

Spontaneous behavior is a succession of self-directed tasks

Highlights

- Spontaneous mouse behavior is organized into task-like behavioral states
- A hierarchical model (shMoSeq) identifies states from behavioral syllables
- The prefrontal cortex preferentially encodes state-related variables relative to low-level movements
- Prefrontal lesions disrupt context-appropriate, long-timescale behavioral organization

Authors

Caleb Weinreb,
Lakshanya Thamarai Kannan,
Alia Newman-Boule, ...
Mohammed Abdal Monium Osman,
Scott W. Linderman,
Sandeep Robert Datta

Correspondence

srdatta@hms.harvard.edu

In brief

Weinreb et al. reveal a hierarchy of timescales in mouse behavior, including low-level syllables and high-level behavioral states. States and syllables are encoded in different brain areas. Prefrontal encoding of behavioral states identifies principles common to both structured tasks and spontaneous behavior.

Article

Spontaneous behavior is a succession of self-directed tasks

Caleb Weinreb,¹ Lakshanya Thamarai Kannan,¹ Alia Newman-Boule,¹ Tim Sainburg,¹ Winthrop F. Gillis,¹ Alex Plotnikoff,¹ Sofia Makowska,¹ Jonah E. Pearl,¹ Mohammed Abdal Monium Osman,¹ Scott W. Linderman,^{2,3} and Sandeep Robert Datta^{1,4,*}

¹Department of Neurobiology, Harvard Medical School, Boston, MA 02115, USA

²Wu Tsai Neurosciences Institute, Stanford University, Stanford, CA 94305, USA

³Department of Statistics, Stanford University, Stanford, CA 94305, USA

⁴Lead contact

*Correspondence: srdatta@hms.harvard.edu

<https://doi.org/10.1016/j.neuron.2025.11.021>

SUMMARY

Animals achieve high-level goals by sequencing low-level actions. This transformation is best understood in structured tasks that impose a specific mapping between goals and actions. However, it remains unclear whether spontaneous behavior is similarly organized in the service of identifiable goals or how it might be supported by brain regions responsible for goal-oriented behavior, such as the prefrontal cortex (PFC). Here, we show that low-level actions in freely exploring mice are hierarchically organized into seconds-long behavioral states that correspond to task-like programs of behavior. These persistent states structure neural activity in the PFC, which preferentially encodes the identity of states relative to low-level behavioral features and shapes which states are expressed in a given context. These findings argue that spontaneous behavior is organized as a succession of self-directed tasks and identify principles of neural control that are common to structured tasks and spontaneous exploration.

INTRODUCTION

The brain translates goals into actions. Systems neuroscience typically studies this transformation through structured tasks in which animals generate specific actions (e.g., licking left or right) to earn rewards.¹ In this framework, movements linked to reward are considered purposive.^{2,3} Evidence from such tasks suggests that goal-oriented behaviors rely on the prefrontal cortex (PFC), which encodes key information about context-, task-, and value-related variables.^{4–15} In contrast, when animals act spontaneously—as they do most of the time in the natural world—it is not obvious whether actions are organized around specific goals or whether the PFC similarly encodes task-related information.¹⁶

Inferring the purpose of behavior is challenging because most low-level actions are compatible with many high-level goals.^{17,18} For example, mice exploring an open field can establish home bases, plan escape routes, and attend to novel objects and odors—all of which employ common actions (e.g., running, sniffing, and rearing) that might also occur in the absence of a specific goal.^{19–22} Consistent with the idea that spontaneous behavior is not always goal oriented, PFC lesions have little effect on open-field exploration, at least when measured using conventional metrics.^{23–27}

One strategy for connecting low-level actions to high-level goals arises from ethology, which uses the grouping of actions over time to illuminate their functions.^{28,29} An animal might

crouch, stalk, and sprint while hunting, for example, but burrow and manipulate bedding to build a nest.²⁹ This organization leaves a statistical signature: frequently co-occurring movements are likely to serve the same goal, and sudden shifts between groups of movements can mark when one goal succeeds another.³⁰ While a recent explosion of machine learning tools has given us unprecedented access to animals' low-level movements, including the sub-second behavioral motifs (or “syllables”) that compose mouse behavior,^{31–33} we currently lack tools for identifying behavioral states on longer timescales.³⁴

Given such a tool, several criteria would, if fulfilled, indicate that the identified states reflect goal-directed tasks. First, each state should include movements that are *prima facie* related to a task. Second, the states expressed in a given context should depend on available affordances³⁵ (i.e., elements of the environment that allow some form of interaction). Finally, information about states and corresponding task-related variables should be present in brain regions that support goal-directed behavior, such as the PFC.^{4,11,12,14,15,36} Indeed it is sometimes possible to infer the structure of a task *a priori* from PFC recordings, as recently demonstrated in rats trained to play hide-and-seek.³⁷ Furthermore, in social contexts, PFC neurons represent spatial location, social proximity, and conspecific actions, all of which are key to effective social behavior.^{38–44} However, there have been few recordings of the PFC outside of structured tasks or

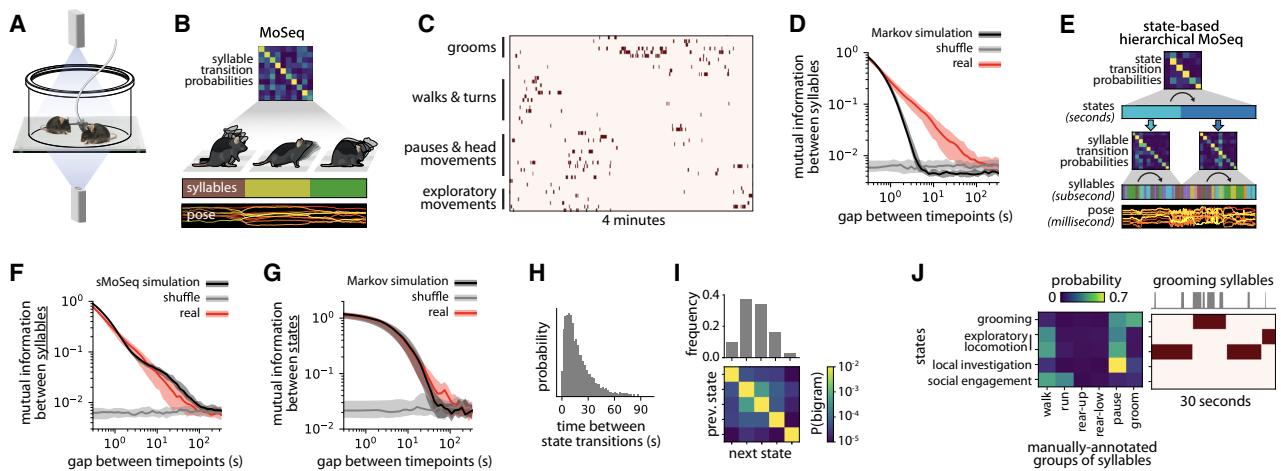


Figure 1. Behavioral states capture the higher-order structure in exploratory behavior

(A) Schematic of the arena used for behavioral recordings.

(B) Schematic describing MoSeq. Pose trajectories (bottom) are determined by syllables (middle), which transition with fixed probabilities (top).

(C) Example syllable sequence. The raster shows which syllable (row) is expressed at each time point (column).

(D) Mutual information (MI) between syllables at a range of temporal lags, shown for real data (red), Markov model simulations (black), and shuffled data (gray). Line and shading indicate median and interquartile interval across recordings ($n = 31$). Syllables were kinematically clustered for MI computation to prevent sparsity-related artifacts.

(E) Schematic describing shMoSeq. A chain of behavioral states (top) specifies transition probabilities between syllables (middle), which govern pose (bottom).

(F) As (D), here comparing real data (red) to shMoSeq simulations (black).

(G) MI between behavioral states at a range of temporal lags, plotted as in (D). States were derived from real data using shMoSeq (red) or simulated from a Markov model (black).

(H) Distribution of state durations.

(I) Top: usage of each behavioral state. Bottom: bigram frequencies, capturing transitions between states.

(J) Left: overlap between states and manually annotated categories of syllables. Right: occurrence of grooming syllables (gray ticks) alongside behavioral states (heatmap) from an example interval.

social interactions (but see Lindsay et al.⁴⁵ and Maisson et al.⁴⁶), and so whether the PFC encodes hierarchical features of spontaneous behavior remains largely unexplored.

Here, we ask whether spontaneous behavior is organized into self-directed tasks that are legible in PFC activity. To do so, we develop a novel hierarchical model that identifies seconds-to-minutes-long behavioral states from patterns of syllables. We find that dorsomedial PFC (dmPFC) activity is dominated by behavioral states rather than low-level movements. Lesions demonstrate that the dmPFC promotes the expression of less-used behavioral states when appropriate in a given context. These results suggest that spontaneous behavior consists of self-initiated tasks whose appropriate selection depends on the dmPFC, and establishes common principles shared by experimenter-defined tasks and self-guided exploration.

RESULTS

To test whether spontaneous behavior is organized into task-like states anchored to neural activity, we performed calcium imaging in the dmPFC⁴⁷ while monitoring behavior with depth camera-based motion sequencing (MoSeq). MoSeq is an unsupervised machine learning approach for segmenting behavior into sub-second syllables, such as rears, turns, or sniffs.^{32,48–50} We recorded male mice in an open-field arena that was alternately empty or supplemented with a male conspecific belonging to a

strain (C57) that is less prone to aggression in novel environments (Figures 1A and S1A–S1C; Videos S1 and S2).⁵¹ In this setting, MoSeq identified 72 behavioral syllables, including various forms of locomotion, rearing, grooming, and investigation (Figures 1B and S1D–S1F).⁵²

Spontaneous behavior is organized into higher-order behavioral states

Ethologists recognize structure in animal behavior by analyzing how actions cluster over time; behavior possesses higher-order structure when this clustering cannot be explained through simple Markovian transition rules.^{30,34} In our recordings, subsets of kinematically related syllables co-occurred in bouts lasting up to a minute, suggesting that spontaneous behavior is structured over long timescales (Figure 1C). To test whether this structure was non-Markovian, we compared observed syllable sequences with synthetic sequences drawn from a Markov model. In the synthetic sequences, mutual information between syllables (which captures the predictability of behavior over time) decayed to chance in just a few seconds (Figure 1D).⁵³ In contrast, real syllable sequences remained predictable for tens of seconds (Figures 1D and S1G). This order-of-magnitude gap persisted across a variety of syllable timescales and resolutions⁵⁴ (Figures S1H–S1M). Thus, simple transition rules cannot explain the persistent syllable use patterns captured by MoSeq, implying that higher-order states may organize behavior on a seconds-to-minutes timescale.

Hierarchical hidden Markov models (HHMMs) offer one approach for identifying such states. In an HHMM, each state specifies the expected usage and transition frequencies between MoSeq syllables. Until now, hierarchical extensions of MoSeq have faltered on account of the large number of syllables and their strong autocorrelations³² (Figures S2A and S2B). To overcome these challenges, we devised a new type of data-efficient HHMM called state-based hierarchical MoSeq (shMoSeq). shMoSeq has three levels corresponding to different timescales in behavior: a chain of multi-second behavioral states, which bias the selection of sub-second syllables, which, in turn, govern the dynamics of pose (Figure 1E). Each behavioral state is defined by a unique matrix describing transitions between syllables, which shMoSeq constructs by combining a baseline transition matrix (that is shared across states) with syllable usage biases that are unique to each state.

shMoSeq effectively captures the range of timescales in behavior; synthetic syllable sequences generated by shMoSeq were virtually identical to real syllable sequences when assessed by mutual information, indicating that the dynamics instantiated by shMoSeq are sufficient to explain the temporal correlations observable in mouse behavior (Figure 1F). Furthermore, in contrast to syllables, shMoSeq-derived state sequences were well approximated by a simple Markov chain, suggesting that additional hierarchical layers are not necessary to explain behavior within the 1-h time span of our recordings (Figure 1G).

During our initial experiments (in which mice explored an open field and interacted with a C57 conspecific), shMoSeq identified five behavioral states, including grooming, investigation (of either the mouse's immediate surroundings or the conspecific), and two forms of exploratory locomotion in which mice circumnavigated the arena in a clockwise or counterclockwise pattern (Figures 1H, 1I, and S2C–S2J; Video S3). Consistent with the premise that similar movements can serve different high-level goals,³⁰ the mapping between syllables and states was not one-to-one; except for grooming, most syllables were used in most states, and even grooming syllables were neither limited to—nor continuous during—the grooming state (Figure 1J). Consequently, there was often an interval of uncertainty as one state transitioned to the next, and up to several seconds of behavior were needed to classify the current state (Figures S2K–S2M). Nevertheless, the output of shMoSeq was largely consistent across model fits, suggesting that states are readily identifiable (Figure S2G). Together, these findings demonstrate that exploratory behavior is organized at three hierarchically nested timescales (millisecond poses, sub-second syllables, and seconds-long states) that are effectively captured by shMoSeq.

Affordances sculpt the distribution of behavioral states

Animal behavior is shaped by affordances, which offer opportunities for interaction and serve as substrates for goal-oriented behaviors.³⁵ Engagement with different affordances might therefore structure the behavioral states captured by shMoSeq. In an empty arena, mice primarily interact with the walls (Figure 2A). These interactions formed the basis for two behavioral states in which mice explored the circular arena clockwise (keeping the wall on

their right) (“exploratory locomotion,” Figures 2B and 2C). When a male conspecific was added to the arena—expanding the space of affordances—a new social state emerged, which included stationary investigation, investigation amid pursuit, and outright pursuit that sometimes tipped into aggression (Figures 2D–2F). Statistically, the social state was distinguished by proximity to the conspecific and high mutual information between the subject's syllables and those of the conspecific, which was driven in part by their tendency to express similar behaviors at the same time (Figures 2G, S3A, and S3B).

These data suggest that behavioral states reflect nearby affordances (although proximity alone does not determine which state will be expressed [Figures S3C and S3D]). To further test this idea, we performed additional experiments in which five unfamiliar objects were added to the arena (Figure 2H). To capture fine-scale movements, we tracked 3D keypoints using an array of high-speed cameras; this setup allowed us to quantify head and limb positioning during object investigation and to capture nose oscillations, which serve as a proxy for sniffing.^{49,55} Despite these changes in the framerate and recording modality, the five behavioral states identified by shMoSeq had similar durations to those derived from depth cameras (Figures S3E and S3F). The states included two forms of investigation, one directed toward the air (“sniffing”) and the other directed toward the ground (“local investigation”) (Figures S3G and S3H). Mice used these states at different frequencies depending on context. For example, novel objects elicited a 3- to 4-fold increase in the local investigation state, which was further characterized by proximity to objects and deployment of object-oriented syllables, such as “hunched investigation” and “slow approach” (Figures 2I–2L, S3I, and S3J).

Although mice preferred recently added objects in the minutes after they first appeared, instances of the investigation state were distributed throughout each session and therefore were mostly self-initiated (Figure S3K). A separate 3D keypoint dataset comparing social and object investigation revealed subtle differences in posture and syllable expression between object- and conspecific-associated behavior states, reinforcing the notion that objects and conspecifics function as distinct affordances that drive unique patterns of behavioral engagement (Figures S3L–S3O). Taken together, these experiments demonstrate that spontaneous behavior is composed of behavioral states in which mice flexibly interact with specific affordances. Mice remain in each state for up to a minute, and transitions between states capture the mouse disengaging from one affordance and engaging with another.

Behavioral states structure neural activity in the dmPFC

The observation that spontaneous behavior is composed of states—each associated with distinct affordances and patterns of movement—raises the possibility that these states correspond to self-motivated tasks. One strong prediction is that if states indeed correspond to tasks, their identity should be apparent in areas such as the dmPFC, where neural activity is known to reconfigure as animals engage in different experimenter-defined tasks.^{4,6,7,37}

In simultaneous recordings of behavior and neural activity in the dmPFC, up to 20% of neurons were activated or inhibited during each behavioral state, with up to 60% of neurons

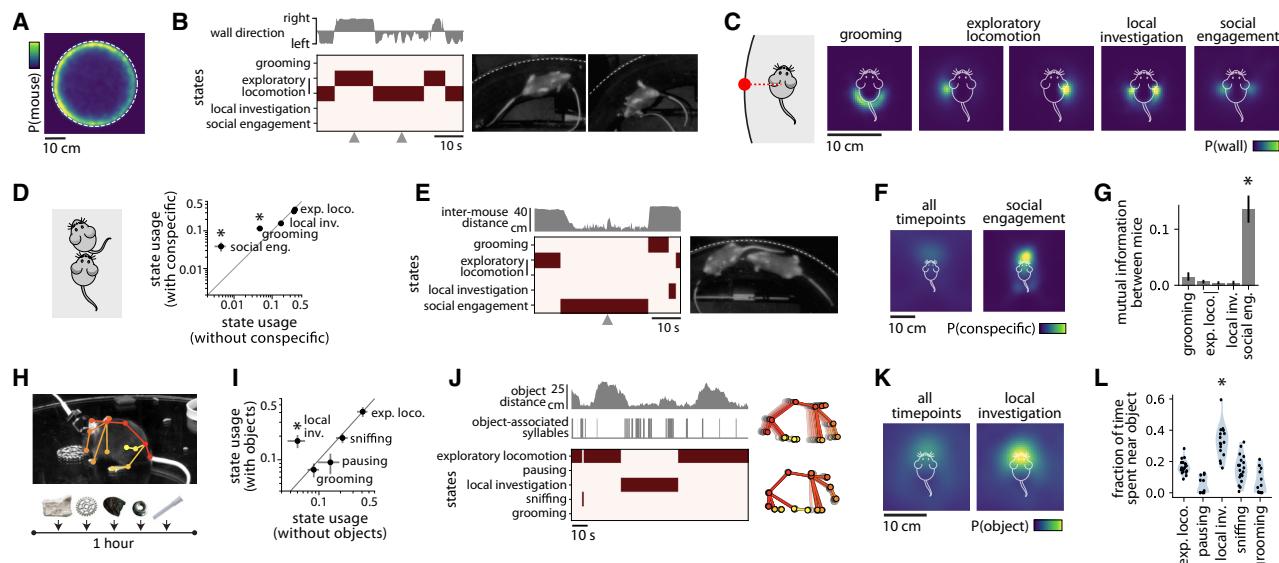


Figure 2. Affordances sculpt the distribution of behavioral states

(A) Distribution of mouse locations across all recordings. Dashed line shows the arena boundary.

(B) Example interval showing relative wall direction (top) and behavioral states (bottom). Video frames correspond to the two time points marked with triangles.

(C) Distribution (for each behavioral state) of the nearest wall location in egocentric coordinates.

(D) Changes in state usage when a C57 conspecific is present, showing mean and standard error across $n = 31$ recordings ($p < 1.5e-4$).

(E) As (B), here showing distance between mice.

(F) As (C), here showing location of the conspecific.

(G) Mutual information between subject and conspecific syllables during each state. Error bars show bootstrap 95% confidence across $n = 31$ recordings ($p < 0.005$). Syllables were kinematically clustered to minimize sparsity-related artifacts.

(H) Design of the novel object experiment. One object was added every 10 min.

(I) Changes in state usage when novel objects are present, plotted as in (D) ($p = 0.001, n = 21$).

(J) Left: example interval showing distance to the nearest object (top), occurrence of object-associated syllables (middle) and behavioral states (bottom). Right: keypoint trajectories for two object-associated syllables.

(K) As (C), here showing location of novel objects.

(L) Fraction of time spent near objects during each state. Dots correspond to recordings ($n = 21, p < 6e-5$).

responding to at least one state (Figures 3A and 3B). Although stringent thresholding suggested that most neurons were tuned to a single state, lifetime sparseness calculations identified many neurons with broad tuning (Figures 3C, S4A, and S4B).⁵⁶ These correlations were reliable enough to decode behavioral states from dmPFC activity (Figures 3D and 3E), with narrowly and broadly tuned neurons yielding similar performance (Figures S4C and S4D). Indeed, state-related neural activity appeared to be the dominant source of structure in the dmPFC, as distinct behavioral states consistently segregated to different regions of a two-dimensional (2D) embedding of neural activity (Figures 3F, 3G, S4E, and S4F); furthermore, the axes of neural activity associated with each behavioral state tended to align with the top handful of neural principal components (Figures S4G–S4I). Consistent with recent reports that PFC activity evolves continuously as mice progress through structured tasks,^{57,58} we observed continuous turnover in neural activity during each state, with a transient increase in turnover rate when one state transitioned to another (Figure S4J). However, unlike in these recent reports, there was no detectable stereotypy in state-associated neural trajectories (Figure S4K).

The accuracy with which states could be decoded from neural activity was higher for shMoSeq models with greater heldout

likelihoods and peaked when state durations respected our original (behavior-only) criterion for parameter selection (Figures S4L and S4M). Thus, models that did a better job of explaining the hierarchical structure of behavior were also better aligned with neural activity. In contrast, there was no clear peak in decoding accuracy when different numbers of behavioral states were tested, suggesting that the neural data are compatible with several different levels of behavioral coarse graining (Figure S4N).

Prefrontal representations privilege states over instantaneous movements

State-related activity in the dmPFC could reflect a succession of naturalistic tasks or be a trivial correlate of ongoing movements that are broadcast throughout the rodent neocortex.^{59–64} However, we find little evidence for this kinematics-centered alternative. In 2D maps of neural activity, syllable representations were scattered more broadly than state representations, suggesting that syllables play a smaller role in structuring dmPFC activity (Figures S4E and S5A). Decoding accuracy was also lower for syllables than for states (Figure S5B). Furthermore, when we trained encoders to predict neural activity from syllables, kinematics, and (optionally) states, the inclusion of states improved

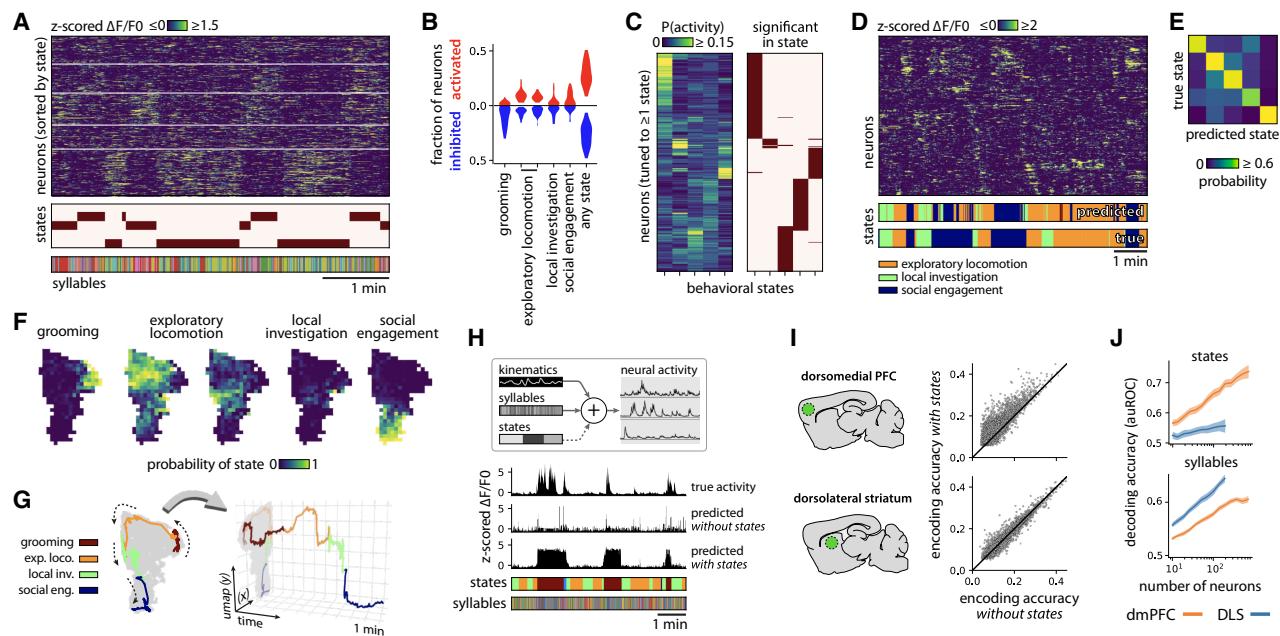


Figure 3. Behavioral states structure neural activity in the dmPFC

- (A) Example of state-related shifts in dmPFC activity.
- (B) Fraction of state-responsive neurons in each recording ($n = 31$, false discovery rate [FDR] < 5%).
- (C) Left: activity of state-responsive neurons across each behavioral state (shown for one recording, “active” defined as Z scored $\Delta F/F_0 \geq 2$). Right: raster indicating whether tuning is statistically significant (FDR < 5%).
- (D) Behavioral state decoding for an example interval.
- (E) Agreement between true and decoded states.
- (F) Usage of each behavioral state across a 2D projection of neural activity from one recording.
- (G) Example trajectory from the neural activity map in (F), colored by the current behavioral state.
- (H) Top: design of encoding model. Bottom: predictions for an example neuron, shown for encoders with and without access to behavioral state information.
- (I) Encoding accuracy with versus without access to behavioral state information, shown for all recorded neurons in the dmPFC (top) and DLS (bottom).
- (J) Decoding accuracy for states (top) and syllables (bottom) using activity from the dmPFC versus the DLS. Neural population sizes were matched via down-sampling. Line and shading show mean and 95% confidence interval (CI) across recordings ($11 \leq n \leq 31$).

accuracy for almost every neuron; this result replicated in an empty arena, with a conspecific, and in recordings with novel objects (Figures 3H, 3I, and S5C–S5E). The inclusion of states also improved accuracy when encoding models had access to world-centric variables such as distance and angle to nearby boundaries and conspecifics, demonstrating that state-related activity does not simply encode spatial relationships to affordances (Figures S5F and S5G).

To test whether privileged encoding of behavioral states is unique to the dmPFC, we performed similar recordings in the dorsolateral striatum (DLS), which encodes syllables and influences their selection during spontaneous behavior.^{48,65,66} Unlike in the dmPFC, the inclusion of shMoSeq states did not impact DLS encoder performance, indicating that states provide no useful information about DLS activity beyond that already contained in syllables and kinematic variables (Figures 3I and S5E). Consistent with this finding, behavior state decoders were more accurate for the dmPFC than the DLS, whereas syllable and kinematic decoders were more accurate for the DLS (Figures 3J, S5H, and S5I). The timescale of neurobehavioral correlates also differed between regions. In the dmPFC, correlations were strongest when kinematic variables were smoothed over several seconds,

indicating that dmPFC activity tracks slow, state-level changes in behavior rather than instantaneous movements. In contrast, sub-second smoothing maximized correlations between DLS activity and behavior (Figure S5J). Fluctuation rates of neural activity (analyzed independently of behavior) also differed between the DLS and dmPFC, with dmPFC neurons remaining autocorrelated ten times longer on average than those in the DLS (Figures S5K and S5L). This separation of neurobehavioral timescales reifies the hierarchical division of behavior into sub-second syllables and multi-second states.

