2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

On Insufficiently Informative Measurements in Bayesian Quickest Change Detection and Identification

Jason J. Ford[1], Jasmin James[2], and and Timothy L. Molloy[3]

***Abstract*— In this paper, we describe an undesirable weak practical super-martingale hallucination phenomenon that can emerge in the Bayesian quickest detection and identification problem. We establish that when measurements are insufficiently informative, a situation described by a relative entropy condition on measurement densities, the Bayesian quickest detection and identification solution can (undesirably) become increasingly confident that a change has occurred, even when it has not. Finally, we illustrate the phenomenon in simulation studies and the vision-based aircraft detection application which illustrates the optimal rule can be unsuitable in the sense of hallucinating a change that has not occurred.**

I. Introduction

The problem of quickly detecting a change in the statistical proprieties of a sequentially observed signal is considered in a diverse range of applications, including fault detection [1], cyber-security [2], vision-based aircraft detection [3]–[5] and more [6]. Quickest change detection and identification (QCDI) is an important sequential decision problem involving quickly detecting and also identifying a possible change as one of several possible anomalous conditions [5], [7], [8]. The QCDI problem is commonly posed in a Bayesian setting as minimising a trade-off between detection delay, probability of false alarm and probability of misidentification. In this paper, we investigate the properties of Bayesian QCDI problems in weak measurement environments (in the sense of being insufficiently informative).

Optimal detection and identification rules for Bayesian QCDI have been established in [5], [8] as the first entry of the change posterior into a union of convex sets. Whilst calculation of the change posterior is relatively easy, determination of these convex stopping regions is computationally challenging. Several numerical techniques for calculating these stopping regions have been proposed [8], including approximation via linear segments [9], or greedy stopping rules [5]. Further, [7] also established that in the absence of a misidentification penalty, the problem simplifies significantly and the optimal detection rule collapses to a simple threshold test on the no change posterior.

Central to this work, [10] established conditions for QCD without identification in which measurements were insufficiently informative to overcome the change time's geometric prior and the no change posterior becomes a weak practical super-martingale hallucination. More recently, this undesirable hallucination phenomenon has also been established in Bayesian QCD for Markov chains [11].

The key contribution of this paper is to extend these works to identify and characterise conditions under which the weak practical super-martingale phenomenon can emerge in the Bayesian QCDI problem; and the QCDI solution undesirably becomes increasingly confident that a change has occurred, even when it has not. Specific contributions of this paper are:

i) Establishing an upper bound on the no change posterior update step based on relative entropy conditions on measurement densities, providing a condition where an undesirable phenomenon can emerge.
ii) Illustrating the undesirable weak practical super-martingale hallucination phenomenon in both simulation and an aircraft detection application.

The rest of this paper is organized as follows: In Section II we revisit the Bayesian QCDI problem and present an optimal detection and identification rule. In Section III we investigate sufficient conditions for insufficiently informative measurements and the emergence of the weak practical super-martingale hallucination phenomenon. In Section IV we illustrate the phenomenon in both simulations and an important application. Finally, some conclusions are presented in Section V.

II. Quickest change detection and identification problem formulation

In this section we revisit the Bayesian QCDI problem and optimal solution [5], [7], [8] and introduce a hidden Markov model representation that is useful for investigating the weak practical super-martingale phenomenon.

A. Bayesian signal change model

For $k > 0$, let $y_k \in \mathcal{Y}$ be an independent and identically distributed (*i.i.d.*) sequence of random variables taking values in the set $\mathcal{Y} \subseteq \mathbb{R}^M$. Initially, when in the no change state, the random variables y_k have a pre-change (marginal) probability density $b^1(\cdot)$ before, at some random change time $\nu \geq 1$, switching to having a post-change (marginal) probability density $b^d(\cdot)$ for some post-change state $d \in \{2, \ldots, N\}$, with $N > 2$. We assume that $b^i(.) < Q$ for some finite $Q < \infty$ for $i \in \{1, \ldots, N\}$.

We will assume densities are absolutely continuous with respect to each other so we can define relative entropy between densities as $D\left(b^i(y_k)\middle|\middle|b^j(y_k)\right) \triangleq$

[1] J. Ford is with the School of Electrical Engineering and Robotics, Queensland University of Technology, Brisbane QLD, 4000 Australia. j2.ford@qut.edu.au. JF's work was supported by the Queensland University of Technology's Centre for Robotics. [2]J. James is with the School of Mechanical and Mining Engineering, University of Queensland, St Lucia, QLD 4072, jasmin.martin@uq.edu.au. [3] T. Molloy is with the School of Engineering, Australian National University (ANU), Acton, ACT 2601, Australia tim.molloy@anu.edu.au.

$\int_{\mathcal{Y}} b^i(y_k) \log\left(\frac{b^i(y_k)}{b^j(y_k)}\right) dy_k$ for all $i, j \in \{1, \ldots, N\}$. Recall that $D\left(\cdot \middle|\middle| \cdot\right) \geq 0$ are non-negative and note that our bounded density assumption gives that these relative entropies are bounded. We will also assume unique densities in the sense that $D\left(b^i(y_k) \middle|\middle| b^j(y_k)\right) > 0$ for $i \neq j$.

