373. Tracking Social Insects

Ashok Veeraraghavan Rama Chellappa
Department of Electrical and Computer Engineering
University of Maryland
College Park, Md 20770

Abstract

Behavior analysis of insects has provided significant insight into understanding evolution and has even led to advances in several fields like control systems, flight navigation etc. Significant time and effort has been invested in following trajectories of insects in videos and manually labelling these insect motions. In this paper, we propose certain general principles that help in automatic tracking and behavior analysis. We also apply these principles for the specific case of tracking a bee in a hive. We estimate the position of the bee and the orientation of the various body parts of the bee in the video. The tracking algorithm is based on a particle filter. We also use a Markovian motion model on the behavior of the bees to incorporate both tracking and behavior analysis in a single unified framework. Experimental results for tracking bees performing the waggle dance are provided.

1. Introduction

Behavorial research in the study of the organisational structure and communication forms in social insects like the ants and bees has garnered tremendous impetus in recent years [1][2][3][4]. Such a study of social insects, apart from providing us with insights about evolution, has also, in several instances provided us with practical models and means for tasks like work organisation, reliable distributed communication, navigation etc [5][6]. Usually, when such an experiment to study these insects is setup, the insects in an observation hive are videotaped. The hours of videotape is then manually studied and hand-labelled. This task of manually labelling the video data takes up the bulk of the time and effort of such experiments. In this paper, we discuss general methodologies for automatic labelling of such videos and provide an example of tracking a bee in a bee hive.

We aim to build a system that will assist researchers in behavorial research study and analyse the motion and behavior of insects. Such an automated system speeds up the analysis of video data obtained from experiments significantly and also prevents manual errors in the labelling of data. Moreover, parameters like the orientation of the various body parts of the insects(which is of great interest to the behavioral researcher) can be automatically extracted in such a framework. But, such a system also has certain limitations. It requires the technical input of a behaviorioal researcher (who would be the end user) regarding the type of behaviors exhibited by the insect being studied. Also, such a system would not be capable of seamlessly generalizing like humans do. For example, if the insect under observation were to exhibit behaviors that have never been observed, (and therefore the system has no knowledge of these behaviors) then, though the system would be able to detect this, it would be left to the researcher to actually analyse this anamolous behavior.

1.1 Prior Work

There has been significant work on tracking objects through video. Most tracking methodologies can be classified as either deterministic or stochastic. Deterministic approaches solve an optimization problem under a prescribed cost function[7][8]. Stochastic approaches estimate posterior distribution of the position of the object in the current frame using a Kalman filter [9][10] or particle filters [11][12][13][14][15][16][17]. Most of these, do not directly adapt well to tracking insects because they exhibit very specific forms of motion(for example, bees can turn by a right angle within 2 or 3 frames). In order to extend such tracking methods, it is important to consider the physical nature (body parts) of these insects and incorporate both their physical nature and the nature of their motions in the tracking algorithm. [17] has proposed the use of a rao-blackwellized particle filter for tracking and has shown results on tracking bees. But this method cannot extract relevant physical parameters like the orientation of the abdomen with respect to the thorax etc. The use of Hidden Markov Models(HMM) for cateogarizing bee movement into dancer, follower and active hive worker has been proposed in [18]. But this method requires the trajectories of bees as input. In this paper, we present a method of automatically tracking and cateogarizing insect behaviors from videos.

The next section describes some general principles that assist while designing algorithms to track and cateogarize insect motions. These general principles include modeling the physical structure and range of motions of these insects. Section 3,4 and 5 discuss the application of these general principles to the specific problem of tracking bees in a hive. Specifically, Section 3 discusses the shape and motion model of a bee in a hive while section 4 presents the tracking algorithm based on a particle filter. Section 5 provides the experimental results. Section 6 provides the summary and conclusions of this work.

