INFO-GAP APPROACH TO MULTI AGENT SEARCH UNDER SEVERE UNCERTAINTY

Itay Sisso, Tal Shima, *Senior Member, IEEE,* and Yakov Ben-Haim

***ABSTRACT*—A robust satisficing approach based on info-gap theory is suggested as a solution for a spatial search planning problem with imprecise probabilistic data. A group of agents are searching predefined patches of land for stationary targets, given an a priori probability map of the targets' locations. This prior probabilistic information is assumed to be severely uncertain and may contain large errors. An analysis of a simplified case shows that in some situations one might prefer a different strategy than the expected-utility maximizing one, in terms of robustness to uncertainty. Deterministic numeric results confirm the theoretical predictions for more complex cases. Finally, stochastic numeric analysis of robust satisficing solutions on a large group of much more complex, randomly generated cases, reveals an interesting behavior of a consolidation of effort in specific cells, and implies the potential of robust satisficing in more realistic scenarios. As the robustness to uncertainty comes at the expense of the expected utility, one must choose its decisions carefully. However, it is shown that, in various circumstances, one obtains results which are superior to the expected-utility maximizing strategy in terms of robustness, while sacrificing almost no expected utility.**

***Index Terms*—Decision-making, Uncertainty, Autonomous Agents, Cooperative systems, Genetic algorithms.**

I. INTRODUCTION

Originated by Bernard Koopman in the U.S. Navy's anti submarine warfare operations research group during World War II, the field of optimal search for targets is constantly evolving. Some of the different approaches to this problem nowadays include classical search theory [1], exhaustive geographical search [2], game theory [3], stochastic modeling [4] and neural networks [5],[6]. The search may be conducted by a single agent or a group of multiple agents. Today, when technological advances enable relatively cheap agents for a large diversity of applications (e.g., uninhabited aerial vehicles, or UAVs), the field of cooperative control is gaining more and more popularity [7], [8], [9]. As the complexity of finding an optimal solution for multiple cooperating searching agents rises exponentially with the number of agents, stochastic solution search methods and decentralized control are often used in the decision-making process.

In order to conduct an effective mission all relevant information must be incorporated into the planning / decision-making process. In a spatial search mission such information might be in the form of an a priori probability map of the targets' locations. A common method of representing the 2-D spatially distributed probability is through a cellular decomposition, i.e., the search area is divided into cells (usually with equal area) while each cell contains a probability value (also known as target occupancy probability) of target existence [1], [5], [10], [11], [12]. That type of information is usually obtained through assessments and various possibly noisy intelligence gathering sensors and therefore prone to severe uncertainties. As a result, much of the recent literature on cooperative search deals with different approaches to handling these uncertainties in order to provide robust search methods. The approaches presented by Boeva and De Baets [13] considered a delimitation of the prior probability between 2 bound values. Although simple, this method assumes further knowledge of the uncertainty. Another method suggested by Krokhmal et al. [14] uses a set of scenarios, each comprising a set of different probability values, and adopts the best decision using the Conditional Value at Risk (CVaR) method. The CVaR method works by averaging a percentage of the worst case scenarios and maximizing the utility for this average value. This method therefore assumes some probability distribution of the scenarios, data which is not necessarily available. Given no such information, one might assume the trivial case of uniform distribution. However, as Shafer pointed out [15], this might yield misleading results as in doing so one fails to distinguish between uncertainty, or lack of knowledge, and equal certainty. Furthermore, adopting the uniform distribution can lead to logical contradictions, as discussed by Keynes [16], and Ben-Haim [17]. A different approach for addressing the uncertainties in the probabilities is presented by Bertuccelli and How [10] where a Beta distribution is used as a distribution of the probability value in each cell. While yielding notable results, one might not expect to possess such information on the probability distribution in cases of qualitative assessment. Moreover, this information would be difficult to obtain relying on past observations when the exact posterior probability distribution remains unknown.

In this article we examine an info-gap approach to the decision-making problem of multiple agents search mission under severe uncertainty in the targets' locations. A previous version of this paper is presented in [18] for low and medium dimension cases. In this paper we further expand the discussion to random-generated large-scale scenarios, thus asserting the contribution of info-gap in more practical situations. Furthermore, in this paper we describe how genetic algorithms (GA) were used to solve these large-scale scenarios. Info-gap decision theory [19] is a method for decision-making under severe uncertainty which models the uncertainty as a family

This work was partially supported by the Arlene and Arnold Goldstein UAVs and Satellites Research Center at the Technion

I. Sisso and T. Shima are with the Faculty of Aerospace Engineering, Technion – Israel Institute of Technology, Haifa, 32000, Israel. (e-mail: sisso1@tx.technion.ac.il; tal.shima@technion.ac.il)

Y. Ben-Haim is with the Faculty of Mechanical Engineering, Technion – Israel Institute of Technology, Haifa, 32000, Israel. (e-mail: yakov@technion.ac.il)

of nested sets. Each set corresponds to a particular degree of knowledge-deficiency, according to its level of nesting. The process of choosing the best strategy is usually done by a robustness analysis. Instead of looking for the expected-utility maximizing strategy (which is a logical choice **if** no uncertainties are involved), we look for the robust satisficing strategy. Notice that we use the term "expected-utility" throughout this paper as the expected utility **assuming the estimates are true**. Robust satisficing strategy, as opposed to the Max-Min strategy, does not ameliorate worst-case scenarios but rather satisfices certain critical performance parameters (i.e., the demand) while maximizing robustness (i.e., the immunity to deficient information). While searching for a strategy which yields the most robustness to uncertainty (given a failure criterion), one ensures a solution which can "tolerate" the most knowledge deficiencies and still satisfice the given criterion. As shown by Ben-Haim [20] for a particular class of problems (the "proxy theorem"), the use of a robust-satisficing method such as info-gap ensures the maximization of the probability of success.

The outline of the paper is as follows: In section II we will formulate the problem and present a generic system model along with an uncertainty model to incorporate the lack of knowledge. In section III we will analytically study a simple degenerate case (1 agent, 1 target, and 2 distant search cells) using a more specific system model (which will be presented along with its assumptions) and show the implications of using a robust satisficing strategy for the decision-making as opposed to an expected-utility maximizing strategy. In section IV we will broaden the discussion to a more complex case of several distant cells while comparing numerical results (using a deterministic search method) of three different methods (expected-utility maximizing vs. local and global robustness maximizing). In section V we further expand the model to a more realistic multiple agents search, show how to solve it using a stochastic search method (GA), and discuss the behavior of the robust satisficing results, extracted from 500 random scenarios. Finally, we will provide some conclusions in section VI.

II. Problem Formulation

In this section we will mathematically formulate the descriptive system model of a group of agents searching for targets in an uncertain environment. We will start with a general model of the problem; continue to the uncertainty model which incorporates the lack of knowledge in the problem; then the definition of robustness will be provided; and conclude with a collection of all the assumptions made along the way.

A. General system model

We will begin with a formulation of the data, the decisions to be made, and the goal of the mission (the reward/utility function and the demand).

Let us consider a set C of N_c cells (e.g., patches of land), each denoted by an index i; and a set X of N_t targets, each denoted by an index j. The cells are physically detached and distant from one another. Each cell is attributed with a certain probability of occupancy by each one of the targets, denoted as $P_{i,j}$. In other words, $P_{i,j}$ represents the **actual** a priori probability that target j is in cell i. Due to its nature, the assessment of the a priori probability, given as $\tilde{P}_{i,j}$, is prone to imprecision and uncertainties. We now define P and $\tilde{P}$ as the actual and estimated probability matrices respectively.

