

The Optimal 2D Multiframe Detector/Tracker

Marcelo G.S. Bruno and José M. F. Moura

Dedicated to Professor Johann F. Böhme on the occasion of his 60th birthday

Abstract We detail in this paper the implementation of the optimal Bayes multiframe detector/tracker for rigid objects moving randomly in two-dimensional (2D) finite grids. We present 2D models for target signature and target motion that build an integrated framework for detection and tracking. We model the background clutter by 2D correlated noncausal Gauss-Markov fields of arbitrary order. By exploring the structure of the signature, motion, and clutter models, we indicate how substantial computational savings can be achieved in the implementation of the algorithm. The detection performance of the proposed Bayes scheme is evaluated through Monte Carlo simulations. The results show significant performance gains of over 6 dB in peak signal-to-noise ratio when the optimal multiframe detector is compared to the optimal single frame likelihood ratio test (LRT) detector.

Keywords Multiframe detection and tracking, Nonlinear stochastic filtering, Noncausal Gauss-Markov random fields

1. Introduction

Improved sensors, such as high resolution radars, sonars, and precision infrared (IR) cameras, make it important to develop and implement better performing signal processing algorithms. In particular, the problem of automatic detection and tracking of targets has recently received increasing attention, e.g. [3, 5, 6]. In [7] and [8], we introduced an optimal one-dimensional (1D) Bayesian algorithm for integrated multiframe detection and tracking of rigid objects in finite discrete grids. In the present paper, we extend the algorithm to two dimensions (2D), describe in detail its computational implementation, and present 2D performance results.

The optimal multiframe Bayes detector/tracker processes a sequence of noisy 2D images, generated for example by a radar or IR imaging sensor. We refer to each image in the sequence as a sensor frame. The frames have a finite resolution, so they are better represented by a finite lattice where each site or pixel represents one sensor

resolution cell. We restrict our study in this paper to the single target case, i.e., at each sensor scan, there is at most one target present in the sensor grid. The target images are cluttered by returns from spurious reflectors plus measurement noise.

The task is to determine whether a target is present or not at a given frame (detection) and, if the target is declared present, to estimate its position (tracking). Prior approaches to this problem separate detection and tracking [2, 4]: a preliminary single frame detector produces an initial estimate of the target state which is subsequently associated as a noisy measurement to a linearized multiframe tracking filter [2]. By contrast, we use nonlinear stochastic filtering to develop the optimal multiframe Bayes detector/tracker that processes the sensor images directly and integrates detection and tracking into a unified framework. The algorithm incorporates the target motion, target signature, and clutter models into both detection and tracking, and uses all past and present available observations to make decisions.

The optimal Bayes detector/tracker provides a bound to the performance that can be achieved by any suboptimal algorithm. In [7], we presented several Monte Carlo-based performance results for the Bayes algorithm in one dimension (1D), and compared against these bounds the performance of other common schemes such as spatial correlators (matched filters), and linearized Kalman-Bucy (KBf) trackers. In this paper, we present 2D detection performance results and compare the optimal multiframe Bayes detector with a conventional single frame likelihood ratio test (LRT) detector.

Section 2 describes the models for the sensor, target signature, target motion, and clutter that underly our integrated framework for detection and tracking. These models extend the 1D models introduced in [8]. We consider 2D extended targets with random translational motion and known deterministic signatures. Other scenarios with rotational motion and random signatures are subject of current research. We assume spatially correlated clutter modeled by noncausal, spatially homogeneous, 2D Gauss-Markov random fields (GMrf) [9, 10] of arbitrary order. GMrfs describe the clutter intensity at one given pixel as a weighted average of the intensity of the neighboring pixels plus a random error term. Such models are realistic in many practical scenarios and can be used to represent a variety of backgrounds, ranging from smooth patterns to highly structured texture [9]. Non-Gaussian clutter models with heavy tail statistics are examined in other papers [8].

Received July 30, 1999.

Marcelo G. S. Bruno is with the Electrical Engineering Department, University of São Paulo, P.O. Box 61548, São Paulo SP 05424-670, Brazil.

José M. F. Moura is with Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 02139, on sabbatical leave from the ECE Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

Using the models described in section 2, we derive in section 3 the 2D multiframe Bayes detector/tracker. Section 4 gives a more detailed description of the computational implementation of the algorithm. We present in section 5 the receiver operating characteristic (ROC) curves for the multiframe Bayes algorithm and compare them with the ROC curves for a conventional single frame LRT detector. The results show that, in a scenario of low signature targets and heavy clutter, optimal multiframe detection provides a significant improvement with peak signal-to-noise ratio gains of up to 6 dB over the optimal single frame detector. Finally, section 6 summarizes the main contributions of the paper.

2. The model

We present in this section the models for the sensor, the target signatures, the target motion, and the clutter. We restrict possible targets to rigid bodies with translational motion. We assume that, at any given sensor scan, the noise-free image of a target that is present is contained in 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. For simplicity, we consider that there is at most one target present at each sensor frame.

2.1 Sensor and target model

We model the region that is scanned by the sensor as the uniform 2D finite discrete lattice $\mathcal{L} = \{(i, j): 1 \leq i \leq L, 1 \leq j \leq M\}$, where L and M are the number of sensor resolution cells in each dimension. We refer to lattice \mathcal{L} as the *sensor lattice*.

It is useful to extend the sensor lattice so that we can model situations when a target moves in and out of the sensor grid or, when a target is absent from the grid. We define the *centroid lattice* $\widehat{\mathcal{L}} = \{(i, j): -r_s + 1 \leq i \leq L + r_i, -l_s + 1 \leq j \leq M + l_i\}$ that collects all possible values of the target centroid position for which at least one target pixel is within the 2D region scanned by the sensor.

We introduce an equivalent 1D representation for the 2D centroid lattice using row *lexicographic ordering*. This operation consists of sequentially stacking all rows of the 2D lattice in one long 1D lattice. The 1D *lexicographed centroid lattice*, $\overline{\mathcal{L}}$, is

$$\overline{\mathcal{L}} = \{l: 1 \leq l \leq (L + r_i + r_s)(M + l_i + l_s)\} . \quad (1)$$

The position of the target centroid at instant n in the lattice $\overline{\mathcal{L}}$ is indicated by the 1D random variable z_n such that, for $1 \leq k \leq L + r_i + r_s$ and $1 \leq p \leq M + l_i + l_s$,

$$z_n = (k - 1)(M + l_i + l_s) + p . \quad (2)$$

The formal equivalence between the 1D lattice $\overline{\mathcal{L}}$ and the 2D lattice $\widehat{\mathcal{L}}$ is established by observing that, if z_n is the target centroid position at instant n in the 1D lattice $\overline{\mathcal{L}}$ and \mathbf{z}_n is the target centroid position at instant n in the equivalent 2D lattice $\widehat{\mathcal{L}}$, then

$$z_n = (k - 1)(M + l_i + l_s) + p \Leftrightarrow \mathbf{z}_n = (k - r_s, p - l_s) \quad (3)$$

for $1 \leq k \leq L + r_i + r_s$ and $1 \leq p \leq M + l_i + l_s$.

