

Automatically Tracking and Analyzing the Behavior of Live Insect Colonies

Tucker Balch, Zia Khan and Manuela Veloso

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

trb || zkhan || mmv@cs.cmu.edu

ABSTRACT

We introduce the study of *live* social insect colonies as a relevant and exciting domain for the development and application of multi-agent systems modeling tools. Social insects provide a rich source of *traceable* social behavior for testing multi-agent tracking, prediction and modeling algorithms. An additional benefit of this research is the potential for contributions to experimental biology — the principled techniques developed for analyzing artificial multi-agent systems can be applied to advance the state of knowledge of insect behavior. We contribute a novel machine vision system that addresses the challenge of tracking hundreds of small animals simultaneously. Fast color-based tracking is combined with movement-based tracking to locate ants in a real-time video stream. We also introduce new methods for analyzing the spatial activity of ant colonies. The system was validated in experiments with laboratory colonies of *Camponotus festinatus* and several example analyses of the colonies' spatial behavior are provided.

1. INTRODUCTION

The behavior of social insects, especially ants, is a growing source of inspiration for computer scientists, especially those investigating multi-agent systems and robotics [3, 13]. “Ant algorithms” are employed in network routing systems, robot navigation and scheduling problems [3]. Most of the work in this area has focused on applying biological models of social insect behavior to information technology tasks.

In contrast, this work is focused on applying and extending research in computer science to the study of biology. Our ultimate goal is to achieve *full automation* of the following processes: (i) simultaneous tracking of multiple ants, (ii) recognition of individual and colony behaviors, (iii) acquisition of new single and multi-agent behavior models, and (iv) application of the acquired models to multi-agent software and robotic systems. We believe this will enable

a wide range of challenging and exciting research in automated multi-agent modeling, in particular for ant colonies, but also for observation and modeling tasks in general. This work will contribute:

- **New multi-agent science:** New multi-agent observing and tracking algorithms will provide a wealth of data for testing and developing multi-agent modeling tools.
- **New biological science:** The algorithms developed for this research domain will substantially advance the state of knowledge of social insect behavior.

Here we describe progress towards these goals, namely: setting up ant colonies for automated observation, machine vision algorithms for effective simultaneous tracking of multiple moving animals, and novel methods of analyzing a colony's spatial behavior. We contribute a fully implemented observation system for ant colonies that is reproducible by other researchers in their labs.

To motivate our interest in ants as a target for multi-agent research we begin with a brief introduction to ant behavior and the techniques used by myrmecologists to study and model it. Next we describe our laboratory setup, including details on keeping captive ant colonies. After that we present our machine vision algorithm for finding ants in images. Finally we show how the system has been used to track and analyze the behavior of captive colonies.

1.1 The Complexity of Ant Society

An ant colony is a complex system of individuals interacting with each other and their environment. Even though colonies have at least one queen, and they appear to act cooperatively and purposefully, there is no leader. Aggregate colony behavior emerges from the complex interaction of chemical cues, physical contact between individuals and external environmental pressures.

Nearly all ant species are *eusocial*; meaning that they care for their young cooperatively, there is a division into reproductive and sterile castes, and generations overlap (older individuals help raise younger generations). In most ant species, a single queen establishes a colony after she leaves her home colony and is fertilized (sometimes by multiple males). The queen establishes a nest and begins to rear non-reproducing workers who, in turn care for her and their siblings. Individual workers rarely live longer than 12 months, but a queen may live for 20 years or more and produce thou-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Autonomous Agents 2001 Montreal, Canada
Copyright 2000 ACM 0-89791-88-6/97/05 ..\$5.00

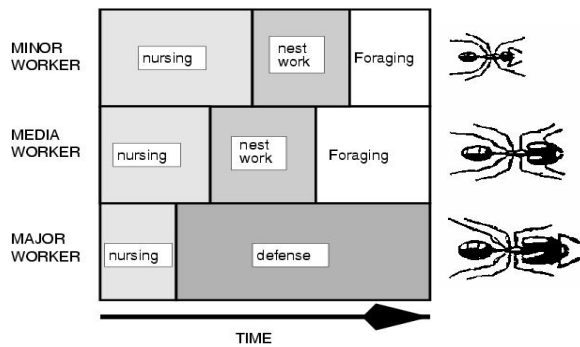


Figure 1: Individuals in a colony assume distinct task-based roles as they age. Individual morphology and environmental pressures affect the speed with which they move from role to role, and which roles they assume. (From Holldobler and Wilson, 1990).

sands to millions of workers. When the queen dies, however, the colony withers.

