Fine decomposition of rodent behavior via unsupervised segmentation and clustering of inertial signals

Romain Fayat[1], Marie Sarraudy[1], Clément Léna[1], Daniela Popa[1], Pierre Latouche[2,3] and Guillaume P. Dugué[1*]

[1]Neurophysiology of Brain Circuits, Institut de Biologie de l'École Normale Supérieure (IBENS), École Normale Supérieure, CNRS, INSERM, Université PSL, 75005 Paris, France
[2]Université Clermont Auvergne, CNRS, Laboratoire LMBP, UMR 6620, 63178X Aubière, France
[3]Institut Universitaire de France (IUF)

Abstract

Decomposing behavior into elementary components remains a central challenge in computational neuroethology. The current standard in laboratory animals involves multi-view video tracking, which, while providing unparalleled access to full-body kinematics, imposes environmental constraints, is data-intensive, and has limited scalability. We present an alternative approach using inertial sensors, which capture high-resolution, environment-independent, and compact three-dimensional kinematic data, and are commonly integrated into rodent neurophysiological devices. Our analysis pipeline leverages unsupervised, computationally efficient change-point detection to break down inertial time series into variable-length, statistically homogeneous segments. These segments are then grouped into candidate behavioral motifs through high-dimensional, model-based probabilistic clustering. We demonstrate that this approach achieves detailed rodent behavioral mapping using head inertial data. Identified motifs, corroborated by video recordings, include orienting movements, grooming components, locomotion, and olfactory exploration. Higher-order behavioral structures can be accessed by applying a categorical hidden Markov model to the motif sequence. Additionally, our pipeline detects both overt and subtle motor changes in a mouse model of Parkinson's disease and levodopa-induced dyskinesia, highlighting its value for behavioral phenotyping. This methodology offers the possibility of conducting high-resolution, observer-unbiased behavioral analysis at minimal computational cost from easily scalable and environmentally unconstrained recordings.

*Corresponding author: guillaume.dugue@cnrs.fr

Introduction

In animal studies, investigations into brain-behavior relationships have typically relied on constrained experimental environments with specialized instruments designed to monitor predefined actions, such as lever presses or nose pokes[1]. Advancing toward more flexible and ethologically relevant experimental designs requires innovative methods for continuously tracking animal activities over time and objectively parsing the underlying action sequences[2,3]. Progress in this field has been largely driven by computer vision techniques[4,5], particularly through the use of convolutional neural networks to detect user-defined anatomical landmarks, or keypoints, in images of behaving animals[6–8]. Applied to multi-view video recordings, keypoint detection can be used to compute three-dimensional pose estimates[9–14], which can then serve as input data for machine learning-based behavioral mapping frameworks[13,15–18].

However, these approaches come with several limitations that impose significant constraints on how they can be effectively deployed. Keypoint detection is particularly vulnerable to occlusions, which occur when body parts are obscured by environmental elements, other individuals, or when animals adopt contorted postures. Additionally, the estimation of keypoint localization is inherently jittery, reducing the spatio-temporal resolution of this technique and complicating subsequent data analysis[12,18]. In larger environments, the trade-off between image resolution and field-of-view leads to reduced spatial accuracy unless a complex, self-orienting video tracking system is employed[19]. Enhancing image resolution and increasing frame rate to minimize motion blur can produce sharper measurements but results in larger data volumes. Three-dimensional pose estimation presents further challenges, including the need for meticulous camera calibration, a strategy to mitigate the effect of potential frame drops and the management of larger data loads, which are often incompatible with long-term recordings.

Wearable inertial measurement units (IMUs), which directly report the instantaneous three-dimensional angular velocity and acceleration of specific body parts, provide a compelling alternative for monitoring behavior. These sensors produce small data files compatible with longitudinal onboard data logging, accounting for their widespread adoption among field ethologists[20,21], animal welfare scientists[22], biomechanists[23] and clinicians[24,25]. Due to their high precision, exquisite temporal resolution, and compact, lightweight form factors, IMUs are also increasingly favored by physiologists for tracking rodent head[26–45], limb[46,47], and body[48–51] movements. Being environment-agnostic, these devices are ideally suited for transitioning toward more naturalistic experimental settings involving large, complex or crowded setups. Reflecting their growing popularity in the field, IMUs are now routinely integrated into both commercial and open source head-mounted electrophysiological and imaging devices[52–54]. Despite these advantages, inertial data have rarely been used as a main source of information for automating behavioral analysis in experimental animal models[26,43,51], in contrast with IMU-based approaches in human activity recognition[55], animal health monitoring[56] and ethological research[57]. Yet, the high sensitivity, data efficiency and versatility of IMUs are well-suited for developing advanced analysis frameworks aimed at unraveling the architecture of naturalistic behavior, offering the enhanced resolution required by computational neuroethology[2] and deep behavioral phenotyping[58].

To unlock the potential of inertial recordings for mapping rodent behavior, we developed a fully unsupervised analysis pipeline called DISSeCT (Decomposing Inertial Sequences by Segmentation and Clustering in Tandem), which parses IMU signals into discrete activity motifs. Our pipeline employs a two-step process: first, transitions between statistically homogeneous segments are detected using an optimal, computationally efficient change-point detection method, taking full advantage of the low noise, high sampling rate, and artifact-free nature of IMU recordings. The resulting segments are then grouped into candidate behavioral motifs using high-dimensional Gaussian mixture modeling. We ran our pipeline on head inertial recordings obtained from freely moving rodents, leveraging the fact that these animals rely heavily on head movements to explore their environment[31,37,59] and that many of their activities exhibit distinct head inertial signatures[26,30,43,60]. We show that DISSeCT effectively decomposes inertial data into meaningful activity motifs, validated by concurrent video recordings, parsing the full repertoire of orienting movements, isolating rearing and locomotion episodes, distinguishing the main components of self-grooming, and highlighting episodes of olfactory exploration. Using a mouse model of Parkinson's disease and levodopa-induced dyskinesia, we also illustrate how our pipeline can be used to detect and characterize behavioral anomalies.

Results

Pipeline architecture

Inertial measurements of head kinematics were conducted on four Long-Evans rats placed inside a 70×70 cm arena. A custom multi-angle video capture system, operating synchronously, was used to independently document the animals' activity and compute three-dimensional pose estimation (see Methods). The headborne IMU contained a 3-axis accelerometer and gyroscope, providing head-referenced three-dimensional vector data for acceleration (denoted as $\mathbf{a}$) and angular velocity (denoted as $\boldsymbol{\omega}$), respectively. An estimation of the gravitational component of acceleration (denoted as $\mathbf{a_G}$) was computed using a previously validated method that combines $\mathbf{a}$ and $\boldsymbol{\omega}$[42] (see Methods), fully accounting for head orientation relative to gravity (hereafter referred to as head tilt). Subtracting $\mathbf{a_G}$ from $\mathbf{a}$ provided an estimate of the non-gravitational component of acceleration (denoted as $\mathbf{a_{nG}}$), representing acceleration generated by the animal's movements (Fig. 1A, left).

The initial stage of DISSeCT aims to break down the multivariate time series composed of $\mathbf{a_G}$ and $\boldsymbol{\omega}$ into non-overlapping segments with stationary statistics, in accordance with the notion that behavior can be conceptualized as a succession of elementary motifs[61,62]. To achieve this, we opted for kernel change point detection[63], implemented through an algorithm that ensures optimal segmentation with a linear computational cost[64,65] (see Methods). Unlike other techniques that focus solely on detecting shifts in mean values[66] or changes in variances[67], kernel change point detection identifies changes in the entire distribution of the multivariate time series[68]. This allows for the detection of a broader range of segment types, effectively generalizing most traditional change point detection methods. The algorithm was calibrated to produce a median segment duration of 300 to 400 ms, aligning with the timescale of previously identified mouse behavioral modules[61]. Using this methodology, our data set, totalling 8 h and 38 min of IMU recordings sampled at 300 Hz, was parsed into 44,322 segments (median duration: 377 ms, $q_{25} = 203$ ms, $q_{75} = 820$ ms) in about 5 min on a conventional workstation (see Table S3).

The second stage of DISSeCT aims to conduct a cluster analysis of segments, treating each segment as an individual observation. It begins by extracting a catalog of 240 features that characterize segment properties. These features comprise multiple statistics on $\mathbf{a_G}$, $\mathbf{a_{nG}}$ and $\boldsymbol{\omega}$, as well as time-frequency attributes of $\mathbf{a_{nG}}$ and $\boldsymbol{\omega}$ along with directional information computed from $\mathbf{a_G}$ and $\boldsymbol{\omega}$ (see Supplementary Methods). One option, as explored in previous studies[13,15–17,69], would have been to implement the clustering step by partitioning a two-dimensional representation of the feature space, albeit at the expense of losing potentially valuable information[70]. Instead, our approach involves applying a more conservative dimension reduction to maintain a high-dimensional space for clustering while achieving manageable runtimes. To accomplish this, standardized feature data were subjected to principal component analysis, retaining the first 30 principal components (PCs) as determined optimally by cross validation (see Methods). After an outlier removal step, a Gaussian mixture model (GMM) was then fitted to the data in the PC space. Contrary to standard clustering algorithms like K-means[71], approaches based on GMM can uncover clusters with different shapes, and not just spheres in high dimensions[72]. Moreover, no hypothesis is made regarding the sizes of the clusters such that clusters with heterogeneous weights can be retrieved[73]. Finally, as a statistical model, GMM allows existing statistical tools to be employed. In particular, the Bayesian information criterion (BIC), which comes with theoretical backing, can be used to estimate the number of clusters present in the data[74]. Thus, considering a grid over the number of clusters and computing BIC for each GMM fit[75], we obtained 48 mixture components (Fig. S4B). The corresponding 48 clusters were identified by assigning each segment to the mixture component to which it most likely belonged to. For visualisation purposes, a supervised UMAP (S-UMAP) algorithm was chosen to project segments onto a two-dimensional embedding of the standardized feature space, facilitating a quick examination of the distribution of clusters relative to each other. The resulting map was color-coded according to the original feature values, revealing a global polarizing influence of head pitch and a more localized influence of other kinematic features (Fig. 1B).

In summary, DISSeCT provides a unified framework for pre-processing and parsing IMU data (Fig. 1A). It essentially relies on a fully unsupervised two-step process, wherein time series segmentation is used to isolate elementary kinematic motifs that are subsequently clustered in a high-dimensional space.

Primary behavioral categories are parsed into distinct families of clusters

To assess the behavioral significance of IMU data decomposition using DISSeCT, we proceeded as follows: for each cluster, an observer was tasked with labeling the top one hundred segments with the highest posterior probabilities of belonging to it, and a "consistency score" was computed to highlight the prevalence of the most frequent label (see Methods and Table S4). With the exception of three sparse clusters, which represented highly scattered and underpopulated mixture components (see Methods and Fig. S7), all clusters exhibited a consistency score of at least 0.5 (with two-thirds of them displaying scores above 0.8, Fig. 1F) and could be clearly associated with a predominant behavioral category (orienting, rearing, grooming, locomoting or idle; see Fig. 1C,D) occupying a distinct sub-region of the supervised UMAP embedding (Fig. 1E). This analysis shows that the decomposition of IMU data obtained with our pipeline largely aligns with the major classes of rodent activities. The fact that each behavioral category was subdivided into several clusters, with three of them (orienting, rearing and grooming) comprising no fewer than 12 clusters each, indicated a higher level of granularity where each category is further broken down into multiple components. This prompted us to conduct a comprehensive examination of the GMM output in order to assess the level of detail achieved with DISSeCT.

Identification of clusters mapping the repertoire of rapid changes in head orientation

Considering that our primary data included direct measurements of three-dimensional head angular velocity, we anticipated that DISSeCT would be able to discern motor patterns associated with rapid changes in head orientation. A preliminary examination of the output of the change point detection step indeed suggested that such movements were isolated as individual segments (see Movie S1). Among the features extracted from all segments, we focused on net angular speed, calculated by dividing the net head angular displacement (from segment start to finish) by the segment's duration. The distribution of per-cluster average net angular speed values exhibited a bimodal pattern, with a cutoff of about $150\,^{\circ}/\mathrm{s}$ (Fig. 2A). The upper half of this distribution contained 18 clusters representing brief (179 ± 45 ms) head rotations associated with the "orienting", "rearing" and "grooming" categories (Fig. 2A,B) and accounting for 37% of all segments (11% of total recording duration). With the exception of the rearing termination phase, all these clusters corresponded to rotations of 30–50° (mean: $42.8 \pm 8.8^{\circ}$, Fig. 2B), regardless of the net angular speed ($r = 0.033$, $p = 0.90$). This can be explained by the fact that the average segment duration decreased as a function of the average net angular speed ($r = -0.86$, $p = 1.6\mathrm{e}{-5}$), suggesting that rapid head movements might be tuned to produce relatively invariant head angular displacement.

Behavioral categories could be distinguished based on the head pitch angle and net azimuthal speed (Fig. 2C): grooming-associated rotations, corresponding to single strokes or brief head repositioning between grooming bouts (Movie S1), occurred with the head pitched downward ($< -45^{\circ}$) and showed limited change in azimuth; rearing-associated rotations, mainly corresponding to rearing initiation or termination, took place with the head pitched upward ($> 0^{\circ}$); "orienting" movements predominantly fell within an intermediate range of pitch angles, accompanied by either minimal or significant changes in azimuth.

These observations show that our pipeline offers an effective method for isolating and sorting rapid three-dimensional changes in head orientation. A visually informative overview depicting the full repertoire of these movements was obtained using custom two-dimensional projected maps showing the segment-wise spherical trajectories of $\mathbf{a_G}$ in the head reference frame and of the animal's heading in an absolute (external) reference frame (Fig. 2D$_3$, see Supplementary Methods and Movie S2).

We noticed that some clusters contained segments belonging to different categories (Fig. 2C), interpreting this as a consequence of distinct behavioral motifs sometimes generating similar head movements. For example, rearing initiation and upward head rotation may share common head kinematic features, differing primarily in whether the forepaws are leaving the ground or not. This point is further investigated in a later section where we show that such ambiguities can be resolved by taking into account the sequences of behavioral motifs.

Locomotion splits into different regimes, including a mobile olfactory exploration module

As shown in Fig. 1C,D, the "locomotion" and "idle" categories, which accounted for 18 % of all segments and 36 % of the total recording duration, contained fewer clusters compared to other behavioral categories, suggesting a higher degree of kinematic homogeneity. When considering actual animal displacement, three of the "locomotion" clusters emerged as those associated with the highest speed (Fig. 3A), representing 68 % of the time when the animal's speed exceeded 15 cm/s. Two of the "idle" clusters emerged as those with the maximal fraction of time spent immobile (defined using a head angular speed threshold; see Methods; Fig. 3A). One of these (cluster 6) corresponded to true periods of immobility, while the other (cluster 38) contained longer duration segments during which the animal was nearly immobile but exhibited small head movements (Fig. 3A,B).

Given that the primary input data to DISSeCT did not contain any information on animal location, we did not necessarily anticipate that our pipeline would isolate locomotion phases. To understand what caused the clustering algorithm to do so, we examined the head kinematic features associated with the three "locomotion" clusters highlighted in Fig. 3A.

Two clusters (46 and 13), corresponding to second-scale fast locomotion bouts, were notably characterized by sawtooth-like patterns in the z component of non-gravitational acceleration (Fig. 3B,C) synced with each step cycle, as previously described from head acceleration measurements in running rats[28]. This signature likely reflects a bobbing head motion associated with the successive cycles of storage and release of kinetic and potential energy in the animal's skeletal muscles[76]. Continuous wavelet transform revealed that this oscillatory regime took place in the 5–7 Hz band, consistent with the known step cycle duration during trotting in rats[77].

The third cluster (37) corresponded to a strikingly different locomotor regime in which most of the power was concentrated around 10 Hz in the x and y axes of non-gravitational acceleration and angular velocity, respectively (Fig. 3D). This signature, reflecting coordinated fore-aft motion and pitch rotation, has already been described as linked with sniffing in rodents[31,78]. Compared to the other two "locomotion" clusters, this locomotor regime was also characterized by a larger downward head pitch angle of about 40° (Fig. 3E). Video inspection clearly showed that this third locomotion cluster corresponded to episodes of olfactory exploration, during which animals slowly moved forward maintaining their snout close to the floor (Movie S3).

Head pitch was not the only feature distinguishing mobile olfactory exploration from rapid locomotion. While the distribution of head roll was unimodal in the first case, it appeared to be bimodal in the second (Fig. 3F), peaking at 5–10° toward the left and right sides. This phenomenon did not result from variability in IMU placement on the head, as segments from the same animal, and even from the same session, showed both orientations. During rapid locomotion, per-segment average roll angles showed a weak correlation with net azimuthal speed ($r = 0.15, p < 1\text{e}{-4}$). This suggests that head roll tends to occur when the animal follows a curved path, with the head tilting towards the concave side of the trajectory.

