

Decision-Theoretic Cooperative Sensor Planning *

Diane J. Cook, Piotr Gmytrasiewicz and Lawrence B. Holder
Department of Computer Science Engineering
University of Texas at Arlington
Email: {cook, piotr, holder}@cse.uta.edu

Abstract

This paper describes a decision-theoretic approach to cooperative sensor planning between multiple autonomous vehicles executing a military mission. For this autonomous vehicle application, intelligent cooperative reasoning must be used to select optimal vehicle viewing locations and select optimal camera pan and tilt angles throughout the mission. Decisions are made in such a way as to maximize the value of information gained by the sensors while maintaining vehicle stealth. Because the mission involves multiple vehicles, cooperation can be used to balance the work load and to increase information gain. This paper presents the theoretical foundations of our cooperative sensor planning research and describes the application of these techniques to ARPA's Unmanned Ground Vehicle program.

1 Introduction

Traditionally, research in multi-agent planning and research in image understanding have been pursued independently. The ARPA Unmanned Ground Vehicle (UGV) program is pulling together these two technologies in order to autonomously execute multi-vehicle missions. In this paper we describe our component of ARPA's UGV program which combines decision theory, sensor planning, and multi-agent planning with cooperation to direct and focus sensors on-board a unit of military vehicles.

We develop two capabilities for a unit of autonomous vehicles based on these ideas of cooperative sensor planning. The first capability selects points along a path or in a bounded re-

gion that provide optimal locations for vehicles to observe a specified objective area. The second capability selects optimal pan/tilt angles, or a *field of view*, for each vehicle's camera as it moves in formation with the vehicle's military unit.

For both of these capabilities, we pursue a decision-theoretic approach to the guidance of sensing based on goals of the military mission. These goals include such possibly-conflicting priorities as maximizing the expected value of information obtained, avoiding exposure of one's presence to the enemy, and balancing sensor planning and processing work between cooperating vehicles. Static and dynamic sensor planning can be achieved in a principled manner using the decision-making techniques developed in the area of multiattribute utility and decision theories.

In this paper, we first describe ARPA's Unmanned Ground Vehicle program, list the capabilities of the vehicles and sensors used in our project, and review the principles of multiattribute decision theory. Next, we describe our Observation Point refinement system that selects optimal vehicle positions for observation of a specified area. Section 5 introduces a method of using multi-agent planning techniques with decision theory to dynamically focus sensor attention during a military mission. Finally, we review related work to date and conclude with directions for future work.

2 Sensor Planning for Unmanned Ground Vehicles

The United States armed forces continuously seek to increase soldier effectiveness and surviv-

*Supported by ARPA contract DAAHO4-93-G-0423.



Figure 1: ARPA's Autonomous Vehicle

ability in the face of increasingly lethal battlefields. The need exists to operate in environments that are hazardous because of enemy actions, to increase survivability, to enhance continuous operation, and to expand the radius of reconnaissance / surveillance units. Unmanned semi-autonomous ground vehicles can meet these needs.

The goal of ARPA's Unmanned Ground Vehicle program is to develop and demonstrate field-deployable semi-autonomous ground vehicles incorporating ARPA-sponsored technologies. The capabilities that are currently being demonstrated by these vehicles include autonomous on-road and off-road navigation, obstacle avoidance, path planning, formation control, target detection and recognition, and cooperative sensor planning.

A typical UGV task would be to cooperatively plan a reconnaissance, surveillance, and target acquisition (RSTA) task in which the vehicles move in formation along a specified route. At some point in the mission the vehicles may split up and move to pre-computed spots from which they can maximally view a specified objective area while still maintaining stealth. The vehicles would then reconvene and move in formation to the objective area.

In this problem domain, planning for vehicle movement and planning for sensor movement must be performed in harmony. One reason for

this synchronization is the need to maintain 360 degree security around the platoon unit formation. Maintaining this security requires planning sensor directions to cover the entire area surrounding the vehicles while minimizing duplicated work among separate vehicles.


Another motivation for controlling vehicle and sensor movement together is active RSTA, or coordinated recognition of enemy targets. Executing such a coordinated effort requires planning of both vehicle positions and camera angles for maximum coverage of the target while maintaining formation security. A third reason that plan generation is needed for both vehicle and sensor movement is stealth navigation through unknown terrain, and a fourth reason is that both vehicle and sensor capabilities must be considered to cooperatively compute optimal observation point locations for individual vehicles.

