Automated Behavioral Analysis of Asleep Fruit Fly

Ali Osman Berk Şapcı
Advisor: Oznur Tastan[1]
Co-advisor: Sündüz Keleş[2,3]
1 : Faculty of Engineering and Natural Sciences, Sabancı University, Turkey;
2 : Department of Statistics, University of Wisconsin, Madison, WI, USA;
3 : Department of Biostatistics and Medical Informatics, University of Wisconsin, Madison, WI, USA.

Table of contents

Behavior

Definition

Behaviour is the internally coordinated responses (actions or inactions) of whole living organisms to internal and/or external stimuli.

Why biologists are interested in behavior?

Because it is one of the most robust outputs of the brain.

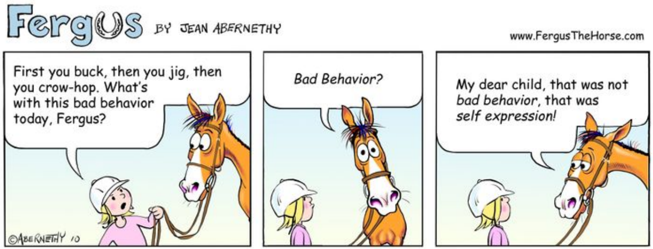

Phenotyping behavior to decipher sleep

Motivation

Despite its importance and ubiquity, the function of sleep remains unknown.

Phenotyping behavior to decipher sleep

Motivation

Despite its importance and ubiquity, the function of sleep remains unknown.

Goal

Quantification of fruit flies' behavioral repertoire of sleep.

The ultimate goal is to discover behaviorally defined states of sleep, similar to REM sleep.

NEUROPHYSIOLOGY

A deep sleep stage in *Drosophila* with a functional role in waste clearance

Bart van Alphen[1*], Evan R. Semenza[1†], Melvyn Yap[2], Bruno van Swinderen[2], Ravi Allada[1*]

Figure: van Alphen et al. (2021)

A call to action: *Computational Ethology*

- Emerging field of computational ethology.
- The aim is to automate the quantification of animal behavior.
- It is made possible by recent technical advances, especially in machine learning.

Toward a Science of Computational Ethology

David J. Anderson[1,2,4,*] and Pietro Perona[3,4,*]

(a) (Anderson and Perona, 2014)

Computational Neuroethology: A Call to Action

Sandeep Robert Datta,[5,9,*] David J. Anderson,[1,2,3] Kristin Branson,[4] Pietro Perona,[8] and Andrew Leifer[6,7,*]

(b) Datta et al. (2019)

Quantifying behavior to understand the brain

Talmo D. Pereira[1], Joshua W. Shaevitz[2,3] and Mala Murthy[1]✉

(c) Pereira (2020)

Quantifying the dynamics of behavior

- Deep learning have enabled pose estimation and tracking of animals.

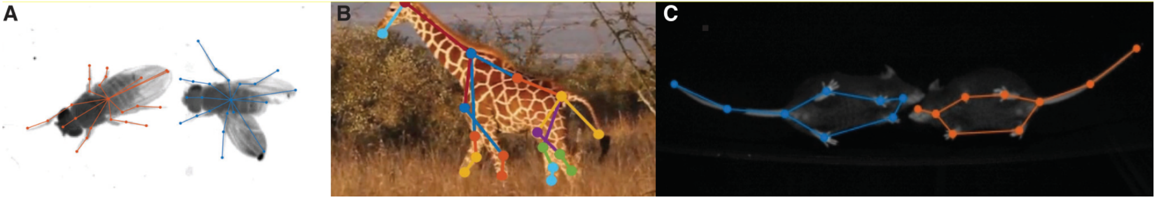

Figure: Fly, giraffe and mouse tracking examples shown from Pereira et al. (2022).

Quantifying the dynamics of behavior

- Deep learning have enabled pose estimation and tracking of animals.
- However, a set of coordinate values are not sufficient to describe behaviors.

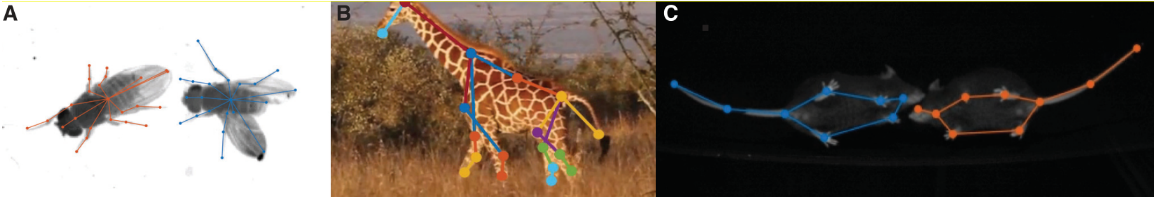

Figure: Fly, giraffe and mouse tracking examples shown from Pereira et al. (2022).

Quantifying the dynamics of behavior

- Deep learning have enabled pose estimation and tracking of animals.
- However, a set of coordinate values are not sufficient to describe behaviors.
- Behavior is complex!
  - There are thousands of unique postures.
  - Behaviors occur at different time scales.
  - There are no rigid boundaries between behavioral categories.

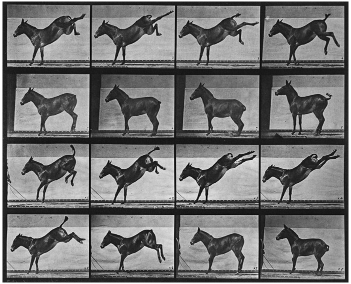

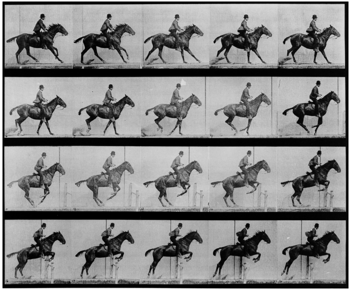

Figure: Muybridge's motion studies, 1887.

Problem: Mapping behavioral repertoire of dormant fruit fly

Problem

Mapping video recordings of sleep experiments to the defined behavioral categories.

Problem: Mapping behavioral repertoire of dormant fruit fly

Problem

Mapping video recordings of sleep experiments to the defined behavioral categories.

- Our input data is a set of body part coordinates, i.e., output of pose estimation tool.

Problem: Mapping behavioral repertoire of dormant fruit fly

Problem

Mapping video recordings of sleep experiments to the defined behavioral categories.

- Our input data is a set of body part coordinates, i.e., output of pose estimation tool.
- 6 annotated categories (5 behaviors and one quiescent & other).

Problem: Mapping behavioral repertoire of dormant fruit fly

Problem

Mapping video recordings of sleep experiments to the defined behavioral categories.

- Our input data is a set of body part coordinates, i.e., output of pose estimation tool.
- 6 annotated categories (5 behaviors and one quiescent & other).

Problem: Mapping behavioral repertoire of dormant fruit fly

Problem

Mapping video recordings of sleep experiments to the defined behavioral categories.

- Our input data is a set of body part coordinates, i.e., output of pose estimation tool.
- 6 annotated categories (5 behaviors and one quiescent & other).

Challenges

- Behaviors occur sparsely and rarely during long sleep cycles.

Problem: Mapping behavioral repertoire of dormant fruit fly

Problem

Mapping video recordings of sleep experiments to the defined behavioral categories.

- Our input data is a set of body part coordinates, i.e., output of pose estimation tool.
- 6 annotated categories (5 behaviors and one quiescent & other).

Challenges

- Behaviors occur sparsely and rarely during long sleep cycles.
- Behavioral repertoire of sleep is defined by unobtrusive changes.

Problem: Mapping behavioral repertoire of dormant fruit fly

Problem

Mapping video recordings of sleep experiments to the defined behavioral categories.

- Our input data is a set of body part coordinates, i.e., output of pose estimation tool.
- 6 annotated categories (5 behaviors and one quiescent & other).

Challenges

- Behaviors occur sparsely and rarely during long sleep cycles.
- Behavioral repertoire of sleep is defined by unobtrusive changes.
- Pose estimation data is highly noisy.

Problem: Mapping behavioral repertoire of dormant fruit fly

Problem

Mapping video recordings of sleep experiments to the defined behavioral categories.

