2014 22nd International Conference on Pattern Recognition

Automated Social Behaviour Recognition At Low Resolution

Tanmay Nath
iMinds-Vision Lab
University of Antwerp
Antwerp, Belgium
Email: tanmay.nath@uantwerpen.be

Guangda Liu
Vlaams Instituut voor Biotechnologie(VIB),
Katholieke Universiteit
Leuven,Belgium
Email:guangda.liu@cme.vib-kuleuven.be

Bassem Hassan
Vlaams Instituut voor Biotechnologie(VIB),
Katholieke Universiteit
Leuven,Belgium
Email: bassem.hassan@cme.vib-kuleuven.be

Barbara Weyn
DCILabs
Antwerp,Belgium
Email:barbara.weyn@dcilabs.com

Steve De Backer
DCILabs
Antwerp,Belgium
Email: steve.debacker@dcilabs.com

Paul Scheunders
iMinds-Vision Lab
University of Antwerp,Belgium
Email:paul.scheunders@uantwerpen.be

***Abstract*—Automated behaviour recognition is a challenging problem and it has recently gained momentum in biological behaviour studies. This paper describes a framework for tracking and automatical classification of the behaviour of multiple freely interacting Drosophila Melanogaster (fruit flies) in a low resolution video. The movements of interacting flies are recorded by Flyworld, a dedicated imaging platform. Each individual fly is identified in every frame and tracked over the complete video without losing its identity. The orientation of the flies is tracked as well, by defining their head and tail positions. From the obtained tracks, temporal features for every pair of fly are derived, allowing quantitative analysis of the fly behaviour. In order to derive information of the fly social activity, we concentrate on 2 specific behaviours: 'sniffing' and 'chasing'. Experimental results show that the classifier is able to classify the correct behaviour with an average overall accuracy of 95.46%.**

I. Introduction

The analysis of behaviour has recently been a very active research topic [2], [5], [6], [14]. Behaviour classification finds applications in a wide range of area's, including video surveillance, neuroscience, social robotics. Animal behaviour classification is a challenging problem for computer vision researchers. There are a lot of practical advantages of studying animal behaviour. Firstly, segmentation and tracking of animal is easy in laboratory setups, where the background does not change with time. Secondly, animal behavior is relatively simple and easy to understand as compared to human behaviour [2]. Thirdly, data acquisition is practical and reproducible [2] and it can be collected by multiple recording sessions. Thus, studying behaviour in animals presents an opportunity for making progress on modeling and classifying behaviour, especially social behaviour, which is difficult to study in humans [2].

Animal behaviour recognition also helps biologists to understand the neurobiological basis of social abnormalities in psychiatric disorders [5]. A central issue in analysis of these behaviour is reliable recognition of specific behavioural parameters. In such investigations, a manual scoring of the social interactions is still the preponderant experimental bottleneck [1]. Indeed, manual scoring could be an ideal solution, but it becomes almost impossible when it comes to analyse videos with several thousand frames. Also, manual scoring would suffer from lack of reproducibility and standardization. Moreover, it becomes even more challenging when multiple animals are involved.

Hence, there is an increasing interest in development of systems with automated behaviour analysis from video. Researchers have worked on automating the recognition of Rodent behaviour, but no substantial work has been done on analysis of fly behaviour, which is equally important for biologists. Flies are the perfect model for genetic manipulation since they grow fast, exhibit simple behaviour and are easy to handle in different ways, compared to the months of waiting for rodent models which show complex behaviour during interaction. Automating the analysis of fly behaviour is challenging. Firstly, building a system that accurately tracks the moving flies is technically difficult. Secondly, lack of accurate and standard definition of a behaviour makes an automated classification method more difficult. Lastly, inadequate availability of benchmark dataset restrains the evaluation of the algorithm [2].

Keeping all these challenges in mind, we present a framework of an automated technique for behaviour classification of freely interacting Drosophila Melanogaster (fruit flies).

