

ACTIVE ANOMALY DETECTION WITH SWITCHING COST

Da Chen[†], Qiwei Huang[†], Hui Feng[†], Qing Zhao^{*}, Bo Hu[†]

[†]Research Center of Smart Networks and Systems, School of Information Science and Engineering, Key Lab of EMW Information, Fudan University, Shanghai 200433, China

^{*}School of Electrical and Computer Engineering, Cornell University

Emails: [†]{dachen16, qwhuang16, hfeng, bohu}@fudan.edu.cn, ^{*}qz16@cornell.edu

ABSTRACT

The problem of anomaly detection among multiple processes is considered within the framework of sequential design of experiments. The objective is an active inference strategy consisting of a selection rule governing which process to probe at each time, a stopping rule on when to terminate the detection, and a decision rule on the final detection outcome. The performance measure is the Bayes risk that takes into account not only sample complexity and detection errors, but also costs associated with switching across processes. While the problem is a partially observable Markov decision process to which optimal solutions are generally intractable, a low-complexity deterministic policy is shown to be asymptotically optimal and offer significant performance improvement over existing methods in the finite regime.

Index Terms— Sequential design of experiment, Active hypothesis testing, Anomaly detection.

1. INTRODUCTION

We consider the problem of detecting an anomalous process among M processes. Following the terminology in target detection, we refer to the processes as cells and the anomalous process as the target. At each time, one cell can be probed and a noisy observation is obtained from the cell. The objective is an active inference strategy consisting of a selection rule governing which cell to probe at each time, a stopping rule on when to terminate the detection, and a decision rule on the final detection outcome. The performance measure is the Bayes risk that takes into account of not only detection errors and detection delay (i.e., sample complexity), but also costs associated with switching across cells.

The problem is an instance of active hypothesis testing pioneered by Chernoff [1] under the term of sequential design of experiments, where the objective is on *asymptotic* optimality in the regime of vanishing detection errors. Chernoff's original formulation, however, does not consider switching cost. The test developed by Chernoff (referred to as Chernoff test) is a *randomized* test that generates random actions based on adaptively designed distributions. The advantage of the randomized test is its amenability to asymptotic analysis. Specifically, after an initial phase with

a bounded duration, the action distribution chosen by the Chernoff test no longer changes over time. The actions, consequently the test statistic of log likelihood ratio, thus become i.i.d. over time, allowing easy asymptotic analysis of the stopping time on the test statistic. However, choosing action independently of past actions is fundamentally incompatible with the objective of reducing the switching cost, especially when the switching cost is comparable or even higher than the observation cost. Furthermore, the Chernoff test is given only inexplicitly, as the solution to a maximin problem, which can be computationally expensive to solve.

In this paper, we propose a low-complexity deterministic test for the above active hypothesis testing problem with switching cost. Referred to as the Deterministic Bounded Switching (DBS) policy, the proposed policy explicitly specifies the probing action at each time with little computation. Specifically, the policy is based on a key criterion that integrates all parameters affecting the Bayes risk: the number M of cells, the switching cost s , the observation cost c , and the rates at which the target cell and normal cells can be identified as given by the Kullback-Liebler (KL) divergences between the corresponding observation distributions. This criterion partitions the problem space into two cells. In one cell, the DBS policy probes the cell most likely to be the target. In the other, DBS probes cells that are likely to be normal and eliminates them one at a time to reduce the number of switchings. The DBS policy is simple, intuitively appealing, yet enjoys asymptotic optimality and strong performance in the finite regime as demonstrated in the simulation examples. Compared to the Chernoff test, establishing the asymptotic optimality of DBS is much more involved, due to the strong temporal and spatial (i.e., across cells) dependencies in the actions and test statistics.

Incorporating switching cost into Bayes risk is motivated by a number of applications. For example, in many robotics applications, relocating the robot (or other autonomous decision makers such as UAVs) incurs considerable cost in terms of energy or delay. Another example is medical diagnostics, where frequent and fast switching across drugs and medical procedures carries high risk and side effects.

Related Work:

Chernoff's work have been extended in various directions. Bessler [2] considered the M -ary hypothesis in 1960.