The vigor of olfactory exploration can be tracked by analyzing its spectral signature

Having observed a head kinematic signature linked with sniffing during slow locomotion, we wondered if a similar pattern could be detected in other behavioral contexts. Visual inspection of segments and their associated video snippets indeed suggested that $\simeq$ 10 Hz rhythmicity in the fore-aft non-gravitational acceleration and pitch angular velocity was a prevalent feature of head IMU recordings also manifesting in other clusters. To identify these clusters, we devised a "sniffing score" reflecting the magnitude of this spectral signature based on a per-segment ratio of continuous wavelet transform coefficients in specific frequency bands (see Fig. 4A and Methods).

In addition to cluster 37, associated with the slow locomotor regime identified in Fig. 3D, this approach revealed six clusters exhibiting a high sniffing score (Fig. 4B) and associated with the "orienting" and "rearing" categories. The strongest scores were observed for three "orienting" clusters not part of the rapid head movements identified in Fig. 2B, corresponding to small upward pitch movements or rightward or leftward head rotations with the head pitched down (Fig. 4C), and displaying net angular speed values in the 35–60 °/s

range. The two clusters with lower scores corresponded to rearing initiation and high rear with a vertically stretched body (Fig. 4C). These observations, along with the bimodal distribution of average head pitch angles associated with these clusters (Fig. 4C), suggest that animals preferentially engaged in robust sniffing either when investigating the floor or during rearing.

The average peak frequencies observed for these eight clusters with high sniffing score (9.1 ± 2.1 Hz and 8.2 ± 1.7 Hz for fore-aft non-gravitational acceleration and pitch angular velocity, respectively) were in the upper range of sniffing rates reported for rats during odor discrimination in stationary or head fixed conditions (6–9 Hz)[31,79], and were consistent with those observed in freely-moving animals tracking an odor trail (8–12 Hz)[80]. Overall, these results indicate that head IMU signals can be leveraged to obtain a detailed and quantitative readout of olfactory exploration, using a simple metric that can be applied independently of our pipeline.

The main types of grooming activities distribute into separate clusters with distinctive spectral signatures

Self-grooming is a major component of the rodent behavioral repertoire, encompassing a variety of actions often performed sequentially[81]. Monitoring grooming activities through video recordings is however challenging due to the animal's contorted posture (e.g. during body licking) or the rapid movements involved (e.g. during scratching). To assess DISSeCT's ability to parse this family of actions, we conducted a detailed examination of the 14 clusters assigned to the "grooming" category during segment inspection (Fig. 1C), which accounted for 19 % of all segments and 23 % of the total recording duration.

These clusters were characterized by a specific range of head tilts, with the less ambiguous ones (corresponding to body licking) exhibiting a unique combination of head pitch and roll angles (see cluster 17 in Fig. 5A). Unexpectedly, their most distinctive feature was a clear deviation from a linear relationship observed for all other clusters between the correlation of x and z-axes gyroscope measurements and the head pitch angle ($r = 0.95$, $p = 2.2\text{e}{-16}$; Fig. 5A). In other words, the coupling between roll and yaw head rotations during grooming appears to depart from a general biomechanical rule linking it with head pitch in all other conditions.

To assess whether "grooming" clusters corresponded to ethologically distinct actions, we relied on the labels assigned during the inspection phase (Fig. 5B; see Methods). The less ambiguous clusters (where more than 85 % of inspected segments received the same label) corresponded to body licking (clusters 0, 17 and 26), facial grooming (cluster 43), head and neck scratching (cluster 45) and hindpaw licking (cluster 30), each exhibiting a characteristic spectral signature (Fig. 5C). Another group of clusters had less homogeneous labels, with short bouts of facial grooming mixed with non-grooming related head orienting movements (cluster 8), or individual body licking strokes mixed with brief head movements occurring when the animal investigated the corners of the arena with a tilted head posture (cluster 21). The non grooming-related segments included in these clusters likely correspond to head motion patterns whose kinematics resemble those encountered during grooming. Kinematic ambiguities may also explain why ethologically distinct grooming actions were mixed in the same clusters (see clusters 43 and 25, Fig. 5B). Finally, a last group of clusters (clusters 19, 29, 32 and 33) contained various head orienting movements executed in the context of grooming episodes (at episode onset or offset or between adjacent grooming bouts).

These results show that our pipeline successfully isolated the main components of grooming as a family of clusters with distinctive postural, kinematic and spectral signatures. They also point out kinematic ambiguities which could be taken into account to design strategies aiming at achieving an even greater separation of grooming actions.

Behavioral patterns can be unveiled from cluster sequences

As detailed in the previous sections, DISSeCT effectively parsed head IMU signals into various behaviorally relevant components. To assess how these components assemble over time and potentially form sequential patterns, we fitted a simple hidden Markov model (HMM) with categorical emissions to the cluster sequence

(i.e. where categories correspond to clusters). Hyperparameter optimization resulted in 16 states, each associated with a vector of emission probabilities for the different clusters (see Methods; Fig. 6A). We observed that states often appropriately gathered clusters corresponding to the same movement executed in opposite leftward vs rightward directions (Fig. S8). Additionally, the highest emission probabilities for a given state most often corresponded to clusters belonging to the same behavioral category, allowing each state to be linked to a specific category (Fig. 6A).

The state transition matrix revealed that the most likely transitions occurred within the "rearing", "grooming" and "orienting" categories (Fig. 6B), indicating that these behaviors tended to manifest as multi-state sequences rather than isolated events. States associated with the "rearing" and "grooming" categories, in particular, almost never transitioned to one another, while the "locomoting", "rearing" and "orienting" categories did exhibit specific temporal relationships. These observations were further confirmed by examining the cluster transition matrix (Fig. S9).

Characteristic intra-category sequences included back-and-forth switching between states 1 and 3, representing an alternation between head scratching or hindpaw licking episodes and brief head orienting movements (Fig. 6C,D_1), and a 3-state chain (states 5, 8 and 10) representing rearing episodes (Fig. 6C). Prominent inter-category sequences included an alternation between rearing episodes and locomotion bouts (Fig. 6C,D_2, Fig. S10A), and a 3-state sequence (states 6, 2 and 7) representing episodes during which the animal lifted its head up while keeping its front paws on or close to the ground (Fig. 6C,D_3, Fig. S10B).

One key property of HMMs is that the probability of being in a given state for a particular observation changes over time depending of the local context. As a result, segments assigned to a given cluster by the GMM but occurring in different behavioral contexts might be associated with different states, potentially reflecting different behavioral motifs. To examine this, we plotted posterior probabilities for each cluster (Fig. S11A) and found clusters linked to multiple states. A detailed analysis of two dual-state clusters revealed that they could indeed be split into kinematically similar but ethologically distinct head movements based on the posterior probability of each segment. Cluster 12 was parsed into rearing segments (linked with state 8) and segments during which the head was pitched-up but the front paws remained on or near the ground (linked with state 2), as shown in Fig. S11B. Cluster 19 was parsed into brief orienting movements executed between head scratching and hindpaw licking bouts (linked with state 3) and single strokes during facial grooming or body licking (linked with state 11), as shown in Fig. S11C.

These observations demonstrate that a simple categorical HMM can be effectively used as a curation step after DISSeCT to merge ethologically-similar clusters and resolve certain kinematic ambiguities by splitting clusters containing similar head movements performed in different contexts.

Unsupervised analysis of behavioral alterations in a mouse model of Parkinson's disease and levodopa-induced dyskinesia

One advantageous feature of GMMs is their probabilistic nature, which makes them suitable for estimating how well a new data set conforms to a previously-established model and thus for detecting anomalies. To test the applicability of DISSeCT for quantifying behavioral changes, we used a classical toxin-induced mouse model of Parkinson's disease (PD). 6-OHDA lesioned mice (PD mice, $n = 12$) were obtained through unilateral nigrostriatal dopaminergic denervation using 6-OHDA injection (see Methods) while a control group received saline injection (sham mice, $n = 4$). PD mice exhibited decreased locomotion, impaired movement initiation and increased ipsiversive rotations relative to the lesion side, as previously described[82]. PD mice were then treated with daily injections of levodopa (L-DOPA), the first-line drug for the management of PD motor symptoms, along with sham animals. L-DOPA induced abnormal involuntary movements (AIMs) in PD mice, mimicking the common complication of this treatment in PD patients, known as levodopa-induced dyskinesia (LID). These AIMs, which affect the musculature on the side contralateral to the lesion[82], were accompanied by contraversive rotations.

Weekly IMU recordings were performed on PD and sham mice before and during L-DOPA treatment (Fig. 7A). DISSeCT was then run on sham data to generate a GMM representing the behavior of healthy animals (sham-GMM, see Methods). A joint S-UMAP embedding, computed across all animals in all conditions

using cluster labels from the sham-GMM, revealed a similar global structure for sham and PD mice, while entirely new regions emerged for PD mice during the LID phase (Fig. 7B). To quantify how much the motor repertoire of PD mice deviated from the one of healthy animals, we calculated the average log-likelihood of segments under the sham-GMM (see Methods). This "score under sham model" (SUSM) was only slightly lower in PD mice compared to sham mice before L-DOPA administration (Fig. 7C). After L-DOPA injection, the SUSM of PD mice decreased substantially (Fig. 7C), indicating that the sham-GMM poorly explained their behavior during the LID phase. Given the characteristic rotational behavior in this model and its inversion during the LID phase, we quantified rotations using IMU data as previously described[42]. PD mice displayed mild but significant ipsiversive rotations relative to the lesion side (-1.73 ± 0.80 turns per minute, vs -0.12 ± 0.10 for sham mice; $p = 1.1\text{e}{-3}$, Wilcoxon test) before L-DOPA injection and a pronounced circling behavior in the opposite direction after injection (17 ± 9 turns per minute). Rotations during the LID phase negatively correlated with the SUSM ($r = -0.89, p = 9.0\text{e}{-5}$; Fig. 7D), suggesting that most of the L-DOPA-induced behavioral remapping captured by our pipeline is linked with an overrepresentation of contraversive movements. Notably, these rotations also negatively correlated with the SUSM measured before injection ($r = -0.71, p = 9.5\text{e}{-3}$; Fig. 7E), suggesting that changes occurring in the PD state are predictive of the overall response to L-DOPA.

To further characterize behavioral changes in PD mice, we examined how frequently sham-GMM clusters were used in these mice in comparison with sham mice (Fig. 8A). Before L-DOPA injections, PD mice predominantly employed two clusters: one corresponding to episodes of immobility and the other to slow, mobile olfactory exploration with ipsiversive turns. During the LID phase, cluster usage shifted to a series of five clusters, all associated with contraversive movements (Fig. 8B,C). Among clusters whose usage did not change drastically, we focused on one representing upward pitch movements, which exhibited a slight positional shift in the embedding space across groups and conditions (Fig. 8D). This shift occurred along a gradient of head roll angles, with PD mice segments concentrated on opposite sides of the gradient before and after L-DOPA injection. This rearrangement corresponded to a mild ipsiversive bias of upward pitch movements in PD mice, which reversed to the opposite side under L-DOPA (Fig. 8E). This example illustrates how DISSeCT can be used to single out subtle changes in motor patterns.

Discussion

IMUs are becoming an increasingly popular tool for motion tracking in neurophysiological research, yet their potential for parsing rodent behavior remains largely untapped. Unsupervised approaches, which do not rely on human labels to uncover data structure, are essential to fully assess this potential. Drawing inspiration from previous work on accelerometer or GPS data collected from free-roaming animals[83–85] and egocentric videos recorded in a clinical context[86], we developed DISSeCT, a two-step unsupervised segmentation and clustering pipeline specifically designed for rodent IMU data.

Similar to current video-based unsupervised behavioral mapping procedures[13,15,16], DISSeCT was able to distinguish the main categories of rodent activities (Fig. 1), while also delivering fine-grained information on their various components. Our pipeline mapped the full repertoire of rapid three-dimensional head orienting movements (Fig. 2), offering a convenient alternative to methods based on threshold crossing detection in gyroscope measurements[39]. Interestingly, the duration and angular speed of the corresponding motifs co-varied to produce comparable head angular displacements (Fig. 2B), suggesting a default angular set point for orienting movements. DISSeCT also successfully isolated episodes of locomotion, characterized by a rhythmic $\simeq 6$ Hz pattern in head-vertical acceleration (Fig. 3). A similar signature at 8–10 Hz, scaling with animal speed (50–100 cm/s) was previously described in running rats[28]. The lower frequency observed here is probably related to the lower speed (10–15 cm/s) imposed by the restricted size of our arena. Inspection of a slow locomotion cluster revealed a $\simeq 10$ Hz signature of sniffing distributed along a specific combination of axes (Fig. 3D), and whose angular speed component had already been reported[31]. This signature, found in several other clusters (Fig. 4), could be used in the future as a standalone metric to assess the level of olfactory exploration. Finally, grooming activities emerged as a family of clusters with diverse time-frequency

properties, uniquely characterized by an uncoupling of the correlation between roll and yaw rotations from the head pitch angle (Fig. 5A). These results demonstrate that DISSeCT is indeed a relevant approach for in-depth IMU data mining applied to rodent behavioral mapping.

In this study, deciphering cluster content required visual inspection of concurrently acquired video recordings, allowing us to identify the head inertial signatures associated with various behavioral activities. The present mapping shall serve as a foundation for future studies, reducing the need for extensive video checks. Nonetheless, maintaining a parallel video stream should remain valuable, particularly when analyzing IMU data obtained under new experimental conditions. DISSeCT's companion software, INSpECT, is specifically designed to facilitate browsing through video snippets corresponding to individual segments, making it a versatile and essential add-on to our pipeline.

Cluster inspection revealed that ethologically-distinct motifs may sometimes share similar head kinematic properties, such as upward head movements and rearing initiation. Using a post-hoc categorical HMM to model the symbolic sequence resulting from segment clustering (Fig. 6), we show that some of these ambiguities can be resolved if the corresponding motifs occur within distinct identifiable sequences (Fig. S11). Other solutions to improve the separation of similar motifs include fitting GMMs to the corresponding subsets of segments using tailored features or employing hierarchical clustering methods. Missing body kinematic information could be obtained using additional IMUs, either subcutaneously implanted[47,51] or embedded into jackets worn by animals, keeping in mind that sensor position and orientation may vary over time[87]. Since maintaining a minimal concurrent video stream is a logical option for enabling rapid cluster inspection, as discussed above, head kinematic ambiguities could also potentially be resolved by integrating basic information extracted from a single-view video with IMU data in DISSeCT. This hybrid approach could enhance the discriminative power of our pipeline without significantly increasing the complexity of the experimental setup.

Our pipeline differs significantly from existing frameworks for video-based unsupervised analysis of rodent behavior. A commonly used approach involves extracting features from fixed-length video segments and embedding them into two-dimensional maps, where predominant behavioral patterns manifest as high density areas[13,15–17]. While effective for broad behavioral categorization, this method is not designed to precisely delineate behavioral motifs in time. Additionally, its drastic dimension reduction step prior to clustering likely diminishes its discriminative power[70]. In contrast, DISSeCT uses change point detection to accurately identify behavioral transitions and performs clustering in a high-dimensional space, minimizing information loss. Another technique employs an HMM variant to model time series derived from videos as a sequence of discrete elementary modules of varying length[18,61], optimizing segmentation and clustering simultaneously. We opted for a different strategy for two reasons. First, this HMM-based approach proved impractical for modeling IMU data acquired at several hundred hertz, leading to computation times that were at least an order of magnitude longer than those achieved by our pipeline (Table S3). Second, while HMMs jointly uncover hidden state boundaries and identities, we aimed to decouple segmentation from clustering to create a modular and adaptable platform.

DISSeCT's modular design allows for the integration of alternative clustering methods or iterative processes, such as additional rounds of feature extraction and segment clustering for finer-grained mapping of specific data subspaces. Automated scoring of user-defined activity patterns could also be achieved using a supervised learning algorithm downstream of the change point detection step. Regardless of whether segments are analyzed using supervised or unsupervised methods, feature selection remains a critical component of the pipeline. Instead of relying on principal components derived from a broad set of user-defined features, an interesting alternative would be to explore whether automatic feature engineering using variational autoencoders[62] could provide significant performance gains.

The combined resolution and versatility of IMU recordings open up a wide range of potential applications for our pipeline. By accurately detecting abrupt behavioral transitions and parsing brief behavioral motifs, DISSeCT provides the temporal resolution required for neuroethological studies aimed at uncovering the neuronal mechanisms underlying natural behavior[2]. One particular area where our pipeline could significantly advance the state-of-the-art is in analyzing behaviors that are challenging to accurately assess via video. Self-grooming, for example, is a complex behavior associated with contorted postures where IMUs offer considerable advantages over video recordings for in-depth analysis[16,26,51]. DISSeCT successfully parsed

the main components of self-grooming, shedding light on their distinct head kinematic signatures. While further refinement is needed to enhance the separation of grooming subcomponents, our approach introduces a powerful new method for automating the analysis of this elaborate syntactic behavior[81], with potential applications in the study of rodent models of neurological and psychiatric disorders[88].

By parsing high resolution kinematic data into an unsupervised probabilistic model suited for the quantification of behavioral anomalies, DISSeCT also meets the requirements of a robust deep phenotyping tool. In our use-case scenario of a mouse model of Parkinson's disease and levodopa-induced dyskinesia, our pipeline provided both a global measure of behavioral changes (Fig. 7) and a means to detect subtle motor alterations (Fig. 8). As such, DISSeCT holds potential both as a generalist platform for detecting the behavioral effects of experimentally induced perturbations, and as a tool for scrutinizing specific motor deficits.

