

Comparing methods for UAV-based autonomous surveillance

Michael Freed, Robert Harris, Michael Shafto

NASA Ames Research Center

Abstract: *We describe an approach to evaluating algorithmic and human performance in directing UAV-based surveillance. Its key elements are a decision-theoretic framework for measuring the utility of a surveillance schedule and an evaluation testbed consisting of 243 scenarios covering a well-defined space of possible missions. We apply this approach to two example UAV-based surveillance methods, a TSP-based algorithm and a human-directed approach, then compare them to identify general strengths and weaknesses of each method.*

UAV-based Surveillance

Aerial reconnaissance, surveillance, and other observation tasks have been primary aircraft applications since the early days of powered flight. They remain key activities in domains from military and security operations to land management and scientific research. However, airborne observation is typically a deadly dull process that strains the vigilance and morale of human pilots and makes poor use of their costly, hard-won skills. Thus, following the rule of “dull, dirty or dangerous,” it is considered an excellent application for autonomous vehicles. Unmanned aerial vehicles (UAVs) have been employed in this capacity for decades, though almost exclusively for reconnaissance (DoD 2002). Technological improvements combined with increasing investment and interest in UAVs promise to increase their capabilities and availability, thus enabling more diverse and demanding missions. Of particular interest to several operational communities are missions using UAVs to maintain “situation awareness” by continuous or periodic surveillance.

Autonomous surveillance of spatially separated sites raises issues beyond those related to reconnaissance at a single site. In particular, since a given UAV can only be one place at a time, it must be treated as a limited resource that needs to be allocated as effectively as possible. Effectiveness, in this case, means providing the best possible information to the user at the best possible time – i.e. maximizing the value of returned information. For any surveillance agent, airborne or otherwise, this entails a variety of interlinked choices about which sites to visit over the course of a mission, how often to visit each site, what paths to take, how long to spend observing, and what kind of measurements to take (cf. Sacks (2003) for a related discussion on police patrol, Carbonell (1969) regarding human visual scanning of instruments and Koopman (1956) regarding submarine-based search).

Factors specific to aerial vehicles affect what kind of algorithms can most effectively make these decisions. For instance, Massios et al. (2001) have studied the problem of optimizing surveillance for autonomous ground

vehicles (UGVs) operating inside buildings. Given their assumption that every space in the building is worth observing, the problem of deciding where to go next is highly constrained by the structure of the building. The problem of how to get to a location not immediately adjacent requires path-planning. With UAVs, sites of interest may all be accessible by a direct path, reducing the need for path-planning but offering weaker constraints on where to go next. A second factor, wind, usually has little effect on UGVs, but has a large effect on UAVs, increasing or reducing required traverse time between almost any two sites. Algorithms for UAV-based surveillance should thus treat wind as a critical parameter and, ideally, should enable execution-time adaptation to changes in wind speed or direction.

Differences in kinematics and vantage together create a third significant difference between UGV- and UAV-based surveillance. Because of its altitude, a UAV will frequently be able to observe a site from a distance without obstruction and thus may not have to travel the full distance to that site. And, due to the low friction on an air vehicle in aerodynamic flight, a UAV making “snapshot” observations may be able to retain most of its speed when transitioning between approach to one site and approach to the next. A surveillance algorithm that takes advantage of these aviation-specific factors should perform significantly better than one that does not.

Our work on UAV-based surveillance represents one part of a larger project to develop a practical and flexible UAV observation and data-delivery platform. The Autonomous Rotorcraft Project (Whalley et al. 2003) is an Army/NASA collaborative effort combining advanced work on avionics, telemetry, sensing, and flight control software in addition to software for high-level autonomous control. The base platform selected for the project, a Yamaha RMAX helicopter, has been enhanced in a variety of ways that increase its potential effectiveness as a surveillance vehicle. Flight control software allowing it to fly aerodynamically extends the vehicle’s speed and improves its fuel-efficiency, thus extending both operating range and base flight duration (60 minutes hovering with full payload). The vehicle includes a range of sensors and the capacity to integrate and control additional sensors as demanded by particular missions. Its high-level autonomy component, Apex (Freed 1998), incorporates reactive planning and scheduling capabilities needed for mission-level task execution, navigation, response to health/safety contingencies and interaction with human users. To enable the system to become highly effective for surveillance, scheduling capabilities must be extended based on

algorithms of demonstrated effectiveness in diverse mission scenarios relevant to the Army and to NASA.

