TRACK-BEFORE-DETECT WITH CENSORED OBSERVATIONS

Emanuele Grossi, Marco Lops, and Luca Venturino

e.grossi@unicas.it, lops@unicas.it, l.venturino@unicas.it
DAEIMI – Università degli Studi di Cassino, via G. di Biasio 43, Italy

ABSTRACT

In this work, we consider the problem of detecting the presence of a prospective moving target from a set of censored measurements produced by a generic sensor and propose a novel track-before-detect algorithm which exploits the inherent sparse nature of the measurement data to achieve complexity efficiency. The impact of the censoring stage on the system complexity and performance is investigated.

Index Terms— Track-before-detect (TBD), censored observation, generalized likelihood ratio test (GLRT).

1. INTRODUCTION

Track-Before-Detect (TBD) has been proposed to detect the presence of a moving object and possibly estimate its position in a suitably defined state space [1–5]. The returns collected by a sensor in a number of consecutive epochs are jointly processed (typically in a Generalized-Likelihood Ratio Test, GLRT), and target detection is reduced to searching for the path with maximum metric in a state trellis and comparing the cumulated metric to a detection threshold. However, this strategy hardly leads to real-time implementable schemes when the cardinality of the state space in *large*, even resorting to dynamic programming algorithms, such as the Viterbi algorithm [6]. The main reason for such an un-affordable complexity is that all of the observations are retained at each epoch and processed. In this paper we propose and analyze a different approach, based on retaining only the *significant* returns, i.e., those whose modula exceed a pre-assigned primary threshold: thus all censored observations collapse into a unique state, and at each epoch the state space may shrink or expand based on the intensity of the observed returns. Higher primary thresholds may reduce the complexity of the algorithm used to compute the test statistics but may also cause some loss with respect to the un-censored case: however, we produce experimental evidence that the complexity-performance trade-off can be advantageous.

The contributions of this study are: a) We outline the observation model for censored observations, deriving the GLRT (which generalizes the one of [5]); b) Based on a), we propose a new TBD algorithm, offering a closed-form complexity analysis; c) The system performance is assessed in term of probability of detection (PD) versus probability of false alarm (PFA), showing the inherent trade-off between the achievable PD gain, the system complexity (tied to the number of integrated frames), and the censoring threshold.

This works was carried out under a research contract sponsored by Selex Galileo, Nerviano (MI), Italy.

2. PROBLEM FORMULATION

We consider a sensor consisting of $N_x \times N_y$ resolutions cells (or pixels) of side length Δ_x and Δ_y along the x and y dimensions, respectively. The center of cell (i,j) is at $(i\Delta_x, j\Delta_y)$, for $i = 1,\dots,N_x$ and $j = 1,\dots,N_y$. Measurements are recorded at discrete time instants ℓT, where T is the sampling interval and $\ell \in \mathbb{Z}$. We assume that at most one target is present in the scene, and that the target signature appears on at most one pixel at each sampling instant.

The input signal $\zeta_\ell(i,j)$ captured by resolution element (i,j) at epoch ℓ is

$$\zeta_\ell(i,j) = \begin{cases} s_\ell(i,j) + n_\ell(i,j), & \text{if the target} \\ & \text{is present in } (i,j) \\ n_\ell(i,j), & \text{otherwise} \end{cases}$$

where $n_\ell(i,j)$ and $s_\ell(i,j)$ represent the noise and the target component, respectively, modeled as independent Gaussian circularly-symmetric complex random variables with variances 1 and ρ, ρ being the signal-to-disturbance (SDR) ratio. The measurement generating process consists of a square-law envelope detector followed by a censoring stage operating with a threshold γ_1. The sensor measurements are

$$z_\ell(i,j) = \begin{cases} |\zeta_\ell(i,j)|^2, & \text{if } |\zeta_\ell(i,j)|^2 \geq \gamma_1 \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

for $i = 1,\dots,N_x$ and $j = 1,\dots,N_y$. We refer to the collection of data measurements $Z_\ell = \{z_\ell(i,j),\ i = 1,\dots,N_x,\ j = 1,\dots,N_y\}$ at epoch ℓ as the ℓ-th data frame.

