Social Behavior in a Team of Autonomous Sensors

Yasuaki Sakamoto and Jeffrey V. Nickerson
Center for Decision Technologies
Howe School of Technology Management
Stevens Institute of Technology
Hoboken, NJ 07030 USA

Abstract— **Probabilities of physical attack are often determined by various environmental factors. As the environment changes, the probability of attack associated with an area changes. In such dynamic environments, autonomous sensors are potentially useful to optimally cover regions that have high probabilities of attack. We present results from agent-based simulations, in which autonomous sensors "forage" a space to find areas with high attack probabilities. Simple heuristics often resulted in optimal coverage of the attack regions, without a centralized control. By varying how quickly sensors respond to a threat, we can encourage some sensors to cover some areas, and others to hang back and defend different areas, allowing them to distribute optimally as a team. The idea of making team members hang back may seem counterintuitive. In fact, people often converge all at once to respond to an immediate threat. Our results show that it is useful to have some agents remain behind, in case the environment changes.**

I. Introduction

Probabilities of physical attacks are often determined by various environmental factors. For example, divers cannot swim against strong currents, and thus the probability of such an attack is low. Furthermore, the probability of attack associated with an area changes as the environment changes. For instance, the time of day and the weather conditions can affect the strength of currents.

There is another reason to pay attention to the environment. In many security situations, the probability of an attack at any particular instant is small, and the time between attacks is long. Sensors, then, are rarely detecting target signals. But the environment is changing. The sensors can monitor the environment and learn how to respond to it. Learning about the always-surrounding environment is a useful proxy for learning about the rarely-appearing target.

In dynamic environments, autonomous sensors are potentially useful for optimally covering regions that have high probabilities of attack. A problem faced by mobile agents is how to search the environment for resources (e.g., potential attack regions). When the environment consists of other agents that are also searching for resources, the optimal strategy for an agent is no longer a simple function of the distribution of resources but is also a function of the strategies adopted by other agents.

In the current work, we present initial results from agent-based simulations, in which autonomous agents or sensors "forage" a space to find areas that have high probabilities of attacks. Instead of following a centralized control, the agents all obey the same simple heuristics that control their behaviors. The interactions of individual agents give rise to the emergent group-level behavior, without leaders ordering the organization.

We examined two sets of rules in the current simulations. In one condition, all agents moved toward the region with a high attack probability as quickly as possible. In the other condition, the agent closest to the area with high attack probability moved toward the area more quickly than the other agents.

A. *Drawbacks of All Sensors Moving Quickly*

The strategy to move all sensors to the attack region as quickly as possible seems intuitive. Rather than having sensors in areas where attacks are highly unlikely, we want to move more sensors to the region that has high probabilities of attacks. People do tend to converge all at once to respond to an immediate threat.

However, significant problems can arise when everyone in a team responds to a threat at the same speed. One problem is that responding to a new threat can be quite difficult. For example, if all team members responded to an alert in area X, it may take a while for everyone to respond to a new alert in area Y. Another problem is that attending to two threats at the same time will be impossible if all team members move together. These issues are important because an attack is often followed by more attacks and simultaneous attacks are not uncommon.

B. *Benefits of Some Team Members Hanging Back*

By making the agent closest to the attack area move toward the area more quickly than the other agents, we can encourage some sensors to cover some areas, and others to hang back and defend different areas. This hold-back strategy deals with the shortcomings of all team members moving at the same speed; it allows the agents to be ready for new situations and provides more coverage when there are multiple areas that are likely to be attacked. As an example, imagine children playing soccer. Children all run to the ball and congregate around the ball, analogous to all team members responding to a threat together. When the ball escapes the crowd, it may take a while for everyone to run to the ball again. If some children remain behind, the escaped ball may fall in front of them, allowing a quick recovery of the ball.

Although people often do not hang back in emergency responding, they do in other situations. For example, in addition to using their knowledge of resource density to

1-4244-1330-3/07/$25.00 ©2007 IEEE.

move to profitable regions, people use their knowledge of the locations of other foragers to distance themselves from the crowd [1]. By hanging back, agents have the opportunities to discover new areas, allowing them to distribute optimally as a team. Researchers covering the domain of inquiry, neurons specializing for certain information, and ants foraging for food are some examples of group-level, or social, behaviors that emerge from some agents remaining behind [2]. One reason why people tend to converge all at once in emergency situations, instead of some people hanging back, may be that time pressure and emotional factors associated with crises interfere with people's cognitive abilities, such as quickly but carefully analyzing the situations and thinking ahead.