- Our input data is a set of body part coordinates, i.e., output of pose estimation tool.
- 6 annotated categories (5 behaviors and one quiescent & other).

Challenges

- Behaviors occur sparsely and rarely during long sleep cycles.
- Behavioral repertoire of sleep is defined by unobtrusive changes.
- Pose estimation data is highly noisy.
- Large data: more than one and a half million frames for a single experiment.

Problem: Mapping behavioral repertoire of dormant fruit fly

Problem

Mapping video recordings of sleep experiments to the defined behavioral categories.

- Our input data is a set of body part coordinates, i.e., output of pose estimation tool.
- 6 annotated categories (5 behaviors and one quiescent & other).

Challenges

- Behaviors occur sparsely and rarely during long sleep cycles.
- Behavioral repertoire of sleep is defined by unobtrusive changes.
- Pose estimation data is highly noisy.
- Large data: more than one and a half million frames for a single experiment.
- Limited supervision: number of annotated experiments is not many.

Existing Methods: Why do we need a different approach?

- Because existing methods are not designed for dormant animals!

Existing Methods: Why do we need a different approach?

- Because existing methods are not designed for dormant animals!
- Focus has been on behaviors that are defined by major postural or positional changes.
  - MotionMapper (Berman et al., 2014)
  - B-SOID (Hsu and Yttri, 2021)

Existing Methods: Why do we need a different approach?

- Because existing methods are not designed for dormant animals!
- Focus has been on behaviors that are defined by major postural or positional changes.
  - MotionMapper (Berman et al., 2014)
  - B-SOID (Hsu and Yttri, 2021)
- A few studies have developed ad-hoc solutions for specific behaviors such as grooming and feeding.
  - Qiao et al. (2018)
  - Itskov et al. (2014)

Existing Methods: Why do we need a different approach?

- Because existing methods are not designed for dormant animals!
- Focus has been on behaviors that are defined by major postural or positional changes.
  - MotionMapper (Berman et al., 2014)
  - B-SOID (Hsu and Yttri, 2021)
- A few studies have developed ad-hoc solutions for specific behaviors such as grooming and feeding.
  - Qiao et al. (2018)
  - Itskov et al. (2014)

Existing Methods: Why do we need a different approach?

- Because existing methods are not designed for dormant animals!
- Focus has been on behaviors that are defined by major postural or positional changes.
  - MotionMapper (Berman et al., 2014)
  - B-SOID (Hsu and Yttri, 2021)
- A few studies have developed ad-hoc solutions for specific behaviors such as grooming and feeding.
  - Qiao et al. (2018)
  - Itskov et al. (2014)

Detailed phenotyping of behavioral repertoire of sleep has not been addressed before.

How do we approach the problem?

- We develop a multi-stage and end-to-end pipeline: `basty`.
  - Automated Behavioral Analysis of Asleep Fruit Fly
  - Configurable, open source, and easy-to-use software package.

How do we approach the problem?

- We develop a multi-stage and end-to-end pipeline: `basty`.
  - Automated Behavioral Analysis of Asleep Fruit Fly
  - Configurable, open source, and easy-to-use software package.
- `basty` focuses on subtle movements exhibited during sleep.

How do we approach the problem?

- We develop a multi-stage and end-to-end pipeline: `basty`.
  - Automated Behavioral Analysis of Asleep Fruit Fly
  - Configurable, open source, and easy-to-use software package.
- `basty` focuses on subtle movements exhibited during sleep.
- Our approach consists of:
  - Computing meaningful spatio-temporal representations,
  - Detecting activities in long sleep experiments,
  - Mapping activities to defined behavioral categories.

How do we approach the problem?

- We develop a multi-stage and end-to-end pipeline: `basty`.
  - Automated Behavioral Analysis of Asleep Fruit Fly
  - Configurable, open source, and easy-to-use software package.
- `basty` focuses on subtle movements exhibited during sleep.
- Our approach consists of:
  - Computing meaningful spatio-temporal representations,
  - Detecting activities in long sleep experiments,
  - Mapping activities to defined behavioral categories.
- We evaluate our pipeline with a dataset of 11 sleep experiments.
  - Sleep experiment: 14 hours of video recordings of flies during the night.

Conducting sleep experiments and collecting data

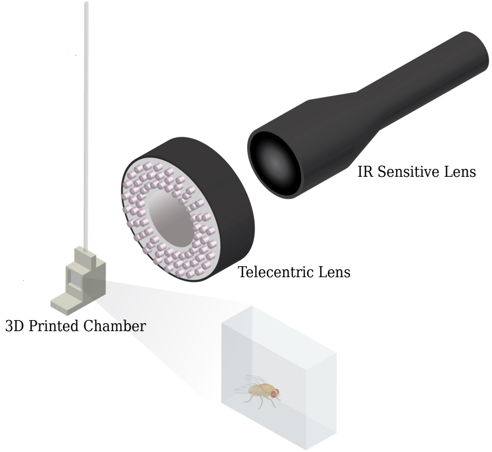

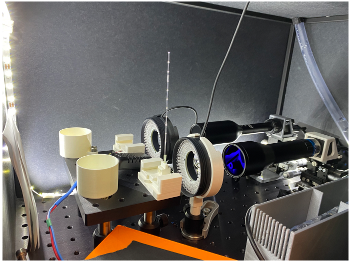

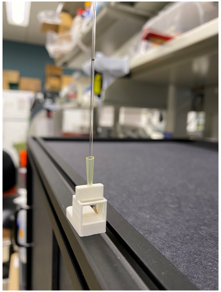

Performed by Dr. Mehmet Fatih Keles.

Conducting sleep experiments and collecting data

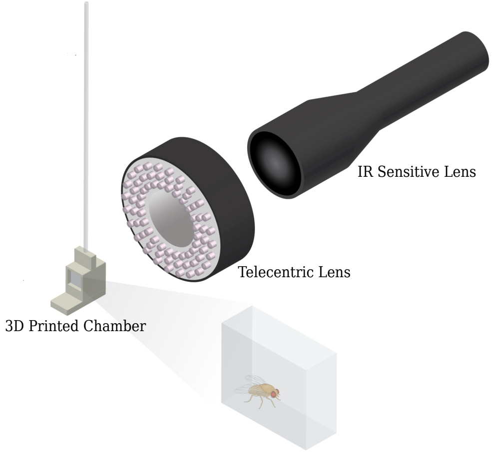

- Recorded between ZT10-ZT2 for wild-type sleep experiments, and between ZT10-ZT6 for sleep-deprived experiments.
- 30 FPS camera; resulting data contains more than one and a half million frames.

Performed by Dr. Mehmet Fatih Keles.

Pose estimation and body-part tracking

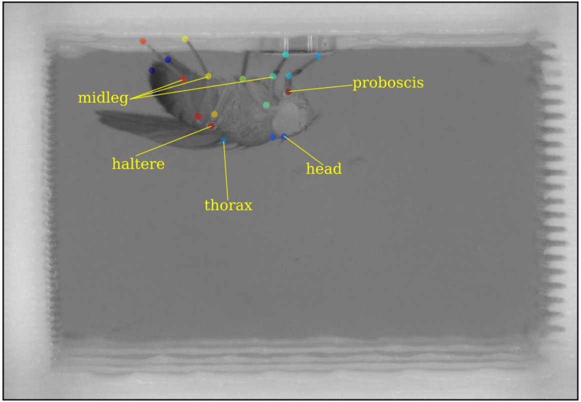

- Markerless pose estimation is performed.
- DeepLabCut (Mathis et al., 2018), a ResNet-50 neural network based pose estimation tool.
- 26 body parts are tracked in the videos.

