medicine. We used 8-25 week old male and female mice and 23-62 week old male Long Evans rats (Taconic Farms). No statistical methods were used to pre-determine sample sizes; sample sizes are similar to previous publications. Mouse strains were acquired from the Jackson Laboratory, Allen Brain Institute, or generated at Cold Spring Harbor Laboratory. The mouse room had a relative humidity of 30-70%, and room temperature ranging from 69-78°F.

General surgical procedures

Surgeries were performed under 1-2% isoflurane in oxygen anesthesia. After induction of anesthesia, 1.2 mg/kg meloxicam was injected subcutaneously and sterile lidocaine ointment was applied topically to the skin incision site. After making a midline cranial incision, the skin was retracted laterally and

fixed in position with tissue adhesive (Vetbond, 3M). We then built an outer wall using dental cement (C&B Metabond, Parkell; Ortho-Jet, Lang Dental) along the lateral edge of the dorsal cranium (frontal and parietal bones). A custom titanium skull post was then attached to the dental cement. For skull clearing, the skull was thoroughly cleaned followed by the application of a thin layer of cyanoacrylate (Zap-A-Gap CA+, Pacer technology)[25].

Mouse behavioral training

The behavioral setup was controlled with a microcontroller-based (Arduino Due) finite state machine (Bpod r0.5, Sanworks) using custom Matlab code (2015b, Mathworks) running on a Linux PC. Servo motors (Turnigy TGY-306G-HV) and touch sensors were controlled by microcontrollers (Teensy 3.2, PJRC) running custom code. Mice were trained on a delayed, spatial discrimination task. Mice initiated trials by placing their forepaws on at least one of the two handles, which were mounted on servo motors that rotated out of reach during the inter-trial period. To initiate trials, animals placed their forepaws on the handles and, after a variable duration of 0.25-0.75 s of continuous contact, the auditory stimulus was presented. Auditory stimuli consisted of a sequence of Poisson-distributed, 3-ms long auditory click sounds, presented from either a left and/or right speaker for a variable duration between 1 and 1.5 s. A subset of 5 mice (Figure 4A,C,D) completed a visual version of this task in which circular grating stimuli were presented on either the left or right side of a screen. For all mice, the stimulus period was followed by a variable delay (up to 1 s), after which servo motors moved two lick spouts close to the animal's mouth. If the animal licked twice on the side where more sensory events were presented, a drop of water reward was dispensed. After a spout was licked twice, the contralateral spout moved out of reach to force the animal to commit to its decision. The water volume rewarded per trial (typically 1.5 to 3 μL) was constant within a single session but was sometimes adjusted daily based on the animal's body weight.

Data was collected from multiple sensors in the behavioral setup. Touch sensors using a grounding circuit on handles and lick spouts detected contact with the animal's forepaws and tongue, respectively. A piezo sensor (1740, Adafruit LLC) below the animal's trunk was used for monitoring body and hindlimb movements. Two webcams (C920 and B920, Logitech) were positioned to capture the animal's face (side view) and the ventral surface of the body (ventral view). Pupil data in Figure 4 (A,C and D) were acquired with CM3-U3-13Y3M-C S cameras (FLIR) at 60 Hz, synchronized using digital input and output pins.

Trained mice were housed in groups of two or more under reverse light cycle (12-hour dark and 12-hour light) and trained during their active dark cycle. Animals were trained over the course of approximately 30-60 days. After 2-3 days of restricted water access, animals began habituation to head fixation and received water from spouts in the behavior chamber. During these sessions, unilateral auditory stimuli were presented followed by a drop of water from the ipsilateral water spout. After several habituation sessions, animals were required to touch the handles to trigger stimulus presentation. Once mice could reliably reach for the handles, the required touch duration was progressively increased to 0.75 s. During the next training stage, both spouts moved within reach of the animal following stimulus presentation. An animal was considered trained when its detection performance across two or more sessions was >80%.