The diversity of possible surveillance missions poses particular challenges. First, an algorithm that performs well in certain kinds of missions may perform poorly in others. For instance, an algorithm that does well optimizing observations for a small number of sites may not scale well to missions involving a large number of sites. Similarly, an algorithm that assumes that information obtained at different sites becomes obsolete at equal rates or that the value of making an observation at one site necessarily equals that at another will not perform well when such assumptions do not hold. It is not yet well-understood which attributes are most significant in distinguishing one mission from another. While the number of sites to be observed is clearly an important factor, the importance of other factors, e.g. the centrality of the takeoff/land location with respect to the set of target sites, is less clear. Finally, for a single system to provide autonomous surveillance capability for a broad range of missions requires an underlying theory of surveillance. If users need to communicate mission goals in terms of that theory, its generality is likely to pose difficulties for most users (Freed et al. 2004). For instance, a theoretical foundation based on mathematics unfamiliar to most users (as will be described below) may require them to specify the mission in terms of seemingly exotic mathematical parameters.

These challenges lay out three areas of work: (1) developing methods for measuring the effectiveness of a given algorithm and for comparing the performance of an algorithm to that of human operators (i.e. to current practice); (2) creating planning and scheduling algorithms that perform surveillance effectively in significant parts of the space of possible missions; and (3) addressing issues of usability in the specification of missions by non-expert users. In this paper, we describe our work in the first of these areas to create a framework for evaluating algorithm performance and human performance at surveillance tasks. We then describe the application of the framework to two cases useful as benchmark surveillance approaches – a modified Traveling Salesman Problem (TSP) algorithm and human-directed surveillance.

Measuring Surveillance Performance

The first issue in devising an evaluation framework is to define what it means to do a good job at surveillance. Intuitively, the purpose of surveillance is to return information on a set of targets to some user or set of users. Performance at the surveillance task will depend on the information's quantity, accuracy, importance and timeliness. As will be discussed, there are many variations on the general problem. To accommodate the diversity of surveillance missions, we start with a very general, decision-theoretic formulation of the overall goal:

to maximize the utility of returned information over a defined interval.

Like Massios et al. (2001), we characterize information value in the negative – i.e. in terms of the cost of not having observed a target for a given interval rather than the benefit of having observed the target at a given time. Consider the example of maintaining surveillance over a set of buildings, any of which might catch fire at any time. Observing the building allows us to call the fire department if necessary, and thus limit the amount of damage. The longer we go without observing, the more likely it is that a fire will have occurred (though the probability may still be very small) and the more damage any such fire is likely to have inflicted. Thus, the expected cost of not observing the building (and thus remaining ignorant of its state) for a given interval depends on the fire's probability and expected cost of occurrence. Specifically, the expected cost of ignorance (ECI) for having not observed a target τ during the interval t_1 to t_2 is:

$$ECI_{\tau}(t_1, t_2) = \int_{t=t_1}^{t_2} p(t) \cdot cost(t_2 - t) dt$$

where $p(t)$ is probability density function for the occurrence of some cost-imposing event E (e.g. a fire breaking out) and $cost(d)$ is a function describing the expected cost imposed by E as a function of the time from occurrence to intervention. In other words, the cost of ignorance is the sum, for all points in the interval, of the probability of the event occurring at that point¹ times the expected cost if it occurs at that point. If more than one kind of event can occur at a target, and the event-types are uncorrelated, the expected cost of ignorance is simply the sum of the ECI values for each.

Over the course of a surveillance mission, an interval running from t_{start} to t_{end} , expected cost accumulates at each target². If the target is never observed during that period, the total mission ECI for that target is determined by the above equation for ECI_{τ} with $t_1=t_{start}$ and $t_2=t_{end}$. Otherwise, observations divide the target's mission timeline into a sequence of intervals I_{τ} where the target's total mission ECI equals the sum of ECIs for each interval.

$$ECI_{\tau-mission} = ECI_{\tau}(t_{start}, t_{end}) = \sum_i^{I_{\tau}} ECI(i_{start}, i_{end})$$

¹ Here we assume expected detectability latency $l_0 = 0$ and refer to the time of occurrence of an event rather than the time it becomes detectable. Values of $l_0 > 0$ can be accommodated by integrating from $\max(0, t_1 - l_0)$ rather than from t_1 .

² Time t_{start} represents a reference start time at which costs begin accruing.