The detection function, $P_d(i,\tau_i)$, expresses the conditional probability of detecting a given target within cell i after searching that cell for time τ_i (whether by a single agent or the sum of multiple agents search times, given that they satisfy the assumptions for that detection function, e.g., coordinate their search regions inside the cell in case of a methodic search), given that the target is in cell i.

The detection function used in this paper is the random search formula of Koopman [21]. This detection function assumes a random search inside each cell, but, as stated by Stone [12] (p.25) it is also useful as a lower bound on the effectiveness of a systematic search under most cases of path placement errors. The detection function is therefore given by the expression:

$$P_d(i,\tau_i) = 1 - e^{-\frac{wv\tau_i}{A_i}} = 1 - e^{-a_i\tau_i} \tag{1}$$

where w is the sweep width of the sensor, v is the speed of the agent (constant) and A_i is the area of cell i. For convenience, we define a_i as the exponential convergence rate, $a_i = \frac{wv}{A_i}$, which will be used later on.

As the cells are distant from each other, a cell switch time matrix is considered and denoted as c'. The elements of c', c'_{i_n,i_m}, represent the transition time from cell i_n to cell i_m (calculated as the distance between the centers of the 2 cells divided by v).

The payoff function for the i^{th} cell represents the expected number of targets detected in that cell after searching it for time τ_i and is denoted:

$$e\left(i,\tau_i\right) = P_d\left(i,\tau_i\right) \cdot \sum_{j=1}^{N_t} P_{i,j} \tag{2}$$

The search plan, denoted as θ, incorporates all the decision variables in it. In a task assignment problem these decision variables may include the cell switch times (at what times should the agents switch between cells) and to which cells (in case of a variety of cells). In a path planning problem these decision variables might include control commands, given to the agents at constant time intervals. Notice that in a multi-agent centralized approach, a single θ will consist of a search plan for each one of the agents. On the other hand, in a decentralized approach, each agent might hold a separate θ of itself.

Let us assume, without loss of generality, that the search process starts at cell number 1. We now formulate θ, the decision parameter, as follows: We define a search plan for a single (k^{th}) agent denoted as θ_k:

$$\theta_k = \left\{\left(t_{k,1}, t_{k,2}, ..., t_{k,N_s}\right), \left(i_{k,1}, i_{k,2}, ..., i_{k,N_s}\right)\right\} \tag{3}$$

where $t_{k,n}$ represents the times at which the agent is assigned a new search cell and $i_{k,n}$ represents the cell assigned in time $t_{k,n}$. N_s represents the number of steps (i.e. new cell assignments) in the entire search plan. A search plan for multiple agents will be formulated as a set of N_u single agent search plans (N_u being the number of agents):

$$\theta = \{\theta_1 \ \dots \ \theta_{N_u}\} \tag{4}$$

We now define Θ as the set of all the feasible values for the decision variable (in this case, all possible search plans). This set might include limitations such as fuel constraints (implemented as the maximum time of a search plan), kinematic constraints (if the search plan is given as a certain path to follow), etc. For the discussed case, let us formulate Θ as:

$$\begin{aligned}\Theta = \{&\theta : t_{k,n+1} \geq t_{k,n} \ \forall \ 1 \leq n \leq N_s - 1, k \in K;\\ &0 \leq t_{k,n} \leq T \ \forall \ 1 \leq n \leq N_s, k \in K;\\ &i_{k,1} ... i_{k,N_s} \in C \ \forall \ k \in K\} \qquad \text{where } K = \{1, ..., N_u\}\end{aligned} \tag{5}$$

where T is the upper bound for the total amount of time, defined by the limited fuel capacity and fuel consumption rate of the agents. In other words, in an acceptable search plan the switch times are arranged chronologically for each agent, the total time is bound to a constraint T, and the cells are within the set of participating cells - C. In the case of unvisited cells in a search plan, the cells' matching switch times will have the value of T (as switching at the end of the search period is equivalent to not switching).

Let us now define the overall payoff function representing the expected number of targets to be found in the entire search region:

$$E(\theta, P) = \sum_{i=1}^{N_c} e(i, \tau_i) \tag{6}$$

where τ_i is bound to limitations such as fuel constraints and is the sum of all search times in that cell (by the multiple agents).

We next define a demand that the overall payoff will be no less than a certain critical value - E_c (e.g., $0.7 \cdot N_t$, in case we demand to find, on average, 70% of the targets):

$$E(\theta, P) \geq E_c \tag{7}$$

The use of such a parameter in a search mission is practical for cases such as a tank squadron hunt, where the search is distributed between a few distant sites such as refueling facilities, hideouts, and several strategic points (where the best estimate of the probability distribution could be derived from past observations of the known duration of each stage according to the alert state, etc.). A demand (E_c) in that case could be the number of tanks needed to be destroyed in order to neutralize the squadron (a criterion used in modern combats, and is of course significantly lower than the number of tanks in the squadron).

E_c is a parameter expressing the analyst's aspirations. The analyst may have preferences on values of E_c before the analysis. However, the analysis of robustness to uncertainty for achieving any specified average number of target-finds (E_c) may lead the analyst to choose a different value of E_c – either larger or smaller – than the value initially preferred.

B. *Uncertainty model*

In order to incorporate the lack of knowledge into the model, let us consider an info-gap model of uncertainty on the a priori probability distributions of the targets:

$$\begin{aligned}U\left(\alpha, \tilde{P}\right) = \Big\{&P : \left|P_{i,j} - \tilde{P}_{i,j}\right| \leq \alpha\omega_{i,j},\\ &\sum_{i=1}^{N_c} P_{i,j} \leq 1, \ P_{i,j} \geq 0 \quad \forall \quad i \in C, \ j \in X\Big\}; \ \alpha \geq 0\end{aligned} \tag{8}$$

α is a non-negative number representing the uncertainty horizon and $\omega_{i,j}$ is a non-negative uncertainty weighting parameter intended for cases of unequal degree of (lack of) information regarding the targets' whereabouts. Note that this uncertainty model is suitable for the case where a probability distribution assessment for each target exists and, for a specific uncertainty horizon, the relative deviation from the actual distribution is defined for each cell and target using $\omega_{i,j}$. Of course, the actual magnitude of deviation of $P_{i,j}$ from $\tilde{P}_{i,j}$ is unknown since the horizon of uncertainty is unknown. If the assessment $P_{i,j}$ is more certain for a specific target $j \in T$ in cell $i \in C$, $\omega_{i,j}$ would be chosen as a smaller number. For example, if $\tilde{P}_{i,j}$ is the result of a statistical estimation, the $\omega_{i,j}$ could be the statistical error of this estimate. In case all the information regarding the probability distribution is equally uncertain then one might choose $\omega_{i,j} = 1 \ \forall \ i, j$.

Also notice that a constraint is given on the sum of probabilities ($P_{i,j}$) as an **inequality** condition, rather than an equality. In other words, we "allow" the targets to be out of the entire search region. If this is not the case, and the targets are known for certain to be somewhere in the search region, an equality condition should be imposed. Although at first glance it might seem like a minor adjustment, such a modification of the constraints radically changes the uncertainty model and might yield different recommendations of the robust maximizing strategy (which is acceptable, considering the fact that it holds more information on the situation). Such a choice of constraints would greatly complicate the calculation method and is not considered in this paper.