In conclusion, given the unique advantages of IMUs over other methods of behavioral data collection, DISSeCT offers access to a new region of the "behavioral science space"[89], where in-depth and high-throughput behavioral analysis can be conducted with limited technical and environmental constraints, while keeping computational and data management costs minimal.

Material and Methods

Animals

Four adult male Long-Evans rats (Janvier Labs) and sixteen RJOrl:SWISS mice (CD-1, eight males and eight females, Janvier Labs) were used in this study. They were group-housed under standard conditions (12-hour non-inverted light/dark cycle, 21–24 °C, 50–60 % humidity) and provided with unrestricted access to food and water along with appropriate cage enrichment such as cardboard tunnels, shredded paper, and dental cotton rolls.

Surgical procedures

Wired IMU fixation in rats A wired IMU measuring $5.0 \times 7.5 \times 3.0$ mm and weighting 0.3 g (mptek.fr), housing a 9-axis digital inertial motion sensor (MPU-9250, Invensense) and operated via a dedicated USB acquisition board (mptek.fr) was used for rat recording sessions. The IMU featured a 2×4 pin header (TSM-104-01-S-DV, Samtec Inc.) used to connect the device to its counterpart female socket header (NPPC042KFMS-RC, Sullins Connector Solutions), which was permanently affixed to the animal's skull through a surgical procedure outlined below. Animals underwent this procedure at six weeks of age and were allowed one week to recover before starting the recording sessions.

The connector implantation surgery was carried out as follows. Preemptive analgesia was ensured by subcutaneously injecting buprenorphine (0.01 mg/kg) thirty minutes before the procedure. Anesthesia was induced with 3% isoflurane in oxygen delivered at 3 L/min. The animal was then positioned in a stereotaxic apparatus, with its eyes protected by ophtalmic gel. Throughout the remainder of the procedure, body temperature was kept at 37 °C using a homeothermic monitoring system, while anesthesia was maintained with 1% isoflurane delivered at 0.5 L/min. The animal's scalp was shaved and cleansed with povidone-iodine followed by 70% ethanol. Prior to incising the scalp, a subcutaneous injection of lidocaine (2%) was administered. The periosteum of the skull was gently scraped using a scalpel blade, followed by application of a 3% hydrogen peroxide solution. The female socket header, held upright using one arm of the apparatus, was then affixed to the cranial surface using self-curing dental adhesive (Super-Bond C&B, Sun Medical). Finally, the skin around the connector was sutured both anteriorly and posteriorly, and the animal was placed on a heating blanket while its recovery was monitored. Once fully alert and mobile, the rat was returned to its home cage.

6-OHDA lesion and wireless IMU fixation in mice A wireless IMU (WEAR wireless sensor v2.1, Champalimaud Scientific Hardware Platform) measuring $14 \times 6 \times 26$ mm and weighting 1.85 g, housing a 9-axis digital inertial motion sensor (MPU-9250, Invensense) and operated via a dedicated USB acquisition board (WEAR basestation, Champalimaud Scientific Hardware Platform), was used for mouse recording sessions. The IMU

was mounted on the head using a 1×4 pin header (M52-040023V0445, Harwin Inc.) on the IMU side and its counterpart female socket header (M52-5000445, Harwin Inc.) on the animal side (see surgical procedure below).

Unilateral nigrostriatal dopaminergic denervation was achieved at 9–10 weeks of age by injecting 6-hydroxydopamine hydrochloride (6-OHDA, Sigma-Aldrich) into the medial forebrain bundle (MFB) as previously described[90,91]. A subcutaneous dose of buprenorphine (0.1 mg/kg) and an intraperitoneal dose of desipramine (25 mg/kg, Sigma-Aldrich) were administered thirty minutes before initiating the procedure. The initial steps of the surgery, from anesthesia induction to skull preparation, were similar to the rat connector implantation protocol described above. For the Parkinsonian model (PD mice), the left MFB (−1.2 mm AP, 1.25 mm ML and −4.7 mm DV relative to Bregma) was injected with 1 µL of a 6-OHDA solution (3.2 µg/µL free-base concentration in 0.02% ascorbic acid, Sigma-Aldrich); whereas for control animals (sham mice), an equivalent volume of the vehicle solution (0.02% ascorbic acid) was administered. The injection was performed with a Hamilton syringe at a rate of 0.3 µL/min, with the syringe maintained in situ for two minutes before and four minutes after injection. The female socket header used for IMU fixation, held upright using one arm of the apparatus, was then affixed to the cranial surface using self-curing dental adhesive (Super-Bond C&B, Sun Medical). The skin around the connector was sutured both anteriorly and posteriorly, and the animal received a subcutaneous injection of Meloxicam (5 mg/kg, Metacam, Boehringer Ingelheim) and of a glucose-saline solution (Glucose 5%, 10 mL/kg, Osalia). Throughout the recovery phase, the mouse was placed on a heating blanket while being closely monitored. Once fully alert and mobile, it was returned to its home cage. To mitigate the risk of post-surgical complications and distress, mice received daily injections of glucose-saline solution and were nourished with a blend of Blédine (Blédina, Danone) and concentrated milk for a period spanning 6 to 10 days post-surgery, resulting in a 100% survival rate.

Data acquisition

Synchronized inertial and video data acquisition for quantifying rat behavior Rat recordings started one week after connector implantation. Once equipped with the wired IMU, rats were placed in a 70×70 cm arena with transparent glass floor and transparent PMMA walls, allowing for video monitoring from multiple angles, including from underneath, similarly to previous studies[47,92,93]. Inertial signals were acquired at 300 Hz with accelerometer and gyroscope ranges set at $\pm 2\,g$ and $\pm 1000\,^{\circ}/\mathrm{s}$, respectively. Video recordings were captured using five off-the-shelf industrial cameras (DMK 37BUX273, The Imaging Source) connected via USB to an acquisition computer. Frame acquisition was triggered at 30 Hz by an Arduino Uno running a custom pulse-width modulation script. These triggers were simultaneously recorded by the IMU acquisition board to precisely timestamp video frame acquisition within the IMU data stream. In our conditions, occasional frame drops were observed, at a frequency which gradually increased during continuous acquisition. We found that delivering triggers in one-minute chunks (consisting of 1800 frames), separated by three-second pauses, and saving the corresponding videos as separate files, drastically reduced the occurrence of frame drops to 0.023 % (0.39 ± 0.65 missing frame per chunk).

Synchronized inertial and video acquisition for quantifying behavior in a mouse model of Parkinsons's disease and levodopa-induced dyskinesia Mouse recordings started one week after 6-OHDA injection and connector implantation,

Considering that 6-OHDA injection and connector implantation were conducted at day 0, IMU recordings started on day 10. Once equipped with the wireless IMU, mice were placed in a 30×35 cm arena with transparent PMMA floor and walls. Video recordings were captured using two off-the-shelf industrial cameras (DMK 37BUX273, The Imaging Source) positioned below and on the side of the arena. Synchronized acquisition of video and IMU data was carried out using dedicated acquisition hardware (WEAR basestation, Champalimaud Scientific Hardware Platform) controlled via a custom pipeline programmed in Bonsai[94]. Videos were acquired at a rate of 50 Hz, while IMU data was recorded at 200 Hz with accelerometer and gyroscope ranges set at $\pm 2\,g$ and $\pm 1000\,^{\circ}/\mathrm{s}$, respectively.

Recordings were conducted over a five-week period, with one recording session per week (Fig. 7A). The first three sessions (days 10, 17 and 24) lasted 30 minutes each. Starting on the 24th day, mice received daily

intraperitoneal injections of L3,4-dihydroxyphenylalanine methyl (L-DOPA or levodopa, 3 days at 3 mg/kg and then 6 mg/kg, Sigma-Aldrich) and benserazide hydrochloride (12 mg/kg, Sigma-Aldrich), a peripheral decarboxylase inhibitor. The final two recording sessions were conducted on days 31 and 38, starting 10 minutes before and extending 60 minutes after levodopa injection.

Markerless 3D pose estimation

In the case of rat video recordings, the use of five different viewpoints enabled 3D pose estimation. Using the DeepLabCut toolbox[6], a unique deep neural network was trained to recognize twelve anatomical keypoints from all viewpoints, providing an independent 2D pose estimate for each video. The training dataset consisted of 1950 hand-labeled frames selected to encompass all different viewpoints and a wide range of animal positions and postures within the arena. Additional details regarding the parameters employed can be found in Table S1. Before each session, the intrinsic and extrinsic parameters of the cameras were estimated by recording a video of a ChArUco board moved throughout the arena, and applying the calibration pipeline available in Anipose[12]. This pipeline combines checkerboard detection using OpenCV with an optimization procedure aimed at minimizing reprojection errors.

Spatio-temporal filtering is a crucial step in obtaining high quality estimates of 3D pose dynamics[12,47,95]. Fot this purpose, we used the filtering pipelines integrated into Anipose. Median filtering was first applied to the five independent 2D pose estimates to alleviate the impact of sudden keypoint jumps. Triangulation was then performed using camera calibration and spatio-temporal regularization, taking advantage of the continuity and relative smoothness of animal movements, as well as the stability of bodypart lengths over time. Finally, a 3D median filter was applied to the resulting 3D pose estimate to reduce keypoint jitter. A summary of the parameters used for triangulation and filtering with Anipose is provided in Table S2.

Pre-processing of IMU data

Subtraction of sensor offsets The estimation of sensor offsets followed a procedure described previously[42]. For each IMU, data was collected while placing the device in a set of 50 different static positions using a custom robotic arm. Keeping only the periods of immobility, accelerometer offsets were estimated by minimizing the average squared difference between the Euclidean norm of offset-corrected acceleration measurements and $1\,g$. Gyroscope offsets were computed as the median gyroscope values observed during these periods of immobility. Accelerometer and gyroscope offsets were then subtracted from acceleration and angular velocity measurements, respectively.

Estimation of head tilt Head orientation relative to gravity (head tilt) can be accurately estimated from rodent IMU data using the appropriate Attitude and Heading Reference System (AHRS) filtering algorithm. For each IMU data point, an extended Kalman filter[96] with a previously validated set of parameters ($var_{acc} = 0.002$, $var_{gyr} = 0.75$)[42] was used to compute a unit quaternion describing sensor orientation relative to an arbitrary, gravity-polarized Earth reference frame (absolute reference frame). This quaternion represents a 3D rotation which, when applied to this absolute reference frame, aligns it with the sensor's orientation. To estimate head orientation relative to gravity (head tilt, i.e. the orientation of the gravitational acceleration vector in the head reference frame), the inverse rotation, corresponding to the conjugate quaternion, was applied to the unit vector $\begin{pmatrix} 0 & 0 & 1 \end{pmatrix}$.

Change point detection

Algorithm Pre-processed IMU data were segmented by applying the PELT (Pruned Exact Linear Time) change point detection method to z-scored gravitational acceleration and angular velocity time series (non-gravitational acceleration was not used here because of its lower signal-to-noise ratio). This algorithm leverages dynamical programming to find the optimal number and location of change points in time series by minimizing a cost function (see Supplementary Methods), while maintaining a linear computation time[64]. Specifically, we used a kernel-based change point detection algorithm[63] relying on a Gaussian kernel, which offers the advantage of

being characteristic; meaning that changes in the mean of the embedded data correspond to changes in the data distribution[68,97,98] (See Supplementary methods).

Selection of the penalty parameter Penalty for segmenting IMU data was set to $\lambda = 14$, yielding segments with a median duration of $\simeq 350$ ms, to align with the typical duration of individual elements identified by MoSeq, a depth video-based unsupervised method used to parse mouse behavior[61]. This duration also matched the timescale of the finest behavioral motifs identifiable by a human observer (see Fig. S2).

Segment clustering

Feature extraction A total of 240 features were extracted for each segment, effectively capturing postural and kinematic aspects of the animal's activity within the segment. These features included statistics on the signals, parameters derived from head tilt and changes in head azimuth, as well as coefficients from continuous wavelet transforms of angular velocity and non-gravitational acceleration data (see Supplementary Methods).

Outlier removal A small subset of segments exhibited outlier feature values, leading to convergence issues during clustering. To address this, a threshold of 35 was applied to absolute robustly rescaled feature values, obtained through median subtraction and division by the interquartile range. Segments exceeding this threshold were excluded from the analysis. These outliers consisted of extended periods of immobility (17 %) and high magnitude events corresponding either to vigorous movements (21 %, typically $> 1\,g$ of non-gravitational acceleration and $> 500\,^{\circ}/\mathrm{s}$ of angular speed) or to wet dog shakes[99] (with the norm of non-gravitational acceleration and angular velocity exceeding 2 g and $1000\,^{\circ}/\mathrm{s}$, respectively). The latter type of outlier is characterized by a high amplitude 16–20 Hz oscillatory regime whose wavelet transform signature artificially contaminates neighboring segments due to edge effects (Fig. S5). Consequently, segments immediately before and after wet-dog shake events were also excluded from the analysis. This curation procedure resulted in the exclusion of 1.7% of all segments, representing 3.7% of the total recording duration.

Gaussian Mixture Model clustering After having excluded outlier segments as described above, features were z-scored and projected in a lower-dimensional space via PCA. The number of principal components was determined through cross-validation[100] (Fig. S4A). A clustering step was then executed on the resulting projection using scikit-learn's Gaussian Mixture Model implementation. Model selection and hyperparameter tuning were conducted through grid search with the aim of maximizing the Bayesian information criterion (BIC)[101]. The grid search encompassed covariance type (full, diagonal, tied, spherical) and number of clusters (ranging from 5 to 100), employing 10 random initialization for each hyperparameter set (Fig. S4B), yielding 48 mixture components with full covariance matrices. Clusters were identified by assigning each segment to the mixture component to which it most likely belonged to. More information on the procedure is available in the Supplementary Methods.

Two-dimensional embedding of clustered data For visualization, z-scored features associated with individual segments were projected onto a two-dimensional space using supervised UMAP embedding[102] (with parameters: $n_neighbors = 8$, $target_weight = 1e-9$, $min_dist = 1e-3$). This method incorporates cluster label information to guide the embedding process, ensuring that points belonging to the same cluster are placed closer together, while still preserving the overall structure of the data. Feature values computed on outlier segments, excluded from the clustering step, were also included in this embedding step, albeit ignoring the associated label.

Sparse clusters The dispersion of individual mixture components was evaluated by calculating the trace of their covariance matrix. Clusters associated with the three mixtures exhibiting the highest trace values were considered as sparse clusters (Fig. S7A). Collectively, these clusters accounted for 2.1% of the total number of segments and for 1.0% of the total recording duration. Examination of these clusters revealed that they comprised heterogeneous segments, corresponding to occasional brief and vigorous movements (approaching $2\,g$ of non-gravitational acceleration and $1000\,^{\circ}/\mathrm{s}$ of angular velocity), and which appeared highly dispersed on the UMAP embedding (Fig. S7B). The presence of such clusters is expected, as GMMs tend to assign rare and/or outlier observations to a few highly dispersed, poorly informative clusters[103].

Ad hoc metrics

Per-segment average head roll and pitch angles Within the head's reference frame, gravitational acceleration can be conceptualized as a unit position vector, more specifically as a point on a head-centered sphere of radius $1\,g$ that fully represents head tilt. Due to the non-linearity of trigonometric functions, calculating the average head tilt over a certain period of time is not as straightforward as averaging the coordinates of this point. Instead, the average head tilt can determined as the centroid (estimated mean direction) of a von Mises-Fisher distribution fitted to the various positions assumed by this point on the unit sphere. Noting this centroid $\boldsymbol{\mu} = \begin{bmatrix}\mu_x & \mu_y & \mu_z\end{bmatrix}$, the corresponding average head roll and pitch angles are computed as the opposite values of the polar coordinates of $\boldsymbol{\mu}$:

- $\mu_{roll} = -arctan2(\mu_y, \mu_z)$
- $\mu_{pitch} = -arctan2(-\mu_x, \sqrt{\mu_y^2 + \mu_z^2})$

Detection of immobility Angular speed was computed as the Euclidean norm of offset-corrected gyroscope measurements and then smoothed using a Gaussian kernel ($\sigma = 100\,\text{ms}$). The animal was considered immobile when this value was below $20\,^\circ/\text{s}$.