Naghshvar and Javidi studied Active Sequential Hypothesis Testing (ASHT) under different scenarios from a Bayesian cost minimization perspective in [3] and [4]. Nitinawarat and Veeravalli [5] proposed a universal sequential scheme for optimal search and stop using only the knowledge of the absence distribution. In [6] and [7], more general models which considered Markovian observations and non-uniform costs on actions were proposed. Recently, Cohen and Zhao studied ASHT for anomaly detection in [8]–[11]. In contrast to the random policies advocated in other works, they introduced a simple deterministic model and showed that it is asymptotically optimal.

There are few studies on ASHT considering the switching cost. Vaidhiyan developed a modified Chernoff test (referred to as Sluggish policy) in [12] and introduced a switching parameter η which determines the switching probability, their policy can be seen as an ε -greedy strategy. They claimed that the Sluggish policy approaches the asymptotic performance of Chernoff test as $\eta \rightarrow 0$. But the policy will result in a higher observation cost in the finite regime as demonstrated in the simulations.

2. PROBLEM FORMULATION

Consider the problem of detecting a target among M cells. Let H_m denote the hypothesis that the target is in cell m . At each time, only one cell can be probed. When cell m is probed at time n , an observation $y_m(n)$ is drawn independently from a known distribution. If hypothesis H_m is false, $y_m(n)$ follows distribution $f(y)$; if hypothesis H_m is true, $y_m(n)$ follows distribution $g(y)$. Let \mathbf{P}_m be the probability measure under hypothesis H_m and \mathbf{E}_m the operator of expectation concerning the measure \mathbf{P}_m .

Let $\Gamma = (\phi, \tau, \delta)$ denote an active inference policy, where $\phi = \{\phi(n)\}_{n \geq 1}$ is the sequence of selection rules, τ is the stopping rule, and δ is the decision rule at the time of stopping. Let $P_e(\Gamma) = \sum_{m=1}^M \pi_m \alpha_m(\Gamma)$ be the probability of error under policy Γ , where π_m is the prior probability that H_m is true and $\alpha_m(\Gamma) = \mathbf{P}_m(\delta \neq m | \Gamma)$ is the probability of declaring $\delta \neq m$ when H_m is true. Let $\mathbf{E}(\tau | \Gamma) = \sum_{m=1}^M \pi_m \mathbf{E}_m(\tau | \Gamma)$ be the expected detection delay under Γ . Let τ_s be the total number of switchings at the time of stopping and $\mathbf{E}(\tau_s | \Gamma) = \sum_{m=1}^M \pi_m \mathbf{E}_m(\tau_s | \Gamma)$ be its expected value.

Let c denote the observation cost and s be the switching cost. We consider the case where the switching cost is of no greater order than the observation cost in the asymptotic regime of $c \rightarrow 0$, i.e., $\limsup_{c \rightarrow 0} \frac{s}{c}$ is bounded.

The Bayes risk of Γ under hypothesis H_m is given by

$$R_m(\Gamma) \triangleq \alpha_m(\Gamma) + c\mathbf{E}_m(\tau | \Gamma) + s\mathbf{E}_m(\tau_s | \Gamma). \quad (1)$$

The Bayes risk is

$$R(\Gamma) = \sum_{m=1}^M \pi_m R_m(\Gamma) = P_e(\Gamma) + c\mathbf{E}(\tau | \Gamma) + s\mathbf{E}(\tau_s | \Gamma). \quad (2)$$

The objective is a policy that minimizes the Bayes risk.

3. THE DBS POLICY

In this section, we present the proposed DBS policy. The test statistic is the log-likelihood ratio (LLR) of each cell m denoted as

$$l_m(n) \triangleq \log \frac{g(y_m(n))}{f(y_m(n))}. \quad (3)$$

The sum LLR of cell m at time n is given by

$$S_m(n) \triangleq \sum_{t=1}^n l_m(t) \mathbf{1}_m(t), \quad (4)$$

where $\mathbf{1}_m(n)$ is the indicator function on whether cell m is probed at time n . $\mathbf{1}_m(n) = 1$ if cell m is probed at time n , and $\mathbf{1}_m(n) = 0$ otherwise.

