

# B-SOiD: An Open Source Unsupervised Algorithm for Discovery of Spontaneous Behaviors

Alexander I. Hsu<sup>1</sup> and Eric A. Yttri<sup>1</sup> ✉

<sup>1</sup>Department of Biological Sciences, Carnegie Mellon University, Pittsburgh, PA USA

The motivation, control, and selection of actions comprising naturalistic behaviors remains a tantalizing but difficult field of study. Detailed and unbiased quantification is critical. Interpreting the positions of animals and their limbs can be useful in studying behavior, and significant recent advances have made this step straightforward (1, 2). However, body position alone does not provide a grasp of the dynamic range of naturalistic behaviors. Behavioral Segmentation of Open-field In DeepLabCut, or B-SOiD ("B-side"), is an unsupervised learning algorithm that serves to discover and classify behaviors that are not pre-defined by users. Our algorithm segregates statistically different, sub-second rodent behaviors with a single bottom-up perspective video camera. Upon DeepLabCut estimating the positions of 6 body parts (snout, the 4 paws, and the base of the tail), our software performs novel expectation maximization fitting of Gaussian mixture models on t-Distributed Stochastic Neighbor Embedding (t-SNE). The original features taken from dimensionally-reduced classes are then used build a multi-class support vector machine classifier that can decode millions of actions within seconds. We demonstrate that the highly reproducible, independently-classified behaviors can be used to extract kinematic parameters of individual actions as well as broader action sequences. This open-source platform enables the efficient study of the neural mechanisms of spontaneous behavior as well as the performance of disease-related behaviors that have been difficult to quantify, such as grooming and stride-length in Obsessive-Compulsive Disorder (OCD) and stroke research.

Behavioral Analysis | Open-field | DeepLabCut | Unsupervised Learning  
Correspondence: [eyttri@andrew.cmu.edu](mailto:eyttri@andrew.cmu.edu)

## Introduction

The brain has evolved to support generation of action sequences in animals, particularly for surviving constantly changing environments. The core of surviving depends on an animal's ability to correctly anticipate external events, and such decisions can be visualized through sequence of actions that have been adapted to those events for optimal outcome (3). Comprehensive tracking of behavior permits behavioral scientists to not only to quantify behavioral output (4); but also to identify unforeseen motions, such as "tap dancing" in songbirds(5). However, processing such data from recorded behavior is an extremely time and labor intensive process, even with the limited and expensive commercial options available (6).

Recent advances in computer vision and machine learning has accelerated automatic tracking of geometric estimates

of body parts (1, 2). Although establishing the location of body part can be informative given the right experimental configuration, the behavioral interpretability is quite low. For instance, the estimated position of where each paw is in an Open-field does not capture what the animal is doing. Moreover, the angle criterion ascribed to whether a turning movement has occurred, or the minimum duration of an ambulatory bout, are all subjective and could even depend on animal size and/or video capture technique. This uncertainty in behavioral outcome only makes dissecting the neural mechanisms underlying the behavior more complicated. The challenges presented above motivated our investigation into how partnering dimensionality reduction with pattern recognition could create a viable tool in automatic classification of an animal's behavioral repertoire.

In a recent behavioral study, Markowitz and colleagues were able to automatically identify subgroups of animal's behaviors using depth sensors. Their algorithm, MoSeq, utilized the principal components of "spinograms" to identify action groups (7). The authors uncovered the striatal neural correlates of action, consistent with the notion that population of medium spiny neurons (MSNs) encodes certain behavior, and the organization of (8–12). In another seminal study, Klaus et al. employed a unique dimensionality reduction method, t-Distributed Stochastic Neighbor Embedding (t-SNE) (13), to visualize clusters behaviors based on time-varying signals such as speed and acceleration (12). There, they discovered that the t-SNE clusters are easily interpreted as mouse behaviors and was correlated with distinct striatal neural ensembles. The powerful rationale behind selecting t-SNE to project high dimensional features onto low-dimensional space is to preserve the contrast in distributions of the original features, or so called "local structure". Together, these studies suggest that data-driven algorithms can provide reasonable clusters of animal behavior that in turn provide insight into how thoughts generate actions.

Building on this progression, we were motivated to investigate the low-dimensional clusters of a composite of high-dimensional features - such as body length, distance and angle between body parts, occlusion of a body part from view, and speed. In conjunction with an open source platform, DeepLabCut, we provide an open-source unsupervised learning algorithm that integrates the three key layers of understanding (dimensionality reduction, pattern recognition, and machine learning) to enable autonomous

multi-dimensional behavioral classification. We found that behavioral-Segmentation of Open-field in DeepLabCut (B-SOiD) automatically uncovers reasonable behaviors with excellent accuracy (tested against both held out data and blind human observers). It is easily applied across behavioral studies. Moreover, in testing the power of the tool, we have discovered a broader and dynamic behavioral structure - an organization of "what to do next" - exists in mice exploring a novel environment. Finally, we have released our repository B-SOiD on open-source platform GitHub, and have included three versions of our work: manual thresholding of pre-defined behaviors, unsupervised discovery of action classes, and action class modeling, to suit a variety of behavioral research needs.

## Results

A schematic flowchart describing the proposed unsupervised learning algorithm B-SOiD (Fig.1). Upon DeepLabCut estimating the body parts outlining the animal (snout, 4 paws, and the proximal tail), features such as body length, speed and angle can be computed. To visualize high-dimensional feature clusters, we opt to follow a recent study(12) and use t-SNE, a particular dimensionality reduction algorithm that preserves local structure. These feature clusters appears to carry topology from the high-dimensional space; therefore, we hypothesize that the clusters can be grouped using expectation maximization (EM) algorithm by fitting parameters of Gaussian mixture models (GMM). High-dimensional features from each GMM class will then be used to train learners from multi-class support vector machine (SVM) classifier. To validate our performance, we propose to test held-out data against the original cluster assignments over many iterations.

We provide this as an open source toolbox for the neuroscience community who uses DeepLabCut to study animal behavior.

<https://github.com/YttriLab/B-SOiD>

**A. Selection of high dimensional features.** An animal behavior can be parsed into a sequence of changes in physical features. For feature selection, we performed Hierarchical clustering analyses on 20 features and identified 3 major classes. The classes can be generalized as pose-estimate speed, angle, and length. We tested multiple different combinations and strategically selected a combination of 7 that will most likely help isolate the different rodent behaviors. Feature 1 examines the body length of the mouse, dissociates stationary from ambulatory states, and identifies actions that consist of elevated body parts (mouse will appear shorter with rearing, grooming or scrunching with a flat bottom-up 2D perspective) (Eqn.1). Feature 2 subtracts front paws to base of tail distance from body length, or whether the animals front paws are further away from base of tail than that of the snout (Eqn.2). Feature 3 subtracts back paws to base of tail distance from body length, or how extended/contracted the snout is from back paws using body length as a reference (Eqn.3). Feature 4 finds the distance between two front paws,

as the proximity of the two paws can be an important marker for various behaviors (Eqn.4). Feature 5 looks at the speed of snout, or how far the snout has moved per unit time (Eqn.5). Feature 6 looks at the speed of base of tail, or how far the base of tail has moved per unit time (Eqn.6). Finally, feature 7 is concerned with body angle, whether the animal is orienting to the right or left, and how big of a angular change there is (Eqn.7). Upon extracting these 7 features, we visualized the clusters in a dimensionally-reduced space.

**B. t-SNE clustering and visualization of high-dimensional features.** To visualize the clusters in a 7-dimensional feature space, we performed a particular dimensionality reduction clustering algorithm, t-SNE, onto a 3-dimensional space. On this 3-dimensional t-SNE space, we saw distinct nodes of data clusters that appear to be interconnected, albeit much fewer data points (Fig. 2). The spectral nature of low-dimensional features align with prior work (1, 12). We further examined this 3-D space and uncovered that each individual cluster appears to represent an interpretable action, as opposed to similar snapshots between two very different actions. This is consistent with the notion that a behavior is usually inferred by a composite of high dimensional features. In other words, a rear and groom will both contain elevation (away from the camera), but only groom would have consistent changes in inter-fore-paw distances. Together, these results suggest that we conserved the contrast in 7-dimensional space that are important for behavioral segmentation even in a lower-dimensional space.

**C. Expectation maximization fitting of Gaussian mixture models.** Unsupervised grouping will be required to differentiate one group from another autonomously. Since t-SNE utilizes Gaussian kernels for joint distribution, we chose to classify the clusters using expectation maximization (EM) algorithms for fitting of Gaussian mixtures models (GMM) (14). The "E-step" in EM requires a set of parameters to be initialized, then subsequently updated with the "M-step". Since we do not have a priori in mouse action classes in a naturalistic setting, we initialized model mean, covariance, and priors with random values. The danger of randomly initializing GMM parameters is getting a sub-optimal local minimum log-likelihood. To attempt at escaping poor initializations, we chose to perform this method iteratively over many times and kept the initialized parameter set that gave rise to the lowest log-likelihood. Additionally, we allowed the EM algorithm to identify up to 30 classes, of which 15 unique classes (colored) were pulled out (Fig.2). Interestingly, a couple classes (8 and 15), can only be found in a few animals. These results argue that our algorithm adapts to inter-animal variability in naturally occurring behaviors if sufficient data are collected.