Finally, to build an integrated framework for detection and tracking, we augment the 1D lexicographed lattice with an additional dummy state that represents the absence of the target from the sensor image. For convenience, we assign to the absent state the index $(L + r_i + r_s)(M + l_i + l_s) + 1$. The final 1D *extended lattice* is

$$\widetilde{\mathcal{L}} = \{l: 1 \leq l \leq (L + r_i + r_s)(M + l_i + l_s) + 1\} . \quad (4)$$

Target Model Let $\mathcal{B} = \{0, 1\}$ and let I be the 2D finite lattice $I = \{(k, l): -r_i \leq k \leq r_s, -l_i \leq l \leq l_s\}$. In this paper, we refer to a 2D sequence indexed on a discrete 2D lattice as a *field*. We use the notations $x_{k,l}$ and $x(k, l)$ interchangeably to denote the individual elements of a 2D field.

Target Signature: The target signature is determined by the product of the shape parameters, which represent the target geometry, and the intensity parameters, which describe the actual intensity of each pixel in the target image.

The target shape at instant n is specified by the field $\{c_{k,l}^n\}, (k, l) \in I$ such that, for each $(k, l) \in I$, $c_{k,l}^n \in \mathcal{B}$. The target intensity is in turn described by the real-valued field $\{\phi_{k,l}^n\}, (k, l) \in I$, such that, for each $(k, l) \in I$, $\phi_{k,l}^n \in \mathfrak{R}$. Finally, we define the target signature at instant n as the field $\{a_{k,l}^n\}, (k, l) \in I$ such that, $a_{k,l}^n = c_{k,l}^n \phi_{k,l}^n$.

For rigid bodies, the shape parameters $c_{k,l}^n$ are time-invariant and, accordingly, we can drop the superscript. We assume that the target shape is known to the detector/tracker. The intensity parameters, on the other hand, may be time-variant, possibly deterministic but unknown and, in general, may be random as a result of stochastic fluctuations in the reflectivity and illumination of the targets or due to random variations in the channel characteristics. For simplicity, we consider as a first approximation in this paper that the intensity parameters are time-invariant, deterministic, and known. Since the shape is also assumed time-invariant and known, the target signature at each scan n is perfectly specified by a known, deterministic and finite 2D field $\{a_{k,l}\}, (k, l) \in I$. We introduce then the *target signature template* field $G(k, l)$ such that, for $-r_s \leq k \leq r_i, -l_s \leq l \leq l_i$,

$$G(k, l) = a(-k, -l) . \quad (5)$$

This definition will be useful to detail the structure of the optimal Bayes algorithm in subsequent sections.

2D Extended Target Model: Let z_n be the position of the target centroid in the 1D extended lattice $\tilde{\mathcal{L}}$ defined in equation (4). We represent the noise-free image of a target centered at z_n during the n th sensor scan as the nonlinear mapping

$$\begin{aligned} \mathbf{F}_n: \tilde{\mathcal{L}} &\mapsto \Re^{L \times M} \\ z_n &\rightarrow \mathbf{F}(z_n) \end{aligned} \quad (6)$$

where, for $z_n = (i_n + r_s - 1)(M + l_i + l_s) + (j_n + l_s)$, $(i_n, j_n) \in \tilde{\mathcal{L}}$, the function $\mathbf{F}(z_n)$ is given by

$$\mathbf{F}[z_n(i_n, j_n)] = \sum_{k=-r_i}^{r_s} \sum_{l=-l_i}^{l_s} a_{k,l}^n \mathbf{E}_{i_n+k, j_n+l} \cdot \quad (7)$$

In (7), for $1 \leq i \leq L$, $1 \leq j \leq M$, $\mathbf{E}_{i,j}$, is an $L \times M$ matrix whose entries are all zero, except for the element (i, j) which is one. For any $(i, j) \notin \mathcal{L}_1 \times \mathcal{L}_2$, where $\mathcal{L}_1 = \{l: 1 \leq l \leq L\}$ and $\mathcal{L}_2 = \{l: 1 \leq l \leq M\}$, we define $\mathbf{E}_{i,j}$ to be identically zero. Finally, if $z_n = (L + r_i + r_s)(M + l_i + l_s) + 1$, we make

$$\mathbf{F}(z_n) = \mathbf{0}_{L \times M} \cdot \quad (8)$$

The noise-free target model in (7) is a function that associates to each possible centroid position of a target that is present a spatial distribution of target signature parameters centered at that respective centroid position. Since $\mathbf{E}_{i,j}$ is defined as zero whenever $(i, j) \notin \mathcal{L}_1 \times \mathcal{L}_2$, the model in (7) automatically takes into consideration the fact that, as a target moves in and out of the surveillance space, portions of the target that lie outside the sensor grid are not visible in the sensor image. Whenever a target is physically absent from the sensor scan, the corresponding target image is an identically zero matrix, as indicated by equation (8).

2.2 Observations and clutter model

The measurements at the n th sensor assuming a single target are the $L \times M$ matrix

$$\mathbf{Y}_n = \mathbf{F}(z_n) + \mathbf{V}_n \quad (9)$$

where z_n is the position of the target centroid in the equivalent 1D extended lattice (including the absent state), $\mathbf{F}(\cdot)$ is the 2D extended target model described by equations (7) and (8), and \mathbf{V}_n is the background clutter matrix, also referred to as the background clutter frame. We assume that the clutter frames \mathbf{V}_n , $n = 0, 1, \dots$, are independent, identically distributed (i.i.d.).

Gaussian clutter Let \mathbf{v}_n be the equivalent 1D row lexicographed representation of the 2D clutter matrix \mathbf{V}_n . Under the assumption of Gaussianity, vector \mathbf{v}_n has the probability density function (pdf), $p(\mathbf{v}_n) = N(\mathbf{0}, \mathbf{R})$, where \mathbf{R} is the clutter spatial covariance, $\mathbf{0}$ is the mean, and $N(\cdot)$ stands for the normal (multivariate Gaussian) pdf. The zero mean assumption assumes a pre-processing stage that removes the mean. A non-zero mean can be accounted for trivially.