This approach to the data uncertainty determines that the amount of uncertainty affects each cell individually (relatively to $\omega_{i,j}$), and is therefore of a rather local manner. One might think of different ways to express the uncertainty in the a priori data. Another approach which will be examined in this article is a global robustness approach in which the uncertainty horizon, α, determines the maximal **sum** of data errors. The uncertainty model formulation using the global robustness approach is therefore:

$$\begin{aligned}U\left(\alpha, \tilde{P}\right) = \Bigg\{&P : \sum_{i=1}^{N_c} \left|\frac{P_{i,j} - \tilde{P}_{i,j}}{\omega_{i,j}}\right| \leq \alpha,\\ &\sum_{i=1}^{N_c} P_{i,j} \leq 1, \ P_{i,j} \geq 0 \quad \forall \quad i \in C, \ j \in X\Bigg\}; \ \alpha \geq 0\end{aligned} \tag{9}$$

Notice that $\omega_{i,j}$ is now dividing $\left|P_{i,j} - \tilde{P}_{i,j}\right|$ so each cell and target might contribute in a different way (by their level of knowledge) to the overall uncertainty. $\omega_{i,j}$ is therefore non-zero too (in case of an absolutely definite value of a certain

$\tilde{P}_{i,j}$ (i.e., $w_{i,j} = 0$) this specific $\tilde{P}_{i,j}$ should be removed from the uncertainty model). This type of uncertainty might be useful when considering a scenario in which a certain source divides a certain amount of total research effort between the cells (in an unknown manner) while obtaining the prior probability distribution - $\tilde{P}$, so that decreasing the likelihood of an error in one of the cells would increase the likelihood of an error in the others.

C. Robustness

An important parameter in the field of decision under severe uncertainty is the robustness. In order to minimize the effect of uncertainty on the result, one might choose the strategy which maximizes the robustness to uncertainty. The robustness to uncertainty in info-gap theory is defined as the maximum uncertainty horizon for which no failure is possible (given the failure criterion - E_c). Let us formulate the robustness of a search plan θ given a failure criterion E_c:

$$\hat{\alpha}(\theta, E_c) = \max\left\{\alpha : \left[\min_{P \in U(\alpha, \tilde{P})} E(\theta, P)\right] \geq E_c\right\} \tag{10}$$

Note that $\hat{\alpha}$ is the least upper bound of a set of α values. When this set is empty, then we define $\hat{\alpha} = 0$.

In order to choose the best search plan under uncertainties in the a priori probabilities one might aspire to maximize the robustness of the search plan, as defined above. The robust maximizing search plan, denoted as $\hat{\theta}$ is then formulated as:

$$\hat{\theta} = \arg\max_{\theta \in \Theta} \hat{\alpha}(\theta, E_c) \tag{11}$$

The complexity of solving equation 11 is dependent on the system model. It is possible to have multiple solutions for equations 11 and 10.

We now have a general mathematical formulation of the problem at hand including a model which incorporates uncertainties in the preliminary data and a method of choosing the best decision.

III. Analytical Investigation of a Low Order Problem

In this section we will demonstrate the use of info-gap decision theory on a low order degenerate case of the problem of agents search. We will specify the system model furthermore using the definitions from section II. For simplicity, only the local uncertainty model will be examined.

A. Specific system model

Let us consider a simplified low dimensional case of 1 agent searching in 2 cells, distant from one another, for a single target. Fuel constraint limits the search time to T. The search is performed randomly in each cell and a decision has to be made regarding the time of switch to the second cell.

The search plan decision is therefore reduced to choosing t, the time of switch to cell number 2, which is limited by the fuel constraint:

$$\theta = \{t\} \qquad \Theta = \{t : 0 \leq t \leq T\} \tag{12}$$

The switch time matrix, c', in this case is reduced to a single parameter, c, which represents the travel time from cell 1 to cell 2 and vise versa.

B. Maximum robustness strategy

Let us now mathematically analyze the maximum robustness strategy of choosing the best decision in the degenerate 1-agent 2-cells case. Later on a maximum expected-utility strategy will be presented and a comparison will be made between the two.

In order to find the decision which maximizes the robustness to uncertainty, $\hat{\alpha}$, we must first find an expression for the robustness of a decision (in this simplified case t, the time of switch between the cells) as a function of the demand (E_c, the critical expected utility).

Let us start by finding an explicit presentation of the utility function E as a function of the decision θ (or simply t in this case) and the state P. For the case of 1 target the payoff function can be simply derived from equations 2, 1, and 6 as:

$$E(\theta, P) = P_1 \cdot \left(1 - e^{-\frac{wv}{A_1}t}\right) + P_2 \cdot \left(1 - e^{-\frac{wv}{A_2}\max\{T-t-c,0\}}\right) \tag{13}$$

where P_1 and P_2 are defined as the actual probabilities of the target being in cell number 1 and 2 respectively.

Let us assume for simplification that the uncertainty weighting parameters $\omega_{1,1}$ and $\omega_{2,1}$ both equal to 1 (the relative uncertainty in the two cells is identical). Using the local uncertainty model presented in equation 8 and the payoff function in equation 13 we derive the following equation:

$$\begin{aligned}\min_{P \in U(\alpha, \tilde{P})} E(\theta, P) &= \max\left\{\tilde{P}_1 - \alpha, 0\right\} \cdot \left(1 - e^{-\frac{wv}{A_1}t}\right) \\ &+ \max\left\{\tilde{P}_2 - \alpha, 0\right\} \cdot \left(1 - e^{-\frac{wv}{A_2}\max\{T-t-c,0\}}\right)\end{aligned} \tag{14}$$

This function is monotonically decreasing with α, therefore, the robustness $\hat{\alpha}$ (the maximal α that satisfices the requirement in equation 10) can be extracted through the solution for $\hat{\alpha}$ of:

$$\begin{aligned}E_c = &\max\left\{\tilde{P}_1 - \hat{\alpha}, 0\right\} \cdot \left(1 - e^{-\frac{wv}{A_1}t}\right) \\ &+ \max\left\{\tilde{P}_2 - \hat{\alpha}, 0\right\} \cdot \left(1 - e^{-\frac{wv}{A_2}\max\{T-t-c,0\}}\right)\end{aligned} \tag{15}$$

The robustness function can be explicitly formulated as:

$$\hat{\alpha} = \begin{cases} F & , \ 0 \leq F \leq \tilde{P}_{min} \\ \tilde{P}_{max} - \frac{E_c}{P_d(1,\tau_1)} & , \ \tilde{P}_{min} \leq F \leq \tilde{P}_{max} \\ 0 & , \ \text{else} \end{cases} \tag{16}$$

where:

$$F = \frac{\tilde{E} - E_c}{P_d(1, \tau_1) + P_d(2, \tau_2)} \tag{17}$$

$$\tilde{P}_{min} = \min\left\{\tilde{P}_1, \tilde{P}_2\right\}; \quad \tilde{P}_{max} = \max\left\{\tilde{P}_1, \tilde{P}_2\right\} \tag{18}$$

$$\tau_1 = t; \quad \tau_2 = \max\{T - t - c, 0\} \tag{19}$$

$$\tilde{E} = E\left(\theta, \tilde{P}\right) = \tilde{P}_1 \cdot P_d(1, \tau_1) + \tilde{P}_2 \cdot P_d(2, \tau_2) \tag{20}$$

F cannot be greater than both $\tilde{P}_1$ and $\tilde{P}_2$.

