

Multiframe Detector/Tracker: Optimal Performance

Marcelo G. S. Bruno⁽¹⁾ and José M. F. Moura⁽²⁾

(1) Elect. Eng. Depart., University of São Paulo, São Paulo, SP 05508-090, Brazil

ph: (55-11) 3818-5290; fax: (55-11) 3818-5128; email:bruno@lcs.poli.usp.br

(2) Depart. Elect. and Comp. Eng., Carnegie Mellon University, Pittsburgh, PA, 15213

ph: (412) 268-6341; fax: (412) 268-3890; email:moura@ece.cmu.edu

Abstract

We develop in this paper the optimal Bayes multiframe detector/tracker for rigid extended targets that move randomly in clutter. The performance of this optimal algorithm provides a bound on the performance of any other suboptimal detector/tracker. We determine by Monte Carlo simulations the optimal performance under a variety of scenarios including spatially correlated Gaussian clutter and non-Gaussian (K and Weibull) clutter. We show that, for similar tracking performance, the optimal Bayes tracker can achieve peak signal-to-noise ratio gains possibly larger than 10 dB over the commonly used combination of a spatial matched filter (spatial correlator) and a linearized Kalman-Bucy tracker. Simulations using real clutter data with a simulated target suggest similar performance gains when the clutter model parameters are unknown and estimated from the measurements.

Keywords

Multiframe Detection and Tracking; Nonlinear Stochastic Filtering; Noncausal Gauss-Markov Random Fields.

I. INTRODUCTION

The paper studies integrated detection and tracking of randomly moving targets in clutter using as input data a sequence of noisy images. The images may be collected by electromagnetic sensors such as high resolution radars, or optical sensors such as infrared (IR) devices. At each sensor scan, an image or frame is produced. If one or more targets are present during a scan, the corresponding image contains the returns from the targets plus the returns from the background clutter. Otherwise, if no target is present, the sensor return consists exclusively of clutter. The clutter accounts for spurious reflectors, which may appear as false targets, and for measurement noise.

In the case when multiple, at most M , targets of interest are present, the multitarget detector decides from the noisy data how many targets $(0,1,2,\dots,M)$ are present in each frame. Once a target is declared present by the detector, a subsequent tracker estimates its position in the surveillance space. The interpolation across successive scans of the estimated positions of a target forms a track for that target. Due to clutter, false detections, known as false alarms, may occur, and false tracks may be estimated. Conversely, actual targets may fail to be detected. This situation is known as a miss. Even if correct detections (i.e., no misses or false alarms) occur, the background clutter can still cause the tracker to produce a wrong estimate of the target's position, i.e., a tracking error.

The ultimate goal is to estimate the target state, typically a collection of kinematic components such as position, velocity, or acceleration. In most existing algorithms, e.g., [1], detection and tracking are two separate stages. The measurements of interest to the tracker are not the raw sensor images, but the outputs of preliminary detection subsystems. The detection stage involves the thresholding of the raw data, usually one single sensor frame. After further preprocessing, validated detections provide measurements that, for targets that are declared present, are treated as noise-corrupted observations of the target state such as, for example, direct estimates of position (range, azimuth, and elevation). Due to the occurrence of random false alarms in the detection process, or due to clutter coming from spurious reflectors, interfering targets, or man-made decoys, validated measurements may actually be false measurements that do not originate from true targets.

Multitarget trackers generally assume [1], [8] that the targets are pointwise and associate a linear (or linearized) dynamic model to the state of each target of interest. A tracking filter, usually a variation

on the Kalman-Bucy filter, combines the validated measurements with the dynamic model, providing an estimate of the state of the target. An important issue arising from the decoupling of detection and tracking is the problem of deciding which set of measurements or weighted combination of measurements should be associated to each target state estimator or to clutter. This problem is known as data association. The most common data association algorithms, see [8], compute posterior probabilities of association conditioned on the measurements and use them throughout the estimation process.

Brief review of the literature References concerned only with target detection, not tracking, include [4], [5], [6]. In [5], Pohlig introduces an algorithm for detection of constant velocity objects such as meteors, asteroids, and satellites, in fixed stellar backgrounds. The measurements are obtained by a staring sensor with an array of charged coupled device (CCD) sensors in the focal plane of a telescope. The focal plane image is integrated and sampled in space and time, resulting in a three-dimensional (two spatial dimensions and one temporal dimension) discrete model, where the optical intensity of both targets and the background are modeled as Poisson distributions with different means that reflect the different photon counts arising from targets and clutter. Pixel intensities under both hypotheses of presence and absence of target are assumed spatially uncorrelated. The detection algorithm in [5] is a 3D generalized likelihood ratio test (GLRT) based on batch processing: all available sensor frames are stacked in a data volume, and then the GLRT decides on the presence or absence of a target anywhere in that volume.

The work by Reed, Gagliardi, and Shao [6] is similar in nature to Pohlig's approach and introduces a 3D (again space plus time) matched filter for detection of known, moving targets within a Gaussian background clutter with known spectral density. However, unlike reference [5], reference [6] considers the case of continuous (non-sampled) data; it is best suited for optical rather than digital processing. A different problem is considered by Chen and Reed in [4]. The goal in [4] is to introduce a constant false alarm rate (CFAR) algorithm to solve the problem of detection of a known target signal in a given scene, using a set of K correlated reference scenes that contain no target or, alternatively, very weak target returns. The reference scenes are obtained either from different frequency bands of the main scene (multispectral or hyperspectral imagery) or from sequential observations in time. The proposed detection algorithm is a generalized likelihood ratio test (GLRT) that tests for the presence or absence of a target in the main scene using as data the entire collection of reference scenes plus the main scene itself. The underlying model assumes that, after pre-processing (essentially removal of the local variable mean), each individual scene is a zero-mean, Gaussian, white random vector, i.e., the spatial correlation between the pixels in each individual image is neglected. However, the model assumes a cross-correlation between pixels at the same spatial location in different scenes. An alternative modeling for multispectral imagery that incorporates both interframe and intraframe correlation was proposed in [7].

In this paper, instead of decoupling detection and tracking as in [1], or considering detection-only of moving objects as in [5], [6], we develop the optimal, multiframe, Bayes detector/tracker that processes directly the sensor images and integrates detection and tracking into a unified framework. The Bayesian strategy involves the computation at each scan of the posterior probability of the unknown target states conditioned on the observations. In [3], the author uses a dynamic programming approach and the Viterbi

algorithm to study target detection. We postpone to section III-E a detailed discussion comparing the Bayes algorithm with the dynamic programming approach in [3].

In our approach, we integrate detection and tracking into the same framework by augmenting the target state space with additional dummy states that represent the absence of targets. The posterior probability of a given target being absent is propagated in time together with the posterior probabilities of the other “present target” states. In contrast to Pohlig’s *batch* detector [5], we develop a *recursive* framework where we still process all frames available in an optimal way, but these frames are processed one by one and discarded as we finish processing them. As a new frame is available, we simply update the posterior probabilities for the target states by running one more iteration of the algorithm.

Modeling assumptions The optimal Bayesian algorithm takes full advantage of all prior information on the clutter, target signature, and target motion models, and allows multiframe detection and tracking with recursive processing across all observed sensor scans. We consider in this paper both pointwise (single pixel) and extended (multipixel) targets. We present detection results for targets with deterministic signatures and for targets with time-varying random signatures. The random signatures are described by multivariate, spatially correlated Gaussian distributions. We assume translational motions, and we define as the target state the spatial coordinates of the target’s geometric centroid. Since practical sensors have a finite resolution, we restrict the target centroid positions to a finite grid where each pixel represents a resolution cell of the sensor. We describe motions by finite state machines (FSMs) obtained by discretizing the continuous differential equations that describe the target dynamics. The dummy states that represent the absence of a target are incorporated into the FSM model that also specifies the transition probabilities between the absence and the presence of a target, and vice-versa.

We consider two classes of clutter models: spatially correlated clutter with Gaussian statistics, and uncorrelated non-Gaussian clutter with heavy tail amplitude (envelope) statistics. The spatial correlation of the clutter is captured by using noncausal, spatially homogeneous, Gauss-Markov random fields (GMrfs) of arbitrary order [25]. GMrfs are statistical models that capture the locality properties of the clutter, namely, the clutter at a given spatial location is strongly dependent on the clutter intensity in neighboring locations. This assumption is intuitively realistic in many practical scenarios. Regarding non-Gaussian clutter, we represent it by spherically invariant-random vectors (SIRVs) [13], [14], [15], which have been shown to generate a variety of envelope statistics of practical interest, including the Weibull, K, Rician [18], and G [19] envelopes.

Performance studies This paper focuses on performance results for the optimal multiframe Bayes detector/tracker in a variety of scenarios, including, as we mentioned before, both deterministic and random signature targets, observed in both Gaussian and non-Gaussian clutter. We test the proposed algorithm primarily on synthetic data with known clutter and target models. The optimal performance curves, obtained through extensive Monte Carlo simulations, provide an upper bound to the performance of suboptimal algorithms. We benchmark against these bounds the performance of competing suboptimal schemes such as the association of a single frame spatial correlator (matched filter) with a multiframe linearized Kalman-Bucy filter (KBf) tracker. These studies show that there is a significant margin of

improvement to be had over existing detectors and trackers.

In practice, the situation of perfect match between the data and the model is not realistic. In order to assess the robustness of the algorithm to mismatches between the measurements and the model, we present an example of detection/tracking with real clutter data, obtained by a laser radar mounted to the bottom of an aircraft. We fit the model to the real clutter by estimating its parameters from the data. The experimental results confirm that there is a significant improvement in performance over conventional algorithms such as a plain single frame image correlator associated with a KBF.

Summary of the paper The paper is divided into 6 sections. Section I is this introduction. Section II presents the models for sensor, target, motion, and clutter that underly our integrated approach to detection and tracking. Section III examines the derivation of the optimal Bayesian detector/tracker based on the models from section II. Sections IV and V quantify respectively the detection and tracking performances of the algorithm through comprehensive Monte Carlo simulations assuming a single target scenario. Both correlated Gaussian clutter and non-Gaussian clutter situations are considered, and performance comparisons with alternative suboptimal detection and tracking algorithms are provided. Finally, section VI summarizes the contributions of the paper.

We omit in this paper specific details on the implementation of the Bayes detector/tracker. These can be found in reference [24] for the particular case of a single, deterministic 2D target observed in GMrf clutter.

II. THE MODEL

At each sensor scan, there are at most M targets present in the surveillance space. Each target is a rigid body with translational motion belonging to one of M possible classes characterized by their signature parameters, and by the dimensions of their noise-free image. We restrict our discussion to the situation where all targets are distinct. For simplicity of notation, we restrict this section to one-dimensional (1D) spaces. A brief discussion on the corresponding 2D models and a comprehensive investigation of 2D detection/tracking performance are found in section V (see also [24] for further details on modeling and implementation of the 2D detector/tracker algorithm).

A. Surveillance Space and Target Model

We first model the surveillance space of the sensor. Given the sensor's finite resolution, we discretize the 1D space by the uniform finite discrete lattice

$$\mathcal{L} = \{l: 1 \leq l \leq L\} \tag{1}$$

where L is the number of resolution cells and l is an integer. We refer to the lattice \mathcal{L} as the *sensor lattice*. The resolution cells are also referred to as pixels.

To develop an integrated framework for detection and tracking, it is useful to extend the lattice \mathcal{L} with additional states that will be used to represent the absence of targets and to account for the fact that target images extend over more than one pixel in the sensor lattice. We introduce first the vector $\mathbf{Z}_n = [z_n^1 \dots z_n^M]^T$, which collects the positions of the geometric centroids of the M possible targets in the sensor image at scan n .

