

How Multirobot Systems Research Will Accelerate Our Understanding of Social Animal Behavior

Researchers are tracking movements of ants and monkeys using robotics algorithms; they hope to automatically recognize animal behavior and to simulate it using robots.

By TUCKER BALCH, Member IEEE, FRANK DELLAERT, Member IEEE, ADAM FELDMAN, ANDREW GUILLORY, CHARLES L. ISBELL, JR., ZIA KHAN, STEPHEN C. PRATT, ANDREW N. STEIN, AND HANK WILDE

ABSTRACT | Our understanding of social insect behavior has significantly influenced artificial intelligence (AI) and multirobot systems' research (e.g., ant algorithms and swarm robotics). In this work, however, we focus on the opposite question: "How can multirobot systems research contribute to the understanding of social animal behavior?" As we show, we are able to contribute at several levels. First, using algorithms that originated in the robotics community, we can track animals under observation to provide essential quantitative data for animal behavior research. Second, by developing and applying algorithms originating in speech recognition and computer vision, we can automatically label the behavior of animals under observation. In some cases the automatic labeling is more accurate and consistent than manual behavior identification. Our ultimate goal, however, is to automatically create, from observation, *executable* models of behavior. An execut-

able model is a control program for an agent that can run in simulation (or on a robot). The representation for these executable models is drawn from research in multirobot systems programming. In this paper we present the algorithms we have developed for tracking, recognizing, and learning models of social animal behavior, details of their implementation, and quantitative experimental results using them to study social insects.

KEYWORDS | Multirobot systems; social animals; tracking

I. INTRODUCTION

Our objective is to show how robotics research in general and multirobot systems research in particular can accelerate the rate and quality of research in the behavior of social animals. Many of the intellectual problems we face in multirobot systems research are mirrored in social animal research. And many of the solutions we have devised can be applied directly to the problems encountered in social animal behavior research.

One of the key factors limiting progress in all forms of animal behavior research is the rate at which data can be gathered. As one example, Gordon reports in her book that two observers are required to track and log the activities of one ant. One person observes and calls out what the ant is doing, while the other logs the data in a notebook [1]. In the case of *social* animal studies, the problem is compounded by the multiplicity of animals interacting with one another.

Manuscript received June 1, 2005; revised June 1, 2006. This work was supported by the National Science Foundation under NSF Award IIS-0219850. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the National Science Foundation.

T. Balch, F. Dellaert, A. Feldman, A. Guillory, C. L. Isbell, Jr., and H. Wilde are with Interactive and Intelligent Computing, Georgia Institute of Technology, Atlanta, GA 30308 USA (e-mail: tucker@cc.gatech.edu; dellaert@cc.gatech.edu; storm@cc.gatech.edu; guillory@cc.gatech.edu; isbell@cc.gatech.edu; hwilde@aethon.com).

Z. Khan is with Sarnoff Corporation, Princeton, NJ 08543-5300 USA (e-mail: zkhan@sarnoff.com).

S. C. Pratt was with the Department of Ecology and Evolutionary Biology, Princeton University, Princeton, NJ 08544 USA. He is now with the School of Life Sciences, Arizona State University, Tempe, AZ 85287 USA.

A. Stein is with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: astein@cs.cmu.edu).

Digital Object Identifier: 10.1109/JPROC.2006.876969

One way robotics researchers can help is by applying existing technologies to enhance traditional behavioral research methodologies. For instance computer vision-based tracking and gesture recognition can automate much of the tedious work in recording behavior. However, in developing and applying existing technologies, we have discovered new, challenging, research problems. Vision-based multitarget tracking, for instance, is far from being a solved problem. Algorithms capable of tracking several targets existed when we began our work in this area, but research with ants calls for software that can track dozens or hundreds of ants at once. Similarly, algorithms for labeling the behavior of one person (or animal) existed several years ago, but now we must label social behavior between two or more animals at once. The point being that the application of robotics algorithms to a new domain also drives new research in robotics.

Now, in addition to helping animal behavior researchers speed up their traditional modes of research, we can also provide them with entirely new and powerful tools. In particular, mobile robotics can offer new ways to represent and model animal behavior [2]. What is the best model of behavior? It is our position that an *executable* model provides the most *complete* explanation of an agent's behavior. By “executable” we mean that the model can run in simulation or on a mobile robot. By “complete” we mean that all aspects of an agent's behavior are described: from sensing, to reasoning, to action. Executable models provide a powerful means for representing behavior because they fully describe how perception is translated into action for each agent. Furthermore, executable models can be tested experimentally on robots or in a multiagent simulation to evaluate how well the behavior of a simulated “colony” matches to the behavior of the actual animals. Other types of models cannot be tested in this way.

Robotics researchers are well positioned to provide formalisms for expressing executable models because the programs used to control robots are in fact executable models. Furthermore, researchers who use behavior-based approaches to program their robots use representations that are closely related to those used by biologists to describe animal behavior. Our objective in this regard is to create algorithms that can learn executable models of social animal behavior directly from observations of the social animals themselves.

In this paper we report on our progress in this area over the last six years. A primary contribution is the idea that we ought to apply artificial intelligence (AI) and robotics approaches to these problems in the first place. However, additional, specific contributions include our approaches to solving the following challenges.

- **Tracking multiple interacting targets:** This is an essential first step in gathering behavioral data. In many cases, this is the most time-consuming step for a behavioral researcher.

- **Automatically labeling behavior:** Once the locations of the animals are known over time, we can analyze the trajectories to automatically identify certain aspects of behavior.
- **Generating executable models:** Learning an executable model of behavior from observation can provide a powerful new tool for social behavior research.

The rest of the paper follows the order listed above. We review the challenges of multitarget tracking and our approach to the problem. We show how social behavior can be recognized automatically from trajectory logs. Then we close with a presentation of our approach to learning executable models of behavior. Each topic is covered at a high level, with appropriate citations to more detailed descriptions. We also review the relevant related work in each area in the beginning of each corresponding section.

II. TRACKING MULTIPLE INTERACTING TARGETS

Tracking multiple targets is a fundamental task for mobile robot teams [3]–[5]. An example from RoboCup soccer is illustrated in Fig. 1. In the small-size league, 10 robots and a ball must be tracked using a color camera mounted over the field. In order to be competitive, teams must find the location of each robot and the ball in less than 1/30 of a second.

One of the leading color-based tracking solutions, CMVision was developed in 2000; it is used now by many teams competing at RoboCup [6]. We used CMVision for our first studies in the analysis of honeybee dance behavior [7]. However, for bees, color-based tracking requires marking all the animals to be observed (Fig. 1)—this is something of a tricky operation, to say the least. Even if the




Fig. 1. Multitarget tracking in robot soccer and social animal research. In the RoboCup small-size league robots are tracked using an overhead camera (image courtesy Carnegie Mellon).

animals are marked, simple color-based tracking often gets confused when two animals interact.

We now use a probabilistic framework to track multiple targets by an overhead video camera [8]. We have also extended the approach for multiple observers (e.g., from the point of view of multiple sensors or mobile robots). Comprehensive details on these approaches are reported in [8]–[11].

Our objective in this first stage (tracking) is to obtain a record of the trajectories of the animals over time, and to maintain correct, unique identification of each target throughout. We have focused our efforts on tracking social insects in video data, but many of the ideas apply to other sensor modalities (e.g., laser scanning) as well. The problem can be stated as follows.

- **Given:**
 - 1) examples of the appearance of the animals to be tracked; and
 - 2) the initial number of animals.
- **Compute:** trajectories of the animals over time with a correct, unique identification of each target throughout.
- **Assume:**
 - 1) the animals to be tracked have similar or identical appearance;
 - 2) the animals' bodies do not deform significantly as they move; and
 - 3) the background does not vary significantly over time.

A. Particle Filter Tracking

Traditional multitarget tracking algorithms approach this problem by performing a target detection step followed by a track association step in each video frame. The track association step solves the problem of converting the detected positions of animals in each image into multiple individual trajectories. The multiple hypothesis tracker [12] and the joint probabilistic data association filter (JPDAF) [13], [14] are the most influential algorithms

in this class. These multitarget tracking algorithms have been used extensively in the context of computer vision. Some example applications are the use of nearest neighbor tracking in [15], the multiple hypothesis tracker in [16], and the JPDAF in [17]. A particle filter version of the JPDAF was proposed in [18].

Traditional trackers (e.g., the extended Kalman–Bucy filter) rely on physics-based models to track multiple targets through merges and splits. A merge occurs when two targets overlap and provide only one detection to the sensor. Splits are those cases where a single target is responsible for more than one detection. In radar-based aircraft tracking, for instance, it is reasonable to assume that two planes that pass close to one another can be tracked through the merge by predicting their future locations based on their past velocities. That approach does not work well for targets that interact closely and unpredictably. Ants, for instance, frequently approach one another, stop suddenly, then move off in a new direction. Traditional approaches cannot maintain track in such instances. Our solution is to use a joint particle filter tracker with several novel extensions.