II. Simulation

In our simulation, each autonomous sensor was randomly assigned a location on the 370 pixels by 320 pixels grid world. We used three mobile sensors. In Fig. 1, sensors are represented by black dots and regions of potential attacks are represented by green patches. For example, static sensors can be distributed across a river that transmit signals based on environmental data. For instance, the static sensors can send stronger signals when the current is weaker. Green patches represent the signals from these static sensors, where a larger green patch indicates stronger signals, and hence a higher probability of attack. Thus, we want to have more mobile sensors in larger green patches. Moreover, we want to have all green patches covered when there are enough mobile sensors to cover all of the patches.

A. Simple Rules Governing Sensor Behavior

In our proposed scheme, the following three simple steps determined the behavior of the autonomous sensors.

1. One green pixel (i.e., signal) was randomly selected from all the green pixels present. Thus, when the probability of attack was higher, the signal was more likely to be selected. We assumed that all the mobile sensors received a signal from the same source at a given time, which can be achieved by making the static sensors send signals asynchronously.
2. The mobile sensor that was closest to the selected green pixel moved toward the green pixel with one speed, determined by the value of the slider labeled *Closest Rate* in Fig. 1. We assumed that the distance was related to the strength of a signal received by a sensor and that the sensors were able to communicate with one another to determine who was closest to the selected pixel (i.e., who received the strongest signal).
3. All of the other sensors moved toward the selected green pixel with another speed, determined by the value of the slider labeled *Not Closest Rate* in Fig. 1.

Each autonomous sensor, i, updated its position, Pos_i, according to

$$\Delta Pos_i = Speed_i \times Direction_i \quad (1)$$

where *Speed* determined how fast the sensor moved as described above and *Direction* determined to which direction the sensor moved. We assumed that the sensors were able to determine the directions of the signal source (e.g., by way of triangulation). In this simulation, each autonomous sensor selected one of 8 possible directions at a given time that reduced its distance to the signal source: North, North East, East, South East, South, South West, West, or North West.

B. Results

The simple heuristics proposed in the current work often led to optimal coverage of the attack regions. By changing speed, we could encourage some sensors to cover some areas, and others to hang back and defend different areas.

Each box in Fig. 1 displays a snapshot of a continuous simulation at a different time, after the sensor movements became stable under each condition. The same patterns of simulation results were obtained when the attack regions (i.e., green patches) changed shapes and positions from time to time or moved continuously (i.e., target chasing). These simulations (Java Applets) can be found at http://personal.stevens.edu/~ysakamot/.

1) All Sensors Moving Quickly: In the top part of Fig. 1, the closest sensor and the other sensors moved at the same speed, resulting in every sensor going to the same place. The sensors converged all at once. This strategy led to a quick coverage of an attack region when there was a single region to cover or when one region had a higher probability of an attack than the other. However, the sensors failed to cover each region when there were multiple regions to cover. As can be seen in the rightmost condition in the top part of Fig. 1, the team became stuck in the middle of the two regions when the two regions had the same probability of attacks.

2) Closest Moving Faster: In the bottom part of Fig. 1, the closest sensor moved faster than all the other sensors. This parameter setting allowed the sensors to become specialized to certain regions and as a group optimally cover a set of regions; each region was covered and more sensors were allocated to larger regions. When there were three autonomous sensors and two regions with the same attack probability, the extra sensor stayed in the middle of the regions, as can be seen in the leftmost box in the bottom part of Fig. 1. This behavior could be useful when communication distance is limited; the middle sensor can serve as a bridge.

III. Conclusions and Future Research

In many security situations, the probability of an attack at any particular instant is small, and the time between attacks is long. Sensors, then, are rarely detecting target signals. But the environment is changing, and the probability of a physical attack associated with a region changes as the environment changes. The sensors can monitor the environment and learn how to respond to it. Learning about the always-surrounding environment is a useful proxy for learning about the rarely-appearing target.