Spatially correlated homogeneous GMrf model We model the spatial correlation of the clutter using Gauss-Markov random fields (GMrf) [9, 10]. We present next the formal definition of a 2D GMrf. For simplicity of notation, we restrict the discussion to square lattices, i.e., $L = M$.

For each pixel $(i, j) \in \mathcal{L} = \{(i, j): 1 \leq i, j \leq L\}$, we define its neighborhood system η_{ij} . Under the assumption of spatial homogeneity, $\eta_{ij} = \eta$, $\forall (i, j) \in \mathcal{L}$. The clutter field, \mathbf{V}_n , is a finite order, spatially homogeneous GMrf if it is the output of the 2D finite difference equation [9]

$$V_n(i, j) = \sum_{(i-k, j-l) \in \eta} -\alpha_k^l V_n(i-k, j-l) + U_n(i, j) \quad (10)$$

for $(i, j) \in \mathcal{L}$. In (10), we used the assumption of spatial homogeneity to make the coefficients α_k^l independent of (i, j) . We also added the assumption of diagonal symmetry such that the coefficients α_k^l satisfy the equalities $\alpha_{-k}^l = \alpha_k^l$, $\alpha_k^{-l} = \alpha_k^l$ and $\alpha_{-k}^{-l} = \alpha_k^l$.

We add boundary conditions (bc's) to equation (10) so that it can extend to the boundaries of the lattice. A common choice of bc's is to make $V_n(i, j) = 0$ for any $(i, j) \notin \mathcal{L}$. These are called *Dirichlet* boundary conditions. Alternatively, other boundary conditions could be used [9, 10].

Second Order Statistics of GMrfs: Let \mathbf{v}_n and \mathbf{u}_n be the 1D row lexicographed representations of \mathbf{V}_n and \mathbf{U}_n respectively. Equation (10) is written compactly in matrix format as

$$\mathbf{A} \mathbf{v}_n = \mathbf{u}_n \cdot \quad (11)$$

It follows from the application of the orthogonality principle that the covariance matrix of the driving noise \mathbf{u}_n is [9]

$$\Sigma_u = \sigma_u^2 \mathbf{A} \quad (12)$$

and the covariance matrix of the clutter \mathbf{v}_n is

$$\Sigma_v = \mathbf{R} = \sigma_u^2 \mathbf{A}^{-1} \Leftrightarrow \mathbf{A} = \sigma_u^2 \mathbf{R}^{-1} \quad (13)$$

where $\sigma_u^2 = E[v_n(l)u_n(l)]$, for all l in the lexicographed equivalent of the sensor lattice \mathcal{L} . The quantity σ_u^2 is spatially invariant due to the spatial homogeneity assumption. In the particular case of Dirichlet bc's, σ_u^2 is also equal to $E[u_n(l)^2]$ [9, 10]. The symbol $E[\cdot]$ stands for expected value or ensemble average.

We refer to the matrix \mathbf{A} in (11) as the *potential matrix* [10]. It follows from (13) that the inverse of the covariance matrix, Σ_v , is proportional to the potential matrix. The potential matrix is in turn a highly structured matrix described by a limited number of parameters. In section 4, we explore this matrix structure to design computationally efficient detection and tracking algorithms. We describe next the structure of \mathbf{A} in detail.

Structure of the potential matrix: Let W_p be the maximum Euclidean distance between sites (i, j) and $(i-k, j-l)$ such that $(i-k, j-l) \in \eta$. The quantity W_p associated to a particular neighborhood system η defines the *order* of

the Markov model. Assuming Dirichlet boundary conditions, spatial homogeneity and diagonal symmetry, we see by inspection that the potential matrix \mathbf{A} has a symmetric, *block banded*, *block Toeplitz* structure of the form [9, 11]

$$\mathbf{A} = \mathbf{I} \otimes \mathbf{D}_0 + \sum_{j=1}^m (\mathbf{K}_1^j + \mathbf{K}_2^j) \otimes \mathbf{D}_j. \quad (14)$$

where m is the largest integer such that $m \leq W_p$. The matrix \mathbf{K}_1 , referred to as the backward shift matrix [1], has all its entries equal to zero except for the first upper diagonal whose entries are all equal to 1. The matrix \mathbf{K}_2 , called the forward shift matrix, is the transpose of \mathbf{K}_1 . The symbol \otimes denotes the Kronecker (or tensor) product [13]. The blocks D_i , $1 \leq i \leq m$, in (14) are themselves structured with a *Toeplitz and banded* structure of the form

$$\begin{aligned} \mathbf{D}_0 &= \mathbf{I} + \sum_{t=1}^{l(0)} \alpha_0^t (\mathbf{K}_1^t + \mathbf{K}_2^t) \\ \mathbf{D}_j &= \alpha_j^0 \mathbf{I} + \sum_{t=1}^{l(j)} \alpha_j^t (\mathbf{K}_1^t + \mathbf{K}_2^t) \quad 1 \leq j \leq m \end{aligned} \quad (15)$$

In (15), m is the largest integer such that $m \leq W_p$, $l(0) = m$, and $l(j)$, $1 \leq j \leq m$, is the largest integer such that $l^2(j) + j^2 \leq W_p^2$. The parameters α_0^r , $1 \leq r \leq m$ and α_j^t , $1 \leq j \leq m$, $0 \leq t \leq l(j)$, are defined on a parameter space \mathcal{P} such that the potential matrix \mathbf{A} is positive definite. A comprehensive analysis of the eigenstructure of the potential matrix in the 2D case for different choices of bc's and its relation to fast sinusoidal orthogonal transforms is found in [11].

GMrf Model Template: Given the potential matrix \mathbf{A} associated to a particular noncausal GMrf model, we define the GMrf *model template* as the field $H(i, j)$ such that

$$\begin{aligned} H_{i,j} &= \delta_{i,j} + \sum_{r=1}^m \alpha_0^r [\delta_{i,j-r} + \delta_{i,j+r}] \\ &+ \sum_{t=1}^m \left\{ \alpha_t^0 \delta_{i+t,j} + \sum_{r=1}^{l(t)} \alpha_t^r [\delta_{i+t,j+r} + \delta_{i+t,j-r}] \right\} \\ &+ \sum_{t=1}^m \left\{ \alpha_t^0 \delta_{i-t,j} + \sum_{r=1}^{l(t)} \alpha_t^r [\delta_{i-t,j+r} + \delta_{i-t,j-r}] \right\} \end{aligned}$$

In the preceding equation, $\delta_{i,j}$ is the general element of the 2D delta Kronecker field. This definition of model template will be useful to detail the structure of the Bayes detector/tracker in section 4.

