D. R. Morrell and Y. Xue, “Analysis of and Heuristics for Sensor Configuration in a Simple Target Localization Problem,” *Conference Record of the 35th Asilomar Conference on Signals, Systems, and Computers,* November 2001, pp. 1391–1395.

Analysis of and Heuristics for Sensor Configuration in a Simple Target Localization Problem

Darryl Morrell and Ya Xue*

Department of Electrical Engineering
Arizona State University
E-mail: morrell@asu.edu, ya.xue@asu.edu

Abstract

We investigate Bayesian methods and heuristics for management of a configurable sensor in a simple target localization problem. A target is located in one of M cells. A sensor, characterized by probabilities of correct detection and false alarm, repeatedly chooses a cell to interrogate; the resulting observations are used to update the posterior probability distribution of target location. Interrogations are repeated either a fixed number of times or until the probability of error drops below a pre-selected threshold. The Bayes optimal solution is exponentially complex, motivating the use of heuristics. Four heuristic rules are characterized using Monte Carlo simulation. Of these heuristics, choosing the most probable cell minimizes the number of observations, and the myopic Bayes optimal rule minimizes the probability of error.

1 Introduction

The management of configurable sensors to optimize overall system performance is an important issue in target detection, localization, and tracking systems. Bayesian decision theory provides a mathematical framework by which the sensor management problem can be addressed. In this paper, we apply Bayesian decision theory to the problem of choosing a sequence of sensor controls for a steerable sensor in a simple target localization problem. Our goal is to find and characterize optimal and good sub-optimal rules for steering the sensor.

We consider the following simple abstract scenario. A target is known to be located in one of M cells. At each time epoch, a sensor interrogates a single cell and returns a binary-valued observation indicating whether the target was detected in the cell. Our objective is to correctly locate the target from a sequence of observations collected by the sensor.

We represent this scenario as a sequential Bayesian decision problem. Initial knowledge about the target location is represented by a prior probability distribution. The sensor is characterized by a probability of target detection and a probability of false alarm; these probabilities are independent of the cell that is interrogated. Observations are collected either a fixed number of times or until a given probability of error is reached. After each observation is obtained, the distribution of the target location is updated using Bayes theorem. The interrogated cell at a given time thus depends on previously collected observations.

We seek a decision rule that chooses which cell to interrogate in order to optimize the system performance. A *myopic optimal* decision rule is one that optimally chooses the cell given that a single observation will be collected. A *sequential optimal* decision rule is one that optimally chooses cells for the collection of an entire sequence of observations. We measure optimality by one of two criteria: probability of error given a fixed number of observations, or number of observations needed to guarantee a fixed probability of error.

Sequential optimal decision rules can be derived for this problem; unfortunately, their complexity increases exponentially with the length of the observation sequence. Thus, in this paper we develop myopic decision rules, but we evaluate the performance of these rules on sequences of observations.

This paper is structured as follows. Section 2 discusses the theoretical background of this problem and introduces related prior work. In Section 3, we give a mathematical formulation of the problem, and in Section 4, we derive the myopic optimal decision rule. We present several heuristics in Section 5, and evaluate their performance through Monte Carlo simulation in Section 6. We present conclusions and areas for future work in Section 7.

*This work supported by AFOSR under award number F49620-00-1-0124

2 Previous Work

The investigation of this simple sensor scheduling problem was motivated by a dual-mode sensor problem posed in [4, 9, 8]. In this work, a myopic Bayesian optimal decision rule to minimize an expected loss function was developed for a dual-mode sensor that could either interrogate a single cell or the entire search region. The decision rule was repeatedly applied until the probability of error was less than 0.05. The average number of observations needed to obtain this probability of error was characterized through Monte Carlo simulation for a few combinations of sensor performance and loss function parameters.

This work raised several questions. One is how the Bayesian myopic optimal decision rule, derived to minimize the expected loss, performs in terms of the number of observations needed to achieve a given probability of error. Another is the relative performance of repeated application of a myopic optimal decision rule compared to a sequential optimal decision rule.

