A Bayesian Risk Approach to Multi-Mode Detection

Dana Sinno, Douglas Cochran, and Darryl R. Morrell
Arizona State University, Tempe, AZ 85287-7206, U.S.A.

This paper considers a situation in which a detection system is configurable in such a way as to provide multiple modes of operation that differ in their detection performance and geographical coverage. A technique for optimal mode selection based upon minimizing Bayesian risk is formulated and demonstrated for the case of a two-mode system.

0. PROLOGUE

In the first Joint Australia/U.S. Workshop on Defence Signal Processing, held in Adelaide in 1997, two of the authors (Sinno and Cochran) presented a paper involving estimation using a configurable sensor system [6]. During the Workshop, Dr. Paul Miller of the Australian Defence Science and Technology Organisation told us about a real-world scenario in which searches for ground vehicles are carried out over vast uninhabited areas by helicopters outfitted with dual-mode radar systems. We understood the operating modes of the radar system to be such that they could be loosely described as "broad search" and "focused" modes and that the strategy for switching between modes during a search was left to the helicopter crew.

We subsequently proposed a mathematical model of this kind of scenario and a Bayesian approach to choosing mode switching strategies [1]. This formulation made use of a payoff function consisting of two terms, one of which captures the performance of the sensing strategy in detecting the presence of a target in the search area and the other measuring its effectiveness at localizing the target.

When we were invited to participate in the second Australia/U.S. Workshop, revisiting this problem and exploring an alternative approach that addresses a shortcoming of our technique in [1] seemed a natural way to connect our contribution with the previous Workshop and the many fine technical interactions it seeded. The approach in this paper is based upon Bayesian risk analysis and it eliminates concerns regarding correlation between the terms of the payoff function arising in our previous treatment of the problem.

1. INTRODUCTION

This paper considers a situation in which a detection system is configurable in such a way as to provide multiple modes of operation that differ in their detection performance and geographical coverage. The development that follows focuses on the case of a detector with two operating modes: a "broad search" mode that provides wide coverage and a "focused" mode that provides better detection performance but covers less area. The system is invoked in a sequence of tests to detect and localize a target within a framework that is formulated precisely in the following section.

As noted in the Prologue, this problem was addressed in an earlier paper [1] using a cost functional approach in conjunction with a Bayesian method for incorporating the results of earlier tests in deciding which mode to use in each test. The approach presented here uses a more classical approach, based on minimization of Bayesian risk, that allows more precise designation of the priority of correct detection relative to that of correct localization.

2. PROBLEM SETUP

The situation described above is modeled as follows. The entire region of interest C is partitioned into N disjoint cells $C_1, ..., C_N$. Operating in the broad search mode (*Mode A*), the detector tests for the presence of a signal source in C. In the focused mode (*Mode B*), however, the test may be applied to exactly one cell C_n.

To account for difference in detector performance in the two operating modes, detector performance is modeled as arising from the problem of detecting a known signal in white gaussian noise of known variance. This model provides a well understood solution (i.e., the matched filter) in each test, admits several straightforward generalizations, and allows detection performance in *Mode B* to be distinguished from that in *Mode A* by simply raising the signal-to-noise ratio (SNR). More specifically, in each mode of operation the detector encounters a problem of the form

$$\begin{aligned} H_0 &: \mathbf{X} = \mathbf{N} \\ H_1 &: \mathbf{X} = S + \mathbf{N} \end{aligned} \tag{1}$$

where S is a known signal M-vector with energy $||S||^2 = 1$ and $\mathbf{N}$ is a zero-mean white gaussian M-vector having known variance σ^2; i.e., $\mathbf{N} \sim \mathcal{N}[\mathbf{0}, \sigma^2 \mathbf{I}]$ where $\mathbf{I}$ is the $n \times n$ identity matrix. Since $||S||$ is fixed, the SNR (and hence the performance of the detector) in each mode can be adjusted by varying σ^2.

Assuming at most one signal source is present, denote by H_1 and H_0 the events that the signal source is, respectively, present in and absent from C. Let $h_0 = H_0$ and, for $n = 1, ..., N$, denote by h_n the event that the signal source is present in cell C_n. With these definitions, $H_1 = \cup_{n=1}^{N} h_n$. Regardless of whether it is operating in *Mode A* or *Mode B*, the system yields a decision $\rightarrow h_n$ with $0 \leq n \leq N$.