First we review the operation of a basic particle filter tracker, then we introduce the novel aspects of our approach. The general operation of the tracker is illustrated in Fig. 2. Each particle represents one hypothesis regarding a target's location and orientation. For ant tracking in video, the hypothesis is a rectangular region approximately the same size as the ant targets. In the example, each target is tracked by five particles. In actual experiments we use hundreds of particles per target.

We assume we start with particles distributed around the target to be tracked. A separate algorithm (not covered here) addresses the problem of initializing the tracker by finding animals in the first image of the sequence. After initialization, the principal steps in the tracking algorithm include the following.

- 1) **Score:** each particle is scored according to how well the underlying pixels match an appearance model.




Fig. 2. Particle filter tracking. (a) The appearance model used in tracking ants. This is an actual image drawn from the video data. (b) A set of particles (white rectangles), are scored according to how well the underlying pixels match an appearance model. (c) Particles are resampled according to the normalized weights determined in the previous step. (d) The estimated location of the target is computed as the mean of the resampled particles. (e) The previous image and particles. A new image frame is loaded. (f) Each particle is advanced according to a stochastic motion model. The samples are now ready to be scored and resampled as above.

- 2) **Resample:** the particles are “resampled” according to their score. This operation results in the same number of particles, but very likely particles are duplicated while unlikely ones are dropped.
- 3) **Average:** the location and orientation of the target is estimated by computing the mean of all the associated particles. This is the estimate reported by the algorithm as the pose of the target in the current video frame.
- 4) **Load new image:** read the next image in the sequence.
- 5) **Apply motion model:** each particle is stochastically repositioned according to a model of the target’s motion.
- 6) **Go to Step 1.**

B. Joint Particle Filter Tracking

The algorithm just described is suitable for tracking an individual ant, but it fails in the presence of many identical targets [8]. The typical mode of failure is “particle hijacking” whereby the particles tracking one animal latch on to another animal when they pass close to one another.

Three additional extensions are necessary for successful multitarget tracking: First, each particle is extended to include the poses of all the targets (i.e., they are *joint* particles). An example of joint particles for ant tracking is illustrated in Fig. 3. Second, in the “scoring” phase of the algorithm, particles are penalized if they represent hypotheses that we know are unlikely because they violate known constraints on animal movement (e.g., ants seldomly walk on top of each other). Blocking, in which particles that represent one ant on another are penalized, is illustrated in Fig. 3. Finally, we must address the exponential complexity of joint particle tracking (this is covered below). These extensions, and evaluations of their impact on tracking performance, are reported on in detail in [8] and [10].

The computational efficiency of joint particle tracking is hindered by the fact that the number of particles and the amount of computation required for tracking is exponential in the number of targets. For instance, ten targets tracked by 200 hypotheses each would require $200^{10} = 10^{23}$ particles in total. Such an approach is clearly intractable. Independent trackers are much more efficient, but as mentioned above, they are subject to tracking failures when the targets are close to one another.

We would like to preserve the advantages of joint particle filter tracking, but avoid an exponential growth in complexity. In order to reduce the complexity of joint particle tracking we propose that targets that are in difficult to follow situations (e.g., interacting with one another) should be tracked with more hypotheses, while those that are isolated should be tracked with fewer hypotheses. We use Markov chain Monte Carlo (MCMC) sampling to accomplish this in our tracker [8], [10]. The approach is rather complex, but the effective result is that we can achieve the same tracking quality using orders of magnitude fewer particles than would be required otherwise. Fig. 4 provides an example of how the approach can focus more hypotheses on challenging situations while using an overall smaller number of particles. An example result from our tracking system using MCMC sampling is illustrated in Fig. 5. In this example, 20 ants are tracked in a rectangular arena.

After trajectory logs are gathered they are checked against the original video. In general, our tracker has a very low error rate (about one of 5000 video frames contains a tracking error [8]). However, to correct such errors and ensure nearly perfect data, we verify and edit the trajectories using TeamView, a graphical trajectory editing program also developed in our laboratory [19].

C. Using Tracking for Animal Experiments

Most of our efforts in tracking are focused on social insect studies. These are examined in more detail in the




Fig. 3. Joint particles and blocking. When tracking multiple animals, we use a joint particle filter where each particle describes the pose of all tracked animals (left). In this figure there are two particles—one indicated with white lines, the other with black lines. Particles that overlap the location of other tracked targets are penalized (right).




Fig. 4. MCMC sampling. (a) Each ant is tracked with three particles, or hypotheses. With independent trackers this requires only 12 particles, but failures are likely. (b) With a joint tracker, to represent three hypotheses for each ant would require 81 particles altogether. However, MCMC sampling (c) enables more selective application of hypotheses, where more hypotheses are used for the two interacting ants on the right side, while the lone ants on the left are tracked with only one hypothesis each.




Fig. 5. Example tracking result. 20 ants are tracked in a rectangular arena. The white boxes indicate the position and orientation of an ant. The black lines are trails of their recent locations.

following sections. However, we have also been deeply involved with the Yerkes Primate Research center to apply our tracking software to the study of monkey behavior. In

particular, we are helping them evaluate spatial memory in rhesus monkeys by measuring the paths the monkeys take as they explore an outdoor three dimensional arena over repeated trials. In this study we use two cameras to track the movements of monkeys inside the arena and generate three-dimensional (3-D) data (see Fig. 6). Over the course of this study, we have collected over 500 h of trajectory data—this represents the largest corpus of animal tracking data we are aware of. The design of our tracking system for this work is reported in [20].


In the initial study, our system only tracked one monkey at a time. We are now moving towards an experiment where we hope to track as many as 60 monkeys at once in a 30 m by 30 m arena. To accomplish this we are experimenting with a new sensor system: scanning laser range finders (Fig. 7). Each laser scans a plane out to 80 m in 0.5° increments. Example data from a test in our lab is shown in the figure on the right. In this image the outline of the room is visible, as well as five ovals that represent the cross sections of five people in the room. We plan to extend and adapt our vision-based tracking software to this new sensor modality. The laser-based tracks will be limited



(a)



(b)



(c)

Fig. 6. Left: experimental setup at Yerkes Primate Research Center for tracking monkeys in an outdoor arena. Center: a monkey in the arena. Right: an example 3-D trajectory.




Fig. 7. Left: scanning laser range finders mounted on portable boxes for easy field deployment. Right: data from an experiment in our lab. The outline of the room is visible, along with 5 oval objects—people in the room.

to two-dimensional (2-D) (versus 3-D for the vision-based tracking). Maja Mataric's group at USC is also using laser range-finders to track people [21], [22].

D. Limitations

To date, our vision-based multitarget tracking algorithms have been successful primarily in laboratory environments in which the lighting is controlled or the background provides high contrast. Fortunately, the behavior of some social insects (e.g., *Temnothorax curvispinosus* and *Apis mellifera*) is not significantly affected by a laboratory environment. However, it would certainly be preferable to track and assess behavior in the natural (outdoor) environment of the animal under study.

We have been successful tracking monkeys outdoors using computer vision, but in those experiments we are only concerned with tracking one monkey at a time. In our future investigations, we will use laser-based tracking to overcome the limitations of vision-based solutions.

III. AUTOMATIC RECOGNITION OF SOCIAL BEHAVIOR

Behavior recognition, plan recognition, and opponent modeling are techniques whereby one agent estimates the internal state or intent of another agent by observing its actions. These approaches have been used in the context of robotics and multiagent systems research to make multi-robot teams more efficient and robust [23]–[25]. Recognition helps agents reduce communications bandwidth requirements because they can infer the intent of their team members rather than having to communicate their intent. In other cases the objective is to determine the intentions of an opposing agent or team in order to respond more effectively. Related problems are studied in the context of speech recognition, gesture recognition, and surveillance [26]–[28]. Social animal research offers a similar challenge: how can we automatically recognize the behavior of animals under observation?

As an example, in ants, encounters between animals are thought to mediate individual and colony-level

behavior, including for instance, forager recruitment, response to external disturbances, and nest selection for the colony [29]–[32]. Typically, behaviors such as these are investigated experimentally by having human observers record observed or videotaped activity. Generating such records is a time-consuming task for people, however, especially because the numbers of interactions can increase exponentially in the number of animals.

Computer vision-based tracking and machine analysis of the trajectories offers an opportunity to accelerate such research. Once we have trajectories of the animals' movements, we can analyze them to infer various aspects of their behavior. We seek to recognize *social* behavior, so we focus on aspects of the animals' movement that reveal social interaction.

We have explored two approaches to this problem. In the first approach, we take advantage of a model of the geometry of the animal's sensory apparatus to estimate when significant interactions occur. In the second approach, we use example labels provided by a human expert to train the system to label activities automatically.

A. Using Sensory Models to Recognize Social Behavior

When ants approach one another, they often (perhaps deliberately) meet and tap one another with their antennae. Ants of the genus *Temnothorax* for example, use such encounters to detect the density of nestmates in a potential nest site, a crucial cue for their collective decision about whether to move the whole colony to the site [32]. These encounters can be readily divided into three types, based on the posture of one ant (the focal ant) relative to the other: head-to-head (HH), head-to-body (HB), and body-to-head (BH). Body-to-body (BB) encounters may also be important, so we include them for completeness. It is not known for certain the extent to which the various encounter types affect behavior differently—in fact, that is one purpose of the research we are conducting. Formally, we seek to address the following problem.

