Chalmers Publication Library

Asymmetric Threat Modeling Using HMMs: Bernoulli Filtering and Detectability Analysis

(article starts on next page)

Asymmetric Threat Modeling Using HMMs: Bernoulli Filtering and Detectability Analysis

Karl Granström, *Member, IEEE*, Peter Willett, *Fellow, IEEE*, and Yaakov Bar-Shalom, *Fellow, IEEE*

Abstract—There is good reason to model an asymmetric threat (a structured action such as a terrorist attack) as an HMM whose observations are cluttered. Within this context this paper presents two important contributions. The first is a Bernoulli filter that can process cluttered observations and is capable of detecting if there is an HMM present, and if so, estimate the state of the HMM. The second is an analysis of the problem that, for a given HMM model, is able to make statements regarding the minimum complexity that an HMM would need to involve in order that it be detectable with reasonable fidelity, as well as upper bounds on the level of clutter (expected number of false measurements) and probability of miss of a relevant observation. In a simulation study the Bernoulli filter is shown to give good performance provided that the probability of observation is larger than the probability of an irrelevant clutter observation. Further, the results show that the longer the delays are between the HMM state transitions, the larger the probability margin must be. The feasibility prediction shows that it is possible to predict the boundary between poor performance and good performance for the Bernoulli filter, i.e., it is possible to predict when the Bernoulli filter will be useful, and when it will not be.

I. INTRODUCTION

The term *asymmetric threat* refers to tactics employed by, e.g., terrorist groups to carry out attacks on a superior opponent, while trying to avoid direct confrontation. Terrorist groups are elusive, secretive, amorphously structured decentralized entities that often appear unconnected. Analysis of prior terrorist attacks suggests that a high magnitude terrorist attack requires certain enabling events to take place.

In this paper terrorist activites are modeled using Hidden Markov Models (HMMs). In previous work HMMs have been shown to provide powerful statistical techniques, and they have been applied to various problems such as speech recognition, DNA sequence analysis, robot control, fault diagnosis, and signal detection, to name a few. Excellent tutorials on HMMs can be found in [16], [17]. The applicability of HMMs for terrorist activity modeling and other national security problem situations has been illustrated in previous work, see e.g. [5], [21], [23]–[25], [27]. For example, [5] uses HMMs to identify groups with suspicious behaviour, and [21] uses HMMs for pattern recognition of international crises.

A number of different terrorist plan HMMs are proposed in [23]–[25], [27], including models for a truck bombing, a

Karl Granström is with the Department of Signals and Systems, Chalmers University of Technology, Gothenburg, Sweden. E-mail: karl.granstrom@chalmers.se.

Yaakov Bar-Shalom and Peter Willett are with the Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT, USA. E-mail: {ybs,willett}@engr.uconn.edu.

Fig. 1. Markov chain network modeling the planning of a truck bombing. S_1: Selection of targets and reconnaissance. S_2: Set up A_1 cell. S_3: Set up A cell. S_4: Acquire money for operation. S_5: Gather resources. S_6: Expert arrives to assemble bombs. S_7: Target reconnaissance. S_8: Communications and final setup. S_9: Attack. The HMM is taken from [24]

plane hijacking, and production of weapons grade material. A relatively simple HMM that models a truck bombing is shown in Figure 1. These HMMs include multiple paths from plan initiation to plan completion, following the intuition that there are multiple ways to, e.g., hijack a plane. An empirical HMM can be constructed using available prior information, or with the help from experienced intelligence analysts [24]. For example, the HMM for *development of a nuclear weapons program* (DNWP) in [23] is gleaned using the open sources [3], [14], [22], [26], [28].

The basic motivation for modeling terrorist activities via HMMs is twofold. Firstly, carrying out a terrorist activity requires planning and preparations, following steps that form a pattern. This pattern of actions can be modeled using a Markov chain. Secondly, the terrorists leave detectable clues about these enabling events in the observation space. The clues are not direct observations of the planning and preparations, but are rather related to them, meaning that the states in the Markov model are hidden. For example, an observation of a purchase of chemicals could be indicative of intentions to produce a chemical weapon. However, a purchase of chemicals could very well be a benign event, which motivates inclusion of a model of observations that are unrelated to the HMM. Following the target tracking literature, see e.g. [2], such observations are here designated as clutter observations.

Ultimately the task is to find out if there is an activity being planned, and if so, find what stage the planning is in. To this end, in this paper we present a simple and efficient filter that estimates simultaneously the probability of HMM existence and the probability mass function for the HMM state, and we also analyze the detectability of the problem. In doing so we draw upon both signal detection theory, see e.g. [15], and target tracking theory, see e.g. [2]. Note that in [23]–[25], [27] (especially [25]) a CUSUM test was used to declare existence of an ongoing activity. This is reasonable, but does not scale well when multiple activities (perhaps some of them benign) are extant. In fact, that situation is reminiscent of multiple-

target tracking (MTT), and the detection problem suggestive of track-before-detect (TBD). As such, in this paper we are more interested in MTT that uses TBD: the approach we proffer *is* scalable to multiple-HMM adjudication.

In the next section we give a problem formulation and list the contributions. The rest of the paper is organized as follows: Asymmetric threat modeling using HMMs is presented in Section III. A Bernoulli filter (for MTT with TBD) is presented in Section IV, and a detectability analysis is presented in Section V. The results from a simulation study are presented in Section VIII. The paper is concluded in Section IX.

II. PROBLEM FORMULATION AND CONTRIBUTIONS

Consider a sequence of time steps t_1 to t_N. For some of the time steps there are observations (at most one per time step), for others there are not (denoted $\emptyset$), e.g.:

$$
\begin{array}{llccccccc}
\text{Time:} & \ldots & t_{k-2} & t_{k-1} & t_k & t_{k+1} & t_{k+2} & \ldots \\
\text{Obs.:} & \ldots & \mathbf{z}_{k-2} & \emptyset & \emptyset & \mathbf{z}_{k+1} & \emptyset & \ldots
\end{array}
\tag{1}
$$

The observations $\mathbf{z}$ are intelligence data; specifically they are intelligence reports that belong to a set of a priori known categories. These categories can, e.g., correspond to statements such as the following:

1) Persons seen repeatedly around sensitive location; possible reconnaissance of area/object of interest
2) Transfer of funds from country X to country Y
3) A person has signed a lease for an appartment or house in city X in country Y
4) Purchase of material that can be used to assemble bomb(s)
5) Arrival of Improvised Explosive Device (IED) expert
6) Collect information on design of bombs
7) Recruitment/training of new members
8) Setup base/cell in location X in country Y
9) Forged documents arranged for entry to country Y

The specific form (1) of the observation sequence is motivated further in Section III-B3.

We consider the following two hypotheses

H_0: The observations were generated by a clutter process. This means that there is no structure in the sequence of observations, it is random.

H_1: The observations were generated by an HMM-in-clutter process. This means that among the random clutter observations, there are observations caused by the HMM that has some degree of structure.

The problem is to decide between the two hypotheses, and, if an HMM exists (H_1), to estimate the HMM state probability mass function (pmf). To this end the paper contains the following two main contributions:

1) A Bernoulli filter that can process a sequence of observations and detect if there is an activity being planned and organized, and if so, what stage of planning the activity is in. An early version of this work is presented in [7].
2) A detectability analysis that is used to make "back-of-the-envelope" predictions about when the problem is feasible at all. An early version of this work is presented in [8].

The scope of the paper is limited by the assumption that the transition function and observation function of the HMM are known. Note however that no assumptions are made regarding the HMM's existence, nor regarding measurement origin (HMM or clutter). Work regarding the modeling of the HMM transition and observation functions can be found in, e.g., [5], [21], [23]–[25], [27].

III. ASYMMETRIC THREAT MODELING

We cast detection and estimation of an HMM using cluttered observations as a single target detection and tracking problem. Single target detection and state estimation using cluttered observations is well studied in the target tracking literature. In this work we will use Finite Set Statistics (FISST) and Random Finite Set (RFS) theory to model the problem, specifically the so called Bernoulli RFSs. Tutorials of random set methods are given in [9], [12], [29], with in-depth descriptions of FISST and RFS found in [11], [13]. A tutorial introduction to Bernoulli filters is given in [18]. In previous work these methods have typically been applied to problems where both the state and the observations are continuous random variables, in contrast to the work here where the states and observations are discrete.

A. State, transitions, observations

Let $\mathbf{s}_k \in \mathcal{S}$ denote the HMM state at time t_k, where $\mathcal{S}$ is a discrete state space with N_s states, $\mathcal{S} = \{S_1, S_2, \ldots, S_{N_s}\}$. Further, let $\mathbf{t}_k \in \mathcal{T} = \{0, 1\}$ denote the transition state, defined as $\mathbf{t}_k = 1$ if $\mathbf{s}_k \neq \mathbf{s}_{k-1}$ and $\mathbf{t}_k = 0$ otherwise. The state transitions are important because in the variant of HMMs used here the observations become available only upon state transitions. The auxiliary transition variable $\mathbf{t}$ is used because the authors found that it simplifies mathematical analysis and implementation. Let $\zeta_k = (\mathbf{s}_k, \mathbf{t}_k)$ denote the joint variable.

For the joint transition probability $\pi(\zeta_k|\zeta_{k-1}) = \pi(\mathbf{s}_k, \mathbf{t}_k|\mathbf{s}_{k-1}, \mathbf{t}_{k-1})$ the following holds

$$
\pi(\zeta_k|\zeta_{k-1}) = \pi(\mathbf{t}_k|\mathbf{s}_k, \mathbf{s}_{k-1}, \mathbf{t}_{k-1})\pi(\mathbf{s}_k|\mathbf{s}_{k-1})
\tag{2}
$$

The HMM state transitions follow a first order Markov chain with transition probability $\pi(\mathbf{s}_k|\mathbf{s}_{k-1})$. For the transition state $\mathbf{t}_k$ the transition matrix is

$$
\Pi = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix} \text{ if } \mathbf{s}_k \neq \mathbf{s}_{k-1}; \quad \Pi = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix} \text{ otherwise.}
\tag{3}
$$

The observations $\mathbf{z}_k \in \mathcal{Z}$ are discrete random variables, where $\mathcal{Z}$ is a discrete measurement space with N_z measurements, $\mathcal{Z} = \{Z_1, Z_2, \ldots, Z_{N_z}\}$. The observations Z_i correspond to the different intelligence data categories; a list of examples was given in Section II. With a state dependent probability of detection

$$
p_{\mathrm{D}}(\zeta_k) = \begin{cases} p_{\mathrm{D}}^0 \in (0,1) & \text{if } \mathbf{t}_k = 1, \\ 0 & \text{otherwise,} \end{cases}
\tag{4}
$$

the HMM generates an observation $\mathbf{z}_k$. The observation process is defined by the likelihood

$$
h(\mathbf{z}_k|\zeta_k) = g_s(\mathbf{z}_k|\mathbf{s}_k)
\tag{5}
$$

There are also clutter observations (false alarms) super-imposed on the true HMM observations. In each time-step, with probability $0 < p_{\text{FA}} < 1$ a clutter observation is generated as a random sample from a process with probability mass function (pmf) $g_{\text{FA}}(\mathbf{z}_k)$.

