ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) | 978-1-7281-6327-7/23/$31.00 ©2023 IEEE | DOI: 10.1109/ICASSP49357.2023.10096555

DATA-DRIVEN QUICKEST CHANGE DETECTION IN MARKOV MODELS

Qi Zhang   *Zhongchang Sun*   *Luis C. Herrera*   *Shaofeng Zou*

University at Buffalo, the State University of New York
Email:qzhang48@buffalo.edu   zhongcha@buffalo.edu   lcherrer@buffalo.edu   szou3@buffalo.edu

ABSTRACT

The problem of quickest change detection in Markov models is studied. A sequence of samples are generated from a Markov model, and at some unknown time, the transition kernel of the Markov model changes. The goal is to detect the change as soon as possible subject to false alarm constraints. The data-driven setting is investigated, where neither the pre- nor the post-change Markov transition kernel is known. A kernel based data-driven algorithm is developed, which applies to general state space and is recursive and computationally efficient. Performance bounds on the average running length and worst-case average detection delay are derived. Numerical results are provided to validate the performance of the proposed algorithm.

Index Terms— Sequential Change Detection, Maximum Mean Discrepancy, Computationally Efficient, CuSum-Type Test, Second-Order Markov Chain.

1. INTRODUCTION

Suppose we have a sequence of observations taken from a stochastic process. At some unknown time (change-point), an event occurs and causes a change in the data-generating distribution of the observations. The goal is to detect the change as quickly as possible subject to false alarm constraints. This problem is referred to as quickest change detection (QCD) and has been widely studied in the literature [1–5]. It finds a wide range of applications, e.g., fault detection in DC micro-grids [6], quality control in online manufacturing systems [7] and spectrum monitoring in wireless communications [8].

Existing studies mostly focus on the setting where samples are independent and identically distributed (i.i.d.) before and after the change. But this i.i.d. assumption may be too restrictive in practice, e.g., in photovoltaic systems [9], samples are not independent over time. For the general non-i.i.d. setting, a generalized CuSum algorithm was developed in [10] and was shown to be asymptotically optimal. The case when samples are generated according to a (hidden) Markov model has also been widely studied in the literature e.g., [7, 10–16]. In these studies, it is usually assumed that the data generating distributions before and after the change are known exactly, and therefore those algorithms may not be applicable if such knowledge is unavailable. In this paper, we focus on the data-driven setting of QCD in Markov models. Specifically, the samples are generated according to a Markov model and at some unknown time, the transition kernel changes. Neither the pre- nor the post-change transition kernel is unknown. The goal is to detect this change as quickly as possible subject to false alarm constraints.

In [17], the problem of QCD in Markov models with unknown post-change transition kernel was studied. Its basic idea is to use the maximum likelihood estimate (MLE) of the unknown post-change transition kernel to construct a generalized CuSum-type test. This approach however requires finite state Markov chain or parameterized Markov transition kernel so that the MLE approach can be applied. In [18], a kernel based data-driven test is proposed for the same problem and the basic idea is to use a kernel method of maximum mean discrepancy (MMD) [19] [1] to construct a CuSum-like test statistic. Specifically, the MMD between a batch of most recent samples and a reference batch of pre-change samples is computed and then used in the test statistic. The average detection delay (ADD) and the average run length (ARL) to false alarm were theoretically characterized. The kernel MMD method was also used in the literature, e.g., [20, 21], for data-driven QCD problem under the i.i.d. setting, however, their analyses cannot be directly applied to the Markov model in this paper.

In this paper, we develop a data-driven QCD algorithm for Markov models, where neither the pre- nor the post-change transition kernels are known. To address the issue that for two different transition kernels, their stationary distributions may still be the same, we derive a second-order Markov model so that as long as two transition kernels are different, their induced stationary distributions (for the second-order Markov model) are different. We maintain a buffer of size m that consists of the most recent samples. Once the buffer is full, we compute the MMD between the second-order samples in the buffer and a reference batch of pre-change samples, and then leverage it to construct a CuSum-type test. After that, we empty the buffer and continue to collect samples. We theoretically show that the ARL of our algorithm is lower bounded exponentially in the threshold and the ADD is upper bounded

[1] See section 2.1 for an introduction to MMD.

linearly in the threshold. Combined with the universal lower bound on the ADD in [10], our method achieves order-level optimal performance. The major challenges in the proof lie in that the samples are from Markov models, and therefore, the bias in the MMD estimate needs to be explicitly characterized in order to analyze the ARL and ADD. Note that in [18], the ARL lower bound is linear in the threshold, which is not ideal since it implies frequent false alarms. In terms of computational complexity, our algorithm only incurs a complexity of $\mathcal{O}(m^2)$ for every m samples, whereas the computational complexity in [18] is $\mathcal{O}(m^2)$ at each time step. Our simulation results also show that our algorithm has a lower ADD for a fixed ARL, and is more computationally efficient.