- **Given:**
 - 1) time-stamped trajectories of the pose of multiple interacting animals; and
 - 2) a geometric model of their sensory field.
- **Compute:** appropriate labels for the trajectories, including encounters, automatically.
- **Assume:** the geometric models accurately reflect the sensory apparatus of the animals.

To accomplish this we approximate an ant's antennal and body sensory fields with a polygon and circle, respectively (Fig. 8). An encounter is inferred when one of these regions for one ant overlaps a sensory region of another ant. Our model is adapted from the model introduced for army ant simulation studies by Couzin and Franks [33]. Note that for this work we do not consider the visual field of the ant. While the visual field is




Fig. 8. Left: we have examined the problem of detecting different types of encounters between ants: HH, HB, BH, BB, or no interaction (X). Right: in one approach, we use a geometric model of an ant's sensory system. Parameters of the model include R_b , the radius of the animal's "body" sensor, R_a the range of antenna sensing, and θ the angular field of view of the antennae.

certainly important for guiding ants we consider that it is not critical for determining whether physical contact occurs between them.

The following are details of our sensory model for *Aphaenogaster cockerelli*. Recall that the tracking software reports the location of the center of the ant and its orientation. We estimate the front of the head to be a point half a body length away from the center point along the centerline. From the head, the antennae project to the left and right at 45° . We construct a polygon around the head and antennae as illustrated, with one point at the center of the ant's thorax, two at the outer edges of the antennae and an additional point, one antenna length away, directly along the ant's centerline in front of the head. The inferred "antennal field of view" is the resulting polygon defined by these points. We assume a body length of 1 cm and antenna length of 0.6 cm. We estimate an ant's "body" sensory field to include a circle centered on the ant with a radius of 0.5 cm.

We assume that any object within the head sensory field will be detected and touched by the animal's antennae, and any object within the range of the body can be detected by sense organs on the legs or body of the ant. By considering separate sensory fields for the body and the antennae, we are able to classify each encounter into one of the four types described above. To determine when an interaction occurs, we look for overlaps between the polygonal and circular regions described above.

Computationally, determining an overlap consists of checking for intersections between line segments and circles. Additionally, we require that two ants' orientations

must differ by at least 5° before an interaction can count as "head-to-head." Without this constraint, we noticed many spurious "encounters" logged by the system for ants that were simply standing next to one another without actually interacting.

In order to evaluate the encounter detection system, we had a human observer and our detection software label interactions on the same experimental data (a videotape of ants interacting in an arena). In cases where the logs differed, we reviewed the video carefully at that time using slow-motion playback to determine whether or not an encounter actually occurred. We categorized each difference into one of four classes: 1) computer false positives in which the computer identified an encounter that did not occur; 2) computer false negatives in which the computer failed to identify an encounter that actually occurred; 3) human false positives in which the human observer logged an encounter that did not actually occur; and 4) human false negatives in which the human did not record interactions that actually occurred. We repeated this examination for three different videotaped experiments. Quantitative results with standard deviations are provided in the table at the bottom of this page.

We can clearly benefit from automatic labeling in this application. Behavior labeling is faster and more accurate than human labeling. The primary weakness is that some detected encounters have not actually occurred as indicated by a 7% false positive rate.

B. Using Trainable Models to Recognize Social Behavior

The geometric sensory model-based method described above performs well, but it is difficult to apply in general because a new model would have to be devised for each application. Accordingly, we have also developed recognition methods that take advantage of input provided by human experts. In these systems, a human labels some of the experimental data (referred to as the training set). A model of the human's labeling is created, then used as a reference to label new examples. The power of this approach is that it does not require the human expert to elucidate the considerations involved in the classification of an event into one of several categories, it only requires examples of correctly classified observations. This is of course a well-studied problem. Variations of the problem are studied in computer vision for gesture recognition, face recognition, and activity recognition [27], [34], [35]. In our application, again, we are specifically interested in

Computer		Human	
false positives	false negatives	false positives	false negatives
7%(2%)	0%(0%)	15%(2%)	29%(12%)




Fig. 9. Ants are tracked and encounters between them are detected automatically in video images. (a) A vision-based multitarget tracker locates and tracks ants in video. (b) Geometric models of animal sensory fields are checked for overlap. (c) A head-to-head encounter is detected and indicated by a circle around the interacting animals. (d) Example video in which encounters are detected and highlighted.

social behavior, so the problem for human-trainable recognition algorithms is specified as follows.

- **Given:**
 - 1) time-stamped trajectories of the pose of multiple interacting animals; and
 - 2) a set of quantitative features defined over (and between) the trajectories; and
 - 3) labeled example trajectories.
- **Compute:** appropriate labels for the trajectories, including encounters, automatically.
- **Assume:** the features are suitable for discriminating between the behaviors that should be identified, and the training examples provided are correct.

For the problem of detecting different types of interactions between ants, we selected three features of a potential encounter that can be easily determined from the time-stamped trajectory data. The features are illustrated and described in Fig. 10. We use the following steps to learn and apply behavior labels.

- 1) **Track:** animals are tracked using the algorithms described above.
- 2) **Label:** at each frame, the behavior of each animal is labeled by a human observer (e.g., HH, BH, HB, or X).




Fig. 10. Three features used to classify interactions between ants.

- 3) **Cluster:** the data gathered above is processed in order to determine disjoint regions in feature space that enclose each type of interaction.
- 4) **Classify:** in online operation, we determine in which region of feature space a particular observation lies and label it accordingly.

In our current system, the third step of the procedure is a hill-climbing optimization. For each type of interaction (set of thresholds) we seek to maximize the number of data points with that label within the bounding box *minus* the number of data points with different labels inside the box. Each of the six thresholds is optimized in turn, with iterations continuing until there is no improvement between subsequent loops. The resulting thresholds determined from a 5-min labeled sequence are illustrated in Fig. 10.

When the system is run on new (unlabeled) data, it detects and labels interactions by computing the quantitative features at each time stamp, then finding which bounding box (if any) the point is enclosed by. It is possible that one ant may experience several interactions at once with more than one other ant (the system examines the three closest other ants). If there are multiple interactions, we select the highest priority type according to the order: HH, HB, BH, X. (Note that in these experiments we did not include BB interactions).

Observe that this is a rather simple approach. There are more sophisticated algorithms that could also be applied including, for instance, kernel regression or decision trees [36], [37]. We compared the quality of our results with kernel regression and found no significant advantage for one over the other. We suspect that the feature space in this application lends itself well to the bounding box approach.

We evaluated the system by having a human expert label two 5-min videos of *Temnothorax curvispinosus* searching a new habitat. In each segment ants enter the field of view, interact in various ways, then depart. There are hundreds of interactions in each video. We used the labeling of one video as training data for the system, which then labeled the other video. We then compared the




Fig. 11. This image shows an example frame of video of *Temnothorax curvispinosus* labeled automatically by our human-trainable system. The colored triangles over the animals are coded for the different types of interaction that animal is experiencing (cyan: BH, yellow: HB, magenta: HH, blue: X).

automatic labeling of the test data to the human labeling of the same data. The automatic method correctly identified 94% of the interactions. We then reversed the order of training: we trained the system on the second video, then tested it on the first. In that case the system correctly identified 87% of the interactions. Example frames from video labeled by our system are reproduced in Fig. 11.

This performance is good, and it compares favorably with similar systems used for gesture recognition (e.g., [34]). However, as many as 43% of the interactions reported by the system are false positives; meaning that the system detects many interactions that never occurred.

All is not lost, however. In similar work in which we automatically labeled honeybee dance behavior, we achieved a similar 93% accuracy, but with a substantially lower false positive rate [7]. The difference is that in the honeybee work we used a hidden Markov model (HMM) in addition to the feature-based labeling. HMMs provide a “smoothing” influence that tends to suppress brief “noisy” detections. As this work with ants progresses, we will incorporate HMMs as well.

C. Using Automatic Recognition to Study Behavior in Animals

Our objective in this section is to show how multitarget tracking and automatic social behavior recognition can be used in the context of an animal study. We present some preliminary results from our own experiments, but we expect even more impact from our collaborative work with biologists and psychologists.

We are using our behavior recognition methods to investigate interaction rates in two ant species. At Georgia Tech we are researching the behavior of colonies of *Aphaenogaster cockerelli*. At Princeton, we are exploring the means by which *Temnothorax curvispinosus* selects a new nest site. For *T. curvispinosus* the colony’s choice depends on

whether the ants achieve a quorum of nestmates at a candidate site (Pratt *et al.* [32]). The ants appear to detect the quorum through their rate of encounters with each other at the site [32]. Our work is focused on discovering the parameters of this decision process. Here we report on our experiments with *A. cockerelli* because the work with *T. curvispinosus* is still underway.