Pose estimation and body-part tracking

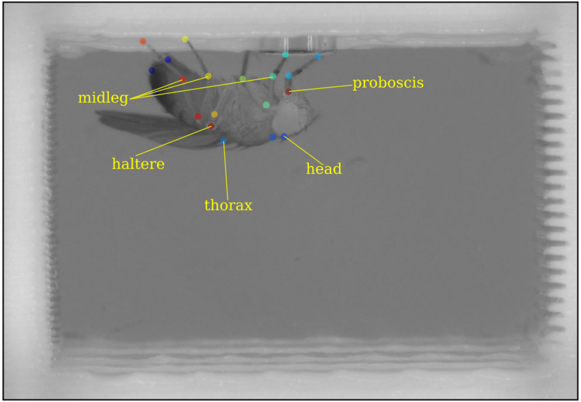

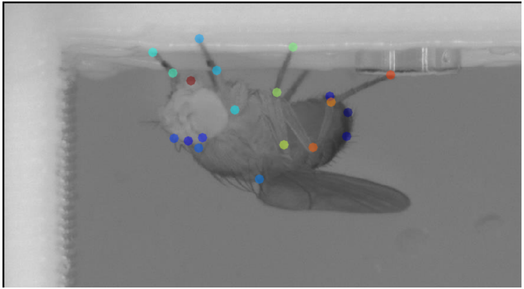

Pumping-like movement of proboscis

VIDEO

Switch-like movement of haltere

VIDEO

Annotating behavioral categories

5 behavioral categories are annotated: feeding, grooming, haltere switch, postural adjustment, and proboscis pumping.

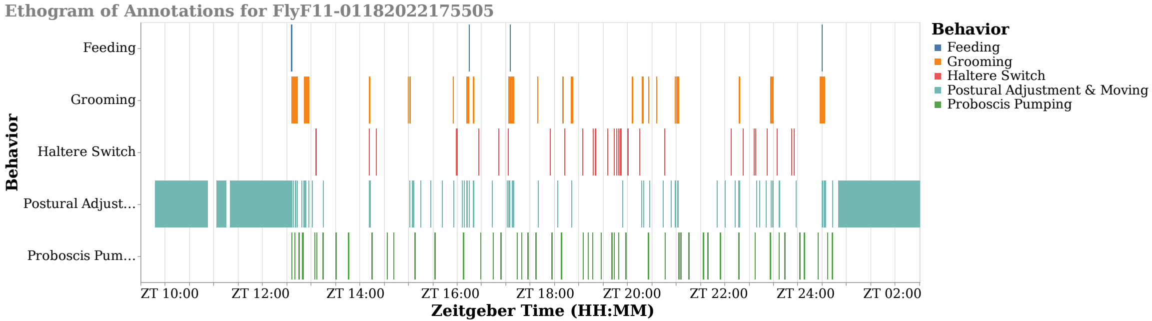

Figure: An example ethograms of annotated behaviors observed during the sleep experiments. The dark period starts at ZT12 and ends at ZT0. Behaviors are exhibited non-uniformly during sleep.

Annotating behavioral categories

Behavioral repertoires vary among experiments.

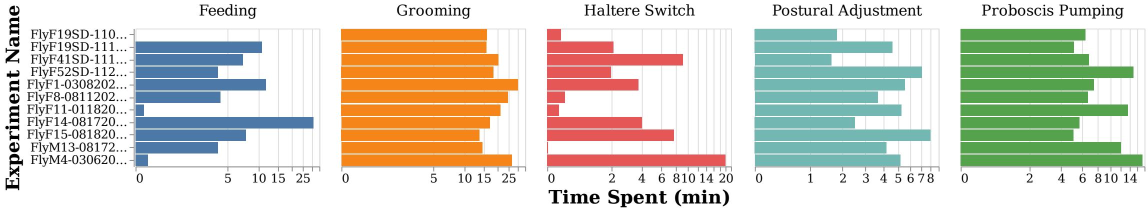

Figure: Time spent while exhibiting each behavioral category for all experiments.

Overview of the Pipeline

- **Feature Extraction**: Computing spatio-temporally meaningful behavioral representations.
- **Activity Detection**: Detecting micro-activities exhibited during sleep.
- **Behavior Mapping**: Mapping micro-activities to defined behavioral categories.

Computing meaningful spatio-temporal features

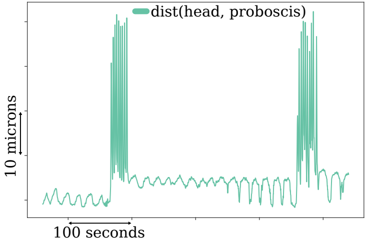

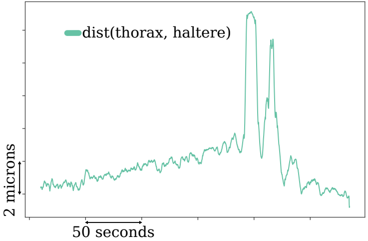

- A set of coordinate values are not descriptive enough.

Computing meaningful spatio-temporal features

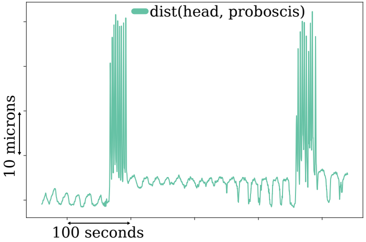

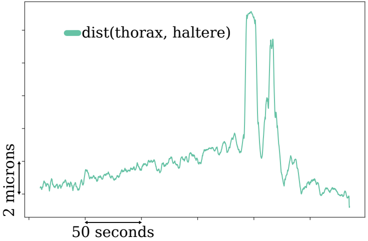

- A set of coordinate values are not descriptive enough.
- Meaningful spatio-temporal features need to be computed:

Computing meaningful spatio-temporal features

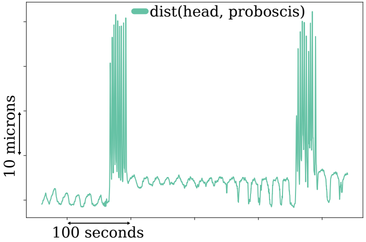

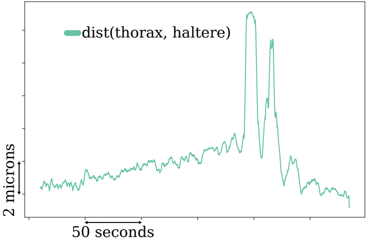

- A set of coordinate values are not descriptive enough.
- Meaningful spatio-temporal features need to be computed:
  - distances between body parts and corresponding velocities,
  - angles between body parts and corresponding angular velocities,
  - cartesian components, x and y, of coordinate values.

Computing meaningful spatio-temporal features

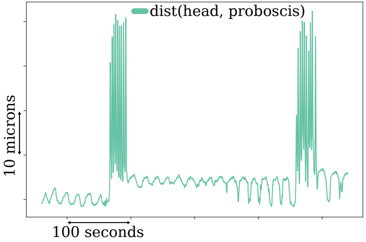

(a) Pumping-like movement of the proboscis.

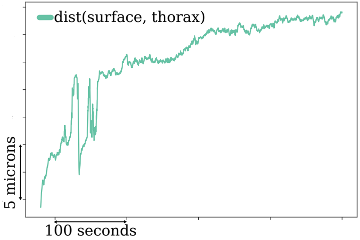

(b) Postural adjustment and movement of the thorax.

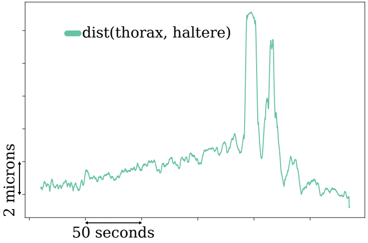

(c) Switch-like movement of the haltere.

Figure: Three spatio-temporal feature examples describing the kinematics of different behaviors.

Using a spectrogram representation

Wavelet domain is useful for representation of postural dynamics (Berman et al., 2014).

- Describing dynamics over multiple timescales simultaneously.

Using a spectrogram representation

Wavelet domain is useful for representation of postural dynamics (Berman et al., 2014).

- Describing dynamics over multiple timescales simultaneously.
- Capturing periodic movements and actions without precise temporal alignment.

Using a spectrogram representation

Wavelet domain is useful for representation of postural dynamics (Berman et al., 2014).

- Describing dynamics over multiple timescales simultaneously.
- Capturing periodic movements and actions without precise temporal alignment.

Using a spectrogram representation

Wavelet domain is useful for representation of postural dynamics (Berman et al., 2014).

- Describing dynamics over multiple timescales simultaneously.
- Capturing periodic movements and actions without precise temporal alignment.

Wavelet transformation

Morlet wavelet, and 20 frequency channels dyadically spaced between 1 Hz and 20 Hz.

Computing high-dimensional behavioral representations

Example

15 spatio-temporal features and 20 frequency channels result in a 15×20 matrix for each data point.