Rat behavioral training

We trained freely moving rats on a 2-choice auditory/visual decision task[34]. Experiments took place in dark, sound-proof rigs (Industrial Acoustics, Bronx, NY). Trials with visual, auditory, and multisensory stimuli were interleaved. In the behavior setup (composed of one speaker, one LED board, and three ports), the rat initiated the trial by putting its nose in the central port. Next, the speaker and/or LED board started to play the stimulus sequence (sound clicks, white light flashes, or both). The number of stimulus events, presented over 1000 ms, ranged from 7 to 17. The decision boundary (12 Hz) was abstract and learned with experience by the animal. When the number of stimulus events was above 12 Hz, right choices were rewarded; when the number was below 12 Hz, left choices were rewarded. 12 Hz stimuli were rewarded randomly. The duration of the inter-event interval followed a Poisson distribution to ensure the stimuli' randomness. Thus, the rats must make each decision using only the rate information on that trial, rather than remembering a set of fixed sequences.

The stimulus window was followed by a short delay period. The duration of the delay period was randomly assigned from an exponential distribution (mean = 0.12 s). After this delay period, an auditory go cue informed the animal that it could leave the central port and make the choice. After the rat poked its nose into the left or right port to report the choice, a reward (24 μL water) or a punishing sound (2-s long sharp noise) was delivered based on whether the choice was correct. A 2-s long white noise sound was delivered if the rat left the central port too early (i.e., during the stimulus or delay period).

A webcam (CM3-U3-13Y3M-CS 1/2" Chameleon®3 Monochrome Camera) was used to monitor animal movements. The camera was positioned above the animal and sampled 80 Hz.

Decoding task variables based on DLC labels

Support vector machine (SVM) decoders were trained to categorize task variables based on the DLC labels of the head-fixed mouse and freely moving rat datasets. Four task variables were selected: choice, outcome, the choice of previous trial, and the outcome of previous trial. After balancing the number of sample trials, the position of all DLC labels in each frame was used as the input to predict the task variable. SVM was trained with function *fitclinear* in MATLAB. Because balancing the number of trials requires downsampling, the SVM training and prediction were replicated ten times to avoid potential biases induced by the downsampling. All the predictions were 5-fold cross-validated.

Relating movements to decision-making performance.

Our new metric, TIM, quantifies the magnitude of movements that are independent of task events. To calculate TIM, we first used linear regression to capture the variance in DLC label positions that can be explained by task variables. 7 task variables were used in the regression model to predict the position of DLC labels: stimulus strength, choice, reward, the choice-reward interaction in the current trial, and choice, outcome, and the choice-reward interaction in the previous trial. Separate regressions were performed for each DLC label at each time point in the trial. The regression model was fitted with the *fitlm* function in MATLAB. Before doing the regression, we subtracted the mean position of each DLC label of each session to reduce the effect of camera shift across sessions.

The equation to predict DLC label positions as a function of task variables is as follows:

$$\vec{y}(t) = \beta_{1t}x_{stim} + \beta_{2t}x_{choice} + \beta_{3t}x_{reward} + \beta_{4t}x_{interaction} + \beta_{5t}x_{choice-1} + \beta_{6t}x_{reward-1} + \beta_{7t}x_{interaction-1} + \beta_{0t}$$

After quantifying the variance dependent on the task variables (Supp. Fig 1A), we defined TIM as the Euclidean distance between the actual DLC label position and the model-predicted position. We calculated the mean value of TIM of all DLC labels across all time points in each trial, and then used this value as the TIM of the corresponding trial. To ensure that TIM was not dominated by the body parts that have large movement amplitude, the TIM value of each DLC label was standardized (z-score) across time before the mean value was calculated.

Notably, animals showed stereotypical licking behavior on rewarded trials, which would reduce the value of TIM on correct trials in the response window. To avoid this confound, the linear regression and TIM calculation was performed within the stimulus and delay windows, before rewards (or punishment tones) were delivered.

Before calculating the correlation coefficient between TIM and performance/probability of being in the engaged state, TIM was smoothed with a 50-trial long window with function *smoothdata* in MATLAB. The smoothing filter was a Gaussian-weighted moving average filter. Correlations from data in Figs. 3G and Fig. 4B,C were calculated on downsampled data to eliminate confounds from the sliding window.