Due to the space limitation, we omitted the full proof in this paper. The full proof can be found in [22].

2. PROBLEM FORMULATION

Consider a Markov chain $\{X_n\}_{n=1}^{\infty}$ defined on a probability space $(\mathcal{X}, \mathcal{F}, \mathbb{P})$, where $\mathcal{X}$ is bounded i.e., $|x| < \zeta$, for any $x \in \mathcal{X}$. Let $P : \mathcal{X} \to \mathcal{P}(\mathcal{X})$ be the transition kernel, where $\mathcal{P}(\mathcal{X})$ denotes the probability simplex on $\mathcal{X}$. At some unknown time τ, the transition kernel changes into $Q : \mathcal{X} \to \mathcal{P}(\mathcal{X})$. The goal is to detect the change as quickly as possible subject to false alarm constraints. Denote a stopping time by $\hat{T}$ and use $\mathbb{P}_\tau$ ($\mathbb{E}_\tau$) to denote the probability measure (expectation) when the change happens at time τ and $\mathbb{P}_\infty$ ($\mathbb{E}_\infty$) to denote the probability measure (expectation) when there is no change. We assume access to a reference sequence of samples $\{Y_n\}_{n=1}^{\infty}$ generated from the pre-change transition kernel P. Let $\mathcal{F}_t$ be the σ-field generated by $\{X_i, Y_i\}_{i=1}^{t}$. We define the average running length (ARL) and the worst-case average detection delay (WADD) for $\hat{T}$ as follows:

$$\text{ARL}(\hat{T}) = \mathbb{E}_\infty[\hat{T}], \tag{1}$$

$$\text{WADD}(\hat{T}) = \sup_{\tau \geq 1} \operatorname{esssup} \mathbb{E}_\tau[(\hat{T} - \tau)^+ | \mathcal{F}_{\tau-1}]. \tag{2}$$

Here the ARL measures how often we see a false alarm, and the WADD measures how many samples after the change-point are needed to raise an alarm. The goal is to minimize the WADD subject to a constraint on the ARL:

$$\min_{\hat{T}} \text{WADD}(\hat{T}), \text{s.t. } \text{ARL}(\hat{T}) \geq \psi, \tag{3}$$

where $\psi > 0$ is some pre-specified constant. For technical convenience, we make the following assumption.

Assumption 1. *The Markov chains with transition kernel P and Q are uniformly ergodic. Specifically, denote by $\{Z_n\}_{n=1}^{\infty}$ and $\{W_n\}_{n=1}^{\infty}$ Markov chains with transition kernels P and Q, respectively, where π_P and π_Q are their invariant measures. For any $z, w \in \mathcal{X}$, $A \subseteq \mathcal{X}$ and $t > 0$, we have $\sup_A |\mathbb{P}(Z_{t+i} \in A | Z_i = z) - \pi_P(A)| \leq R_P \lambda_P^t$ and $\sup_A |\mathbb{P}(W_{t+i} \in A | W_i = w) - \pi_Q(A)| \leq R_Q \lambda_Q^t$, for some $R_P, R_Q < \infty$ and $0 < \lambda_P, \lambda_Q < 1$.*