We would like to ascertain whether ants deliberately seek or avoid encounters with other ants, and whether they prefer certain types of encounters over others. From casual observation it seems that the ants are more effective at exchanging information when they meet head-to-head. It also seems that when a head-to-head encounter is feasible, the ants approach one another along a trajectory to ensure a successful meeting. We are interested in discovering whether the ants make these choices deliberately or whether their actions are essentially random. The importance of encounter rates has been explored by others, especially Gordon [1], [30], [31], but to our knowledge no one has examined whether certain types of encounters are preferred over others. In order to explore this question, we consider the following null hypothesis.

Null Hypothesis: the rate and type of encounters between ants arises from random motion of the animals.

If our data contradict this hypothesis, the assertion that ants move deliberately is supported.

For the experiments, foragers (ants outside the brood chambers, on or near food), were selected randomly with forceps and placed in an otherwise empty 10 cm by 15 cm arena with vertical walls. The walls were treated with fluon to prevent escape. Two 30-min episodes were recorded, the first with 10 ants, the second with 20 ants. Both sets of ants were kept separated from the rest of the colony until the experiment was completed. In both trials the ants moved actively around the arena during the entire 30 min with no apparent degradation in their enthusiasm for exploration. Additional details regarding the animals and their habitats are included in Appendix I.

The system counted the total number of each type of encounter for trials with 10 and 20 ants. There were a total of 5055 encounters detected: 1382 encounters with 10 ants and 3673 with 20 ants. The average number of encounters *per ant* is reported in Table 1. Observe that as the density of ants doubles from 10 to 20 ants the total number of interactions increases by only 33%. Gordon asserts that if

Table 1 Automatically Detected Encounters for 10- and 20-Ant Experiments

Encounter type	Encounters per ant		
	10 ants	20 ants	Change
head-to-head	12.6	24.3	+93%
head-to-body	36.4	48.4	+33%
body-to-body	50.0	62.8	+26%
body-to-head	39.2	48.2	+23%
total	138.2	183.7	+33%




Fig. 12. Simulated ant trajectories (left) compared with real ant trajectories (right).

the ants moved according to Brownian motion, we should expect the number of interactions to double. Head-to-body, body-to-body and body-to-head encounters each increase by 20%–30%. However, the average number of head-to-head encounters nearly doubles, growing from 12.6 to 24.3 encounters per ant. All of these proportional increases are statistically significant (the 95% confidence intervals do not overlap).

Observe the following regarding these results.

- Although the density of ants increases by 100% from the first experiment to the second the total number of encounters increases by only one-third.
- Of the four various types of interactions, head-to-head interactions make up a small proportion (9% for 10 ants and 13% for 20 ants).
- In sharp contrast to the other types of interaction, the number of head-to-head interactions *nearly doubles* as the density of ants doubles from the first experiment to the second.

The first result confirms earlier work by Gordon in which she explored a number of variations affecting ant density, and similarly observed that the number of interactions did not increase in proportion to the density of animals in the arena [30]. The result that head-to-head interactions make up a relatively small proportion of the interactions raises more questions, namely: Does this represent a deliberate choice on the part of the ants, or is it merely a consequence of geometry and random motion? And why does the number of head-to-head interactions *double* when the other types only increase by about one third?

We created a simulation of ant movement to explore these questions. (Details of the simulation are provided in Appendix II.) In particular we wanted to examine the assertion that we would observe the same distribution of encounters for agents that just move randomly. Our objective for the simulation was to create antlike (but *random*) trajectories, where the simulated ants move without regard to the location and orientation of other ants. The simulated ants do not turn towards each other to initiate an interaction, or turn away in avoidance. In all other respects, the simulation mimics the movement of

real ants, based on statistics gathered from observation. Fig. 12 illustrates the types of trajectories generated by the simulation in comparison to those created by live ants.

We ran ten simulation experiments for both 10- and 20-ant cases. We then ran the same analysis on this data as we ran on the real ant data. Quantitative results from the simulation experiments are reported in Table 2. We also compare the quantitative behavior of real ants with the behavior of simulated ants in Table 3.

Consider Table 3 first. This data shows how well the behavior of the real and simulated ants corresponds when 10 ants are in the arena. Overall there is only an 11% difference in the total number of interactions, but there is a 75% difference in the number of head-to-head interactions. This indicates that there is a substantial difference in the way real ants and “random” ants interact head-to-head. Now consider Table 2. For simulated ants, when the number of ants doubles from 10 to 20, we see a 68% increase in the total number of interactions. In the real ant experiments, however, we see only a 33% increase in the number of interactions. We suggest that the data supports the following conclusions.

- Ants seek head-to-head encounters at statistically significant higher rate than if they moved randomly.
- As the number of ants increases, they deliberately avoid encounters to maintain a lower encounter rate than if they moved randomly.

This experiment illustrates how automatic behavior recognition can facilitate social animal studies. The results arise from a few minutes of video, yet they represent over 5000 interactions of four different types between closely

Table 2 Quantitative Results From Simulation for 10 and 20 Ants

Encounter type	Encounters per ant in simulation		
	10 ants	20 ants	Difference
head-to-head	3.1	4.8	+55%
head-to-body	36.0	62.4	+73%
body-to-body	46.7	76.2	+63%
body-to-head	36.9	62.7	+70%
total	122.7	206.1	+68%

Table 3 Comparison of Simulation Experiments With Animal Experiments. Automatically Detected Encounters for 10 Ants

Encounter type	Encounters per ant		Difference
	10 ants (real)	10 ants (simulation)	
head-to-head	12.6	3.1	-75%
head-to-body	36.4	36.0	-1%
body-to-body	50.0	46.7	-7%
body-to-head	39.2	36.9	-6%
total	138.2	122.7	-11%

interacting animals. Our use of computer vision-based tracking and automatic recognition of social interactions enabled a statistical analysis that would have been very difficult or impossible to conduct otherwise. We expect to report additional and more comprehensive results like this in the future.

IV. LEARNING EXECUTABLE MODELS OF BEHAVIOR

It is our position that an *executable* model provides the most *complete* explanation of an agent’s behavior. By “executable” we mean that the model can run in simulation or on a mobile robot. By “complete” we mean that all aspects of an agent’s behavior are described: from sensing, to reasoning, to action. In this section we will propose a formalism for representing executable models, and support it with examples from biology and robotics. Later we will introduce algorithms for learning executable models and illustrate their use.

Behavioral ecologists seek to explain complex behavior by developing and testing predictive models. There are a number of methods for modeling social animal behavior at the individual and colony level. Ethograms are among the most effective and frequently used approaches [38]. A

sample ethogram of ant behavior is provided in Fig. 13 [39]. The nodes of this diagram represent the behavioral acts of individual animals. The links between the nodes show how behaviors are sequenced. The frequency of observed transitions is also recorded and represented.

Computer scientists will recognize a similarity between the diagram in Fig. 13 and a Markov process (MP). The nodes representing behavioral acts in an ethogram correspond to states in an MP. Transitions between behaviors correspond to the probabilistic transitions of an MP. Researchers are already investigating methods for automatically learning Markov models, including some who apply the approach to learning models of behavior [40]. A goal of this work is to employ a similar approach to the task of learning animal behavior models.

The MP is also quite similar to methods for programming sequences of behavior for mobile robots in the behavior-based paradigm, and we will take advantage of that similarity. The central idea of behavior-based robotics is to closely couple a robot’s sensors and actuators so as to avoid the trouble of maintaining a map of the world or deliberating over it. Brooks’ subsumption architecture and Arkin’s motor schemas are the best-known examples [41], [42]. Both of these approaches were inspired by biological models.

In the behavior-based approach, *primitive* behaviors express separate goals or constraints for a task. As an example, important behaviors for a navigational task would include primitive behaviors like **avoid_obstacles** and **move_to_goal**. Behaviors may be grouped to form more complex, emergent behaviors. Groups of behaviors are referred to as *behavioral assemblages*. One way behavioral assemblages may be used in solving complex tasks is to develop an assemblage for each subtask and to execute the assemblages in an appropriate sequence. The steps in the

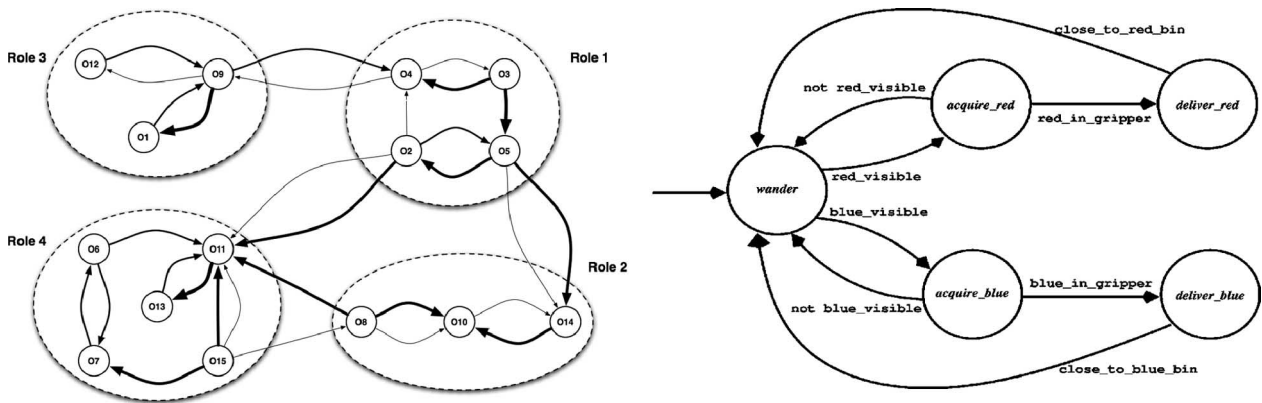


Fig. 13. Representations of behavior in ants (left) and robots (right). Left: this representation of individual ant behavior (called an ethogram) is quite similar to a stochastic MP. The MP is one formulation used by AI researchers to model the behavior of intelligent agents. In this ethogram, behavioral acts, a_i , are linked by arcs indicating transitions from act to act. Thicker lines indicate higher probability transitions. Right: this diagram illustrates a programmed sequence of behaviors for a foraging robot.