Let the pixel length S^p of a class p target, $1 \leq p \leq M$, be $S^p = (l_i^p + l_s^p + 1)$, where l_i^p and l_s^p are the maximum extent in pixels of the target, respectively to the left and to the right of its centroid. These parameters are assumed known and time-invariant in the paper. If S^p is odd, we make $l_i^p = l_s^p = (S^p - 1)/2$. Otherwise, if S^p is even, we adopt the convention that $l_s^p = S^p/2$ and $l_i^p = l_s^p - 1$. To account for the situations when targets move in and out of the sensor range, we define the extended *centroid lattice*,

$$\bar{\mathcal{L}}_p = \{l: -l_s^p + 1 \leq l \leq L + l_i^p\} \quad (2)$$

which corresponds to the set of all possible centroid positions z_n^p such that at least one pixel of the target is still visible in the sensor image. Finally, to include the possibility of absence of a target, we introduce an additional dummy state. We adopt the convention that, whenever a class p target is not present at the n th scan, z_n^p takes the value $L + l_i^p + 1$. With the addition of this dummy absent target state, we define the *augmented lattice*,

$$\tilde{\mathcal{L}}_p = \{l: -l_s^p + 1 \leq l \leq L + l_i^p + 1\}. \quad (3)$$

Extended Target Model When a class p target is present, i.e., $z_n^p \in \bar{\mathcal{L}}_p$, the noise free target image is simply the spatial distribution of the target pixel intensities, a_k^p , $-l_i^p \leq k \leq l_s^p$, centered at the centroid lattice cell z_n^p . Otherwise, if $z_n^p = L + l_i^p + 1$, meaning the target is absent, the sensor image corresponding to target-only returns reduces to a null image. These intuitive ideas are formalized mathematically by expressing the noise free image of a class p target at the n th sensor scan as the nonlinear function

$$\mathbf{t}_p(z_n^p) = \sum_{k=-l_i^p}^{l_s^p} a_k^p \mathbf{e}_{z_n^p+k} \quad z_n^p \in \bar{\mathcal{L}}_p \quad (4)$$

$$\mathbf{t}_p(z_n^p) = \mathbf{0}_L \quad z_n^p = L + l_i^p + 1 \quad (5)$$

where \mathbf{e}_l , $1 \leq l \leq L$, is an L -dimensional vector whose entries are all zero, except for the l th entry which is one. If $l < 1$ or $l > L$, \mathbf{e}_l is defined as the identically zero vector. This particular definition for \mathbf{e}_l outside the original sensor grid \mathcal{L} is adopted to guarantee that the target model in (4) will accurately describe the disappearance of portions of the target from the sensor image as the target's centroid moves closer to the boundaries of the surveillance space.

The pixel intensity coefficients a_k^p in (4) are also known as the *target signature* coefficients. They may be deterministic and known, deterministic and unknown, or random. Random signatures account for fluctuations in the reflectivity, or in the conditions of illumination of the target, as well as random variations in channel characteristics such as fading. For simplicity, we assume in most of this paper that the signature coefficients are known and time-invariant. An extension of the detection/tracking algorithms to targets with random signature in section III-D. Monte Carlo simulations with synthetic spatially correlated/temporally uncorrelated Gaussian targets are presented in section IV.

B. Multitarget Observations and Clutter Models

We consider a multitarget scenario with M possible targets, and collect the L sensor readings at each pixel of the n th scan in the L -dimensional column vector \mathbf{y}_n . Due to the presence of spurious reflectors and background, \mathbf{y}_n consists of the superposition of the various noise-free target images plus clutter. Using the extended target model introduced in subsection II-A, the observation vector or n th sensor frame, \mathbf{y}_n , is given by

$$\mathbf{y}_n = \mathbf{t}_1(z_n^1) + \mathbf{t}_2(z_n^2) + \dots + \mathbf{t}_M(z_n^M) + \mathbf{v}_n \quad (6)$$

where \mathbf{v}_n is the background clutter vector, also referred to as the n th clutter frame, and $\mathbf{t}_p(z_n^p)$, $1 \leq p \leq M$, is given by equations (4) and (5), depending on whether the p th target is present or absent at the n th scan. The clutter \mathbf{v}_n is assumed to be statistically independent of $\mathbf{t}_p(z_n)$, $1 \leq p \leq M$.

At each frame, the clutter at a given spatial location (pixel) may be statistically correlated with the clutter at another spatial location. The clutter intensity may also have Gaussian or non-Gaussian statistics. We adopt one of three models for \mathbf{v}_n : spatially white Gaussian clutter; spatially correlated Gaussian clutter; and spatially white non-Gaussian clutter. These models allow us to assess how clutter spatial correlation or non-Gaussian clutter statistics affect the performance of the detection/tracking algorithms. **Gaussian clutter** under the assumption of Gaussianity, the vector \mathbf{v}_n has a multivariate normal probability density function (pdf), $p(\mathbf{v}_n) = N(\mathbf{0}, \mathbf{R})$, where \mathbf{R} is the clutter spatial covariance, and $\mathbf{0}$ is the mean. The zero mean assumption assumes a pre-processing stage that removes the possibly spatially variant local mean. A non-zero mean can be accounted for trivially. We distinguish two cases for the covariance matrix \mathbf{R} .

White spatially homogeneous Gauss clutter: With spatially uncorrelated (white) clutter, the covariance matrix \mathbf{R} is diagonal. Assuming spatial homogeneity, $\mathbf{R} = \sigma_v^2 \mathbf{I}$, where \mathbf{I} is the identity matrix and σ_v^2 is the variance (power) of the clutter.

Spatially correlated homogeneous Gauss-Markov clutter: We model spatially correlated clutter as a Gauss-Markov random field (GMrf) [25]. This model simply states that the clutter intensity at a given pixel of the sensor image is a weighted average of the clutter intensity in neighboring pixels plus an error term. We assume in this paper a *noncausal* neighborhood region for each pixel. If we add the assumption of spatial homogeneity, an m th order 1D noncausal GMrf model for the n th clutter frame is given by the spatial difference equation

$$v_n(l) = \sum_{j=1}^m \alpha_j [v_n(l-j) + v_n(l+j)] + u_n(l) \quad 1 \leq l \leq L \quad (7)$$

where $u_n(l)$ is a zero-mean, correlated prediction error such that

$$E[v_n(l)u_n(k)] = 0 \quad \forall k \neq l \quad (8)$$

and the symbol $E[\cdot]$ stands for expectation or ensemble average. In order to completely define equation (7) at all pixel locations, we specify boundary conditions (bc's) outside the sensor lattice \mathcal{L} . Common boundary conditions are simply $v_n(l) = 0$ for $l < 1$ or $l > L$. These are known as Dirichlet bc's. Other bc's can be alternatively used, see for example [23], [25].

Second-Order Statistics of GMrfs: The GMrf model is very attractive because it provides a simple parameterization for the inverse of the covariance matrix of the background clutter \mathbf{v}_n . Collecting the clutter samples $v_n(l)$ and the error samples $u_n(l)$, $1 \leq l \leq L$, in two L -dimensional vectors \mathbf{v}_n and \mathbf{u}_n , an equivalent matrix representation for the difference equation in (7) is

$$\mathbf{A} \mathbf{v}_n = \mathbf{u}_n \quad (9)$$

where \mathbf{A} is a sparse and highly structured matrix usually referred to as the *potential matrix*. For the 1D m th order homogeneous model in (7), the potential matrix is an m -banded, Toeplitz, symmetric matrix

with structure [25]

$$\mathbf{A} = \begin{bmatrix} 1 & -\alpha_1 & -\alpha_2 & \dots & -\alpha_m & 0 & 0 & \dots & 0 & 0 \\ -\alpha_1 & 1 & -\alpha_1 & \dots & -\alpha_{m-1} & -\alpha_m & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -\alpha_{m-1} & -\alpha_{m-2} & -\alpha_{m-3} & \dots & 1 & -\alpha_1 \\ 0 & 0 & 0 & \dots & -\alpha_m & -\alpha_{m-1} & -\alpha_{m-2} & \dots & -\alpha_1 & 1 \end{bmatrix}. \quad (10)$$

We now derive the second-order statistics of \mathbf{u}_n , which is referred to as the prediction error, and of the clutter field \mathbf{v}_n . Combining the orthogonality condition in (8) with the matrix equation (9), we note that

$$E[\mathbf{u}_n \mathbf{u}_n^T] = E[\mathbf{A} \mathbf{v}_n \mathbf{u}_n^T] = \sigma_u^2 \mathbf{A} \quad (11)$$

where the superscript ‘‘T’’ denotes the transpose of a vector or matrix. In (11), we used the assumption of spatial homogeneity (roughly speaking, the spatial ‘‘equivalent’’ of stationarity) to make $E[v_n(l)u_n(l)] = \sigma_u^2$, for all l , $1 \leq l \leq L$.

Finally, since \mathbf{A} is nonsingular and symmetric, then $\mathbf{v}_n = \mathbf{A}^{-1} \mathbf{u}_n$ and

$$\begin{aligned} \Sigma_v^{-1} &= (E[\mathbf{v}_n \mathbf{v}_n^T])^{-1} = (\mathbf{A}^{-1} E[\mathbf{u}_n \mathbf{u}_n^T] \mathbf{A}^{-T})^{-1} \\ &= (\sigma_u^2 \mathbf{A}^{-1} \mathbf{A} \mathbf{A}^{-T})^{-1} = \mathbf{A} / \sigma_u^2. \end{aligned} \quad (12)$$

Equation (12) gives for free, with no matrix inversion required, the inverse of the clutter covariance in terms of the highly structured matrix \mathbf{A} given in equation (10). This structure is used to design computationally efficient detection and tracking algorithms when the clutter is correlated as a Gauss-Markov random field. Finally, for our simulation studies, we use equation (9) and a technique based on the upper Cholesky factorization of the potential matrix \mathbf{A} , [25], to generate samples of the GMrf clutter \mathbf{v}_n .

Non-Gaussian clutter When dealing with non-Gaussian clutter, we assume that the sensor measures, at each resolution cell, the in-phase and quadrature returns of the clutter and targets echoes. The clutter measurements at instant n correspond to a sampling of the returned clutter complex envelope and are given by the even-sized vector

$$\mathbf{v}_n = [v_{c_n}^1 v_{s_n}^1 \dots v_{c_n}^L v_{s_n}^L] \quad (13)$$

where L is the number of resolution cells. We assume that the double-sized vector \mathbf{v}_n has a joint pdf with *non-Gaussian* statistics such that the sequence of random variables

$$e_k = \sqrt{(v_{c_n}^k)^2 + (v_{s_n}^k)^2} \quad 1 \leq k \leq L \quad (14)$$

is identically distributed with a probability density function different from a Rayleigh distribution.

K and Weibull envelope statistics

We are interested in analyzing how the tracker performs against a background clutter whose envelope at each resolution cell has heavier tails than a Rayleigh envelope. Useful clutter envelope statistics are the K and Weibull models that are frequently used in the literature to represent the amplitude statistics of clutter returns [20], [21], [22]. The corresponding pdfs for the two models are [18]

1. K pdf: $p_E(e) = \frac{b^{\nu+1} e^\nu}{2^{\nu-1} \Gamma(\nu)} K_{\nu-1}(be) \quad e \geq 0$

where ν is a shape parameter, $\Gamma(\cdot)$ is the Eulerian function, $K_{\nu-1}(\cdot)$ is a modified Bessel function of the second kind and b is related to σ^2 by $b^2 = (2\nu)/\sigma^2$.

2. Weibull pdf: $p_E(e) = a c e^{c-1} \exp(-a e^c) \quad e \geq 0$

where c is a shape parameter and a relates to the average power σ of the quadrature components by $2\sigma^2 = a^{-2/c} \Gamma(1 + \frac{2}{c})$.

Simulation of K and Weibull clutter samples Rayleigh envelope statistics correspond to a multivariate joint Gaussian distribution of the in-phase and quadrature clutter returns. Similarly, heavy-tailed envelope statistics such as the Weibull and K distributions correspond to a generalized spherically-invariant random vector (SIRV) model in the backscatter domain [13], [14]. Techniques to simulate heavy-tailed clutter using SIRV models have been discussed extensively in the literature [14], [15], [16]. In particular, we used the algorithms in [16] to generate the samples of uncorrelated K and Weibull clutter that were used in the Monte Carlo simulations in section IV-B. We omit the simulation details here for lack of space and refer the reader instead to the literature, particularly [16].

C. Target Motion

Assuming that the targets are rigid bodies with translational motion, the target motion is completely specified by the dynamics of the target centroid. We adopt a first order statistical model for the centroid dynamics. Given the sensor finite resolution, we model the motion of a class p target in the corresponding augmented lattice $\tilde{\mathcal{L}}_p$ by a set of transition probabilities

$$\{P(z_n^p = k \mid z_{n-1}^p = j)\} \quad k, j \in \tilde{\mathcal{L}}_p. \quad (15)$$

The transition probabilities $P(z_n^p \mid z_{n-1}^p)$, $1 \leq p \leq M$, represent the likelihood of displacement of a class p target between two consecutive frames. The transition probabilities in (15) define a finite state machine (FSM) that specifies the dynamics of the centroid for the class p target.

Example: Targets with constant mean velocity

A particular dynamic model of interest is a target where the nominal velocity is constant. We perturb this mean or nominal velocity by an m th order random walk fluctuation. The target centroid position at instant n is then given by

$$z_{n+1} = z_n + d + \varepsilon_n \quad (16)$$

where d is the mean velocity, and ε_n is a discrete-valued, zero-mean white noise component that is independent of the centroid position and takes values on the discrete set $\mathcal{S} = \{-m, \dots, -1, 0, 1, \dots, m\}$ for some $m \geq 1$. Figure 1 shows the central section of the FSM that corresponds to the model in (16) when $d = 0$ and $m = 1$. This FSM is simply a first order discrete Markov chain. As mentioned




Fig. 1. Example of Finite State Machine Diagram

before, a target that is present moves to the absent state whenever its centroid is outside the lattice $\bar{\mathcal{L}}_p = \{l: -l_s^p + 1 \leq l \leq L + l_i^p\}$. When no target of a given class is present, we assume that there is

a non-zero probability p_a of a new target from that same class appearing randomly at the next sensor scan. We assume that the target centroid may appear at any pixel of the centroid lattice $\bar{\mathcal{L}}_p$, with equal probability $p_a/(L + l_i^p + l_s^p)$. This assumption is a *worst case* scenario, when the detector/tracker has no a priori information about initial position of a new target. Other more elaborate distributions for the probability of reappearance are easily taken into account.