The effect of observation occurring at t_2 is to reduce the maximum expected cost of an event occurring at $t < t_2$ from $\text{cost}(t_{\text{end}} - t)$ to $\text{cost}(t_2 - t)$. This reduces the total mission ECI and also constrains its maximum. For example, $\text{cost}(t)$ may asymptote at \$5M, corresponding in our example to the building burning to the ground. If, e.g., the building is observed every 30 minutes and $\text{cost}(30 \text{ minutes})$ is \$1M, the ECI over the course of the mission for that target cannot exceed \$1M.

With this way of determining the mission ECI for a target, the total mission ECI can be defined simply as the sum of mission ECIs for all surveillance targets. The performance of a surveillance algorithm in a given mission is thus measured by its success in minimizing this total expected cost. We define ECI_{max} as the total mission ECI if no targets are observed during the course of a given mission and $\text{ECI}_{\text{method}}$ as the total mission ECI resulting from an observation schedule generated by a particular method. Thus:

$$\text{value}_{\text{method}} = \text{ECI}_{\text{max}} - \text{ECI}_{\text{method}}$$

Modeling a Mission

The choice of what probability function and what cost function to use to model ignorance cost at a given target depends on the kind of cost-imposing event(s) that may occur there. Some events are once-only, meaning that we assume they can occur at most once during the course of a mission (e.g. theft of an item). Others can re-occur serially (e.g. a security gate left open which can be closed and then left open again) or in parallel (e.g. an individual entering an area illegally). Event probability may vary with some regular event (e.g. rush hour, nighttime), contingent upon some other event (e.g. rain) or may remain constant. For the work described here, we have assumed that all events are once-only and that occurrence probability (hazard function) is constant assuming no prior occurrence. Thus, the exponential function $1 - e^{-at}$ describes the probability that event E has occurred by time t (assuming the start of the mission $t_{\text{start}} = 0$); its derivative yields the probability density function $p(t) = ae^{-at}$.

The cost function combines a number of factors. Most important is how the physical process initiated by an event unfolds and how cost accrues as a result. For instance, a building fire may start out slowly, at some point begin increasing rapidly in intensity, then eventually taper off as flammable material runs out and the cost of the fire approaches the total value of the building. This suggests an s-shaped cost function such as a sigmoid. Other factors include the initial cost c_0 of the event (e.g. from an explosion that causes a fire), the maximum cost m that may accrue from an event (e.g. the cost of the building plus fire cleanup costs), the expected intervention latency l_1 (e.g. how much time it takes firefighters to get to the site and

put out the fire) and expected reporting latency l_2 (e.g. how long it takes to get in range to alert firefighters). The work described here assumes that all events are modeled using a sigmoid normalized to intercept the y-axis (cost) at c_0 and to asymptote at m .

$$\text{cost}(d) = c_0 + \left(\frac{2}{1 + e^{-k(d+l_1+l_2)}} - 1 \right) (m - c_0)$$

Multiplying the probability and cost functions with initial-cost and latency factors factored out for simplicity ($c_0 = l_1 = l_2 = 0$), we get the ECI equation below for evaluating the expected cost of not observing a specified target (associated with parameters a , k and m) during the interval t_1 to t_2 (each a displacement from the mission start time $t_0=0$). We omit discussion of the closed-form solution for the integral and of how best to compute ECI values.

$$\text{ECI}_\tau(t_1, t_2, a, k, m) = \int_{t=t_1}^{t_2} ae^{-at} m \left(\frac{2}{1 + e^{-k(t_2-t)}} - 1 \right) dt$$

From this framework, a clear process emerges for how a user can specify mission parameters, apply a surveillance decision method and then evaluate the output of that method with respect to the mission. The first step is to specify the mission. This involves defining a start/end location, mission duration, surveillance vehicle (with range, kinematics, sensors and other characteristics) and set of target locations. Each target is associated with one or more events, and each event with parameterized probability density and cost functions. Given our previously described assumptions about these functions, users would specify three parameters for each event: a , k and m . The value m is simply the maximum (asymptotic) cost of the event. To determine the probability rate parameter a , a user should specify some reference probability interval for the event. For instance, the user may specify that the probability of the event is .2 during a 60 minute interval. Solving for a yields the value .00372. To determine the cost rate parameter k , the user should specify a reference cost interval such as \$1M during the first 30 minutes following occurrence. Solving for k yields the value .0135.

Second, after specifying all elements of the mission, this information is made available to the agent (algorithm or person) responsible for generating a surveillance schedule. The agent's output may take the form of a repeatable sequence that must be translated into a schedule. For example, the sequence ABCAB denotes that targets A, B and C will be visited repeatedly and in order, skipping C on alternate circuits and breaking off just in time to return to the start location before the mission end time. A schedule specifying at what times each target is observed over the course of the mission can be generated

by simulation based on vehicle characteristics, weather and map information. The resulting schedule is then used to compute $value_{method}$ as described in the previous section, providing a measurement of the expected benefit of performing surveillance using a given method.