2. Tracking insects:General Principles

Tracking insects in a video is a very challenging task because of several important factors. Firstly, we are usually interested in tracking insects in their hives where there will be hundreds or thousands of other insects of the same species. It is a very demanding task to track a particular individual insect in the presence of others that are very similar in appearance. Secondly, the range of motions exhibited by insects is very diverse and differs from species to species. Also, insects are capable of making surprisingly fast movements where naive motion models like the constant velocity model might not prove sufficient. Moreover, social insects are often capable of communicating to each other, and therefore there are strong interactions between different insects. Though different species of insects differ in the nature of their actual motions, the problems encountered in automatic tracking of these insects is similar in nature across species. Therefore, in this section we develop some general strategies that assist in reliable tracking and behavior analysis of insects.

2.1 Anatomical model

Modeling the anatomy of insects is very important for reliable tracking, because the structure of their body parts and their relative positions present some physical limits on their possible relative orientations. Thus the range of motions exhibited by insects can be easily captured by an anatomically correct model of insects. Moreover, their anatomical structure also results in significant correlations among the orientations of their various body parts. These can also be captured by such structural models. Also, the structural similarity between insects implores algorithms to exploit this structural similarity.

Inspite of their great diversity, the anatomy of all insects is surprisingly very similar. An insect body has a hard exoskeleton protecting a soft interior. The body is divided into three main parts- the head, thorax and the abdomen each of which is in turn divided into several smaller segments. Figure 1 shows the image of a bee, an ant and a beetle. Though there are individual differences in their body structure, the three main parts of the body are evidently visible. Each

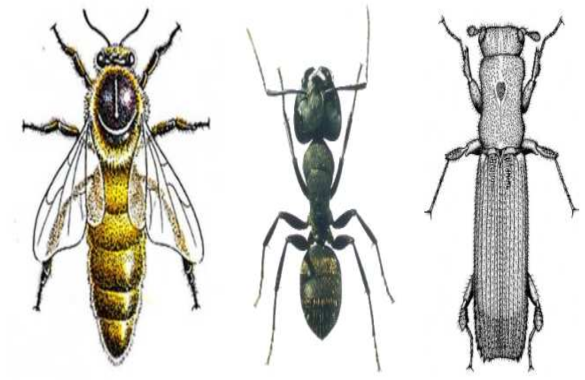

Figure 1: A Bee, an Ant and a Beetle

of these three parts can be regarded as rigid body parts for the purposes of video based tracking. The interconnection between parts provide some physical limits for the relative movement of these parts. Most insects also move towards the direction of their head. Therefore, during specific movements such as turning, the orientation of the abdomen usually follows the orientation of the head and the thorax with some lag. Such interactions between body parts can be easily captured using a structural model for insects.

In this paper, we use three ellipses for modeling these three body parts(refer Figure 4). Each ellipse is assumed to be rigid. Each of the two joints (head-thorax and thorax-abdomen) has one rotational degree of freedom. The wings of insects have been ignored in this model. We will discuss the specific structural model in detail in Section 3.

2.2 Behaviorial model

Insects, especially social insects like bees and ants, exhibit extremely rich behaviors. These behaviors are dependent on the insect species, the nature of work they are currently performing, their surroundings etc. For example, foraging bees, when they return to the hive execute a certain form of dance called waggle dance as a means of communication with other bees. We will see more about this in Section 3 and 4. The foraging patterns of ants has also been studied in detail and explicit modeling of this foraging pattern can significantly help in tracking ant colonies. Modeling such behaviors explicitly go a significant distance in accurate and robust tracking. Moreover, explicitly modeling such behaviors also leads to algorithms where position tracking and behavior analysis are tackled in a unified framework.

Several algorithms use motion models (like constant velocity model, random walk model etc) for tracking[12][17][19][16]. We propose the use of behaviorial models for the problem of tracking insects. Such behavioral models have been used for the problem of identifying behaviors from tracked trajectories in [18]. The difference

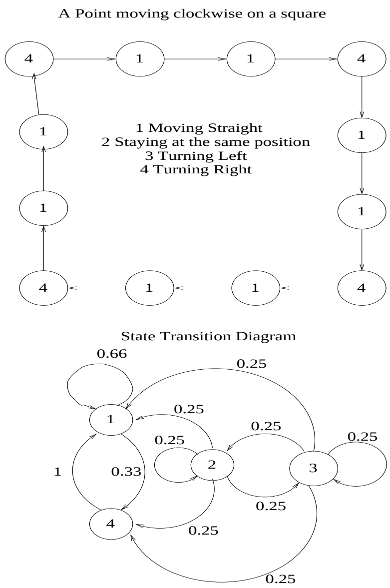

Figure 2: A point moving clockwise on a square and the corresponding Markov behavior model.

