
The Multi-Agent Behavior Dataset: Mouse Dyadic Social Interactions

Jennifer J. Sun
Caltech

Tomomi Karigo
Caltech

Dipam Chakraborty
AICrowd Research

Sharada P. Mohanty
AICrowd Research

Benjamin Wild
Freie Universität Berlin

Quan Sun
OPPO Research Institute

Chen Chen
OPPO Research Institute

David J. Anderson
Caltech

Pietro Perona
Caltech

Yisong Yue
Caltech

Ann Kennedy
Northwestern University
ann.kennedy@northwestern.edu

Dataset Website: <https://sites.google.com/view/computational-behavior/our-datasets/calms21-dataset>

Abstract

Multi-agent behavior modeling aims to understand the interactions that occur between agents. We present a multi-agent dataset from behavioral neuroscience, the Caltech Mouse Social Interactions (CalMS21) Dataset. Our dataset consists of trajectory data of social interactions, recorded from videos of freely behaving mice in a standard resident-intruder assay. To help accelerate behavioral studies, the CalMS21 dataset provides benchmarks to evaluate the performance of automated behavior classification methods in three settings: (1) for training on large behavioral datasets all annotated by a single annotator, (2) for style transfer to learn inter-annotator differences in behavior definitions, and (3) for learning of new behaviors of interest given limited training data. The dataset consists of 6 million frames of unlabeled tracked poses of interacting mice, as well as over 1 million frames with tracked poses and corresponding frame-level behavior annotations. The challenge of our dataset is to be able to classify behaviors accurately using both labeled and unlabeled tracking data, as well as being able to generalize to new settings.

1 Introduction

The behavior of intelligent agents is often shaped by interactions with other agents and the environment. As a result, models of multi-agent behavior are of interest in diverse domains, including neuroscience [55], video games [26], sports analytics [74], and autonomous vehicles [8]. Here, we study multi-agent animal behavior from neuroscience and introduce a dataset to benchmark behavior model performance.

Traditionally, the study of animal behavior relied on the manual, frame-by-frame annotation of behavioral videos by trained human experts. This is a costly and time-consuming process, and cannot easily be crowdsourced due to the training required to identify many behaviors accurately. Automated behavior classification is a popular emerging tool [29, 2, 18, 44, 55], as it promises to reduce human annotation effort, and opens the field to more high-throughput screening of animal behaviors. However, there are few large-scale publicly available datasets for training and benchmarking social behavior classification, and the behaviors annotated in those datasets may not match the set of behaviors a particular researcher wants to study. Collecting and labeling enough training data to

reliably identify a behavior of interest remains a major bottleneck in the application of automated analyses to behavioral datasets.

We present a dataset of behavior annotations and tracked poses from pairs of socially interacting mice, the **Caltech Mouse Social Interactions 2021 (CalMS21) Dataset**, with the goal of advancing the state-of-the-art in behavior classification. From top-view recorded videos of mouse interactions, we detect seven keypoints for each mouse in each frame using Mouse Action Recognition System (MARS) [55]. Accompanying the pose data, we introduce three tasks pertaining to the classification of frame-level social behavior exhibited by the mice, with frame-by-frame manual annotations of the behaviors of interest (Figure 1), and additionally release video data for a subset of the tasks. Finally, we release a large dataset of tracked poses without behavior annotations, that can be used to study unsupervised learning methods.

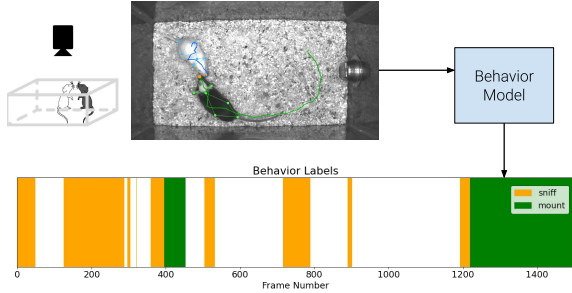


Figure 1: Overview of behavior classification. A typical behavior study starts with extraction of tracking data from videos. We show 7 keypoints for each mouse, and draw the trajectory of the nose keypoint. The goal of the model is to classify each frame (30Hz) to one of the behaviors of interest from domain experts.

As part of the initial benchmarking of CalMS21, we both evaluated standard baseline methods as well as solicited novel methods by having CalMS21 as part of the Multi-Agent Behavior (MABe) Challenge 2021 hosted at CVPR 2021. To test model generalization, our dataset contains splits annotated by different annotators and for different behaviors.

In addition to providing a performance benchmark for multi-agent behavior classification, our dataset is suitable for studying several research questions, including: How do we train models that transfer well to new conditions (annotators and behavior labels)? How do we train accurate models to identify rare behaviors? How can we best leverage unlabeled data for behavior classification?

2 Related Work

Behavior Classification. Automated behavior classification tools are becoming increasingly adopted in neuroscience [55, 19, 29, 2, 44]. These automated classifiers generally consists of the following steps: pose estimation, feature computation, and behavior classification. Our dataset provides the output from our mouse pose tracker, MARS [55], to allow participants in our dataset challenge to focus on developing methods for the latter steps of feature computation and behavior classification. We will therefore first focus our exploration of related works on these two topics, specifically within the domain of neuroscience, then discuss how our work connects to the related field of human action recognition.

Existing behavior classification methods are typically trained using tracked poses or hand-designed features in a fully-supervised fashion with human-annotated behaviors [27, 55, 5, 19, 44]. Pose representations used for behavior classification can take the form of anatomically defined keypoints [55, 44], fit shapes such as ellipses [27, 12, 46], or simply a point reflecting the location of an animal’s centroid [45, 51, 68]. Features extracted from poses may reflect values such as animal velocity and acceleration, distances between pairs of body parts or animals, distances to objects or parts of the arena, and angles or angular velocities of keypoint-defined joints. To bypass the effort-intensive step of hand-designing pose features, self-supervised methods for feature extraction have been explored [59]. Computational approaches to behavior analysis in neuroscience have been recently reviewed in [50, 13, 17, 2].

Relating to behavior classification and works in behavioral neuroscience, there is also the field of human action recognition (reviewed in [71, 75]). We compare this area to our work in terms of models and data. Many works in action recognition are trained end-to-end from image or video data [38, 57, 63, 7], and the sub-area that is most related to our dataset is pose/skeleton-based action recognition [10, 37, 60], where model input is also pose data. However, one difference is that these




Figure 2: **Behavior classes and annotator variability.** A. Example frames showing some behaviors of interest. B. Domain expert variability in behavior annotation, reproduced with permission from [55]. Each row shows annotations from a different domain expert annotating the same video data.

models often aim to predict one action label per video, since in many datasets, the labels are annotated at the video or clip level [7, 76, 56]. More closely related to our dataset is the works on temporal segmentation [34, 58, 33], where one action label is predicted per frame in long videos. These works are based on human activities, often goal-directed in a specific context, such as cooking in the kitchen. We would like to note a few unique aspects animal behavior. In contrast to many human action recognition datasets, naturalistic animal behavior often requires domain expertise to annotate, making it more difficult to obtain. Additionally, most applications of animal behavior recognition are in scientific studies in a laboratory context, meaning that the environment is under close experimenter control.

Unsupervised Learning for Behavior. As an alternative to supervised behavior classification, several groups have used unsupervised methods to identify actions from videos or pose estimates of freely behaving animals [3, 32, 69, 64, 39, 28, 41] (also reviewed in [13, 50]). In most unsupervised approaches, videos of interacting animals are first processed to remove behavior-irrelevant features such as the absolute location of the animal; this may be done by registering the animal to a template or extract a pose estimate. Features extracted from the processed videos or poses are then clustered into groups, often using a model that takes into account the temporal structure of animal trajectories, such as a set of wavelet transforms [3], an autoregressive hidden Markov model [69], or a recurrent neural network [39]. Behavior clusters produced from unsupervised analysis have been shown to be sufficiently sensitive to distinguish between animals of different species and strains [25, 39], and to detect effects of pharmacological perturbations [70]. Clusters identified in unsupervised analysis can often be related back to human-defined behaviors via post-hoc labeling [69, 3, 64], suggesting that cluster identities could serve as a low-dimensional input to a supervised behavior classifier.

Related Datasets. The CalMS21 dataset provides a benchmark to evaluate the performance of behavior analysis models. Related animal social behavior datasets include CRIM13 [5] and Fly vs. Fly [19], which focus on supervised behavior classification. In comparison to existing datasets, CalMS21 enables evaluation in multiple settings, such as for annotation style transfer and for learning new behaviors. The trajectory data provided by the MARS tracker [55] (seven keypoints per mouse) in our dataset also provides a richer description of the agents compared to single keypoints (CRIM13). Additionally, CalMS21 is a good testbed for unsupervised and self-supervised models, given its inclusion of a large (6 million frame) unlabeled dataset.

While our task is behavior classification, we would like to note that there are also a few datasets focusing on the related task of multi-animal tracking [48, 22]. Multi-animal tracking can be difficult due to occlusion and identity tracking over long timescales. In our work, we used the output of the MARS tracker [55], which also includes a multi-animal tracking dataset on the two mice to evaluate pose tracking performance; we bypass the problem of identity tracking by using animals of differing coat colors. Improved methods to more accurately track multi-animal data is another direction that can help quantify animal behavior.

Other datasets studying multi-agent behavior include those from autonomous driving [8, 61], sports analytics [73, 14], and video games [54, 23]. Generally, the autonomous vehicle datasets focus on tracking and forecasting, whereas trajectory data is already provided in CalMS21, and our focus is on behavior classification. Sports analytics datasets also often involves forecasting and learning player strategies. Video game datasets have perfect tracking and generally focus on reinforcement learning or imitation learning of agents in the simulated environment. While the trajectories in CalMS21 can be used for imitation learning of mouse behavior, our dataset also consist of expert human annotations of behaviors of interest used in scientific experiments. As a result, CalMS21 can be used to benchmark supervised or unsupervised behavior models against expert human annotations of behavior.

3 Dataset Design

The CalMS21 dataset is designed for studying behavior classification, where the goal is to assign frame-wise labels of animal behavior to temporal pose tracking data. The tracking data is a top-view pose estimate of a pair of interacting laboratory mice, produced from raw 30Hz videos using MARS [55], and reflecting the location of the nose, ears, neck, hips, and tail base of each animal (Figure 3).

We define three behavior classification tasks on our dataset. In Task 1 (Section 3.1), we evaluate the ability of models to classify three social behaviors of interest (attack, mount, and close investigation) given a large training set of annotated videos; sample frames of the three behaviors are shown in Figure 2A. In Task 2 (Section 3.2), models must be adjusted to reproduce new annotation styles for the behaviors studied in Task 1: Figure 2B demonstrates the annotator variability that can exist for the same videos with the same labels. Finally, in Task 3 (Section 3.3), models must be trained to classify new social behaviors of interest given limited training data.




Figure 3: **Pose keypoint definitions.** Illustration of the seven anatomically defined keypoints tracked on the body of each animal. Pose estimation is performed using MARS [55].

In Tasks 1 and 2, each frame is assigned one label (including "other" when no social behavior is shown), therefore these tasks can be handled as multi-class classification problems. In Task 3, we provide separate training sets for each of seven novel behaviors of interest, where in each training set only a single behavior has been annotated. For this task, model performance is evaluated for each behavior separately: therefore, Task 3 should be treated as a set of binary classification problems. Behaviors are temporal by nature, and often cannot be accurately identified from the poses of animals in a single frame of video. Thus, all three tasks can be seen as time series prediction problems or sequence-to-sequence learning, where the time-evolving trajectories of 28-dimensional animal pose data (7 keypoints x 2 mice x 2 dimensions) must be mapped to a behavior label for each frame. Tasks 2 and 3 are also examples of few-shot learning problems, and would benefit from creative forms of data augmentation, task-efficient feature extraction, or unsupervised clustering to stretch the utility of the small training sets provided.

To encourage the combination of supervised and unsupervised methods, we provide a large set of unlabeled videos (around 6 million frames) that can be used for feature learning or clustering in any task (Figure 4).

3.1 Task 1: Classical Classification

Task 1 is a standard sequential classification task: given a large training set comprised of pose trajectories and frame-level annotations of freely interacting mice, we would like to produce a model to predict frame-level annotations from pose trajectories on a separate test set of videos. There are 70 videos in the public training set, all of which have been annotated for three social behaviors of interest: close investigation, attack, and mount (described in more detail in the appendix). The goal is for the model to reproduce the behavior annotation style from the training set.

