

Tracking the Swimming Motions of *C. elegans* Worms with Applications in Aging Studies

Christophe Restif and Dimitris Metaxas

Rutgers, the State University of New Jersey
christophe.restif@centraliens.net

Abstract. Quantitative analysis of the swimming motions of *C. elegans* worms are of critical importance for many gene-related studies on aging. However no automated methods are currently in use. We present a novel training-based method that automatically tracks and segments multiple swimming worms, in challenging imaging conditions. The position of each worm is predicted by comparing its latest motion with a set of previous observations, and then adjusted laterally and longitudinally to fit the image. After segmentation, a variety of measures can be used to assess the evolution of swimming patterns over time, allowing a quantitative comparison of worm populations over their lifetime. The complete software is being evaluated for mass processing in biology laboratories.

1 Introduction

The nematodes *C. elegans*, short for *Caenorhabditis elegans*, have been the focus of numerous studies in the biology of aging [1]. Their motion patterns can be studied to analyze the evolution of their strength and coordination, when crawling on solid surfaces and swimming in liquid. To assess the influence of stochastic and genetic factors in aging, quantitative individual-based studies are needed, over the complete life-span of large populations of worms, based on non-invasive measures [2]. Optical microscopy is a promising tool, and many computer vision methods have been published to study the worm crawling motion [3, 4, 5]. Yet to the best of our knowledge, no satisfactory equivalent for swimming motion is in use. In this article we present a complete automated tool to track swimming worms on videos in challenging conditions, extract measures from the segmentation, and use them to compare individuals or mutant populations.

Tracking multiple swimming worms is a challenging computer vision task, and significantly different from crawling worms. Although optical microscopy is a method well adapted to the biological constraints of worm imaging, the images produced are not optimal for segmentation and tracking. Videos typically show a dozen worms, about four pixel wide, all with similar appearances and behaviors (see Fig. 1). The worms frequently overlap, and they need to be tracked robustly. Generic segmentation methods adapted to tracking, such as dynamic graph cuts [6] or level sets [7], are not robust enough in those conditions. Also, the following features restrict the use of explicit motion models. As invertebrate, swimming worms can show a large variety of body shapes, and change over only

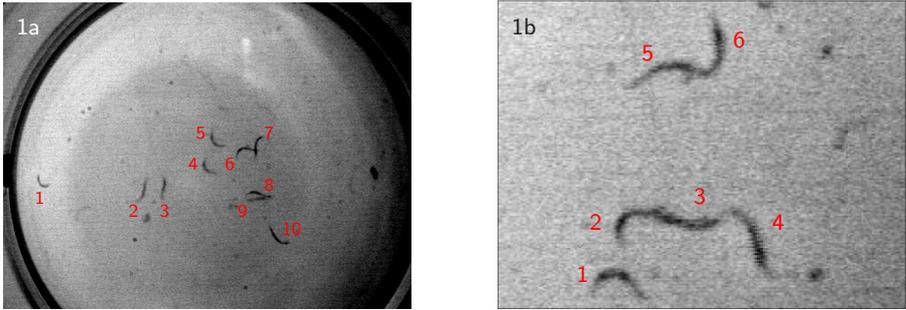


Fig. 1. A: typical frame showing ten worms. Notice the variations of background intensity and texture, the variations of worm appearances, and the partial overlaps. b: close-up of another frame showing six overlapping worms.

a few frames. They are evolving in three dimensions and imaged from above. As a result their length can visually change suddenly as a result of a body bend. Finally, in large-scale studies, the motion patterns vary significantly over time and over individuals, from vigorous and flexible to stiff and uncoordinated. Such motion patterns are to be the output of the tracking, and the use of explicit specific worm motion models for tracking [3, 8] would induce a high bias.

In contrast, crawling worms move in two dimensions. They can be in contact but less frequently overlap. They also exhibit very different motions. From frame to frame, most of their body keeps the same location, only the head and tail move; and they do not show the high-frequency body bending, informally described as alternate I-C and S-Z shapes, typical of swimming. Also, as their motions are slower and more limited in space, they can be imaged at higher resolution and with better contrast. Hence, intensity thresholding and binary morphological operations can be enough for segmentation [4, 5, 3], but cannot be used in our context. The work closest to ours is [9], which addresses the tracking swimming motion. Yet it is a semi-automated tool, requiring manual input on each video on worms and background, and not fast enough for large-scale processing.

