

SENSOR SELECTION IN ADVERSARIAL SETTING

Emre Ertin

Department of Electrical and Computer Engineering
The Ohio State University
Columbus, OH 43210, USA

ABSTRACT

We consider the problem of sensor selection for a binary hypothesis testing problem when the conditional density of the sensor readings can be affected by an adversary. A typical application of the proposed setup is surveillance with spatially distributed sensors, where the adversary is changing locations to evade detection. We consider a zero-sum game model where the primary and the adversary are choosing open loop randomized strategies and the payoffs are specified by the asymptotic detection probability under a false alarm constraint. We prove the existence of the Nash equilibrium of this surveillance game and characterize the optimal min-max strategies and the value of the game. A computed example of decentralized detection with sensors providing binary valued observations is given to illustrate the results.

Index Terms— Decision Theory, Game Theory, Sensor Management

1. INTRODUCTION

Distributed sensor systems use a multitude of sensors to obtain information for making inferences about the scene under observation. Constraints on power, communication bandwidth, computational complexity results can limit the number of sensors that can be activated at any time instance. Sensor management problems concern selection of active sensors for optimal detection and estimation. In particular, various information metrics have been used to guide sensor selection strategies [1, 2]. In [3, 4] mutual information criteria were proposed to select informative measurements in a greedy fashion one sensor at a time. In [5] this approach was extended to the general finite horizon case through an approximate dynamic programming method. Alternative approaches to sensor management include stochastic optimization for error covariance [6], geometric sensor selection schemes for bounded error sensor models [7] and convex optimization heuristics for error covariance minimization [8]. The use of open control randomized strategies have been suggested for sequential hypothesis testing with sensors [9]. Typically, in

all these previous work the target under the surveillance is modeled as a random process leading to a decision problem where the observer takes an expectation of the performance metric to optimize sensor selection.

In this work, we consider an intelligent adversary that is aware of the sensor characteristics and can affect the conditional density of the sensor observations. A typical application of the proposed setup is surveillance with spatially distributed sensors, where the adversary is changing locations to evade detection. We model the sensor selection problem as a game between two players with opposing objectives. The observer is choosing an open loop randomized strategy to choose sensor observations that maximizes probability of detection, whereas the target is using an open loop randomized control strategy over the available evading actions to minimize the probability of being detected. Figure 1 depicts the surveillance game between the observer choosing sensors and the target choosing evading actions.




Fig. 1. Surveillance game between the observer choosing sensors (k) and the target choosing evading actions (m)

The Nash equilibrium of this zero-sum game provides optimal strategies for surveillance and evasion and the value of the game quantifies the guaranteed performance of the surveillance system. In the next section, we present the system model and define the zero sum game between the observer and the target. In Section 3, we characterize the Nash equilibrium of this game and prove its existence. We con-

This work is in part sponsored by ARO MURI W911NF-11-1-0391 and DARPA/ONR grant N66001-10-1-4090.

clude with a computed example of distributed detection using two sensors with binary observations and a target with two evading actions. For brevity, we only provide only outlines of the proofs of each result.

2. SYSTEM MODEL

We assume a binary detection problem with two hypotheses H_0 and H_1 . We assume there are K sensors with observations drawn from a finite measurement space \mathcal{X} with J elements $\mathcal{X} = \{x_1, x_2, \dots, x_J\}$. The adversarial target has M actions. If H_1 is true, q_j^{km} gives the probability that x_j will be the observation of the k 'th sensor under the m 'th action of the target; if H_0 is true then p_j^k gives the probability that x_j will be the observation of the k 'th sensor. We assume that two technical conditions holds true. First, $q_j^{km} \neq 0$ and $p_j^k \neq 0$ for all j, k and m ; and second $\|\sum_m \theta_m q^{km} - p^k\| > 0$ for all k . The first condition guarantees that the sensors have the same support in \mathcal{X} and no sensor observation can decide on H_0 or H_1 with certainty. The second condition states that by randomizing its actions the target cannot identically create sensor density under H_0 .