B. Random finite set modeling

An RFS is a random variable whose realizations are sets with a finite cardinality (number of elements). The cardinality, and each element, are all random variables. Specifically, a Bernoulli RFS $\mathbf{X}$ is either an empty set, with probability $1 - q$, or has a single element, with probability q. In case there is an element $\mathbf{x}$, it is distributed over the state space $\mathcal{X}$ according to the probability mass function (pmf) $P(\mathbf{x})$. The FISST probability density function (pdf) of $\mathbf{X}$ is

$$f(\mathbf{X}) = \begin{cases} 1 - q & \text{if } \mathbf{X} = \emptyset \\ q \cdot P(\mathbf{x}) & \text{if } \mathbf{X} = \{\mathbf{x}\} \\ 0 & \text{if } |\mathbf{X}| \geq 2 \end{cases} \quad (6)$$

The state space for $\mathbf{X}$ is $\emptyset \cup \sigma(\mathcal{X})$, where $\sigma(\mathcal{X})$ is the set of all singletons $\{\mathbf{x}\}$ such that $\mathbf{x} \in \mathcal{X}$. A singleton is a set with cardinality one. For a Bernoulli pdf a set integral is defined as follows [11],

$$\int f(\mathbf{X})\delta\mathbf{X} = f(\emptyset) + \int f(\{\mathbf{x}\})\mathrm{d}\mathbf{x} \quad (7a)$$

$$= 1 - q + \int q P(\mathbf{x})\mathrm{d}\mathbf{x} = 1, \quad (7b)$$

and it follows that $f(\mathbf{X})$ as defined in (6) is indeed a proper pdf. Note that integrals over the discrete random variable $\mathbf{x}$ are sums, e.g.

$$\int m(\mathbf{x})P(\mathbf{x})\mathrm{d}\mathbf{x} = \sum_{X \in \mathcal{X}} m(\mathbf{x} = X)P(\mathbf{x} = X) \quad (8)$$

for a function $m(\mathbf{x})$, however for brevity we will use the integral notation rather than the sum notation.

1) HMM state model: The joint HMM state ζ_k at time t_k is modeled as a Bernoulli RFS $\mathbf{S}_k$. The state space is $\emptyset \cup \sigma(\mathcal{S} \times \mathcal{T})$, where $\sigma(\mathcal{S} \times \mathcal{T})$ is the set of all singletons $\{\mathbf{s}, \mathbf{t}\}$ such that $\mathbf{s} \in \mathcal{S}$ and $\mathbf{t} \in \mathcal{T}$. Let $\varepsilon_k \in \{0, 1\}$ denote the existence of the HMM: if $\varepsilon_k = 1$ the HMM exists at time t_k, i.e., $\mathbf{S}_k \neq \emptyset$, and if $\varepsilon_k = 0$ the HMM does not exists at time t_k, i.e., $\mathbf{S}_k = \emptyset$. Note that $\varepsilon_k = 0$ corresponds to H_0 being true, and $\varepsilon_k = 1$ corresponds to H_1 being true.

The posterior FISST pdf at time step t_k for the Bernoulli RFS $\mathbf{S}_k$, denoted as

$$f_{k|k}(\mathbf{S}_k|\mathbf{Z}^k) = \begin{cases} 1 - q_{k|k} & \text{if } \mathbf{S}_k = \emptyset \\ q_{k|k} \cdot P_{k|k}(\zeta) & \text{if } \mathbf{S}_k = \{\zeta\} \\ 0 & \text{if } |\mathbf{S}_k| \geq 2 \end{cases} \quad (9)$$

is specified by the posterior probability of HMM existence and the posterior pmf of the joint HMM state of ζ_k,

$$q_{k|k} = P(|\mathbf{S}_k| = 1|\mathbf{Z}^k), \quad (10a)$$

$$P_{k|k}(\zeta) = P(\zeta_k|\mathbf{Z}^k). \quad (10b)$$

2) Dynamics model: The dynamics of HMM existence ε_k are modeled as a first order Markov chain with transition probability matrix

$$P^\varepsilon_{k|k-1} = \begin{bmatrix} (1 - p_b) & p_b \\ (1 - p_s) & p_s \end{bmatrix}. \quad (11)$$

The probability of HMM birth

$$p_b = P(\varepsilon_k = 1|\varepsilon_{k-1} = 0) \quad (12)$$

models the probability that at time t_k a plan is initiated. The probability of HMM survival

$$p_s = P(\varepsilon_k = 1|\varepsilon_{k-1} = 1) \quad (13)$$

models the probability that an initiated plan is not cancelled. If an HMM is intiated at time t_k the initial pmf is denoted $P^b_{k|k-1}(\zeta)$.

The dynamic model of the RFS $\mathbf{S}$ is a Markov process with transition density $P^{\mathbf{S}}_{k|k-1}(\mathbf{S}_k|\mathbf{S}_{k-1})$,

$$P^{\mathbf{S}}_{k|k-1}(\mathbf{S}_k|\emptyset) = \begin{cases} 1 - p_b & \text{if } \mathbf{S}_k = \emptyset \\ p_b \cdot P^b_{k|k-1}(\zeta_k) & \text{if } \mathbf{S}_k = \{\zeta_k\} \\ 0 & \text{if } |\mathbf{S}_k| \geq 2 \end{cases} \quad (14a)$$

$$P^{\mathbf{S}}_{k|k-1}(\mathbf{S}_k|\{\zeta_{k-1}\}) = \begin{cases} 1 - p_s & \text{if } \mathbf{S}_k = \emptyset \\ p_s \cdot \pi(\zeta_k|\zeta_{k-1}) & \text{if } \mathbf{S}_k = \{\zeta_k\} \\ 0 & \text{if } |\mathbf{S}_k| \geq 2 \end{cases} \quad (14b)$$

$$P^{\mathbf{S}}_{k|k-1}(\mathbf{S}_k|\mathbf{S}_{k-1}) \text{ is undefined if } |\mathbf{S}_{k-1}| \geq 2 \quad (14c)$$

where the joint transition probability $\pi(\zeta_k|\zeta_{k-1})$ was defined in (2). The transition density (14) can be understood as follows:

a) If at time $k - 1$ there is no HMM ($\mathbf{S}_{k-1} = \emptyset$), then
   - with probability $1 - p_b$ there will not be any HMM at time k ($\mathbf{S}_k =$); this is the first alternative in (14a)
   - with probability p_b an HMM will come into existence ($\mathbf{S}_k = \{\zeta_k\}$), and the pmf will be the birth pmf $P^b_{k|k-1}(\zeta_k)$; this is the second alternative in (14a)
   - the probability of more than one HMM is zero; this is the third alternative in (14a). This means that if no HMM exists, at most one new HMM can appear.

b) If at time $k - 1$ there is an HMM ($\mathbf{S}_{k-1} = \{\zeta_{k-1}\}$), then
   - with probability $1 - p_s$ the HMM will cease to exist ($\mathbf{S}_k = \emptyset$); this is the first alternative in (14b).
   - with probability p_s the HMM will remain in existence ($\mathbf{S}_k = \{\zeta_k\}$), and the state evolution from ζ_{k-1} to ζ_k is described by the joint transition probability; this is the second alternative in (14b).
   - the probability of more than one HMM is zero; this is the third alternative in (14b). This means that a new target cannot appear if the HMM persists into the next time step.

c) There cannot be more than one HMM at the same time, so the RFS transition density $P^{\mathbf{S}}_{k|k-1}(\mathbf{S}_k|\mathbf{S}_{k-1})$ is undefined for the case $|\mathbf{S}_{k-1}| \geq 2$, see (14c).

3) Observations: We model the sequence of observations as a sequence of RFS observations $\mathbf{Z}_k$, where k refers to time t_k. For the example in (1) we then have

$$\begin{aligned} \mathbf{Z}_{k-2} &= \{\mathbf{z}_{k-2}\} \\ \mathbf{Z}_{k+1} &= \{\mathbf{z}_{k+1}\} \\ \mathbf{Z}_{k-1} &= \mathbf{Z}_k = \mathbf{Z}_{k+2} = \emptyset \end{aligned} \tag{15}$$

Let $\mathbf{Z}^k$ denote all such observations from time t_1 to t_k, $\mathbf{Z}^k = \{\mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_k\}$. If $\varepsilon_k = 0$ then $\mathbf{Z}_k = \mathbf{C}_k$ and if $\varepsilon_k = 1$ then $\mathbf{Z}_k$ is modeled as the union of two independent RFS,

$$\mathbf{Z}_k = \mathbf{W}_k \cup \mathbf{C}_k, \tag{16}$$

where $\mathbf{W}_k$ is HMM generated observations and $\mathbf{C}_k$ is clutter observations. The model (16) allows both $\mathbf{W}_k$ and $\mathbf{C}_k$ to be non-empty simultaneously, which goes against the assumption in (1) that there is at most one observation per time step. We comment on this modeling choice in the remark at the end of this section.

The clutter observations are modeled as a Bernoulli RFS with FISST pdf

$$\kappa(\mathbf{Z}) = \begin{cases} 1 - p_{\text{FA}} & \text{if } \mathbf{Z} = \emptyset \\ p_{\text{FA}} \cdot g_{\text{FA}}(\mathbf{z}) & \text{if } \mathbf{Z} = \{\mathbf{z}\} \\ 0 & \text{if } |\mathbf{Z}| \geq 2 \end{cases} \tag{17}$$

At timesteps for which an HMM exists ($\varepsilon_k = 1$, $\mathbf{S}_k = \{\zeta_k\}$), the HMM generated observations are modeled as a Bernoulli RFS with FISST pdf

$$\eta(\mathbf{Z}|\{\zeta\}) = \begin{cases} 1 - p_{\text{D}}(\zeta) & \text{if } \mathbf{Z} = \emptyset \\ p_{\text{D}}(\zeta) \cdot g_s(\mathbf{z}|\zeta) & \text{if } \mathbf{Z} = \{\mathbf{z}\} \\ 0 & \text{if } |\mathbf{Z}| \geq 2 \end{cases} \tag{18}$$

The observation likelihood function is denoted $\varphi(\mathbf{Z}|\mathbf{S})$ and has two forms: one for $\mathbf{S} = \emptyset$ and one for $\mathbf{S} = \{\zeta\}$. In the former case we have the FISST pdf

$$\varphi(\mathbf{Z}|\emptyset) = \kappa(\mathbf{Z}) \tag{19}$$

and in the latter case, for a union of independent RFS as in (16), the FISST pdf is

$$\varphi(\mathbf{Z}|\{\zeta\}) = \sum_{\mathbf{W} \subseteq \mathbf{Z}} \eta(\mathbf{W}|\{\zeta\})\kappa(\mathbf{Z}\backslash\mathbf{W}) \tag{20}$$

where $\backslash$ denotes set difference. For both the clutter observations and the HMM observations the set can be either empty or singletons, and thus the union (16) can have zero, one or two elements. The summation then has three different cases

$$\varphi(\mathbf{Z}|\{\zeta\}) = \tag{21}$$
$$\begin{cases} \eta(\emptyset|\{\zeta\})\kappa(\emptyset) & \text{if } \mathbf{Z} = \emptyset, \\ \eta(\emptyset|\{\zeta\})\kappa(\mathbf{z}) + \eta(\mathbf{z}|\{\zeta\})\kappa(\emptyset) & \text{if } \mathbf{Z} = \{\mathbf{z}\}, \\ \eta(\mathbf{z}^1|\{\zeta\})\kappa(\mathbf{z}^2) + \eta(\mathbf{z}^2|\{\zeta\})\kappa(\mathbf{z}^1) & \text{if } \mathbf{Z} = \{\mathbf{z}^1, \mathbf{z}^2\} \end{cases}$$

Remark: The observation sequence in (1) allows at most one observation per time step; this implies that the HMM observation set $\mathbf{W}_k$ and clutter set $\mathbf{C}_k$ are correlated. However, the observation model (16) allows the two sets to be simultaneously non-empty; this implies sets that are independent, which is a slightly inaccurate representation of (1). The model (16) is adopted because it allows statistical independence between

$\mathbf{W}_k$ and $\mathbf{C}_k$ to be assumed. If the modeling is restricted to observation set cardinality $|\mathbf{Z}| \leq 1$, then the observation probabilities have to fulfill $p_{\text{D}}^0 + p_{\text{FA}} \leq 1$. By allowing $|\mathbf{Z}| = 2$ the observation probabilities p_{D}^0 and p_{FA} can be chosen freely in the interval $(0,\ 1)$, and the modeling becomes more general. Further, the model (16) subsumes (1) as a special case.

