Cell Reports

Article

Hierarchical action encoding in prefrontal cortex of freely moving macaques

Graphical abstract

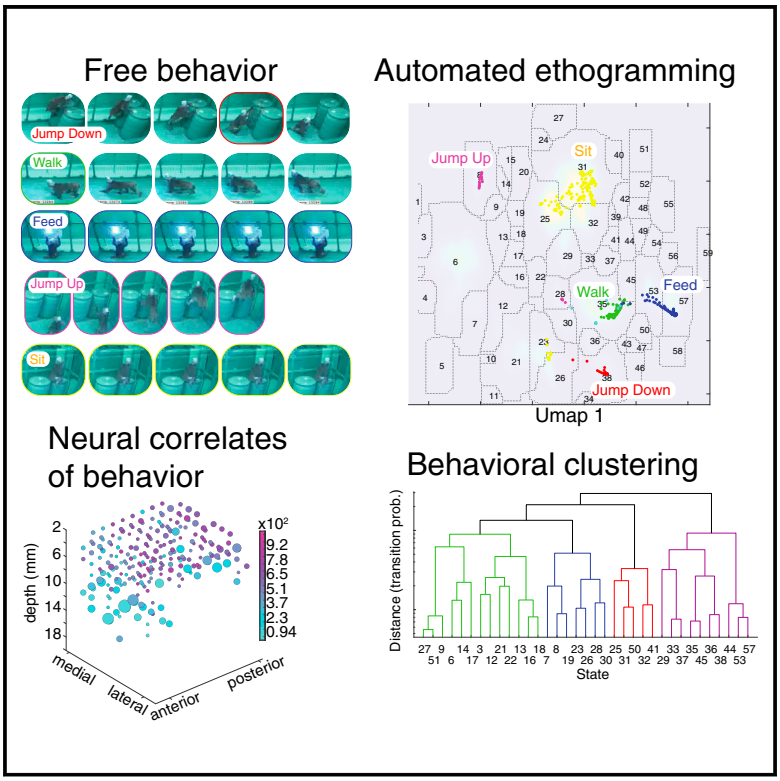

Authors

Benjamin Voloh, David J.- N. Maisson, Roberto Lopez Cervera, Indirah Conover, Mrunal Zambre, Benjamin Hayden, Jan Zimmermann

Correspondence

janz@umn.edu

In brief

Voloh et al. show that the behavior of freely moving macaques can be automatically categorized into states and that prefrontal brain areas in freely moving rhesus macaques encode these states. The strength of encoding increases along a ventral-to-dorsal gradient, and an abstract switching variable shows the reverse pattern.

Highlights

- Behavior of freely moving monkeys can be automatically classified into states
- Multiple prefrontal regions encode automatically identified behavioral states
- The same regions predict behavioral switches, with more abstract signals more ventrally

Voloh et al., 2023, Cell Reports *42*, 113091
September 26, 2023
https://doi.org/10.1016/j.celrep.2023.113091

CellPress

Cell Reports

CellPress
OPEN ACCESS

Article

Hierarchical action encoding in prefrontal cortex of freely moving macaques

Benjamin Voloh,[1] David J.- N. Maisson,[1] Roberto Lopez Cervera,[1] Indirah Conover,[1] Mrunal Zambre,[1] Benjamin Hayden,[1,3,4] and Jan Zimmermann[1,2,4,5,*]

[1]Department of Neuroscience, University of Minnesota, Minneapolis, MN 55455, USA
[2]Center for Magnetic Resonance Research, University of Minnesota, Minneapolis, MN 55455, USA
[3]Department of Neurosurgery, Baylor College of Medicine, Houston, TX 77030, USA
[4]These authors contributed equally
[5]Lead contact
*Correspondence: janz@umn.edu
https://doi.org/10.1016/j.celrep.2023.113091

SUMMARY

Our natural behavioral repertoires include coordinated actions of characteristic types. To better understand how neural activity relates to the expression of actions and action switches, we studied macaques performing a freely moving foraging task in an open environment. We developed a novel analysis pipeline that can identify meaningful units of behavior, corresponding to recognizable actions such as sitting, walking, jumping, and climbing. On the basis of transition probabilities between these actions, we found that behavior is organized in a modular and hierarchical fashion. We found that, after regressing out many potential confounders, actions are associated with specific patterns of firing in each of six prefrontal brain regions and that, overall, encoding of action category is progressively stronger in more dorsal and more caudal prefrontal regions. Together, these results establish a link between selection of units of primate behavior on one hand and neuronal activity in prefrontal regions on the other.

INTRODUCTION

A common view in neuroscience sees the expression of behavior as the result of computations in a circumscribed set of dorsal prefrontal regions, especially the motor and premotor cortices. In this view, other more rostral and ventral regions have qualitatively distinct contributions, such as value representation and working memory. A contrasting view sees the prefrontal cortex as a hierarchy and emphasizes the quantitative but not qualitative difference between regions.[1–7]

We were especially interested in understanding the answers to these questions in the domain of natural behavior. Most of the research on the neural basis of action is focused on simple movements, especially reach movements with the arm and saccadic movements in the oculomotor system. Research using these methods have given us important knowledge; indeed, many of the successes in motor control come from such research, including work that identified important roles for the primary (M1) and dorsal premotor (PMd) cortices, as well as roles for the supplementary motor area (SMA) and even the dorsal anterior cingulate cortex (dACC) in action planning and execution.[8–11] However, it remains unclear whether these ideas will apply to more naturalistic behaviors of the type that characterizes real-world behavior.

Indeed, there is a good deal of research looking into naturalistic behavior.[12–15] For example, recent work shows that the behavior of mice is organized into a grammatical structure and that the dorsolateral striatum mediates this organization.[16] Likewise, selection and change of behavioral state in flies and zebrafish is mediated by specific patterns of neural activity in specific neurons.[17,18]

The study of the neuroscience of natural behavior has depended critically on the development of high-quality video-based motion-tracking systems. These systems have led to the ability to track positions of body landmarks in small animals, including worms, flies, and mice.[12,17,19–23] This problem is much more difficult in primates; although, even here, significant progress has been made.[24–31] In smaller animals, the behaviors identified by tracking systems allow for the automated identification of specific meaningful behavioral units (sometimes called "ethogramming."[23,25,30,32–34] Behavior in these organisms consists of structured sets of behaviors that are organized hierarchically. The success of these methods raises hopes that we can develop parallel methods for quantifying behavior in macaques.[30,31]

One question we were particularly interested in is that of how the hierarchy of behavior is reified in the prefrontal cortex. On one hand, a modular viewpoint would predict that different anatomical regions would have qualitatively different roles and, thus, that when looking at the neural basis of state, some regions would show much stronger effects than others or even qualitatively different roles. On the other hand, a hierarchical, or gradient-based, viewpoint would predict that these different regions would have largely similar effects qualitatively, but they

would differ quantitatively with their position in an anatomical hierarchy.[5,35,36]

We examined the behavior of two macaques performing a foraging task moving around a large (2.45 × 2.45 × 2.75 m) open field. We made use of Open Monkey Studio, a system that can perform detailed three-dimensional (3D) behavioral tracking in rhesus macaques with high spatial and temporal precision.[27,28] We used this system to track the positions of 15 body landmarks at high temporal and spatial resolution as our subjects performed a foraging task. We recorded brain activity in six brain regions: orbitofrontal cortex (OFC), dACC, SMA, ventrolateral prefrontal cortex (vlPFC), dorsolateral prefrontal cortex (dlPFC), and PMd. In all regions, we found that neural activity varies systematically with the action of the subject. We found that the strength of action coding was systematically greater in more caudal and more dorsal structures. We also found that neurons have non-specific signals associated with the switch between actions and that the strength of these signals is—in a reverse of the pattern for action signals—progressively stronger in more ventral prefrontal structures.

RESULTS

Behavioral and neural recordings

We studied the behavior of two male rhesus macaques (*Macaca mulatta*) performing a depleting-rewards freely moving foraging task (see STAR Methods; Videos S1, S2, S3, and S4) in a large open cage (2.45 × 2.45 × 2.75 m cage with four barrels) that allowed for free movement (Figures 1A and 1B). The enclosure contained reward stations (typically four, but on a few occasions two or three; see STAR Methods). These reward stations consisted of touch screens, juice reservoirs, dispenser tubes, and levers (Figure 1B). Each station had the same arrangement, but individual stations were positioned at fixed locations of varying heights. Subject behavior was otherwise physically unconstrained within this large environment.

The task was straightforward; the computer displays were initially all blue. If the subject pressed the lever at a reward station, the juice tube provided an immediate aliquot of preferred liquid reward (typically water, always 1.5 mL) and a green cross appeared while a 2-s tone played (Figure 1C). When the tone stopped, the screen turned off. Following another 2-s interval, the screen returned to blue to indicate that the lever had been reactivated. The subject could repeat this process four times to obtain four (identically sized) rewards (Figure 1D). Following the fourth lever press, the patch was inactivated, meaning it would no longer activate with time. Instead, to reactivate a patch, the subject was required to press the lever again (that is, a fifth time); this press would not be rewarded, but instead would initiate a 3-min-long waiting period, following which the patch would reset. Note that, during early training, we found that in the absence of a task, macaque subjects will typically sit in the cage without much movement. Thus, the major goal of the task for this study was to elicit variegated behavior.

The length of each daily recording session was set by the experimenter (average: 100.2 min). We analyzed 81 sessions for subject Y and 86 for subject W (see STAR Methods for exclusion criteria). Subjects were tracked with 62 high-resolution machine vision cameras (Figure 1B), and the pose (3D positions of 13 body landmarks, see STAR Methods) was determined for each frame (30 frames per second) using our pose-tracking system, Open Monkey Studio[27,28] (Figure 1E). This system provides estimates of the positions of 13 body landmarks in every frame of video in three dimensions. As a result, we obtained an average of 180,307 frames of pose per session and a total of 35,881,005 frames. We simultaneously recorded neural activity using a locally mounted data logger attached to a multielectrode (n = 128 electrodes) array with independently movable electrodes (Figure 1F; STAR Methods), targeting multiple regions in the prefrontal cortex (Figure 1G).

Embedding recovers actions

To identify behaviors, we used an embedding approach similar to one that has been used successfully in rodents and flies.[32,33] A high-level description of the pipeline is visualized in Figure 2A. In brief, we extracted kinematic features from the 13-landmark reconstructed pose, re-projected these samples onto a common 2D embedding space, and clustered the samples to discover repeated, recognizable behaviors. We extracted a total of 91 pose-related features. We selected these features with the goal of providing a fairly full accounting of the diversity of behavior. This list of features included limb and joint configurations, periodicity in limb configurations, and gross dynamics of the subject; see the STAR Methods for a full list of features and Figure 2E for an example.

To cluster samples robustly and efficiently, we employed nonlinear dimensionality reduction (UMAP[37]) on these 91 features to re-project samples. We trained the embedder on a training set from subject Y. We chose to construct our basis set from the subject with greater behavioral range; had we chosen the other subject, we would have obtained similar, albeit slightly noisier, results (data not shown). This process resulted in a re-projection of the 91D behavioral space onto two dimensions, in which periods of similar kinematics are adjacent to one another in the resultant lower-dimensional space.

We then performed a kernel density estimation to approximate the probability density of embedded poses at equally interspersed points (Figure 2B). The embedding density maps clearly reveal a clustered organization, as evinced by visually separable peaks (Figure 2B). Each cluster reflects a set of similar poses that are relatively distinct from others. We formally identified clusters using a watershed algorithm on the density map from the training set (technically, on its inverse, so that peaks became troughs).[33] This algorithm treats troughs in the maps as sinks and draws optimal boundaries separating them. The resulting embedding space contained 59 distinct clusters. We then used the learned embedding and cluster assignment to re-project all data samples from both subjects and assign a cluster to each sample (Figure 2C).