Recall that the optimal solution, in terms of minimal probability of error, to a detector problem of the form (1) is a test on the inner product $S^T \mathbf{X}$ where the detection threshold is a function of the *a priori* probability that a signal is present [2, 5]. The probabilities of detection and false alarm for each test are given by error functions of the detection thresholds. In particular, the tests applied in both operating modes will be of this form, but their detection thresholds and probabilities of detection and false alarm will all be different (even when *Mode B* is applied to different cells) because of their dependence on $\Pr(h_n)$, $n = 0, ..., N$.

3. A BAYESIAN RISK FORMULATION

Using the notation of [5], define a random "state of nature" parameter $\boldsymbol{\theta}$ by $\boldsymbol{\theta} = n$ if h_n is true, $n = 0, 1, ..., N$. A prior distribution for $\boldsymbol{\theta}$ is assumed and a test (i.e., a *Mode A* test or a *Mode B* test on a particular cell C_n) is chosen and performed yielding a binary

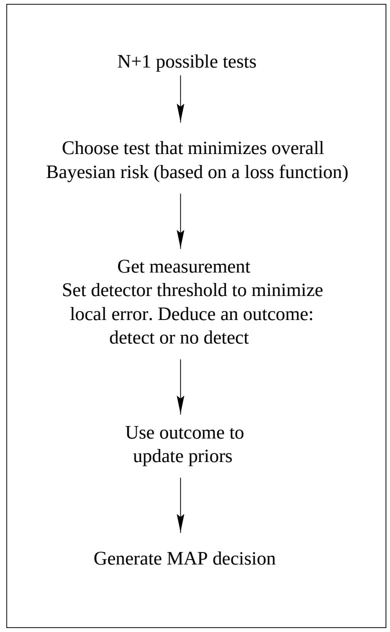

FIG. 1. Detection/localization algorithm.

outcome $\mathbf{o}$. If $\Pr(H_0|\mathbf{o}) > \Pr(H_1|\mathbf{o})$, the system decides for H_0. Otherwise, the system decides in favor of the hypothesis h_n having the largest posterior probability $\Pr(h_n|\mathbf{o})$; i.e., in this case the system decision rule ϕ takes the value n if h_n has the largest posterior probability. As shown in [1], these posterior probabilities are straightforward to compute using the detection and false alarm probabilities of the chosen test, which follow from the (prior) distribution of $\boldsymbol{\theta}$. The overall algorithm is depicted schematically in FIG. 1. Note that, once a test is selected, the rule ϕ for choosing a hypothesis h_n based on the test's outcome is well defined.

The approach to mode (i.e., test) selection is to choose the one that minimizes Bayes risk with respect to a pre-defined loss functional. Since the overall goal of the system is to both detect the signal source *and* localize it, with these two subgoals possibly being of unequal importance, a loss functional of the following form is used:

$$L(\theta, \phi) = \begin{cases} 0 & \theta = \phi \\ 1 & \theta \neq \phi,\ \phi \neq 0, \text{ and } \theta \neq 0 \\ c_1 & \theta \neq \phi \text{ and } \theta = 0 \\ c_2 & \theta \neq \phi \text{ and } \phi = 0 \end{cases}$$

With $c_1 > 1$, this functional imposes a greater penalty for a false alarm (i.e., deciding in favor of H_1 when H_0 is true) than for correct detection with incorrect localization (i.e., H_1 is correctly chosen, but the wrong cell is picked). With $c_2 > c_1$, an even greater penalty is levied if the system decides in favor of H_0 when a target is actually present. Depending on the application, the weights c_1 and c_2 can be chosen to adjust the relative importance of detection and classification in an intuitively appealing way.

With this loss functional, the risk is

$$R(\boldsymbol{\theta}, \phi) = \begin{cases} c_1 \sum_{n \neq 0} \Pr(\rightarrow h_n | h_0) & \boldsymbol{\theta} = 0 \\ c_2 \Pr(\rightarrow h_0 | h_k) + \sum_{n=1, n \neq k}^{N} \Pr(\rightarrow h_n | h_k) & \boldsymbol{\theta} \neq 0 \end{cases}$$

and the Bayes average risk of the decision rule ϕ is thus

$$\mathsf{E}[R(\boldsymbol{\theta}, \phi)|\boldsymbol{\theta}] = c_1 \Pr(h_0) \sum_{n=1}^{N} \Pr(\rightarrow h_n | h_0)$$
$$+ \sum_{n=1}^{N} \Pr(h_n) \left[c_2 \Pr(\rightarrow h_0 | h_n) + \sum_{k=1, k \neq n}^{N} \Pr(\rightarrow h_k | h_n) \right]$$

Since the decision rule ϕ (and hence each decision $\rightarrow h_n$) depends on the probabilities of the hypotheses h_k posterior to the test, this quantity depends on which test (mode) is selected.