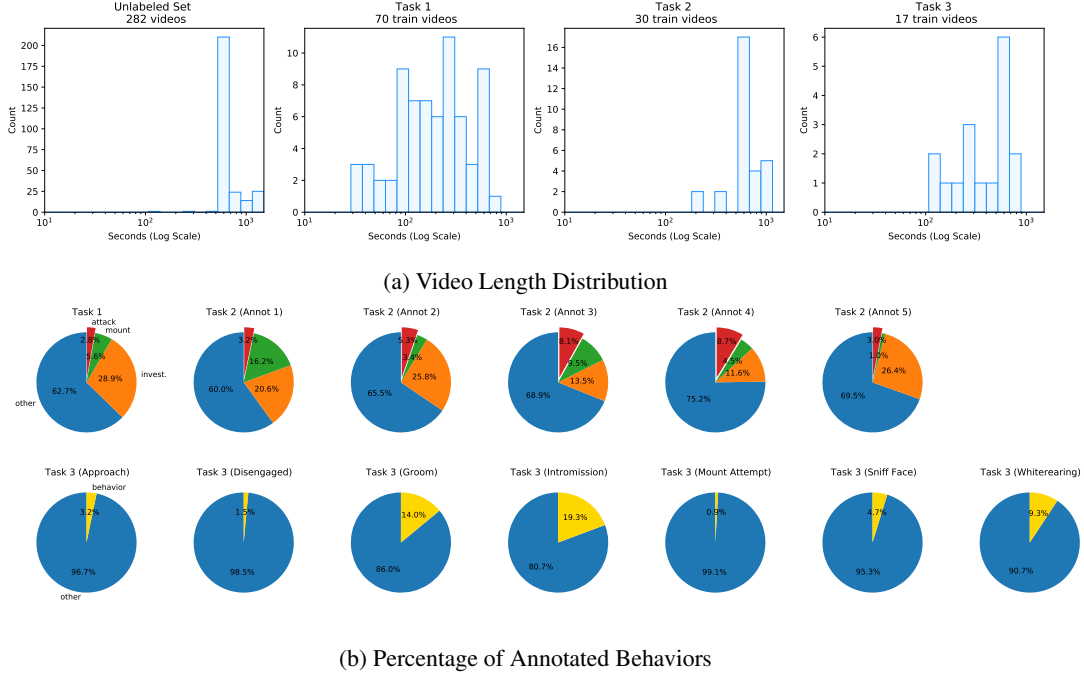


Figure 4: **Available data for each task in our challenge.** Our dataset consists of a large set of unlabeled videos alongside a set of annotated videos from one annotator. Annot 1, 2, 3, 4, 5 are different domain experts, whose annotations for attack, mount, and investigation are used in Task 2. Bottom row shows new behaviors used in Task 3.

Sequential classification has been widely studied, existing works use models such as recurrent neural networks [11], temporal convolutional networks [35], and random forests with hand-designed input features [55]. Input features to the model can also be learned with self-supervision [59, 9, 36], which can improve classification performance using the unlabeled portion of the dataset.

In addition to pose data, we also release all source videos for Task 1, to facilitate development of methods that require visual data.

3.2 Task 2: Annotation Style Transfer

In general, when annotating the same videos for the same behaviors, there exists variability across annotators (Figure 2B). As a result, models that are trained for one annotator may not generalize well to other annotators. Given a small amount of data from several annotators, we would like to study how well a model can be trained to reproduce each individual’s annotation style. Such an "annotation style transfer" method could help us better understand differences in the way behaviors are defined across annotators and labs, increasing the reproducibility of experimental results. A better understanding of different annotation styles could also enable crowdsourced labels from non-experts to be transferred to the style of expert labels.

In this sequential classification task, we provide six 10-minute training videos for each of five annotators unseen in Task 1, and evaluate the ability of models to produce annotations in each annotator’s style. All annotators are trained annotators from the David Anderson Lab, and have between several months to several years of prior annotation experience. The behaviors in the training datasets are the same as Task 1. In addition to the annotator-specific videos, competitors have access to a) the large collection of task 1 videos, that have been annotated for the same behaviors but in a different style, and b) the pool of unannotated videos, which could be used for unsupervised clustering or feature learning.

This task is suitable for studying techniques from transfer learning [62] and domain adaptation [65]. We have a source domain with labels from task 1, which needs to be transferred to each annotator in task 2 with comparatively fewer labels. Potential directions include learning a common set

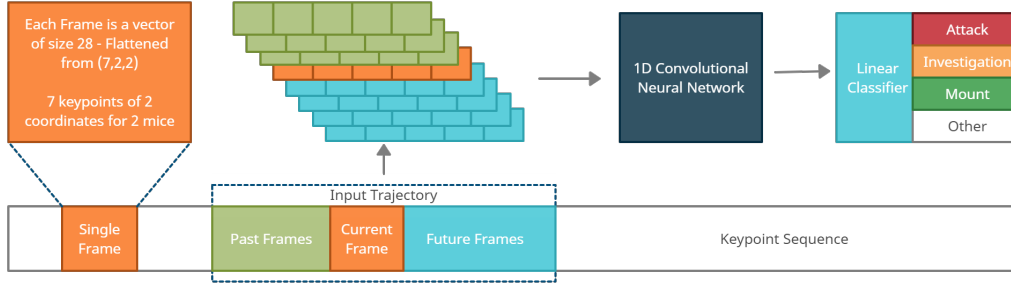


Figure 5: **Sequence Classification Setup.** Sequence information from past, present, and future frames may be used to predict the observed behavior label on the current frame. Here, we show a 1D convolutional neural network, but in general any model may be used.

of data-efficient features for both tasks [59], and knowledge transfer from a teacher to a student network [1].

3.3 Task 3: New Behaviors

It is often the case that different researchers will want to study different behaviors in the same experimental setting. The goal of Task 3 is to help benchmark general approaches for training new behavior classifiers given a small amount of data. This task contains annotations on seven behaviors not labeled in Tasks 1 & 2, where some behaviors are very rare (Figure 4).

As for the previous two tasks, we provide a training set of videos in which behaviors have been annotated on a frame-by-frame basis, and evaluate the ability of models to produce frame-wise behavior classifications on a held-out test set. We expect that the large unlabeled video dataset (Figure 4) will help improve performance on this task, by enabling unsupervised feature extraction or clustering of animal pose representations prior to classifier training.

Since each new behavior has a small amount of data, few-shot learning techniques [66] can be helpful for this task. The data from Task 1 and the unlabeled set could also be used to set up multi-task learning [77], and for knowledge transfer [1].

4 Benchmarks on CalMS21

We develop an initial benchmark on CalMS21 based on standard baseline methods for sequence classification. To demonstrate the utility of the unlabeled data, we also used these sequences to train a representation learning framework (task programming [59]) and added the learned trajectory features to our baseline models. Additionally, we presented CalMS21 at the MABe Challenge hosted at CVPR 2021, and we include results on the top performing methods for each of the tasks.

Our evaluation metrics are based on class-averaged F1 score and Mean Average Precision (more details in the appendix). Unless otherwise stated, the class with the highest predicted probability in each frame was used to compute F1 score.

4.1 Baseline Model Architectures

Our goal is to estimate the behavior labels in each frame from trajectory data, and we use information from both past and future frames for this task (Figure 5). To establish baseline classification performance on CalMS21, we explored a family of neural network-based classification approaches (Figure 6). All models were trained using categorical cross entropy loss [21] on Task 1: Classic Classification, using an 80/20 split of the train split into training and validation sets during development. We report results on the full training set after fixing model hyperparameters.

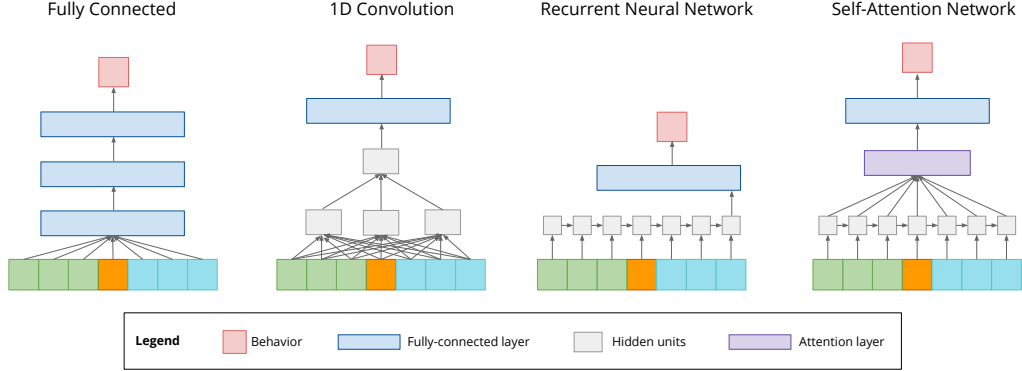


Figure 6: **Baseline models.** Different baseline setups we evaluated for behavior classification. The input frame coloring follows the same convention as Figure 5: past frames in green, current frame in orange, and future frames in cyan.

Among the explored architectures, we obtained the highest performance using the 1D convolutional neural network (Table 1). We therefore used this architecture for baseline evaluations in all subsequent sections.

Hyperparameters we considered includes the number of input frames, the number of frame skips, the number of units per layer, and the learning rate. Settings of these parameters may be found in the project code and the appendix. The baseline with task programming models use the same hyperparameters as the baseline. The task programming model is trained on the unlabeled set only, and the learned features are concatenated with the keypoints at each frame.

4.2 Task 1 Classic Classification Results

Baseline Models. We used the 1D convolutional neural network model outlined above (Figure 5) to predict attack, investigation, and mount behaviors in two settings: using raw pose trajectories, and using trajectories plus features learned from the unlabeled set using task programming (Table 2). We found that including task programming features improved model performance. Many prediction errors of the baseline models were localized around behavior transition boundaries (Figure 7). These errors may arise in part from annotator noise in the human generated labels of behaviors. An analysis of such intra- (and inter-) annotator variability is found in [55].

Task1 Top-1 Entry. We also include the top-1 entry in Task 1 of our dataset at MABe 2021 as part of our benchmark (Table 2). This model starts from an egocentric representation of the data; in a preprocessing stage, features are computed based on distances, velocities, and angles between coordinates relative to the agents' orientations. Furthermore, a PCA embedding of pairwise distances of all coordinates of both individuals is given as input to the model.

The model architecture is based on [47] with three main components. First, the embedder network consists of several residual blocks [24] of non-causal 1D convolutional layers. Next, the contexter network is a stack of residual blocks with causal 1D convolutional layers. Finally, a fully connected residual block with multiple linear classification heads computes the class probabilities for each behavior. Additional inputs such as a learned embedding for the annotator (see Section 4.3) and absolute spatial and temporal information are directly fed into this final component.

The Task 1 top-1 model was trained in a semi-supervised fashion, using the normalized temperature-scaled cross-entropy loss [47, 9] for all samples and the categorical cross-entropy loss for labeled samples. During training, sequences were sampled from the data proportional to their length with a 3:1 split of unlabeled/labeled sequences. Linear transformations that project the contexter component's

Method	Average F1	MAP
Fully Connected	.659 \pm .005	.726 \pm .004
LSTM	.675 \pm .011	.712 \pm .013
Self-Attention	.610 \pm .028	.644 \pm .018
1D Conv Net	.793 \pm .011	.856 \pm .010

Table 1: Class-averaged results on Task 1 (attack, investigation, mount) for different baseline model architectures. The value is shown of the mean and standard deviation over 5 runs.

Method	Data Used During Training			Average F1	MAP
	Task 1 (train split)	Unlabeled Set	All Tasks (all splits)		
Baseline	✓			.793 ± .011	.856 ± .010
Baseline w/ task prog	✓	✓		.829 ± .004	.889 ± .004
MABe 2021 Task 1 Top-1	✓	✓	✓	.864 ± .011	.914 ± .009

Table 2: Class-averaged results on Task 1 (attack, investigation, mount; mean \pm standard deviation over 5 runs.) See appendix for per class results. The “All Tasks” column indicates that the model was jointly trained on all three tasks, and “all splits” indicates that both labeled train set and trajectory from unlabeled test set are used.




Figure 7: Example of errors from a sequence of behaviors from Task 1.

outputs into the future are learned jointly with the model, as described in [47]. This unsupervised loss component regularizes the model by encouraging the model to learn a representation that is predictive of future behavior. A single model was trained jointly for all three tasks of the challenge, with all parameters being shared among the tasks, except for the final linear classification layers. The validation loss was monitored during training, and a copy of the parameters with the lowest validation loss was stored for each task individually.

4.3 Task 2 Annotation Style Transfer Results

Baseline Model. Similar to Task 1, Task 2 involves classifying attack, investigation, and mount frames. However, in this task, our goal is to capture the particular annotation style of different individuals. This step is important in identifying sources of discrepancy in behavior definitions between datasets or labs. Given the limited training set size in Task 2 (only 6 videos for each annotator), we used the model trained on Task 1 as a pre-trained model for the baseline experiments in Task 2, to leverage the larger training set from Task 1. The performances are in Table 3, with per-annotator results in the appendix.

Task2 Top-1 Entry. The top MABe submission for Task 2 re-used the model architecture and training schema from Task 1, described in Section 4.2. To address different annotation styles in Task 2, a learned annotator embedding was concatenated to the outputs of the contexter network. This embedding was initialized as a diagonal matrix such that initially, each annotator is represented by a one-hot vector. This annotator matrix is learnable, so the network can learn to represent similarities in the annotation styles. Annotators 3 and 4 were found to be similar to each other, and different from annotators 1, 2, and 5. The learned annotator matrix is provided in the appendix.