The censored observations are sent to the detector, whose task is to determine if a target is present (hypothesis H_1) or not (hypothesis H_0) at epoch ℓ. The detector is a causal filter which jointly elaborates the current data frame Z_ℓ along with the $L-1$ past frames $Z_{\ell-1},\dots,Z_{\ell-L+1}$. If a target is declared, the detector sends its position and signal strength to the successive tracking stage for further processing. The detector cannot declare a target in cell (i, j) at epoch ℓ if $z_\ell(i, j) = 0$, since there would be no signal strength information to be sent to the tracking stage. The focus of this work is on the design of the detection stage.

3. GENERALIZED LIKELIHOOD RATIO TEST

To simplify exposition, let us assume that Z_L is the current data frame, so that the observations taken at epochs $\ell = 1, \dots, L$ are jointly processed. Also, indicate by (x_ℓ, y_ℓ) the cell occupied by the target in the ℓ-th frame under the hypothesis H_1, and define $\boldsymbol{x} = (x_1 \ \cdots x_L)^T$ and $\boldsymbol{y} = (y_1 \ \cdots \ y_L)^T$. The sequence of pixels $(\boldsymbol{x}, \boldsymbol{y})$ defines the target trajectory over the L processed frames.

Assuming cell-to-cell and frame-to-frame independence, the densities of the observations under the two hypotheses are

$$\begin{cases} \displaystyle\prod_{\ell=1}^{L} f_1\big(z_\ell(x_\ell, y_\ell)\big) \prod_{\substack{i=1 \\ i \neq x_\ell}}^{N_x} \prod_{\substack{j=1 \\ j \neq y_\ell}}^{N_y} f_0\big(z_\ell(i,j)\big), & \text{under } H_1 \\ \displaystyle\prod_{\ell=1}^{L} \prod_{i=1}^{N_r} \prod_{j=1}^{N_a} f_0\big(z_\ell(i,j)\big), & \text{under } H_0 \end{cases} \tag{2}$$

where

$$\begin{aligned} f_1(z) =& (1 - \mathrm{pd}_1)\, \mathbb{1}_{\{0\}}(z) + \frac{e^{-z/(1+\rho)}}{1+\rho}\, \mathbb{1}_{[\gamma_1, \infty)}(z) \\ f_0(z) =& (1 - \mathrm{pfa}_1)\, \mathbb{1}_{\{0\}}(z) + e^{-z}\, \mathbb{1}_{[\gamma_1, \infty)}(z) \end{aligned}$$

$\mathbb{1}_B$ is the indicator function for the set B, and[1]

$$\mathrm{pfa}_1 = \int_{\gamma_1}^{\infty} e^{-z} dz = e^{-\gamma_1} \tag{3a}$$

$$\mathrm{pd}_1 = \int_{\gamma_1}^{\infty} \frac{e^{-z/(1+\rho)}}{1+\rho} dz = \mathrm{pfa}_1^{1/(1+\rho)}. \tag{3b}$$

From (2), the GLRT is obtained by maximizing the log-likelihood ratio over the set of admissible target trajectories, say $\mathcal{R}$, and comparing it with a secondary threshold, i.e.,

$$\max_{(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{R}} \sum_{\ell=1}^{L} \ln \frac{f_1\big(z_\ell(x_\ell, y_\ell)\big)}{f_0\big(z_\ell(x_\ell, y_\ell)\big)} \underset{H_0}{\overset{H_1}{\gtrless}} \gamma_2'. \tag{4}$$

After some manipulations, the GLRT in (4) can be recast as

$$\max_{(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{R}} \sum_{\ell=1}^{L} \big[z_\ell(x_\ell, y_\ell) - c_{\text{miss}}\big] \mathbb{1}_{[\gamma_1, \infty)}\big(z_\ell(x_\ell, y_\ell)\big) \underset{H_0}{\overset{H_1}{\gtrless}} \gamma_2 \tag{5}$$

[1]The measurements in (1) are mixed (discrete/continuous) random variables, and the densities f_1 and f_0 are computed with respect to the measure defined as the sum of the Dirac measure centered in 0 and the Lebesgue measure.