Finding an optimal fixed length sequence of interrogations to minimize the probability of error can be approached from several perspectives, including Partially Observed Markov Decision Problems (POMDP)[2], dynamic programming[1], and discrete multistage decision networks[5, Chapter 8]. Each of these approaches result in solutions that are mathematically equivalent. Unfortunately, the complexity of solutions generally increases exponentially with the length of the sequence of observations.

Some previous work applies directly to this problem of minimizing the probability of error for a fixed number of observations. Dynamic programming has been used to find optimal observation sequences in the context of dynamic search strategies[3] and fault diagnosis[7]. It is shown that the myopic optimal strategy always chooses between the most and second most probable cells. When the probability of detection and false alarm are symmetric (i.e. the probabilities of false alarm and missed detection are equal), the myopic optimal strategy is to choose either the most or second most probable cell, and the sequential optimal strategy is to repeatedly apply the myopic optimal strategy[3].

Heuristic approaches to minimize the probability of error have also been suggested. In [7], a greedy approach is introduced, in which the myopic optimal strategy is repeatedly applied. In [6], a cell is chosen to maximize a discrimination function that is essentially the Kulback-Liebler distance between prior and posterior distributions.

We believe that the problem of finding a sequence of cell interrogations that minimizes the expected number of observation necessary to guarantee a given probability of correct decision can be formulated using dynamic programming. To our knowledge, it has not been solved.

3 Problem Formulation

We now formulate the sensor control problem mathematically. A target is located in one of M possible cells; the target location is a random variable denoted $X \in \{1, \ldots, M\}$. Our knowledge of the initial target location is represented by a prior probability distribution $p_X(x)$.

At each time epoch, the sensor can choose to interrogate one cell; the cell chosen is denoted $d \in \{1, \ldots, M\}$. This interrogation results in a single observation consisting of a detection or non-detection. The observation is denoted $Y \in \{0, 1\}$; a value of 1 indicates that the target was detected, while a value of 0 indicates that the target was not detected. The sensor performance is known and characterized by a probability of correct detection p_D and false alarm p_F. For notational convenience, we define $\overline{p_D} = 1 - p_D$ and $\overline{p_F} = 1 - p_F$. The probability distribution of Y conditioned on X and d is

$$p_{Y|X,d}(y|x,d) = \begin{cases} p_D, & y = 1 \text{ and } x = d \\ \overline{p_D}, & y = 0 \text{ and } x = d \\ p_F, & y = 1 \text{ and } x \neq d \\ \overline{p_F}, & y = 0 \text{ and } x \neq d \end{cases} \quad (1)$$

The observed value y is used to compute the posterior probability distribution $p_{X|Y,d}(x|y,d)$ of the target location via Bayes theorem:

$$p_{X|Y,d}(x|y,d) = \frac{p_{Y|X,d}(y|x,d)p_X(x)}{\sum_{x'} p_{Y|X,d}(y|x',d)p_X(x')}$$

The process of choosing a cell, obtaining an observation from this cell, and updating the probability distribution is repeated either a fixed number of times or until the posterior probability of the target being located in a particular cell given all observations exceeds a threshold close to one. At this point, the estimated target location is the cell whose posterior probability is largest; this cell is denoted $\hat{x}$.

4 The Myopic Optimal Rule

We derive the myopic optimal rule to choose the cell to interrogate to minimize the probability of error; the optimal cell is denoted d_0. This derivation leads to a decision rule that is mathematically equivalent to those obtained in [3, 7]. It can be easily extended to a more complex dual-mode sensor problem similar to that in [9, 8]. This derivation may also lead to significant reductions in the cost of computing sequential optimal decision rules (although the growth in their complexity with the sequence length is still exponential), which may lead to feasible Monte Carlo evaluation of the sequential optimal rule for moderate sequence lengths and small values of M.