III. OPTIMAL BAYES MULTITARGET DETECTOR/TRACKER

We assume that at each scan n an unknown number of targets ranging from zero to M may be present. The targets that are present belong to distinct classes (i.e., in the context of this model, they have different signatures). We collect the observation scans from instant 0 up to instant n in the long observation vector $\mathbf{Y}_0^n = [\mathbf{y}_0^T \dots \mathbf{y}_n^T]^T$. Given \mathbf{Y}_0^n , we want to perform three tasks at instant n : (1) determine how many targets are present/absent (detection); (2) assign the detected targets to a given class (data association); (3) estimate the positions of the detected targets (tracking).

A. Nonlinear Stochastic Filtering Approach

As mentioned in section II, the vector

$$\mathbf{Z}_n = [z_n^1 \dots z_n^M]^T, \quad (17)$$

collects the positions of the centroids of the M possible targets in the sensor image. If all $z_n^p = L + l_i^p + 1$, $1 \leq p \leq M$, then no target is present in the surveillance space at the n th scan. The optimal Bayes solution to the joint detection/tracking problem is obtained by computing at each scan the joint posterior probability, $P(\mathbf{Z}_n | \mathbf{Y}_0^n)$, i.e., the conditional probability of the vector

The Bayes detector/tracker that we present processes the observations as they become available. It computes recursively $P(\mathbf{Z}_n | \mathbf{Y}_0^n)$ at each scan, thus avoiding having to store all the measurements from instant zero up to the present. The recursion is divided into two steps. The first step is the prediction step: it uses the statistical description of the target motion between two consecutive scanned frames to predict the current position of the targets based on all past observations. Once a new sensor frame is available, a second step, known as the filtering step, uses the new measurements to correct the prediction. The incoming sensor data is processed using the information in the clutter and target signature models. In the sequel, we describe both steps in further detail.

The following assumptions are made in the derivation of the algorithm:

1. In each frame, only one target from each of the M possible classes may be present.
2. The sequence of clutter frames $\{\mathbf{v}_n\}$, $n \geq 1$, is independent, identically distributed (i.i.d.).
3. The sequence of target states $\{\mathbf{Z}_k\}$, $k \geq 0$, is statistically independent of the sequence of clutter frames $\{\mathbf{v}_k\}$, $k \geq 0$.
4. Targets from different classes move independently and the translational motions for targets from each class are described by first-order discrete Markov Chains completely specified by the transition probabilities $P(z_n^p | z_{n-1}^p)$, $1 \leq p \leq M$, $z_n^p \in \tilde{\mathcal{L}}_p$.
5. In all observed frames, the target signatures are deterministic and known (but not necessarily time-invariant) for each target class.

We make the following remarks regarding the previous assumptions:

- a) The detector/tracker algorithm can be easily modified to account for unknown, random target signatures in each sensor frame. We discuss the necessary modifications in subsection III-D.
- b) Instead of assuming that at most one target from each class is present in each frame, we could have used an alternative problem setup in which there is a *known* maximum number of targets, $\mathcal{N}_p \geq 1$, for each target class p . In this paper, for convenience, and without loss of generality, we make $\mathcal{N}_p = 1$, $\forall p$, $1 \leq p \leq M$.
- c) The assumption that the sequence $\{\mathbf{v}_n\}$, $n \geq 0$, is i.i.d. is equivalently to ignoring all interframe statistical correlation between the clutter pixels. The 2D GMrf model in section II-B assumes however an intraframe or spatial clutter correlation.

We now detail the derivation of the algorithm. In the subsequent derivation, we denote the probability mass function of discrete-valued random variables by the capital letter P , whereas the probability density function of continuous-valued random variables is denoted by lowercase p .

Prediction Step This step computes the *prediction posterior probability*

$$P(\mathbf{Z}_n | \mathbf{Y}_0^{n-1}) = P(z_n^1, \dots, z_n^M | \mathbf{Y}_0^{n-1}) \quad z_n^p \in \tilde{\mathcal{L}}_p \quad 1 \leq p \leq M. \quad (18)$$

From $P(\mathbf{Z}_n | \mathbf{Y}_0^{n-1})$, we can obtain the marginal posterior probabilities of the centroid position of each target conditioned on the past frames from instant 0 to instant $n - 1$. We also obtain the posterior probabilities of absence of each target conditioned on the past observations.

Combining the theorem of Total Probability with Bayes law, we write

$$\begin{aligned} P(\mathbf{Z}_n | \mathbf{Y}_0^{n-1}) &= \sum_{\mathbf{Z}_{n-1}} P(\mathbf{Z}_n, \mathbf{Z}_{n-1} | \mathbf{Y}_0^{n-1}) \\ &= \sum_{\mathbf{Z}_{n-1}} P(\mathbf{Z}_n | \mathbf{Z}_{n-1}, \mathbf{Y}_0^{n-1}) P(\mathbf{Z}_{n-1} | \mathbf{Y}_0^{n-1}). \end{aligned} \quad (19)$$

Since the sequence of target centroid positions $\{\mathbf{Z}_k\}$, $k \geq 1$, is, by assumption, a first-order Markov process, then, conditioned on \mathbf{Z}_{n-1} , the current state \mathbf{Z}_n is statistically independent of the sequence $\{\mathbf{Z}_k\}$, $0 \leq k \leq n - 2$. If we add the assumption that \mathbf{Z}_n is also independent of the sequence of previous clutter frames $\{\mathbf{v}_k\}$, $0 \leq k \leq n - 1$, $n \geq 1$, we conclude that, conditioned on \mathbf{Z}_{n-1} , \mathbf{Z}_n is statistically independent of the previous observations, \mathbf{Y}_0^{n-1} , i.e.,

$$P(\mathbf{Z}_n | \mathbf{Z}_{n-1}, \mathbf{Y}_0^{n-1}) = P(\mathbf{Z}_n | \mathbf{Z}_{n-1}). \quad (20)$$

Replacing (20) in (19), we get

$$P(\mathbf{Z}_n | \mathbf{Y}_0^{n-1}) = \sum_{\mathbf{Z}_{n-1}} P(\mathbf{Z}_n | \mathbf{Z}_{n-1}) P(\mathbf{Z}_{n-1} | \mathbf{Y}_0^{n-1}). \quad (21)$$

Finally, assuming that the different targets move according to statistically independent Markov chains,

$$P(\mathbf{Z}_n | \mathbf{Z}_{n-1}) = P(z_n^1 | z_{n-1}^1) \dots P(z_n^M | z_{n-1}^M) \quad (22)$$

and we write

$$P(\mathbf{Z}_n | \mathbf{Y}_0^{n-1}) = \sum_{z_{n-1}^1} \dots \sum_{z_{n-1}^M} P(z_n^1 | z_{n-1}^1) \dots P(z_n^M | z_{n-1}^M) P(\mathbf{Z}_{n-1} | \mathbf{Y}_0^{n-1}). \quad (23)$$

Filtering Step We now compute the *filtering posterior probability*, $P(\mathbf{Z}_n | \mathbf{Y}_0^n)$. From Bayes' law,

$$P(\mathbf{Z}_n | \mathbf{Y}_0^n) = P(\mathbf{Z}_n | \mathbf{y}_n, \mathbf{Y}_0^{n-1}) \quad (24)$$

$$= C_n p(\mathbf{y}_n | \mathbf{Z}_n, \mathbf{Y}_0^{n-1}) P(\mathbf{Z}_n | \mathbf{Y}_0^{n-1}) \quad (25)$$

$$= C_n p(\mathbf{y}_n | \mathbf{Z}_n) P(\mathbf{Z}_n | \mathbf{Y}_0^{n-1}) \quad (26)$$

where C_n is a normalization constant. To write equation (26), we used the fact that the sequence of clutter vectors $\{\mathbf{v}_n\}_{n=0}^\infty$ is i.i.d and independent of the sequence of state vectors $\{\mathbf{Z}_n\}_{n=0}^\infty$. Hence, given \mathbf{Z}_n , vector \mathbf{y}_n is independent of \mathbf{Y}_0^{n-1} . The term $p(\mathbf{y}_n | \mathbf{Z}_n)$ is referred to in the nonlinear stochastic filtering literature as the *observations kernel* [29], [31] and specifies the conditional statistics of the observed data assuming that the targets' states (positions) are known. The analytical expression for the observation kernel depends on the clutter and target models. We present next the optimal detection and tracking algorithms.

B. Minimum probability of error Bayes detector

For each of the M possible targets, there are two possible detection states during the n th scan: absent or present. The detection algorithm is therefore a statistical test that, based on all present and past observed data, \mathbf{Y}_0^n , chooses one among 2^M possible hypotheses H_m , $0 \leq m, \leq 2^M - 1$. In this notation, hypothesis H_0 stands for “all M possible targets are absent”. Conversely, hypothesis H_{2^M-1} means that all M possible targets are present. Hypotheses H_m , $m \neq 0$ and $m \neq 2^M - 1$, represent all other combinations in between of presence and absence of the multiple targets.

Given $P(\mathbf{Z}_n | \mathbf{Y}_0^n)$, we compute the posterior probabilities of the detection hypothesis H_m , $0 \leq m \leq 2^M - 1$. The minimum probability of error detector decides that hypothesis H_m is true if [32]

$$P(H_m | \mathbf{Y}_0^n) > P(H_k | \mathbf{Y}_0^n) \quad \forall k \neq m, \quad 0 \leq m, k \leq 2^M - 1 \quad (27)$$

where $P(H_m | \mathbf{Y}_0^n)$ is the posterior probability of hypothesis H_m . We describe two illustrative examples.

Example 1: Single Target

With a single target, there are only two possible hypotheses at each sensor scan:

1. H_0 : {target absent}.
2. H_1 : {target present}.

The minimum probability of error detector assuming equal cost assignment to misses and false alarms and zero cost assignment to correct decisions reduces to

$$\frac{P(H_0 | \mathbf{Y}_0^n)}{P(H_1 | \mathbf{Y}_0^n)} \underset{H_1}{\overset{H_0}{>}} 1. \quad (28)$$

Introducing the posterior probability vector, $\mathbf{f}_{n|n}$, such that its k th component is

$$f_{n|n}(k) = P(z_n^1 = k | \mathbf{Y}_0^n) \quad k \in \tilde{\mathcal{L}}_1 \quad (29)$$

then

$$P(H_0 | \mathbf{Y}_0^n) = f_{n|n}(L + l_i^1 + 1) \quad (30)$$

$$P(H_1 | \mathbf{Y}_0^n) = 1 - f_{n|n}(L + l_i^1 + 1). \quad (31)$$

Remark: The decision rule in (28) minimizes the total probability of decision errors, misses, and false alarms. Alternatively, if we change the threshold in (28) and vary it over a wide range, the detection algorithm operates as a Neyman-Pearson detector [32] that maximizes the probability of detection for a given probability of false alarm.

Example 2: Two Targets

We illustrate next how to compute the quantities $P(H_m | \mathbf{Y}_0^n)$ from the filtering posterior probability $P(\mathbf{Z}_n | \mathbf{Y}_0^n)$ when there are two targets, i.e., $M = 2$. With two targets, there are 4 possible hypotheses for the presence or absence of targets at the n th sensor scan:

$H_0 = \{\text{Both targets absent}\}$.

$H_1 = \{\text{Target 1 absent and Target 2 present}\}$.

$H_2 = \{\text{Target 1 present and Target 2 absent}\}$.

$H_3 = \{\text{Both targets present}\}$.

