Motion Recognition of Physical Agents Constrained to Two Dimensions

Rande Shern RANDE@CS.CMU.EDU
Center for Automated Learning and Discovery, Carnegie Mellon University

Abstract

We have devolped a method for the recognition of a physical agent's behavior while observing its two dimensional motion. Our data is acquired via a vision based tracking system and converted into an encoded sequence of observations. Based upon these observations we select a model that characterizes the agents current behavior. Hidden Markov models for this purpose. Our test subjects are social insects and we are able to select between different models of behavior with a Bayesian inference scheme.

1. Introduction

Social insect behavior has become an interesting paradigm in computer science as it offers a wealth of ideas on how to organize computer software systems. Information technology people utilize these ideas to develop new algorithms for such things as the coordination and control multi-agent systems. Robotics researchers interested in the design and implementation of robust mutli-robotic systems can also draw deep insights from social-biology. Most of the work thus far has been to take models developed by biologists and apply them to computer science systems. In this work we are trying to do something different. We are using computer science and robotic techniques in an attempt to help the biologists generate characteristic models of social insects.

Our system is centered around an ant colony made up of the species *Camponotus festinatus* and the first task is to be able to determine what an individual ant is doing while watching it with a vision system. This paper describes the design of a *Recognizer* capable of taking visual information as input and identifying behavioral characteristics of an ant, a physical agent, as it moves.

The vision system returns a streaming sequence of coordinate positions of an ant. From this and other information about the environment (domain knowledge) we group together data into meaningful features that encode an abstraction of an ant's motion. Our features encode aspects of the physical nature of an ant's motion including velocity and current bearing as well as aspects that can be measured relative to the environment such as the relative position to other objects. This new sequence is the observed trace sequence which is the input of the *Recognizer*.

We are treating the recognition of ant behaviors as a statistical signal processing problem. Our assumption is that the observed trace sequence can be well characterized as a parametric random process. Hidden Markov Models (HMM) are the chosen statistical models. The basic theory of hidden Markov models is well documented.

The task is to calculate the probability of a Behavior Model given the observations. When we have multiple behaviors to choose from we pick the posterior with the highest probability. It should be noted that the denominator, the probability of the observations, is a constant and ignoring it simplifies the calculations.

$$Pr(Behavior|Observation) = \frac{Pr(Observation|Behavior)Pr(Behavior)}{Pr(Observation)} \quad (1)$$

The calculation of $Pr(Behavior|Observation)$ has two pieces. $Pr(Behavior)$ is the Ant Model. An ant has a finite number of behaviors that we are able to recognize. The Ant Model tells us a-priori the probability that an ant may be in anyone of these behavior modes. The other piece, $Pr(Observations|Behavior)$, is the Behavior Model. It defines which sequence of observations characterize a particual Behavior. Each model of a behavior has a signature entailed by some sequence of observations and different behaviors will have different signatures. Thus, by calculating the probability of an observed ant trace given a model we determine how close the observations match that model. If two or more traces score high in probability for more than one behavor the behavior that the ant would a prior be more likely to be executing is chosen. For example, if two different behaviors are defined by a trace sequence that happen to be identical over a certian period, then multiplying by the probability of the ant model, $P(Behavior)$, the decision is made based on prior knowledge of the animals habits.

The Behavior Model is the trestle of the recognition system. By continually feeding observations exctraced from the tracking system to it we can calculate $P(Behavior|Observations)$ for some time period, T. Determinging the proir probability distribution of the possible ant behaviors is very difficult. At the moment, we are making the assumption of a flat prior, so that each behavior is equally likely in the Ant Model.

By using the above methods we are able to select between two different ant behavioral models using a simulated ant. The ant is moving in a rectangular arena with designated areas for a food source, a nest enterence, and regions associated with these objects so the ant's current heading can be calculated.

2. Approach

The approach we use to build a *Recognizer* capable of selecting between multiple behavior models has two phases. The first part is to extract features from the visual input stream of the ant moving around the environment. The visual data is comprised of time-stamped (x,y) positions. From this we calculate the ant's velocity and bearing along various parts of its trace. The second part is to use the trace sequence that has been generated and predict which model best explains the visual input. This is done by calculating the probability of a sequence taken over all the models in the system using Bayes' Rule and choosing the one with the best score. Each part will be described below.