Examples of behaviors discovered through this pipeline are visualized in Figures 3A–3C. Figure 3A shows example behaviors delineated by an observer (such as jumping, walking, and sitting) and the corresponding cluster assignment. Each of these behaviors was associated with a unique fingerprint of kinematics (Figure 3B), corresponding to separate clusters in the low-level embedded space (Figure 3C). We confirmed that clusters were

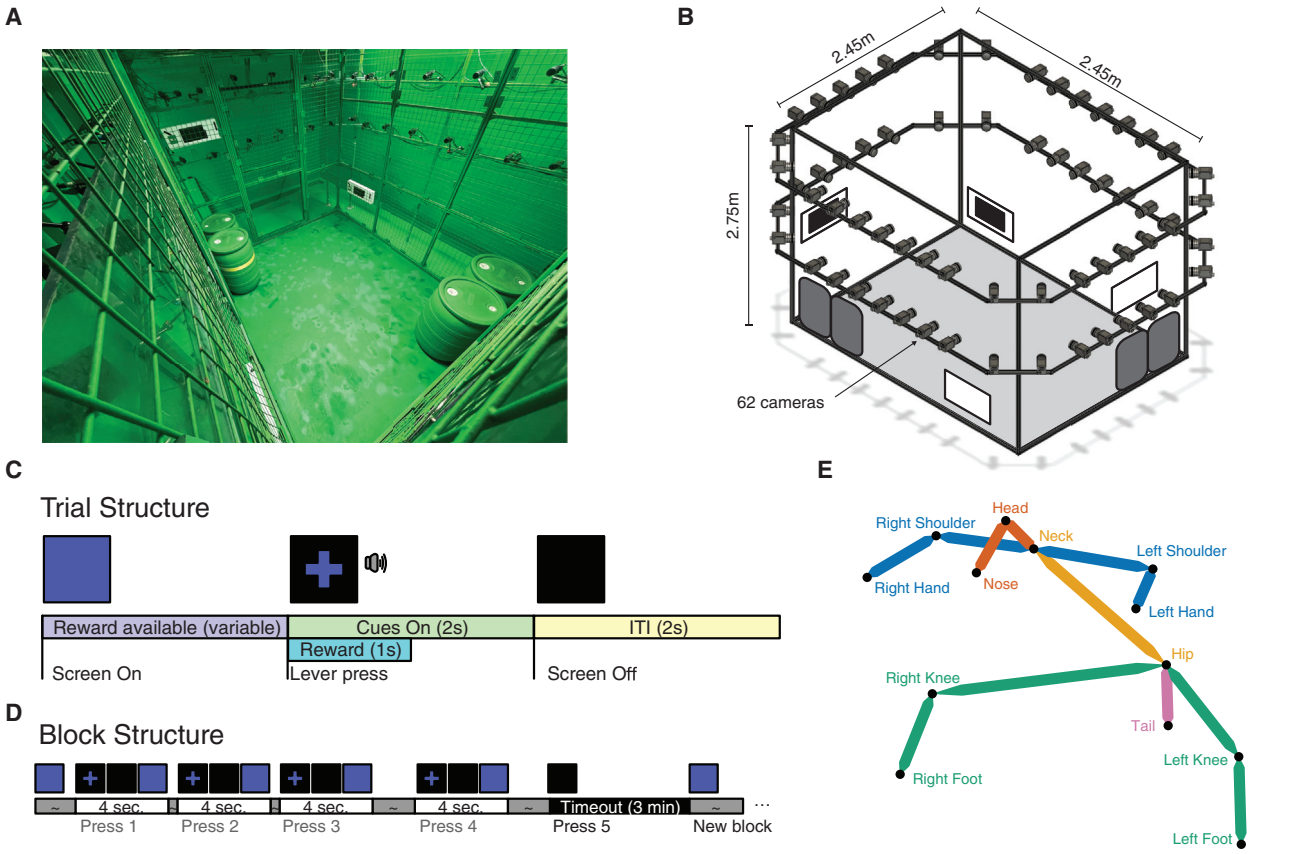

Figure 1. Environment, task, and electrophysiology for freely moving macaques

(A) Photograph of inside of enclosure permitting freely moving foraging. Up to four feeders were present for the experiment, located at different corners and requiring different postures to reach.

(B) Schematic of the cage. Sixty-two machine vision cameras provided multiview coverage of every part of the cage.

(C) Structure and timing of one trial. Following a blue initialization screen, subjects pressed a lever and received a fluid reward. After 2 s, the trial reset.

(D) Structure of one block of trials. Four rewards were available. A fifth lever press initiated a 3-min time-out period.

(E) A 3D reconstruction of 13 landmarks using the Open Monkey Studio pipeline.

(F) A 3D model of the recording system superimposed on a subject's cranium.

(G) A 3D rendering of the prefrontal areas from which neural data were recorded (left), and the corresponding number of recorded neurons (right). Y, subject Y; W, subject W.

well separated by computing the Davies-Bouldin index (DBI), a ratio of the within-cluster scatter to between-cluster distances. In both subjects, the DBI was lower than chance, indicating that samples within a cluster were more similar than between clusters and could thus be statistically separable (subject Y mean $DBI_{observed}$ = 2.24 ± 0.023, DBI_{rand} = 65.1 ± 1.8, $p < 0.001$; subject W mean $DBI_{observed}$ = 2.9 ± 0.029, DBI_{rand} = 36.1 ± 0.41, $p < 0.001$). Indeed, even when considering only pairs of clusters, the vast majority of pairs were separable (subject Y, 99% ± 0.0003%; subject W, 86% ± 0.03%). This provides confirmation that the extracted clusters corresponded to unique kinematics that are separable in the embedding space. As such, for

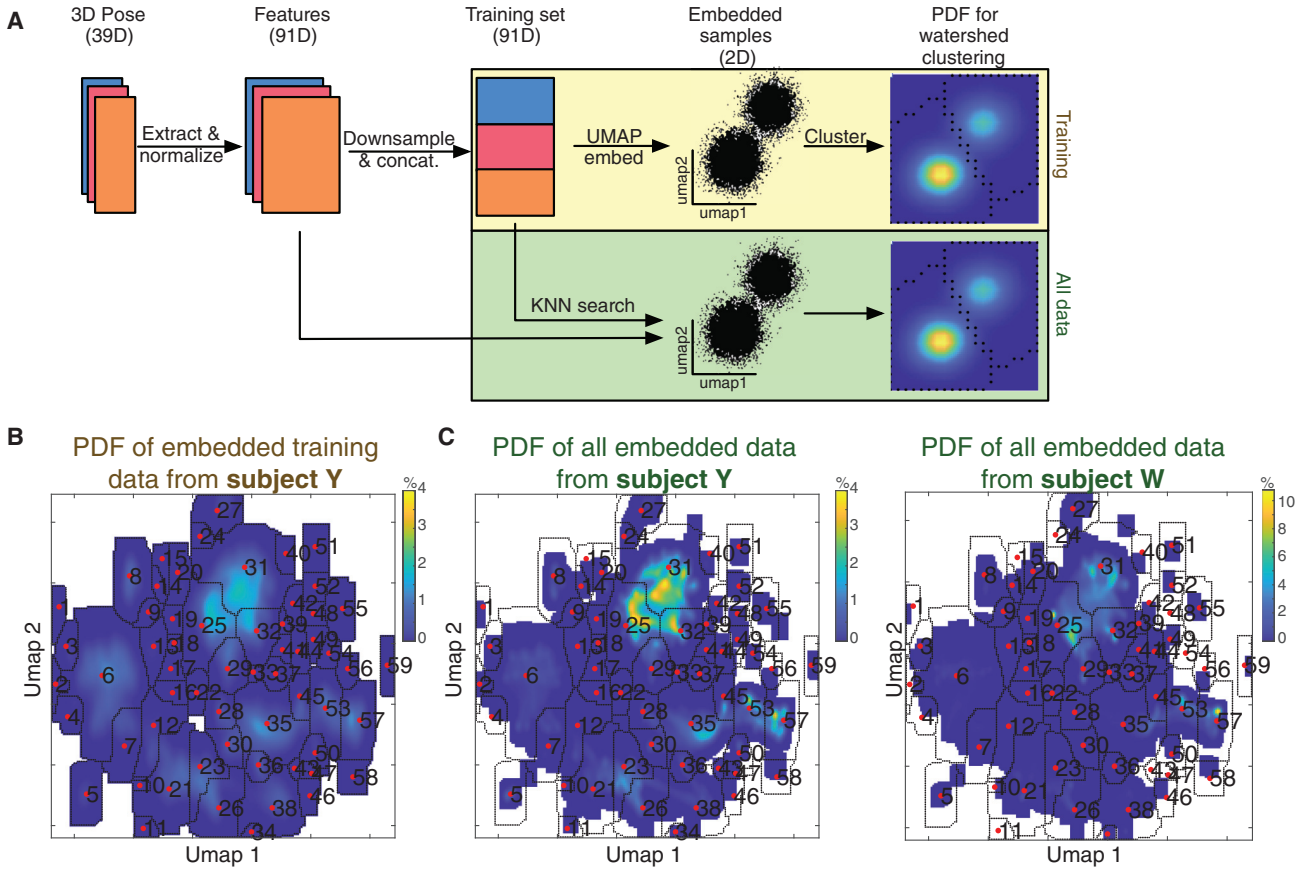

Figure 2. Embedding pose kinematics to identify actions

(A) Pipeline for unsupervised behavioral segmentation. After preprocessing, each dataset of 13 landmark positions was used to extract 91 salient features. We generated a 2D embedding, which was used for clustering via watershed. Using the learned embedding and cluster labels, we used the training set to re-embed and re-label all data.

(B) Probability density map for subject Y. Peaks in the heatmap indicate samples with similar postural dynamics. Dotted lines, cluster boundaries from the watershed algorithm. Numbers, cluster labels.

(C) Same as (B), but for the (re-embedded) testing data for subjects Y (left) and W (right).

the remainder of this paper, we will refer to these clusters of behavior as "actions." Each action lasted on average 1.63 ± 0.009 s in subject Y and 0.67 ± 0.0025 s in subject W. Note that the ranges given here are SEM; for SD, the respective values are 4.96 and 2.14. We acknowledge that different actions have different durations and that different categories of actions have different mean durations. We hope that future studies will productively explore the variation within and across categories.

Action organization is modular and hierarchical

It has long been supposed that natural behavior in primates has an inherently hierarchical organization. Indeed, previous work, including from our own lab, has shown that function is often better described as a hierarchical progression rather than as a set of functionally distinct regions.[1–3,35,38–40] We therefore wondered whether this was also true in this task and, importantly, extended to unconstrained tasks in which the animals freely choose to engage in behavior. We next sought to characterize this hierarchy in our subjects. Formally speaking, we asked if sets of actions tend to systemically co-occur, which would indicate the presence of higher-level organization. To do this, we computed transition probability matrices for each possible pair of actions using the actions discovered through embedding (see above).

The transition probability matrix this process produces is, in graph theoretic terms, a directed graph. In the parlance of graph theory, nodes that form strong links between one another are referred to as modules or communities. Because of this, we refer to sets of actions that have a high probability of co-transition as "action modules." The identification of these modules allows us to sort the transition probability matrix so that the actions that are part of the same module are adjacent; graphically, the result of this is that there will be conspicuous blocks on the diagonal (Figure 4A). These blocks correspond to action modules.

We then used a recently developed algorithm named Paris,[41] which performs hierarchical clustering on the graph derived from transition probabilities and returns a tree (we will examine the assumption of hierarchy using the Dasgupta score below). To determine the number of modules, we cut the tree at several levels and computed a modularity score for the resulting

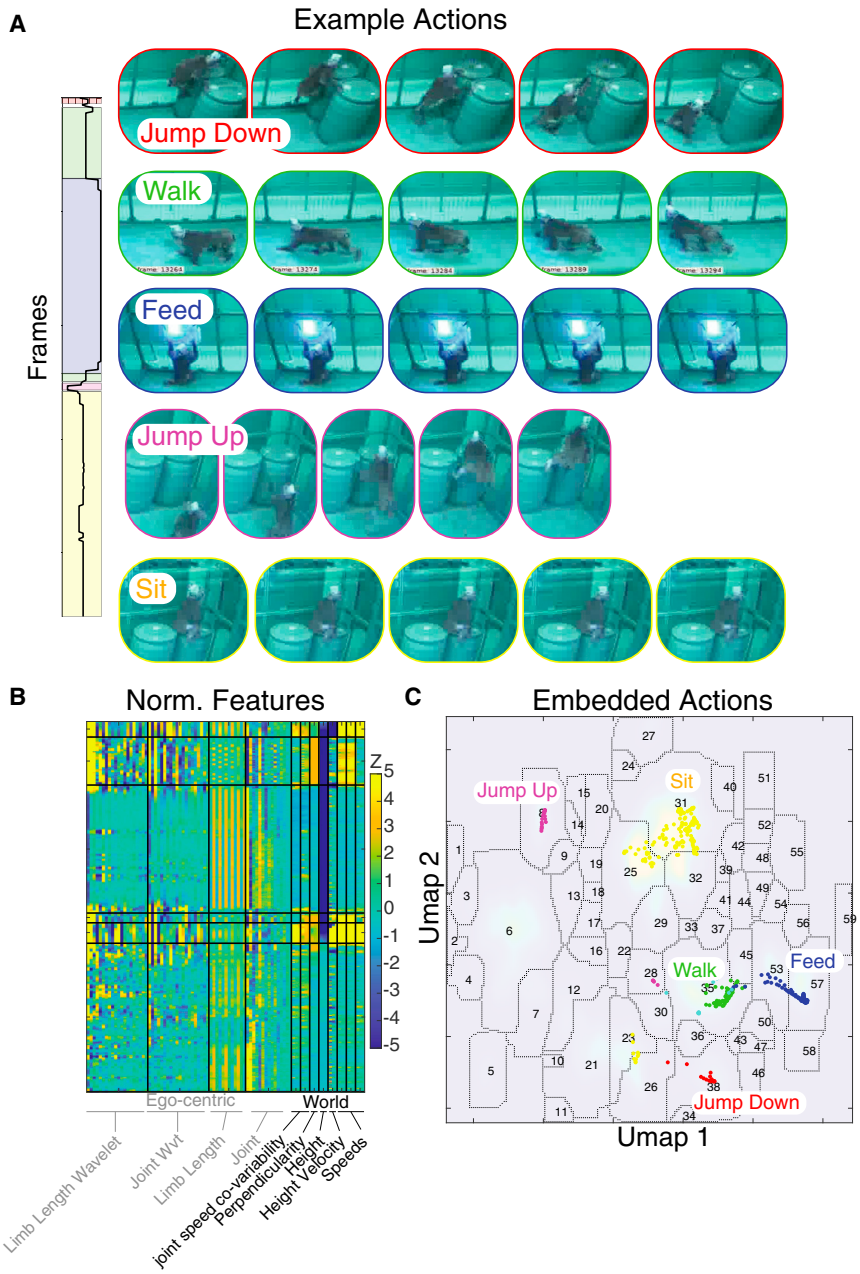

Figure 3. Example actions, their pose features, and their embedding

(A) Example actions from one dataset with an ethogram of actions (left). Colored backgrounds represent different actions as determined by an experimenter. Black line represents the learned cluster label.