We introduce the filtering posterior probability matrix, $\mathbf{F}_{n|n}$, whose (k, j) element is the conditional probability that target 1 is at pixel k and target 2 is at pixel j , conditioned on the observation path \mathbf{Y}_0^n , i.e.,

$$F_{n|n}(k, j) = P(z_n^1 = k, z_n^2 = j \mid \mathbf{Y}_0^n) \quad k \in \tilde{\mathcal{L}}_1, j \in \tilde{\mathcal{L}}_2. \quad (32)$$

The posterior probabilities of the different hypothesis are computed as follows:

$$\begin{aligned} P(H_0 \mid \mathbf{Y}_0^n) &= F_{n|n}(L + l_i^1 + 1, L + l_i^2 + 1) \\ P(H_1 \mid \mathbf{Y}_0^n) &= \sum_{j=-l_s^2+1}^{L+l_i^2} F_{n|n}(L + l_i^1 + 1, j) \\ P(H_2 \mid \mathbf{Y}_0^n) &= \sum_{k=-l_s^1+1}^{L+l_i^1} F_{n|n}(k, L + l_i^2 + 1) \\ P(H_3 \mid \mathbf{Y}_0^n) &= \sum_{k=-l_s^1+1}^{L+l_i^1} \sum_{j=-l_s^2+1}^{L+l_i^2} F_{n|n}(k, j). \end{aligned} \quad (33)$$

The posterior probability of the two targets being present can be alternatively calculated as

$$P(H_3 \mid \mathbf{Y}_0^n) = 1 - \sum_{r=0}^2 P(H_r \mid \mathbf{Y}_0^n). \quad (34)$$

C. Tracking: maximum a posteriori (MAP) tracker

We examine next the solution to the tracking (localization) problem. We use a *maximum a posteriori* (MAP) strategy that gives optimal localization in a Bayesian sense, with respect to a cost function that assigns uniform penalty to any tracking error regardless of the magnitude of the error [32].

If, after detection, hypothesis H_m , $1 \leq m \leq 2^M - 1$, is declared true, we introduce the conditional probability tensor $\mathbf{\Pi}_{n|n}^m$ defined as

$$\mathbf{\Pi}_{n|n}^m(\mathbf{Z}_n) = P(\mathbf{Z}_n \mid H_m, \mathbf{Y}_0^n) = \frac{P(\mathbf{Z}_n, H_m \mid \mathbf{Y}_0^n)}{P(H_m \mid \mathbf{Y}_0^n)}. \quad (35)$$

The MAP Bayes tracker looks for the maximum of $\mathbf{\Pi}_{n|n}^m$ over \mathbf{Z}_n to estimate the positions of the targets that are assumed present under hypothesis H_m .

Example 1: Single Target

In the single target case, the tensor $\mathbf{\Pi}_{n|n}^1$ reduces to a vector whose general element is

$$\begin{aligned} \mathbf{\Pi}_{n|n}^1(k) &= P(\mathbf{z}_n = k \mid \text{target is present}, \mathbf{Y}_0^n) \\ &= \frac{f_{n|n}(k)}{1 - f_{n|n}(L + l_i^1 + 1)}. \end{aligned} \quad (36)$$

When the target is present, the maximum a posteriori (MAP) estimates of the actual target position are

$$\widehat{z_{n|n}^1} = \arg \max_{-l_s^1+1 \leq k \leq L+l_i^1} \mathbf{\Pi}_{n|n}^1(k). \quad (37)$$

Example 2: Two Targets

In the case of two targets, the conditional probability tensors $\mathbf{\Pi}_{n|n}^k$, $k = 1, 2, 3$ are matrices. Let $\mathbf{F}_{n|n}$ be the filtering posterior probability matrix defined in (32) and let H_1 , H_2 and H_3 be the three possible ‘‘target present’’ hypotheses as described before. We have three cases:

1. Target 1 is declared absent and target 2 is declared present: In this case, we find the optimal MAP estimate of the centroid position of the class 2 target, denoted by $\widehat{z_{n|n}^2}$, using the expression

$$\widehat{z_{n|n}^2} = \arg \max_{j \in \overline{\mathcal{L}}_2} \frac{F_{n|n}(L + l_i^1 + 1, j)}{P(H_1 | \mathbf{Y}_0^n)}. \quad (38)$$

2. Target 1 is declared present and target 2 is declared absent: This situation is the dual of the previous case. The optimal MAP estimate of the centroid position of the class 1 target, denoted by $\widehat{z_{n|n}^1}$, is given by

$$\widehat{z_{n|n}^1} = \arg \max_{k \in \overline{\mathcal{L}}_1} \frac{F_{n|n}(k, L + l_i^2 + 1)}{P(H_2 | \mathbf{Y}_0^n)}. \quad (39)$$

3. Targets 1 and 2 are declared present: when both targets are declared present, the optimal MAP centroid estimates are

$$(\widehat{z_{n|n}^1}, \widehat{z_{n|n}^2}) = \arg \max_{k \in \overline{\mathcal{L}}_1, j \in \overline{\mathcal{L}}_2} \frac{F_{n|n}(k, j)}{P(H_3 | \mathbf{Y}_0^n)} \quad k \in \overline{\mathcal{L}}_1, j \in \overline{\mathcal{L}}_2. \quad (40)$$

D. Detection/Tracking of targets with random signature

In the previous subsections, we considered the situation where the targets' signatures are deterministic and known. We now extend the algorithm to account for targets with random pixel intensity.

For simplicity, assume that the targets have equal size, i.e., $l_i^p = l_i$ and $l_s^p = l_s$, $1 \leq p \leq M$, where M is the number of targets. Let $\bar{l} = l_i + l_s + 1$ and define the \bar{l} -dimensional column vector of the signature parameters \mathbf{a}_n^p such that its i th component is

$$a_n^p(i) = a_{i,n}^p \quad -l_i \leq i \leq l_s \quad 1 \leq p \leq M. \quad (41)$$

Now stack these signature parameters in the column vector

$$\Theta_n = [(\mathbf{a}_n^1)^T \dots (\mathbf{a}_n^M)^T]^T \quad (42)$$

Assume that the sequence $\{\Theta_n\}$ is i.i.d. and independent of $\{\mathbf{Z}_n\}$ and $\{\mathbf{v}_n\}$, for $n \geq 1$. After a few algebraic steps, it is easy to show that

$$P(\mathbf{Z}_n | \mathbf{Y}_0^n) = C_n \left[\int p(\mathbf{y}_n | \Theta_n, \mathbf{Z}_n) p(\Theta_n) d\Theta_n \right] P(\mathbf{Z}_n | \mathbf{Y}_0^{n-1}). \quad (43)$$

Equation (43) shows that, under the assumptions that the sequence of target signatures $\{\Theta_n\}$ is i.i.d. and statistically independent of both the sequence of target positions and the sequence of clutter frames, we can obtain the observations kernel for each possible state vector \mathbf{Z}_n at instant n by averaging the conditional pdf of the measurements $p(\mathbf{y}_n | \Theta_n, \mathbf{Z}_n)$ over all possible realizations of the vector of target signatures.

E. Comparison with dynamic programming approaches

In this subsection, we contrast the nonlinear stochastic filtering approach to target tracking with previous work by Barniv [3]. We contrast Barniv's paper [3] with ours with respect to two issues: (i) the Viterbi algorithm used in [3] versus Bayes' law as used by us; (ii) setup of the problem and other modeling assumptions. We also make some brief comments on computational complexity.

(i) Viterbi algorithm versus Bayes' law For simplicity of notation and to follow Barniv's model, we consider a single target scenario and assume initially that the target is present in the surveillance space in all observed sensor frames, so that no hard detection decision between presence or absence of target has to be made at each sensor scan. Reference [3] applies Bellman's dynamic programming [12] and its

implementation by Viterbi [10] and Larson [11] to solve the target trajectory estimation problem. Let z_k and \mathbf{y}_k denote respectively the unknown target state and the observations at instant k . Define the target state path, \mathbf{Z}_0^k , and the observation path, \mathbf{Y}_0^k , such that

$$\mathbf{Z}_0^k = [z_0, z_1, \dots, z_k] \quad (44)$$

$$\mathbf{Y}_0^k = [\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_k] . \quad (45)$$

Denote by $\hat{z}_{k+1|k+1}$ the estimate of the unknown target state z_{k+1} at instant $k+1$, based on the observations \mathbf{Y}_0^{k+1} . Barring some minor differences in the indexing of the state and observation paths, Barniv's estimate of the target state is given by

$$\hat{z}_{k+1|k+1} = \arg \max_{z_{k+1}} I(z_{k+1}) \quad (46)$$

where $I(z_k)$ is a quantity that is proportional to

$$\max_{\mathbf{Z}_0^{k-1}} P(\mathbf{Z}_0^k | \mathbf{Y}_0^k) \quad (47)$$

and is computed using the recursion

$$I(z_{k+1}) = p(\mathbf{y}_{k+1} | z_{k+1}) \max_{z_k} [P(z_{k+1} | z_k) I(z_k)] \quad k \geq 0 . \quad (48)$$

When the observation \mathbf{y}_0 at instant zero is available, equation (48) is initialized with

$$I(z_0) = p(\mathbf{y}_0 | \mathbf{z}_0) P(z_0) . \quad (49)$$

Barring some minor differences in the initialization of the algorithm due to the availability of the observation \mathbf{y}_0 , equation (48) corresponds essentially to the forward recursion step of the Viterbi algorithm, see [10]. By contrast, our tracking algorithm is an MAP estimator based on Bayes' law, i.e., our estimate for the unknown target state at instant $k+1$ is

$$\hat{z}_{k+1|k+1} = \arg \max_{z_{k+1}} P(z_{k+1} | \mathbf{Y}_0^{k+1}) . \quad (50)$$

Combining the prediction and filtering steps of our algorithm in one equation, the posterior probability mass function on the righthand side of equation (50) is obtained by the recursion

$$P(z_{k+1} | \mathbf{Y}_0^{k+1}) = C_{k+1} p(\mathbf{y}_{k+1} | z_{k+1}) \left[\sum_{z_k} P(z_{k+1} | z_k) P(z_k | \mathbf{Y}_0^k) \right] \quad k \geq 0 \quad (51)$$

where C_{k+1} is a normalization constant that is independent of z_{k+1} . We initialize (51) with

$$P(z_0 | \mathbf{y}_0) = C_0 p(\mathbf{y}_0 | z_0) P(z_0) \quad (52)$$

where C_0 is a normalization constant that is independent of z_0 .

Equations (48) and (51) clearly define two different recursive algorithms. We now show that equations (46) and (50) correspond to two different maximization problems and may lead to different state estimates. Write

$$\begin{aligned} P(\mathbf{Z}_0^{k+1} | \mathbf{Y}_0^{k+1}) &= P(\mathbf{Z}_0^k, z_{k+1} | \mathbf{Y}_0^{k+1}) \\ &= P(\mathbf{Z}_0^k | z_{k+1}, \mathbf{Y}_0^{k+1}) P(z_{k+1} | \mathbf{Y}_0^{k+1}) . \end{aligned} \quad (53)$$

The second factor in equation (53) is the conditional pmf of the current state z_{k+1} given the path of observations \mathbf{Y}_0^{k+1} up to instant $k+1$. This is what our proposed nonlinear stochastic filter computes at each instant. The first factor can be simplified to

$$P(\mathbf{Z}_0^k | z_{k+1}, \mathbf{Y}_0^{k+1}) = P(\mathbf{Z}_0^k | z_{k+1}, \mathbf{Y}_0^k) . \quad (54)$$

Recall that Barniv's state estimate is given by

$$\begin{aligned} \hat{z}_{k+1|k+1} &= \arg \max_{z_{k+1}} \left\{ \max_{\mathbf{Z}_0^k} P(\mathbf{Z}_0^{k+1} | \mathbf{Y}_0^{k+1}) \right\} \\ &= \arg \max_{z_{k+1}} \left\{ \max_{\mathbf{Z}_0^k} P(\mathbf{Z}_0^k | z_{k+1}, \mathbf{Y}_0^k) P(z_{k+1} | \mathbf{Y}_0^{k+1}) \right\} . \end{aligned} \quad (55)$$

Note now that

$$\max_{z_{k+1}} \left\{ \max_{\mathbf{Z}_0^k} P(\mathbf{Z}_0^k | z_{k+1}, \mathbf{Y}_0^k) P(z_{k+1} | \mathbf{Y}_0^{k+1}) \right\} \neq \left[\max_{\mathbf{Z}_0^k} P(\mathbf{Z}_0^k | \mathbf{Y}_0^k) \right] \max_{z_{k+1}} P(z_{k+1} | \mathbf{Y}_0^{k+1}). \quad (56)$$

If the factorization in (56) were possible, then Barniv's estimate $\widehat{z}_{k+1|k+1}$ and ours would coincide. However, because the maximization on the lefthand side of (56) does not factor as the expression on the righthand side of the same equation, the two estimates may be different. Also note that we provide in our paper only the *filtering* estimate for the unknown state path, i.e., our algorithm computes the sequence

$$\widehat{z}_{k|k} \quad \text{for } k \geq 0.$$

Reference [3] on the other hand provides the *smoothed* state path estimate, i.e., the sequence

$$\widehat{z}_{i|k} \quad \text{for } k \geq 0, i \leq k.$$

The smoothed estimates in [3] are obtained using the backward retrieval step of the Viterbi algorithm, see [10]. In terms of applications, Barniv's algorithm provides a *batch* estimate of the state path,

$$\widehat{\mathbf{Z}}_0^k = \arg \max_{\mathbf{Z}_0^k} P(\mathbf{Z}_0^k | \mathbf{Y}_0^k),$$

whereas ours is an *on-line* algorithm that is similar in nature to Kalman-Bucy filtering, i.e., whenever a new state estimate is available at instant k , we do not go back and reestimate the previous states z_i for $i \leq k$.