Notice that $\hat{\alpha} = 0$ for every $E_c \geq \tilde{E}$. In other words, when the demand is larger (or equal) to the expected utility, the robustness is 0. This is a logical statement considering the robustness definition from subsection II-C - the maximum uncertainty horizon for which no failure is possible. When the expected utility is below the demand even an uncertainty horizon of 0 will allow failure.

Note that the slope of $\hat{\alpha}$ vs. E_c is always negative. This negative slope is the mathematical manifestation of the robustness-demand tradeoff. As the demand gets higher one can naturally assure less immunity to uncertainty and the robustness therefore decreases.

C. Curve crossing example

In this subsection we will show and explain the interesting phenomenon of crossing robustness curves. When robustness curves ($\hat{\alpha}(t, E_c)$ vs. E_c) for two different choices of t cross each other one concludes that the best decision (between these two values of t) depends on the value of E_c that one requires. When one of these curves is the one of the expected-utility maximizing decision one might say that, depending on E_c, there is a more robust decision than the expected-utility maximizing one.

Figure 1 shows robustness curves for different values of the decision variable, t, in a specific example ($a = 1$, $c = 1$, $T = 3.5$, $P_1 = 0.65$, and $P_2 = 0.35$). The striped and dotted lines show the robustness curves for $t = T$ and $t = 0$ respectively, while the solid line shows the robustness curve of the decision for optimal expected utility ($t = t^*$). Notice the crossing of the striped and solid curves around $E_c = 0.41$. While the decision $t = t^*$ (solid) maximizes the expected utility, a choice of $t = T$ (striped) will yield more robustness for $E_c < 0.41$.

Notice that the value of E_c at $\hat{\alpha} = 0$ for every curve is the expected utility of the decision, as evident from equation 15. One may logically understand that fact as for a demand which equals the **expected** utility there is no room for error, as the smallest negative deviation from the given data will result in a failure (as it will lower the reward beneath the demand), and therefore $\hat{\alpha} = 0$. The decision of $t = T$ therefore yields less expected utility than $t = t^*$ (as expected).

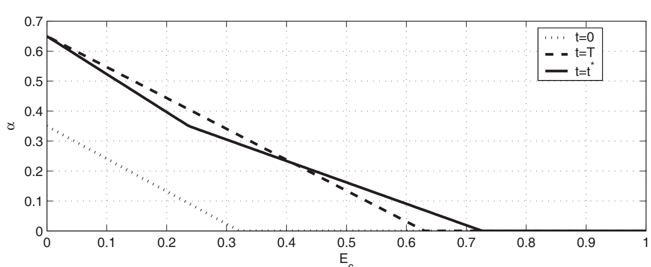

Fig. 1. Curve crossing- $\hat{\alpha}(t, E_c)$ vs. E_c for different values of t ($a = 1$, $c = 1$, $T = 3.5$, $P_1 = 0.65$, and $P_2 = 0.35$)

The piecewise linear form of the graph can be intuitively explained as follows: As the uncertainty horizon increases (the α axis in Figure 1), the minimal P_i value in this uncertainty horizon decreases linearly (due to the chosen uncertainty model described in eq. 8), and with it the expected utility (and hence the demand this solution can fulfill - E_c). The change in the linear slope (i.e., the kink in the solid line) is caused due to the "depletion" of P_i in one of the cells (in this case - cell 2), as from this point on, as the uncertainty horizon keeps on increasing, only P_1 decreases (as $P_2 = 0$).

A detailed analysis in [18] shows that for the case of 2 cells, 1 agent and 1 target problem, a robustness curve of a different decision will **always** cross the robustness curve of $t = t^*$ when $\tilde{P}_1 \neq \tilde{P}_2$, and $t^* \neq T$ (while $\tilde{P}_1 > \tilde{P}_2$) or $t^* \neq 0$ (while $\tilde{P}_2 > \tilde{P}_1$).

IV. Deterministic Numerical Investigation of a Medium Order Problem

After mathematically analyzing the behavior of the optimal decision in terms of maximum robustness and maximum expected utility in a low order case, let us now generalize and analyze a multi-dimensional case by increasing the number of cells in the problem. The decision to be made now is the switching time between cells and the cells to switch to, under a fuel constraint which again limits the search time of each agent to T. Due to its mathematical complexity, this case will be inspected through a numerical analysis using a deterministic search method of finding the optimal solution.

Unlike in section III, in this section both the local and global robustness approaches (see subsection II-B on page 3) will be examined and compared.

A. Solution method

In order to remove any doubts regarding the optimality of the results, the numeric investigation in this section will be performed using a deterministic search algorithm. Where in the infamous traveling salesman problem (TSP) one is required to find only a sequence of cells to be visited, in our problem one is required to find a sequence and a duration of staying in each cell. As the problem at hand is clearly more complex than TSP, it is clearly of a NP-hard complexity. This fact limits the dimensionality of the problem to a relatively small number, in order to be solved in a reasonable time, and hence the medium order type of problem in this section.

In order to obtain a numerical result of the chosen strategies (robust maximizing and expected-utility maximizing) we discretize the time (t) dimension using a certain resolution, i.e., the cell switch times can only get specific equally spaced values. The resolution of this discretization greatly affects the calculation time of the solution.

The nature of this problem does not allow using common tree-search methods, such as A^* and Dijkstra. A branch and bound method was considered, however, it requires an efficient upper bound calculation method, in order to work efficiently and reduce the number of branches in the search tree of the problem, which was yet to be found. The search for solution is therefore done by a brute force approach of checking all the possible search plans (after discarding the ones in which a new cell is assigned before reaching the next one) and choosing the one with the most robustness or expected utility. Due to the great number of cases this process is very time consuming. The number of cells is therefore kept low as well as the number of time steps, and, as mentioned before, the problem is limited to a single agent search.

B. Example

We will consider a problem of four equally spaced search cells arranged in the corners of two adjacent equilateral triangles, each with an edge length of 0.2. In one of these cells resides a single target. We would like to find a search strategy for a single agent searching these cells. The best estimate for the occupancy probability distribution for this target is: $\tilde{P} = \{0.45, 0.05, 0.15, 0.35\}$.

For simplicity, we will assume $\omega_{i,1} = 1 \; \forall \, i \in C$ (the amount of uncertainty is even in all the cells). Let us now examine the best strategy chosen by the expected-utility maximizing approach and the robustness maximizing approach. After running the program for a search time (T) value of 2, detection function exponential convergence rate (a, or $\frac{wv}{A}$) of 2, and expected utility demand (E_c) of 0.4 (at a time resolution of 0.025, or 80 time steps) the following results were obtained:

1) Largest expected utility result: The largest expected utility was achieved by a search plan which divided the search time between the two "high probability" cells (1 and 4). The search times in the selected strategy (calculated with a time resolution of 0.025) are- $\tau_1 = 0.975, \tau_4 = 0.825, \tau_2 = \tau_3 = 0$ (which, along with the 0.2 transition time from cell 1 to cell 4, adds to the total $T = 2$). Notice that the higher probability cell (1) got more search time allocated to it. The expected utility of this strategy (the probability of finding the target given the prior probability distribution and the described search time allocation) is 0.67.