To obtain a minimum probability of error decision rule, we use an objective function $l(x,\hat{x})$ that rewards correct decisions and punishes errors equally:

$$l(x,\hat{x}) = \begin{cases} 1, & x = \hat{x} \\ 0, & x \neq \hat{x} \end{cases} \quad (2)$$

We obtained the following Bayes optimal myopic decision rule using the decision network approach in [5, Chapter 8]; this rule could also be obtained from a straight-forward application of basic Bayesian decision theory (c.f. [8]):

$$d_0 = \arg\max_d \sum_{y=0}^{1} \max_{\hat{x}} \sum_{x=1}^{M} p_{Y|X,d}(y|x,d) p_X(x) l(x,\hat{x}) \quad (3)$$

This decision rule could be implemented directly to determine d_0; however, when M is large, the number of computations is prohibitive.

We can significantly simplify (3). From the structure of $l(x,\hat{x})$,

$$d_0 = \arg\max_d \sum_{y=0}^{1} \max_{\hat{x}} p_{Y|X,d}(y|\hat{x},d) p_X(\hat{x}) \quad (4)$$

For notational convenience, we assume without loss of generality that $p_X(1) \geq p_X(2) \geq \ldots \geq p_X(M)$. We first consider maximizing $p_{Y|X,d}(y|\hat{x},d)p_X(\hat{x})$ over $\hat{x}$ for a given value of d and $y = 0$:

$$p_{Y|X,d}(0|\hat{x},d) p_X(\hat{x}) = \begin{cases} \overline{p_D}\; p_X(\hat{x}), & \hat{x} = d \\ \overline{p_F}\; p_X(\hat{x}), & \hat{x} \neq d \end{cases} \quad (5)$$

Since $p_X(1) \geq p_X(2) \geq \ldots \geq p_X(M)$,

$$\max_{\hat{x} \neq d} p_{Y|X,d}(0|\hat{x},d) p_X(\hat{x}) = \begin{cases} \overline{p_F}\; p_X(2), & d = 1 \\ \overline{p_F}\; p_X(1), & d \neq 1 \end{cases} \quad (6)$$

Combining (5) and (6) gives

$$\begin{aligned} &\max_{\hat{x}} p_{Y|X,d}(0|\hat{x},d) p_X(\hat{x}) \\ &= \begin{cases} \max\left[\overline{p_F}\; p_X(2), \overline{p_D}\; p_X(1)\right], & d = 1 \\ \max\left[\overline{p_F}\; p_X(1), \overline{p_D}\; p_X(d)\right], & d \neq 1 \end{cases} \end{aligned} \quad (7)$$

When $y = 1$, a similar line of reasoning gives

$$\begin{aligned} &\max_{\hat{x}} p_{Y|X,d}(1|\hat{x},d) p_X(\hat{x}) \\ &= \begin{cases} \max\left[p_F\; p_X(2), p_D\; p_X(1)\right], & d = 1 \\ \max\left[p_F\; p_X(1), p_D\; p_X(d)\right], & d \neq 1 \end{cases} \end{aligned} \quad (8)$$

Substituting (7) and (8) into (4) gives the following:

$$d_0 = \arg\max_d \begin{cases} \max\left[\overline{p_F}\; p_X(2), \overline{p_D}\; p_X(1)\right] + \\ \max\left[p_F\; p_X(2), p_D\; p_X(1)\right], & d = 1 \\ \max\left[\overline{p_F}\; p_X(1), \overline{p_D}\; p_X(d)\right] + \\ \max\left[p_F\; p_X(1), p_D\; p_X(d)\right], & d \neq 1 \end{cases} \quad (9)$$

Consider the term in (9) for $d \neq 1$. Since $p_X(2) \geq \ldots \geq p_X(M)$, this term will always be maximized by $d = 2$. Note that this implies that d_0 will always be either 1 or 2; the optimal cell to interrogate is always the most probable or the second most probable. We can use this with (9) to get the following decision rule.

