Model-based Detection and Analysis of Animal Behaviors using Signals Extracted by Automated Tracking

Gennady Denisov, Tomoko Ohyama, Tihana Jovanic and Marta Zlatic
Janelia Farm Research Campus, Howard Hughes Medical Institute, Ashburn, Virginia, 20147, U.S.A.

Keywords: Animal Behavior, Automated Tracking, High-throughput Detection, Model, Algorithm, Software, Signal Processing, Statistical Analysis, Feature Extraction, Hit Detection.

Abstract: Analysis of behaviors of model organisms has a number of applications, particularly to determination of the function of genes and neurons. Drosophila larva is an especially convenient model system for this kind of study because of availability of powerful genetic analysis tools and of automated tracking software that allows high-throughput recording of animal's shape and position characteristics as time-dependent signals. We have developed an open source software that allows a high-throughput detection and analysis of a comprehensive set of meaningful behaviors of this species. Using the recorded signals as input variables and a set of processing thresholds as parameters, the software employs model-based algorithms to detect the behavioral actions with high accuracy, typically 1-5%. For each detected action it extracts and stores meaningful quantitative features that allow statistical discrimination of mutants from wild type animals and set stage for subsequent application of machine learning techniques to classification of the mutants.

1 INTRODUCTION

Insights into the function of a gene or neuron can be gained in multiple ways, including the loss-of-function screening of genes and neurons underlying observed behavioral phenotypes. Drosophila larva is an especially tractable and convenient model system for this kind of study because of the relative simplicity of its nervous system, availability of powerful genetic analysis tools (Pfeiffer et al., 2008), and availability of automated tracking software, such as Multi-Warm Tracker (MWT) (Swierczek et al., 2011). Given a contour of a tracked animal object, the MWT makes a guess about the position of its spinal cord (Fig. 1), assigns a "center of mass" to the object, computes a variety of other metrics and stores the results as time series data.

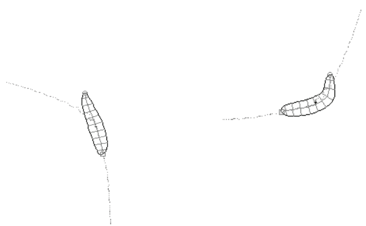

Figure 1: Drosophila larva objects tracked by Multi-Worm Tracker (MWT).

Automated analysis of behaviors of Drosophila larva has been primarily focused on study of different types of taxis (Gomez-Marin et al., 2011), (Gomez-Marin and Louis, 2012), (Gershow et al., 2012), (Luo et al., 2010). The most high-throughput approach, employing custom machine vision software, was used in (Gershow et al., 2012). Machine vision was also applied to analysis of social and sex behaviors in adult Drosophila melanogaster (Branson et al., 2009), where a machine learning classifier was used for automatic detection of animal behaviors with high accuracy. A typical machine learning classifier, which takes tracking movies as input, may employ from several hundreds to several thousand features. This has two drawbacks. First, the number of features that has to be processed reduces the efficiency of behavior detection. Screening of thousands of mutant lines, with each line being represented by a sufficiently large group of animals for making statistical conclusions, under a variety of experimental conditions/stimuli in order to explore different behavioral actions, will require efficient data analysis algorithms. Second, it is hard to know which of the used features are really important. Knowing which motor patterns are actually altered in specific mutants is a key to gaining biological insight into the function of a gene or neuron.

For this reason, we explored a different approach,

Denisov G., Ohyama T., Jovanic T. and Zlatic M..
Model-based Detection and Analysis of Animal Behaviors using Signals Extracted by Automated Tracking.
DOI: 10.5220/0004235101750181
In *Proceedings of the International Conference on Bio-inspired Systems and Signal Processing* (BIOSIGNALS-2013), pages 175-181
ISBN: 978-989-8565-36-5