Figure 2 illustrates several unit formations that may be employed throughout a mission. At planning time, each vehicle's sensor space is divided into individual *fields of regard* and weighted to ensure complete 360 degree security around the unit formation while keeping the major focus of attention toward the objective area.


The camera suite on board the vehicles can be controlled in terms of pan, tilt, zoom, and focus. All of these parameters are controlled by our systems as described in the remainder of this paper.

3 Decision-Theoretic Foundations


Our approach to intelligent sensing behavior during a military scouting mission is based on the utility and decision theories. Our basic premise is that the guidance of sensing behavior be based on the rational trade-off between two conflicting priorities (or attributes) of a military scouting mission: maximize the expected value of information obtained, and avoid exposing one's presence to the enemy. One of the basic results in the multiattribute utility theory states that if the attributes are utility-independent then the global utility function is a multiplicative function over the attributes considered [Keeney and Raiffa, 1976]. Thus, for our case of a military scouting mission, the global




Column



Diamond



Line



Wedge

Figure 2: Vehicle Fields of Regard

function that the agent is attempting to maximize, $U(A, S, P)$, of scanning an area A using a sensor S from the position P , can be postulated to have the following form:

$$U(A, S, P) = k_1 U_{Scan}(A, S, P) + k_2 U_{Stealth}(P, S) + k_1 k_2 U_{Scan}(A, S, P) U_{Stealth}(P, S), \quad (1)$$

where $U_{Scan}(A, S, P)$ is the expected value of information obtained during the sensing action, and $U_{Stealth}(P, S)$ is the expected utility of maintaining stealth, i.e., remaining hidden from the enemy while occupying the position P and using the sensor S . Let us note that this value includes not only the danger of being discovered due to occupying an exposed location P , but also the danger of using the sensor S that could itself be detected by the enemy. The values of the constants k_1 and k_2 determine the relative weight with which the desirable attributes of gaining more information and remaining undetected can be traded off against one another. These values can be made to reflect the parameters of a particular military scouting mission at hand.

The following two subsections briefly describe our proposed approach to calculating the values of U_{Scan} and $U_{Stealth}$ during a military scouting mission.

3.1 Determining the Utility of a Sensing Action

The value of information, in general, is equal to the difference between the expected value of action when the system has the information and the expected value of action without the information. The resulting desirable behavior will be an agent's attempt to look for objects of greatest concern and relevance to the overall military objective. The calculation of the expected value of a sensing action has to include the likelihood that the interesting object is located within the area scanned, and that the sensor can successfully recognize the object at that location. Following Feldmann and Sproull [Feldman and Sproull, 1977], we propose that the general expression for

the value of scanning the area A , using sensor S , from the position P , be:

$$U_{Scan}(A, S, P) = \int_A \sum_k P1(x, y) P2_k(x, y) V I_k dx dy, \quad (2)$$

where $P1(x, y)$ is the probability that an object located at (x, y) will be correctly identified from the position P with the sensor S , $P2_k(x, y)$ is the prior probability that an object of type k is located at (x, y) , and the $V I_k$ is the value of information about the object of type k .

Intuitively, the probability $P1(x, y)$ contains information about the sensing ability of sensor S given the current conditions (day, night, fog, smoke, etc.), and the distance involved, i.e., how well and how far the sensor can currently "see". The probability $P2_k(x, y)$ contains the prior information as to where various kinds of objects are likely to be located. This probability can be initially derived from the prior Intelligence Preparation of the Battlefield (IPB), commonly used to assign a degree of interest to regions of the terrain under consideration. This probability is also continually updated as the sensing actions are performed, so that it correctly represents the up-to-date state of the system's knowledge about the environment. The updating function on which we will concentrate is Bayes' rule; however, other methods, including the Kalman filter technique, will also be investigated.

The value of information $V I_k(t)$ about the object of type k , reflects the importance of knowing the location of an object of type k . The kinds of objects the system may be interested in include the various kinds of enemy forces in the area (tanks, HMMWVs, trucks, etc.), the locations of friendly forces and other scouting vehicles, and the locations of navigational obstacles and other relevant elements of the terrain.