Myopic Optimal Decision Rule: Choose $d_0 = 2$ if

$$\begin{aligned} &\max\left[\overline{p_F}\; p_X(1), \overline{p_D}\; p_X(2)\right] \\ &+ \max\left[p_F\; p_X(1), p_D\; p_X(2)\right] \\ &\quad > \max\left[\overline{p_F}\; p_X(2), \overline{p_D}\; p_X(1)\right] \\ &\quad\quad + \max\left[p_F\; p_X(2), p_D\; p_X(1)\right] \end{aligned}$$

Otherwise, choose $d_0 = 1$

5 Application of Heuristics

The form of the myopic optimal decision rule suggests the following three heuristics that could be applied to obtain a cell to interrogate:

1. choose the most probable cell,

2. choose the second most probable cell,

3. choose the myopic optimal cell.

A fourth heuristic has been suggested to address a similar sensor allocation problem[6]. In this heuristic, the cell is chosen to maximize an expected discrimination gain between the posterior and prior distributions on target location. The expected discrimination gain $\Delta D(d)$ is defined as

$$\Delta D(d) = E_Y \left\{ D\left[p_{X|Y,d}(x|Y,d), p_X(x)\right] \right\}$$

where $D\left[p_{X|Y,d}(x|Y,d), p_X(x)\right]$ is the Kulbach-Liebler distance between $p_X(x)$ and $p_{X|Y,d}(x|Y,d)$.

Note that all of these heuristics are myopic-they choose only the next cell to interrogate. We apply them to the problem of choosing a sequence of cells by repeatedly using the myopic heuristic to choose a cell, obtaining the observation from that cell, and then computing the posterior distribution of the target position.

6 Simulation Results

The performance of these four heuristics has been evaluated using Monte Carlo simulations. In these simulations, we investigated a number of combinations of p_F and p_D. The relative performance of the different heuristics over the range of investigated p_D values is similar; thus, we only present simulation results for values of $p_D = 0.9$. We also investigated several different prior distributions for X. The

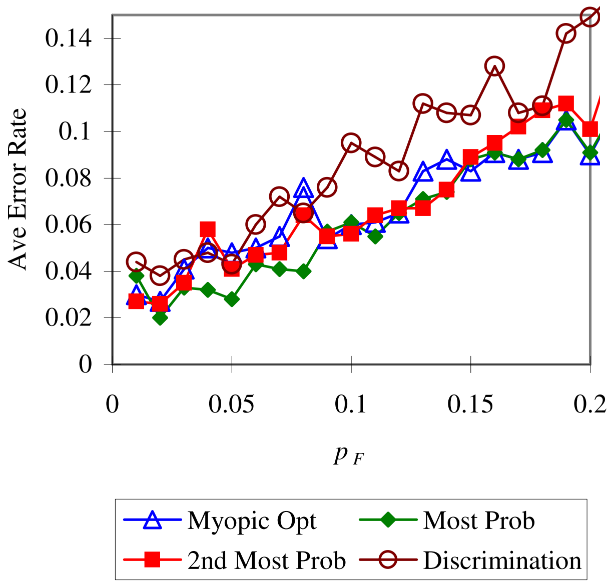

Figure 1. Average probability of error as a function of p_F for $p_D = 0.9$, $M = 10$, and fifteen observations (1000 simulation runs).