Sniffing score To identify segments displaying the $\simeq 10\,\text{Hz}$ sniffing signature primarily observed in the naso-occipital non-gravitational acceleration (noted $a_{nG,x}$) and pitch angular velocity (noted ω_y), we designed an *ad hoc* sniffing score (noted Φ). This metric was constructed empirically using the maximal wavelet power (noted W) across all available frequencies (1–20 Hz) or frequencies largely above ($> 14\,\text{Hz}$), below ($< 6\,\text{Hz}$) or within the band of interest (8.5–11 Hz):

$$\Phi = \alpha \cdot \frac{W^{a_{nG,x}}_{8.5-11\,\text{Hz}} + W^{\omega_y}_{8.5-11\,\text{Hz}}}{W^{\omega_x}_{fullrange} + W^{a_{nG,x}}_{<6\,\text{Hz}} + W^{a_{nG,x}}_{>14\,\text{Hz}} + W^{a_{nG,y}}_{fullrange} + W^{a_{nG,z}}_{fullrange} + 100}$$

where α is a factor that penalizes Φ for segments associated with rapid head reorientation or large roll angles, which often exhibit non-sniffing related 8–10 Hz signatures. Brief reorientations can indeed last approximately 100 ms, thereby displaying a $\simeq 10\,\text{Hz}$ time-frequency signature, while substantial roll angles are linked to body grooming, an activity with a typical $\simeq 4\,\text{Hz}$ pattern that may manifest a notable harmonic at $\simeq 8\,\text{Hz}$. This factor was computed as follows:

$$\alpha = \frac{1}{1+\eta} \cdot \frac{1}{1+|\mu_{roll}|}$$

with η representing the heading change rate in $^\circ/\text{s}$ (the angle between heading estimates at the beginning and end of a segment, divided by its duration) and μ_{roll} the roll of the average tilt in radians.

2-dimensional visualization of per-segment head attitude and heading trajectories

To visualize 3D head orientation during individual segments, we overlaid the trajectories of head-tilt (pitch and roll) and 3D-heading (pitch and azimuth). Head-tilt was represented using the coordinates of the Earth's vertical in the head reference frame, estimated by the gravitational component of acceleration signals obtained through AHRS filtering of IMU data (see Methods). The estimate of 3D-heading was derived by applying the rotation from the head reference frame to a gravity-polarized Earth reference frame, achieved via AHRS filtering, to the x-axis of the sensor's reference frame. The resulting vector was then rotated along the vertical axis to set the first sample of the segment to an azimuth of 0 in the Earth reference frame (see Supplementary Methods).

These 3-dimensional time series of unit-norm vectors were projected to two dimensions using a Lambert Azimuthal-Equal Area projection, centered either on the *North Pole* (i.e., $(0, 0, 1)$) or *South Pole* (i.e., $(0, 0, -1)$) of the respective reference frames. The trajectories were then superimposed on a single plot using the cartography module cartopy.

Post hoc categorical HMM

A categorical hidden Markov model was fitted to the sequence of clusters obtained from GMM modeling, using the implementation provided by hmmlearn. Segments deemed outliers, as defined earlier, were treated as an independent cluster. Hyperparameter optimization was conducted via BIC maximization, employing ten independent fits per hyperparameter value, to determine the optimal number of states. More information on the procedure is available in the Supplementary Methods.

Inspection and labeling of video segments

Software To expedite the inspection and labeling of video snippets corresponding to selected segments, we contracted AquiNeuro to develop custom software. This software, named INSpECT (Interface for Navigating Spates of video Excerpts and Categorizing Them), features a graphical user interface and keyboard shortcuts for navigating a series of multi-camera video snippets and assigning user-defined labels to each. The code of INSpECT is available at https://github.com/rfayat/INSpECT.

Cluster inspection and labeling To evaluate the content of any particular cluster, we began by isolating segments assigned to that cluster with a duration suitable for visual inspection, using a threshold of 233 ms (equivalent to 7 video frames). From this selection, we then identified the top one hundred segments with the highest posterior probabilities of belonging to the cluster and imported the corresponding video snippets into INSpECT. An observer was tasked with assigning a unique main category to each one (choosing from the following options: orienting, rearing, grooming, locomoting, idle or other), along with an additional unconstrained label providing a more detailed description of the ongoing behavioral activity. The category "orienting" was used for rapid changes in head orientation that did not coincide with rearing or grooming episodes. "Rearing" referred to situations during which the animal stood on its hind legs with both forelegs above the ground. "Grooming" encompassed all types of self-cleaning activities. The category "idle" applied to situations where the animal stood immobile, occasionally accompanied by small twitches or head sways. For each cluster, we calculated a "consistency score" ranging between 0 and 1, and representing the occurrence of the predominant category within inspected segments.

Mapping of healthy and pathological mouse behavior

Pipeline implementation Change point detection was conducted on all available mouse IMU data. Similar to the approach taken with rat data, the penalty level was set to a value ($\lambda = 9$ here) resulting in a median segment duration of $\simeq 350$ ms. The same set of features used for rats was extracted for each segment, and the identical procedure for detecting outlier segments was applied, leading to the exclusion of 0.7% of all segments from the clustering step. These segments primarily consisted of extended periods of immobility and accounted for 10.6% of the total recording duration.

Parameters for feature rescaling, PCA dimension reduction and Gaussian mixture modeling were solely fitted on sham (healthy) mouse recordings before levodopa treatment. Hyperparameter optimization based on BIC maximization yielded 33 clusters with full covariance matrices. Based on this model, cluster assignment was then carried out for all other non-outlier segments (from sham mice under levodopa treatment and PD mice in all conditions).

The final supervised UMAP embedding was fitted using only one out of three segments for computational efficiency and using the same hyperparameters as for rat data, and was then applied for all mice and conditions.

Score under a sham model The Gaussian mixture model fitted on segments from sham mice prior to levodopa treatment served as the reference model for computing the log-likelihood of all available segments. For any given segment, this value indicates how likely this segment is under the assumption that the mouse is healthy. The average log-likelihood across a set of segments, referred to hereafter as the *score under a sham model* (SUSM) was considered a measure of how accurately the corresponding IMU recording bout could be described by this reference model. In essence, the SUSM provided a practical means of evaluating behavioral changes, with a decrease in this value indicating a discrepancy between the healthy model and the examined data.

Overlay of behavioral maps To overlay UMAP scatter plots (Fig. 8D), we first transformed them into two-dimensional histograms using a 500×500 grid. The bin values were then rescaled to fall between 0 and 1 by dividing each value by the 99.8^{th} percentile of the histogram, with any values exceeding one clipped to one to mitigate the impact of outliers. The resulting histograms were visualized as an image using the red, green and blue channels for PD-LDOPA, PD and sham mice, respectively.

References

1. Sousa, N., Almeida, O. F. & Wotjak, C. T. A hitchhiker's guide to behavioral analysis in laboratory rodents. *Genes, Brain and Behavior* **5,** 5–24 (2006).
2. Datta, S. R., Anderson, D. J., Branson, K., Perona, P. & Leifer, A. Computational Neuroethology: A Call to Action. *Neuron* **104,** 11–24. ISSN: 10974199 (2019).
3. Kennedy, A. The what, how, and why of naturalistic behavior. *Current Opinion in Neurobiology* **74,** 102549 (2022).
4. Berman, G. J. Measuring behavior across scales. *BMC Biology* **16,** 23–34 (2018).
5. Von Ziegler, L., Sturman, O. & Bohacek, J. Big behavior: challenges and opportunities in a new era of deep behavior profiling. *Neuropsychopharmacology* **46,** 33–44 (2021).
6. Mathis, A. *et al.* DeepLabCut: Markerless Pose Estimation of User-Defined Body Parts with Deep Learning. *Nature Neuroscience* **21,** 1281–1289 (2018).
7. Pereira, T. D. *et al.* Fast Animal Pose Estimation Using Deep Neural Networks. *Nature methods* **16,** 117–125 (2019).
8. Graving, J. M. *et al.* DeepPoseKit, a software toolkit for fast and robust animal pose estimation using deep learning. *Elife* **8,** e47994 (2019).
9. Nath, T. *et al.* Using DeepLabCut for 3D markerless pose estimation across species and behaviors. *Nature Protocols* **14,** 2152–2176 (2019).
10. Deepfly3D, a deep learning-based approach for 3D limb and appendage tracking in tethered, adult Drosophila. *eLife* **8,** 1–23 (2019).
11. Automated markerless pose estimation in freely moving macaques with OpenMonkeyStudio. *Nature Communications* **11,** 1–12 (2020).
12. Karashchuk, P. *et al.* Anipose: A toolkit for robust markerless 3D pose estimation. *Cell Reports* **36,** 109730.
13. Dunn, T. W. *et al.* Geometric Deep Learning Enables 3D Kinematic Profiling across Species and Environments. *Nature Methods,* 1–10 (2021).
14. Schneider, A. *et al.* 3D pose estimation enables virtual head fixation in freely moving rats. *Neuron* **110,** 2080–2093.e10 (2022).
15. Marshall, J. D. *et al.* Continuous Whole-Body 3D Kinematic Recordings across the Rodent Behavioral Repertoire. *Neuron* **109,** 420–437.e8 (2021).
16. Mimica, B. *et al.* Behavioral decomposition reveals rich encoding structure employed across neocortex in rats. *Nature Communications* **14,** 3947 (2023).
17. Hardcastle, K. *et al.* Differential kinematic coding in sensorimotor striatum across species-typical and learned behaviors reflects a difference in control. *bioRxiv,* 2023.10.13.562282 (2023).
18. Weinreb, C. *et al.* Keypoint-MoSeq: parsing behavior by linking point tracking to pose dynamics. *Nature Methods* **21,** 1329–1339 (2024).
19. Nourizonoz, A. *et al.* EthoLoop: automated closed-loop neuroethology in naturalistic environments. *Nature methods* **17,** 1052–1059 (2020).
20. Brown, D. D., Kays, R., Wikelski, M., Wilson, R. & Klimley, A. Observing the unwatchable through acceleration logging of animal behavior. *Animal Biotelemetry* **1,** 20 (2013).
21. Kays, R. *et al.* The Movebank system for studying global animal movement and demography. *Methods in Ecology and Evolution* **13,** 419–431 (2022).
22. Chapa, J. M., Maschat, K., Iwersen, M., Baumgartner, J. & Drillich, M. Accelerometer systems as tools for health and welfare assessment in cattle and pigs – A review. *Behavioural Processes* **181,** 104262 (2020).

23. Garcia-De-Villa, S., Casillas-Perez, D., Jimenez-Martin, A. & Garcia-Dominguez, J. J. Inertial Sensors for Human Motion Analysis: A Comprehensive Review. *IEEE Transactions on Instrumentation and Measurement* **72** (2023).
24. Onnela, J. P. Opportunities and challenges in the collection and analysis of digital phenotyping data. *Neuropsychopharmacology* **46,** 45–54 (2021).
25. Gu, C., Lin, W., He, X., Zhang, L. & Zhang, M. IMU-based motion capture system for rehabilitation applications: A systematic review. *Biomimetic Intelligence and Robotics* **3,** 100097 (2023).
26. Venkatraman, S., Jin, X., Costa, R. M. & Carmena, J. M. Investigating Neural Correlates of Behavior in Freely Behaving Rodents Using Inertial Sensors. *Journal of Neurophysiology* **104,** 569–575 (2010).
27. Chen, C. Y., Yang, C. C., Lin, Y. Y. & Kuo, T. B. Locomotion-induced hippocampal theta is independent of visual information in rats during movement through a pipe. *Behavioural Brain Research* **216,** 699–704 (2011).
28. Ledberg, A. & Robbe, D. Locomotion-related oscillatory body movements at 6–12 Hz modulate the hippocampal theta rhythm. *PloS one* **6,** e27575 (2011).
29. Structured Variability in Purkinje Cell Activity during Locomotion. *Neuron* **87,** 840–852 (2015).
30. Pasquet, M. O. *et al.* Wireless inertial measurement of head kinematics in freely-moving rats. *Scientific reports* **6,** 35689 (2016).
31. Kurnikova, A., Moore, J. D., Liao, S.-M., Deschênes, M. & Kleinfeld, D. Coordination of orofacial motor actions into exploratory behavior by rat. *Current Biology* **27,** 688–696 (2017).
32. Dugué, G. P., Tihy, M., Gourévitch, B. & Léna, C. Cerebellar re-encoding of self-generated head movements. *Elife* **6,** e26179 (2017).
33. Klaus, A. *et al.* The Spatiotemporal Organization of the Striatum Encodes Action Space. *Neuron* **95,** 1171–1180.e7 (2017).
34. Dhawale, A. K. *et al.* Automated long-Term recording and analysis of neural activity in behaving animals. *eLife* **6,** 1–40 (2017).
35. Wilson, J. J., Alexandre, N., Trentin, C. & Tripodi, M. Three-dimensional representation of motor space in the mouse superior colliculus. *Current Biology* **28,** 1744–1755 (2018).
36. Menardy, F., Varani, A. P., Combes, A., Léna, C. & Popa, D. Functional alteration of cerebello–cerebral coupling in an experimental mouse model of Parkinson's disease. *Cerebral Cortex* **29,** 1752–1766 (2019).
37. Meyer, A. F., O'Keefe, J. & Poort, J. Two distinct types of eye-head coupling in freely moving mice. *Current Biology* **30,** 2116–2130 (2020).
38. Bouvier, G., Senzai, Y. & Scanziani, M. Head movements control the activity of primary visual cortex in a luminance-dependent manner. *Neuron* **108,** 500–511 (2020).
39. Guitchounts, G., Masis, J., Wolff, S. B. & Cox, D. Encoding of 3D head orienting movements in the primary visual cortex. *Neuron* **108,** 512–525 (2020).
40. Angelaki, D. E. *et al.* A gravity-based three-dimensional compass in the mouse brain. *Nature Communications* **11** (2020).
41. Mallory, C. S. *et al.* Mouse entorhinal cortex encodes a diverse repertoire of self-motion signals. *Nature communications* **12,** 671 (2021).
42. Fayat, R. *et al.* Inertial Measurement of Head Tilt in Rodents: Principles and Applications to Vestibular Research. *Sensors* **21** (2021).
43. Demars, F. *et al.* Post-Trauma Behavioral Phenotype Predicts the Degree of Vulnerability to Fear Relapse after Extinction in Male Rats. *Current Biology* **32,** 3180–3188.e4 (2022).
44. Younk, R. & Widge, A. Quantifying defensive behavior and threat response through integrated headstage accelerometry. *Journal of Neuroscience Methods* **382,** 109725 (2022).
45. Parker, P. R., Abe, E. T., Leonard, E. S., Martins, D. M. & Niell, C. M. Joint coding of visual input and eye/head position in V1 of freely moving mice. *Neuron* **110,** 3897–3906.e5 (2022).
46. Silasi, G., Boyd, J. D., Ledue, J. & Murphy, T. H. Improved methods for chronic light-based motor mapping in mice: automated movement tracking with accelerometers, and chronic EEG recording in a bilateral thin-skull preparation. *Frontiers in neural circuits* **7,** 123 (2013).

47. Monsees, A. *et al.* Estimation of skeletal kinematics in freely moving rodents. *Nature Methods* **19,** 1500–1509 (2022).
48. Shih, Y. H. & Young, M. S. Integrated digital image and accelerometer measurements of rat locomotor and vibratory behaviour. *Journal of Neuroscience Methods* **166,** 81–88 (2007).
49. Lu, J. *et al.* Development of implantable wireless sensor nodes for animal husbandry and MedTech innovation. *Sensors* **18,** 1–12 (2018).
50. Wright, J. P. *et al.* A fully implantable wireless bidirectional neuromodulation system for mice. *Biosensors and Bioelectronics* **200,** 113886 (2022).
51. Chen, M. *et al.* Wireless AI-Powered IoT Sensors for Laboratory Mice Behavior Recognition. *IEEE Internet of Things Journal* **9,** 1899–1912 (Feb. 2022).
52. Massot, B. *et al.* ONEIROS, a new miniature standalone device for recording sleep electrophysiology, physiology, temperatures and behavior in the lab and field. *Journal of Neuroscience Methods* **316,** 103–116 (2019).
53. De Groot, A. *et al.* Ninscope, a versatile miniscope for multi-region circuit investigations. *eLife* **9,** 1–24 (2020).
54. Jonathan P. Newman *et al.* A unified open-source platform for multimodal neural recording and perturbation during naturalistic behavior. *bioRxiv* **44** (2023).
55. Gomaa, W. & Khamis, M. A. *A perspective on human activity recognition from inertial motion data* **28,** 20463–20568 (Springer London, 2023).
56. Kleanthous, N. *et al.* A survey of machine learning approaches in animal behaviour. *Neurocomputing* **491,** 442–463 (2022).
57. Wang, G. Machine learning for inferring animal behavior from location and movement data. *Ecological Informatics* **49,** 69–76 (2019).
58. Bermudez Contreras, E., Sutherland, R. J., Mohajerani, M. H. & Whishaw, I. Q. Challenges of a small world analysis for the continuous monitoring of behavior in mice. *Neuroscience and Biobehavioral Reviews* **136,** 104621 (2022).
59. Parker, P. R. *et al.* Distance estimation from monocular cues in an ethological visuomotor task. *eLife* **11** (2022).
60. Meyer, A. F., Poort, J., O'Keefe, J., Sahani, M. & Linden, J. F. A Head-Mounted Camera System Integrates Detailed Behavioral Monitoring with Multichannel Electrophysiology in Freely Moving Mice. *Neuron* **100,** 46–60.e7 (2018).
61. Wiltschko, A. B. *et al.* Mapping Sub-Second Structure in Mouse Behavior. *Neuron* **88,** 1121–1135 (2015).
62. Luxem, K. *et al.* Identifying behavioral structure from deep variational embeddings of animal motion. *Communications Biology* **5** (2022).
63. Arlot, S., Celisse, A. & Harchaoui, Z. A kernel multiple change-point algorithm via model selection. *Journal of machine learning research* **20,** 1–56 (2019).
64. Killick, R., Fearnhead, P. & Eckley, I. A. Optimal Detection of Changepoints With a Linear Computational Cost. *Journal of the American Statistical Association* **107,** 1590–1598 (2012).
65. Truong, C., Oudre, L. & Vayatis, N. ruptures: change point detection in Python. *arXiv:1801.00826,* 1–5 (2018).
66. Lebarbier, É. Detecting multiple change-points in the mean of Gaussian process by model selection. *Signal processing* **85,** 717–736 (2005).
67. Cappello, L. & Padilla, O. H. M. Bayesian variance change point detection with credible sets. *arXiv:2211.14097* (2022).
68. Celisse, A., Marot, G., Pierre-Jean, M. & Rigaill, G. New efficient algorithms for multiple change-point detection with reproducing kernels. *Computational Statistics & Data Analysis* **128,** 200–220 (2018).
69. Berman, G. J., Choi, D. M., Bialek, W. & Shaevitz, J. W. Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of The Royal Society Interface* **11,** 20140672 (2014).
70. Todd, J. G., Kain, J. S. & De Bivort, B. L. Systematic exploration of unsupervised methods for mapping behavior. *Physical Biology* **14** (2017).
71. Lloyd, S. Least squares quantization in PCM. *IEEE transactions on information theory* **28,** 129–137 (1982).