between motion models and behavioral models is the range of time scales at which modeling is done. Motion models typically model the probability distribution (pdf) of the position in the next frame as a function of the position in the current frame. Instead, behavioral models capture the probability distribution of position over time as a function of the behavior that the tracked object is exhibiting. We believe that the use of behavioral models presents a significant layer of abstraction that is able to capture the variety and complexity of the motions exhibited by the insects.

In this paper, we advocate the use of Markovian models to model behavior. We define probability distributions for some basic motions such as moving straight ahead, turning left, turning right and hovering at the same location. These basic motions are regarded as states and behaviors are modelled as being Markovian on this motion state space. Figure 2 shows an illustration of a point object moving on the perimeter of a square and shows how a Markov model may be built to capture the structure of such a motion. Such a Markovian model for behavior does not capture all the structure present in certain specific behaviors, yet we believe that it has the ability to capture most of the structure present. Moreover, it can also generalize to motions that have previously not been seen but similar to any of the behaviors modelled. Once each specific behavior has been modelled as a Markov process, then our tracking system can simultaneously track the position and the behavior of insects in videos. Learning these behavior models can be done from example trajectories of each behavior provided by the expert.

2.3 Modeling Interactions

Behavioral models capture the nature of the behaviors of individual insects. But, as has been pointed out earlier, social insects have very strong interactions among the behaviors of individuals in a colony. For example, when a foraging bee returns to perform a waggle dance, there are other bees who follow the dancing bee without waggling (refer Figure 5. Modeling the interactions between behaviors of different individuals will definitely improve both tracking performance and also help in behavior analysis. Moreover, there is another subtle but important reason why modeling these interactions is important. In insects, behavior is often a means of communication [1]. Behavioral researchers are interested in extracting correlations between behaviors of individuals and thereby learn how behaviors lead to communication amongst insects. Modeling interactions may help us recognize some strong correlations that have not yet been noticed (and therefore not modelled) which is one of the goals of behavioral analysis.

These are some general principles that improve tracking performance and also have certain other benefits for the behavioral researcher. Structural modeling implies that the state space in which the insects position is being tracked has direct relationship with the actual body parts. Behaviorial modeling ensures that tracking and behavior analysis can be performed in a unified framework while modeling interactions is essential to study social insects. In the next section we will take the specific problem of studying the behavior of a bee in a hive and apply these principles to the problem. We have not explicitly modelled the interactions between different bees in the current implementation.

3. Bee in a hive

In this section, we will discuss a very specific problem in behavioral analysis of insects and provide a solution for this problem following the general principles outlined in the previous section. Honeybee is an insect that can communicate with other honeybees through behaviors [1] and also exhibit social characteristics like division of labor. These make the honeybee a very interesting subject for analysis. We have been able to simultaneously track the position and the orientation of the various body parts of the bee. The relative orientation of the body parts is important for behavioral analysis. Moreover, we were able to perform both tracking and behavioral analysis in a unified framework.

3.1 Shape Model

Bees like all other insects are made of three parts, -the head, the thorax and the abdomen. There are individual differences between the various types of bees. Figure 3 shows similarities and differences in the anatomy of a drone, queen