The key feature of the proposed DBS policy is to partition the problem space into two regions by comparing the order of $D(f||g)/(M-1)$ and $D(g||f) + \Delta$, where $D(\cdot||\cdot)$ is the KL divergence between two distributions, and the offset Δ is given by

$$\Delta \triangleq \frac{s(M-2)D(g||f)D(f||g)}{-c(M-1)\log c}. \quad (5)$$

We thus have the following two cases given below,

$$\begin{aligned} \text{Case I: } & D(g||f) + \Delta \geq \frac{D(f||g)}{M-1}, \\ \text{Case II: } & D(g||f) + \Delta < \frac{D(f||g)}{M-1}, \end{aligned}$$

which are dealt with by DBS with drastically different selection rules. We specify below the DBS policy in each of the two cases separately.

In Case I, DBS probes the cell most likely to be the target. The selection rule, stopping rule, and decision rule are as follows.

$$\phi(n) = m^1(n), \quad (6)$$

$$\tau = \min \{n : S_{m^1(n)}(n) > -\log c\}, \quad (7)$$

$$\delta = m^1(\tau), \quad (8)$$

where $m^1(n) = \arg \max_m S_m(n)$ is the index of the cell with the largest observation sum LLRs (cells with the same sum LLRs can be ordered arbitrarily) among all the cells.

In Case II, DBS probes the cells that are likely to be normal and eliminates them one by one. Specifically, let $\mathcal{B}(n)$ denote the set of cells that can be reliably determined as normal at time n , i.e.,

$$\mathcal{B}(n) = \{m : S_m(n) < \log c\} \quad (9)$$

The selection, stopping, and decision rules of DBS in Case II are given by

$$\phi(n) = \tilde{m}^{-1}(n),, \quad (10)$$

$$\tau = \min \{n : |\mathcal{B}(n)| = M-1\}, \quad (11)$$

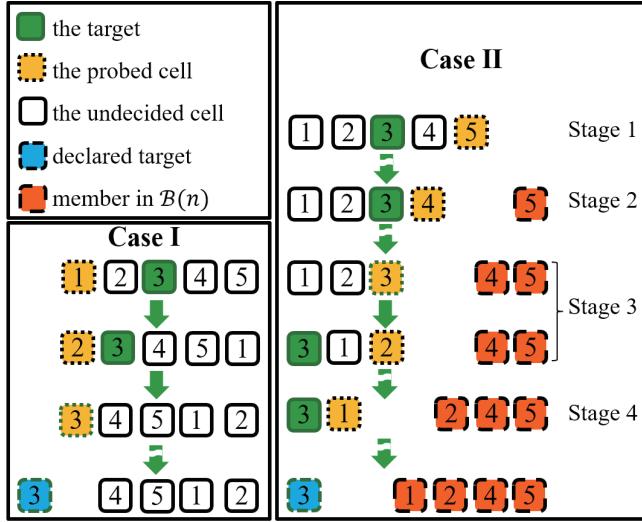


Fig. 1. Illustration of the DBS Policy ($M = 5$).

$$\delta = \mathcal{M} \setminus \mathcal{B}(\tau), \quad (12)$$

where $\tilde{m}^{-1}(n) = \arg \min_{m \notin \mathcal{B}(n)} S_m(n)$ is the index of the cell with the smallest observation sum LLRs among all the cells that can not be reliably determined as normal at time n , and $\mathcal{M} = \{1, 2, \dots, M\}$ is the set of all cells.

We illustrate the DBS policy in Fig. 1, juxtaposing these two cases. Consider hypothesis H_m is true. Under H_m , $S_m(n)$ is a random walk with positive expected increment $\mathbf{E}_m(l_m(n)) = D(g||f) > 0$, while $S_j(n)$, for $j \neq m$ is a random walk with negative expected increment $\mathbf{E}_m(l_j(n)) = -D(f||g) < 0$.

In Case I, DBS probes the cell with the largest sum LLR. The asymptotic detection time approaches $-\log c / D(g||f)$ since the target can be probed with higher probability than other cells at each time. The test is finalized once sufficient information is gathered from the target, when the largest sum LLR exceeds the threshold of $-\log c$. The number of switching is limited because the targeted cell will quickly become the $\tilde{m}^1(n)$ cell, and the probability of switching will become smaller and smaller with the increase of n .