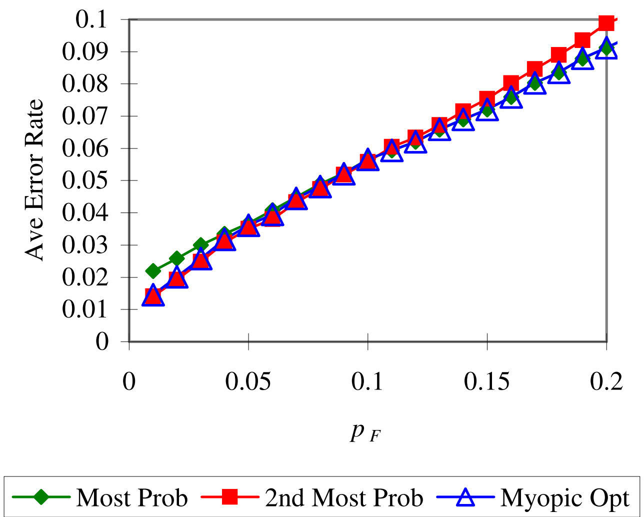

Figure 2. Average probability of error as a function of p_F for $p_D = 0.9$, $M = 10$, and fifteen observations (500,000 simulation runs).

results shown below are for distributions in which the probability of the target being in cell x was proportional to x; again, the relative performance of the heuristics for other initial distributions is similar. Finally, we also investigated the effect of the number of cells on the relative performance of heuristics, and found no difference between small problems ($M = 10$ cells) and large problems ($M = 100$ cells).

Figure 1 shows the average error rate for the four heuristics as a function of p_F with $p_D = 0.9$, $M = 10$, and fifteen observations. This error rate is averaged over 1000 Monte Carlo simulation runs. Note that the discrimination gain heuristic has a higher error rate than the other heuristics. The relative performance of the other heuristics cannot be determined from these results due to the small number of Monte Carlo runs; the differences in performance between the three heuristics is smaller than the 95% confidence interval for these values.

To more accurately characterize the performance of the three better performing heuristics, we repeated the Monte Carlo simulations with 500,000 trials for the most probable, second most probable, and myopic optimal heuristics. We did not continue investigating the discrimination gain heuristic because of its poorer initial performance and its computational complexity. Figure 2 shows the simulation results obtained. Note that the myopic optimal rule gives the lowest probability of error for all values of p_F. Choosing the most probable cell performs as well as the myopic optimal rule for values of $p_F > 1 - p_D$. This is because the myopic optimal rule usually chooses the most probable cell when $p_F > 1 - p_D$. Similarly, choosing the second most probable cell performs about as well as the myopic optimal for values of $p_F < 1 - p_D$, since the myopic optimal rule usually chooses the second most probable cell when $p_F < 1 - p_D$. Choosing the most probable or second most probable cell is computationally less complex than choosing the myopic optimal cell.

We also ran Monte Carlo simulations to determine which heuristic required fewer observations to achieve a probability of error less than 0.05. Figure 3 shows the number of observations averaged over 1000 Monte Carlo runs for the four heuristics. In these simulations, we capped the number of observations at 30; thus, this plot gives an optimistic assessment of the heuristics' performance. The discrimination gain is not the best performing heuristic; because of its computational complexity, it was not considered further.

Figure 4 shows the number of observations averaged over 500,000 Monte Carlo runs for the three heuristics. In this simulation, we did not cap the number of observations. From the graph, it is clear that choosing the most probable cell is the best heuristic. Again, the myopic optimal rule performs as well as choosing the most probable cell when $p_F > 1 - p_D$ because both rules have essentially the same behavior for these values of p_F. Since choosing the most probable cell has a lower computational cost and gives

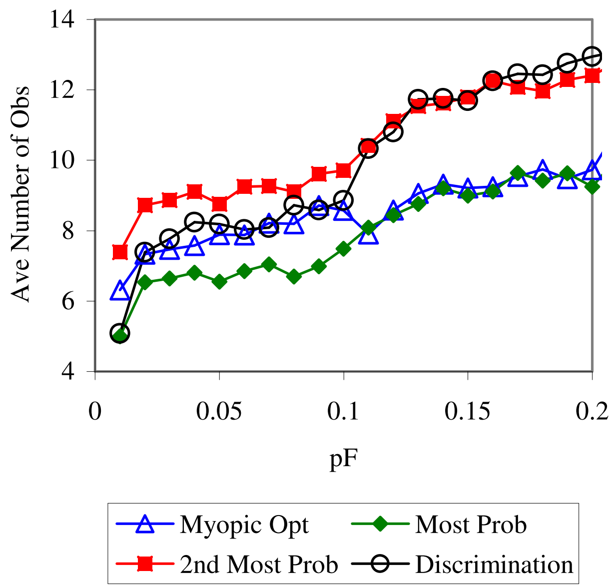

Figure 3. Average number of observations needed to achieve a probability of error less than 0.05 as a function of p_F for $p_D = 0.9$ and $M = 10$ (1000 simulation runs).