where

$$c_{\text{miss}} = \frac{1+\rho}{\rho} \ln \frac{(1+\rho)\big(1 - \mathrm{pfa}_1^{1/(1+\rho)}\big)}{1 - \mathrm{pfa}_1}. \tag{6}$$

Notice that, if no censoring is performed on data (i.e., $\gamma_1 = 0$), (5) reduces to the GLRT presented in [5]. If $\gamma_1 > 0$, there is a correcting term, c_{miss}, in the test statistic that accounts for the number of censored data measurements along each candidate trajectory. Finally, from (3a) and (6), it can be shown that $0 \leq c_{\text{miss}} \leq \gamma_1$.

3.1. Admissible target trajectories

A target trajectory $(\boldsymbol{x}, \boldsymbol{y})$ is admissible if it does not have a zero measurement in the L-th (current) frame, *and* it complies with the physical constraints on the target kinematic. In this work we just force a physical constraint on the maximum target speed $v_{\max}$, whereby we have:

$$\begin{aligned} \mathcal{R} = \Big\{ (\boldsymbol{x}, \boldsymbol{y}) : z_L(x_L, y_L) > 0 \text{ and } \Delta_x^2 (x_q - x_h)^2 \\ + \Delta_y^2 (y_q - y_h)^2 \leq (v_{\max} T)^2 (q-h)^2, \quad \forall\, q, h \Big\}. \end{aligned}$$

3.2. Complexity issues

Brute-force solution to (5) requires evaluating the test statistic for all admissible target trajectories in set $\mathcal{R}$, which entails a complexity $\mathcal{O}\left((N_x N_y)^L\right)$, i.e., exponential in the number of integrated frames. As shown in [1–5], a possible alternative is recasting the maximization (5) as the problem of estimating the best state sequence of a discrete-time Markov chain whose state elements are the $N_x \times N_y$ pixels and taking advantage of the Viterbi algorithm [6]: in this case, the complexity is $\mathcal{O}(L N_x N_y)$, i.e., linear in the number of integrated frames and of resolution elements. In many real-time applications even this complexity can be unaffordable. Suppose for example that $N_x = 3000$ and $N_y = 600$, and that the target is allowed a transition of ± 20 resolution elements in each dimension; the Viterbi algorithm still has to compute a maximum between 400 elements (40×40 cells along the two dimensions), for each of the $N_x N_y = 1.8 \cdot 10^6$ resolution elements of each frame. A novel approach to solve (5), which exploits the inherent sparse nature of the censored data, is presented in the next section.

4. PROPOSED TBD ALGORITHM

The proposed algorithm takes advantage of the fact that the censoring stage reduces the number of non-zero data measurements in each frame, which can be even much smaller than the number of resolution elements $N_x N_y$.

4.1. Alarm lists

To explain the algorithm, we organize the non-zero data measurements (also called here alarms) at epoch ℓ into an alarm list $\mathcal{S}_\ell = \{\mathcal{S}_{1,\ell}, \dots, \mathcal{S}_{D_\ell,\ell}\}$, where D_ℓ is the number of alarms, and $\mathcal{S}_{k,\ell} = (\bar{x}_{k,\ell}\ \bar{y}_{k,\ell}\ \bar{z}_{k,\ell})$. The trajectory of a prospective target from epoch 1 to epoch p, with $1 \le p \le L$, can now be specified by a p-dimensional vector, say $(t_1, \dots, t_p)$, with $t_\ell \in \{0, 1, \dots, D_\ell\}$ for $\ell = 1, \dots, p$. Specifically, $t_\ell = m$ means that the target is observed at epoch ℓ, and the corresponding alarm is $\mathcal{S}_{m,\ell}$, while $t_\ell = 0$ that there is a censored detection at epoch ℓ. The sequence of positions indexed by $(t_1, \dots, t_p)$ is

$$\big\{(\bar{x}_{t_1,1}, \bar{y}_{t_1,1}), \dots, (\bar{x}_{t_p,p}, \bar{y}_{t_p,p})\big\}$$