As an example, let us examine in more detail the value of information about the locations of obstacles. The value of this information represents the effectiveness of the vehicle moving through a terrain when the locations of all of the obstacles are known a priori vs. the vehicle




Figure 3: Vehicle paths chosen with (solid line) and without (dashed line) prior knowledge of obstacle locations

having to deal with the obstacles as they are encountered (see Figure 3). In the first case (the solid line in the figure), the vehicle can plan the best path beforehand. The length of this path is likely to be greater than the straight-line distance, with the ratio of the two reflecting the extent to which the environment is cluttered with obstacles. Let us call the average value of this ratio the obstruction coefficient, OC . If, on the other hand, the vehicle is moving through unknown territory, it has to circumvent the obstacles as they are encountered, and the length of the generated path will likely be greater than in the case above (the dashed line in the figure). Let us call the average ratio of such a path to the length of the straight-line distance a maze coefficient, MC . The obstruction and maze coefficients, therefore, statistically summarize the effectiveness of the vehicle movement through known and unknown territory, respectively. The value of information about the obstacles is thus proportional to the difference $MC - OC$. Using these statistical factors, the vehicle can maintain its current estimate of the value of information about the obstacles.

Typically, the value of information about other kinds of objects can be input to the system directly by the operators. The example detailed in this section describes the value of information with respect to a path-planning goal. In the UGV mission, information value also relates to the goal of locating and identifying enemy targets. For the UGV project, specifications could serve to guide the particular scouting mission, detailing what information is of interest and to what extent.

3.2 Utility of Maintaining Stealth

We can postulate that the utility of maintaining the stealth of the scouting vehicle is the negative of the expected cost, $EC_{Disc}(P, S)$, of being discovered by the enemy:

$$U_{Stealth}(P, S) = -EC_{Disc}(P, S). \quad (3)$$

The expected cost of being discovered is the probability $p_{Disc}(P, S)$ of being discovered while

at position P and using sensor S , multiplied by the cost itself:

$$EC_{Disc}(P, S) = p_{Disc}(P, S)C_{Disc}(P). \quad (4)$$

The probability of the vehicle being discovered, $p_{Disc}(P, S)$, depends on the location of the enemy forces relative to the location P and their line of sight, and on the detectability of the sensor used. Thus, the same location can be safe or dangerous, depending on whether the enemy’s expected position allows for a clear line of sight of this location. Also, some sensors can be safer due to the fact that they are more difficult to detect by the enemy. Thus, the probability $p_{Disc}(P, S)$ can be computed from prior information about the possible locations of the enemy forces, and from an estimate of detectability of the sensor used.

The cost of being discovered by the enemy at the given location is also a combination of several factors. There is the immediate danger that the given vehicle will become a target of enemy forces, as well as the danger associated with how the information about the location of the vehicle could be used by the enemy to harm and obstruct the mission of other friendly forces.

The exact algorithms for computing $p_{Disc}(P)$ and $C_{Disc}(P)$ remain among the active research areas within the UGV project. The obvious question is the appropriate level of detail that should be included in these calculations. Our strategy in answering this question will begin by examining the elements of the very rigorous and detailed view, as outlined above, and proceed by successive approximations to find the sensitivity of the resulting performance to the details neglected, given the real-time control requirements of the modern battlefield.

4 Decision-Guided Observation Point Refinement

The observation-point refinement algorithm (OP) is used to select optimal observation points from which vehicles can observe a specified area of interest. OP is provided with polygonal descriptions of the area to be observed (*area of in-*

terest) and the selected regions where the vehicles will perform observation. OP will return a sorted list of observation-point sets (a set contains a single observation point for each vehicle), where each set is rated by its visibility/stealth measure.

The OP framework applies decision-theoretic techniques to guide observation point refinement. Here, we are applying the basic idea to select an observation point so as to optimize the expected value of the information obtained using the sensors from that location, while maintaining stealth. When selecting an observation point, therefore, an agent should trade off the expected value of information acquired on one hand, and the dangers of revealing the agent’s presence, on the other hand. In accordance with the multiattribute utility formulation discussed in the previous section, we use Equation 2 to rate the suitability of candidate vehicle locations, given an area of interest to be scanned, and given the suit of sensors on board the vehicle.

In our implementation we assume that a sensor can successfully recognize an object if there exists a clear line of sight between the sensor and object and if the object is within the sensor’s range. The likelihood of recognizing the object decreases with the increase in obstructions (such as vegetation) between the sensor and the object. The likelihood of recognizing an object increases as the number of vehicles that can view the object increases. Prior information as to where various kinds of objects are likely to be located can be initially derived from the prior Intelligence Preparation of the Battlefield (IPB).