In this framework, the first issue to be tackled is "multiple fly tracking", i.e. automatic detection of positions of multiple flies in each frame and linking of the flies found in a previous frame with the flies in a current frame without losing their identity. The second issue to be addressed is the automatic analysis of fly behaviour. This involves recognition and classification of certain fly behaviours as temporal actions, from a video of multiple freely interacting flies.

With the development of these machine learning and computer vision applications, we intend to help the biologists in automatic monitoring of the behaviour of Drosophila

1051-4651/14 $31.00 © 2014 IEEE
DOI 10.1109/ICPR.2014.403

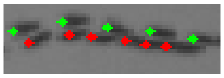

(a)

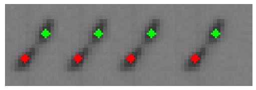

(b)

Fig. 1. The two behaviour events. Figure 1(a) shows the *chase* event. The green and red dots indicate the presence of two different flies. The flies position is overlaid for every third frame to get a proper view. Figure 1(a) shows a red fly chasing a green fly. Figure 1(b) shows the *Sniff* event. The red fly is sniffing the green fly.

Melanogaster.

In this work, we will concentrate on 2 different categories of social behaviour: *Chase* and *Sniff*. Figure 1 shows instances of the *Chase* and *Sniff* categories. We will describe the tracking, the calculation of behaviour features and the classification of behaviours into one of these 2 types of social behaviour, or the 'no-behaviour' class.

The structure of the paper is as follows. In Section II we discuss the related work. We define the behaviour classes in Section III. In Section IV we present our proposed method. Section V describes the experiments. Section VI shows the result and discussion. Finally, we conclude in Section VII.

II. Related Work

Tracking and behaviour recognition are often related problems and many approaches of behaviour recognition are based on tracking models of varying sophistication, from paradigms that use explicitly shape models in either 2D or 3D to those that rely on tracked features [9]. On a general overview, the time series trajectory obtained from the tracking acts as a descriptor for behaviour recognition.

Heiko Dankert et. al. [14] introduced a machine learning method to study the aggression and courtship behavior in Drosophila Melanogaster. Their tracking algorithm, *QTrack* and analysis package, *Caltech Automated Drosophila Aggression-Courtship Behavioral Repertoire Analysis* (CADABRA) is able to monitor interacting pair of flies and extract features which are used to detect behaviors exhibited during aggression and courtship. Therefore, the study is restricted to behavior analysis of only a pair of fly. Thus, their tracking method and analysis could not be implemented in studying the social behavior of more than a pair of flies.

Burgos-Artizzu et al.[2] adopted an approach of classifying pair wise behaviour of a mouse. They adopted an approach based on machine learning, where the behaviour is learned from examples given by the human annotators. Initially, their undisclosed algorithm tracks a mouse and extracts the temporal features. A machine learning model (AdaBoost [12]) is trained to classify the behaviour of the mouse. Although, it provides an extensive work on behaviour classification, yet it fails to generalise over multiple mice interactions or under different experimental conditions[5].

Their behaviour classification results were further improved by Giancardo et al [5]. They identify each mouse using a temporal watershed and mice matching module, which is an approximation of the Hungarian algorithm. It is able to track multiple mice in experimental conditions. Pair-wise spatio-temporal features from the trajectories of each mouse act as descriptors for the behaviour classification. However, these works focused on rodent models.

JAABA developed by Kabra et.al. [6], the current state of the art, proposed an interactive machine learning tool for automated annotation of animal behaviour including fruit flies. The behaviour classification was built on top of the results of their tracking algorithm CTrax [1]. In order to track, the flies in every frame were segmented and their positions in past and future frames were linked as a linear assignment problem using the Hungarian algorithm. A set of high resolution features define the behaviour of flies and the Gentleboost algorithm [13], a modification of the AdaBoost [12] algorithm, was used to classify the behaviour. However, it was built as a binary classifier. The biologists manually annotate any one class at a time and observe the predictions of the classifier. All the false positive events are manually corrected and the classifier is retrained. This way of iteratively training the classifier reduces the number of false positives and thus increases the accuracy. However, it is time consuming and only one behaviour class can be classified at a time.