There are a number of tasks for workers in the colony. In general, but depending on the species, workers are dedicated to brood care (nursing), nest maintenance, foraging and colony defense. However, individuals are not committed to a single task for their entire life. In fact, they switch from task to task as they mature [9] (Figure 1). Newly eclosed (hatched) ants start their lives as brood care workers. Later, they move on to nest maintenance, and finally, foraging. Some species include castes morphologically suited for combat; these individuals become soldiers rather than foragers.

Myrmecologists have developed a number of methods for modeling ant behavior at the individual and colony level. An example of one kind of model, referred to as an ethogram, is provided in Figure 2. 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. Similar models have been developed for colony-level behavior as well (e.g. [14]).

Computer scientists will recognize a similarity between the diagram in Figure 2 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 (e.g. [11]). A goal of this work is to employ a similar approach to the task of learning ant behavior models.

The spatial aspects of ant behavior are also quite interesting and important. There is evidence, for example, that *Linepithema humile* colonies employ a strategy for exploring new spaces that is similar to what would result from Bayesian updating of expectations for success [12]. Other researchers have investigated how encounters between ants regulate behavior of the colony [5]. Outside the nest, ants organize efficient 2-dimensional foraging pathways [6], negotiate foraging territories with neighboring colonies [5], and wage war [14]. Of particular interest is a mechanism whereby colonies adjust their foraging strategies according




Figure 2: An ethogram of individual ant behavior. Behavioral acts, a_i , are linked by arcs indicating transitions from act to act. Thicker lines indicate higher probability transitions. (From Holldobler and Wilson, 1990).




Figure 3: The shape of foraging paths of *Linepithema humile* change as the density of ants increases (left to right). When the density is low, paths are straighter, but they are more convoluted when density is high. (From Gordon, 1999).

to the density of ants [5] (Figure 3).

1.2 Technical Barriers to Ant Research

All of the research outlined above relies on careful observation and recording of animal activity. Observation in the field can be especially arduous; researchers are often in place before dawn, and remain until dusk for weeks at a time to monitor the activities of their subjects [10, 5].

Even for laboratory experiments, collecting behavioral data is a time-consuming operation. Researchers must sit patiently for hours at a time to observe and record the actions of their subjects with pen and paper. There are several obvious limitations to this approach. Any lapse in attention by the observer, for instance, may result in missing a potentially significant event. Also, because colony activity is often distributed spatially, either one observer must split her attention between areas, or multiple observers must attend to the same experiment.

To address this, researchers are adopting technological methods (e.g. videotape) for collecting data. However, even when video recording is employed, a human observer must watch and record events manually. In this paper we describe a novel video processing algorithm that reliably tracks multiple (up to hundreds) of ants simultaneously. Each ant can be tracked individually, and its movement recorded digitally. Automating the monitoring and recording of animal activity can significantly impact the accuracy and breadth of research in ant behavior.

In the work most closely related to our approach, Gordon reports using image processing techniques to track ants in




Figure 4: A laboratory colony of *Camponotus pennsylvanicus* consisting of a queen (large), seven workers and brood. This colony was reared from a single, locally captured queen. The ants live in test tubes, moving from tube to tube according to their preferences for humidity.

her study of foraging patterns in *Linepithema humile* [5]. However, to our knowledge, details of the system she used were not published, so we are unable to compare it with our approach.