(B) Normalized features computed from pose kinematics from the examples in (A). Color limits have been scaled to aid in visualization between features.

(C) Embedding density map from Figure 2B, with data from the example actions (A and B) overlaid. Actions were well separated from one another.

subtrees and chose the number of modules that maximized the modularity score. The modules that result from this cut give the highest average within-module action-transition probability and the lowest average across-module action-transition probability. These modules maximize the difference between these two measures. From this process we can identify the most likely behavioral modules (that is, the ones that fit the data the best).

An example of this process is visualized in Figure 4A. For the transition probability matrix (Figure 4Ai), modularity was maximized for five modules (Figure 4Aii). The modular nature of transitions is evident in the sorted transition probability matrix (Figure 4Aiii). Note that not all actions were present in this session; we performed clustering and modularity computations with this in mind.

We tested for the modular organization of behavior by computing the modularity score for each session. We found that, in all sessions, both subjects exhibited significant modularity (Figure 4B; randomization test, $p < 0.001$). In other words, behavior was organized such that specific modules of actions tended to transition between one another, but not to actions outside of the module. The average number of modules was 4.6 ± 0.08 in subject Y and 4.83 ± 0.07 in subject W. These results indicate that subjects' behavior was organized into modules, each of which consisted of stereotyped actions.

We then asked if action modules exhibit higher-level organization. Examples of hierarchically organized actions are shown in Figure 4C. Higher-level connections in this dendrogram show how different action modules are related. For example, the green branch on this tree refers to a set of closely related actions that can be labeled as different variants of "sitting," while the red branch refers to actions that can be labeled as variants of "walking." Of course, each of the green subbranches is a slightly different form of sitting, and so on. At the higher levels of the hierarchy are groups of actions that cluster together. Thus, for example, the algorithm treats different forms of "passively hanging" (blue branch) as being similar to "sitting" (green branch) and more similar to each other than either is to "walking" (red branch). To quantify the degree of hierarchical organization, we computed the Dasgupta score, which quantifies the quality of hierarchical clustering on the transition probability graph.[42] A score above chance indicates the observed tree has high-level components that are distinctly related to one another. The Dasgupta score was above chance in all sessions in both subjects (Figure 4D; randomization test, $p < 0.001$), indicating hierarchical organization of behavior. Note that, to assess

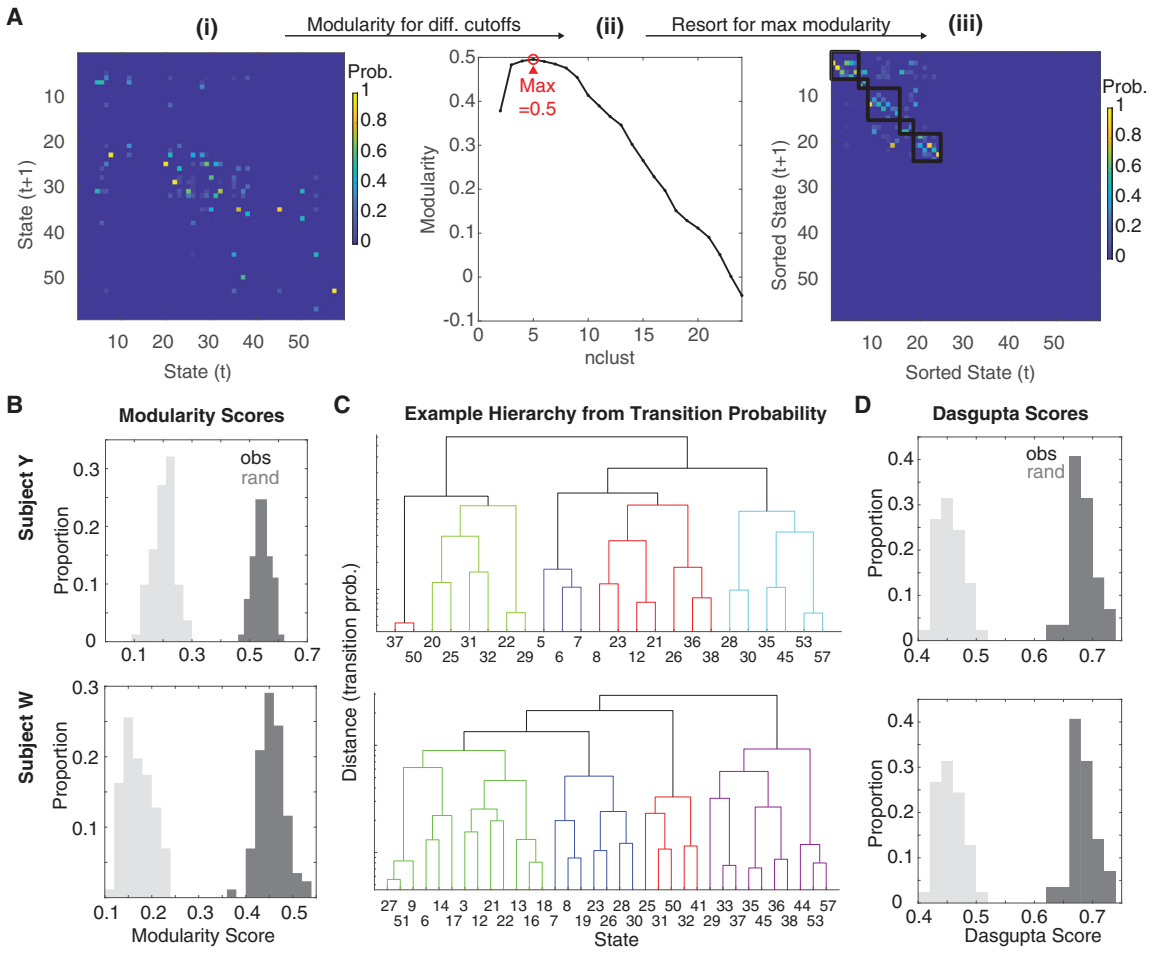

Figure 4. Actions show modular and hierarchical organization

(A) Example modularity for one dataset. (i) Original transition probability matrix with a transition lag of 1. (ii) Modularity score for different cuts of the behavioral hierarchy. The maximum modularity was for five modules. (iii) Transition matrix sorted to provide maximum modularity. It is evident that transitions tend to occur within, rather than between, modules.

(B) Distribution of maximum modularity (e.g., peak in Aii) for each dataset, for subjects Y (top) and W (bottom). Observed modularity scores (dark) were greater than chance for all datasets (light).

(C) Example dendrograms showing action organization for one example session for each subject. Each branch on the dendrogram corresponds to a set of closely related actions. For example, branches can correspond to actions like "sitting," "walking," "approaching," or "climbing." Lower-level branches correspond to more elementary actions, such as different smaller movements that, together, constitute climbing. Higher levels are groups of actions that cluster together. Thus, for example, the algorithm treats different forms of "passively hanging" (blue branch) as being similar to "sitting" (green branch) and more similar to each other than either is to "walking" (red branch).

(D) Dasgupta scores for each dataset, measuring the degree of hierarchy as evidenced by transition probabilities. Observed Dasgupta scores were greater for observed (dark) rather than randomized (light) transition matrices, indicative of the existence of a hierarchical structure.

significance, we used a standard randomization approach. Specifically, we randomized the time series of cluster labels and computed the new Dasgupta score. Scores greater than the distribution of randomized scores were taken as statistically significant.

One possibility is that the modular and hierarchical organization we observed was primarily driven by transitions between highly similar periods of postural dynamics that were too finely segmented during embedding. To control for this possibility, we merged the fine action clusters into progressively coarser clusters (4–58 clusters) according to the distance of their medoids in the embedding space. We found that even coarse clusters exhibited modular and hierarchical organization above chance (Figure S1). Modularity and hierarchy scores were greatest using all 59 clusters, and thus all clusters were taken into consideration for further analysis.

Neurons across the prefrontal cortex encode actions

We next sought to understand how actions are encoded across six prefrontal regions. We recorded a total of 10,502 neurons over 196 sessions. Of these, 2,818 neurons were excluded from the following analyses because preliminary investigations indicated that their sessions had tracking that was too noisy for our purposes or because regressing out extraneous variables

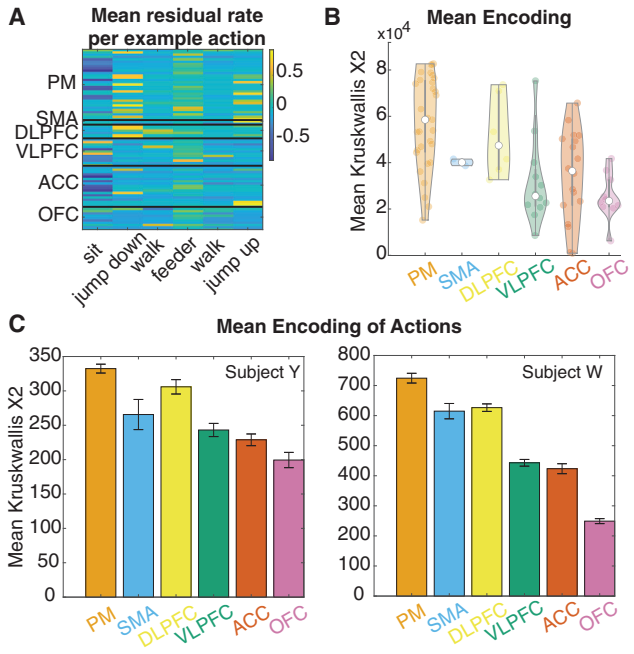

Figure 5. Actions encoding across brain areas

(A) Mean (residualized) firing rates for six example actions from one session, split by brain area. Here, residualized means after accounting for potential confounders (see STAR Methods).

(B) Mean encoding strength for each area for the entire session, operationalized as the Kruskal-Wallis chi-squared value, across all recorded areas. More dorsal areas exhibit higher encoding strength.

(C) Mean encoding strength across all sessions and areas, in subjects Y (left) and W (right). More dorsal areas exhibited stronger encoding strength.

failed (see STAR Methods for exclusion criteria). The remaining 7,684 neurons over 167 sessions were recorded from six structures in the prefrontal cortex, OFC, vlPFC, dACC, SMA, dlPFC, and PMd (see STAR Methods and Figure 1E).

We examined the average firing rate of each neuron during each of the 59 identified actions by averaging each neuron's responses across the entire action (see STAR Methods). First, we computed the spike density time series of each neuron and downsampled it to match the pose time series (30 Hz). Next, to isolate activity related to actions, we fit a Poisson generalized linear model (GLM) to identify neuronal responses related to task events (such as lever presses, rewards, and cues) and subject position (in the xyz dimensions). We then used the fit model to residualize the neuronal time series. All neural analyses were performed on these residualized and downsampled time series.