In Case II, the observation process will be divided into $M - 1$ stages. The DBS policy eliminates normal cells one by one and get the target finally. At each stage, the cell $\tilde{m}^{-1}(n)$ is probed. The asymptotic observation time approaches $-(M - 1) \log c / D(f||g)$ since one of the $M - 1$ normal cells is probed at each given time with high probability. The test is finalized once $|\mathcal{B}(n)| = M - 1$.

There will be more switching times in Case II since it has $M - 2$ more stages than Case I. Considering the impact of switching on the strategy, we introduced an offset Δ . It is worthy noting that the value of Δ changes as c changes, and it approaches zero when $c \rightarrow 0$, which means Δ does not affect the asymptotic property of the policy. But the selected case is affected by Δ in the finite regime, we will analyze it in the formulation.

4. ASYMPTOTIC ANALYSIS

The following theorem shows that the DBS policy is asymptotically optimal regarding minimizing the Bayes risk as c approaches zero.

Theorem 1. (Asymptotic Optimality of the DBS Policy): Let R^* and $R(\Gamma)$ denote the Bayes risks in the DBS policy and any other policy Γ respectively. Under the assumption that $s = O(c)$, we have¹

$$R^* \sim \frac{-c \log c}{I^*(M)} \sim \inf_{\Gamma} R(\Gamma) \text{ as } c \rightarrow 0, \quad (13)$$

where

$$I^*(M) \triangleq \begin{cases} D(g||f), & \text{if Case I,} \\ D(f||g)/(M - 1), & \text{if Case II.} \end{cases} \quad (14)$$

Proof: For a detailed proof see [13]. We provide here a sketch of the proof due to the space limit. Firstly, we show that $\frac{-c \log c}{I^*(M)}$ is an asymptotic lower bound on the Bayes risk. Secondly, we prove that the Bayes risk R^* under the DBS policy approaches the asymptotic lower bound as $c \rightarrow 0$. Compared to the proof of the DGF policy [10], the challenges in our policy is to prove that the number of switching is smaller relative to other policy and the asymptotic expected detection time approaches $\frac{-\log c}{I^*(M)}$ as $c \rightarrow 0$.

There are only one stage in Case I, but in Case II, we can split the testing into multiple stage, each stage is defined according to the time when a certain cell is declared and placed into set \mathcal{B} . By analyzing the three last passage times at every stage k , denote as $\tau_1^k, \tau_2^k, \tau_3^k$, we can get the final conclusion. \square

5. NUMERICAL RESULTS

To validate the effectiveness of the proposed algorithm, We present numerical examples to illustrate the performance of the DBS policy compared with the Chernoff test [1], DGF policy [10] and Sluggish policy [12].

Let $R(\Gamma)$ be the Bayes risks under the policy Γ , and $R_{LB} = \frac{-c \log c}{I^*(M)}$ be the asymptotic lower bound on the Bayes risk as $c \rightarrow 0$. Define $L(\Gamma) \triangleq (R(\Gamma) - R_{LB})/R_{LB}$ as the relative loss regarding Bayes risk under policy Γ compared to the asymptotic lower bound, which serve as performance measures of the tests in the finite regime.

We set a single target located in one of M cells with the following parameters: When cell m is probed at time n , an observation $y_m(n)$ is independently drawn from a Poisson distribution $f \sim \text{Pois}(\lambda_f)$ or $g \sim \text{Pois}(\lambda_g)$, depending on whether the cell is the target or normal. We consider the case where $M = 5$. The switching cost is set to $s = 10c$. For the Sluggish policy [12], low switching probability will result in the waste of observations, so we set switching probability to $p = 0.1$. Firstly, we set $\lambda_f = 10$, $\lambda_g = 1$ and obtain $D(g||f) \approx 6.7$, $D(f||g)/(M - 1) \approx 3.5$, the DBS policy is