72. Bishop, C. M. *Pattern recognition and machine learning* (Springer, 2006).
73. McLachlan, G. J., Lee, S. X. & Rathnayake, S. I. Finite mixture models. *Annual review of statistics and its application* **6,** 355–378 (2019).
74. Bouveyron, C., Celeux, G., Murphy, T. B. & Raftery, A. E. *Model-based clustering and classification for data science: with applications in R* (Cambridge University Press, 2019).
75. Dempster, A. P., Laird, N. M. & Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society: series B (methodological)* **39,** 1–22 (1977).
76. Cavagna, G. A. Storage and utilization of elastic energy in skeletal muscle. *Exercise and sport sciences reviews* **5,** 89–130 (1977).
77. Skinner, R. & Garcia-Rill, E. The mesencephalic locomotor region (MLR) in the rat. *Brain research* **323,** 385–389 (1984).
78. Findley, T. M. *et al.* Sniff-synchronized, gradient-guided olfactory search by freely moving mice. *Elife* **10,** e58523 (2021).
79. Uchida, N. & Mainen, Z. F. Speed and accuracy of olfactory discrimination in the rat. *Nature neuroscience* **6,** 1224–1229 (2003).
80. Khan, A. G., Sarangi, M. & Bhalla, U. S. Rats track odour trails accurately using a multi-layered strategy with near-optimal sampling. *Nature communications* **3,** 703 (2012).
81. Berridge, K. C., Fentress, J. C. & Parr, H. Natural syntax rules control action sequence of rats. *Behavioural Brain Research* **23,** 59–68 (1987).
82. Lundblad, M., Picconi, B., Lindgren, H. & Cenci, M. A model of L-DOPA-induced dyskinesia in 6-hydroxydopamine lesioned mice: relation to motor and cellular parameters of nigrostriatal function. *Neurobiology of disease* **16,** 110–123 (2004).
83. Patin, R., Etienne, M. P., Lebarbier, E., Chamaillé-Jammes, S. & Benhamou, S. Identifying stationary phases in multivariate time series for highlighting behavioural modes and home range settlements. *Journal of Animal Ecology* **89,** 44–56 (2020).
84. Sur, M. *et al.* Supervised versus unsupervised approaches to classification of accelerometry data. *Ecology and Evolution* **13,** 1–11 (2023).
85. Liu, J., Bailey, D. W., Cao, H., Cao Son, T. & Tobin, C. T. Development of a Novel Classification Approach for Cow Behavior Analysis Using Tracking Data and Unsupervised Machine Learning Techniques. *Sensors* **24,** 4067 (2024).
86. Perochon, S. & Oudre, L. *Unsupervised Action Segmentation of Untrimmed Egocentric Videos* in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2023), 1–5.
87. Venkatraman, S., Long, J. D., Pister, K. S. J. & Carmena, J. M. Wireless inertial sensors for monitoring animal behavior. *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference* **2007,** 378–81 (2007).
88. Kalueff, A. V. *et al.* Neurobiology of rodent self-grooming and its value for translational neuroscience. *Nature Reviews Neuroscience* **17,** 45–59 (2016).
89. Gomez-Marin, A., Paton, J. J., Kampff, A. R., Costa, R. M. & Mainen, Z. F. Big behavioral data: Psychology, ethology and the foundations of neuroscience. *Nature Neuroscience* **17,** 1455–1462 (2014).
90. Sebastianutto, I., Maslava, N., Hopkins, C. R. & Cenci, M. A. Validation of an improved scale for rating l-DOPA-induced dyskinesia in the mouse and effects of specific dopamine receptor antagonists. *Neurobiology of disease* **96,** 156–170 (2016).
91. Coutant, B. *et al.* Cerebellar stimulation prevents Levodopa-induced dyskinesia in mice and normalizes activity in a motor network. *Nature Communications* **13,** 1–17 (2022).
92. Hsu, A. I. & Yttri, E. A. B-SOiD, an open-source unsupervised algorithm for identification and fast prediction of behaviors. *Nature Communications* **12,** 5188 (Aug. 31, 2021).
93. Sarkar, I. *et al.* Evaluation of deep lift pose models for 3D rodent pose estimation based on geometrically triangulated data. *arXiv:2106.12993* (June 24, 2021).
94. Lopes, G. *et al.* Bonsai: An event-based framework for processing and controlling data streams. *Frontiers in Neuroinformatics* **9,** 1–14 (2015).

95. Zhang, L., Dunn, T., Marshall, J., Olveczky, B. & Linderman, S. *Animal pose estimation from video data with a hierarchical von Mises-Fisher-Gaussian model* in *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics* (eds Banerjee, A. & Fukumizu, K.) **130** (PMLR, Apr. 2021), 2800–2808.
96. Sabatini, A. M. Kalman-filter-based orientation determination using inertial/magnetic sensors: Observability analysis and performance evaluation. *Sensors* **11,** 9182–9206 (2011).
97. Harchaoui, Z. & Cappé, O. *Retrospective mutiple change-point estimation with kernels* in *2007 IEEE/SP 14th Workshop on Statistical Signal Processing* (2007), 768–772.
98. Truong, C., Oudre, L. & Vayatis, N. Selective review of offline change point detection methods. *Signal Processing* **167,** 107299 (2020).
99. Dickerson, A. K., Mills, Z. G. & Hu, D. L. Wet mammals shake at tuned frequencies to dry. *Journal of the Royal Society, Interface* **9,** 3208–3218 (Dec. 2012).
100. Josse, J. & Husson, F. Selecting the number of components in principal component analysis using cross-validation approximations. *Computational Statistics & Data Analysis* **56,** 1869–1879 (2012).
101. Schwarz, G. Estimating the dimension of a model. *The annals of statistics,* 461–464 (1978).
102. McInnes, L., Healy, J. & Melville, J. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426* (2018).
103. Fraley, C. & Raftery, A. E. Model-based clustering, discriminant analysis, and density estimation. *Journal of the American statistical Association* **97,** 611–631 (2002).
104. Paxinos, G. & Franklin, K. B. J. *The Mouse Brain in Stereotaxic Coordinates, Compact, Third Edition: The coronal plates and diagrams* 3rd ed., 256 (Academic Press, 2008).

Code availability

The code of DISSeCT is available at https://github.com/rfayat/DISSeCT. The code of INSpECT, a companion software used to examine and label video snippets, is available at https://github.com/rfayat/INSpECT.

Acknowledgements

We extend our gratitude to Ahmed Boughdiri and Stefano Bettani for their assistance in setting up the multi-camera video acquisition system, and to Idriss Tsayem Ngueguim for his help in rat implantation surgeries. We also thank Ana Margarida Pinto for her help in establishing the acquisition system used for mouse recordings, and Adja Sissokho for her assistance with mouse experiments. Additionally, we are grateful to Laurent Oudre, Éric Burguière, Ashesh Dhawale and Bence Ölveczky for their valuable feedback on the project. Our thanks also go to the IBENS computing platform, fablab, and animal facility for their support.

Author contribution

R.F. assembled and programmed IMU and video recording setups, designed and performed rat experiments, co-designed the architectures of DISSeCT and INSpECT, wrote the DISSeCT code, performed the analysis, generated the figures and drafted the manuscript. M.S. implemented the mouse model of Parkinson's disease and levodopa-induced dyskinesia and conducted mouse recordings. D.P. and C.L. contributed to project design and edited the manuscript. P.L. and G.D. co-supervised R.F.'s work, co-designed the DISSeCT pipeline and wrote the manuscript.

Funding

This work was supported by grants from the Agence Nationale de la Recherche (ANR-20-CE37-0016 to G.D. and ANR-19-CE37-0007 to D.P.) and the Fondation pour la Recherche Medicale (FRM-EQU202103012770 to C.L.).

Competing interests

The authors declare no competing interests.

Ethical approval

This project has undergone review and approval by the ethical review board and the department responsible for overseeing the use of animals for scientific purposes within the French Ministry of Higher Education, Research, and Innovation. Approval was granted under reference APAFIS #29793-202102121752192 v3.

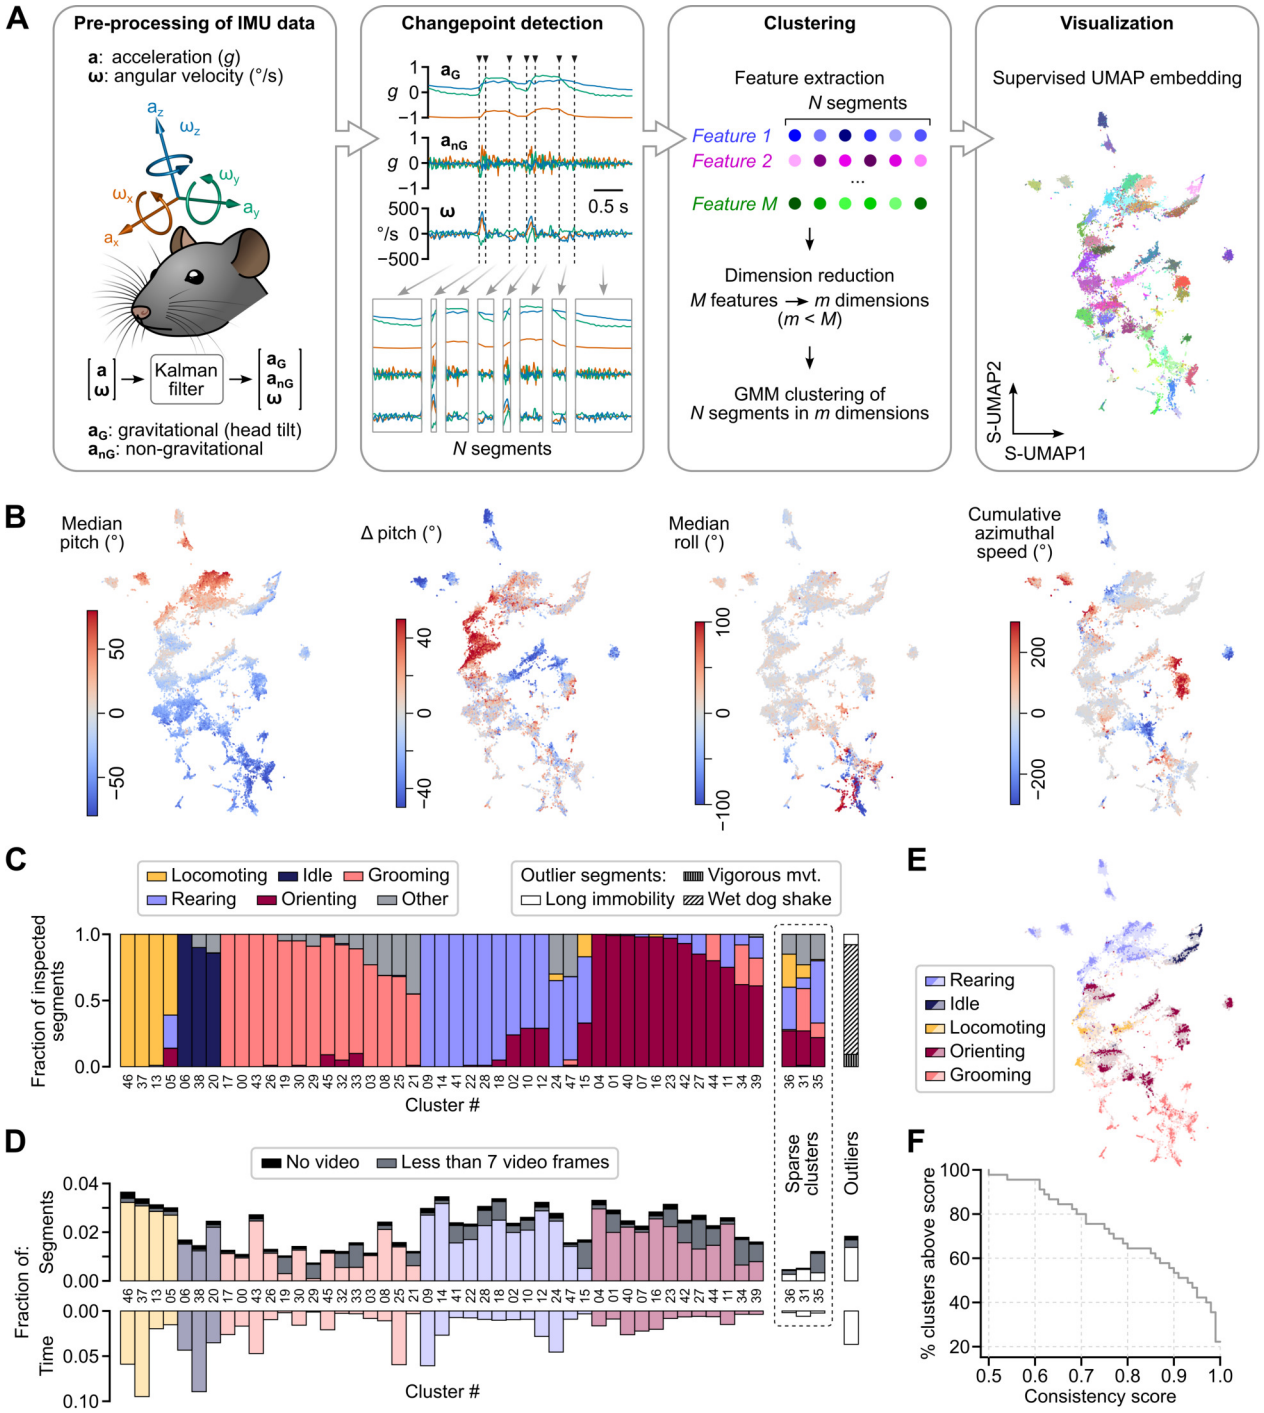

Figure 1: Decomposition of rat behavior using DISSeCT. A, Pipeline architecture. **B**, Example supervised UMAP segment embeddings with color-coded feature values. **C**, Stacked bar plot depicting the assignment of main behavioral categories to inspected segments within each cluster. Cluster are grouped by their dominant main category and sorted by decreasing consistency score. **D**, Top: fraction of the total number of segments represented by each cluster. For each cluster, the proportion of segments whose corresponding video snippet was either lacking or comprised less than 7 frames is highlighted in black and grey, respectively. Bottom: fraction of the total recording duration represented by each cluster (all segments included). **E**, Supervised UMAP embedding showing segments color-coded according to their cluster's dominant main category (dark colors: inspected segments; light colors: other segments). **F**, Fraction of clusters whose consistency score exceeds a certain value, plotted as a function of this value.