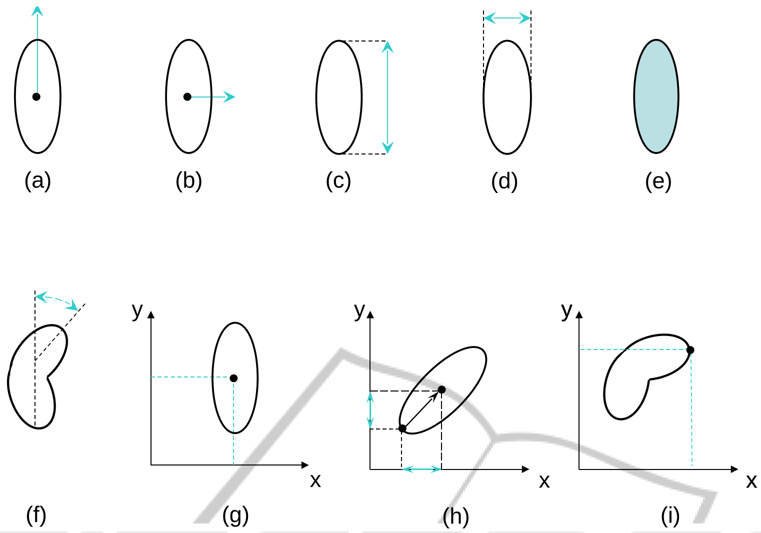

Figure 2: Signals used as input by SALAM: (a) (crawling) speed, mm/s; (b) crabspeed (or rolling speed), mm/s; (c) midline, mm; (d) morpwidth, mm; (e) area of the tracked object, mm^2; (f) (head) cast, degrees; (g) x and y, mm, (h) tailvecx and tailvecy, mm; and (i) headx and heady, mm.

which makes use of the signals extracted by MWT, rather than of tracking movies. This approach has been implemented in an open source software package SALAM (http://sourceforge.net/projects/salam-hhmi), written in Python and R programming languages. By properly combining the input signals, employing a set of specifically tuned thresholds and model-based algorithms, our software allows detection of a comprehensive set of meaningful behavioral actions of Drosophila larva with high accuracy and throughput. The detected actions include, but are not limited to, the crawling runs, head casts and the earlier studied types of taxis.

SALAM has been used by our recent study (Ohyama et al., 2012). This paper provides an overview of its capabilities and algorithms.

2 APPROACH OVERVIEW: SIGNALS, BIOLOGICALLY MEANINGFUL ACTIONS AND DATA PROCESSING WORKFLOW

Figure 2 schematically represents several signals output by Choreography software of MWT package (http://sourceforge.net/projects/mwt) and used by SALAM as input. Speed, or crawling speed (a), and crabspeed, or rolling speed (b), are defined as the speed of the center of mass in the direction parallel and perpendicular to the spine, respectively; (c) midline is the spine length; (d) morpwidth is measured at the center of mass perpendicular to the spine; (f) (head) cast is defined as the angle formed between the straight lines drawn for the top 20% and the bottom 80% of the spine; (g) (x, y) is the position of the center of mass, (h) (tailvecx, tailvecy) are components of the vector joining the bottom and middle points of the spine; and (i) (headx, heady) is the position of an animal nose, i.e. the top point of the spine.

Figure 3 schematically outlines behavioral actions observed in Drosophila larva and detected by SALAM, using its built-in algorithms, given the signals listed in Fig. 2. Crawling, Fig. 3(a), and casting, Fig. 3(b), are the most typical actions: these are what an animal is doing most of the time. Other actions usually occur in response to a certain kind of stimulus. For example, hunch, or contraction of animal's body, Fig. 3(c), and rolling, Fig. 3(d), usually occur in response to a time-dependent stimulus, such as IR light or ultrasound. Digging, whereby an animal pops up in an attempt to dig a hole in the support, Fig. 3(e), and following, Fig. 3(f), typically occur on supports possessing scratches. Animals may be prevented from digging and/or following by other stimuli, for example by a strong air flow. Finally, taxis action, Fig. 3(g), may be taken when there is a spatial gradient of a stimulus. Such an action may include two components: 1) turning an animal's body towards the source of stimulus (positive taxis) or from it (negative taxis), and 2) navigating (crawling) towards the area with higher or lower intensity of the stimulus, respectively.

Fig. 4 overviews a scope of operations and a workflow performed by SALAM.