**D. GMM classes represent distinct actions.** The 15 classes identified using GMM was based on data in the low-dimensional space. This does not warrant a proper segregation of actual behavior. The classes could very well either subdivide a sequence of turns, group a whole spectrum

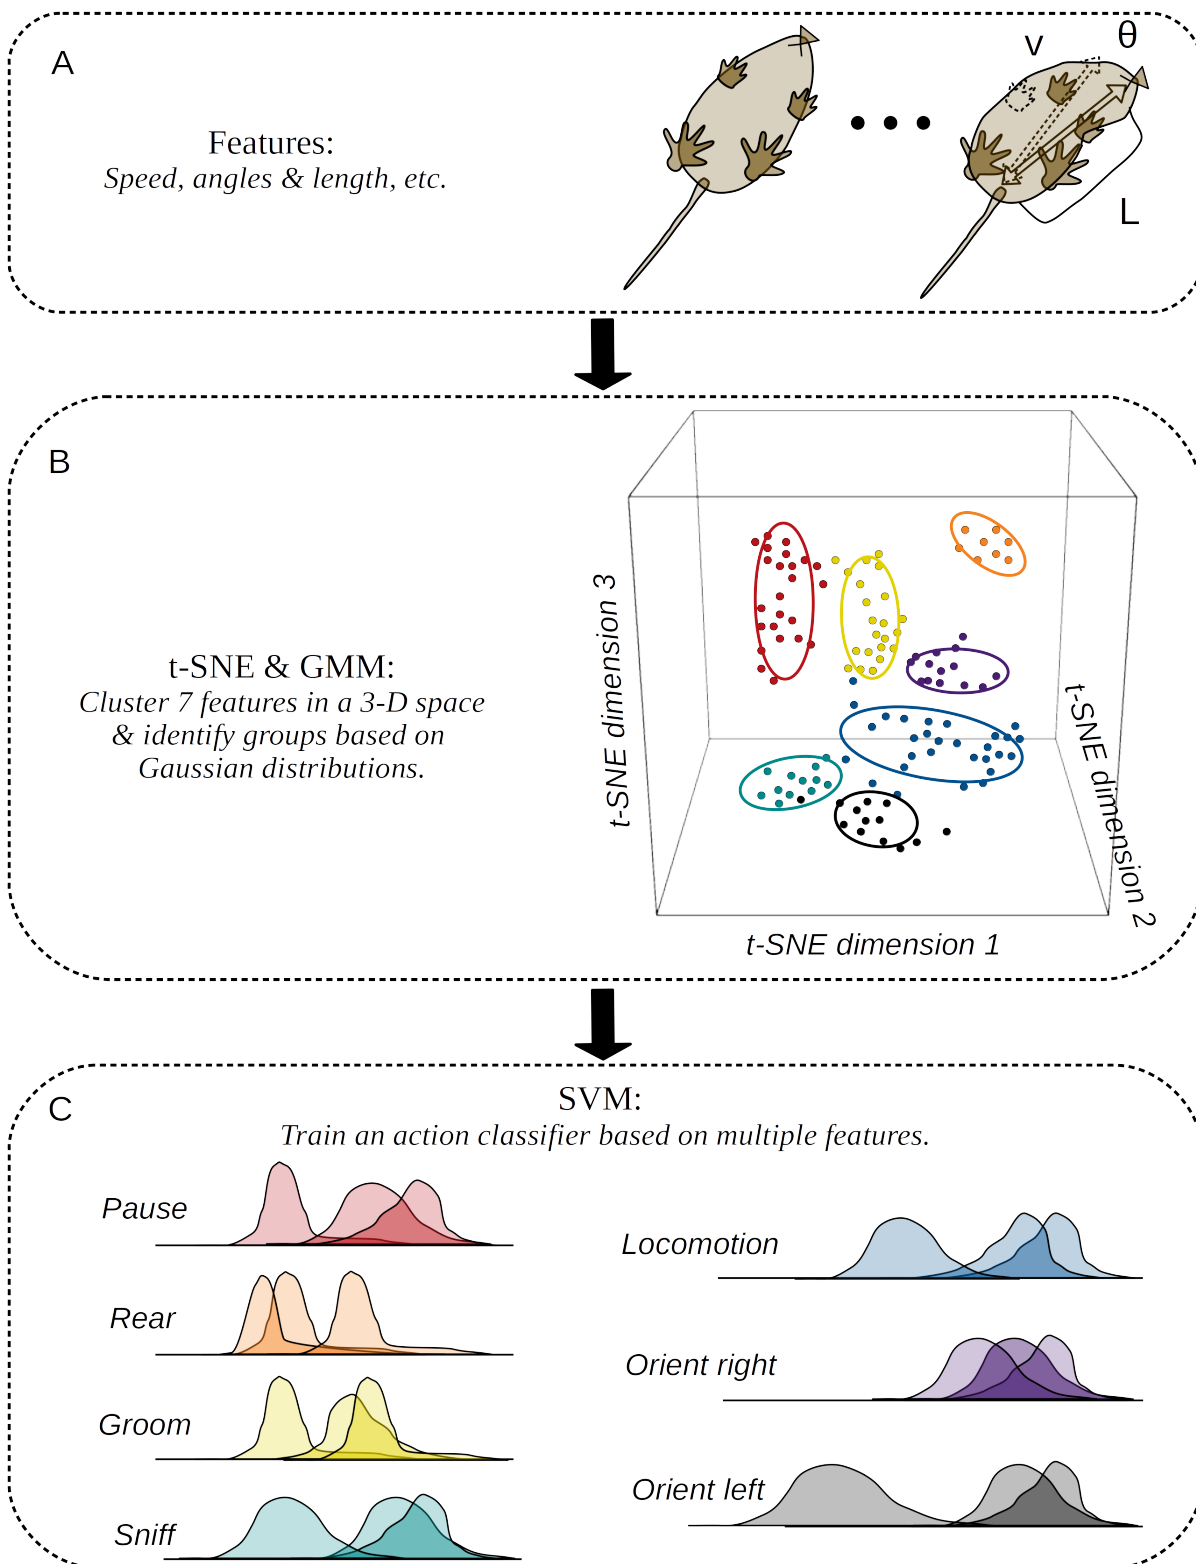


Fig. 1. Flow-chart of our proposed algorithm, B-SOiD. (A) Obtain physical features that pertain to the three different class (speed, angle and length) that define movement. (B) Cluster high-dimensional input features in a low-dimensional space utilizing t-distributed stochastic neighbor embedding (t-SNE) and fit Gaussian mixture models (GMM) to segment behaviors. (C) Train a support vector machine (SVM) classifier based on the distribution of all 7 features.

of behavior into one, or even both. To answer this, we randomly isolated short videos based on classes (See <https://github.com/YttriLab/B-SOiD> for details). We found that these 15 classes were very differentiable from one another. Since we could potentially carry bias in examining these action classes, we further investigated the distribution of low-dimensional (physical) features. Our results indicated that the 15 distinct Gaussian mixtures in the high-dimensional space, were also different in the compiled physical feature space ( $p < 0.01$ , KS-test). Given no two classes were alike, we can confidently implement a simple Support Vector Machine (SVM) classifier that deals with learning to distinguish more than two classes. All seven features will be used to train the learners 2.

**E. SVM classifier design for multi-class actions.** So far, all analyses were done post-hoc (after collecting hours of data and running dimensionality reduction on close to 1 million samples). Although this automation can already drastically improve neurobehavioral correlation, we went a step further and incorporated machine learning for a more real-time solution. Since we output more than 2 classes that we would need to differentiate, a simple GLM would not be sufficient for encoding. Recent computational advances have enabled SVMs to significantly improve decoding accuracy using error correcting output codes (ECOC) (15). Based on our utilization of normal distribution with t-SNE and GMM, we hypothesize that using Gaussian kernel functions for classifier training would supply the most robust and accurate decoding results. To test our hypothesis, we trained our classifier with three types of kernel tricks: linear, Gaussian, and polynomial. To benchmark model performance, we performed cross-validation on various partition size of held-out data. With 100 iterations for each partition size, we found that SVM with Gaussian kernel function predicts the classes most accurately and with least variation given sufficient training data (75,000 frames, or 70 % of 3 hours of video) (Fig.3). Since our expectation matched our observation, it suggests that we understand the technique that we chose to parcellate behavior.