Affordance-related variables are selectively emphasized during specific behavioral states

In conventional experiments with explicitly defined tasks, the dmPFC not only encodes the identity of the current task (e.g., that reward is contingent on stimulus A and not stimulus B) but also information about task-relevant variables (e.g., the current value of stimulus A).^{6,8} If spontaneous behavior is analogously composed of self-initiated tasks that involve interacting with specific affordances, then affordance-related variables should similarly be emphasized during the states in which they are relevant. Consistent with this possibility, we found that diverse

spatial and affordance-related variables were represented in the dmPFC at the single-neuron level, and neurons tuned to these variables were also often modulated by specific states (Figures 4A and S6A). For example, neurons encoding distance and direction to the nearest wall were preferentially active during the exploratory locomotion states, which are reminiscent of a behavioral strategy called thigmotaxis, in which walls are used as a guide for navigation and a source of safety (Figures 2B, 2C, 4A, 4B, S6B, and S6C). This state-specific modulation remained even when controlling for the mouse's distance, angle, and movement relative to the wall (Figures 4C, S6D, and S6E). As a consequence, wall position could be decoded with greatest accuracy during the thigmotactic states—i.e., during the putative task in which wall position is most relevant (Figure 4D).

Rather than specifically supporting thigmotaxis, wall-related activity in the dmPFC could alternatively play a more general role in spatial navigation. Indeed, such a role has been attributed to egocentric boundary vector (EBV) cells in areas such as the retrosplenial cortex (RSC),^{67,68} which similarly represent the distance and relative direction of nearby walls. However, comparing our dmPFC data with RSC recordings from a previous study⁶⁸ revealed important coding differences between the two brain areas. Whereas EBV cells in the RSC tiled a range of preferred distances to the boundary, most wall-tuned neurons in the dmPFC preferred close boundaries and monotonically increased their activity as boundaries grew closer (Figures 4E and S6F). Neurons in the RSC were also more strongly modulated by wall angle than those in the dmPFC (Figure S6F). Furthermore, whereas RSC neurons exhibited a spectrum of angle preferences that tiled the full 0°–360° range, neurons in the dmPFC formed two discrete clusters tuned to walls on the left or right of the animal, respectively (Figures 4G and S6F). Thus, dmPFC representations appear specialized for wall-guided exploration rather than generic spatial coding, consistent with the dmPFC being enriched for neurons encoding specific aspects of the task at hand.

State-dependent coding of specific affordances was also evident during social engagement and object investigation. In recordings with a conspecific, up to 6% of dmPFC neurons were tuned to social proximity, and these responses were much more pronounced during the social engagement state—an effect that persisted even after excluding bouts of aggression and controlling for the relative position and angle of the conspecific (Figures 4H–4J). This enhancement may reflect the slow time-scale of dmPFC activity, which required longer and more frequent social encounters to fully develop (Figures S6G–S6L). Population decoding of social proximity was also more accurate during intervals that encompassed the social engagement state, even though outside these intervals subject mice spanned the full distribution of conspecific distances (Figures 4K and S6M). Similarly, in experiments with novel objects, decoding of object proximity and identity were most accurate during local investigation, even when mice were a full body-length away from the object itself (Figures 4L–4O). The activity of object-responsive neurons was also highest during the local investigation state, even after controlling for distance and angle to the nearest object (Figures S4N and S4O). Thus, representations of walls, objects, and conspecifics are each enhanced during the behavioral states in which they are most relevant.

Changes in task relevance alter representations of affordance-related variables

If behavioral states correspond to self-initiated tasks, then making a task especially salient should enhance the strength of task-related representations in the dmPFC. To test this hypothesis, we subjected mice to attack by a conspecific, which prompted them to use the wall as a defensive affordance. Subject mice were recorded across 5-min blocks with an aggressive conspecific (belonging to the CD1 strain) who attacked frequently, a non-aggressive CD1 conspecific who never attacked, or no partner mouse (Figures 5A and S6P). During attacks, mice entered a shMoSeq-identified state that consisted of flights, submissive postures, and periods of immobility near the wall. This defensive state—which did not occur in our earlier experiments—often continued after attacks were over and even persisted during subsequent interactions with the benign conspecific, albeit at a lower rate (Figures 5B, 5C, and S6Q–S6S). During the defensive state, mice hugged the wall and dmPFC neurons tuned to wall direction or proximity became more active (Figures 5D, 5E, and S6T). Consequently, wall-tuned neurons were more active during aggression blocks (even after controlling for the position of the wall), and the relative wall position could be decoded with greater accuracy during aggression blocks than during blocks alone or with the benign conspecific (Figures 5F–5I, S6U, and S6V).

In the aggression experiments, information related to a particular affordance (the wall) was encoded more prominently in a relevant context (conspecific attack). Conversely, might affordance representations diminish in less-relevant contexts? Indeed, during our original social interaction experiments (with a neutral C57 conspecific) dmPFC responses to social proximity declined over time—as did the accuracy of social distance decoding—suggesting that one goal of social engagement is to gain familiarity and that conspecific representations become de-emphasized as familiarity is gained (Figures S6W–S6Y). Similarly, during novel object exploration, there was a progressive decrease in the activity of object-responsive neurons during later object encounters (Figure S6Z); this effect was especially pronounced when we introduced five identical objects instead of five unique objects.

Thus, environmental affordances—including walls, objects, and conspecifics—drive the expression of specific, long-lasting behavioral states in which affordance-related variables—such as distance to a wall or the identity of an object—are represented more prominently in the dmPFC. Transitions between states occur both spontaneously and in response to perturbations that render certain affordances more relevant (such as walls during aggression). This pattern is consistent with the idea that behavioral states capture affordance-related tasks that arise naturally during self-directed behavior.

Neural dynamics are a lagging indicator of behavior

To understand the causal processes that underlie state and affordance coding in the dmPFC, we asked whether changes in neural activity occurred before or after changes in behavior, reasoning that if the dmPFC were driving behavior—as predicted by some models of hierarchical action planning^{69–71}—then neural activity should precede behavior initiation. We instead observed a systematic delay in neural activity, with state-tuned

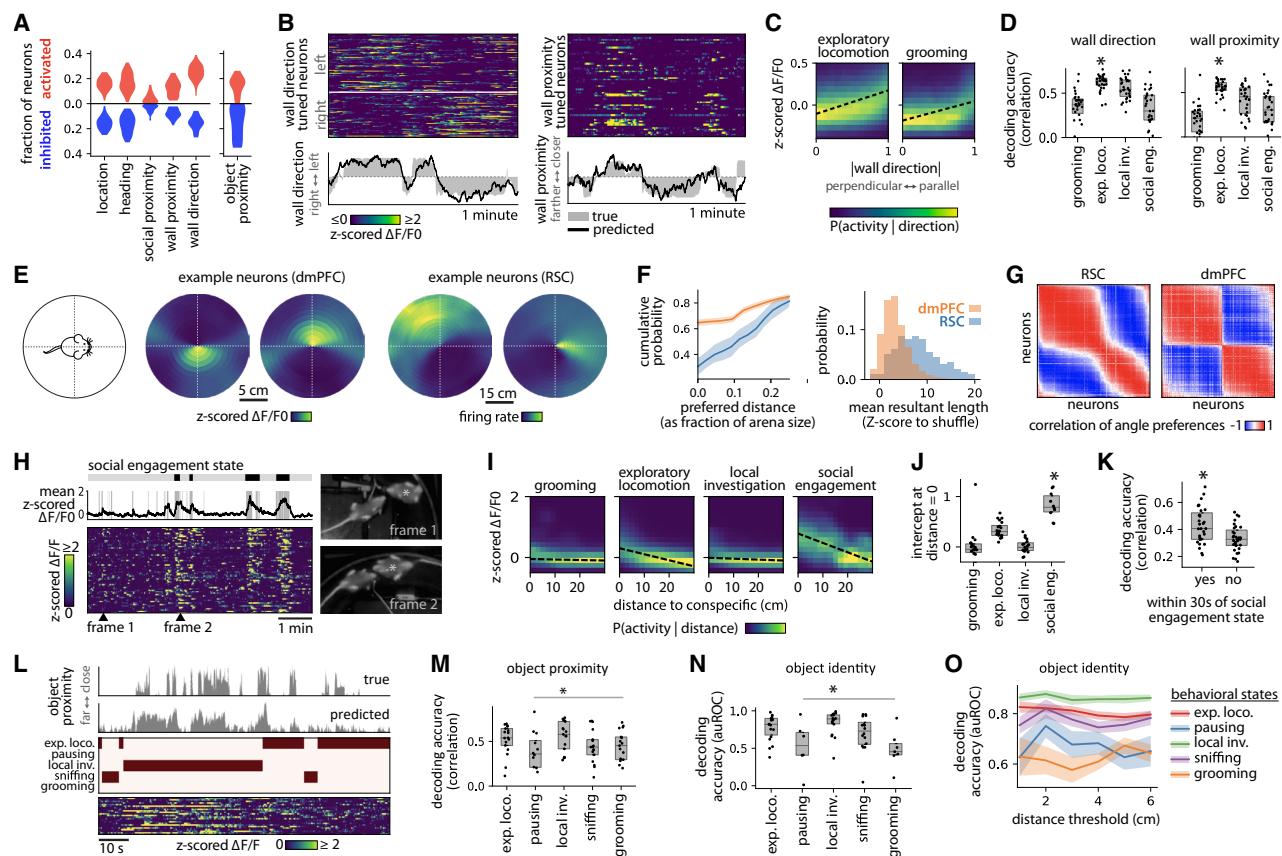


Figure 4. Encoding of affordance-related variables depends on behavioral state

(A) Fraction of neurons modulated by a range of spatial- and affordance-related variables, showing distributions across recordings ($n = 31$, FDR < 5%).

(B) Top: activity of dmPFC neurons modulated by wall direction (left) or proximity (right). Bottom: true and decoded values of each variable over an example interval.

(C) Activity of wall-direction-tuned neurons as a function of wall direction, stratified by behavioral state.

(D) Decoding accuracy for wall direction (left) and proximity (right) during each behavioral state ($n = 31$ recordings, left: $p < 0.03$, right: $p < 6e-4$).

(E) Wall-related activity for example neurons in the dmPFC (left) and retrosplenial cortex (RSC, right). Heatmaps show average activity when boundaries are present at a given distance and angle relative to the mouse.

(F) Left: cumulative distribution of preferred wall distances for neurons in the dmPFC versus the RSC. Arena diameters were, respectively, 125 cm (RSC) and 40 cm (dmPFC). Line and shading show mean and bootstrap 95% CI. Right: strength of wall angle tuning (Z scored mean resultant length) for neurons in the dmPFC and RSC, respectively.

(G) Pairwise correlation of angle preferences for neurons in the dmPFC versus the RSC (neural population sizes were matched via down-sampling). The two blocks for dmPFC correspond to left-wall- and right-wall-prefering neurons, respectively.

(H) State-dependent responses to social proximity. Top-left: instances of the social engagement state. Middle-left: intervals when mice are close together (gray, <5 cm) and mean activity of social-proximity-tuned neurons (black). Bottom-left: activity of individual neurons tuned to social proximity. Right: video frames at the two marked time points.

(I) Activity of social-proximity-tuned neurons as a function of social distance, stratified by behavioral state.

(J) Intercepts of the best-fit lines shown in (I), here calculated for each recording with ≥ 10 social-proximity-tuned neurons ($n = 20$, $p < 3e-7$).

(K) Decoding accuracy for social proximity, comparing time points near to or far from the social engagement state ($n = 31$ recordings, $p = 0.003$).

(L) Example of state-dependent responses to nearby objects, showing true versus decoded proximity (top), behavioral states (middle), and activity of object-proximity-tuned neurons (bottom).

(M) Accuracy of object proximity decoding across behavioral states. Each dot represents one recording in which the state occurred at least once ($11 \leq n \leq 18$). Correlations are highest during the local investigation state; differences are significant for all states except exploratory locomotion ($p < 0.02$).

(N) Accuracy of object identity decoding (quantified as area under the receiver operating characteristic curve [auROC]). Analysis was restricted to time points when the mouse's nose was within 1 cm of exactly one object and the mouse was facing the object ($6 \leq n \leq 18$, $p < 0.05$).

(O) Median accuracy of object identity decoding for a range of nose-object distances. Shading shows the standard deviation across 20 random train/test splits. Accuracy is highest during the local investigation state; this difference is significant with respect to train/test splits ($p < 7e-7$).

neurons becoming active (or inactive) after state onset (or offset) (Figures 6A and S7A). A similar lag was evident when decoding behavioral states from population-wide activity (Figures S7B

and S7C) and was undiminished in look-ahead decoders that were trained to predict future states from current neural activity (Figure S7D).

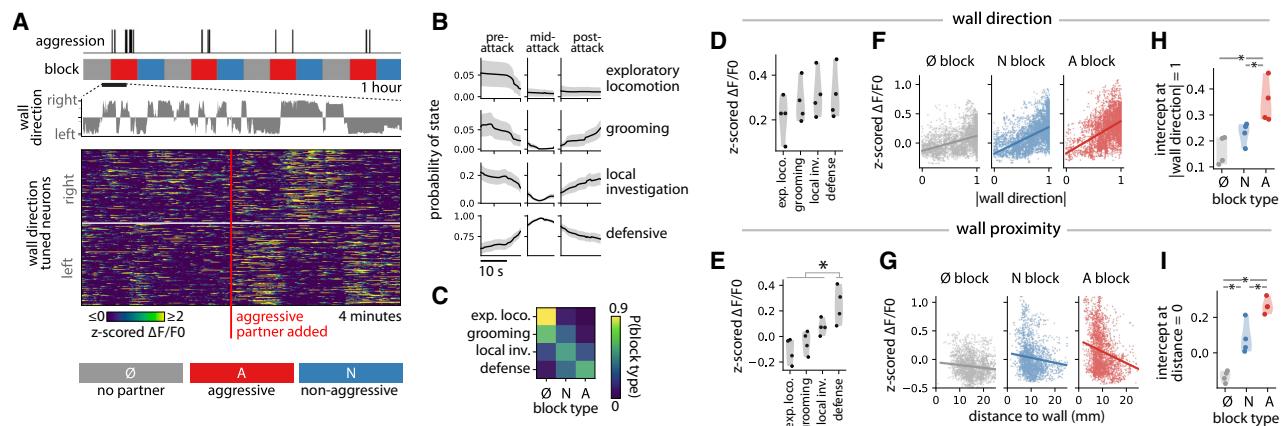


Figure 5. Changes in task relevance alter the encoding of affordance-related variables

(A) Design of the conspecific aggression experiment. Zoom-in shows sudden intensification of wall encoding during the aggression block.
 (B) Behavioral state probabilities aligned to aggression bouts, showing mean and standard error across aggressive episodes ($n = 48$).
 (C) Overlap between behavioral states and experimental blocks.
 (D) Average activity of wall-direction-tuned neurons when mice are close and parallel to the wall ($n = 4$ recordings).
 (E) As (D), here showing activity of wall-proximity-tuned neurons ($p < 0.057$).
 (F) Activity of wall-direction-tuned neurons as a function of wall angle, stratified by block type. Dots represent time points from one example recording.
 (G) Activity of wall-proximity-tuned neurons versus distance to the wall, plotted as in (F).
 (H) Intercept at $x = 1$ (parallel to the wall) of the best-fit lines shown in (F), here shown for all recordings ($n = 4$, $p = 0.028$).
 (I) Intercept at distance = 0 of the best-fit lines shown in (G), here shown for all recordings ($n = 4$, $p = 0.028$).

The delay in social-state encoding is somewhat surprising because dmPFC stimulation can encourage social approach.⁴¹ However, some delay might be expected if state onsets are frequently triggered by actions of the conspecific. To distinguish these possibilities, we reanalyzed published dmPFC recordings from the “tube test” of social dominance, in which pairs of mice enter a tube from opposite ends and compete over who can move forward and who must back out (Figure 6B); prior examination of these data found that the dmPFC represents self-behavior and other behavior but did not report the relative timing of these signals.⁴⁰ Our reanalysis revealed delays common to both subject and opponent behaviors. For example, neural activity associated with self-approach was virtually absent before behavior onset and did not peak until ~ 2 s after onset (Figure 6C); this manifested as a 1-s delay in the cross-correlation between true and decoded behavior (Figure S7E). Cross-correlation peaked at similar offsets for other annotated behaviors, with no statistical difference between self and other actions (Figures 6D and S7E). Thus, dmPFC representations of social interactions appear to lag behavior in multiple contexts.

The apparent lag between dmPFC and ongoing behavior could be exaggerated by our use of calcium indicators to measure neural activity. We therefore recorded from the dmPFC using chronically implanted neuropixel probes as mice explored an open-field arena.⁷² As observed earlier, behavioral states could be decoded from population activity and occupied distinct territories in 2D maps derived from the spiking data (Figures S7F and S7G). To assess the relative timing of dmPFC activity and behavior, we constructed peristimulus time histograms (PSTHs) aligned to state transitions. Most PSTHs showed increases in activity after state onset (Figures 6E and S7H). There were instances where the rise started earlier, but these may have

arisen from uncertainty in the timing of state transitions, as when we focused on transitions with more definite timing (see STAR Methods), we no longer observed pre-transition increases in neural activity (Figures 6F, S7I, and S7J). Neurobehavioral cross-correlations were similarly negative (i.e., neural activity lagging behavior) or not significantly different from zero when measured with respect to state transitions (Figure 6G). Although changes in dmPFC activity led to changes in some non-state behavioral variables (such as height and wall direction), they lagged behind the derivatives of those variables, suggesting that the changes in activity may reflect immediately preceding movements (such as rearing up or turning), rather than triggering those movements per se (Figure 6I). Together, these data suggest that shifts in dmPFC activity generally lag changes in behavioral state, arguing against a model in which dmPFC exerts continual, moment-by-moment control over self-directed behavior.

dmPFC is required for appropriate selection and sequencing of behavioral states

What other roles might the coding of state and affordance play in spontaneous exploratory behavior? The canonical function of the PFC is to exert top-down, context-dependent control over behavior, which, during spontaneous behavior, could involve biasing the selection, duration, or content of behavioral states. To explore these possibilities, we lesioned the dmPFC (including the prelimbic area and the margins of the infralimbic, anterior cingulate, and medial orbital areas) and recorded open-field exploration and novel object interaction (Figures 7A and S7K). Lesioned and control mice were indistinguishable in raw behavioral videos and had nearly identical distributions of wall distances and object distances, respectively (Video S4),

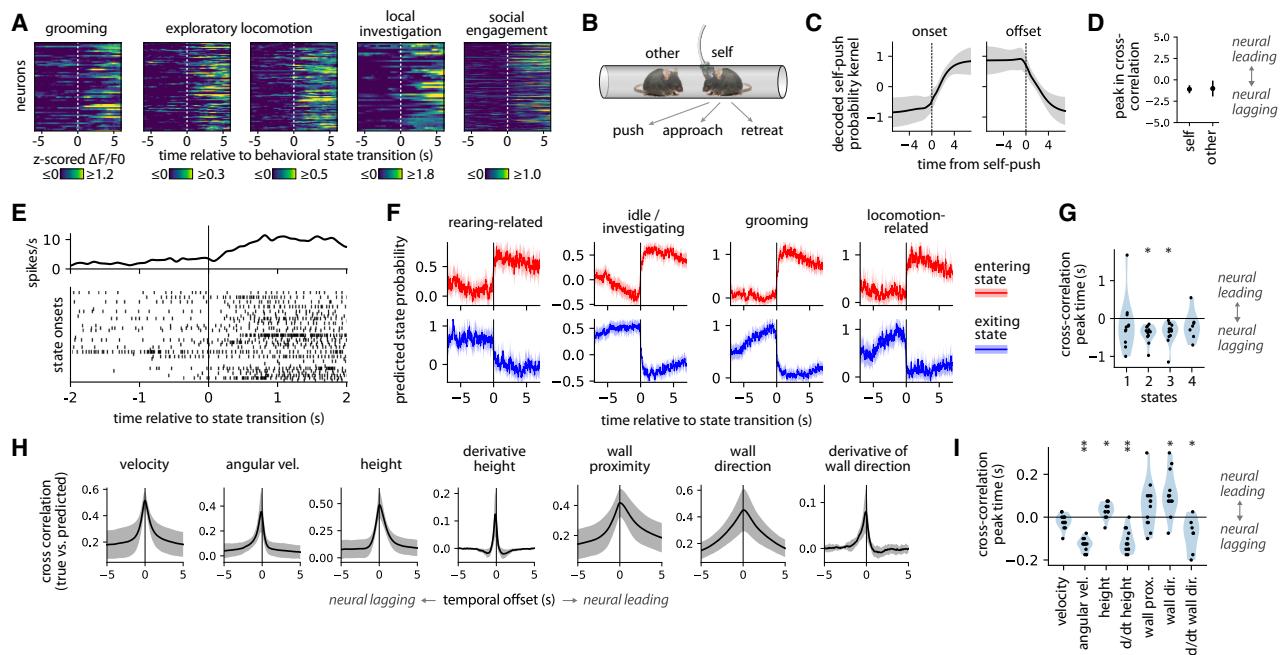


Figure 6. Neural dynamics lag behavioral transitions

(A) Average onset-aligned activity of state-tuned neurons from an example recording.

(B) Illustration of tube test assay.

(C) Decoded “self-push” probability at onsets and offsets of self-push bouts (estimated via kernel regression), showing mean and 95% CI across $n = 10$ recordings.

(D) Timing of peak cross-correlation between true and decoded annotations of self versus other behaviors. Dots and lines represent mean and 95% CI across recordings and behaviors ($15 \leq n \leq 20$, $p = 0.22$).

(E) Raster plot for an example neuron aligned to onset of the grooming state.

(F) State probabilities predicted from spiking activity, aligned to state onsets (top) and offsets (bottom). Transitions were excluded from analysis if uncertainty in their timing was >1 s. Line and shading show mean and 95% CI across recordings ($n = 13$).

(G) Timing of peak cross-correlation between true and decoded state probabilities, restricted to state/recording pairs for which the peak correlation exceeded the 95th percentile of a shuffle distribution ($8 \leq n \leq 13$, $p \leq 6e-4$).

(H) Cross-correlation between true and decoded behavioral variables (mean and 95% CI, $n = 13$).

(I) Timing of peak cross-correlation between true and decoded behavioral variables ($7 \leq n \leq 13$, $*p \leq 0.03$, $**p \leq 7e-6$).

Figure S7L). Lesions also had no effect on the duration or usage of syllables within each behavioral state (Figures S7L and S7M). However, lesions did systematically contract the timescales over which behavior was organized. Mutual information between syllables decayed faster in lesioned mice than in controls, reaching chance levels in less than half the time (Figures 7B and 7C). This difference was especially stark during open-field recordings without novel objects (Figure 7D). Similar effects were evident in a second cohort of animals in which the gap between surgery and recording was shortened to reduce compensatory effects (Figure S7N).

Consistent with the possibility that the dmPFC influences the higher-order structure of behavior, dmPFC lesions altered behavioral state usage, reducing the expression of less commonly used states, such as grooming and local investigation (Figures 7E and S7O–S7S); in contrast, the overall distribution of state durations was unaltered (Figure 7F). These shifts in state usage likely explain the contraction of behavioral timescales. Time spent in the grooming and local investigation states correlated session-by-session with the timescale over which behavior remained predictable (Figure 7G), likely because these specific

states are themselves especially well structured over time and hence contribute disproportionately to the sustained mutual information between syllables that we observe during spontaneous behavior (Figure S7T). These results argue that the dmPFC contributes to the higher-order structure of behavior by differentially influencing the usage of behavioral states.

The observation that dmPFC lesions promote common states, such as exploratory locomotion,⁷³ at the expense of less common states, for example, grooming and local investigation (Figure 7E), could be explained either as a generic increase in the most common types of behavior or as a more specific loss of context-specific behavioral control. Consistent with the latter idea, loss of the investigation state was most pronounced in recordings with novel objects (Figure 7H). Furthermore, whereas control mice groomed more in the vicinity of novel objects, this effect almost disappeared in lesioned mice (Figure 7I). Together, these lesion results suggest that the dmPFC imparts hierarchical structure to behavior by facilitating the expression of a broader set of behavioral states—especially rare states that depend on context—without substantively altering their granular (syllabic) contents; this stands in contrast to DLS lesions, which perturb

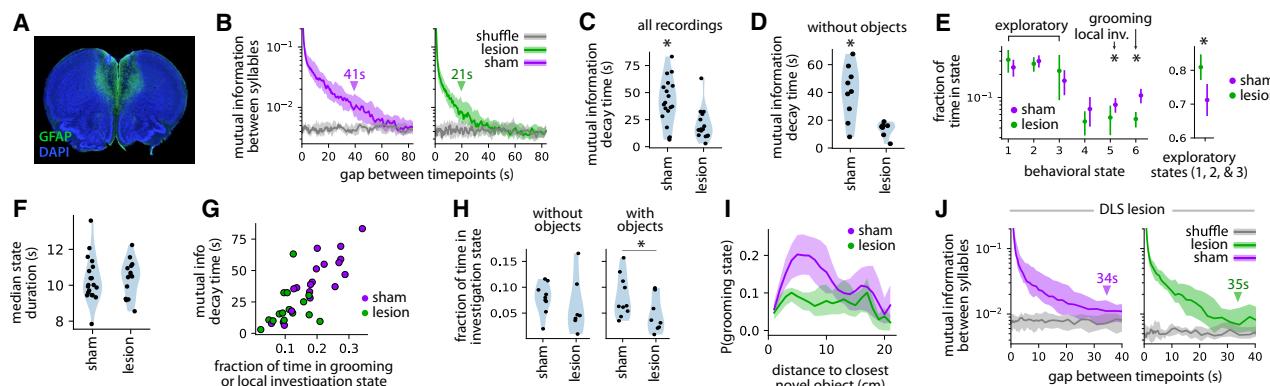


Figure 7. dmPFC influences the usage of behavioral states

(A) Coronal section from a lesioned mouse, stained with an astrocytic antibody (glial fibrillary acidic protein [GFAP]) that highlights the lesioned area.

(B) Mutual information (MI) between syllables at a range of temporal lags (median and interquartile interval across recordings, $n = 19$ sham, $n = 15$ lesion). Pointers show the mean decay time, defined as time to reach the 99th percentile of the shuffle distribution.

(C) MI decay times for each recording, defined as in (B) ($p = 0.006$).

(D) As (C), here restricted to recordings without novel objects ($n = 9$ sham, $n = 7$ lesion, $p = 0.01$).

(E) Left: behavioral state usage (mean and 95% CI) for lesion versus sham recordings (FDR < 5%). Right: combined usage of the exploratory locomotion states.

(F) Median duration of behavioral states for sham and lesion recordings ($p = 0.66$).

(G) MI decay time versus usage of the grooming and local investigation states ($n = 34$ recordings, $r = 0.79$, $p = 2.5e-8$).

(H) Usage of the investigation state in lesion versus sham recordings, either with (right) or without (left) novel objects ($p = 0.034$).

(I) Probability of the grooming state as a function of distance to the closest novel object, showing mean and 95% CI across recordings with novel objects ($n = 10$ sham, $n = 8$ lesion).