We will treat the same problem of automated behaviour classification as a multi-class classification where we train the classifier with three different classes at the same time. Based on the features, the classifier predicts the classes, thus reducing the time and manual efforts of annotating video with different classes.

All of the related work has been done on a high resolution video. However, not all biologists require high resolution and sophisticated mechanisms to study the fly behaviour. Social biologists looking at the social index [7] of flies or studying the effect of social isolation of flies [10] are not concerned about the high level behaviour. These studies do not require specific physical details of flies e.g. wing extension angle, antennae angle, position of legs etc. Moreover, the study involves detection of interactions and quantification of social network based on these interactions. The study requires classification of social and non social behaviour that can even be detected at low resolution. Another motivation for using low resolution camera's is to provide larger fields of view in experiments studying e.g. effects of social isolation in fly behaviour.

In our previous work, [8] we state the critical problem with using low resolution data. We have shown the limitations of usage of the current state of art, CTrax [1] on low resolution video. Moreover, we also state the motivation behind developing a method which works for low resolution data was the inability of standard algorithm used in CTrax [1] to provide desired result on low resolution data. In our current work, we provide a computational framework for behaviour analysis on low resolution video, integrating our tracking algorithm, FlyTracker [8] and a new method for classifying behaviours of a group of freely interacting fruit flies.

III. Behaviour Class Definition

Data is acquired using the FlyWorld [8], a dedicated imaging platform. Typically, 20 interacting files are recorded, videos last

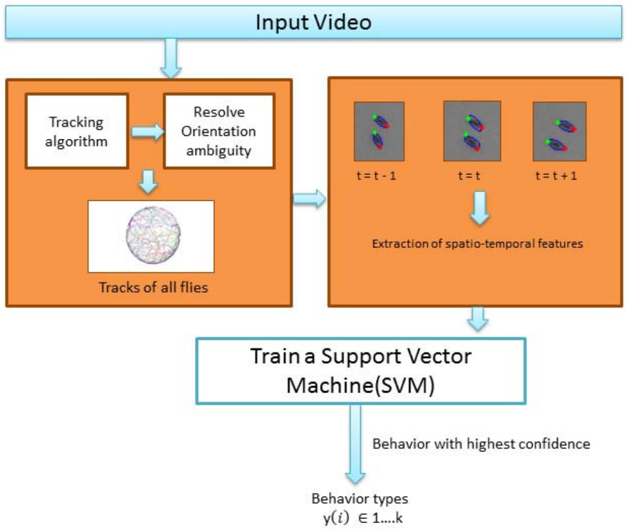

Fig. 2. Overview of our approach. The flies in the input video are tracked and their orientation ambiguity is resolved. Pair wise spatio-temporal features are calculated from their trajectory. A classifier is trained on these features. The green and the red points indicate the result of the orientation ambiguity. The green points indicate the nose while the red points indicate the tail of each fly at different time points.

around 20 minutes and are recorded at 30 fps with a field of view of 410x410 pixels.

The behaviour classes are defined in collaboration with biologists involved in the study of the effect of social isolation in fly behaviour. The touching event is the most occurring social behaviour amongst the fly behaviour and intuitively all subsequent behaviour follows from this. The behaviours are classified in a set of three different mutually exclusive action categories. 2 behaviours are the *Chase* and *Sniff*. The third class is the *None* which is defined as a pair of flies that does not display any social behaviour. These classes were chosen since these behaviours make an obvious choice of studying the effect of social isolation on fly behaviour. In every video, the flies start with exploring the FlyWorld chamber and then start interacting with other flies. The fly which starts to interact by approaching the other fly and makes a nose-to-tail touch is known as a principal interactor fly. This behaviour is defined as a *Sniff*. If the principal interactor fly finds out that the interactee fly is a female fly, then the former tries to copulate with the latter. During this event, the interactor fly walks behind the interactee. This behaviour is defined as *Chase*. Figure 1 shows example video frames for each behaviour in top view. All other behaviours are defined as *None*.