2.1. Preliminaries: Maximum Mean Discrepancy

In this section, we review the kernel mean embedding and the maximum mean discrepancy (MMD) [19, 23]. Denote by k: $\mathcal{X} \times \mathcal{X} \to \mathbb{R}$ the kernel function of a reproducing kernel Hilbert space (RKHS) $\mathcal{H}_k$. Denote by $\langle \cdot, \cdot \rangle_{\mathcal{H}_k}$ the inner product in the RKHS. For any $x, y \in \mathcal{X}$, $k(x, \cdot) \in \mathcal{H}_k$ and $k(x, y) = \langle k(x, \cdot), k(y, \cdot) \rangle_{\mathcal{H}_k}$. For any function $f \in \mathcal{H}_k$, $f(x) = \langle f, k(x, \cdot) \rangle_{\mathcal{H}_k}$ due to the reproducing property. In this paper, we consider a positive definite kernel function k: $\mathcal{X} \times \mathcal{X} \to \mathbb{R}$. We assume the kernel is bounded, i.e., $0 \leq k(x, y) \leq 1$. Denoted by $\mu_F = \mathbb{E}_{X \sim F}[k(X, \cdot)]$ the kernel mean embedding of measure F. For a characteristic kernel k and two measures F and G, $\mu_F = \mu_G$ if and only if $F = G$ [19]. The MMD between F and G are defined as:

$$D(F, G) = \sup_{f \in \mathcal{F}} \Big| \mathbb{E}_{X \sim F}[f(X)] - \mathbb{E}_{Y \sim G}[f(Y)] \Big|, \tag{4}$$

where $\mathcal{F} = \{f \in \mathcal{H}_k : \|f\|_{\mathcal{H}_k} \leq 1\}$. The square MMD can be equivalently written as [24]:

$$D^2(F, G) = \mathbb{E}_{X, X' \sim F}[k(X, X')] + \mathbb{E}_{Y, Y' \sim G}[k(Y, Y')] - 2\mathbb{E}_{X \sim F, Y \sim G}[k(X, Y)]. \tag{5}$$

3. MAIN RESULTS

Note that Markov chains with different transition kernels may induce the same invariant measure [18]. Therefore, direct generalization of existing QCD approaches for the i.i.d. setting [20, 21], which estimates the MMD between invariant measures π_P and π_Q, may fail even if $P \neq Q$. To solve this problem, we employ the second-order Markov chain [18]. Denote the product σ-algebra on $\mathcal{X} \times \mathcal{X}$ generated by the collection of all measurable rectangles as $\mathcal{F} \otimes \mathcal{F} = \sigma\{A \times B : A \subseteq \mathcal{X}, B \subseteq \mathcal{X}\}$. Define a probability measure F_P on measurable space $(\mathcal{X} \times \mathcal{X}, \mathcal{F} \otimes \mathcal{F})$ as $F_P(A \otimes B) = \int_A \pi_P(dz) \mathbb{P}(Z_{i+1} \in B | Z_i = z)$, for any $A, B \in \mathcal{F}$. From Lemma 1 in [18], if $P \neq Q$, then $F_P \neq F_Q$. The kernel function and MMD in Section 2.1 can be easily extended to the space of $\mathcal{X} \times \mathcal{X}$.

Lemma 1. *Consider a uniformly ergodic Markov chain $\{Z_n\}_{n=1}^{\infty}$, then its corresponding second-order Markov chain $\{\tilde{Z}_i = (Z_i, Z_{i+1})\}_{n=1}^{\infty}$ is also uniformly ergodic.*

We then construct the test statistic and the stopping rule using the second-order Markov chain. The basic idea is to estimate the MMD between the invariant measures of the second-order Markov chains, and then use that estimate to construct a CuSum-type test. Firstly, we partition the samples into non-overlapping blocks with size m. Denote by $\tilde{X}_i = (X_i, X_{i+1})$, $\tilde{Y}_i = (Y_i, Y_{i+1})$ the second-order Markov chains. For $\{\tilde{X}_i\}_{i=1}^{\infty}$ the invariant measure is F_P before the change and F_Q after the change. For $\{\tilde{Y}_i\}_{i=1}^{\infty}$ the invariant measure is F_P. For the t-th block, $t = 0, 1, 2, ...$, define $F_{\tilde{X}}^t$

its empirical measure as $F^t_{\tilde{X}} = \frac{1}{m-1}\sum_{i=1}^{m-1}\delta_{\tilde{X}_{mt+i}}$, where $\delta_{\tilde{X}_i}$ is the Dirac measure. Similarly, we can define $F^t_{\tilde{Y}}$. The MMD $D(F^t_{\tilde{X}}, F^t_{\tilde{Y}})$ can then be written as:

$$D(F^t_{\tilde{X}}, F^t_{\tilde{Y}}) = \frac{1}{(m-1)}\Bigg(\sum_{mt<i,j<m(t+1)} k(\tilde{X}_i, \tilde{X}_j) + \sum_{mt<i,j<m(t+1)} k(\tilde{Y}_i, \tilde{Y}_j) - 2\sum_{mt<i,j<m(t+1)} k(\tilde{X}_i, \tilde{Y}_j)\Bigg)^{\frac{1}{2}}.$$