4.4 Task 3 New Behaviors Results

Baseline Model. Task 3 is a set of data-limited binary classification problems with previously unseen behaviors. Although these behaviors do occur in the Task 1 and Task 2 datasets, they are not labeled. The challenges in this task arise from both the low amount of training data for each new behavior and the high class imbalance, as seen in Figure 4.

For this task, we used our trained Task 1 baseline model as a starting point. Due to the small size of the training set, we found that models that did not account for class imbalance performed poorly. We therefore addressed class imbalance in our baseline model by replacing our original loss function with a weighted cross-entropy loss in which we scaled the weight of the under-represented class by the number of training frames for that class. Results for Task 3 are provided in Table 4. We found classifier performance to depend both on the percentage of frames during which a behavior was observed, and on the average duration of a behavior bout, with shorter bouts having lower classifier performance.

Method	Data Used During Training			Average F1	MAP
	Task 1 & 2 (train split)	Unlabeled Set	All Tasks (all splits)		
Baseline	✓			.754 ± .005	.813 ± .003
Baseline w/ task prog	✓	✓		.774 ± .006	.835 ± .005
MABe 2021 Task 2 Top-1	✓	✓	✓	.809 ± .015	.857 ± .007

Table 3: Class-averaged and annotator-averaged results on Task 2 (attack, investigation, mount; mean \pm standard deviation over 5 runs). The “All Tasks” column indicates that the model was jointly trained on all three tasks, and “all splits” indicates that both labeled train set and trajectory from unlabeled test set are used. See appendix for per class and per annotator results.

Method	Data Used During Training			Average F1	MAP
	Task 1 (train split)	Task 3 (train split)	Unlabeled Set		
Baseline	✓	✓		0.338 ± .004	.317 ± .005
Baseline w/ task prog	✓	✓	✓	.328 ± .009	.320 ± .009
MABe 2021 Task 3 Top-1		✓		.319 ± .025 (.363 ± .020)	.352 ± .023

Table 4: Class-averaged results on Task 3 over the 7 behaviors of interest (mean \pm standard deviation over 5 runs.) The average F1 score in brackets corresponds to improvements with threshold tuning. See appendix for per class results.

Task3 Top-1 Entry. The model for the top entry in Task 3 of the MABe Challenge was inspired by spatial-temporal graphs that have been used for skeleton-based action recognition algorithms in human pose datasets. In particular, MS-G3D [37] is an effective algorithm for extracting multi-scale spatial-temporal features and long-range dependencies. The MS-G3D model is composed of a stack of multiple spatial-temporal graph convolution blocks, followed by a global average pooling layer and a softmax classifier [37].

The spatial graph for Task 3 is constructed using the detected pose keypoints, with a connection added between the necks of the two mice. The inputs are normalized following [72]. MS-G3D is then trained in a fully supervised fashion on the train split of Task 3. Additionally, the model is trained with data augmentation based on rotation.

5 Discussion

We introduce CalMS21, a new dataset for detecting the actions of freely behaving mice engaged in naturalistic social interactions in a laboratory setting. The released data include over 70 hours of tracked poses from pairs of mice, and over 10 hours of manual, frame-by-frame annotation of animals’ actions. Our dataset provides a new way to benchmark the performance of multi-agent behavior classifiers. In addition to reducing human effort, automated behavior classification can lead to more objective, precise, and scalable measurements compared to manual annotation [2, 16]. Furthermore, techniques studied on our dataset can be potentially applied to other multi-agent datasets, such as those for sports analytics and autonomous vehicles.

In addition to the overall goal of supervised behavior classification, we emphasize two specific problems where we see a need for further investigation. The first of these is the utility of behavior classifiers for comparison of annotation style between different individuals or labs, most closely relating to our Task 2 on annotation style transfer. The ability to identify sources of inter-annotator disagreement is important for the reproducibility of behavioral results, and we hope that this dataset will foster further investigation into the variability of human-defined behavior annotations. A second problem of interest is the automated detection of new behaviors of interest given limited training data. This is especially important for the field of automated behavior analysis, as few-shot training of behavior classifiers would enable researchers to use supervised behavior classification as a tool to rapidly explore and curate large datasets of behavioral videos.

Alongside manually annotated trajectories provided for classifier training and testing, we include a large set of unlabeled trajectory data from 282 videos. The unlabeled dataset may be used to improve the performance of supervised classifiers, for example by learning self-supervised representations of trajectories [59], or it may be used on its own for the development of unsupervised methods for behavior discovery or trajectory forecasting. We note that trajectory forecasting is a task that is of interest to other fields studying multi-agent behavior, including self-driving cars and sports analytics. We hope that our dataset can provide an additional domain with which to test these models. In addition, unsupervised behavior analysis may be capable of identifying a greater number of behaviors than a human annotator would be able to annotate reliably. Recent single-animal work has shown that unsupervised pose analyses can enable the detection of subtle differences between experimental conditions [70]. A common problem in unsupervised analysis is evaluating the quality of the learned representation. Therefore, an important topic to be addressed in future work is the development of appropriate challenge tasks to evaluate the quality of unsupervised representations of animal movements and actions, beyond comparison with human-defined behaviors.

Broader Impact. In recent years, animal behavior analysis has emerged as a powerful tool in the fields of biology and neuroscience, enabling high-throughput behavioral screening in hundreds of hours of behavioral video [4]. Prior to the emergence of these tools, behavior analysis relied on manual frame-by-frame annotation of animals’ actions, a process which is costly, subjective, and arduous for domain experts. The increased throughput enabled by automation of behavior analysis has seen applications in neural circuit mapping [53, 6], computational drug development [70], evolution [25], ecology [16], and studies of diseases and disorders [40, 27, 70]. In releasing this dataset, our hope is to establish community benchmarks and metrics for the evaluation of new computational behavior analysis tools, particularly for social behaviors, which are particularly challenging to investigate due to their heterogeneity and complexity.

In addition to behavioral neuroscience, behavior modeling is of interest to diverse fields, including autonomous vehicles, healthcare, and video games. While behavior modeling can help accelerate scientific experiments and lead to useful applications, some applications of these models to human datasets, such as for profiling users or for conducting surveillance, may require more careful consideration. Ultimately, users of behavior models need to be aware of potentially negative societal impacts caused by their application.

Future Directions. In this dataset release, we have opted to emphasize behavior classification from keypoint-based animal pose estimates. However, it is possible that video data could further improve classifier performance. Since we have also released accompanying video data to a subset of CalMS21, an interesting future direction would be to determine the circumstances under which video data can improve behavior classification. Additionally, our dataset currently focuses on top-view tracked poses from a pair of interacting mice. For future iterations, including additional organisms, experimental settings, and task designs could further help benchmark the performance of behavior classification models. Finally, we value any input from the community on CalMS21 and you can reach us at mabe.workshop@gmail.com.

6 Acknowledgements

We would like to thank the researchers at the David Anderson Research Group at Caltech for this collaboration and the recording and annotation of the mouse behavior datasets, in particular, Tomomi Karigo, Mikaya Kim, Jung-sook Chang, Xiaolin Da, and Robert Robertson. We are grateful to the team at AICrowd for the support and hosting of our dataset challenge, as well as Northwestern University and Amazon Sagemaker for funding our challenge prizes. This work was generously supported by the Simons Collaboration on the Global Brain grant 543025 (to PP), NIH Award #K99MH117264 (to AK), NSF Award #1918839 (to YY), NSERC Award #PGSD3-532647-2019 (to JJS), as well as a gift from Charles and Lily Trimble (to PP).

References

- [1] Sungsoo Ahn, Shell Xu Hu, Andreas Damianou, Neil D Lawrence, and Zhenwen Dai. Variational information distillation for knowledge transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9163–9171, 2019.
- [2] David J Anderson and Pietro Perona. Toward a science of computational ethology. *Neuron*, 84(1):18–31, 2014.
- [3] Gordon J Berman, Daniel M Choi, William Bialek, and Joshua W Shaevitz. Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of The Royal Society Interface*, 11(99):20140672, 2014.
- [4] Kristin Branson, Alice A Robie, John Bender, Pietro Perona, and Michael H Dickinson. High-throughput ethomics in large groups of drosophila. *Nature methods*, 6(6):451–457, 2009.
- [5] Xavier P Burgos-Artizzu, Piotr Dollár, Dayu Lin, David J Anderson, and Pietro Perona. Social behavior recognition in continuous video. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1322–1329. IEEE, 2012.
- [6] Jessica Cande, Shigehiro Namiki, Jirui Qiu, Wyatt Korff, Gwyneth M Card, Joshua W Shaevitz, David L Stern, and Gordon J Berman. Optogenetic dissection of descending behavioral control in drosophila. *Elife*, 7:e34275, 2018.
- [7] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [8] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8748–8757, 2019.
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *ICML*, 2020.
- [10] Vasileios Choutas, Philippe Weinzaepfel, Jérôme Revaud, and Cordelia Schmid. Potion: Pose motion representation for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7024–7033, 2018.
- [11] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [12] Heiko Dankert, Liming Wang, Eric D Hoopfer, David J Anderson, and Pietro Perona. Automated monitoring and analysis of social behavior in drosophila. *Nature methods*, 6(4):297–303, 2009.
- [13] Sandeep Robert Datta, David J Anderson, Kristin Branson, Pietro Perona, and Andrew Leifer. Computational neuroethology: a call to action. *Neuron*, 104(1):11–24, 2019.
- [14] Tom Decroos, Jan Van Haaren, and Jesse Davis. Automatic discovery of tactics in spatio-temporal soccer match data. In *Proceedings of the 24th acm sigkdd international conference on knowledge discovery & data mining*, pages 223–232, 2018.
- [15] Aaron Defazio and Samy Jelassi. Adaptivity without compromise: A momentumized, adaptive, dual averaged gradient method for stochastic optimization. *arXiv preprint arXiv:2101.11075*, 2021.
- [16] Anthony I Dell, John A Bender, Kristin Branson, Iain D Couzin, Gonzalo G de Polavieja, Lucas PJJ Noldus, Alfonso Pérez-Escudero, Pietro Perona, Andrew D Straw, Martin Wikelski, et al. Automated image-based tracking and its application in ecology. *Trends in ecology & evolution*, 29(7):417–428, 2014.
- [17] SE Roian Egnor and Kristin Branson. Computational analysis of behavior. *Annual review of neuroscience*, 39:217–236, 2016.
- [18] Eyrún Eyjolfssdóttir, Kristin Branson, Yisong Yue, and Pietro Perona. Learning recurrent representations for hierarchical behavior modeling. *ICLR*, 2017.
- [19] Eyrún Eyjolfssdóttir, Steve Branson, Xavier P Burgos-Artizzu, Eric D Hoopfer, Jonathan Schor, David J Anderson, and Pietro Perona. Detecting social actions of fruit flies. In *European Conference on Computer Vision*, pages 772–787. Springer, 2014.
- [20] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. Datasheets for datasets. *arXiv preprint arXiv:1803.09010*, 2018.
- [21] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Machine learning basics. *Deep learning*, 1:98–164, 2016.
- [22] Jacob M Graving, Daniel Chae, Hemal Naik, Liang Li, Benjamin Koger, Blair R Costelloe, and Iain D Couzin. Deepposekit, a software toolkit for fast and robust animal pose estimation using deep learning. *Elife*, 8:e47994, 2019.