2. APPARATUS

Keeping laboratory colonies of hardy ant species requires perseverance and a watchful eye, but overall it is not too difficult. Primary considerations are: containment, temperature, humidity, and food. In our lab we keep two colonies of *Camponotus festinatus*, a species native to the southwestern U.S., and eight colonies of *Camponotus pennsylvanicus* (carpenter ants). The one year old *C. festinatus* colonies include one queen and about 250 workers each. The recently founded carpenter ant colonies number 3 to 15 workers each. These colonies were all raised from captured queens.

The colonies are housed in open 10mm by 75mm test tubes (Figure 4). The animals are allowed to move between the test tubes freely. Some test tubes are filled about $\frac{1}{3}$ with distilled water and fitted with a cotton stopper, others are dry. The inside of the test tubes with water is considerably more humid than the outside environment — this is important for insects who usually live in moist earth. The ants move from tube to tube according to their humidity preferences.

The test tubes are mounted in 6 inch diameter petri dishes with modeling clay. All of our colonies are small enough to live comfortably in these covered petri dishes. However, because we are interested in studying their exploration and foraging activities, we place the open petri dishes in larger containers (10 gallon aquaria) for observation. The walls of the aquaria are treated with machine oil to prevent escape.

Color video cameras with wide angle lenses are mounted above the the observation aquaria looking downward to provide full coverage of the experimental arena (Figure 5). The camera output is connected to a Pentium computer that captures images using an off-the-shelf video capture card. The capture card provides 640 by 480 pixel images at 30Hz to the image processing algorithms. The video stream is processed in real-time to calculate the locations of ants in




Figure 5: Experimental arena. The colony is housed in a petri dish on the right, three food items are placed on the left.

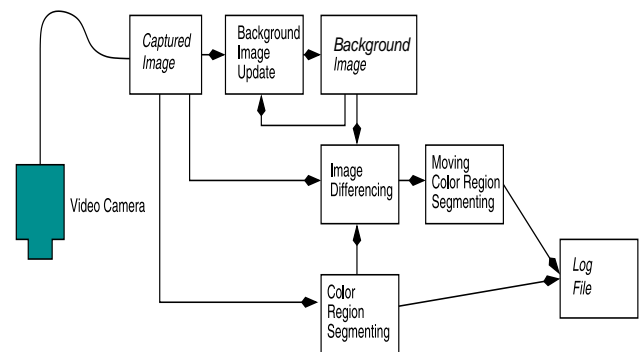


Figure 6: Block diagram of image processing operations for finding moving objects in the video stream.

the arena.

We describe the details of image processing to locate the ants in the next section.

3. FINDING ANTS IN IMAGES

One of the main contributions of this work is an algorithm for finding ants in images and tracking them over time. Even though the laboratory arena provides a high-contrast background that aids image processing, the task is complicated by several factors:

- **Small targets:** As Figure 5 illustrates, the ants are rather small in the image.
- **Color ambiguities:** Because ants are dark (nearly black) their color is ambiguous with respect to other items in the arena (e.g. food, shadows and waste).
- **Noise:** Dark areas in the image are noisy.

To address these issues we use a hybrid approach combining color-based classification and movement-based classification techniques. The basic idea is to use color classification (a fast operation) to identify regions in the image that should be further scrutinized for indications of movement using more costly image differencing.

An overview of the system is provided in Figure 6. The output of a video camera overlooking the arena is fed into

a capture card that provides digital color images at 30Hz. The image stream is used to build the *background image* — an approximation of the arena with all moving objects removed.

Incoming images are also processed by a color region segmentation algorithm. Pixels that match specific color specifications are grouped together and identified by *bounding boxes*. A bounding box is a rectangular region in the image that contains the specified color.

The color region bounding boxes are used to identify regions in the incoming image to be further analyzed for movement. Pixels that contain moving objects are detected using image differencing (described below). Finally, bounding boxes describing the colored regions of interest, and colored regions that contain moving objects are written to a log file for later examination. Now we describe the processing steps in more detail.

3.1 Tracking by Color

Images are initially processed using the CMVision algorithm [1]. CMVision offers fast and reliable color-based classification and has been successfully applied to a variety of robot vision tasks. Color classes are specified using six threshold parameters in a three-dimensional color space (YUV or RGB). The six thresholds for each color correspond to upper and lower bounds in each dimension. In addition to classifying pixels by color, CMVision groups adjacent pixels of the same color together as segmented regions. On a medium power workstation CMVision can classify hundreds of objects of up to 32 different colors in images at 30Hz.