2) Largest global robustness result: As mentioned in subsection II-B, the robustness calculation method is divided into two approaches - global (eq. 9) and local (eq. 8). Let us compare these two approaches for the case at hand. In the **globally** robust maximizing (at $E_c = 0.4$) search plan, the search effort is distributed differently between the cells. As presented in Figure 2, the search effort is now distributed between 3 cells (1, 3, and 4), while cells 1 and 4 are allocated an identical search time ($\tau_1 = \tau_4 = 0.65$), and cell 3 is allocated less time ($\tau_3 = 0.3$, which, along with the 2 passes which take 0.4 time units in total, adds up to $T = 2$). Two possible ways of this time allocation are shown in the figure. Notice that these two options, although different from the perspective of the control system of the agent, are identical in terms of time allocation (τ), and are therefore treated as one in the following analysis. Figure 3 shows a comparison of the robustness (calculated in the global method) vs. E_c for the global robustness maximizing strategy (at $E_c = 0.4$) and the expected-utility maximizing one. The expected utility of this strategy is 0.64.

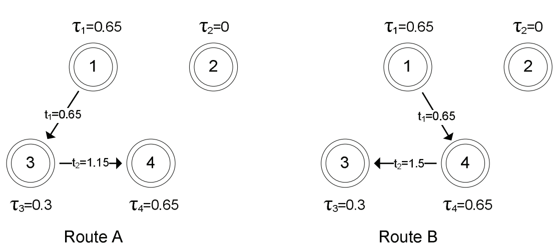

Fig. 2. Global robustness maximizing strategy (at $E_c = 0.4$) results for the case of $\tilde{P} = \{0.45, 0.05, 0.15, 0.35\}$, $T = 2$, $a = 2$

Let us now compare the two different strategies (the expected-utility maximizing and robustness maximizing strategies). As can be seen in Figure 3, when looking at the robustness curves of these two strategies as a function of E_c the two curves cross at about $E_c = 0.53$. Therefore, while striving to maximize the global robustness to uncertainty in the probability distribution, and for an E_c value of less than 0.53, one would prefer a choice of the 2^{nd} strategy (the one which maximizes the robustness at $E_c = 0.4$) over the 1^{st} one despite the fact that the latter yields greater expected utility (as can be seen by the E_c value of the two curves at $\hat{\alpha} = 0$).

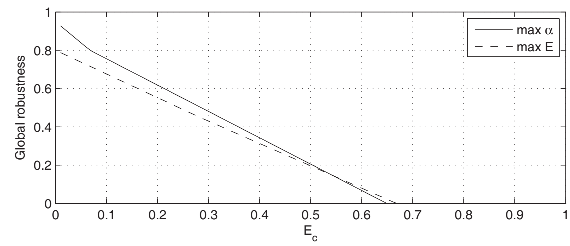

Fig. 3. Global robustness vs. E_c of robust maximizing strategy (at $E_c = 0.4$) and expected-utility maximizing strategy for $\tilde{P} = \{0.45, 0.05, 0.15, 0.35\}$, $T = 2$, $a = 2$

The robustness difference between the two strategies is not very significant for the given case. The expected utility change is very minor too. For a drop of 0.03 in the expected utility (4.5%, from 0.67 to 0.64, as can be seen in the crossing of the x axis) one gains 0.08 in global robustness for $E_c = 0.2$ (14.5%, from 0.55 to 0.63). In other words, the prior estimation of probability, $\tilde{P}$, can error (in total) in 0.08 more when using the strategy presented in subsection IV-B2 than when using the one presented in subsection IV-B1, and still satisfice the 0.2 utility demand (while decreasing the probability of finding the target given no data errors, by only 0.03).

3) Largest local robustness result: As can be seen in Figure 4, the local robustness (at $E_c = 0.4$) maximizing strategy is the same as the expected-utility maximizing strategy. As we will see in the next subsection, a change in one or more of the problem's variables might change this situation and lead to cases where the local robustness method is more robust (locally speaking) than the expected-utility maximizing strategy. One might notice, by comparing Figure 3 and 4, that the values of the local robustness are lower than those of the global one. This fact is clear as the local robustness refers to a maximal error in each one of the cells while the global approach refers to the sum of errors in all the cells together.

C. Demand effect

Let us now examine the behavior of the robust maximizing solution in comparison to the expected-utility maximizing one, while changing the value of the demand - E_c. Due to severe complexity issues involved in the global robustness model for a large number of cells, we will only analyze the local robustness model. Figure 5 shows the change in the search time allocated for each cell, in both the local robust maximizing strategy and expected-utility maximizing strategy, with the change in E_c

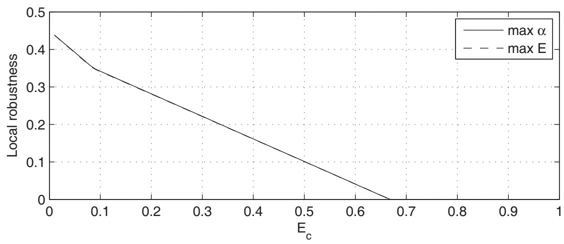

Fig. 4. Local robustness vs. E_c of robust maximizing strategy (at $E_c = 0.4$) and expected-utility maximizing strategy for $\tilde{P} = \{0.45, 0.05, 0.15, 0.35\}$, $T = 2$, $a = 2$. The two robustness curves merge to one in this case

(top subfigure), as well as the change in the local robustness (middle subfigure) and expected utility (bottom subfigure) of both solutions. The thick lines represent the robust maximizing solution while the thin lines represent the expected-utility maximizing one.

Notice that, unlike Figures 3 and 4, in Figure 5 each point in E_c represents a different solution in the robust maximizing lines (the one maximizing the robustness for that E_c value). One can see the robust maximizing lines in Figure 5 as the maximum of a collection of many robustness curves, such as the ones presented in Figures 3 and 4 for a specific E_c (0.4). Looking at the search time allocation graph (top subfigure) it is evident, as expected, that the maximum expected utility solution is unaffected by the change of E_c.

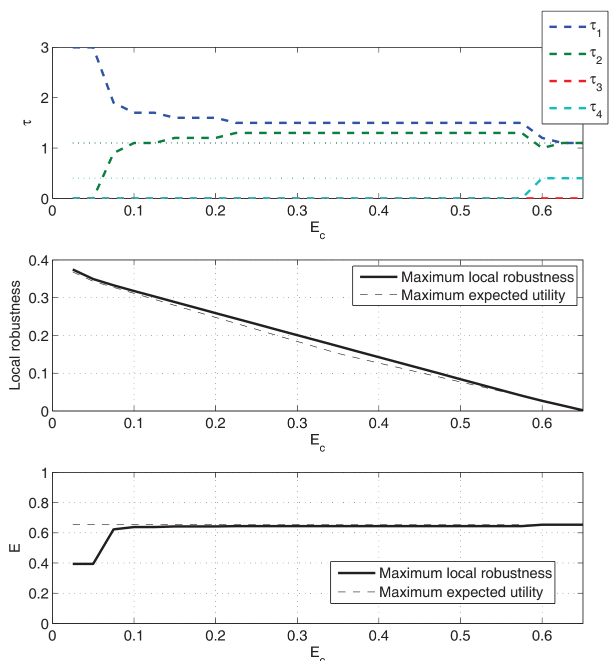

Fig. 5. Cell search time allocation (top), local robustness (middle), and expected utility (bottom) vs. E_c, of local robustness maximizing strategy (at each E_c) (thick lines) and expected-utility maximizing strategy (thin lines), for $\tilde{P} = \{0.4, 0.35, 0.1, 0.15\}$, $T = 3$, $a = 1.4$