Every possible pair of flies in the video frame is labelled as one of the three categories, resulting in segmentation of the videos into action intervals. The start and end of each category has to be accurately defined.

A GUI was developed in MATLAB to annotate these behaviours. This tool provides a frame by frame view of the movie with a play back feature, thus providing the human annotator to carefully analyse the videos.

IV. Proposed Method

Figure 2 shows an overview of our approach. Our approach comprises two challenges: tracking all flies at low resolution without losing their identity and extracting temporal features to train the classifier.

In our previous work [8] we have shown that CTrax tracking results [1] lose the identity of flies in low resolution video. This makes the classification solution as described in [6] inadequate for low resolution videos.

In [8], we proposed a tracking algorithm for low resolution video. FlyWorld, the imaging setup allows the videos to be recorded with constant background. In our setup, flies appear as dark dots on a light background. The background is modeled by selecting the frames at a reduced sampling rate and taking the pixel-wise maximum over all selected frames. We used a two step Hungarian algorithm. While the first step of the Hungarian algorithm provides an estimate of the number of flies in a group of interacting flies, the second step helps in linking the flies from a previous frame to a current frame. Figure 3 shows the procedure. This has greatly improved the tracking results with respect to CTrax, as is shown in Table I.

A. Body Points

In the current work, we discuss the feature extraction and classification steps. The tracking results yield weak trajectory features such as velocity, acceleration etc. which are derived from the position of fly centroid at every consecutive frame. The low resolution of the video restricts a clear and distinguishing view of the fly orientation. In order to describe the wanted behaviour classes, more anatomical points than only the body centroid are required. We will define the *nose* and the *tail* as body points. The definition of these body points is required to resolve the orientation ambiguity. In order to resolve this ambiguity, we make an assumption that the orientation of a fly body and the direction of the fly's velocity will approximately match [1]:

Let θ_t, ϕ_t and v_t be the orientation, velocity direction and speed of a fly at frame *t*, and *T* the total number of frames in the video. An indicator function $s_t \in \{0,1\}$ indicates whether π needs to be added to the orientation or not i.e. $\theta' = \theta + \pi s_t$. The indicator function s_t is obtained by minimizing the following criterion:

$$J(s_{1:T}) = \sum_{t=1}^{T} [J_1(s_t) + J_2(s_t, s_{t-1})], \text{Where} \tag{1}$$

$$J_1(s_t) = w(v_t) \mid (\theta_t + \pi s_t - \phi)_{(-\pi,\pi]} \mid \tag{2}$$

$$J_2(s_t, s_{t-1}) = (1 - w(v_t)) \mid (\theta_t + \pi s_t - \theta_{t-1} - \pi s_{t-1})_{(-\pi,\pi]} \mid \tag{3}$$

The equation 2 penalizes the orientations $\theta_t + \pi s_t$ that differ from the velocity direction ϕ_t. Since for flies that are sitting still, the direction of their velocity does not relate to their orientation, weights $w(v_t)$ are introduced to make the error term proportional to the magnitude of velocity of the fly. The equation 3 penalizes the orientation at frame *t*, $\theta_t + \pi s_t$,