The stealth of a vehicle is inversely proportional to the ability of an enemy, assumed to be located inside the area of interest, to sense and identify the vehicle at the observation point. While visibility of a potential target inside the area of interest is measured from the top of the vehicle (where the sensor is located), the visibility of a vehicle is measured with respect to the center of the vehicle. Figure 4 shows sample observation points for two vehicles looking at a single point inside an area of interest. Although the bottom vehicle can see the point, the vehicle itself is also clearly seen, yielding a high visibil-




Figure 4: Optimal Observation Point

ity value and a low stealth value. On the other hand, the top vehicle positions itself behind the crest of a hill, so it cannot be seen well from the area of interest point. The point can be viewed from the vehicle, though a tree along the path impedes visibility.

Three types of operations are executed by OP. All three operations rely on the formulae described earlier.

Selection of observation points within an area. For this operation, observation regions are selected by the user for each vehicle. OP then collects sample points within each region, and computes visibility/stealth measures for each *combination* of observation points. The sorted list of sets is returned to the user to be integrated into the mission plan.

Selection of observation points along a path. Although an explicit observation task is often integrated into the mission plan, the operator may desire for vehicles to scan an area of interest throughout the entire plan, from a variety of vantage points along the specified mission route. As a result, OP can be used to sample points along a specified path and select a set of these points that maximizes information gain and information value. The entire formation would stop at these locations to scan the area of interest.

Division of objective area among vehicles. The current efficiency of target recognition algorithms prevents a quick perusal of the objective area by unmanned ground vehicles. Therefore, the work must be divided as efficiently as possi-

ble among the vehicles involved in the mission. Given a polygonal description of the area of interest, OP divides the polygon into separate contiguous regions, one per vehicle. These regions are selected to maximize visibility of the region by the vehicles while also trying to balance the work load between vehicles.

To date, the objective area is split into separate regions along lines parallel to the x or y axis. In the future, piecewise linear regression algorithms [Weiss and Kulikowski, 1990] can be used to yield a more effective split.

All of these operations can be computationally expensive because of the large number of rays that must be traced. The procedure can be made more efficient by reducing the number of sample points considered within the observation regions and within the area of interest. The number of sample points is reduced by either utilizing a coarse-grain sampling technique or by selecting sample points with the highest prior probability. Once a few potential observation points are selected, the selected points can be refined recursively using a more detailed sampling. Parallel processing techniques can also be used to improve efficiency by calculating visibility/stealth measures for individual points in parallel.

5 Dynamic Zone Security

The cooperative RSTA planning system is used to cooperatively maintain 360-degree camera security around a moving unit formation and to cooperatively search for targets. Both capabilities use a decision-theoretic approach to select the current camera field of view throughout mission execution. In this section we describe the field of view selection method and cooperative reasoning methods that are central to cooperative sensor planning.

5.1 Decision-Based Field Of View Selection

Zone security is accomplished by defining multiple fields of regard for each vehicle according to the current formation (line, wedge, diamond, column) and the vehicle's position in the formation. Weights assigned to each field of regard

allow the vehicle to spend more time looking in higher-interest areas.

Each vehicle's field of regard is divided into individual *fields of view* (*FOV*). With each field of view is associated a location (in terms of pan/tilt angles or a world coordinate focus location), a weight indicating the priority of the FOV, and the desired angular width of the FOV. Each time the camera is ready to move to a new location, the weights of the fields of view are updated and the current FOV is selected.

The weights of each individual field of view are adjusted dynamically. As we mentioned, the value of a particular field of view area A viewed with sensor S from position P can be calculated as

$$U_{Scan}(A, S, P) = \int_A \sum_k P1(x, y) P2_k(x, y) VI_k dx dy, \quad (5)$$

where $P1(x, y)$ is the probability that an object located at (x, y) will be correctly identified from the position P with the sensor S , $P2_k(x, y)$ is the prior probability that an object of type k is located at (x, y) , and VI_k is the value of information about the object of type k .

There are several factors that affect $P1$ and $P2$, including security, continuity, focus of attention, and terrain reasoning. A priority of a multi-vehicle reconnaissance mission is to maintain 360-degree security around the unit at all times. The computational demands of current target detection and recognition algorithms prevent rapid processing of each image frame. Because each vehicle cannot quickly scan the entire area around the unit, vehicles must make use of teamwork to effectively divide the work. Each vehicle is assigned the portion of the unit security that can best be handled from that position in the unit formation.