Inferring latent states with the GLM-HMM

To assess the correlation between TIM and states of engagement, we fit a 3-state hidden Markov model with Bernoulli generalized linear model observations (GLM-HMM)[13]. Model hyperparameters (including the number of latent states) were first selected with a grid search, where hyperparameter performance was assessed via 10-fold cross validation. The GLM-HMM parameters were fit as previously described[13] using the SSM Python package[35]. We used maximum a posteriori estimation to estimate the HMM transition matrix and GLM weights. Models were fit to the combined data from all mice. Posterior state probabilities were then inferred via the forward-backward algorithm, also implemented within the SSM package.

Psychometric function fitting

Data from all mice was combined to fit psychometric functions for different states (Fig. 3D). Psychometric functions were modeled with a four-parameter cumulative Gaussian (bias, threshold, left lapse, right lapse). Parameters were then estimated using Nelder-Mead Optimization.

Widefield imaging

Widefield imaging was done as reported previously[15,25,36] using an inverted tandem-lens macroscope and an sCMOS camera (Edge 5.5, PCO) running at 30 fps. The top lens had a focal length of 105 mm

(DC-Nikkor, Nikon); the bottom lens 85 mm (85M-S, Rokinon). The field of view was 12.5 x 10.5 mm^2 and the imaging resolution was 640 x 540 pixels after 4x spatial binning, resulting in a spatial resolution of ~20 µm per pixel. To capture fluorescence, a 525 nm band-pass filter (#86-963, Edmund optics) was placed in front of the camera. Using excitation light at two different wavelengths, we isolated Ca^{2+}-dependent fluorescence and corrected for intrinsic signals (e.g., hemodynamic responses). Excitation light was projected on the cortical surface using a 495 nm long-pass dichroic mirror (T495lpxr, Chroma) placed between the two macro lenses. The excitation light was generated by a collimated blue LED (470 nm, M470L3, Thorlabs) and a collimated violet LED (405 nm, M405L3, Thorlabs) that were coupled into the same excitation path using a dichroic mirror (#87-063, Edmund optics). We alternated illumination between the two LEDs from frame to frame, resulting in one set of frames with blue and the other with violet excitation at 15 fps each. Excitation of GCaMP at 405 nm results in non-calcium dependent fluorescence, allowing us to isolate the true calcium-dependent signal by rescaling and subtracting frames with violet illumination from the preceding frames with blue illumination. Subsequent analyses were based on this differential signal. Imaging data was rigidly aligned to the Allen common coordinate framework (CCF), using four anatomical landmarks: the left, center and right points where the anterior cortex meets the olfactory bulbs, and the medial point at the base of the retrosplenial cortex.

Preprocessing of neural data

We used a rigid-body image registration method implemented in the frequency domain[37] to align each imaging frame to the median over all frames in the first trial. To reduce the computational cost of subsequent analyses, we then computed the 200 highest-variance components using singular value decomposition (SVD). These components accounted for at least 95% of the total variance in each recording, whereas computing 500 components accounted for little additional variance (~0.15%). SVD reduces the raw imaging data Y to a matrix of ‘spatial components’ U (of size pixels by components), ‘temporal components’ VT (of size components by frames) and singular values S (of size components by components) to scale temporal components to the original data. The resulting decomposition has the form Y = USVT. All subsequent analysis in the time domain (such as trial averages and the linear encoding model) were performed on the product SVT and the respective results were later multiplied with U, to recover results for the original pixel space. To avoid slow drift in the imaging data, SVT was high-pass filtered above 0.1 Hz using a zero-phase, second-order Butterworth filter.

To compute trial averages and perform encoding model analysis (see below), imaging data in individual trials were aligned to the four trial periods, each marked by a specific event. This was required because the duration of different trial events was randomized to reduce temporal correlations, e.g. between trial initiation, the stimulus presentation and subsequent lick responses. The first period (Initiate) was aligned to the time when animal initiated a trial by touching the handles, the second (Stimulus) was aligned to the stimulus onset, the third (Delay) to the end of the stimulus sequence, and the fourth (Response) to the time when spouts were moved in to allow a lick response.