Our justification for the "single-observation-per-time-step" assumption is that typical surveillance systems have observations that are entered one after the other, i.e., time has been discretized in such a way that in each "time-bin" there is at most one observations. However, the Bernoulli filter formulation would seem to be amenable to relaxation of this (perhaps even to be receptive to it) and exploration of a system where observations are taken (say) once per minute will be explored later. The time step must, however, be chosen such that only single HMM transitions are feasible. $\qquad\square$

IV. BERNOULLI FILTER

In this section, using the models defined above, we propose a filter within the RFS framework such that a joint estimate of the probability of HMM existence and of the HMM state distribution (pmf) are obtained.

Assume that the predicted FISST pdf $f_{k|k-1}(\mathbf{S}_k|\mathbf{Z}^{k-1})$ is known. The predicted FISST pdf at time step t_k for the Bernoulli RFS $\mathbf{S}_k$ is specified by the posterior probability of HMM existence and the posterior pmf of the joint HMM state of ζ_k,

$$q_{k|k-1} = P(|\mathbf{S}_k| = 1|\mathbf{Z}^{k-1}), \tag{22a}$$
$$P_{k|k-1}(\zeta) = P(\zeta_k|\mathbf{Z}^{k-1}). \tag{22b}$$

The Bernoulli filter (BF) propagates both quantities over time using an update equation

$$f_{k|k}(\mathbf{S}_k|\mathbf{Z}^k) = \frac{\varphi(\mathbf{Z}_k|\mathbf{S}_k)f_{k|k-1}(\mathbf{S}_k|\mathbf{Z}^{k-1})}{f(\mathbf{Z}_k|\mathbf{Z}^{k-1})}, \tag{23}$$

and a prediction equation

$$\begin{aligned} f_{k+1|k}&(\mathbf{S}_{k+1}|\mathbf{Z}^k) \\ &= \int P_{k+1|k}^{\mathbf{S}}(\mathbf{S}_{k+1}|\mathbf{S})f_{k|k}(\mathbf{S}|\mathbf{Z}^k)\delta\mathbf{S} \tag{24a} \\ &= P_{k+1|k}^{\mathbf{S}}(\mathbf{S}_{k+1}|\emptyset)(1 - q_{k|k}) \\ &\quad + q_{k|k} \int P_{k+1|k}^{\mathbf{S}}(\mathbf{S}_{k+1}|\{\zeta'\})P_{k|k}(\zeta')\mathrm{d}\zeta' \tag{24b} \end{aligned}$$

The likelihood $f(\mathbf{Z}_k|\mathbf{Z}^{k-1})$ in the update (23) is defined as

$$\begin{aligned} f(\mathbf{Z}_k|\mathbf{Z}^{k-1}) &= \int \varphi(\mathbf{Z}_k|\mathbf{S})f_{k|k-1}(\mathbf{S}|\mathbf{Z}^{k-1})\delta\mathbf{S} \tag{25a} \\ &= \kappa(\mathbf{Z}_k)(1 - q_{k|k-1}) \tag{25b} \\ &\quad + q_{k|k-1} \int \varphi(\mathbf{Z}_k|\{\zeta\})P_{k|k-1}(\zeta)\mathrm{d}\zeta \end{aligned}$$

A. Update

If the observation is an empty set $\mathbf{Z}_k = \emptyset$ we have

$$q_{k|k} = \frac{1 - \Delta_{k|k-1}^1}{1 - q_{k|k-1}\Delta_{k|k-1}^1}q_{k|k-1} \tag{26a}$$

$$P_{k|k}(\zeta) = \frac{1 - p_D(\zeta)}{1 - \Delta_{k|k-1}^1}P_{k|k-1}(\zeta) \tag{26b}$$

where $\Delta_{k|k-1}^1$ is defined in Appendix A. If the observation is a singleton $\mathbf{Z}_k = \{\mathbf{z}\}$ we have

$$q_{k|k} = \frac{1 - \Delta_{k|k-1}}{1 - q_{k|k-1}\Delta_{k|k-1}}q_{k|k-1} \tag{26c}$$

$$P_{k|k}(\zeta) = \frac{1 - p_D(\zeta)}{1 - \Delta_{k|k-1}}P_{k|k-1}(\zeta) \tag{26d}$$
$$+ \frac{1 - p_{FA}}{p_{FA}g_{FA}(\mathbf{z})}\frac{p_D(\zeta)g_s(\mathbf{z}|\zeta)}{1 - \Delta_{k|k-1}}P_{k|k-1}(\zeta)$$

where $\Delta_{k|k-1}$ is defined in Appendix A. For the sake of completeness, if the observation is $\mathbf{Z}_k = \{\mathbf{z}^1, \mathbf{z}^2\}$ we have

$$q_{k|k} = 1 \tag{26e}$$

$$P_{k|k}(\zeta) = \frac{g_{FA}(\mathbf{z}^1)p_D(\zeta)g_s(\mathbf{z}^2|\zeta)P_{k|k-1}(\zeta)}{g_{FA}(\mathbf{z}^1)G_{k|k-1}^s(\mathbf{z}^2) + g_{FA}(\mathbf{z}^2)G_{k|k-1}^s(\mathbf{z}^1)} \tag{26f}$$
$$+ \frac{g_{FA}(\mathbf{z}^2)p_D(\zeta)g_s(\mathbf{z}^1|\zeta)P_{k|k-1}(\zeta)}{g_{FA}(\mathbf{z}^1)G_{k|k-1}^s(\mathbf{z}^2) + g_{FA}(\mathbf{z}^2)G_{k|k-1}^s(\mathbf{z}^1)}$$

where $G_{k|k-1}^s(\mathbf{z})$ is defined in Appendix A. Derivation details are given in Appendix B

B. Prediction

The predicted probability of HMM existence and the predicted joint HMM state pmf are

$$q_{k+1|k} = p_b(1 - q_{k|k}) + p_s q_{k|k}, \tag{27a}$$

$$P_{k+1|k}(\zeta) = \frac{p_b(1 - q_{k|k})}{q_{k+1|k}}P_{k+1|k}^b(\zeta) \tag{27b}$$
$$+ \frac{p_s q_{k|k}}{q_{k+1|k}}\int \pi(\zeta|\zeta')P_{k|k}(\zeta')d\zeta'$$

Derivation details are given in Appendix C.

V. DETECTABILITY ANALYSIS

Given the above Bernoulli filter, an important question to consider is: When is it even a good idea to use this filter? Or, more generally put, when is the HMM-in-clutter problem even feasible? In this section we will perform a detectability analysis of the problem, with the aim of being able to make statements about maximum levels of clutter allowable; maximum intervals between relevant observations; and a minimum level of complexity.

To decide between the hypotheses H_0 and H_1 we employ a decision rule

$$\widetilde{\delta}_\ell = \begin{cases} 1 & > \\ \gamma & \text{if } \ell = \tau \\ 0 & < \end{cases} \tag{28}$$

where ℓ is the natural logarithm of the likelihood ratio L, $\ell = \log(L)$. To analyze the detectability we will focus on the two conditional error probabilities P_F (probability of threat false alarm) and P_M (probability of threat miss) defined by

$$P_F(\widetilde{\delta}) = P_0(\widetilde{\delta} \text{ chooses } H_1) \tag{29a}$$
$$P_M(\widetilde{\delta}) = P_1(\widetilde{\delta} \text{ chooses } H_0) \tag{29b}$$

Given probability density functions (pdfs) $p(\ell(\mathbf{Z}^k)|H_0)$ and $p(\ell(\mathbf{Z}^k)|H_1)$ it is easy to compute the errors for given model parameters and a given threshold τ. However, expressing the pdfs $p(\ell(\mathbf{Z}^k)|H_i)$ analytically is prohibitively difficult and complex in the general case.

We will solve this problem by first making some simplifications and approximations such that the pdf $p(\ell(\mathbf{Z}^k)|H_1)$ can be approximated by a Gaussian density. Using this Gaussian approximation the probability of threat miss $P_M(\widetilde{\delta})$ can be computed exactly, and for the probability of threat false alarm $P_F(\widetilde{\delta})$ an upper bound can be found.

In the following subsections we will first give expressions for the likelihood ratio. We then present some model simplifications and likelihood approximations that are necessary to model the log-likelihood pdf under H_1. Finally, we give expressions for the conditional error probabilities: an exact probability of miss and an upper bound for the probability of false alarm.

A. Likelihood ratio

The log-likelihood ratio (LLR) is denoted

$$\ell(\mathbf{Z}^k) = \sum_{j=1}^{k} \ell(\mathbf{Z}_j) = \sum_{j=1}^{k}(\ell_1(\mathbf{Z}_j) - \ell_0(\mathbf{Z}_j)) \tag{30}$$

where $\ell_j(\cdot)$ is the log-likelihood function under hypothesis j. Let $L_j(\cdot) = \exp(\ell_j(\cdot))$ denote the likelihood of hypothesis j, defined in (25). For H_0 (an HMM does not exist) at time step t_j we have probability of HMM existence $q_{j|j-1} = 0$ and observation FISST pdf (19), and it follows that the likelihood (25) is

$$L_0(\mathbf{Z}_j) = \varphi(\mathbf{Z}_j|\emptyset) = \kappa(\mathbf{Z}_j) \tag{31}$$

For H_1 (an HMM exists) at time step t_j we have probability of HMM existence $q_{j|j-1} = 1$ and observation FISST pdf (20), and it follows that the likelihood (25) is

$$L_1(\mathbf{Z}_j) = \int \varphi(\mathbf{Z}_j|\zeta)P_{j|j-1}(\zeta)d\zeta \tag{32}$$

For the likelihood ratio

$$L(\mathbf{Z}_j) = \frac{L_1(\mathbf{Z}_j)}{L_0(\mathbf{Z}_j)} = \frac{\int \varphi(\mathbf{Z}_j|\zeta)P_{j|j-1}(\zeta)d\zeta}{\kappa(\mathbf{Z}_j)} \tag{33}$$

we have two cases. If the detection set is an empty set $\mathbf{Z}_j = \emptyset$ we have

$$L_j^\emptyset = \frac{\int \eta(\emptyset|\{\zeta\})\kappa(\emptyset)P_{j|j-1}(\zeta)d\zeta}{\kappa(\emptyset)} \tag{34a}$$
$$= 1 - p_D^0 P_{j|j-1}(\mathbf{t} = 1) \tag{34b}$$

where $L_j^{\emptyset} = L(\mathbf{Z}_j = \emptyset)$. If the detection set is a singleton $\mathbf{Z}_j = \{\mathbf{z}\}$ we have

$$L_j^{\mathbf{z}} = \frac{\int [\eta(\emptyset|\{\zeta\})\kappa(\mathbf{z}) + \eta(\mathbf{z}|\{\zeta\})\kappa(\emptyset)]\, P_{j|j-1}(\zeta)\mathrm{d}\zeta}{\kappa(\mathbf{z})} \quad (35a)$$

$$= 1 - p_{\mathrm{D}}^0 P_{j|j-1}(\mathbf{t} = 1)$$
$$+ \frac{1 - p_{\mathrm{FA}}}{p_{\mathrm{FA}} g_{\mathrm{FA}}(\mathbf{z})} p_{\mathrm{D}}^0 \int g_s(\mathbf{z}|\mathbf{s}) P_{j|j-1}(\mathbf{s}, \mathbf{t} = 1)\mathrm{d}\mathbf{s} \quad (35b)$$

where $L_j^{\mathbf{z}} = L(\mathbf{Z}_j = \{\mathbf{z}\})$.