The second factor is continuity. Continuity ensures smooth motion of the camera pan/tilt, because fields of view that neighbor the current FOV are more heavily weighted than other fields of view. Continuity is important when the camera pan/tilt motion is slow, when information shared between neighboring fields of view

is useful for the target detection/recognition algorithms, and when pan/tilt limits prevent easy movement between non-neighboring views.

The third factor affecting FOV weights is focus of attention. Prior probabilities can be established for target locations based on database information, and these probabilities can be updated as vehicles detect potential targets throughout the mission. Any region identified as a high-probability region will receive increased priority in the FOV-selection algorithm.

Finally, terrain reasoning can be used to dynamically update FOV weights. Military vehicles often make use of terrain for maximal stealth as well as ease of navigation. These terrain features will thus affect the probability that a target is located in a given region. For example, target vehicles are more likely to be found along tree lines than in an open field. In addition, vehicles are less likely to be found on a steep slope, where the vehicle would be unstable. Figure 5 demonstrates this type of terrain reasoning.

On the other hand, terrain obstructions can limit the amount of information obtainable from a FOV, and thus the corresponding weight should be decremented. Figure 5 demonstrates a case in which a tree obstructs the view from the vehicle, and the weight of the corresponding FOV is decremented accordingly. Because the terrain surrounding the unit changes as the unit moves along the specified route, the terrain-based weights must be updated dynamically throughout the mission. Each of these four factors can contribute much or little to the overall values of $P1$ and $P2$, depending on how heavily the user weights each parameter.

When a new field of view is requested, a *biased roulette wheel* is spun. Roulette wheel methods have proven to be effective in a variety of adaptive algorithm applications [Goldberg, 1989]. Using this selection method, potential fields of view are assigned a portion of the wheel corresponding to their fraction of the total possible weight. The probability of selecting a given FOV is proportional to the FOV's share of the roulette wheel.

Once a field of view is selected, the width of the field of view is dynamically computed to fit




Figure 5: FOV adjustment based on terrain

the corresponding terrain. A sampling of rays is collected that centers around the FOV direction. Each ray is traced to the point where the ray intersects the ground, and average distance to the ground is calculated for the FOV. If the average ray length is greater than the sensor range, the camera elevation is adjusted slightly or the view is re-selected, based on current terrain features. The width of the FOV is then computed in such a way that the amount of pixel information per FOV is constant, as given by the formula

$$\text{ViewWidth} = \tan^{-1}\left(\frac{\text{NumDesiredPixels}}{\text{AverageRayLength}}\right).$$

Figure 5 demonstrates how the hill forces the ray lengths for the top FOV to be short and thus the width will be greater than for the other fields of view.

5.2 Multi-Vehicle Cooperative Reasoning

For a military RSTA mission, multiple vehicles must be utilized. The use of multiple vehicles increases the chance of a successful mission because of the increased robustness, increased security, and increased number of observation points for scouting an area.

Multi-agent planning and coordination is a focus of much attention in AI reasoning [Briggs and Cook, 1995; Durfee and Montgomery, 1991; Ephrati *et al.*, 1995; Moses and Tennenholtz, 1990; Moses and Tennenholtz, 1991; Rosenschein

and Zlotkin, 1994; von Martial, 1992]. As automation of intelligent tasks increases, the need arises for heterogeneous agents to work in a common environment. While the need for multi-agent planning algorithms is apparent, the development of algorithms which meet each agent's goals in a timely fashion, avoid deadlock, and do not incur heavy communication costs provides a challenging task.

Cook [Cook, 1994] describes three types of multi-agent control schemes: central control, distributed control, and local control (no communication). Central control is shown to be effective when communication is reliable, and local control is effective if no communication is needed; otherwise, distributed control is necessary.

Our cooperative sensor planning system has elements of all three control schemes. The initial partitioning of the fields of regard and the weighting of each of the four selection factors is controlled by a central force: the mission leader. The selection of each FOV and the dynamic updating of weights and view widths is performed at a local level. Distributed coordination schemes, while often the most robust, are also very complex. This section describes the distributed cooperation that is necessary to our cooperative sensor planning system.

There are several tasks that required distributed decision-making and coordination, including:

- target confirmation,
- security handoff, and
- health checks.