As in [7], [8], we assume the random change time $\nu \geq 1$ has a geometric prior of $\pi_k \triangleq (1-\rho)^{k-1}\rho$ (with $\pi_k \triangleq 0$, $k < 1$), where $\rho \in (0, 1)$ corresponds to the probability of a change at each time step given no change has yet occurred, and we assume that change is equally likely to one of the $d \in \{2, \ldots, N\}$ post-change states.

Let us now introduce some probability measure spaces. Let $\mathcal{F}_k = \sigma(y_{[1,k]})$ denote the filtration generated by $y_{[1,k]}$. We will assume the existence of probability spaces, for change at $\nu \geq 1$ to post-change state $d \in \{2, \ldots, N\}$, $(\Omega, \mathcal{F}, P_\nu^d)$ where Ω is a sample space of sequences $y_{[1,\infty]}$, σ-algebra $\mathcal{F} = \cup_{k=1}^{\infty} \mathcal{F}_k$ with the convention that $\mathcal{F}_0 = \{0, \Omega\}$, and P_ν^d is the probability measure constructed using Kolmogorov's extension on the joint probability density function of the observations $p_\nu^d(y_{[1,k]}) = \Pi_{m=1}^{\nu-1} b^1(y_m) \Pi_{n=\nu}^{k} b^d(y_n)$ where we define $\Pi_{j=\nu}^{k} b^d(y_j) = 1$ when $k < \nu$. We will let E_ν^d denote expectation under P_ν^d and use the probability measure P_∞ and expectation E_∞ to denote the special case when there is no change event.

Under our geometric change prior assumption, we can construct a new averaged measure $P_\pi(G) = \frac{1}{N-1}\sum_{d=2}^{N}\sum_{k=1}^{\infty} \pi_k P_\nu^d(G)$ for all $G \in \mathcal{F}$ and we let E_π denote the corresponding expectation operation.

Similar to [8], we will represent the change time and pre-change or post-change state at each time through a process X_k, for $k \geq 0$, taking values in space S_X where $S_X \triangleq \{e_1, \ldots, e_N\}$ with $e_i \in \mathbb{R}^N$ being indicator vectors with 1 as the ith element and 0 elsewhere. A change at time $\nu > 0$ to post-change state $d \in \{2, \ldots, N\}$ would be represented by the process with $X_k = e_1$ for $k < \nu$ and $X_k = e_d$ for $k \geq \nu$. Here e_1 indicates being in the pre-change (no change) state.

It will be useful to note that under the assumption of a geometric prior for the change time ν, the change process X_k corresponds to a first-order time-homogeneous Markov chain, with initial condition $X_0 = e_1$, and a transition probability matrix $A \in \mathbb{R}^{N \times N}$ with i,jth elements $A^{i,j} \triangleq \mathbb{P}\left(X_{k+1} = e_i \middle| X_k = e_j\right)$ for all $i, j \in \{1, \ldots, N\}$, where $A^{1,1} = 1 - \rho$, $A^{i,1} = \rho/(N-1)$ for all $i \in \{2, \ldots, N\}$, $A^{i,i} = 1$ for all $i \in \{2, \ldots, N\}$ and $A^{i,j} = 0$ elsewhere. That is, at each time step before the change event, $\rho/(N-1)$ is the probability of transitioning from e_1 into any e_i, $i \in \{2, \ldots, N\}$ and $(1-\rho)$ is the probability of remaining in e_1. Once the change occurs, the change process stays in this e_d for all future time.

Finally, let $\Pi = \{\lambda \in \mathbb{R}^N : \underline{1}'\lambda = 1, 0 \leq \lambda^i \leq 1 \text{ for all } i \in \{1, \ldots, N\}\}$ denote a probability simplex, and note that $S_X \subset \Pi$, where $\underline{1} \in \mathbb{R}^N$ is the vector of all ones, and $'$ denotes the transpose operation.

B. Quickest change detection and identification problem

The Bayesian QCDI problem's goal is the quickest detection and identification of a change event. For this purpose, we seek to design a stopping time $\tau \geq 0$ with respect to the $\mathcal{F}_k$ filtration and a $\{2, \ldots, N\}$ valued $\mathcal{F}_\tau$-measurable random variable identification decision $\hat{d} \in \{2, \ldots, N\}$ that minimises the following detection and identification cost criterion

$$J(\tau, \hat{d}) \triangleq E_\pi\left[\sum_{k=0}^{\tau-1} \mathcal{C}(X_k) + \mathcal{S}_{\hat{d}}(X_\tau)\right]. \tag{1}$$

Here the continuing and stopping costs, for any $X \in \Pi$, are given by

$$\mathcal{C}(X) \triangleq \bar{c}_1' X, \tag{2}$$

$$\mathcal{S}_i(X) \triangleq c_2 X^1 + c_3(1 - X^i). \tag{3}$$

where $\bar{c}_1 \in \mathbb{R}^N$ ($\bar{c}_1^1 = 0$, $\bar{c}_1^i = c_1 > 0$ for $i \in \{2, \ldots, N\}$) denotes the delay penalty, $c_2 > 0$ denotes the false alarm penalty and c_3 denotes the misidentification penalty.