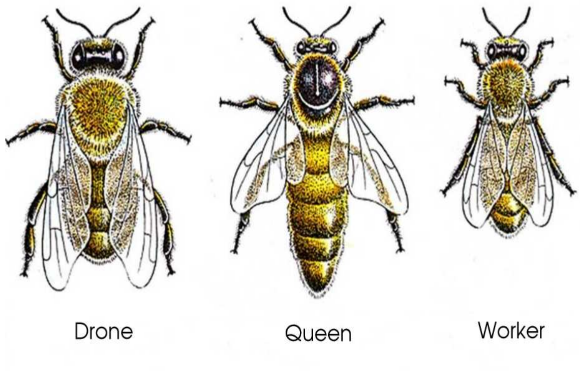

Figure 3: A Drone, A Queen bee and a Worker

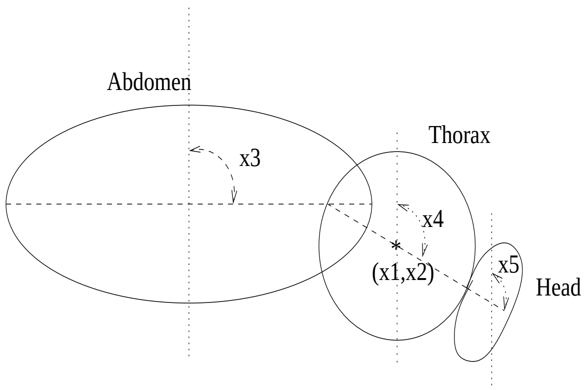

Figure 4: The shape model of a Bee

and a worker. We note that all of these bees can be adequately modelled by modelling the abdomen, thorax and head as rigid parts with interconnections.

The shape model of the bee is made of three ellipses, one for each body part. Since during flight the legs are barely visible, neglecting them will not affect results. Wings are visible in certain frames during flight, but they beat at such a high rate compared to frame rate of capture that it would be difficult to ascertain their position in a video frame. Therefore, we have neglected the effect of wings in our shape model. Figure 4 shows the shape model of a bee. The dimensions of the various ellipses are fixed during initialization. Currently the initialization for the first frame is manual. It consists of clicking two points to indicate the enclosing rectangle for each ellipse. Automatic initialization in such videos with hundreds of bees is a challenging problem in itself and is outside the scope of our current work.

The location of the bee and its parts in any frame can be given by five parameters- namely, the location of the center of the thorax(2 parameters), the orientation of the head, the orientation of the thorax and the orientation of the abdomen (refer Figure 4. Tracking the bee over a video essentially amounts to estimating these five model parameters($\mathbf{X} = [x_1\ x_2\ x_3\ x_4\ x_5]'$) for each frame. In our current approach we have assumed that the actual sizes of these ellipses do not change with time. This would of course be the case as long as the bee remains at the same distance from the camera. Since the behaviors we attempted to study in our work (like the waggle dance) are performed on a vertical plane along the beehive, the bees projected the same part sizes during the entire length of video captures. Nevertheless, it is very easy to incorporate the effect of distance from the camera in our shape model, by introducing a scale factor as one more parameter in our state space.

3.2 Behavioral model

We model the probability distributions of location parameters **X** for certain basic motions($m_1 - m_4$). In our implementation, we modelled four different motions- 1) Moving straight ahead, 2) Turning, 3) Waggle, and 4) Motionless. We use appropriate Gaussian pdfs (p_{m1}, p_{m3}, p_{m4}) for the straight, waggle and motionless while we use a mixture of two Gaussians for modeling the turning motion.

$$p_{mi}(\mathbf{X}_t) = X_{t-1} + N(\vec{\mu}_{mi}, \Sigma_{mi}); fori = 1, 3, 4. \quad (1)$$

$$p_{m2}(\mathbf{X}_t) = X_{t-1} + 0.5(N(\vec{\mu}_{m2}, \Sigma_{m2}) + N(-\vec{\mu}_{m2}, \Sigma_{m2})) \quad (2)$$

The means $(\vec{\mu_{m1}} - \vec{\mu_{m4}})$ and the covariance matrices $(\Sigma_{m1} - \Sigma_{m4})$ were chosen carefully in order to account for the physical limitations imposed by the structural model and to account for the strong correlation between the orientation of the various body parts.