We found that neurons exhibited specific firing to individual actions. An example of rates of neurons recorded in one session (n = 79 neurons) to six different actions is shown in Figure 5A. In this session, neurons had distinct firing patterns for each action; different areas showed distinct patterns as well (Figure 5B). We operationalized the strength of action encoding of each individual neuron by performing a Kruskal-Wallis test and extracting the χ^2 value. This value reflects how well actions can be dissociated from one another by the (residualized) firing rate. We opted for a non-parametric measure as rates were highly non-uniformly distributed. In this session, posterior regions, such as the PMd, had the highest average strength of action encoding, while more ventral and anterior regions, such as dACC, vlPFC, and OFC, had weaker encoding (Figure 5B).

Indeed, across all sessions and in both subjects, we found evidence of significant action encoding in all six regions tested (Figure 5C). Specifically, in all six regions, the median χ^2 was significantly greater than chance (randomization test, $p < 0.001$ in all areas and for both subjects). Furthermore, the mean strength of encoding differed between regions (ANOVA; subject Y, $F = 32.2$, $p < 0.001$; subject W, $F = 170$, $p < 0.001$). These results suggest that action encoding is prevalent across the prefrontal expanse, even after controlling for task-related activity and subject position.

Action-encoding strength grows along a dorsal-ventral gradient

A visual inspection of the data suggests that action coding grows systematically with the dorsoventral position of the recording site (Figure 5C). We next sought to formally test this hypothesis by modeling encoding strength as a function of electrode position (Figure 6A). For each electrode, we extracted its depth (z position, along the vertical axis) and its x (along the coronal plane) and y (along the sagittal plane) position on the recording grid. We then fit a linear model, predicting encoding strength (specifically, χ^2 value) as a function of 3D spatial position (see STAR Methods).

We found that encoding strength varied systematically with the position of the electrodes (subject W, $F = 293$, $p = 0$, $R^2 = 0.18$; subject Y, $F = 70$, $p = 0$, $R^2 = 0.05$). Looking specifically at each dimension (Figure 6B), we found a significant positive relationship between encoding strength and electrode elevation (stronger coding in more dorsal recording sites; subject W, $p < 0.0001$; subject Y, $p < 0.001$). We also found a significant positive relationship between encoding strength and electrode anterior-posterior position (stronger coding in more posterior/caudal locations; subject W, $p < 0.0001$; subject Y, $p < 0.001$). There was an inconsistent correspondence of encoding with the mediolateral axis (subject Y, $p = 0.28$; subject W, $p < 0.001$).

Some locations had much better coverage than others, which could lead to spurious fits due to large imbalances in the data. To control for this possibility, we binned encoding values according to spatial location (for 5, 10, 15, and 20 bins along each dimension) and then repeated this analysis. We consistently found that the anterior-posterior and depth dimensions predicted encoding strength in both subjects ($p < 0.02$). On the other hand, in contrast to the results above, the mediolateral axis consistently did not predict encoding strength ($p > 0.25$).

In both subjects, recording position significantly predicted encoding strength (subject W, $F = 293$, $p = 0$, $R^2 = 0.18$; subject Y, $F = 130$, $p = 0$, $R^2 = 0.096$). Encoding was stronger in shallow rather than deep locations (subject W, $p = 0$; subject Y, $p = 0$) and posteriorly rather than anteriorly (subject W, $p = 0$; subject Y, $p = 0$). As before, encoding along the mediolateral axis was inconsistent (subject W, $p = 0$; subject Y, $p = 0.62$). In our original analysis, encoding strength ranged from ~60 to ~600 in subject Y and from ~90 to ~900 in subject W; in this analysis, encoding

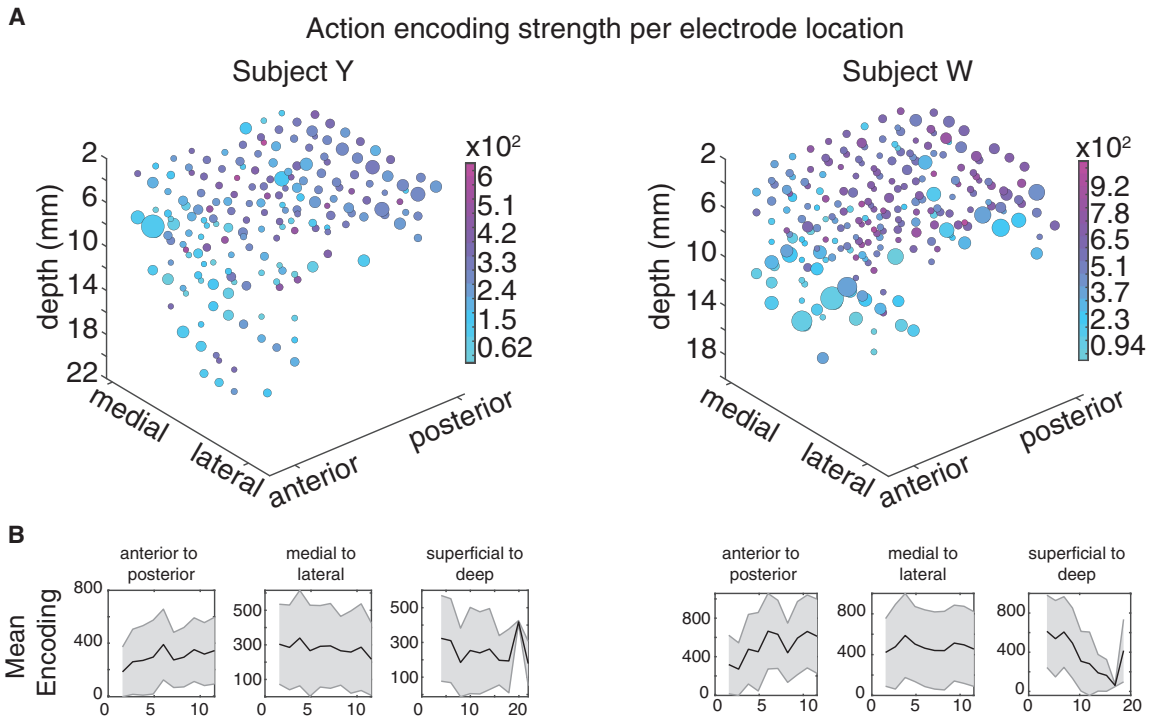

Figure 6. A dorsal-ventral gradient of action encoding

(A) Average encoding strength as a function of electrode position for subject Y (left) and subject W (right). Bubbles correspond to bins (n = 20) of neurons. Color indicates strength of encoding, size indicates number of neurons in bin. Stronger encoding (darker color) is more common in dorsal and posterior locations; weaker encoding (lighter color) is more common in ventral and anterior locations.

(B) Average encoding strength after collapsing along the anterior-posterior (left), medial-lateral (middle), and superficial-deep (right) axes.

strength ranged from ~30 to ~480 in subject Y and from ~30 to ~420 in subject W.

As a further control, we also tested whether the action-encoding gradient arose due to the residualization of rates. We found an identical pattern of results as before (stronger encoding in more dorsal and posterior regions, rather than ventral and anterior); however, encoding strength was approximately half that of the analysis above, highlighting that residualization effectively isolated activity related to actions. We conclude that, at least for our analysis methods, actions are most strongly encoded in posterior and superficial areas, such as the PMd, and weakest in deep anterior structures, such as OFC.

Systematic relationship between switching-related patterns and hierarchical position

We next wondered how neurons implement switches between actions: do they simply identify the action being implemented, or do they signal the switch as well? We therefore examined the patterns of neural activity associated with switches between actions. To this end, we selected a 2-s time window centered on the moment of switching around each action transition. We restricted our analysis to this short temporal window to account for the variable nature of temporal dependencies between actions. Since for this analysis we considered peri-switch activity, for any one neuron, we considered segments only where the action before or after the switch lasted longer than 200 ms, thus ensuring that the 400 ms around the switch were not contaminated by other switches. We then normalized activity within and across neurons (see STAR Methods). We adopted this normalization procedure because, for the hypothesis in question, the absolute levels of firing rate pre- and post-switch are irrelevant. That is, we are interested only in the shape of the temporal envelope, not in the absolute levels. For this reason, we eliminated that information through averaging. We acknowledge that these processing steps may appear complex. However, they are carefully designed to avoid confounding factors.

The average segment-normalized activity is shown in Figure 7. We found that in all six cortical regions, and in both subjects individually, neuronal activity changed after action switches. In the OFC, which saw the strongest switch effects, the switch produced a reliable increase in neural activity in both subjects (0.02–0.06 times greater than the preswitch mean; Figure 7B). Note that, due to our normalization procedure, this apparent rise in firing reflects what is more likely the rate of return to a higher between-switch baseline level of firing. Indeed, additional analyses (not shown) suggest that the switch-related change is best described as a switch-evoked suppression in activity followed by a return to a variable but generally higher between-trial baseline rate.

The critical result is that the type of change differed somewhat between subjects. PMd, SMA, and dlPFC were inconsistent between subjects, showing relatively increased activity in subject Y

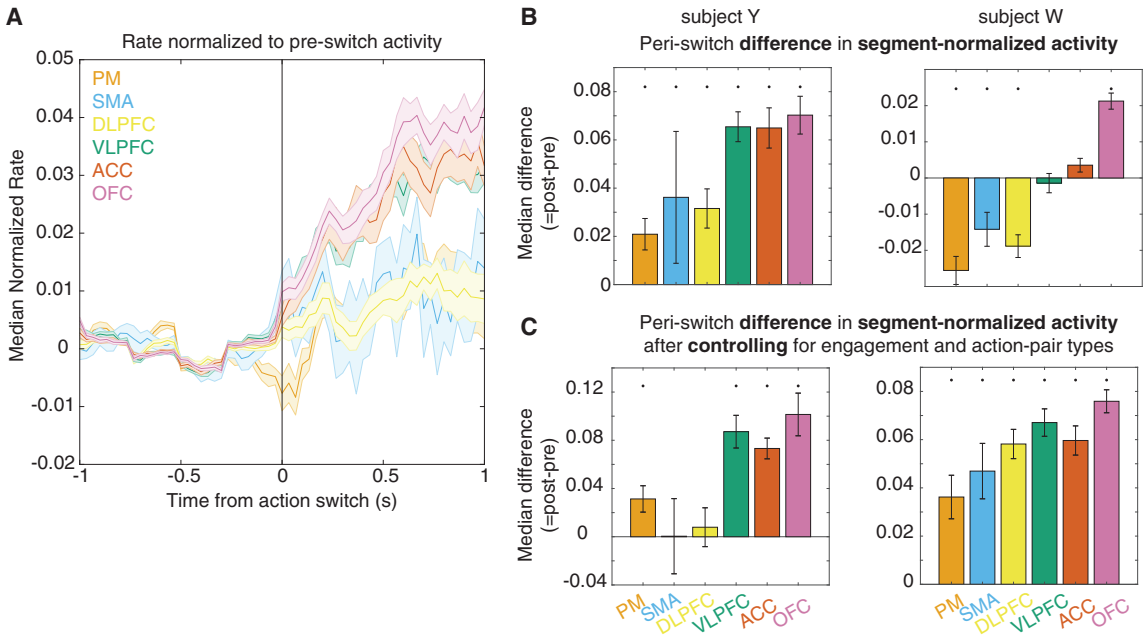

Figure 7. Neural correlates of action switches

(A) Average segment-normalized activity aligned to time of switch. Colors indicate brain areas.

(B) Average (median) difference of post- vs. presegment-normalized activity, split by area (color). Black dots indicate significant difference ($p < 0.05$). Both subjects show an increase with more ventral structures.

(C) Same as (B) but after controlling for possible reward and action-pair type effects.

and decreased activity in subject W. On the other hand, OFC, vlPFC, and dACC showed a similar trend to increase activity in both subjects. However, in both subjects, the average difference between pre- and post-switch activity differed by area and in the same way (Kruskal-Wallis test; subject Y, $\chi^2 = 48$, $p < 0.001$; subject W, $\chi^2 = 176$, $p < 0.001$). Furthermore, the difference in activity could be predicted by the position of the electrodes (linear model; subject Y, $F = 10.3$, $p < 0.001$; subject W, $F = 49.3$, $p < 0.001$). The switching difference was greater for deeper depths (subject Y, $p < 0.001$; subject W, $p < 0.001$) and for more anterior (near significant in subject Y, $p = 0.06$; subject W, $p < 0.001$) locations. In subject W, the difference was also greater for more lateral locations ($p < 0.001$); this effect was not significant in subject Y, although we observed a trend in the same direction ($p = 0.2$). Thus, the relative change in activity after action switches exhibited the opposite trend compared with action encoding per se, with more ventral regions showing greater activity post-switch and more dorsal regions showing less or relatively less activity post-switch.