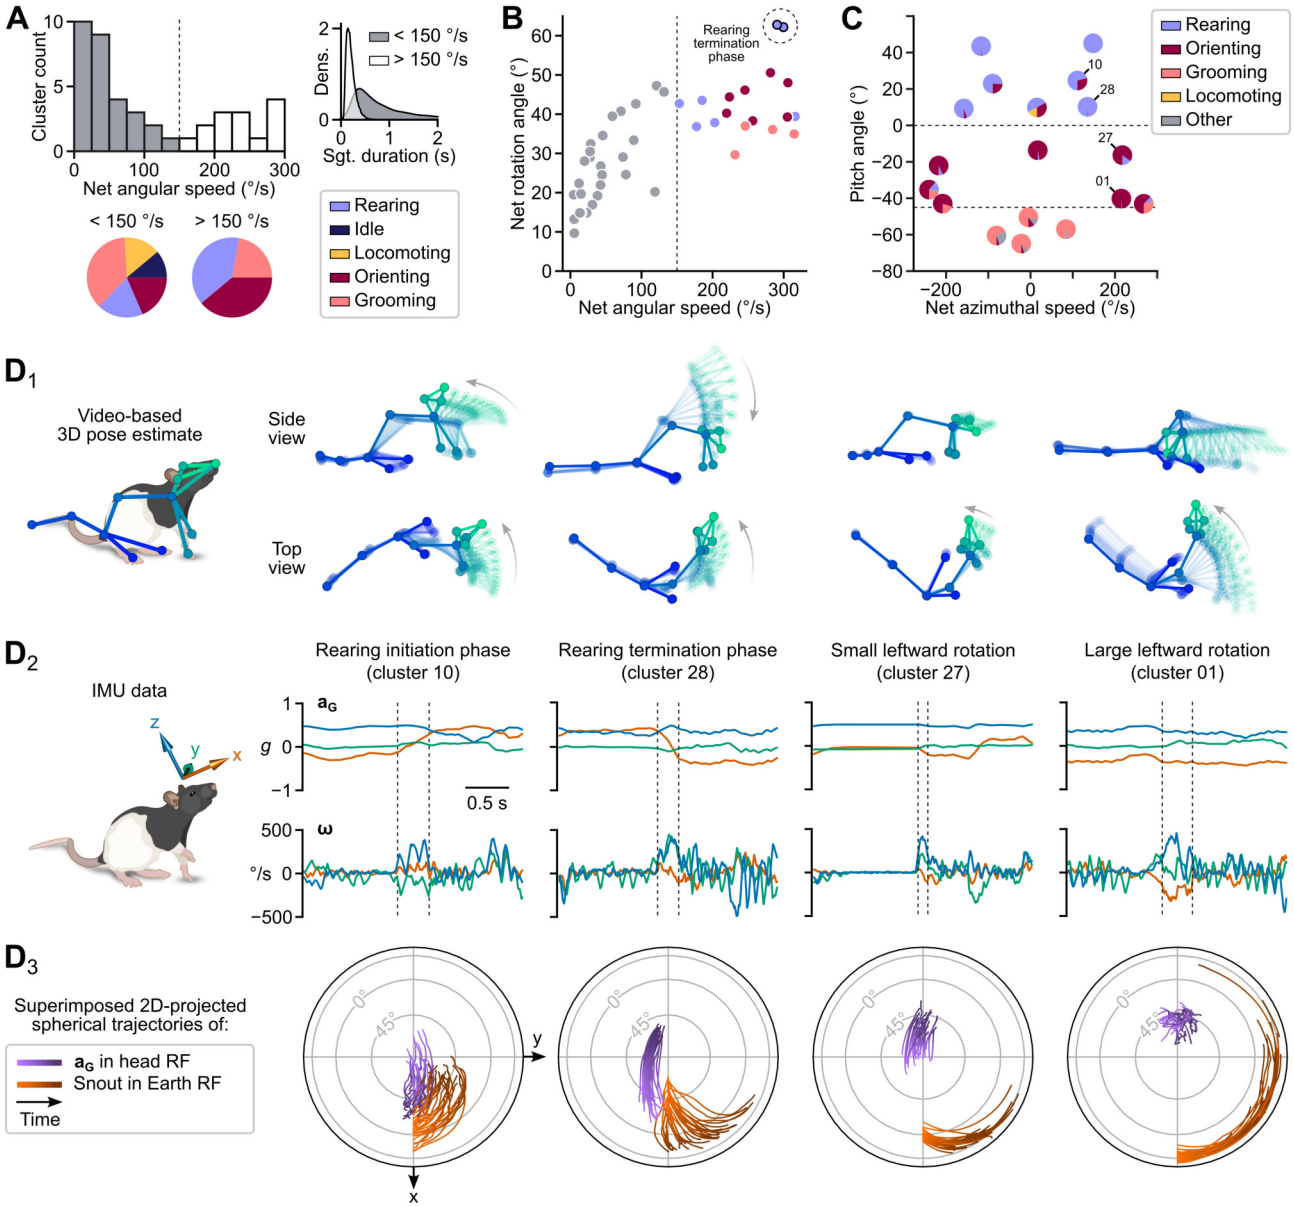

Figure 2: Diversity and fine dynamics of rapid changes in head orientation. **A**, Distribution of net angular speed values (median of per-segment average values) for all clusters. The dashed vertical line at 150 °/s corresponds to the threshold for selecting clusters representing rapid head rotations (highlighted in white in the histogram). Pie charts illustrate the proportion of cluster-associated behavioral categories (one unique category per cluster) for each side of the distribution. The upper right inset shows a density estimate of segment duration for each side of the distribution. **B**, Net rotation angle plotted against net angular speed (median of per-segment average values) for each cluster. Clusters are colored according to their associated behavioral category (using the same color code as in A and Fig. 1C). **C**, Head pitch angle (see Methods) plotted as a function of net azimuthal speed (median of per-segment average values) for each cluster. Each cluster is displayed as a pie chart showing the behavioral categories attributed to the top 100 segments of the cluster (cf. Fig. 1C). **D_{1-3}**, Detailed examination of four example clusters (identified in C). Each column corresponds to one cluster. D_1: Skeletons representing the animal's 3D pose dynamics during one example segment. Zero transparency corresponds to the last video frame of the segment. D_2: Gravitational acceleration and angular velocity signals ($\mathbf{a_G}$ and $\boldsymbol{\omega}$, respectively) for the same segment. Dashed lines represent segment boundaries. D_3: Joint visualization of the trajectories of $\mathbf{a_G}$ in the head reference frame (purple) and of the animal's heading in the Earth reference frame (orange) for the top 30 segments of each cluster (see Supplementary Methods; RF: reference frame).

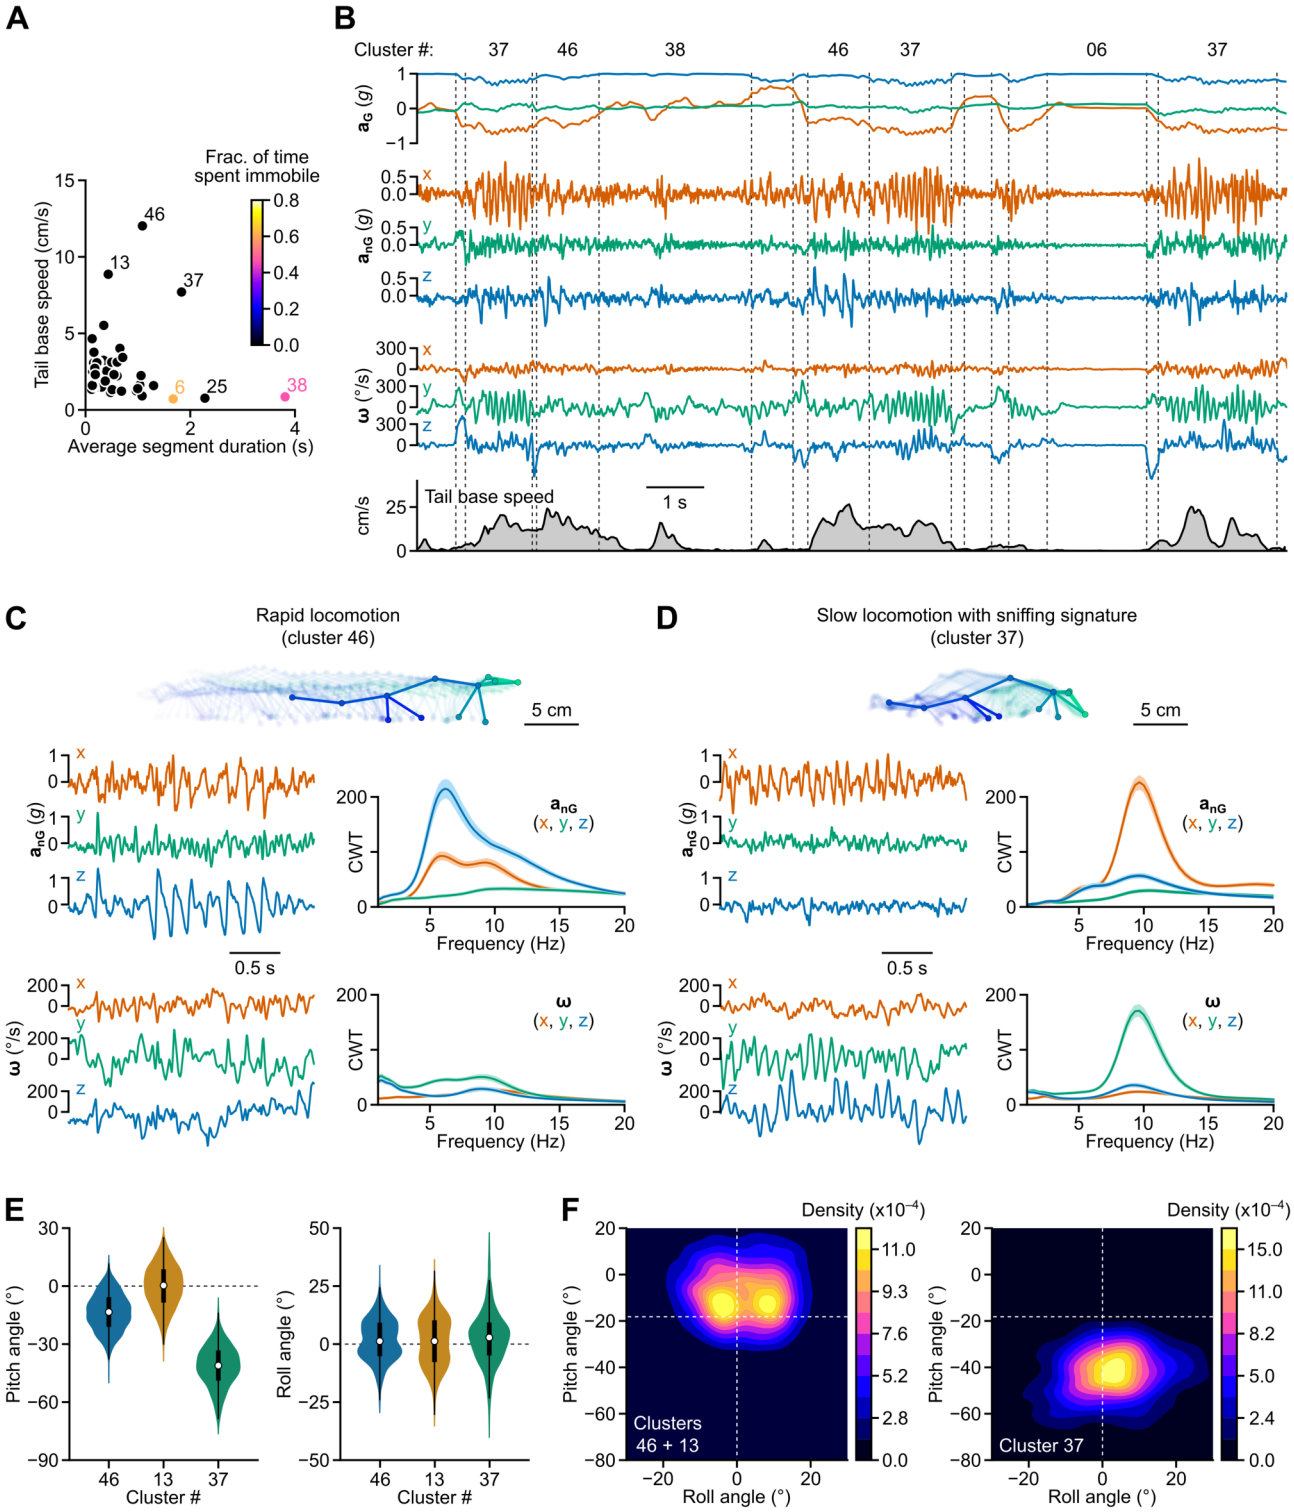

Figure 3: Identification of two locomotor regimes (caption on next page).

Figure 3: Identification of two locomotor regimes. A, Tail base speed (median of per-segment average values) plotted against average segment duration for each cluster. The color scale represents the median fraction of time spent immobile, computed for each segment based on head angular speed (see Methods). Cluster numbers are indicated for three "locomotion" (13, 37 and 46) and two "idle" clusters (6 and 38). **B**, Upper panels: Example traces of head gravitational and non-gravitational acceleration and angular velocity ($\mathbf{a_G}$, $\mathbf{a_{nG}}$ and $\boldsymbol{\omega}$, respectively). Lower panel: Concurrent horizontal tail base speed computed from 3D pose estimates. Dashed vertical lines indicate segment boundaries. Corresponding cluster numbers are shown above the traces. **C**, Top: Skeletons representing the animal's 3D pose dynamics during 30 video frames (1 s) extracted from the segment with the highest posterior probability of belonging to cluster 46 (associated with rapid locomotion). Zero transparency corresponds to the last frame. Left: Example traces of head non-gravitational acceleration and angular velocity ($\mathbf{a_{nG}}$ and $\boldsymbol{\omega}$, respectively) for the same segment. Right: Continuous wavelet transform (CWT) coefficients for all three axes of $\mathbf{a_{nG}}$ and $\boldsymbol{\omega}$. Solid lines represent the average of per-segment median log-magnitude values (shared area: 95% confidence interval). **D**, Similar to panel **C**, but for cluster 37 (corresponding to mobile olfactory exploration). **E**, Distribution of per-segment average head pitch and roll angles for the three "locomotion" clusters identified in **A**. **F**, 2D histograms showing the distribution of per-segment average head pitch and roll angles for clusters associated with rapid locomotion (left) and mobile olfactory exploration (right). Dashed lines indicate the resting pitch and roll angles, computed as the median of per-segment average pitch and roll angles for all segments.