Each behavior is now modelled as a Markov process on these motions, i.e., each behavior is characterized by a prior probability for each of these motions and a transition probability matrix between these motions. For illustration we will discuss the manner in which the waggle dance is modelled. Foraging bees, when they return to the hive perform a certain 'waggle dance' in order to recruit other bees and inform them about this feeding location. This waggle dance, as a form of communication, has interested researchers for a long time [1] [20]. Figure 5 shows the trajectory followed by a bee during a single run of the waggle dance. It also shows some followers who follow the dancer but do not waggle. As can be seen in the figure, this dance is characterized by the central waggling portion which is immediately followed by a turn, a straight run another turn and a return to the waggling section. The bee visits the two sides of the waggle section alternatively. The returning forager communicates the distance of the food source, the direction and the sweetness of the nectar to the other recruits. The angle of the waggle portion of the dance (waggle axis) from the vertical indicates the orientation of the food source from the sun. The intensity of the dance serves as a measure of the sweetness and abundance of nectar in the food source,

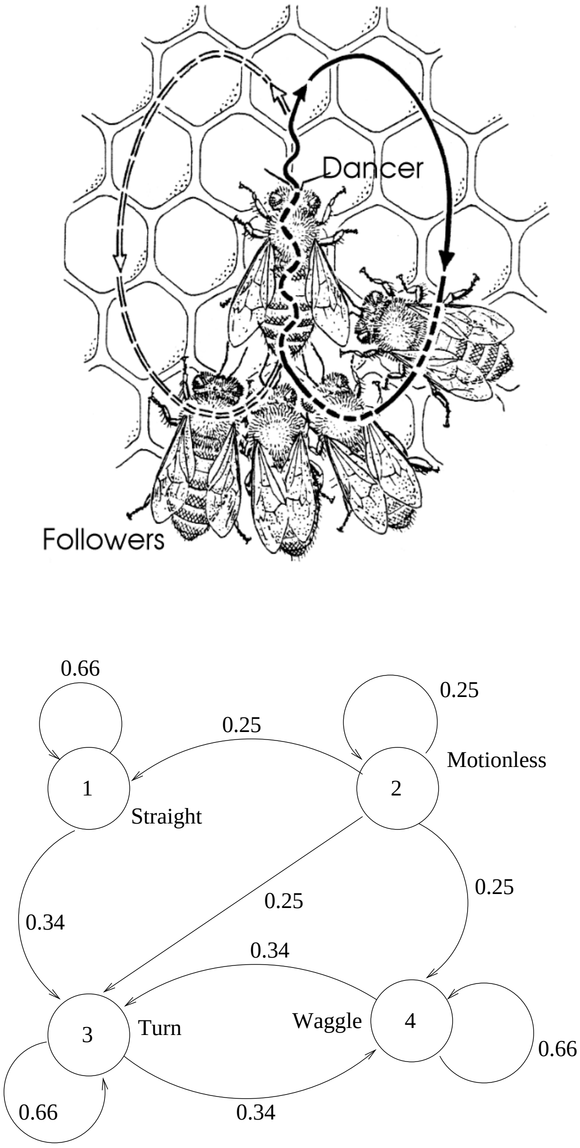

Figure 5: A Bee performing a waggle dance and the Behavioral model for the Waggle dance

while the frequency of waggling indicates the distance of the food source. Refer [1] for details regarding the waggle dance as a means of communication. We modelled this waggle dance as a Markov model. A typical Markov model for the waggle dance is also shown in Figure 5.

4. Shape and Motion encoded Particle Filter

We address the tracking problem as a problem of estimating the state X_1^t given the observations Z_1^t. Since both the state transition model and the observation model are non-linear, filtering procedures like the Kalman filter are inadequate. This Bayesian estimation of the state given observations can be done recursively using particle filters. Particle filters [13][11][12] provide a method of recursively updating the posterior pdf $P(X_t/Z_1^t)$ as a set of N weighted particles $\{X_t^{(i)}, \pi_t^{(i)}\}_{i=1}^N$.