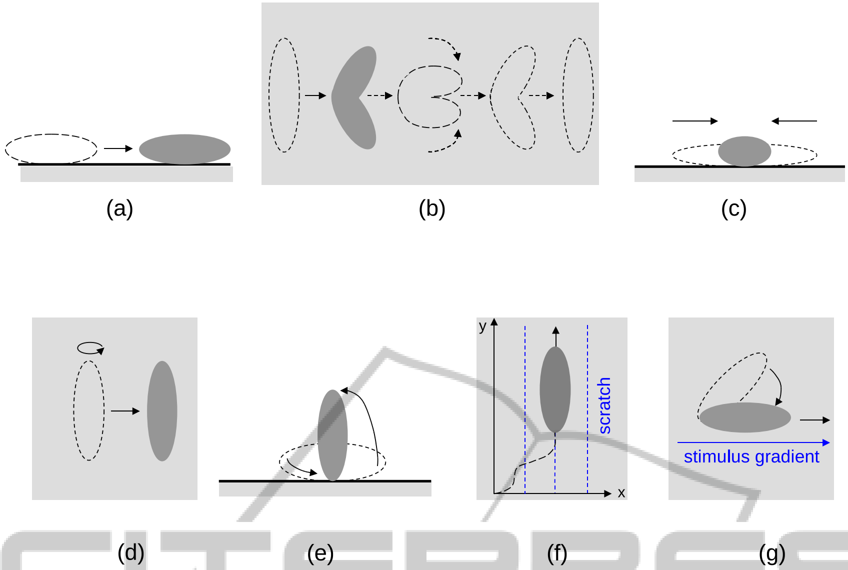

Figure 3: Meaningful behavioral actions detected by SALAM: (a) (peristaltic) crawling, (b) cast, where the middle chart represents the possibility of a *strong* cast, (c) hunch, (d) rolling, (e) digging, (f) following, i.e. crawling along a scratch on a support, and (g) taxis, where the stimulus may be an odor (chemotaxis), a temperature (thermotaxis), a visible light (phototaxis), or an air pressure, resulting in air flow (thigmotaxis).

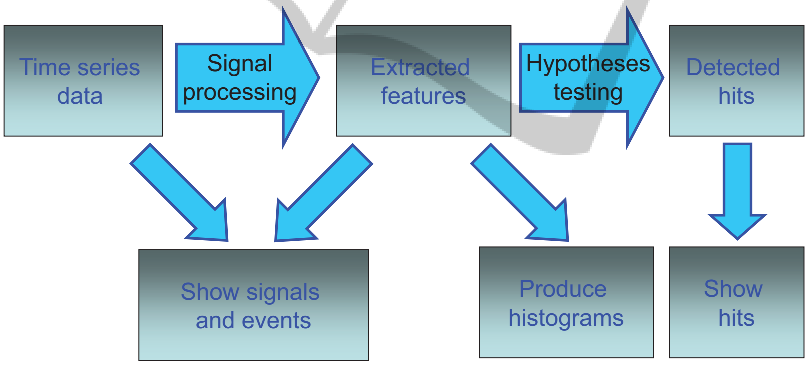

Figure 4: Flowchart of data processing. The upper row of boxes represents the core, high-throughput processing. The lower row of boxes represents visualization tools, which are typically applied to a limited amount of data with the purposes of debugging the algorithms or reviewing the results.

3 DETECTING BEHAVIORAL ACTIONS AND EXTRACTING FEATURES

Analysis of behavioral actions by SALAM starts from detecting events in the individual signals used.

3.1 Detection of Individual Signal Events

For non-oscillating signals, such as cast, crabspeed, midline or morpwidth, events are simply the significant peaks or wells, with amplitude above the specified thresholds. Detection of events in oscillating signals (e.g. speed, area) resulting from the peristaltic nature of larval crawling, takes a different approach. In either case, event detection is completed by computing an event signal, which is nonzero at events (equals the event amplitude) and zero otherwise.

Event detection in an oscillating signal is initiated by identifying peaks. Peak positions are the local maxima of a signal. Peak amplitude is a signal value at the peak position. Peak boundaries are set at the minima of signal on both sides of a peak. A peak is considered good if its amplitude exceeds a specified threshold. Positions of good peaks in speed signal are shown as dashed green lines in the left plot of Fig. 5. A speed event is a sequence of at least three adjacent good peaks. The event amplitude is the mean height of all the good peaks comprising the event. Event boundaries are set at the boundaries of the first and last peak comprising the event. (When detecting crawling actions, the algorithm additionally requires that a crawling event is terminated by any event in cast or crabspeed signal, Fig. 5).