- [23] William H Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden Codel, Manuela Veloso, and Ruslan Salakhutdinov. Minerl: A large-scale dataset of minecraft demonstrations. *arXiv preprint arXiv:1907.13440*, 2019.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [25] Damián G Hernández, Catalina Rivera, Jessica Cande, Baohua Zhou, David L Stern, and Gordon J Berman. A framework for studying behavioral evolution by reconstructing ancestral repertoires. *arXiv preprint arXiv:2007.09689*, 2020.
- [26] Katja Hofmann. Minecraft as ai playground and laboratory. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, pages 1–1, 2019.
- [27] Weizhe Hong, Ann Kennedy, Xavier P Burgos-Artizzu, Moriel Zelikowsky, Santiago G Navonne, Pietro Perona, and David J Anderson. Automated measurement of mouse social behaviors using depth sensing, video tracking, and machine learning. *Proceedings of the National Academy of Sciences*, 112(38):E5351–E5360, 2015.
- [28] Alexander I Hsu and Eric A Yttri. B-soid: An open source unsupervised algorithm for discovery of spontaneous behaviors. *bioRxiv*, page 770271, 2020.
- [29] Mayank Kabra, Alice A Robie, Marta Rivera-Alba, Steven Branson, and Kristin Branson. Jaaba: interactive machine learning for automatic annotation of animal behavior. *Nature methods*, 10(1):64, 2013.
- [30] Tomomi Karigo, Ann Kennedy, Bin Yang, Mengyu Liu, Derek Tai, Iman A Wahle, and David J Anderson. Distinct hypothalamic control of same-and opposite-sex mounting behaviour in mice. *Nature*, 589(7841):258–263, 2021.
- [31] Gary King and Langche Zeng. Logistic regression in rare events data. *Political analysis*, 9(2):137–163, 2001.
- [32] Ugne Klibaite, Gordon J Berman, Jessica Cande, David L Stern, and Joshua W Shaevitz. An unsupervised method for quantifying the behavior of paired animals. *Physical biology*, 14(1):015006, 2017.
- [33] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 780–787, 2014.
- [34] Anna Kukleva, Hilde Kuehne, Fadime Sener, and Jurgen Gall. Unsupervised learning of action classes with continuous temporal embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12066–12074, 2019.
- [35] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 156–165, 2017.
- [36] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Zhaoyu Wang, Li Mian, Jing Zhang, and Jie Tang. Self-supervised learning: Generative or contrastive. *arXiv preprint arXiv:2006.08218*, 1(2), 2020.
- [37] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 143–152, 2020.
- [38] Diogo C Luvizon, David Picard, and Hedi Tabia. 2d/3d pose estimation and action recognition using multi-task deep learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5137–5146, 2018.
- [39] Kevin Luxem, Falko Fuhrmann, Johannes Kürsch, Stefan Remy, and Pavol Bauer. Identifying behavioral structure from deep variational embeddings of animal motion. *bioRxiv*, 2020.
- [40] Ana S Machado, Dana M Darmohray, Joao Fayad, Hugo G Marques, and Megan R Carey. A quantitative framework for whole-body coordination reveals specific deficits in freely walking ataxic mice. *Elife*, 4:e07892, 2015.
- [41] João C Marques, Simone Lackner, Rita Félix, and Michael B Orger. Structure of the zebrafish locomotor repertoire revealed with unsupervised behavioral clustering. *Current Biology*, 28(2):181–195, 2018.
- [42] Rafael Müller, Simon Kornblith, and Geoffrey Hinton. When does label smoothing help? *arXiv preprint arXiv:1906.02629*, 2019.
- [43] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016.
- [44] Simon RO Nilsson, Nastacia L Goodwin, Jia J Choong, Sophia Hwang, Hayden R Wright, Zane Norville, Xiaoyu Tong, Dayu Lin, Brandon S Bentzley, Neir Eshel, et al. Simple behavioral analysis (simba): an open source toolkit for computer classification of complex social behaviors in experimental animals. *BioRxiv*, 2020.

- [45] Lucas PJJ Noldus, Andrew J Spink, and Ruud AJ Tegelenbosch. Ethovision: a versatile video tracking system for automation of behavioral experiments. *Behavior Research Methods, Instruments, & Computers*, 33(3):398–414, 2001.
- [46] Shay Ohayon, Ofer Avni, Adam L Taylor, Pietro Perona, and SE Roian Egnor. Automated multi-day tracking of marked mice for the analysis of social behaviour. *Journal of neuroscience methods*, 219(1):10–19, 2013.
- [47] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [48] Malte Pedersen, Joakim Bruslund Haurum, Stefan Hein Bengtson, and Thomas B Moeslund. 3d-zef: A 3d zebrafish tracking benchmark dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2426–2436, 2020.
- [49] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [50] Talmo D Pereira, Joshua W Shaevitz, and Mala Murthy. Quantifying behavior to understand the brain. *Nature neuroscience*, pages 1–13, 2020.
- [51] Alfonso Pérez-Escudero, Julián Vicente-Page, Robert C Hinz, Sara Arganda, and Gonzalo G De Polavieja. idtracker: tracking individuals in a group by automatic identification of unmarked animals. *Nature methods*, 11(7):743–748, 2014.
- [52] Joelle Pineau, Philippe Vincent-Lamarre, Koustuv Sinha, Vincent Larivière, Alina Beygelzimer, Florence d’Alché Buc, Emily Fox, and Hugo Larochelle. Improving reproducibility in machine learning research (a report from the neurips 2019 reproducibility program). *arXiv preprint arXiv:2003.12206*, 2020.
- [53] Alice A Robie, Jonathan Hirokawa, Austin W Edwards, Lowell A Umayam, Allen Lee, Mary L Phillips, Gwyneth M Card, Wyatt Korff, Gerald M Rubin, Julie H Simpson, et al. Mapping the neural substrates of behavior. *Cell*, 170(2):393–406, 2017.
- [54] Mikayel Samvelyan, Tabish Rashid, Christian Schroeder De Witt, Gregory Farquhar, Nantas Nardelli, Tim GJ Rudner, Chia-Man Hung, Philip HS Torr, Jakob Foerster, and Shimon Whiteson. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043*, 2019.
- [55] Cristina Segalin, Jalani Williams, Tomomi Karigo, May Hui, Moriel Zelikowsky, Jennifer J Sun, Pietro Perona, David J Anderson, and Ann Kennedy. The mouse action recognition system (mars): a software pipeline for automated analysis of social behaviors in mice. *bioRxiv*, 2020.
- [56] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1010–1019, 2016.
- [57] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *arXiv preprint arXiv:1406.2199*, 2014.
- [58] Sebastian Stein and Stephen J McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 729–738, 2013.
- [59] Jennifer J Sun, Ann Kennedy, Eric Zhan, David J Anderson, Yisong Yue, and Pietro Perona. Task programming: Learning data efficient behavior representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2876–2885, 2021.
- [60] Jennifer J Sun, Jiaping Zhao, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, and Ting Liu. View-invariant probabilistic embedding for human pose. In *European Conference on Computer Vision*, pages 53–70. Springer, 2020.
- [61] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2446–2454, 2020.
- [62] Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. A survey on deep transfer learning. In *International conference on artificial neural networks*, pages 270–279. Springer, 2018.
- [63] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.
- [64] Joshua T Vogelstein, Youngser Park, Tomoko Ohshima, Rex A Kerr, James W Truman, Carey E Priebe, and Marta Zlatic. Discovery of brainwide neural-behavioral maps via multiscale unsupervised structure learning. *Science*, 344(6182):386–392, 2014.

- [65] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018.
- [66] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM Computing Surveys (CSUR)*, 53(3):1–34, 2020.
- [67] Joe H Ward Jr. Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, 58(301):236–244, 1963.
- [68] Benjamin Wild, David M Dormagen, Adrian Zachariae, Michael L Smith, Kirsten S Traynor, Dirk Brockmann, Iain D Couzin, and Tim Landgraf. Social networks predict the life and death of honey bees. *Nature communications*, 12(1):1–12, 2021.
- [69] Alexander B Wiltchko, Matthew J Johnson, Giuliano Iurilli, Ralph E Peterson, Jesse M Katon, Stan L Pashkovski, Victoria E Abaira, Ryan P Adams, and Sandeep Robert Datta. Mapping sub-second structure in mouse behavior. *Neuron*, 88(6):1121–1135, 2015.
- [70] Alexander B Wiltchko, Tatsuya Tsukahara, Ayman Zeine, Rockwell Anyoha, Winthrop F Gillis, Jeffrey E Markowitz, Ralph E Peterson, Jesse Katon, Matthew J Johnson, and Sandeep Robert Datta. Revealing the structure of pharmacobehavioral space through motion sequencing. *Nature Neuroscience*, 23(11):1433–1443, 2020.
- [71] Di Wu, Nabin Sharma, and Michael Blumenstein. Recent advances in video-based human action recognition using deep learning: A review. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 2865–2872. IEEE, 2017.
- [72] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [73] Yisong Yue, Patrick Lucey, Peter Carr, Alina Bialkowski, and Iain Matthews. Learning fine-grained spatial models for dynamic sports play prediction. In *2014 IEEE international conference on data mining*, pages 670–679. IEEE, 2014.
- [74] Eric Zhan, Albert Tseng, Yisong Yue, Adith Swaminathan, and Matthew Hausknecht. Learning calibratable policies using programmatic style-consistency. *ICML*, 2020.
- [75] Hong-Bo Zhang, Yi-Xiang Zhang, Bineng Zhong, Qing Lei, Lijie Yang, Ji-Xiang Du, and Duan-Sheng Chen. A comprehensive survey of vision-based human action recognition methods. *Sensors*, 19(5):1005, 2019.
- [76] Weiyu Zhang, Menglong Zhu, and Konstantinos G Derpanis. From actemes to action: A strongly-supervised representation for detailed action understanding. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2248–2255, 2013.
- [77] Yu Zhang and Qiang Yang. A survey on multi-task learning. *arXiv preprint arXiv:1707.08114*, 2017.

Appendix for CalMS21

The sections of our appendix are organized as follows:

- **Section A** contains dataset hosting and licensing information.
- **Section B** contains dataset documentation and intended uses for CalMS21, following the format of the Datasheet for Datasets[20].
- **Section C** describes the data format (.json).
- **Section D** describes how animal behavior data was recorded and processed.
- **Section E** shows the evaluation metrics for CalMS21, namely the F1 score and Average Precision.
- **Section F** contains additional implementation details of our models.
- **Section G** provides additional evaluation results.
- **Section H** addresses benchmark model reproducibility, following the format of the ML Reproducibility Checklist[52].

A CalMS21 Hosting and Licensing

The CalMS21 dataset is available at <https://sites.google.com/view/computational-behavior/our-datasets/calms21-dataset> (DOI: <https://doi.org/10.22002/D1.1991>), and is distributed under a Creative Commons Attribution/Non-Commercial/Share-Alike license (CC-BY-NC-SA).

CalMS21 is hosted via the Caltech Research Data Repository at [data.caltech.edu](https://data.caltech.edu/records/1991). This is a static dataset, meaning that any changes (such as new tasks, new experimental data, or improvements to pose estimates) will be released as a new entity; these updates will typically accompany new iterations of the MABe Challenge. News of any such updates will be posted both to the dataset website <https://sites.google.com/view/computational-behavior/our-datasets/calms21-dataset> and on the data repository page at <https://data.caltech.edu/records/1991>.

Code for all baseline models is available at <https://gitlab.aicrowd.com/aicrowd/research/mab-e/mab-e-baselines>, and is distributed under the MIT License. We as authors bear all responsibility in case of violation of rights.

Code for the top entry for Tasks 1 & 2 will be released by BW under the MIT License. The code for Task 3 will not be made publicly available.

B CalMS21 Documentation and Intended Uses

This section follows the format of the Datasheet for Datasets[20].

Motivation

For what purpose was the dataset created? Was there a specific task in mind? Was there a specific gap that needed to be filled? Please provide a description.

Automated animal pose estimation has become an increasingly popular tool in the neuroscience community, fueled by the publication of several easy-to-train animal pose estimation systems. Building on these pose estimation tools, pose-based approaches to supervised or unsupervised analysis of animal behavior are currently an area of active research. New computational approaches for automated behavior analysis are probing the detailed temporal structure of animal behavior, its relationship to the brain, and how both brain and behavior are altered in conditions such as Parkinson's, PTSD, Alzheimer's, and autism spectrum disorders. Due to a lack of publicly available animal behavior datasets, most new behavior analysis tools are evaluated on their own in-house data. There are no established community standards by which behavior analysis tools are evaluated, and it is unclear how well available software can be expected to perform in new conditions, particularly in cases where training data is limited. Labs looking to incorporate these tools in their experimental pipelines therefore often struggle to evaluate available analysis options, and can waste significant effort training and testing multiple systems without knowing what results to expect.

The Caltech Mouse Social 2021 (CalMS21) dataset is a new animal tracking, pose, and behavioral dataset, intended to a) serve as a benchmark dataset for comparison of behavior analysis tools, and establish community standards for evaluation of behavior classifier performance b) highlight critical challenges in computational behavior analysis, particularly pertaining to leveraging large, unlabeled datasets to improve performance on supervised classification tasks with limited training data, and c) foster interaction between behavioral neuroscientists and the greater machine learning community.

Who created this dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?

The CalMS21 dataset was created as a collaborative effort between the laboratories of David J Anderson, Yisong Yue, and Pietro Perona at Caltech, and Ann Kennedy at Northwestern. Videos of interacting mice were produced and manually annotated by Tomomi Karigo and other members of the Anderson lab. The video dataset was tracked, curated and screened for tracking quality by Ann Kennedy and Jennifer J. Sun, with pose estimation performed using version 1.7 of the Mouse Action Recognition System (MARS). The dataset tasks (Figure 8) were designed by Ann Kennedy and Jennifer J. Sun, with input from Pietro Perona and Yisong Yue.

Who funded the creation of the dataset? If there is an associated grant, please provide the name of the grantor and the grant name and number.

Acquisition of behavioral data was supported by NIH grants R01 MH085082 and R01 MH070053, Simons Collaboration on the Global Brain Foundation award no. 543025 (to DJA and PP), as well as a HFSP Long-Term Fellowship (to TK). Tracking, curation of videos, and task design was funded by NIMH award #R00MH117264 (to AK), NSF Award #1918839 (to YY), and NSERC Award #PGSD3-532647-2019 (to JJS).

Any other comments?

None.

Composition

What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)? Are there multiple types of instances (e.g., movies, users, and ratings; people and interactions between them; nodes and edges)? Please provide a description.