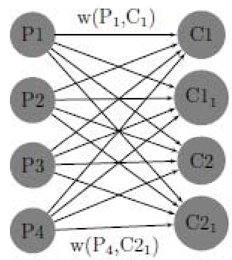

(a) Estimate the number of interacting flies, k

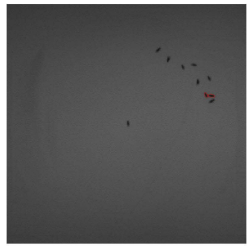

(b) Segmentation with $k = 2$

Fig. 3. Result of an estimate of the number of flies using the Hungarian Algorithm. It is assumed that only the nearby flies will have the tendency to interact with each other. Thus a linear assignment problem will give an exact estimate of the number of interacting flies. The vertices P_1,..,P_4 represent the individual flies in a previous frame, while C1 and C2 represent the candidate blobs with multiple flies. $C1_1$, $C2_1$ are the newly appended vertices, which represent the nearby flies in previous frames. The weights w represents the distance between the flies in two subsequent frames.

differing from the orientation at frame t - 1,$\theta_{t-1} + \pi s_{t-1}$.

Once the orientation ambiguity is resolved, points representing the *nose* and the *tail* are defined.

B. Temporal Features

In order to capture the different behaviours of flies, the feature vector needs to describe the relative motion between pairs of flies. A feature vector for every possible pair of flies, $X_t^{\alpha \Rightarrow \beta}$ (α is the principal interactor fly and β is the interactee fly) is constructed using a set of temporal features. A sliding window centered at each frame is used to extract these temporal features. The optimal sliding window size is directly related to the duration of behaviours[2]. A series of experiments with different window sizes provided an optimal window size of 11 frames ($\approx$ 0.36 sec).
The feature vector is a 23-dimensional vector, which measures following two categories:

Relative Position: The Euclidean distances measured between each pair of flies. This distance is measured between the corresponding nose of the principal interactor fly and the tail of the interactee fly. This feature does not represent action dynamics.

Movement: These spatio-temporal features represent the fly's action dynamics in past and future frames. A window size of 11 frames is defined as described above to measure the action dynamics. Specifically two distances are measured in this temporal window: one describing the pairwise distance between the body of the principal interactor fly and the body of the interactee fly, and another describing the displacement of the principal interactor fly over the time window, t-5, t-4,...t+4, t+5.

C. Classification

A classifier which can differentiate between the behaviours is developed using a Support Vector Machine (SVM) [4] using Gaussian radial basis function kernels (RBF). SVM maps the input feature vectors into some high dimensional feature space through non-linear mapping [4]. In this space SVM tries to find the optimal hyperplane to separate the classes. SVM's are known to optimally handle the case when the relation between class labels and attributes is nonlinear. A RBF kernel was chosen because it has less parameters and provides better accuracy. The kernel is defined by:

$$\exp(-\gamma * |u - v|^2) \tag{4}$$

where $\gamma = \frac{1}{2\sigma^2}$ of a Gaussian with variance σ.
The standard library, *LIBSVM* [3] was used to implement the classifier. There are two parameters to be estimated for the classifier, namely the penalty parameter, C and the kernel parameter, γ. The estimation is accomplished using a grid search on C and γ using cross-validation. An exponential growing sequence of pairs of (C,γ) are tried and the one with the best cross-validation accuracy is picked [3].

TABLE I
PERFORMANCE OF FLYTRACKER AND CTRAX

Videos	Frame Analyzed	Flies Present	Tracks Detected by FlyTracker	Tracks detected by CTrax
1	28480	10	10	31
2	31829	15	15	47
3	29105	49	49	850

V. EXPERIMENTS

A video of 20 interacting flies was recorded using the FlyWorld [8] for 20 minutes which comprised of approximately 36000 frames. 5 pairwise fly interactions were studied. This was done because not all flies were active and most of them either do not get involved in pairwise interactions or they prefer to remain unsocial. This class of unsocial behaviour was considered as *No Class*.

The training data is composed of a set of 2 different pairwise interaction events. Since the results of FlyTracker [8] are accurate, we can maintain the identity of all the flies. A set of fly Ids {5,9} and {7,9} was used as training data. The number of training events in the training data is 1000.