B. Model simplifications and likelihood approximations

1) Daisy-chain HMM*:* The HMM is in the form of a "daisy-chain", i.e. the HMM state $\mathbf{s}$ can only transition to the next state or remain in the same state. Expressed in terms of the transition probability $\pi(\mathbf{s}_k|\mathbf{s}_{k-1})$, if $\mathbf{s}_k = S_i, 1 \le i < N_s$ then

$$\pi(\mathbf{s}_k = S_j|\mathbf{s}_{k-1} = S_i) = \begin{cases} 1 - P_T & j = i \\ P_T & j = i + 1 \\ 0 & \text{otherwise} \end{cases} \quad (36)$$

and if $\mathbf{s}_k = S_{N_s}$ then

$$\pi(\mathbf{s}_k = S_j|\mathbf{s}_{k-1} = S_i) = \begin{cases} 1 & \text{if } j = i \\ 0 & \text{otherwise} \end{cases} \quad (37)$$

2) Diagonal HMM *observation pmf:* The size of the observation space is equal to the HMM state space, i.e. $N_z = N_s$. The observation pmf is

$$g_s(\mathbf{z} = Z_j|\mathbf{s}_k = S_i) = \begin{cases} P_{obs} & \text{if } j = i \\ \frac{1 - P_{obs}}{N_z - 1} & \text{otherwise} \end{cases} \quad (38)$$

where $0 \ll P_{obs} \lesssim 1$ (i.e. P_{obs} is close to one).

Remark: This simplification means that each state has its own unique type of detection, and it is very unlikely that – given that there is a detection – a state would give the "wrong" type of detection. For example, let there be two states representing 1) that an apartment has been rented, and 2) that a large quantity of fertilizer has been bought. Given that the state is detected, we assume that it is unlikely that apartment rental will produce a true detection that fertilizer was bought, or vice versa. However, note that we do not make any assumptions regarding the probability that a state is detected. $\qquad\square$

3) Uniform clutter: The clutter is uniformly distributed

$$g_{\mathrm{FA}}(\mathbf{z}) = N_z^{-1} = N_s^{-1} \quad (39)$$

i.e. it is equiprobable for all the N_z possibilities.

4) Probability of state transition: Under the assumption that the HMM is a daisy-chain, the predicted marginal probability of state transition is

$$P_{j|j-1}(\mathbf{t} = 1) = P_T \left(1 - P_{j|j-1}(\mathbf{s} = S_{N_s})\right) \quad (40)$$

We approximate this as

$$P_{j|j-1}(\mathbf{t} = 1) \approx P_T. \quad (41)$$

i.e. the approximation is $P_{j|j-1}(\mathbf{s} = S_{N_s}) \approx 0$. Under this approximation it follows that the likelihood ratio at time steps for which there is no detection (34) is approximated as

$$L_j^{\emptyset} \approx 1 - p_{\mathrm{D}}^0 P_T \quad (42)$$

Remark: The approximation $P_{j|j-1}(\mathbf{s} = S_{N_s}) \approx 0$ is typically valid when the true state is not close to the last state S_{N_s}. In general the approximation is more accurate the more states the HMM has, i.e. the larger N_s is. $\qquad\square$

5) Observation likelihood: Assume that at time step t_m there is a detection $\mathbf{z}_m = Z_i$. For the observation likelihood function (38) the integral in the likelihood ratio (35) is

$$\int g_s(\mathbf{z} = Z_i|\mathbf{s}) P_{m|m-1}(\mathbf{s}, \mathbf{t} = 1)\mathrm{d}\mathbf{s}$$
$$= P_{obs} P_{m|m-1}(\mathbf{s} = S_i, \mathbf{t} = 1)$$
$$+ \frac{1 - P_{obs}}{N_z - 1} \left(1 - P_{m|m-1}(\mathbf{s} = S_i, \mathbf{t} = 1)\right) \quad (43)$$

We approximate

$$\frac{1 - P_{obs}}{N_z - 1} \approx 0 \quad (44)$$

and it follows that the likelihood ratio is approximated

$$L_m^{\mathbf{z}} \approx 1 - p_{\mathrm{D}}^0 P_T \quad (45)$$
$$+ \frac{1 - p_{\mathrm{FA}}}{p_{\mathrm{FA}} N_s^{-1}} p_{\mathrm{D}}^0 P_{obs} P_{m|m-1}(\mathbf{s} = S_i, \mathbf{t} = 1)$$

Remark: The approximation used here is valid under the model simplification above that P_{obs} is almost one. Further, in general the approximation is more accurate the larger N_s is, i.e. the more states there are in the HMM. $\qquad\square$

C. Approximation of log-likelihood pdf under H_1

When H_1 is true, an HMM exists and the detections $\mathbf{Z}^k$ are generated by both the HMM and the clutter process. For either type of detection, the likelihood ratio can be further simplified.

1) Clutter detection: Empirically it was found that if the detection is a clutter detection, the typical case is that $\mathbf{s}_m = S_i$ has low predicted probability ($P_{m|m-1}(\mathbf{s} = S_i, \mathbf{t} = 1) \approx 0$). In this case the likelihood ratio (45) can be further approximated as

$$L_m^{\mathbf{z}} \approx 1 - p_{\mathrm{D}}^0 P_T \quad (46)$$

Remark: The approximation $P_{m|m-1}(\mathbf{s} = S_i, \mathbf{t} = 1) \approx 0$ is more accurate the more states there are, i.e. the larger N_s is. Note that $(1 - p_{\mathrm{FA}})/p_{\mathrm{FA}}$ in (45) increases for decreasing p_{FA}, and this (together with an increaseing number of states N_s) will make the approximation less accurate. However, a very low p_{FA} ($< 1\%$) is of little practical interest, because in most realistic scenarios the probability of a clutter detection will not be close to zero. $\qquad\square$

2) HMM *detection:* Empirically it was found that if the detection $\mathbf{z}_m = Z_i$ was caused by the HMM, the approximation $P_{m|m-1}(\mathbf{s}_k = S_i, \mathbf{t}_k = 1) \approx 0$ typically does not hold. For this case the predicted probability is simplified as follows.

Assume that at time step $n < m$ there was a detection $\mathbf{z}_n = Z_{i-1}$ such that the posterior pmf $P_{n|n}(\zeta)$ indicates a high probability that the HMM is in state $\mathbf{s}_n = S_{i-1}$. Further, assume that in between time steps n and m there were no detections. Given the pmf $P_{n|n}(\mathbf{s}_n, \mathbf{t}_n)$ the sought after predicted probability $P_{m|m-1}(\mathbf{s}_m = S_i, \mathbf{t}_m = 1)$ is then given by first iterating prediction and no-detection-measurement update

for $N = m - n$ time steps, and then evaluating for $\mathbf{s} = S_i$ and $\mathbf{t} = 1$.

By approximating $P_{n|n}(\cdot)$ as follows,

$$P_{n|n}(\mathbf{s} = S_s, \mathbf{t} = t) \approx \begin{cases} 0.9 & \text{if } s = i-1, \quad t = 1 \\ 0.09 & \text{if } s = i-2, \quad t = 0 \\ 0.01 & \text{if } s = i-3, \quad t = 0 \\ 0 & \text{otherwise} \end{cases} \quad (47)$$

the probability $P_{m|m-1}(\mathbf{s}_m = S_i, \mathbf{t}_m = 1)$ can be approximately expressed as a function of the number of time steps N since the last measurement update. We denote this probability as $\bar{P}(N)$. Thus, the likelihood ratio (45) is approximated by

$$L_m^{\mathbf{z}} \approx 1 - p_{\text{D}}^0 P_T + \frac{1 - p_{\text{FA}}}{p_{\text{FA}} N_s^{-1}} p_{\text{D}}^0 P_{obs} \bar{P}(N) \quad (48)$$

Remark: The specific numerical values in (47) are motivated as follows: the majority of the probability mass is concentrated in the same state as the detection indicates. Some probability mass is contained in the two previous states — a reflection of the probability that the detection was a false alarm and the state $\mathbf{s}$ has not transitioned to S_{i-1} after all. Empirically we have found that these values are accurate for $p_{\text{FA}} > 10\%$. For $p_{\text{FA}} \approx 1\%$ almost all probability is concentrated in the state S_{i-1}, however such low p_{FA} are probably of little to no practical interest. Further, empirically we have found that the values are accurate for all p_{D}^0. □

3) Log-likelihood pdf approximation: Let N_t be the total number of time steps that it takes for the HMM to transition from the first to the last state. The HMM has to pass through each state in the daisy chain, meaning that $N_t \geq N_s$. The number of time steps the HMM state will remain in a specific state S_i (i.e. no state transition) is well known to be a random variable that is geometrically distributed with parameter P_{T}. The total number of time steps the HMM state will remain in the same state (i.e. no state transition) is therefore the sum of $N_s - 1$ geometrically distributed random variables, each with parameter P_{T}. A sum of $N_s - 1$ identically geometrically distributed random variables is well known to be negative binomial distributed with parameters $N_s - 1$ and P_{T}. Thus, for N_t we have the following pmf

$$P(N_t) = \begin{cases} 0 & N_t < N_s \\ \mathcal{NBIN}(N_t - N_s; N_s - 1, P_{\text{T}}) & \text{otherwise} \end{cases} \quad (49)$$

where

$$\mathcal{NBIN}(k; r, p) = \binom{k + r - 1}{k} (1 - p)^r p^k \quad (50)$$

Given N_t, let $N_d \in \{0, \dots, N_t\}$ be the number of "true" detections (if the state just transitioned to S_i, the "true" detection is Z_i). The probability of a true detection is $p_{\text{D}}^0 P_{obs}$. The detections are assumed to be independent of each other, and N_d given N_t is binomial distributed

$$P(N_d | N_t) = \mathcal{BIN}(N_d; N_t, p_{\text{D}}^0 P_{obs}) \quad (51)$$

$$= \binom{N_t}{N_d} (p_{\text{D}}^0 P_{obs})^{N_d} (1 - p_{\text{D}}^0 P_{obs})^{N_t - N_d} \quad (52)$$

For a given N_t and N_d we assume that the N_d detections are uniformly distributed over the N_t time steps, i.e. there are

N_t / N_d time steps between each of the "true" detections. This is a reasonable model of the average number of time steps between the detections. Under this assumption, and using the approximations above, the LLR ℓ is approximated by

$$\hat{\ell}(N_t, N_d) = (N_t - N_d) \log(1 - p_{\text{D}}^0 P_{\text{T}}) \quad (53)$$
$$+ N_d \log(1 - p_{\text{D}}^0 P_{\text{T}} + \frac{1 - p_{\text{FA}}}{p_{\text{FA}} N_s^{-1}} P_{obs} \bar{P}(N_t / N_d))$$

For a given measurement sequence $\mathbf{Z}^k$, the resulting LLR $\ell(\mathbf{Z}^k)$ is deterministic. The probability density of ℓ, conditioned on N_t and N_d is approximated as

$$p(\ell | N_t, N_d) \approx \delta(\ell = \hat{\ell}(N_t, N_d)) \quad (54)$$

where $\delta(\cdot)$ is the Dirac delta function. The cth moment of the LLR is approximated as follows

$$\mathbb{E}[\ell^c | H_1] = \int \ell^c p(\ell) \mathrm{d}\ell \quad (55a)$$

$$\approx \int \sum_{N_t} \sum_{N_d} \ell^c p(\ell | N_t, N_d) P(N_d | N_t) P(N_t) \mathrm{d}\ell \quad (55b)$$

$$\approx \sum_{N_t} \sum_{N_d} \left(\hat{\ell}(N_t, N_d) \right)^c P(N_d | N_t) P(N_t) \quad (55c)$$

We approximate the true pdf over ℓ with a Gaussian pdf

$$p(\ell | H_1) \approx \mathcal{N}(\ell; \hat{\mu}_1, \hat{\sigma}_1) \quad (56)$$

where the mean and standard deviation are given as

$$\hat{\mu}_1 = \mathbb{E}[\ell | H_1] \quad (57a)$$
$$\hat{\sigma}_1 = \mathbb{E}[\ell^2 | H_1] - \mathbb{E}[\ell | H_1]^2 \quad (57b)$$

D. Conditional error probabilities

Assume that we have a Gaussian density (or an approximation (56)) for the log-likelihood under H_1. The probability of threat miss $P_M(\widetilde{\delta})$ for a given threshold τ is then given by

$$P_M(\widetilde{\delta}) = \mathrm{P}(\ell < \tau) = F_\ell(\tau) \quad (58)$$

where $F_\ell(\cdot)$ is the Gaussian cumulative distribution function (cdf).