(J) As (B), here showing DLS lesion data from Markowitz et al.⁴⁸ There is no significant difference in MI decay times ($p = 0.3$, $n = 14$ sham, $n = 34$ lesion).

both syllable usage and sequencing⁴⁸ but have no effect on the overall timescale at which behavior is organized (Figure 7J).

DISCUSSION

Animals act toward many ends—to gain information; escape danger; find food, water, and shelter; and much more.^{29,74,75} Sometimes one need predominates, as when researchers deprive animals of food and then offer food rewards during a task. However, in less-structured contexts, the purpose of behavior is often mysterious; animals may have competing needs that they pursue at different times, each corresponding to a kind of self-motivated task.⁷⁶ In principle, these emergent tasks could have profound effects on behavior and neural dynamics,^{36,77} yet the hypothesis that self-paced exploratory behavior entails a set of tasks has not been explicitly explored in the lab. Here, we define criteria that such task-related states should fulfill: they should organize low-level movements, correspond to recognizable behaviors, reflect available affordances, and be visible in neural activity, especially in areas (such as the dmPFC) that are important for goal-directed behavior. Given these criteria, we find that spontaneous behavior is structured as a succession of self-directed tasks. These observations required the development and validation of shMoSeq, a hierarchical model that bridges the gap between microscopic behaviors, identified at the level of poses and syllables, and the macroscopic behaviors that enable animals to solve problems and reach goals.

Comparisons between real and simulated data show that the hierarchy instantiated by shMoSeq—millisecond poses, sub-second syllables, and multi-second states—can fully explain the spectrum of timescales apparent in a 1-h recording of exploratory behavior. Of course, this does not preclude the existence of structure at other timescales: circadian rhythms and internal states such as hunger and thirst (which are distinct from the seconds-long “behavioral states” we describe here) would likely become evident in longer recordings,⁷⁸⁻⁸⁰ and substates that are intermediate between states and syllables—though not strictly necessary to explain the statistics of our data—might be practically useful for segmenting behavior (e.g., by splitting the social state into separate components corresponding to investigation, pursuit, and aggression).⁸¹ That said, it is notable that state decoding accuracy peaks at the seconds-to-minutes timescale nominated by shMoSeq, consistent with the role of states in scaffolding neural activity.

The correspondence between behavioral states and self-initiated tasks is supported by their representation in the dmPFC, which, in several ways, resembles that observed during experimenter-defined tasks. When experimenter-defined tasks involve shifting between distinct contexts, the representation of each context is typically stable over time and abstracted from moment-to-moment action selection.^{4,7,9,13,57,82} Similarly, during spontaneous behavior, dmPFC encodes long-lasting behavioral states while abstracting over low-level movements. When the rules of experimenter-defined tasks change over time, the PFC preferentially encodes stimuli that are salient given the current rules.^{6,8,82-85} We found that representations of environmental affordances (walls, conspecifics, or novel objects) became more prominent during the behavioral states in which they were relevant. These changes suggest that behavioral state transitions correspond to shifts in attention⁸⁶ and hint at the animal’s underlying goals. For example, the fact that object identity can be decoded better during the local investigation state

10 Neuron 114, 1–16, March 4, 2026

suggests that this state is associated with the goal of seeking information about objects.

What causal role does the dmPFC play in structuring behavior? Animals rely on the PFC to adjust their behavior in a context-specific manner, especially when there are conflicts between lower-level systems of behavioral control.⁴ Indeed at focal moments of conflict, the dmPFC can participate directly in action initiation; examples include time-locked activity prior to active avoidance during a shuttle-crossing task⁸⁷ and instantaneous induction of winning by optogenetic activation during social competition.⁸⁸ It is unclear whether this instructive role for the dmPFC persists across time or during self-directed exploratory behavior. One influential hypothesis places PFC at the apex of a network charged with action planning and execution.^{69,70,89} In the strongest of these proposals, high-level behavioral states correspond to metastable attractors in the PFC, and behavioral transitions arise from stochastic switching between attractors.⁷¹ However, our findings suggest a more permissive role during spontaneous behavior; rather than actively initiating new states, the dmPFC may construct abstract representations of ongoing behavior (e.g., by passively monitoring incoming sensory and motor information) that help specify the distribution of states in a given context. This mode of regulation could reflect the influence of the dmPFC on brain-wide neural dynamics^{90,91} and is compatible with a wide range of brain areas instigating behavioral state transitions.

The attribution of goals to living organisms has a fraught history.⁹² “Goal” can refer broadly to any process of control resulting from natural selection (i.e., “teleonomy”⁹³) or, more narrowly, to directed programs of behavior that are under cognitive control.⁹⁴ Whereas goal-directedness in traditional neuroscience is typically tested through the manipulation of task rewards (as during devaluation⁹⁵), other approaches are required for naturalistic behavior, which often lacks task structure and explicit rewards.⁶⁵ Here, we addressed this problem by articulating a set of observational criteria for goal-directedness involving both behavior and neural activity. We note that much of spontaneous behavior remains intact after dmPFC lesions, suggesting either redundancy in the systems that support the expression of task-related states during unstructured exploration or a more prominent role for sensorimotor habits. Longer-timescale behavioral recordings among a wider variety of affordances may provide a clearer view.

To identify behavioral states, we designed a data-efficient HHMM (shMoSeq). Unlike approaches that cluster syllables into mutually exclusive groups,^{34,54} shMoSeq does not enforce a strict mapping between syllables and states. Rather, states are defined by probabilistic changes in syllable usage, meaning that syllables can be reused across states in a context-specific manner. This framework effectively captures the long-timescale (i.e., non-Markovian) structure of syllable sequences and was ultimately ratified by its close correspondence with neural dynamics in the dmPFC. Several other methods for inferring behavioral hierarchy have also been reported; these employ methods such as compression,⁹⁶ repeated sequence search,⁹⁷ or context-free-grammars⁹⁸ and, as such, are more geared toward short, repeated action sequences than the prolonged behavioral states highlighted here. Future elaborations of

shMoSeq that take advantage of self-supervised behavioral embeddings might reveal states with greater richness or specificity.^{99,100} ShMoSeq analysis of uninstructed movements during experimenter-defined tasks could also be useful for uncovering changes in task engagement or strategy.¹⁰¹

We have released shMoSeq as a python package with documentation and tutorials (<https://state-moseq.readthedocs.io/en/latest/>). Information about runtime, compute requirements, parameter selection, and the types of data for which shMoSeq may be useful are available in the **STAR Methods**. It is important to note that shMoSeq affords some choice in the total number of identified states. Although we heuristically chose this number based on the stability of model outputs, more than one level of coarse graining may be reasonable for a given dataset, and the relevant level will doubtless vary based on the total amount of data and complexity of behavior therein.

Whatever form they take, hierarchical behavior models will be an important future tool for exploring cognition during unstructured behavior. For decades, researchers have constrained animals’ movements and narrowed their motivations in order to isolate specific cognitive processes and unspool their neural causes. The result has been to minimize the agency and spontaneity that characterize most behavior outside the lab.^{1,31,102} Ultimately, we would like the best of both worlds—to probe well-defined computations using normative models even as animals set their own goals and pursue them freely.^{103,104} Here, we have taken a small step in that direction by identifying the behavioral states that emerge during unstructured exploration in mice. By bracketing intervals of relative stability in the animal’s behavior, these states expose the possible succession of tasks that structure cognition over time—a hypothesis that is borne out in our analysis of dmPFC activity. This connection between states and tasks—if it proves durable—would allow researchers to employ during naturalistic behavior the same normative frameworks that they typically seek through experimental reductionism.

RESOURCE AVAILABILITY

Lead contact

Requests for further information and resources should be directed to, and will be fulfilled by, the lead contact, Sandeep Robert Datta (srdatta@hms.harvard.edu).

Materials availability

This study did not generate new materials.

Data and code availability

- Original data have been deposited at Zenodo at <https://doi.org/10.5281/zenodo.17488068> and are publicly available as of the date of publication.
- All original code has been deposited at Zenodo at <https://doi.org/10.5281/zenodo.17488068> and is publicly available as of the date of publication.
- Any additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

ACKNOWLEDGMENTS

S.R.D. is supported by NIH grants U24NS109520 and R01AG073625, the Simons Collaboration on the Global Brain, and the Simons Collaboration on

Plasticity and the Aging Brain. C.W. is supported by the Jane Coffin Childs Memorial Fund for Medical Research and the King Trust Postdoctoral Research Fellowship Program, Bank of America Private Bank, Trustee. S.R.D. and S.W.L. are supported by NIH grant U19NS113201. J.E.P. is supported by NIH grant F31DC022154. W.F.G. is supported by NIH grant F31NS113385. T.S. is supported by the Schmidt Science Fellowship and the Leonard and Isabelle Goldenson Fellowship. We thank J. Araki and J. Martz for administrative support; the HMS Research Instrumentation Core, which is supported by the Bertarelli Program in Translational Neuroscience and Neuroengineering, and by NEI grant EY012196; the HMS Neurobiology Imaging Facility for microscopy support; and members of the Datta laboratory for useful comments on the paper. Portions of this research were conducted on the O2 High Performance Compute Cluster at Harvard Medical School. Mouse illustrations were downloaded from <https://www.scidraw.io/>. We also thank M. Hasselmo and A. Alexander for sharing raw data related to egocentric boundary vector tuning in retrosplenial cortex.

AUTHOR CONTRIBUTIONS

C.W. and S.R.D. conceived the project and designed the experiments. C.W. and S.W.L. designed the mathematical model. C.W. and J.E.P. wrote the pre-processing pipeline for depth videos. C.W., W.F.G., and T.S. collected data, with assistance from L.T.K., A.N.-B., A.P., and S.M. C.W. performed analysis, with assistance from M.A.M.O. C.W. and S.R.D. wrote the manuscript, with input from all authors. S.R.D. supervised the project.

DECLARATION OF INTERESTS

S.R.D. and C.W. have patent registrations and applications related to the use of MoSeq and shMoSeq for analysis of animal behavior.

STAR METHODS

Detailed methods are provided in the online version of this paper and include the following:

- **KEY RESOURCES TABLE**
- **EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS**
- **METHOD DETAILS**
 - Behavior Assays
 - Stereotactic surgery procedures
 - Recording setups
 - Preprocessing of neural data
 - MoSeq analysis from 3D keypoints
 - Syllable coarse-graining
 - Analysis of behavioral timescales
 - ShMoSeq: generative model and fitting algorithm
 - ShMoSeq: practical application
 - Breadth of neural tuning
 - Decoding models
 - Analysis of single-neuron behavior associations
 - Neural encoding models
 - Neural activity manifolds
 - Mapping behavioral states to neural principal components
 - Analysis of stereotypy in neural trajectories
 - Analysis of relative wall position during open field exploration
 - Analysis of aggression experiments
 - Analysis of novel object experiments
 - Analysis of social engagement state
 - Tuning to allocentric spatial variables
 - Analysis of tube test dataset
 - Analysis of DLS lesion experiments
 - Comparative analysis of egocentric boundary vector tuning
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
- **ADDITIONAL RESOURCES**

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.neuron.2025.11.021>.

Received: January 22, 2025

Revised: October 6, 2025

Accepted: November 20, 2025

REFERENCES

1. Miller, C.T., Gire, D., Hoke, K., Huk, A.C., Kelley, D., Leopold, D.A., Smear, M.C., Theunissen, F., Yartsev, M., and Niell, C.M. (2022). Natural behavior is the language of the brain. *Curr. Biol.* 32, R482–R493. <https://doi.org/10.1016/j.cub.2022.03.031>.
2. Rangel, A., and Hare, T. (2010). Neural computations associated with goal-directed choice. *Curr. Opin. Neurobiol.* 20, 262–270. <https://doi.org/10.1016/j.conb.2010.03.001>.
3. Dickinson, A., and Balleine, B. (1994). Motivational control of goal-directed action. *Anim. Learn. Behav.* 22, 1–18. <https://doi.org/10.3758/BF03199951>.
4. Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 24, 167–202. <https://doi.org/10.1146/annurev.neuro.24.1.167>.
5. Aoi, M.C., Mante, V., and Pillow, J.W. (2020). Prefrontal cortex exhibits multidimensional dynamic encoding during decision-making. *Nat. Neurosci.* 23, 1410–1420. <https://doi.org/10.1038/s41593-020-0696-5>.
6. Stokes, M.G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., and Duncan, J. (2013). Dynamic coding for cognitive control in prefrontal cortex. *Neuron* 78, 364–375. <https://doi.org/10.1016/j.neuron.2013.01.039>.
7. Asaad, W.F., Rainer, G., and Miller, E.K. (2000). Task-specific neural activity in the primate prefrontal cortex. *J. Neurophysiol.* 84, 451–459. <https://doi.org/10.1152/jn.2000.84.1.451>.
8. Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84. <https://doi.org/10.1038/nature12742>.
9. Rich, E.L., and Shapiro, M. (2009). Rat prefrontal cortical neurons selectively code strategy switches. *J. Neurosci.* 29, 7208–7219. <https://doi.org/10.1523/JNEUROSCI.6068-08.2009>.
10. Le Merre, P., Ährlund-Richter, S., and Carlén, M. (2021). The mouse prefrontal cortex: Unity in diversity. *Neuron* 109, 1925–1944. <https://doi.org/10.1016/j.neuron.2021.03.035>.
11. Ragozzino, M.E., Kim, J., Hassert, D., Minniti, N., and Kiang, C. (2003). The contribution of the rat prelimbic-infralimbic areas to different forms of task switching. *Behav. Neurosci.* 117, 1054–1065. <https://doi.org/10.1037/0735-7044.117.5.1054>.
12. Balleine, B.W., and O'Doherty, J.P. (2010). Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology* 35, 48–69. <https://doi.org/10.1038/npp.2009.131>.
13. Woodward, D.J., Chang, J.Y., Janak, P., Azarov, A., and Anstrom, K. (1999). Mesolimbic neuronal activity across behavioral states. *Ann. N. Y. Acad. Sci.* 877, 91–112. <https://doi.org/10.1111/j.1749-6632.1999.tb09263.x>.
14. Basu, R., Gebauer, R., Herfurth, T., Kolb, S., Golipour, Z., Tchumatchenko, T., and Ito, H.T. (2021). The orbitofrontal cortex maps future navigational goals. *Nature* 599, 449–452. <https://doi.org/10.1038/s41586-021-04042-9>.
15. Klein-Flügge, M.C., Bongianni, A., and Rushworth, M.F.S. (2022). Medial and orbital frontal cortex in decision-making and flexible behavior. *Neuron* 110, 2743–2770. <https://doi.org/10.1016/j.neuron.2022.05.022>.

16. Balleine, B.W. (2019). The Meaning of Behavior: Discriminating Reflex and Volition in the Brain. *Neuron* 104, 47–62. <https://doi.org/10.1016/j.neuron.2019.09.024>.
17. Flash, T., and Hochner, B. (2005). Motor primitives in vertebrates and invertebrates. *Curr. Opin. Neurobiol.* 15, 660–666. <https://doi.org/10.1016/j.conb.2005.10.011>.
18. Anderson, D.J., and Perona, P. (2014). Toward a Science of Computational Ethology. *Neuron* 84, 18–31. <https://doi.org/10.1016/j.neuron.2014.09.005>.
19. Botta, P., Fushiki, A., Vicente, A.M., Hammond, L.A., Mosberger, A.C., Gerfen, C.R., Peterka, D., and Costa, R.M. (2020). An Amygdala Circuit Mediates Experience-Dependent Momentary Arrests during Exploration. *Cell* 183, 605–619.e22. <https://doi.org/10.1016/j.cell.2020.09.023>.
20. Eilam, D., and Golani, I. (1990). Home base behavior in amphetamine-treated tame wild rats (*Rattus norvegicus*). *Behav. Brain Res.* 36, 161–170. [https://doi.org/10.1016/0166-4328\(90\)90170-J](https://doi.org/10.1016/0166-4328(90)90170-J).
21. Gordon, G., Fonio, E., and Ahissar, E. (2014). Emergent exploration via novelty management. *J. Neurosci.* 34, 12646–12661. <https://doi.org/10.1523/JNEUROSCI.1872-14.2014>.
22. Claudi, F., Campagner, D., and Branco, T. (2022). Innate heuristics and fast learning support escape route selection in mice. *Curr. Biol.* 32, 2980–2987.e5. <https://doi.org/10.1016/j.cub.2022.05.020>.
23. Baran, S.E., Armstrong, C.E., Niren, D.C., and Conrad, C.D. (2010). Prefrontal cortex lesions and sex differences in fear extinction and perseveration. *Learn. Mem.* 17, 267–278. <https://doi.org/10.1101/lm.1778010>.
24. Alsina-Llanes, M., and Olazábal, D.E. (2021). NMDA lesions in the prefrontal cortex delay the onset of maternal, but not infanticidal behavior in pup-naïve adult mice (C57BL/6). *Behav. Neurosci.* 135, 402–414. <https://doi.org/10.1037/bne0000427>.
25. Yashima, J., Uekita, T., and Sakamoto, T. (2023). The prelimbic cortex but not the anterior cingulate cortex plays an important role in social recognition and social investigation in mice. *PLoS One* 18, e0284666. <https://doi.org/10.1371/journal.pone.0284666>.
26. Schneider, M., and Koch, M. (2005). Behavioral and morphological alterations following neonatal excitotoxic lesions of the medial prefrontal cortex in rats. *Exp. Neurol.* 195, 185–198. <https://doi.org/10.1016/j.exneurol.2005.04.014>.
27. Ortega, L.A., Glueck, A.C., Uhelski, M., Fuchs, P.N., and Papini, M.R. (2013). Role of the ventrolateral orbital cortex and medial prefrontal cortex in incentive downshift situations. *Behav. Brain Res.* 244, 120–129. <https://doi.org/10.1016/j.bbr.2013.01.029>.
28. Baerends, G.P. (1976). The functional organization of behaviour. *Anim. Behav.* 24, 726–738. [https://doi.org/10.1016/S0003-3472\(76\)80002-4](https://doi.org/10.1016/S0003-3472(76)80002-4).
29. Tinbergen, N. (1951). *The Study of Instinct* (Clarendon Press).
30. Dawkins, R. (1976). Hierarchical organisation: A candidate principle for ethology. In *Growing Points in Ethology* (Cambridge University Press).
31. Datta, S.R., Anderson, D.J., Branson, K., Perona, P., and Leifer, A. (2019). Computational Neuroethology: A Call to Action. *Neuron* 104, 11–24. <https://doi.org/10.1016/j.neuron.2019.09.038>.
32. Wiltschko, A.B., Johnson, M.J., Iurilli, G., Peterson, R.E., Katon, J.M., Pashkovski, S.L., Abraira, V.E., Adams, R.P., and Datta, S.R. (2015). Mapping Sub-Second Structure in Mouse Behavior. *Neuron* 88, 1121–1135. <https://doi.org/10.1016/j.neuron.2015.11.031>.
33. Berman, G.J., Choi, D.M., Bialek, W., and Shaevitz, J.W. (2014). Mapping the stereotyped behaviour of freely moving fruit flies. *J. R. Soc. Interface* 11, 20140672. <https://doi.org/10.1098/rsif.2014.0672>.
34. Berman, G.J., Bialek, W., and Shaevitz, J.W. (2016). Predictability and hierarchy in *Drosophila* behavior. *Proc. Natl. Acad. Sci. USA* 113, 11943–11948. <https://doi.org/10.1073/pnas.1607601113>.
35. Gibson, J.J. (1986). *The Ecological Approach to Visual Perception* (Lawrence Erlbaum Associates).
36. Allen, W.E., Kauvar, I.V., Chen, M.Z., Richman, E.B., Yang, S.J., Chan, K., Gradinari, V., Deverman, B.E., Luo, L., and Deisseroth, K. (2017). Global Representations of Goal-Directed Behavior in Distinct Cell Types of Mouse Neocortex. *Neuron* 94, 891–907.e6. <https://doi.org/10.1016/j.neuron.2017.04.017>.
37. Bagi, B., Brecht, M., and Sanguinetti-Scheck, J.I. (2022). Unsupervised discovery of behaviorally relevant brain states in rats playing hide-and-seek. *Curr. Biol.* 32, 2640–2653.e4. <https://doi.org/10.1016/j.cub.2022.04.068>.
38. Lee, E., Rhim, I., Lee, J.W., Ghim, J.W., Lee, S., Kim, E., and Jung, M.W. (2016). Enhanced Neuronal Activity in the Medial Prefrontal Cortex during Social Approach Behavior. *J. Neurosci.* 36, 6926–6936. <https://doi.org/10.1523/JNEUROSCI.0307-16.2016>.
39. Murugan, M., Jang, H.J., Park, M., Miller, E.M., Cox, J., Taliaferro, J.P., Parker, N.F., Bhave, V., Hur, H., Liang, Y., et al. (2017). Combined Social and Spatial Coding in a Descending Projection from the Prefrontal Cortex. *Cell* 171, 1663–1677.e16. <https://doi.org/10.1016/j.cell.2017.11.002>.
40. Kingsbury, L., Huang, S., Wang, J., Gu, K., Golshani, P., Wu, Y.E., and Hong, W. (2019). Correlated Neural Activity and Encoding of Behavior across Brains of Socially Interacting Animals. *Cell* 178, 429–446.e16. <https://doi.org/10.1016/j.cell.2019.05.022>.
41. Kingsbury, L., Huang, S., Raam, T., Ye, L.S., Wei, D., Hu, R.K., Ye, L., and Hong, W. (2020). Cortical Representations of Conspecific Sex Shape Social Behavior. *Neuron* 107, 941–953.e7. <https://doi.org/10.1016/j.neuron.2020.06.020>.
42. Franch, M., Yellapantula, S., Parajuli, A., Kharas, N., Wright, A., Aazhang, B., and Dragoi, V. (2024). Visuo-frontal interactions during social learning in freely moving macaques. *Nature* 627, 174–181. <https://doi.org/10.1038/s41586-024-07084-x>.
43. Zhang, W., and Yartsev, M.M. (2019). Correlated Neural Activity across the Brains of Socially Interacting Bats. *Cell* 178, 413–428.e22. <https://doi.org/10.1016/j.cell.2019.05.023>.
44. Rose, M.C., Styr, B., Schmid, T.A., Elie, J.E., and Yartsev, M.M. (2021). Cortical representation of group social communication in bats. *Science* 374, eaba9584. <https://doi.org/10.1126/science.aba9584>.
45. Lindsay, A.J., Caracheo, B.F., Grewal, J.J.S., Leibovitz, D., and Seamans, J.K. (2018). How Much Does Movement and Location Encoding Impact Prefrontal Cortex Activity? An Algorithmic Decoding Approach in Freely Moving Rats. *Eur. Neurol.* 5, ENEURO.0023-18.2018. <https://doi.org/10.1523/ENEURO.0023-18.2018>.
46. Maisson, D.J.N., Cervera, R.L., Voloh, B., Conover, I., Zambre, M., Zimmermann, J., and Hayden, B.Y. (2023). Widespread coding of navigational variables in prefrontal cortex. *Curr. Biol.* 33, 3478–3488.e3. <https://doi.org/10.1016/j.cub.2023.07.024>.
47. Ghosh, K.K., Burns, L.D., Cocker, E.D., Nimmerjahn, A., Ziv, Y., Gamal, A.E., and Schnitzer, M.J. (2011). Miniaturized integration of a fluorescence microscope. *Nat. Methods* 8, 871–878. <https://doi.org/10.1038/nmeth.1694>.
48. Markowitz, J.E., Gillis, W.F., Beron, C.C., Neufeld, S.Q., Robertson, K., Bhagat, N.D., Peterson, R.E., Peterson, E., Hyun, M., Linderman, S.W., et al. (2018). The Striatum Organizes 3D Behavior via Moment-to-Moment Action Selection. *Cell* 174, 44–58.e17. <https://doi.org/10.1016/j.cell.2018.04.019>.
49. Weinreb, C., Pearl, J.E., Lin, S., Osman, M.A.M., Zhang, L., Annapragada, S., Conlin, E., Hoffmann, R., Makowska, S., Gillis, W.F., et al. (2024). Keypoint-MoSeq: parsing behavior by linking point tracking to pose dynamics. *Nat. Methods* 21, 1329–1339. <https://doi.org/10.1038/s41592-024-02318-2>.
50. Akiti, K., Tsutsui-Kimura, I., Xie, Y., Mathis, A., Markowitz, J.E., Anyoha, R., Datta, S.R., Mathis, M.W., Uchida, N., and Watabe-Uchida, M. (2022). Striatal dopamine explains novelty-induced behavioral dynamics and individual variability in threat prediction. *Neuron* 110, 3789–3804.e9. <https://doi.org/10.1016/j.neuron.2022.08.022>.