The testing data is composed of a set of 3 different pairs of interacting flies. The classifier was never trained on this dataset and its annotations were used as ground-truth to validate the results of the classifier.

VI. RESULT

The purpose of this section is twofold: 1) to analyse the tracking results and compare with the current state of art, Ctrax [1] and 2) test the performance measure of the classifier.

Videos were analyzed using FlyTracker [8] and Ctrax[9]. Video sequences were recorded with the maximum frame rate of 30 frames per second, at a resolution of 2.4 pixel/mm. Table I shows the comparative results of FlyTracker and Ctrax. The result of the experiments show that FlyTracker is able to track the flies in low resolution, and gives the correct number of flies present in the chamber. Various videos were acquired and

TABLE II
COMPARISON OF HUMAN PERFORMANCE WITH THE CLASSIFIER RESULTS.

Annotator II

Ground truth	Chase	Sniff	No Class
Chase	38	8	5
Sniff	0	6	0
No Class	0	0	943

Our Method

Ground truth	Chase	Sniff	No Class
Chase	44	7	0
Sniff	4	2	0
No Class	18	2	923

Annotator II

Ground truth	Chase	Sniff	No Class
Chase	35	0	0
Sniff	11	0	15
No Class	2	6	931

Our Method

Ground truth	Chase	Sniff	No Class
Chase	20	15	0
Sniff	3	23	0
No Class	40	7	892

Annotator II

Ground truth	Chase	Sniff	No Class
Chase	47	1	22
Sniff	0	13	18
No Class	0	8	891

Our Method

Ground truth	Chase	Sniff	No Class
Chase	54	16	0
Sniff	8	22	1
No Class	12	3	884

analyzed with increasing number of flies and it was found that FlyTracker could track up to 49 flies.

On the other hand, to test the current state of the art, JAABA [6], an attempt was made to use it on our data set. Since JAABA is built on CTrax, there were many events with identity loss and identity swapping, resulting in tracking errors. Consequently, it was not possible to manually label and classify a specific pair of fly behaviour using JAABA, which clearly indicates the limited use of CTrax and JAABA in a low resolution environment.

For this reason, we applied the classification only on tracking results obtained by FlyTracker. To test the performance of the classifier, the ground truth data was compared to the result of the classifier. A separate comparison was made with a second expert's annotation on the same video. This comparison indicates the difficulty for humans to accurately and precisely classify certain fly behaviour.

The Table II shows the confusion matrix for three different pairs of fly behaviour classification. On the left hand, the results of the human performance are shown, while the tables on the right side show the classifier performances. The classifier is able to obtain a maximum overall accuracy of 96.90% with an average overall accuracy of 95.46% as compared to the average overall accuracy of 96.8% of human annotations. However, the overall accuracy is misleading as the classes are unbalanced (there are far more samples of the No Class than of the social behaviour classes). To show the complexity of accurately and

TABLE III
AVERAGE (ACROSS THE THREE FLY PAIRS) OF RECALL AND PRECISION(IN PARENTHESIS) FOR THE BEHAVIOUR CLASSES DEFINED.

Class	Human Performance (in %)	Our Method (in %)
Chase	80.54(90.97)	73.51(53.73)
Sniff	47.3(33.98)	63.17(41.88)
No Class	99.4(97.86)	97.4(99.92)
Average	75.74(74.27)	78.02(65.17)

precisely annotating a particular behaviour, we need to compare the recall and precision [11] values for both the methods. If *TP* represents the True Positives, *FN* the False Negatives and *FP* the False Positives, then recall and precession is defined as:

$$Recall = \frac{TP}{TP + FN} \quad (5) \qquad Precision = \frac{TP}{TP + FP} \quad (6)$$

Table III shows the average (across the three fly pairs) of recall and precision(in parenthesis) for the behaviour classes defined. The results indeed highlight the complexity that even human performance is error prone and does not necessarily achieve a high success rate. On the other hand results of our approach are comparable to human performance.