Finally, in a multitarget scenario where targets are not assumed a priori to be always present, a multitarget detection step must be added to the tracking algorithm. In Barniv's work, the Viterbi forward recursion is run as if only one single target were present and multitarget detection is done simply by thresholding the function $I(z_k)$ at the last stage of the recursion. All states z_k for which $I(z_k)$ exceeds a certain threshold are assumed to be the final state of one possible target. The state trajectories for each detected target are then retrieved by moving backwards along the path of corresponding surviving nodes in the Viterbi trellis, see [10] and [3] for details. Since this procedure leads to a large number of false detections (roughly 40 detections per target [3]), a post-processing clustering step is used to merge nearby estimated trajectories. In our approach, we expand the state space to include dummy "absent target" states and propagate the joint posterior probability mass function of all target states, including the dummy states. Multitarget detection is then accomplished using a minimum probability of error M-ary Bayes hypotheses test.

(ii) Setup of problem and modeling assumptions In the sequel, we contrast briefly the state and observation models used in our work with the models introduced in [3]. In our paper, for targets that are present, the corresponding states are the pixel locations at each sensor frame of the target centroids in the discrete centroid grid. In [3], the states are defined as straight line trajectory segments across a group of $G > 1$ sensor frames that define the stages (instants) for the Viterbi forward recursion. The corresponding observation (measurements) model in [3] involves a differential pre-processing of the original sensor images. After pre-processing, it is assumed in [3] that all residual measurement noise is Gaussian and white. In our work, the measurements are the raw sensor frames themselves, with no pre-processing except for a possible removal of the moving local mean (as explained in section V-B). Instead of using a white Gaussian measurement noise assumption, we take full advantage of the real statistics of the background clutter

to improve detection/tracking performance. That includes exploring both the clutter spatial (intraframe) correlation and the clutter's possibly non-Gaussian amplitude statistics.

(iii) Computational complexity We compare the computational complexity of the Viterbi algorithm to our proposed Bayesian tracker. Let $z_n \in \tilde{\mathcal{L}}$ be the hidden variable, with the number of elements in $\tilde{\mathcal{L}}$ being denoted by L_1 . Define the L_1 -dimensional vectors $\mathbf{f}_{k|k}$ and \mathbf{i}_k , such that $f_{k|k}(j) = P(z_k = j \mid \mathbf{Y}_0^k)$, and $i_k(j) = I(z_k = j)$, $j \in \tilde{\mathcal{L}}$, where $I(z_k)$ is the function defined in (47). Introduce also the $L_1 \times L_1$ transition probability matrix, \mathbf{P}_T , such that $P_T(n, j) = P(z_{k+1} = n, \mid z_k = j)$, $(k, j) \in \tilde{\mathcal{L}} \times \tilde{\mathcal{L}}$. The recursion in equation (51), that corresponds to the Bayes tracker, can be rewritten in matrix notation as

$$\mathbf{f}_{k+1|k+1} = C_{k+1} \mathbf{S}_{k+1} \odot [\mathbf{P}_T \mathbf{f}_{k|k}] \quad (57)$$

where \odot denotes the pointwise multiplication operator and \mathbf{S}_{k+1} is a $L_1 \times 1$ vector such that $S_{k+1}(j) = p(\mathbf{y}_{k+1} \mid z_{k+1} = j)$, $j \in \tilde{\mathcal{L}}$. On the other hand, Viterbi's forward recursion in equation (48) is written as

$$\mathbf{i}_{k+1} = \mathbf{S}_{k+1} \odot [\max \{ \mathbf{P}_T^l \odot \mathbf{i}_k \}]_{1 \leq l \leq L_1} \quad (58)$$

where \mathbf{P}_T^l is the l th row of the transition matrix \mathbf{P}_T , i.e., $P_T^l(j) = P(z_{k+1} = l \mid z_k = j)$, $j \in \tilde{\mathcal{L}}$. The bracketed expression on the righthand side of equation (58) reads as follows: for each l , $1 \leq l \leq L_1$, do the pointwise multiplication of the l th row of the transition probability matrix, \mathbf{P}_T , by the previous filtering vector, \mathbf{i}_k , resulting in a L_1 -dimensional vector, \mathbf{i}_k^l . Then, look for the maximum of the entries $i_k^l(j)$ over the range $1 \leq j \leq L_1$ and assign this maximum to the l th entry of the bracketed vector.

A comparison between equations (57) and (58) shows that the two recursions differ basically in the computation of the bracketed vector on the righthand side. The Bayesian tracker involves the multiplication of an $L_1 \times L_1$ matrix by a $L_1 \times 1$ vector, which requires L_1^2 floating point multiplications and $L_1(L_1 - 1)$ floating point additions. On the other hand, the forward recursion of the Viterbi algorithm requires L_1^2 floating point multiplications and L_1 global maximum searches over an L_1 -dimensional vector. Those maximum searches require in turn $L_1(L_1 - 1)$ comparisons. The two algorithms therefore trade arithmetic (addition) computational complexity for logic (comparison) computational complexity. We make two additional remarks:

Remark 1: The Viterbi smoother requires that, in addition to the forward propagation of $I(z_k)$ using recursion (58), we must also store the indices of the maxima over j , $1 \leq j \leq L_1$, of $i_k^l(j)$, for all $k > 0$ and all l , $1 \leq l \leq L_1$. This table of stored indices is necessary for the implementation of Viterbi's backward retrieval step, see [10].

Remark 2: In most applications, the transition probability matrix, \mathbf{P}_T , is not a full $L_1 \times L_1$ matrix, as transitions are only allowed between adjacent target states. As a result of the sparse nature of \mathbf{P}_T , the number of floating point multiplications required in the prediction step for both the Bayes tracker and the Viterbi recursion falls in practice from $O(L_1^2)$ to $O(\alpha L_1)$, where $\alpha \ll L_1$. The required number of floating point sums in the prediction step of the Bayes tracker and the complexity of the maxima searches in the Viterbi recursion are also reduced accordingly.

F. Flowchart summary of the Bayes detector/tracker

We present in Table I a flowchart summary of the proposed optimal Bayes detector/tracker.

- 1) Initialize $P(\mathbf{Z}_n | \mathbf{Y}_0^{n-1})$ with the given prior $P(\mathbf{Z}_0)$.
 - 2) Compute $P(\mathbf{Z}_0 | \mathbf{Y}_0)$ using equation (26) with $n = 0$.
 - 3) For $n = 1$ up to the total number of available frames:
 - Compute $P(\mathbf{Z}_n | \mathbf{Y}_0^{n-1})$ using equation (23).
 - Compute $P(\mathbf{Z}_n | \mathbf{Y}_0^n)$ using equation (26).
 - Do M-ary detection using the hypotheses test (27).
 - If hypothesis H_m is declared true, compute $\Pi_{n|n}^m(\mathbf{Z}_n)$ using (35) and look for its maximum over \mathbf{Z}_n to estimate the centroid positions of detected targets.
- End of for-loop.

TABLE I

FLOWCHART SUMMARY OF THE MULTITARGET, MULTIFRAME BAYES DETECTOR/TRACKER.

G. Illustrative Example: Two Extended Targets in Gaussian Noise

Finally, we close this section with an illustrative example of application of the optimal Bayes detector/tracker in a multitarget scenario with overlapping targets. We track/detect two extended targets in a 1D finite grid, against a white Gaussian background clutter with covariance matrix $\sigma_w^2 \mathbf{I}$. Both class 1 and class 2 targets extend over 9 resolution cells with $l_i = r_i = r_s = l_s = 4$, but have different (deterministic) signatures. Class 1 targets have a rectangular-shaped signature, whereas class 2 targets have a triangular-shaped signature.

In any given sensor frame, either two targets (one from class 1, the other from class 2) are present, or just one target (either class 1 or 2) is present, or no target is present. When two targets are present, the corresponding sensor returns may be apart from each other, as shown in figure 2. Otherwise, they may overlap in the sensor, causing their signatures to be added in the sensor image, as shown in figure 3.




Fig. 2. (a) Noise-free sensor scan with two targets. (b) Observed (noisy) sensor scan, PSNR=3 dB

The targets have translational motion with the position of the targets centroids in the 1D grid described by known first order discrete Markov chains with deterministic drifts $d_1 = 2$ and $d_2 = 3$ for class 1 and class 2 targets, respectively. Once a target belonging to a given class disappears from the sensor range, another target of the same class may appear randomly at any resolution cell with a probability $p_a = 0.3$. The simulation was conducted for a total of 100 frames, with 100 resolution cells per frame. A target that is estimated to be absent is indicated by a ‘+’ mark on the horizontal axis, while a true absence of target




Fig. 3. (a) Noise-free sensor scan with superimposed targets. (b) Observed sensor scan, PSNR=3 dB

is indicated by a ‘o’ mark on the same axis.

Figures 4 (a) and (b) show the tracked positions of the centroids of class 1 and class 2 targets respectively between frames 20 and 70, with peak signal-to-noise ratio ($\text{PSNR} = 10 \log_{10}(1/\sigma_w^2)$) equal to 10 dB. Notice that, between frames 65 and 70, class 1 targets are absent from the sensor view, which is correctly indicated by the detector/tracker as a series of superimposed ‘o’ and ‘+’ marks on the horizontal axis in figure 4 (a). During the same time period, a class 2 target is present and is correctly tracked as indicated in the bottom right corner of figure 4 (b). The opposite situation occurs between frames 36 and 42 when the class 2 target is absent and is correctly declared not present by the detector/tracker (as seen in figure 4 (b)), but the class 1 target is present and accurately tracked as indicated in figure 4 (a). Finally, notice that between frames 43 and 54, not only are both targets present, but also their images overlap since the difference between their centroid positions is less than $2l_s + 1$. Despite the superposition, figures 4 (a) and (b) show that the centroids of the two targets are accurately tracked independently. The algorithm is capable of performing data association with a high degree of accuracy at the same time that it is able to reject false alarms and prevent misses.

IV. DETECTION PERFORMANCE

We study the detection performance of the optimal nonlinear detector/tracker assuming a single target scenario. The performance is evaluated through the receiver operating characteristic (ROC) curves for the Neyman-Pearson detector, obtained by varying the threshold in (28). The experimental ROC’s presented in this section are generated using *Monte Carlo* simulations.

A. Correlated Gaussian targets in correlated Gaussian clutter

We consider first the case when the background clutter is Gaussian and correlated. We simulate 1D Gaussian targets with dimensions $l_i = l_s = 4$ moving in a grid of size $L = 100$. The synthetic targets are samples of a correlated first-order GMrf model with mean $m_a = 1$ and covariance parameters α_a and σ_a . The target image is cluttered by a first order GMrf clutter with parameters α_c and σ_c . The synthetic targets move in a 1D range resolution grid with a mean drift of 2 cells/scan. There is a fluctuation probability of one cell around the mean displacement equal to 0.4. As one target disappears from the

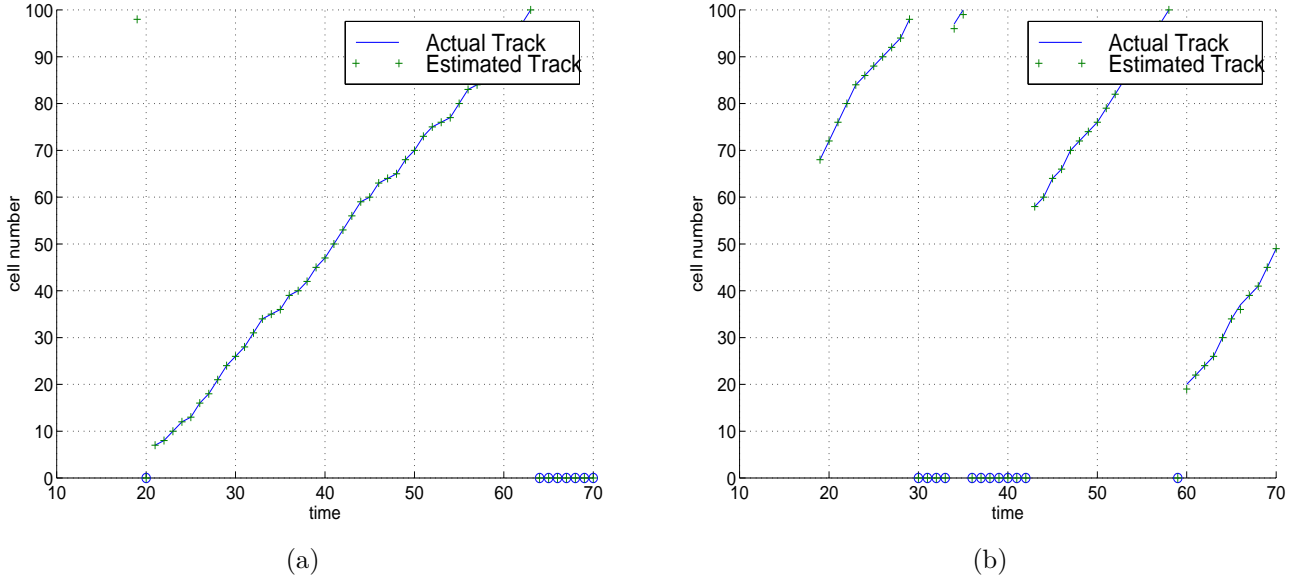


Fig. 4. (a) Centroid tracking for class 1 targets. (b) Centroid tracking for class 2 targets, PSNR = 10 dB

surveillance space, there is a 20 % probability of a new target reappearing at any arbitrary position in the grid $\mathcal{L} = \{l:1 \leq l \leq L\}$, i.e., $p_a = 0.2$. This assumption corresponds to a worst case scenario when new tracks can be initialized with uniform probability at any cell in the sensor grid.