We present a complete automated tracking and segmentation method, for an unlimited number of worms. For tracking, we use the most relevant exemplars from the training set to predict each worm’s position on the next frame, and adjust the prediction with segmentation methods based on standard motion detection techniques and worm-specific appearances. We model a worm as a central body line and body borders, so that subsequent measures such as the frequencies of head bending, body shapes, or total body curvature, are straightforward to compute after the segmentation. Statistics on worm populations are then computed and can be used for further phenotype analysis.

2 Method

Worm model. We model the central body line of a worm w at time t as an open polygon with n vertices $B^w(t) = \{B_i^w(t)\}_{1 \leq i \leq n}$, and use a constant width

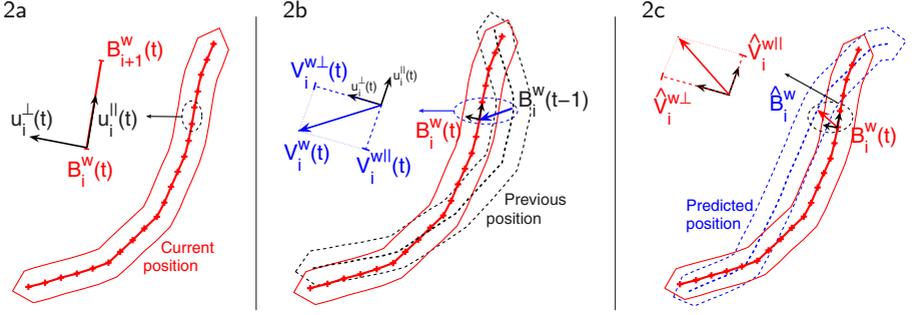


Fig. 2. Worm model for tracking prediction. 2a: Base vectors $u_i^{\parallel}(t)$ and $u_i^{\perp}(t)$ at vertex $B_i^w(t)$. 2b: velocity $V_i^w(t)$ and its decomposition $V_i^{w\parallel}(t)$ and $V_i^{w\perp}(t)$ from Eq. (2). 2c: predicted position computed with $\hat{V}_i^{w\parallel}(t+1)$ and $\hat{V}_i^{w\perp}(t+1)$ from Eq. (4).

of 2 pixels around $B^w(t)$ to delineate the worm body. At each vertex $B_i^w(t)$, we define the local base vectors $(u_i^{\parallel}(t), u_i^{\perp}(t))$ as the unit vectors parallel and perpendicular to the edge $(B_i^w(t), B_{i+1}^w(t))$, dropping the w in the notation for clarity, as illustrated in Fig. 2a, and define for the last vertex $(u_n^{\parallel}(t), u_n^{\perp}(t)) = (u_{n-1}^{\parallel}(t), u_{n-1}^{\perp}(t))$. This model follows from the morphology of the worms, which are symmetric along their central body lines; their flexibility are reflected by the evolution of the relative angles between consecutive vertices. With the local bases used, movements along $u_i^{\parallel}(t)$ corresponds to body stretching and compressions, and along $u_n^{\perp}(t)$, to body side-bendings. When body lengths vary and vertices are added or removed during segmentation, we resample $B^w(t)$ to $n = 100$ regularly-spaced vertices. From frame $t - 1$ to frame t , we define the velocity of each vertex i of worm w and decompose it in the local base as (see Fig. 2b.):

$$V_i^w(t) = B_i^w(t) - B_i^w(t-1) = V_i^{w\parallel}(t) \cdot u_i^{\parallel}(t) + V_i^{w\perp}(t) \cdot u_i^{\perp}(t) \quad (1)$$

We then define the $2n$ -dimensional *velocity vector* of worm w at time t as:

$$V^w(t) = \left(V_1^{w\parallel}(t), \dots, V_n^{w\parallel}(t), V_1^{w\perp}(t), \dots, V_n^{w\perp}(t) \right) \quad (2)$$

Given a worm w segmented at times $t = 1 \dots t_0$, we compute the worm's position on the next frame $t_0 + 1$ in two steps. First, we predict its position by predicting its next velocity vector, noting those predictions $\hat{B}^w(t_0 + 1)$ and $\hat{V}^w(t_0 + 1)$. Then we adjust $\hat{B}^w(t_0 + 1)$ to the contents of frame $t_0 + 1$, and obtain the final segmentation $B^w(t_0 + 1)$.