The mode is chosen to minimize the conditional expectation of the Bayes average risk. The posterior probabilities of each h_n given a particular test and its outcome can be calculated using Bayes' rule, the prior probabilities of the hypotheses, and the detectors' probabilities of detection and false alarm. These calculations are given explicitly in [1]. Once a mode is selected and the test T_m is performed, the decision $\rightarrow h_n$ is completely determined by its outcome $o_m \in \{0, 1\}$; prior to performing the test, the only uncertainty about the decision arises because the outcome of the test is not yet known. The conditional expectation of the Bayes average risk for a chosen test T_m is $R_0 \Pr(\mathbf{o}_m = 0) + R_1 \Pr(\mathbf{o}_m = 1)$ where

$$R_0 = \mathsf{E}[R(\boldsymbol{\theta}, \phi)|\boldsymbol{\theta}, \mathbf{o}_m = 0] = c_1 \Pr(h_0 | \mathbf{o}_m = 0) \sum_{n=1}^{N} \Pr(\rightarrow h_n | h_0, \mathbf{o}_m = 0)$$
$$+ \sum_{n=1}^{N} \Pr(h_n | \mathbf{o}_m = 0) \left[c_2 \Pr(\rightarrow h_0 | h_n, \mathbf{o}_m = 0) + \sum_{k=1, k \neq n}^{N} \Pr(\rightarrow h_k | h_n, \mathbf{o}_m = 0) \right]$$

and

$$R_1 = \mathsf{E}[R(\boldsymbol{\theta}, \phi)|\boldsymbol{\theta}, \mathbf{o}_m = 1] = c_1 \Pr(h_0 | \mathbf{o}_m = 1) \sum_{n=1}^{N} \Pr(\rightarrow h_n | h_0, \mathbf{o}_m = 1)$$
$$+ \sum_{n=1}^{N} \Pr(h_n | \mathbf{o}_m = 1) \left[c_2 \Pr(\rightarrow h_0 | h_n, \mathbf{o}_m = 1) + \sum_{k=1, k \neq n}^{N} \Pr(\rightarrow h_k | h_n, \mathbf{o}_m = 1) \right]$$

The complexity of these expressions for R_0 and R_1 belie their relatively simple nature. There are two cases:

- **Case 1**: $\phi = 0$

$$\begin{aligned} R_0 &= c_2 \Pr(H_1 | \mathbf{o}_m = 0) \\ R_1 &= c_2 \Pr(H_1 | \mathbf{o}_m = 1) \end{aligned}$$

- **Case 2**: $\phi = k > 0$

$$\begin{aligned} R_0 &= c_1 \Pr(H_0|\mathbf{o}_m = 0) + \Pr(H_1|\mathbf{o}_m = 0)[\Pr(H_1|\mathbf{o}_m = 0) - \Pr(h_k|\mathbf{o}_m = 0)] \\ R_1 &= c_1 \Pr(H_0|\mathbf{o}_m = 1) + \Pr(H_1|\mathbf{o}_m = 1)[\Pr(H_1|\mathbf{o}_m = 1) - \Pr(h_k|\mathbf{o}_m = 1)] \end{aligned}$$

The conditional probabilities in these expressions are exactly the post-test probabilities computed in [1]; the probabilities of the test outcomes are computed as follows. For a *Mode A* test,

$$\begin{aligned} \Pr(\mathbf{o}_n = 0) &= (1 - P_{d,A}) \Pr(H_1) + (1 - P_{f,A}) \\ \Pr(\mathbf{o}_n = 1) &= P_{d,A} \Pr(H_1) + P_{f,A} \Pr(H_0) \end{aligned}$$

and for a *Mode B* test,

$$\begin{aligned} \Pr(\mathbf{o}_n = 0) &= (1 - P_{d,B}^{(n)}) \Pr(h_n) + (1 - P_{f,B}^{(n)})(1 - \Pr(h_n)) \\ \Pr(\mathbf{o}_n = 1) &= P_{d,B}^{(n)} \Pr(h_n) + P_{f,B}^{(n)}(1 - \Pr(h_n)) \end{aligned}$$

In these expressions, $P_{d,B}^{(n)}$ and $P_{f,B}^{(n)}$ are the probabilities of detection and false alarm, respectively, of the *Mode B* detector used on cell n. $P_{d,A}$ and $P_{f,A}$ are the corresponding probabilities for the *Mode A* detector.