Empirically it was found to be difficult to approximate $p(\ell | H_0)$ analogously to (56). When the pdf for the log-likelihood under H_0 is unknown the probability of threat false alarm cannot be directly computed. However, using the pdf under H_1 we can derive the Chernoff bound for the probability of threat false alarm.

We can write, see e.g. [15, Section III.C.2],

$$P_F(\widetilde{\delta}_\ell) \leq \exp(\mu_{\ell,0}(s) - s\tau) \quad (59)$$

for all $s > 0$, where $\mu_{\ell,i}$ is the cumulant generating function[1] of ℓ under H_i defined as

$$\mu_{\ell,0}(s) = \log\left(\mathbb{E}\left[e^{s\ell} \middle| H_0\right]\right) \quad (60a)$$
$$\mu_{\ell,1}(t) = \log\left(\mathbb{E}\left[e^{t\ell} \middle| H_1\right]\right) \quad (60b)$$

for $s > 0$ and $t < 0$. The bound (59) can be minimized over $s > 0$ to find the tightest bound.

---

[1] The natural logarithm of the moment generating function.

We assume that P_j has density p_j for $j = 0$ and 1, and we have $\ell = \log(L)$, where $L = p_1/p_0$. In this case we have

$$\mu_{\ell,1}(t) = \log\left(\int_\Gamma e^{t\log L} p_1 \mathrm{d}\mu\right) = \log\left(\int_\Gamma L^t p_1 \mathrm{d}\mu\right) \quad (61a)$$

$$\mu_{\ell,0}(s) = \log\left(\int_\Gamma L^s p_0 \mathrm{d}\mu\right) = \log\left(\int_\Gamma L^{s-1} p_1 \mathrm{d}\mu\right) \quad (61b)$$

$$= \mu_{\ell,1}(s-1) \quad (61c)$$

We can now rewrite the bound (59) as

$$P_F(\widetilde{\delta}_\ell) \le \exp\left(\mu_{\ell,1}(s-1) - s\tau\right) \quad (62)$$

and with a variable substitution $t = s - 1$ we get

$$P_F(\widetilde{\delta}_\ell) \le \exp\left(\mu_{\ell,1}(t) - (t+1)\tau\right) \quad (63)$$

for all $t > -1$.

For the Gaussian density (56) the moment generating function is

$$\mathbb{E}\left[e^{t\ell} \,\middle|\, H_1\right] = \exp\left(t\hat{\mu}_1 + \frac{1}{2}t^2\hat{\sigma}_1^2\right) \quad (64)$$

and it follows that the cumulant generating function is

$$\mu_{\ell,1} = t\hat{\mu}_1 + \frac{1}{2}t^2\hat{\sigma}_1^2 \quad (65)$$

The error bound (59) becomes

$$P_F(\widetilde{\delta}_\ell) \le \exp\left(t\hat{\mu}_1 + \frac{1}{2}t^2\hat{\sigma}_1^2 - (t+1)\tau\right) \quad (66)$$

and achieves its minimum at value t_F

$$t_F = \max\left\{-1 \,,\, \frac{\tau - \hat{\mu}_1}{\hat{\sigma}_1^2}\right\} \quad (67)$$

The minimum error bound is

$$P_F(\widetilde{\delta}_T) \le \exp\left(t_F\hat{\mu}_1 + \frac{1}{2}t_F^2\hat{\sigma}_1^2 - (t_F+1)\tau\right) \quad (68)$$

Note that a property of the Chernoff bound, see [15], is that it may be trivial, i.e. for some $\hat{\mu}_1$, σ and τ the error bound is larger than one.

VI. Predicting the detectability performance

In general HMMs designed for asymmetric threats are not daisy chains, nor do they have equal transition probability for all states. However, the detectability analysis can still be used, as will be shown here. In this section we explain how a general HMM can be approximated by a daisy chain, such that the above detectability analysis is applicable, and we also show how to find a prediction of the probability of detection given an upper bound on the probability of false alarm.

A. Daisy chain approximation of complex HMM

The detectability analysis above is derived for the special case of daisy chain HMM, however it is still applicable to more complex HMMs.

If the HMM of interest is not a daisy chain, there is no direct correspondence between the number of states of the HMM, and the length of the daisy chain. Empirically we have found that in the detectability analysis N_S should be set to the expected value of the number of states that the HMM passes through from first to last state. The reason for this is that, in order to pass from the first to the last state, it is not necessary to pass through each state. For example, in Figure 1 there are four different ways to go from state 1 to state 9. The expected value of the number of states that are passed is 6.5, which can be rounded.

Additionally, the HMM of interest may not have uniform transition probabilities, which is an assumption in the detectability analysis. Note that if P_T is uniform, then so is the probability of remaining in the same state $1 - P_\mathrm{T}$. If this is the case, the transition probability used in the detectability analysis can be set as follows

$$P_T = 1 - \frac{1}{N_S}\sum_{S \in \mathcal{S}} \pi\left(\mathbf{s}_k = S | \mathbf{s}_{k-1} = S\right) \quad (69)$$

In other words, $1 - P_\mathrm{T}$ is set equal to the average probability of remaining in the same state.

B. Detection at desired false alarm level

Using the upper bound (68) for the probability of threat false alarm it is possible to find a likelihood threshold τ_α that gives a probability of threat false alarm less than or equal to α. Letting $t_F = (\tau - \hat{\mu})/\hat{\sigma}^2$, setting the right hand side of (68) equal to α, and solving for τ we get

$$\tau_\alpha = \left(\hat{\mu} - \hat{\sigma}^2 + \sqrt{\hat{\sigma}^4 - 2\hat{\mu}\hat{\sigma}^2 - 2\hat{\sigma}^2\log(\alpha)}\right) \quad (70)$$

For some values of α, $\hat{\mu}$ and $\hat{\sigma}$ the solution will be an imaginary number. This will happen when $(\tau - \hat{\mu}_1)/\hat{\sigma}_1 < -1$. In this case τ is trivially given by setting $(\tau - \hat{\mu}_1)/\hat{\sigma}_1 = -1$, which gives $\tau = \hat{\mu} - \hat{\sigma}^2$. The corresponding probability of miss is $F_\ell(\tau_\alpha)$, and the probability of detection is $1 - F_\ell(\tau_\alpha)$.

This means that for a given combination of the parameters N_S, P_T, p_D^0 and p_FA, we can compute a prediction of the detectors properties, i.e. for a given upper bound for the probability of threat false alarm we can predict what the probability of threat detection will be.

VII. Simulation setup

Intelligence observation data of the kind considered here is inherently secret, and for this reason results for real observation data records are unavailable, and could not be published if they were. Instead we present results for simulated data.

In this section we present the simulated models and the performance measures, in the next section we give the results.
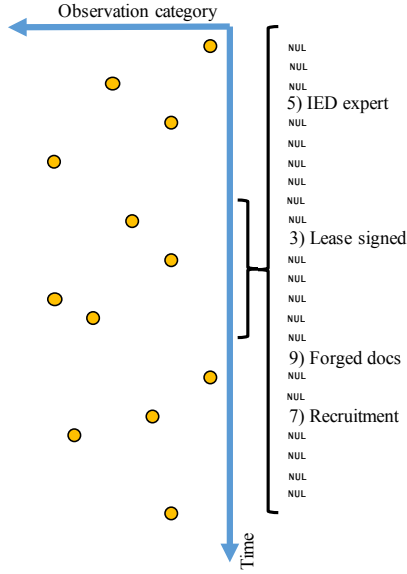
Fig. 2. Example observations. Left: Time on vertical axis, observation category on horizontal axis, observations illustrated by orange circles. Right: zoom of part of time line. Note that some time steps are empty (NUL), while others contain a single observation. Example categories shown here are abbreviations of examples given in Section II.

A. Simulated HMMs

We have simulated both daisy chain HMMs and more complex HMMs. For all the simulated models the HMM observation pmf (38) and clutter pmf (39) were used. To generate a sequence of observations, a threat birth time is first randomly sampled, and then the Markov transition model is then simulated to obtain a state sequence. For each state transition the HMM observation model is simulated, and to the full time sequence (i.e., also before threat birth) clutter is added by simulating the clutter model. Note that to simulate the observation process it is not necessary to simulate intelligence reports, such as the examples given in Section II. Instead the observation category is simulated by, e.g., generating integers that correspond to the observation category. An example of the simulated observation process is given in Figure 2.

Different probabilities of HMM observation and probabilities of clutter observation were simulated with $P_{obs} = 0.99$. Empirically we have found that it is not necessary to also simulate multiple values for P_{obs} because it is the product $P_{obs}p_D^0$ that is important, i.e. it is sufficient to simulate different values of p_D.

1) Daisy chain HMMs: Daisy chain results are primarily used to evaluate the density approximation (56), and for detectability performance prediction. We generated "daisy-chain" HMMs with parameters

$$N_s \in \{30, 40, 50\} \qquad p_D^0 \in \{0.1, 0, 2, \ldots, 0.9\}$$
$$P_T \in \{0.05, 0.1, 0.30\} \qquad p_{FA} \in \{0.1, 0, 2, \ldots, 0.9\} \quad (71)$$

which gives 729 different HMMs. Each daisy chain HMM was simulated 100 times.

2) More complex HMMs: The Bernoulli filter, Section IV, and the detectability performance prediction, Section VI, were evaluated using the following five models:


Fig. 3. Markov chain network modeling the production of weapons grade material, taken from [25, Fig. 6].

1) Planning of a truck bombing. This model has $N_s = 9$ states and uniform transition probabilities, see Figure 1 and [24] for details.
2) Production of weapons grade material, which is necessary for a country or group that wishes to develop nuclear weapons. This model has $N_s = 27$ states and uniform transition probabilities, see Figure 3 and [25] for details.
3) Planning and strategy. This model has $N_s = 9$ states and non-uniform transition probabilities, see [27] for details.
4) Collect resources. This model has $N_s = 8$ states and non-uniform transition probabilities, see [27] for details.
5) Preparations for a hijacking. This model has $N_s = 9$ states and non-uniform transition probabilities, see [27] for details.

The last three models are sub-parts of a larger model of the hijacking of IA Flight IC-814, flying from Kathmandu to New Delhi, on December 24, 1999. In the remainder of this paper we will refer to the models as HMM 1, HMM 2, and so on. None of the models is a simple daisy chain; on the contrary all five have more complex structure, as shown for HMM 1 in Figure 1, and for HMM 2 in Figure 3. Because of page length constraints, illustrations of remaining three HMMs are omitted.

HMM 1 and HMM 2 have uniform transition probabilities. The last three models do not have this property, instead the transition probabilities are specified in the models, see [27]. For the HMM state transitions in models 1 and 2, three different probabilities of transition were simulated

$$P_T \in \{0.10, 0.20, 0.30\}$$

Different probabilities of HMM observation and probabilities of clutter observation were simulated

$$p_D^0 \in \{0.05, 0.10 \ldots 0.95\}, \quad p_{FA} \in \{0.05, 0.10 \ldots 0.95\}$$

B. Performance evaluation

1) Bernoulli filtering: In each time step t_k, if $q_{k|k} > \tau_\varepsilon$ a maximum a posteriori (MAP) HMM state estimate was computed,

$$\hat{\mathbf{s}}_{k|k} = \arg\max_{\mathbf{s} \in \mathcal{S}} P_{k|k}(\mathbf{s}), \qquad (72)$$

where $P_{k|k}(\mathbf{s}) = \int P_{k|k}(\mathbf{s}, \mathbf{t})d\mathbf{t}$ is the marginal posterior distribution. We did not evaluate estimates of the transition state $\mathbf{t}_k$, because knowing the HMM state $\mathbf{s}_k$ is more important than knowing whether or not the HMM just transitioned to that state.