Linear encoding model

The linear encoding model included task- and movement-related variables as described previously[9]. Each variable consisted of multiple regressors that were combined into a larger design matrix. Binary regressors contained a single pulse that signaled the occurrence of specific events, such as the stimulus onset, and additional regression copies that were shifted forward or backward in time to account for changes in cortical activity before or after the respective event. For auditory stimuli, the time-shifted copies spanned all frames from the onset of the auditory sequence until the end of the trial. Individual click sounds were also captured by an additional regressor set that spanned the 2 seconds after click onset. For licks and whisks, the time-shifted copies spanned the frames from 0.5 s before until 5 s after each event. For some variables, e.g., previous choice, the time-shifted copies spanned the whole trial. Other variables were analog, such as measures from the piezo sensor, pupil diameter, and the 200 highest temporal components of video information from both cameras (using SVD as described above). This ensured that the model could account for animal movements and accurately isolate task-related activity. Movement and task variables were additionally decorrelated due to the variable durations of the initiation, stimulus and delay period. The model was fit using ridge regression to allow for similar contributions from different correlated variables. To determine the regularization penalty λ for each column of the widefield data, we used marginal maximum likelihood estimation (MLE)[38]. MLE expresses the encoding model as a Bayesian linear model and determines the ridge penalty λ by maximizing the marginal likelihood $\pi(D \mid \lambda)$ of the model, given data D. This was done iteratively by testing different λ values to determine a global minimum for the negative log-likelihood $-\log \pi(D \mid \lambda)$. The main advantage of this approach is that λ can be determined without computationally expensive cross-validation procedures, resulting in a ~50-fold decrease in required compute time on a regular workstation. Moreover, the faster MLE approach allows adjusting λ values for individual widefield data components which results in higher cross-validated explained variance of the encoding model, compared to a regular cross-validation approach.

Encoding models were fit to data within single sessions. We fit the encoding model twice for every session: first for trials assigned by the GLM-HMM to the engaged state and then for trials assigned by the GLM-HMM to the disengaged state. Because there were typically more trials in the engaged state, these were randomly downsampled so that the trial numbers in each state for a session were always matched. Sessions with fewer than 32 engaged and 32 disengaged trials were discarded to eliminate unreliable model fits.

References

1. Wiltschko, A. B. *et al.* Mapping Sub-Second Structure in Mouse Behavior. *Neuron* **88**, 1121–1135 (2015).
2. Heinrichs, S. C. Nonexercise muscle tension and behavioral fidgeting are positively correlated with food availability/palatability and body weight in rats. *Physiol. Behav.* **79**,

199–207 (2003).

3. Erisken, S. *et al.* Effects of Locomotion Extend throughout the Mouse Early Visual System. *Curr. Biol.* **24**, 2899–2907 (2014).
4. Niell, C. M. & Stryker, M. P. Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. *Neuron* **65**, 472–479 (2010).
5. Mineault, P. J., Tring, E., Trachtenberg, J. T. & Ringach, D. L. Enhanced Spatial Resolution During Locomotion and Heightened Attention in Mouse Primary Visual Cortex. *J. Neurosci.* **36**, 6382–6392 (2016).
6. Fu, Y. *et al.* A Cortical Circuit for Gain Control by Behavioral State. *Cell* **156**, 1139–1152 (2014).
7. Dolensek, N., Gehrlach, D. A., Klein, A. S. & Gogolla, N. Facial expressions of emotion states and their neuronal correlates in mice. *Science* **368**, 89–94 (2020).
8. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255–255 (2019).
9. 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).
10. Salkoff, D. B., Zagha, E., McCarthy, E. & McCormick, D. A. Movement and Performance Explain Widespread Cortical Activity in a Visual Detection Task. *Cereb. Cortex N. Y. N 1991* **30**, 421–437 (2020).
11. Tremblay, S., Testard, C., DiTullio, R. W., Inchauspé, J. & Petrides, M. Neural cognitive signals during spontaneous movements in the macaque. *bioRxiv* (2022) doi:10.1101/2022.09.05.506681.
12. Wu, H. G., Miyamoto, Y. R., Castro, L. N. G., Ölveczky, B. P. & Smith, M. A. Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat. Neurosci.* **17**, 312–321 (2014).
13. Ashwood, Z. C. *et al.* Mice alternate between discrete strategies during perceptual

decision-making. *Nat. Neurosci.* **25**, 201–212 (2022).