VII. CONCLUSIONS

In this work we have presented an approach for automated social behaviour recognition of Drosophila Melanogaster(fruit fly) in low resolution continuous video. The approach is based on the extraction of temporal and trajectory features of every pair of flies and classification of 2 classes of social behaviour. The biological motivation is the study the effect of isolation on fly behaviour.

The current state of art lacks good tracking results and provides poor classification on low resolution video. We show that our tracking results are better than the current state of art [1]. We proposed a new classification strategy that is able to classify different types of social behaviour. The performance of our technique is comparable to how humans would perform. The proposed classification strategy allows to process multiple videos in an automated way.

However, there is still room for improvement in classification results. In the future, we will try to solve the low precision due to the imbalance in the different classes. We will also try to define other relevant classes of social behaviour, given the limitation of low resolution video. A limitation of the current work is that it is unable to classify higher level complex behaviours such as wing rowing, copulation, wing extension etc.

ACKNOWLEDGEMENT

The authors would like to thank the European Commission FP7-PEOPLE-2011-ITN project FliACT (289941) for funding this project.

REFERENCES

[1] K. Branson, A. A. Robie, J. Bender, P. Perona, and M. Dickinson. High-throughput ethiomics in large groups of drosophila. *Nature Methods*, 4:3099–3104, 2009.

[2] X. Burgos-Artizzu, P. Dollar, D. Lin, D. Anderson, and P. Perona. Social behaviour recognition in continuous video. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 1322–1329, 2012.

[3] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at http://www.csie.ntu.edu.tw/ cjlin/libsvm.

[4] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.

[5] L. Giancardo, D. Sona, H. Huang, S. Sannino, F. Managò, D. Scheggia, F. Papaleo, and V. Murino. Automatic visual tracking and social behaviour analysis with multiple mice. *PLoS ONE*, 8(9):e74557, 09 2013.

[6] M. Kabra, A. Robie, A. M. Rivera, S. Branson, and K. Branson. Jaaba: interactive machine learning for automatic annotation of animal behaviour. *Nature Methods*, 10:6467, 2013.

[7] A. F. Simon, M.-T. Chou, E. D. Salazar, T. Nicholson, N. Saini, S. Metchev, and D. E. Krantz. A simple assay to study social behaviour in drosophila: measurement of social space within a group1. *Genes, Brain and behaviour*, 11(2):243–252, 2012.

[8] T. Nath, G. Liu, B. Weyn, B. Hassan, S. De.Backer and P. Scheunders. Tracking for quantifying social network of drosophila melanogaster. In *Computer Analysis of Images and Patterns*, volume 8048, pages 539–545. Springer, 2013.

[9] P. Dollar, V. Rabaud, G. Cottrell, and S. Belongie. behaviour recognition via sparse spatio-temporal features. In *Visual Surveillance and Performance Evaluation of Tracking and Surveillance, 2005. 2nd Joint IEEE International Workshop on*, pages 65–72, 2005.

[10] F. De Melo Sene. Effect of social isolation on behaviour drosophila silvestris from hawaii. In *Proceedings, Hawaiian Entomological Society*, 1977.

[11] M. Sokolova and G. Lapalme. A systematic analysis of performance measures for classification tasks. *Information Processing & Management*, 45(4):427 – 437, 2009.

[12] Y. Freund and R. E. Schapire. A Decision-Theoretic Generalization of on-Line Learning and an Application to Boosting.1995

[13] J. Friedman, T. Hastie and R. Tibshirani. Additive Logistic Regression: a Statistical View of Boosting. *Annals of Statistics*, 2000(28):337–407, 1998.

[14] Heiko Dankert, Liming Wang, Eric D Hoopfer, David J Anderson, and Pietro Perona, “Automated monitoring and analysis of social behavior in drosophila,” *Nature Methods*, vol. 6, pp. 297–303, 2009.