The HMM state estimation is evaluated using the following performance measures:

- Threat estimation rate: for $\varepsilon_k = 1$ and $q_{k|k} > \tau_\varepsilon$, the % time steps for which $\hat{\mathbf{s}}_{k|k} = \mathbf{s}_k$.
- Average estimated pmf: assuming the true state is $\mathbf{s}_k = S_i$, then ideally the estimated pmf is $P_{k|k}(\mathbf{s})$ has a "peak" around $\mathbf{s} = S_i$. For each state S_i, we compute the Monte Carlo average estimated pmf.

The HMM detection is evaluated using the following performance measures:

- Threat false alarm rate: for $\varepsilon_k = 0$ the % time steps t_k for which $q_{k|k} > \tau_\varepsilon$.
- Threat detection rate: for $\varepsilon_k = 1$ the % time steps t_k for which $q_{k|k} > \tau_\varepsilon$.

2) Log-likelihood pdf approximation: Each of the 729 daisy chains were simulated 100 times, and the empirical data was used to construct an empirical distribution of $p(\ell|H_1)$. The density approximation (56) is compared to a Gaussian fitted to the empirical data $\mathcal{N}(\ell;\ \lambda, \nu)^2$ using the Kullback-Leibler divergence (KL-div) [10]. Defined for two pdfs $p(\ell)$ and $q(\ell)$ as

$$\mathrm{KL}\left(p(\ell)||q(\ell)\right) = \int p(\ell)\log\left(p(\ell)/q(\ell)\right)\mathrm{d}\ell \qquad (73)$$

the KL-div is a measure of the information lost when the distribution $q(\ell)$ is used to approximate $p(\ell)^3$, see e.g. [4]. When it comes to approximating distributions in a maximum likelihood sense, the KL-div is often considered the optimal difference measure, see e.g. [1], [6], [19], [20], [30]. For the Gaussian fit $\mathcal{N}(\ell;\ \lambda, \nu)$ to the empirical distribution and the approximation $\mathcal{N}\left(\ell;\ \hat{\mu}_1, \hat{\sigma}_1\right)$ the KL-div is

$$\mathrm{KL}\left(\mathcal{N}(\ell;\ \lambda,\nu)||\mathcal{N}\left(\ell;\ \hat{\mu}_1, \hat{\sigma}_1\right)\right)$$
$$= \log\left(\frac{\hat{\sigma}_1}{\nu}\right) + \frac{\nu^2 + (\lambda - \hat{\mu}_1)^2}{2\hat{\sigma}_1^2} - \frac{1}{2} \qquad (74)$$

The proof of (74) is straightforward: input $p(\ell) = \mathcal{N}(\ell;\ \lambda, \nu)$ and $q(\ell) = \mathcal{N}\left(\ell;\ \hat{\mu}_1, \hat{\sigma}_1\right)$ into (73) and calculate the integral.

3) Detectability prediction: Ideally, for a given threshold, the threat false alarm rate is low and the threat detection rate is high. However these two objectives are often difficult to obtain simultaneously, and a tradeoff is necessary. The receiver operating characteristic (ROC) curve shows the trade off between the threat false alarm rate and the threat detection rate, and it can be used to find the threat detection rate at a specific threat false alarm rate. The detectability analysis/prediction is evaluated using the following performance measure:

- ROC area: the area under the ROC curve computed using the threat false alarm rate and threat detection rate, compared to the area under the ROC curve computed using the detectability analysis/prediction.

Note that a perfect detector has ROC area equal to 1, and a random-guess-detector has ROC area equal to 0.5.

VIII. SIMULATION RESULTS

A. Bernoulli filtering

Here we will highlight results from HMM 2, which has 27 states and models the *production of weapons grade material*

$^2\lambda$ is the sample mean, and ν is the sample variance.

^{3}Here $q(\cdot)$ is the Gaussian approximation (56) and $p(\cdot)$ is the empirical distribution.

Fig. 4. Bernoulli filtering results for HMM 2, which models production of weapons grade material, see Figure 3. Shown from left to right are the threat false alarm rate, threat detection rate, and threat estimation rate. The colormap is shown beneath each respective plot.
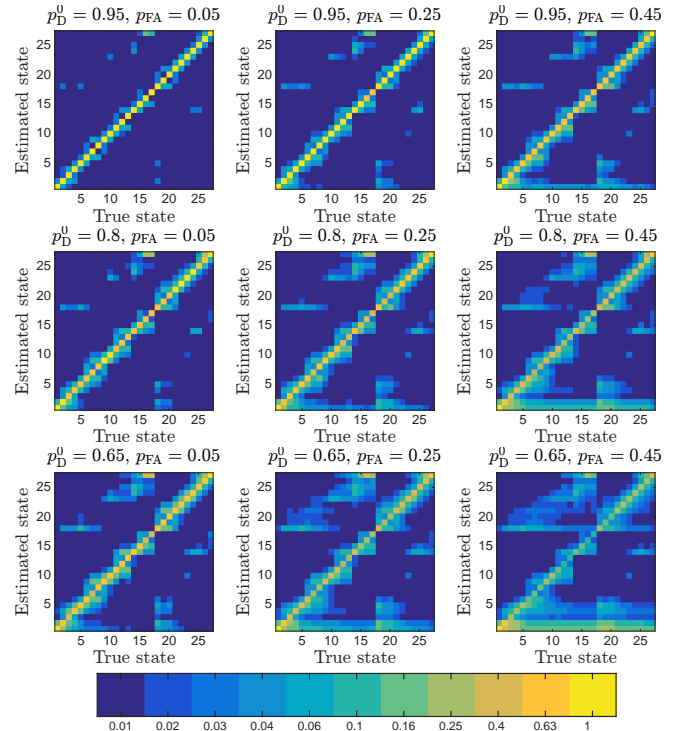

Fig. 5. Bernoulli filtering results showing estimated pmf $P_{k|k}(\mathbf{s})$ for different values of the true state. Probability of detection $p_D^0 \in \{0.65, 0.80, 0.95\}$ and $p_{\mathrm{FA}} \in \{0.05, 0.25, 0.45\}$. The colormap, shown at the bottom, has a logarithmic scale to increase clarity.

(PWGM) [25, Fig. 6], see Figure 3. The BF's birth and survival probabilities were set to $p_b = 0.01$ and $p_s = 0.99$, and the transition probability was $P_T = 0.2$. As noted above different probabilities of HMM observation p_D^0 and probabilities of clutter observation p_{FA} were tested.

For each p_D, p_{FA} pair, the BF was evaluated as follows. First we simulated 10^5 clutter observations and determined which existence probability thresholds $\tau_\varepsilon \in [0, 1]$ that gave 1% empirical false alarm rates. Next the PWGM-HMM was simulated 1000 times; in each simulation HMM birth time, state transitions, HMM observations, and clutter observations were all randomly simulated.

The threat detection rate, threat false alarm rate, and threat estimation rate, are summarized in Figure 4. Average estimated pmfs for nine selected combinations of p_D^0 and p_{FA} are shown Figure 5.
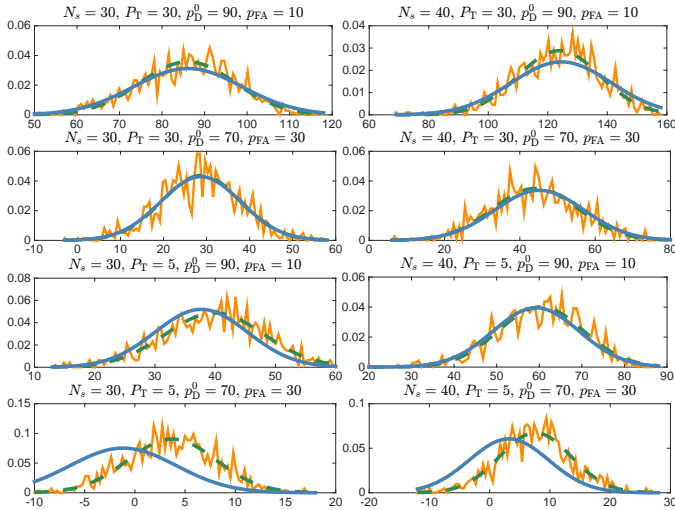
Fig. 6. Empirical distributions (solid orange) compared to Gaussians fitted to the data (dashed green) and Gaussian approximations (dashed blue).

The results confirm that the existence probability threshold indeed can be computed using sequences of clutter observations. It is more difficult to estimate the existence of an HMM when p_{D}^0 is lower and p_{FA} is higher. A comparison shows that the threat estimation rates are lower than the threat detection rates. This is mostly a result of missed HMM observations: a missed HMM observation makes it more difficult for the BF to estimate the state transition, and subsequently the estimated state $\hat{\mathbf{s}}_{k|k}$ is incorrect for a couple of time steps following the missed HMM observation. As expected, the estimated pmf is "peakier" around the true state when p_{D}^0 is higher and p_{FA} is lower, see Figure 5.

In a MATLAB implementation run on a desktop computer with two 2.66 GHz processors and 4 GB RAM, the median time for a single iteration (prediction and update) is 0.1ms, indicating that the proposed BF is capable of real-time performance.

B. Log-likelihood pdf approximation

In Figure 6 we show a comparison of the empirical distribution, the Gaussian fitted to the empirical data $\mathcal{N}(\ell\,;\,\lambda,\nu)$, and the Gaussian approximation $\mathcal{N}(\ell\,;\,\hat{\mu}_1,\hat{\sigma}_1)$, for a selection of eight parameter combinations out of the 729 combinations (71). From these results we see that the Gaussian approximation is more accurate for larger number of states N_S and/or higher transition probability P_{T}.

In Figure 7 we show the KL-div (74) for all 729 parameter combinations. The results show that the approximation (56) is most accurate when $p_{\mathrm{D}}^0 > p_{\mathrm{FA}}$ and for higher N_s and/or higher P_{T}. Importantly, the results show that the approximation derived in Section V is least accurate for parameter combinations that are of little practical interest (low p_{D}^0 and low p_{FA}).

These results indicate that it is reasonable to predict the detectability performance (as outlined in Section VI) based on the Gaussian approximation (56).
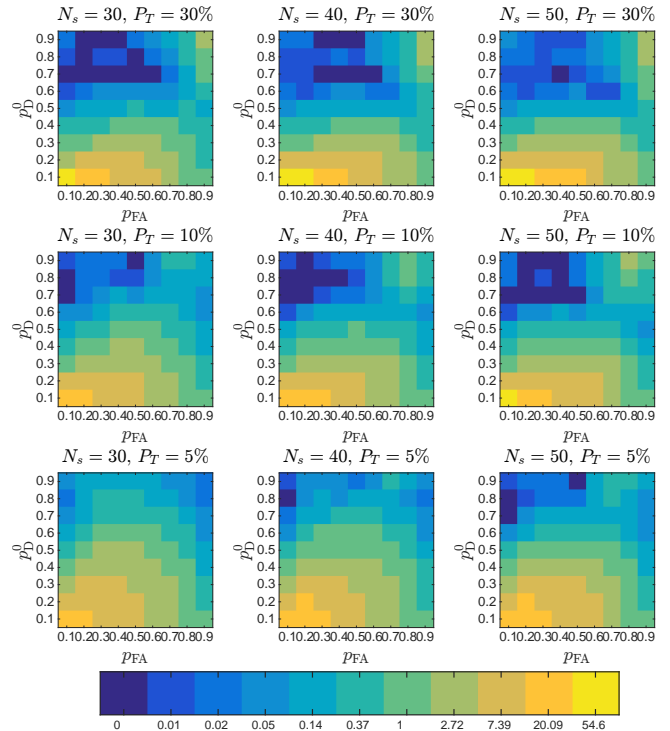

Fig. 7. Results showing the KL-div (74). Each of the nine subfigures show a heatmap of the KL-div, the colormap is given at the bottom. Note that the colormap has logarithmic scale.

C. Detectability prediction

Empirically we found that the detectability analysis is not accurate enough to make reliable predictions of what the threat detection rate will be at a given threat false alarm rate α. However, the detectability analysis is accurate enough to make predictions about the overall performance of the HMM detection, as measured by the ROC area.