Target confirmation. For every target that vehicle A detects, vehicle A asks for confirmation from all other vehicles. All other vehicles that are available to help and are in line of sight of the potential target interrupt their work to focus on the target region. All target detection information is passed to the requesting vehicle.

Security handoff. If a detected target is stationary, target detection and recognition can be down quickly enough to maintain the integrity

of the unit security camera movements. If a detected target is moving, all other camera work is abandoned while the vehicle tracks the detected target. Unit security is a cooperative task and the responsibility is shared by all vehicles. If vehicle A needs to track a target, vehicle A hands off its security work to another available vehicle (vehicle B). After tracking is complete, vehicle A finds the owner of the shared security fields of regard (vehicle B may have handed the work over to yet another vehicle) and resumes its original security work.

Sharing security work is fairly straightforward. A description of the fields of regard and their weights must be communicated from vehicle A to vehicle B. Vehicle B adds the corresponding fields of view to the existing list, in effect dividing the roulette wheel into a greater number of pieces. When the security work is returned, the shared fields of view are removed from vehicle B's list.

Health checks. Although it is not desirable for a vehicle to malfunction or be destroyed during a mission, the success of the mission should not depend entirely on the health of any one vehicle. To ensure that the goals of the mission are met, the leader of the unit periodically performs health checks on the other vehicles. If a vehicle does not respond in a timely manner, the sensor plan is reconfigured for one less vehicle and new work is partitioned among the unit. In this way, no work is lost because of a missing vehicle. As long as one vehicle is remaining, the mission can be accomplished. If the vehicle comes back to life, the plan can be reconfigured to include the revived vehicle.

6 Conclusions

In this paper we present a cooperative sensor planning system that unifies research from active vision and sensor planning, decision theory, and multi-agent planning and communication.

Work on active vision has been well established in the field. Bajcsy [Bajcsy, 1988] introduced the idea of active perception as applied to controlling a sensor at any level of abstraction, from controlling the focus of the physical device (as we are doing) to controlling

the semantic interpretation of information returned from the sensor. Active vision research has also made use of decision-theoretic techniques to measure the utility of gathering information [Aloimonos *et al.*, 1988; Bajcsy, 1988; Krotkov, 1989], but has not been applied to this type of military application where both static and dynamic decision making is necessary to achieve the overall mission objectives. Literature in the cognitive science community shows that this decision-theoretic approach to sensor planning is also demonstrated in human perception and information-gathering [Boff *et al.*, 1986; Shaw and Shaw, 1977].

Although active vision techniques have been used to make centralized decisions about sensor movements, there is very little work that allows decentralized sensor planning decisions. Multi-agent planning and negotiation techniques are common in the AI literature, but have not been integrated into computer vision work. This paper describes a project that uses ideas from multi-agent planning and from sensor planning to allow centralized allocation of major tasks, local control of individual sensor direction, and distributed cooperation between intelligent perceptual agents.

The work described in this paper has been implemented in the context of the ARPA Unmanned Ground Vehicle program and successfully used on-board a unit consisting of two HMMWVs. Although the work to date has demonstrated effective cooperative decision-guided sensor planning, there are a number of avenues we plan to pursue in the future.

Ballard and Brown [Ballard and Brown, 1992] and Bajcsy and Campos [Bajcsy and Campos, 1992] both emphasize that learning is an important part of active perception. To date, the probabilities associated with each aspect of terrain reasoning, security, and stealth have been hard-coded. Future extensions of this project will learn the value of information and probabilities of each aspect of the mission from experience. One method of learning probability values is through the use of adaptive probabilistic networks, a subset of belief nets that can learn individual probability values and dis-

tributions using gradient descent [Pearl, 1988; Russell and Norvig, 1995; Spiegelhalter *et al.*, 1993].

Another important extension of this project is to add temporal reasoning to the sensor planning algorithm. In general, the probability of a given object existing at a specified location should increase if the object has been captured with a sensor, but should decrease as time passes since the object was last perceived. Thus, if a vehicle detects a target at one point in the mission, the target is not guaranteed to remain at the same location through the rest of the mission, and the associated probabilities should be continuously adjusted.