CMVision is quite effective in identifying marked objects. It is possible in many robotic applications to mark the environment or relevant objects for color identification. However, it is difficult and sometimes dangerous (for the ants) to mark ants in a similar manner. To complicate matters, most species of ants are black or brown, and these colors are common even in a controlled laboratory environment. The classifier cannot distinguish the black pixels of an ant from the black pixels along the edge of a petri dish.

Fortunately, ants can be distinguished from other objects of the same color by their movement. By filtering most of the image using color first, we are able to focus the more costly search for movement on regions that are likely to contain an ant.

3.2 Tracking by Movement

Frame differencing is a standard technique for finding movement in a series of images. Pixels in the current camera image are compared with the corresponding pixels in the previous camera image. If a pixel has changed sufficiently it is classified as containing movement. Typically, and also in our approach, the differencing is only applied to the intensity, or Y dimension of the images.

Frame differencing is effective, but it is subject to several limitations. First, because the time scale over which movement can be detected is short, (usually only one or two image frames or 33 to 66 milliseconds), only rapidly moving objects are detected. Furthermore, frame differencing is generally not effective at extracting all of the relevant moving pixels if the object is uniform in color; only the pixels along the edges of the object will be classified as moving.

In the approach we use, referred to as *adaptive background subtraction*, an image representing the scene devoid of mov-



Figure 7: Use of background differencing to find moving objects. Left: Background image computed using a running locally weighted average. Center: Sample “live” image. Right: Difference between background and live image reveals the locations of two ants.




Figure 8: A case where greedy association fails. Two ants, a1 and a2, are shown at times t and $t+1$. The arrows show the incorrect greedy association that matches one ant at a time (ant a2 is matched before ant a1) based on minimum distance.

ing objects is computed by averaging camera images over time [2]. We subtract the *current* image from the *background* image to find movement. The process is illustrated in Figure 7. Because the background image is stable and computed over a long period, it avoids most limitations of simple frame differencing.

The background image is computed using a locally weighted running average as follows:

$$B_{ij} = (\alpha - 1)B_{ij} + \alpha I_{ij}$$

where B_{ij} is the pixel in the i th column and j th row of the background image, I_{ij} is the corresponding pixel in the current camera image, and α is a parameter specifying the “learning rate” or speed at which the background image adapts to changes.

When α is set to a very low value, new objects in the scene only become part of the background image if they remain for a long time (typically α is set to a value near 0.0005). The idea is that moving objects (e.g. ants) will only occupy a particular location for a short period, and will therefore have little impact on the corresponding pixels in the background image at that location.

In the final steps of the algorithm, the intensity values of pixels matching the “ant color” specification are subtracted from the background image. If this difference is greater than a specified threshold value the pixel is considered “moving ant color.” The threshold may be set to any value between 0 and 255; 35 was used in the experiments here. Next, connected pixels of “moving ant color” are grouped into regions; sufficiently large regions are recorded as containing an ant.

Summarizing the approach, first we search for regions of “ant color” in the image, then we examine those regions us-