We considered multiple potential confounds to this analysis. First, we may have failed to fully regress the effects of task and reward. More ventral regions such as OFC and dACC have been particularly strongly associated with reward processing in the literature.[11,43,44] Thus, the relatively higher activity evident after action switches may have been due to engagement with the task. Second, there is a large imbalance in possible action transitions (see Figure 4). Although this is why we opted for the segment-normalized activity, large imbalances may nevertheless affect the result. To control for these possibilities, we did two things. First, we considered action transitions only in periods of time when subjects were not engaged with the task for at least 1 s. Second, we calculated a weighted average of activity, whereby we first calculated *action-specific* peri-switch activity for each possible post-switch action and then averaged across these. After performing these controls, we found a similar pattern of results, namely, that there was a greater increase in post-switch activity in more ventral regions such as the OFC and a more modest (or non-existent) change in more dorsal/posterior regions such as the PMd (Figure 7C). We also found that the directionality of results was more consistent between subjects, particularly for PMd, SMA, and dlPFC. Thus, although there was some variability in the degree of per-switch activity across subjects and controls, the relative pattern of results—with more ventral regions showing increased activity after the switch—was consistent. Finally, we note that the specific pattern observed may be driven by a variety of factors, such as a reduction or increase in velocity at the time of switching; however, the main finding here, the difference between the different brain areas, and the systematic relation between hierarchical position and firing rate pattern, does not suffer from this confound.

DISCUSSION

We developed an automated behavioral classification approach and used it to investigate the freely moving behavior and neural correlates of behavior in two rhesus macaques performing a foraging task. Our analysis pipeline used a set of kinematic features calculated over short time windows and a dimensionality

reduction approach to cluster and identify distinct actions. We found that behavior in this task comprises 59 distinct actions. On the basis of transition probabilities between actions, we were then able to determine that the actions were organized into action modules, and we were able to further delineate the full hierarchical structure of this space of actions. We showed that responses of neurons in six brain regions encoded actions, with greater encoding in more dorsal structures. We found that neurons show a characteristic signature of action switching (a dip in activity aligned to the time of the switch) and that the pattern of switch-aligned response varies systematically with dorsoventral position, in the direction opposite the action encoding.

Perhaps the most striking finding is the two oppositely pointed hierarchies in the prefrontal cortex. Several observers have proposed that the prefrontal cortex is characterized by a ventral-to-dorsal functional gradient in which information is transformed from abstract to concrete and from conceptual to motor.[1–7,38–40,45,46] Partially supporting this idea, we have shown stronger task,[35] navigational,[36] and action-specific (this paper) representations in more dorsal structures. However, we have not previously identified any variables that have stronger patterns of modulation in more ventral structures. We do that here. Specifically, we show that more ventral regions have a stronger task-switch signal. This signal is common across different switch types, suggesting it either is amodal or has an amodal component. This result, then, is consistent with other proposed task-switching signals, such as those observed in prefrontal and parietal circuits.[47–53] One explanation for these results may be that the switch reflects the coordination of activity across multiple regions, but with the bulk of the relevant cognitive processing most ventrally; another explanation would be that the switch signals are more latent in more dorsal structures, perhaps because the switch signals regulate ongoing hierarchically intermediate neural activity. The present results also relate to many ongoing debates in the literature, including those relating to the OFC. We believe that our results are most aligned with emerging theories that emphasize the functional continuity between the OFC and the rest of the prefrontal cortex, rather than its functional distinctness. Such studies typically emphasize the role of the OFC as the start or head of a hierarchical process that transforms goals into actions.

Some recent work has begun to unravel the neural processes that occur during free movement.[16–18,54] Our work builds on this past work in several ways. First, we were able to measure neural correlates of actions in an animal with a richer behavioral repertoire (59 distinct actions). Second, we are able to draw conclusions about the organization of the primate prefrontal cortex, whose homology to other species remains debated.[55–57] Third, and relatedly, because of our ability to record in six different regions across most of the prefrontal cortex, we were able to link encoding of action and action switching to a specific prefrontal hierarchy. Fourth, our work is performed in macaques, which have several features that distinguish them from rodents, including greater relevance to humans due to similar behavioral repertoires and more homologous prefrontal anatomy. For example, behavior in macaques is likely to be more directly interpretable with regard to human behavior than behavior in rodents. Moreover, macaques are especially important in this regard because of their pivotal role as a model organism for biomedical research.[58,59]

To what extent is the hierarchical pattern of behavior we observe an artifact of the fact that our subjects are performing a structured task? A previous paper by our group suggests that the hierarchical pattern is intrinsic to macaque behavior and is not imposed by the task itself.[30] Specifically, in that study, we contrasted behavior in two different tasks and in a passive "no task" condition. We found clear evidence of organization in all three task contexts, although, interestingly, the particular organization was different in the three. Ultimately, these results indicate that hierarchical organization of action is unlikely to be a result of our tasks, but the particular organization we observe may be.

Recently, several studies of simpler animals have established the "physics of behavior," namely, the core elements from which behavior is composed and the rules that govern how those elements combine.[14] So far, this approach has not been applied to non-human primates; instead, the study of the behavior of non-human primates has largely been limited to highly constrained and simple motoric experiments (such as reaching and saccades) or otherwise relied on simple heuristics to delineate gross motivational state.[60] Our study builds on this work to discover specific actions via an automatic and unsupervised approach and delineates their organizational principles. While it is not surprising that behavior is hierarchically organized, it is important to have a specific map of the important behaviors.

One of the greatest potential benefits for statistical analysis of highly quantified behavior is in the prospect of automated ethogramming.[20,31,61] By ethogramming, we mean the classification of pose sequences into specific behavior into ethologically meaningful categories such as walking, foraging, grooming, and sleeping. Currently, constructing an ethogram requires the delineation of ethogrammatical categories, which involves the time-consuming and careful annotation of behavior by highly trained human observers.[54] Human-led ethogramming is slow, extremely costly, error prone, and susceptible to characteristic biases.[61–64] For these reasons, it is simply impractical for even moderately large datasets, collected either in an open environment or in the home cage.[65] These kinds of datasets require automated alternatives. Automated ethogramming requires both high-quality behavioral tracking and novel methods applied to tracked data that result in detection of meaningful categories. Such techniques have not, until recently, existed for primates.[31] Our methods take the raw information needed for ethogramming—pose data—and infer actions and higher-level categories from it. As such, they provide the first step toward automated ethogramming in primates. There are a number of steps that may be taken in the future to increase and refine the repertoire of discovered behaviors, including improvements in image capture resolution, more accurate and robust models of pose, and more varied experimental designs that induce more varied behaviors. We are particularly optimistic about the potential benefits of ethogramming for systems neuroscience. Relating behavior to neural circuits and networks is an important goal in the field, so being able to quantify behavior more rigorously,

without sacrificing freedom of movement or naturalness, is likely to be invaluable for future studies.

Limitations of the study

One limitation of this work is that, while we find neural correlates of action, we do not investigate the specific type of information carried in these signals. It may be that neural activity specifies the details of action in a very detailed way, such as describing where each limb should be and what joint angle should be generated.[66] Conversely, it may be very abstract; it may encode the action, or even entire sequences of action, and let downstream structures specify the details. For example, in a study of *C. elegans*, Marques and colleagues found that sensorimotor tuning accuracy changed markedly with the animal's state.[18] These are critical questions that should be addressed in further study of unrestrained and spontaneous behavior (ideally in a setup such as in the present study), but are distinct from the points of interest of the present work. Specifically, we did not seek to tackle how actions are planned, executed, and monitored, but rather sought to obtain a heuristic comparison of action encoding across the brain. By delineating that actions are more strongly encoded in dorsal structures, whereas more ventral ones have a stronger signal for action switching, this work suggests that future studies may focus on ventral structures for action planning and monitoring and dorsal ones for execution. Another limitation is that we chose only one animal as the focus, to generate embeddings. In doing so, we limit our ability to identify specific idiosyncratic subject-specific actions. We emphasize here that our goal is not to identify such actions or their neural correlates. Instead, our goal is to identify to what extent brain areas encode actions through their neural activity.

STAR★METHODS

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

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
  - Lead contact
  - Materials availability
  - Data and code availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
  - Animal model
- METHOD DETAILS
  - Surgical procedures
  - Electrophysiology
  - Task
  - Pose acquisition
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Pose reconstruction
  - Pose preprocessing
  - Feature engineering
  - Action identification via embedding and clustering
  - Assessing cluster quality
  - Assessing behavioral modularity and hierarchy
  - Matching neural and pose data
  - Action and action module encoding
  - Assessing peri-switch change in firing rate
  - Assessing the hierarchy of encoding across cortex

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.celrep.2023.113091.

ACKNOWLEDGMENTS

We thank the Hayden/Zimmermann lab for valuable discussions. This work was supported by NIH grants R01 MH128177 (to J.Z.), P30 DA048742 (J.Z. and B.H.), R01 MH125377 (B.H.), and NSF 2024581 (J.Z. and B.H.) and a UMN AIRP award (J.Z. and B.H.) from the Digital Technologies Initiative (J.Z. and B.H.) of the Minnesota Institute of Robotics (J.Z.).

AUTHOR CONTRIBUTIONS

Formal analysis, B.V.; investigation, D.J.-N.M., R.L.C., B.V., I.C., and M.Z.; data curation, B.V.; writing, B.V.; supervision, J.Z. and B.H.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: October 28, 2022
Revised: June 23, 2023
Accepted: August 18, 2023

REFERENCES

1. Fuster, J.M. (2001). The prefrontal cortex—an update: time is of the essence. Neuron *30*, 319–333.
2. Fuster, J.M. (2000). Executive frontal functions. Exp. Brain Res. *133*, 66–70.
3. Koechlin, E., Ody, C., and Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. Science *302*, 1181–1185.
4. Yoo, S.B.M., and Hayden, B.Y. (2018). Economic choice as an untangling of options into actions. Neuron *99*, 434–447.
5. Fine, J.M., and Hayden, B.Y. (2022). The whole prefrontal cortex is premotor cortex. Philos. Trans. R. Soc. Lond. B Biol. Sci. *377*, 20200524.
6. Merel, J., Botvinick, M., and Wayne, G. (2019). Hierarchical motor control in mammals and machines. Nat. Commun. *10*, 5489.
7. Wilson, C.R.E., Gaffan, D., Browning, P.G.F., and Baxter, M.G. (2010). Functional localization within the prefrontal cortex: missing the forest for the trees? Trends Neurosci. *33*, 533–540.
8. Gallego, J.A., Makin, T.R., and McDougle, S.D. (2022). Going beyond primary motor cortex to improve brain–computer interfaces. Trends Neurosci. *45*, 176–183.
9. Ebbesen, C.L., and Brecht, M. (2017). Motor cortex—to act or not to act? Nat. Rev. Neurosci. *18*, 694–705.
10. Shenoy, K.V., Sahani, M., and Churchland, M.M. (2013). Cortical control of arm movements: a dynamical systems perspective. Annu. Rev. Neurosci. *36*, 337–359.
11. Heilbronner, S.R., and Hayden, B.Y. (2016). Dorsal anterior cingulate cortex: a bottom-up view. Annu. Rev. Neurosci. *39*, 149–170.
12. Calhoun, A.J., and Murthy, M. (2017). Quantifying behavior to solve sensorimotor transformations: advances from worms and flies. Curr. Opin. Neurobiol. *46*, 90–98.
13. Gomez-Marin, A., and Ghazanfar, A.A. (2019). The life of behavior. Neuron *104*, 25–36.

14. Brown, A.E.X., and De Bivort, B. (2018). Ethology as a physical science. Nat. Phys. *14*, 653–657.

15. Mao, D., Avila, E., Caziot, B., Laurens, J., Dickman, J.D., and Angelaki, D.E. (2021). Spatial modulation of hippocampal activity in freely moving macaques. Neuron *109*, 3521–3534.e6.

16. 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., and Datta, S.R. (2018). The striatum organizes 3D behavior via moment-to-moment action selection. Cell *174*, 44–58.e17.

17. Calhoun, A.J., Pillow, J.W., and Murthy, M. (2019). Unsupervised identification of the internal states that shape natural behavior. Nat. Neurosci. *22*, 2040–2049.

18. Marques, J.C., Li, M., Schaak, D., Robson, D.N., and Li, J.M. (2020). Internal state dynamics shape brainwide activity and foraging behaviour. Nature *577*, 239–243.

19. Mathis, M.W., and Mathis, A. (2020). Deep learning tools for the measurement of animal behavior in neuroscience. Curr. Opin. Neurobiol. *60*, 1–11.

20. Pereira, T.D., Shaevitz, J.W., and Murthy, M. (2020). Quantifying behavior to understand the brain. Nat. Neurosci. *23*, 1537–1549.