Tracking prediction. The swimming motion of worms consists of high-frequency body bends, where a worm takes alternative S, I, Z, C, and O shapes, reflected in $V^{w\perp}$, and lower-frequency body stretching and contracting, reflected in $V^{w\parallel}$. Although V^w exhibits oscillatory patterns over time, strictly speaking it is non-periodic. Modeling its evolutions on the long term with a periodic function is not

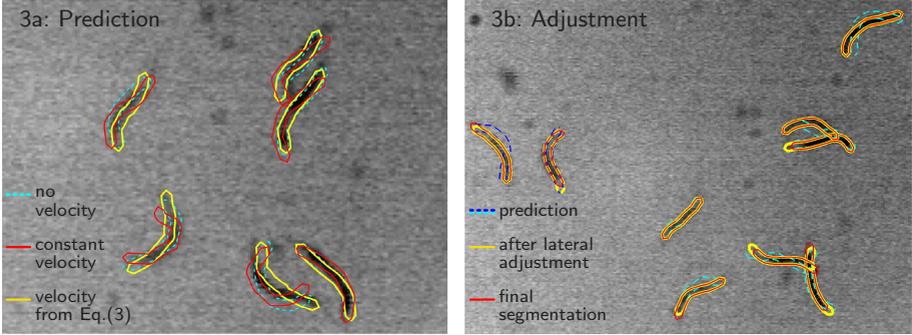


Fig. 3. Prediction and adjustment of the segmented worms. 3a: comparison of three models for the predicted position; in dotted blue, $\hat{V}^w(t_0 + 1) = 0$; in red, $\hat{V}^w(t_0 + 1) = V^w(t_0)$; in yellow, $\hat{V}^w(t_0 + 1)_i$ from Eq. (3). 3b: adjustments; in dotted blue, the predicted position from Eq. (4); in red: the result of the lateral adjustment by gradient ascent; in yellow: the final result after longitudinal adjustment.

reliable. Instead we use a non-parametric model for the short-term evolutions of V^w , making use of its regularity. $V^w(t_0)$ is compared to all previous observations $V^w(t)$, for $1 \leq t \leq t_0 - 1$, using the L^2 distance. Intuitively, the closer $V^w(t_0)$ is to a particular $V^w(t)$, the closer the prediction $\hat{V}^w(t_0 + 1)$ will be to $V^w(t + 1)$. Instead of selecting only the closest vector $V^w(t)$, we use all the previous observations. $\hat{V}^w(t_0 + 1)$ is defined as a weighted average of the successors $V^w(t + 1)$. Let $d^w(t) = \|V^w(t) - V^w(t_0)\|$ using L^2 norm, the predicted velocity vector is defined as: (the starting value $t = 0$ is explained in the following paragraph)

$$\hat{V}^w(t_0 + 1) = \frac{\sum_{t=0}^{t_0-1} e^{-d^w(t)} \cdot V^w(t + 1)}{\sum_{t=0}^{t_0-1} e^{-d^w(t)}} \quad (3)$$

This approach shares similarities with particle filtering. However, we use all the previous observations of a worm's velocity vector as samples, as opposed to a probability-based random selection. To account for the noise in this model, we have included two terms in the computation of the average: $V^w(0) = V^w(1) = 0$. In case none of the previous velocity vectors are close to $V^w(t_0)$, the zero vector $V^w(1)$ included in Eq. (3) lowers the predicted velocity. This keeps the predicted position closer to the previous position: in effect, this reduces the significance of the prediction in case the model is not adapted. Noting the components of the predicted velocity vector $\hat{V}^w(t_0 + 1)$ as $\hat{V}_i^{w\parallel}(t_0 + 1)$ and $\hat{V}_i^{w\perp}(t_0 + 1)$, as done in Eq. (2), the predicted position of vertex i is:

$$\hat{B}_i^w(t_0 + 1) = B_i^w(t_0) + \hat{V}_i^{w\parallel}(t_0 + 1) \cdot u_i^{\parallel}(t_0) + \hat{V}_i^{w\perp}(t_0 + 1) \cdot u_i^{\perp}(t_0) \quad (4)$$

as illustrated in Fig. 2c. Examples of predictions are shown in red in Fig. 3a, with two alternative models shown in blue and green, respectively $\hat{V}^w(t_0 + 1) = 0$ and $\hat{V}^w(t_0 + 1) = V^w(t_0)$. This predicted body position now has to be adjusted to the frame $t_0 + 1$.