- We flatten feature tensor $T \times 15 \times 20$ to a generate a feature matrix $T \times 300$.

Computing high-dimensional behavioral representations

Example

15 spatio-temporal features and 20 frequency channels result in a 15×20 matrix for each data point.

- We flatten feature tensor $T \times 15 \times 20$ to a generate a feature matrix $T \times 300$.
- Data normalization is required (different orientations, sleep deprivation, sex).

Computing high-dimensional behavioral representations

Example

15 spatio-temporal features and 20 frequency channels result in a 15×20 matrix for each data point.

- We flatten feature tensor $T \times 15 \times 20$ to a generate a feature matrix $T \times 300$.
- Data normalization is required (different orientations, sleep deprivation, sex).
- We apply L_1 normalization to each frame.

Computing high-dimensional behavioral representations

Example

15 spatio-temporal features and 20 frequency channels result in a 15×20 matrix for each data point.

- We flatten feature tensor $T \times 15 \times 20$ to a generate a feature matrix $T \times 300$.
- Data normalization is required (different orientations, sleep deprivation, sex).
- We apply L_1 normalization to each frame.

Computing high-dimensional behavioral representations

Example

15 spatio-temporal features and 20 frequency channels result in a 15×20 matrix for each data point.

- We flatten feature tensor $T \times 15 \times 20$ to a generate a feature matrix $T \times 300$.
- Data normalization is required (different orientations, sleep deprivation, sex).
- We apply L_1 normalization to each frame.

High-dimensional behavioral representations

Each frame is represented by a probability distribution of features.

Spatio-temporal characteristics of annotated behaviors

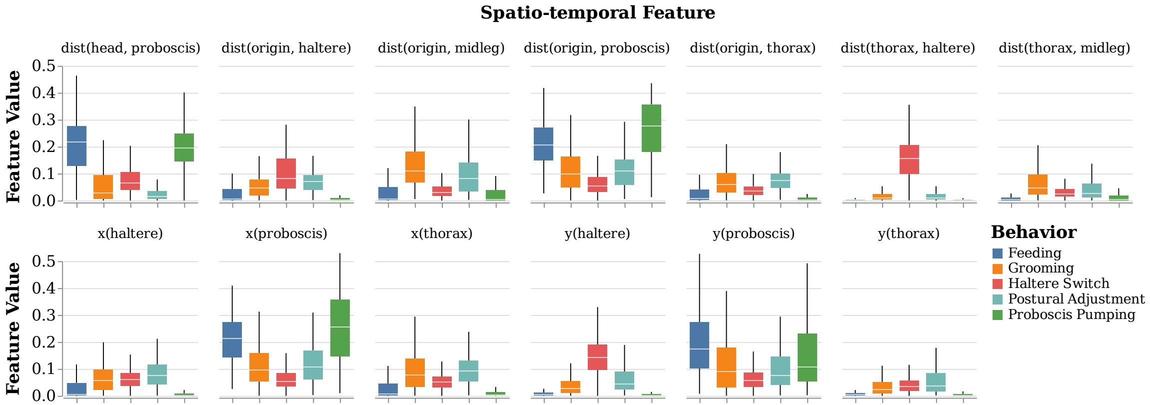

Figure: Summation of all frequency channels of behavioral representation value for each spatio-temporal feature.

Quantifying and detection micro-activity bouts

Dormancy

- Dormancy is characterized by displacement.

Quantifying and detection micro-activity bouts

Dormancy

- Dormancy is characterized by displacement.
- Displacement is quantified as the moving average of velocities.

$$v_t = \sum_{\tau \in T_G} \sum_{i=1}^{N_G} \mu_\tau \left(g_i, t\right)$$

Quantifying and detection micro-activity bouts

Dormancy

- Dormancy is characterized by displacement.
- Displacement is quantified as the moving average of velocities.

$$v_t = \sum_{\tau \in T_G} \sum_{i=1}^{N_G} \mu_\tau (g_i, t)$$

We end up with a single scalar value.
Use a Gaussian Mixture Model to detect a threshold.

Quantifying and detection micro-activity bouts

Dormancy

- Dormancy is characterized by displacement.
- Displacement is quantified as the moving average of velocities.

$$v_t = \sum_{\tau \in T_G} \sum_{i=1}^{N_G} \mu_\tau (g_i, t)$$

We end up with a single scalar value.
Use a Gaussian Mixture Model to detect a threshold.

Micro-activity

- Micro-activities are characterized by unobtrusive changes.

Quantifying and detection micro-activity bouts

Dormancy

- Dormancy is characterized by displacement.
- Displacement is quantified as the moving average of velocities.

$$v_t = \sum_{\tau \in T_G} \sum_{i=1}^{N_G} \mu_\tau \left(g_i, t\right)$$

Micro-activity

- Micro-activities are characterized by unobtrusive changes.
- Such changes are quantified as the sum of all scales in the power spectrum, separately for each body part.

$$u_{t,i} = \sum_{1/f \in T_G} W_f^0 \left(s_i, t\right)$$

We end up with a single scalar value.

Use a Gaussian Mixture Model to detect a threshold.

Quantifying and detection micro-activity bouts

Dormancy

- Dormancy is characterized by displacement.
- Displacement is quantified as the moving average of velocities.

$$v_t = \sum_{\tau \in T_G} \sum_{i=1}^{N_G} \mu_\tau (g_i, t)$$

We end up with a single scalar value.

Use a Gaussian Mixture Model to detect a threshold.

Micro-activity

- Micro-activities are characterized by unobtrusive changes.
- Such changes are quantified as the sum of all scales in the power spectrum, separately for each body part.

$$u_{t,i} = \sum_{1/f \in T_G} W_f^0 (s_i, t)$$

We end up with a scalar value for each feature.

Use a random forest ensemble of decision trees to classify.

Computation of behavioral embedding spaces

- High-dimensional behavioral representations and correlation between values.

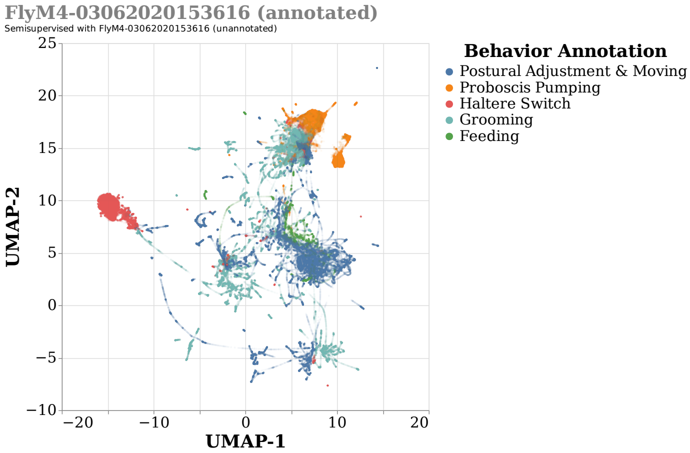

Figure: Semi-supervised pair embeddings with an annotated and unannotated experiment.

Computation of behavioral embedding spaces

- High-dimensional behavioral representations and correlation between values.
- We use Hellinger distance.

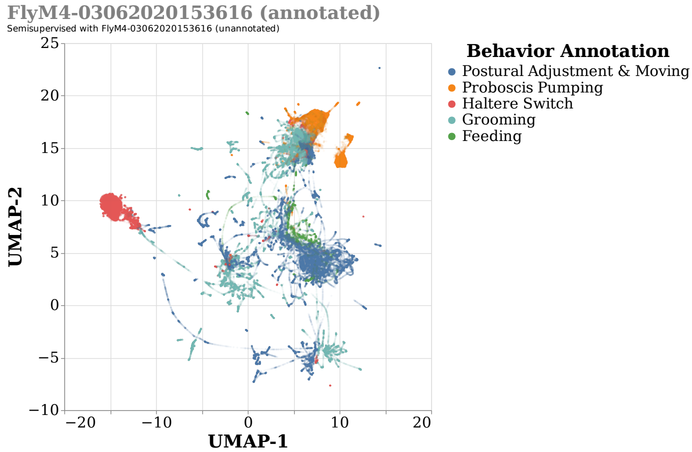

Figure: Semi-supervised pair embeddings with an annotated and unannotated experiment.