14. Brunton, B. W., Botvinick, M. M. & Brody, C. D. Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. *Science* **340**, 95–98 (2013).

15. Musall, S. *et al.* Pyramidal cell types drive functionally distinct cortical activity patterns during decision-making. *Nat. Neurosci.* (2023) doi:10.1038/s41593-022-01245-9.

16. Mathis, A. *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* **21**, 1281–1289 (2018).

17. Harris, K. D. *Nonsense correlations in neuroscience*. http://biorxiv.org/lookup/doi/10.1101/2020.11.29.402719 (2020) doi:10.1101/2020.11.29.402719.

18. Raposo, D., Kaufman, M. T. & Churchland, A. K. A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* **17**, 1784–1792 (2014).

19. McGinley, M. J. *et al.* Waking State: Rapid Variations Modulate Neural and Behavioral Responses. *Neuron* **87**, 1143–1161 (2015).

20. Nassar, M. R. *et al.* Rational regulation of learning dynamics by pupil-linked arousal systems. *Nat. Neurosci.* **15**, 1040–1046 (2012).

21. Granholm, E. & Steinhauer, S. R. Pupillometric measures of cognitive and emotional processes. *Int. J. Psychophysiol.* **52**, 1–6 (2004).

22. Hulsey, D., Zumwalt, K., Mazzucato, L., McCormick, D. A. & Jaramillo, S. *Decision-making dynamics are predicted by arousal and uninstructed movements*. http://biorxiv.org/lookup/doi/10.1101/2023.03.02.530651 (2023) doi:10.1101/2023.03.02.530651.

23. McGinley, M. J., David, S. V. & McCormick, D. A. Cortical Membrane Potential Signature of Optimal States for Sensory Signal Detection. *Neuron* **87**, 179–192 (2015).

24. Aydın, Ç., Couto, J., Giugliano, M., Farrow, K. & Bonin, V. Locomotion modulates specific functional cell types in the mouse visual thalamus. *Nat. Commun.* **9**, 4882 (2018).

25. Couto, J. *et al.* Chronic, cortex-wide imaging of specific cell populations during behavior. *Nat. Protoc.* **16**, 3241–3263 (2021).

26. Bimbard, C. *et al.* Behavioral origin of sound-evoked activity in mouse visual cortex. *Nat. Neurosci.* **26**, 251–258 (2023).

27. Tremblay, S., Testard, C., DiTullio, R. W., Inchauspé, J. & Petrides, M. *Neural cognitive signals during spontaneous movements in the macaque*. http://biorxiv.org/lookup/doi/10.1101/2022.09.05.506681 (2022) doi:10.1101/2022.09.05.506681.

28. Urai, A. E., Braun, A. & Donner, T. H. Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nat. Commun.* **8**, 14637 (2017).

29. Roth, M. M. *et al.* Thalamic nuclei convey diverse contextual information to layer 1 of visual cortex. *Nat. Neurosci.* **19**, 299–307 (2016).

30. Allen, W. E. *et al.* Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science* **364**, eaav3932 (2019).

31. Reitman, M. E. *et al.* Norepinephrine links astrocytic activity to regulation of cortical state. *Nat. Neurosci.* **26**, 579–593 (2023).

32. Talluri, B. C. *et al. Activity in primate visual cortex is minimally driven by spontaneous movements*. http://biorxiv.org/lookup/doi/10.1101/2022.09.08.507006 (2022) doi:10.1101/2022.09.08.507006.

33. Liska, J. P. *et al. Running modulates primate and rodent visual cortex via common mechanism but quantitatively distinct implementation*. http://biorxiv.org/lookup/doi/10.1101/2022.06.13.495712 (2022) doi:10.1101/2022.06.13.495712.