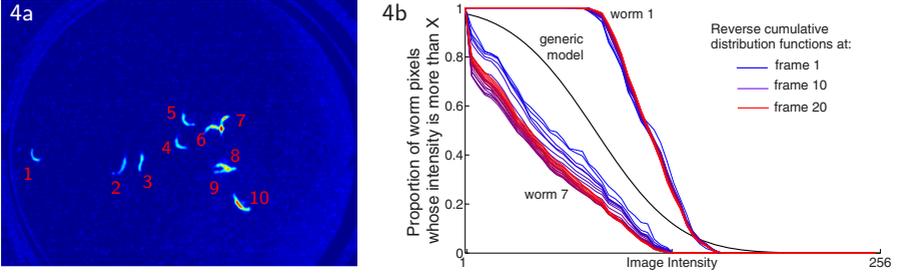


Fig. 4. Image preprocessing and worm appearances. 4a: result of the pre-processing in Eq. (5) on the image shown in Fig. 1a, using the generic appearance function. 4b: evolutions of the appearance of two worms. In black, the generic appearance function used for the first frame, computed from the training data. From blue to red: the reverse cumulative mass functions of intensities for two worms, shown for 20 successive frames.

Tracking adjustment. To overcome the image quality issues described in the introduction, images are pre-processed as follows. For each pixel (x, y) on 40 consecutive images, the median $m(x, y)$ and the standard deviation $\sigma(x, y)$ of the intensity distribution are computed. Thus for each image I , two maps are defined: the difference to the median $I - m$ and the standard deviation σ . The latter is very robust to the flickering intensity noise affecting the background, and only enhances pixels where worms are moving, have moved or will move, but is common for the 40 consecutive images. The former map is less robust to noise, but is image-specific, and enhances pixels where worms are currently moving. A third map is computed, to account for each worm’s specific appearance $\mathcal{A}^w(i)$, which is defined as the proportion of pixels in worm w whose intensities are less than i . Intuitively, the darker a pixel, the more likely it is to belong to a worm, yet that likeliness depends on the worm (compare worms 1 and 7 in Fig. 1). Initially, a generic appearance function is used for all worms, shown in black in Fig. 4b; it is defined as the reverse cumulative function of a Gaussian, whose mean and standard deviation are computed from the training set using the histograms of manually segmented worms. Then, the reverse cumulative histogram of each worm is computed and incremented at each frame, and normalized accordingly, as illustrated for two worms in Fig. 4b in colors. This third map $\mathcal{A}^w(I(x, y))$ is multiplied with the other two to define the pre-processed image $\hat{I}^w(x, y)$:

$$\hat{I}^w(x, y) = \sigma(x, y) \times (I(x, y) - m(x, y)) \times \mathcal{A}^w(I(x, y)) \quad (5)$$

To initialize the tracking, the first frame is pre-processed using the generic appearance function, as shown on Fig. 4a. A greedy line-growing method is then used on this image to fit the central body lines of worms, until all the regions of a value higher than a training-based threshold can be used to segment worms 4 pixel wide and at least 20 pixel long. In successive frames, once each worm’s position is predicted using Eq. 4, only a subpart of the frame is pre-processed, 20 pixels around the predicted position, using Eq. 5 with the relevant $\mathcal{A}^w(i)$ function. The adjustment of the predicted position \hat{B}^w on this pre-processed map is

Table 1. Evaluation of segmentation results against ground truth on 400 worms

| <i>Model</i> | <i>True positive</i> | <i>Average distance</i> | <i>Lost in tracking</i> |
|-----------------------------------|----------------------|-------------------------|-------------------------|
| $\hat{V}^w(t_0 + 1) = 0$ | 77.4 % | 3.4 pixels | 17 % |
| $\hat{V}^w(t_0 + 1) = V^w(t_0)$ | 84.7 % | 2.5 pixels | 12 % |
| $\hat{V}^w(t_0 + 1)$ from Eq. (3) | 92.4 % | 1.1 pixels | 4 % |

performed in two steps. First, the position is adjusted laterally: each vertex \hat{B}_i^w is moved along its local base vector u_i^\perp by gradient ascent on $\hat{I}^w(x, y)$. Then, longitudinally, the length of the worm is adjusted to fit $\hat{I}^w(x, y)$. Extreme vertices 1 and n are removed until they stand on regions of \hat{I}^w above a threshold defined from the training set, and new vertices are added at the extremities with the same greedy line-growing method as during initialization. The two steps of the adjustment are shown in Fig. 3b.