Computation of behavioral embedding spaces

- High-dimensional behavioral representations and correlation between values.
- We use Hellinger distance.
- Semi-supervised UMAP embeddings for each pair of annotated and unannotated experiments, providing views of the unannotated experiments.

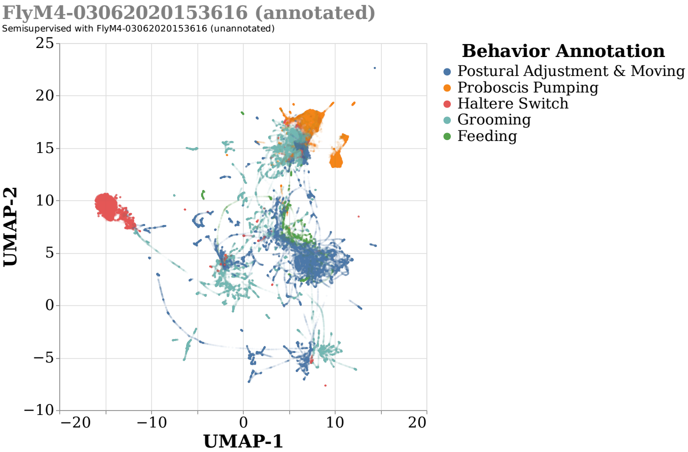

Figure: Semi-supervised pair embeddings with an annotated and unannotated experiment.

Computation of behavioral embedding spaces

- High-dimensional behavioral representations and correlation between values.
- We use Hellinger distance.
- Semi-supervised UMAP embeddings for each pair of annotated and unannotated experiments, providing views of the unannotated experiments.
- Only *Micro-activity* is embedded.

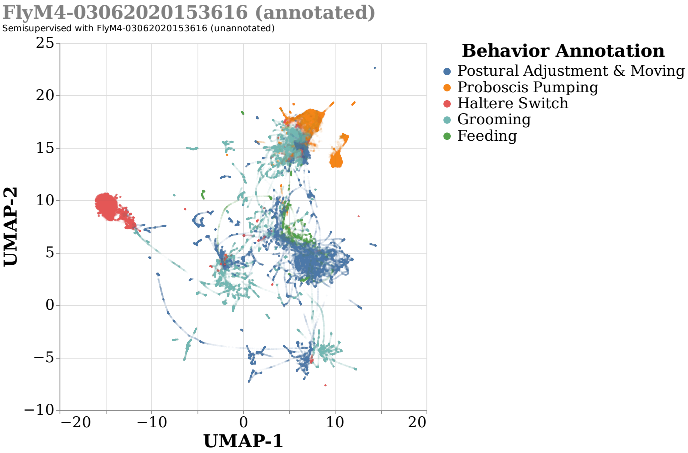

Figure: Semi-supervised pair embeddings with an annotated and unannotated experiment.

Computation of behavioral embedding spaces

- High-dimensional behavioral representations and correlation between values.
- We use Hellinger distance.
- Semi-supervised UMAP embeddings for each pair of annotated and unannotated experiments, providing views of the unannotated experiments.
- Only *Micro-activity* is embedded.

For R^- unannotated expt., and R^+ annotated expt., semi-supervised pair embeddings will be generated for each $R^- \times R^+$ pair.

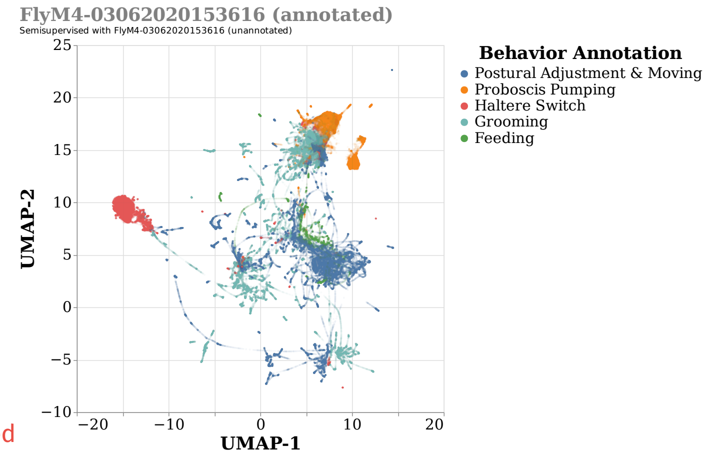

Figure: Semi-supervised pair embeddings with an annotated and unannotated experiment.

Nearest neighbor analysis and computing behavioral weights

Given behavioral embeddings

For a pair of annotated expt^{+} and unannotated expt^{-}; their semi-supervised pair embeddings are $\mathbf{E}^{+} = \left[\mathbf{e}_1^{+}, \cdots \mathbf{e}_{F^{+}}^{+}\right]$ and $\mathbf{E}^{-} = \left[\mathbf{e}_1^{-}, \cdots \mathbf{e}_{F^{-}}^{-}\right]$.

Nearest neighbor analysis and computing behavioral weights

Given behavioral embeddings

For a pair of annotated expt^{+} and unannotated expt^{-}; their semi-supervised pair embeddings are $\mathbf{E}^{+} = \left[\mathbf{e}^{+}_{1}, \cdots \mathbf{e}^{+}_{F^{+}}\right]$ and $\mathbf{E}^{-} = \left[\mathbf{e}^{-}_{1}, \cdots \mathbf{e}^{-}_{F^{-}}\right]$.

For frame f, and behavioral category i, neighbor weight is given by;

$$b_{f,i} = \sum_{f' \in k\text{-NN}_i(f)} \frac{1}{\mathrm{d}\left(\mathbf{e}^{-}_{f}, \mathbf{e}^{+}_{f'}\right) + \epsilon}$$

where $k\text{-NN}_i\,(f) = \left\{ f' \;:\; y^{+}_{f'} = i, \text{ and } f' \in k\text{-NN}\,(f) \right\}$ nearest neighbors of f annotated as behavior i.

Nearest neighbor analysis and computing behavioral weights

For frame f, and behavioral category i, neighbor weight is given by;

$$b_{f,i} = \sum_{f' \in k\text{-NN}_i(f)} \frac{1}{\mathrm{d}\left(\mathbf{e}^{-}_{f}, \mathbf{e}^{+}_{f'}\right) + \epsilon}.$$

To incorporate imbalance into the formulation, weights are normalized with the number of occurrences:

$$b'_{f,i} = \frac{b_{f,i}}{\log_2(1 + N^{+}_{i})}$$

where $N^{+}_{i} = \left|\left\{ f \;:\; y^{+}_{f} = i \right\}\right|$ is the number of frames annotated as behavioral category i.

Nearest neighbor analysis and computing behavioral weights

For frame f, and behavioral category i, neighbor weight is given by;

$$b_{f,i} = \sum_{f' \in k\text{-NN}_i(f)} \frac{1}{\mathrm{d}\left(\mathbf{e}^-_f, \mathbf{e}^+_{f'}\right) + \epsilon}.$$

To incorporate imbalance into the formulation, weights are normalized with the number of occurrences:

$$b'_{f,i} = \frac{b_{f,i}}{\log_2(1 + N^+_i)}.$$

$b'_{f,i}$ are not comparable among expt^+s. Using $\mathrm{softmax}$, $b'_{f,i}$ values are mapped to $[0, 1]$ by:

$$\hat{b}_{f,i} = \frac{\exp\left\{b'_{f,i}\right\}}{\sum_{j=1}^{K} \exp\left\{b'_{f,j}\right\}},$$

where behavioral weight vector $\hat{\mathbf{b}}_f \in [0, 1]^K$ can be considered as a probability distribution.

Forming a committee of experiments by voting

Given behavioral weights

Each $\mathrm{expt}^{+}\ k$ provides a view on the behavioral repertoire, and forming a committee by voting combines different views.

Forming a committee of experiments by voting

Given behavioral weights

Each expt^{+} k provides a view on the behavioral repertoire, and forming a committee by voting combines different views.

For frame f, and the behavioral category i, contribution of $exptAnn\ k$ to committee is given by:

$$vote^{k}_{f,i} = \left(1 - \max_{1 \le j \le K} \hat{b}^{k}_{f,j}\right) \hat{b}^{k}_{f,i},$$

where we penalize the uncertainty of the highest weight.