1. The notion $f \sim g$ as $c \rightarrow 0$ refers to $\lim_{c \rightarrow 0} f/g = 1$.

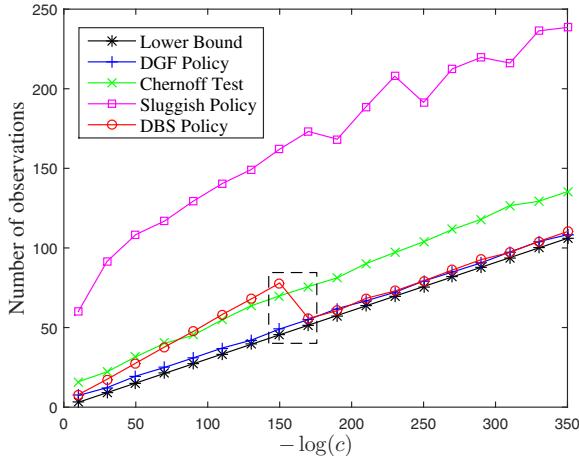


Fig. 2. The number of observations.

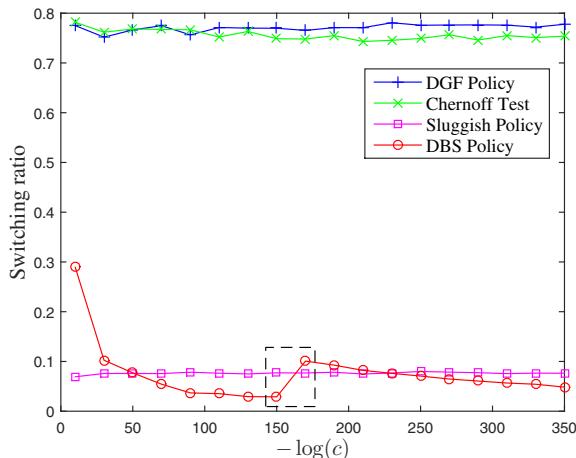


Fig. 3. The proportion of switching times relative to the number of observations.

in Case I for all values of c . All algorithms perform similarly in this case because they all observe the $m^1(n)$ cell, so we omit the diagrams of results here.

Then we set $\lambda_f = 2$, $\lambda_g = 0.001$ and obtain $D(g||f) \approx 1.99$, $D(f||g)/(M-1) \approx 3.3$. The performance of all algorithms is presented in Fig.2-Fig.4, which are the average of 100 trials. It is showed that the Sluggish policy has the highest number of observations and the DGF policy is the lowest in Fig.2. In Fig.3, the proportion of switching times relative to the number of observations of all algorithms are displayed as a function of $-\log c$. We can find that the switching ratio of the DBS policy is less than DGF policy and Chernoff test. The dashed rectangular area of Fig.2 and Fig.3 shows that the observation times and the switching ratio of the DBS policy change suddenly when $-\log c = 150$. The reason is that $D(g||f) + \Delta \geq \frac{D(f||g)}{M-1}$ and $D(g||f) < \frac{D(f||g)}{M-1}$ when $-\log c$ is less than 150. Our algorithm is in Case I, the DBS policy chooses to observe

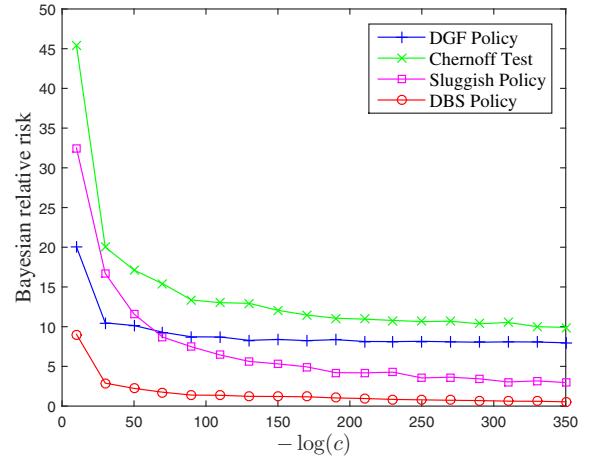


Fig. 4. Relative loss in terms of Bayes risk.

the cell with the largest sum LLRs. Although the policy will result in more observations, which can reduce the number of switching. However, the DBS policy changes to Case II and chooses to observe the $M-1$ normal cells when $-\log c > 150$, which will increase the number of switching and reduce observations, so the switching ratio rise suddenly. After this point, the switching ratio of DBS policy will continuously decrease and approach 0 as $c \rightarrow 0$.