Define S_t as: $S_t = D(F^t_{\tilde{X}}, F^t_{\tilde{Y}}) - \sigma$, where $\sigma > 0$ is some positive constant to be specified later. Then, our stopping time is defined as

$$T(c) = \inf\left\{mt : \max_{1\le i\le t}\sum_{j=i}^{t} S_j > c\right\}, \tag{6}$$

where $c > 0$ is a pre-specified threshold. This algorithm can be updated recursively: $\max_{1\le i\le t}\sum_{j=i}^{t} S_j = \max\big\{0, \max_{1\le i\le t-1}\sum_{j=i}^{t-1} S_j + S_t\big\}$. Moreover, the computational complexity for MMD is $\mathcal{O}(m^2)$ for every m samples. At time n, there are $\lfloor\frac{n}{m}\rfloor$ non-overlapping blocks in total. Hence the overall computational complexity up to time n is $\mathcal{O}(mn)$. Nevertheless, the algorithm in [18] needs to calculate the MMD for $n - m + 1$ times due to the use of overlapping blocks. Thus the total computational complexity in [18] is $\mathcal{O}(m^2 n)$. This shows that our algorithm is more computationally efficient.

Before the change, for large m, $D(F^t_{\tilde{X}}, F^t_{\tilde{Y}}) \approx D(F_P, F_P) = 0$ (which will be proved later). After the change, for large m, $D(F^t_{\tilde{X}}, F^t_{\tilde{Y}}) \approx D(F_P, F_Q) > 0$. If we choose $0 < \sigma < D(F_P, F_Q)$, then before the change the test statistic has a negative drift and fluctuates around 0, while after the change, the test statistic has a positive drift and goes above the threshold quickly.

In the following theorem, we theoretically characterize the upper bound on the WADD of our test in (6).

Theorem 1. *The WADD for the stopping time in* (6) *can be bounded as follows:*

$$\begin{aligned}&\mathrm{WADD}(T(c))\\ \le&\frac{2\sqrt{ad}mc}{(a-d)^2} + \frac{(a+d)mc}{(a-d)^2} + 2m + \frac{a+\sqrt{ad}}{d-a}m\\ =&\mathcal{O}(mc).\end{aligned} \tag{7}$$

where $a_P = \sqrt{\frac{2-2\lambda_P+4R_P}{(m-1)(1-\lambda_P)}}$, $a_Q = \sqrt{\frac{2-2\lambda_Q+4R_Q}{(m-1)(1-\lambda_Q)}}$, $a = a_P + a_Q$ *and* $d = D(F_P, F_Q) - \delta$.

Proof sketch. Denote by τ' the last time index of the block where the change happens. Then, samples after τ' follow the transition kernel of Q. Let $\xi = \frac{a_P + a_Q + \sqrt{(D(F_P,F_Q)-\sigma)(a_P+a_Q)}}{D(F_P F_Q)-\sigma-a_P-a_Q}$ such that $D(F_P, F_Q) - \sigma - \frac{\xi+1}{\xi}(a_P + a_Q) > 0$. Denote by $n_c = \lceil\frac{(\xi+1)c}{D(F_P,F_Q)-\sigma}\rceil$. Firstly, we have that

$$\mathbb{E}_\tau\Big[\frac{(T(c)-\tau')^+}{mn_c}\Big|\mathcal{F}_{\tau-1}\Big] \le \sum_{j=0}^{\infty}\mathbb{P}_\tau\Big(\frac{T(c)-\tau'}{mn_c} \ge j\Big|\mathcal{F}_{\tau-1}\Big).$$

By the definition of $T(c)$, it follows that

$$\begin{aligned}&\mathbb{P}_\tau\Big(\frac{T(c)-\tau'}{mn_c} \ge j\Big|\mathcal{F}_{\tau-1}\Big)\\ &\le \mathbb{P}_\tau\Big(\sum_{i=(j-1)n_c+1}^{jn_c} S_i < c\Big|\mathcal{F}_{\tau-1}, \frac{T(c)-\tau'}{mn_c} \ge j-1\Big)\\ &\times \mathbb{P}_\tau\Big(\frac{T(c)-\tau'}{mn_c} \ge j-1\Big|\mathcal{F}_{\tau-1}\Big).\end{aligned} \tag{8}$$

