Chapter 12

Information-theoretic approach to fully adaptive radar resource management

Kristine Bell,[1] Chris Kreucher,[2] and Muralidhar Rangaswamy[3]

[1]*Metron, Inc., Reston, VA, USA*
[2]*KBR Wyle Services, LLC, Ann Arbor, MI, USA*
[3]*Air Force Research Laboratory, Dayton, OH, USA*

*Corresponding Author: Kristine Bell; kristine.bell@ieee.org

Abstract: The recent emergence of agile software-controlled waveform generation provides an opportunity to dramatically improve radar system performance. This chapter describes an integrated algorithm for estimating the state of a surveillance region and using this estimate to design future radar transmissions. Our fully adaptive radar resource management approach emulates the perception-action cycle of cognition by first constructing a probabilistic estimate of the surveillance region and then using an information-theoretic objective function to select and transmit the collection of waveforms that is expected to maximally improve this estimate. We illustrate our approach in simulation using a model of an agile multimode radar which is charged with tracking and classifying multiple target aircraft.

Keywords: fully adaptive radar, cognitive radar, information theory, resource management, tracking, classification

12.1. Introduction

The arrival of agile software-controlled waveform generation and selection in modern radar systems enables new algorithms which select radar transmission parameters to optimize performance. In this chapter, we describe and illustrate via simulation an information-theoretic approach to fully adaptive radar (FAR) resource management (RM) which exploits this new flexibility. FAR, sometimes referred to as "cognitive radar" is a closed-loop process which includes the steps of sensing the environment and responding by selecting and executing the next action(s) for the radar. It takes inspiration from the neurobiological processes of animal cognition to allow radar systems to adapt to their environments [1, 2, 3, 4, 5]. Neuropsychologists believe that animal cognition relies on dynamic feedback loops known as perception-action cycles (PACs) to understand and respond to sensory information [6].

Past work in radar resource management (RRM) which has applied perception-action approaches includes [7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18]. While recent research in this area refers to the techniques as "cognitive radar resource management" [15, 16, 17, 18], older related work simply uses "sensor management" and/or "resource allocation" [19, 20, 8, 9, 10, 11, 12, 13, 14]. Regardless of the name, the algorithms apply a two-step process of perception and action selection. First, the algorithms capture (perceive) the surveillance region using a probabilistic estimate of its state. In the multi-target detection, tracking, and identification context we focus on, this includes estimating probability

density functions (PDFs) on target states for each target and probability mass functions (PMFs) on target class for each target. Second, this probabilistic description is used to drive the selection of future sensing actions through an optimization process where the available actions are considered and the action set that maximizes future utility is selected.

The primary challenge of RRM is to balance multiple competing objectives in the presence of constraints on resource usage, such as time and power. In our use case, a single radar system is to perform multi-target track initialization, tracking, and classification. It is given a finite time window and must balance the demands of all the tasks using the available time. This balance is achieved through a single global objective function which calculates the total utility of a set of waveform selections over the time window. This objective function is numerically optimized, subject to the time constraint, to select the future actions the radar is to take.

Developing a single global objective function that captures the goals of the system in a mathematical form amenable to optimization is critical for successful implementation of a fully adaptive radar resource management (FARRM) system. As the number of tasks and actions available for adaptation increases, the complexity of the optimization becomes increasingly more difficult. Broadly speaking, there are two main approaches to defining this objective function: task-driven [18, 21] and information-theoretic [8]. Some work has appeared which has compared these approaches [10, 22].

The task-driven approach first enumerates specific task-level components of the objective function, such as the root mean square tracking error (RMSE) and correct classification percentage. A user establishes the goal for each of these tasks and defines the utility of being above or below that goal. Then,

a multi-objective optimization problem is solved. One method is to define a global task-driven objective function as the weighted sum of the utility of the individual tasks (see [23] for a discussion of multi-objective optimization methods for fully adaptive radar). The main benefits of this approach are that the objective function is tied to explicit target-level goals and that it directly lays out the relative importance of the various tasks. However, specifying task requirements, costs, and the relative per-target and per-task weightings over the incommensurate task-level measures requires significant domain knowledge [18, 24, 23].