21. Sturman, O., von Ziegler, L., Schläppi, C., Akyol, F., Privitera, M., Slominski, D., Grimm, C., Thieren, L., Zerbi, V., Grewe, B., and Bohacek, J. (2020). Deep learning-based behavioral analysis reaches human accuracy and is capable of outperforming commercial solutions. Neuropsychopharmacology *45*, 1942–1952.

22. Hsu, A.I., and Yttri, E.A. (2020). B-SOiD: An Open Source Unsupervised Algorithm for Discovery of Spontaneous Behaviors. Preprint at bioRxiv. https://doi.org/10.1101/770271.

23. Bohnslav, J.P., Wimalasena, N.K., Clausing, K.J., Dai, Y.Y., Yarmolinsky, D.A., Cruz, T., Kashlan, A.D., Chiappe, M.E., Orefice, L.L., Woolf, C.J., and Harvey, C.D. (2021). DeepEthogram, a machine learning pipeline for supervised behavior classification from raw pixels. Elife *10*, e63377.

24. Marks, M., Qiuhan, J., Sturman, O., von Ziegler, L., Kollmorgen, S., von der Behrens, W., Mante, V., Bohacek, J., and Yanik, M.F. (2021). Deep-learning Based Identification, Pose Estimation and End-To-End Behavior Classification for Interacting Primates and Mice in Complex Environments. Preprint at bioRxiv. https://doi.org/10.1101/2020.10.26.355115.

25. Bain, M., Nagrani, A., Schofield, D., Berdugo, S., Bessa, J., Owen, J., Hockings, K.J., Matsuzawa, T., Hayashi, M., Biro, D., et al. (2021). Automated audiovisual behavior recognition in wild primates. Sci. Adv. *7*, eabi4883. https://doi.org/10.1126/sciadv.abi4883.

26. Dunn, T.W., Marshall, J.D., Severson, K.S., Aldarondo, D.E., Hildebrand, D.G.C., Chettih, S.N., Wang, W.L., Gellis, A.J., Carlson, D.E., Aronov, D., et al. (2021). Geometric deep learning enables 3D kinematic profiling across species and environments. Nat. Methods *18*, 564–573.

27. Bala, P.C., Eisenreich, B.R., Yoo, S.B.M., Hayden, B.Y., Park, H.S., and Zimmermann, J. (2020). Automated markerless pose estimation in freely moving macaques with OpenMonkeyStudio. Nat. Commun. *11*, 4560.

28. Bala, P., Zimmermann, J., Park, H.S., and Hayden, B.Y. (2023). Self-supervised Secondary Landmark Detection via 3D Representation Learning. Int. J. Comput. Vis. *131*, 1980–1994.

29. Labuguen, R., Matsumoto, J., Negrete, S.B., Nishimaru, H., Nishijo, H., Takada, M., Go, Y., Inoue, K.I., and Shibata, T. (2020). MacaquePose: A novel "in the wild" macaque monkey pose dataset for markerless motion capture. Front. Behav. Neurosci. *14*, 581154.

30. Voloh, B., Eisenreich, B.R., Maisson, D.J.N., Ebitz, R.B., Park, H.S., Hayden, B.Y., and Zimmermann, J. (2023). Hierarchical organization of rhesus macaque behavior. Oxf. Open Neurosci. *2*, kvad006.

31. Hayden, B.Y., Park, H.S., and Zimmermann, J. (2022). Automated pose estimation in primates. Am. J. Primatol. *84*, e23348.

32. Marshall, J.D., Aldarondo, D.E., Dunn, T.W., Wang, W.L., Berman, G.J., and Ölveczky, B.P. (2021). Continuous whole-body 3D kinematic recordings across the rodent behavioral repertoire. Neuron *109*, 420–437.e8.

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.

34. 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.

35. Maisson, D.J.N., Cash-Padgett, T.V., Wang, M.Z., Hayden, B.Y., Heilbronner, S.R., and Zimmermann, J. (2021). Choice-relevant information transformation along a ventrodorsal axis in the medial prefrontal cortex. Nat. Commun. *12*, 4830.

36. 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. *In press*.

37. McInnes, L., Healy, J., and Melville, J. (2018). Umap: Uniform Manifold Approximation and Projection for Dimension Reduction. Preprint at arXi. https://doi.org/10.48550/arXiv.1802.03426.

38. Hunt, L.T., and Hayden, B.Y. (2017). A distributed, hierarchical and recurrent framework for reward-based choice. Nat. Rev. Neurosci. *18*, 172–182.

39. Hunt, L.T., Behrens, T.E.J., Hosokawa, T., Wallis, J.D., and Kennerley, S.W. (2015). Capturing the temporal evolution of choice across prefrontal cortex. Elife *4*, e11945.

40. Badre, D., and D'esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? Nat. Rev. Neurosci. *10*, 659–669.

41. Bonald, T., Charpentier, B., Galland, A., and Hollocou, A. (2018). Hierarchical Graph Clustering Using Node Pair Sampling. Preprint at arXiv.:1806.01664.

42. Dasgupta, S. (2016). A cost function for similarity-based hierarchical clustering. In Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, pp. 118–127.

43. Wallis, J.D. (2007). Orbitofrontal cortex and its contribution to decision-making. Annu. Rev. Neurosci. *30*, 31–56.

44. Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model. Annu. Rev. Neurosci. *34*, 333–359.

45. Badre, D. (2008). Cognitive control, hierarchy, and the rostro–caudal organization of the frontal lobes. Trends Cognit. Sci. *12*, 193–200.

46. Choi, E.Y., Drayna, G.K., and Badre, D. (2018). Evidence for a functional hierarchy of association networks. J. Cognit. Neurosci. *30*, 722–736.

47. Dove, A., Pollmann, S., Schubert, T., Wiggins, C.J., and von Cramon, D.Y. (2000). Prefrontal cortex activation in task switching: an event-related fMRI study. Brain Res. Cogn. Brain Res. *9*, 103–109.

48. Sohn, M.H., Ursu, S., Anderson, J.R., Stenger, V.A., and Carter, C.S. (2000). Inaugural article: the role of prefrontal cortex and posterior parietal cortex in task switching. Proc. Natl. Acad. Sci. USA *97*, 13448–13453.

49. Nakahara, K., Hayashi, T., Konishi, S., and Miyashita, Y. (2002). Functional MRI of macaque monkeys perform- ing a cognitive set-shifting task. Science *295*, 1532–1536.

50. Rushworth, M.F.S., Hadland, K.A., Paus, T., and Sipila, P.K. (2002). Role of the human medial frontal cortex in task switching: a combined fMRI and TMS study. J. Neurophysiol. *87*, 2577–2592.

51. Brass, M., Ullsperger, M., Knoesche, T.R., von Cramon, D.Y., and Phillips, N.A. (2005). Who comes first? The role of the prefrontal and parietal cortex in cognitive control. J. Cognit. Neurosci. *17*, 1367–1375.

52. Crone, E.A., Wendelken, C., Donohue, S.E., and Bunge, S.A. (2006). Neural evidence for dissociable components of task-switching. Cerebr. Cortex *16*, 475–486.

53. Hyafil, A., Summerfield, C., and Koechlin, E. (2009). Two mechanisms for task switching in the prefrontal cortex. J. Neurosci. *29*, 5135–5142.

54. Testard, C., Tremblay, S., Parodi, F., DiTullio, R.W., Acevedo-Ithier, A., Gardiner, K., Kording, K., and Platt, M. (2023). Neural Signatures of Natural Behavior in Socializing Macaques. Preprint at bioRxiv. https://doi.org/10.1101/2023.07.05.547833.

55. Heilbronner, S.R., Rodriguez-Romaguera, J., Quirk, G.J., Groenewegen, H.J., and Haber, S.N. (2016). Circuit-based corticostriatal homologies between rat and primate. Biol. Psychiatr. *80*, 509–521.

56. Laubach, M., Amarante, L.M., Swanson, K., and White, S.R. (2018). What, if anything, is rodent prefrontal cortex? eneuro *5*. ENEURO.0315.

57. Passingham, R.E., and Wise, S.P. (2012). The Neurobiology of the Prefrontal Cortex: Anatomy, Evolution, and the Origin of Insight (Oxford: OUP).

58. Buffalo, E.A., Movshon, J.A., and Wurtz, R.H. (2019). From basic brain research to treating human brain disorders. Proc. Natl. Acad. Sci. USA *116*, 26167–26172.

59. Rudebeck, P.H., Rich, E.L., and Mayberg, H.S. (2019). From bed to bench side: Reverse translation to optimize neuromodulation for mood disorders. Proc. Natl. Acad. Sci. USA *116*, 26288–26296.

60. Shahidi, N., Schrater, P., Wright, A., Pitkow, X., and Dragoi, V. (2021). Population Coding of Strategic Variables during Foraging in Freely-Moving Macaques. Preprint at BioRxiv. 811992.

61. Anderson, D.J., and Perona, P. (2014). Toward a science of computational ethology. Neuron *84*, 18–31.

62. Kardish, M.R., Mueller, U.G., Amador-Vargas, S., Dietrich, E.I., Ma, R., Barrett, B., and Fang, C.C. (2015). Blind trust in unblinded observation in ecology, evolution, and behavior. Front. Ecol. Evol. *3*, 51.

63. Holman, L., Head, M.L., Lanfear, R., and Jennions, M.D. (2015). Evidence of experimental bias in the life sciences: Why we need blind data recording. PLoS Biol. *13*, e1002190.

64. Tuyttens, F.A.M., de Graaf, S., Heerkens, J., Jacobs, L., Nalon, E., Ott, S., Stadig, L., Van Laer, E., and Ampe, B. (2014). Observer bias in animal behaviour research: Can we believe what we score, if we score what we believe? Anim. Behav. *90*, 273–280.

65. Womelsdorf, T., Thomas, C., Neumann, A., Watson, M.R., Banaie Boroujeni, K., Hassani, S.A., Parker, J., and Hoffman, K.L. (2021). A Kiosk Station for the Assessment of Multiple Cognitive Domains and Cognitive Enrichment of Monkeys. Front. Behav. Neurosci. *15*, 721069.

66. Fetz, E.E. (1992). Are movement parameters recognizably coded in the activity of single neurons? Behavioral and brain sciences *15*, 679–690.

67. Azab, H., and Hayden, B.Y. (2018). Correlates of economic decisions in the dorsal and subgenual anterior cingulate cortices. Eur. J. Neurosci. *47*, 979–993.

68. Yoo, S.B.M., Tu, J.C., Piantadosi, S.T., and Hayden, B.Y. (2020). The neural basis of predictive pursuit. Nat. Neurosci. *23*, 252–259.

69. Blanchard, T.C., Wolfe, L.S., Vlaev, I., Winston, J.S., and Hayden, B.Y. (2014). Biases in preferences for sequences of outcomes in monkeys. Cognition *130*, 289–299.

70. Wang, M.Z., Hayden, B.Y., and Heilbronner, S.R. (2022). A structural and functional subdivision in central orbitofrontal cortex. Nat. Commun. *13*, 3623.

71. Yoo, S.B.M., Hayden, B.Y., and Pearson, J.M. (2021). Continuous decisions. Philos. Trans. R. Soc. Lond. B Biol. Sci. *376*, 20190664.

72. Schonberger, J.L., and Frahm, J.M. (2016). Structure-from-motion revisited. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4104–4113.

73. Davies, D.L., and Bouldin, D.W. (1979). A cluster separation measure. IEEE Trans. Pattern Anal. Mach. Intell. *1*, 224–227.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Preprocessed data	Zenodo	https:doi.org/10.5281/zenodo.8067596
Software and algorithms		
MATLAB analysis code	Zenodo	https://doi.org/10.5281/zenodo.8067596

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, Jan Zimmermann (janz@umn.edu).

Materials availability

The study did not generate new unique reagents.

Data and code availability

- The processed data reported in this paper has been deposited on Zenodo and is publicly available as of the date of publication. The DOI is listed in the key resources table.
- All original code has been deposited on Zenodo and is publicly available as of the date of publication. The DOI is listed in the key resources table.
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Animal model

Two male rhesus macaques (*Macaca mulatta*) served as subjects. Subjects were habituated to laboratory conditions, trained to enter and exit an open arena, and then trained to operate water dispensers. The University Committee on Animal Resources at the University of Minnesota approved all animal procedures. Animal procedures were designed and conducted in compliance with the Public Health Service's Guide for the Care and Use of Animals and approved by the institutional animal care and use committee (IACUC) of the University of Minnesota.