Comparative Evaluation Testbed

In the previous section, we addressed the question of how to measure the performance of a surveillance method in a given mission. The next step is to make it possible to compare different methods so as to learn their relative strengths and weaknesses. Such comparisons serve two important practical purposes. First, the process of developing and refining surveillance algorithms depends on knowing what weaknesses should be addressed and on being able to measure the effect of intended improvements. Second, this kind of analysis might allow a system to automatically select the best method for a newly defined mission by matching to the most appropriate method.

Comparative analysis requires testing surveillance methods against a set of significantly different mission types. This raises the question of what features are likely to differentially affect the performance of different methods. A set of such features would provide a basis for classifying missions into different types and thus for creating a stable testbed mission set. Unfortunately, it is not altogether clear which are important. It is not clear, for example, what features should be considered at all, what tradeoffs exist in the design of algorithms that are likely to impact sensitivity to a given feature and what features tend to vary significantly in missions arising in real operations.

We have created an initial testbed mission set consisting of 243 missions based on 5 feature types (dimensions), each with 3 values. Feature types include: **N**, the number of targets to be observed, with possible values 4, 8 and 16; **spatial scale**, representing the size of the map in which the mission takes place, with possible values .002, .02 and .2 of the range of the vehicle; **spatial distribution**, the degree to which targets are clustered, with possible values of uniform, globular and 2-cluster; **maxcost distribution**, representing the variability across targets of the parameter m , with possible values of fixed, uniform, and 2-cluster; and **cost rate distribution**, the variability across targets of the cost rate parameter k ,³ with possible values fixed, uniform, and 2-cluster. All missions use the mission modeling framework described above and all have the following features in common: mission duration is fixed at 60 minutes (the worst-case flight duration of our RMAX helicopter); start/end point is

located at the centroid of mission targets; the probability of occurrence of all events is fixed at .2 per hour; and initial cost (c_0) = detection latency (l_0) = response latency (l_1) = reporting latency (l_2) = 0.

Because we expect to enlarge and refine the testbed repeatedly as our understanding of user needs and algorithm design tradeoffs grows, we have created software that lets us easily create and modify testbeds, and run evaluation experiments with both algorithms and human subjects. The software includes a model of the flight characteristics of the RMAX, allowing us to accurately compute travel time between targets. This is likely to be especially important for evaluating the impact of spatial scale, particularly where targets are relatively near one another, since turn rate in aerodynamic flight, acceleration to cruise speed and other UAV characteristics are likely to have large and varying effects on travel time.

Case Study: TSP vs. Human Performance

To illustrate the described evaluation framework, we describe its application to two surveillance methods. The first method is based on a 2-opt exchange algorithm (Reinelt 1994) for the Traveling Salesman Problem (TSP). The algorithm has been modified in a number of ways in order to (a) generate a repeating cycle of visits that start and end on a given location but do not visit it in the interim and (b) make use of a flight dynamics model requiring that travel time between locations is not a constant, but instead varies with initial speed, initial turn angle and end turn angle. Though these modifications make the algorithm more applicable to our surveillance problem, any TSP-based approach is likely to perform poorly in many of our testbed missions. We selected this method as a simple way to illustrate the general approach and because it might reasonably perform well in some cases.

The second method we evaluated was human-directed surveillance. A human subject selected surveillance paths for each of the 243 mission scenarios in our testbed. Each mission was represented graphically as a map showing all relevant dimensions. Targets were represented as icons colored to indicate cost rate (urgency) and with shape varied to represent maximum cost (importance). The start/end (home) point was displayed as a distinctive icon and spatial scale as a dotted-line circle centered on the home point whose radius represented .002 of the vehicle's specified flight range. Subjects used a mouse to select and modify a route and were allowed as much time as they wished on each mission. In contrast to the TSP-based method which always attempted to visit all targets, humans were allowed to exclude targets from the surveillance route if they wished.

Our initial expectations were that performance would vary significantly between the methods based on certain strengths and weaknesses. In particular, the TSP

³ Specific values of m are {10,20,30,40} with 30 used when maxcost distribution = fixed. Specific values of k are based on {20,40,60,80} minutes to reach .9 maxcost with 60 minutes used when cost rate distribution = fixed.