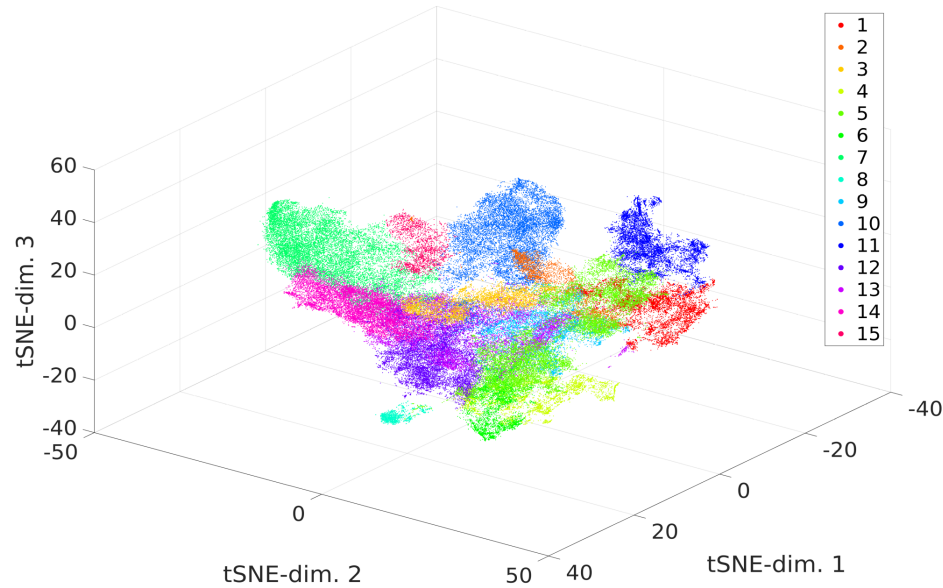
**F. SVM classifier generalizes to new dataset.** For a model to be applicable, an essential feature to carry is generalizability across variable datasets. If our behavioral model truly understands behavior in a high-dimensional space, a mouse of very different physique shall not impact performance. After we trained our model, we used it to predict the actions of a mouse that is noticeably different. Within a few seconds, the classifier categorized all 18000 actions (30-minute video, 10 frames/second (fps)). The randomly sampled predicted classes showed very similar behavior to the classes in training (Fig.4). Although there is no good measure to benchmark such similarity, our results do suggest that the behavioral model we built did not over-fit our training mice. In the following sections, we provide more evidence to increase the probability that our data-driven algorithm is valid for behavioral quantification.

**G. Classifier performance comparable to human observers.** Certainly, we want to exemplify the coherence between our data-driven technique with human observed behaviors. Up to this point, our results have mostly demonstrated statistical consistency in B-SOiD; However, part of the motivation was to build a toolbox that discovers user-defined behaviors for DeepLabCut. Transition probability matrix is a common method for analyzing similarity between states based on the predictability of the next states. For example, if actions A and B, but not C, have a high probability going into action D at the next time-step, we would consider actions A and B are more similar to each other than them to C. Based on the transition probability matrix, our 15 classes can be categorized into three major groups: stationary/rear, groom, and non-stationary 6. We ran one 10 minute video at same temporal resolution (10 fps) against 4 blind human observers and found that the inter-grader coherence ( $80.97 \pm 3.96\%$ ) was similar to machine-grader coherence ( $72.04 \pm 4.53\%$ ). Our findings not only provided evidence that human defined behaviors are typically statistical inferences of multi-dimensional patterns, but also raised an awareness for biases in subjective behavioral quantification, which takes 1000x longer. To further support the validity of our model, we looked into the 15 individual classes.

**H. Unsupervised learning reasonably parses behaviors and aligns with expectation.** To address the robustness of the 15 classes modeled, we collected a data from 10 different animals with 19 total half-hour sessions (in addition to a training set of 4 animals comprising a total of 6 sessions). After combining training and prediction data, we can analyze a few characteristics of individual actions to see if those match our expectations. First, we examine bout durations for each individual class. We found that across animals, the variability in time is allotted per action class was conserved when exploring a novel open-field (Fig.5). In addition, when grouped based on the commonly agreed upon states (rest/pause, sniff/nose-poke, rear, groom, gather/scrunch, locomotion/orientation), the relative distribution of durations are in alignment with what we have observed. For instance: rearing against the wall generally will last longer than unsupported rears that are out in the open; a groom that has a smaller range of motion may last over 4 seconds, while a faster and coarser "scratch" may not; orientating movements are much briefer than locomotion given the size of our open-field.

Second, we analyzed the transitional probabilities of all classes. We can test the validity of our model based on the predictability of the action followed. If the subdivided grooms are true classes, the three classes that we ascribed "groom" should be interconnected given the stereotypical sequence of grooming in rodents (?). Indeed, the mean across-animal transition matrix exhibits an oscillating structure around the intersection of current and next states of grooms (Fig.6). We also observed that pause and sniff states have a higher probability of transitioning into non-stationary states (locomotion and orientation) than the inverse, consistent with the notion of novelty-seeking behavior, particularly

A



B

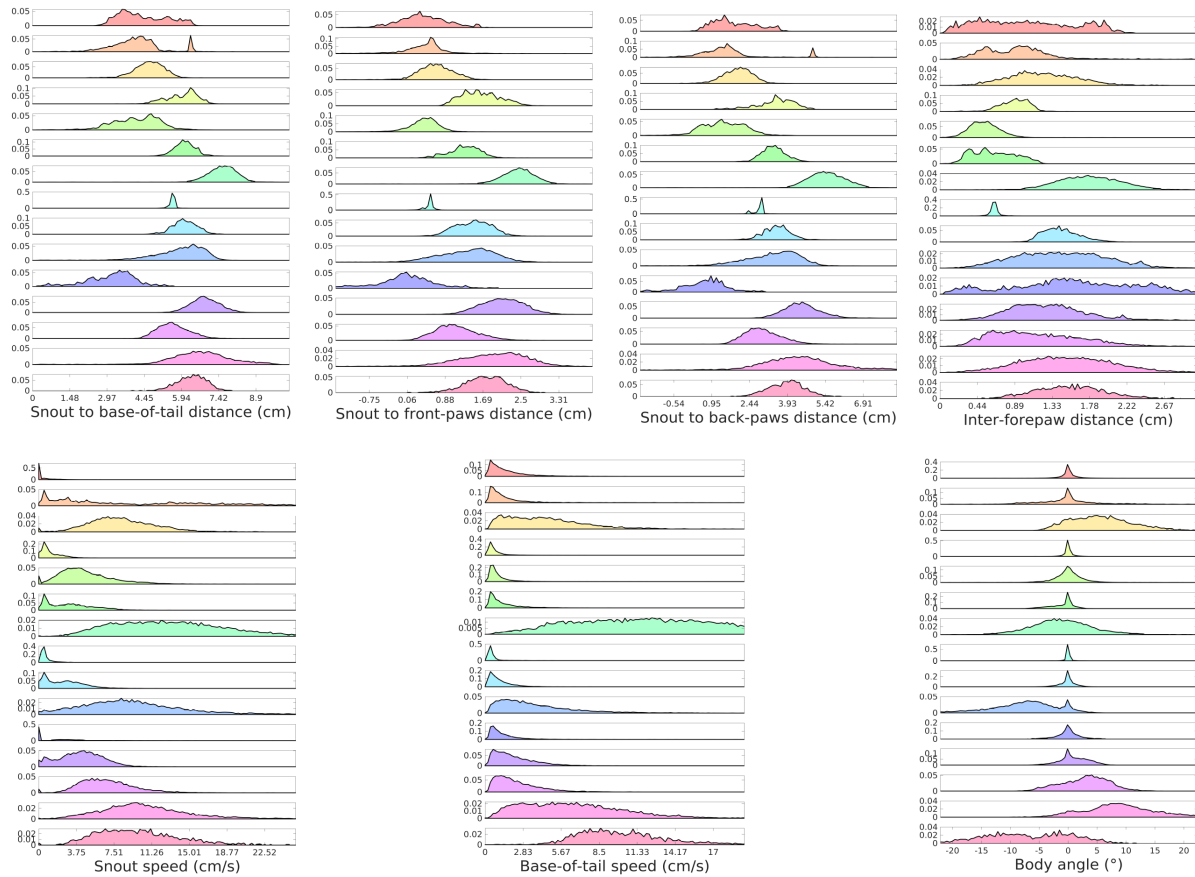


Fig. 2. Gaussian mixture models (GMM) classified clusters into unique actions. (A) Expectation maximization for fitting of GMMs on 3-dimensional t-SNE space identifies 15 classes. (B) All 7 low-dimensional (physical) features per class, each represent the corresponding colors from (A).

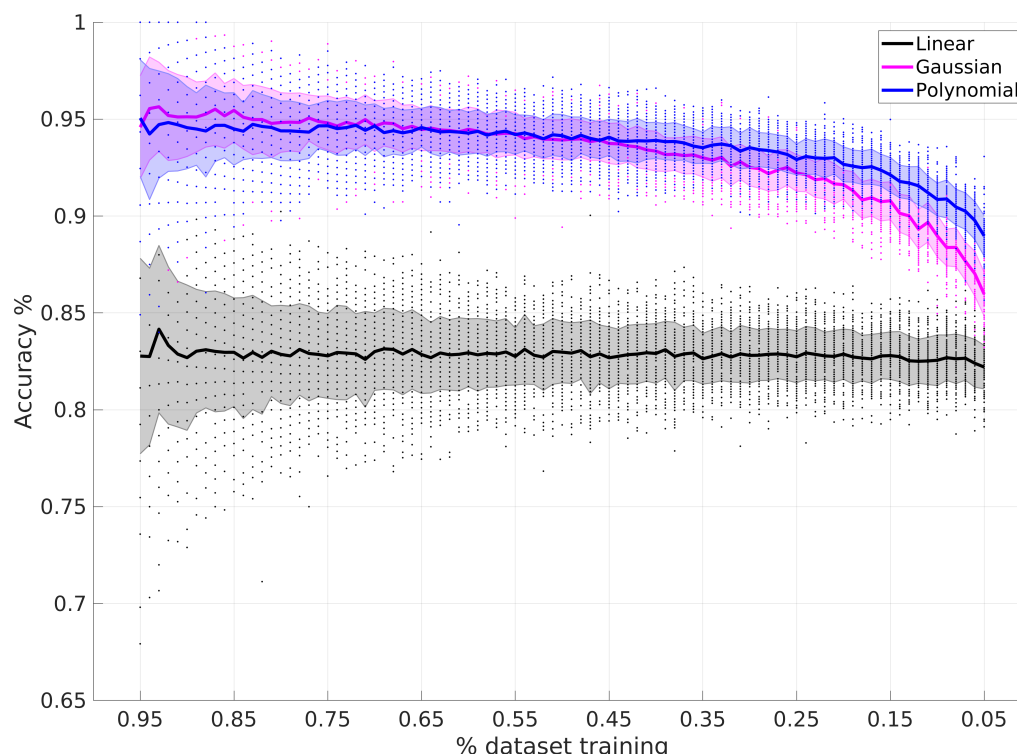


Fig. 3. A. Support vector machine classifier robustness with three different kernel tricks. Given sufficient training data, it appears Gaussian kernel function (magenta) works best and most robust, albeit polynomial works better for small training data (blue).