Examples of Sensor Frames In the sequel, we present an illustrative example of a synthetic sensor frame that was simulated using the target and clutter models just introduced. Under the assumption of deterministic and known signatures, without loss of generality, we make the target pixel intensities constant and equal to 1. Figure 1(a) shows the noise-free image of a rectangular-shaped target in a 2D

100 x 100 sensor grid, when the target is centered at coordinates (50, 50). Figure 1(b) shows the same target image plus 2D spatially correlated clutter. The clutter is a first order, noncausal GMrf with $\alpha_0^1 = -\beta_h = \alpha_1^0 = -\beta_v = 0.24$. The signature-to-clutter ratio is measured using the peak signal-to-noise ratio (PSNR) defined as

$$\text{PSNR} = 10 \log_{10} \left(\frac{1}{\sigma_u^2} \right). \quad (16)$$

The PSNR in figure 1 is equal to 0 dB. The background clutter is simulated from white noise using the equivalent recursive (one-sided) representation of equation (11) obtained through the upper Cholesky factorization of the potential matrix \mathbf{A} . For simulation details, see [10].

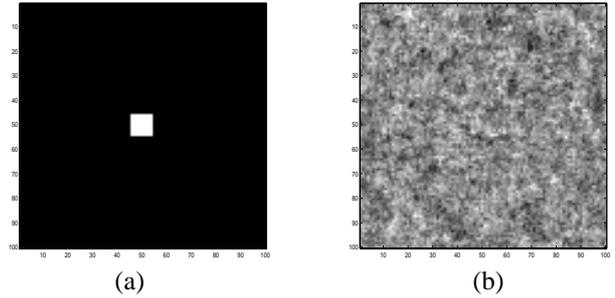


Fig. 1. (a) Noise-free rectangular target image. (b) Simulated sensor scan with target plus correlated GMrf clutter, PSNR=0 dB

2.3 Motion model

With translational motion, it suffices to model the dynamics of the target centroid to describe the target motion. The changes between two consecutive frames in the position of the target centroid in the extended lattice $\tilde{\mathcal{L}}$ are specified by a *transition probability matrix*, \mathbf{P}_T , whose general element $P_T(k, r)$, $(k, r) \in \tilde{\mathcal{L}} \times \tilde{\mathcal{L}}$, is

$$P_T(k, r) = \text{Prob}(z_n = k \mid z_{n-1} = r). \quad (17)$$

2D drifting targets We focus our discussion on 2D randomly drifting targets, i.e., targets that move with a constant nominal drift disturbed by a 2D random walk fluctuation. Let $\mathbf{z}_n = (i_n, j_n)$ be the position in the centroid lattice $\tilde{\mathcal{L}}$ of a target at the n th scan. Let η_c be a bounded 2D finite discrete lattice limited by a maximum Euclidean distance \mathcal{D} , i.e., $\eta_c = \{(t, s) : t^2 + s^2 \leq \mathcal{D}^2, t, s \in \mathcal{Z}\}$. We refer to η_c as the random walk *fluctuation region* in the 2D plane. We model the evolution in time of the centroid position (i_n, j_n) by the equations

$$\begin{aligned} i_n &= i_{n-1} + d_1 + \varepsilon_n^1 \\ j_n &= j_{n-1} + d_2 + \varepsilon_n^2 \end{aligned} \quad (18)$$

where $[(\varepsilon_n^1)^2 + (\varepsilon_n^2)^2 \leq \mathcal{D}^2]$, $\varepsilon_n^m \in \mathcal{Z}$, $m = 1, 2$. For a first order noncausal fluctuation region, $\mathcal{D} = 1$ and ε_n^m , $m = 1, 2$, may take values in the set $\mathcal{S} = [-1, 0, 1]$. A

2D first order random walk is shown in figure 2 with the respective transition probabilities. In the figure, under the zero mean random walk assumption, $s = t$ and $r = q$. In equation (18), we assume that ε_n^1 and ε_n^2 are mutually

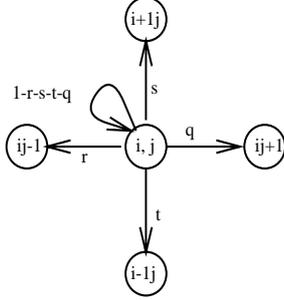


Fig. 2. Probabilities of fluctuation of the centroid position around pixel (i, j) for a 2D random walk model

independent and also statistically independent of the centroid position (i_n, j_n) .

To complete the motion model, we need to specify the transition probabilities to and from the absent target state. Whenever the target centroid crosses the boundaries of the centroid lattice $\tilde{\mathcal{L}}$, the target is declared absent and the random variable z_n takes the absent target state value, i.e., $z_n = (L + r_i + r_s)(M + l_i + l_s) + 1$. Conversely, when no target is present, there is a non-zero probability of a new target appearing randomly at the next sensor scan. We assume that there is an equal probability of the new target appearing centered at any pixel of the centroid lattice $\tilde{\mathcal{L}}$.

In [1], we detail the analytic structure of the transition probability matrix \mathbf{P}_T for the 2D drifting target model. This matrix is highly structured and sparse, consisting basically of a shifted block tridiagonal structure, where the blocks themselves are structured. We explore this structure to achieve further computational savings in the implementation of the optimal Bayes detector/tracker. We omit this discussion here for lack of space and refer the reader to [1].

3. Optimal multiframe detector/tracker

We develop in this section the solution to the problem of integrated detection and tracking of a single target in a 2D finite grid using the models from section 2.

Let z_n be the position of the target centroid at the n th sensor scan defined on the extended lattice $\tilde{\mathcal{L}}$ and denote by \mathbf{y}_n the row lexicographed version of the 2D sensor image \mathbf{Y}_n at the n th scan. We introduce first the *filtering posterior probabilities vector*, $\mathbf{p}_{n|n}$, whose l th component is

$$p_{n|n}(l) = P(z_n = l \mid \mathbf{Y}_0^n) \quad l \in \tilde{\mathcal{L}}. \quad (19)$$

The vector $\mathbf{Y}_0^n = [\mathbf{y}_0^T \mathbf{y}_1^T \dots \mathbf{y}_n^T]^T$ collects all the observations from the initial time 0 up to instant n . Like-

wise, we define the *prediction posterior probabilities vector*, $\mathbf{p}_{n|n-1}$, whose l th component is

$$p_{n|n-1}(l) = P(z_n = l \mid \mathbf{Y}_0^{n-1}) \quad l \in \tilde{\mathcal{L}}. \quad (20)$$

We derive next an algorithm for the recursive computation of the filtering posterior probabilities vector. The algorithm is divided into two steps.