Pct. Adv.			N Spatial								
			4			8			16		
Scale	Rate	Cost	2-Cluster	Globular	Uniform	2-Cluster	Globular	Uniform	2-Cluster	Globular	Uniform
0.002	Fixed	Fixed	0	0	0	0	14	0	6	31	-3
		Clustered	0	0	0	0	16	5	7	29	7
		Uniform	0	0	-1	42	24	5	2	10	10
	Clustered	Fixed	0	0	0	0	0	11	0	0	11
		Clustered	0	0	0	0	0	0	0	0	1
		Uniform	0	0	0	0	0	0	0	0	1
	Uniform	Fixed	143	47	0	130	19	47	10	93	47
		Clustered	22	28	23	0	12	6	3	8	0
		Uniform	61	20	-2	0	18	7	2	5	9
0.02	Fixed	Fixed	0	0	0	1	1	0	0	1	0
		Clustered	0	0	0	1	1	0	0	1	0
		Uniform	0	0	0	1	1	0	1	2	0
	Clustered	Fixed	0	0	0	1	0	0	1	6	0
		Clustered	0	0	0	1	0	0	0	9	1
		Uniform	0	0	0	1	16	0	0	6	0
	Uniform	Fixed	0	0	0	1	0	0	2	2	0
		Clustered	0	0	0	1	0	0	2	2	1
		Uniform	0	0	0	1	0	0	0	2	1
0.2	Fixed	Fixed	1	0	0	0	-1	0	-2	-22	4
		Clustered	1	0	0	-1	5	-10	-6	-15	3
		Uniform	0	0	-1	-2	4	-8	-5	-19	4
	Clustered	Fixed	0	0	0	8	0	14	9	-1	10
		Clustered	0	0	0	9	-1	-11	11	-14	5
		Uniform	0	0	-1	16	8	3	5	-10	4
	Uniform	Fixed	3	0	23	14	-7	22	2	-7	10
		Clustered	16	0	23	9	-3	-12	9	-16	8
		Uniform	23	0	31	15	4	2	-3	-26	5

Table 1. Percentage difference in performance between TSP-based and human-directed surveillance

method, with a computer's advantages in speed and precision, would presumably do well on small scale maps where aerodynamic factors would favor complex paths that minimize turn angle rather than (only) distance between targets. It would likely perform poorly on maps with varying max-cost and rate parameters since it could not reason about that information. Humans, with natural visual-spatial capabilities that exceed any computer-based technique, might perform well when targets are spatially grouped. And, allowed to exclude targets from the surveillance schedule, people would likely perform well on maps where non-fixed distributions of max-cost and rate make some targets worth skipping and on large scale maps where the importance of being selective is especially great.

Table 1 shows data for all 243 missions. Data entries represent the percentage difference in performance between the two methods, with positive values indicating TSP advantage and negative values human advantage. Values outside the range -10% to 10% are in boldface to indicate where the greatest differences in performance lie.

Overall performance was comparable, with TSP doing 4.9% better on average. In the human best case, the subject outperformed TSP by 26%, whereas the best TSP case had a 146% advantage. The latter was almost certainly due to human error, as the mission in which it

occurred was similar to others where the subject performed well. This may indicate a phenomenon favoring algorithmic methods in general: human tendency to err when making surveillance decisions.

Across the five independent variables, scale and cost-distribution stood out as especially significant in differentiating human from TSP performance (standard deviations of 6.3 and 5.6 respectively). In all 9 cases where humans outperformed TSP by at least 10%, scale was large (.2). In 24 of 36 (66%) cases where TSP was better by at least 10%, scale was small (.002). N was least significant (s.d. = 1.3), though 7 of the 9 cases with human advantage $\geq 10\%$ were with N=16. The ability to exclude least-important targets is most likely to prove valuable in large scale maps with large numbers of targets. That human performance was best in those cases suggests that this ability was the principal human advantage.

Contrary to expectations, the TSP performed relatively well with non-fixed cost and rate distributions. It performed particularly well when the rate distribution was uniform, performing at least 10% better in 24 cases. Human advantage $\geq 10\%$ occurred with uniform rate in 3 cases. Confirming prior expectations, human performance was better in cases with spatial structure (globular and 2-cluster), especially when N was high – 7 cases with 10%

advantage vs. 1 for TSP. This suggests a constraint on the conditions in which people will be able to judge which targets to exclude.