Figure 5 (a) shows the experimental receiving operating characteristic curves (ROCs) for the multi-frame Bayes detector, with $\sigma_a = 0.2$, $\alpha_a = 0.16$, and $\alpha_c = 0.24$, for two levels of the average SNR = $20 \log_{10}(m_a/\sigma_c)$, respectively 3 and 0 dB (i.e., $\sigma_c = 0.7$ and $\sigma_c = 1$). Figure 5(b) shows the ROCs when $\sigma_a = 0.4$ (i.e, increasing the variance of the target pixels). The experimental curves were estimated from a total of 8,000 Monte Carlo runs. The plots in Figure 5 indicate good detection performance, even in the




Fig. 5. Performance of the multiframe Bayes detector with correlated Gaussian targets in correlated GMrf clutter: (a) $\sigma_{\text{targ}} = 0.2$, (b) $\sigma_{\text{targ}} = 0.4$

adverse conditions of heavy clutter. For example, the algorithm reaches a 90 % probability of detection

for false alarm rate of 10^{-3} . Figure 5(b) shows that there is a slight deterioration in performance when we increase the target variance.

B. Non-Gaussian clutter

To evaluate the detection performance for non-Gaussian clutter, we ran Monte Carlo simulations with a succession of single targets moving in uncorrelated complex clutter with K and Weibull envelope statistics. For simplicity, the successive targets are pointwise with a unit signature in the in-phase component. We use as figure of merit the peak signal-to-noise ratio defined as $\text{PSNR} = 10 \log_{10}(1/\sigma_w^2)$, where σ_w^2 is the variance of the white noise sequence from which the corresponding non-Gaussian SIRV is simulated, see reference [16] for details. The motion parameters are the same as in the Gaussian simulation, except that, in the case of complex clutter, a speed of d resolutions cells/scan corresponds to a speed of $2d$ in the double-sized sensor image. The probabilities of detection and false alarm for each value of the threshold are obtained from statistics collected from 10,000 sensor frames, where each frame corresponds to 64 resolution cells or 128 complex quadrature returns.

Figure 6 (a) shows the superposition of the optimal Bayes ROCs with Rayleigh, Weibull, and K clutters, respectively from top to bottom, for $\text{PSNR} = 6$ dB. The simulated K clutter has shape parameter $\nu = 1.5$. The parameters for the simulated Weibull clutter are $c = 1.5$ and $a = 0.6777$. The plot shows that, even when the correct clutter statistics are incorporated into the detector's structure, there is still a slight deterioration in performance under non-Rayleigh clutter, probably due to the heavier tails of the K and Weibull statistics, which increase the likelihood of false alarms. Figure 6 (b) shows the superposition of the ROC under Rayleigh clutter for $\text{PSNR} = 4$ dB and the ROC under K clutter, $\text{PSNR} = 6$ dB. The two curves in Figure 6 (b) roughly coincide, thus suggesting an approximate 2 dB gain in PSNR under Rayleigh clutter compared to the spiky K clutter.

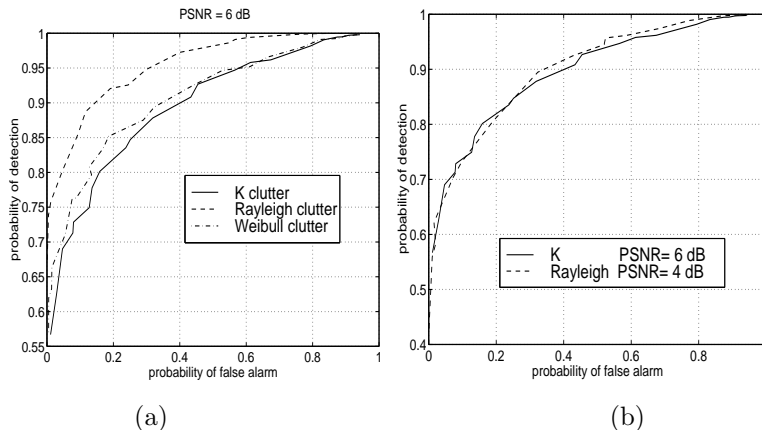


Fig. 6. Comparison of optimal Bayes detector performance (a) under Rayleigh, Weibull ($c = 1.5$, $a = 0.6777$), and K ($\nu = 1.5$) clutter, $\text{PSNR} = 6$ dB; (b) under K ($\nu = 1.5$) clutter ($\text{PSNR} = 6$ dB) and Rayleigh clutter ($\text{PSNR} = 4$ dB)

V. 2D DETECTION/TRACKING PERFORMANCE

We investigate in this section the performance of the proposed Bayes detector/tracker with two-dimensional (2D) extended targets moving in digital images corrupted by heavy clutter. We carry out two sets of experiments: one with 2D synthetic data and the other with a real data clutter intensity image recorded by an airborne infrared laser radar [17]. In the second set of experiments, a simulated target (military vehicle) template is inserted into the real data background.

A. 2D simulations with synthetic clutter

Clutter: The background clutter is a 2D noncausal first order GMrf. The clutter intensity at pixel (i, j) during the n th scan, $v_n(i, j)$ is modeled by its minimum mean square error (MMSE) representation [30]

$$v_n(i, j) = \beta_h [v_n(i, j-1) + v_n(i, j+1)] + \beta_v [v_n(i-1, j) + v_n(i+1, j)] + u_n(i, j) \quad (59)$$

where $u_n(i, j)$ is the driving noise term. We collect the clutter samples, $v_n(i, j)$, and the error samples $u_n(i, j)$, in two long row-lexicographed vectors, respectively \mathbf{v}_n and \mathbf{u}_n , thus obtaining an equivalent matrix representation to the difference equation (59). The matrix representation is analogous to the 1D case, i.e.,

$$\mathbf{A}\mathbf{v}_n = \mathbf{u}_n \quad (60)$$

where \mathbf{A} is the potential matrix of the 2D GMrf. Using the orthogonality between the field $\{v_n(i, j)\}$ and the driving noise, $\{u_n(i, j)\}$, the clutter covariance matrix is proportional [25] to the inverse of the potential matrix \mathbf{A} .

The parametric structure of the potential matrix extends naturally to the 2D case [23], [25]. The corresponding inverse of the clutter covariance for the 2D clutter background is a block-Toeplitz, block-banded matrix where each of the individual blocks is itself Toeplitz and banded [25]. A comprehensive study of the eigenstructure of perturbed Toeplitz and block-Toeplitz matrices and their relation to 1D and 2D GMrf models of *arbitrary order* with different choices of boundary conditions is found in [23].

Target and Observations: For simplicity, we limit our discussion in this section to a single target scenario. The target is a rigid body with a 2D translational motion. We assume that, at any given frame, the target's clutter-free image is contained inside a 2D rectangular region of size $(r_i + r_s + 1) \times (l_i + l_s + 1)$. In this notation, r_i and r_s denote the maximum vertical pixel distances in the target image when we move away, respectively up and down, from the target centroid. Analogously, l_i and l_s denote the maximum horizontal pixel distances in the target image when we move away, respectively left and right, from the target centroid. Let I be the 2D finite lattice $I = \{(k, l): -r_i \leq k \leq r_s, -l_i \leq l \leq l_s\}$. The 2D target signature at frame n is given by the the signature coefficients

$$a^n(k, l) = c^n(k, l)\phi^n(k, l) \quad (k, l) \in I \quad (61)$$

where the term $\phi^n(k, l) \in \mathfrak{R}$ specifies the target's *pixel intensity* whereas $c^n(k, l) \in \mathcal{B} = \{0, 1\}$ is a binary shape coefficient. In the Monte Carlo simulations presented in this section, we assume that the targets have a rectangular template and that their signatures are deterministic, time-invariant and known. Without loss of generality, we make the target pixel intensities constant and equal to 1. The targets are contained in a square region of size 9×9 , and are cluttered by a first order, highly correlated GMrf




Fig. 7. (a) Clutter-free target image, (b) Simulated sensor image, PSNR = 0 dB.

background with $\beta_h = \beta_v = 0.24$. Figures 7 (a) and (b) show respectively the clutter-free target image and a random sample of the target plus clutter image when the target is centered at pixel (100, 50) with peak-signal-to-noise ratio (PSNR) equal to 0 dB.

Motion We describe here the 2D translational motion model used in the simulations presented in the paper. The simulated targets have mean velocity of 2 resolution cells per frame in both the horizontal and vertical directions, and move in a grid of size $L \times L$ pixels. The actual displacement is a random walk fluctuation around the average displacement. In other words, if (i, j) is the predicted centroid position according to the deterministic velocity, the actual position lies in a 2D spatial region around (i, j) , with fluctuation probabilities r , s , t , and q shown in Figure 8. In our simulations, all fluctuation probabilities of one cell were set to 20 %. Like in subsection II-A, we expand the 2D sensor lattice to build a centroid




Fig. 8. 2D random walk fluctuation around average target drift

lattice that accounts for boundary effects. We define then an equivalent 1D representation of the 2D centroid lattice that is obtained by sequentially stacking the rows of the 2D lattice into a long vector, see reference [24]. The equivalent 1D centroid lattice is denoted as $\bar{\mathcal{L}}$. Finally, like in the 1D case, we build the augmented lattice $\tilde{\mathcal{L}}$ by adding to $\bar{\mathcal{L}}$ a dummy state that represents the absence of the target. The unknown state at instant n is a 1D random variable, z_n , defined on $\tilde{\mathcal{L}}$. We denote by \mathbf{P}_T the transition probability matrix that collects the transition probabilities $\{P(z_n = i | z_{n-1} = j)\}$, $(i, j) \in \tilde{\mathcal{L}} \times \tilde{\mathcal{L}}$.

Multiframe Bayes detector/tracker Let \mathbf{y}_n be the observed $L \times L$ sensor frame at instant n . Define the filtering posterior probability vector, $\mathbf{f}_{n|n}$ such that $f_{n|n}(l) = P(z_n = l | \mathbf{Y}_0^n)$, $l \in \tilde{\mathcal{L}}$, where $\tilde{\mathcal{L}}$ now denotes the 1D row lexicographed centroid lattice augmented by the dummy absent target state. Similarly, introduce the prediction posterior probability vector, $\mathbf{f}_{n|n-1}$. Finally, define the observations kernel vector

\mathbf{S}_n such that $S_n(l) = p(\mathbf{y}_n | z_n = l)$, $l \in \tilde{\mathcal{L}}$. For simplicity of notation, make $r_i = l_i$ and $r_s = l_s$ and introduce $L_1 = L + l_i + l_s$. Let \mathbf{I}_p denote the $p \times p$ identity matrix and \mathbf{H}_p be a $p \times p$ matrix such that $H(i, j) = 1$ for $|i - j| = 1$ or zero otherwise. We use the symbol \odot to denote pointwise multiplication and the symbol \otimes to denote the Kronecker or tensor product [27]. The symbol vec denotes the operator that converts an $P \times Q$ into a PQ -dimensional column vector by sequentially stacking all the rows of the matrix. In the particular case of a single, deterministic target with known and time-invariant signature coefficients $a_{k,l}$, and a first-order 2D GMrf clutter model as described by equation (60), the optimal 2D Bayes detector/tracker with a total number of frames equal to N , is implemented by the pseudocode in Table II (for a detailed derivation, see [24]).