2.1 Observation Sequence

The vision tracking system can identify and track multiple ants marching over a one-meter square area. This data can also be logged to a file for later analysis. We have to manipulate the raw data to extract features from it so that we can pass an observation sequence on to the classifier. The ant data is a series of time-stamped (x,y) coordinates. Each ant is given a unique ID number. Since the data are (x,y) coordinates on a plane it is natural to choose a vector representation for each time step that the ant moves. This vector will be 'from' the previous time step 'to' the current time step. It is possible to calculate the straight Euclidean distance between any two points, and since there is a time associated with each step the length of this vector corresponds directly to the velocity with which the animal moved during that increment. As a result, a single ant moving across the surface of the lab space has its motion decomposed into a chain of tiny vectors. Each beginning at one point and terminating at the next point. Given any two vectors in the sequence there is an angle between them. This angle is the ant's *Turning Angle*. As the ant moves about the environment a series of *Turning Angle* are generated. Therefore, the raw (x,y) data gives us two series of numbers (signals) we can use to create features with the *Turning Angle* and the veloctiy vector between every point.

The tracking system will also let us identify points in the environment that correspond to stationary objects. These items may be food, such as sugar water, or the enterence to the nest. As an ant patrols an area another observable, dynamic quantity of motion is current bearing and absolute position. Both of these quantities have to be measured with respect to the environment. Two more features are extracted from these. The event that the ant is heading towards a specific object and also the duration it stays on that bearing.

(missing figures)

In summary, the features used for the classification presented here are the current bearing which is split into different classes based upon different objects in the environment (e.g. heading toward the food source.) The animal's absolute position in the environment is also split into different classes based upon particular regions of interest (e.g. the nest enterence.)

2.2 Behavior Model Recognition

As the animal moves around the observing arena it will be engaged in goal directed action. In order to carry out this action it must move around on its legs and interact with the world with manidbles and antenna. By watching what the animal has been doing in the past, and what it is doing now, we can match what is being observed to a probabilistic model of ant behavior. We are using hidden Markov Models for our pattern matching approach.

2.2.1 Hidden Markov Models

Since the 1970s the theory of HMMs have usually been discussed in the context of speech recognition systems. In this paper our use of HMMs is similar to the basic speech recognition process. In that process, after recording, a speech signal is divided into pieces and processed so that a sequence of predifined catagories is generated. From this sequence the recognizer then has to determine which sequence of sounds where spoken, and as a result recognize which words where spoken.

A hidden Markov model is a double embedded stochastic process. It has an underlying stochastic process that is not directly observable, it is hidden, and can be observed only through another stochastic processes that produce the sequence of observations. A hidden Markov model is defined as:

$$\lambda = (A, B, \pi)$$

where A is the state-transition probability distribution A $= \{a_{ij}\}$. B is the observation symbol probability B$=\{b_j(k)\}$ with $b_j(k) = Pr(oberation_k|state_j)$. π is the initial state distribution $\pi = \{\pi_i\}$.

Given this parameterization we see that this defines a probability for a given sequence of observations, O, i.e. $Pr(O|\lambda)$. The calculation of this measure is the basis used to select between multiple ant behavior models. Our method sees this as $Pr(Observations|Behavior)$.

2.2.2 Forward Algorithm

Our basic evaluation problem when given a model and a sequence of observations is how do we compute the probability that the observed sequence was produced by an HMM? View this as scoring how well a given model matches the observation trace of an ant. The most straightforward way of doing this is through enumerating every possible state sequence of length T (the number of observations). However, the computational overhead of this approach makes this task daunting. A more efficient way to perform the calculation $Pr(O|Behavior)$ is by use of the Forward Algorithm. This algorithm is an induction rule. It uses the insight that because there are only a finite number of states in an HMM, calculating the probability of being in a state at a particular time only involves the probabilities of the possible transitions into it from other states at the previous time step.

The calculation is done as follows. Consider the forward variable $\alpha_t(i)$ defined as

$$\alpha_t(i) = Pr(\mathbf{o}_1\mathbf{o}_2 \ldots \mathbf{o}_t, q_t = i|\lambda) \tag{2}$$

that is the probability of the partial observation sequence $\mathbf{o}_1\mathbf{o}_2 \ldots \mathbf{o}_t$ (until time t) and state i at time t, given the model λ. We can solve for $\alpha_t(i)$ inductively, as follows:

1. Initialization

$$\alpha_1(i) = \pi_i B_i(\mathbf{o}_1), \qquad 1 \le i \le N. \tag{3}$$

2. Induction

$$\alpha_{t+1}(j) = \left[\sum_{i=1}^{N} \alpha_t(i)a_{ij}\right] b_j(\mathbf{o}_{t+1}) \qquad \begin{matrix} 1 \le t \le T-1 \\ 1 \le j \le N \end{matrix} . \tag{4}$$

3. Termination

$$Pr(\mathbf{O}|\lambda) = \sum_{\mathbf{i=1}}^{\mathbf{N}} \alpha_{\mathbf{T}}(\mathbf{i}) \tag{5}$$

Step 1 intializes the forward probabilities as the joint probability of state i and inital observation $\mathbf{o}_1$. The induction step is the heart of the calculation. It is by this method that we calculate the $Pr(Observations|Model)$.

2.2.3 Model Selection

We may never know the true plans that are inside an ant's brain, however Myrmecologists have made extensive study of their behavior at the individual and colony level. Multiple behavior models for an individual ant have been designed using HMMs as the representation. The task of the *Recognizer* is to select between these models. To do this we use a consistent classifier. We select the model, at any given time, that has the highest probability from all model classes. The probability is calculated by equation (1), $Pr(Behavior_i|Observations)$.

A trace window, T, of n observations is taken as a sliding window along the ant's path. These n observations compose the observation trace sequence and are the Observation in the probability calculations.

$$\text{Behavior } = \max_{i \in L}[Pr(Behavior_i|Observations)] \tag{6}$$

L is the set of all behavior models.

3. Results

Having a computer system that can observer social insects and automatically model their behaviors and collect data to support new theories would be make a huge impact on the way biology is done in the world today. The development done so far is a step in that direction. At this point we can watch an insect crawl around and from only that data determine with some probability what behavior model that insect is performing. The system does choose between different models by selecting the model with highest probability. The validity of the model is still in question however, and that answer will not be answered until we can use supervised learning to train the system.

The way we evaluate the current system is with empirical measures. We have humans watch what the ant is doing and at what the system is doing and look for an association between them. We have two models in the system at the moment. The first model is what we are calling the foraging model. This model is defined by the ant being at the nest enterence, then wandering, then at a food source, and then returning to the nest enterence. The second model is of a patrolling ant. The model is defined by the ant moving through regions out in the foraging area. The work so far is seen by us as a first step to completing a larger system.

(missing figures)

4. Related Work

Multi-Agent Modeling is the process of building models of multiple interacting agents based upon observations of these agents over time. Such models allow us to identify the state of agents and the organizations of which they are a part. There has been work done in adversarial domains such as robot soccer where the need to infer the opposing teams current tactical intent is useful in planning a correct response in a robot soccer game. There have been interesting advancements in the area of gesture recognition, but this work is being done on humans and not insects.

Ants perform a variety of tasks outside their nests. The one that comes readily to mind is foraging. Ants leave their nest in search of food and once they have located a food item they maneuver their way back to their nest. Usually, at the same time, there are scouts outside the nest walking around to detect items of interest for the colony or fighting to defend it. Trying to determine the intentions of the insects are what some biologists are working on when they study social insects. Previous work done in the mid 1990's have the researchers out walking around on the desert floor watching what ants do during their daily activities and logging them in a notebook. Other researches have used a video camera in a laboratory and filmed ants in their foraging arenas, and then later went back and point-by-point annotated the ant's movements so that mathematical calculations could be done.

What we are trying to do is something different. We want a computer to automatically track and characterize what it is seeing. Currently, our method is to pick a behavior model based on observation that the system is seeing. This is a step toward have a system automatically collect data which may be useful to biology and other systems sciences. Computers will be much faster at collecting data and will have the advantage of being more accurate once a system is tuned.

5. Conclusion

We have developed a method to recognize a simple behavior for a physical agent moving in two dimensions. We are treating the recognition of ant behaviors as a statistical signal processing problem. Our assumption is that the observation trace sequence can be well characterized as a parametric random process. We are using hidden Markov models for our statistical models. A lot of the work so far in the gap between computer science and biology has been to take models developed by biologists and apply them to computer science systems. In this work we are trying to do something different. We are using computer science and robotic techniques in an attempt to help the biologists observe social insects.

The observations, or signal, are collected from a vision system. The raw data is convereted into a sequence of features which is the input of a classifier. The classifier selects between multiple behavior models by calculating the posterior probabilities of observation sequences given a hidden Markov model. The first part of the probability calculation is done with the forward algorithm.