while the corresponding measurements are $\bar{z}_{t_1,1}, \dots, \bar{z}_{t_p,p}$, with the understanding that $\bar{z}_{t_\ell,\ell} = 0$ and $(x_{t_\ell,\ell}, y_{t_\ell,\ell})$ is not defined if $t_\ell = 0$. Let $\mathcal{T}_{k,p}$ be the set of p-dimensional vectors indexing the admissible trajectories ending in $\mathcal{S}_{k,p}$ at epoch p, i.e.,

$$\begin{aligned}\mathcal{R}_{k,p} = \Big\{&(t_1, \dots, t_p) : t_p = k \text{ and } \Delta_x^2(\bar{x}_{t_q,q} - \bar{x}_{t_h,h})^2 \\ &+ \Delta_y^2(\bar{y}_{t_q,q} - \bar{y}_{t_h,h})^2 \le (v_{\max}T)^2(q-h)^2, \\ &\forall\, q, h \text{ such that } t_q, t_h \ne 0\Big\}\end{aligned}$$

The GLTR (5) can now be recast as

$$\max_{k\in\{1,\dots,D_L\}} \underbrace{\max_{(t_1,\dots,t_L)\in\mathcal{R}_{k,L}} \sum_{\ell=1}^{L} \bar{z}_{t_\ell,\ell}}_{F_{k,L}} \underset{H_0}{\overset{H_1}{\gtrless}} \gamma_2.$$

An algorithm to compute $\{F_{k,L}\}_{k=1}^{D_L}$ is presented next.

4.2. Track formation

Let $T_{k,p}$ be the p-dimensional integer vector indexing the best (i.e., with the largest statistic) trajectory ending in $\mathcal{S}_{k,p}$ at epoch p, and let $F_{k,p}$ be the corresponding statistic:

$$T_{k,p} = \underset{(t_1,\dots,t_p)\in\mathcal{R}_{k,p}}{\arg\max} \sum_{\ell=1}^{p} \bar{z}_{t_\ell,\ell}$$

$$F_{k,p} = \max_{(t_1,\dots,t_p)\in\mathcal{R}_{k,p}} \sum_{\ell=1}^{p} \bar{z}_{t_\ell,\ell}.$$

Also, let $\mathcal{M}_{p,k}(q)$ denote the set of alarm indices at epoch q compatible with alarm k at epoch p, i.e.,

$$\begin{aligned}\mathcal{M}_{k,p}(q) = \Big\{&j \in \{1, \dots, D_q\} :\ \Delta_x^2(\bar{x}_{j,q} - \bar{x}_{k,p})^2 \\ &+ \Delta_y^2(\bar{y}_{j,q} - \bar{y}_{k,p})^2 \le (v_{\max}T)^2(q-p)^2\Big\}.\end{aligned}$$

Then, the following algorithm iteratively computes $\{T_{k,p}, F_{k,p},\ \ k = 1, \dots, D_p\}$ from $\{T_{k,q}, F_{k,q}, k = 1, \dots, D_q, q = 1, \dots, p-1\}$.

Algorithm 1.

1: ***Initialization:*** $p = 1$
2: ***for*** $k = 1, \dots, D_1$ ***do***
3: $\quad F_{k,1} = \bar{z}_{k,1}$
4: $\quad T_{k,1} = (k)$
5: ***end for***
6: ***Iteration:*** $p = 2, \dots, L$
7: ***for*** $k = 1, \dots, D_p$ ***do***
8: $\quad m(q) = \max_{j\in\mathcal{M}_{k,p}(p-q)} F_{j,p-q},\ q = 1, \dots, p-1$ $(\max \emptyset = 0)$
9: $\quad$ ***if*** $m(q) = 0\,\forall\, q$ ***then***
10: $\qquad F_{k,p} = \bar{z}_{k,p}$
11: $\qquad T_{k,p} = (\underbrace{0 \dots 0}_{p-1} k)$
12: $\quad$ ***else***
13: $\qquad F_{k,p} = \bar{z}_{k,p} + \max_{q\in\{1,\dots,p-1\}} m(q)$
14: $\qquad u = \arg\max_{q\in\{1,\dots,p-1\}} m(q)$
15: $\qquad w = \arg\max_{j\in\mathcal{M}_{k,p}(p-u)} F_{j,p-u}$
16: $\qquad T_{k,p} = (T_{w,p-u}\,\underbrace{0, \dots, 0}_{u-1}, k)$
17: $\quad$ ***end if***
18: ***end for***