Measures. As the result of the segmentation, the central body line of a worm, defined as an open polygon, can be used to measure various features related to the worm’s configuration and coordination. Here we focus on three of the measures. Let $\alpha_i^w(t)$ be the relative signed angle between the edges $(B_{i-1}^w(t), B_i^w(t))$ and $(B_i^w(t), B_{i+1}^w(t))$, measured between $-\pi$ and π . We define the head, the tail, and the absolute curvatures of worm w at time t as follows:

$$\mathcal{HC}^w(t) = \sum_{i=2}^{n/2} \alpha_i^w(t) \quad , \quad \mathcal{TC}^w(t) = \sum_{i=n/2}^{n-1} \alpha_i^w(t) \quad , \quad \mathcal{AC}^w(t) = \sum_{i=2}^{n-1} |\alpha_i^w(t)| \quad (6)$$

The terms ‘head’ and ‘tail’ are convenient to refer to each end of a worm, yet there is no guarantee that they correspond to the actual head and tail of the worm. Their evolution over time can still be used to estimate the number of actual head and tail bends. \mathcal{HC}^w is positive when the head is turned to the left, and negative for the right. As a worm bends its head from side to side, \mathcal{HC}^w oscillates between positive and negative values. The number of these oscillations over time is used to compute the average frequency of head bending over the entire video. The frequency of tail bending is defined similarly, and the higher of these two frequencies is considered as the actual worm head bend frequency. The absolute curvature \mathcal{AC}^w indicates the vigor of a worm’s motion. Although it oscillates over time, its distribution over the entire video sequence indicates the amplitude of bending achieved by the worm, and how much time it spent in each curvature. Results are presented below.

3 Results and Discussion

In the data set we used, images are either 640×480 or 700×520 , and the frame rate is either 24 or 30 frames per second. The tracking and segmentation results were evaluated using a testing set of 400 manually segmented worms. We measured the true positive rate, the average distance between central body lines,

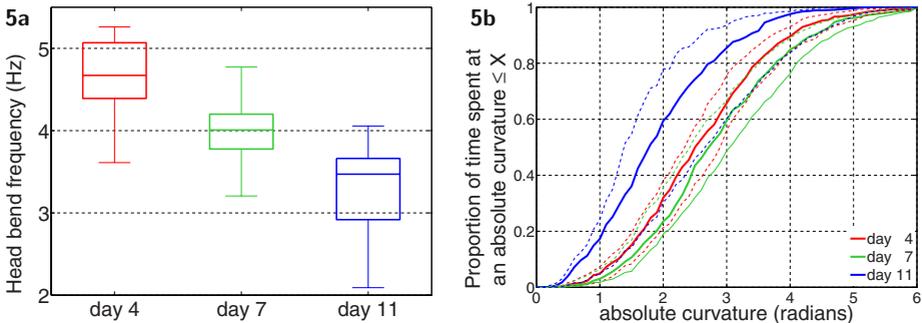


Fig. 5. Evolution of the swimming patterns of a population of mutant worms over their lifetime. 5a: head bend frequency. In each plot, the central line shows the median value of the population, the top and bottom of the box show the lower and upper quartiles, and the external lines show the extent of the remaining values. 5b: Cumulative distribution functions of the absolute curvature. For each color, the central thick line shows the median value, and the top and bottom dotted lines show the lower and upper quartiles. See page 40 for details on the measures' definition.

and the proportion of worms lost during tracking. We compared the method presented in the previous section with two other models: $\hat{V}^w(t_0 + 1) = 0$ (assuming no velocity) and $\hat{V}^w(t_0 + 1) = V^w(t_0)$ (assuming constant velocity). Results are listed in Table 1. Our method improves all three results significantly. It is to be noted that because of the imaging conditions, the manual labeling uncertainty is about one pixel around the worms' borders. Thus, the average distance to the ground truth obtained with our method, 1.1 pixels, is very accurate given the image resolution. A frame with 10 worms is segmented in 1.2 seconds on average, on a dual 3-GHz PC running Matlab 2006, which brings the complete processing of a 900-frame video under 20 minutes, and allows large-scale processing. The main current limitation of our method is that the worms to be tracked are fixed by the results of the line-growing method on the first frame.