To determining the best Behavior Model we have to deal with two problems. Given a set of observations and a model how do we best choose the corresponing state sequence to encode the desired behavior, and how do we adjust the parameters of an HMM to maximize $Pr(Observed|Behavior)$. The Ant Model, $P(Behavior)$, is the proir probability distribution of the possible ant behaviors we can recognize with our system. Determining this is very difficult. At the moment, we are making the assumption of a flat prior, so that each behavior is equally likely. By collecing data of ants performing the behaviors we want to encode we can use supervised learning techniques, EM, and the Viterbi Algorithm, to find the single best parameterized state sequence. This will be the major subject of future work. Also, for our future work we are looking into clustering a data base of ant traces to make a better attempt at estimate a prior probability distribution.

6. Acknowledgements

The author would like to thank Tucker Balch, Sebastian Thrun, and Tom Minka for their guidance and input in the creation of this work.

This work was funded by DARPA under the Control of Agent Based Systems Program, and by Cargenie Mellon and the Robotics Institute.

7. References

Balch, T. *A Computer Vision-based TestBed for Observing and Modeling Insect Behavior*. The Multi-Robot Lab, Carnegie Mellon University

Balch, T., Khan, Z., Veloso, M (2001). *Automatically Tracking and Analyzing the Behavior of Live Insect Colonies* To appear in Autonomous Agents.

Bonadeau, E., Theraulaz, G., Fourcassie, V., Deneubourg, J. *The phase-ordering kinetics of cemetary organization in ants* Santa Fe Institute.

Bruce, J., Balch, T. and Veloso, M. *Fast and inexpensive color image segmentation for interactive robots* In Proc. IROS-2000.

Collins, Lipton, Kanade, Fujiyoshi, Duggins, Tsin, Tolliver, Enomoto, and Hasegawa. (2000) *A system for video surveillance and monitoring* Techinical Report CMU-RI-TR-00-12 Robotics Institute, Carnegie Mellon University, 2000.

Dorigo, M., Caro, G. (1999). *The and colony optimization meta-heuristic* New Ideas in Optimization, pages 11-32. McGraw-Hill.

Durbin, R., Eddy, S., Krogh, A., Mitchison, G. (1989) *Biological sequence analysis: Probabilistic models of proteins and nucleic acids* Cambridge University Press.

Goel, P., Zellner, A. *Bayesian Inference and Decision Techniques* North-Holland, 1986.

Gordon, D. (1999) *Ants at Work* The Free Press, New York.

Gordon, D., Adler, F (1992) *Information collection and spread by networks of patrolling ants* The American Naturalist Vol. 140, No. 3

Gordon, D. (1995) *The expandable network of ant exploration* Animal Behaviour, 50, 4.

Galata, A., Johnson, N., Hogg, D. (1999) *Learning Structured Behaviour Models Using Variable Length Markov Models* IEEE.

Haefner, Crist (1994) *Spatial Model of Movement and foraging in harvester ants*

Han, K., Veloso, M. (1999) *Automated Robot Behavior Recognition* In Proc. of ISSR, pages 199-204. Snowbird, Utah.

Holldobler, Wilson (1990) *The Ants* Harvard Press.

Jensen, R., Veloso, M. *OBDD-based Universal Planning for Multiple Synchronized Agents in Non-Deterministic Domains* Carnegie Mellon University

Lee, Xu. *Online, interactive learning of gestures for human/robot interfaces*

Mitchell, T. (1997) *Machine Learning* WCB/McGraw-Hill Companies.

Pedro Miguel Martins da Silva Aparicio (2000) *Designing and Implementation of a Population of Cooperative Autonomous Robots.* Thesis

Rabiner, L. (1989) *A tutorial on Hidden Markov Models and Selected Applications in Speech Recognition* IEEE.

Riley, P. (1999) *Classifying Adversarial Behaviors in a Dynamic Inaccessible Multi-Agent Environment* CMU-CS-99-175

Tambe, M. (1997) *Agent architectures for flexible, practical teamwork*

Vidal, I. (1998) *Computational Agents that Learn About Agents: Algorithms for their Design and a predictive Theory of their Behavior* Thesis.

Whitehouse, J. (1996) *Ant Wars: Combat strategies, territory and nest defense in the leaf-cutting ant att laevigata* Animal Behaviour, vol 51.

Yang, et al. (1993) *Hidden Markov model approach to skill learning and its application to telerobotics* ech. report CMU-RI-TR-93-01, Robotics Institute, Carnegie Mellon University, January, 1993.