It is worthwhile giving some comments on the iterative step of the above algorithm. In order to compute $F_{k,p}$ the algorithm searches in the set $M_{k,p}(p-q)$ for the best alarm at epoch $p-q$ that can be linked with the alarm $\mathcal{S}_{k,p}$ at epoch p (line 8): the largest statistic from epoch $p-q$, namely, $\max_{j\in\mathcal{M}_{k,p}(p-q)} F_{j,p-q}$, is stored in $m(q)$ for $q = 1, \dots, p-1$. If no past alarm can be linked to $\mathcal{S}_{p,k}$ at epoch p, $F_{k,p}$ is initialized with the current measurement $\bar{z}_{k,p}$ (line 10), and the corresponding trajectory has $p-1$ trailing zeros and k as the last entry (line 11). Otherwise the best admissible past alarm is linked with $\mathcal{S}_{k,p}$: $F_{k,p}$ is computed by adding the measurement $\bar{z}_{k,p}$ to the best previous metric (line 13), and the corresponding trajectory $T_{k,p}$ is updated accordingly (lines 14-16).

4.3. Complexity analysis

$\{D_\ell\}_{\ell=1}^{L}$ is a sequence of independent and identically distributed random variables. Under hypothesis H_K, $K \in \{0, 1\}$, each D_ℓ is the sum of two independent Binomial random variables with parameters $(N_x N_y - K, \mathrm{pfa}_1)$ and (K, pd_1). At iteration 1, no operation is required to evaluate the quantities $\{F_{k,1}, T_{k,1}, k = 1, \dots, D_1\}$. For $p \ge 2$, the kinematic constraints between alarms $\mathcal{S}_{k,p}$ and $\mathcal{S}_{j,p-q}$ must be checked for $j = 1, \dots, D_{p-q}$ and for $q = 1, \dots, p-1$; thus, the number of required operations is on the order of $D_p \sum_{q=1}^{p-1} D_{p-q}$. Since the iterative step of the algorithm runs from 2 to L, the average number of required operations is on the order of

$$\mathbb{E}_K\left[\sum_{p=2}^{L} D_p \sum_{q=1}^{p-1} D_{p-q}\right] = \frac{(L-1)L}{2}\left[(N_xN_y - K)\mathrm{pfa}_1 + K\mathrm{pd}_1\right]^2$$

where $\mathbb{E}_K$ denotes statistical expectation under hypothesis H_K. Hence, the average complexity of the algorithm is $\mathcal{O}\big(L^2[(N_xN_y - K)\mathrm{pfa}_1 + K\mathrm{pd}_1]^2\big)$ under H_K. Recall now that the complexity of the Viterbi-based routine is $\mathcal{O}(LN_1N_2)$, whereby the proposed procedure is preferable under both hypotheses if $N_xN_y\mathrm{pfa}_1 + 1 < \sqrt{N_xN_y/L}$, i.e., if

$$\begin{cases} N_xN_y > L \\ \mathrm{pfa}_1 < \frac{1}{\sqrt{LN_xN_y}} - \frac{1}{N_xN_y}. \end{cases}$$

The first inequality is usually met since, in general, the number of processed frames is much smaller than the number of pixels, while the second holds when γ_1 is set large enough. (i.e., for sparse data measurements).