It is evident that for the very low values of E_c (around 0.05) the local robustness maximizing solution substantially differs from the expected-utility maximizing one. This results in a significant decrease in the expected utility of the robust maximizing solution, while yielding only an insignificant increase in robustness. A mission planner must take into account these possible scenarios and ask oneself - how much is this demand (E_c) more important than higher values of expected utility? If, for example, an expected utility lower than the demand would be considered a failure while a higher one be considered as sheer success, and no extra benefit is gained from a higher value of expected utility - then one should implement the robust maximizing strategy. However, this situation is unreasonable in many cases, such as the tank search scenario described earlier. If one benefits more from an expected utility higher than the demand (as expected in the example), then this fact should be incorporated in the analysis. In such cases, the robustness to the demand should be weighted along with the robustness to higher values of demands and the expected utility of the strategy (through a strategies robustness curve comparison). Keep in mind that, while increasing the expected utility seems to be desirable, it always comes at the expense of decreased robustness to uncertainty in the prior data (in this case, the occupancy probability distribution).

V. Stochastic Numerical Investigation of a High Order Problem

In the last two sections, we demonstrated an analytical and a numerical analysis of one-dimensional and medium-order multi-dimensional case studies of an agent search mission with uncertainty on the prior information. Let us now perform an analysis of a more realistic higher-order multi-dimensional case. Deterministic solving methods, such as the one presented in the previous section, are inappropriate in such cases due to the very high complexity level. We therefore suggest a GA as a stochastic search method. Being a heuristic search method, GAs allow one to find adequate (although usually sub-optimal) solutions to complex problems in reasonable time. Although one cannot assure the optimality of the solutions using a stochastic search algorithm, a sufficiently long evolution of the solutions (depending on the fitness convergence rate) can usually assure very good solutions.

We will now briefly describe the GA modeling used in this work. For convenience, the chosen chromosome structure, which incorporates the cell visit order and times, is in the form of 2 vectors for each agent, consisting of the chronological visit order in one (a series of cell numbers), and the time to perform each visit in the other. A few limitations are applied to qualify as a legitimate chromosome. As the first vector in the sub-chromosome of each agent represent the chronological visit order, the second vector must be sorted as increasing time values (between 0 and T). The time spacing between each two cell visits must be large enough to include the travel time between the cells (stored in the travel time matrix - c). Furthermore, to speed up the convergence process, no repetition to a formerly visited cell (by the same agent) is allowed (as that would be a waste of time on travels between cells). Another limit forced in the generation of new (or new parts) of chromosomes is that the next cell to visit is usually chosen from a reduced number of optional cells in the vicinity of the current one. In the performed runs, the optional next cells for each step are the 5 closest ones (out of 15 cells in general). This "neighborhood choosing" method is used in the generation of each new step (in the initialization, crossover, and mutation processes) with a predefined probability (0.9 in our case), rather than always, to allow diversity (when it is not used, the next cell is chosen randomly from all the other cells). Figure 6 shows an example of a chromosome of a 2-agents 10-cell task-assignment problem. Again, the starting condition, consisting of the first cell of each agent, are stored in a separate vector. Notice that the last few steps in every sub-chromosome are done at time $t = T$. This is a method for not including some of the cells in a certain agent's mission.

		step 1	step 2	step 3	step 4	step 5	step 6	step 7	step 8	step 9
UAV1	Cell #	2	6	3	9	5	4	7	8	10
	Time [min]	0.2	0.7	1.5	2.1	2.5	3	3	3	3
UAV2	Cell #	5	10	4	3	8	2	9	7	6
	Time [min]	0.1	0.4	0.9	1.3	1.8	2.3	2.7	3	3

Fig. 6. Chromosome structure example for a 2 agents, 10 cells problem ($T = 3$)

Due to all the special limitations on the chromosome's structure, a simple crossover operation is not applicable in this case. Since the problem has some resemblance to the well-known traveling-salesman problem (TSP), some GA encoding methods for TSP have been considered. However, these methods assume TSP conditions which are not true in this case, such as the necessity to visit each city exactly once. Moreover, these methods fall short of addressing the time dedicated to searching each cell. As a result, no crossover operator is encoded in the GA code for this model in the sense of combining 2 solutions to create 2 new ones, but rather a random slicing of solutions, while the rest of the solution is completed randomly step by step using the "optional close cells" rule. The times are set randomly in a rising order. A special module then processes the solutions to fix illegal cases (i.e., cells visited more than once, and insufficient time gaps).

Let us compare the results of the GA using the two fitness calculation methods - the expected utility and the local robustness. We now consider a problem of a single target, searched for in 15 disconnected equally sized search cells, randomly placed in a 2000x2000[m] area, by 3 identical searching agents. The a-priori probability distribution ($\tilde{P}$) is also randomly distributed. Again, for simplicity, we assume $\omega_{i,1} = 1 \forall i \in C$ (the amount of uncertainty is even in all the cells). In all the calculated scenarios, the final time is: $T = 3$[min], the speed of the agents: $v = 100[m/s]$, the sensor's sweep width: $w = 20[m]$, and the area of each cell: $A_i = 60,000[m^2]\ \forall\ i \in C$ (which makes the detection function convergence rate: $a = \frac{vw}{A} = 2\left[\frac{1}{\text{min}}\right]$). The shown results were taken for each run as the best of 3 runs, each one with a population size of 100, evolved for 200 generations. As can be seen in Figure 7, when examining the progression of the mean maximum fitness (calculated from 50 random scenarios), or MMF, with the generations, the three types of fitness (expected utility (solid), robustness at $E_c = 0.5$ (dashed), and robustness at $E_c = 0.1$ (dotted)) all seem to be well converged after 200 generations. Computationally-wise, the robust-satisficing strategy is harder to calculate, however, it is not by an order of magnitude. While the calculation of the expected utility maximizing solution, using the GA, took around 20 seconds (on an average PC), the calculation of the robustness maximizing one took around 75 seconds.

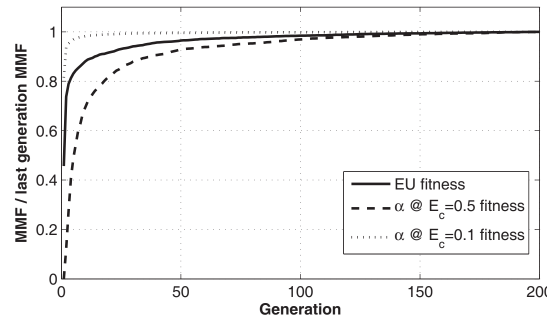

Fig. 7. The MMF divided by the MMF value at the last generation to compare the convergence of the three fitness types after 200 generations

We will begin analyzing a single scenario, and continue with some average results obtained from a group of randomly generated scenarios. The random variables are the locations of the cells, and the initial probability distribution.

A. Example

The following results were obtained using the GA for a specific case. Figure 8 shows the evolved courses of 3 different fitness types: expected utility (left), robustness for $E_c = 0.5$, and robustness for $E_c = 0.1$.