METHOD DETAILS

Surgical procedures

We placed a cranium adherent form-fitted Gray Matter (Gray Matter Research) recording chamber and a 128-channel microdrive recording system over the area of interest. We used CT scans which were compared to corresponding CT studies performed following surgical implantation and placement of the electrodes. The hyperdense appearance of electrodes using the post-recording CT allowed us to verify that electrodes followed the trajectory of the pre-operative plan. In both cases (that is, for both subjects), the results of the CT show that this was the case. Specifically, these results indicate that our procedures provided co-registration with an error of less than 0.5 mm. Since electrodes were advanced incrementally not all locations are verified using this approach, but positions could be validly inferred using interpolation. In addition, we performed a second complementary method that also confirmed placement of the electrodes. Specifically, as we have done in many past studies, as we moved the electrodes down into the brain, we made note of the auditorily detectable change from gray to white matter, and reconciled this information with our preoperative CT scans. In all cases, the two matched. Animals received analgesics and antibiotics after all procedures. Procedures were designed and conducted in compliance with the Public Health Service's Guide for the Care and Use of Animals and approved by the institutional animal care and use committee (IACUC) of the University of Minnesota.

Electrophysiology

Recordings were made with a 128-channel microdrive system (Gray Matter), targeting a wide swath of the prefrontal cortex ranging from OFC to PMd. Each electrode was independently moveable along the depth dimension. Neural recordings were acquired with a wireless datalogger (HH128; SpikeGadgets). The datalogger was triggered to start recording with a wireless RF transceiver (SpikeGadgets), and periodically received synchronization pulses. Data were recorded at 30 kHz, stored on a memory card for the duration of the experiment, and then offloaded after completion of the session. Each feeder had local code running the experiment. Task events triggered a TTL pulse, as well as a wireless event code. A dedicated PC running custom code controlled all feeders, and aggregated event codes. Syncing of all data sources as accomplished via the Main Control Unit (MCU; SpikeGadgets), which received dedicated inputs from the pose acquisition system (see below), and feeders. Recording sessions were initiated and controlled by Trodes software (SpikeGadgets). After neural recordings were offloaded, they were synced with other sources of data via the DataLogger GUI (SpikeGadgets).

Recordings were performed for 4-6 days weekly for a period of 4-6 months. For an initial period of 2-4 weeks, we lowered up to 10 electrodes in each session until each had punctured the dura and their position was well-within cortex as seen from the fMRI reconstruction. Subjects still performed experiments during this time, but as the signal was noisy, no recordings were performed during this time.

A typical recording day consisted of multiple stages, including electrode adjustment, an experimental session, and extraction of the recorded signal. For the duration of the experiment, on each day, we tracked yields on each electrode and visually assessed the quality of the signal. If an electrode had poor yields for up to 5 days in a row, we would lower it up to 1 mm (or more if it was intended to move to a new area).

Spike sorting was performed in a semi-supervised fashion. After an automatic phase which extracted putative units, lab members then curated the output. Units were retained based on waveform shape, inter-spike interval distribution, and temporal stability. All units recorded over days on the same electrodes were treated as independent in subsequent analysis of this manuscript. While there is a risk that this approach slightly inflates the number of independent neurons measured there is currently no way to definitively test whether the signal recorded over consecutive days stems from the same or different neurons.

Subdivisions of the brain were collapsed to anatomical areas, listed below as defined in the D99 parcellation of the NMT atlas (Saleem et al., 2021):

- dACC: 24a', 24a, 24b, 24b', 24c, 24c'
- vlPFC: 45a, 45b, 46d, 46v, 46f, 12r
- dlPFC: 8ad, 8bd, 8av, 8bs, 9d, 8bm, 9m
- SMA: F3, F6
- PMd: F1, F2, F5, F7, F4
- OFC: 13b, 13m, 13l, 12l, 12m, 12o, 11l, 11m

Task

There were four water dispensing stations (''patches'') available with programmed delivery schedules. Each patch delivered a fixed amount of 1.5 mL per lever press. The first four presses, regardless of patch sequencing, were rewarded with water delivery. The fifth lever press was unrewarded and led to a 3-minute deactivation of the patch. That is, animals could press fewer than five times, leave to engage with a different patch, and return to the same patch in the state they left it. No reset or deactivation was applied if the animal left the patch. A patch was only reset if the subject pressed the lever a fifth time and waited 3 minutes for it to reactivate. Each rewarded press followed the same programmed sequence.

An activated patch was indicated by a fully blue display. A lever press changed the display to white with a green plus-sign in the center, an auditory cue was played, and a solenoid opened to dispense reward. After dispensing the solenoid closed, the auditory cue ended and the green plus-sign disappeared. The screen remained white for two additional seconds before the screen turned blue again to indicate the availability of another lever press. The fifth lever press was instead followed by the screen immediately turning white, with no visual or auditory reward cue and no water delivery. However, it is important to note that task engagement was not required; a subject could choose not to engage with the patches for the entirety of the session.

Otherwise, the measured behavior was simply the free movement of the subject through the arena. On average, subjects pressed levers 136 ± 8 times (subject Y: 166 ± 7, subject W: 107 ± 9), amounting to roughly 204 mL of water, per session, received for interacting with patches. Prior and concurrent chaired task training of these subjects included several standard chaired laboratory tasks.[67–71]

Pose acquisition

Images were captured with 62 cameras (Blackfly, FLIR), synchronized via a high-precision pulse generator (Aligent 33120A) at a rate of 30 Hz. The cameras were positioned to ensure coverage of the entire arena, and specifically, so that at least 10 cameras captured the subject with high-enough resolution for subsequent pose reconstruction, regardless of the subject's position and pose. Images were streamed to one of 6 dedicated Linux machines. The entire system produced about six TB of data for a two hour session. After data acquisition, the data were copied to an external drive for processing on a dedicated Linux server (Lambda Labs).

To calibrate the camera's geometries for pose reconstruction, a standard recording session began with a camera calibration procedure. A facade of non-repeating visual patterns (mixed art works and comic strips) was wrapped around two columns of barrels placed at the center of the room, and images of this calibration scene were taken from all 62 cameras. These images were used to calibrate the camera geometry (see below). This setup was then taken down, and the experiment began.

QUANTIFICATION AND STATISTICAL ANALYSIS

Pose reconstruction

We first extracted parameters relating to the cameras' geometry for the session. To this end, we used a standard structure-from-motion algorithm (*colmap*[72]) to reconstruct the space containing the 3D calibration object and 62 cameras from the calibration images, as well as determine intrinsic and extrinsic camera parameters. We first prepared images by subtracting the background from each image in order to isolate the subject's body. Then, 3D center-of-mass trajectories were determined via random sample consensus (RANSAC). Finally, the 3D movement and subtracted images were used to select and generate a set of maximally informative cropped images, such that the subject's entire body was encompassed. To reduce the chance that the tire swing would bias pose estimation, we defined a mask of pixels to ignore that encompassed the tire's swinging radius.

Next, we inferred 3D landmark positions using a trained convolutional pose machine (CPM[27]) and an additional approach that allowed to augmentation of the original thirteen landmarks with two additional ones.[28] The augmented reconstruction resulted in 15 annotated landmarks for each image. We then used these landmark positions to identify poses (see above). We used a loss function that incorporated physical constraints (such as preserving limb length, and temporal smoothness) to refine landmark localization. We found residual variability in limb length across subjects after reconstruction, between subjects, particularly for the arm, resulting in poses that were highly specific to individual subjects. To prevent subject-specific limb lengths from biasing subsequent behavior identification, we augmented the original 13 inferred landmarks to include two new ones (positions of left and right elbows) using a supplementary trained CPM model.[27] Thus, the augmented reconstruction resulted in 15 annotated landmarks for each image.

Pose preprocessing

We applied a number of smoothing and transformation steps to the 3D pose data. First, we transformed the reconstructed space to a reference space that was measured using the Optitrack system.[27] Then, we ignored any frame where a limb was outside the bounds of the cage due to poor reconstruction, or residual frames where subject poses were still subject to collapse (defined as where the mean limb length < 10 cm). Next, we interpolated over any segments of missing data (lasting at most 10 frames, or 0.33 sec) using a piecewise cubic interpolation. After this, we also removed any segments lasting less than 30 seconds. Note that only a small number of frames were removed after this whole procedure; specifically, 0.64% of frames on average were ignored. Finally, we rescaled any limbs in frames where the limb length was >3 standard deviations above the session mean to be at most 3 SDs long. Finally, data was smoothed with a Hampel (median) filter over 5 samples.

Feature engineering

To discover behaviors, we next extracted a number of informative kinematic features, calculated over short time windows. Relevant parameters were calculated for one session from subject Y that exhibited a wide range of behaviors (similar results were obtained using other sessions), and then applied to every other session from both subjects.

Some features were in the *world-reference frame*. Each of these features, unless otherwise noted, was smoothed with gaussian of length 10, 30, and 60 frames (0.3, 1, and 2 sec.), resulting in a total of 8x3=24 features.

- *Speed*: Calculated as the absolute of the numerical derivative of the center-of-mass (COM; defined as the midpoint between the hip and neck landmark) of the subject.
- *Ground Speed*: COM speed but just along the ground plane
- *Height Speed*: COM speed but just along the gravity dimension
- *Height Velocity*: Numerical derivative of COM along the gravity dimension
- *Perpendicularity*: scalar representing how vertically oriented the animal was. Calculated as the norm of the cross product of the spine (hip to neck vector) and gravity dimension. A value of 0 means the subject was vertically oriented, while 1 represents horizontal orientation.
- *Height*
- *Limb-speed variability (height)*: helps to differentiate ballistic from non-ballistic movements (e.g jumping). Obtained by calculating the speed of major landmarks (right hand, left hand, right foot, left foot, hip, neck), and then, for each frame, the standard deviation among these.
- *Limb-speed variability (ground)*: same as above, but calculating speed along the ground plane.

Next, we calculated a set of features in the *reference frame of the subject*. To this end, we normalized the orientation of poses on individual frames in a two step procedure; first, by aligning poses to face the same direction, and second, by adding back rotation corresponding to the perpendicularity of the spine relative to gravity. First, we transformed each pose to face a common direction. To

do this rotation, we first defined two vectors, one corresponding to the spine (neck to hip landmarks), and the other to the expanse of the shoulders (left and right shoulder landmarks, which was then centered on the neck landmark). Poses were then rotated such that the plane defined by these vectors faced the same direction (in essence, so that the torso faced the same direction). Next, we rotated poses such that the spine had the same original orientation; in other words, if the first step aligned poses where the spine was either perpendicular or parallel to face the same direction, this step undid that rotation. We found the angle of rotation by comparing the spine (hip to neck) vector with that of gravity, and then rotated poses along the sagittal plane (ie splitting the subjects body into left and right). After this procedure, poses were aligned to the same direction, but with their original orientation of the spine (e.g if sitting or walking, poses would face in the same direction, but in the first case, the spine is vertical, while in the second, it is horizontal). From these poses, we calculated the following 67 features:

- Landmarks: we calculated the major sources of variation of the major landmarks (right hand, left hand, right foot, left foot, hip, neck). We performed a principal component analysis (PCA) of the xyz coordinates of each of these landmarks and extracted the top 15 PCs.
- Landmark periodicity: to ascertain periods of periodicity in poses, we obtained a time-frequency decomposition of each of the PCs. We first whitened the time series by performing a smooth differentiation using a Savitsky-Golay filter of order 3 and length 5 (matlab function: *sgolay*). Then, we obtained the time-frequency decomposition using Morlet wavelets via the Fieldtrip toolbox (https://www.fieldtriptoolbox.org/) function *ft_freqanalysis*, 5 cycles and 3 standard deviations of the gaussian. We obtained 100 frequencies, equally and logarithmically spaced between 0.1 and 15 Hz. We then obtained power values, and clipped them at a maximum of the 99.9th percentile the entire session. Finally, we reduced the dimensionality of this representation using PCA with 20 components.
- Segment Length: Segments were defined thus: right hand to right shoulder, left hand to left shoulder, hip to right foot, hip to left foot, neck to hip, right hand to left hand, left hand to left foot, left foot to right foot, and right foot to right hand. Segment lengths were smoothed with gaussian windows of length 10, 30, and 60 frames.
- Segment Length Periodicity: We then obtained the time-frequency representation of the (unsmoothed) segment lengths. Spectral decomposition and dimensionality reduction was performed as described above.

This process resulted in a 91-dimensional time-resolved feature set of pose kinematics. To avoid imbalance effects in embedding due to differently scaled features, we normalized each set of features described above. For each whole set of features, we calculated a robust z-score (using the median and median absolute deviation, instead of the mean and standard deviation).

Action identification via embedding and clustering

We created behavioral maps by embedding the extracted kinematic features into two dimensions. As a first step, for computational efficiency, we first constructed a training set using data from subject Y. Specifically, for each session in subject Y, we constructed the training set by sampling every 6 points (i.e. every 200 ms). We oversampled rare events, either where the instantaneous speed was > 2, or the height of the subject's was > 3. Roughly speaking, these corresponded to ballistic movements (such as jumping), and climbing.