when first exposed to a new environment. We postulate that this is also the case for individual animals. We split hour-long sessions into early versus late exploration blocks for individual animals ( $N=9$ , 30 minutes each block). We discovered that such asymmetry into non-stationary versus stationary was present across individual animals in early exploration (left panels) when compared to that in late exploration (right panels; Fig. 6). Our findings suggest that, aside from defining individual actions, B-SOiD serves as a useful tool for uncovering the larger structure of action sequences.

**I. Exploration versus exploitation.** The exploration-exploitation dilemma is a widely established phenomenon in which an agent, in this case a mouse, learns about the environment through series of transitions between obtaining information by exploring and exploiting previously obtained knowledge. With the B-SOiD algorithm, we have isolated potential behavioral strategies going from early to late exploration and suspect a transition from exploration to exploitation. We hypothesize that the temporal inter-class intervals (ICIs) distribution will converge for stationary versus non-stationary behaviors. The empirical cumulative distribution function (eCDF) for time between the same actions demonstrates that the difference between non-stationary and stationary ICIs distributions diminished during late exploration (Fig. 7), consistent with the idea that exploration-exploitation trade-off depends on experience in

the environment (16). More surprisingly, we observed that the our extracted stereotypical grooming sequence appear to be affected by experience as well. These results suggest that our algorithm can dissect out a finer behavioral predictor for exploration-exploitation trade-off.

## J. Behavioral organization in naturalistic exploration.

We found it interesting that mice revisit grooming behaviors more often in the later stages of exploration. We further investigated the baseline-subtracted history-dependent transition matrix (HDTM). This measure reveals how likely an action is to occur above and beyond the baseline probability of it occurring at all (Fig. 8). In the mean across-animal HDTM, a rhythm for grooming in particular is observable. For individual animals, we see that the pattern is more pronounced in the late exploration phase (mouse 2 experienced the open field for an additional 15 minutes prior to the 'early' block shown). This may be due to many causes, including, but not limited to, comfort with new surroundings, decreased benefit in exploration, or even merely energy expenditure. Although more in-depth analyses are needed, we are encouraged by B-SOiD's ability to capture a full range of behavioral responses.

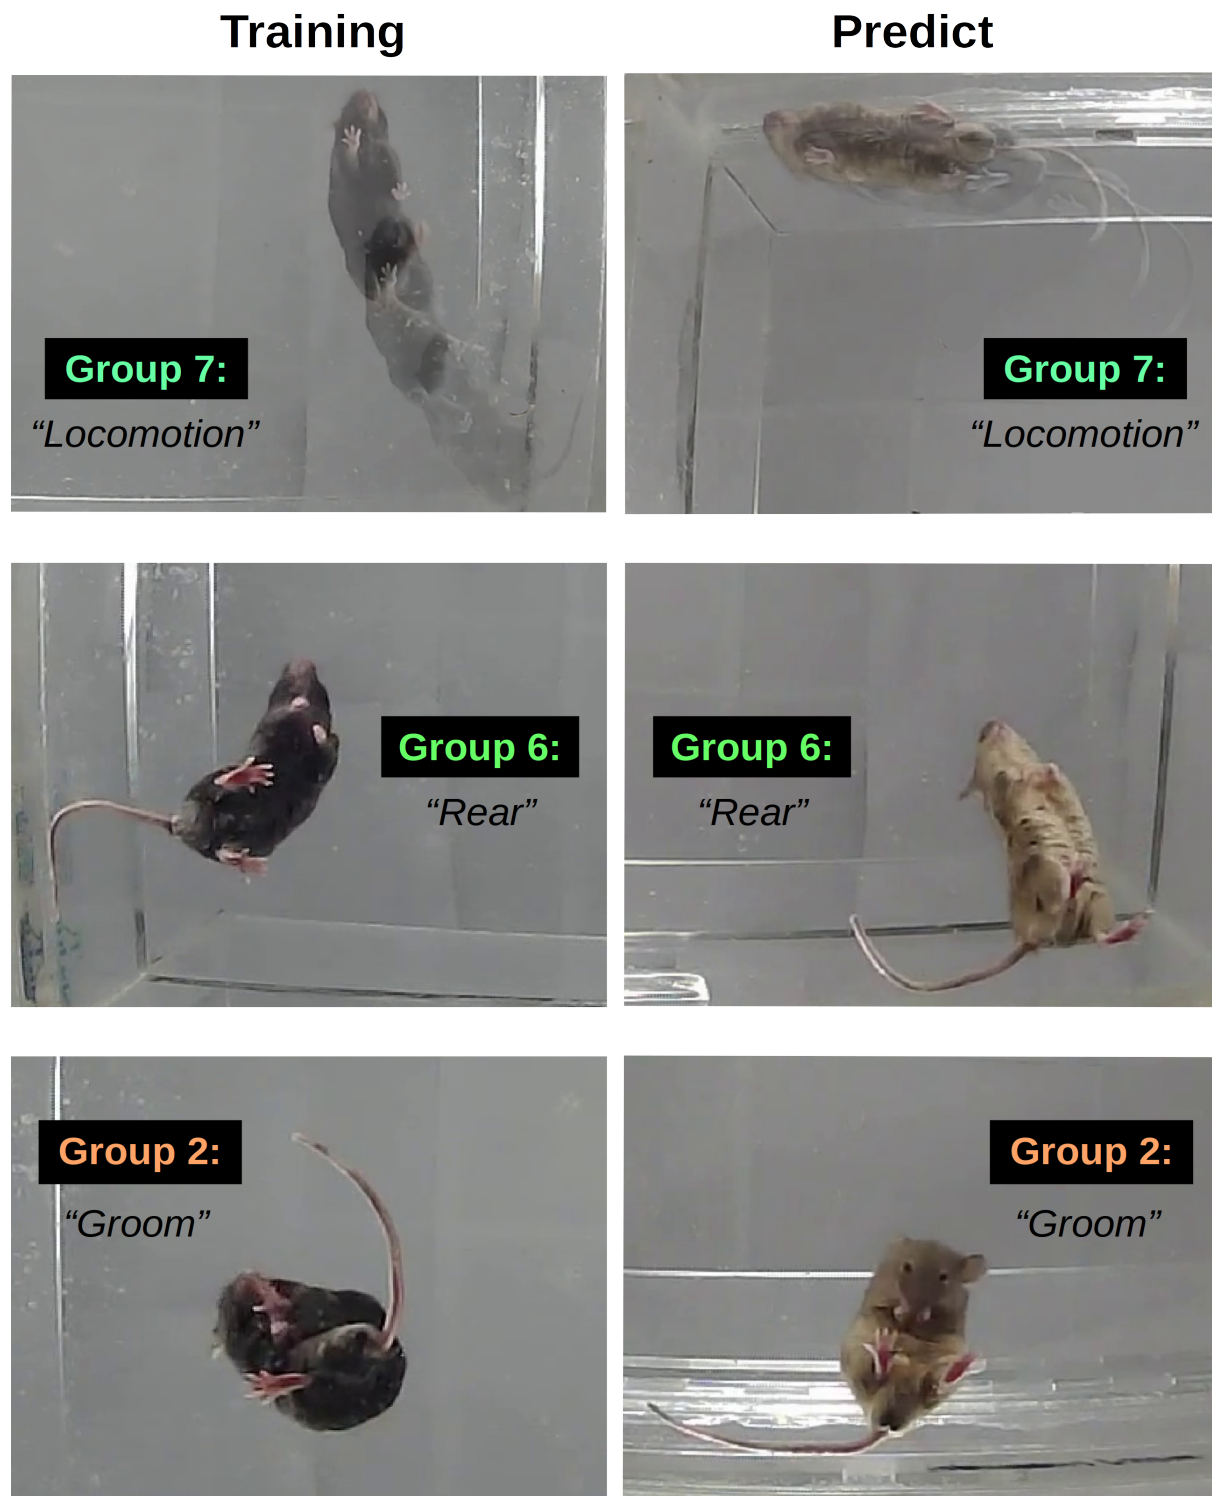


Fig. 4. Classifier predicted actions of a novel animal (brown mouse, right) matched those of the trained animal (black mouse, left). Plotted locomotion was the beginning 500ms (going from shaded gray to solid colors) of an ambulation bout that continued on for more than 500ms; whereas grooms and rears were snapshots of 300ms into a continuous behavior that lasted at least 500ms.