Finally, we used our program to assess quantitatively the evolution of swimming patterns as worms age. Five different mutant worm populations were imaged at three stages of their adult lives, at day 4, 7, and 11. As a result, 30 videos, each showing 10 worms over 900 frames, were produced and processed with our program. After tracking and segmentation, the measures described above were generated. Two examples of resulting graphs for one mutant population are shown in Fig. 5. The head bend frequencies are shown to decrease gradually for all individuals. Also, the range of frequencies at day 11 is wider than at younger ages, suggesting that individuals of that mutant population age differently, some keeping more active than others. Fig. 5b shows the proportion of time spent at certain absolute curvatures. It appears that worms have similar trends on day 4 and 7, while on day 11, they spend less time at higher curvatures, showing an increase of body stiffness. The same remark on the wider range of values at day 11 applies, indicating that stiffness affects some individuals more than others at that age. Although these trends were observed qualitatively by the biologists providing the data, our program allowed

the automatic generation and display of quantitative results. Similar graphs were generated to compare individual worms, the same worms at different ages, and different mutant populations. Further analysis of those results is a subject for studies on the biology of aging, and is beyond the scope of this article.

4 Conclusion

We presented an automated method to track and segment the swimming motions of *C. elegans* worms under challenging imaging conditions, which have not been addressed despite their significance for research, in particular on the biology of aging. The method is based on a two-step processing. First each worm's position is predicted, by comparing its latest motion to a set of previous observations. Then this prediction is adjusted, sideways and longitudinally, so that the segmented worm fits a region of high intensity variations, indicating motion, and intensity similar to its previous intensity distributions. Based on this segmentation, quantitative measures are defined to analyze the evolution of the vigor and swimming patterns of worms. Such analysis is possible at the level of an individual worm, of a population, and across populations. The complete software is currently being used by biologists for research on the aging of *C. elegans*.

The authors gratefully thank Prof. Monica Driscoll and her team, in particular Dr Carolina Ibañez-Ventoso and Mehul Vora, for providing the data and for their useful comments. This work is funded by NIH grant number 4-22567.

References

1. Kenyon, C.: Environmental factors and gene activities that influence life span. In: *Caenorhabditis elegans*, vol. II. Cold Spring Harbor Laboratory Press (1997)
2. Herndon, L.A., Schmeissner, P.J., Dudaronek, J.M., Brown, P.A., Listner, K.M., Sakano, Y., Paupard, M.C., Hall, D.H., Driscoll, M.: Stochastic and genetic factors influence tissue-specific decline in ageing *C. elegans*. *Nature* 419, 808–814 (2002)
3. Fontaine, E., Barr, A.H., Burdick, J.W.: Model-based tracking of multiple worms and fish. In: *ICCV Workshop on Dynamical Vision*, Rio, Brazil (October 2007)
4. Geng, W., Cosman, P., Berry, C.C., Feng, Z., Schafer, W.R.: Automatic tracking, feature extraction and classification of *C. elegans* phenotypes. *IEEE Trans. on Biomedical Engineering* 10(51), 1811–1820 (2004)
5. Roussel, N., Morton, C.A., Finger, F.P., Roysam, B.: A computational model for *C. elegans* locomotory behavior: Application to multiworm tracking. *IEEE Trans. on Biomedical Engineering* 54(10), 1786–1797 (2007)
6. Kohli, P., Torr, P.H.: Dynamic graph cuts for efficient inference in Markov Random Fields. *IEEE Pattern Analysis and Machine Intelligence* 29(12), 2079–2088 (2007)
7. Cremers, D., Osher, S., Soatto, S.: Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. *Int. J. Comp. Vis.* 69(3), 335–351 (2006)
8. Brackenbury, J.: Swimming kinematics and wake elements in a worm-like insect: the larva of the midge *Chironomus plumosus*. *Journal of Zoology* 260, 195–201 (2003)
9. Tschepnakis, G., Bianchi, L., Driscoll, M., Metaxas, D.: Tracking *C. elegans* populations in fluid environments for the study of different locomotory behaviors. In: *Microscopic Image Analysis with Applications in Biology*, Piscataway, NJ (September 2007)