4.1 Particle Filter

The particle filter provides a method for recursively estimating the unknown state **X**, from a collection of noisy observations $\mathbf{Z}_1^t$. We employ a state space model for performing such an estimation. The state parameters(**X**) to be estimated are the position and orientation of the bee in the current frame. The observation is the appearance of the bee ($\mathbf{Z}_t$). The state transition and the observation models are given by,

$$\textit{State Transition Model: } X_t = F_B(X_{t-1}, N_t) \quad (3)$$

$$\textit{Observation Model: } Z_t = G(X_t, W_t) \quad (4)$$

where, N_t is the system noise and W_t is the observation noise. The state transition function F_B characterizes the state evolution for a certain behavior B. In usual tracking problems, the motion model is used to characterize the state transition function. In our current algorithm, the behavioral model described in Section 3.2 is used as the state transition function. Therefore, the state at time t, (X_t) depends upon the state at the previous frame (X_{t-1}), the behavioral model and the system noise. The observation function G models the observation (Z_t) as a function of its current position(state X_t) and observation noise. Once such a description for the state evolution has been made, the particle filter provides a method for representing and estimating the posterior pdf $P(X_t/Z_1^t)$ as a set of N weighted particles $\{X_t^{(i)}, \pi_t^{(i)}\}_{i=1}^N$. Then the state X_t can be estimated either as the maximum likelihood estimate or as the minimum mean square error estimate as given below,

$$\hat{X}_t^{ML} = \arg\max_{X_t} \pi_t^{(i)} \quad (5)$$

$$\hat{X}_t^{MMSE} = E(X_t/Z_1^t) = N^{-1}\sum_{i=1}^{N}\pi_t^{(i)}X_t^{(i)} \qquad (6)$$

The complete algorithm for the estimation of the posterior pdf as a set of weighted particles is given below

1. **Initialize** the tracker with a sample set according to a prior distribution $p(X_0)$.

2. **For** Frame $= 1, 2, ...$

   (a) **For** sample $i = 1, 2, 3, ...N$

   - **Resample** $\{X_{t-1}^{(i)}\}$ to obtain $\{X_t^{(i)}\}$.
   - **Predict** sample $X_t^{(i)}$ by sampling from $F_B(X_{t-1}^{(i)}, N_t)$ where F_B is a Markov model for behavior estimated in the previous frame.
   - **Compute Observation** i.e., compute the appearance of the Abdomen, the Thorax and the Head in the current frame for each particle $Z_t^{(i)} = G(Y_t, X_t^{(i)})$
   - **Compute Weights** for the particle using the likelihood model i.e., $\pi_t^{(i)} = p(Z_t^{(i)}/X_t^{(i)})$.

   (b) **Normalize** the weights using $\pi_t^{(i)} = \pi_t^{(i)}/\sum_{i=1}^{N}\pi_t^{(i)}$ so that the particles represent a probability mass function.

   (c) **Estimate** the ML or MMSE estimate of the state X_t. using the particles and their weights.

   (d) **Estimate** the behavior of the bee using a ML estimate from the various behavior models as $B = \arg\max_j P(X_1^t/B_j)$., where $B_j for j = 1, 2, ..$ indicate the behaviors modelled.

4.2 Prediction and Likelihood Model

In typical tracking applications it is customary to use motion models for prediction [12][17][19][16]. As has been discussed previously, we use behavioral models in addition to motion models. Behavioral models are capable of capturing both the nature of motions exhibited by the bee and also incorporate the limitations imposed by the structure of the bee on the relative orientation of its body parts. Therefore, the use of such models for prediction improves tracking performance significantly.

In the algorithm for the shape and motion encoded particle filter, we have not yet described the exact nature of the likelihood model used for updating the weights for the predicted samples. In our current implementation, we have manually selected five frames from the video and extracted the appearance of the bee (i.e., the intensity within the Abdomen, Thorax and Head) for these five frames. Since the appearance of the bee changes drastically over the video sequence, we use an appearance model consisting of multiple color exemplars($A_1, A_2, .., A_5$). The RGB components of color are treated independently and identically. The appearance of the bee in any given frame is assumed to be Gaussian centered around one of these five exemplars. Therefore, given the appearance of a particle ($Z_t^{(i)}$), the weight for the particle ($\pi_t^{(i)}$) is updated as

$$\pi_t^{(i)} = p(Z_t^{(i)}/X_t^{(i)}); \qquad (7)$$

$$p(Z_t^{(i)}/X_t^{(i)}) = \min_{j=1,2,..,5} p(Z_t^{(i)}/X_t^{(i)}, A_i); \qquad (8)$$

$$p(Z_t^{(i)}/X_t^{(i)}, A_i) = L(Z_t^{(i)}; A_i, \sigma^2) \qquad (9)$$

where, $L(Z_t^{(i)}; A_i, \sigma^2)$ is a normal density with mean A_i and diagonal covariance matrix $\sigma^2 * I_{3dX3d}$ (3d since there are d pixels in the appearance model and the RGB components are treated independently).