Forming a committee of experiments by voting

Given behavioral weights

Each expt[+] k provides a view on the behavioral repertoire, and forming a committee by voting combines different views.

For frame f, and the behavioral category i, contribution of $exptAnn\ k$ to committee is given by:

$$vote^{k}_{f,i} = \left(1 - \max_{1 \leq j \leq K} \hat{b}^{k}_{f,j}\right) \hat{b}^{k}_{f,i}.$$

A score value for each category is computed by combining votes by:

$$score_{f,i} = \frac{\sum_{k=1}^{R^{+}} vote^{k}_{f,i}}{\sum_{i=1}^{K} \sum_{k=1}^{R^{+}} vote^{k}_{f,i}}.$$

Forming a committee of experiments by voting

For frame f, and the behavioral category i, contribution of $exptAnn\ k$ to committee is given by:

$$vote_{f,i}^{k} = \left(1 - \max_{1 \leq j \leq K} \hat{b}_{f,j}^{k}\right) \hat{b}_{f,i}^{k}.$$

A score value for each category is computed by combining votes by:

$$score_{f,i} = \frac{\sum_{k=1}^{R^{+}} vote_{f,i}^{k}}{\sum_{i=1}^{K} \sum_{k=1}^{R^{+}} vote_{f,i}^{k}}.$$

Finally, prediction can be done by assigning the category with the highest score:

$$\hat{y}_{f} = \arg\max_{1 \leq i \leq K} score_{f,i}.$$

`basty` robustly detects micro-activities.

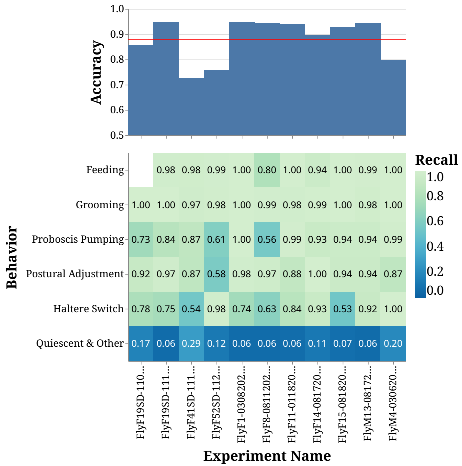

Figure: Performance summary of supervised micro-activity detection.

Definition

Micro-activity: Positive class predictions (union of annotated behavioral categories).
Quiescence: Negative class predictions (quiescent and other).

`basty` robustly detects micro-activities.

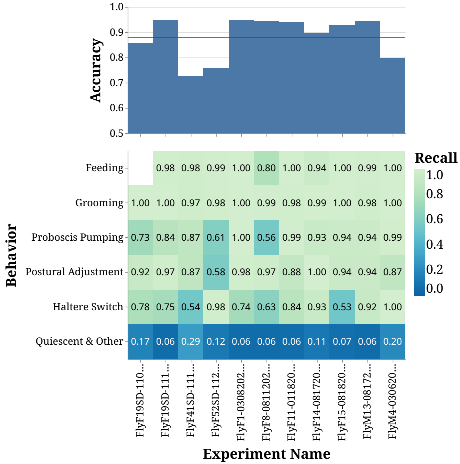

Figure: Performance summary of supervised micro-activity detection.

Definition

Micro-activity: Positive class predictions (union of annotated behavioral categories).
Quiescence: Negative class predictions (quiescent and other).

- High recall for annotated behaviors.

basty robustly detects micro-activities.

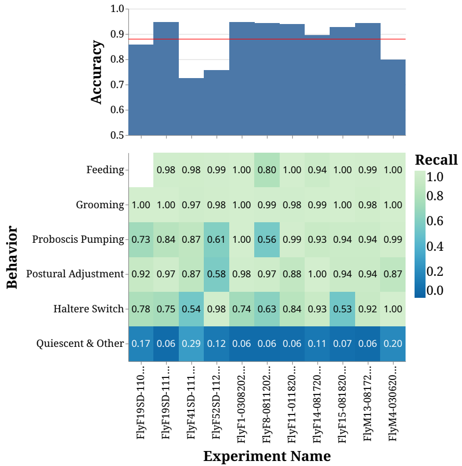

Figure: Performance summary of supervised micro-activity detection.

Definition

Micro-activity: Positive class predictions (union of annotated behavioral categories).
Quiescence: Negative class predictions (quiescent and other).

- High recall for annotated behaviors.
- Low recall for quiescent and other.

`basty` robustly detects micro-activities.

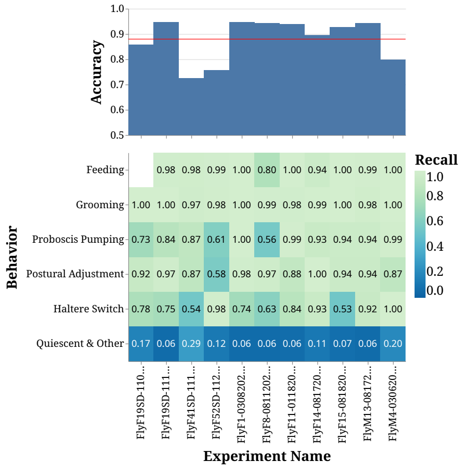

Figure: Performance summary of supervised micro-activity detection.

Definition

Micro-activity: Positive class predictions (union of annotated behavioral categories).
Quiescence: Negative class predictions (quiescent and other).

- High recall for annotated behaviors.
- Low recall for quiescent and other.
- False positive micro-activities negatively behavior affects mapping performance.

Score distributions reveal robust characterization of behaviors.

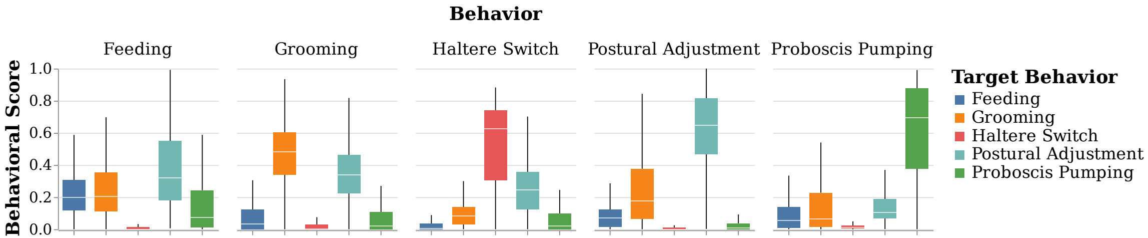

Figure: Distributions of behavioral score values of each behavioral category for all splits.

- Except for feeding, the median score for the correct category is greater than 0.5.

Score distributions reveal robust characterization of behaviors.

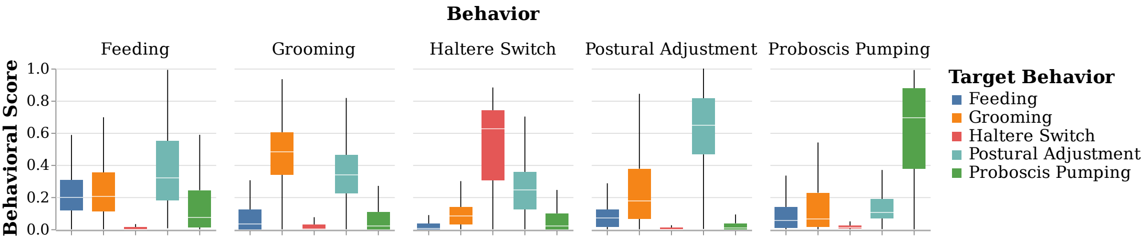

Figure: Distributions of behavioral score values of each behavioral category for all splits.

- Except for feeding, the median score for the correct category is greater than 0.5.
  - Behavioral scores reflect similarities between categories.

Score distributions reveal robust characterization of behaviors.