Event detection in a non-oscillating signal. Our procedure is an extension of the "Schmitt trigger" approach previously used for detection of movement events in flies (Robie et al., 2010). It employs four adjustable thresholds, which are specifically tuned for each type of signal. The thresholds are: 1) the upper and 2) the lower amplitude threshold, shown as the solid and dashed horizontal green line, respectively, in Fig. 5 (cast and crabspeed plots) and in Fig. 6; 3) the width threshold; and 4) the gap threshold. An event starts when the absolute value of a signal, while increasing as a function of time, crosses the upper amplitude threshold. An event ends when the absolute value of a signal, while decreasing as a function of time, crosses the lower amplitude threshold. Event duration is the difference between the event end and event start times. Event amplitude is the highest absolute value of a signal during the event. A single event of duration less than the width threshold will not be detected (i.e., will be ignored). However, if two or more adjacent events of the same type (all peaks or all wells) are less than the gap threshold apart one from another, and the time duration between the start of the first event and the end of the last event exceeds the width threshold, then all the events will be merged into a single detected event, as illustrated by the crabspeed plot in Figure 5.

3.2 Detection of Behavioral Actions

Detection of the most of actions listed in Table 1, with the only exception for rolling, requires simultaneous processing of more than one signal. This requirement can be illustrated by comparing the strong cast (Fig. 3(b), middle chart), hunch, Fig. 3(c), and digging action, Fig. 3(e). During any of these actions, a drop in the midline signal is observed, so the shape of a tracked animal object is close to a ball (and therefore the MWT may be unable to properly identify the spine line). Thus, while the midline signal can be used for detecting these actions, it may not be sufficient for their reliable discrimination, so additional signals must be used.

Our model algorithm for detection of a strong cast action makes use of the cast, midline and morpwidth signals, as illustrated in Fig. 6. First, we note that a strong cast must be accompanied not only by a well in the midline signal, but also by a peak in the morpwidth signal. (The morpwidth peak is not expected to be observed neither during hunch nor during digging action.) Second, a strong cast can only appear as a part of the sequence of states schematically depicted in Fig. 3(b). Thus, one should always observe cast signal events (peaks or wells) on both sides of a midline well/morpwidth peak during the strong cast action, as illustrated by Fig. 6.

Detection of a hunch action makes use of the same three variables as detection of cast action, but employs additional thresholds to make sure there are no significant cast and morpwidth events in a close vicinity of the midline well.

Our algorithm for detection of a digging action, in addition to the presence of a deep well in a midline signal, requires that the values of both x and y signal stay constrained to a certain small region within a certain period of time, since the tracked animal object practically does not move while the animal digs.

A following action is detected similarly to the detection of digging action, but does not require the presence of a midline well and requires that $\Delta x/\Delta y$ stays below a certain small threshold for a certain period of time, since in all the experiments we used scratches were parallel to the y axis.

Detection of taxis actions is more sophisticated. For example, during chemotaxis, an animal usually performs two or more head casts, both left and right ones, and compares the local concentrations of an odorant at the position (headx, heady) of its nose during each cast (Gomez-Marin et al., 2011). It then changes its orientation, described by the tail vector with components (taivecx, tailvecy), and subsequently crawls for a certain time in the newly chosen direction.

3.3 Feature Extraction

The features defined for detected actions can be subdivided into two categories: those computed per-event, e.g. event amplitude, and those computed per-animal, e.g. an event frequency, i.e. number of events detected per unit time interval. Table 1 lists the features extracted for detected actions. For crawling action, only per-event features are defined: the mean peak amplitude (height), frequency of oscillation (determined by applying the Fourier transform to the portion of speed signal comprising the action event, and the event duration. For head cast, hunch and rolling actions, the features are the event amplitude (maximal value of a signal during the event) and event duration, as well as the event frequency. For digging, the only