At each frame we also need to estimate the current behavior exhibited by the bee. We do a maximum likelihood estimation for the same ($B = \arg\max_j P(X_1^t/B_j)$.). The maximum likelihood estimate is done in two steps. For each frame, the basic motion exhibited by the bee is estimated as $\arg\max_i p(\hat{X}_t/\hat{X}_{t-1}, M = m_i)$ where m_i represents the various basic motions modeled. Once this is done the behavior is estimated from the estimates of basic motions for each frame as $B = \arg\max_j P(M_1^t/B_j)$., where $B_j for j = 1, 2, ..$ indicate the behaviors modelled.

5. Results

We conducted tracking experiments on two video sequences of bees in a hive. In both video sequences there was a forager performing the waggle dance.We tracked the dancer in both videos. Similar results were obtained for both video sequences. In all our simulations we used 300 to 600 particles. Both sequences were over 250 frames long. We were able to track the dancer for the entire length of the video sequence without any missed tracks. Moreover, we were able to extract parameters like the orientation of the various body parts during each frame. We used these parameters to automatically identify the waggle portion of the dance. We also verified this estimate manually and found it to be robust and accurate. Automatic extraction of the orientation of the abdomen during the waggle dance is important because the waggle portion of the dance encodes information such as the distance and the direction of food source. Specifically, the orientation of the waggle axis indicates the direction of the food source.

Figure 6 shows the structural model of the tracked bee superimposed on the original image frame. The results are best viewed in color since the tracking algorithm had color images as observations. The figure shows the top five tracked particles(blue being the best particle and red being the fifth best particle). The first row shows frames 20,34 and 45 from the first video sequence. The subsequent three rows show some frames from the second video sequence. As is apparent from the sample frames the appearance of the dancer varies significantly within and across the videos. There is also significant occlusions in some frames. Frame 34 of video sequence 1 and frames 170,172 and 187 of video sequence 2 show the ability of the tracker to maintain track during partial occlusions. In fact, in frame 172(video sequence 2), occlusion forces the posterior pdf to become bimodal (another bee in close proximity). But we see that the track is regained when the bee emerges out of occlusion in frame 175. In frame 187(video sequence 2), we see that the thorax and the head of the bee is occluded while the abdomen of the bee is seen. Therefore the estimate of the abdomen is very precise (all five particles shown indicate the same orientation of abdomen). Since the thorax is not seen we see that there is high variance in the estimate of the orientation of the thorax and the head. Structural modeling has ensured that inspite of occlusion, only physically realizable orientations of the thorax and the head are maintained. Frames 30-34(video sequence 2) show the bee executing a waggle dance. Notice that the abdomen of the bee waggles from one side to another. We were able to estimate the direction of the waggle axis from the estimates of the orientation of the bee during the waggle dance. This axis of this waggling is useful since it encodes information about the direction of the food source.

6. Summary and Conclusions

In this paper, we have presented some general principles for efficiently tracking insects in their hives. We have also applied these principles in the framework of a shape and motion encoded particle filter to perform reliable tracking of bees in a hive. The problem of tracking a bee in its hive presents significant challenges like occlusion, clutter, drastic movements and a background that is very similar to the foreground. The structural and behavioral model cast in a shape and motion encoded particle filter is able to tackle such difficulties. In this paper, we have not modelled the interactions between the various bees in the hive. We are currently looking at methods to model these interactions while keeping the computational load tractable.

References

[1] Von Frisch, *The Dance Language and orientation of bees*, Cambridge MA:Harvard University Press.