The core element of this dataset, called a *sequence*, captures the tracked postures and actions of two mice interacting in a standard resident-intruder assay filmed from above at 30Hz and manually annotated on a frame-by-frame basis for one or more behaviors. The resident in these assays is always a male mouse from strain C57Bl/6J, or from a transgenic line with C57Bl/6J background. The intruder is a male or female BALB/c mouse. Resident mice may be either group-housed or single-housed, and either socially/sexually naive or experienced (all factors that impact the types of social behaviors animals show in this assay.)

The core element of a *sequence* is called a *frame*; this refers to the posture of both animals on a particular frame of video, as well as one or more labels indicating the type of behavior being performed on that frame (if any).

The dataset is divided into four sub-sets: three collections of sequences associated with Tasks 1, 2, and 3 of the MABe Challenge, and a fourth "Unlabeled" collection of sequences that have only the *keypoint* elements with no accompanying *annotations* or *annotator-id* (see "What data does each instance consist of?" for explanation of these values.) Tasks 1-3 are split into train and test sets. Tasks 2 and 3 are also split by *annotator-id* (Task 2) or behavior (Task 3).

How many instances are there in total (of each type, if appropriate)?

Instances for each dataset are shown in table 5, divided into train and test sets. Number of instances is reported as both *frames* and *sequences*, where frames within a sequence are temporally contiguous and sampled at 30Hz (and hence not true statistically independent observations).

Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?

If the dataset is a sample, then what is the larger set? Is the sample representative of the larger set (e.g., geographic coverage)? If so, please describe how this representativeness was validated/verified. If it is not representative of the larger set, please describe why not (e.g., to cover a more diverse range of instances, because instances were withheld or unavailable).

Task	Category	Training set		Test set	
		Frames	Sequences	Frames	Sequences
Task 1	–	507,738	70	262,107	19
Task 2	Annotator 1	139,112	6	286,810	13
	Annotator 2	135,623	6	150,919	6
	Annotator 3	107,420	6	77,079	4
	Annotator 4	108,325	6	76,174	4
	Annotator 5	92,383	6	364,007	20
Task 3	Approach	20,624	3	126,468	25
	Disengaged	35,751	2	19,088	1
	Grooming	45,174	2	156,664	13
	Intromission	19,200	1	55,218	10
	Mount-attempt	46,847	4	85,836	12
	Sniff-face	19,244	3	251,793	47
	White rearing	36,181	2	17,939	1

Table 5: Training and test set instances counts for each task and category.

The assembled dataset presented here was manually curated from a large, unreleased repository of mouse behavior videos collected across several years by multiple members of the Anderson lab. Only videos of naturally occurring (not optogenetically or chemogenetically evoked) behavior were included. Selection criteria are described in the "Collection Process" section.

As a result of our selection criteria, the videos included in the Tasks 1-3 datasets may not be fully representative of mouse behavior in the resident-intruder assay: videos with minimal social interactions (when the resident ignored or avoided the intruder) were omitted in favor of including a greater number of examples of the annotated behaviors of interest.

What data does each instance consist of? "Raw" data (e.g., unprocessed text or images) or features? In either case, please provide a description.

Each sequence has three elements. 1) *Keypoints* are the locations of seven body parts (the nose, left and right ears, base of neck, left and right hips, and base of tail) on each of two interacting mice. Keypoints are estimated using the Mouse Action Recognition System (MARS). 2) *Annotations* are manual, frame-wise labels of an animal's actions, for example attack, mounting, and close investigation. Depending on the behaviors annotated, only between a few percent and up to half of frames will have an annotated action; frames that do not have an annotated action are labeled as *other*. The *other* label should not be taken to indicate that no behaviors are happening, and it should not be considered a true label category for purposes of classifier performance evaluation. 3) *Annotator-id* is a unique numeric ID indicating which (anonymized) human annotator produced the labels in *Annotations*. This ID is provided primarily for use in Task 2 of the MABe Challenge, which pertains to annotator style capture.

Note that this dataset does not include the original raw videos from which pose estimates were produced. This is because the objective of releasing this dataset was to determine the accuracy with which animal behavior could be detected using tracked keypoints alone.

Is there a label or target associated with each instance? If so, please provide a description.

In the Task 1, Task 2, and Task 3 datasets, the *annotation* field for a given behavior sequence consists of frame-wise labels of animal behaviors. Note that only a minority of frames have behavior labels; remaining frames are labeled as *other*. Only a small number of behaviors were tracked by human annotators (most typically *attack*, *mount*, and *close investigation*), therefore frames labeled as *other* are not a homogeneous category, but may contain diverse other behaviors.

The "Unlabeled" collection of sequences has no labels, and instead contains only keypoint tracking data.

Is any information missing from individual instances? If so, please provide a description, explaining why this information is missing (e.g., because it was unavailable). This does not include intentionally removed information, but might include, e.g., redacted text.

There is no missing data (beyond what was intentionally omitted, eg in the Unlabeled category.)

Are relationships between individual instances made explicit (e.g., users' movie ratings, social network links)? If so, please describe how these relationships are made explicit.

Each instance (*sequence*) is to be treated as an independent observation with no relationship to other instances in the dataset. In almost all cases, the identities of the interacting animals are unique to each sequence, and this information is not tracked in the dataset.

Are there recommended data splits (e.g., training, development/validation, testing)? If so, please provide a description of these splits, explaining the rationale behind them.

The dataset includes a recommended train/test split for Tasks 1, 2, and 3. In Tasks 2 and 3, the split was designed to provide a roughly consistent, small amount of training data for each sub-task. In Task 1, the split was manually selected so that the test set included sequences from a range of experimental conditions and dates.

Are there any errors, sources of noise, or redundancies in the dataset? If so, please provide a description.

Pose keypoints in this dataset are produced using automated pose estimation software (the Mouse Action Recognition System, MARS). While the entire dataset was manually screened to remove sequences with poor pose estimation, some errors in pose estimation and noise in keypoint placement still occur. These are most common on frames when the two animals are in close contact or moving very quickly.

In addition, manual annotations of animal behavior are inherently subjective, and individual annotators show some variability in the precise frame-by-frame labeling of behavior sequences. An investigation of within- and between-annotator variability is included in the MARS pre-print.

Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)? If it links to or relies on external resources, a) are there guarantees that they will exist, and remain constant, over time; b) are there official archival versions of the complete dataset (i.e., including the external resources as they existed at the time the dataset was created); c) are there any restrictions (e.g., licenses, fees) associated with any of the external resources that might apply to a future user? Please provide descriptions of all external resources and any restrictions associated with them, as well as links or other access points, as appropriate.

The dataset is self-contained.

Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor-patient confidentiality, data that includes the content of individuals non-public communications)? If so, please provide a description.

No.

Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety? If so, please describe why.

No such material; dataset contains only tracked posture keypoints (no video or images) and text labels pertaining to mouse social behaviors.

Does the dataset relate to people? If not, you may skip the remaining questions in this section.

No.

Does the dataset identify any subpopulations (e.g., by age, gender)? If so, please describe how these subpopulations are identified and provide a description of their respective distributions within the dataset.

n/a

Is it possible to identify individuals (i.e., one or more natural persons), either directly or indirectly (i.e., in combination with other data) from the dataset? If so, please describe how.

n/a

Does the dataset contain data that might be considered sensitive in any way (e.g., data that reveals racial or ethnic origins, sexual orientations, religious beliefs, political opinions or union memberships, or locations; financial or health data; biometric or genetic data; forms of government identification, such as social security numbers; criminal history)? If so, please provide a description.

n/a

Any other comments?

A subset of videos in Task 1 and the Unlabeled dataset are from animals that have been implanted with a head-mounted microendoscope or optical fiber (for fiber photometry.) Because the objective of this dataset is to learn to recognize behavior in a manner that is invariant to experimental setting, the precise preparation of the resident and intruder mice (including age, sex, past experiences, and presence of neural recording devices) is not provided in the dataset.

Collection Process

How was the data associated with each instance acquired? Was the data directly observable (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indirectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for age or language)? If data was reported by subjects or indirectly inferred/derived from other data, was the data validated/verified? If so, please describe how.

Sequences in the dataset are derived from video of pairs of socially interacting mice engaged in a standard resident-intruder assay. In this assay, a black (C57Bl/6J) male "resident" mouse is filmed in its home cage, and a white (BALB/c) male or female "intruder" mouse is manually introduced to the cage by an experimenter. The animals are then allowed to freely interact for between 1-2 and 10 minutes. If there is excessive fighting (injury to either animal) the assay is stopped and that trial is discarded. Resident mice typically undergo several (3-6) resident-intruder assays per day with different intruder animals.

Poses of both mice were estimated from top-view video using MARS, and pose sequences were cropped to only include frames where both animals were present in the arena. Manual, frame-by-frame annotation of animals' actions were performed from top- and front-view video by trained experts.

What mechanisms or procedures were used to collect the data (e.g., hardware apparatus or sensor, manual human curation, software program, software API)? How were these mechanisms or procedures validated?

Video of the resident-intruder assay was captured at 30Hz using top- and front-view cameras (Point Grey Grasshopper3) recorded at 1024x570 (top) and 1280x500 (front) pixel resolution. Manual annotation was performed using custom software (either the Caltech Behavior Annotator ([link](#)) or Bento ([link](#))) by trained human experts. All annotations were visually screened to ensure that the full sequence was annotated.

If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?

The Task 1 dataset was chosen to match the training and test sets of behavior classifiers of MARS. These training and test sets, in turn, were sampled from among unpublished videos collected and annotated by a member of the Anderson lab. Selection criteria for inclusion were high annotation quality (as estimated by the individual who annotated the data) and annotation completeness; videos with diverse social behaviors (mounting and attack in addition to investigation) were favored. The Tasks 2 and 3 datasets were manually selected from among previously collected (unpublished) datasets, where selection criteria were for high annotation quality, annotation completeness, and sufficient number of behavior annotations. The Unlabeled dataset consists of videos from a subset of experiments in a recent publication[30]. The subset of experiments included in this dataset was chosen at random.

Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?

Behavioral data collection and annotation was performed by graduate student, postdoc, and technician members of the Anderson lab, as a part of other ongoing research projects in the lab. (No videos or annotations were explicitly generated for this dataset release.) Lab members are full-time employees of Caltech or HHMI, or are funded through independent graduate or postdoctoral fellowships, and their compensation was not dependent on their participation in this project.

Over what timeframe was the data collected? Does this timeframe match the creation timeframe of the data associated with the instances (e.g., recent crawl of old news articles)? If not, please describe the timeframe in which the data associated with the instances was created.

Data associated with this dataset was created and annotated between 2016 and 2020, with annotation typically occurring within a few weeks of creation. Pose estimation was performed later, with most videos processed in 2019-2020. This dataset was assembled from December 2020 - February 2021.

Were any ethical review processes conducted (e.g., by an institutional review board)? If so, please provide a description of these review processes, including the outcomes, as well as a link or other access point to any supporting documentation.

All experiments included here were performed in accordance with NIH guidelines and approved by the Institutional Animal Care and Use Committee (IACUC) and Institutional Biosafety Committee at Caltech. Review of experimental design by the IACUC follows the steps outlined in the NIH-published Guide for the Care and Use of Laboratory Animals. All individuals performing behavioral experiments underwent animal safety training prior to data collection. Animals were maintained under close veterinary supervision, and resident-intruder assays were monitored in real time and immediately interrupted should either animal become injured during aggressive interactions.

Does the dataset relate to people? If not, you may skip the remaining questions in this section.

No.

Did you collect the data from the individuals in question directly, or obtain it via third parties or other sources (e.g., websites)?

n/a

Were the individuals in question notified about the data collection? If so, please describe (or show with screenshots or other information) how notice was provided, and provide a link or other access point to, or otherwise reproduce, the exact language of the notification itself.

n/a

Did the individuals in question consent to the collection and use of their data? If so, please describe (or show with screenshots or other information) how consent was requested and provided, and provide a link or other access point to, or otherwise reproduce, the exact language to which the individuals consented.

n/a

If consent was obtained, were the consenting individuals provided with a mechanism to revoke their consent in the future or for certain uses? If so, please provide a description, as well as a link or other access point to the mechanism (if appropriate).

n/a

Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a data protection impact analysis) been conducted? If so, please provide a description of this analysis, including the outcomes, as well as a link or other access point to any supporting documentation.

n/a

Any other comments?

None.

Preprocessing/cleaning/labeling

Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)? If so, please provide a description. If not, you may skip the remainder of the questions in this section.

No preprocessing was performed on the *sequence* data released in this dataset.

Was the "raw" data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)? If so, please provide a link or other access point to the "raw" data.

n/a

Is the software used to preprocess/clean/label the instances available? If so, please provide a link or other access point.

n/a

Any other comments?

None.

Uses

Has the dataset been used for any tasks already? If so, please provide a description.