C. Optimal change detection and identification solution

Let the pre-change and post-change posterior probabilities be denoted $\hat{X}_k^i \triangleq P_\pi(X_k = e_i | y_1, \ldots, y_k)$ for $i \in \{1, \ldots, N\}$, and define the change posterior vector $\hat{X}_k \triangleq [\hat{X}_k^1, \ldots, \hat{X}_k^N]' \in \Pi$. Solutions to the QCDI problem described by (1) are given in [5], [8]. Proposition 11 of [8] and Theorem 1 of [5] give that optimal quickest decision for our QCDI problem is given by (τ^*, d^*) with

$$\tau^* \triangleq \inf\{k : \hat{X}_k \in \mathcal{R}_S\} \tag{4}$$

where the optimal stopping region $\mathcal{R}_S \triangleq \cup_{i \in \{2, \ldots, N\}} \mathcal{R}_S^i$ is the union of $N-1$ convex regions $\mathcal{R}_S^i$ containing e_i, and the optimal identification decision is given by

$$d^* \triangleq \text{argmin}_i(\mathcal{S}_i(\hat{X}_{\tau^*})). \tag{5}$$

As studied in [7], when there is no penalty on misidentification $c_3 = 0$, the optimal stopping region becomes one connected region $\mathcal{R}_S$, and the stopping rule becomes a threshold test on the (pre- or) no change posterior $\hat{X}_k^1$ (see a $N = 3$ example in left of Figure 1).

We highlight that including a misidentification penalty, i.e. $c_3 > 0$, significantly changes the nature of the problem. As shown in [5], [8], solutions to the QCDI problem containing a penalty on misidentification should not expect to have an individually connecting stopping region (see $\mathcal{R}_S^2$ and $\mathcal{R}_S^3$ in a $N = 3$ example in right of Figure 1).

D. Hidden Markov Model Filter for $\hat{X}_k$

We now present the hidden Markov model (HMM) filter for efficiently calculating $\hat{X}_k$ used in the detection (4) and identification rule (5).

At time $k > 0$, we let $B(y_k) = \text{diag}(b^1(y_k), \ldots, b^N(y_k))$ denote the diagonal matrix of output probability densities. We can now calculate $\hat{X}_k$ via the HMM filter [12]

$$\hat{X}_k = N_k B(y_k) A \hat{X}_{k-1}, \tag{6}$$

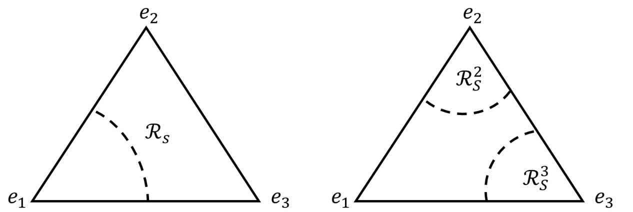

Fig. 1. An example of the stopping regions when there is no penalty on misidentification (left) and when there is a penalty (right).

with initial condition $\hat{X}_0$ and where N_k are scalar normalisation factors defined by

$$N_k^{-1} \triangleq \langle \underline{1}, B(y_k) A \hat{X}_{k-1} \rangle. \quad (7)$$

We highlight that we can write the (pre- or) no change posterior as $\hat{X}_k^1 = N_k(1-\rho)b^1(y_k)\hat{X}_{k-1}^1$. This no change posterior form will be used in the following section to establish some important properties.

III. Conditions for the super martingale hallucination phenomenon

This section begins by defining the concepts of a weak practical super-martingale and a weak post-change density bound. We then establish that when there is a weak post-change density bound and the corresponding post-change posterior is close enough to 1, then the (pre- or) no change posterior trends downwards in a conditional expectation sense. That is, the no change posterior may become increasingly confident that a change has occurred even if it has not.

In this paper, we are concerned with the following phenomenon.

Definition 1. (Weak practical super-martingale) The no change posterior probability $\log(\hat{X}_k^1)$ is called a weak practical super-martingale, if for any arbitrarily small $\delta_p > 0$ there exists a $m \in \{2, \ldots, N\}$ and a $h_s \in (0,1)$ such that when $\hat{X}_k^m > 1 - h_s$ then

$$P_\pi \left(\text{for all } n \geq k, E_\pi[\log(\hat{X}_{n+1}^1)|\hat{X}_n] < \log(\hat{X}_n^1) \right) > 1 - \delta_p. \quad (8)$$

This weak practical super-martingale phenomenon was recently established to occur in weak measurement environments in the Bayesian QCD problems, see [10] (QCD without identification) and [11] (QCD for Markov chains).