It then suffices to bound $\mathbb{P}_\tau\Big(\sum_{i=(j-1)n_c+1}^{jn_c} S_i < c\Big|\mathcal{F}_{\tau-1}, \frac{T(c)-\tau'}{mn_c} \ge j-1\Big)$ and to apply (8) recursively. By the triangle inequality and Markov inequality, we have that

$$\begin{aligned}&\mathbb{P}_\tau\left(\sum_{i=(j-1)n_c+1}^{jn_c} S_i < c\middle|\mathcal{F}_{\tau-1}, \frac{T(c)-\tau'}{mn_c} \ge j-1\right)\\ \le&\mathbb{P}_\tau\left(\sum_{i=(j-1)n_c+1}^{jn_c}(D(F^i_{\tilde{X}}, F_P) + D(F^i_{\tilde{Y}}, F_Q))\right.\\ \ge& \left. n_c(D(F_P, F_Q)-\sigma) - c\middle|\mathcal{F}_{\tau-1}, \frac{T(c)-\tau'}{mn_c} \ge j-1\right)\\ \le&\sum_{i=(j-1)n_c+1}^{jn_c}\\ &\frac{\mathbb{E}_\tau\Big[D(F^i_{\tilde{X}}, F_P) + D(F^i_{\tilde{Y}}, F_Q)\Big|\mathcal{F}_{\tau-1}, \frac{T(c)-\tau'}{mn_c} \ge j-1\Big]}{n_c(D(F_P, F_Q)) - \sigma) - c}\\ \le&\frac{a_P + a_Q}{D(F_P, F_Q) - \sigma}\left(1 + \frac{1}{\xi}\right).\end{aligned} \tag{9}$$

The last inequality in (9) follows from Proposition 1, which will be given later. □

Note that $a_P, a_Q = \mathcal{O}\Big(\frac{1}{\sqrt{m}}\Big)$. Therefore, the upper bound is linear in the threshold c and the batch size m.

Recall that F_P and $F^t_{\tilde{Z}}$ are the invariant measure and empirical measure for $\{\tilde{Z}_n\}_{n=1}^{\infty}$, respectively. The expectation of their MMD can be bounded in the following proposition.

Proposition 1. $\forall z_{mt} \in \mathcal{X}$, $\mathbb{E}\Big[D(F^t_{\tilde{Z}}, F_P)|Z_{mt} = z_{mt}\Big] \le a_P$.

The samples are Markovian, hence the empirical measures of different blocks are dependent. Nonetheless, Proposition 1 provides an upper bound on the conditional expectation

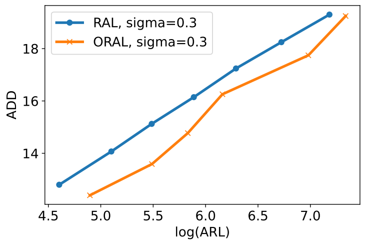

Fig. 1. Comparison of the two algorithms: $\sigma = 0.3$.

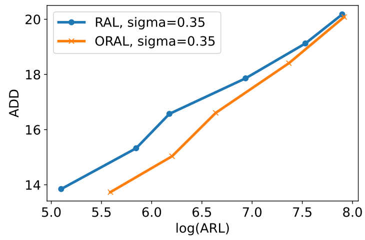

Fig. 2. Comparison of the two algorithms: $\sigma = 0.35$.

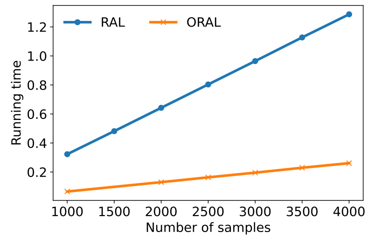

Fig. 3. Number of samples v.s. computational complexity.

of $D(F^t_{\tilde{Z}}, F_P)$. It generalizes the results in [25], where in [25] it was assumed that Z_{mt} follows the invariant distribution.

We then provide a lower bound on the ARL, which is exponential in the threshold c. Before that, we first present a useful proposition that provides a high-probability bound on the sum of two MMDs between the invariant measure and the empirical measure.