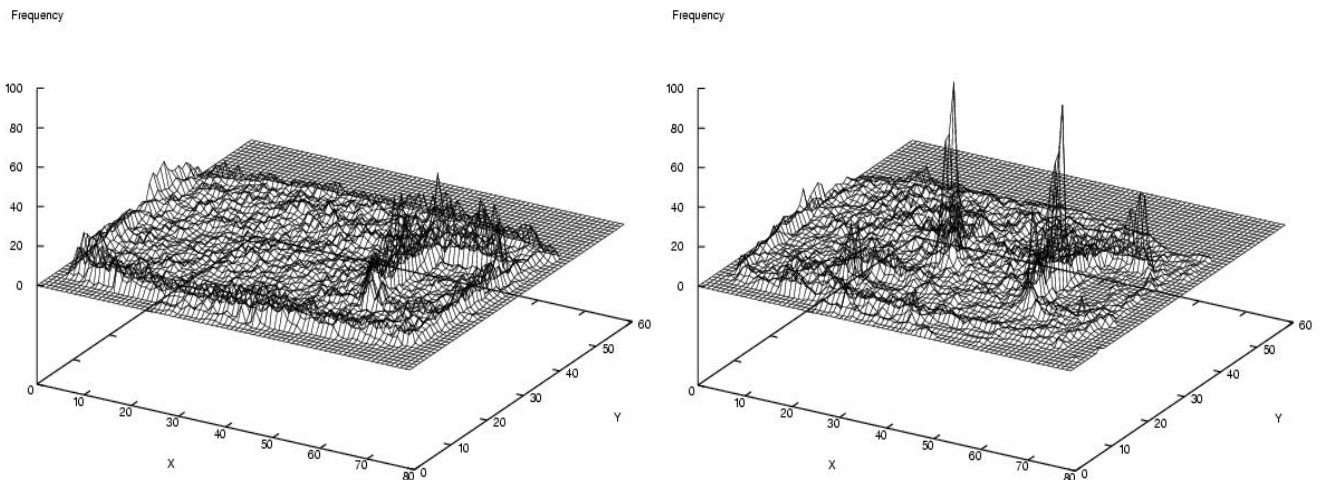


Figure 9: A new method for evaluating the spatial behavior of multi-agent systems. This experiment compares activity in the arena depicted in Figure 5 with and without the presence of food. Both graphs depict the number of visits by ants to each location in the arena over a 30 minute period. Left: Foraging activity with no food present. Right: Activity with food placed in the center of the arena. A distinct peak is evident at the location of the food.

ing background differencing to find regions of “moving ant color.” Sufficiently large regions of “moving ant color” are classified as ants. Our approach to building a background image as a means of detecting motion was developed independently, but is similar to the method proposed by Collins, *et al* [2]. A key difference is our use of color filters to limit the area examined for movement.

3.3 Associating Observations with Individuals

The data gathered through observation corresponds to multiple observed agents, in this case, ants. To identify colony behaviors, it is crucial for our automated system be able to track *individual* ants. When tracking artificial creatures, such as robotic agents, it is possible to add a pattern on the robot for the purpose of identification. When using live creatures like ants, however, adding identification patterns to the moving animals is not always feasible. In earlier work we developed a data association algorithm that is capable of identifying and tracking multiple soccer playing robots without any specific identification [7]. We have now applied and extended this algorithm to the problem of individual ant tracking.

Formally, data association addresses the problem of retaining the ant identification in subsequent frames gathered by the observation system. Our algorithm retains association based on the spatial locations of the ants. We assume that the starting positions of all the ants are known. We then use a minimum distance scheme to retain association between consecutive frames under the assumption that the ants will move only within a circle of some maximum distance over consecutive captured frames.

Given two consecutive frames gathered at times t and $t+1$, the points corresponding to each ant detected at time $t+1$ are matched with the *closest* positions of each detected ant of the frame at time t . This greedy association algorithm is computationally effective but it can generate an incorrect match, as shown in Figure 8.

There is an improved algorithm that can solve this prob-

lem. This algorithm generates all possible sets of observed matching points between two consecutive frames and then calculates the total fitness of each of the sets globally according to a least square criteria:

$$\sum_{i=1}^N (dist(prev_i, cur_i))^2,$$

where $(prev_i, cur_i)$ are the i^{th} matching pair. And the function $dist(x, y)$ is the Euclidean distance. The set of matches that minimizes the above criteria is selected. Even so, this algorithm does not guarantee perfect association, in particular with our cluttered environment, but the implementation has shown to be robust.

4. RESULTS

We have assessed the accuracy and utility of the system in several experimental observations of *Camponotus festinatus* colonies.