51. Lidster, K., Owen, K., Browne, W.J., and Prescott, M.J. (2019). Cage aggression in group-housed laboratory male mice: an international data crowdsourcing project. *Sci. Rep.* 9, 15211. <https://doi.org/10.1038/s41598-019-51674-z>.
52. Lin, S., Gillis, W.F., Weinreb, C., Zeine, A., Jones, S.C., Robinson, E.M., Markowitz, J., and Datta, S.R. (2024). Characterizing the structure of mouse behavior using Motion Sequencing. *Nat. Protoc.* 19, 3242–3291. <https://doi.org/10.1038/s41596-024-01015-w>.
53. Sainburg, T., Theilman, B., Thielk, M., and Gentner, T.Q. (2019). Parallels in the sequential organization of birdsong and human speech. *Nat. Commun.* 10, 3636. <https://doi.org/10.1038/s41467-019-11605-y>.
54. Voloh, B., Maisson, D.J.N., Cervera, R.L., Conover, I., Zambre, M., Hayden, B., and Zimmermann, J. (2023). Hierarchical action encoding in prefrontal cortex of freely moving macaques. *Cell Rep.* 42, 113091. <https://doi.org/10.1016/j.celrep.2023.113091>.
55. Findley, T.M., Wyrick, D.G., Cramer, J.L., Brown, M.A., Holcomb, B., Attey, R., Yeh, D., Monasevitch, E., Nouboussi, N., Cullen, I., et al. (2021). Sniff-synchronized, gradient-guided olfactory search by freely moving mice. *eLife* 10, e58523. <https://doi.org/10.7554/eLife.58523>.
56. Willmore, B., and Tolhurst, D.J. (2001). Characterizing the sparseness of neural codes. *Network* 12, 255–270. <https://doi.org/10.1080/net.12.3.255.270>.
57. El-Gaby, M., Harris, A.L., Whittington, J.C.R., Dorrell, W., Bhomick, A., Walton, M.E., Akam, T., and Behrens, T.E.J. (2024). A cellular basis for mapping behavioural structure. *Nature* 636, 671–680. <https://doi.org/10.1038/s41586-024-08145-x>.
58. Zhang, Y., Denman, A.J., Liang, B., Werner, C.T., Beacher, N.J., Chen, R., Li, Y., Shaham, Y., Barbera, G., and Lin, D.T. (2022). Detailed mapping of behavior reveals the formation of prelimbic neural ensembles across operant learning. *Neuron* 110, 674–685.e6. <https://doi.org/10.1016/j.neuron.2021.11.022>.
59. Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C.B., Carandini, M., and Harris, K.D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science* 364, 255. <https://doi.org/10.1126/science.aav7893>.
60. Musall, S., Kaufman, M.T., Juavinett, A.L., Gluf, S., and Churchland, A.K. (2019). Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* 22, 1677–1686. <https://doi.org/10.1038/s41593-019-0502-4>.
61. Steinmetz, N.A., Zatka-Haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of choice, action and engagement across the mouse brain. *Nature* 576, 266–273. <https://doi.org/10.1038/s41586-019-1787-x>.
62. Salkoff, D.B., Zagha, E., McCarthy, E., and McCormick, D.A. (2020). Movement and Performance Explain Widespread Cortical Activity in a Visual Detection Task. *Cereb. Cortex* 30, 421–437. <https://doi.org/10.1093/cercor/bhz206>.
63. Kauvar, I.V., Machado, T.A., Yuen, E., Kochalka, J., Choi, M., Allen, W.E., Wetzstein, G., and Deisseroth, K. (2020). Cortical Observation by Synchronous Multifocal Optical Sampling Reveals Widespread Population Encoding of Actions. *Neuron* 107, 351–367.e19. <https://doi.org/10.1016/j.neuron.2020.04.023>.
64. Wang, Z.A., Chen, S., Liu, Y., Liu, D., Svoboda, K., Li, N., and Druckmann, S. (2023). 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>.
65. Markowitz, J.E., Gillis, W.F., Jay, M., Wood, J., Harris, R.W., Cieszkowski, R., Scott, R., Brann, D., Koveal, D., Kula, T., et al. (2023). Spontaneous behaviour is structured by reinforcement without explicit reward. *Nature* 614, 108–117. <https://doi.org/10.1038/s41586-022-05611-2>.
66. Klaus, A., Martins, G.J., Paixao, V.B., Zhou, P., Paninski, L., and Costa, R.M. (2017). The Spatiotemporal Organization of the Striatum Encodes Action Space. *Neuron* 95, 1171–1180.e7. <https://doi.org/10.1016/j.neuron.2017.08.015>.
67. van Wijngaarden, J.B., Babil, S.S., and Ito, H.T. (2020). Entorhinal-retrosplenial circuits for allocentric-egocentric transformation of boundary coding. *eLife* 9, e59816. <https://doi.org/10.7554/eLife.59816>.
68. Alexander, A.S., Carstensen, L.C., Hinman, J.R., Raudies, F., Chapman, G.W., and Hasselmo, M.E. (2020). Egocentric boundary vector tuning of the retrosplenial cortex. *Sci. Adv.* 6, eaaz2322. <https://doi.org/10.1126/sciadv.aaz2322>.
69. Fuster, J.M. (2001). The prefrontal cortex—an update: time is of the essence. *Neuron* 30, 319–333. [https://doi.org/10.1016/s0896-6273\(01\)00285-9](https://doi.org/10.1016/s0896-6273(01)00285-9).
70. Fine, J.M., and Hayden, B.Y. (2022). The whole prefrontal cortex is pre-motor cortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 377, 20200524. <https://doi.org/10.1098/rstb.2020.0524>.
71. Mazzucato, L. (2022). Neural mechanisms underlying the temporal organization of naturalistic animal behavior. *eLife* 11, e76577. <https://doi.org/10.7554/eLife.76577>.
72. Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., Bauza, M., Barbarits, B., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydin, Ç., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. *Nature* 551, 232–236. <https://doi.org/10.1038/nature24636>.
73. Deacon, R.M.J., Penny, C., and Rawlins, J.N.P. (2003). Effects of medial prefrontal cortex cytotoxic lesions in mice. *Behav. Brain Res.* 139, 139–155. [https://doi.org/10.1016/s0166-4328\(02\)00225-5](https://doi.org/10.1016/s0166-4328(02)00225-5).
74. Gomez-Marin, A., and Ghazanfar, A.A. (2019). The Life of Behavior. *Neuron* 104, 25–36. <https://doi.org/10.1016/j.neuron.2019.09.017>.
75. Dennis, E.J., El Hady, A., Michael, A., Clemens, A., Tervo, D.R.G., Voigts, J., and Datta, S.R. (2021). Systems Neuroscience of Natural Behaviors in Rodents. *J. Neurosci.* 41, 9119–9119. <https://doi.org/10.1523/JNEUROSCI.1877-20.2020>.
76. Pezzulo, G., and Cisek, P. (2016). Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition. *Trends Cogn. Sci.* 20, 414–424. <https://doi.org/10.1016/j.tics.2016.03.013>.
77. Siegel, M., Buschman, T.J., and Miller, E.K. (2015). Cortical information flow during flexible sensorimotor decisions. *Science* 348, 1352–1355. <https://doi.org/10.1126/science.aab0551>.
78. Weissbrod, A., Shapiro, A., Vasserman, G., Edry, L., Dayan, M., Yitzhaky, A., Hertzberg, L., Feinerman, O., and Kimchi, T. (2013). Automated long-term tracking and social behavioural phenotyping of animal colonies within a semi-natural environment. *Nat. Commun.* 4, 2018. <https://doi.org/10.1038/ncomms3018>.
79. Sharma, A., Johnson, R., Engert, F., and Linderman, S. (2018). Point process latent variable models of larval zebrafish behavior. In *NeurIPS*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds. (Curran Associates, Inc.).
80. McKenzie-Smith, G.C., Wolf, S.W., Ayroles, J.F., and Shaevitz, J.W. (2025). Capturing continuous, long timescale behavioral changes in *Drosophila melanogaster* postural data. *PLoS Comput. Biol.* 21, e1012753. <https://doi.org/10.1371/journal.pcbi.1012753>.
81. Stagkourakis, S., Spigolon, G., Marks, M., Feyder, M., Kim, J., Perona, P., Pachitariu, M., and Anderson, D.J. (2023). Anatomically distributed neural representations of instincts in the hypothalamus. Preprint at bioRxiv. <https://doi.org/10.1101/2023.11.21.568163>.
82. Miller, E.K., Erickson, C.A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J. Neurosci.* 16, 5154–5167. <https://doi.org/10.1523/JNEUROSCI.16-05-154.1996>.
83. Rainer, G., Asaad, W.F., and Miller, E.K. (1998). Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature* 393, 577–579. <https://doi.org/10.1038/31235>.
84. Barbosa, J., Proville, R., Rodgers, C.C., DeWeese, M.R., Ostojic, S., and Boubenec, Y. (2023). Early selection of task-relevant features through

population gating. *Nat. Commun.* 14, 6837. <https://doi.org/10.1038/s41467-023-42519-5>.

85. McKee, J.L., Riesenhuber, M., Miller, E.K., and Freedman, D.J. (2014). Task dependence of visual and category representations in prefrontal and inferior temporal cortices. *J. Neurosci.* 34, 16065–16075. <https://doi.org/10.1523/JNEUROSCI.1660-14.2014>.

86. Courtney, S.M. (2004). Attention and cognitive control as emergent properties of information representation in working memory. *Cogn. Affect. Behav. Neurosci.* 4, 501–516. <https://doi.org/10.3758/CABN.4.4.501>.

87. Jercog, D., Winke, N., Sung, K., Fernandez, M.M., Francioni, C., Rajot, D., Courtin, J., Chaudun, F., Jercog, P.E., Valerio, S., et al. (2021). Dynamical prefrontal population coding during defensive behaviours. *Nature* 595, 690–694. <https://doi.org/10.1038/s41586-021-03726-6>.

88. Zhou, T., Zhu, H., Fan, Z., Wang, F., Chen, Y., Liang, H., Yang, Z., Zhang, L., Lin, L., Zhan, Y., et al. (2017). History of winning remodels thalamo-PFC circuit to reinforce social dominance. *Science* 357, 162–168. <https://doi.org/10.1126/science.aak9726>.

89. Botvinick, M.M. (2008). Hierarchical models of behavior and prefrontal function. *Trends Cogn. Sci.* 12, 201–208. <https://doi.org/10.1016/j.tics.2008.02.009>.

90. Deco, G., Sanz Perl, Y., Ponce-Alvarez, A., Tagliazucchi, E., Whybrow, P.C., Fuster, J., and Krriegelbach, M.L. (2023). One ring to rule them all: The unifying role of prefrontal cortex in steering task-related brain dynamics. *Prog. Neurobiol.* 227, 102468. <https://doi.org/10.1016/j.pneurobio.2023.102468>.

91. Krriegelbach, M.L., and Deco, G. (2024). Prefrontal cortex drives the flexibility of whole-brain orchestration of cognition. *Curr. Opin. Behav. Sci.* 57, 101394. <https://doi.org/10.1016/j.cobeha.2024.101394>.

92. Rosenblueth, A., Wiener, N., and Bigelow, J. (1943). Behavior, Purpose and Teleology. *Philos. Sci.* 10, 18–24. <https://doi.org/10.1086/286788>.

93. Vane-Wright, R.I., and Corning, P.A. (2023). Teleonomy in living systems: an overview. *Biol. J. Linn. Soc.* 139, 341–356. <https://doi.org/10.1093/biolinnean/blad037>.

94. Botvinick, M., and Braver, T. (2015). Motivation and cognitive control: from behavior to neural mechanism. *Annu. Rev. Psychol.* 66, 83–113. <https://doi.org/10.1146/annurev-psych-010814-015044>.

95. Pierce-Messick, Z., Shipman, M.L., Desilets, G.L., and Corbit, L.H. (2024). Outcome devaluation as a method for identifying goal-directed behaviors in rats. *Nat. Protoc.* 20, 518–538. <https://doi.org/10.1038/s41596-024-01054-3>.

96. Gomez-Marin, A., Stephens, G.J., and Brown, A.E.X. (2016). Hierarchical compression of *Caenorhabditis elegans* locomotion reveals phenotypic differences in the organization of behaviour. *J. R. Soc. Interface* 13, 20160466. <https://doi.org/10.1098/rsif.2016.0466>.

97. Reddy, G., Desban, L., Tanaka, H., Roussel, J., Mirat, O., and Wyart, C. (2022). A lexical approach for identifying behavioural action sequences. *PLoS Comput. Biol.* 18, e1009672. <https://doi.org/10.1371/journal.pcbi.1009672>.

98. Gupta, S., and Gomez-Marin, A. (2019). A context-free grammar for *Caenorhabditis elegans* behavior. Preprint at bioRxiv. <https://doi.org/10.1101/708891>.

99. Stoffl, L., Bonnetto, A., d'Ascoli, S., and Mathis, A. (2024). Elucidating the Hierarchical Nature of Behavior with Masked Autoencoders. Preprint at bioRxiv. <https://doi.org/10.1101/2024.08.06.606796>.

100. Azabou, M., Mendelson, M.J., Ahad, N., Sorokin, M., Thakoor, S., Urvay, C., and Dyer, E.L. (2024). Relax, it doesn't matter how you get there: a new self-supervised approach for multi-timescale behavior analysis. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* (Curran Associates Inc.).

101. Hulsey, D., Zumwalt, K., Mazzucato, L., McCormick, D.A., and Jaramillo, S. (2024). Decision-making dynamics are predicted by arousal and unstructured movements. *Cell Rep.* 43, 113709. <https://doi.org/10.1016/j.cellrep.2024.113709>.

102. Gomez-Marin, A., Paton, J.J., Kampff, A.R., Costa, R.M., and Mainen, Z.F. (2014). Big behavioral data: psychology, ethology and the foundations of neuroscience. *Nat. Neurosci.* 17, 1455–1462. <https://doi.org/10.1038/nn.3812>.

103. Cisek, P., and Green, A.M. (2024). Toward a neuroscience of natural behavior. *Curr. Opin. Neurobiol.* 86, 102859. <https://doi.org/10.1016/j.conb.2024.102859>.

104. Kennedy, A. (2022). The what, how, and why of naturalistic behavior. *Curr. Opin. Neurobiol.* 74, 102549. <https://doi.org/10.1016/j.conb.2022.102549>.

105. Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295–300. <https://doi.org/10.1038/nature12354>.

106. Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B.L., Koay, S.A., Taxidis, J., Najafi, F., Gauthier, J.L., Zhou, P., et al. (2019). CalmAn an open source tool for scalable calcium imaging data analysis. *eLife* 8, e38173. <https://doi.org/10.7554/eLife.38173>.

107. Buccino, A.P., Hurwitz, C.L., Garcia, S., Magland, J., Siegle, J.H., Hurwitz, R., and Hennig, M.H. (2020). SpikeInterface, a unified framework for spike sorting. *eLife* 9, e61834. <https://doi.org/10.7554/eLife.61834>.

108. Pachitariu, M., Sridhar, S., Pennington, J., and Stringer, C. (2024). Spike sorting with Kilosort4. *Nat. Methods* 21, 914–921. <https://doi.org/10.1038/s41592-024-02232-7>.

109. Birman, D., Yang, K.J., West, S.J., Karsh, B., Browning, Y., International Brain Laboratory, Siegle, J.H., and Steinmetz, N.A. (2023). Pinpoint: trajectory planning for multi-probe electrophysiology and injections in an interactive web-based 3D environment. Preprint at bioRxiv. <https://doi.org/10.1101/2023.07.14.548952>.

110. Sun, K., Xiao, B., Liu, D., and Wang, J. (2019). Deep High-Resolution Representation Learning for Human Pose Estimation. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5686–5696.

111. McInnes, L., Healy, J., Saul, N., and Großberger, L. (2018). UMAP: Uniform Manifold Approximation and Projection. *J. Open Source Softw.* 3, 861. <https://doi.org/10.21105/joss.00861>.

112. Gerfen, C.R., Paletzki, R., and Heintz, N. (2013). GENSAT BAC cre-recombinase driver lines to study the functional organization of cerebral cortical and basal ganglia circuits. *Neuron* 80, 1368–1383. <https://doi.org/10.1016/j.neuron.2013.10.016>.

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

114. Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., and Liang, J. (2018). UNet++: A Nested U-Net Architecture for Medical Image Segmentation. *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018* (Springer-Verlag).

115. Cheng, B., Xiao, B., Wang, J., Shi, H., Huang, T.S., and Zhang, L. (2020). HigherHRNet: Scale-Aware Representation Learning for Bottom-Up Human Pose Estimation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5385–5394.

116. Zhou, Q.-Y., Park, J., and Koltun, V. (2018). Open3D: A Modern Library for 3D Data Processing. Preprint at arXiv.

117. Zhang, L., Dunn, T., Marshall, J., Olveczky, B., and Linderman, S. (2021). Animal pose estimation from video data with a hierarchical von Mises-Fisher-Gaussian model. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, B. Arindam and F. Kenji, eds. (PMLR).

118. Bonald, T., Charpentier, B., Galland, A., and Hollocou, A. (2018). Hierarchical Graph Clustering using Node Pair Sampling. Preprint at arXiv.
119. Hsu, A.I., and Yttri, E.A. (2021). B-SOI_D, an open-source unsupervised algorithm for identification and fast prediction of behaviors. *Nat. Commun.* **12**, 5188. <https://doi.org/10.1038/s41467-021-25420-x>.
120. Luxem, K., Mocellin, P., Fuhrmann, F., Kürsch, J., Miller, S.R., Palop, J.J., Remy, S., and Bauer, P. (2022). Identifying behavioral structure from deep variational embeddings of animal motion. *Commun. Biol.* **5**, 1267. <https://doi.org/10.1038/s42003-022-04080-7>.
121. Hubert, L., and Arabie, P. (1985). Comparing partitions. *J. Classif.* **2**, 193–218. <https://doi.org/10.1007/BF01908075>.
122. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830.
123. Li, Y., Mathis, A., Grewe, B.F., Osterhout, J.A., Ahanonu, B., Schnitzer, M.J., Murthy, V.N., and Dulac, C. (2017). Neuronal Representation of Social Information in the Medial Amygdala of Awake Behaving Mice. *Cell* **171**, 1176–1190.e17. <https://doi.org/10.1016/j.cell.2017.10.015>.
124. Newman, M.E. (2006). Modularity and community structure in networks. *Proc. Natl. Acad. Sci. USA* **103**, 8577–8582. <https://doi.org/10.1073/pnas.0601602103>.
125. Jocher, G., Chaurasia, A., Stoken, A., Borovcik, J., NanoCode012, Kwon, Y., Michael, K., TaoXie, Fang, J., imyhxy, et al. (2022). ultralytics/yolov5: v7.0 - YOLOv5 SOTA Realtime Instance Segmentation. Zenodo. <https://doi.org/10.5281/zenodo.7347926>.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Antibodies		
Chicken Anti-GFAP	Abcam	Cat# ab4674, RRID:AB_304558
Alexa Fluor 488, Donkey Anti-Chicken IgY	Jackson Laboratory	Cat# 703-545-155; RRID_AB_2340375
Bacterial and virus strains		
AAV1.Syn.GCaMP6f.WPRE.SV40	Douglas Kim & GENIE Project; Chen et al. ¹⁰⁵	Addgene viral prep # 100837-AAV1
AAV1.Syn.Flex.GCaMP6f.WPRE.SV40	Penn Vector Core	Cat# AV-1-PV2819
AAV1.EF1a.DIO.GCaMP6s.P2A.nls.dTomato	Jonathan Ting	Addgene viral prep # 51082-AAV1
AAV9.CBA.DO(Fas).GCaMP6s	UNC Vector Core	N/A
Chemicals, peptides, and recombinant proteins		
N-Methyl-D-Aspartate	Sigma	Cat# M3262-25MG
Vybrant™ CM-Dil Cell-Labeling Solution	ThermoFisher	Cat# V22888
Deposited data		
Raw and processed behavior and neural activity	Zenodo	10.5281/zenodo.17488068
Experimental models: Organisms/strains		
Mouse: CD1	Charles River	Strain code: 022
Mouse: B6.FVB(Cg)-Tg(Drd1-cre)EY262Gsat/Mmucd	MMRRC-UCD	Stock# 030989-UCD
Mouse: B6.FVB(Cg)-Tg(Adora2a-cre)KG139Gsat/Mmucd	MMRRC-UCD	Stock# 036158-UCD
Mouse: BALB/cJ	Jackson Laboratory	Jax stock #000651
Mouse: C57BL/6J	Jackson Laboratory	Jax stock #000664
Software and algorithms		
Custom depth-MoSeq pipeline	Zenodo	10.5281/zenodo.17488068
Multi-camera calibration pipeline	Zenodo	10.5281/zenodo.17488068
Keypoint-MoSeq	Weinreb et al. ⁴⁹	https://github.com/dattalab/keypoint-moseq
ShMoSeq	Zenodo	10.5281/zenodo.17488068
CalmAn	Giovannucci et al. ¹⁰⁶	https://github.com/flatironinstitute/CalmAn
Open Ephys (v0.6)	Open Ephys	https://open-ephys.org
SpikeInterface	Buccino et al. ¹⁰⁷	https://spikeinterface.readthedocs.io
Kilosort 4	Pachitariu et al. ¹⁰⁸	https://github.com/MouseLand/Kilosort
Pinpoint	Birman et al. ¹⁰⁹	https://github.com/VirtualBrainLab/Pinpoint
HRNet	Sun et al. ¹¹⁰	https://github.com/HRNet
UMAP	McInnes et al. ¹¹¹	https://umap-learn.readthedocs.io

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

All experiments were carried out in accordance with Harvard Medical School institutional animal care and use committee (IACUC) protocol number IS00000138. Subject mice included wildtype C57BL6/J mice (Jackson Laboratory stock no. 000664), C57BL6/J mice harboring either the Drd1a-Cre allele (B6.FVB(Cg)-Tg(Drd1-cre)EY262Gsat/Mmucd; MMRRC #030989-UCD) or the A2a-Cre allele (B6.FVB(Cg)-Tg(Adora2a-cre)KG139Gsat/Mmucd; MMRRC #036158-UCD),¹¹² CD1 mice (Charles River, #022), and BALB/cJ mice (Jackson Laboratory stock no. 000651). All mice were recorded between 3 and 7 months of age and maintained in a 12 h:12h light/dark cycle with food and water ad libitum. Individual housing was used following surgery and group-housing otherwise.

We recorded dorsomedial prefrontal cortex (dmPFC) activity in 6 male mice across 31 sessions in the initial round of solitary and social open field recordings, 4 male mice across 4 sessions during the aggression experiments, 9 female mice across 18 sessions during the object investigation experiments with unique objects, and the same 9 female mice across 9 sessions during the object investigation experiments with identical objects. We recorded dorsolateral striatum (DLS) activity in 21 mice across 30 sessions. For the lesion experiments, we recorded 10 male mice (5 lesion, 5 sham) across 34 sessions (18 with novel objects, 16 without),

followed by a second cohort of 8 male mice (4 lesion, 4 sham) across 16 sessions (without objects). We also tracked 3D keypoints in 5 additional male C57 mice interacting with objects (21 sessions) or with female BALB/cJ conspecifics (19 sessions). For neuropixels recordings in dmPFC, we recorded 3 male C57 mice across 13 sessions.

METHOD DETAILS

Behavior Assays

All mice were habituated to handling, head-fixation, and the recording arena for before experiments began. On recording days, mice were brought to the laboratory and habituated in darkness for at least 20 minutes before recording. All experiments were performed during the dark cycle under infrared illumination.

Open field exploration

During the initial set of social and solitary open field recordings with a C57 conspecific, subject mice were recorded for 80 minutes in alternating blocks: 10 minutes alone, 30 minutes with an unfamiliar conspecific, 10 minutes alone, and then 30 minutes with the same conspecific. Some conspecifics bore a commutated patch cord that was not used during the experiment.

Conspecific aggression

For recordings of aggressive social interaction, subject mice were recorded in an open field arena for one hour in alternating blocks: 5 minutes alone, 5 minutes with an aggressive conspecific (CD1), 5 minutes with a non-aggressive conspecific (CD1), and so on (same pair of conspecifics each time). Prior to the initiation of experimental trials, CD1 mice were pre-screened to identify highly aggressive or non-aggressive individuals using the following procedure: CD1 were placed alone in an open field arena and allowed to habituate for 10 minutes. We then performed a sequence of four probe trials in which a C57 mouse was placed in the arena for 10 minutes or until the onset of aggression (as identified by a trained observer). CD1 mice with the shortest average latency to aggression were defined as highly-aggressive. CD1 mice who never engaged in aggressive behavior were defined as non-aggressive.

Novel object exploration

For the novel object experiments, mice were recorded for one-hour sessions in the 3D keypoint tracking arena. A new object was added to the arena every ten minutes, resulting in five (unique) objects by the end of the recording; objects were randomly selected from the following set: 15 mL Falcon tube cap, pencil nub, rolled up piece of lab tape, small rubber band, crumpled twist tie, cube of packing foam, origami star, gear, syringe cap, tin foil ball, bubble wrap clipping, hexnut, glove clipping. We also recorded sessions with five identical objects, which in every case were hex nuts.

Lesion recordings

Lesioned mice were recorded in the 3D keypoint tracking arena. In a subset of recordings, novel objects (hex nuts) were added to the arena at ten minute intervals as described above. In another set, mice were allowed to explore the arena for one hour without interruption. Among this full set of recordings, we identified a handful of extreme outlier recordings in which mice remained in one small part of the arena for almost the entire session. We therefore excluded from analysis all sessions where the mouse was immobile (< 1 mm/s centroid velocity) at least 75% of the time; a total of 4 out of 56 sessions were excluded.

Stereotactic surgery procedures

For all stereotactic surgeries, mice were anaesthetized using 1–2% isoflurane in oxygen at a flow rate of 1 L/min and injected with bupivacaine (1.25 mg/kg) under the scalp. All coordinate axes were zeroed relative to bregma (including dorsal/ventral, which was zeroed relative to the skull surface), and coordinates are in units of mm. All injections were performed using a Nanoject II or Nanoject III (Drummond) at 60 nL/min. Incisions were closed using Vetbond (3M). Postoperative care included a subcutaneous injection of buprenorphine SR (1 mg/kg, given 1 hour prior to surgery start) and carprofen (5 mg/kg) administered through drinking water.

Calcium imaging in dorsomedial prefrontal cortex (dmPFC)

To record neural activity in dmPFC, virus injection and gradient index (GRIN) lens implantation were performed across two different surgeries. In the first surgery, we injected AAV1.Syn.GCaMP6f.WPRE.SV40¹⁰⁵ (Addgene #100837, titer: 0.8e13 – 1.2e13) unilaterally into the left hemisphere at two different depths (2.0 mm AP, 0.38 mm ML, -1.7 & -1.9 mm DV, 350 nL each). Incisions were closed and mice were allowed 4–8 weeks for recovery and viral expression. In the second surgery, a circular craniotomy was opened to the left of the injection site (center: 2.0 mm AP, 0.75 mm ML, diameter: 1.4 mm). Brain tissue exposed by the craniotomy was aspirated to an approximate depth of 1 mm from the skull surface at the margin of the craniotomy. The head was then rotated 10 degrees clockwise and a GRIN lens (1 mm diameter, 4 mm length, Inscopix part 1050-004623) was lowered through the center of the craniotomy to an approximate depth of 1.7 mm from the skull surface at bregma. The GRIN lens and medical-grade titanium headbars were then anchored to the skull using biocompatible cyanoacrylate superglue (Loctite 454). After 2–4 weeks, a baseplate (Inscopix part 1050-004638) was added with additional superglue and exposed glue was light-sealed using black nail polish.

Calcium imaging in dorsolateral striatum (DLS)

To record neural activity in DLS, virus injection and GRIN lens implantation were performed in a single surgery, which varied slightly depending on the population being imaged (direct pathway, indirect pathway, or both at once). For pathway-specific imaging, 500–600 nL of AAV1.Syn.Flex.GCaMP6f.WPRE.SV40 virus (Penn Vector Core) was injected into the right DLS (AP 0.5, ML 2.25, DV 2.4) of Drd1a-Cre (N=7) or A2a-Cre (N=7) mice respectively and a GRIN lens (1 mm diameter, 4 mm length; Inscopix part 130-000143) was implanted 200uM above the injection site immediately following virus injection. Pathway-independent imaging followed a similar

procedure, except we injected Drd1a-Cre mice (N=5) with a 1:1 mixture containing AAV1.EF1a.DIO.GCaMP6s.P2A.nls.dTomato (Cre-On; Addgene #51082) and AAV9.CBA.DO(Fas).GCaMP6s (Cre-Off; UNC Vector Core). Data from all pathways were combined for analysis since the key results did not differ by pathway (Figure S5i).

Electrophysiological recordings in dmPFC

To record spiking activity in dmPFC, we implanted Neuropixels 1.0 probes (IMEC) chronically into the right medial prefrontal cortex (AP +2.3 mm, ML +0.5 mm, DV 3.5–4.5 mm from the cortical surface). Prior to implantation, probes were coated with the lipophilic tracer Dil (Fisher Scientific, catalog #V22888) for histological verification.

dmPFC lesions

To lesion the dmPFC, we injected male C57BL/6J mice (8 weeks old) with N-Methyl-D-Aspartate (NMDA; Sigma, M3262-25MG) dissolved in saline (10 mg/mL). Control (sham) mice were injected with saline alone. Two injections of 120 nL were made in each hemisphere (i.e., four total injections per mouse) at coordinates AP 2.3, ML ±0.4, DV –2.3 and AP +1.7, ML ±0.4, DV –2.0. Mice were recorded 5–12 weeks after surgery.