We next define the concept of a weak post-change density bound which facilitates analysis of this super-martingale phenomenon in Bayesian QCDI. This concept is sufficient for analysis but expected to be stronger than needed for the super-martingale hallucination phenomenon to emerge.

Definition 2. (Weak post-change density bound) If there is some $m \in \{2, \ldots, N\}$, such that $D\left(b^i(y_k)\middle|\middle|b^m(y_k)\right) < D\left(b^i(y_k)\middle|\middle|b^1(y_k)\right)$ holds for all $i > 1$, then we term $b^m(\cdot)$ a weak post-change density bound.

Definition 2 is related to the weak stochastic boundedness concepts arising in minimax robust QCD, see [13, Def. 4.1] which is a generalisation of the θ-Pythagorean concept of [14, Def. 3.1] (when post-change models can be parametrized) and is a relaxation of the joint stochastic boundedness concepts of [15, Def. 1].

Similar to those works, a weak post-change density bound need not exist; however, a sufficient condition that ensures existence is that the post-change densities have smaller relative entropies between themselves than with the pre-change densities in the sense $\max_{i\neq 1, j\neq 1} D\left(b^i(y_k)\middle|\middle|b^j(y_k)\right) < \min_{i1} D\left(b^i(y_k)\middle|\middle|b^1(y_k)\right)$; which does not seem an overly restrictive condition. Due to the next proposition result, it can be useful to check if the post-change density $b^m(\cdot)$ closest to the pre-change density $b^1(\cdot)$ in the (forward) relative entropy sense that $D\left(b^1(y_k)\middle|\middle|b^m(y_k)\right) < \min_{i\neq 1,m} D\left(b^1(y_k)\middle|\middle|b^i(y_k)\right)$ also meets the (reverse) relative entropy conditions of Definition 2.

For the following investigation, let us define the no-change posterior ratio $M_k^1 \triangleq \hat{X}_k^1/\hat{X}_{k-1}^1$, and note that we can write the posterior update as $\log(\hat{X}_k^1) = \log(M_k^1) + \log(\hat{X}_{k-1}^1)$. Hence the sign of $\log(M_k^1)$ describes the direction of update.

Proposition 1. ($\hat{X}_k^m$ dependent bound on $\log(M_k^1)$) Assume there is a $m \in \{2, \ldots, N\}$ such that $b^m(\cdot)$ is a weak post-change density bound (as Definition 2). For any $\delta > 0$, there is a $h_\delta \in (0,1)$ such that $\hat{X}_{k-1}^m > 1 - h_\delta$ and we have

$$E_\pi \left[\log(M_k^1) \middle| \hat{X}_{k-1} \right] < \log(1-\rho) + D\left(b^1(y_k)\middle|\middle|b^m(y_k)\right) + \delta.$$

Proof. This proof follows a similar strategy to Lemma 2 of [10]. For the $m \in \{2, \ldots, N\}$ selected in the proposition statement, let us defined $\gamma_k^m \triangleq \log(b^m(y_k)) - \log(N_k^{-1})$ and note we can write $\log(N_k) = -\log(b^m(y_k)) + \gamma_k^m$. We then write

$$\begin{aligned} &E_\pi \left[\log(N_k) \middle| \hat{X}_{k-1} \right] \\ &\quad = -E_\pi \left[\log\left(b^m(y_k)\right) \middle| \hat{X}_{k-1} \right] + E_\pi \left[\gamma_k^m \middle| \hat{X}_{k-1} \right]. \end{aligned} \quad (9)$$

It then follows from the definition of M_k^1, and recalling we can write $\hat{X}_k^1 = N_k(1-\rho)b^1(y_k)\hat{X}_{k-1}^1$, that

$$\begin{aligned} &E_\pi[\log(M_k^1)|\hat{X}_{k-1}] \\ &= \log(1-\rho) + E_\pi\left[\log(N_k)\middle|\hat{X}_{k-1}\right] + E_\pi\left[\log(b^1(y_k))\middle|\hat{X}_{k-1}\right] \\ &= \log(1-\rho) + E_\pi\left[\log(b^1(y_k))\middle|\hat{X}_{k-1}\right] \\ &\quad - E_\pi\left[\log\left(b^m(y_k)\right)\middle|\hat{X}_{k-1}\right] + E_\pi\left[\gamma_k^m\middle|\hat{X}_{k-1}\right] \\ &= \log(1-\rho) + E_\pi\left[\log\left(\frac{b^1(y_k)}{b^m(y_k)}\right)\middle|\hat{X}_{k-1}\right] + E_\pi\left[\gamma_k^m\middle|\hat{X}_{k-1}\right]. \end{aligned} \quad (10)$$