[2] C.T. David, "Compensation for height in the control of groundspeed by drosophila in a new, 'barbers's pole' wind tunnel," *Journal of computational physiology*, 1982.

[3] M.V. Srinivasan, S.W. Zhang, M. Lehrer, and T.S. Collett, "Honeybee navigation en route to the goal: visual flight control and odometry," *The Journal of Experimental Biology*, vol. 199, pp. 237–244, 1996.

[4] H. Esch and J.E. Burns, "Distance estimation by foraging honeybees," *Journal of Experimental Biology*, vol. 199, pp. 155–162, 1996.

[5] T.R. Neumann and H.H. Bulthoff, "Insect inspired visual control of translatory flight," *Proceedings of the 6th European Conference on Artificial Life ECAL 2001*, pp. 627–636, 2001.

[6] F. Mura and N. Franceschini, "Visual control of altitude and speed in a flight agent," *Proceedings of 3rd International Conference on Simulation of Adaptive Behaviour: From Animal to Animats*, pp. 91–99, 1994.

[7] G.D. Hager and P.N. Belhumeur, "Efficient region tracking with parametric models of geometry and illumination," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 20, pp. 1025–1039, 1998.

[8] D. Comaniciu, V. Ramesh, and P. Meer, "Real-time tracking of non-rigid objects using mean-shift," *Conference on Computer Vision and Pattern Recognition*, vol. 2, pp. 142–149, 2000.

[9] T.J. Broida, C. Shekhar, and R. Chellappa, "Recursive techniques for the estimation of 3-d translation and rotation parameters from noisy image sequences," *IEEE Transactions on Aerospace and Electronic systems*, vol. AES 26, pp. 639–656, 1990.

[10] A. Azarbayejani and A. Pentland, "Recursive estimation of motion structure and focal length," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 17, pp. 562–575, 1995.

[11] A. Doucet, N.D. Freitas, and N. Gordon, *Sequential Monte Carlo methods in practice*, Springer-Verlag, New York, 2001.

[12] M. Isard and A. Blake, "Contour tracking by stochastic propogation of conditional density," *European Conference on Computer Vision*, pp. 343–356, 1996.

[13] N.J. Gordon, D.J. Salmond, and A.F.M Smith, "Novel approach to non-linear/non-gaussian babayesian state estimation," *IEE Proceedings on Radar and Signal Processing*, vol. 140, pp. 107–113, 1993.

[14] M. Isard and A. Blake, "Icondensation: Unifying low-level and high-level tracking in a stochastic framework," *European Conference on Computer Vision*, vol. 1, pp. 767–781, 1998.

[15] J.S. Liu and R. Chen, "Sequential monte carlo for dynamical systems," *Journal of the American Statistical Association*, vol. 93, pp. 1031–1041, 1998.

[16] S.K. Zhou, V. Krueger, and R. Chellappa, "Probabilistic recognition of human faces from video," *Computer Vision and Image Understanding*, vol. 91, pp. 214–245, 2003.

[17] Z. Khan, T. Balch, and F. Dellaert, "A Rao-Blackwellized particle filter for eigen tracking," *Conference on Computer Vision and Pattern Recognition*, 2004.

[18] A. Feldman and T. Balch, "Automatic identification of bee movement using human trainable models of behavior," *Mathematics and Algoritms of Social Insects*, Dec 2003.

[19] M.J. Black and A.D. Jepson, "A probabilistic framework for matching temporal trajectories: Condensation based recognition of gestures and expressions," *European Conference on Computer Vision*, vol. 1, pp. 909–924, 1998.

[20] M.V. Srinivasan and S.W. Zhang, "Visual motor computations in insects," *Annu. Rev. Neurosci. 2004*, vol. 27, pp. 679–696, 2004.

Figure 6: Sample Frames from two tracked sequences of a bee in a beehive. (Shows the top five particles in each frame) Top Row:Video Sequence 1: Frame Numbers 20,34 and 45. Subsequent Rows:Video Sequence 2: Frame Numbers 30,31,32,33,34,170,172,175,187.