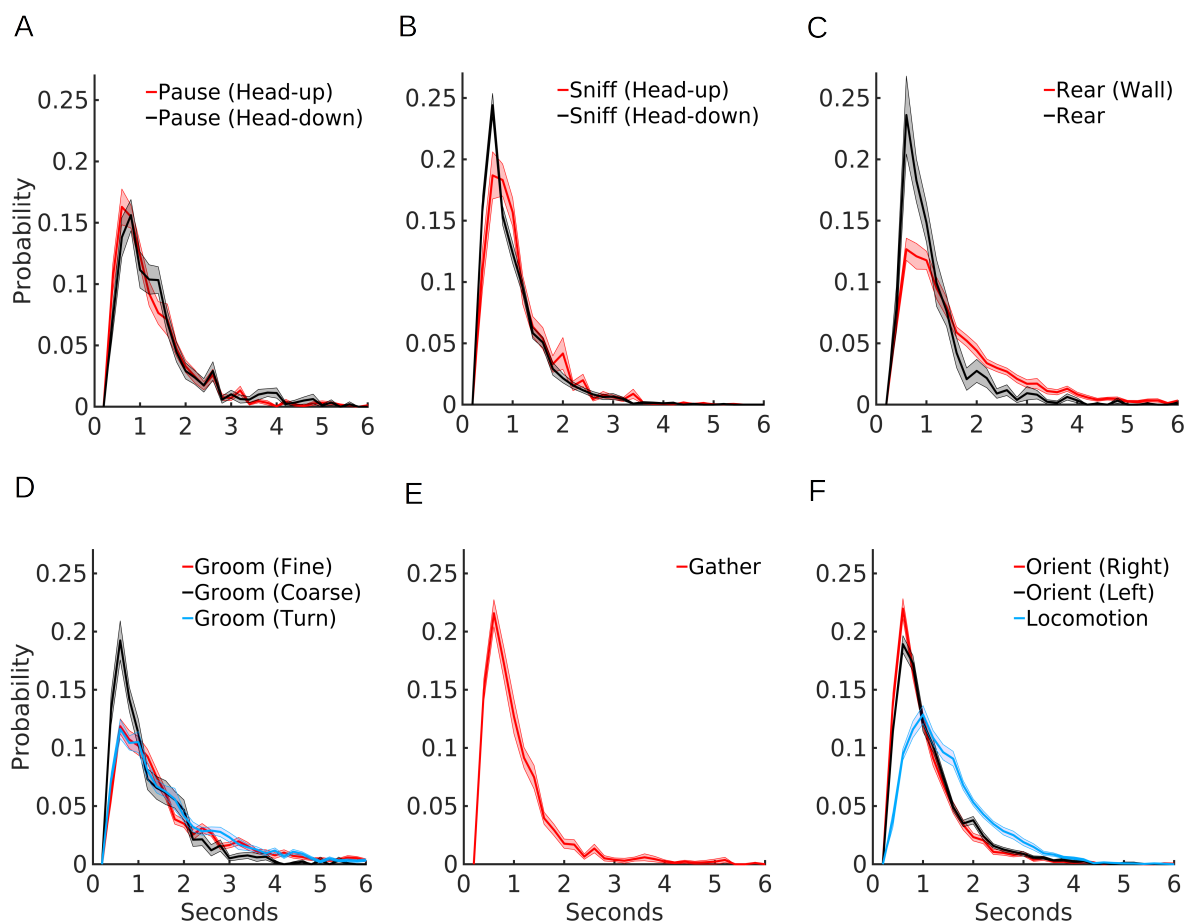


Fig. 5. Duration mean and standard error for all behavioral bouts from 10 mice, consisting of 19 total 30-minute sessions ( $n=2$  for mouse 1-9 and  $n=1$  for mouse 10). A-B. Two types of unique pause and sniff were identified, similar bout durations. C. The two classified rears, either against the wall or without support, show that rearing lasts longer when it is against the wall. D. The sub-divided grooming sequence shows that the animals make large range of motions (coarse) for shorter periods of time compared to that of fine grooming or when it is orienting to either side to groom. E. Gathering or scrunching, has majority of action bouts shorter than 2s. F. Orienting to the right or to the left has very similar bout duration distribution, all appears to be shorter than locomotion.

## Discussion

Naturalistic behavior provides a rich account of an animal's motor plans and decisions, largely unfettered by experimental constraints. Until recently, capturing the complexity of these behaviors in the open field has proven prohibitively taxing. Still, new methods are difficult to implement and/or offer an incomplete account of the movements of the actual effectors (e.g. limbs, head). Our unsupervised algorithm, B-SOiD, offers the opportunity to capture limb and action dynamics through the use of the popular DeepLabCut software. This tool also serves as the vital bridge between knowing the location of body parts (provided by DeepLabCut) and the actions that those body parts perform in concert with one another. It also demonstrates the utility of artificial intelligence, specifically the integration of multi-dimensional embedding, iterative expectation maximization, and a multi-learner design coding matrix, in classifying behavior. Paired with the insight that the presence of missing information is itself information, these algorithms even allow the extraction

of three-dimensional movements from two-dimensional data. Finally, our approach has enabled the initial study of action sets and how transitions between actions change with context - in this case, something as simple as the passage of time. The described unsupervised learning algorithm allows users to automate detection of various classes of actions so as to understand the dynamics of Open-field, naturalistic behavior and the neural mechanisms governing their selection and performance.

Also important is the ability to decompose behaviors into their constituent movements. By using limb position, we can extract not only whether an animal was walking or grooming, but also determine the contributing stride length, speed of arm extension, etc. While we have previously benefited from access to such performance parameters (17), this may prove to be a potent advantage in the study of disease models. The study of obsessive compulsive disorder in particular has long sought improved identification and quantification of grooming behavior. B-SOiD enables the detection of grooming, it's