Note that we can write, letting $\bar{\rho} \triangleq (1-\rho)$,

$$\begin{aligned}
N_k^{-1} &= b_1(y_k)\bar{\rho}\hat{X}_{k-1}^1 + \sum_{i=2}^N b_i(y_k)[\bar{\rho}\hat{X}_{k-1}^1 + \hat{X}_{k-1}^i] \\
&= b_1(y_k)\bar{\rho}\hat{X}_{k-1}^1 + \sum_{i=2,\neq m}^N b_i(y_k)[\bar{\rho}\hat{X}_{k-1}^1 + \hat{X}_{k-1}^i] \\
&+ b_m(y_k)[\bar{\rho}\hat{X}_{k-1}^1 + \hat{X}_{k-1}^m] \\
&= b_1(y_k)\bar{\rho}\hat{X}_{k-1}^1 + \sum_{i=2,\neq m}^N b_i(y_k)[\bar{\rho}\hat{X}_{k-1}^1 + \hat{X}_{k-1}^i] \\
&+ b_m(y_k) + b_m(y_k)(\bar{\rho}\hat{X}_{k-1}^1 + \hat{X}_{k-1}^m - 1) \\
&= b_m(y_k) + b_m(y_k)(\hat{X}_{k-1}^m - 1) + \sum_{i\neq m}^N c_i(y_k)\hat{X}_{k-1}^i
\end{aligned}$$

for some $c_i(\cdot)$ that are finite as we assumed $b^i(.) < Q$ for some finite $Q < \infty$ for all $i \in \{1, \ldots, N\}$. Then note that for a $h_\delta > 0$ such that for any $\hat{X}_{k-1}^m > 1 - h_\delta$, we also have $\hat{X}_{k-1}^i < h_\delta$ for $i \neq m$. Hence, noting log is monotonic increasing and the definition of γ_k^m gives that for any $\delta > 0$ there is a $h_\delta > 0$ so we have $E_\pi\left[\gamma_k^m \middle| \hat{X}_{k-1}\right] \leq \delta$. Therefore (10) gives that

$$\begin{aligned}
&E_\pi[\log(M_k^1)|\hat{X}_{k-1}] \\
&\quad < \log(1-\rho) + E_\pi\left[\log\left(\frac{b^1(y_k)}{b^m(y_k)}\right) \middle| \hat{X}_{k-1}\right] + \delta. \quad (11)
\end{aligned}$$

Then using the law of total probability we note that

$$\begin{aligned}
&P_\pi(y_k|\hat{X}_{k-1}) \\
&\quad = \sum_{i=1}^N P_\pi(X_k = e_i|\hat{X}_{k-1})P_\pi(y_k|X_k = e_i, \hat{X}_{k-1}). \quad (12)
\end{aligned}$$

As a shorthand, let us write $\hat{X}_{k-1}^{i+} \triangleq P_\pi(X_k = e_i|\hat{X}_{k-1})$ for $i \in \{1, \ldots, N\}$. The definition of the expectation operation and noting $P_\pi(y_k|X_k = e_i, \hat{X}_{k-1}) = b^i(y_k)$ then gives

$$\begin{aligned}
&E_\pi\left[\log\left(\frac{b^1(y_k)}{b^m(y_k)}\right) \middle| \hat{X}_{k-1}\right] \\
&\quad = \hat{X}_{k-1}^{1+}\int_{\mathcal{Y}} b^1(y_k)\log\left(\frac{b^1(y_k)}{b^m(y_k)}\right) dy_k \\
&\quad\quad + \sum_{i=2}^N \hat{X}_{k-1}^{i+}\int_{\mathcal{Y}} b^i(y_k)\log\left(\frac{b^1(y_k)}{b^m(y_k)}\right) dy_k \quad (13)
\end{aligned}$$

Define $\Delta_k^m \triangleq \sum_{i=2,\neq m}^N \hat{X}_{k-1}^{i+}\int_{\mathcal{Y}} b^i(y_k)\log\left(\frac{b^1(y_k)}{b^m(y_k)}\right) dy_k$, and we note from the definition of relative entropy, and algebraic manipulation of log properties that

$$\begin{aligned}
\int_{\mathcal{Y}} b^i(y_k)\log\left(\frac{b^1(y_k)}{b^m(y_k)}\right) dy_k &= D\left(b^i(y_k)\middle|\middle|b^m(y_k)\right) \\
&- D\left(b^i(y_k)\middle|\middle|b^1(y_k)\right) < 0
\end{aligned}$$

where negative follows because $b^m(\cdot)$ is a weak post-change density bound (as Definition 2), and hence $\Delta_k^m < 0$ as $\hat{X}_{k-1}^{i+} \geq 0$ for all $i \in \{1, \ldots, N\}$. Therefore we can write

$$\begin{aligned}
&E_\pi\left[\log\left(\frac{b^1(y_k)}{b^m(y_k)}\right) \middle| \hat{X}_{k-1}\right] \\
&\quad = \hat{X}_{k-1}^{1+}D(b^1(y_k)||b^m(y_k)) \\
&\quad\quad + \hat{X}_{k-1}^{m+}\int_{\mathcal{Y}} b^m(y_k)\log\left(\frac{b^1(y_k)}{b^m(y_k)}\right) dy_k + \Delta_k^m \\
&\quad \leq \hat{X}_{k-1}^{1+}D(b^1(y_k)||b^m(y_k)) \\
&\quad\quad - \hat{X}_{k-1}^{m+}D(b^m(y_k)||b^1(y_k)) \\
&\quad < D\left(b^1(y_k)\middle|\middle|b^m(y_k)\right) \quad (14)
\end{aligned}$$

where in the 2nd last line we have used the definition of relative entropy and $\Delta_k^m < 0$, and the last line has used that $\hat{X}_{k-1}^{1+} \leq 1$ and $\hat{X}_{k-1}^{m+} > 0$. Substitution of (14) into (11) gives the proposition result. □