Histological verification

Following completion of behavioral tests, a subset of mice was anaesthetized using 1–2% isoflurane in air and perfused with cold PBS followed by 4% paraformaldehyde. Coronal brain sections (60 μm) were sliced on a Leica VT1000 vibratome. All slices were stained with DAPI, and slices from lesioned mice were additionally stained with antibodies for glial fibrillary acidic protein (GFAP; Abcam; ab4674, 1/1000 dilution). Slices were imaged with an Olympus VS120 Virtual Slide Microscope.

Recording setups

Open field arena

Neural and behavioral recordings in DLS were performed as previously described.⁴⁸ Recordings in dmPFC (excluding the novel object experiments) used an updated open field arena with a transparent floor that allowed simultaneous depth and infrared (IR) acquisition from above and below the animal. A pair of Kinect Azure cameras acquired synchronous 30Hz video data while neural calcium transients were recorded at 15Hz using the nVista 3.0 platform from Inscopix. To synchronize the two recording modalities, the nVista trigger port was connected to an Arduino that updated an array of IR LEDs in tandem with the neural data acquisition.

High-speed 3D keypoint tracking

Electrophysiological recordings, novel object recordings, and lesion recordings were captured using an array of 6 Basler ace acA1300-200um Monochrome USB 3.0 Cameras (Edmund Optics 33-978) as previously described.⁴⁹ The cameras were triggered at 120Hz using an Arduino, and a copy of the trigger signal was sent to the Inscopix or neuropixels data acquisition system for synchronization. The arena was illuminated with 32 near-infrared high power LED stars (LEDSupply, CREEXPE-FRD-3). To avoid reflections and saturations effects, the bottom camera was triggered slightly out of phase with the top cameras and the LEDs were split into two groups: one group below the arena that turned on during the bottom camera's exposure, and one group above the arena that turned on during the top and side cameras' exposures.

Preprocessing of neural data

Miniscope recordings

Calcium-fluorescence videos were recorded at 15Hz using the nVista 3.0 platform from Inscopix. Videos were spatially down sampled by a factor of 4 and motion corrected using CalmAn.¹⁰⁶ We used the constrained non-negative matrix factorization for microendoscopic data (CNMF-e) algorithm^{106,113} to extract regions of interest (ROIs) corresponding to putative cells and their corresponding activity traces. ROIs were filtered using quality metrics output by CalmAn (CNN prediction > 0.05; signal-to-noise ratio > 2.5). Recordings with fewer than 50 detected neurons were rejected. Raw CNMF-e traces were Z-scored before downstream analysis.

Electrophysiology recordings

Neuropixels signals were recorded with Open Ephys (V0.6). Data was then preprocessed with SpikeInterface¹⁰⁷; we applied a band-pass filter between 300 and 6000Hz, followed by a global common median reference. We then spikersorted each recording independently using Kilosort 4.¹⁰⁸ In Kilosort 4, we set three hyperparameters: 'amplitude_cutoff_thresh' was set to 0.1, 'isi_violations_ratio_thresh' was set to 1, and 'presence_ratio_thresh' was set to 0.9; the remaining hyperparameters were left to default values. Spike times were aligned to video frames and then binned using a 100ms sliding window. Unit coordinates were determined based on the position of the Neuropixels channel with the maximum amplitude for that unit. That channel's position within the brain was then estimated using the Allen Institute brain atlas using the Pinpoint library.¹⁰⁹ Only units mapping to the prelimbic cortex were used for downstream analysis.

MoSeq analysis from top/bottom depth and infrared

To estimate syllables from top-down and bottom-up depth/infrared cameras in the presence of occluders (e.g., during social interaction), we developed a novel pipeline that included segmentation, multi-camera registration, missing data imputation, dimensionality reduction and MoSeq analysis (Videos S1 and S2). A version of this pipeline for single-animal recordings is available online (<https://github.com/calebweinreb/top-bottom-moseq>).

Segmentation

We first used a neural network with a UNET++ architecture¹¹⁴ to segment the two interacting mice as well as the miniscope and its attached cable. To reliably distinguish between mice, we marked the tail base of one with a black sharpie and performed multi-animal keypoint tracking using higherHRNET.¹¹⁵ Gaussian activations corresponding to tracked keypoints were included as additional channels when training and applying the segmentation pipeline (Figures S1A and S1B).

Multi-camera registration

Intrinsic parameters for the top-down and bottom-up depth cameras were extracted using the Kinect Azure API. Extrinsic parameters were calculated from videos of a checkerboard calibration object, as demonstrated in the “calibrate.ipynb” notebook from the top-bottom-moseq repository. After calibration, 3D point clouds of segmented mice were derived from each camera and then embedded in a common coordinate system. Outlier points were removed using a nearest neighbor algorithm from Open3D.¹¹⁶ The embedded point clouds were then rendered as a pair of reconstructed depth-maps from virtual orthographic top-down and bottom-cameras. IR reflectance signals were propagated along with the point clouds. The resulting depth and IR images were similar to the original camera outputs, but now registered to a common coordinate system. Areas that were out of view from the original (non-orthographic) camera angles (e.g., along the far flank of the mouse) were registered as missing data (see [Video S1](#) for an example). These steps can be reproduced using the “orthographic_reprojection” function from the top-bottom-moseq repository.

Missing data imputation

Imputation of missing pixel values was performed in three steps: a forward pass, a backward pass, and a merging step (Figure S1C; [Video S2](#)). Each pass was performed using a convolutional neural net (CNN) with a U-net architecture. During the forward pass, images were imputed based on masked observations from the current timepoint and imputation outputs from the previous timepoint. The purpose of recurrence here was to propagate relevant pose information across frames and to ensure temporal continuity of the imputed videos. Depth and IR channels from both cameras were imputed simultaneously using a single network. The backward pass was the same except that frames were fed in the opposite order. Another neural net was used to merge the forward and backward pass outputs. The resulting consensus images for each frame thus incorporated relevant pose information from past and future timepoints. The above steps can be reproduced using the “inpaint_session” function from the top-bottom-moseq repository.

Dimensionality reduction and MoSeq

After imputation, top-down and bottom-up depth/IR images were centered and aligned and then projected into low-dimensional pose space. Rather than using principal components analysis (PCA) for dimensionality reduction – as is done in the standard MoSeq pipeline⁵² – we opted for a shallow convolutional autoencoder due to its greater representational capacity ([Video S2](#)). The autoencoder included a 10-dimensional bottleneck layer and activations in this layer were used as the final pose representation. Pose trajectories were then modeled using the standard MoSeq pipeline.⁵² Separate MoSeq models were fit for the initial set of open field dmPFC recordings (with a C57 conspecific), the DLS open field recordings, and the aggression recordings with a CD1 conspecific.

MoSeq analysis from 3D keypoints

For the high-speed multi-camera setup, 3D keypoint tracking and syllable inference were performed as previously described.⁴⁹ Briefly, 2D keypoint detection from each of the six camera angles was performed using an HRNet.¹¹⁰ We then performed 3D triangulation using a custom multi-camera calibration pipeline (https://github.com/calebweinreb/multicam-calibration/tree/main/multicam_calibration). Keypoint trajectories were refined using GIMBAL,¹¹⁷ which is a model-based approach that leverages anatomical constraints and motion continuity. The refined keypoint estimates were then fed to keypoint-MoSeq⁴⁹ for syllable inference. For the latter step, we used a 6-dimensional latent space, set the maximum syllable count of 50, and targeted a median syllable duration of 400ms. Separate models were fit for the lesion dataset, the miniscope recordings with novel objects, and the neuropixels dataset.

Syllable coarse-graining

Clustering syllables based on kinematic similarity

To cluster syllables by kinematic similarity, we calculated the pairwise distance between their mean pose trajectories and then performed hierarchical clustering using complete linkage. For syllables derived from depth-MoSeq, the pose trajectory for a single syllable instance was defined as the sequence of dimensionally reduced poses spanning 5 frames before to 15 frames after the syllable start time and the mean trajectory was calculated as by averaging 1000 instances. The similarity between mean trajectories was computed using Pearson correlation. Keypoint-derived syllables, mean trajectories were estimated using the “get_typical_trajectories” function from the keypoint-moseq python package (version 0.4.7), with parameters set as follows: pre=20, post=60, min_duration=12, density_sample=False. The similarity between mean trajectories was computed using cosine distance.

Clustering syllables based on transition probabilities

To cluster syllables based on transition probabilities (Figures S1J and S1K), we represented the transition matrix as a directed graph. We then applied the Paris algorithm for hierarchical graph clustering,^{54,118} which outputs a dendrogram describing the successive merges of syllables into clusters. To select levels of the hierarchy at which to partition the dendrogram, we examined the differences in merge distances across consecutive steps. Specifically, we calculated the first derivative of the dendrogram linkage values and identified prominent local maxima using a peak-finding algorithm (minimum threshold = 0.02, minimum spacing = 3 merges). Three maxima were identified corresponding to $n = 13$, $n = 9$, and $n = 4$ clusters respectively.

Analysis of behavioral timescales

To assess the predictability of syllable sequences, we computed mutual information (MI) at a range of temporal lags. Given offset Δt and syllable sequence z_1, \dots, z_T , we defined:

$$\text{predictability}(\Delta t) = \text{MI}(\{z_t, z_{t+\Delta t} \mid \text{for all } t \text{ such that } 1 \leq t \leq T - \Delta t\})$$

For Markov model comparisons, we calculated an empirical transition matrix for each recording and then sampled a Markovian sequence with the same length as the real sequence. For comparisons to shMoSeq, we fit the parameters of a shMoSeq model using real syllable sequences and then sampled synthetic state and syllable sequences by conditioning on those parameters (see “Inferring behavioral states” for details of model fitting and simulation).

We used a shuffle procedure to estimate the null distribution of MIs corresponding to a complete lack of predictability in behavior. The shuffle MIs were calculated by pairing independent syllable sequences from separate recordings. Specifically, given offset Δt , syllable sequence z_1, \dots, z_T , and a second syllable sequence z'_1, \dots, z'_T from another recording, we defined:

$$\text{shuffle predictability}(\Delta t) = \text{MI}(\{z_t, z'_{t+\Delta t} \mid \text{for all } t \text{ such that } 1 \leq t \leq T - \Delta t\})$$

this procedure was repeated once for each syllable sequence z_1, \dots, z_T , with the second sequence z'_1, \dots, z'_T sampled uniformly at random from the dataset. To estimate the temporal horizon over which a syllable sequence remains predictable, we calculated minimum Δt required for its MI to reach a noise floor (“mutual information decay time”). The noise floor was defined as the 99th percentile of the shuffle MIs (calculated over all shuffled sequences and Δt ’s).

In [Figures 1D, 1F, 7B, 7J, S7N, and S7T](#), we used coarse-graining to improve the accuracy of MI estimation. Our reasoning was that MI becomes unreliable when there are too many syllable categories, as some categories occur only rarely. With limited samples, these rare categories inflate apparent correlations simply by chance, leading to an overestimate of predictability. To mitigate this sampling noise, we grouped syllables into 8 clusters based on kinematic similarity and relabeled them before computing MI (i.e. after the Markov and shMoSeq simulations mentioned above).

Coarse-graining was also applied in [Figures S1K and S1M](#), but here to test whether non-Markovian sequence statistics were due to over-fragmentation of behavior. Syllables in this case were grouped into $4 \leq N \leq 13$ kinematic clusters ([Figure S1K](#)) or $3 \leq N \leq 16$ transition-based clusters ([Figure S1M](#)) and relabeled before both MI computation and Markov simulation.

To examine predictability within individual states ([Figure S7T](#)), we calculated MI over pairs of timepoints that fell within contiguous instances of each state. Since MI calculations are sensitive to the total amount of data used, we sampled an equal number of instances of each state ($n = 200$ instances) and then used these to calculate per-state MI curves. This procedure was repeated 100 times to generate a bootstrap distribution of MI values for each temporal offset.

ShMoSeq: generative model and fitting algorithm

Model motivation

We sought a model that could (i) identify behavior states corresponding to different patterns of syllable usage over time; (ii) efficiently use the data available in a typical experiment (i.e. avoid over-parameterization). To that end, we created a hierarchical hidden Markov model (HHMM) called state-based hierarchical MoSeq (shMoSeq). In the model, each behavioral state specifies a transition matrix over syllables. Crucially, shMoSeq parameterizes each state’s transition matrix by combining a baseline transition matrix (that is shared across behavioral states) with a vector of syllable biases (that is specific to each state) ([Figure S2A](#)). The baseline matrix captures general first order statistics – such as self-transitions among syllables or the tendency of mice rear down after they rear up – whereas the bias term state-dependent shifts in overall syllable usage. By combining these two terms, shMoSeq can express a wide variety of sequence statistics while limiting the total number of parameters.

shMoSeq has several advantages over a standard HHMM. First, it allows the dominant feature of syllable transitions – their auto-correlation – to be represented non-redundantly in the baseline transition matrix, rather than being separately ‘relearned’ by each state. Second, it allows for robust fitting given limited data; whereas a standard HHMM would have roughly $n_{\text{states}} * n_{\text{syllables}}^2$ total parameters, shMoSeq has on the order of $n_{\text{states}} * n_{\text{syllables}} + n_{\text{syllables}}^2$ parameters. To illustrate why this is important, note that a typical experiment includes several hours of data with roughly 5000 syllable transitions per hour. If we assume that there are 50-100 unique syllables and 5-10 behavioral states, then the number of parameters in a standard HHMM (10^4 – 10^5) would rival or exceed the total number of independent observations in the dataset. By contrast, the number of parameter in a shMoSeq model (10^3 – 10^4) would be well below the number of observations.

To formally test this intuition, we simulated data from a shMoSeq model instantiated with random parameters and then performed model fitting using either shMoSeq or a standard HHMM. The simulated datasets had 250,000 frames (~2.5 hours at 30 frames per second) and 5 states with average durations of ~10 seconds. For the underlying syllable sequences, we scanned over a range of syllable counts (5, 10, 20, or 50) and durations (1–100 frames). We found that for small numbers of syllables and/or short syllable durations, shMoSeq and the standard HHMM performed equally well ([Figure S2B](#)). However, in the more realistic scenario of 50 syllables with durations of ≥ 10 frames, shMoSeq continued to perform well while the standard HHMM failed catastrophically.

We note that an alternative approach for reducing the parameter count is to use a categorical hidden Markov model (CatHMM), where each behavioral state specifies a categorical distribution over syllables. CatHMMs have the virtue of being parameter efficient,

but they introduce a new problem by assuming that syllables are conditionally independent given states. Because of this false assumption, CatHMMs are liable to overestimate the certainty of state sequences. This is illustrated by the following scenario: imagine there is a “walk” syllable that typically lasts half a second (15 frames) and occurs with 30% probability in state A and 60% probability in state B. Whenever this syllable occurs, the CatHMM will treat it as ~ 15 independent observations. If states A and B are equally likely a priori, the posterior probability of state B would be:

$$P(B \mid 15 \text{ frames of walk}) = \frac{P(\text{walk} \mid B)^{15}}{P(\text{walk} \mid A)^{15} + P(\text{walk} \mid B)^{15}} = 0.99997$$

whereas the correct estimate would be closer to:

$$P(B \mid \text{one instance of walk}) = \frac{P(\text{walk} \mid B)}{P(\text{walk} \mid A) + P(\text{walk} \mid B)} = 0.67$$

Thus, compared to alternatives such as standard HHMMs or CatHMMs, the parameterization instantiated in shMoSeq is well-suited to the syllable counts, autocorrelations, and dataset sizes typical of MoSeq datasets.

Model definition

The generative model implemented in shMoSeq is defined as follows, where w_t denotes the current behavioral state and z_t denotes the current syllable.

$$\begin{aligned} w_t &\sim \text{Cat } \pi_{w_{t-1}} && \text{(state sequence)} \\ z_t &\sim \text{Cat } T_{z_{t-1}}^{(w_t)} && \text{(syllable sequence)} \end{aligned}$$

To prevent rapid switching between states, the transition matrix π has a sticky Dirichlet prior (defined below), in which each row π_w is drawn from a Dirichlet distribution that is boosted in the w 'th position by a stickiness hyperparameter κ .

$$\pi_w \sim \text{Dirichlet}(\beta, \dots, \beta + \kappa, \dots, \beta)$$

The size of κ determines the bias toward self-transitions, hence the durations of behavioral states. The syllable transition matrices $T^{(w)}$ are constructed using a generalized linear model that adds state-specific syllable biases B_w to a baseline matrix A that is shared between states.

$$T_z^{(w)} = \text{Softmax}(A_z + B_w)$$

Because softmax is invariant to constant shifts of its inputs, we require that the rows of A and B are centered. We also assume that the columns of B are centered, since any uniform shift in a column of B can be countered by an equal and opposite shift in the same column of A without affecting the sums $\{A_z + B_w \mid \forall w, z\}$. We therefore define

$$\begin{aligned} A &= A' C_n & A'_{ij} &\sim \mathcal{N}(0, \sigma_A^2) \\ B &= C_m^\top B' C_n & B'_{ij} &\sim \mathcal{N}(0, \sigma_B^2) \end{aligned}$$

where n is the total number of syllables, m is the total number of states, and C_n is an orthonormal matrix that embeds \mathbb{R}^{n-1} into the space of centered vectors in \mathbb{R}^n .

Parameter inference

Syllable sequences were derived using MoSeq and treated as observed. The shMoSeq model parameters and behavioral state sequence were fit using Gibbs sampling. We used a Laplace approximation to sample (A', B') . The mode of the Laplace approximation was approximated by gradient descent and the Hessian was computed through auto differentiation. All the other variables were sampled exactly. The behavioral state sequence sampled during the final Gibbs step was carried forward for analysis. Marginal probabilities of behavioral states were computed exactly using the forward-backward algorithm.

ShMoSeq: practical application

Software availability and compute requirements

shMoSeq can be run on any operating system that supports JAX (<https://github.com/jax-ml/jax>). A CUDA-capable GPU substantially reduces runtime but is not required. For reference, a 90-hour dataset takes ~ 1 minute to fit on a GPU and requires ~ 17 GB of memory. The same dataset takes ~ 6 minutes on a CPU. Memory and runtime scale approximately with dataset duration, the number of syllables, and the number of states. shMoSeq can in principle be applied to any sequence of behavioral labels, not just those derived from MoSeq. Examples include labels generated by other unsupervised algorithms (e.g., B-SOI¹¹⁹, VAME¹²⁰) or from supervised classifiers, provided the sequences exhibit non-Markovian structure. We provide an online tutorial for mutual information-based analysis of non-Markovian statistics (as in [Figures 1F and 1G](#)).

To understand how much data is needed for model fitting, we fit shMoSeq to randomly down-sampled copies of our initial open field dataset. The quality of each fit was assessed using (1) log probability of the model (when applied to the full dataset; [Figure S2N](#)); (2) similarity to the final model that was fit using the full dataset ([Figure S2O](#)); this latter metric tests whether the down-sampled

dataset is functionally equivalent to the full dataset with respect to model outputs. Both metrics rose steeply from 0 - 10 hours of data and then plateaued (or rose gradually) beyond 20 hours. To test the generalizability of these results, we repeated the same experiment using a larger dataset [200 hours of open field exploration from a previous paper⁶⁵] and a range of values for the number of states (Figures S2P and S2Q). This time model quality plateaued after 20-40 hours of data, although some improvement beyond this point was evident for higher state counts. Based on these results, we recommend that datasets be ≥ 10 hours at minimum and ideally ≥ 40 hours.

When applying shMoSeq to a new dataset, it is often necessary to scan over stickiness values, state counts, and random seeds. To reduce the burden of hyperparameter search, we analyzed how the optimal stickiness varies with dataset size and number of states. Specifically, we fit shMoSeq to down-sampled copies of a large open field dataset from a previous paper⁶⁵, we down-sampled to a range of dataset sizes and fit models with a range of state numbers, stickiness parameters and random seeds. After determining the optimal stickiness for each state count and dataset size, we observed a power-law scaling relationship (i.e. linear scaling in log space), suggesting that users can extrapolate the optimal stickiness from a smaller set of test fits (Figures S2R-S2T). The scaling relationship can be expressed as follows; coefficients will vary by dataset.

$$\log(\text{optimal stickiness}) \approx \beta_0 + \beta_1 \log(\text{number of states}) + \beta_2 \log(\text{amount of data})$$

Fitting and hyperparameter selection

We fit separate MoSeq models and thus separate shMoSeq models for each dataset in the paper (see Table 2 for a list of datasets). In each case, we set $\sigma_A = \sigma_B = 1$, $\beta = 1$ and performed 500 Gibbs iterations, at which point we always observed a plateau in the held out log likelihood. We note that while 500 Gibbs iterations is enough to obtain a reasonable point estimate of the model parameters, we do not observe full mixing, so the outcome of fitting should not be viewed as a representative sample from the posterior distribution of the model. To determine the stickiness parameter and the total number of states for each dataset, we performed a two-parameter grid scan and fit 5 to 10 models for each parameter combination. Models were fit using 75% of the data and held-out probabilities were computed using the remaining 25% of the data. Across datasets, we found that the held-out probability tended to peak for intermediate values of the stickiness parameter; in each case, we used this peak to choose a stickiness value for the final model.

Choosing the number of shMoSeq states

When scanning over the number of states, held-out likelihood typically rose monotonically with state number and therefore could not be used to pick a final state count. This is similar to MoSeq, where likelihood increases monotonically with the number of syllables. We therefore turned to cluster stability – which measures the consistency of state assignments across independent model fits. Our reasoning was that under-counting states would force shMoSeq to arbitrarily merge different behaviors, and this might occur differently on different model runs. Similarly, over-counting states would force shMoSeq to arbitrarily split behaviors, which again might lead to instability across model runs. Stability was calculated as follows. For each candidate number of states, we fit 5 to 10 independent models at the optimal stickiness and then calculated pairwise similarity (adjusted rand index¹²¹) between the resulting state sequences. A consensus sequence was selected based on average similarity to other sequences in the ensemble, and its average similarity was reported as the “cluster stability.” We selected the consensus sequence with greatest cluster stability for downstream analysis (Figure S2F). After selecting a final model, we excluded states that occupied fewer than 1% of all timepoints. States were named *post hoc* based on inspection of videos and analysis of syllable usage.

It is important to note that for most datasets, there will be multiple valid choices for the number of states. Cluster stability provides one heuristic, but other considerations may be appropriate given the specific biological question. It may be a useful exercise to vary the number of states and observe how behaviors are split and merged as a result. For example, in our initial open field dataset (for which 5 was the optimal number of states), a 4-state model merged social engagement with exploratory locomotion, whereas a 6-state model split social engagement into two states with different velocity distributions (Figures S2H and S2I). Importantly, however, the timing of state transitions was very consistent across a wide range of state counts (Figure S2J), which suggests that the underlying notion of long-lived behavioral states does not depend on any specific level of coarse-graining.

Breadth of neural tuning

Lifetime sparseness

Lifetime sparseness (Figure S4B), which measured the breadth of neural tuning⁵⁶ was calculated as follows, where r_{ij} denotes the average activity of neuron i during timepoints assigned to state j ($1 \leq j \leq m$):

$$\text{sparseness(neuron } i) = 1 - \frac{\text{mean}(r_{i1}, \dots, r_{im})^2}{\text{mean}(r_{i1}^2, \dots, r_{im}^2)}$$

State decoding from top- N sparsest or least sparse neurons

To test the relative contribution of sparsely tuned versus broadly tuned neurons to behavioral state decoding, we ranked neurons by sparseness and then computed decoding accuracy using only the top N most sparse neurons or bottom N least sparse neurons for varying N (Figures S4C and S4D). In each of these cases, neurons were re-ranked based on the specific training data from each step of k-fold cross-validation to prevent leakage between the training and testing datasets.

State decoding from top-N most informative neurons

To test the relative contribution of highly informative versus weakly informative neurons to behavioral state decoding, we ranked neurons based on the MI between neural activity and behavioral state and then computed decoding accuracy using only the top N neurons for varying N (Figure S4D). To calculate MI for a given neuron, its activity was discretized into quartiles. As above, MI was recalculated based on the specific training data from each step of k-fold cross-validation to prevent leakage between the training and testing datasets.

Decoding models

Decoding analyses were performed using either ridge regression (for continuous variables) or logistic regression (for discrete variables) with Z-scored neural activity traces as input. Models were implemented using scikit-learn version 1.2.2.¹²² Measurements of accuracy were performed using 5-fold cross validation. To prevent leakage between neighboring frames, recordings were split into non-overlapping blocks of 1-5 minutes each and each block was randomly assigned to one of five groups; this yielded a 5-way partition of the recording that was used for cross-validation. For continuous variables, accuracy was quantified as the Pearson correlation between true and predicted values of the decoded variable. For discrete variables, accuracy was quantified using area under the receiver operating characteristic curve (auROC).

Analysis of single-neuron behavior associations

Significant associations between individual neurons and behavior variables were computed using an auROC-based approach.^{40,123} To compute auROC, we treated behavior as a binary “outcome” and neural activity as a continuous “predictor”. To determine statistical significance, the auROC was recalculated across 1000 cyclic permutations. P-values were computed relative to this simulated null distribution and then transformed using the Benjamini-Hochberg procedure to control the false discovery rate.

Neural encoding models

To understand if neurons contain information about states that is independent of instantaneous kinematics, we constructed encoder models that predicted neural activity from kinematics, syllables, and states or just kinematics and syllables (Figures 3H, 3I, and S5C–S5E). Prediction was performed using ridge regression. The kinematic variables included height, centroid velocity, and angular velocity. Syllables were encoded one-hot. States were represented by their marginal probabilities. Thus the final design matrices had dimension (3 + num syllables) or (3 + num syllables + num states). In one case (Figures S5F and S5G), we also included four world-centric variables in both encoders (distance to wall, egocentric direction to wall, proximity of conspecific, egocentric angle to conspecific). Accuracy was assessed using the correlation between true and predicted neural activity. A neuron’s accuracy was considered significant when it exceeded the 95th percentile of a shuffle distribution generated via cyclic permutation. Neurons were only included in downstream plotting and analysis if predictions from both encoder models were significant. For the recording- and mouse-level comparisons in Figure S5E, we only included recordings where at least 10 neurons met this criterion.