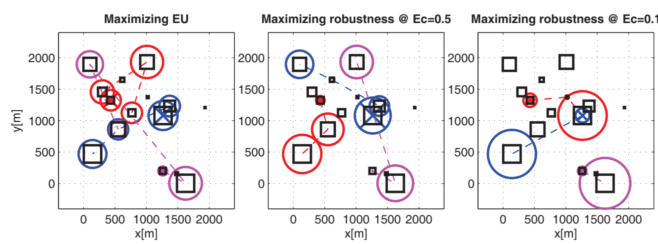

Fig. 8. From left to right - the expected-utility maximizing, the robustness maximizing at $E_c = 0.5$, and the robustness maximizing at $E_c = 0.1$ courses, calculated by the GA for the same scenario. The black squares represent the locations of the cells and their sizes – the $\tilde{P}_i$ value. The circles represent the time invested by each agent (different color for each one) in each cell. The starting cell for each agent is marked with an X inside the circle and the transitions between the cells by the crushed lines.

Looking at these graphs, one might see a distinct difference between them: as E_c is lowered (remember that maximizing the expected utility stands for the highest E_c), fewer cells are visited and the search effort is invested mainly more and more in the higher probability cells (those with a relatively high value of $\tilde{P}_i$). This search-effort allocation behavior makes sense considering the local uncertainty model used in the robustness calculation, given in equation 8. As E_c decreases, the robustness increases, allowing the worst-case probability under the uncertainty horizon, to be zero in more cells. This fact leads to the strategy of allocating more time to the higher probability cells, which will still yield some utility in relatively large uncertainty situations. For a very small E_c value, the strategy of visiting only the very high prior-probability cells could withstand a great deal of uncertainty, as these will still “hold” enough probability to satisfice the requirement.

Let us now look at the robustness curves (robustness vs. E_c) for the three strategies presented in Figure 8. Figure 9 shows the three robustness curves of the expected-utility maximizing, robustness maximizing at $E_c = 0.5$, and robustness maximizing at $E_c = 0.1$ curves.

First, let us compare the strategy of maximizing the robustness at $E_c = 0.5$ (dashed), and the strategy of maximizing the robustness at $E_c = 0.1$ (dotted). Notice that while using the latter, the resulting robustness at $E_c = 0.1$ is higher, but merely by 1%. However, the former strategy enjoys a very clear advantage in the robustness at higher E_c values. Moreover, remembering that the robustness curve’s crossing of the E_c axis represents the course’s expected utility, notice that the course which maximizes the robustness at $E_c = 0.5$ (dashed) has a much larger expected utility (about 50% more!) than that of the one maximizing the robustness at $E_c = 0.1$ (dotted). Therefore, in a realistic situation (such as the tank search example), it is hard to imagine a decision maker who would prefer the strategy which maximizes the robustness at $E_c = 0.1$ over the strategy which maximizes the robustness at $E_c = 0.5$, even if the failure criterion is $E_c = 0.1$.

Let us now compare the expected-utility maximizing strategy (solid) to the robustness maximizing at $E_c = 0.5$ strategy (dashed). The maximum expected utility (0.782) is only a fraction of a percent higher than that of the robustness maximizing at $E_c = 0.5$ strategy (0.78). However, the robustness gain from the robustness maximizing at $E_c = 0.5$ strategy is quite significant (as much as 20% gain at $E_c = 0.53$!).

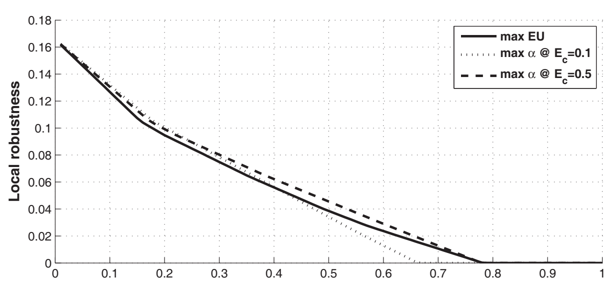

Fig. 9. The robustness curves of the expected-utility maximizing (solid line), robustness maximizing at $E_c = 0.1$ (dotted line), and robustness maximizing at $E_c = 0.5$ (dashed line) solutions (the scenario from Figure 8).

B. Analyzing average results

In order to understand the general behavior of these solutions, an analysis of 500 randomly generated scenarios was made. Figure 10 shows the robustness (at $E_c = 0.5$) and expected utility gain histograms. The upper frame shows the robustness of the robust-maximizing (RM) solution divided by the robustness of the expected utility maximizing (EUM) solution. The lower frame shows the expected utility of the EUM solution divided by the expected utility of the RM solution. The dashed line in the upper frame shows an average 5.3% gain in robustness of the RM strategy over the EUM

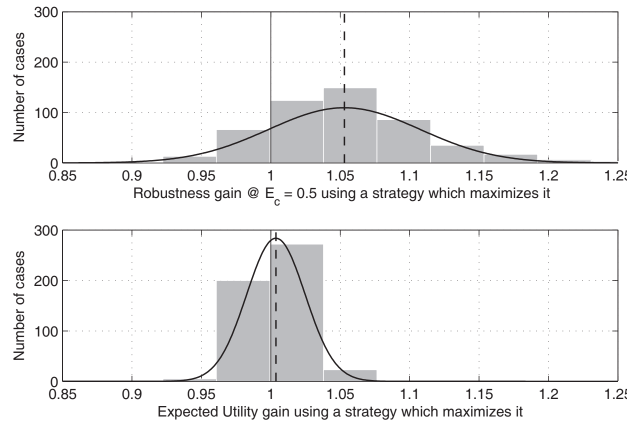

Fig. 10. The robustness (at $E_c = 0.5$)) gain histogram (top) and the expected-utility gain histogram (bottom), along with a fitted (multiplied) normal distribution PDFs (solid lines), calculated from 500 random scenarios.

strategy. The standard deviation is 0.055. A t test would strongly reject the hypothesis that the true mean of this distribution is 1 (implying no difference between RM and EUM), against the hypothesis that the mean is 1.053 ($p < 0.001$). The lower frame shows that this robust gain entails no loss in expected utility, since the ratio of the expected utilities of the two strategies is essentially one. Notice in the upper histogram that 14% of the cases are below 1, meaning that the expected utility maximizing strategy yielded more robustness than the robustness maximizing strategy 14% of the time. This is perfectly normal due to the stochastic nature of the GA. The histogram clearly shows that these 14% are the tail of the PDF, and the mean is significantly larger than 1. For the expected utility gain the number of cases below 1 is more than 40%, and the average gain is 0.4%.

Let us now examine the observation regarding the number of visited cells, which was decreasing as E_c was lowered, as can be seen in Figure 8. Figure 11 shows the average number of visited cells (by all three agents, calculated from 50 randomly generated scenarios) using the robustness maximizing strategy at each E_c value (dashed), and using the expected-utility maximizing one (solid). The expected-utility maximizing strategy is independent of E_c and the solid line is therefore constant. It can be seen clearly that the average number of visited cells increases with E_c, and that the average number of visited cells by the expected-utility maximizing strategy is higher.