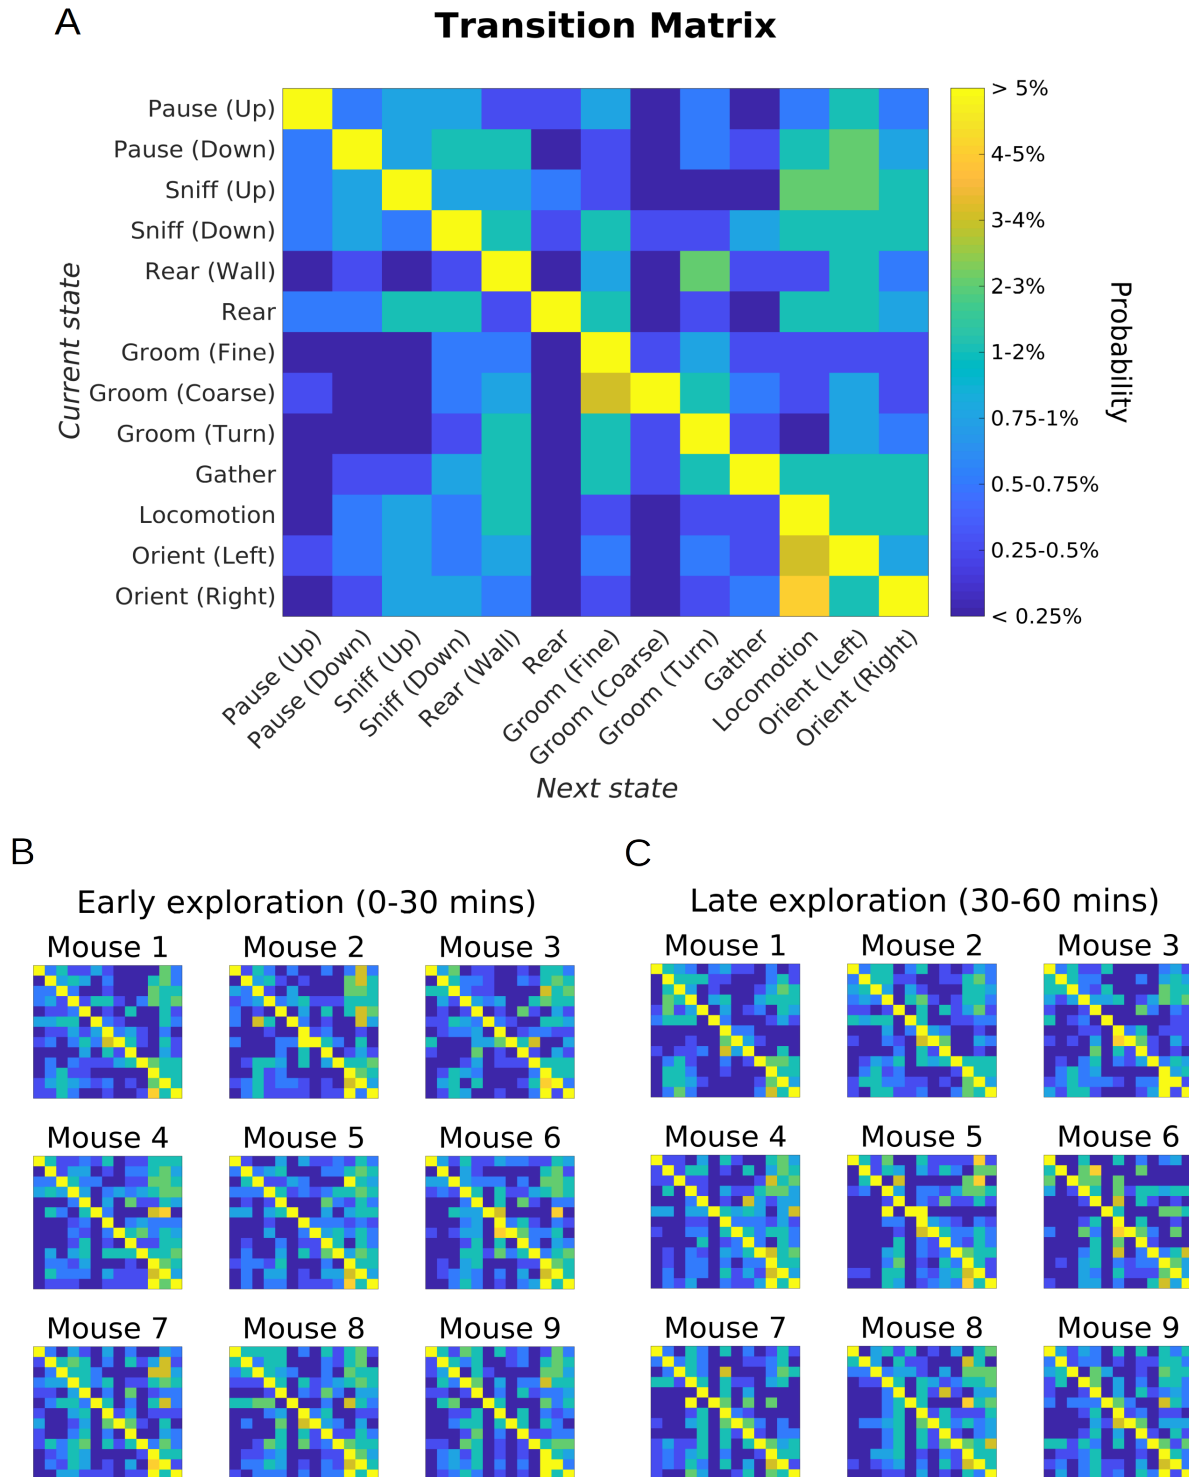


Fig. 6. Frame-by-frame transition probabilities, each row describes a current action, corresponding columns represent the range of upcoming behaviors, and colors indicate the computed probability. A. Average Frame-by-frame transition matrix from 10 different animals, 19 total 30-minute long sessions, predicts reasonable transitions between stationary segments, between segregated grooming sequence, and between ambulatory actions. B.-C. Individual mouse's transition matrix appears to be slightly variable; however, gross structural change in behavior appears to be similar going from early exploration (0:00:00.16 to 00:30:00.00, or 0-30 mins, B.) to late exploration (0:30:00.16 to 1:00:00.00, or 0-60 mins, C.).

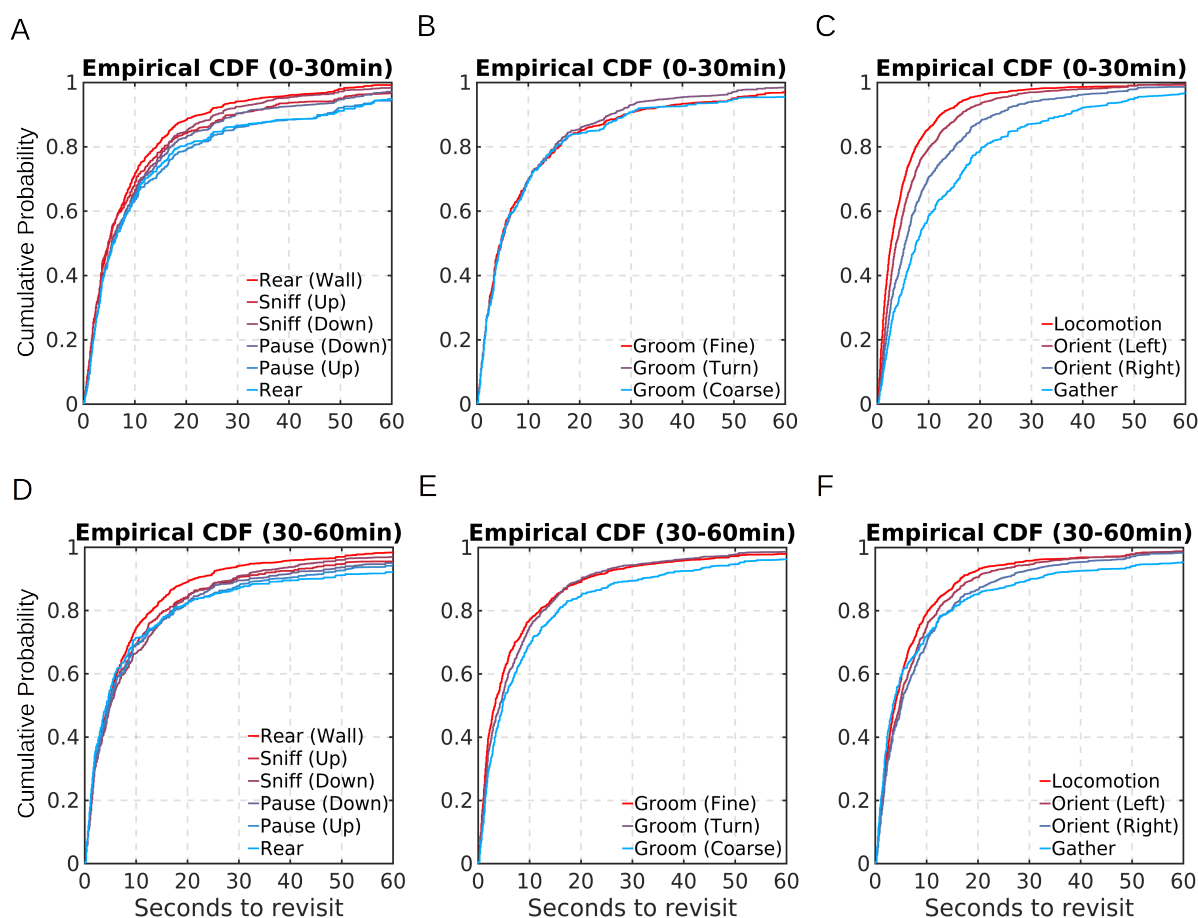


Fig. 7. Empirical cumulative distribution function of recurrent behavioral intervals from 10 mice, first 30 minutes versus last 30 minutes Open-field exploration. A.-C. In the first 30 minutes of Open-field navigation, all animal appears to revert back to actions that assist in exploration of environment faster, i.e. locomotion, orientation, and perhaps rearing against the wall. D.-F. In the last 30 minutes of the experiment, the recurrent behavioral intervals appear to shift in majority of the actions, showing a shorter interval between grooms and pauses, and a longer interval between locomotion.

relation to other behavior, and how vigorously each grooming bout is performed (18). Additionally, neurobehavioral deficits such as diminished locomotor speed may be the result of shorter stride length or slower stride – with important differences between the two causes. Therefore, understanding both the structure and substructure of actions increases the potential of research.

Though commonly used, and perhaps as a result of its common usage, the term action has a broad range of applications ranging from sub-second muscle activations (e.g. "snapping your fingers") to a prolonged series of motor commands (e.g. "going home" or reorienting, walking, and engaging a different behavioral port). Our approach, while initially focused on what we believed to be behavioral building blocks (grooming, rearing, walking, etc) was susceptible to priors and anthropomorphizing. Based upon the the described dimensionality reduction applied to length, distance, change in position and angle, we discovered that, indeed, we may be over-simplifying animal behaviors, that may generate unnecessary challenges in neurobehavioral correlation. In the software package, we preserve the user-defined categories of "ac-

tions", along with the unsupervised clustering and a hybrid classifier format. Therefore, B-SOiD provides a unique advantage. It allows the experimenter to automatically or manually build classifiers of all types of behavior. Given the input data of control or any disease tetrapod model - which spans rodents to non-human primates, the algorithm derives classes of behavior based solely upon what is present in the data.

## Methods

**Animal and Open-field set-up.** 10 C57/BL6 adult mice (5 males and 5 females) were placed in a clear 15x12 inch rectangular Open-field for one hour while a 1280x720p video-camera captured video at 60Hz. Video was acquired from the bottom-up, 19 inches away from the diagonal midpoint of the container. The videos were then divided into first and last 30 minutes for analyses purposes.