Fig. 14. A team of two robots executes a behavior-based foraging program coded as an FSA.

sequence are separate *behavioral states*. Perceptual events that cause transitions from one behavioral state to another are called *perceptual triggers*. A resulting task solving strategy can be represented as a finite state automaton (FSA). This technique is referred to as *temporal sequencing*.

An example FSA for executing a foraging task is illustrated on the right in Fig. 13. A team of robots executing this program are illustrated in Fig. 14. The robots step through a sequence of assemblages including a search of the environment (i.e., **wander**), acquisition of a detected food object, then delivery of the object to homebase. It is not necessary to understand the particulars of the behaviors at this point. What is important, however, is the idea that transitions from one behavioral assemblage to another are triggered by perceptual events. As an example, the **deliver_blue** behavior, in which a blue item is carried by the robot back to the homebase (or “nest”) is activated when the robot has a blue item in its gripper.

Note that these representations of sequenced behavior are strongly similar to the representations used by biologists to model behavior. This is at least partly because the approaches were inspired by biological mechanisms in the first place.

A. Learning Executable Models of Behavior From Observations

Consider the task for an ecologist seeking to describe the behavior of an ant or bee. Typically, in studying a particular behavior, a scientist will first define a set of primitive behaviors that can be observed and recorded. For example, in social insect foraging, such primitive actions include searching for and exploiting a food source [29]. Having defined these actions, the scientist must then determine the high-level structure of the behavior; that is, what causes the animal to switch between low-level behaviors. More formally we define our problem as follows.

- **Given:**
 - 1) a set of behavioral assemblages (i.e., controllers) an agent may use;
 - 2) a model of the agent’s perception apparatus;

- 3) a set of perceptions that may cause the agent to switch behaviors; and
- 4) example trajectories of the agent interacting in its environment.

- **Compute:** an executable model composed of the behavioral assemblages and the perceptual triggers that cause transitions between them.

- **Assume:**
 - 1) the agent acts according to an MP with inputs; and
 - 2) we have an accurate representation of the primitive behaviors the agent uses and the perceptions that cause transitions.

Stated another way: Assume we already know the constituent behaviors that reside at the nodes (states) of an FSA (like the one depicted on the right in Fig. 13). The task then is to learn the transitions between the nodes by observing animals executing the behaviors.

Our assumptions may sound rather strong; how can we know the details of the behavioral assemblages that ants use? As it turns out, in related work we have developed a means for learning such low level controllers. The approach has been demonstrated by learning controllers from live ant data and running them in simulation and on a mobile robot [43]. Our eventual goal is to combine these two approaches in order to learn the entire control system. At present, however, we will focus on learning the high-level switching, given the lower level controllers.

This work is in contrast to work in *activity recognition* where learning is used to build models of activity. These models are often generative, but their goal is recognition, not re-creation. If domain knowledge is assumed, it is typically high-level domain knowledge of the structure of the activity being recognized, for example, in the form of topology restrictions or high-level grammars. Such techniques have been used to classify bee behavior [7], [44], for example. Here we assume that we have knowledge of detectable low-level actions that we can use to understand the structure of a behavior sequence and reproduce the behavior in new situations.

Furthermore, we consider that it is important for the learned model to be understandable by humans—thereby making the learned models useful to biologists. In this way our work is similar to imitation learning. The notion of primitives in [45] is similar to our notion of behavioral assemblages.

1) *The IOHMM Formalism:* Hidden Markov models (HMMs) [26], [46] represent the probability of an observation sequence $P(\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T)$ where \mathbf{y}_t is the observation vector for time t . Input/output hidden Markov models (IOHMMs) [47] are a generalization of this model that represent the conditional probability of an observation sequence given an input sequence $P(\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T | \mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_T)$ where \mathbf{u}_t is the input vector for time t . For each state $i = 1, \dots, n$ in an HMM there is a transition distribution $P(x_t | x_{t-1} = i)$ and an output distribution $P(\mathbf{y}_t | x_t = i)$ where x_t is the discrete state at time t . In an IOHMM there is a conditional transition distribution $P(x_t | x_{t-1} = i, \mathbf{u}_t)$ and a conditional output distribution $P(\mathbf{y}_t | x_t = i, \mathbf{u}_t)$. The standard algorithm for training IOHMMs is an expectation–maximization (EM) algorithm that is a straightforward extension of Baum–Welch for HMMs. After training, the model can be interactively executed on an input sequence in a manner similar to a finite state machine.

In our case, \mathbf{u}_t is our state, including the position of the agent. The corresponding output \mathbf{y}_t is the position of the agent after moving. We think of the inputs as corresponding to the agent’s sensors and the output as the agent’s low-level actions expressed in terms of observables. By using the model interactively in a simulator environment, we can recreate the global behavior of the agent. We incorporate domain knowledge by modeling the output distributions as mixtures over the known low-level actions

$$P(\mathbf{y}_t | x_t = i, \mathbf{u}_t) = \sum_j c_{ij} P(\mathbf{y}_t | a_t = j, \mathbf{u}_t)$$

where a_t is the action at time t and c_{ij} is the mixing weight for state i and action j .

The actions are known, so only the mixing weights need to be estimated. Assuming we have calculated $P(\mathbf{y}_t | a_t = j, \mathbf{u}_t)$ for each time step and action—a task we discuss in the next section—the mixing weights can be reestimated using the standard formulas for mixtures within HMMs [6]. These mixing weights provide a soft mapping between the states of the model and the known actions. The mapping is not necessarily one-to-one, admitting models with several states corresponding to the same low level action and with states that correspond to a mixture of multiple low-level actions.

Our variation of IOHMMs is less of a discriminant model than the standard IOHMM formulations because the actual input–output mappings are fixed. Our primary

reason for selecting IOHMMs over, say, HMMs is that the conditional transition distributions allow us to more easily represent behaviors where switching between actions is triggered by sensors.

In our experiments we condition transition distributions on a set of binary sensory features (e.g., “bumped” or “see_target”), using the standard formulas for reestimation [47]. We represent the transition distributions as simple look up tables (as opposed to the neural networks used in some IOHMM work). We also experimented with an encoding where we ranked the features and took as the input the index of highest ranking nonzero feature. The ranking method greatly reduces the number of parameters and allows for simple diagrams to be drawn of the resulting models, but choosing a valid feature ranking can be difficult and requires domain knowledge of the relative importance of the features.

2) *Detecting Low-Level Actions:* To estimate the mixing weights during EM we need $P(\mathbf{y}_t | a_t, \mathbf{u}_t)$ at each time step; this is equivalent to detecting which behavioral assemblage (or blend of them) is active at each step. Unfortunately motion is complicated by interactions with the environment, especially near walls and irregularly shaped objects, so detection is not easy or unambiguous.

However, because we have executable versions of the behavioral assemblages, we can approximate the distribution through sampling, observing that sampling from $P(\mathbf{y}_t | a_t, \mathbf{u}_t)$ is equivalent to simulating a single time step of the agent’s motion. Formally we estimate the probability values using kernel density estimation

$$P(\mathbf{y}_t | a_t = j, \mathbf{u}_t) \approx \frac{1}{m} \sum_{k=1}^m \frac{1}{V_{Dh}} K\left(\frac{\mathbf{y}_t - \mathbf{y}_{tk}}{h}\right)$$

where \mathbf{y}_{tk} is sampled from $P(\mathbf{y}_t | a_t = j, \mathbf{u}_t)$, K is the kernel function, h is the bandwidth of the kernel, and V_{Dh} is the volume of the kernel function for the bandwidth and the dimension of the data D . In our experiments we used a Gaussian kernel. At each step and for each behavioral assemblage, we run a simulation many times to produce a distribution of points predicting the next position of the agent if that behavior were active. We then compare these point distributions to the actual next position. One advantage of this method is that it assumes nothing concerning the representation of the behavioral assemblages, as long as they can be executed in simulation. The method fully captures interactions with the environment to the extent that the simulation used can recreate them.