Prediction Step Predict the next state of the target from the previous one using the transition probability matrix, \mathbf{P}_T . By the Theorem of Total Probability and using the fact that, conditioned on z_{n-1} , z_n is independent from \mathbf{Y}_0^{n-1} , we get

$$P(z_n \mid \mathbf{Y}_0^{n-1}) = \sum_{z_{n-1}} P(z_n \mid z_{n-1})P(z_{n-1} \mid \mathbf{Y}_0^{n-1}) \quad (21)$$

In matrix notation, equation (21) is written as

$$\mathbf{p}_{n|n-1} = \mathbf{P}_T \mathbf{p}_{n-1|n-1}. \quad (22)$$

Filtering Step Correct the prediction with the new information given by the observations. From Bayes' Law and using the fact that, conditioned on z_n , \mathbf{y}_n is independent from \mathbf{Y}_0^{n-1} , we can write

$$P(z_n \mid \mathbf{Y}_0^n) = C_n p(\mathbf{y}_n \mid z_n) P(z_n \mid \mathbf{Y}_0^{n-1}). \quad (23)$$

In matrix notation, equation (23) is rewritten as the point-wise vector multiplication

$$\mathbf{p}_{n|n} = C_n \mathbf{S}_n \odot \mathbf{p}_{n|n-1} \quad (24)$$

where

$$S_n(l) = p(\mathbf{y}_n \mid z_n = l) \quad l \in \tilde{\mathcal{L}}. \quad (25)$$

The constant C_n is a normalization factor such that

$$\sum_{l \in \tilde{\mathcal{L}}} p_{n|n}(l) = 1. \quad (26)$$

We now consider detection and tracking.

Detection Let $L_2 = (L + r_i + r_s)(M + l_i + l_s)$. The probability of the target being absent at instant n conditioned on the observations is given by $p_{n|n}(L_2 + 1)$. Representing by H_0 the hypothesis that the target is absent and by H_1 , the hypothesis that the target is present, the minimum probability of error detector, assuming equal cost for misses and false alarms and zero cost for correct decisions, is the test [12]

$$\frac{P(H_0 \mid \mathbf{Y}_0^n)}{P(H_1 \mid \mathbf{Y}_0^n)} \underset{H_1}{\overset{H_0}{>}} 1 \Leftrightarrow \frac{p_{n|n}(L_2 + 1)}{1 - p_{n|n}(L_2 + 1)} \underset{H_1}{\overset{H_0}{>}} 1. \quad (27)$$

Alternative tests are obtained by changing the threshold in (27). Each value of the threshold will correspond to a fixed probability of false alarm. By varying the threshold over a wide range, the performance of the test is assessed through a receiver operating characteristic (ROC) curve that plots probability of detection versus probability of false alarm.

MAP Tracking Introduce now the conditional probability

$$\begin{aligned} Q_l^f [n] &= P(\mathbf{z}_n = l \mid \text{target is present}, \mathbf{Y}_0^n) \quad l \in \bar{\mathcal{L}} \\ &= \frac{p_{n|n}(l)}{1 - p_{n|n}(L_2 + 1)} \end{aligned} \quad (28)$$

where $\bar{\mathcal{L}}$ is the lexicographed centroid lattice, see section 2. The maximum a posteriori (MAP) estimate of the target's centroid position assuming that the target is present is [12]

$$\hat{z}_{\text{map}} [n] = \arg \max_{l \in \bar{\mathcal{L}}} Q_l^f [n] . \quad (29)$$

4. Implementation of the multiframe Bayes detector/tracker

We detail in this section the implementation of the optimal Bayes detector/tracker for a single target moving in GMrf clutter. We focus our discussion on the implementation of the filtering step of the algorithm, see section 3. The implementation of the prediction step is detailed in [1].

4.1 Filtering step: correlated clutter

The filtering step in equation (24) involves the computation at each instant n of the observations kernel \mathbf{S}_n . Assume, for simplicity of notation, but without loss of generality, that $L = M$, $l_i = r_i$, and $l_s = r_s$. Let $L_1 = (L + r_i + r_s) = (M + l_i + l_s)$. Let \mathbf{y}_n be the lexicographed $L^2 \times 1$ representation of the $L \times L$ sensor image \mathbf{Y}_n such that

$$y_n((i-1)L + j) = Y_n(i, j) \quad 1 \leq i, j \leq L . \quad (30)$$

The observations kernel is written as the $(L_1^2 + 1) \times 1$ vector \mathbf{S}_n such that, for $1 \leq i, j \leq L_1$,

$$S_n((i-1)L_1 + j) = p(\mathbf{y}_n \mid z_n = (i-1)L_1 + j) . \quad (31)$$

The entry corresponding to the absent target state is

$$S_n(L_1^2 + 1) = p(\mathbf{y}_n \mid z_n = L_1^2 + 1) . \quad (32)$$

Computation of the Observations Kernel When the background clutter \mathbf{v}_n is a zero mean Gaussian vector with spatial covariance \mathbf{R} , the observations kernel is

$$S_n((i-1)L_1 + j) = N_{\mathbf{y}_n}(\mathbf{f}_{ij}, \mathbf{R}) \quad 1 \leq i, j \leq L_1 \quad (33)$$

where \mathbf{f}_{ij} is the lexicographed vector obtained from the noise-free 2D image of a target centered at (i, j) as modeled in (6). The notation $N(\cdot)$ in (33) stands for the multivariate Gaussian pdf

$$N_{\mathbf{a}}(\mathbf{b}, \mathbf{R}) = k \exp\left(-\frac{(\mathbf{a} - \mathbf{b})^T \mathbf{R}^{-1} (\mathbf{a} - \mathbf{b})}{2}\right) \quad (34)$$

with $k = (2\pi)^{-L^2/2} |\mathbf{R}|^{-1/2}$. On the other hand, for the absent target state,

$$S_n(L_1^2 + 1) = N_{\mathbf{y}_n}(\mathbf{0}, \mathbf{R}) . \quad (35)$$

The computation under the Gaussian assumption of the observations kernel requires the evaluation of the quadratic form $(\mathbf{y}_n - \mathbf{f}_{ij})^T \mathbf{R}^{-1} (\mathbf{y}_n - \mathbf{f}_{ij})$ for each pair (i, j) . For a generic $L^2 \times L^2$ full inverse covariance matrix, \mathbf{R}^{-1} , and a generic $L^2 \times 1$ mean vector \mathbf{f}_{ij} , $1 \leq i, j \leq L_1$, this operation requires $O(L^6)$ floating point multiplications [1]. These computational requirements can be reduced between three to four orders of magnitude if we explore the structures of \mathbf{R}^{-1} , \mathbf{f}_{ij} , and of the multivariate Gaussian pdf.