4.1 Quantifying Spatially Distributed Activity

One of the important contributions of this work is a new method for evaluating the distributed spatial activity of insect colonies. This capability is critical to the investigation of distributed behavior, for instance to evaluate how neighboring colonies establish boundaries between their foraging areas. Our method represents spatial activity as a three-dimensional surface. The two-dimensional arena is divided into an array of “bins.” Each time an ant enters the area corresponding to a bin, the bin’s value is incremented. A three-dimensional surface is constructed where the height at each position indicates the number of visits to that area. The surface shows peaks in areas of concentrated activity.

Figure 9 illustrates an example analysis. In this experiment we evaluate how the presence of food impacts the spatial activity of one of our laboratory colonies. The two activity plots generated in this experiment are quite different, indicating that the colony changes its exploration strategy

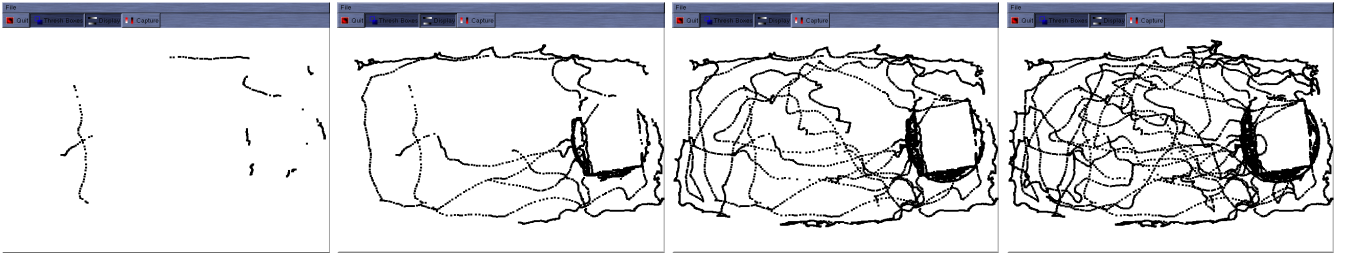


Figure 11: Tracking the movement of multiple ants simultaneously. This sequence of images depicts the paths of multiple ants patrolling the arena depicted in Figure 5. The snapshots, sequenced from left to right, were taken at 60 second intervals; the entire sequence was recorded in three minutes. The ants do not often visit the area on top of their nest (a cardboard square).




Figure 10: Spatially distributed interactions between ants. The plot shows the number of times ants interacted with one another across the two-dimensional arena. This type of analysis could be employed to study “skirmishes” between nearby colonies sharing a foraging arena to determine how they establish territories.

depending on whether food is present.

The surface on the left of Figure 9 shows activity over 30 minutes when no food is available in the arena. Note that the entire arena is covered fairly evenly by the exploring ants. The outline of the petri dish where the ants live, and the boundary of the arena are evident as slightly raised “hills.” This may indicate that the ants use vertical walls as references for navigation [5].

The image on the right in Figure 9 also shows activity over 30 minutes. However during this experiment, food was placed in the center of the arena. A large peak is clearly evident in the center of the surface corresponding to the location of the food. Two other characteristics of this plot are worth noting. First, activity behind and near the nest is reduced, indicating that there is less interest in exploring those areas when food is available elsewhere. Also, during these experiments, we noticed that the ants often interact at the entrance to the nest as they return from gathering food. These interactions are reflected in the activity peak on the side of the nest towards the food.

In related work, Goldberg and Mataric suggest quantification of the interactions between agents as a tool for evaluating the effectiveness of foraging strategies [4]. We are able to conduct a similar analysis of the interactions between ants under observation. Each time two ants are present in the same small area, the corresponding bin is incremented. Figure 10 illustrates the results of this evaluation in the arena over 30 minutes after food is placed in the middle. Increased numbers of interactions are evident near the nest and at the location of the food.

4.2 Tracking Multiple Ants Simultaneously

A long-range goal of this work is to infer the behavioral state of individual ants by evaluating traces of their movement. To support that goal, our system must be able to recognize and record the movement of multiple individual ants simultaneously.

Our system is able to trace the movement of multiple individuals simultaneously. An example of this capability is provided in Figure 11. In this experiment we recorded the activity of ants in the arena over a three-minute period. During this time, four to six ants explored the entire arena.