<p>a) Initialization.</p> <p>b) For $n = 1$ to N</p> <ul style="list-style-type: none"> • Prediction step: $\mathbf{f}_{n n-1} = \mathbf{P}_T \mathbf{f}_{n-1 n-1}$. • for $i, j = 1$ to L <ul style="list-style-type: none"> Differential operator: $\mu_{i,j} = y_n(i, j) - \beta_h [y_n(i, j+1) + y_n(i, j-1)] - \beta_v [y_n(i+1, j) + y_n(i-1, j)]$, with $y_n(i, j) = 0$, for $i, j < 1$ or $i, j > L$. end of loop • for $I, J = 1$ to L_1 <ul style="list-style-type: none"> $i = I - l_s$; $j = J - l_s$; Matched filter: $\lambda_{i,j} = \sum_k \sum_l a_{k,l} \mu_{i+k, j+l}$, with the limits for the summations given in Table III. $\mathbf{a}_{i,j} = \{a_{k,l}\}$, with k and l in the ranges assigned to each pair (i, j) in Table III. $\mathbf{A}_t = \mathbf{I}_{\bar{T}} \otimes (\mathbf{I}_{\bar{T}} - \beta_h \mathbf{H}_{\bar{T}}) - \beta_v \mathbf{H}_{\bar{T}} \otimes \mathbf{I}_{\bar{T}}$ where $(\bar{T}, \bar{l}) = \text{size}(\mathbf{a}_{i,j})$. Energy term: $\rho_{i,j} = (\text{vec}[\mathbf{a}_{i,j}])^T \mathbf{A}_t (\text{vec}[\mathbf{a}_{i,j}])$. Observations kernel: $S_n((I-1)L_1 + J) = \exp \left[(2\lambda_{I-l_s, J-l_s} - \rho_{I-l_s, J-l_s}) / 2\sigma_u^2 \right]$. end of loop. • Normalized kernel entry for absent target state: $S_n(L_1^2 + 1) = 1$. • Filtering step: $\mathbf{f}_{n n} = C_n \mathbf{S}_n \odot \mathbf{f}_{n n-1}$ where C_n is a normalization constant such that $\sum_l f_{n n}(l) = 1$. • Binary Detection: $f_{n n}(L_1^2 + 1) \underset{H_1}{\overset{H_0}{>}} 1 - f_{n n}(L_1^2 + 1)$. • MAP estimation: If hypothesis H_1 (target present) declared true, $\hat{z}_{n n} = \arg \max_{l \in \tilde{\mathcal{L}}} f_{n n}(l)$. • $\mathbf{f}_{n-1 n-1} = \mathbf{f}_{n n}$. End of outer for-loop <p>c) End of program.</p>

TABLE II

PSEUDOCODE FOR THE 2D BAYES DETECTOR/TRACKER

$\lambda(i, j)$	$-l_s + 1 \leq j \leq l_i$	$l_i + 1 \leq j \leq L - l_s$	$L - l_s + 1 \leq j \leq L + l_i$
$-l_s + 1 \leq i \leq l_i$	$\sum_{k=-i+1}^{l_s} \sum_{l=-j+1}^{l_s} (\cdot)$	$\sum_{k=-i+1}^{l_s} \sum_{l=-l_i}^{l_s} (\cdot)$	$\sum_{k=-i+1}^{l_s} \sum_{l=-l_i}^{L-j} (\cdot)$
$l_i + 1 \leq i \leq L - l_s$	$\sum_{k=-l_i}^{l_s} \sum_{l=-j+1}^{l_s} (\cdot)$	$\sum_{k=-l_i}^{l_s} \sum_{l=-l_i}^{l_s} (\cdot)$	$\sum_{k=-l_i}^{l_s} \sum_{l=-l_i}^{L-j} (\cdot)$
$L - l_s + 1 \leq i \leq L + l_i$	$\sum_{k=-l_i}^{L-i} \sum_{l=-j+1}^{l_s} (\cdot)$	$\sum_{k=-l_i}^{L-i} \sum_{l=-l_i}^{l_s} (\cdot)$	$\sum_{k=-l_i}^{L-i} \sum_{l=-l_i}^{L-j} (\cdot)$

TABLE III

COMPUTATION OF THE DATA TERM λ_{ij}

Remark The actual implementation of the matrix multiplication $\mathbf{P}_T \mathbf{f}_{n-1|n-1}$ in Table II explores the sparse and block-banded structure of the transition probability matrix \mathbf{P}_T . Note also that the energy term ρ is constant for the range $l_i + 1 \leq i, j \leq L - l_s$ and, therefore, can be computed off-line. In general, for an $L \times L$ sensor grid, it can be shown that, using the GMrf clutter model, the Markov chain motion model, and the small extended target models, we reduce total number of required floating point multiplications from $O(L^6)$ to $O(\alpha L^2)$ in the filtering step of the algorithm and, from $O(L^4)$ to $O(\gamma L^2)$ in the prediction step, where $\gamma \ll L$.

We now discuss two suboptimal trackers whose performance we later compare to the performance of our optimal Bayes tracker: the maximum likelihood tracker, and the linearized Kalman-Bucy tracker.

Maximum Likelihood Tracker The correlation or memoryless maximum likelihood (ML) tracker ignores the information on the dynamics of the target motion and makes tracking decisions at each sensor scan based solely on the present observed data. In the single target case, let z_n denote as before the position of the centroid of a target that is assumed present in the n th sensor scan. The memoryless ML estimate of the centroid position is given by

$$\hat{z}_{ML} = \arg \max_{i \in \bar{\mathcal{L}}} p(\mathbf{y}_n | z_n = i) \quad (62)$$

where $\bar{\mathcal{L}}$ is the equivalent 1D centroid lattice.

Linearized Kalman-Bucy Tracker Although the single frame, memoryless ML tracker is generally accurate in scenarios of high SNR, its performance deteriorates noticeably as the power level of the clutter increases. An improvement to the ML estimator in heavy clutter is obtained by further post-processing its output with a linearized Kalman-Bucy filter (KBf). This is a suboptimal multiframe tracking scheme that reintroduces information about the target dynamics into the tracking algorithm. It resembles common algorithms presented in the literature [1].

Remark: Computational Complexity We compare the computational cost of the Bayes detector/tracker to the association ML-KBf of the maximum likelihood tracker with the Kalman-Bucy filter. The ML step in equation (62), in the case of deterministic targets observed in correlated Gauss-Markov random clutter, reduces to maximizing over all the target centroid positions (i, j) the quantity $\lambda_{i,j} - \rho_{i,j}/2$, where $\lambda_{i,j}$ and $\rho_{i,j}$ are the data and energy terms described in Table II. For an $L \times L$ sensor grid with $L \gg l_i + l_s$, this cost is dominated by the computation of $\lambda_{i,j}$ that is an operation of order $O(\beta L^2)$ with $\beta \ll L$. The computational cost of the KBf stage is negligible in comparison. For a scalar dynamic model in both dimensions, the cost is only 4 floating point multiplications and 4 floating point additions. Hence, the overall ML-KBf association has cost $O(\beta L^2)$. On the other hand, as discussed before, the Bayes detector/tracker requires $O(\alpha L^2)$ floating point multiplications in the filtering step, and $O(\gamma L^2)$ floating point multiplications in the prediction step, with $\alpha \approx \beta \approx \gamma \ll L$. In summary, the Bayes detector/tracker costs roughly twice the cost of the ML-KBf association. The ML-KBf association saves the prediction step of the Bayes detector/tracker.

Tracking Performance in Correlated Clutter We study first the tracking performance of the Bayes algorithm using synthetic data. The simulated targets are 2D rectangular objects with constant pixel intensity shown in Figure 7(a). At each sensor scan, we assume that at most one target is present. The target

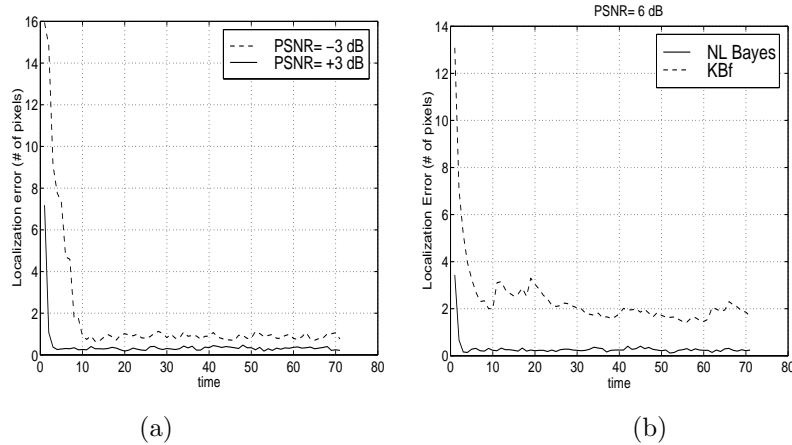


Fig. 9. (a) Performance of the nonlinear Bayes tracker in correlated GMrf clutter for PSNR= -3 and $+3$ dB; (b) Performance of the nonlinear Bayes tracker vs the linearized KBf, for PSNR= 6 dB

starts from an *unknown* random location in the 50×50 upper left corner of the image and is subsequently tracked over 70 consecutive sensor frames. Figure 9(a) shows the evolution over time of the standard deviation of the error in the centroid’s vertical position estimate given by the nonlinear Bayes tracker. The standard deviation is expressed in number of pixels and evaluated by averaging the errors over 150 Monte Carlo runs. We repeat the experiment for two values of PSNR, respectively $+3$ dB and -3 dB. The corresponding curves for the horizontal position estimate are qualitatively similar and are omitted for conciseness. Figure 9(a) shows that there is an initial localization error which declines over time as new measurements become available. The target acquisition time (i.e., the number of sensor scans for the error to reach its steady state) increases as the PSNR decreases. Likewise, the initial and steady state localization errors also increase with decreasing PSNR.

Next, we compare the nonlinear Bayes tracker with the alternative suboptimal association of the spatial matched filter and a linearized KBf. Figure 9(b) plots the standard deviation over time of the error in the vertical position estimate for both trackers, with PSNR equal to 6 dB. We see from the plot that the KBf tracker has higher initial and steady state position estimate errors and a longer target acquisition time. Figure 10 shows again the vertical position estimate error curves for both trackers. This time, the KBf curve is obtained for PSNR equal to 6 dB, while, for the nonlinear Bayes tracker, we lower the PSNR to -5 dB. The two curves in Figure 10 show that the steady state performance of the nonlinear Bayes tracker and the linearized KBf are very similar, despite an 11 dB difference in PSNR.

B. Detection/Tracking Example with Real Clutter

In order to have a qualitative assessment of the effect of model/data mismatch on the performance of the algorithm, we ran a small-scale simulation with real clutter data. We used real-world intensity imagery of a snow-covered field in Stockbridge, NY, obtained by a $0.85 \mu m$ down-looking laser radar [17] mounted to the bottom of a Gulfstream G-1 aircraft. The imagery is from the Infrared Airborne Radar (IRAR) collection at MIT Lincoln Laboratory and was obtained through the website of the Center for

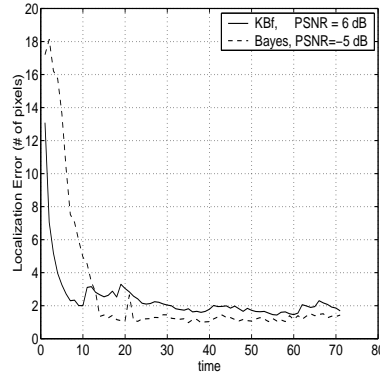


Fig. 10. Comparison of the tracking performance of the Kbf tracker with PSNR = 6 dB and the nonlinear Bayes tracker with PSNR = - 5 dB




Fig. 11. Simulated real target in heavily cluttered real background (IRAR imagery): (a) target plus clutter, (b) target template.

Imaging Science at Johns Hopkins University [33]. Figure 11(a) shows a 120×120 real clutter gray-level intensity image with a heavily cluttered military vehicle (tank) model superimposed on it. Figure 11(b) shows the tank template alone as a binary image with target intensity equal to 1 and background intensity equal to zero. The tank template (shape) was extracted from a real image of the vehicle taken at the same field with the same sensor. The target pixel intensity was set arbitrarily to achieve the desired low level of contrast between target and clutter. In order to assess tracking performance, we simulated a random trajectory for the target template and detected/tracked it over 27 frames using the Bayes algorithm. The target starts from an unknown location in the 120×120 image and moves in the real clutter background according to a 2D random walk model whose parameters are known to the tracker.

Since the clutter background is real data, we initially preprocess each frame in the image sequence. The pre-processing consists of the segmentation of the original images and the subsequent removal of the spatially-variant local mean in each subimage so that the pixel intensity histogram approaches a zero-mean Gaussian distribution [4]. We then adjust a first-order GMrf model to the Gaussianized data, estimating the corresponding parameters β_h , β_v , and σ_u for each frame. The 2D GMrf parameters were estimated

using a simplified version of the maximum likelihood estimation algorithm introduced in reference [26].