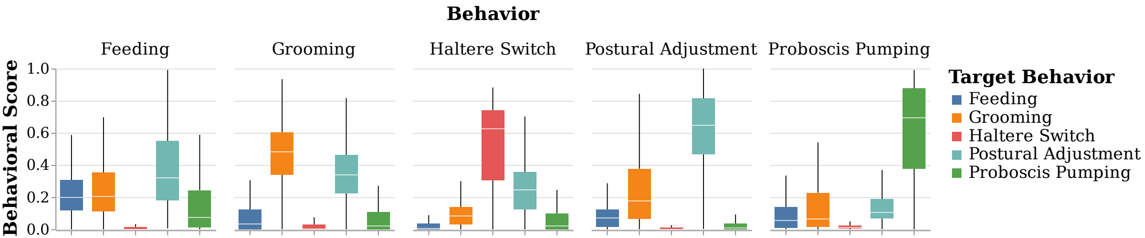

Figure: Distributions of behavioral score values of each behavioral category for all splits.

- Except for feeding, the median score for the correct category is greater than 0.5.
- Behavioral scores reflect similarities between categories.
- Feeding has the poorest performance.

`basty` achieves excellent discrimination.

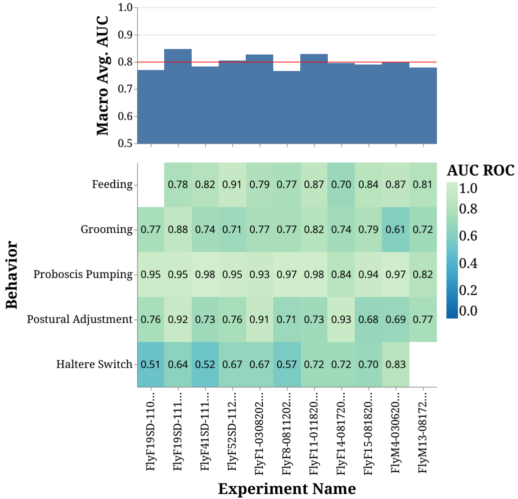

Figure: AUC scores for *Micro-activity* set.

- Considering all splits together, we achieve a macro average of 0.80 AUC score.

`basty` achieves excellent discrimination.

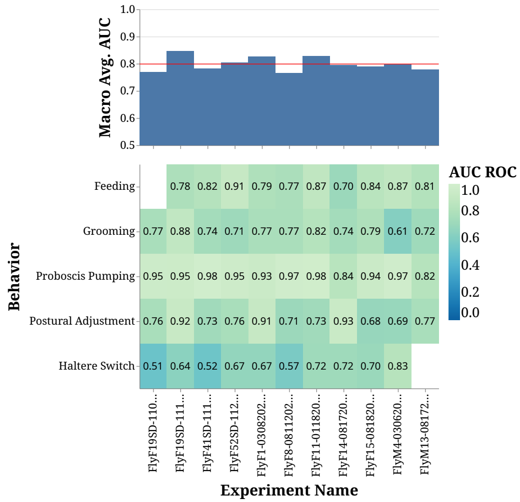

Figure: AUC scores for *Micro-activity* set.

- Considering all splits together, we achieve a macro average of 0.80 AUC score.
- Performance varies greatly among different experiments and categories.

`basty` achieves excellent discrimination.

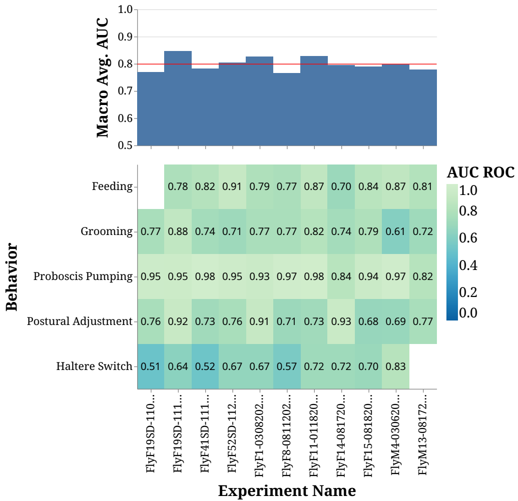

- Considering all splits together, we achieve a macro average of 0.80 AUC score.
- Performance varies greatly among different experiments and categories.
- False-positive micro-activities affect haltere switch most.

Figure: AUC scores for *Micro-activity* set.

ROC and precision-recall curves show successful prediction.

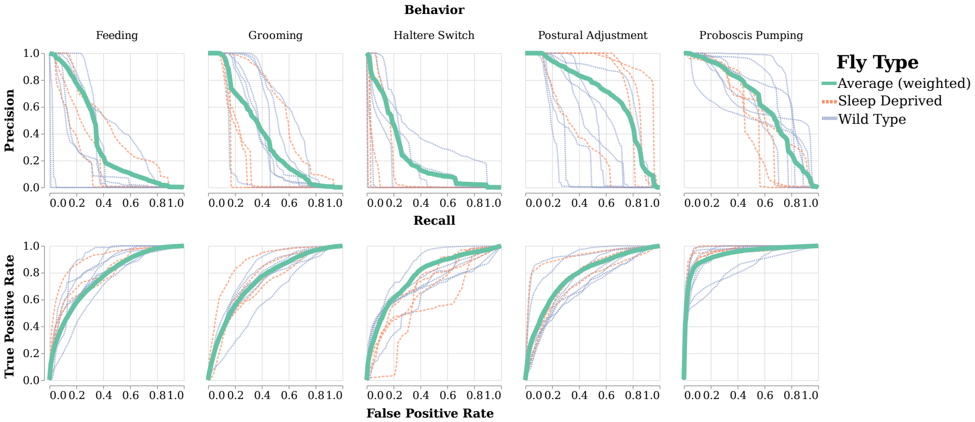

Figure: ROC and precision-recall curves for frames detected as micro-activity.

Our pipeline is able to detect unseen behavioral categories.

Question

What if the behavioral repertoires of annotated and unannotated experiments differ?

Our pipeline is able to detect unseen behavioral categories.

Question

What if the behavioral repertoires of annotated and unannotated experiments differ?

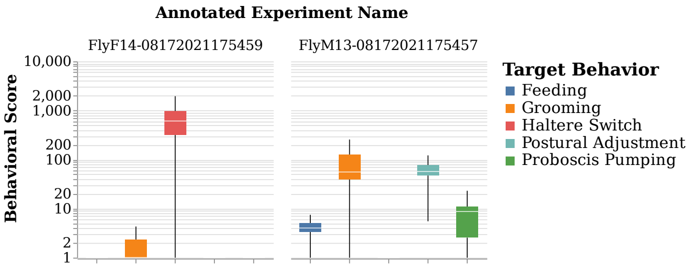

(a) Behavioral score values, before L_1 normalization.

- Mapping an experiment with many haltere switches.
- Comparison of behavioral scores computed separately.
- As opposed to FlyF14, FlyM13 lacks haltere switches.

Our pipeline is able to detect unseen behavioral categories.

What if the behavioral repertoires of annotated and unannotated experiments differ?

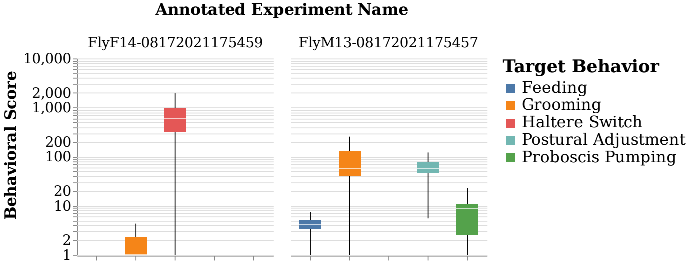

(a) Behavioral score values, before L_1 normalization.

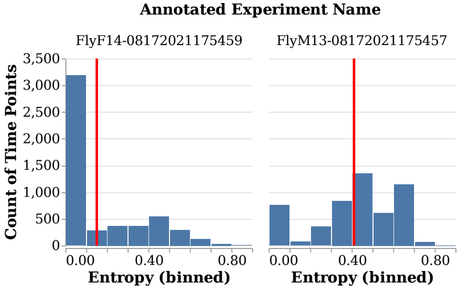

(b) Entropy values. The red line is the mean.

Figure: Histogram of entropy values and box-plot of the behavioral scores computed using one unannotated and two annotated experiments with varying behavioral repertoires.

Thank you!

Questions?

Details of extracted spatio-temporal features

	Snapshot Features	Gradient Features
Distance between	haltere and origin proboscis and origin thorax and origin head and proboscis haltere and thorax (average) midlegs and origin (average) midlegs and thorax	head and proboscis thorax and proboscis thorax and origin
Cartesian components of	haltere head thorax	head proboscis thorax

Table: Extracted spatio-temporal features.