**Low-pass likelihood filter.** DeepLabCut estimates the likeliest position of all body parts for all frames, even when it is completely occluded. Since we are recording from bottom-up, it is often the case that the mouse's snout or forelimbs

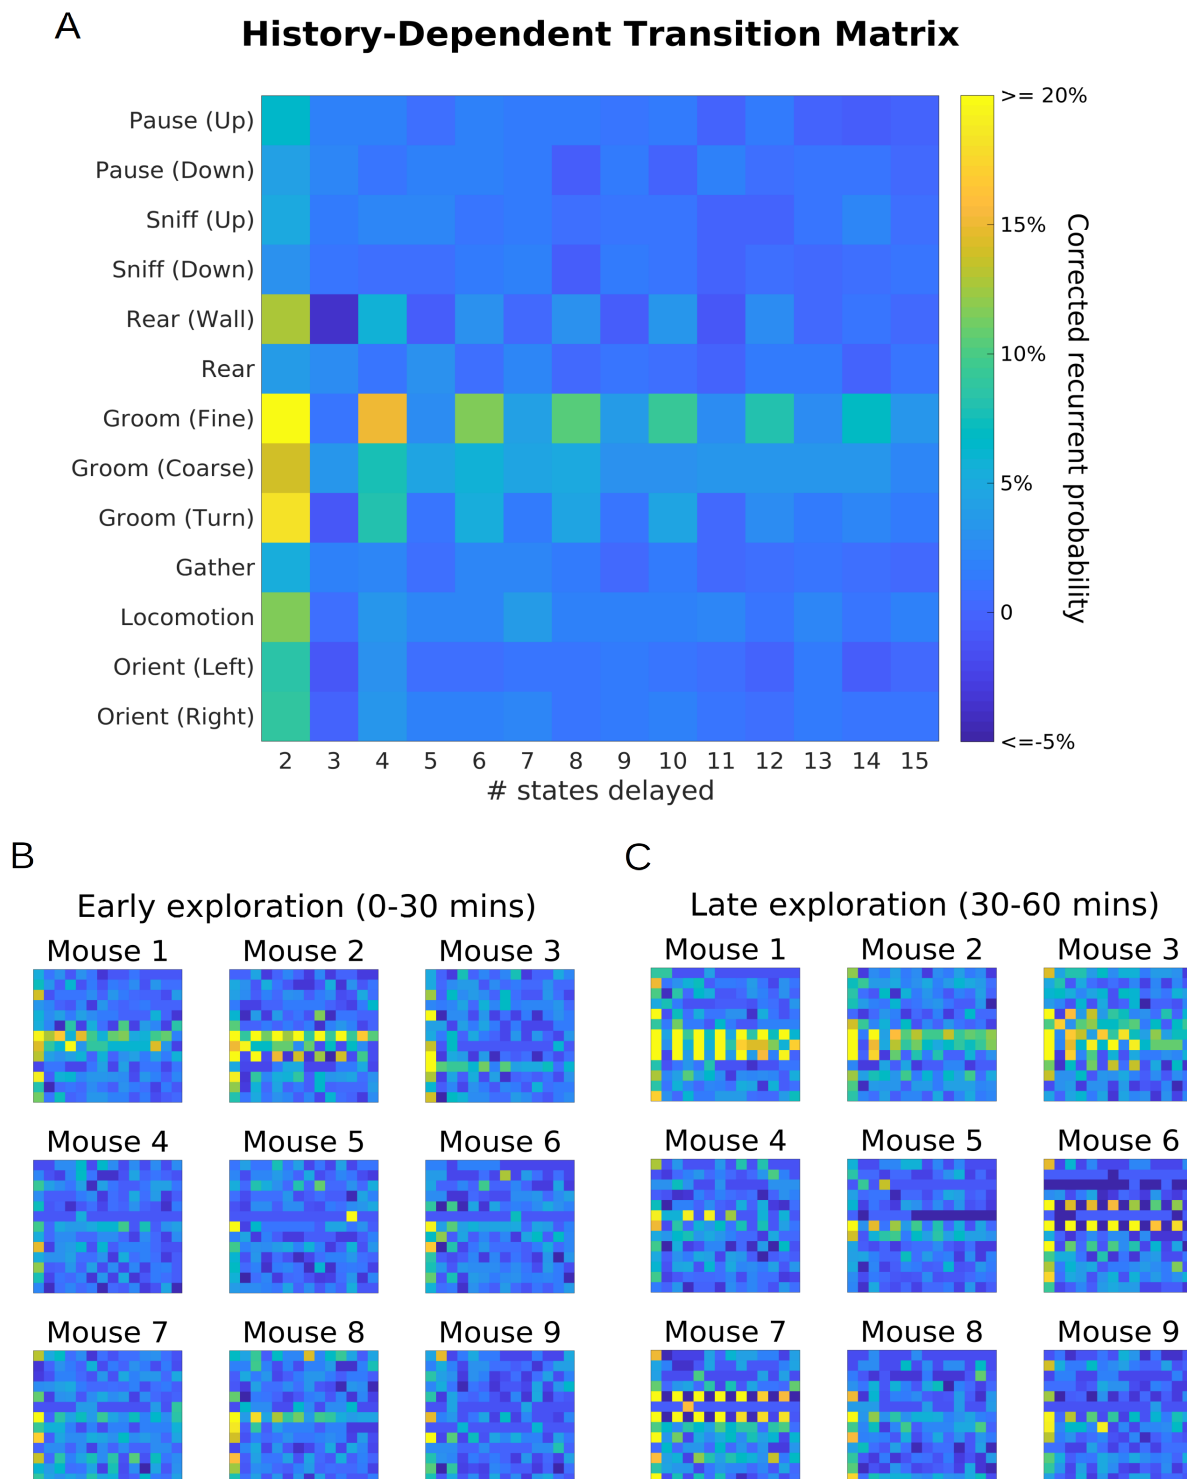


Fig. 8. Surfaced behavioral structure to exploration. A. Average recurrent matrix aligned to end of most recent behavior X from all sessions revealed that various grooms identified are recurring more often than baseline, whereas other behaviors do not. B.-C. Individual mouse's recurrent matrix appears to be slightly variable; however, majority of animals show an increase in probability of recurring grooms from early exploration (0:00:00.16 to 00:30:00.00, or 0-30 mins, B.) to late exploration (0:30:00.16 to 1:00:00.00, or 0-60 mins, C.). Note that mouse 2 was pre-exposed to the open-field for 15 minutes prior to recording.

will not be seen during rearing, grooming, and the likes of them. To take advantage of this situation, instead of replacing the estimated position with unworkable variables, we apply a low-pass filter and convert all sub-threshold estimated positions as the previous likeliest position, if exceeds our cutoff ( $p < 0.2$ ). This particular workaround allows the machine to treat it as a unique signal (no signal).

**Feature extraction.** Out of the 20 possible features we studied, we distilled it down to 7 based on hierarchical clustering analyses (HCA) and features that capture absence of signals best. The body length, or distance from snout to base of tail, or  $d_{ST}$ , is formulated as follows,

$$||d_{ST}|| = \sqrt{\sum_{D=1}^2 (S_D - T_D)^2} \quad (1)$$

, where  $S_D$  and  $T_D$  represent the likeliest position of snout and base of tail, respectively, and  $D$  denoting x or y dimension.

The front paws to base of tail distance relative to body length, or  $d_{SF}$ , is computed with the equation,

$$d_{SF} = ||d_{ST}|| - ||d_{FT}|| \\ = \sqrt{\sum_{D=1}^2 (S_D - T_D)^2} - \sqrt{\sum_{D=1}^2 (F_D - T_D)^2} \quad (2)$$

, where  $d_{FT}$  is the distance between front paws and base of tail,  $F_D$  is the mean x and y position of the two front paws. The back paws to base of tail distance relative to body length, or  $d_{SB}$ , is calculated using the formula,

$$d_{SB} = ||d_{ST}|| - ||d_{BT}|| \\ = \sqrt{\sum_{D=1}^2 (S_D - T_D)^2} - \sqrt{\sum_{D=1}^2 (B_D - T_D)^2} \quad (3)$$

, where  $d_{BT}$  is the distance between back paws and base of tail,  $B_D$  is the mean x and y position of the two back paws. The distance between two front paws,  $d_{FP}$ , is derived from,

$$||d_{FP}|| = \sqrt{\sum_{D=1}^2 (FR_D - FL_D)^2} \quad (4)$$

, where  $FR_D$  and  $FL_D$  are the likeliest positions of right and left front paws, respectively.

The snout speed,  $v_S$ , or displacement over period of 16 ms, uses the following equation,

$$||v_S|| = \sqrt{\sum_{D=1}^2 (S_D^{t+1} - S_D^t)^2} \quad (5)$$

, where  $S_D^{t+1}$  and  $S_D^t$  refer to the current and past likeliest snout positions, respectively.

The base of tail speed,  $v_T$ , or displacement over period of 16 ms, similar to the formula above, as follows,

$$||v_T|| = \sqrt{\sum_{D=1}^2 (T_D^{t+1} - T_D^t)^2} \quad (6)$$

, where  $T_D^{t+1}$  and  $T_D^t$  refer to the current and past likeliest base of tail positions, respectively.

The snout to base of tail change in angle,  $\Delta\theta_A^A$ , is formulated as follows,

$$\Delta\theta_A^A = \text{sign}(A'_x A_y - A_x A'_y) \left( \frac{180}{\pi} \right) \\ \arctan \left( \frac{A'_x A_y - A_x A'_y}{A'_x A_x + A'_y A_y} \right) \left( \frac{180}{\pi} \right) \\ \text{sign}(A'_x A_y - A_x A'_y) (1 - \text{sign}(A_x A'_x + A_y A'_y)) \quad (7)$$

, where  $A$  and  $A'$  represent body length vector at past ( $t$ ) and current ( $t+1$ ) time steps, respectively,  $\text{sign}$  equals 1 for positive, -1 for negative, 0 for 0, and  $x \cdot ||x||$  for complex numbers. Note that the Cartesian product and dot product are necessary for four-quadrant inverse tangent and that the sign is flipped to determine left versus right in terms of animal's perspective.

In addition, the features are also smoothed over, or averaged across, a sliding window of size equivalent to 60 ms (30 ms prior to and after the frame of interest).