34. Raposo, D., Sheppard, J. P., Schrater, P. R. & Churchland, A. K. Multisensory Decision-Making in Rats and Humans. *J. Neurosci.* **32**, 3726–3735 (2012).

35. Linderman, S., Antin, B., Zoltowski, D. & Glaser, J. SSM: Bayesian Learning and

Inference for State Space Models. (2020).

36. 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).

37. Reddy, B. S. & Chatterji, B. N. An FFT-based technique for translation, rotation, and scale-invariant image registration. *IEEE Trans. Image Process.* **5**, 1266–1271 (1996).

38. Karabatsos, G. Marginal maximum likelihood estimation methods for the tuning parameters of ridge, power ridge, and generalized ridge regression. *Commun. Stat. - Simul. Comput.* **47**, 1632–1651 (2018).

Supplementary Figure 1

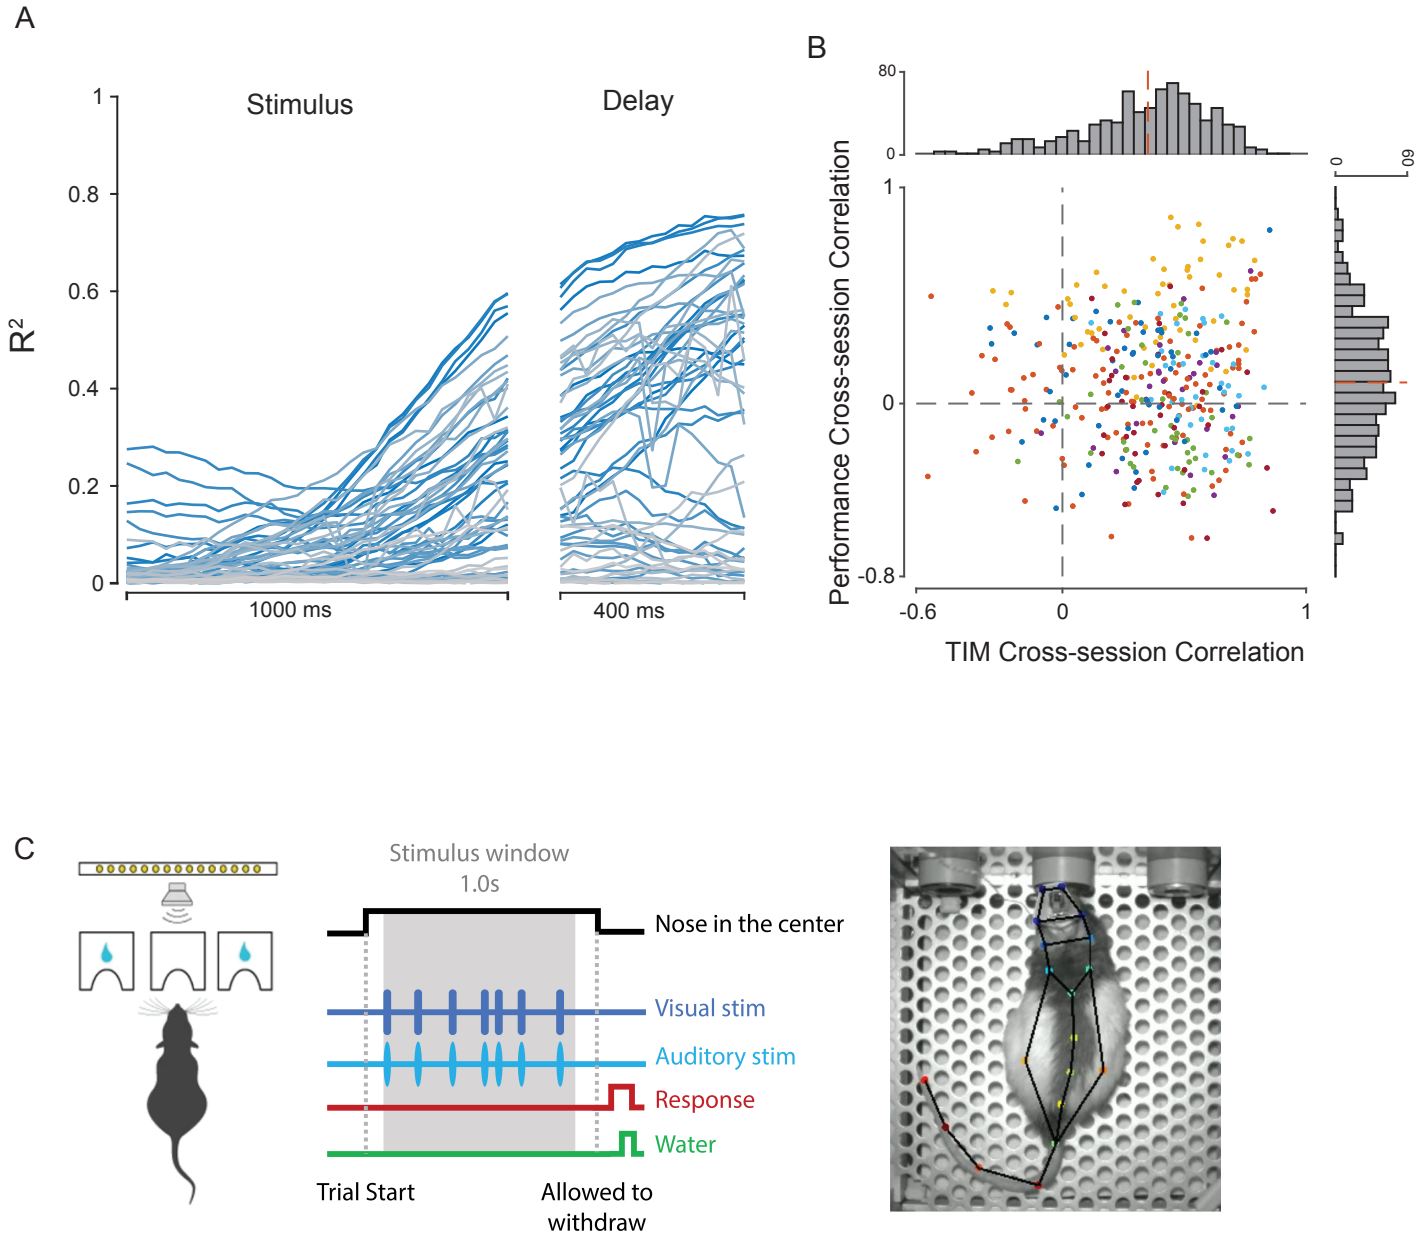

A. R^2 of the behavioral model for each from an example session in one mouse. Each line shows the R^2 for a single DeepLabCut label (x- and y-values indicated