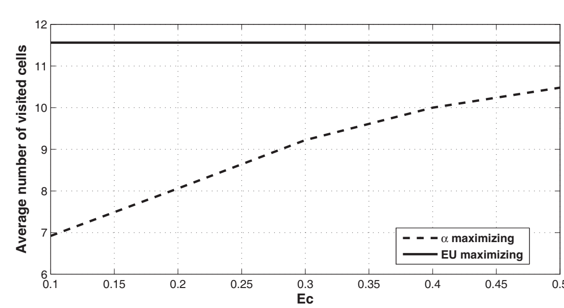

Fig. 11. Average number of visited cells using expected-utility maximizing strategy (solid) and robust maximizing strategy (dashed) at each E_c

VI. Conclusions

An analytical investigation of a simplified 2-cell case was performed. It is shown that in some cases robustness curves cross, which means that, depending on the demand, one might prefer a robust satisficing strategy which is different from an expected-utility maximizing one. Numerical results obtained by a deterministic search method confirm this for more complex 4-cell cases. The effects of different parameters of the problem on the resulting solution, including two different uncertainty models (local and global), were studied. Later on, a large-scale problem was solved using a stochastic numeric method (GA) after dealing with some encoding issues for the chosen system model. A study of 50 random scenarios revealed a typical consolidation of effort behavior (in high initial-probability cells) of the robust maximizing solutions, which becomes more distinct as the value of E_c is lowered. The behavior was explained using the local uncertainty model which was in use.

Similar to the Pareto principal, and the No Free Lunch principal in optimization, the increase in robustness comes at the expense of the expected utility. While the equilibrium point must be chosen by the mission planner, it is shown that, in various circumstances, one obtains results which are superior to the expected-utility maximizing strategy (in terms of robustness at a broad range of E_c values), while sacrificing almost no expected utility. In light of these results, the recommended strategy is robust-satisficing: satisfice the expected utility and maximize the robustness.

The presented method of decision-making might be useful in various scenarios, as it incorporates different types of data, such as a critical expected utility value, and uncertainty weighting parameters, reflecting different quality of information for different search cells and targets. That said, it is important to implement it on problems in which such data exists. When the prior data is known to be very reliable, one might rightfully choose to maximize the expected utility.

References

[1] M. Baum and K. Passino, "A search-theoretic approach to cooperative control for uninhabited air vehicles," in *Guidance, Navigation, and Control Conference and Exhibit*. AIAA, 2002.

[2] S. V. Spires and S. Y. Goldsmith, "Exhaustive geographic search with mobile robots along space-filling curves," in *International workshop on collective robotics, Paris*, vol. 1456, 1998, pp. 1–12.

[3] P. B. Sujit and D. Ghose, "Multiple agent search of an unknown environment using game theoretical models," in *American Control Conference*, vol. 6, 2004, pp. 5564–5569.

[4] M. Flint, E. Fernandez-Gaucherand, and M. Polycarpou, "Stochastic modeling of a cooperative autonomous uav search problem," *Military Operations Research*, 2003.

[5] M. M. Polycarpou, Y. Yang, and K. M. Passino, "A cooperative search framework for distributed agents," in *IEEE International Symposium on Intelligent Control*. IEEE, 2001, pp. 1–6.

[6] Y. Yang, M. M. Polycarpou, and A. A. Minai, "Opportunistically cooperative neural learning in mobile agents," in *International Joint Conference on Neural Networks*, vol. 3. IEEE, 2002, pp. 2638–2643.

[7] T. Shima, S. Rasmussen, A. Sparks, and K. M. Passino, "Multiple task assignment for cooperating uninhabited aerial vehicles using genetic algorithms," *Computers & Operations Research*, vol. 33, no. 11, pp. 3252–3269, 2005.

[8] T. Shima and S. Rasmussen, *UAV cooperative decision and control: challenges and practical approaches*. SIAM's Advances in Design and Control series, 2008.

[9] E. Edison and T. Shima, "Genetic algorithms for cooperative uav task assignment and path optimization," in *Guidance, Navigation and Control conference*. AIAA, 2008.

[10] L. F. Bertuccelli and J. P. How, "Robust uav search for environments with imprecise probability maps," in *44th IEEE Conference on Decision and Control, 2005 and 2005 European Control Conference (CDC-ECC '05)*. IEEE, 2005, pp. 5680–5685.

[11] Y. Yang, A. A. Minai, and M. M. Polycarpou, "Decentralized cooperative search by networked uavs in an uncertain environment," in *American Control Conference*, vol. 6, 2004, pp. 5558–5563.

[12] L. D. Stone, *Theory of optimal search*. New York: Academic Press, 1975.

[13] V. Boeva and B. De Baets, "A new approach to admissible alternatives in interval decision making," in *2nd International IEEE Conference on Intelligent Systems*, vol. 1. IEEE, 2004, pp. 110–115.

[14] P. Krokhmal, R. Murphey, P. Pardalos, S. Uryasev, and G. Zrazhevski, "Robust decision making: addressing uncertainties in distributions," Cooperative Control: Models, Applications, and Algorithms, 2003.

[15] G. Shafer, *A mathematical theory of evidence*. Princeton, New Jersey: Princeton Univ. Press, 1976.

[16] J. M. Keynes, *A treatise on probability*. London: Macmillan and Co., 1921.

[17] Y. Ben-Haim, "Uncertainty, probability and information-gaps," *Reliability Engineering and System Safety*, vol. 85, pp. 249–266, 2004.

[18] I. Sisso, T. Shima, and Y. Ben-Haim, "Info-gap approach to spatial search under severe uncertainty," in *Proceedings of the 49th Israel Annual Conference on Aerospace Sciences*, Tel-Aviv/Haifa, Israel, 2009.

[19] Y. Ben-Haim, *Info-gap decision theory: decisions under severe uncertainty*, 2^{nd} ed. Elsevier, 2006.

[20] ——, "Info-gap forecasting and the advantage of sub-optimal models," *European Journal of Operational Research*, no. 197, pp. 203–213, 2009.

[21] B. O. Koopman, "The theory of search. ii. target detection," *The Journal of the Operations Research Society of America*, vol. 4, pp. 503–531, 1956.

Itay Sisso Received his B.Sc., and M.Sc. degrees in Aerospace Engineering from the Technion - Israel Institute of Technology, in 2004, and 2009, respectively. During his M.Sc. Itay serviced in the Israeli Air Force as an officer in the special projects branch, and later on in the Air Force Battle Laboratory. He is now serving as a captain in the Air Force Operations Research department. Itay won the Hillel prize, and an excellence scholarship, for his M.Sc. thesis which led to this paper.

Tal Shima Received his B.Sc., MA, and Ph.D. degrees, all in Aerospace Engineering, from the Technion - Israel Institute of Technology, in 1992, 1998, and 2001, respectively. He also received the MBA degree from the Tel-Aviv University in 1997. After his Ph.D. Tal worked in the Air-to-Air and Air-Defense Directorate of RAFAEL ltd. in Israel. From 2004-2006 he was a National Research Council scholar at the US Air Force Research Labs, Wright-Patterson Air Force Base, performing research on UAV cooperative decision and control. From 2006 Dr. Shima is a senior lecturer with the Department of Aerospace Engineering at the Technion. His research interests include guidance and control of single and cooperating autonomous aerial vehicles. Dr. Shima has published 70 refereed papers in these research areas. He is an associate fellow of AIAA, and a senior member of IEEE.

Yakov Ben-Haim Initiated and developed info-gap decision theory for modelling and managing severe uncertainty. Info-gap theory is applied in engineering, biological conservation, economics, project management, homeland security, medicine, and other areas. He has been a visiting scholar in Australia, Canada, England, France, Germany, Italy, Japan, Korea, Netherlands, Norway, and the US. He has lectured at universities, medical and technological research institutions and central banks around the world. He has published more than 70 articles and 4 books. He is a professor of mechanical engineering and holds the Yitzhak Moda'i Chair in Technology and Economics at the Technion - Israel Institute of Technology.