Neural activity manifolds

To generate low-dimensional representations of neural activity, we binned neural activity in non-overlapping half-second intervals, normalized with a rolling Z-score (using a 20 minute window) and then performed PCA to 5 dimensions. Projection into 2D was then performed using Uniform Manifold Approximation and Projection (UMAP).¹¹¹ To visualize behavioral features in the 2D map while minimizing the impact of crowding/plot-order, we superimposed a square lattice and generated a heatmap by averaging the values within each lattice cell. To quantify neural segregation of different behavioral partitions, we computed the modularity¹²⁴ of this partition with respect to the k-nearest-neighbor graph of neural activity 5D PCA space. Specifically, for each timepoint we counted the proportion of its 500 nearest neighbors that shared the same label and then computed the average of this proportion across all timepoints. We repeated this procedure for a shuffled partition and finally defined a “segregation index” as the difference between real and shuffled values. To minimize the effect of autocorrelation in the neural signal, we restricted analysis to neighbors that were far apart in time. Specifically, each recording was split into two subsets (“A” and “B”) consisting of alternating 10 min blocks separated by one-minute gaps. For each timepoint in “A”, we only considered nearest neighbors from “B” (and vice versa). This ensured that all edges linked timepoints separated by at least one minute.

Mapping behavioral states to neural principal components

To determine whether behavioral states constitute the dominant axes of neural activity, we compared the neural activity vectors associated with each state to the top-ranking neural principal components (PCs) (Figures S4G–S4I). State-associated vectors were defined through least squares regression as follows. Let $Z_{n,t}$ represent the activity of neuron n at timepoint t , and let $P_{i,t}$ represent the marginal probability of state i and timepoint t . The vector associated with state i was defined by

$$\hat{v}_i = \operatorname{argmin}_v \sum_t \left(P_{i,t} - \sum_n Z_{n,t} v_n \right)^2$$

Analysis of stereotypy in neural trajectories

To determine if neurons were preferentially active during a certain phase of a behavioral state (Figure S4K), we used the following time-warping procedure. For each neuron that was tuned to a particular state, we collected all instances of the state between 5 and 60 seconds in duration and placed them on a common timeline using linear time-warping. This yielded a matrix $Z_{i,t}$ of neural activity across instances $i = 1, \dots, N$ and warped timepoints t . If a neuron is preferentially active at particular times within each state instance, then the column-wise mean of this matrix should significantly deviate from a flat line. To assess this possibility, we computed the following test statistic for real data and for shuffle data in which the rows of $Z_{i,t}$ were cyclically permuted.

$$\text{test statistic} = \sum_t \left| \sum_i Z_{i,t} \right|$$

This yielded a P-value for each neuron/state pair. The P-values were then transformed using the Benjamini-Hochberg procedure to control the false discovery rate.

Analysis of relative wall position during open field exploration

Definition of relative wall angle and direction

At each timepoint, we calculated a “wall-vector”, defined as the vector that points toward the mouse’s centroid and originates at the nearest point on the wall. We then defined the “egocentric wall angle” as the angle between the wall-vector and mouse’s heading vector; this angle was 0° when the mouse pointed away from the wall, 180° when the mouse pointed toward the wall, and $\pm 90^\circ$ when the mouse was parallel to the wall. Relative “wall direction” was then defined as the cosine of the egocentric wall angle.

Wall distance and proximity

To capture how close the mouse was to the wall, we defined two metrics: “wall distance” and “wall proximity”. Wall distance was defined as the distance from the mouse’s centroid to the nearest point on the wall. “Wall proximity” was a function of wall distance:

$$\text{wall proximity} = -\tanh\left(\frac{\text{wall distance} - d_0}{d_0}\right)$$

where d_0 was 2.5 cm (Figure S6B); this function was designed to magnify variation near 0 and disregard fluctuations in wall distance when the mouse was far from the wall.

Definition of direction-tuned and proximity-tuned neurons

Neurons tuned to wall proximity or wall direction were identified using the auROC-based method described in “[analysis of single-neuron behavior associations](#)”. This entailed defining binary variables and testing whether neurons were more or less active than chance during the on-times for each variable. Wall proximity was represented as a binary variable by thresholding wall distance at 5 cm. Wall direction was represented as a pair of binary variables by thresholding below -0.8 or above 0.8.

Analysis of aggression experiments

Annotation of aggressive episodes

We used manual labeling to identify instances of aggression in behavioral videos. Labeling was performed on a frame-by-frame basis. Frames were included if they contained any of the following behaviors. Aggression bouts were defined as continuous blocks, ignoring gaps of less than 500ms.

- **Chasing:** One mouse actively pursues another at high speed.
- **Pouncing:** One mouse lunges or leaps towards another.
- **Mounting:** One mouse climbs onto the back of another.
- **Tumbling:** Mice become entangled and roll or tumble together at high speed.
- **Boxing:** Mice rear and face each other, usually with one in a defensive posture.

Visualization of pre- and post-aggression behavior state probabilities

For visualization of pre-, mid-, and post-aggression behavioral state probabilities (Figure 5B), we included all aggression bouts where no aggression occurred in the preceding or succeeding 15 seconds respectively. Since aggression bouts have variable duration, mid-aggression state probabilities were linearly time-warped to a common duration of 10 s.

Regressing neural activity against wall direction and proximity (aggression experiment)

When assessing whether neural responses to wall direction or proximity were altered by the presence of an aggressor (Figures 5F–5I), we tried to minimize the effect of correlations between behavioral variables (e.g. mice tended to be close to the wall and parallel to the wall at the same times). Specifically, we isolated the effects of distance to the wall by restricting to timepoints when mice were already parallel to it ($|\cos(\text{egocentric wall angle})| \geq 0.9$) and isolated the effects of angle to the wall by restricting to timepoints when mice were already close to it (within 2 cm). Furthermore, to emphasize state-dependent (rather than motor-dependent) responses, we excluded from analysis all timepoints when the subject mouse was actively being attacked.

Wall-related activity in each behavior state (aggression experiment)

To assess wall-related activity in each behavioral state (Figures 5D and 5E), we restricted to timepoints when mice were near the wall (within 2 cm) and parallel to it ($|\cos(\text{egocentric wall angle})| \geq 0.9$) and excluded timepoints when mice were being actively attacked. Mean activity of wall-direction-tuned neurons was calculated using the subset of neurons that preferred the current wall direction at each timepoint.

Analysis of novel object experiments

Object tracking

To track object bounding boxes, we trained a yolov5 object detection neural network.¹²⁵ Post processing of object detections varied depending on the dataset type. For videos where every object was unique, we: (1) filtered out detections with confidence below 0.7; (2) denoised the remaining detections using a 4-second median filter; (3) performed linear interpolation to fill in missing detections. For videos where all the objects were identical, we: (1) filtered out detections with confidence below 0.5; (2) filtered out frames where the number of detections differed from the true number of objects; (3) tracked object identities over time by maximizing pairwise intersection-over-union with the Hungarian algorithm; (4) performed linear interpolation to fill in missing detections. In both cases, we defined object locations by the centroids of the resulting bounding boxes.

Definition of object-associated syllables

To identify syllables that occurred with greater than chance frequency in the vicinity of objects (Figure S3I), we calculated – for each syllable – what fraction of instances occurred while the mouse's nose was within 6 cm of the object centroid. Syllables were deemed significant if this fraction exceeded the 95th percentile of a shuffle distribution generated through cyclic permutation of the mouse's tracking data. "Object-associated syllables" were defined as the set of 5 syllables exceeding this significance threshold, with the exclusion of one syllable that consisted entirely of grooming.

Definition of object distance and proximity

We defined distance to an object as the planar distance between the x/y coordinates of the mouse's nose and the x/y coordinates of the object centroid. When multiple objects were present, we used the distance to the closest object. "Object proximity" was a function of object distance:

$$\text{object proximity} = -\tanh\left(\frac{\text{object distance} - d_0}{d_0}\right)$$

where d_0 was 2.5 cm (Figure S5B); this function was designed to magnify variation near 0 and disregard fluctuations in object distance when the mouse was far from an object. Because tanh has an asymptote at -1, we set object proximity = -1 when no objects were present.

Neural activity related to object proximity

Neurons tuned to object proximity were identified using the auROC-based method described in "analysis of single-neuron behavior associations". As the target variable, we used a binary indicator that was 1 when any object was within 3 cm and 0 otherwise. When regressing neural activity against object proximity (Figures S6N and S6O), we excluded recordings with fewer than 5 object-tuned neurons and state/recording pairs that contained very few object interactions (i.e. where object distance fell below 6 cm in fewer than 10% of frames). The number of recordings that passed this filter varied from 8 to 16 depending on the behavioral state.

Decoding of object proximity

When comparing the accuracy of object proximity decoding across states (Figure 4M), we only included recordings where object proximity could be decoded with above chance accuracy when assessed using all frames. Specifically, the correlation between true and predicted object proximity had to exceed the 95th percentile of a shuffle distribution (N=17 recordings satisfied this criterion).

Decoding of object identity

When training and testing decoding models for object identity (Figures 4N and 4O), we restricted to timepoints when mice were within threshold distance to one and only object, and also facing that object. The distance threshold varied from 1 cm to 6 cm; Figure 4N shows results for 1 cm, and Figure 4O shows results across the full range of thresholds. We classified a mouse as 'facing' a target object if it deviated from the mouse's heading angle by less than 30 degrees. Analysis was restricted to the first three objects added to the arena, since the last two objects were absent for most of the recording.

Analysis of social engagement state

Our initial open field dmPFC recordings were 80 minutes long and contained two 30-minute blocks when a C57 conspecific was present. All the analyses below were restricted to these 60 minutes.

Mutual information between subject and conspecific syllables

To characterize the degree of social interaction across behavioral states, we computed the MI between subject and conspecific behavior (Figure 2G). To mitigate the effects of sparsity, behavior was represented using coarse-grained syllable labels (N=8 clusters, defined as described in "kinematic clustering of syllables"). We used bootstrap resampling to quantify uncertainty (N=200 bootstrap samples). For each resampling round, recordings were sampled with replacement, concatenated, and then filtered by each

behavioral state respectively for calculation of MI. The same 8-way clustering of syllables was used to calculate the frequency with which mice performed kinematically similar behaviors at the same time (Figure S3B).

Social distance and proximity

Because subject mice primarily interacted in a ‘nose-first’ manner, social distance was defined as the distance from the nose of the subject mouse to the centroid of the conspecific mouse. “Social proximity” was a function of social distance:

$$\text{social proximity} = -\tanh\left(\frac{\text{social distance} - d_0}{d_0}\right)$$

where d_0 was 10 cm. Neurons tuned to social proximity were identified using the auROC-based method described in “[analysis of single-neuron behavior associations](#)”. As the target variable, we used a binary indicator that was 1 when the social distance was ≤ 5 cm and 0 otherwise.

Time-dependent changes in neural activity

To capture temporal changes in the strength of activity elicited by social contact (Figure S6W), we divided each recording into 2-minute bins that tiled the first 10 minutes of each social interaction block. The bins were overlapping and spaced at 30s intervals. Within each bin, we calculated the average activity of social proximity-tuned neurons during “close” timepoints (social distance ≤ 5 cm) and “far” timepoints (social distance ≥ 20 cm). Bins with no “close” or “far” timepoints were counted as missing data. Recordings were only included if the number of social proximity-tuned neurons was ≥ 10 . Since some decline in activity is expected due to photobleaching, we also generated a baseline curve (Figure S6X) by calculating the 90th percentile of activity across neurons at each timepoint and then reporting its average within each 2-minute bin.

Regressing neural activity against distance to conspecific

When regressing the activity of social proximity-tuned neurons against distance to the conspecific (Figures 4I and 4J), we restricted analysis to recordings with at least 10 social proximity-tuned neurons. Additionally, for each behavioral state, we only included recordings where the state occurred for a total of 10s or more. To emphasize state-dependent (rather than trivial motor- or sensory-dependent) effects, we excluded bouts of aggression and restricted to timepoints when the subject mouse was facing the conspecific.

Tuning to allocentric spatial variables

To determine whether neurons were tuned to absolute location (Figures 4A and S6A), we divided the circular open field into 8 sectors (like slices of a pie) and then identified neurons that were significantly active in one slice using the auROC-based method described in “[analysis of single-neuron behavior associations](#)”. To determine whether neurons were tuned to absolute heading, we similarly discretized the mouse’s heading angle into 8 bins and then identified neurons that were significantly active in one bin.

Analysis of tube test dataset

Neural activity and behavior annotations from Kingsbury et al.⁴⁰ were obtained upon request to the authors. The dataset consisted of six sessions, each with multiple tube test trials per session. Each session included simultaneous behavior and neural recordings from both mice. Three behaviors were annotated: “push”, “retreat”, and “approach”. We treated each mouse/session pair as an independent recording and labeled behavior as “self” and “other” accordingly. All analyses were limited to session/behavior pairs where the behavior decoding had above chance accuracy ($P < 0.05$); this was assessed by computing an auROC for the true predictions and 100 cyclic permutations. Cross-correlations were between decoded behavior probabilities and binary indicator variables representing each behavior.

We used kernel regression to characterize neural activity associated with onset versus offset of the “self-push” behavior (Figure 6C). Decoding models were first trained to predict the probability of “self-push” from heldout neural activity. The decoded behavior probabilities were then Z-scored and modeled as follows

$$P_{\text{self-push}}(t) = \sum_{ij} \alpha_j K_j(t - t_i^{\text{onset}}) + \sum_{ij} \beta_j K_j(t - t_i^{\text{offset}}) + \text{error}$$

Where $[t_i^{\text{onset}}, t_i^{\text{offset}}]$ represent on/off times for the i ’th “self-push” bout and α_j / β_j are onset/offset coefficients for the j ’th kernel basis function. Basis functions were defined by $K(t) = \tanh((t - t_0)\tau)$ for a range of shifts ($-0.5s \leq t_0 \leq 0.5s$) and slopes ($-0.5 \leq \tau \leq 0.5$). Parameters were fit using Ridge regression.

Analysis of DLS lesion experiments

Timescale analysis of DLS lesion data (Figure 7J) was performed using the procedure described in “Mutual information analysis of behavioral timescales”. Data (including syllable annotations) were obtained from Markowitz et al.⁴⁸ We restricted analysis to recordings in an empty open arena, i.e. without a trimethylthiazoline (TMT) or control odor source.

Comparative analysis of egocentric boundary vector tuning

We analyzed egocentric boundary representations in retrosplenial cortex (RSC) using data kindly provided by the authors of Alexander et al.,⁶⁸ which included rat trajectories, headings, boundary coordinates, spike times, and unit annotations. The data included

117 recordings with 555 units total, of which 134 were classified as egocentric boundary vector (EBV) cells. Analysis was restricted to wall-tuned neurons in dmPFC and EBV cells in RSC.⁶⁸

We assessed neural tuning using egocentric boundary ratemaps (EBRs), which capture responses to boundaries at an array of locations in egocentric coordinates (Figure 4E). Each map was parameterized as a grid in polar coordinates, and the value at each grid location represented the expected activity of a neuron given that a boundary was present at that particular distance and angle relative to the animal (0° = in front, 90° = left, 180° = behind, 270° = right). The grids were spaced with angle increments of 3° and distance increments of either 2.5 cm for RSC or 1 cm dmPFC (finer spacing for dmPFC reflects the smaller size of mice relative to rats). We only considered boundary points within a threshold distance equal to one-quarter the width of the area (30 cm for the RSC dataset, 10 cm for the dmPFC dataset). EBRs were smoothed with a two-dimensional Gaussian kernel ($\sigma = 5$ bins, truncation = 0.5σ) prior to downstream analysis.

Preferred boundary distance (Figure 4F) was defined as the distance with the highest average activity (i.e. the peak in the EBR). Angular tuning was quantified using the mean resultant length (MRL), computed as the magnitude of the vector sum across distance/angle bins weighted by their EBR values. Significance was assessed using a shuffle distribution of MRLs obtained by cyclically permuting neural activity; observed MRLs were Z-scored against this null distribution (Figure 4E).

QUANTIFICATION AND STATISTICAL ANALYSIS

All statistical analyses were conducted with Python using scikit-learn, statsmodels, and NumPy/SciPy. Statistical tests, P-values, measures of uncertainty, and sample sizes (N) for each panel are reported in the corresponding figure legends. Unless noted in the legends, the independent unit of analysis (N) is a recording/session. All hypothesis tests were two-sided with $\alpha=0.05$ unless stated otherwise. Mann-Whitney U tests were used to compare the means of populations. Multiple comparisons were controlled using the Benjamini–Hochberg false discovery rate (FDR) procedure where families of tests are performed (e.g., per-neuron screening or per-state comparisons); the legend specifies when FDR is applied. In all boxplots, the box and center line correspond to the median and interquartile interval respectively.

ADDITIONAL RESOURCES

ShMoSeq documentation page with installation instructions, example datasets, and step-by-step tutorials for fitting models and selecting hyperparameters: <https://state-moseq.readthedocs.io>

Supplemental information

Spontaneous behavior is a succession of self-directed tasks

Caleb Weinreb, Lakshanya Thamarai Kannan, Alia Newman-Boule, Tim Sainburg, Winthrop F. Gillis, Alex Plotnikoff, Sofia Makowska, Jonah E. Pearl, Mohammed Abdal Monium Osman, Scott W. Linderman, and Sandeep Robert Datta

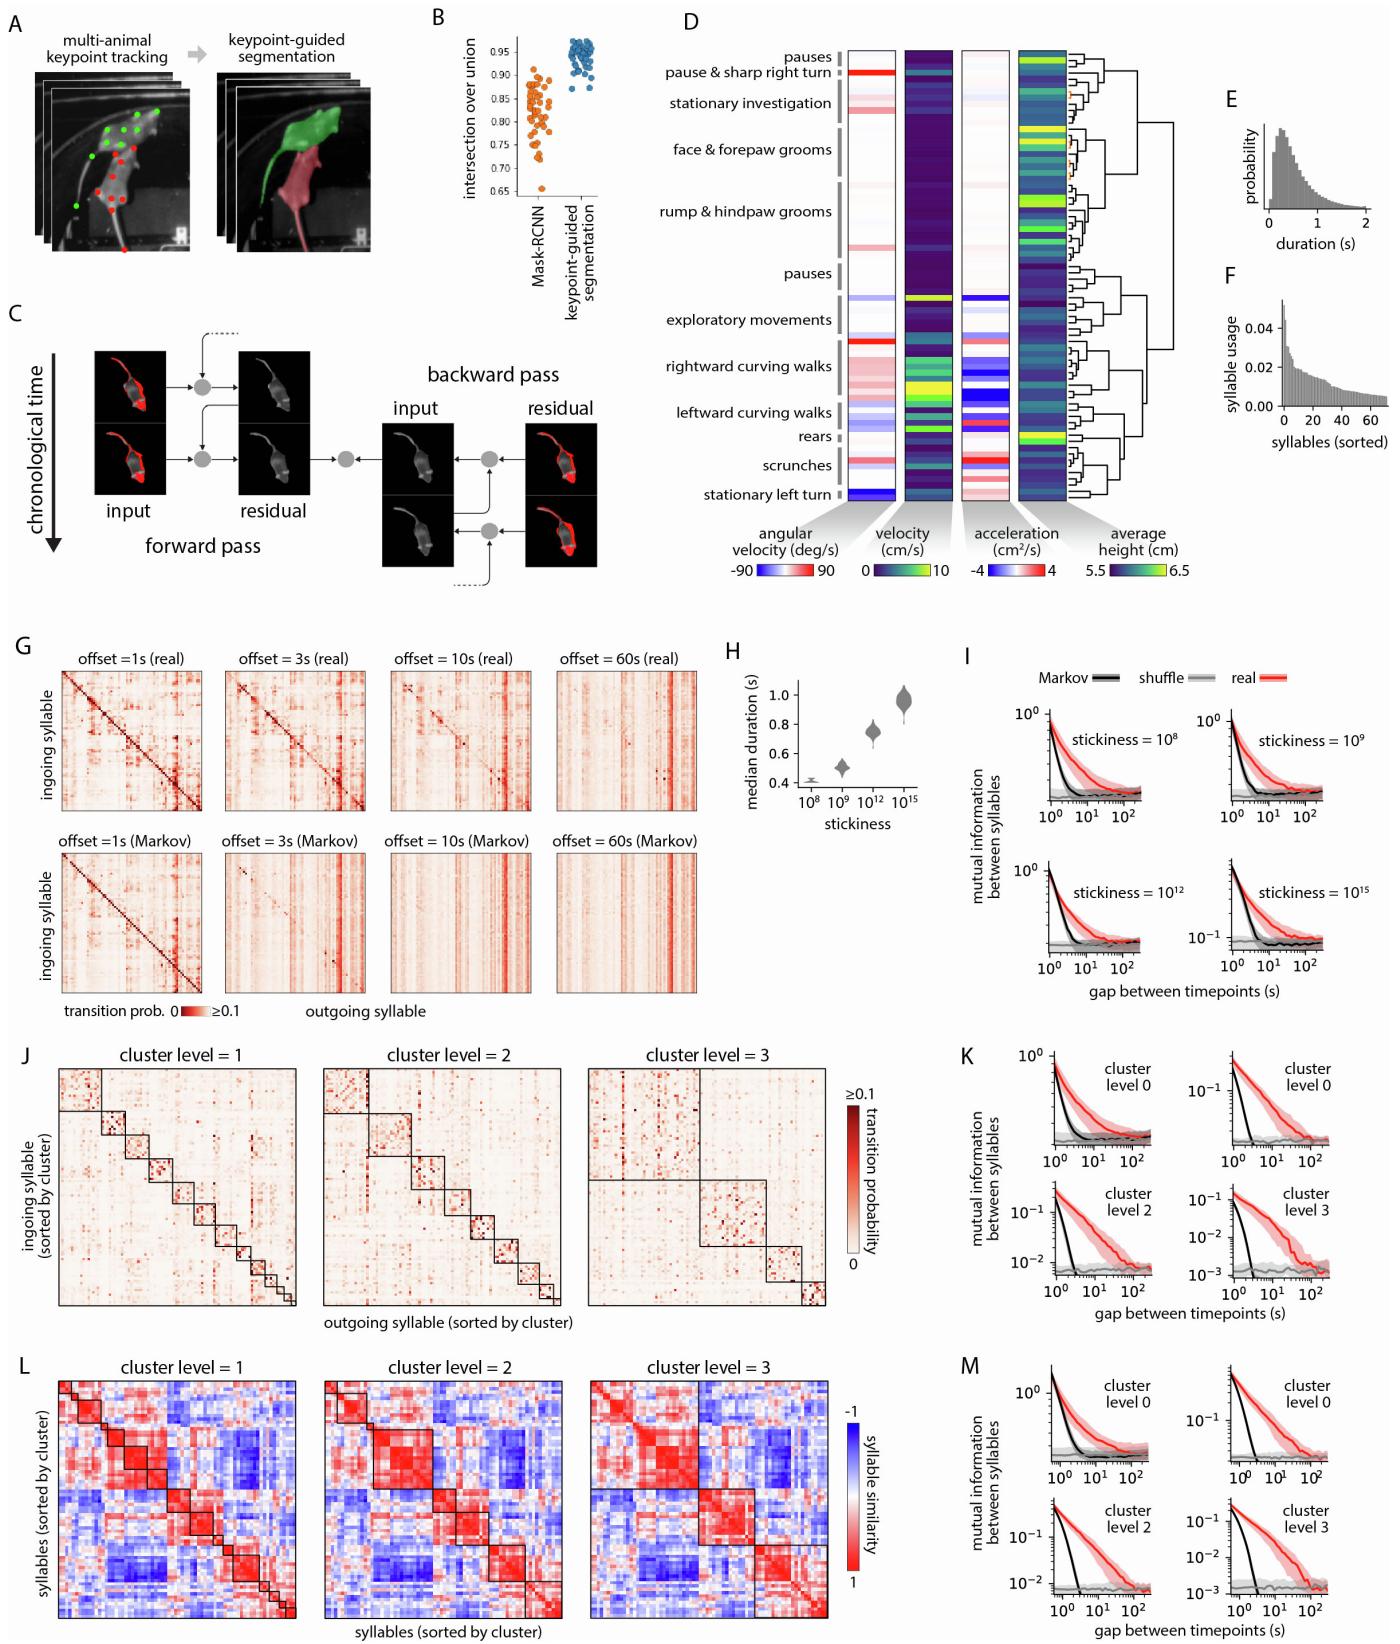


Figure S1: MoSeq analysis reveals long-range structure in behavior, related to Figure 1.

(A) Illustration of keypoint-guided segmentation.

(B) Comparison between our keypoint-guided segmentation and Mask-RCNN, a state-of-the-art algorithm for instance segmentation of images. Accuracy is quantified as intersection-over-union between predicted and human annotated masks in a set of held out video frames. Each dot represents a recording.

- (C) Illustration of the pipeline for missing data imputation. Only the top-camera infrared (IR) channel is shown for simplicity. Red pixels denote missing data.
- (D) Overview of MoSeq syllables derived from solitary and social open field recordings with a C57 conspecific. Heatmaps show averages for various kinematic parameters during each syllable.
- (E) Distribution of syllable durations.
- (F) Overall usage of each syllable (sorted for ease of visualization).
- (G) Syllable transition matrices at a range of temporal offsets, shown for real data (top) and for simulated Markov chains with the same transition probabilities (bottom)
- (H) Median syllable durations (shown as distributions across recordings) for MoSeq models fit with a range of stickiness parameters.
- (I) Mutual information (MI) between syllables at a range of temporal lags, shown for real syllable sequences (red), synthetic sequences drawn from a Markov model with matched transition probabilities (black), and a shuffle noise floor (gray). Panels correspond to MoSeq models with different stickiness parameters, and therefore different syllable durations. Line and shading show the median and interquartile interval across recordings ($N = 31$).
- (J) Hierarchical clustering to test whether non-Markovian dynamics depend on the degree to which syllables are coarse-grained. In this plot, syllables were clustered based on their transition probabilities, i.e. to maximize the probability of within-cluster transitions as opposed to between-cluster transitions. Results are shown for three different dendrogram cutoffs. Each heatmap shows a syllable transition matrix, with syllables sorted by cluster membership. Black squares mark cluster boundaries.
- (K) Mutual information (MI) between transition-based syllable clusters at a range of temporal lags, plotted as in (I). Each panel corresponds to a different level of clustering as shown in (J); level 0 corresponds to unclustered syllables.
- (L) Hierarchical clustering of syllables based on kinematic similarity, as in (J).
- (M) Mutual information (MI) between kinematics-based syllable clusters at a range of temporal lags, plotted as in (I). Each panel corresponds to a different level of clustering as shown in (L), level 0 corresponds to unclustered syllables.