In experimental practice, this method of detecting which behavior is active provides 95% accuracy (detailed results are provided in the next subsection). Despite perfect knowledge of the behavioral assemblages, we still have errors because there are many situations where the

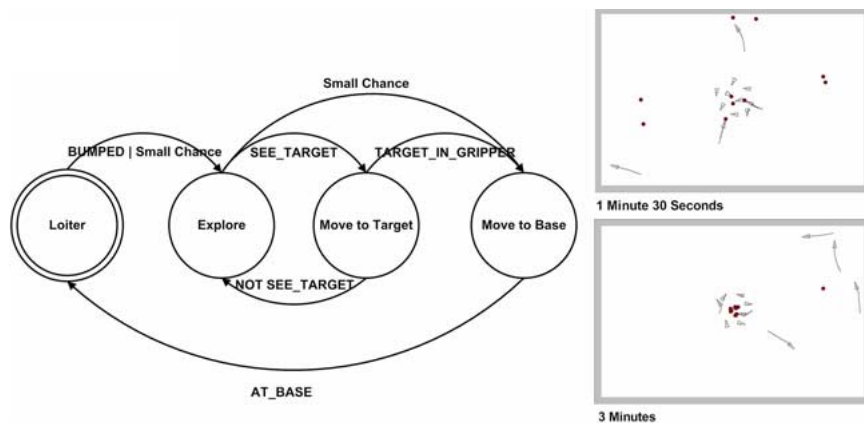


Fig. 15. Left: the original foraging model. Right: two frames from the running simulation that show the agents gathering food items and dropping them at homebase in the center of the arena.

outputs of several behaviors are indistinguishable. For example, when the agent bumps into a wall, most behaviors direct it to move away from the wall. In this case action probabilities become almost uniform, indicating that little can be said about the actions at such times. We expect situations like these to be typical in real-world data.

B. Using IOHMMs to Learn an Executable Model of Foraging

In order to test the effectiveness of our model learning strategy, we applied it to the task of learning an executable model of foraging from simulation data. Because we have the original program used to control the simulated agents, we are able to verify whether the system that was learned is correct by comparison.

We used the TeamBots simulation platform and motor schema-based control to program the agents. In this approach, (which is similar to a hybrid switching controller) agents can be in one of several behavioral states. Fig. 15 shows the state diagram for the original model and simulation screenshots. In each state a distinct set of primitive behaviors is activated. These groupings of behaviors and their corresponding weights and parameters are referred to as behavioral assemblages. Agents switch between states on the basis of perceptual events or triggers. For more detail on motor schema-based robot navigation, please see [41].

The agents are provided four behavioral assemblages: loitering at the base (the center), exploring for targets, moving toward the closest target, and moving back to base. There are four binary features that trigger transitions: bumping into something, seeing a target, being near the base, and holding an object. If an agent is exploring and bumps into a loitering agent, that agent also begins to explore. Some transitions are also triggered randomly (e.g., agents eventually return to the base if they do not find a target).

The simulation includes 12 agents and ten objects to collect. We ran the model in simulation for 5 min with 12 agents at 33 frames per second, waiting for all of the targets to be carried to base, recording at each frame, the position and orientation of all agents as well as the position of all targets. (Note that the output of the simulation log file is equivalent in format to the output of our vision-based tracking software.) Our objective is to learn a duplicate of the original model by examining the activities of the simulated agents recorded in the log file.

We ran 100 trials using EM to learn 100 executable models (for both ranked and unranked input). Some of the learned models correspond well with the original models, but some do not. In order to distinguish poor models from good ones, we calculated the likelihood scores for all the models using the standard forward probability recursion values [47]. Fourteen ranked and 20 unranked input trials yielded models with likelihood scores higher than the original model's. We noted that all models that successfully learned appropriate mixing weights had high likelihood scores. The highest scores for the ranked and unranked input models were essentially identical to the original model. We executed the models with the lowest likelihood scores that were still greater than the likelihood score for the original model (in some sense the worst of the successful trials), and found these models were able to recreate the behavior of the original model. As seen in Fig. 16, these models also recovered the structure of the original model.

There were flaws. In some models, bumping into an object while exploring would sometimes cause the agent to transition prematurely into the move-to-base state, slowing foraging. This appears to be a result of detection inaccuracies from bumping. Similarly, agents would drop a target early if they bumped into something while returning to base. Also if an agent was moving towards a target but another agent picked it up first, the first agent would loiter in place. One

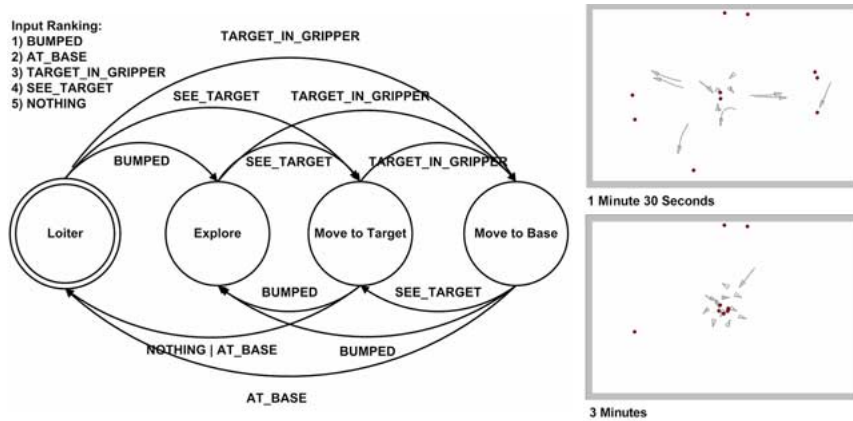


Fig. 16. Left: a learned foraging model. Note that there are more transitions between the behavioral states than in the original model. These transitions, in concert with the ranked inputs result in functional equivalence between this learned model and the original. Right: two frames from the running simulation that show the agents gathering food items and dropping them at homebase in the center of the arena. The behavior is nearly identical to the original.

reason for these apparent failures is due to the fact that some of these situations did not occur in the training run.

Models with low likelihood scores did not perform well. Their behavior was not at all similar to the original model. The most common failure mode was the absence of an important behavioral assemblage as a state in the model (this is more or less equivalent to learning the wrong mixing weights). One example “bad” model is illustrated in Fig. 17. Note that it includes two loiter states and it lacks a move-to-target state.

So far we have applied this method only to data from simulation (versus live ant data). One important reason for exploring simulation results first is that by using a simulation with agents that we program, we are able to directly compare the learned models with the ground-truth

original. We will not be able to do this with ant data because it is impossible to know for sure how ants are really “programmed.”

Nevertheless, we expect to apply this learning technique to the problem of learning models of real ant behavior in the near future. We will also augment the approach with low-level behavior learning so we can relax our assumption that the low-level behavioral assemblages are given.

V. CONCLUSION AND LIMITATIONS

The central assertion of this paper is that robotics has much to offer social animal researchers. Many of the problems studied in robotics, and multirobot systems

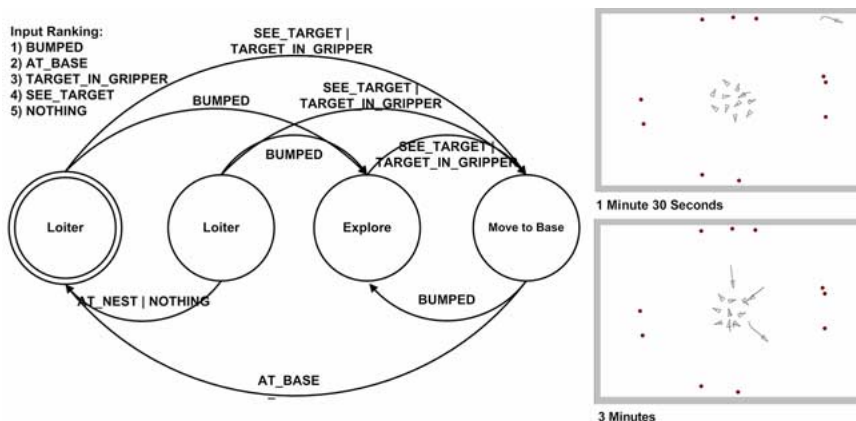


Fig. 17. Left: a low-likelihood learned foraging model. This model includes two loiter states and lacks a move-to-target state. Right: two frames from the running simulation that show that the agents running this model perform poorly.

research, have a dual in social animal research. We focused on three problems: multitarget tracking, automatic behavior recognition, and learning of executable models.

We have developed and demonstrated capable multitarget trackers, but we have not yet addressed a number of remaining challenges in this area. Ant tracking, for instance, is easier than bee tracking because the field of view is generally more sparse in the case of ants, and the background provides high contrast. In the case of bees, however, the field of view is covered by bees, with very little background visible. And the background that is visible is the same color as the bees. We are developing new algorithms to address these problems. We are also beginning to investigate new sensor modalities besides vision. One sensor in particular that we want to explore is the scanning laser range finder, and in particular, multiple distributed sensors.