The DGF policy does not consider the impact of switching, it always probes the cell with the second sum LLRs. The DGF policy can reduce the number of observations effectively, but the observed sum LLRs of cell $m^2(n)$ at time n will become smaller with a greater probability after current observation, and the probed cell will always change, which cause a large switching cost. For the Sluggish policy, the switching ratio will always be close to $p = 0.1$ and it result in a higher number of observations. In Fig.4, the Bayesian relative risk $L(\Gamma)$ of all algorithms are displayed. It can be seen that the DBS policy significantly outperforms the other algorithms in the finite regime for all values of c .

6. CONCLUSION

The problem of anomaly detection with switching cost is studied. We propose a deterministic policy and proved that our policy is asymptotically optimal. Future directions include extensions to cases with multiple targets and simultaneous probing of multiple cells.

Acknowledgment

This work of Fudan team was supported by the National Key Research and Development Program of China (No.213). The work of Qing Zhao was supported in part by the National Science Foundation under Grant CCF-1815559 and the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 754412 during her visit at Chalmers University, Sweden.

7. REFERENCES

- [1] Herman Chernoff, “Sequential design of experiments,” *The Annals of Mathematical Statistics*, vol. 30, no. 3, pp. 755–770, 1959.
- [2] Stuart A Bessler, “Theory and applications of the sequential design of experiments, k-actions and infinitely many experiments. part i. theory,” Tech. Rep., Stanford Univ CA Applied Mathematics and Statistics Labs, 1960.
- [3] Mohammad Naghshvar, Tara Javidi, et al., “Active sequential hypothesis testing,” *The Annals of Statistics*, vol. 41, no. 6, pp. 2703–2738, 2013.
- [4] Mohammad Naghshvar and Tara Javidi, “Sequentiality and adaptivity gains in active hypothesis testing,” *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 5, pp. 768–782, 2013.
- [5] Sirin Nitinawarat and Venugopal V. Veeravalli, “Universal scheme for optimal search and stop,” *Bernoulli*, vol. 23, no. 3, pp. 1759–1783, 2017.
- [6] Sirin Nitinawarat, George K Atia, and Venugopal V Veeravalli, “Controlled sensing for multihypothesis testing,” *IEEE Trans. Autom. Control*, vol. 58, no. 10, pp. 2451–2464, 2013.
- [7] Sirin Nitinawarat and Venugopal V Veeravalli, “Controlled sensing for sequential multihypothesis testing with controlled markovian observations and non-uniform control cost,” *Sequential Analysis*, vol. 34, no. 1, pp. 1–24, 2015.
- [8] Kobi Cohen, Qing Zhao, and Ananthram Swami, “Optimal index policies for quickest localization of anomaly in cyber networks,” in *Global Conference on Signal and Information Processing (GlobalSIP)*. IEEE, 2013, pp. 221–224.
- [9] Kobi Cohen, Qing Zhao, and Ananthram Swami, “Optimal index policies for anomaly localization in resource-constrained cyber systems,” *IEEE Trans. Signal Process.*, vol. 62, no. 16, pp. 4224–4236, 2014.
- [10] Kobi Cohen and Qing Zhao, “Active hypothesis testing for anomaly detection,” *IEEE Trans. Inf. Theory*, vol. 61, no. 3, pp. 1432–1450, 2015.
- [11] Chao Wang, Kobi Cohen, and Qing Zhao, “Active hypothesis testing on a tree: Anomaly detection under hierarchical observations,” in *International Symposium on Information Theory (ISIT)*. IEEE, 2017, pp. 993–997.
- [12] Nidhin Koshy Vaidhiyan and Rajesh Sundaresan, “Active search with a cost for switching actions,” in *Information Theory and Applications Workshop (ITA)*. IEEE, 2015, pp. 17–24.
- [13] Da Chen, Qiwei Huang, Hui Feng, Qing Zhao, and Bo Hu, “Active anomaly detection with switching cost,” <https://arxiv.org/abs/1810.11800>, 2018.