To summarize, the decision rule ϕ depends on the outcome of the test and the posterior probabilities of h_n, $n = 1, ..., N$. These can be computed *before any test is actually run*. Thus, for each candidate test, the expected risk may be calculated using $\Pr(\mathbf{o}_m = 0)$ and $\Pr(\mathbf{o}_m = 1)$ (which come from the detector performance figures) before running any tests. This allows the selection of the test of lowest Bayes risk, as proposed.

4. EXAMPLES

The following two examples show the behavior of the two-mode detection/localization system operating in a five-cell (i.e., $N = 5$) scenario. The test signal and white gaussian noise vectors are of length $M = 10$ and the SNRs in the two modes are -3.1 dB and -6.1 dB. The cost values are $c_1 = 1.2$ and $c_2 = 2$. In the first example (FIG. 2), the initial prior probabilities are $\Pr(h_1) = .1472$, $\Pr(h_2) = .0749$, $\Pr(h_3) = .0935$, $\Pr(h_4) = .1178$, and $\Pr(h_5) = .0667$. The posterior probabilities of the first test which are used as the prior probabilities in the second test, appear in the first column of the grid – and so forth. In this example, a signal source is actually present in cell 4 (indicated by a triangle in the upper right corner). The system chooses *Mode A* for the initial test (indicated by shading of the cells in the first column), does not detect (per the annotation beneath the column), and decides for H_0 (indicated by lack of highlighted frame around any cell). *Mode A* is selected again in the second test, the system detects and chooses cell 1. Following test 2, the system runs in *Mode B* on cell 1, does not detect but decides for cell 4 because it has the highest posterior (and $p_1 > 0.5$). In test 4, the detector runs again in *Mode B* but on cell 4, detects and decides for cell 4.

In the second example (FIG. 3), the initial prior probabilities are $\Pr(h_1) = .2385$, $\Pr(h_2) = .1006$, $\Pr(h_3) = .1315$, $\Pr(h_4) = .0239$, and $\Pr(h_5) = .0056$. In this example, a signal source is actually present in cell 3. It is interesting to note how the system switches from *Mode B* back to *Mode A* in test 4.

	Test 1 outcome 0	Test 2 outcome 1	Test 3 outcome 0	Test 4 outcome 1
Cell 5	.0179	.1047	.1313	.0151
Cell 4	.0316	.1849	.2320	.9119
Cell 3	.0251	.1468	.1841	.0211
Cell 2	.0201	.1176	.1475	.0169
Cell 1	.0395	.2312	.0353	.0041

FIG. 2. Behavior of the two-mode detector in a five-cell scenario (Example 1).

	Test 1 outcome 0	Test 2 outcome 1	Test 3 outcome 0	Test 4 outcome 1	Test 5 outcome 1
Cell 5	.0015	.0088	.0133	.0178	.0020
Cell 4	.0064	.0375	.0571	.0761	.0087
Cell 3	.0353	.2064	.3142	.4190	.9337
Cell 2	.0270	.1579	.2403	.3205	.0366
Cell 1	.0640	.3745	.0479	.0639	.0073

FIG. 3. Behavior of the two-mode detector in a five-cell scenario (Example 2).

5. DISCUSSION AND CONCLUSIONS

Since beginning this work, the authors have become aware of some fine research on related problems involving mode-switchable sensors, most notably by K. Kastella and his colleagues (see, e.g., [3, 4]).

Work currently underway is examining the mean time to correct decision of the approach presented here for various operating parameters, choice of detection thresholds for the individual tests to minimize this mean time, and possible applications outside the context of the original problem.

REFERENCES

1. D. Cochran, D. Sinno, and A. Clausen, "Source Detection and Localization Using a Multi-Mode Detector: A Bayesian Approach," *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Phoenix, March 1999.
2. C.W. Helstrom, *Elements of Signal Detection and Estimation.* Prentice-Hall, 1995.
3. K. Kastella, "Discrimination Gain to Optimize Detection and Classification," *IEEE Transactions on Systems, Man, and Cybernetics – Part A*, vol. 27, no. 1, January 1997.
4. W. Schmaedeke and K. Kastella, "Event-Averaged Maximum Likelihood Estimation and Information Based Sensor Management," *Proceedings of SPIE*, vol. 2232, pp. 91–96, 1994.
5. L.L. Scharf, *Statistical Signal Processing.* Addison Wesley, 1991.
6. D. Sinno and D. Cochran, "Estimation with Configurable and Constrained Sensor Systems," *Signal Processing* (to appear).