by gray and blue traces) during the stimulus and delay periods. B. Scatter plot: the cross-session correlation of correct rate (Performance, vertical axis) as a function of the cross-session correlation for TIM (horizontal axis). Each point shows the value for one animal and one session. Marginal histograms reflect TIM (top) or performance (right). Red dashed lines: mean value across all animals/sessions. C. Decision-making task in freely moving rats. Left: Animals initiate trials and report choices about lights/sounds by interacting with nose pokes. Middle: Schematic of task events. Right: DeepLabCut was used to label and track 19 body parts.

Supplementary Figure 2

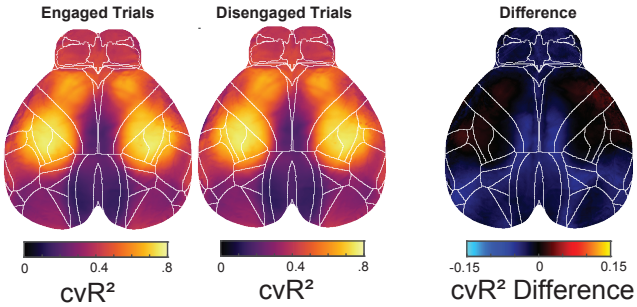

Cross validated R^2 of the encoding model fit to fluroescence data. Left: data was fit with trials from the engaged state. Middle: data was fit with trials from the disengaged state. Right: The difference between the two tended to be close to 0.