We restrict the background clutter to be a spatially homogeneous, noncausal GMrf of arbitrary order characterized by a potential matrix with structure as in (14). We recall that, from equation (13) in section 2.2, $\mathbf{A} = \sigma_u^2 \mathbf{R}^{-1}$, where \mathbf{A} is the potential matrix and σ_u is the power of the driving noise in the 2D finite difference equation model for the GMrf clutter as in (10).

For notational convenience, we introduce the normalized quadratic form, Q_{ij} , such that, for $1 \leq i, j \leq L_1$,

$$\begin{aligned} Q_{ij} &= \sigma_u^2 (\mathbf{y}_n - \mathbf{f}_{ij})^T \mathbf{R}^{-1} (\mathbf{y}_n - \mathbf{f}_{ij}) \\ &= (\mathbf{y}_n - \mathbf{f}_{ij})^T \mathbf{A} (\mathbf{y}_n - \mathbf{f}_{ij}) . \end{aligned} \quad (36)$$

With this normalization, and absorbing all terms that not depend on (i, j) into the constant C_n in (24), we write the entries of the observations kernel, for $1 \leq i, j \leq L_1$, as

$$S_n((i-1)L_1 + j) = \exp\left(-\frac{Q_{ij}}{2\sigma_u^2}\right) . \quad (37)$$

For the absent state,

$$S_n(L_1^2 + 1) = \exp\left(-\frac{\mathbf{y}_n^T \mathbf{A} \mathbf{y}_n}{2\sigma_u^2}\right) . \quad (38)$$

Let $C = \mathbf{y}_n^T \mathbf{A} \mathbf{y}_n$. We rewrite the entries of the normalized quadratic form as

$$Q_{ij} = C - 2\lambda_{ij} + \rho_{ij} \quad 1 \leq i, j \leq L_1 \quad (39)$$

where $\lambda_{ij} = \mathbf{y}_n^T \mathbf{A} \mathbf{f}_{ij}$ is referred to as the *data term*, and $\rho_{ij} = \mathbf{f}_{ij}^T \mathbf{A} \mathbf{f}_{ij}$ is the so-called *energy term*. It follows then that

$$\begin{aligned} S_n((i-1)L_1 + j) &= \exp\left(-\frac{Q_{ij}}{2\sigma_u^2}\right) = \\ &= \exp\left(-\frac{C}{2\sigma_u^2}\right) \exp\left(\frac{2\lambda_{ij} - \rho_{ij}}{2\sigma_u^2}\right) . \end{aligned} \quad (40)$$

when $1 \leq i, j \leq L_1$. For the absent state, on the other hand,

$$S_n(L_1^2 + 1) = \exp\left(-\frac{C}{2\sigma_u^2}\right) . \quad (41)$$

The term $\exp(-C/2\sigma_u^2)$ is independent of (i, j) and, therefore, can be also absorbed into the normalization con-

stant C_n in (24). We are left then with the *normalized* observations kernel, $\bar{\mathbf{S}}$, given by

$$\bar{S}_n((i-1)L_1 + j) = \exp\left(\frac{2\lambda_{ij} - \rho_{ij}}{2\sigma_u^2}\right) \quad 1 \leq i, j \leq L_1$$

$$\bar{S}_n(L_1^2 + 1) = 1. \quad (42)$$

Computation of the Data Term We present in the sequel an efficient algorithm for the computation of the data term λ_{ij} . This result expresses the data term for 2D extended targets as two linear convolution operations in cascade.

Let \mathbf{f}_{ij} be as before the lexicographed vector obtained from the noise-free 2D image of a target centered in (i, j) as in (6). Let $G(i, j)$ be the general term of the target signature template defined in (5), and let $H(i, j)$ be the general term of the GMrf model template defined in section 2.2. With this notation, and assuming that $l_i + l_s > m$ and $L \gg (l_i + l_s)$, the data term for $l_i + m + 1 \leq i, j \leq L - l_s - m$ is written as

$$\lambda_{ij} = \mathbf{y}_n^T \mathbf{A} \mathbf{f}_{ij} = G(i, j) * [Y_n(i, j) * H(i, j)]. \quad (43)$$

The symbol '*' denotes the 2D convolution operator. The proof of equation (43) involves some algebraic doing, we refer the reader to [1]. We omit the proof here for lack of space.

Boundary Conditions: Equation (43) shows how to compute the data term away from the boundaries of the sensor image. Close to the boundaries, both the target template and the GMrf model template must be changed to account for boundary conditions. Hence, for any (i, j) , $-l_s + 1 \leq i, j \leq L + l_i$, the data term λ_{ij} is given by the expressions in table 1 where (\cdot) is the product of $a_{k,l}$ by the output of the noncausal differential operator, $Y_n(i, j) * H(i, j)$ computed at $(i+k, j+l)$ with Dirichlet bc's, i.e., if $i \notin \{1 \leq i \leq L\}$ or $j \notin \{1 \leq i \leq L\}$, then $Y_{ij} = 0$. It can be shown [1] that that the raw computation of table 1 for a first order GMrf clutter requires a number of floating point multiplications that varies from $O(mL^2)$, $m \ll L$, when $(l_i + l_s + 1) \ll \sqrt{L}$, to $O(L^3)$ when $(l_i + l_s + 1) \approx \sqrt{L}$. That represents computational savings of between 3 and 4 orders of magnitude in comparison with the raw computation of the quadratic form in the multivariate Gaussian pdf.