Yes: this dataset was released to accompany the three tasks of the 2021 Multi-Agent Behavior (MABe) Challenge, posted here. The challenge tasks are summarized as follows:

- **Task 1, Classical Classification:** train supervised classifiers to detect instances of close investigation, mounting, and attack from labeled examples. All behaviors were annotated by the same individual.
- **Task 2, Annotation Style Transfer:** given limited training examples, train classifiers to reproduce the annotation style of five additional annotators for close investigation, mounting, and attack behaviors.
- **Task 3, Learning New Behavior:** given limited training examples, train classifiers to detect instances of seven additional behaviors (names of these behaviors were anonymized for this task.)

Is there a repository that links to any or all papers or systems that use the dataset? If so, please provide a link or other access point.

Papers that use or cite this dataset may be submitted by their authors for display on the CalMS21 website at <https://sites.google.com/view/computational-behavior/our-datasets/calms21-dataset>

What (other) tasks could the dataset be used for?

In addition to MABe Challenge Tasks 1-3, which can be studied with supervised learning, transfer learning, or few-shot learning techniques, the animal trajectories in this dataset could be used for unsupervised behavior analysis, representation learning, or imitation learning.

Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses? For example, is there anything that a future user might need to know to avoid uses that could result in unfair treatment of individuals or groups (e.g., stereotyping, quality of service issues) or other undesirable harms (e.g., financial harms, legal risks) If so, please provide a description. Is there anything a future user could do to mitigate these undesirable harms?

At time of writing there is no precise, numerical consensus definition of the mouse behaviors annotated in this dataset (and in fact even different individuals trained in the same research lab and following the same written descriptions of behavior can vary in how they define particular actions such as attack, as is evidenced in Task 2.) Future users should be aware of this limitation, and bear in mind that behavior annotations in this dataset may not always agree with the behavior annotations produced by other individuals or labs.

Are there tasks for which the dataset should not be used? If so, please provide a description.

None.

Any other comments?

None.

Distribution

Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created? If so, please provide a description.

Yes- the dataset is publicly available for download by all interested third parties.

How will the dataset will be distributed (e.g., tarball on website, API, GitHub) Does the dataset have a digital object identifier (DOI)?

The dataset is available on the Caltech public data repository at <https://data.caltech.edu/records/1991>, where it will be retained indefinitely and available for download by all third parties. The data.caltech.edu posting has accompanying DOI <https://doi.org/10.22002/D1.1991>.

The dataset as used for the MABe Challenge (anonymized sequence and behavior ids) is available for download on the AICrowd page, located at (link).

When will the dataset be distributed?

The full dataset was made publicly available on data.caltech on June 6th, 2021.

Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)? If so, please describe this license and/or ToU, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms or ToU, as well as any fees associated with these restrictions.

The CalMS21 dataset is distributed under the CreativeCommons Attribution-NonCommercial-ShareAlike license (CC-BY-NC-SA). The terms of this license may be found at <https://creativecommons.org/licenses/by-nc-sa/2.0/legalcode>.

Have any third parties imposed IP-based or other restrictions on the data associated with the instances? If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms, as well as any fees associated with these restrictions.

There are no third party restrictions on the data.

Do any export controls or other regulatory restrictions apply to the dataset or to individual instances? If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any supporting documentation.

No export controls or regulatory restrictions apply.

Any other comments?

None.

Maintenance

Who will be supporting/hosting/maintaining the dataset?

The dataset is hosted on the Caltech Research Data Repository at data.caltech.edu. Dataset hosting is maintained by the library of the California Institute of Technology. Long-term support for users of the dataset is provided by Jennifer J. Sun and by the laboratory of Ann Kennedy.

How can the owner/curator/manager of the dataset be contacted (e.g., email address)?

The managers of the dataset (JJS and AK) can be contacted at mabe.workshop@gmail.com, or AK can be contacted at ann.kennedy@northwestern.edu and JJS can be contacted at jjsun@caltech.edu.

Is there an erratum? If so, please provide a link or other access point.

No.

Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)? If so, please describe how often, by whom, and how updates will be communicated to users (e.g., mailing list, GitHub)?

Users of the dataset have the option to subscribe to a mailing list to receive updates regarding corrections or extensions of the CalMS21 dataset. Mailing list sign-up can be found on the CalMS21 webpage at <https://sites.google.com/view/computational-behavior/our-datasets/calms21-dataset>.

Updates to correct errors in the dataset will be made promptly, and announced via update messages posted to the CalMS21 website and data.caltech.edu page.

Updates that extend the scope of the dataset, such as additional data sequences, new challenge tasks, or improved pose estimation, will be released as new named instantiations on at most a yearly basis. Previous versions of

the dataset will remain online, but obsolescence notes will be sent out to the CalMS21 mailing list. In updates, dataset version will be indicated by the year in the dataset name (here 21). Dataset updates may accompany new instantiations of the MABe Challenge.

If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were individuals in question told that their data would be retained for a fixed period of time and then deleted)? If so, please describe these limits and explain how they will be enforced.

N/a (no human data.)

Will older versions of the dataset continue to be supported/hosted/maintained? If so, please describe how. If not, please describe how its obsolescence will be communicated to users.

Yes, the dataset will be permanently available on the Caltech Research Data Repository (data.caltech.edu), which is managed by the Caltech Library.

If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so? If so, please provide a description. Will these contributions be validated/verified? If so, please describe how. If not, why not? Is there a process for communicating/distributing these contributions to other users? If so, please provide a description.

Extensions to the dataset will take place through at-most-yearly updates. We welcome community contributions of behavioral data, novel tracking methods, and novel challenge tasks; these may be submitted by contacting the authors or emailing mabe.workshop@gmail.com. All community contributions will be visually reviewed by the managers of the dataset for quality of tracking and annotation data; for new challenge tasks, new baseline models will be developed prior to launch to ensure task feasibility. Community contributions will not be accepted without a data maintenance plan (similar to this document), to ensure support for future users of the dataset.

Any other comments?

If you enjoyed this dataset and would like to contribute other multi-agent behavioral data for future versions of the dataset or MABe Challenge, contact us at mabe.workshop@gmail.com!

C Data Format

Our dataset is released in the json format. Each sequence (video) has associated keypoints, keypoint confidence scores, and behavior annotations, all stored as lists in a dictionary, as well as a dictionary of associated metadata, which is nested within the sequence dictionary (see sample below). The unlabeled set is an exception, as it only contains keypoints and scores, with no annotations or metadata fields. For each task, there is one train file and one test file. The train file is used during development and a held out validation set can be used for hyperparameter tuning. The results are reported on the test file.

For all CalMS21 data, the json format is shown in Listing 1. Note that the number of frames are the number of frames of each video, so this number could vary across sequences.

The layer GROUPNAME groups together sequences with a similar property, such as a common annotator id. In Task 1, GROUPNAME is annotator_id-0, and there is only one GROUP in the file. Task 2, GROUPNAME is annotator_id-X, and there are five groups for $X \in (1, 2, 3, 4, 5)$. In Task 3, GROUPNAME the name of a behavior.

The keypoints field contains the (x,y) position of anatomically defined pose keypoints tracked using MARS[55]. The dimensions ($2 \times 2 \times 7$) correspond to the mouse ID (mouse 0 is the resident and mouse 1 is the intruder), image (x,y) coordinates in pixels, and keypoint ID. For keypoint ID, there are seven tracked body parts, ordered (nose, left ear, right ear, neck, left hip, right hip, tail base).

The scores field corresponds to the confidence from the MARS tracker [55], and its dimensions in each frame (2×7) corresponds to the mouse ID and keypoint ID.

The annotations field contains the frame-level behavior annotations from domain experts as a list of integers.

```

1  {
2      "<GROUPNAME>": {
3          "<sequence_id-1>": {
4              "keypoints" : a list of shape (frames, 2, 2, 7),
5              "scores" : a list of shape (frames, 2, 7),
6              "annotations" : a list of shape (frames),
7              "metadata" : {
8                  "annotator_id": a number identifying the annotator
9                  "vocab": a dictionary of behavior names
10             }
11         },
12         "<sequence_id-2>": {
13             "keypoints" : a list of shape (frames, 2, 2, 7),
14             "scores" : a list of shape (frames, 2, 7),
15             "annotations" : a list of shape (frames),
16             "metadata" : {
17                 "annotator_id": a number identifying the annotator
18                 "vocab": a dictionary of behavior names
19             },
20             ...
21         },
22         ...
23     }

```

Listing 1: Json file format.

The metadata dictionary for all tasks (except the unlabeled data) contains two fields, an integer `annotator_id` and a dictionary `vocab` which gives the mapping from behavior classes to integer values for the annotations list. For example, in Task 1 `vocab` is `attack: 0, investigation: 1, mount: 2, other: 3`.

The dataset website <https://sites.google.com/view/computational-behavior/our-datasets/calms21-dataset> also contains a description of the data format and code to load the data for each task.

D Dataset Preparation

D.1 Behavior Video Acquisition

This section is adapted from [55]. Experimental mice ("residents") were transported in their homecage to a behavioral testing room, and acclimatized for 5-15 minutes. Homecages were then inserted into a custom-built hardware setup[27] where behaviors are recorded under dim red light condition using a camera (Point Grey Grasshopper3) located 46cm above the homecage floor. Videos are acquired at 30 fps and 1024x570 pixel resolution using StreamPix video software (NorPix). Following two further minutes of acclimatization, an unfamiliar group-housed male or female BALB/c mouse ("intruder") was introduced to the cage, and animals were allowed to freely interact for a period of approximately 10 minutes. BALB/c mice are used as intruders for their white coat color (simplifying identity tracking), as well as their relatively submissive behavior, which reduces the likelihood of intruder-initiated aggression.

D.2 Behavior Annotation

Behaviors were annotated on a frame-by-frame basis by a trained human expert. Annotators were provided with simultaneous top- and front-view video of interacting mice, and scored every video frame for close investigation, attack, and mounting, defined as follows (reproduced from [55]):

1. **Close investigation:** resident (black) mouse is in close contact with the intruder (white) and is actively sniffing the intruder anywhere on its body or tail. Active sniffing can usually be distinguished from passive orienting behavior by head bobbing/movements of the resident's nose.
2. **Attack:** high-intensity behavior in which the resident is biting or tussling with the intruder, including periods between bouts of biting/tussling during which the intruder is jumping or running away and the resident is in close pursuit. Pauses during which resident/intruder are facing each other (typically while rearing) but not actively interacting should not be included.
3. **Mount:** behavior in which the resident is hunched over the intruder, typically from the rear, and grasping the sides of the intruder using forelimbs (easier to see on the Front camera). Early-stage copulation is accompanied by rapid pelvic thrusting, while later-stage copulation (sometimes annotated separately as intromission) has a slower rate of pelvic thrusting with some pausing: for the purpose of this analysis, both behaviors should be counted as mounting, however periods where the resident is climbing on the intruder but not attempting to grasp the intruder or initiate thrusting should not. While most bouts of mounting are female-directed, occasional shorter mounting bouts are observed towards males; this behavior and its neural correlates are described in [30].

Annotation was performed either in BENTO[55] or using a custom Matlab interface. In most videos, the majority of frames will not include one of these three behaviors (see Table 6): in these instances, animals may be apart from each other exploring other parts of the arena, or may be close together but not actively interacting. These frames are labeled as "other". Because this is not a true behavior, we do not consider classifier performance in predicting "other" frames accurately.

D.3 Pose Estimation

The poses of mice in top-view recordings are estimated using the Mouse Action Recognition System (MARS,[55]), a computer vision tool that identifies seven anatomically defined keypoints on the body of each mouse: the nose, ears, base of neck, hips, and tail (Figure 3). MARS estimates animal pose using a stacked hourglass model [43] trained on a dataset of 15,000 video frames, in which all seven keypoints were manually annotated on each of two interacting mice (annotators were instructed to estimate the locations of occluded keypoints.) To improve accuracy, each image in the training set was annotated by five human workers, and "ground truth" keypoint locations were taken to be the median of the five annotators' estimates of each point. All videos in the CalMS21 Dataset were collected in the same experimental apparatus as the MARS training set [27].

E Evaluation

For all three Tasks, we evaluate performance of trained classifiers in terms of the F1 score and Average Precision for each behavior of interest. Because of the high class imbalance in behavior annotations, we use an unweighted average across behavior classes to compute a single F1 score and Mean Average Precision (MAP) for a given model, omitting the "other" category (observed when a frame is not positive for any annotated behavior) from our metrics.

Behavior	Percent of Frames
attack	2.76
investigation	28.9
mount	5.64
other	62.7

Table 6: The percentage of frames labeled as attack, investigation, mount, and other in the Task 1 training set.