4.3 Performance: Accuracy and Efficiency

Many investigations of ant behavior rely on counting the number of animals in specific regions over time. Our objective in this experiment was to evaluate the accuracy of the




Figure 12: The number of insects recognized in the experimental arena by the automated tracking system (dotted line) compared with the number of insects actually present. Average error is 1.2 ants per observation. The evaluation was conducted over one hour, with samples taken every 30 seconds.

vision system in counting the number of ants outside the nest, exploring the arena. To do this, we compared human observations with the automated system’s.

Using the arena illustrated in Figure 6, a human observer counted the number of ants outside the nest at 30 second intervals over one hour. The human’s results were compared with the tracking system’s count at each of 120 samples. The two counts are compared in the graph in Figure 12. There were an average of 10.5 ants present in each sample. On average the human and our vision system differed by 1.2 ants (about 11%) at each sample. We repeated this experiment twice with approximately the same results.

To evaluate the efficiency of our algorithms tests were conducted on a 700MHz Pentium III computer running the Linux operating system. On average, it takes 42ms to capture an image and process it. When allowed to run as fast as possible, the system is able to process 24 frames per second. At a moderate 10 frames per second, processing utilizes 35% of the CPU resource, leaving the rest available for additional analysis tasks.

Compressed log files consume about 3MB per hour of observation. This enables 20 days of observation data to be saved on a single CDROM.

The scalability of the system was tested by evaluating its ability to track 100 targets simultaneously. There was no noticeable degradation in performance. This is because the fixed costs of searching the entire image for color classification predominates. The small targets cover only a small percentage of the image, and the algorithm only spends additional time on pixels of the appropriate color. The cost of associating tracks with observations is a separate post-processing step that is not included in this performance evaluation.

5. LIMITATIONS

Our tracking software has proven to be quite reliable, but there are still a few limitations to be addressed, including:

- **Occlusion:** Ants are sometimes occluded by the walls of the petri dish they live in.
- **Clumping:** When two ants are very close, or on top of one another, the system may count them as only one ant.

- **Splitting:** In some cases, the bounding box for one ant may split into multiple bounding boxes (e.g. a specular reflection may confuse the system). In these cases the system will count more ants than are actually present.
- **Motionless ants:** If an ant remains motionless for too long of a time, it merges into the background image and can no longer be tracked.

In our continuing work we are investigating methods for eliminating these sources of error. Clumping and splitting can be addressed in a post-processing step that evaluates the sizes of bounding boxes. Boxes that are very large, for instance, are likely to enclose multiple ants. Occlusion can be addressed using a memory of tracked objects — when an object disappears, it should not be immediately dropped from the list of tracked objects.

Because stationary objects become part of the background, stationary ant colored objects (e.g. refuse, food objects, shadows) are not classified as ants. However, when an ant bumps an ant colored object, the object will be classified as “moving” for some time, but it will eventually be classified as background again. Similarly, an ant that remains stationary for a long time may also become background and no longer be tracked. The rate at which this occurs depends on α . In our experiments we set α so that an ant must remain stationary for at least 15 minutes before it is no longer tracked. Even so, there is sufficient information in the logged data to recover the positions of ants that remain motionless for long periods.

6. CONCLUSION

We propose the study of live insect colonies as an interesting and challenging domain for multi-agent systems modeling research. Ant societies are complex systems of interacting individuals regulated by chemical signals, physical contact and external environmental pressures. Computer science already draws from biology — specifically by applying “ant algorithms” to many information processing tasks. In contrast, we argue for the application of computer science techniques to the study of biology.

We establish the feasibility of this research direction by describing how ant colonies can be kept successful in the lab

and by contributing a novel computer vision algorithm capable of reliably and accurately tracking the activities of hundreds of insects simultaneously. The hybrid vision algorithm uses a combination of color-based tracking and movement-based tracking to find ants in an image.

Two novel techniques for accessing the spatial activity of ant colonies are presented. They are illustrated using example analyses of the foraging behavior of a captive colony of *C. festinatus*.