With regard to behavior recognition, we have focused on recognizing social behavior—those activities that occur in the context of multiple animals interacting. Using geometric models of an animal's sensor system and labeled data, we are able to correctly label behavior with an accuracy of 90% to 95%. In some cases, however, our success is diminished by a high false positive rate—we identify actions that did not take place. We believe these failures can be reduced using HMMs and similar techniques.

We showed how IOHMMs can be used to represent executable models of behavior. We also showed how standard EM techniques can be used to learn executable models from observational data. Our system successfully learned a model of foraging behavior that was run in simulation. The learned model also runs in simulation and it emulates the original very well. Currently the weakness in the approach lies in the fact that we assume we know the low-level behaviors already, and that we need only learn the higher-level sequencing. However, in other work, we have been able to learn low-level models, and we expect to integrate the two approaches in the future.

Finally, we observe that the accuracy of our models are ultimately limited by the quality and breadth of our observations. For instance, if we do not observe an ant in a specific situation, its behavior in this situation will not be modeled. We will explore these limitations in future work. ■

APPENDIX I DETAILS REGARDING ANT EXPERIMENTS

A colony of *Aphaenogaster cockerelli* was raised from a fertilized queen captured in September 2000 near Tucson, AZ. At the time of experiments the colony was composed of about 400 workers and a queen. The colony did not include any alates of either sex. The colony was housed in

two 10-cm-square plastic brood chambers, each with a single 0.5-cm-diameter entrance hole. The chambers were kept in a larger, 25 cm by 40 cm container in which the ants were free to roam. The brood chambers include floors made of hydrostone into which water was occasionally squirted through a tube to maintain humidity. The colony was fed a diet of fresh apple slices once per week and two or three crickets each day. Fresh water was always available from a test tube with a cotton stopper. Temperature was maintained at about 21 °C throughout the experiment.

For the experiments, foragers (ants outside the brood chambers, on or near food), were selected randomly with forceps and placed in an otherwise empty 10 cm by 15 cm arena with vertical walls. The walls were treated with fluon to prevent escape. Two 30-min episodes were recorded, the first with 10 ants, the second with 20 ants. Both sets of ants were kept separated from the rest of the colony until the experiment was completed. In both trials the ants moved actively around the arena during the entire 30 min with no apparent degradation in their enthusiasm for exploration.

Video was captured by a JVC camcorder onto mini-DV tape. The camera was set to capture images at 30 Hz in a progressive scan (non-interlaced) mode. A 345-s (5 min, 45 s) clip was selected at random from each 30-min experiment. These clips were analyzed to produce the results reported below. The video segment used to validate the encounter detection software (outlined above) was from the same experiment, but from a different time period than was used to gather the quantitative data.

APPENDIX II DETAILS REGARDING SIMULATED ANT EXPERIMENTS

The simulated ants are initialized at random locations in a simulated arena the same size as the real ant experiments. Because our video data is acquired at 30 Hz, the simulation also progresses in discrete steps of 1/30th of a second. The simulated ants are constrained within the same size arena as in the real ant experiments. We did not model the sensor systems of the ants, because the simulated behavior only requires detection of collisions with other ants and the boundaries of the arena. The simulated ants act according to the following simple “program.”

- 1) Set `heading` to a random value.
- 2) Set `speed` and `move_duration` to random values.
- 3) Set `timer` = 0.
- 4) While `timer` < `move_duration` do
 - a) Move ant in direction `heading` for distance `speed/30`.
 - b) If collision with wall, set `heading` directly away from the wall, then left or right by a random amount.

- c) If collision with another ant, set heading directly away from other ant, then left or right by a random amount.
- d) Increment timer by 1/30 s.
- 5) Change heading by a random value.
- 6) Go to 2.

In order to make appropriate random selections for the variables `speed`, `heading`, and `move_duration`, we used data from the movements of the real ants. Ignoring encounters with other ants, we observed that the animals tended to move for some distance in a more or less straight line, then they would stop, occasionally pause, then move off in another slightly different direction. We calculated a

mean and standard deviation for the time, speed and change in direction for each of these movement segments. These parameters were then used to draw random numbers for the variables according to a Gaussian distribution. For `speed` and `move_duration`, negative values were clipped at zero. An example of trajectories generated by the simulation is illustrated in Fig. 12.

Acknowledgment

The authors are grateful for the suggestions and contributions of J. Bartholdi, D. Gordon, T. Seeley, M. Veloso, K. Wallen, and R. Herman.

REFERENCES

- [1] D. Gordon, *Ants at Work*. New York: The Free Press, 1999.
- [2] B. Webb, "What does robotics offer animal behaviour," *Animal Behav.*, vol. 60, pp. 545–558, 2000.
- [3] L. Parker, "Distributed algorithms for multi-robot observation of multiple moving targets," *Auton. Robots*, vol. 12, no. 3, 2002.
- [4] A. Stroupe, M. Martin, and T. Balch, "Distributed sensor fusion for object position estimation by multi-robot systems," in *Proc. 2001 IEEE Int. Conf. Robotics and Automation*, pp. 1092–1098.
- [5] V. Isler, J. Spletzer, S. Khanna, and C. Taylor. (2003). Target tracking with distributed sensors: The focus of attention problem. [Online]. Available: <http://www.citeseer.ist.psu.edu/article/isler03target.html>
- [6] J. Bruce, T. Balch, and M. Veloso, "Fast and inexpensive color image segmentation for interactive robots," in *Proc. Intelligent Robots and Systems (IROS) 2000*, vol. 3, pp. 2061–2066.
- [7] A. Feldman and T. Balch, "Representing honey bee behavior for recognition using human trainable models," *Adaptive Behav.*, vol. 12, pp. 241–250, Dec. 2004.
- [8] Z. Khan, T. Balch, and F. Dellaert, "MCMC-based particle filtering for tracking a variable number of interacting targets," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 11, pp. 1805–1819, Nov. 2005.
- [9] T. R. Balch, Z. Khan, and M. M. Veloso, "Automatically tracking and analyzing the behavior of live insect colonies," in *Proc. 5th Int. Conf. Autonomous Agents*, 2001, pp. 521–528.
- [10] Z. Khan, T. Balch, and F. Dellaert, "Efficient particle filter-based tracking of multiple interacting targets using an mrf-based motion model," in *Proc. 2003 IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS'03)*, vol. 1, pp. 254–259.
- [11] M. Powers, R. Ravichandran, and T. Balch, "Improving multi-robot multi-target tracking by communicating negative information," presented at the 3rd Int. Multi-Robot Systems Workshop, Washington, DC, 2005.
- [12] D. Reid, "An algorithm for tracking multiple targets," *IEEE Trans. Automat. Control*, vol. AC-24, no. 6, pp. 84–90, Dec. 1979.
- [13] Y. Bar-Shalom, T. Fortmann, and M. Scheffe, "Joint probabilistic data association for multiple targets in clutter," presented at the Conf. Information Sciences and Systems, Princeton, NJ, 1980.
- [14] T. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Sonar tracking of multiple targets using joint probabilistic data association," *IEEE J. Ocean. Eng.*, vol. 8, no. 3, pp. 173–184, Jul. 1983.
- [15] R. Deriche and O. Faugeras, "Tracking line segments," *Image Vis. Comput.*, vol. 8, pp. 261–270, 1990.
- [16] I. Cox and J. Leonard, "Modeling a dynamic environment using a Bayesian multiple hypothesis approach," *Artif. Intell.*, vol. 66, no. 2, pp. 311–344, Apr. 1994.
- [17] C. Rasmussen and G. Hager, "Probabilistic data association methods for tracking complex visual objects," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 6, pp. 560–576, Jun. 2001.
- [18] D. Schulz, W. Burgard, D. Fox, and A. B. Cremers, "Tracking multiple moving targets with a mobile robot using particle filters and statistical data association," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2001, pp. 1665–1670.
- [19] T. Balch and A. Guillery. Teamview. [Online]. Available: <http://www.borg.cc.gatech.edu/software>
- [20] Z. Khan, R. A. Herman, K. Wallen, and T. Balch, "An outdoor 3-D visual tracking system for the study of spatial navigation and memory in rhesus monkeys," *Behav. Res. Methods*, vol. 37, no. 3, pp. 453–463, Aug. 2005.
- [21] A. Panangadan, M. Mataric, and G. Sukhatme, "Detecting anomalous human interactions using laser range-finders," in *Proc. IEEE Int. Conf. Intelligent Robotics and Systems (IROS 2004)*, vol. 3, pp. 2136–2141.
- [22] A. Fod, A. Howard, and M. Mataric, "Laser-based people tracking," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA 2002)*, pp. 3024–3029.
- [23] K. Han and M. Veloso, "Automated robot behavior recognition applied to robotic soccer," in *Proc. 9th Int. Symp. Robotics Research (ISSR'99)*, pp. 199–204.
- [24] M. Tambe. (1996). Tracking dynamic team activity in Proc. Nat. Conf. Artificial Intelligence (AAAI '96). [Online]. Available: <http://www.citeseer.ist.psu.edu/tambe96tracking.html>
- [25] M. Huber and E. Durfee, "Observational uncertainty in plan recognition among interacting robots," in *Proc. 1993 IJCAI Workshop Dynamically Interacting Robots*, pp. 68–75.
- [26] F. Jelinek, *Statistical Methods for Speech Recognition*. Cambridge, MA: MIT Press, 1997.
- [27] T. Darrell, I. A. Essa, and A. Pentland, "Task-specific gesture analysis in real-time using interpolated views," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 12, pp. 1236–1242, Dec. 1996.
- [28] Y. A. Ivanov and A. F. Bobick, "Recognition of multi-agent interaction in video surveillance," in *Proc. Int. Conf. Computer Vision (ICCV)*, 1999, pp. 169–176.
- [29] D. Sumpter and S. Pratt, "A modeling framework for understanding social insect foraging," *Behav. Ecol. Sociobiol.*, vol. 53, pp. 131–144, 2003.
- [30] D. Gordon, R. Paul, and K. Thorpe, "What is the function of encounter patterns in ant colonies?" *Animal Behav.*, vol. 45, pp. 1083–1100, 1993.
- [31] D. Gordon, "The organization of work in social insect colonies," *Nature*, vol. 380, pp. 121–124, 1996.
- [32] S. Pratt, "Quorum sensing by encounter rates in the ant *temnothorax curvispinosus*," *Behav. Ecol.*, vol. 16, pp. 488–496, 2005.
- [33] I. Couzin and N. Franks, "Self-organized lane formation and optimized traffic flow in army ants," in *Proc. R. Soc. Lond. B*, vol. 270, pp. 139–146, 2002.
- [34] T. Starner, J. Weaver, and A. Pentland, "Real-time american sign language recognition using desk and wearable computer based video," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 20, no. 12, pp. 1371–1375, Dec. 1998.
- [35] R. Brunelli and T. Poggio, "Face recognition: features versus templates," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 15, no. 10, pp. 1042–1052, Oct. 1993.
- [36] A. J. Smola and B. Scholkopf, "On a kernel-based method for pattern recognition, regression, approximation, and operator inversion," *Algorithmica*, vol. 22, no. 1/2, pp. 211–231, 1998.
- [37] E. M. Arkin, H. Meijer, J. S. B. Mitchell, D. Rappaport, and S. Skiena, "Decision trees for geometric models," *Int. J. Comput. Geom. Appl.*, vol. 8, no. 3, pp. 343–364, 1998.
- [38] W. Schleidt, M. Yakalis, M. Donnelly, and J. McGarry, "A proposal for a standard ethogram, exemplified by an ethogram of the bluebreasted quail," *Zietschrift fur Tierpsychologie*, vol. 64, pp. 193–220, 1984.
- [39] B. Holldobler and E. Wilson, *The Ants*. Cambridge, MA: Belknap, 1990.