5. NUMERICAL RESULTS

We discuss here a numerical example where $N_x = N_y = 400$, $\Delta_x = \Delta_y = 5$ m, $v_{\max} = 15$ m/s, and $T = 1$ s. Under H_1, we simulate a target in a random-walk motion with velocity randomly chosen between 0 and $v_{\max}$. The performance of the GLRT is assessed in terms of probability of false alarm (PFA), i.e., accept H_1 under H_0, and probability of correct detection (PD), i.e., accept H_1 under H_1. In Figure 2 we report PD vs PFA for $L = 8$, $\rho = 8, 12$ dB, and $\mathrm{pfa}_1 = \{10^{-i}\}_{i=2}^{5}$. The case $\mathrm{pfa}_1 = 1$ (i.e., uncensored data) is not included, since its computational complexity is too demanding. For the sake of comparison, we also show the performance when $L = 1$ and $\mathrm{pfa}_1 = 10^{-2}$, i.e., when no TBD processing is undertaken. For fixed $\mathrm{pfa}_1 = 10^{-2}$, the detector performance improves for increasing $L = 1, 4, 8$. Clearly, this gain comes at the price of a larger computational complexity, which scales as L^2 for the proposed algorithm. Remarkably, this complexity increment can be balanced by lowering pfa_1 to $10^{-3}, 10^{-4}, 10^{-5}$, i.e., by reducing the number of data measurements to be processed, while still achieving a significant performance improvement with the respect to the case where no TBD processing is undertaken.

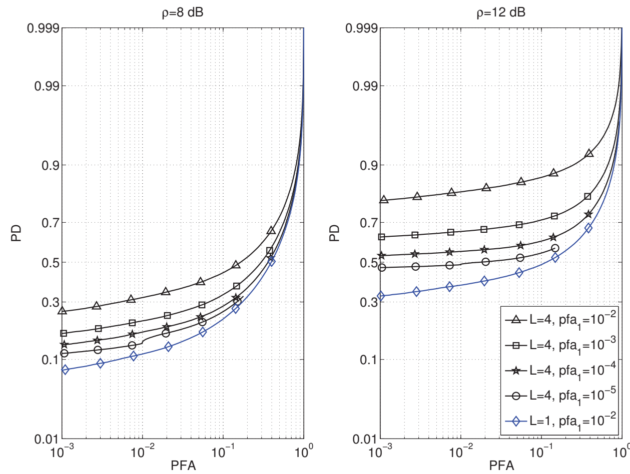

Fig. 1. PFA vs PD for $L = 4$.

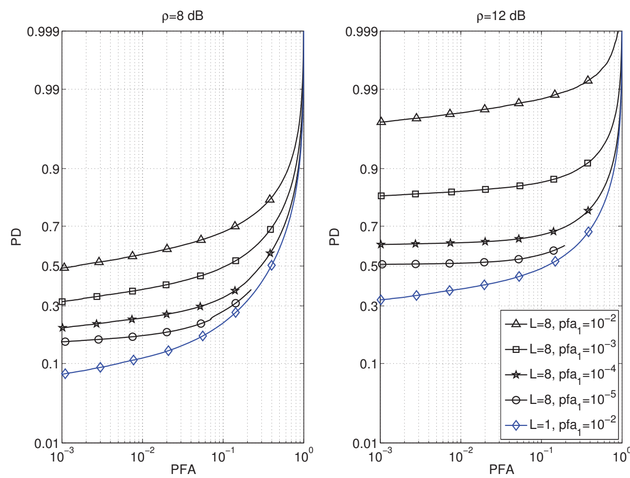

Fig. 2. PFA vs PD for $L = 8$.

6. REFERENCES

[1] Y. Barniv, "Dynamic programming solution for detecting dim moving targets," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 21, no. 1, pp. 144 –156, Jan. 1985.

[2] J. Arnold, S. Shaw, and H. Pasternack, "Efficient target tracking using dynamic programming," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 29, no. 1, pp. 44–56, Jan. 1993.

[3] S. M. Tonissen and R. J. Evans, "Peformance of dynamic programming techniques for track-before-detect," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 32, no. 4, pp. 1440–1451, Oct. 1996.

[4] H. Im and T. Kim, "Optimization of multiframe target detection schemes," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 35, no. 1, pp. 176–186, Jan. 1999.

[5] S. Buzzi, M. Lops, and L. Venturino, "Track-before-detect procedures for early detection of moving target from airborne radars," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 3, pp. 937–954, Jul. 2005.

[6] G. D. J. Forney, "The Viterbi algorithm," *Proc. IEEE*, vol. 29, no. 3, pp. 268–277, Mar. 1973.