Example results for HMM 1 are given in Figure 8. Because the Chernoff bound (68) is used in the detectability analysis, when p_{D} is low enough and p_{FA} is high enough the predicted ROC area is zero. This causes the two heat maps in Figure 8 to look quite different:

- The empirical ROC area transitions from 0.5 (lower right corner of the heat map) to 1 (upper left).
- The predicted ROC area transitions from 0 (lower right) to 1 (upper left).

However, if we compare the boundary along which the heat map transitions from low to high ROC area (0.75 for the empirical heat map, 0.5 for the predicted heat map), we see that the detectability prediction is quite accurate. This is illustrated in the right subfigure in Figure 8.

In Figure 9 the empirical ROC area is shown for all five simulated HMMs. The lines that separate poor detector performance from good detector performance are shown as a white solid line (empirical) and a red dashed line (prediction). Considering all results, we see that the detectability analysis/prediction is fairly accurate at predicting the boundary between low ROC area and high ROC area. This implies that we can predict the boundary between poor detector performance

Fig. 8. Example results for HMM 1. Left: empirical ROC area. Middle: predicted ROC area. Right: comparison of the boundary along which the heat maps transition from low to high. The colormap is given at the bottom.

and good detector performance.

Obtaining the empirical data took almost 10 hours while the corresponding detectability prediction took about one minute to compute. This result is important, as it allows us to predict the detectability performance quickly, and determine if the Bernoulli filter would yield good results or not.

Furthermore, it shows that the detectability analysis can be used to answer general questions such as

- Given P_T, p_D^0 and p_FA, what is the lowest N_s that gives good performance, i.e., what is the minimum level of model complexity?
- Given N_s, P_T and p_D^0, what is the maximum p_FA that gives good performance, i.e., what is the maximum level of clutter observations allowable?
- Given N_s, p_D^0 and p_FA, what is the lowest P_T that gives good performance, i.e., what is the maximum expected interval between relevant HMM observations?

Example results are shown in Figure 10. We see that a higher model complexity (as measured by the number of states N_s) implies better detection performance, see the top plot. Further, a higher model complexity implies a higher tolerance to clutter observations, see the middle plot. Lastly, a higher model complexity also implies that the detection process can handle longer intervals between the relevant HMM observations, see the bottom plot.

IX. CONCLUSIONS AND FUTURE WORK

In previous work the "Adaptive Safety and Monitoring" (ASAM) framework was introduced as a means to model asymmetric threats such as terrorist attacks. The modeling recognizes both the sequential nature of the activity – for example, a planning step must precede a surveillance step which often is followed by (but may be in parallel to) a funding step – and also incorporates statistical uncertainty. The natural framework is a hidden Markov model (HMM). This paper presented two major contributions.

Previous work applied a Page approach to detection of an HMM; but for scalability this is not appropriate. On the other hand, the "tracking" concepts of MTT and TBD are strongly suggested: they are how multiple *traditional* threats are estimated and counted. Hence the first contribution is the derivation of a Bernoulli filter for joint detection and estimation of HMMs. The Bernoulli filter is shown to give


Fig. 9. The heat maps show empirical ROC area for different combinations of the parameters p_D and p_FA; the title of each plot indicates the HMM and the probability of transition. The white lines indicate the transition boundary from low to high empirical ROC area; the red dotted lines indicate the corresponding predicted performance. The colormap is given at the bottom.

good results except in adverse conditions (low probability of detection and high probability of false alarm). Importantly this confirms the prior analysis of the problem.

The second contribution is an answer to the natural question: can these sorts of activities be detected at all? On an intuitive level, for detectability there must be maximum levels of clutter allowable; maximum intervals between relevant observations; and a minimum level of complexity. In this paper we have addressed the issue by approximating the detectability of such a process, and using this approximation we analyzed the error probabilities. The analysis suggested that the HMM can be detected if the probability of HMM detection is larger than the probability of a clutter detection. The lower the probability of state transition is, i.e. the more time steps there are between the state transition, the larger the margin must be between the probability of HMM detection and the probability of clutter detection.

In future work the presented framework can straightforwardly be extended to HMMs that have observations also when there is not a state transition. Other extensions would be to consider multiple observations per time step – e.g., a Poisson distributed number of clutter observations each time step – and tracking of multiple HMM's, using, e.g., a multi-Bernoulli filter.

The presented work assumed that the model structure and model parameters were known. Two different directions for future work is 1) to consider online estimation of the model
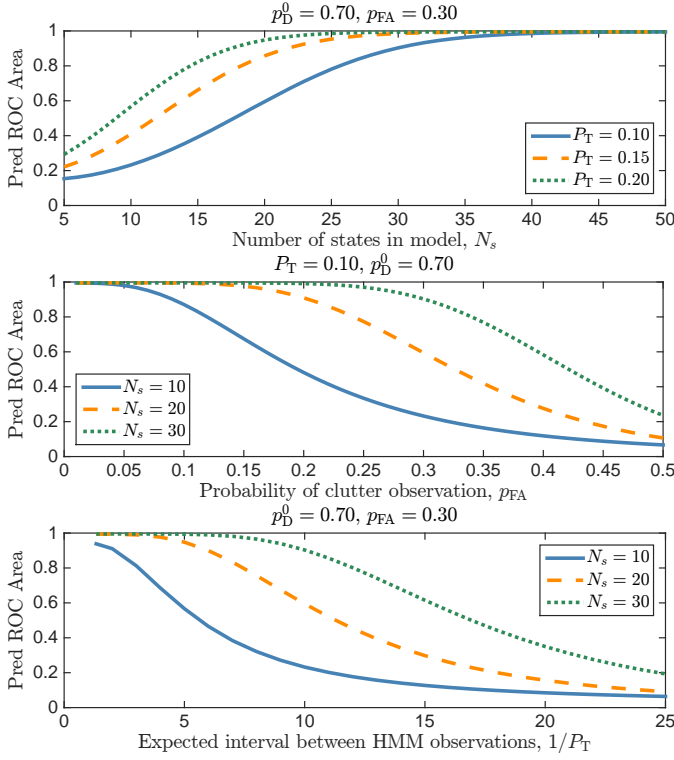
Fig. 10. Detectability prediction results. Top: minimum model complexity. Middle: maximum clutter level. Bottom: maximum expected interval between relevant HMM observations.

structure and the model parameters, and 2) to investigate how model errors affects the performance of the Bernoulli filter.

APPENDIX A
LIKELIHOOD

Here we express the FISST likelihood $f(\mathbf{Z}_k|\mathbf{Z}^{k-1})$ defined in (25). If $\mathbf{Z}_k = \emptyset$ then from (17), (21) and (25) we get

$$
\begin{align}
f(\emptyset|\mathbf{Z}^{k-1}) =&(1-p_{\text{FA}})(1-q_{k|k-1}) \tag{75a} \\
&+ q_{k|k-1}\int (1-p_{\text{FA}})(1-p_{\text{D}}(\zeta))P_{k|k-1}(\zeta)\mathrm{d}\zeta \\
=&(1-p_{\text{FA}})(1-q_{k|k-1}\Delta^1_{k|k-1}) \tag{75b}
\end{align}
$$

where

$$
\begin{align}
\Delta^1_{k|k-1} =& \int p_D(\zeta)P_{k|k-1}(\zeta)\mathrm{d}\zeta \tag{76a} \\
=& p_D^0 P_{k|k-1}(\mathbf{t}=1) \tag{76b}
\end{align}
$$

If $\mathbf{Z}_k = \{\mathbf{z}_k\}$ then from (17), (21) and (25) we get

$$
\begin{align}
f(\{\mathbf{z}\}|\mathbf{Z}^{k-1}) \tag{77a} \\
=p_{\text{FA}}g_{\text{FA}}(\mathbf{z})(1-q_{k|k-1}) + q_{k|k-1} \tag{77b} \\
\times \int [\eta(\emptyset|\{\zeta\})\kappa(\mathbf{z}) + \eta(\mathbf{z}|\{\zeta\})\kappa(\emptyset)] P_{k|k-1}(\zeta)\mathrm{d}\zeta \\
=p_{\text{FA}}g_{\text{FA}}(\mathbf{z})(1-q_{k|k-1}\Delta_{k|k-1}) \tag{77c}
\end{align}
$$

where

$$
\begin{align}
\Delta_{k|k-1} =&\Delta^1_{k|k-1} - \Delta^2_{k|k-1} \tag{78a} \\
\Delta^2_{k|k-1} =&\frac{1-p_{\text{FA}}}{p_{\text{FA}}g_{\text{FA}}(\mathbf{z})}G^s_{k|k-1}(\mathbf{z}) \tag{78b} \\
G^s_{k|k-1}(\mathbf{z}) =&p_D^0 \int g_s(\mathbf{z}|\mathbf{s})P_{k|k-1}(\mathbf{s},\mathbf{t}=1)\mathrm{d}\mathbf{s} \tag{78c}
\end{align}
$$

If $\mathbf{Z}_k = \{\mathbf{z}^1,\mathbf{z}^2\}$ then from (17), (21) and (25) we get

$$
\begin{align}
f\left(\{\mathbf{z}^1,\mathbf{z}^2\}|\mathbf{Z}^{k-1}\right) \tag{79a} \\
=0\cdot(1-q_{k|k-1}) + q_{k|k-1} \tag{79b} \\
\times \int \left[\eta(\mathbf{z}^1|\{\zeta\})\kappa(\mathbf{z}^2) + \eta(\mathbf{z}^2|\{\zeta\})\kappa(\mathbf{z}^1)\right] P_{k|k-1}(\zeta)\mathrm{d}\zeta \\
=q_{k|k-1}p_{\text{FA}} \\
\times \left(g_{\text{FA}}(\mathbf{z}^2)G^s_{k|k-1}(\mathbf{z}^1) + g_{\text{FA}}(\mathbf{z}^1)G^s_{k|k-1}(\mathbf{z}^2)\right) \tag{79c}
\end{align}
$$

APPENDIX B
BERNOULLI FILTER UPDATE

Let $\mathbf{S}_k = \emptyset$. From (23) we get

$$
1-q_{k|k} = \frac{\varphi(\mathbf{Z}_k|\emptyset)(1-q_{k|k-1})}{f(\mathbf{Z}_k|\mathbf{Z}^{k-1})} = \frac{\kappa(\mathbf{Z}_k)(1-q_{k|k-1})}{f(\mathbf{Z}_k|\mathbf{Z}^{k-1})} \tag{80}
$$

Rearranging the terms gives

$$
q_{k|k} = 1 - \frac{\kappa(\mathbf{Z}_k)(1-q_{k|k-1})}{f(\mathbf{Z}_k|\mathbf{Z}^{k-1})} \tag{81}
$$

If $\mathbf{Z}_k = \emptyset$ then from (17), (75) and (81) we get

$$
\begin{align}
q_{k|k} =&1 - \frac{(1-p_{\text{FA}})(1-q_{k|k-1})}{(1-p_{\text{FA}})(1-q_{k|k-1}\Delta^1_{k|k-1})} \tag{82a} \\
=&\frac{(1-\Delta^1_{k|k-1})q_{k|k-1}}{1-q_{k|k-1}\Delta^1_{k|k-1}} \tag{82b}
\end{align}
$$

If $\mathbf{Z}_k = \{\mathbf{z}_k\}$ then from (17), (77) and (81) we get

$$
\begin{align}
q_{k|k} =&1 - \frac{p_{\text{FA}}g_{\text{FA}}(\mathbf{z})(1-q_{k|k-1})}{p_{\text{FA}}g_{\text{FA}}(\mathbf{z})\left(1-q_{k|k-1}\Delta_{k|k-1}\right)} \tag{83a} \\
=&\frac{(1-\Delta_{k|k-1})q_{k|k-1}}{1-q_{k|k-1}\Delta_{k|k-1}} \tag{83b}
\end{align}
$$