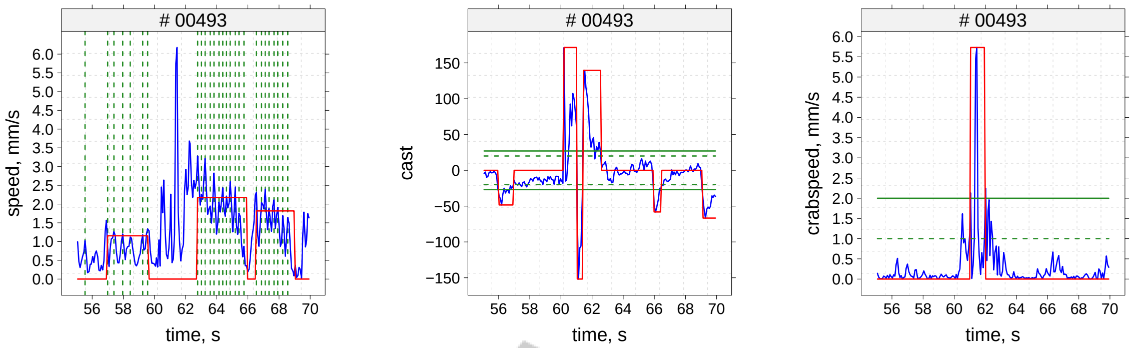

Figure 5: Detection of crawling actions. Shown in blue are three signals recorded for tracked animal #493 in a particular time window: speed (left plot), cast (middle plot) and crabspeed (right plot). The red line indicates an event signal, which is nonzero only at detected events. The green lines provide auxiliary information: positions of detected good peaks in the speed signal plot, and the upper (solid line) and lower (dashed line) amplitude thresholds used for event detection in the cast and crabspeed signal plots.

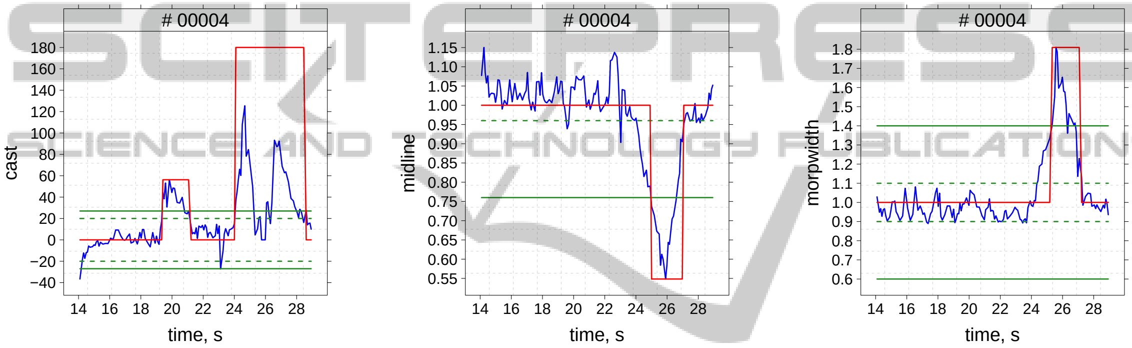

Figure 6: Detection of head cast actions. Shown in blue are three signals extracted by tracking animal #4 in a particular time window: cast (left plot), midline (middle plot) and morpwidth (right plot). Here, the midline and morpwidth signals are normalized to their median values. The meaning of the red and green lines is the same as in Fig. 5. Two detected cast actions have been shown in the left plot: the regular cast (left) and the strong cast (right). The strong cast can only occur as a part of a sequence of events schematically shown in Fig. 3(b). It is assigned a "theoretical" amplitude of 180^0, since MWT is unable to properly determine the value of cast signal for a ball-like object.

feature is the event frequency.

For taxis events, a number of features have been extracted, of which we will here discuss only the navigation index. For experiments where a gradient of a stimulus is constant and parallel to the x-axis, one way to define the navigation index is (Gershow et al., 2012)

$$nind = < speed_x > / < speed > \qquad (1)$$

where $speed_x$ stands for the x-component of the speed signal, and the angle brackets designate averaging over time. According to this formula, the navigation index should be always within the range [-1, +1], reaching the value +1 for animals moving strongly in the direction of a stimulus source and -1 for animals moving strongly in the opposite direction. While this definition can work reasonably well in many cases, it has two limitations. First, if an animal with strongly positive taxis was initially, at time t = 0, located a minimal possible distance away from the stimulus source, then (1) will give $nind = 0$ (since the animal can no longer move toward the source), rather than $nind = +1$. Second, formula (1) is not applicable to the settings with (multiple) point source(s) of stimulus, i.e. with non-flat geometry, which is being the case in studies of chemotaxis (Gomez-Marin and Louis, 2012).

We here present an alternative approach to the computation of navigation index, which attempts to overcome the two limitations mentioned above. Our approach is based on comparing the times an animal spent at different distances away from a stimulus source during experiment.