We compare the tracking results using (a) the proposed Bayes tracker, and (b) a standard 2D image correlator associated to a linearized Kalman-Bucy filter. The corresponding estimated target trajectories are shown in Figure 12. The Bayes tracker assumes a uniform initial target position distribution over the entire sensor grid. The linear filter, on the other hand, is initially favored by using a Gaussian initial position prior that is centered in the vicinity of the true initial position and has a small variance. The real simulated trajectory is shown in solid line. The position estimates generated by the Bayes tracker are indicated by the symbol '+', whereas the estimates generated by the linearized KBf are interpolated using dashed lines. In the first half of the trajectory shown in Figure 12, the simulated tank is going through a

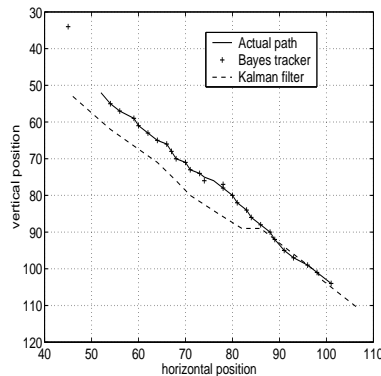


Fig. 12. Nonlinear Bayes detector/tracker versus linearized Kalman-Bucy filter: performance comparison

heavily cluttered section of the background, and the single frame standard image correlator is unable to track the target. The KBf tends to discard the correlator's position estimates and through the inertia in its prediction step, tries to fit a straight line trajectory. In the second half of the simulation, when the tank is on an open field, the image correlator is capable of correctly locating the target and the filtering step of the KBf slowly forces the estimated trajectory to approach the true trajectory. By contrast, the Bayes tracker, which has no prior knowledge of the initial position, makes a large initial localization error (the isolated '+' on the top left corner of Figure 12), but, afterwards, as new frames become available, the tracker immediately acquires the target and tracks it almost perfectly. A comparison shows that, even in steady state, the localization error for the Bayes tracker is lower than for the KBf, while the acquisition time is much shorter.

Remark: The assumption that the sensor frames are uncorrelated in time is unrealistic in practice. However, the good tracking results with real clutter presented in this section lead us to believe that the Bayes detector/tracker exhibits a high degree of robustness to interframe correlation.

VI. CONCLUSION

In this paper, we presented a new optimal recursive algorithm for integrated, multiframe Bayesian detection and tracking of multiple targets that move randomly in heavily cluttered environments. We considered both extended and pointwise targets with deterministic and random signatures. We developed

models for target signature and target motion that take into consideration the finite resolution of the sensors and used these models to build a joint framework for detection and tracking that underlies our proposed Bayesian detector/tracker.

Extensive Monte Carlo simulations in 1D and 2D surveillance spaces determined the performance of the optimal Bayes multiframe detector/tracker under both spatially correlated Gaussian clutter and non-Gaussian clutter with heavy tail (K and Weibull) statistics. The detection and tracking characteristics of the optimal Bayes algorithm is an upper bound to the performance of other suboptimal algorithms. In particular, when tracking stealthy targets (dim targets in heavy clutter) with the optimal Bayes tracker, the Monte Carlo simulations with 2D targets departing from unknown location show substantial steady-state tracking performance gains of up to 11 dB over alternative algorithms found in the literature such as spatial correlators (matched filters) and linearized Kalman-Bucy filters. Examples with real clutter data and a known inserted target also show better detection/tracking performance than the association of a conventional image correlator and a Kalman-Bucy tracker. These examples illustrate the robustness of the algorithm to possible mismatches between the data and the underlying models assumed by the algorithm.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for many useful suggestions and comments.

CONTENTS

I	Introduction	2
II	The Model	5
II-A	Surveillance Space and Target Model	5
II-B	Multitarget Observations and Clutter Models	6
II-C	Target Motion	9
III	Optimal Bayes Multitarget Detector/Tracker	10
III-A	Nonlinear Stochastic Filtering Approach	10
III-B	Minimum probability of error Bayes detector	12
III-C	Tracking: maximum a posteriori (MAP) tracker	13
III-D	Detection/Tracking of targets with random signature	14
III-E	Comparison with dynamic programming approaches	14
III-F	Flowchart summary of the Bayes detector/tracker	17
III-G	Illustrative Example: Two Extended Targets in Gaussian Noise	18
IV	Detection Performance	19
IV-A	Correlated Gaussian targets in correlated Gaussian clutter	19
IV-B	Non-Gaussian clutter	21
V	2D Detection/Tracking Performance	22
V-A	2D simulations with synthetic clutter	22
V-B	Detection/Tracking Example with Real Clutter	26
VI	Conclusion	28

LIST OF FIGURES

1	Example of Finite State Machine Diagram	9
2	(a) Noise-free sensor scan with two targets. (b) Observed (noisy) sensor scan, PSNR=3 dB	18
3	(a) Noise-free sensor scan with superimposed targets. (b) Observed sensor scan, PSNR=3 dB	19
4	(a) Centroid tracking for class 1 targets. (b) Centroid tracking for class 2 targets, PSNR = 10 dB	20
5	Performance of the multiframe Bayes detector with correlated Gaussian targets in correlated GMrf clutter: (a) $\sigma_{\text{targ}} = 0.2$, (b) $\sigma_{\text{targ}} = 0.4$	20
6	Comparison of optimal Bayes detector performance (a) under Rayleigh, Weibull ($c = 1.5$, $a = 0.6777$), and K ($\nu = 1.5$) clutter, PSNR= 6 dB; (b) under K ($\nu = 1.5$) clutter (PSNR = 6 dB) and Rayleigh clutter (PSNR = 4 dB)	21
7	(a) Clutter-free target image, (b) Simulated sensor image, PSNR = 0 dB.	23

8	2D random walk fluctuation around average target drift	23
9	(a) Performance of the nonlinear Bayes tracker in correlated GMrf clutter for PSNR= -3 and +3 dB; (b) Performance of the nonlinear Bayes tracker vs the linearized KBf, for PSNR= 6 dB	26
10	Comparison of the tracking performance of the KBf tracker with PSNR = 6 dB and the nonlinear Bayes tracker with PSNR = - 5 dB	27
11	Simulated real target in heavily cluttered real background (IRAR imagery): (a) target plus clutter, (b) target template.	27
12	Nonlinear Bayes detector/tracker versus linearized Kalman-Bucy filter: performance comparison	28

LIST OF TABLES

I	Flowchart summary of the multitarget, multiframe Bayes detector/tracker.	18
II	Pseudocode for the 2D Bayes detector/tracker	24
III	Computation of the data term λ_{ij}	24

REFERENCES

- [1] Y. Bar-Shalom and X. Li. *Multitarget-Multisensor Tracking: Principles and Techniques*. YBS, Storrs, CT, 1995.
- [2] R. E. Bethel and G. J. Paras. A PDF multitarget tracker. *IEEE Transactions on Aerospace and Electronic Systems* vol.30, n.2, April 1994, pp. 386-403.
- [3] Y. Barniv. Dynamic programming solution for detecting dim moving targets. *IEEE Transactions on Aerospace and Electronic Systems*, Vol 21, n.1, January 1985, pp. 144-156.
- [4] J. Y. Chen and I. S. Reed. A detection algorithm for optical targets in clutter. *IEEE Transactions on Aerospace and Electronic Systems*, vol. 23, January 1987, pp. 46-59.
- [5] S. C. Pohlig. An algorithm for detection of moving optical targets. *IEEE Transactions on Aerospace and Electronic Systems*, vol. 25, January 1989, pp. 56-63.
- [6] I. S. Reed, R. M. Gagliardi and H. M. Shao. Application of three-dimensional filtering to moving target detection. *IEEE Transactions on Aerospace and Electronic Systems*, vol. 19, November 1983, pp. 898-905.
- [7] S. M. Schweitzer and J. M. F. Moura, "Hyperspectral imagery: clutter adaptation in anomaly detection," *IEEE Transactions on Information Theory*, vol. 46, n.5, August 2000, pp. 1855-1871.
- [8] T. E. Fortmann, Y. Bar-Shalom and M. Scheffe. Sonar tracking of multiple targets using joint probabilistic data association. *IEEE Journal of Oceanic Engineering*, OE-8, 3, July 1983, pp. 173-184.
- [9] S. B. Gelfand, T. E. Fortmann and Y. Bar-Shalom. Adaptive detection threshold optimization for tracking in clutter. *IEEE Transactions on Aerospace and Electronic Systems*. Vol 32, n.2, April 1996, pp. 514-522.
- [10] G. D. Forney Jr. The Viterbi algorithm. *Proceedings of the IEEE*, 61, 3, March 1973, pp 268-277.
- [11] R. E. Larson and J. Peschon. A dynamic programming approach to trajectory estimation. *IEEE Transactions on Automatic Control*, vol. 11, July 1966, pp. 537-540.
- [12] R. Bellman. *Dynamic Programming*. Princeton, New Jersey: Princeton University Press, 1957.
- [13] E. Conte and M. Longo. Characterisation of radar clutter as a spherically invariant random process. *IEE Proceedings* Vol. 134, Pt.F, No 2, April 1987, pp. 191-197.
- [14] E. Conte, M. Longo and M. Lops. Modelling and simulation of non-Rayleigh radar clutter. *IEE Proceedings-F* Vol. 138, No 2, April 1991, pp. 121-130.
- [15] M. Rangaswamy, D. Weiner and A. Öztürk. Non-Gaussian random vector identification using spherically invariant random processes. *IEEE Transactions on Aerospace and Electronic Systems*. Vol 29, No. 1, January 1993, pp. 111-123.
- [16] M. Rangaswamy, D. Weiner and A. Öztürk. Computer generation of correlated non-Gaussian radar clutter. *IEEE Transactions on Aerospace and Electronic Systems*. Vol 31, No. 1, January 1995, pp 106-116.
- [17] J. K. Bounds. *The Infrared Airborne Radar Sensor Suite*. RLE Technical Report No. 610, Cambridge, Massachusetts, December 1996.
- [18] J. L. Eaves and E. K. Reedy. *Principles of Modern Radar*. Van Nostrand Reinhold Company, New York, 1987.

- [19] G. Mamic, N. Stitt, R. Iskander. Coherent detection of radar signals in G-distributed clutter. *Proceedings IEEE ICASSP'99*, SPTM2.2, vol.3, pp. 1165-1168.
- [20] M. Sekine, S. Ohtani. Weibull distributed ground clutter. *IEEE Transactions on Aerospace and Electronic Systems*. vol. 17, July 1981, pp 596-598.
- [21] J. K. Jao. Amplitude distribution of composite terrain radar clutter and the K-distribution. *IEEE Transactions on Antennas and Propagation*, AP-32, October 1984, pp. 1049-1062.
- [22] C. J. Oliver. Representation of radar sea clutter. *IEE Proceedings*, 135, Pt. F 6, December 1988, pp. 497-500.
- [23] J. M. F. Moura and M. G. S. Bruno. DCT/DST and Gauss-Markov Fields: Conditions for Equivalence. *IEEE Transactions on Signal Processing*, vol. 46, n. 9, September 1998, pp. 2571-2574.
- [24] M. G. S. Bruno and J. M. F. Moura. The optimal 2D multiframe detector/tracker. *AEÜ International Journal of Electronics and Communications*, 53(6), Special Issue on Statistical Signal Processing, December 1999, pp. 346-355.
- [25] J. M. F. Moura and N. Balram. Recursive structure of noncausal Gauss-Markov random fields. *IEEE Transactions on Information Theory* IT-38(2), March 1992, pp 334-354.
- [26] N. Balram and J. M. F. Moura. Parameter estimation in noncausal Gauss-Markov fields. *IEEE Transactions on Information Theory*, vol.39, n.4, July 1993, pp. 1333-1355.
- [27] A. Graham, *Kronecker Products and Matrix Calculus with Applications*, Ellis Horwood Limited, 1981.
- [28] J. W. Woods. Two-dimensional discrete Markovian fields. *IEEE Transactions on Information Theory*, vol. IT-18, March 1972, pp. 232-240.
- [29] R. S. Bucy, C. Hecht and K. D. Senne. *An Engineer's Guide to Building Nonlinear Filters*. Frank J. Seiler Research Laboratory, SRL-TR-72-0004, May 1972.
- [30] J. W. Woods. Two-dimensional discrete Markovian fields. *IEEE Transactions on Information Theory*, vol. IT-18, March 1972, pp. 232-240.
- [31] A. H. Jazwinski. *Stochastic Processes and Filtering Theory*, Academic Press, New York, 1970.
- [32] H. L. Van Trees. *Detection, Estimation and Modulation Theory, Part I*. John Wiley & Sons, 1968.
- [33] *Center for Imaging Science*, Johns Hopkins University, <http://www.cis.jhu.edu>.