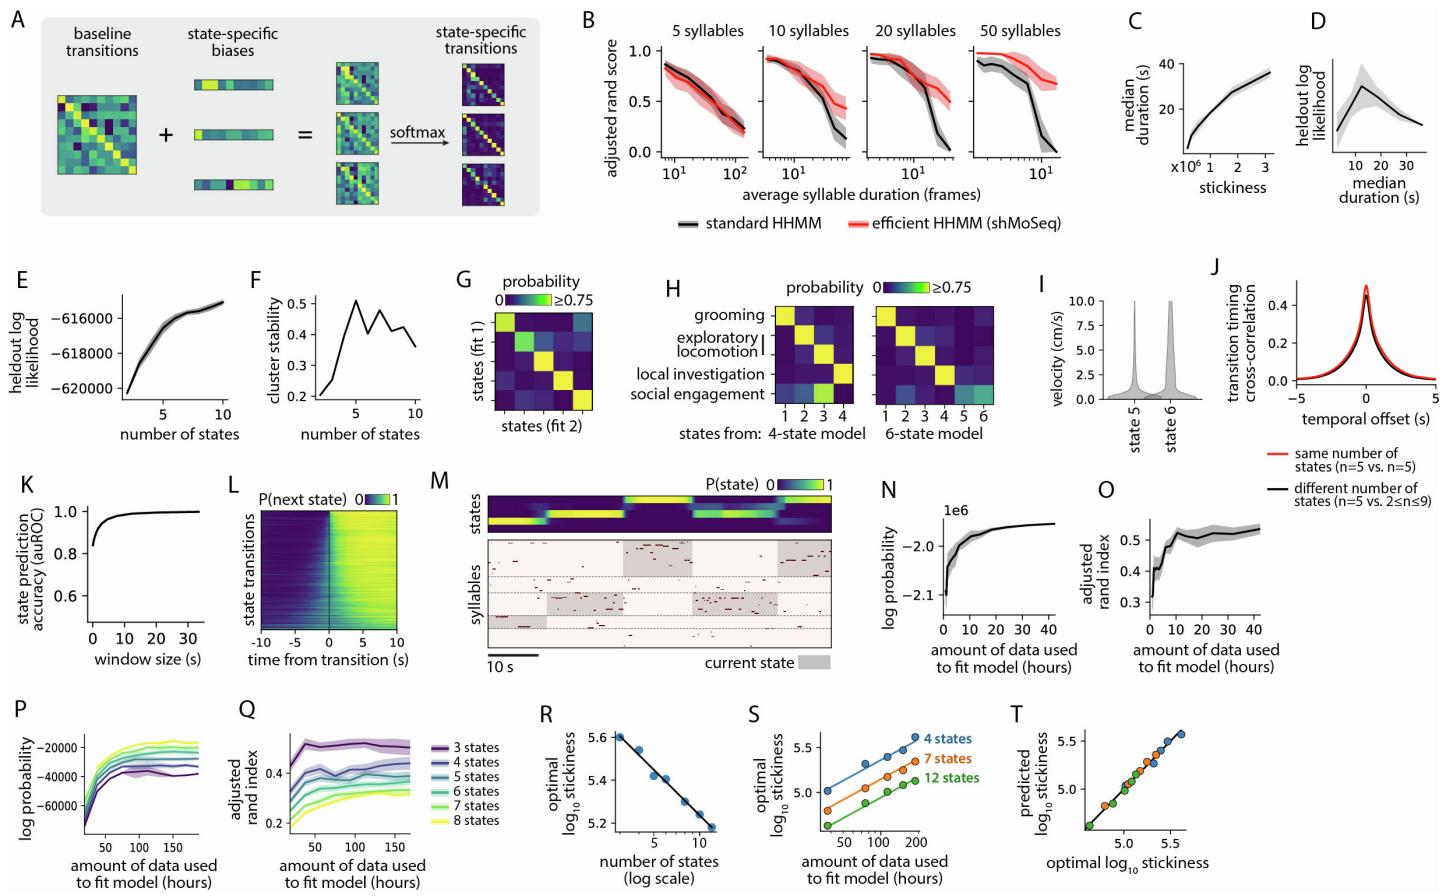


Figure S2: ShMoSeq captures long-range structure in behavior, related to Figure 1.

(A) Schematic showing how shMoSeq combines state-specific syllable biases with a matrix of baseline transition rates to generate the syllable transition matrix for each state.

(B) Performance of standard hierarchical hidden Markov models (HHMMs) versus data efficient HHMMs (as implemented in shMoSeq) when applied to simulated data. Performance is assessed using the adjusted rand score between true and inferred high-level state sequences. Results are shown for a range of syllable numbers and durations. Line and shading show mean and 95% CI across 10 simulated datasets per condition.

(C) Median duration of states output by shMoSeq for different values of the stickiness hyperparameter. Durations are defined as the time intervals between consecutive state transitions. Black line and shading show mean and standard deviation across independent model fits ($N = 10$ fits).

(D) Goodness-of-fit (measured by held-out log likelihood) as a function of median state duration, plotted as in (C).

(E) Held-out log likelihood as a function of number of states, plotted as in (C). The monotonic increase implies that likelihood is less useful as a criterion for choosing the number of states. A similar phenomenon is observed for MoSeq (likelihood increases with the number of syllables).

(F) Consistency of state assignments across independent model fits (i.e. “cluster stability”) as a function of the number of states. Stability peaks at $N=5$ states, which we used for downstream analysis of this dataset.

(G) Confusion matrix showing agreement between independent shMoSeq fits with $N=5$ states. Density along the diagonal is consistent with high cluster stability, as shown in (F).

(H) Confusion matrices comparing state assignments from a 5-state model to those from a 4-state model (left) and 6-state model (right). Social engagement is split into two different states (5 and 6) in the 6-state model, whereas it is merged with exploratory locomotion in the 4-state model.

(I) Velocity distributions for states 5 and 6 from the 6-state model (which correspond to social engagement from the 5-state model). State-6 captures high velocity behaviors.

(J) Cross-correlation in transition timing between models with the same number of states versus those with different numbers of states. We fit 10 independent models per number of states. For each model, we computed per-timepoint transition probabilities by sampling an ensemble of state sequences; these probability time-series were used for calculating cross-correlation. Each line shows the average across all pairwise comparisons with the indicated number of states. The fact that the lines are close together implies that models with different numbers of states still agree on transition timing.

- (K) Classification accuracy when behavioral states are predicted from syllable frequencies in an N-second window, shown here as a function of N. Accuracy is reported using area under the receiver operating curve. Chance accuracy is shown in gray.
- (L) Uncertainty in the timing of state onset. Each row corresponds to a single state transition and shows the marginal probability of the next state.
- (M) Example interval showing uncertainty in transition timing. Top: marginal probability of each behavioral state. Bottom: MoSeq syllables sorted according to the state in which they occur most often. Gray boxes indicate the final state sequence output by shMoSeq.
- (N) Log probability of models that were fit to random subsets of the available data and then applied to the full dataset. Line and shading represent mean and standard deviation across 10 independent model fits. All models shared the same median state duration and number of states (N = 5).
- (O) As (N), but here showing similarity between models fit to subsets of the data and a final model that was fit to the full dataset. Similarity was defined as the adjusted rand index between the models' respective state sequences. The elbow at 10 hours suggests that ≥ 10 hours of data are needed for shMoSeq.
- (P) Modeling results for a separate dataset of open field behavior (Markowitz, 2023), plotted as in (N). Each color represents a different number of states used for fitting.
- (Q) As (P), but here showing similarity between models fit to subsets of the data and a final model that was fit to the full dataset. The plateau around 40 hours [also visible in (P)] suggests that ≥ 40 hours of data is ideal for shMoSeq, although results may continue to improve beyond this point when the state number is high (i.e. ≥ 6 states).
- (R) Optimal value of the stickiness parameter (i.e. value yielding the highest model probability) for models with different numbers of states. All models were fit to the full dataset from (Markowitz, 2023).
- (S) Optimal value of the stickiness parameter for models fit to different-sized subsets of the (Markowitz, 2023) dataset. Results are shown for three different numbers of states.
- (T) Output of a linear regressor that predicts the optimal value of kappa based on the number of states and amount of data used for fitting.

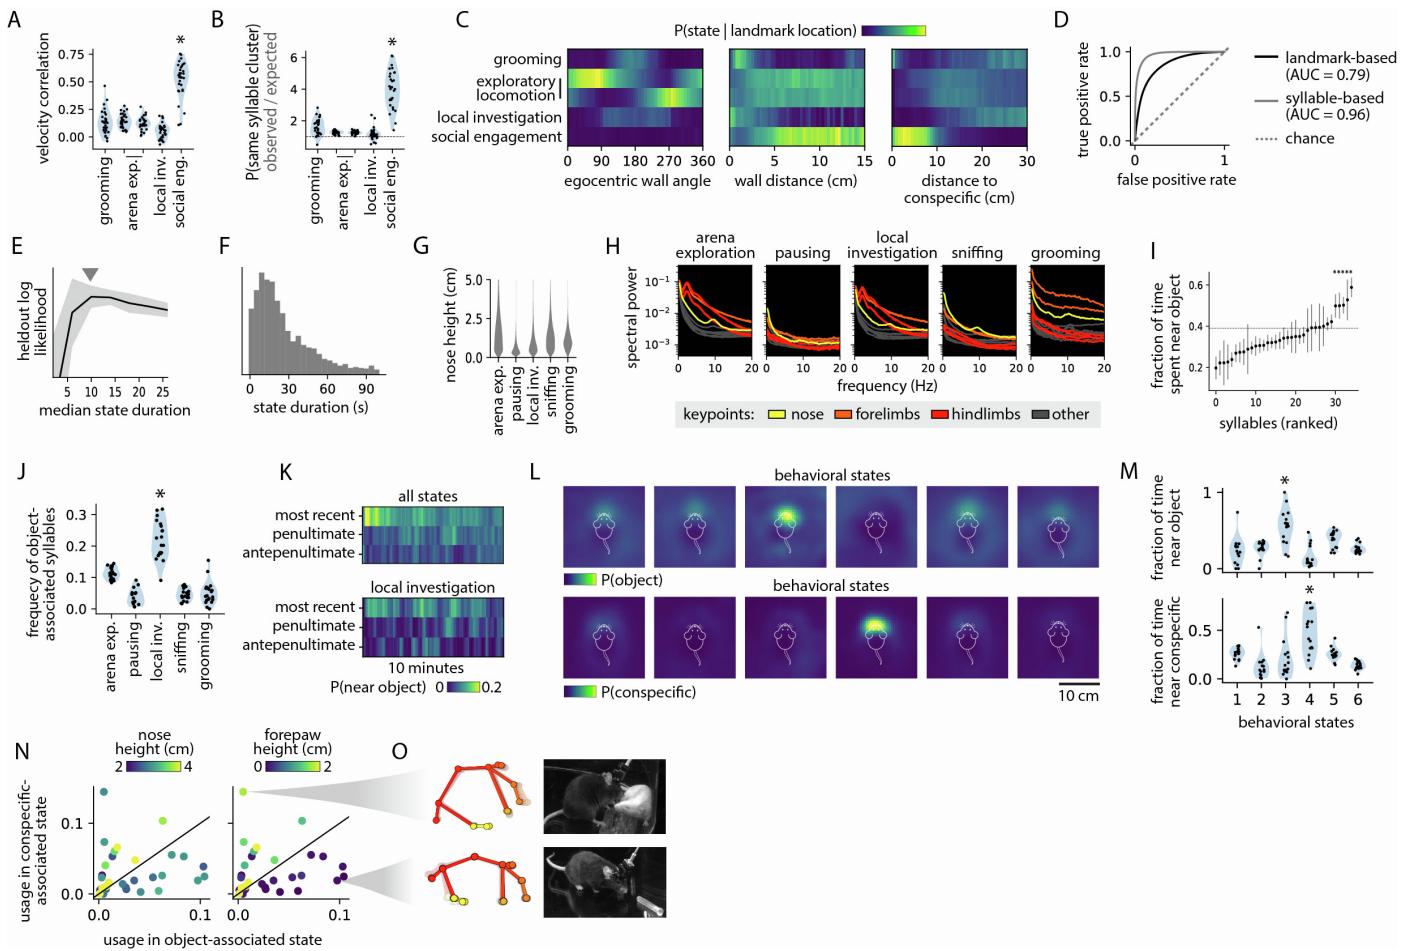


Figure S3: Affordances sculpt behavioral states, related to Figure 2.

- (A) Correlation between velocity of the subject mouse and C57 conspecific, calculated separately for each behavioral state. Each dot represents one recording ($N = 34$, $P < 1.4 \times 10^{-9}$, Mann-Whitney U test).
- (B) Frequency with which the subject and conspecific express syllables in the same kinematic cluster, plotted as in (A).
- (C) Probability of behavioral states given spatial landmarks. Heatmaps show the conditional probability of each state as a function of (left) egocentric wall angle, (middle) wall distance, and (right) distance to a conspecific. In most cases, multiple states are compatible with a given landmark location.
- (D) Receiver operating characteristic (ROC) curves comparing classifiers trained to predict behavioral states from landmark locations (black, AUC = 0.79) versus syllables (gray, AUC = 0.96). Dashed line denotes chance.
- (E) Goodness-of-fit (measured by held-out log likelihood) as a function of median state duration, shown for shMoSeq applied to the novel object experiment. Black line and shading show mean and standard deviation across independent model fits ($N = 10$ fits).
- (F) Distribution of state durations, defined as the time intervals between consecutive state transitions.
- (G) Distribution of nose heights during each behavioral state from the novel object experiments.
- (H) Power spectral density of keypoint velocities, stratified by behavioral state. The 10Hz peak for the nose keypoint represents sniffing. The 2-3Hz peak for the forelimb/hindlimb keypoints corresponds to locomotion.
- (I) Syllables ranked by fraction of time near objects (nose-to-object distance ≤ 6 cm). Dots and error bars show means and 95% CI across recordings ($N = 21$). Dashed lines mark the threshold for significance ($P < 0.05$, permutation test).
- (J) Total usage of object-associated syllables within each behavioral state, plotted as in (A) ($P < 7 \times 10^{-7}$, Mann-Whitney U test).
- (K) Top: time-varying probability of being near the most recently added object (top row) or the two previously added objects (bottom rows), where near is defined as within 3 cm. Bottom: time-varying probability of simultaneously being near an object and expressing the local investigation state.
- (L) Direct comparison between social interaction and object investigation. Mice were recorded in the novel object assay or during hour-long sessions in the presence of a female BALB/cJ conspecific. In both cases, we captured 3D keypoints and characterized behavior using shMoSeq. Heatmaps show the locations of objects (top) and conspecifics (bottom) in the subject mouse's egocentric reference frame. The heatmaps along each row correspond to the six behavioral states identified by shMoSeq.

(M) Top: fraction of time near objects (nose-to-object distance ≤ 5 cm) during each behavioral state. Bottom: fraction of time near conspecifics (nose-to-nose or nose-to-tail distance ≤ 5 cm) during each behavioral state. Each dot represents one recording ($N = 16$). State #3 is enriched near objects and state #4 is enriched near conspecifics ($P < 0.01$, Mann-Whitney U test).

(N) Usage of each syllable in the object-associated state versus the conspecific-associated state. The solid lines represent equal usage in both states. Dots correspond to syllables, with colors representing the average nose height (left) or forepaw height (right) during expression of that syllable.

(O) Left: Average keypoint trajectories for syllables shown in (N), including a conspecific-associated syllable (top) and an object-associated syllable (bottom). Right: example video frames showing each syllable respectively.

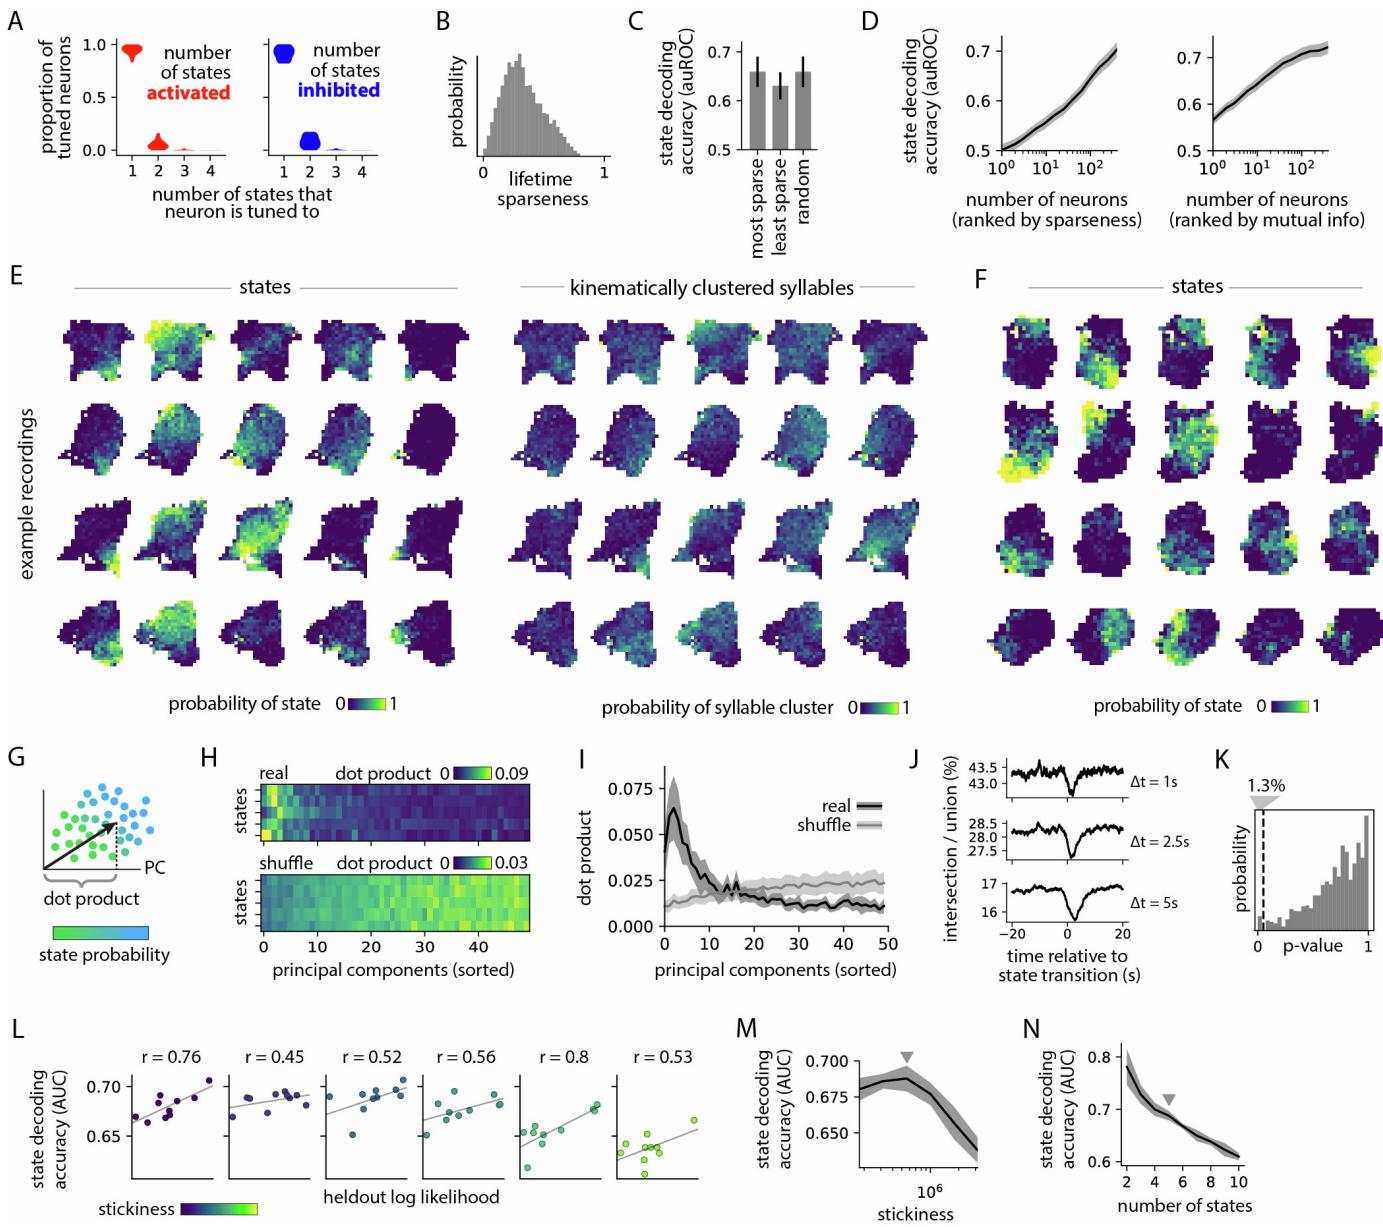


Figure S4: Neural encoding of syllables and behavioral states, related to Figure 3.

(A) Left: proportion of neurons significantly active in one, two, three or four states, normalized by the number significantly active in at least one state; plots show the distributions of these fractions across $N = 31$ recordings. Right: as (left) but here for neurons that are significantly inhibited (FDR $< 5\%$ for significantly active or inhibited neurons).

(B) Distribution of lifetime sparseness across all neurons that were significantly modulated by at least one state (1 = perfectly selective, 0 = completely non-selective).

(C) Accuracy of behavioral state decoding using the 100 sparsest neurons, the 100 least sparse neurons or 100 random neurons from each recording respectively, showing the mean and 95% CI across recordings ($N = 31$).

(D) Left: Accuracy of behavioral state decoding using either the top- N sparsest neurons (left) or the top- N most informative neurons (right) from each recording, shown as a function of N . Informativeness was defined as the mutual information (MI) between neural activity and behavioral states. Line and shading show mean and 95% CI across recordings ($N = 31$).

(E) Distributions of states (left) versus syllables (right) across 2D UMAP projections of neural activity. To ensure a fair comparison, syllables were kinematically clustered to match the number of states. Each row corresponds to a different session from the initial open field recordings with a C57 conspecific. Heatmaps depict probabilities within local regions.

(F) As (E), here showing example recordings from the novel object experiment.

- (G) Procedure for mapping behavioral state-defined axes of neural activity to neural principal components (PCs). Each timepoint is represented by a dot in neural PC space. Dots are colored by the marginal probability of a given behavioral state, and a vector shows the direction along which this probability increases (computed by regressing the probability against neural activity). Projections of this vector onto each PC yield the “dot products” shown in (H) and (I).
- (H) Dot products between each behavioral-state-defined axis (row) and each neural PC (column). Results are shown for true and cyclically permuted state probabilities.
- (I) Summary of dot products shown in (H). Line and shading represent the mean and 95% CI across recordings (N = 31).
- (J) Overlap in neural activity across a fixed time lag, showing the average across state transitions from each recording. Overlap was quantified as the number of neurons active at both timepoints (“active” defined by z-scored $\Delta F/F \geq 1$) divided by the number active at either timepoint.
- (K) Proportion of neurons that are active at a consistent timepoint between state onset and offset. Each neuron was assigned a P-value based on the specificity of its activity compared to a shuffle distribution; the distribution of these P-values is plotted with a cutoff at P = 0.05.
- (L) Accuracy of dmPFC-based state decoding models versus heldout log likelihood (which is computed by shMoSeq independent of neural data). Each dot represents an independent model fit. Colors show the stickiness parameter used during fitting. Partial correlation controls for the stickiness parameter.
- (M) State decoding accuracy as a function of the stickiness parameter (number of states was held constant at N = 5). Line and shading represent mean and 95% CI across 10 independent model fits. Arrowhead marks the optimal stickiness, which we used for downstream analysis.
- (N) State decoding accuracy as a function of the number of states [stickiness was held constant at the optimal value shown in (M)]. Line and shading as in (M). Arrowhead marks the number of states (N = 5) that was used for downstream analysis. Note that a tonic decline in decoder performance is expected when increasing the number of categories to be classified given the same dataset.

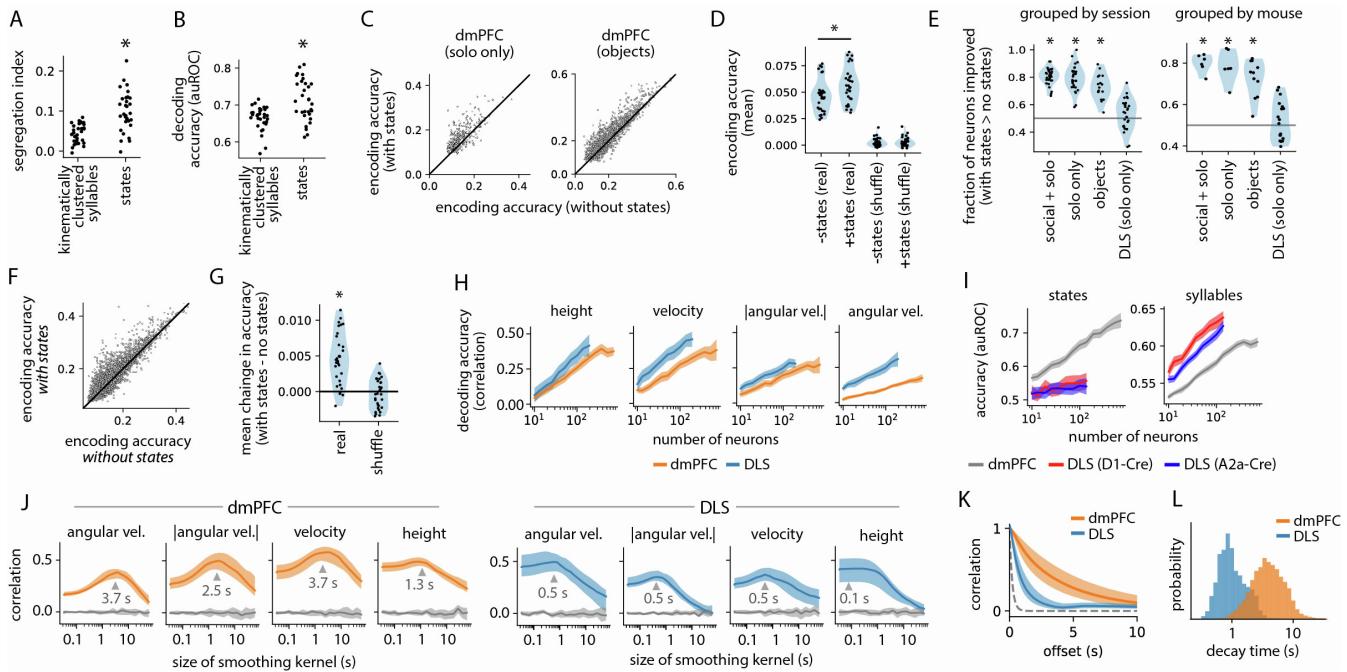


Figure S5: Encoding of affordance-related variables, related to Figure 3.

(A) Segregation in high-dimensional neural activity space (see Methods) for states and syllables respectively. To ensure a fair comparison, syllables were kinematically clustered to match the number of states. Each dot represents a recording ($N = 31$, $P < 4 \times 10^{-5}$, Mann-Whitney U test).

(B) Accuracy of decoders trained to predict states versus syllables using population activity in dmPFC. Syllables were grouped by kinematic similarity to match the number of states. Plotted as in (A) ($P = 6 \times 10^{-4}$, Mann-Whitney U test).

(C) Accuracy of encoders with access to syllables and kinematics (x-axis) versus those with access to syllables, kinematics, and states (y-axis). Left: encoding accuracies for dmPFC neurons during non-social timepoints from our initial set of open field recordings. Right: encoding accuracies for dmPFC neurons from the novel object experiment.

(D) Average encoding accuracy for models with or without access to states, shown for real and circularly permuted data. Each dot represents a recording from our initial open field recordings (all timepoints, $N = 31$). The addition of states increased accuracy for real but not shuffled data ($P = 2 \times 10^{-13}$).