Proposition 2. *For any x_{mt} ,$y_{mt} \in \mathcal{X}$ and $\delta > 0$, there exists a positive constant Γ_P that*

$$\mathbb{P}_\infty\Bigg(D(F^t_{\tilde{X}}, F_P) + D(F^t_{\tilde{Y}}, F_P) \leq 2(1+2\Gamma_P)\sqrt{\frac{\log(\frac{1}{\delta})}{m-1}} + 2a_P\Bigg| X_{mt} = x_{mt}, Y_{mt} = y_{mt}\Bigg) \geq 1-\delta. \tag{10}$$

In [25], it was assumed that X_{mt}, Y_{mt} follow the invariant distribution. Here, our result holds $\forall x_{mt}, y_{mt} \in \mathcal{X}$. More importantly, in [25] it was assumed that Γ_P exists. Here we prove that Γ_P exists and is finite with an explicit expression.

Let $b = \frac{2(1+2\Gamma_P)}{\sqrt{m-1}}$ and $h = \sigma - 2a_P$. Since $a_P = \mathcal{O}(\frac{1}{\sqrt{m}})$, then we can always find a large m so that $h > 0$. Let $\phi(q) = \sqrt{\pi} qb \exp\left(-qh + \frac{q^2 b^2}{4}\right)$. Let $q > 0$ be a constant s.t. $\phi(q) \leq 1$. Note that $\phi(q)$ is continuous in q, $\phi(0) = 0 < 1$ and $\phi(q) \to \infty$ when $q \to \infty$. Therefore, there always exists such a q.

Theorem 2. *The ARL of $T(c)$ in* (6) *can be lower bounded exponentially in the threshold c:*

$$\text{ARL}(T(c)) \geq m\exp(qc). \tag{11}$$

Note that q and m are independent of c. Therefore, the ARL grows exponentially in the threshold c.

The universal lower bound from [10] shows for any stopping time with ARL $\geq \psi$, the detection delay is at least $\mathcal{O}(\log(\psi))$. In our paper, to guarantee (3) is satisfied, the threshold c is chosen to $\frac{\log(\psi)-\log(m)}{q}$. According to Theorem 1, this further implies that our algorithm achieves a detection delay of $\mathcal{O}(\log(\psi))$ while satisfying the false alarm constraint. This matches with (order-level) the universal lower bound in [10] for general non-i.i.d. setting.

4. NUMERICAL RESULTS

In this section, we provide simulation results. We name the kernel based CUSUM algorithm proposed in [18] as RAL (oveRlApping bLocks) and our proposed algorithm as ORAL (nOn-oveRlApping bLocks). We then compare the performance of these two algorithms.

We consider an example that the transition kernel changes from $\mathbf{P} = [0.2, 0.7, 0.1; 0.9, 0.0, 0.1; 0.2, 0.8, 0.0]^\top$ to $\mathbf{Q} = [0.5, 0.5, 0.0; 0.0, 0.5, 0.5; 0.2, 0.3, 0.5]^\top$. We choose the Gaussian kernel function $k(x,y) = \exp\big(-\beta(x-y)^2\big)$, where β is the bandwidth parameter. Set $m = 10$, and $\sigma = 0.3$ and 0.35 respectively. To compare the ADD and the ARL, in Fig 1 and 2, we plot the the ADD as a function of the log of ARL by varying the threshold. To compare the computational complexity, in Fig 3 we plot the running time as a function of the total samples (on Intel W-2295 CPU).

From Fig 1, it can be seen that our method outperforms the method in [18], i.e., for the same level of ARL, our method achieves a smaller ADD. Moreover, for both methods, the ADD grows with the log of ARL linearly, which matches with our bounds in this paper, but contradicts with the bounds in [18]. From Fig 2, it can be seen that our method is more computationally efficient.

5. CONCLUSION

In this paper, we studied the data-driven quickest change detection problem in Markov models. We proposed a kernel based detection algorithm, and investigated the bias in MMD estimate for Markov models. Further, we derived its bounds on ARL and WADD, which achieves the order-level optimal performance. Furthermore, compared to the state-of-the-art study with the same setting, our theoretical bounds are tighter and algorithm is computationally efficient. It is of future interest to generalize our results to the hidden Markov model.

6. ACKNOWLEDGMENT

The work of Q. Zhang, Z. Sun and S. Zou is partially supported by National Science Foundation (NSF) under grants CCF1948165, CCF-2106560, and ECCS-2112693.