6.1 Towards Behavior Recognition and Modeling

Our longer-range goal is to recognize colony behavior by evaluating the trails of multiple individual ants. We have previously developed an extension of Hidden Markov Models, called Behavior Hidden Markov Models (BHMMs) to describe behaviors in robot systems [8]. We also developed an algorithm, using this representation, for automatically recognizing behaviors of single robots [8]. One of the main aspects of this work involves mapping from observations to possible internal states of the robots. We realized that the sequence of actions of a robot left behind a “trace” of their trajectories. The trace provides the observation data. We then mapped the observations into the states of a Hidden Markov Model, where each state captures the transitions between a particular behavior class. We successfully applied our approach to the recognition of a few alternative behaviors (up to 10) for a single robot.

The system we have developed for observing ants provides the same type of observational data used successfully to recognize behavior in robots. We are encouraged by the fact that myrmecologists utilize a representation for behavior (ethograms) that is quite similar to BHMMs. As we continue this research we expect to encode ethograms as BHMMs and use them to implement the same capability for recognizing behavior in ants. Furthermore, in the longer term we plan to develop methods for learning the BHMMs themselves.

7. ACKNOWLEDGMENTS

The authors thank Steve Prchal for information on the care and feeding of *C. festinatus*, Randy Morgan for information on the design of habitats for ant colonies and John Rawlins for advice on experimental technique. Hans Moravec provided useful suggestions on background subtraction and James Bruce assisted with machine vision issues.

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

8. REFERENCES

- [1] J. Bruce, T. Balch, and M. Veloso. Fast and inexpensive color vision for interactive robots. In *Proc. IROS-2000*. IROS, 2000.
- [2] Collins, Lipton, Kanade, Fujiyoshi, Duggins, Tsin, Tolliver, Enomoto, and Hasegawa. A system for video surveillance and monitoring: Vsam final report. Technical Report CMU-RI-TR-00-12, Robotics Institute, Carnegie Mellon University, 2000.
- [3] M. Dorigo and G. Di Caro. The ant colony optimization meta-heuristic. In D. Corne, M. Dorigo, and F. Glover, editors, *New Ideas in Optimization*, pages 11–32. McGraw-Hill, 1999.
- [4] D. Goldberg and M. Mataric. Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings, AAAI-97*, pages 637–642, July 1997.
- [5] D. Gordon. *Ants at Work*. The Free Press, New York, 1999.
- [6] J. Haefner and T. Crist. Spatial model of movement and foraging in harvester ants (*pogonomyrmex*)(i): The roles of memory and communication. *Journal of Theoretical Biology*, 166:299–313, 1994.
- [7] K. Han and M. Veloso. Reactive visual control of multiple non-holonomic robotic agents. In *Proceedings of the International Conference on Robotics and Automation*, Belgium, May 1998.
- [8] K. Han and M. Veloso. Automated robot behavior recognition applied to robotic soccer. In *Proceedings of the 9th International Symposium of Robotics Research (ISSR'99)*, pages 199–204, October 1999. Snowbird, Utah.
- [9] B. Holldobler and E. Wilson. *The Ants*. Harvard University Press, Cambridge, MA, 1990.
- [10] B. Holldobler and E. Wilson. *Journey to the Ants*. Harvard University Press, Cambridge, MA, 1994.
- [11] N. Johnson, A. Galata, and D. Hogg. The acquisition and use of interaction behaviour models, 1998.
- [12] P. Nonacs and J. Soriano. Patch sampling behavior and future foraging expectations in argentine ants, *linepithema humile*. *Animal Behavior*, 55:519–527, 1998.
- [13] R. Vaughan, K. Stoy, G. Sukhatme, and M. Mataric. Whistling in the dark: Cooperative trail following in uncertain localization space. In *Proc. Autonomous Agents 2000*, Barcelona, Spain, 2000.
- [14] M. Whitehouse and K. Jaffe. Ant wars: combat strategies, territory and nest defence in the leaf-cutting ant *atta laevigata*. *Animal Behavior*, 51:1207–1217, 1996.