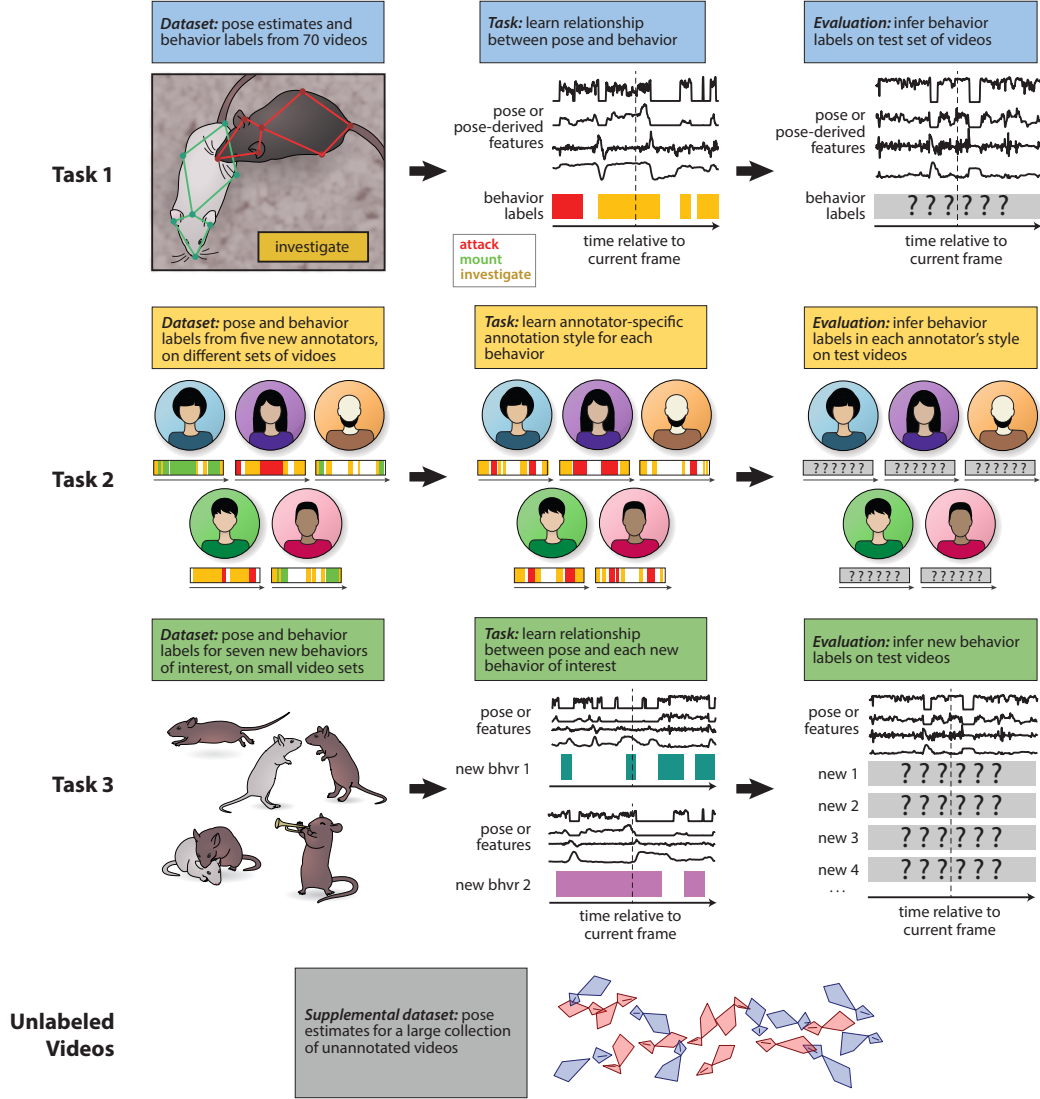


Figure 8: **Summary of Tasks.** Visual summary of datasets, tasks, and evaluations for the three tasks defined in CalMS21.

F1 score. The F1 score is the harmonic mean of the Precision P and Recall R :

$$P = \frac{TP}{TP + FP} \quad (1)$$

$$R = \frac{TP}{TP + FN} \quad (2)$$

$$F1 = \frac{2 \times P \times R}{P + R} \quad (3)$$

Where true positives (TP) is the number of frames that a model correctly labels as positive for a class, false positives (FP) is the number of frames incorrectly labeled as positive for a class, and false negatives (FN) is the number of frames incorrectly labeled as negative for a class.

The F1 score is a useful measure of model performance when the number of true negatives (TN, frames correctly labeled as negative for a class) in a task is high. This is the case for the CalMS21 dataset, where for instance attack occurs on less than 3% of frames.

Average Precision. The AP approximates the area under the Precision-Recall curve for each behavior class. There are a few different ways to approximate AP; here we compute AP using the implementation from Scikit-Learn [49]:

$$AP = \sum_n P_n(R_n - R_{n-1}) \quad (4)$$

where P_n and R_n are the precision and recall at the n -th threshold. This implementation is not interpolated. We call the unweighted class-averaged AP the mean average precision (MAP).

Averaging Across Behaviors and Annotators. Our baseline code, released at <https://gitlab.aicrowd.com/aicrowd/research/mab-e/mab-e-baselines> shows how we computed our metrics. To compare against our benchmarks, the F1 score and MAP should be computed as follows: for Task 1, the metrics should be computed on the entire test set with all the videos concatenated into one sequence (each frame is weighted equally for each behavior); for Task 2, the metrics should be computed separately on the test set of each annotator, then averaged over the annotators (each behavior and annotator is weighted equally); for Task 3, the metrics should be computed separately for the test set of each behavior, then averaged over the behaviors (each behavior is weighted equally). For our evaluation, the class with the highest predicted probability in each frame was used to compute F1 score, but the F1 score will likely be higher with threshold tuning.

F Implementation Details

For more details on the implementation and exact hyperparameter settings, see our code links in Section A.

F.1 Baseline Model Input

Each frame in the CalMS21 dataset is represented by a flattened vector of 28 values, representing the (x,y) location of 7 keypoints from each mouse (resident and intruder). For our baselines, we normalized all (x,y) coordinates by the resolution of the video (1024×570 pixels). Associated with these (x,y) values is a single behavior label per frame: in Tasks 1: Classic Classification and Task 2: Annotation Style Transfer, labels may be “attack”, “mount”, “investigation”, or “other” (i.e. none of the above), while in Task 3: New Behaviors, we provide a separate set of binary labels for each behavior of interest.

We do not require behavior classification models to be causal, so information from both past and future frames can be used for behavior classification. Thus, in the most general form, the input to our model is a stacked set of keypoints from the immediate past, the present frame, and the immediate future, and the model is trained to predict the behavior label only for the present frame (Figure 5). We refer to our input stack of poses across frames as an input *trajectory*, where the number of past and future frames included in the input trajectory is a model hyperparameter.

Neighboring frames in input trajectories are highly correlated. Therefore, to sample a broad temporal window without significantly increasing the dimensionality of model input, we introduced a skip factor as a second hyperparameter. For a given input trajectory, a skip factor of N signifies that only every N^{th} frame is used when sampling past/future frames. Given current frame t , sampling 50 future frames with a skip factor of 2 would produce a stack of frames $\{t, t+2, t+4, \dots, t+(2 \times 50)\}$. We note that more sophisticated compression methods, such as non-uniform downsampling or interpolation, could lead to better representations of temporal data.

We explored model performance as a function of these hyperparameters. For Task 1, we found that models generally performed well when including 100 past frames and 100 future frames, with a skip factor of 2—ie, model input was every other frame from 200 frames (6.667 seconds) before the current frame to 200 frames after the current frame. For Tasks 2 and 3, we included 50 past frames and 50 future frames with a skip factor of 1.

Behavior	% Frames	Weight
approach	3.25	20
disengaged	1.47	50
groom	14.0	5
intromission	19.3	3
mount attempt	0.90	100
sniff face	4.74	20
whiterearing	9.29	10

Table 7: Weighting applied on each class for Task 3.

F.2 Baseline Data Augmentation

Behavior labels should be invariant to certain features of the raw pose data, such as the absolute positions of the agents. We therefore augmented our training data using trajectory transformations, including random rotations, translations, and reflections of the pair of mice. To preserve temporal and relative spatial structure, the same transformation was applied to all frames in a given input trajectory, and to all keypoints from both mice on each frame.

This data augmentation method did not significantly improve model performance for Task 1, although it did improve performance on the more data-limited Tasks 2 and 3. It is possible that a more thorough form of data augmentation, incorporating additional domain-specific knowledge of animal behavior, could further improve model performance. Alternatively, performance on behavior classification tasks could be improved by using domain knowledge to remove non-informative sources of variance, for example by transforming animal trajectories from allocentric to egocentric coordinates.

F.3 Baseline Task 2 Annotation Style Transfer Details

The Task 2 baseline model is fine-tuned on the trained model from Task 1. We found that allowing all weights of our pre-trained model to be modified during fine-tuning often resulted in overfitting of the training set, causing a drop in model performance relative to the pre-trained model. This likely arises in part due to class imbalance, which has a more substantial effect on model performance when the training set is small.

To overcome this, we trained the fine-tuned network in phases. In the first phase of fine-tuning, we froze weights in all layers except for the final, densely connected output layer, and trained for a set number of epochs. The number of initial fine-tuning epochs is another model hyperparameter which we tuned for each annotator. Following this initial period of fine-tuning, we unfroze weights in all layers and train until convergence on a validation set. We then used this set of final hyperparameters and trained on the full Task 2 training set.

F.4 Baseline Task 3 New Behaviors Details

As in Task 2, we fine-tuned our model for each behavior in two steps: first freezing all but the output layer for a fixed number of epochs, and then unfreezing all weights and allowing training to run to completion; the learning rate was hand-tuned on a per-behavior basis based on a held-out validation set (20% of the train split). Once the hyperparameters are chosen, we then trained the final models on the full train split.

To better address the extreme class imbalance in Task 3, we also applied a weight on different behavior samples during training (Table 7).

F.5 Tasks 1 & 2 Top-1 Entry Details

The top-1 entry (Figure 9) for tasks 1 & 2 used joint training on all three tasks, including the unlabeled set. The total input dimensionality at each frame is 52. First, 37 features in total were computed based on velocity, egocentric angles, and distances on top of the trajectory data. In addition, 15 PCA dimensions were computed based on distances between keypoints at each frame.

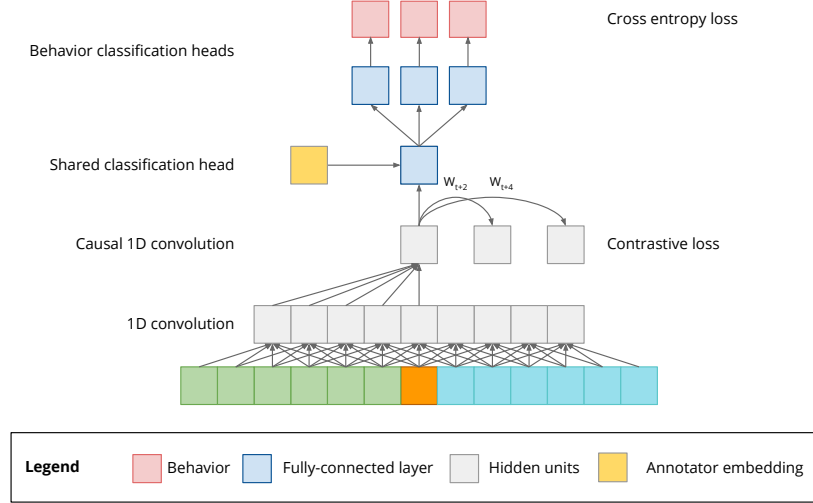


Figure 9: **Challenge Top-1 Model on Tasks 1 & 2.**

For this model, the MADGRAD optimizer [15] with cosine learning rate annealing and a small weight decay of $1e-5$ was used to fit the model. Label smoothing was applied to the labels to regularize the model and improve calibration [42]. Class rescaling weights were used in the categorical cross-entropy loss function because of the unbalanced distribution of behaviors in the data [31].

Cross-validation was used to optimize the hyperparameters of the model manually. Due to the long training time of the model and the large variance in the validation F1 scores on different splits of the dataset, an exhaustive search in the hyperparameter space was not possible, and most parameters were not tuned at all. Batch size and the number of training epochs were evaluated in the ranges of 16 to 64 and 30 to 60 and set to 32 and 40 for the final model. Dropout at various positions in the network architecture was tested but removed for the final model because it did not improve validation performance. The number of embedder residual blocks was tuned because there is likely a tradeoff between the number of future frames the network can use and the effectiveness of the unsupervised loss. Values in the range from 3 to 16 were evaluated and set to 8 for the final model. The learning rate was evaluated for the values 0.1, 0.01, 0.005, 0.001, 0.0005, 0.0001 and set to 0.001 for the final model. Due to computational constraints, the hyperparameters related to the unsupervised loss (e.g., number of linear projections/timestep into the future) were not tuned.

For the MABe challenge, an ensemble of models was trained on cross-validation splits of the data. These models were then used to bootstrap soft labels for the unlabeled sequences by computing their average predictions. These labels were then used during training of the model, i.e., an ensemble of the model was used to bootstrap training data for the next iteration of the model. This approach was used to further regularize the model by utilizing the implicit knowledge in an ensemble of models trained on the full dataset. For evaluating reproducibility, only one model from the ensemble was selected and trained across 5 runs.