We consider an observer using an open control randomized strategy $r = \{r_1, \dots, r_K\}$ with $\sum_k r_k = 1$, where r_k is the probability of choosing the k 'th sensor. The observer chooses a sensor for each observation independently, according to the probability law r to obtain N measurements $y = \{y_1, \dots, y_N\}$. Similarly we assume the target is employing an open control randomized strategy $s = \{s_1, \dots, s_M\}$ with $\sum_m s_m = 1$, where s_i gives the probability of choosing the m 'th action. The target chooses an action for each observation independently according to the probability law s .

The observer uses the measurement vector y collected under the sensor index set $I = \{i_1, \dots, i_N\}$ to make a decision between H_0 and H_1 . The decision rule is specified by disjoint sets \mathcal{U}_0 and \mathcal{U}_1 . If the tuple (y, I) is an element of \mathcal{U}_i then hypothesis H_i is chosen. Type I error $P_N^{FA}(r, s)$ and type II error $P_N^M(r, s)$ are given by:

$$P_N^M(r, s) = \sum_{(y, I) \in \mathcal{U}_0} \prod_{n=1}^N r_{i_n} \sum_m s_m q_{y_n}^{i_n m} \quad (1)$$

$$P_N^{FA}(r, s) = \sum_{(y, I) \in \mathcal{U}_1} \prod_{n=1}^N r_{i_n} p_{y_n}^{i_n} \quad (2)$$

Now we consider all decision rules that satisfies a false alarm constraint:

$$P_N^{FA}(r, s) < \delta \quad (3)$$

for some $\delta \in (0, 1)$. Then we define the payoff function $\mathcal{J}(r, s)$ in the zero-sum game between the observer and the target as the largest achievable error exponent under hypothesis H_1 :

$$\mathcal{J}(r, s) = - \lim_{N \rightarrow \infty} \inf_{\{\mathcal{U}: P_N^{FA}(r, s) < \delta\}} \frac{1}{N} \log(P_N^M(r, s)) \quad (4)$$

The observer and the target have opposing objectives. The observers aims to maximize the exponential decay rate of miss detection probability, whereas the targets objective is to minimize this rate. The Nash-equilibrium of this game is defined as the saddle point (r^*, s^*) satisfying:

$$\mathcal{J}(r, s^*) \leq \mathcal{J}(r^*, s^*) \leq \mathcal{J}(r^*, s) \quad (5)$$

If there exists a Nash-equilibrium (r^*, s^*) then minimax equality is satisfied:

$$V(\mathcal{J}) = \mathcal{J}(r^*, s^*) = \max_r \min_s \mathcal{J}(r, s) = \min_s \max_r \mathcal{J}(r, s)$$

where $V(\mathcal{J})$ defines the value of the game. Using strategy r^* the observer guarantees that the exponential decay rate of P_N^M will be at least as $V(\mathcal{J})$ no matter what strategy that the target employs. Similary, using strategy s^* the target guarantees that the exponential decay rate of P_N^M will be no larger than $V(\mathcal{J})$ irrespective of the strategy of the observer.

3. OPTIMAL STRATEGIES AND THE VALUE OF THE SURVEILLANCE GAME

To gain insight into the two player surveillance game, we first consider the simpler one player decision problem of the observer against a target employing known randomized strategy \bar{s} . We note that the model in which the target is using a known stationary randomized strategy \bar{s} is equivalent to a model with no adversarial actions, where the probability distribution of the observations are replaced with mixture densities $q_j^k(\bar{s})$:

$$q_j^k(\bar{s}) = \sum_m s_m q_j^{km} \quad (6)$$

Result 1. For an observer employing randomized strategy r against a random opponent of known strategy \bar{s} , the largest error exponent under hypothesis H_1 subject to false alarm constraint in (3) is given by:

$$\begin{aligned} \beta(r|\bar{s}) &= - \lim_{N \rightarrow \infty} \inf_{\{\mathcal{U}: P_N^{FA} < \delta\}} \frac{1}{N} \log(P_N^M(r|\bar{s})) \\ &= \sum_k r_k D(p^k | q^k(\bar{s})), \end{aligned}$$

where $D(p|q)$ is the Kullback-Leibler (KL) divergence defined by

$$D(p|q) = \sum_j p_j \log \left(\frac{p_j}{q_j} \right)$$

Proof. (Outline) Since the strategy \bar{s} is known, the observer only has to consider likelihood ratio tests, as shown by the Neyman-Pearson theorem. Using weak law of large numbers the log likelihood of the observed data $\frac{1}{N} \Lambda(y^1, \dots, y^N)$ on the randomly chosen sensor index set I_N converges in probability to its mean of $\sum_k r_k D(p^k | q^k(\bar{s}))$ under H_0 . Then a straightforward variation on Stein's Lemma gives the desired result. \square