In addition, there are a number of ways in which the effectiveness of sensor planning could be improved with additional communication. For example, at each step in the plan one agent could communicate all the information it has gathered to other agents in the unit, thus improving every vehicle's utility database. However, communication is very costly and communicating large amounts of information between spatially-separated entities will decrease the vehicles' stealth as well as slow down mission execution. A future priority of this project is to investigate methods of optimizing the tradeoff between minimizing communication and maximizing cooperation.

Finally, from our existing work in this research area we have found that calculating sensor utility factors for several vehicles over a large area is computationally very expensive. As more and more factors are included in the decision process, this computational burden will increase. However, considering each set of sensor parameters in great detail is not always profitable — some decisions will quickly show themselves to be wasteful or inappropriate for a given task. Directing and limiting the search and calculations involved in cooperative sensor planning is a future research direction for this project.

References

- [Aloimonos *et al.*, 1988] J Aloimonos, I Weiss, and A Bandyopadhyay. Active vision. *International Journal of Computer Vision*, 1(4):333–356, 1988.
- [Bajcsy and Campos, 1992] Ruzena Bajcsy and Mario Campos. Active and exploratory perception. *CVGIP: Image Understanding*, 56(1):31–40, 1992.
- [Bajcsy, 1988] Ruzena Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):996–1005, 1988.
- [Ballard and Brown, 1992] Dana H Ballard and Christopher M Brown. Principles of animate vision. *CVGIP: Image Understanding*, 56(1):3–21, 1992.
- [Boff *et al.*, 1986] Kenneth R Boff, Lloyd Kaufman, and James P Thomas. *Handbook of Perception and Human Performance, Volume 1*. John Wiley and Sons, 1986.
- [Briggs and Cook, 1995] Will Briggs and Diane J. Cook. Flexible social laws. In *Proceedings of the International Joint Conference on Artificial Intelligence (to appear)*, 1995.
- [Cook, 1994] Diane Cook. Reconfiguration of multi-agent planning systems. In *Proceedings of Artificial Intelligence Planning Systems*, pages 225–230. AIPS, 1994.
- [Durfee and Montgomery, 1991] Edmund H. Durfee and Thomas A. Montgomery. Coordination as distributed search in a hierarchical behavior space. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1363–1378, Nov/Dec 1991.
- [Ephrati *et al.*, 1995] Eithan Ephrati, Martha E Pollack, and Sigalit Ur. Deriving multi-agent coordination through filtering strategies. In *Proceedings of the International Joint Conference on Artificial Intelligence (to appear)*, 1995.

- [Feldman and Sproull, 1977] Jerome A Feldman and Robert F Sproull. Decision theory and artificial intelligence ii: the hungry monkey. *Cognitive Science*, 1(2):158–192, April 1977.
- [Goldberg, 1989] David E Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison Wesley, 1989.
- [Keeney and Raiffa, 1976] R L Keeney and H Raiffa. *Decisions with multiple objectives: preferences and value tradeoffs*. John Wiley and Sons, 1976.
- [Krotkov, 1989] E Krotkov. *Active Computer Vision by Cooperative Focus and Stereo*. Springer-Verlay, 1989.
- [Moses and Tennenholtz, 1990] Yoram Moses and M. Tennenholtz. Artificial social systems part i: basic principles. Technical Report CS90-12, Weizmann Institute, 1990.
- [Moses and Tennenholtz, 1991] Yoram Moses and M. Tennenholtz. On formal aspects of artificial social systems. Technical Report CS91-01, Weizmann Institute, 1991.
- [Pearl, 1988] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [Rosenschein and Zlotkin, 1994] Jeffrey S. Rosenschein and Gilad Zlotkin. Designing conventions for automated negotiation. *AI Magazine*, 15(3):29–46, Fall 1994.
- [Russell and Norvig, 1995] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
- [Shaw and Shaw, 1977] Marilyn L Shaw and Peter Shaw. Optimal allocation of cognitive resources to spatial locations. *Journal of Experimental Psychology: Human Perception and Performance*, 3(2):201–211, 1977.
- [Spiegelhalter *et al.*, 1993] D Spiegelhalter, P Dawid, S Lauritzen, and R Cowell. Bayesian analysis in expert systems. *Statistical Science*, 8:219–282, 1993.
- [von Martial, 1992] F. von Martial. *Coordinating Plans of Autonomous Agents*. Springer-Verlag, 1992.
- [Weiss and Kulikowski, 1990] Sholom M Weiss and Casimir A Kulikowski. *Computer Systems That Learn*. Morgan Kaufmann, 1990.