(E) Left: Fraction of neurons that were better-predicted by an encoding model with access to states versus one without. Results are shown for (1) the initial open field recordings in dmPFC with a C57 conspecific (all timepoints); (2) the same experiment but restricted to timepoints with no conspecific present; (3) dmPFC recordings with novel objects; (4) DLS open field recordings. Each dot represents a recording ($21 \leq N \leq 31$). Fractions are significantly above chance in all three dmPFC datasets ($P < 2 \times 10^{-8}$, one sample t-test), but not the DLS dataset. Right: Subject-wise averages of the fractions shown on the left. Each dot represents a subject animal ($6 \leq N \leq 18$, $P < 2 \times 10^{-4}$).

(F) Accuracy of encoders with access to syllables, kinematics, and landmark locations (x-axis) versus those with access to same variables plus states (y-axis). Each dot represents a dmPFC neuron from our initial set of open field recordings.

(G) Mean difference in accuracy for encoding models with versus without access to states [as described in (F)]. Each dot represents a recording ($N = 31$). Data are shown for real and circularly permuted data. The mean change in accuracy for real data is significantly greater than zero ($P < 7 \times 10^{-8}$).

(H) Correlations between various kinematic variables and their values decoded from dmPFC or DLS respectively. Neurons from each recording were down-sampled and the results are plotted as a function of the number of neurons. Lines and shading show mean and 95% CI across down-sampled recordings ($11 \leq N \leq 31$).

(I) Accuracy of decoders trained to predict states (top) or syllables (bottom), plotted as in (I). DLS data are grouped by cell type (direct pathway: D1-Cre, indirect pathway: A2a-Cre).

(J) Timescale of kinematic encoding in dmPFC (left) and DLS (right). Each plot shows the correlation between decoded values of a kinematic variable and true values after smoothing with a Gaussian kernel. Lines and shading show mean and 95% CI across recordings. Pointers mark peak correlations. Shuffle correlations from cyclically permuted data are shown in gray.

(K) Autocorrelation of neural activity, showing median and inter-quartile interval across neurons from dmPFC and DLS respectively. The dotted line shows the decay time of the calcium indicator GCaMP6f (Chen et al., 2013).

(L) Distributions of autocorrelation decay times for dmPFC and DLS neurons respectively, defined as the time required for the autocorrelation to fall below e^{-1} .

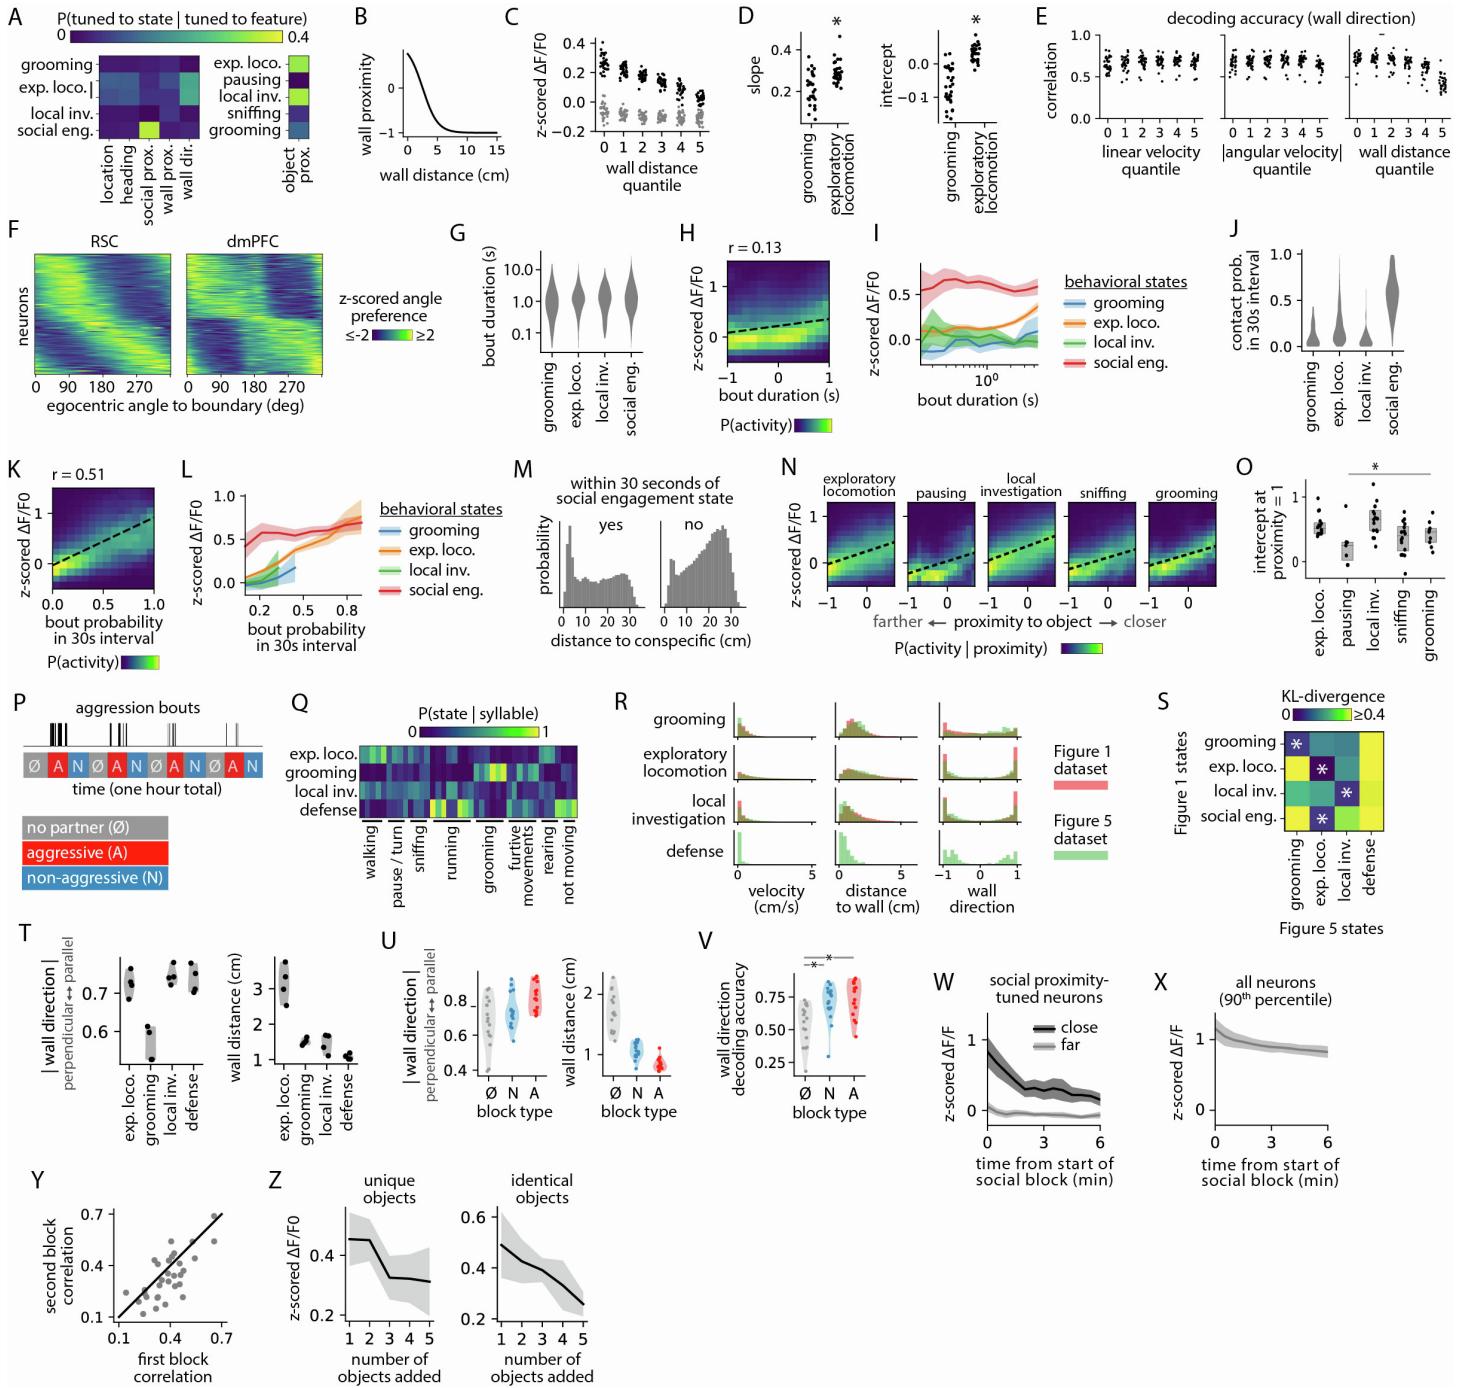


Figure S6: Encoding of affordance-related variables, related to Figures 4 and 5.

(A) Probability that neurons are tuned to each behavioral state given that they are tuned to a particular spatial or affordance-related variable (FDR < 5% for significantly tuned neurons here and throughout figure). Heatmaps are derived from the open field recordings with a C57 conspecific (left) and the novel object recordings (right) respectively.

(B) Relationship between wall distance and wall proximity. Proximity is defined by the function: proximity = $\tanh((\text{distance} - d_0)/d_0)$ where d_0 is 2.5 cm.

(C) Average activity of wall direction-tuned neurons, stratified by distance to the wall, calculated at timepoints when the wall was on each neurons' preferred side (black) or non-preferred side (gray). Each dot represents a recording (N=31).

(D) Slopes and intercepts for best fit lines as shown in (Fig 4C), calculated for each recording separately (one dot per recording, N=31, $P < 5 \times 10^{-5}$, Mann-Whitney U test).

(E) Correlation between true and decoded wall direction, stratified by linear velocity, angular velocity, and wall distance respectively. Each dot represents a recording (N=31).

(F) Neural activity as a function of egocentric boundary angle, shown for egocentric boundary vector cells in RSC (left) and wall-tuned neurons in dmPFC (right). Rows of the heatmap correspond to egocentric rate maps [as shown in (C)] that have been averaged along the distance axis, unwrapped along the angle axis, and then Z-scored. To ease visualization, we used a random subsample of dmPFC neurons when generating the right heatmap.

(G) Duration distributions for social-contact bouts from each behavioral state. Bouts were defined as continuous intervals during which the subject mouse was within 5 cm of the conspecific and facing in its direction (angle < 30°). This comparison addresses the hypothesis that sustained interaction bouts are required for socially-related activity to fully develop.

(H) Activity of socially-tuned neurons as a function of bout duration.

(I) Average activity of socially-tuned neurons across bouts of different durations, stratified by behavioral state. Shading represents 95% CI. The modest rise in activity with increasing bout duration suggests that bout duration contributes (weakly) to state-dependent differences in socially-related neural activity.

(J) Distributions of “contact probability” across bouts from each state. “Contact probability” was defined as the probability of social contact during a 30 second interval centered on a given bout. This comparison addresses the hypothesis that high-frequency bouts are required for socially-related activity to fully develop.

(K) Activity of socially-tuned neurons as a function contact probability.

(L) Activity of socially-tuned neurons across bouts with different contact probabilities, plotted in (L). The convergence of activity levels when bout probability is high suggests that bout probability contributes strongly to state-dependent differences in socially-related neural activity.

(M) Distribution of distance to conspecific for timepoints that are either within 30 seconds of the social engagement state (left) or outside this radius (right).

(N) Average activity of neurons tuned to object proximity (y-axis) as a function of object proximity (x-axis), stratified by behavioral state. Plots show the distribution of activity levels for each proximity-level. In all states, neural activity is higher when mice are closer to an object, but the slope and intercept of this relationship are higher during the object investigation state, as indicated by the dashed best-fit lines.

(O) Intercepts at distance = 0 of best-fit lines shown in (H), but now calculated separately for each recording. Each dot represents a recording with at least 5 object-proximity-tuned neurons and sufficient object interactions within a state (see Methods, 8 ≤ N ≤ 16). Boxes show medians and interquartile intervals. Intercepts are highest for the object investigation state; this difference is significant with respect to all states except locomotion (P < 0.04, Mann-Whitney U test).

(P) Timing of aggressive episodes for an example recording from the CD1-aggression dataset.

(Q) Syllables used in each behavioral state from the aggression dataset.

(R) Distributions of kinematic and spatial variables across shared behavioral states from two different datasets: the initial open field recordings with a C57 conspecific (red, explored first in Figure 1); and the aggression recordings with a CD1 conspecific (green, explored in Figure 5).

(S) Similarity of behavioral states across datasets. For each pair of states, we computed KL divergences for the three variables shown in (R) and the average-of-three in the heatmap. Stars mark the minimum KL divergence (i.e. most similar state) in each row. Uniformly high KL divergences with the defense state indicate that it is unique to the CD1 aggression dataset (featured in Figure 5).

(T) Mean direction (top) and distance (bottom) to the wall during each behavioral state from the aggression dataset. Each dot corresponds to a recording (N = 4).

(U) Median direction (top) and distance (bottom) to the wall during each block type from the aggression dataset. Each dot corresponds to one block (N = 16).

(V) Correlation between true and decoded wall direction, calculated separately for each block, shown as in (U) (P < 0.002, Mann-Whitney U test).

(W) Mean activity of neurons tuned to social proximity (from dmPFC recordings with a C57 conspecific), calculated during times when the conspecific was close (< 5 cm) or far (> 20 cm). Line and shading show mean and 95% CI across all recordings with at least 10 proximity-tuned neurons (N = 20).

(X) 90th percentile of activity across all neurons, plotted as in (W). The shallower decay indicates that the effect in (W) is not due to photobleaching.

(Y) Correlation between true and decoded social proximity for the first 30-minute social block (x-axis) versus second block (y-axis) of each recording with a C57 conspecific. Correlations are significantly lower for the second block (P = 0.008, paired t-test, N = 31).

(Z) Mean activity of object-tuned neurons during frames when the mouse is near an object (nose-to-object distance ≤ 3 cm), stratified by the number of objects added to the arena. Line and shading show mean and 95% CI across N = 6 recordings with at least 5 object-tuned neurons from the recordings with unique objects. Results are shown for the original novel object recordings with unique objects (left) and an additional set of recordings with identical objects (right).

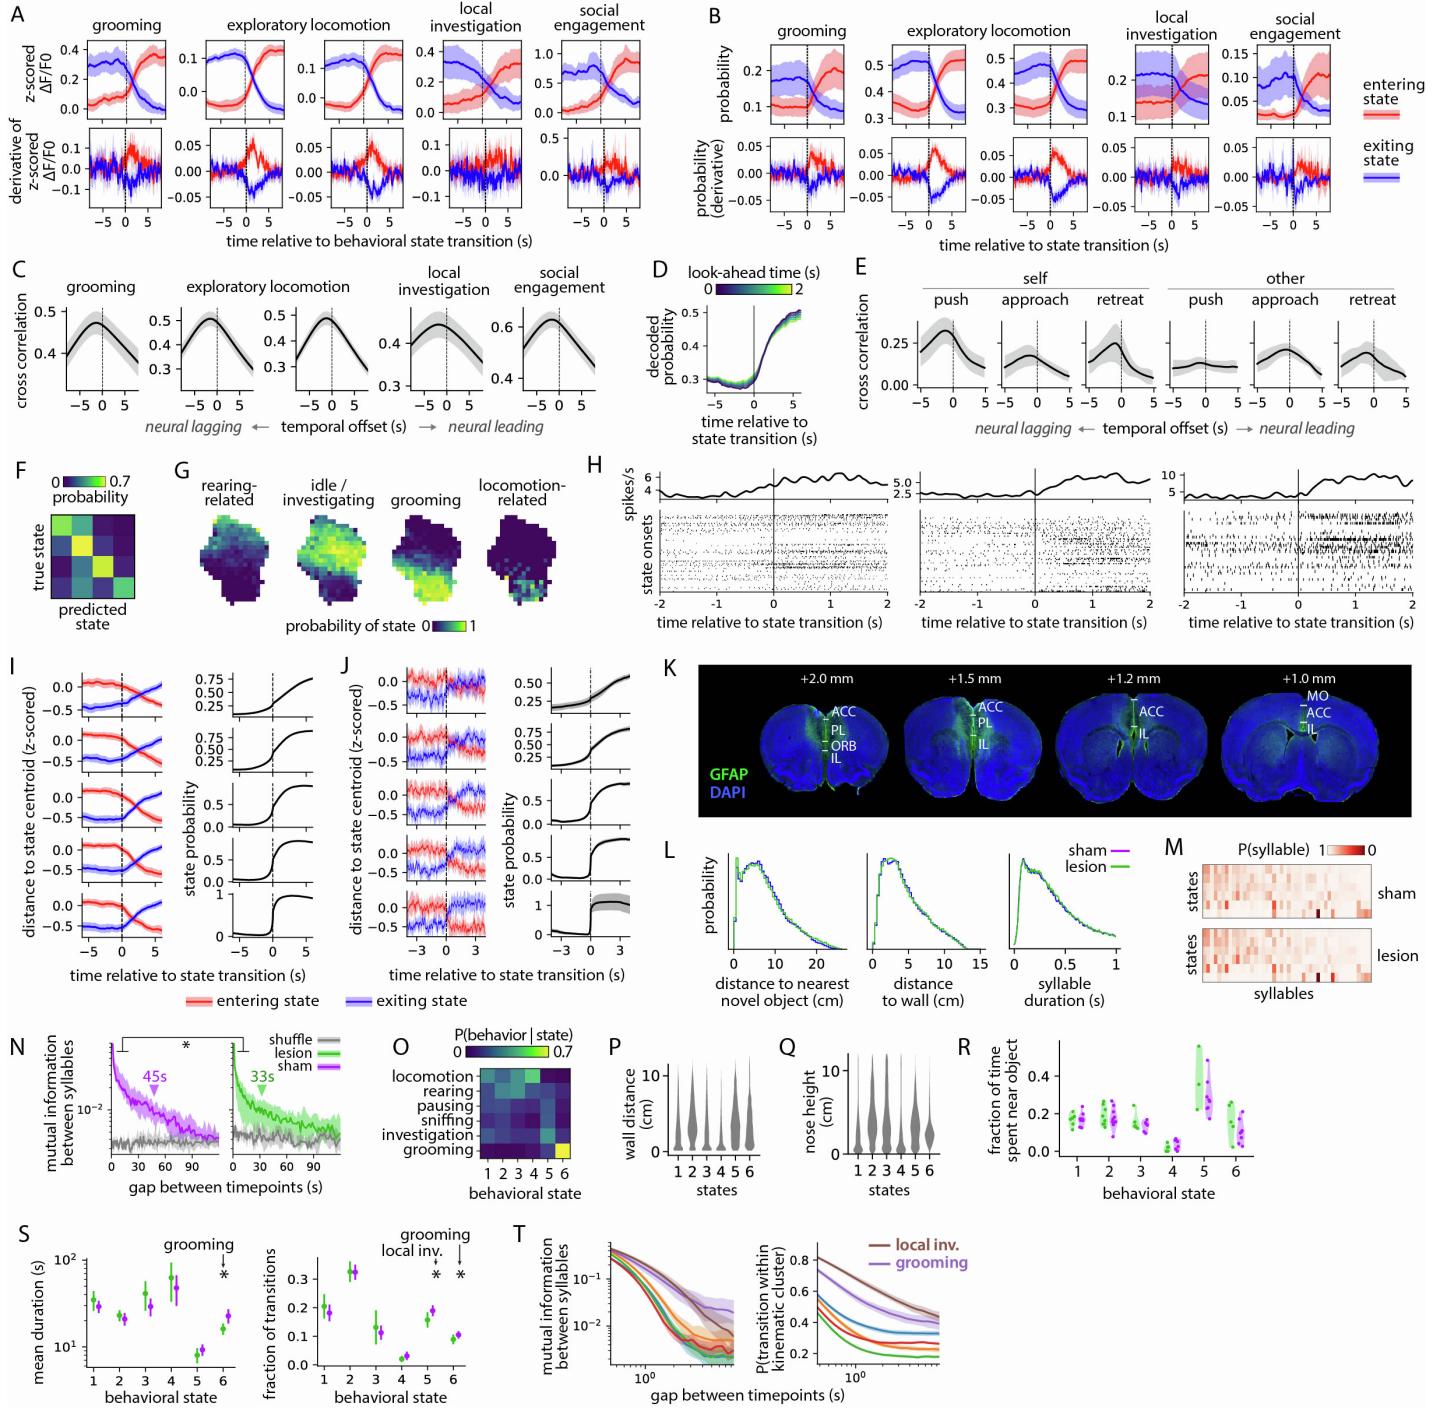


Figure S7: Analysis of timing and excitotoxic lesion effects, related to Figures 6 and 7.

(A) Top: average activity of neurons tuned to each behavioral state, aligned to state onset (red), or offset (blue). For each state, only recordings with at least 10 tuned neurons are included ($7 \leq N \leq 31$). Line and shading show mean and 95% CI. Bottom: derivative of signals shown above.

(B) Top: state probabilities predicted from neural activity, aligned to the onset (red) or offset (blue) of each state. Bottom: rate of change in the predicted probability. Line and shading show mean and 95% CI across recordings ($N = 31$).

(C) Cross correlation between the true and dmPFC-predicted probability of each behavioral state. Line and shading show mean and 95% CI across $N = 31$ recordings.

(D) Decoded state probabilities aligned to state onset, as in (B, red). Each line represents a look-ahead decoder that was trained to predict the state at time $t + \Delta t$ from neural activity at time t .

(E) Cross correlation between behavioral events (rendered as binary time-series) and the decoded probability of those events. Data are from (Kingsbury, 2019). “Self” refers to behaviors performed by the subject mouse and “other” to

behaviors by the opponent. Line and shading show the mean and 95% CI across all recordings where the given behavior was encoded with above-chance accuracy ($3 \leq N \leq 10$).

- (F) Agreement between true and decoded behavioral states across all neuropixels recordings in dmPFC.
- (G) Distribution of behavioral states across a 2D UMAP projection of spiking activity from an example recording. Heatmaps show the probability of behavioral states within local regions of the map.
- (H) Raster plots for three example neurons aligned to behavioral state onsets. Top: peri-event time histogram. Bottom: spike times around each onset event.
- (I) Neural trajectories around state transitions (derived from the initial dataset of calcium-based open field recordings). For each state/recording pair, we calculated a centroid location in neural PCA space. Plots on the left show Z-scored distance to the centroid of the upcoming (red) or previous state (blue) aligned to state transitions. Transitions were stratified into quintiles based on the degree of uncertainty in their timing. Plots on the right show the marginal probability of the upcoming state (derived from shMoSeq); probability increases more sharply when the timing of transitions is more certain. In all plots, line and shading correspond to mean and 95% CI across recordings.
- (J) As (H), here derived from the neuropixels dataset.
- (K) Coronal sections from a lesioned mouse, showing nuclei (DAPI) and an astrocytic antibody (GFAP) that highlights the lesioned area. Numbers indicate anterior-posterior location relative to bregma (ACC, anterior cingulate area; PL, prelimbic area; ORB, orbital area; IL, infralimbic area; MO, secondary motor area).
- (L) Distributions in sham and lesion recordings of (1) distances to the nearest novel object; (2) distances to the wall of the arena; (3) durations of behavioral states.
- (M) Frequency of syllables within each behavioral state from the lesion dataset, calculated separately for control and lesioned mice respectively.
- (N) Results for an additional lesion experiment with a shorter time between surgery and behavior (2 – 3 weeks). The figure shows mutual information (MI) between syllables at a range of temporal lags. Line and shading indicate median and interquartile interval across recordings ($N = 8$ sham, $N = 8$ lesion). MI is significantly lower for lesioned mice at the indicated time-lags ($P < 0.05$). Pointers show the mean decay time for each condition; note that for this smaller dataset, the change in decay time is not statistically significant.
- (O) Distribution of behaviors (defined by hand-labeling of syllables) across the behavioral states.
- (P) Distribution of wall distances for each behavioral state in the lesion dataset.
- (Q) Distribution of nose heights for each behavioral state in the lesion dataset.
- (R) Fraction of time near a novel object (nose-to-object distance ≤ 3 cm) within each behavioral state from the lesion dataset. Each dot represents one recording with novel objects ($N = 10$ sham, $N = 8$ lesion).
- (S) Mean state duration (left) and frequency of state initiation (right) in sham and lesioned animals respectively. Dot and line show mean and 95% CI across recordings ($FDR < 5\%$).
- (T) Left: MI between syllables at a range of temporal lags. Right: probability that syllables from lagged timepoints belong to the same kinematic cluster. Both plots were restricted to pairs of timepoints lying within contiguous instances of each state. Shading shows 95% bootstrap CI over resampled state instances.

Figure panels	Methods section
1D, 1F – 1G, 7B – 7D, 7J, S1L – S1M, S7N, S7T	Analysis of behavioral timescales
3B – 3C, 4A, S4A, S6A	Analysis of single-neuron behavior associations
3D – 3E, 3J, 4B, 4D, 4K – 4O, 6F – 6I, S5B, S5H – S5J, S6E, S6V, S6Y, S7B – S7F	Decoding models
3F, S4E – S4F, S5A, S7G	Neural activity manifolds
3H – 3I, S5C – S5G	Neural encoding models
1D, 1F, 2G, 7B, 7J, S1L – S1M, S3B, S4E, S5A – S5B, S7N, S7T	Clustering syllables based on kinematic similarity
S1J – S1K	Clustering syllables based on transition probabilities

Table 1: Methods used for selected figure panels, related to Figures 1–7. Each row shows representative figure panels that relied on a particular method. Note that this table is not comprehensive. Additional correspondences between methods and figure panels are provided throughout the Methods section in the form of specific figure callouts.

Description	Recording modality	Figure panels
Calcium recordings in dmPFC during solitary and social open field exploration	Depth	1(all), 2A – 2G, 3(all), 4A – 4K, 6A, S1(all), S2A – S2O, S3A – S3D, S4A – S4E, S4H – S4N, S5(all), S6A – S6E, S6G – S6M, S6R – S6S, S6W – S6Y, S7A – S7D, S7I
Calcium recordings in dmPFC during novel object interaction	3D keypoints	2H – 2L, 4A, 4L – 4O, S3E – S3K, S4F, S5C, S5E, S6M – S6O, S6Z
Calcium recordings in dmPFC during exposure to an aggressive conspecific	Depth	5(all), S6P – S6V
Calcium recordings in DLS during solitary open field exploration	Depth	3I – 3J, 7J, S5E, S5H – S5L
Recordings of lesioned and control mice, either in an empty open field or interacting with novel objects	3D keypoints	7B – 7I, S7L – S7M, S7O – S7T
Neuropixels recordings in dmPFC during open field exploration	3D keypoints	6E – 6I, S7F – S7H, S7J
3D keypoint recordings of social interaction and object investigation	3D keypoints	S3L – S3O

Table 2: Datasets that were modeled using shMoSeq, related to Figures 1–7. Each row corresponds to one dataset. To avoid batch effects related to surgical implants, size, sex, and recording modality, separate MoSeq models (and thus shMoSeq models) were fit to each dataset.