The importance of Proposition 1 is if there is a weak post-change density bound (as Definition 2), and it is too close to the pre-change density in the sense $D\left(b^1(y_k)\middle|\middle|b^m(y_k)\right) < \log(1/(1-\rho))$, then there can exist an interval trap $\hat{X}_{k-1}^m > 1 - h_\delta$ where $E_\pi\left[\log(M_k^1)\middle|\hat{X}_{k-1}\right]$ is negative. In this situation, the no change posterior may become increasingly confident that a change has occurred even if it has not (recalling $\log(M_k^1)$ is the direction of the no change posterior update). This leads to concerns about the use of the QCDI detection rule (4) when measurements are insufficiently informative. This weak post-change density bound is sufficient for analysis but the super-martingale hallucination phenomenon is expected to also emerge in other situations.

Note that Proposition 1 does not establish, by itself, that $\hat{X}_n^m$ will remain in the internal $\hat{X}_n^m > 1 - h_\delta$ for $n \geq k$. Future work is the formal establishment of the weak practical super-martingale phenomenon, such as Definition 1, similar to that established for QCD without identification in [10] or for QCD in Markov chains [11]. Conversely, the phenomenon can occur in the absence of a weak post-change density bound which was assumed in Proposition 1 to facilitate analysis. We next investigate the hallucination phenomenon in simulation and experimental studies.

IV. Experimental Investigation of the Hallucination Phenomenon

In this section, we will examine the undesirable super-martingale phenomenon suggested by Proposition 1, and illustrate that the QCDI detector can become increasingly confident that a change has occurred, even when it has not. The potential for this behaviour is an important design consideration in some situations.

In this section, we first examine the behaviour of the no change posterior $\hat{X}_k^1$ in a $N = 4$ QCDI problem with insufficiently informative measurements. Then we illustrate the emergence of the undesirable phenomenon in a practical setting.

A. Illustrative Simulation Example

We simulated a $N = 4$ signal X_k which could switch from a pre-change state e_1 to one of three post-change states e_2, e_3 and e_4 according to the transition probability matrix

$$A = \begin{bmatrix} 1-\rho & 0 & 0 & 0 \\ \rho/3 & 1 & 0 & 0 \\ \rho/3 & 0 & 1 & 0 \\ \rho/3 & 0 & 0 & 1 \end{bmatrix} \qquad (15)$$

with $\rho = 0.05$. The observation measurements y_k are assumed *i.i.d.* with marginal probability densities $b^1(y) = \psi(y - L^1)$, $b^2(y) = \psi(y - L^2)$, $b^3(y) = \psi(y - L^3)$ and $b^4(y) = \psi(y - L^3)$ where $L \triangleq [L^1, L^2, L^3, L^4]' \in R^4$ is a vector of the different means for pre-change and post-change densities and $\psi(\cdot)$ is zero-mean Gaussian probability density function with variance $\sigma^2 = 1$. For simplicity of presentation, we assume the means L^i are increasing in the sense that $L^{i+1} > L^i$, for $i \in \{1, \dots, 3\}$.

Note for this class of models, the post-change density $b^2(\cdot)$ will be closest to the pre-change density $b^1(\cdot)$ in a relative entropy sense and $b^2(\cdot)$ will also be a weak post-change density bound (due to properties of relative entropies between Gaussians). Then note Proposition 5 of [10] can be used to establish that $D\left(b^1(y_k)\middle|\middle|b^2(y_k)\right) < \log(1/(1-\rho))$ whenever $L^2 - L^1 < 0.32$ (that is, when less than $\sqrt{2\log(1/(1-\rho))}$).

We consider an illustrative example with $L = [1.0, 1.1, 1.5, 2]$ which should exhibit the phenomenon, as $L^2 - L^1 = 0.1 < 0.32$.

We generated an example 5000 time step measurement sequence in which no change occurs, which was processed via (6) to compute the change posterior vector. We compare the calculated no change and post-change posteriors in Figure 2 illustrating the undesirable phenomenon in which no change posterior $\hat{X}_k^1$ becomes increasingly confident a change has occurred (with a corresponding increasing post-change posterior $\hat{X}_k^2$) even though no change has occurred.