- [40] N. Johnson, A. Galata, and D. Hogg, "The acquisition and use of interaction behaviour models," in *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, 1998, pp. 866–871.
- [41] R. Arkin, "Motor schema-based mobile robot navigation," *Int. J. Robot. Res.*, vol. 8, no. 4, pp. 92–112, 1989.
- [42] R. Brooks, "A robust layered control system for a mobile robot," *IEEE J. Robot. Autom.*, vol. RA-2, no. 1, p. 14, Mar. 1986.
- [43] M. Egerstedt, T. Balch, F. Dellaert, F. Delmotte, and Z. Khan, "What are the ants doing? Vision-based tracking and reconstruction of control programs," presented at the 2005 IEEE Int. Conf. Robotics and Automation, Barcelona, Spain, 2005.
- [44] S. M. Oh, J. M. R., T. Balch, and F. Dellaert, "Data-driven MCMC for learning and inference in switching linear dynamic systems," in *Proc. Nat. Conf. Artificial Intelligence (AAAI)*, 2005, pp. 944–949.
- [45] D. Bentivegna, A. Ude, C. Atkeson, and G. Cheng. (2002). Humanoid robot learning and game playing using PCased vision. [Online]. Available: <http://www.citeseer.ist.psu.edu/article/isler03target.html>
- [46] Y. Bengio. (1996). Markovian models for sequential data. [Online]. Available: <http://www.citeseer.nj.nec.com/bengio96markovian.html>
- [47] Y. Bengio and P. Frasconi, "Input–output HMMs for sequence processing," *IEEE Trans. Neural Netw.*, vol. 7, no. 5, pp. 1231–1249, Sep. 1996.

ABOUT THE AUTHORS

Tucker Balch (Member, IEEE) received the Ph.D. degree from Georgia Institute of Technology (Georgia Tech), Atlanta, in 1998.

He was a postdoctoral fellow for one year, then a member of the research faculty at Carnegie Mellon for two years before he returned to Georgia Tech in 2001. He now directs the BORG Lab at Georgia Tech with F. Dellaert. He is an Assistant Professor in the School of Interactive and Intelligent Computing at Georgia Tech. He has published over 70 research articles and edited two books, including *Robot Teams* with L. Parker. His research interests include collaboration and communication in multirobot teams, and activity recognition and modeling in live social agent systems.

Prof. Balch is an NSF CAREER Award recipient.



Andrew Guillory received the B.S. degree in computer science from the Georgia Institute of Technology, Atlanta. He is currently working toward the Ph.D. degree at the University of Washington, Seattle.

His research interests include machine learning and learning models of behavior.

Charles L. Isbell, Jr., received the B.S. degree in computer science from the Georgia Institute of Technology, Atlanta, and the Ph.D. degree from the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology, Cambridge.

He is an Assistant Professor at the Georgia Institute of Technology's College of Computing. His research interests include developing machine-learning algorithms and tools to enable lifelong learning in autonomous systems.



Frank Dellaert (Member, IEEE) received the equivalent of the M.Sc. degree in electrical engineering from the Catholic University of Leuven, Belgium, in 1989, the M.Sc. degree in computer science and engineering from Case Western Reserve University, Cleveland, OH, in 1995, and the Ph.D. degree from Carnegie Mellon University, Pittsburgh, PA, in 2001.

He is an Assistant Professor in the College of Computing at the Georgia Institute of Technology, Atlanta. He has published close to 50 papers in journals and refereed conference proceedings, as well as several book chapters. His research focuses on probabilistic methods in robotics and computer vision. He has applied Markov chain Monte Carlo sampling methodologies in a variety of novel settings, from multitarget tracking to the correspondence problem. Before that, with D. Fox and S. Thrun, he introduced the Monte Carlo localization method for robot localization, which is now a standard and popular tool in mobile robotics.




Zia Khan received the B.S. degree in computer science and the B.S. degree in biology from Carnegie Mellon University, Pittsburgh, PA, in 2002.

He was a Research Scientist in the Georgia Institute of Technology's College of Computing between 2002 and 2004. He is currently an Algorithms Developer at Sarnoff Corporation, Princeton, NJ. His research interests include computer vision, machine learning, optimization, and interesting problems at the intersection between computer science and biology.




Adam Feldman received the M.S. degree in computer engineering from Case Western Reserve University, Cleveland, OH, in 2001. He is currently working toward the Ph.D. degree in the School of Interactive and Intelligent Computing at the Georgia Institute of Technology, Atlanta.

His research interests include automatically identifying behaviors and interactions of social agents such as bees, ants, monkeys, and humans.



Stephen C. Pratt received the B.S. degree in biology from Harvard University, Cambridge, MA, and the Ph.D. degree in neurobiology and behavior from Cornell University, Ithaca, NY, in 1997.

He has worked at Harvard University, the Massachusetts Institute of Technology, Princeton University, and the University of Bath, United Kingdom. He is currently an Assistant Professor in the School of Life Sciences at Arizona State University, Tempe. His research focuses on how the complex behavior of ant and bee colonies emerges from the actions and interactions of colony members, without direction from a well-informed central controller.



Andrew N. Stein received the B.S. degree in electrical engineering and the M.S. degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, in 2001 and 2002, respectively, and the M.S. degree in robotics from the Robotics Institute at Carnegie Mellon University, Pittsburgh, PA, in 2004. He is currently working toward the Ph.D. degree at Carnegie Mellon University.



His research experience as a graduate student and as an intern at Intel Research and NASA's Jet Propulsion Laboratory includes occlusion boundary detection, feature-based object recognition, multiobject tracking, segmentation, structure from motion, stereo, and structured light range imaging. He is interested in applications of robotics and computer vision for search and rescue, space exploration, assistive technologies, and fundamental scientific research.

Hank Wilde received the B.S. degree in computer science, specializing in artificial intelligence, from the Georgia Institute of Technology, Atlanta, in 2004.

He is a Systems Engineer working on mobile robotics for healthcare at Aethon, Inc. in Pittsburgh, PA.

Mr. Wilde was awarded the Presidential Undergraduate Research Award while doing research under T. Balch at the BORG Lab, for visualization and simulation software to help biologists investigate social insect behavior.