Energy Term We examine now the computation of the energy term ρ_{ij} for 2D extended targets. As before, we assume for simplicity that $l_i = r_i$ and $l_s = r_s$. Given the target model, we define first the target energy

$$E_t(0, 0) = \sum_{k=-l_i}^{l_s} \sum_{l=-l_i}^{l_s} a(k, l)^2. \quad (44)$$

Similarly, we define the cross-energy coefficients

$$E(t, r) = \sum_{k=-l_i}^{l_s-t} \sum_{l=-l_i}^{l_s-r} a(k, l) a(k+t, l+r)$$

$$E(t, -r) = \sum_{k=-l_i}^{l_s-t} \sum_{l=-l_i+r}^{l_s} a(k, l) a(k+t, l-r)$$

$$E(-t, r) = \sum_{k=-l_i+t}^{l_s} \sum_{l=-l_i}^{l_s-r} a(k, l) a(k-t, l+r)$$

$$E(-t, -r) = \sum_{k=-l_i+t}^{l_s} \sum_{l=-l_i+r}^{l_s} a(k, l) a(k-t, l-r).$$

We can express ρ_{ij} as a linear combination of the target energy and the cross-energy coefficients. The weights in the linear combination are the parameters in the GMrf model defined by the potential matrix \mathbf{A} . For $l_i + m + 1 \leq i, j \leq L - l_s - m$, and $l_i + l_s > m$, we write

$$\rho_{ij} = \mathbf{f}_{ij}^T \mathbf{A} \mathbf{f}_{ij} =$$

$$E(0, 0) + \sum_{r=1}^m \alpha_0^r [E(0, r) + E(0, -r)] +$$

$$\sum_{t=1}^m \left\{ \alpha_t^0 E(t, 0) + \sum_{r=1}^{l(t)} \alpha_t^r [E(t, r) + E(t, -r)] \right\} +$$

$$\sum_{t=1}^m \left\{ \alpha_t^0 E(-t, 0) \right.$$

$$\left. + \sum_{r=1}^{l(t)} \alpha_t^r [E(-t, r) + E(-t, -r)] \right\}.$$

For the proof, we refer the reader to [1].

Remark The expression for the energy term in the preceding result is valid only away from the boundaries of the lattice. Near the boundaries, the target template must be conveniently modified to account for the fact that portions of the target disappear from the image and should not be used in the computation of the energy and cross-energy terms. Except at the boundaries, the energy term ρ_{ij} is constant for all indices i and j .

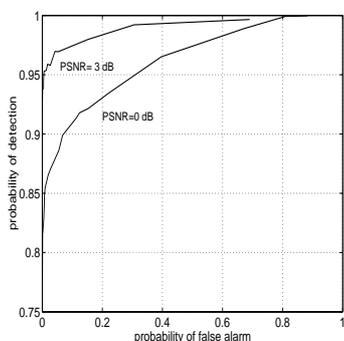
5. Detection performance

We study next the 2D detection performance of the optimal Bayesian algorithm developed in section 3. We plot the ROC curves for the optimal multiframe Bayes detector using synthetic data. At each sensor scan, there may be at most one target present. The simulated targets are rectangular-shaped, unit signature 2D rigid objects as in figure 1(a). Targets that are present move in the sensor grid with a translational motion described by the drifting target model explained in section 2.3. The fluctuation probabilities of one resolution cell about the average 2D drift are $r = s = q = t = 0.2$, see figure 2. The constant nominal velocities in the two dimensions are $d_1 = d_2 = 2$ resolution cells/frame. Once a target becomes absent, a new target can appear randomly in any position of the centroid grid with total probability $p_a = 0.2$. The size of the sensor grid is 100 x 100 and the size of the targets is 9 x 9. The simulated clutter is a spatially correlated, first order noncausal GMrf with parameters $\beta_v = \beta_h = 0.24$, like in figure 1(b).

Table 1. Data term λ_{ij} for $(i, j) \in \hat{\mathcal{L}}$.

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

Figure 3 shows the ROC curves for the optimal Bayes detector with PSNR equal to 3 and 0 dB. These levels of PSNR correspond to low signature/heavily cluttered targets that are hardly if at all distinguishable by a human operator, as shown previously in figure 1(b). The pairs (P_{fa}, P_d) in the ROC curves were estimated using 6,000 Monte Carlo runs for each value of the detection threshold. We see from the plots in figure 3 that, as the PSNR in-

Fig. 3. Optimal Bayes ROC curves in correlated first order GMrf clutter, $\beta_h = \beta_v = 0.24$

creases, the ROC curves tend to a step-like shape, i.e., for low levels of probability of false alarm, we obtain comparatively much higher probabilities of detection. Even at the low PSNR=3 dB, we obtain detection probabilities above 0.95 for false alarm rates in the order of 10^{-2} .

Multiframe and Single Frame Detection: Performance Comparison To study the improvement in performance resulting from the incorporation of motion dynamics information into the detection algorithm, we compare the optimal multiframe detector with a single frame LRT detector that ignores the motion models. The single frame LRT algorithm reduces to the test [12]

$$\frac{p(\mathbf{y}_n | H_0)}{p(\mathbf{y}_n | H_1)} \underset{H_1}{\overset{H_0}{>}} \lambda \quad (45)$$

where λ is a threshold that varies according to the desired probability of false alarm. We ran Monte Carlo simulations using rectangular targets and the same motion, signature and clutter model parameters as used before. The ROC curves, estimated from 3,000 Monte Carlo runs, are shown in figure 4(a) for both the multiframe (optimal)

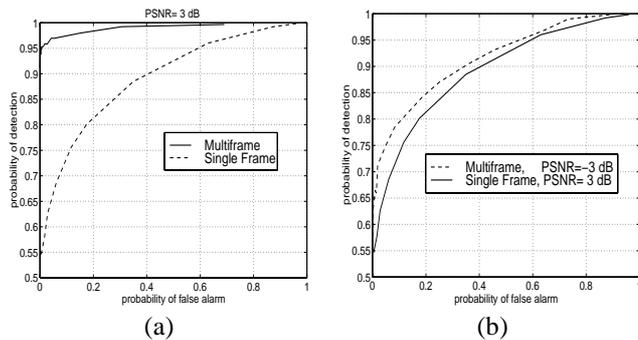


Fig. 4. Performance of the single frame and multiframe (optimal) detectors in correlated GMrf clutter

and single frame (memoryless) detectors in a situation of PSNR equal to 3 dB. The curves in figure 4(a) show that, although the single frame detector may perform well in weak clutter environments, its performance deteriorates significantly in scenarios of heavily cluttered targets.

Figure 4(b) repeats the ROC curves for both single frame and multiframe detectors, except that the PSNR for the multiframe detector is lowered to -3 dB. We note from the plot that the -3 dB multiframe ROC curve is closer to the 3 dB single frame ROC, but still lies slightly above the latter. That indicates a substantial gain in PSNR of over 6 dB when we introduce multiple frames and motion dynamics into the detection process.