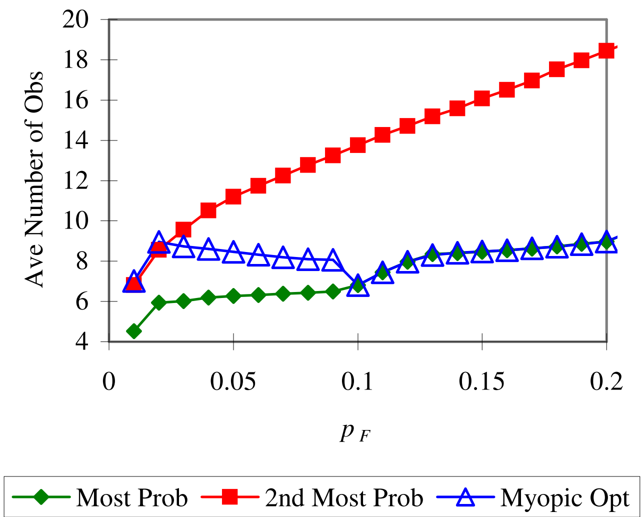

Figure 4. Average number of observations needed to achieve a probability of error less than 0.05 as a function of p_F for $p_D = 0.9$ and $M = 10$ (500,000 simulation runs).

the best performance, it is clearly the best heuristic for this problem.

7 Conclusions

We have formulated a simple sensor control problem and investigated the performance of several myopic control strategies. We have found that the myopic optimal rule results in the lowest average error rate, and that choosing the most probable cell results in the lowest average number of observations.

One important unanswered question is the performance penalty (in terms of probability of error or number of needed observations) incurred by a myopic algorithm compared to an optimal sequential algorithm. For the problem of a fixed number of observations, we conjecture that the performance penalty is not large, particularly when p_D is close in value to p_F. This is an area of on-going research.

References

[1] D. Bertsekas. *Dynamic programming : Deterministic and Stochastic Models*. Prentice-Hall, 1987.

[2] A. Cassandra. *Exact and Approximate Algorithms for Partially Observable Markov Decision Processes*. PhD thesis, Brown University, 1998.

[3] D. Castanon. Optimal search strategies in dynamic hypothesis testing. *IEEE Trans. on Systems, Man, and Cybernetics*, 25(7):1130–1138, July 1995.

[4] D. Cochran, D. Sinno, and A. Clausen. Source detection and localization using a multi-mode detector: a bayesian approach. In *Proceedings of ICASSP99*, volume 3, pages 1173–1176, Mar. 1999.

[5] R. Cowell, A. Dawid, S. Lauritzen, and D. Spiegelhalter. *Probabilistic Networks and Expert Systems*. Springer, 1999.

[6] K. Kastella. Discrimination gain to optimize detection and classification. *IEEE Trans. on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 27(1):112–116, Jan. 1997.

[7] V. Raghavan, M. Shakeri, and K. Pattipati. Test sequencing algorithms with unreliable tests. *IEEE Trans. on Systems, Man, and Cybernetics-Part A: Systems and Human*, 29(4):347–357, July 1999.

[8] D. Sinno. *Attentive Management of Configurable Sensor Systems*. PhD thesis, Arizona State University, May 2000.

[9] D. Sinno, D. Cochran, and D. Morrell. A Bayesian risk approach to multi-mode detection. In *U.S./Australia Joint Workshop on Defense Signal Processing*, Chicago, Aug. 1999.