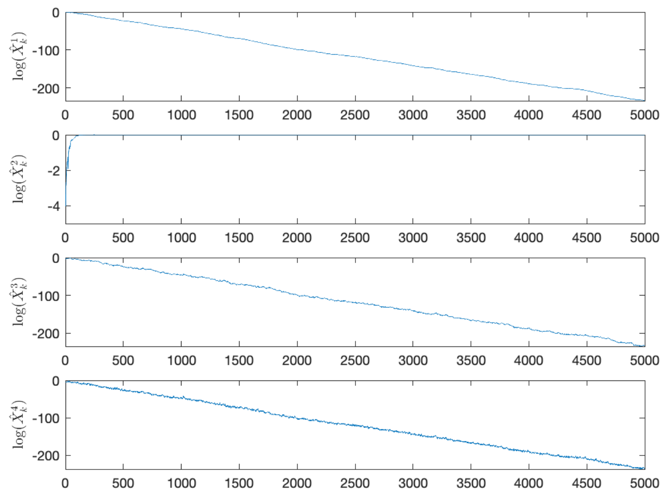

Fig. 2. An illustrative example of undesirable posterior hallucination behaviour in $N = 4$ QCDI. The 4 sub-figures are the calculated no change and post-change posterior on a 5000-point sequence having no change event. The no change posterior $\hat{X}_k^1$ (top sub-figure) has become increasing confidence a change has occurred, even though no change has occurred.

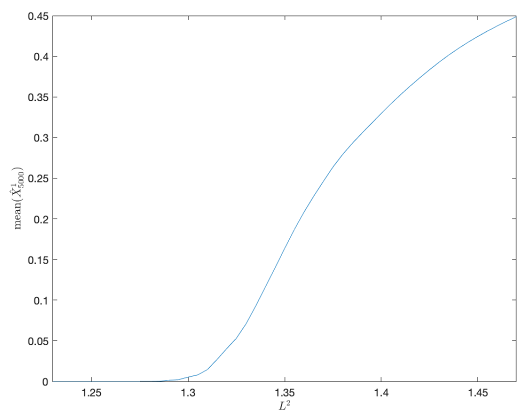

Fig. 3. Monte Carlo study of variation in the mean of the closest post-change density. The mean of $\hat{X}_{5000}^1$ against variation in L^2 the mean of closest post-change density $b^2(\cdot)$.

B. Study of the weak post-change density bound

Next, we examined the nature of the phenomenon via a study of the impact of variation of the mean of the closest post-change density $b^2(\cdot)$. We conducted a Monte Carlo study with 1000 trials of 5000 time steps with no change event, and varied the closest post-change mean L^2 between 1.23 and 1.47. Figure 3 shows the average value of $\hat{X}_{5000}^1$ (no change posterior after 5000 steps), and when the closest post-change mean $L^2 < 1.3$ illustrates the emergence of the hallucination phenomenon of incorrect confidence a change has occurred, consistent with Proposition 1.

C. Application: Vision-based aircraft detection

We now investigate a real-world application that highlights design considerations caused by the phenomenon in a practical setting. As described in [3]–[5], [16] the important vision-based aircraft detection application aims to quickly detect, with low false alarms, an aircraft on a near collision course after it visually emerges, at ranges $> 2km$. Figure 4 gives an illustration of the weak nature of the signal environment at the desired detection ranges. We will investigate the QCDI no change posterior on a data sequence, Case T1, captured during the flight tests described in [16].

Similar to [4], [5], for $k \geq 0$, we let each pixel in the image correspond to a possible change location where the aircraft could emerge and we introduce an extra state to denote when the aircraft is not visually apparent anywhere in the image frame (a no change state). Hence for an image of size $M = 1024 \times 768$ pixels, our Markov chain is $X_k \in \{e_1, \dots, e_N\}$, where $N = M + 1$. We construct our state transition probabilities A such that state transitions from edge pixels that would cross the image boundary self-transition. The no aircraft state has a probability of self-transition of $1 - \rho = 0.9$ (i.e. $\rho = 0.1$) and is equally likely

Fig. 4. Example: weak measurement in the vision-based aircraft detection.

to transition to any of the image pixels.

For $k \geq 0$, y_k are noise corrupted morphologically processed greyscale images of above horizon scene potentially containing an aircraft (raw images are pre-processed by the bottom-hat morphological operation as described in [3], [17]). Similarly to the current state-of-the-art, we use (unnormalised) observation densities $b^1(y_k) = 1$ and $b^i(y_k) = 1 + y_k^i$ for $i \in \{2, \ldots, N\}$ (see [3], [17] for detailed explanation and justification of this observation model). As reported in [4], [16], an aircraft was successfully detected in this Case T1 sequence at ranges exceeding $2km$. Importantly, in [5] an intermittent signal technique is shown to offer performance improvement over state-of-the-art detection by introducing a misidentification cost.

Here we examine the performance of the QCDI technique for the same purpose. Figure 5 shows the QCDI no change posterior $\hat{X}_k^1$ evolution from the start of the Case T1 data sequence (during a period that no aircraft is present). The QCDI's no change posterior exhibits the weak practical super-martingale and the post-change posteriors collected at the pixel states corresponding to the edges of the image.