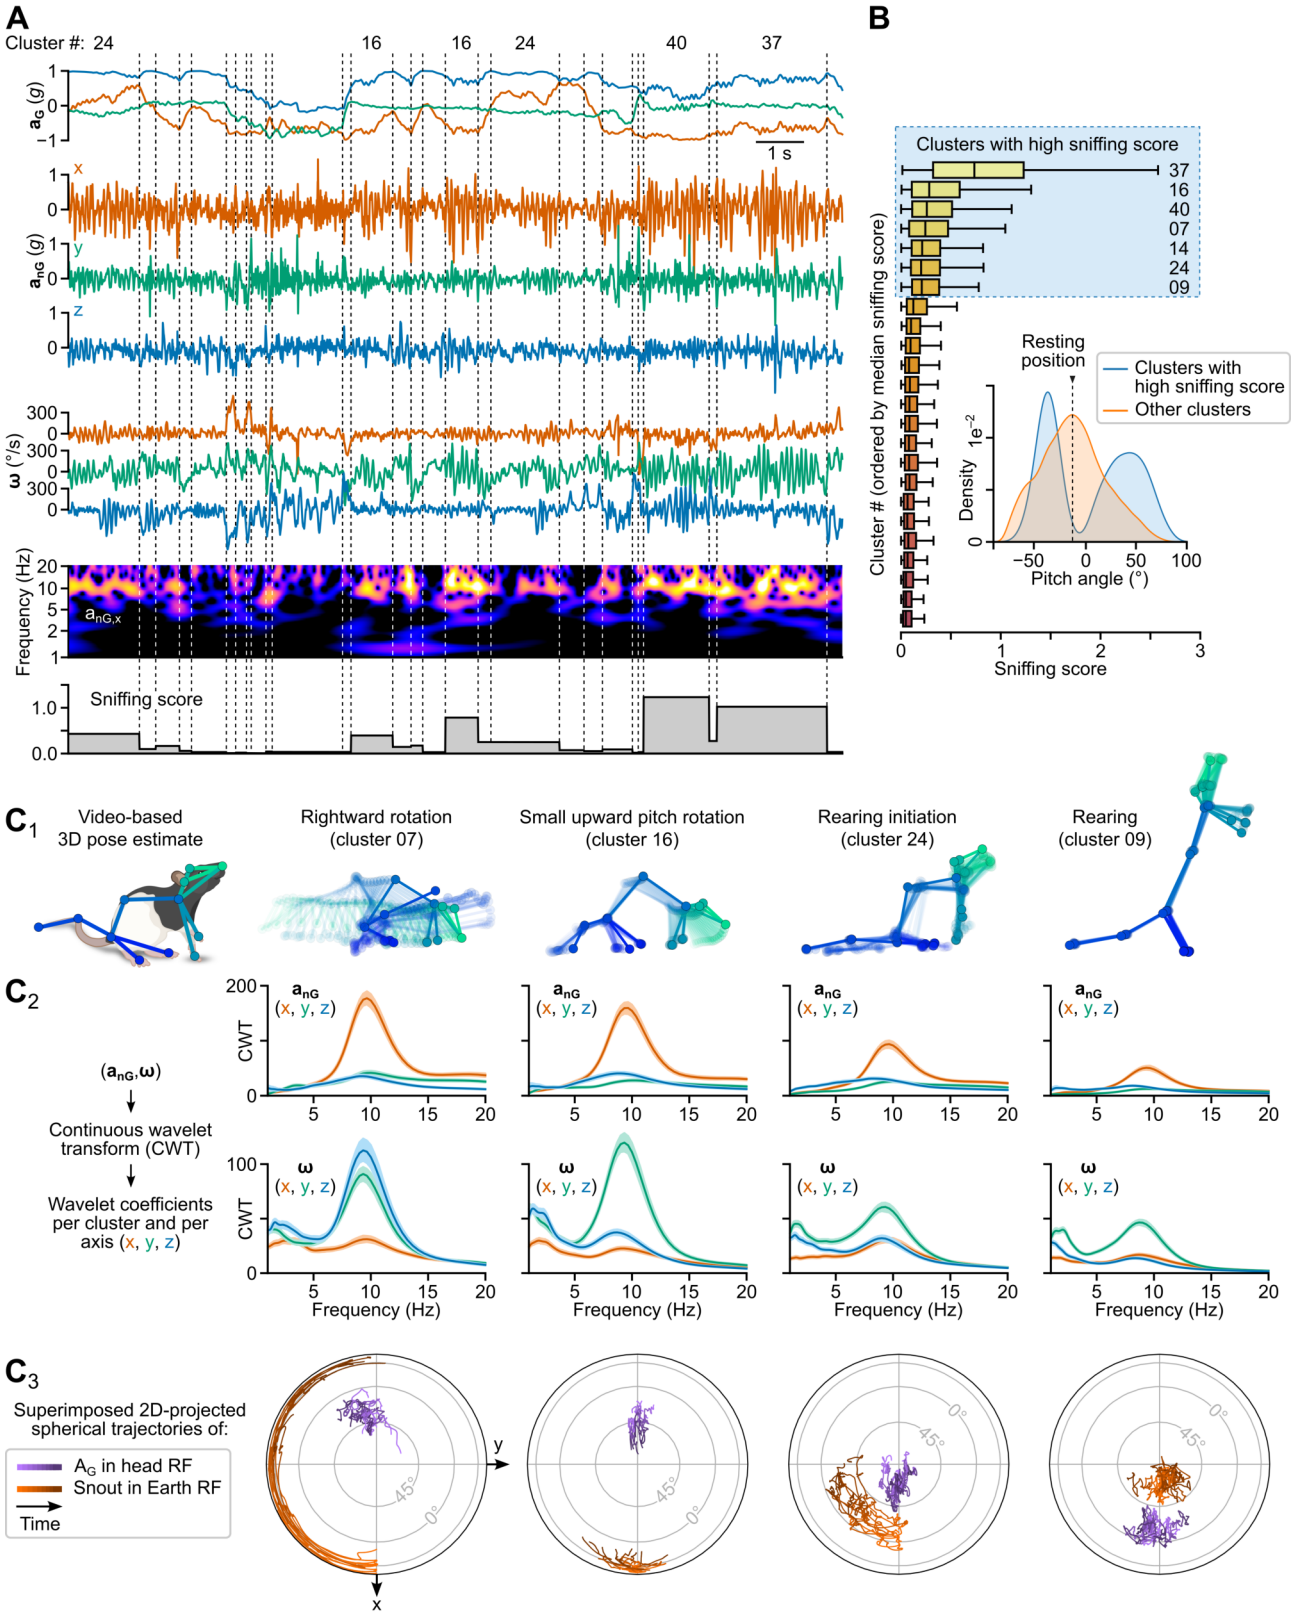

Figure 4: Identification of a sniffing signature (caption on next page).

Figure 4: Identification of a sniffing signature. B, Upper panels: Example traces of head gravitational and non-gravitational acceleration ($\mathbf{a_G}$ and $\mathbf{a_{nG}}$, respectively), and angular velocity ($\boldsymbol{\omega}$). Lower panels: Log-magnitude of continuous wavelet transform coefficients computed for fore-aft non-gravitational acceleration (x-axis component of $\mathbf{a_G}$), and per-segment sniffing score values. Dashed vertical lines indicate segment boundaries. Corresponding cluster numbers are indicated above the traces. **B**, Box plot of sniffing score values shown for 24 clusters. Clusters are sorted by decreasing median sniffing score. Inset: Distribution of per-segment median head pitch angles computed for the seven clusters with highest sniffing scores (blue curve) and for all other clusters (orange curve). $\mathbf{C_{1\text{–}3}}$, Detailed examination of four example clusters with high sniffing scores. Each column corresponds to one cluster. C_1: Skeletons representing the animal's 3D pose dynamics during one example segment. Zero transparency corresponds to the last video frame of the segment. C_2: Log-magnitude of continuous wavelet transform (CWT) coefficients (average of per-segment median values) for all three axes of $\mathbf{a_{nG}}$ and $\boldsymbol{\omega}$, shown for each example cluster (shaded area: 95% confidence interval). C_3: Joint visualization of the trajectories of $\mathbf{a_G}$ in the head reference frame (purple) and of the animal's heading (orange) for the top 10 segments of each cluster (see Supplementary Methods; RF: reference frame).

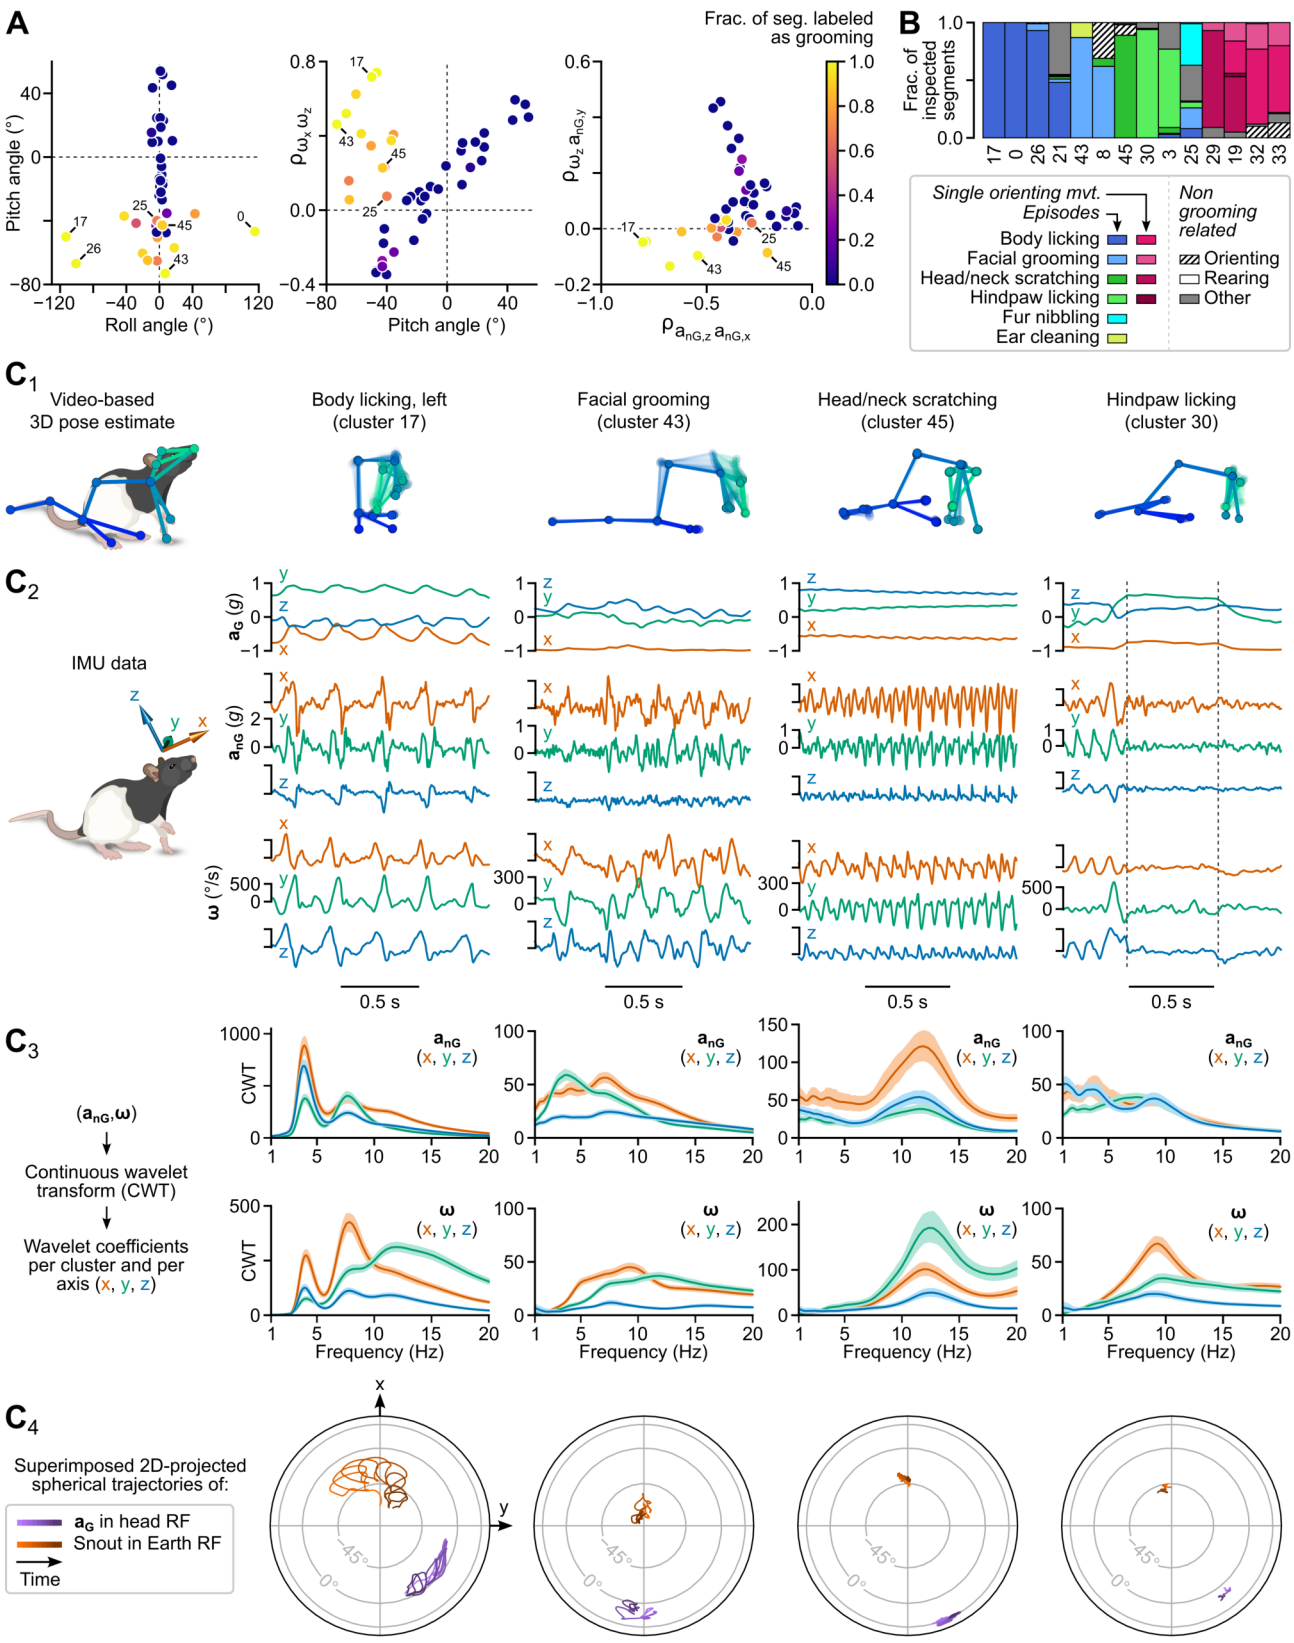

Figure 5: Parsing of grooming components (caption on next page).

Figure 5: Parsing of grooming components. A, Distinctive features of grooming-related clusters. Left: head pitch vs roll angles. Middle: correlation between angular velocities about the x and z axes (ω_x and ω_z, respectively) plotted against head pitch angle. Right: correlation between angular velocity about the z axis (ω_z) and medio-lateral non-gravitational acceleration ($a_{nG,y}$) plotted against the correlation between the head-vertical and fore-aft components of non-gravitational acceleration ($a_{nG,z}$ and $a_{nG,x}$, respectively). These panels show per-cluster median values of per-segment average values. The color scale indicates the fraction of segments identified as belonging to a grooming sequence during inspection of the top 100 segments of each cluster (see Method). **B**, Stacked bar plot depicting the assignment of labels to inspected segments within each grooming-related cluster. $\mathbf{C_{1\text{–}3}}$, Detailed examination of four example clusters. Each column corresponds to one cluster. C_1: Skeletons representing the animal's 3D pose dynamics during one example segment. Zero transparency corresponds to the last video frame of the segment. C_2: Gravitational and non-gravitational acceleration ($\mathbf{a_G}$ and $\mathbf{a_{nG}}$, respectively), and angular velocity ($\boldsymbol{\omega}$) for the same segment. Average segment-wise median log-magnitude of continuous wavelet transform (CWT) coefficients for all three axes of $\mathbf{a_{nG}}$ and $\boldsymbol{\omega}$ (shared area: 95 % confidence interval).