We then found a 2-dimensional embedding of the training set using Uniform Manifold Approximation and Projection (UMAP[37]). We used a euclidean distance metric, and set parameters *min_dist=0.1, n_neighbors=200, and set_op_mix_ratio=0.25*, which we found to be a good balance between separating dissimilar behaviors, while combining similar ones.

To define behavioral clusters, we first estimated the probability density at 200 equally interspersed points both in the first and second UMAP dimensions. This produced a smoothed map of the pose embeddings, with clearly visible peaks (Figure 2B). We then employed the *watershed algorithm* on the inverse of this smoothed map (Berman et al., 2014). This algorithm defines borders between separate valleys in the (inverse of) the embedding space. Thus, the algorithm determines sections of the embedding space with clearly delineated boundaries (i.e. clusters). Samples were then assigned a cluster label according to where they fell within these borders. Using this procedure, we found a total of 59 behavioral clusters.

Finally, we re-embedded and assigned cluster labels to every sample in both subjects. We performed a k-nearest neighbors (KNN), finding the 20 nearest neighbors (again using Euclidean distance) of each sample in both subjects to that of the training set. The position of each sample in the embedding space was the median of the 20 nearest neighbors in the training samples. On the basis of the position in UMAP space, we assigned each sample a cluster label. The KNN search was performed using the *faiss* library with GPU acceleration (Facebook; github.com/facebookresearch/faiss).

Assessing cluster quality

We next assessed the quality of the extracted clusters by leveraging the *Davies-Bouldin index* (DBI[73]). The DBI compares the intra-cluster dispersion to between-cluster distances. Intra-cluster dispersion is determined by computing the average distance of each sample of vectors to their clusters' centroid. These average distances are then compared to the distance between cluster centroid. The DBI is then a ratio of all intra-cluster dispersions vs between-cluster distance. Thus, a lower DBI indicates that samples within a cluster are more tightly related to one another than to samples outside of the cluster.

We leveraged this metric in two ways. First, we determined if on a global scale, the our procedure clustered samples better than expected by chance. To this end, we first computed the DBI for each individual dataset. We then compared it to a chance distribution by randomly permuting cluster labels and calculating a randomized DBI. This was performed 20 times. We determined significance by comparing the observed DBI across datasets to this chance distribution.

Next, we determined if clusters were well-separated locally. We used a similar procedure as above; however, instead of computing DBI (and chance distribution) over all clusters within a dataset, we computed it for each pair of clusters separately. We report on the proportion of such cluster-pairs where observed DBI was greater than the randomized DBI.

Assessing behavioral modularity and hierarchy

To discover how postures are organized, we employed a hierarchical clustering algorithm named *Paris*[41] using the sknetwork library (scikitnetwork.readthedocs.io/). This algorithm employs a distance metric based on the probability of sampling node pairs and performs agglomerative clustering. Paris requires no user-defined parameters (as opposed to another popular graph clustering algorithm, Louvain, which can perform hierarchical clustering according to a user-supplied resolution parameter). It is equivalent to a multi-resolution version of the Louvain algorithm. The result of this algorithm is a dendrogram describing the relation between different action transitions (which we will refer to as the behavioral dendrogram). To segment action transitions into modules, we determined the modularity score (see below) for different cuts of each dendrogram for n=2 to 58 clusters). We then determined module assignment by cutting the behavioral dendrogram where the modularity score was maximized

We leveraged two important graph-theoretic metrics to assess behavioral composition:

- Modularity Score: The modularity score describes the degree to which postures transition within, rather than between, modules. Transition probability matrices with high modularity scores exhibit a high probability of transitions within modules, but not between modules. Modularity was calculated with the matlab function ‘‘modularity.m’’.
- Dasgputa Score: To assess whether the graph defined by posture transitions truly reflected hierarchical organization, we calculated the Dasgputa Score (Dasgupta, 2016). The Dasgupta Score is a normalized version of the Dasgupta Cost, which defines the cost of constructing a particular dendrogram, given distances between nodes. The Dasgupta Score thus provides quantification of the quality of the hierarchical clustering. We calculated this score using the function ‘‘dasgupta_score’’ in the *sknetwork* library.

Matching neural and pose data

To assess how neural activity co-varied with pose, we computed the spike-density function (SDF) of each neuron and downsampled the activity to match the pose time-series. Specifically, for each neuron, we computed the SDF using the *ft_spikedensity* function in the Fieldtrip toolbox, using a gaussian window of [-100 100] ms, sampled at a rate of 30 Hz.

As the environment allowed for free movement in the presence of a task, we sought to isolate neural activity related to postures. To this end, we regressed out neural activity related to task events at each feeder (*Screen On, Lever Press, Reward On, Reward Cue On,* and *Timeout On*), as well as the XYZ coordinates of the subject (determined as the centre-of-mass of the animals, namely the middle point between the hip and neck landmarks). We carefully selected out regressor variables as those identified, *a priori*, to be likely to cause confounding. We note that in regressing out the lever press, we are reducing our sensitivity to action encoding. However, the lever press is highly confounded with rewards, so any putative encoding of lever press actions would not be distinguishable from encoding of reward, which is known to be a major driver of activity in these regions. We believe that a carefully designed - and necessarily less naturalistic - task could be used to identify the neural correlates of the lever press behavior.

Each regressor was normalized to the range [0 1]. The regressor time series was constructed at a sampling rate of 30 Hz, and smoothed with a boxcar filter of [-25 to +25 ms]. We then fit a Poisson generalized linear model (GLM) with a log link function. This model was used to predict the SDF, and the predicted SDF was subtracted from the observed SDF to obtain a residualized SDF. We excluded any cells from analysis where the Deviance was negative, due to poor model fits. Finally, for missing pose frames, activity was deleted and not analyzed. All neural analyses, unless otherwise mentioned, were performed on this residualized activity.

Action and action module encoding

To assess action encoding, we compared the average (residualized) activity of neurons, factorized by action. First, for each action segment in a session, we found averaged activity in that segment. Then, we performed a Kruskal Wallis test, with average segment activity factorized by action type. We extracted the resultant X^2 (chi-squared) value as a metric of the encoding strength, i.e. the degree to which different actions can be separated from one another according to neuronal activity. We opted for the (non-parametric) Kruskal Wallis test as activity was highly non-normal. We assessed inter-areal differences in encoding strength using an Anova, with area as a factor.

To assess whether individual areas showed significant action encoding, we opted for a randomization test. For each unit, we first shifted the action-series of segment activities by a random amount (at least 25%, and at most 75%, of the length of the action series), and then recalculated the encoding strength. We performed this 20 times. To determine significance in each area individually, we compared the (observed) mean areal encoding strength to that of the randomized distribution.

To assess encoding for action modules, we performed the same series of analyses as outlined above but for action modules. First, in each session, we found n=2-58 action modules, as determined by different cuts of the transition probability clustering. We then remapped the action series labels according to which action module they were a part of. We then used these action module labels to find the encoding strength of each neuron.

Assessing peri-switch change in firing rate

To ascertain if neural activity was related to action switches, we computed the average activity around each switch. For each action switch in a session, we extracted neural activity in a [-1 1] segment centered on the action switch. To avoid times with multiple rapid shifts, we only considered segments where there were no other action switches in a [-0.2 0.2] sec window around the current action switch. Then, for each segment, we found the segment-normalized mean. Specifically, for each segment, we found the median and median absolute deviation of activity in the pre-switch period, and used these values to z-score normalize the entire segment. Then, we averaged this normalized activity to obtain the per-switch segment-normalized activity of each neuron. Finally, for visualization purposes (Figure 7A), we averaged this activity for all units according to area. To determine if and how unit activity changed around the switch, we computed a *switching index*. Specifically, for each neuron, we obtained the average segment-normalized activity in the 1 second before and after the switch. The index was defined as the difference between the post and pre activity. To determine if switching activity differed between areas, we performed a Kruskal Wallis test.

We also considered two possible confounds that may drive our results. First, there was a large imbalance in possible action pairs. As such, per-switch differences we observed may be driven by only the most common action switches. Second, although we took care to regress out activity related to task events, some influence may nevertheless have been retained to lagged responses to task events, or poor model fits. Thus, differences in peri-switch activity may have arisen due to encoding of task events (e.g. of reward). We controlled for these possibilities in two ways. First, to address imbalances in action pairs, we computed a weighted segment-normalized activity. Specifically, we segment-normalized activity for every possible pre-switch action, and then averaged across these to obtain the *weighted segment normalized activity*. Second, to address possible residual encoding of task events, we only considered segments where subjects were not engaged with the feeders. To this end, we extracted a time-series of lever presses, and smoothed this using a 1 second boxcar filter. This gave us a time-series of *engagement probability*. Then, we selected segments where engagement probability was strictly zero. This ensured (because of the filtering step) that we only selected periods of time separated from task engagement by at least 1 sec. Finally, we then performed the same series of analyses on these *controlled segment-normalized activities*.

Assessing the hierarchy of encoding across cortex

Our analyses suggested a gradient of postural encoding, with motor areas exhibiting more specific postural tuning, and more anterior prefrontal areas exhibiting more global encoding of the embedding space. To make concrete this observation, we performed a linear regression, predicting encoding strength corrected KL values as a function of the neurons' location. Specifically, the X (anterior-posterior) and Y(lateral-medial) coordinate was defined as the position of the electrode on the electrode grid, while the Z (dorsal-ventral) position was defined as the total turning depth. This was justified as electrodes were equally spaced in a grid.

To account for possible effects driven by imbalances in recording location, we performed the same analysis after binning encoding value. Specifically, we binned encoding values into a 3-dimensional grid, using 5, 10, 15, or 20 bins along each spatial dimension. We then performed the same regression, predicted (binned) encoding strength from the XYZ bin index.

We performed the same series of analyses on the *switching index* values as well, to ascertain whether changes in firing rate as a result of action switches differed according to recording location.

Cell Reports, Volume *42*

Supplemental information

Hierarchical action encoding in prefrontal cortex of freely moving macaques

Benjamin Voloh, David J.- N. Maisson, Roberto Lopez Cervera, Indirah Conover, Mrunal Zambre, Benjamin Hayden, and Jan Zimmermann

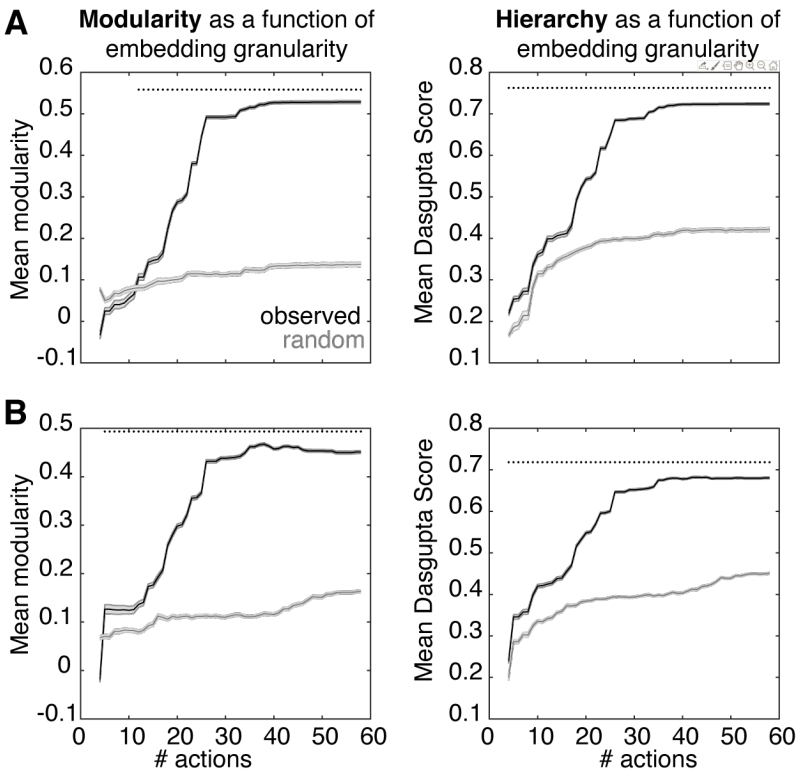

Supplementary Figure 1. Behavioral modularity and hierarchy are above chance even for coarse clustering

(A) Modularity scores after merging actions based on their embedding distances - both for observed (black) and randomized (gray) transition matrices - for subject Y. Dots indicate significance (randomization test, $p<0.05$). Behavioral modularity is above chance even for coarse clustering. (**B)** Same as (A), but plotting the Dasgupta score, a measure of hierarchy. Hierarchy is evident even for coarse clustering. **(C-D)** Same as A-B, but for subject W.