**Data clustering.** With sampling frequency at 60 Hz, 1 frame every 16 ms, we are capturing fragments of movements. Any clustering algorithm will have a difficult time teasing apart the innate spectral nature of action groups. To resolve this issue, we decided to either take the sum over all fragments for time-varying features (features 5-7), or the average across the static measurements (features 1-4) every 6 frames. Due to our sliding window smoothing prior to this step at about double the resolution of the bins, we are not concerned with washing out inter-bin behavioral signals.

Upon t-distributed Stochastic Neighbor Embedding (t-SNE) was performed on our high-dimensional input data to minimize the divergence between the distribution of input objects and the embedded distribution of the low-dimensional space. This algorithm has been preferred over other dimensionality reduction methods simply due to preservation of local structures, allowing behavioral data points to be presented in a continuous fashion. The locations of the embedded points  $y_{ij}$  are determined by minimizing the Kullback-Leibler divergence between joint distributions  $P$  and  $Q$ , which is formulated using the following equations:

$$C(\epsilon) = KL(P||Q) = \sum_{i \neq j} p_{ij} \log \left( \frac{p_{ij}}{q_{ij}} \right) \quad (8)$$

$$\frac{\partial C}{\partial y_i} = 4 \sum_{j \neq i} (p_{ij} - q_{ij}) q_{ij} z(y_i - y_j) \quad (9)$$

In simpler terms, similar objects (think of the object as a mouse action, high values of  $p_{ij}$ ) will retain its similarity visualized in the low-dimensional space (high values of  $q_{ij}$ ), scaled with a normalization constant  $Z$  defined as  $\sum_{k \neq l} (1 + \|y_k - y_l\|^2)^{-1}$ . To accelerate the dimensionality reduction process, we opted to perform Barnes-Hut approximation (13).

**Grouping.** Expectation Maximization based on Gaussian Mixture Models (14) was performed to guarantee convergence to local optimum. We opted to randomly initialize the Gaussian parameters  $\mu_k, \Sigma_k$  and  $\pi_k$ , over number of times to escape a bad local optimum. First, we evaluate the responsibilities using the initialized parameter values, or E-step,

$$\gamma(z_{nk}) = \frac{\pi_k N(x_n | \mu_k, \Sigma_k)}{\sum_{j=1}^k \pi_j N(x_n | \mu_j, \Sigma_j)} \quad (10)$$

Second, we re-estimate the parameters using current responsibilities  $\gamma(z_{nk})$ , or M-step,

$$\mu_k^{new} = \frac{1}{N_k} \sum_{n=1}^N \gamma(z_{nk}) x_n \quad (11)$$

$$\Sigma_k^{new} = \frac{1}{N_k} \sum_{n=1}^N \gamma(z_{nk}) (x_n - \mu_k^{new})(x_n - \mu_k^{new})^T \quad (12)$$

$$\pi_k^{new} = \frac{N_k}{N} \quad (13)$$

Finally, we evaluate the log likelihood,

$$\ln p(X | \mu, \Sigma, \pi) = \sum_{n=1}^N \ln \left\{ \sum_{k=1}^K \pi_k N(x_n | \mu_k, \Sigma_k) \right\} \quad (14)$$

to check whether or not the parameters or log likelihood has converged. If the convergence criterion is not satisfied, then let  $\mu_k, \Sigma_k, \pi_k \leftarrow \mu_k^{new}, \Sigma_k^{new}, \pi_k^{new}$ , and repeat the E and M-steps.

**Classifier design.** Since we are dealing with more than two classes, we performed error correcting output codes (ECOC) (15, 19) to reduce the problem from multi-class discrimination into a set of binary classification problems. To build this SVM classifier, we consider the following exemplar table,

Classes	Learner 1	Learner 2	Learner 3
Class 1	1	1	0
Class 2	-1	0	1
Class 3	0	-1	-1

, where learner 1 learns to differentiate class 1 (1) from class 2 (-1), learner 2 learns that class 1 (1) is different from class 3 (-1), and learner 3 learns to classify class 2 (1) from class 3 (-1). Following constructing the coding design matrix  $M$  with elements  $m_{kl}$ , and  $s_l$  as the predicted classification

score for positive class of learner  $l$ . ECOC algorithm assigns a new observation to the  $\hat{k}^{th}$  class that minimizes the aggregate loss for  $L$  binary learners.

$$\hat{k} = \underset{l=1}{\operatorname{argmin}} \frac{\sum_{l=1}^L |m_{kl}| g(m_{kl}, s_l)}{\sum_{l=1}^L |m_{kl}|} \quad (15)$$

, where  $g$  is the loss of the decoding scheme.

#### ACKNOWLEDGEMENTS

We would like to acknowledge Susanne Ahmari's lab at University of Pittsburgh for validating B-SOiD in OCD mouse models (See <https://github.com/YttriLab/B-SOiD> for videos). We thank Greta Lindemann, Justin Lowe, Grace Lindemann, Jessica Meyers for painstakingly annotating source video manually.

#### References

- Alexander Mathis, Pranav Mamidanna, Kevin M. Cury, Taiga Abe, Venkatesh N. Murthy, Mackenzie Weyand Mathis, and Matthias Bethge. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, 2018. ISSN 15461726. doi: 10.1038/s41593-018-0209-y.
- Tanmay Nath, Alexander Mathis, An Chi Chen, Amir Patel, Matthias Bethge, and Mackenzie Weyand Mathis. Using DeepLabCut for 3D markerless pose estimation across species and behaviors. *Nature Protocols*, 2019. ISSN 1754-2189. doi: 10.1038/s41596-019-0176-0.
- C. R. Gallistel. Representations in animal cognition: An introduction. *Cognition*, 1990. ISSN 00100277. doi: 10.1016/0010-0277(90)90016-D.
- Mayank Kabra, Alice A. Robie, Marta Rivera-Alba, Steven Branson, and Kristin Branson. JAABA: Interactive machine learning for automatic annotation of animal behavior. *Nature Methods*, 10(1):64–67, 1 2013. ISSN 15487091. doi: 10.1038/nmeth.2281.
- Nao Ota, Manfred Gahr, and Masayo Soma. Tap dancing birds: The multimodal mutual courtship display of males and females in a socially monogamous songbird. *Scientific Reports*, 5, 11 2015. ISSN 20452322. doi: 10.1038/srep16614.
- A J Spink, R A Tegelenbosch, M O Buma, and L P Noldus. The EthoVision video tracking system—a tool for behavioral phenotyping of transgenic mice. *Physiology & behavior*, 73(5): 731–44, 8 2001. ISSN 0031-9384. doi: 10.1016/s0031-9384(01)00530-3.
- Jeffrey E Markowitz, Winthrop F Gillis, Celia C Beron, Scott W Linderman, Bernardo L Sabatini, Sandeep Robert, and Datta Correspondence. The Striatum Organizes 3D Behavior via Moment-to-Moment Action Selection. *Cell*, 174:44–49, 2018. doi: 10.1016/j.cell.2018.04.019.
- Barry J. Everitt and Trevor W. Robbins. Neural systems of reinforcement for drug addiction: From actions to habits to compulsion, 11 2005. ISSN 10976256.
- Xin Jin and Rui M Costa. Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature*, 466(7305):457–62, 7 2010. ISSN 1476-4687. doi: 10.1038/nature09263.
- Long Ding and Joshua I. Gold. The basal ganglia's contributions to perceptual decision making, 8 2013. ISSN 08966273.
- Xin Jin, Fatuel Tecuapetla, and Rui M. Costa. Basal ganglia subcircuits distinctively encode the parsing and concatenation of action sequences. *Nature Neuroscience*, 17(3):423–430, 3 2014. ISSN 10976256. doi: 10.1038/nn.3632.
- Andreas Klaus, Gabriela J. Martins, Vitor B. Paixao, Pengcheng Zhou, Liam Paninski, and Rui M. Costa. The Spatiotemporal Organization of the Striatum Encodes Action Space. *Neuron*, 2017. ISSN 10974199. doi: 10.1016/j.neuron.2017.08.015.
- Laurens Van Der Maaten. Accelerating t-SNE using Tree-Based Algorithms. Technical report, 2014.
- M Jordan, J Kleinberg, and B Schölkopf. Pattern Recognition and Machine Learning. Technical report.
- Sergio Escalera, Oriol Pujol, and Petia Radeva. On the decoding process in ternary error-correcting output codes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(1):120–134, 2010. ISSN 01628828. doi: 10.1109/TPAMI.2008.266.
- Oded Berger-Tal, Jonathan Nathan, Ehud Meron, and David Saltz. The exploration-exploitation dilemma: A multidisciplinary framework. *PLoS ONE*, 9(4), 4 2014. ISSN 19326203. doi: 10.1371/journal.pone.0095693.
- Eric A Yttri and Joshua T Dudman. Opponent and bidirectional control of movement velocity in the basal ganglia. *Nature*, 533(7603):402–6, 2016. ISSN 1476-4687. doi: 10.1038/nature17639.
- Ann M Graybiel and Esen Saka. A genetic basis for obsessive grooming. *Neuron*, 33(1): 1–2, 1 2002. ISSN 0896-6273. doi: 10.1016/s0896-6273(01)00575-x.
- Thomas G Dietterich and Ghulam Bakiri. Solving Multiclass Learning Problems via Error-Correcting Output Codes. Technical report, 1995.