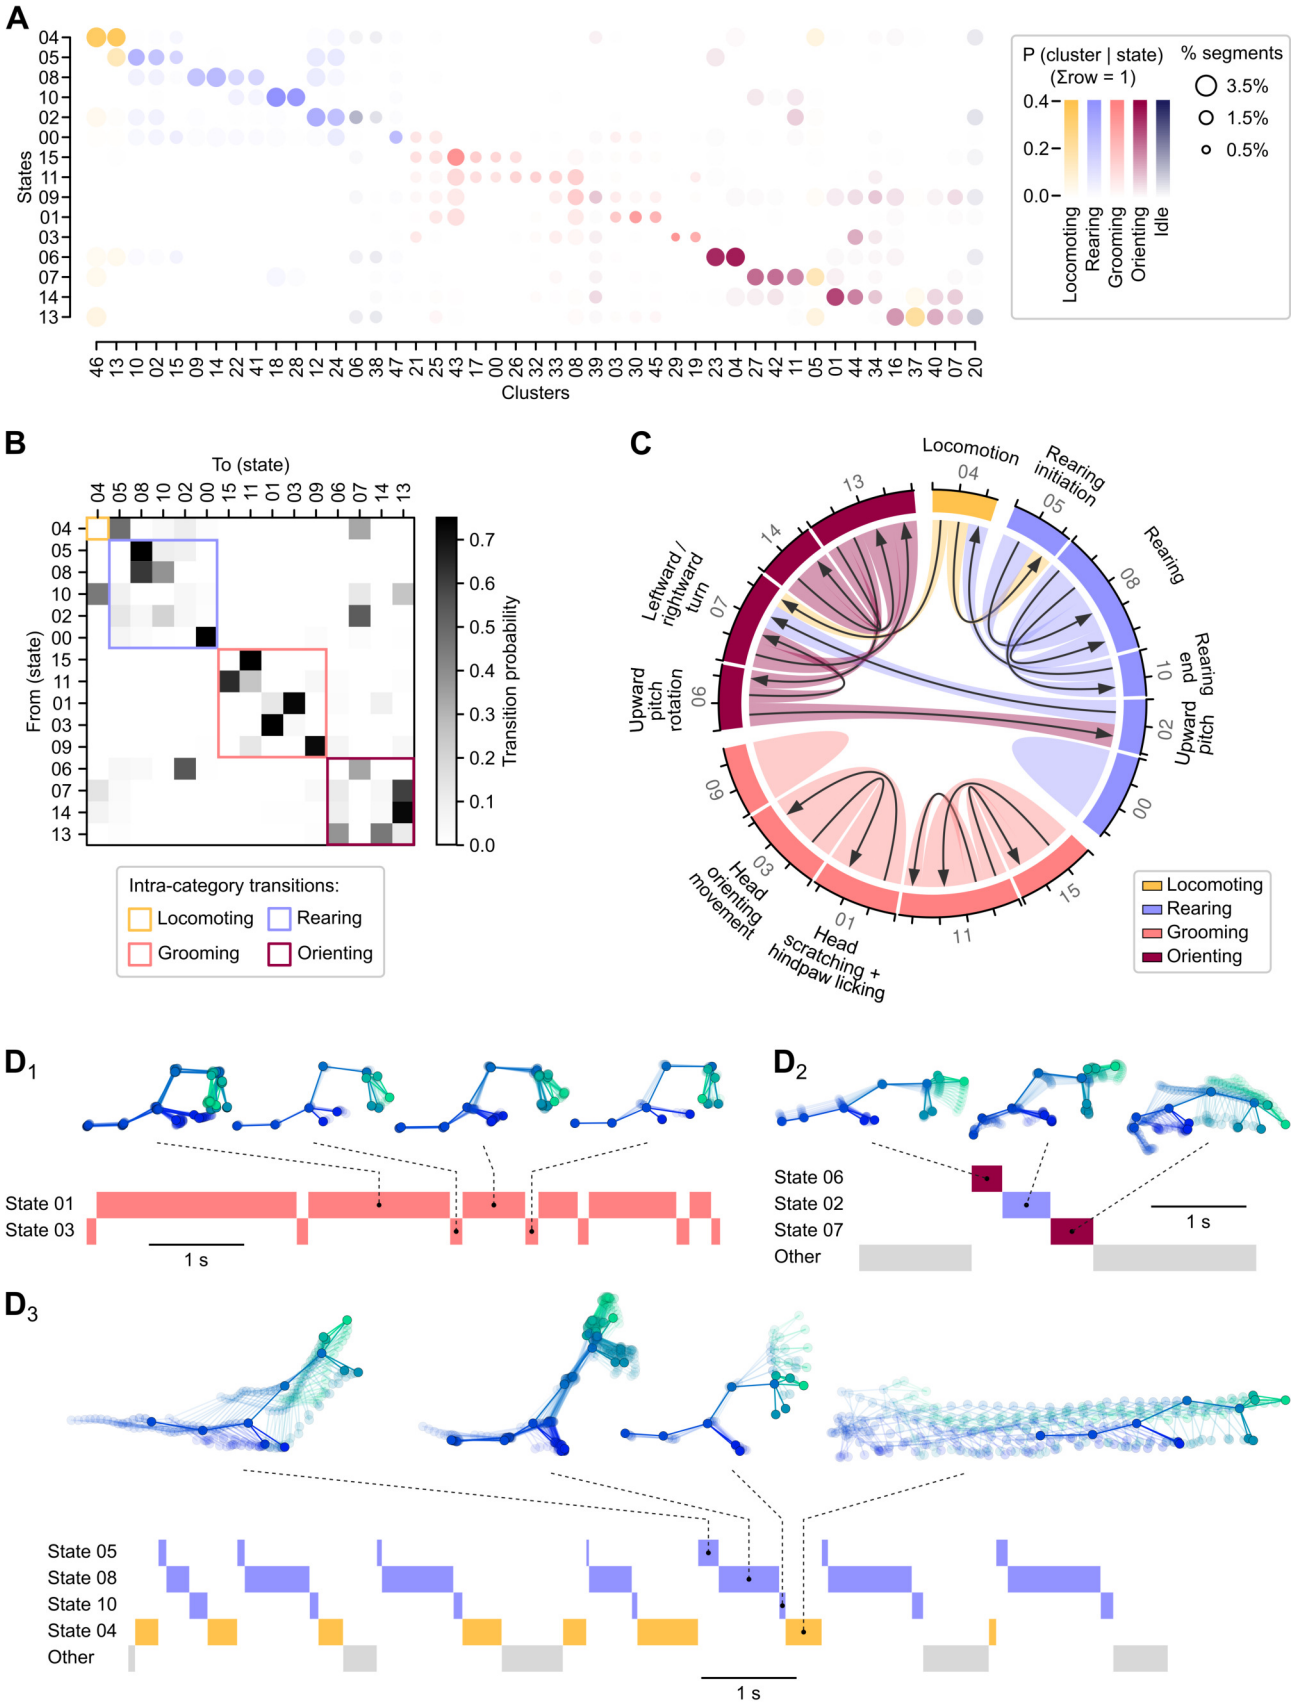

Figure 6: Identification of behavioral sequences using a post-hoc categorical hidden Markov model (caption on next page).

Figure 6: Identification of behavioral sequences using a post-hoc categorical hidden Markov model. **A**, Cluster emission probabilities (columns) for each hidden state (rows). The dots in each column are colored according to the cluster's associated main behavioral category as defined in Fig. 1C, with the opacity representing probability values and dot area representing how much the cluster contributed to the total number of segments. Note that state 12, which contains outlier segments, and sparse clusters (see Methods and Fig. S7) are not represented in this plot. **B**, Hidden state transition matrix. States are ordered according to their dominant main behavioral category. Colored squares highlight intra-category transitions. **C**, Chord plot representing transition probabilities between states. Arrow width represents transition probability (tick spacing: 0.5). States are colored according to their dominant main behavioral category. Only transitions probabilities greater than 0.25 are represented. $\mathbf{D_{1\text{–}3}}$, Representative ethograms for three types of sequences. D_1: Alternation between head scratching or hindpaw licking episodes (state 1) and brief head orienting movements (state 3). D_2: Single episode during which the animal is lifting its head up before executing a turn. D_3: Alternation between rearing episodes (states 5, 8 and 10) and locomotion bouts (state 4). For each ethogram, skeletons representing the animal's 3D pose dynamics are shown for example segments. Zero transparency corresponds to the last video frame of the segment.

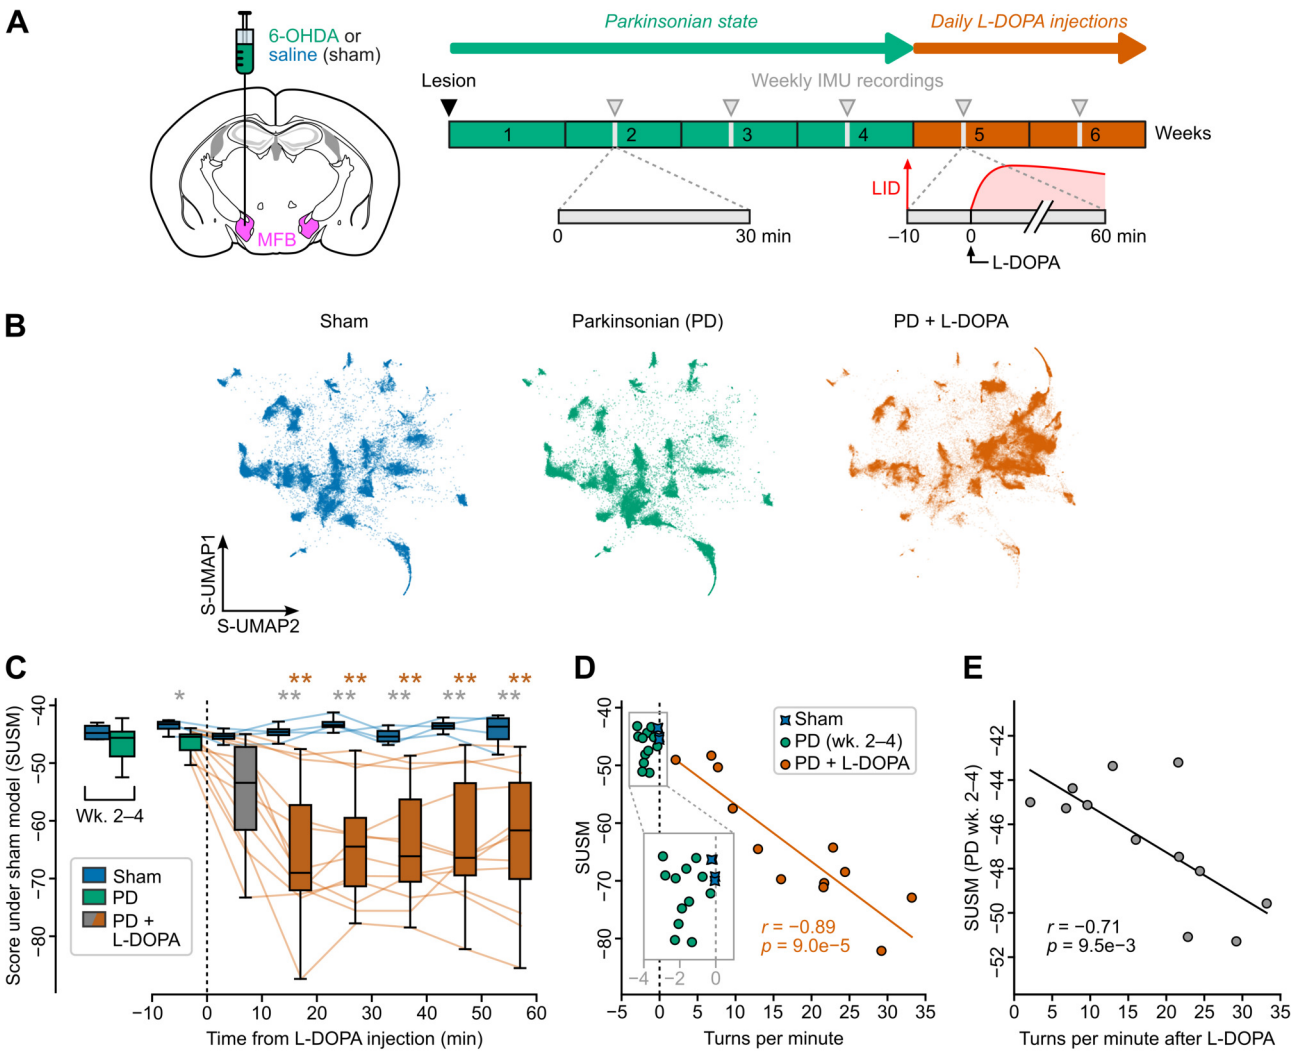

Figure 7: Macroscopic analysis of motor impairments in a mouse model of Parkinson's disease and levodopa-induced dyskinesia. A, Experimental protocol. Left: drawing of a coronal brain slice (1.2 mm posterior to Bregma, adapted from[104]) showing the unilateral injection of 6-OHDA (PD mice, $n = 12$) or saline (sham mice, $n = 4$) into the medial forebrain bundle (MFB, purple). Right: protocol timeline (see Methods for details). The six-week protocol began with 6-OHDA (PD mice) or saline (sham mice) injection and IMU implantation. Weekly IMU recordings (gray arrowheads) were performed during weeks 2 to 4. Daily levodopa (L-DOPA) injections started in week 5, with IMU recording sessions conducted during weeks 5 and 6. Recordings began 10 min before injection. Each L-DOPA injection triggered an episode of abnormal involuntary movements (LID: L-DOPA induced dyskinesia, symbolized by a red curve). **B**, S-UMAP embeddings for sham mice (left, blue) and for PD mice before (middle, green) and during (right, orange) L-DOPA treatment, using a clustering model fitted on sham mice data. **C**, Left: box plots of the score under sham model (SUSM) for sham ($n = 4$) vs PD ($n = 12$) mice before L-DOPA treatment (weeks 2 to 4). Right: box plots of the SUSM computed for the same mice over 10-min bins right before and after L-DOPA injection (weeks 5 and 6). Lines show the timecourse of the SUSM for individual PD (orange) and sham (blue) mice. Gray asterisks denote significant differences between PD vs sham mice SUSM values for the same time bins (unpaired pairwise t-test with Bonferroni correction, **: p-value < 0.01). Orange asterisks denote significant differences between PD mice SUSM values before vs after L-DOPA injection (paired pairwise t-test with Bonferroni correction, **: p-value < 0.01). **D**, Scatter plot of the SUSM vs turns per minutes (estimated from IMU recordings, see Supplementary Methods) for each animal. Along the x axis (turns per minute), negative (resp. positive) values indicate ipsiversive (resp. contraversive) rotations relative to the lesion side. Each PD mouse ($n = 12$) is shown twice: during the Parkinsonian phase (green dots) and during the LID phase (10–60 min period following L-DOPA injection, orange dots). The linear regression was computed for the LID phase. **E**, Comparison of the SUSM before L-DOPA treatment (weeks 2 to 4) and the turns per minute during the LID phase (10–60 minutes following L-DOPA injection) for PD mice ($n = 12$), showing that the SUSM during the Parkinsonian phase is predictive of the amount of contraversive rotations induced by L-DOPA.

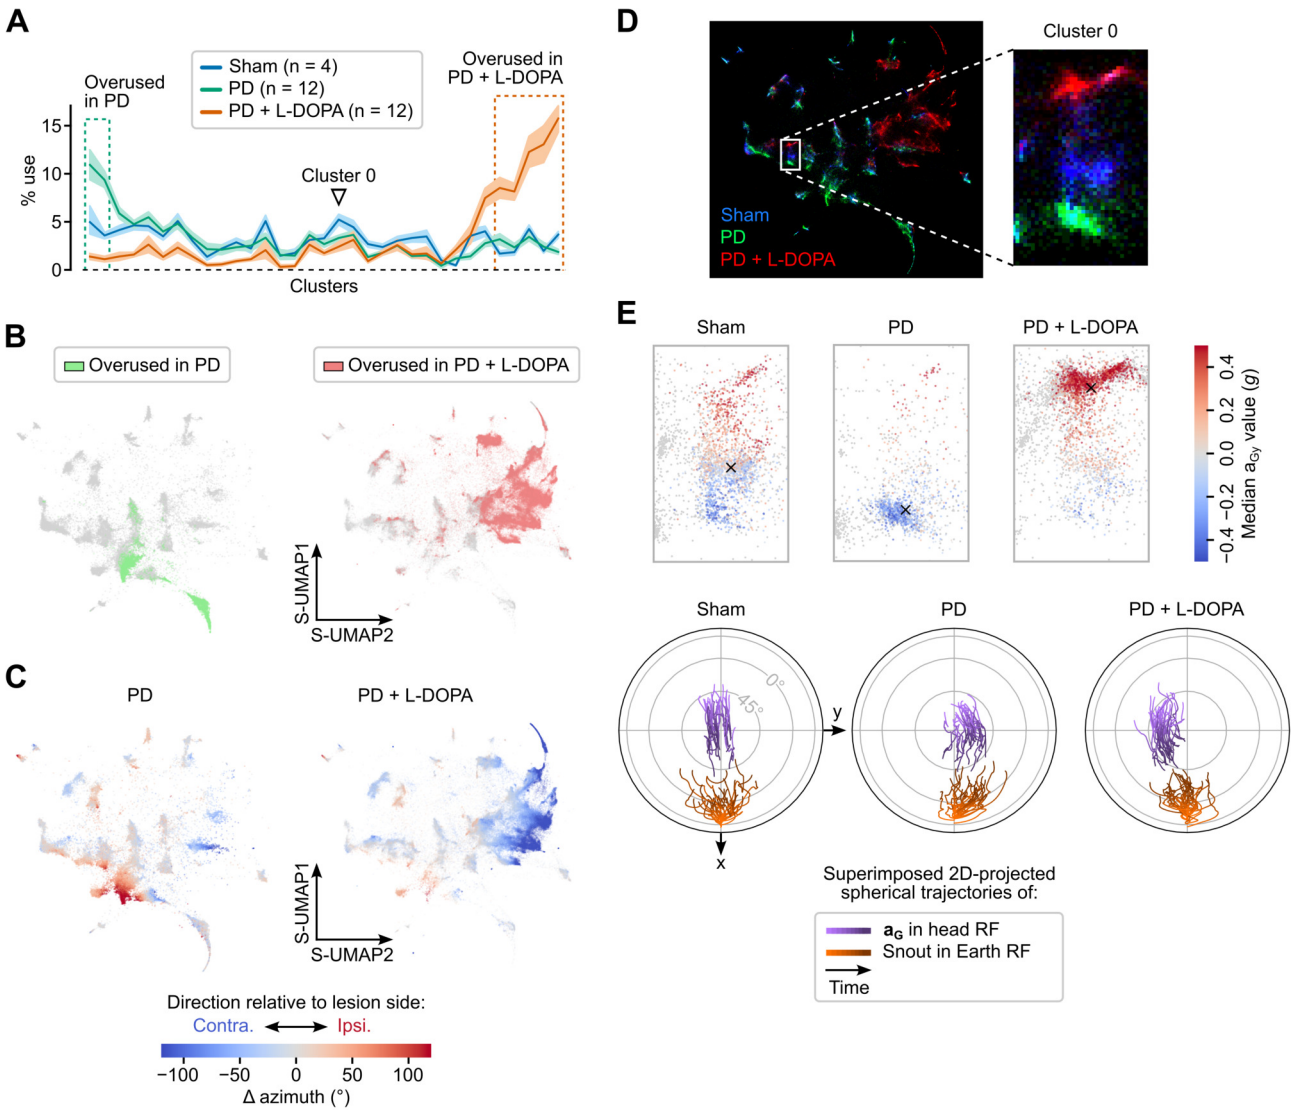

Figure 8: DISSeCT captures both coarse and fine-grained changes in a mouse model of Parkinson's disease and levodopa-induced dyskinesia. **A**, Per-mouse average percentage use of individual clusters in sham (blue curve) vs PD animals before (green curve) and after (orange curve) L-DOPA administration. Shaded areas represent 95% confidence intervals. Dashed rectangles indicate clusters exhibiting an increased use of more than 5 % compared to sham mice (overused clusters) for PD mice before (green rectangle) and after (orange rectangle) L-DOPA injection. **B**, S-UMAP embeddings highlighting segments belonging to overused clusters (defined in panel A) in PD mice before (left, green) and after (right, orange) L-DOPA injection. Gray points represent segments belonging to other clusters. **C**, S-UMAP embeddings showing segments color coded according to the change of head azimuth. Positive (resp. negative) values represent ipsiversive (resp. contraversive) rotations relative to the lesion side. **D**, Right: Superimposed S-UMAP embeddings for sham (blue), PD (green) and PD under L-DOPA (red) mice. Left: close-up view of cluster 0 (indicated by an arrowhead in panel A), which corresponds to upward head movements. **E**, Left: close-up view of cluster 0 in the S-UMAP embedding space, with segments color coded according to the median value of $\mathbf{a_G}$ along the interaural axis. Negative (resp. positive) values indicate a roll towards (resp. away from) the lesion side. Right: corresponding joint visualization of the trajectories of $\mathbf{a_G}$ in the head reference frame (purple) and of the animal's heading in the Earth reference frame (orange) for the top 30 segments of cluster 0.