This comparative evaluation was designed to test and illustrate our technique, not to test true candidate methods for inclusion on a UAV. Both methods could clearly be improved. The TSP-based algorithm could be made more efficient and used for more aggressive optimization. The human-directed method could be improved by training subjects to make better decisions and by adding decision support to the interface, conditions likely in any genuine operational context. Given the limits of our current data, we limit our interpretation of the results to the identification of general patterns that deserve further study.

Next Steps

As described in the first section, a practical and effective UAV-based surveillance capability requires efforts in three areas. The first is to develop means to evaluate and compare different surveillance methods. There are numerous ways to improve the presented approach. The mathematical framework should be extended to include more event types (e.g. sequentially reoccurring), more event features (e.g. detection latencies) and more diverse probability and cost functions. The mission testbed should be refined and extended to include additional features and a greater range of values for each feature type (e.g. $N=100$). The capabilities of human-directed surveillance should be further explored to characterize performance at high levels of expertise. And the whole framework should be extended to accommodate multiple surveillance agents including not only multiple UAVs, but also heterogeneous human and robotic observers.

The second area of work is to develop new and better surveillance algorithms, iteratively refining them based on comparative analyses of their strengths and weaknesses. A particularly important class of algorithms are those that make and/or modify surveillance decisions at execution-time in response to changing conditions (e.g. wind shifts, changes in user information needs). Though our framework has been described as a way to evaluate surveillance schedules prior to execution and without regard to such changes, it applies equally to post-hoc evaluation of schedules generated reactively (at execution-time) in response to unfolding events. As significant changes in physical conditions and user needs are likely to occur frequently in realistic missions, we anticipate that this framework will ultimately be more useful for evaluating reactive surveillance methods than for methods that schedule exclusively in advance. In particular, we anticipate applying it to assess ongoing scheduler enhancements to the Autonomous Rotorcraft Project helicopter's mission-level autonomy component (Apex).

Finally, these approaches must be made "usable" in real operational contexts where limits on time, knowledge and user expertise are likely to constrain interactions with the surveillance agent. One issue of particular concern is to enable users without a background in decision-theory or mathematics to specify mission parameters. Though users may be experts in the operational domain, eliciting the required utility and probability knowledge from them is notoriously difficult, though useful techniques exist (French 1986) and continue to emerge (Wang and Boutilier, 2003).

Acknowledgements

This work was supported by the NASA CICT / Intelligent Systems program. We thank Peter Cheeseman, Jeremy Frank, Roger Remington, Byron Roe, David Smith and Matthew Whalley for many useful discussions.

References

- Wang, T. and Boutilier, C. (2003) Incremental utility elicitation with the minimax regret decision criterion. In *Proceedings of the 2003 International Joint Conference on Artificial Intelligence*. Acapulco, Mexico.
- Carbonell, J.R. (1969) A queueing model for many-instrument visual sampling. *IEEE Transactions on Human Factors in Electronics*, 1966, HFE-7, 157-164.
- DoD (2002) *Unmanned aerial vehicles roadmap 2002-2027*. Report to the Office of the Secretary of Defense.
- Freed, M., Harris, R. and Whalley, M. (2004) Human-interaction challenges in UAV-based autonomous surveillance. In *Proceedings of the 2004 Spring Symposium on Interactions Between Humans and Autonomous Systems Over Extended Operations*, AAAI Press.
- Freed, M. (1998) Managing multiple tasks in complex, dynamic environments. In *Proceedings of the 1998 National Conference on Artificial Intelligence*. Madison, Wisconsin.
- French, S. (1986) *Decision Theory*. Halsted Press, New York.
- Koopman, B.O. (1956) The theory of search; Pt I: Kinematic Bases. *Operations Research*, 1956, 4, 324-346.
- Massios, N., Dorst, L. and Voorbraak, F. A Strategy for Robot Surveillance Using the Hierarchical Structure of the Environment. In *Proceedings of the IJCAI 2001 Workshop on Reasoning with Uncertainty in Robotics*, 43-50, Seattle, USA.
- Reinelt, G. (1994) *The Travelling Salesman. Computational Solutions for TSP Applications*, volume 840 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin Heidelberg New York, 1994.
- Sacks, S. (2003) Evaluation of Police Patrol Patterns. *Working papers 2003-17*, University of Connecticut, Department of Economics.
- Whalley, M., Freed, M., Takahashi, M., Christian, D., Patterson-Hine, A., Schulein, G. and Harris R. (2003) The NASA/Army Autonomous Rotorcraft Project. In *Proceedings of the American Helicopter Society 59th Annual Forum*. Phoenix, Arizona.