In contrast, the information-theoretic approach employs a global measure of information as its objective function. A number of different measures of information have been studied, including mutual information (MI), entropy, Kullback-Leibler divergence (KLD), and Rényi (alpha) divergence [8, 25, 26, 27, 28, 29]. The main benefit of an information-theoretic FARRM objective function is that it implicitly balances the different types of information a radar may acquire – e.g., information about target position and information about target classification are both quantified using a common measure of information [11]. However, an information-theoretic approach does not explicitly optimize task metrics such as RMSE. As such, the information-theoretic objective functions can be difficult for the end-user to understand and connect to specific operational goals [30]. While the information-theoretic approach does allow for separate weightings to give different targets different priorities, and separate weightings to give different types of information different priorities [31], this would require an operator to specify a collection of weights.

The work in this chapter combines and extends our previous work in sensor management [8, 9, 10, 11, 12, 13, 14] and FAR [17, 23, 32, 33, 21, 34, 22, 35]. It is

distinguished from our previous efforts in FARRM [22, 35] in that we develop a more comprehensive approach to information-theoretic methods of scheduling, develop a more general and flexible multiple task tracking and classification model, and provide new simulation examples which use real, collected aircraft trajectories and a tracking and classification scenario where targets enter the surveillance region midway through the vignette.

This chapter proceeds as follows. First, Section 12.2 describes our multiple task FARRM system model, which is based on a stochastic optimization [36] approach and builds on earlier work in [17, 32]. We describe the two main components of the FARRM PAC - the perceptual and executive processors. The perceptual processor is a recursive Bayesian approach to target state estimation, while the executive processor employs an information-theoretic utility function. The FARRM framework itself is agnostic to the specific objective function and has been illustrated with task-based [17], information-based [22], and even hybrid [35] objective function definitions. In Section 12.3 we develop our information-theoretic utility function, which is based on the MI. In Sections 12.4 and 12.5, we provide the models for the tracking and classification tasks, respectively, including explicit descriptions of the perceptual and executive processors. In Section 12.6 we provide airborne radar simulation results using collected aircraft trajectories where we jointly schedule waveforms for multi-target classification, tracking, and track initialization subject to time constraints. Finally, Section 12.7 presents the conclusions.

12.2. FARRM System Model

The FARRM system model is shown in Figure 12.1. It emulates the PAC through interaction between the perceptual and executive processors. The aim of radar scheduling is to select the radar actions (e.g., parameters of the transmitted waveforms), $\boldsymbol{\theta}_k$, to best estimate a system state $\boldsymbol{x}_k$ to achieve the mission of the radar system. At time k, the perceptual processor receives information about the environment (including targets, clutter, noise, and other environmental parameters) in the form of a measurement $\boldsymbol{z}_k$ and uses the observation to refine its perception of the environment. Here, the perception is captured by the posterior density $f(\boldsymbol{x}_k|\boldsymbol{Z}_k;\boldsymbol{\Theta}_k)$, where $\boldsymbol{Z}_k \doteq \{\boldsymbol{z}_1, \cdots, \boldsymbol{z}_k\}$ is the history of measurements up to and including time k and $\boldsymbol{\Theta}_k \doteq \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \cdots, \boldsymbol{\theta}_k\}$ is the history of action vectors. The executive processor is then responsible for the action selection portion of the PAC, i.e., choosing the next value(s) of $\boldsymbol{\theta}_k$ in response to the latest perception. The new $\boldsymbol{\theta}_k$ is then provided to the radar sensor, which takes the action, forms the new measurement $\boldsymbol{z}_k$, and passes the measurement back to the perceptual processor. This completes one iteration of the PAC.