Let $d(t)$ be a distance from an animal object to the closest (flat or point) source of stimulus at time t. We assume that, for a given experimental setting, the minimal and maximal possible values of d are d_{min} and d_{max}, respectively. Then, the navigation index can be computed as

Table 1: Detected behavioral actions, signals used, accuracy of detection (false discovery rate, estimated by comparing the predictions from the software with tracking movies of contours of larvae) and features extracted and stored for each the action. The features are: amplitude, ampl; (crawling) frequency, freq; action event duration, dur; event frequency, efreq; fractions of time and path length spent following, ftime and flength; and the navigation index, nind. See the text for more detailed feature definitions. Accurate estimates of the FD rate for digging, following and taxis actions are not currently available. False negative rate was not estimated, since it does not affect (at least, directly) the statistical characteristics of extracted features, e.g. their mean values.

Action	Signals used	FD rate	Extracted features
crawling	speed, cast, crabspeed	$\leq 5\%$	ampl, freq, dur
head cast	cast, midline, morpwidth	$\leq 1\%$	ampl, dur, efreq
hunch	midline, cast, morpwidth	$\leq 5\%$	ampl, dur, efreq
rolling	crabspeed	$\leq 1\%$	ampl, dur, efreq
digging	x, y, midline		efreq
following	x, y		ftime, flength
taxis	tailvecvx, tailvecy, cast, headx, heady, x, y		nind, ...

$$nind = (d_{max} + d_{min} - 2 \cdot < d >)/(d_{max} - d_{min}) \quad (2)$$

This formula can be applied both to settings with a flat gradient of stimulus (by simply replacing d for x) and to settings with (multiple) point source(s) of stimulus. In the limiting case of an animal with strongly positive taxis, which spent all the time of experiment at a minimal possible distance from the source of stimulus, so that $d(t) = d_{min}$, formula (2) gives $nind$ = +1. In the opposite case of an animal with strongly negative taxis which spent all the experimental time at a maximal possible distance from the source of stumulus, it gives $nind$ = -1. A further adjustment of this formula is possible for non-flat geometries, by taking into account a variation in the experimental space available for animals located at different distances d from the closet source of stimulus.

4 HYPOTHESES TESTING AND HIT DETECTION

To discriminate the behaviors of the wild type and mutant animal groups, we compare their feature distributions/histograms. Our null hypothesis is that the differences between the (possibly, noisy) histograms are statistically insignificant, so that the groups belong to the same population. Whenever evidence is found that the two distributions are different enough so that the groups cannot belong to the same population, we say there is a ”hit”.

The hit detection functionality of SALAM package is currently under development. At this time, the software is only capable of detecting hits based on analysis of one feature at a time. The particular approach to hit detection varies depending on the type of a feature. For per-animal features, e.g. the presence or absence of roll actions in a given animal trace, we employ the statistical tests available for proportions, notably Pearson's χ^2-test or Fisher's exact test. To compare the mean values of features with continuous distributions, binary parametric (t-test) or non-parametric (e.g. Wilcoxon rank sum) tests are available. If the distribution approximately meets a multivariate normality requirement, even more accurate, multivariate test can be used, as illustrated below with an example of crawling speed amplitude. This approach, in addition to the mean, involves comparison of several other characteristics of the distribution, all of them being computed based on statistical moments: the standard deviation, the skewness and the kurtosis. The approach borrows its formalism from Hotelling's theory for T^2-distributions (see, e.g. (Rencher, 2002)), but allows empiric calibration of thresholds, based on existing wild type data, rather than using theoretical thresholds derived for strictly multivariate normal distributions.

Let $U^{(i)}$ be a vector of characteristics (mean, stddev, skewness, kurtosis) representing a histogram of i-th run of a wild type population, $\bar{U} = (1/n) \cdot \sum U^{(i)}$ be the mean of all the wild type runs, and $\hat{\Sigma} = (1/(n-1)) \cdot \sum (U^{(i)} - \bar{U})(U^{(i)} - \bar{U})^T$ be sample (4×4) covariance matrix. Then, given a desired confidence level α, one can build an ellipsoid representing a Mahalanobis vicinity of $\bar{U}$

$$(U^{(i)} - \bar{U})^T \hat{\Sigma}^{-1} (U^{(i)} - \bar{U}) \leq D_{thr} \quad (3)$$

with the threshold distance D_{thr} being calibrated so as to include exactly $(1-\alpha)$ portion of all the wild type runs in the Mahalanobis vicinity. Now, for any run of a mutant population with representative vector of characteristics located within the Mahalanobis vicinity, we say that it is undistinguishable from the wild type population with confidence α, i.e. the null hypothesis is met. Conversely, any mutant run with characteristics located outside this vicinity will represent a hit (Fig. 7).