The achievable error exponent $\beta(r|\bar{s})$ is a linear function of r and therefore the maximum is attained at a corner of the probability simplex. Result 2 follows immediately

Result 2. *The optimal strategy for an observer against a random opponent of known strategy \bar{s} is to use the most informative sensor repeatedly for all observations, i.e. $r^* = \delta(k, k')$, where the optimal sensor k' is defined by the largest KL divergence to the null hypothesis:*

$$k' = \arg \max_k D(p^k | q^k(\bar{s}))$$

The resulting optimal error exponent subject to the false alarm constraint in (3) is given by:

$$\beta(r^*|\bar{s}) = \max_k D(p^k | q^k(\bar{s}))$$

The asymptotic optimality of single sensor measurements has been noted previously in noncomposite Hypothesis testing problems by Tsitsiklis [10], with extensions to M-ary Hypothesis detection problems. This result shows that asymptotically there is also no benefit in randomizing between sensors when faced with a target with known randomized strategy over its available realizations.

In the following, we show that randomization is an essential component in the adversarial setting to achieve min-max optimality in sensor selection. First we characterize the payoffs surveillance game as a function of the random strategy pair (r, s) .

Result 3. *When the observer and target employ open-loop strategies (r, s) in the game defined in (4) the payoffs are given by:*

$$\mathcal{J}(r, s) = \sum_k r_k D(p^k | q^k(s)) \quad (7)$$

Proof. (Outline) This result does not follow immediately from Result 1, since here we cannot assume that the observer has the knowledge of the strategy s of the target. Consequently, the observer cannot form the composite likelihood to perform the hypothesis test. Instead we consider the Hoeffding test [11] given by $\mathcal{U}_0^H = \{y : D(\hat{p}(y)|p^k) < \lambda\}$, where the empirical probability mass function $\hat{p}(y)$ is calculated from y using $\hat{p}_j(y_N) = \frac{1}{N} \sum_i \delta(y_i, j)$.

A variation of the Sanov theorem reveals that the error exponent for the test under hypothesis H_0 is given by $-\lim_{N \rightarrow \infty} \frac{1}{N} \log P_N^{FA} = \lambda$ and the error exponent for the test under hypothesis H_1 is given by $-\lim_{N \rightarrow \infty} \frac{1}{N} \log P_N^M = \sum_k r_k D(p^k | q^k(s)) - \epsilon(\lambda)$ and $\lim_{\lambda \rightarrow 0} \epsilon(\lambda) = 0$. The constant (non-decaying) false alarm constraint (3) will be satisfied by any nonzero λ . Therefore $\epsilon(\lambda)$ can be made arbitrarily small by decreasing λ . Essentially, Hoeffding test with vanishingly small threshold λ achieves the optimal error-exponent $\beta(r|s)$ of the likelihood ratio test, without knowing the target strategy s or equivalently the mixture densities $q^k(s)$ under H_1 . Taking the supremum over the family of Hoeffding tests gives the desired result. \square

Next, we prove the existence of the Nash Equilibrium strategies for the surveillance game. We use the standard tool in game theory for proving existence, Kakutani's fixed point theorem.

Result 4. *For the game defined by the payoff function given in (4) there exists a Nash equilibrium strategy pair (r^*, s^*) that satisfies the saddle point property (5).*

Proof. (Outline) First we construct best response correspondences. The maximum theorem shows the continuity of the best responses and Kakutani's fixed point theorem (KFPT) is used to prove that they intersect. One set of sufficient conditions for the maximum theorem and KFPT are : Strategy space for each player must be a nonempty, compact, convex subset of the Euclidean space and the payoff (utility) functions for each player must be continuous and quasi-convex in its own strategy. Here the strategy spaces are K and M dimensional probability simplexes and therefore compact and convex. The function $\mathcal{J}(r, s)$ is continuous over (r, s) . The observer's payoff function $\mathcal{J}(r, s)$ is linear in r . The target's payoff $-\mathcal{J}(r, s)$ is concave in s . All the conditions of KFPT and the maximum theorem are thus satisfied, so the Nash equilibrium exists. \square

The Nash equilibrium solution typically lies in the interior of the probability simplex when one sensor does not dominate others. In this case, observer alternates between sensors to avoid being exploited by the target and the target is alternates between evading actions to avoid being exploited by the observer.