Correlation between behavioral representations

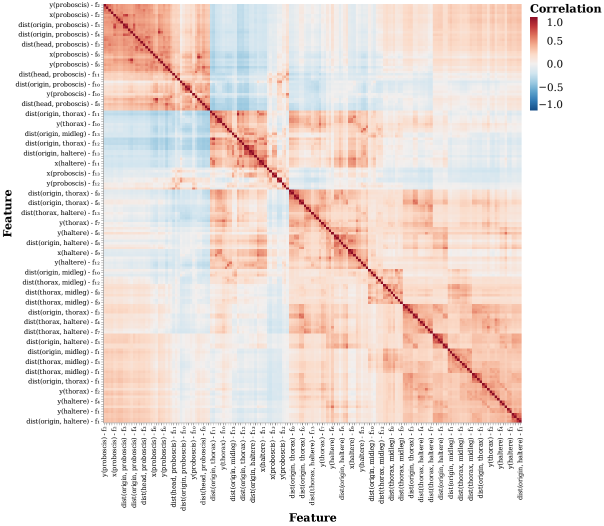

Spearman's correlation coefficient between each feature value of behavioral representations. Frequency channels $f_1, \cdots, f_{20}$ are dyadically spaced between 1 Hz and 20 Hz.

Comparison of supervised and unsupervised activity detection

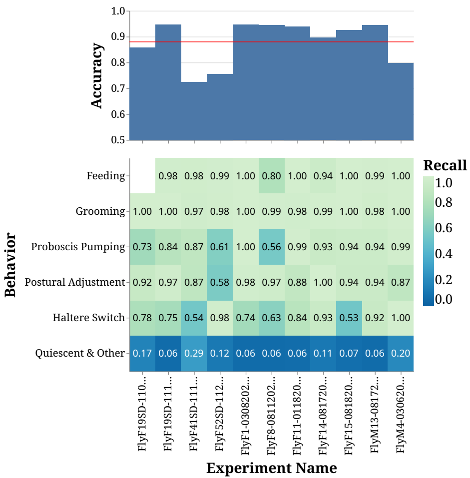

(a) Supervised detection.

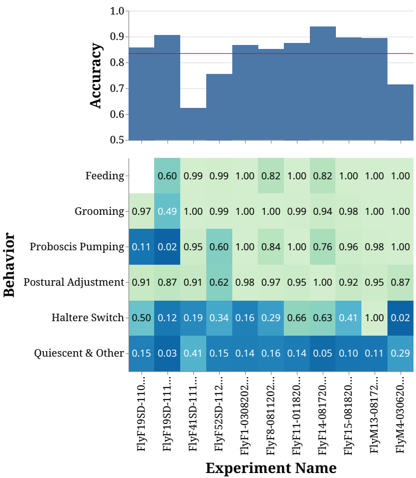

(b) Unsupervised detection.

Temporal organization of behaviors during sleep

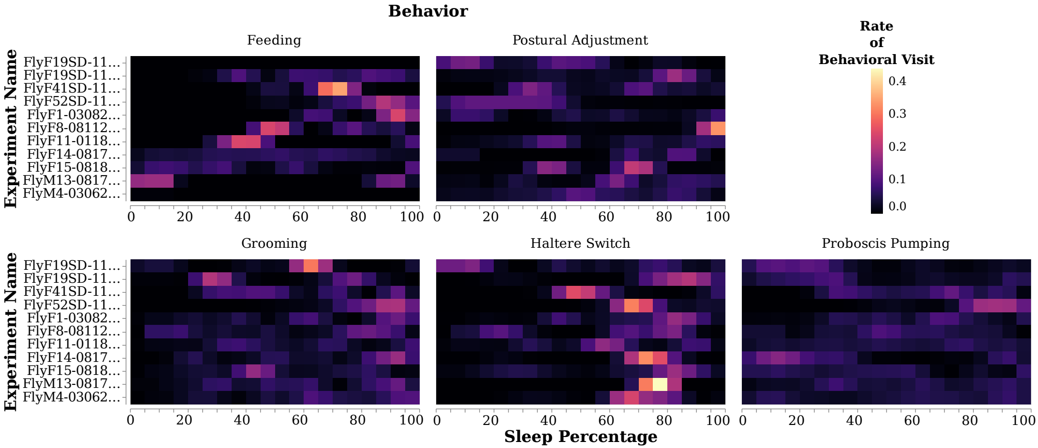

Figure: Rate of behavioral visits during the sleep within each 5 sleep percentage bin. Each row is normalized, to sum up to 1.

Temporal organization of behaviors during sleep

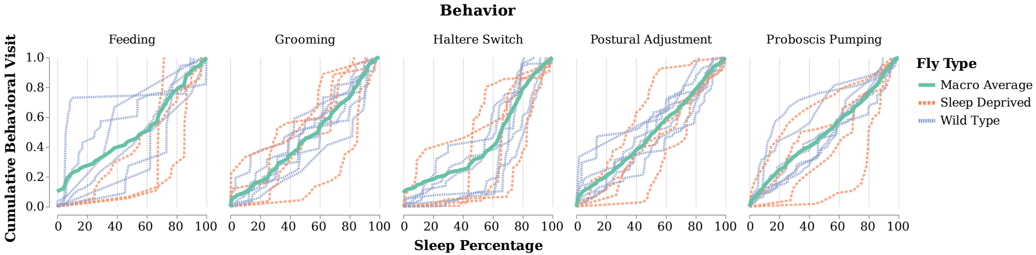

Figure: Cumulative behavioral visits for each behavior. The total number of behavioral visits is normalized by the total number of bouts for each fly and behavioral category, separately.

An interesting observation about haltere switches

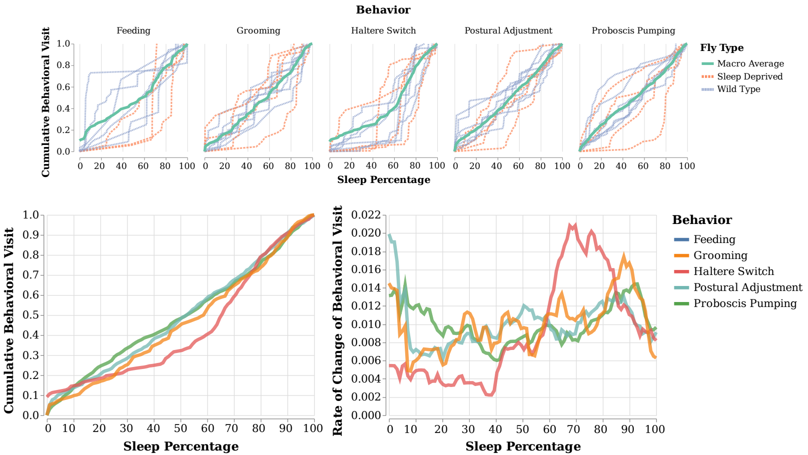

Figure: Occurrences of haltere switches are not uniform along the time, concentrated between 60-80 sleep percentages.

Bibliography

Anderson, D. J. and Perona, P. (2014). Toward a Science of Computational Ethology. *Neuron*, 84(1):18–31.

Berman, G. J., Choi, D. M., Bialek, W., and Shaevitz, J. W. (2014). Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of The Royal Society Interface*, 11(99):20140672. Publisher: Royal Society.

Datta, S. R., Anderson, D. J., Branson, K., Perona, P., and Leifer, A. (2019). Computational Neuroethology: A Call to Action. *Neuron*, 104(1):11–24.

Hsu, A. I. and Yttri, E. A. (2021). B-SOiD, an open-source unsupervised algorithm for identification and fast prediction of behaviors. *Nature Communications*, 12(1):5188.

Itskov, P. M., Moreira, J.-M., Vinnik, E., Lopes, G., Safarik, S., Dickinson, M. H., and Ribeiro, C. (2014). Automated monitoring and quantitative analysis of feeding behaviour in Drosophila. *Nature Communications*, 5(1):4560. Number: 1 Publisher: Nature Publishing Group.

Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, 21(9):1281–1289. Number: 9 Publisher: Nature Publishing Group.

Pereira, T. D. (2020). Quantifying behavior to understand the brain. *Nature Neuroscience*, 23:13.