7. REFERENCES

[1] H. V. Poor and O. Hadjiliadis, *Quickest Detection*. Cambridge University Press, 2009.

[2] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential analysis: Hypothesis testing and changepoint detection*. CRC Press, 2014.

[3] L. Xie, S. Zou, Y. Xie, and V. V. Veeravalli, "Sequential (quickest) change detection: Classical results and new directions," *IEEE J. Sel. Areas Inf. Theory*, vol. 2, no. 2, pp. 494–514, 2021.

[4] V. V. Veeravalli and T. Banerjee, "Quickest change detection," *Academic press library in signal processing: Array and statistical signal processing*, vol. 3, pp. 209–256, 2013.

[5] M. Basseville and I. V. Nikiforov, *Detection of abrupt changes: theory and application*, vol. 104. Prentice Hall Englewood Cliffs, 1993.

[6] K. Gajula, V. Le, X. Yao, S. Zou, and L. Herrera, "Quickest Detection of Series Arc Faults on DC Microgrids," in *Proc. IEEE Energy Conversion Congress and Exposition (ECCE)*, pp. 796–801, IEEE, 2021.

[7] T. L. Lai, "Sequential changepoint detection in quality control and dynamical systems," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 613–658, 1995.

[8] L. Lai, Y. Fan, and H. V. Poor, "Quickest detection in cognitive radio: A sequential change detection framework," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM)*, pp. 1–5, 2008.

[9] L. Chen, S. Li, and X. Wang, "Quickest fault detection in photovoltaic systems," *IEEE Transactions on Smart Grid*, vol. 9, no. 3, pp. 1835–1847, 2016.

[10] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Trans. Inform. Theory*, vol. 44, no. 7, pp. 2917–2929, 1998.

[11] C.-D. Fuh, "SPRT and CUSUM in hidden Markov models," *The Annals of Statistics*, vol. 31, no. 3, pp. 942–977, 2003.

[12] C.-D. Fuh, "Asymptotic operating characteristics of an optimal change point detection in hidden Markov models," *The Annals of Statistics*, vol. 32, no. 5, pp. 2305–2339, 2004.

[13] B. Yakir, "Optimal detection of a change in distribution when the observations form a Markov chain with a finite state space," *Lecture Notes-Monograph Series*, pp. 346–358, 1994.

[14] A. M. Polansky, "Detecting change-points in Markov chains," *Computational statistics & data analysis*, vol. 51, no. 12, pp. 6013–6026, 2007.

[15] B. Darkhovsky, "Change-point problem for high-order Markov chain," *Sequential Analysis*, vol. 30, no. 1, pp. 41–51, 2011.

[16] C.-D. Fuh and A. G. Tartakovsky, "Asymptotic bayesian theory of quickest change detection for hidden Markov models," *IEEE Transactions on Information Theory*, vol. 65, no. 1, pp. 511–529, 2018.

[17] J.-G. Xian, D. Han, and J.-Q. Yu, "Online change detection of Markov chains with unknown post-change transition probabilities," *Communications in Statistics-Theory and Methods*, vol. 45, no. 3, pp. 597–611, 2016.

[18] H. Chen, J. Tang, and A. Gupta, "Change Detection of Markov Kernels with Unknown Post Change Kernel using Maximum Mean Discrepancy," *arXiv preprint arXiv:2201.11722*, 2022.

[19] B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. Lanckriet, "Hilbert space embeddings and metrics on probability measures," *The Journal of Machine Learning Research*, vol. 11, pp. 1517–1561, 2010.

[20] T. Flynn and S. Yoo, "Change detection with the kernel cumulative sum algorithm," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 6092–6099, IEEE, 2019.

[21] S. Li, Y. Xie, H. Dai, and L. Song, "Scan B-statistic for kernel change-point detection," *Sequential Analysis*, vol. 38, no. 4, pp. 503–544, 2019.

[22] Q. Zhang, Z. Sun, L. C. Herrera, and S. Zou, "Data-Driven Quickest Change Detection in Markov Models," *arXiv preprint arXiv:2210.11988*, 2022.

[23] A. Berlinet and C. Thomas-Agnan, *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.

[24] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 723–773, 2012.

[25] B.-E. Chérief-Abdellatif and P. Alquier, "Finite sample properties of parametric MMD estimation: robustness to misspecification and dependence," *Bernoulli*, vol. 28, no. 1, pp. 181–213, 2022.