If $\mathbf{Z}_k = \{\mathbf{z}^1,\mathbf{z}^2\}$ then from (17), (79) and (81) we get

$$
q_{k|k} =1 - \frac{0\cdot(1-q_{k|k-1})}{f(\mathbf{Z}_k|\mathbf{Z}^{k-1})} = 1 \tag{84}
$$

Let $\mathbf{S}_k = \{\zeta\}$. From (23) we get

$$
q_{k|k}P_{k|k}(\zeta) = \frac{\varphi(\mathbf{Z}_k|\{\zeta\})q_{k|k-1}P_{k|k-1}(\zeta)}{f(\mathbf{Z}_k|\mathbf{Z}^{k-1})} \tag{85}
$$

Rearranging the terms gives

$$
P_{k|k}(\zeta) = \frac{\varphi(\mathbf{Z}_k|\{\zeta\})q_{k|k-1}}{f(\mathbf{Z}_k|\mathbf{Z}^k)q_{k|k}}P_{k|k-1}(\zeta) \tag{86}
$$

If $\mathbf{Z}_k = \emptyset$ then from (17), (18), (21), (75) and (86) we get

$$P_{k|k}(\zeta) = \frac{(1-p_{\mathrm{FA}})(1-p_{\mathrm{D}}(\zeta))q_{k|k-1}}{(1-p_{\mathrm{FA}})(1-q_{k|k-1}\Delta^1_{k|k-1})q_{k|k}}P_{k|k-1}(\zeta) \tag{87a}$$

$$= \frac{1-p_{\mathrm{D}}(\zeta)}{1-\Delta^1_{k|k-1}}P_{k|k-1}(\zeta) \tag{87b}$$

If $\mathbf{Z}_k = \{\mathbf{z}\}$ then from (17), (18), (21), (77) and (86) we get

$$P_{k|k}(\zeta) = \frac{[\eta(\emptyset|\{\zeta\})\kappa(\mathbf{z}) + \eta(\mathbf{z}|\{\zeta\})\kappa(\emptyset)]\,q_{k|k-1}}{p_{\mathrm{FA}}g_{\mathrm{FA}}(\mathbf{z})\left(1 - q_{k|k-1}\Delta_{k|k-1}\right)q_{k|k}}P_{k|k-1}(\zeta) \tag{88a}$$

$$= \frac{1-p_{\mathrm{D}}(\zeta)}{1-\Delta_{k|k-1}}P_{k|k-1}(\zeta) \tag{88b}$$
$$+ \frac{1-p_{\mathrm{FA}}}{p_{\mathrm{FA}}g_{\mathrm{FA}}(\mathbf{z})}\frac{p_{\mathrm{D}}(\zeta)h(\mathbf{z}|\zeta)}{1-\Delta_{k|k-1}}P_{k|k-1}(\zeta)$$

If $\mathbf{Z}_k = \left\{\mathbf{z}^1, \mathbf{z}^2\right\}$ then from (17), (18), (21), (79), (86) we get

$$P_{k|k}(\zeta)$$
$$= \frac{\eta(\mathbf{z}^1|\{\zeta\})\kappa(\mathbf{z}^2)q_{k|k-1}}{f(\mathbf{Z}_k|\mathbf{Z}^k)q_{k|k}}P_{k|k-1}(\zeta) \tag{89a}$$
$$+ \frac{\eta(\mathbf{z}^2|\{\zeta\})\kappa(\mathbf{z}^1)q_{k|k-1}}{f(\mathbf{Z}_k|\mathbf{Z}^k)q_{k|k}}P_{k|k-1}(\zeta)$$
$$= \frac{g_{\mathrm{FA}}(\mathbf{z}^2)p_{\mathrm{D}}(\zeta)h(\mathbf{z}^1|\zeta)P_{k|k-1}(\zeta)}{g_{\mathrm{FA}}(\mathbf{z}^2)G^s_{k|k-1}(\mathbf{z}^1) + g_{\mathrm{FA}}(\mathbf{z}^1)G^s_{k|k-1}(\mathbf{z}^2)} \tag{89b}$$
$$+ \frac{g_{\mathrm{FA}}(\mathbf{z}^1)p_{\mathrm{D}}(\zeta)h(\mathbf{z}^2|\zeta)P_{k|k-1}(\zeta)}{g_{\mathrm{FA}}(\mathbf{z}^2)G^s_{k|k-1}(\mathbf{z}^1) + g_{\mathrm{FA}}(\mathbf{z}^1)G^s_{k|k-1}(\mathbf{z}^2)}$$

APPENDIX C
BERNOULLI FILTER PREDICTION

Let $\mathbf{S}_k = \emptyset$. From (14) and (24) we get

$$f_{k+1|k}(\emptyset|\mathbf{Z}^{k-1}) = 1 - q_{k+1|k}$$
$$= P^{\mathbf{S}}_{k+1|k}(\emptyset|\emptyset)(1 - q_{k|k}) \tag{90a}$$
$$+ q_{k|k}\int P^{\mathbf{S}}_{k+1|k}(\emptyset|\{\zeta'\})P_{k|k}(\zeta')\mathrm{d}\zeta'$$
$$= (1-p_b)(1-q_{k|k}) + (1-p_s)q_{k|k}. \tag{90b}$$

Rearranging the terms gives the predicted existence probability

$$q_{k+1|k} = p_b(1-q_{k|k}) + p_s q_{k|k}. \tag{91}$$

Let $\mathbf{S}_k = \{\zeta\}$. From (14) and (24) we get

$$f_{k+1|k}(\{\zeta\}|\mathbf{Z}^k) = q_{k+1|k}P_{k+1|k}(\zeta)$$
$$= P^{\mathbf{S}}_{k+1|k}(\{\zeta\}|\emptyset)(1 - q_{k|k})$$
$$+ q_{k|k}\int P^{\mathbf{S}}_{k+1|k}(\{\zeta\}|\{\zeta'\})P_{k|k}(\zeta')\mathrm{d}\zeta' \tag{92a}$$
$$= p_b P^b_{k+1|k}(\zeta)(1 - q_{k|k})$$
$$+ q_{k|k}\int p_s \pi_{k+1|k}(\zeta|\zeta')P_{k|k}(\zeta')\mathrm{d}\zeta' \tag{92b}$$

Rearranging the terms give the predicted HMM state pmf,

$$P_{k+1|k}(\zeta) = \frac{p_b(1-q_{k|k})}{q_{k+1|k}}P^b_{k+1|k}(\zeta) \tag{93}$$
$$+ \frac{p_s q_{k|k}}{q_{k+1|k}}\int \pi_{k+1|k}(\zeta|\zeta')P_{k|k}(\zeta')\mathrm{d}\zeta'$$

REFERENCES

[1] T. Ardeshiri, K. Granström, E. Özkan, and U. Orguner, "Greedy reduction algorithms for mixtures of exponential family," *IEEE Signal Processing Letters*, vol. 22, no. 6, pp. 676–680, Jun. 2015.

[2] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. YBS Publishing, 2011.

[3] F. Barnaby, *How to Build a Nuclear Bomb and Other Weapons of Mass Destruction*. New York, NY, USA: Nation Books, 2004.

[4] C. M. Bishop, *Pattern recognition and machine learning*. New York, USA: Springer, 2006.

[5] T. Coffman and S. Marcus, "Dynamic classification of groups through social networks and HMMs," in *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, USA, Mar. 2004, pp. 3197–3205.

[6] K. Granström and U. Orguner, "On the Reduction of Gaussian inverse Wishart mixtures," in *Proceedings of the International Conference on Information Fusion*, Singapore, Jul. 2012, pp. 2162–2169.

[7] K. Granström, P. Willett, and Y. Bar-Shalom, "A Bernoulli filter approach to detection and estimation of hidden Markov models using cluttered observation sequences," in *Proceedings of the IEEE Conference on Acoustics, Speech and Signal Processing*, Brisbane, Australia, Apr. 2015.

[8] ——, "Detectability analysis of detection and estimation of structured action from cluttered data," in *Proceedings of the International Conference on Information Fusion*, Washington, DC, USA, Jul. 2015, pp. 173–179.

[9] K. Granström, C. Lundquist, F. Gustafsson, and U. Orguner, "Random Set Methods: Estimation of Multiple Extended Objects," *IEEE Robotics and Automation Magazine*, vol. 21, no. 2, pp. 73–82, Jun. 2014.

[10] S. Kullback and R. A. Leibler, "On information and sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, Mar. 1951.

[11] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2007.

[12] ——, "'statistics 102' for multisensor-multitarget tracking," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 3, pp. 376–389, Jun. 2013.

[13] ——, *Advances in Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2014.

[14] R. Paternoster, "Nuclear weapon proliferation indicators and observables," Los Alamos National Laboratory, Tech. Rep. LA-12430-MS, Dec. 1992.

[15] H. V. Poor, *An Introduction to Signal Detection and Estimation, Second Edition*. Springer, 1994.

[16] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proceedings of IEEE*, vol. 77, no. 2, pp. 257–286, Feb. 1989.

[17] L. Rabiner and B.-H. Juang, "An introduction to hidden Markov models," *IEEE ASSP Magazine*, vol. 3, no. 1, pp. 4–16, Jan. 1986.

[18] B. Ristic, B.-T. Vo, B.-N. Vo, and A. Farina, "A Tutorial on Bernoulli Filters: Theory, Implementation and Applications," *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3406–3430, Jul. 2013.

[19] A. R. Runnalls, "Kullback-Leibler approach to Gaussian mixture reduction," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 3, pp. 989–999, Jul. 2007.

[20] D. Schieferdecker and M. F. Huber, "Gaussian Mixture Reduction via Clustering," in *Proceedings of the International Conference on Information Fusion*, Seattle, WA, USA, Jul. 2009.

[21] P. Schrodt, "Pattern recognition of international crises using hidden Markov models," in *Political Complexity: Nonlinear Models of Politics*, D. Richards, Ed. Ann Arbor, MI, USA: University of Michigan Press, 2000, pp. 296–328.

[22] F. Settle, "Nuclear chemistry, nuclear proliferation."

[23] S. Singh, W. Donat, H. Tu, J. Lu, K. Pattipati, and P. Willett, "An advanced system for modeling asymmetric threats," in *Proceedings of 2006 IEEE International Conference on Systems, Man, and Cybernetics*, Taipei, Taiwan, Oct. 2006.

[24] S. Singh, H. Tu, J. Allanach, J. Areta, P. Willett, and K. Pattipati, "Modeling threats," *IEEE Potentials*, pp. 18–21, Aug./Sept. 2004.

[25] S. Singh, H. Tu, W. Donat, K. Pattipati, and P. Willett, "Anomaly detection via feature-aided tracking and hidden Markov models," *IEEE Transactions on Systems, Man, and Cybernetics–Part A: Systems and Humans*, vol. 39, no. 1, pp. 144–159, Jan. 2009.

[26] L. Spector and J. Smith, *Nuclear Ambitions: The Spread of Nuclear Weapons 1989–1990*. Boulder, CO, USA: Westview Press, 1990.

[27] H. Tu, J. Allanach, S. Singh, K. Pattipati, and P. Willett, "Information integration via hierarchical and hybrid Bayesian networks," *IEEE Transactions on Systems, Man, and Cybernetics–Part A: Systems and Humans*, vol. 30, no. 1, Jan. 2006.

[28] U.S. Congress, Office of Technology Assessment, "Technologies underlying weapons of mass destruction," U.S. Printing Office, Washington, DC, USA, Tech. Rep. OTA-BP-ISC-115, Dec. 1993.

[29] B.-N. Vo, B.-T. Vo, and D. Clark, "Bayesian multiple target filtering using random finite sets," in *Integrated Tracking, Classification, and Sensor Management*, M. Mallick, V. Krishnamurthy, and B.-N. Vo, Eds. New York, NY, USA: Wiley, 2014.

[30] J. L. Williams and P. S. Maybeck, "Cost-Function-Based Gaussian Mixture Reduction for Target Tracking," in *Proceedings of the International Conference on Information Fusion*, Cairns, Queensland, Australia, Jul. 2003.