4. EXAMPLE

In this section we consider a simple example to illustrate the ideas presented in the previous section. Consider detection of a target with $K = 2$ sensors that provides binary observations. Assume the target has also $M = 2$ evading actions. Both sensors have probability of false alarm of 0.1. The probability of detection for each sensor-action pair by

	Action 1	Action 2
Sensor 1	0.8	0.6
Sensor 2	0.6	0.7

Best response functions are given by:

$$r^*(s) = \arg \max_r \mathcal{J}(r, s) \quad s^*(r) = \arg \min_s \mathcal{J}(r, s) \quad (8)$$

Using these response functions Nash equilibrium strategies can be obtained as a solution to max-min optimality

$$r^* = \arg \max_r \mathcal{J}(r, s^*(r)) \quad (9)$$

$$s^* = \arg \min_s \mathcal{J}(r^*(s), s) \quad (10)$$

Figure 2 shows the max-min payoffs for the observer and the target. The optimal strategy for the observer is $r_1 = (1 - r_2) = 1/3$ and the optimal strategy for the target is $s_1 = (1 - s_2) = 1/3$.

Both players alternate between sensors and evading actions to achieve a guaranteed performance level. For example the observer is alternating between Sensor 1 and Sensor 2 with probability 1/3 and 2/3 to guarantee an error exponent of 1.016. In equilibrium, the reason for its alternation is not to maximize the information rate (in fact given the strategy $s_1 = (1 - s_2) = 1/3$ the observer has the same payoff from each sensor) but instead to avoid being exploited. For example, if Sensor 1 is solely employed, it has a worst case payoff of 0.794 since it is open to exploitation by a target employing the strategy $s_1 = 0$.




Fig. 2. Max-min payoffs for the observer (top) and the target (bottom).

5. REFERENCES

- [1] J. Manyika and H. Durrant-Whyte, *Data fusion and sensor management: a decentralized information-theoretic approach*, Prentice Hall PTR, 1995.
- [2] J. Liu, J. Reich, and F. Zhao, "Collaborative in-network processing for target tracking," *EURASIP Journal on Applied Signal Processing*, vol. 2003, pp. 378–391, 2003.
- [3] E. Ertin, J. Fisher, and L. Potter, "Maximum mutual information principle for dynamic sensor query problems," in *Information Processing in Sensor Networks*, 2003.
- [4] J. Liu, M. Chu, J. Liu, J. Reich, and F. Zhao, "Distributed state representation for tracking problems in sensor networks," in *Information Processing in Sensor Networks*, 2004.
- [5] J.L. Williams, J.W. Fisher, and A.S. Willsky, "Approximate dynamic programming for communication-constrained sensor network management," *IEEE Transactions on Signal Processing*, vol. 55, no. 8, pp. 4300–4311, 2007.
- [6] V. Gupta, T.H. Chung, B. Hassibi, and R.M. Murray, "On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage," *Automatica*, vol. 42, no. 2, pp. 251–260, 2006.
- [7] V. Isler and R. Bajcsy, "The sensor selection problem for bounded uncertainty sensing models," *IEEE Transactions on Automation Science and Engineering*, vol. 3, no. 4, pp. 372–381, 2006.
- [8] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 451–462, 2009.
- [9] V. Srivastava, K. Plarre, and F. Bullo, "Randomized sensor selection in sequential hypothesis testing," *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 2342–2354, 2011.
- [10] J.N. Tsitsiklis, "Decentralized detection by a large number of sensors," *Mathematics of Control, Signals, and Systems*, vol. 1, no. 2, pp. 167–182, 1988.
- [11] W. Hoeffding, "Asymptotically optimal tests for multinomial distributions," *The Annals of Mathematical Statistics*, pp. 369–401, 1965.