In changing environments, autonomous sensors can be useful for optimally covering the potential attack regions. In our simulations, the sensors followed simple heuristics

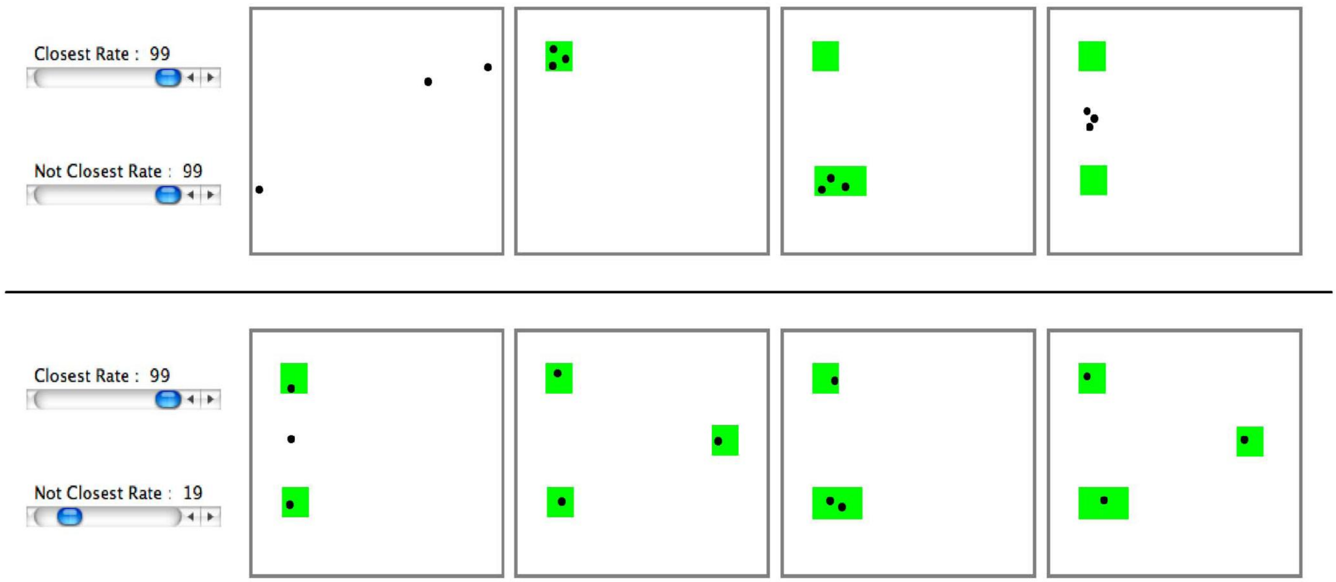

Fig. 1. Black dots are autonomous sensors and green patches are regions of potential attack. Larger green patches have higher probability of attack. In the top part, the sensor closest to a selected pixel and the other sensors moved at the same speed, leading to all sensors converging at once. With this parameter setting, all team members migrated to a region that had higher probability of attack and they became stuck in the middle of the two regions that had the same probability of attack. In the bottom part, the closest sensor moved faster than all the other sensors. This parameter setting allowed the sensors to become specialized to certain regions and as a team optimally cover the areas; each region was covered and more sensors were allocated to larger regions.

and adapted to varying environments by themselves. The interactions of sensors resulting from the simple heuristics often led to social structures in which sensors are allocated optimally, without a centralized instruction set. By varying how quickly sensors respond to a threat, we can encourage some sensors to cover some areas, and others to hang back and defend different areas, allowing them to distribute optimally as a team. The idea of making some team members hang back is important to stress because people often do converge all at once to attend to an immediate threat. It can be useful to make some agents remain behind, in case the environment changes.

A. Exploring Under-Matching

Foraging behavior is often compared with the ideal free distribution model [5], [6], which predicts that a team of foragers will allocate themselves to areas in proportion to the relative resources available at each area. Our learning sensors' emergent behavior was consistent with this prediction. However, experiments on foraging behavior on humans [1] and other animals [7], [8] often report systematical under-matching, in which fewer than expected foragers go to the more profitable region and more than expected foragers go to the less profitable region.

One possibility is that organisms have a tendency to sample pools approximately evenly, leading to under-matching, which may be a useful strategy in an environment in which resource outputs may vary with time [9]. Another possibility is that, unlike our learning sensors, which all had the same abilities, organisms differ in their competitive abilities, and competitively fitter organisms may dominate resources at the richer region and scare away others, leading to under-matching [10]. We are experimenting with how these possibilities might affect the behavior of learning sensors.