6. Conclusions

We presented in this paper details of the implementation of the optimal Bayes multiframe detector/tracker for rigid objects that move randomly in two-dimensional (2D) cluttered images with finite resolution. We developed 2D models for target signature and target motion that build an integrated framework for detection and tracking. The background clutter was described by correlated, non-causal, spatially homogeneous GMrfs of arbitrary order. We explored the structure of the signature, motion, and clutter models, to achieve substantial computational savings in the implementation of the algorithm. The detection performance of the optimal multiframe Bayes detector was evaluated by Monte Carlo simulations using synthetic data. The results show significant performance gains of over 6 dB in peak signal-to-noise ratio in comparison with the suboptimal single frame LRT detector.

Acknowledgement

This work was supported by ONR grant N00014-97-0800. The first autor was also funded by CNPq, Brazil and FAPESP, São Paulo, Brazil.

References

- [1] Bruno, M. G. S.: Joint Detection and Tracking of Moving Targets in Clutter. Ph.D. Thesis, ECE Department, Carnegie Mellon University, 1998.
- [2] Bar-Shalom, Y.; Li, X.: Multitarget-Multisensor Tracking: Principles and Techniques. YBS, Storrs, CT, 1995.
- [3] Bethel, R. E.; Paras, G. J.: A PDF multisensor, multitarget tracker. IEEE Transactions on Aerospace and Electronic Systems vol.34, n.1 (January 1998), 153-168.
- [4] Bethel, R. E.; Paras, G. J.: A PDF multitarget tracker. IEEE Transactions on Aerospace and Electronic Systems vol.30, n.2 (April 1994), 386-403.
- [5] Im, H.; Kim, T.: Optimization of multiframe target detection schemes. IEEE Transactions on Aerospace and Electronic Systems vol.35, n.1 (January 1999), 176-187.
- [6] Sangston, K. J.; Gini, F.; Grego, M. V.; Farina, A.: Structures for radar detection in compound Gaussian clutter. IEEE Transactions on Aerospace and Electronic Systems vol.35, n.2 (April 1999), 445-458.
- [7] Bruno, M. G. S.; Moura, J. M. F.: Performance of the optimal nonlinear detector/tracker in clutter. Proc. IEEE Int. Conf. Acoust. Speech, Sig. Process., Phoenix 1999, 1193-1196.
- [8] Bruno, M. G. S.; Moura, J. M. F.: Multiframe detection and tracking in clutter: optimal performance. Submitted to IEEE Transactions on Aerospace and Electronic Systems, under review.
- [9] Moura, J. M. F.; Balram, N. K.: Statistical algorithms for non-causal Gauss-Markov random fields. Handbook of Statistics, vol. 10, chapter 15, edited by N. K. Bose and C. R. Rao, Elsevier Science Publishers, 1993.
- [10] Moura, J. M. F.; Balram, N. K.: Recursive Structure of Non-causal Gauss Markov Random Fields. IEEE Transactions on Information Theory vol. IT-38.2 (March 1992), 334-354.
- [11] Moura, J. M. F.; Bruno, M. G. S.: DCT/DST and Gauss-Markov Fields: Conditions for Equivalence. IEEE Transactions on Signal Processing vol 46, n.9 (September 1998), 2571-2574.
- [12] Van Trees, H. L.: Detection, Estimation and Modulation Theory, Part I. John Wiley & Sons, 1968.
- [13] Graham, A.: Kronecker Products and Matrix Calculus with Applications. Ellis Horwood Series in Mathematics and its Applications, Ellis Horwood Limited, 1981.

Marcelo G.S. Bruno was born in 1970 in São Paulo, Brazil. He received the bachelor's and master's degree in Electrical Engineering from the University of São Paulo (USP), respectively in 1991 and 1993. He received the Ph.D. degree in Electrical and Computer Engineering from Carnegie Mellon University (CMU), Pittsburgh PA, U.S.A, in 1998.

Dr. Bruno is currently with the Communications and Signals Laboratory (LCS), Electrical Engineering Department, University of São Paulo, where he is a post-doctoral research fellow. In the summer of 1999, he was a visiting research scholar at Carnegie Mellon University. His research interests are primarily on statistical signal processing, particularly optimal detection in non-Gaussian noise, Bayesian and non-Bayesian estimation of multiple parameters, array processing, nonlinear stochastic filtering, and stochastic modeling for optimal image/video processing. Recent research interests also include multiresolution analysis and signal representation theory.

José M. F. Moura received the Engenheiro Eletrotécnico degree in 1969 from Instituto Superior Técnico (IST), Lisbon, Portugal, and the M.Sc., E.E., and D.Sc. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 1973 and 1975, respectively.

He is presently a Professor of Electrical Engineering (visiting) at MIT, on sabbatical leave from Carnegie Mellon University (CMU) where he is Professor of Electrical and Computer Engineering since 1986. Prior to this, he was on the faculty of IST, where he was an Assistant Professor in 1975, Professor Agregado in 1978, and Professor Catedrático in 1979. He has had prior visiting appointments at several institutions, including MIT (Genrad Associate Professor of Electrical Engineering and Computer Science from 1984 to 1986) and the University of Southern California (Research Scholar, Department of Aerospace Engineering, summers between 1978 and 1981). His research interests include statistical signal processing (one- and two-dimensional), digital communications, image and video processing, radar and sonar, and multiresolution techniques. He has more than 190 published technical contributions and is coeditor of two books.

Dr. Moura is Vice-President elect for Publications of the IEEE Signal Processing Society (SPS), the Editor-in-Chief for the IEEE Transactions on Signal Processing, a member of the Board of Governors of SPS, and a member of the SPS Publications Board. He is a member of the Sensor Array and Multichannel Technical Committee and of the Multimedia Signal Processing Technical Committee of the IEEE Signal Processing Society. He was a member of the IEEE Press Board from 1991 to 1995, a Technical Associate Editor for the IEEE Signal Processing Letters from 1993 to 1995, and an Associate Editor for the IEEE Transactions on Signal Processing from 1988 to 1992. He was a program committee member for the IEEE International Conference on Image Processing (ICIP'95) and for the IEEE International Symposium on Information Theory (ISIT'93).

Dr. Moura is a Fellow of the IEEE and a corresponding member of the Academy of Sciences of Portugal (Section of Sciences). He is affiliated with several IEEE Societies, Sigma Xi, AMS, IMS, and SIAM.