F.6 Task 3 Top-1 Entry Details

For Task 3, the top entry in the MABe challenge trained seven binary classifiers in a fully supervised fashion using the MS-G3D Architecture [37] (Figure 10). The model starts from keypoints at each frame, as 2D coordinates (X,Y) in pixel coordinate system. The clips are then represented with tensor of (B, 2, T, 14, 1) dimensions, where B and T respectively denotes batch size and sequence length. All models are trained in 10 epochs with SGD with momentum 0.9, batch size 256 (128 per GPU over 2 GPUs), an initial learning rate 0.2 for 6 and 8 epochs with LR weight decay with a factor 0.0003. All skeleton sequences are sampled by 64 frames with stride 2 to have $T = 32$. Inputs are preprocessed with normalization following [72]. Random keypoint rotations are used for data augmentation.

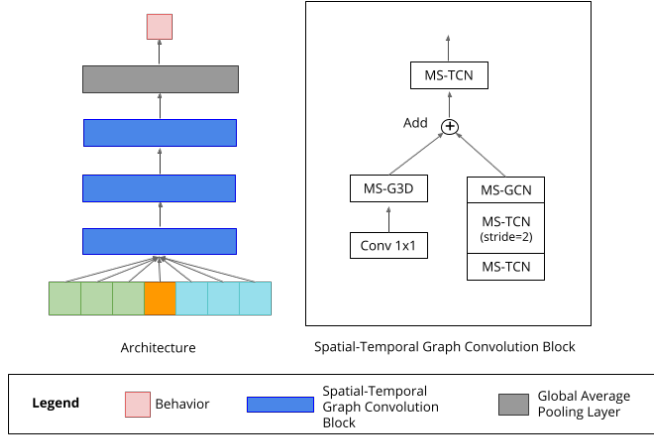


Figure 10: **Challenge Top-1 Model on Task 3.** This model follows the MS-G3D Architecture [37]. “TCN”, “GCN”, “MS-” respectively denotes temporal and graph convolution blocks (STGC), and multi-scale aggregation.

Method	Attack		Investigation		Mount	
	F1	AP	F1	AP	F1	AP
Baseline	.664 ± .031	.724 ± .023	.814 ± .005	.893 ± .005	.900 ± .004	.950 ± .004
Baseline w/ task prog	.789 ± .002	.839 ± .011	.817 ± .003	.889 ± .006	.880 ± .011	.939 ± .009
MABe 2021 Task 1 Top-1	.827 ± .024	.885 ± .012	.852 ± .011	.908 ± .014	.913 ± .019	.950 ± .014

Table 8: Per-class results on Task 1 (mean ± standard deviation over 5 runs.)

During development, since the train split of Task 3 is small, the train split of Task 1 was used during development for validation (a random 20% of the data was held-out). Here, different batch sizes (32/64/128/256), learning rates (0.025/0.05/0.1/0.2), and graph structures (whether the neck keypoints in the spatial graph of the two mice should be connected) was considered. For the final evaluation, all models are trained with the entire train split of task3.

For the MABe challenge, the winning model consisted of an ensemble of models trained in this way. For evaluating reproducibility, only one model from the ensemble was selected and trained across 5 runs.

G Additional Results

G.1 Additional Results for Task 1: Classic Classification

Results for each class on Task 1 is in Table 8. The behavior class “attack” is where we observe the most performance difference on our benchmarks, while model performances on investigation and mount are closer (often within one standard deviation). The performance on the “mount” class is very close to the ceiling, while both “attack” and “investigation” can likely be further improved.

G.2 Additional Results for Task 2: Annotation Style Transfer

Depending on the annotator and behavior, models can have a range of performances from lower F1 scores such as 0.6 to higher ones such as 0.9 (Table 9). We observe that all classifiers generally performs better on annotators 1 and 2, while performing worse on annotators 3, 4, and 5. This demonstrates the importance of studying techniques that can transfer to different annotation styles.

Method	Attack		Investigation		Mount	
	F1	AP	F1	AP	F1	AP
Annotator 1						
Baseline	.713 \pm .027	.819 \pm .017	.804 \pm .010	.878 \pm .008	.889 \pm .012	.969 \pm .002
Baseline w/ task prog	.868 \pm .001	.944 \pm .001	.830 \pm .001	.906 \pm .001	.912 \pm .001	.974 \pm .001
MABe 2021 Task 2 Top-1	.858 \pm .012	.893 \pm .028	.858 \pm .012	.916 \pm .012	.939 \pm .007	.968 \pm .005
Annotator 2						
Baseline	.839 \pm .008	.887 \pm .006	.865 \pm .006	.924 \pm .006	.671 \pm .050	.836 \pm .019
Baseline w/ task prog	.881 \pm .001	.938 \pm .001	.867 \pm .002	.922 \pm .001	.866 \pm .001	.943 \pm .003
MABe 2021 Task 2 Top-1	.865 \pm .019	.937 \pm .020	.904 \pm .010	.954 \pm .004	.849 \pm .041	.914 \pm .016
Annotator 3						
Baseline	.687 \pm .022	.710 \pm .013	.637 \pm .006	.650 \pm .009	.812 \pm .006	.820 \pm .012
Baseline w/ task prog	.710 \pm .004	.755 \pm .002	.583 \pm .002	.637 \pm .001	.773 \pm .004	.866 \pm .003
MABe 2021 Task 2 Top-1	.707 \pm .032	.746 \pm .013	.635 \pm .020	.660 \pm .011	.763 \pm .099	.893 \pm .031
Annotator 4						
Baseline	.736 \pm .012	.721 \pm .017	.686 \pm .010	.732 \pm .011	.882 \pm .032	.947 \pm .009
Baseline w/ task prog	.747 \pm .001	.776 \pm .002	.685 \pm .004	.751 \pm .005	.877 \pm .006	.939 \pm .003
MABe 2021 Task 2 Top-1	.795 \pm .018	.793 \pm .031	.733 \pm .005	.784 \pm .018	.895 \pm .068	.978 \pm .017
Annotator 5						
Baseline	.624 \pm .023	.715 \pm .010	.818 \pm .004	.880 \pm .006	.646 \pm .057	.709 \pm .061
Baseline w/ task prog	.732 \pm .002	.801 \pm .001	.800 \pm .002	.866 \pm .001	.719 \pm .020	.804 \pm .006
MABe 2021 Task 2 Top-1	.698 \pm .051	.693 \pm .063	.838 \pm .010	.874 \pm .024	.791 \pm .060	.848 \pm .039

Table 9: Per-annotator and per-class results on Task 2 (mean \pm standard deviation over 5 runs. The top performing method is in bold.)




Figure 11: Learned Annotator Matrix from Top-1 Model on Tasks 1 & 2. The matrix is initialized as a diagonal matrix and each embedding dimension corresponds to the corresponding annotator id. The lines on the top of the matrix represent the results from hierarchical clustering.

Models that can more accurately classify behavior for all annotators could help annotators better understand the difference in their styles, and could potentially be applicable to more users.

Annotator Embedding Matrix. After training the top entry for the Task 1 & 2 model, the learned annotator embedding matrix was extracted (Figure 11). Euclidean distances were computed for all pairs of annotator embeddings and hierarchical clustering was performed using these distances and Ward’s method [67]. This learned embedding matrix corresponds to learned similarity between the annotators. Here, we see that annotators 3 and 4 are more different from the other annotators, and are more similar to each other.

Method	Approach		Disengaged		Groom	
	F1	AP	F1	AP	F1	AP
Baseline	.338 ± .011	.282 ± .005	.195 ± .005	.237 ± .017	.289 ± .023	.260 ± .028
Baseline w/ task prog	.310 ± .011	.233 ± .015	.198 ± .010	.160 ± .024	.409 ± .008	.404 ± .010
MABe 2021 Task 3 Top-1	.182 ± .039 (.272 ± .031)	.209 ± .016	.101 ± .033 (.175 ± .022)	.086 ± .010	.411 ± .087 (.415 ± .112)	.391 ± .087
Method	Intromission		Mount Attempt		Sniff Face	
	F1	AP	F1	AP	F1	AP
Baseline	.721 ± .010	.746 ± .025	.034 ± .003	.015 ± .001	.358 ± .006	.322 ± .015
Baseline w/ task prog	.609 ± .019	.698 ± .023	.048 ± .009	.021 ± .004	.274 ± .026	.314 ± .022
MABe 2021 Task 3 Top-1	.663 ± .033 (.697 ± .026)	.761 ± .021	.001 ± .002 (.012 ± .011)	.013 ± .007	.304 ± .029 (.361 ± .023)	.311 ± .015
Method	White Rearing					
	F1	AP				
Baseline	.430 ± .006	.355 ± .017				
Baseline w/ task prog	.430 ± .029	.427 ± .025				
MABe 2021 Task 3 Top-1	.569 ± .035 (.606 ± .029)	.699 ± .031				

Table 10: Per-class results on the seven behaviors in Task 3 (mean ± standard deviation over 5 runs). The average F1 score in brackets corresponds to improvements with threshold tuning.

G.3 Additional Results for Task 3: New Behaviors

Based on performance metrics, Task 3 is more difficult than the other two Tasks (Table 10). The challenge of this task is the limited amount of training data for each new behavior, and that some behaviors are very rare (for example, mount attempt occurs in less than 1% of the frames). Here, none of the evaluated models consistently perform well on all the behaviors. We note that threshold tuning can help improve F1 score instead of simply taking the class with max predicted probabilities.

H Reproducibility Checklist

Here we provide additional details based on the ML Reproducibility Checklist.

H.1 Baselines

- **Source code link:** <https://gitlab.aicrowd.com/aicrowd/research/mab-e/mab-e-baselines>
- **Data used for training:** Train split of the corresponding task.
- **Pre-processing:** See Sections 4.1, F.1.
- **How samples were allocated for train/val/test:** CalMS21 provides train and test splits. The val split was held out from a random 20% of the train split.
- **Hyperparameter considerations:** See Section 4.1. In particular, for Task 1, we considered learning rates (0.0001/0.0005/0.005/0.001), frame skip (1/2), window size (50, 100, 200), convolution size (3/5/7/9) for the 1D Conv Net, and channel sizes for each layer (16/32/64/128/256). The hyperparameters for Tasks 2 and 3 are based on the tuned hyperparameters on Task 1.
- **Number of evaluation runs:** 5
- **How experiments where ran:** See Section 4.1.
- **Evaluation metrics:** Average F1 and MAP
- **Results:** See Sections 4.2, 4.3, 4.4, G.
- **Computing infrastructure used:** All baseline experiments were ran on Google Colab on CPU (Intel 2.3 GHz Xeon CPU).

H.2 Baselines with Task Programming Features

- **Source code link:** <https://github.com/neuroethology/TREBA>
- **Data used for training:** Unlabeled set for pre-training the features extraction model with task programming. Train split of the corresponding task for training baseline.
- **Pre-processing:** Same as baselines (see Sections 4.1, F.1) except with task programming features concatenated at each frame.
- **How samples were allocated for train/val/test:** Task programming model was trained on the first 220 sequences of the unlabeled set and validated on last 62 sequences.
- **Hyperparameter considerations:** The task programming model uses the same hyperparameters as [59], except trained for 300 epochs since the training set is smaller. The same hyperparameters as baselines above are used for training the supervised models, except Task 2 requires less epochs to converge (6 epochs instead of 10).
- **Number of evaluation runs:** 5
- **How experiments where ran:** See Section 4.1 except with task programming features concatenated at each frame.
- **Evaluation metrics:** Average F1 and MAP
- **Results:** See Sections 4.2, 4.3, 4.4, G.
- **Computing infrastructure used:** The task programming models are trained on a Amazon p2 instance, with one NVIDIA K80 GPU, and Intel 2.3 GHz Xeon CPU.

H.3 MABe Challenge Tasks 1 & 2 Top-1

- **Data used for training:** All splits of CalMS21, including unlabeled videos.
- **Pre-processing:** See Sections 4.2, F.5.
- **How samples were allocated for train/val/test:** CalMS21 provides train and test splits. See Section F.5 for validation set details.
- **Hyperparameter considerations:** See Section F.5.
- **Number of evaluation runs:** 5
- **How experiments where ran:** See Section F.5.
- **Evaluation metrics:** Average F1 and MAP
- **Results:** See Sections 4.2, 4.3, G.
- **Computing infrastructure used:** A server with 4 x GeForce RTX 2080. Each model for the cross-validation runs was trained on a single GPU. The server has a AMD Ryzen Threadripper 1950X CPU with 64GB RAM.

H.4 MABe Challenge Tasks 3 Top-1

- **Data used for training:** Train split of Task 3.
- **Pre-processing:** See Sections 4.4, F.6.
- **How samples were allocated for train/val/test:** CalMS21 provides train and test splits. The model was validated on 20% of the train split of Task 1.
- **Hyperparameter considerations:** See Section F.6.
- **Number of evaluation runs:** 5
- **How experiments where ran:** See Section F.6.
- **Evaluation metrics:** Average F1 and MAP
- **Results:** See Sections 4.4, G.
- **Computing infrastructure used:** All models are trained on 2 NVIDIA TESLA V100 32GB GPUs.