B. Incorporating Imitation

The hanging back of our learning sensors is analogous to people distancing themselves from the crowds. However, humans also exhibit "bandwagon" behavior, in which they use the appearance of others as evidence that a region might be lucrative [1]. This type of social learning is useful when resources are invisible but other agents are visible.

Imitating others is one of the most powerful methods for quick and effective learning [3] by allowing people to display "no-trial learning" and perform behaviors that they would not have otherwise considered [4]. We assumed in the current simulation that sensors have all the information. However, this is unlikely to be the case in the real environments. Incorporating the capacity for imitation in learning sensors is potentially useful when there are uncertainties in the environment (e.g., created by signal loss).

C. Further Extensions

Our approach using agent-based models is designed to be general. The same principles that we have used in the current simulations can be applied to target chasing [11], [12], in which multiple autonomous sensors attempt to optimally chase multiple moving targets. Simulation results for target chasing are promising. Although we manipulated the speed of movement assuming that the sensors know which direction to move, we can make the sensors learn which direction to move by error-correction and manipulate the learning rate rather than movement speed; the agent closest to the target (or attack region) learns the direction more quickly than the others. Manipulating the learning rate leads to the same patterns of results as manipulating the speed for the simulations of both covering the attack regions and chasing targets.

Another extension of the current work is to use the distance between the sensor and the signal source to determine the speed at which the sensor moves. This method will allow sensors to distribute optimally when they lose their ability to communicate and cannot determine who is closest to the target. Furthermore, the current work can be extended to incorporate memory capacity so that the agents can make predictions about the environment [13] and intercept the targets rather than simply chasing them.

ACKNOWLEDGMENT

This research was supported by the Office of Naval Research, grant #N00014-05-1-00632.

REFERENCES

[1] R. L. Goldstone and B. C. Ashpole, "Human foraging behavior in a virtual environment," *Psychonomic Bulletin & Review*, vol. 11, pp. 508–514, 2004.

[2] M. Resnick, *Turtles, termites and traffic Jams: Exploration in massively parallel microworlds*. Cambridge, MA: MIT Press, 1994.

[3] A. N. Melztoff, "Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli," *Developmental Psychology*, vol. 24, pp. 470–476, 1988.

[4] A. Bandura, "Behavioral modification through modeling procedures," in *Research in behavior modification: New development and implications*, L. Krasner and L. P. Ulmann, Eds. New York: Rinehart and Winston, 1965, pp. 310–340.

[5] S. D. Fretwell and H. L. Lucas, "On territorial behavior and other factors influencing habitat distribution in birds," *Acta Biotheoretica*, vol. 19, pp. 16–36, 1970.

[6] A. K. Seth, "Modeling group foraging: Individual suboptimality, interference, and a kind of matching," *Adaptive Behavior*, vol. 9, pp. 67–91, 2001.

[7] W. Baum and J. R. Kraft, "Group choice: Competition, travel, and the ideal free distribution," *Journal of the Experimental Analysis of Behavior*, vol. 69, pp. 227–245, 1998.

[8] M. Kennedy and R. D. Gray, "Can ecological theory predict the distribution of foraging animals? A critical analysis of experiments on the ideal free distribution," *Oikos*, vol. 68, pp. 158–166, 1993.

[9] K. E. Bell and W. M. Baum, "Group foraging sensitivity to predictable and unpredictable changes in food distribution: Past experience or present circumstances?" *Journal of the Experimental Analysis of Behavior*, vol. 78, pp. 179–194, 2002.

[10] A. I. Houston and J. M. McNamara, "The ideal free distribution when competitive abilities differ: An approach based on statistical mechanics," *Animal Behaviour*, vol. 36, pp. 166–174, 1988.

[11] T. Ishida, "Real-time search for autonomous agents and multiagent systems," *Autonomous Agents and Multi-Agent Systems*, vol. 1, pp. 139–167, 1998.

[12] T. Ishida and R. Korf, "Moving target search: A real-time search for changing goals," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, pp. 609–619, 1995.

[13] A. El-Fakdi, M. Carreras, N. Palomeras, and P. Ridao, "Autonomous underwater vehicle control using reinforcement learning policy search methods," in *IEEE Oceans Europe '05*, Brest (France), 2005.