The observed no change posterior $\hat{X}_k^1$ behaviour is consistent with Proposition 1 and consistent with previous observations about the challenging nature of the vision-based aircraft detection problem [3], [17], and highlights that Bayesian QCDI might be unsuitable for this detection application and other weak signal environments. Furthermore, these observations explain the practical importance of the intermittent signal QCDI problems posed in [5].

V. Conclusion

This paper established that a previously unreported undesirable phenomenon can occur in Bayesian QCDI. The potential phenomenon arises when a post-change density is too close to the pre-change density in a relative entropy sense. That this undesirable hallucination phenomenon can occur in insufficiently informative environments is an important design consideration in some practical settings.

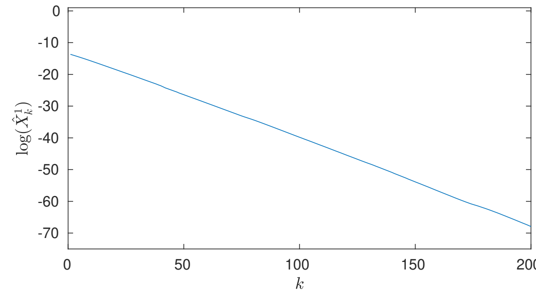

Fig. 5. The no change posterior $\hat{X}_k^1$ exhibits the hallucination phenomenon even though no change has occurred (i.e. aircraft not appeared).

References

[1] I. Hwang, S. Kim, Y. Kim, and C. E. Seah, "A Survey of Fault Detection, Isolation, and Reconfiguration Methods," *IEEE Transactions on Control Systems Technology*, vol. 18, no. 3, pp. 636–653, May 2010.

[2] A. G. Tartakovsky, A. S. Polunchenko, and G. Sokolov, "Efficient computer network anomaly detection by changepoint detection methods," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 1, pp. 4–11, Feb 2013.

[3] J. Lai, J. J. Ford, L. Mejias, and P. O'Shea, "Characterization of sky-region morphological-temporal airborne collision detection," *Journal of Field Robotics*, vol. 30, no. 2, pp. 171–193, Mar 2013.

[4] J. James, J. J. Ford, and T. L. Molloy, "Quickest detection of intermittent signals with application to vision-based aircraft detection," *IEEE Transactions on Control Systems Technology*, Oct 2018.

[5] ——, "A novel technique for rejecting non-aircraft artefacts in above horizon vision-based aircraft detection," in *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2020, pp. 600–606.

[6] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes: Theory and Application*. Prentice Hall, 1993, vol. 2.

[7] D. P. Malladi and J. L. Speyer, "A generalized Shiryaev sequential probability ratio test for change detection and isolation," *IEEE Transactions on Automatic Control*, vol. 44, no. 8, pp. 1522–1534, Aug 1999.

[8] S. Dayanik and C. Goulding, "Sequential detection and identification of a change in the distribution of a Markov-modulated random sequence," *IEEE Transactions on Information Theory*, vol. 55, no. 7, pp. 3323–3345, July 2009.

[9] V. Krishnamurthy, *Partially Observed Markov Decision Processes*. Cambridge University Press, 2016.

[10] J. J. Ford, J. James, and T. L. Molloy, "On the informativeness of measurements in Shiryaev's Bayesian quickest change detection," *Automatica*, vol. 111, p. 108645, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0005109819305060

[11] J. J. Ford, J. M. Kennedy, C. Tompkins, J. James, and A. McFadyen, "Exactly optimal quickest change detection of Markov chains," *IEEE Control Systems Letters*, vol. 7, pp. 2749–2754, 2023.

[12] R. Elliott, L. Aggoun, and J. Moore, *Hidden Markov Models: Estimation and Control*. Springer-Verlag, 1995.

[13] T. L. Molloy and J. J. Ford, "Misspecified and asymptotically minimax robust quickest change detection," *IEEE Transactions on Signal Processing*, vol. 65, no. 21, pp. 5730–5742, November 2017.

[14] ——, "Asymptotic minimax robust quickest change detection for dependent stochastic processes with parametric uncertainty," *IEEE Transactions on Information Theory*, vol. 62, no. 11, pp. 6594–6608, 2016.

[15] J. Unnikrishnan, V. V. Veeravalli, and S. P. Meyn, "Minimax robust quickest change detection," *IEEE Transactions on Information Theory*, vol. 57, no. 3, pp. 1604–1614, 2011.

[16] D. Bratanov, L. Mejias, and J. J. Ford, "A vision-based sense-and-avoid system tested on a ScanEagle UAV," in *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, Jun 2017, pp. 1134–1142.

[17] J. Lai, L. Mejias, and J. J. Ford, "Airborne vision-based collision-detection system," *Journal of Field Robotics*, vol. 28, no. 2, pp. 137–157, Mar 2011.