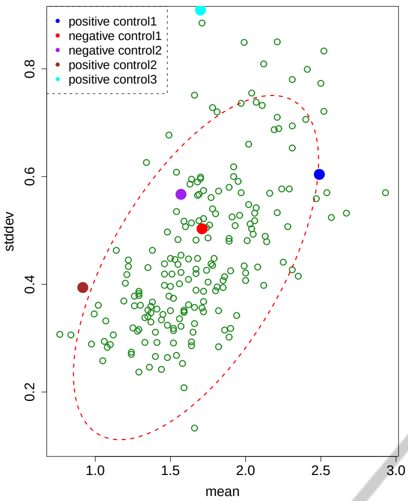

Figure 7: Multivariate hit detection procedure applied to crawling speed amplitude. Shown in the figure are two out of the four used variables: the mean and the standard deviation. The open circles represent the individual runs of a wild type population, with each the run comprising approx. 100 animals. Multivariate normality of this data was confirmed by the test implemented in R package ”energy”. The dashed ellipse, drawn using R package ”chemometrics”, represents a 90%-confidence Mahalanobis vicinity of the center of wild type runs, so that $\alpha = 0.1$. The filled circles represent pooled data for two negative control lines and three positive control lines. A hit is detected whenever a filled circle falls outside the Mahalanobis vicinity, so this chart confirms biological expectations for the indicated confidence level.

5 CONCLUSIONS

We have provided an overview of SALAM, a software package for Statistical Analysis of Larval Motions. The input signals taken by the software, the detected behavioral actions and the processing workflow have been illustrated graphically. The algorithms used for data processing have been discussed. Detection of the two most typical behavioral actions, the crawling and the head cast, as well as a multivariate procedure for detection of hits, have been illustrated with examples. This software is being used as a part of an automated data analysis pipeline for screening mutant lines in order to determine the function of genes and neurons.

ACKNOWLEDGEMENTS

The authors are grateful to Prof. Daniel Naiman for consultation on statistical data analysis.

REFERENCES

Branson, K., Robie, A. A., Bender, J., Perona, P., and Dickinson, M. H. (2009). High-throughput ethomics in large groups of drosophila. *Nature Methods*, 6:451–457.

Gershow, M., Berck, M., Mathew, D., Luo, L., Kane, E. A., Carlson, J. R., and Samuel, A. D. (2012). Controlling airborne cues to study small animal navigation. *Nature Methods*, 9:290–296.

Gomez-Marin, A. and Louis, M. (2012). Active sensation during orientation behavior in the drosophila larva: more sense than luck. *Curr. Opin. Neurobiol.*, 22:208–215.

Gomez-Marin, A., Stephens, G. J., and Louis, M. (2011). Active sampling and decision making in drosophila chemotaxis. *Nature Communications*, 22:441.

Luo, L., Gershow, M., Rosenzweig, M., Kang, K., Fang-Yen, C., Garrity, P. A., and Samuel, A. D. (2010). Navigational decision making in drosophila thermotaxis. *J. Neurosci.*, 30:4261–4272.

Ohyama, T., Jovanic, T., Denisov, G., Hoffmann, D., Kerr, R. A., and Zlatic, M. (2012). Automated high-throughput analysis of stimulus-evoked ethograms in drosophila larva reveals modality-specific escape strategies. *Submitted for publication*.

Pfeiffer, B. D., Jenett, A., Hammonds, A. S., Ngo, T. T., Misra, S., Murphy, C., Scully, A., Carlson, J. W., Wan, K. H., Laverty, T. R., Mungall, C., Svirskas, R., Kadonaga, J. T., Doe, C. Q., Eisen, M. B., Celniker, S. E., and Rubin, G. M. (2008). Tools for neuroanatomy and neurogenetics in drosophila. *Proc. Natl. Acad. Sci. USA*, 105:9715–9720.

Rencher, A. (2002). *Methods of Multivariate Analysis*. John Wiley and Sons, Inc., second edition.

Robie, A. A., Straw, A. D., and Dickinson, M. H. (2010). Object preference by walking fruit flies, drosophila melanogaster, is mediated by vision and graviperception. *J. Exp. Biol.*, 213:2494–2506.

Swierczek, N. A., Giles, A. C., Rankin, C. H., and Kerr, R. A. (2011). High-throughput behavioral analysis in c. elegans. *Nature Methods*, 8:592–598.