12.2.1. Multiple Task PAC

We now develop the explicit mathematical model of the multiple task PAC. Our model problem is of an agile multi-mode radar which must select radar sensing parameters for a set of M tracking and classification tasks.

Let the resource management frame length be T_F seconds, and let k denote the frame index (time step). We assume there are M dwells in each frame, corresponding to the M tasks. Each of the M dwells may be assigned a different

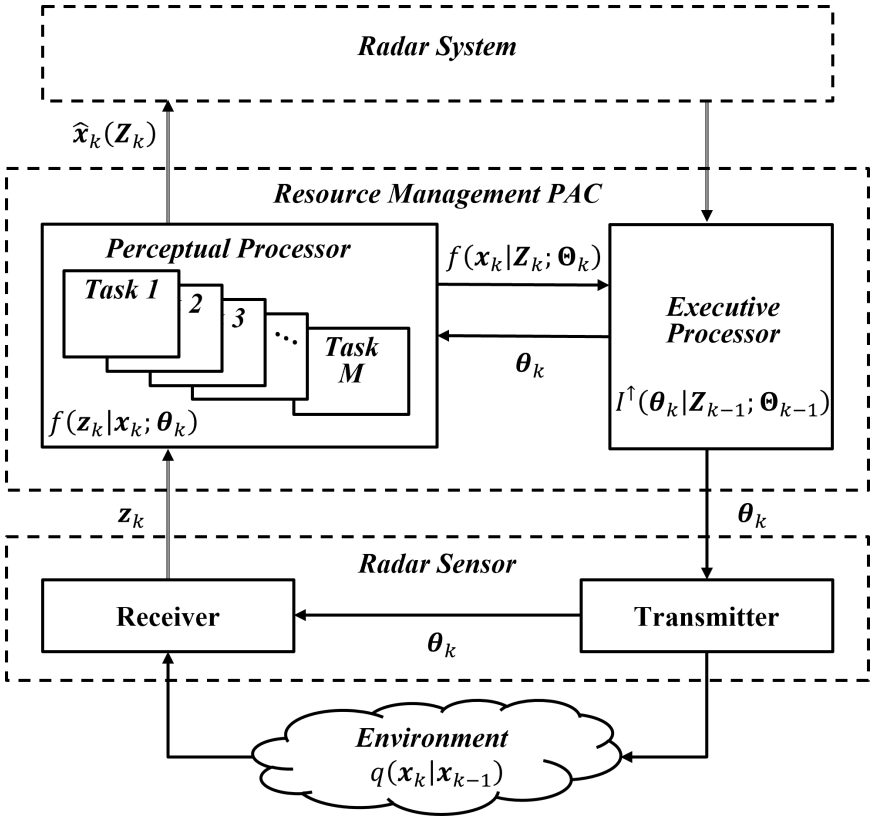

Figure 12.1: FARRM system model

amount of time at each time step. M may change as tasks come and go, for example when a target enters or exits the surveillance region.

For each of the M task dwells, the radar may elect to transmit nothing (taking up no time and providing no utility) or one of L waveforms from a library of waveforms. Let $a_l; l = 0, \cdots, L$ denote each of the possible actions (waveforms), where a_0 indicates no waveform is transmitted, and $\mathcal{A} = \{a_0, \cdots, a_L\}$ denotes the set of actions. Each waveform has a time duration, equal to the coherent processing interval (CPI), as well as specific parameters, such as the bandwidth and pulse repetition frequency (PRF), which characterize the waveform and its utility. When a waveform is chosen, it is fixed for the entire task

dwell/CPI.

The radar resource parameter vector, or action vector, for the kth frame is defined as the $M \times 1$ vector

$$\boldsymbol{\theta}_k = \begin{bmatrix} \theta_{1,k}, & \theta_{2,k}, & \cdots, & \theta_{M,k} \end{bmatrix}^{\mathrm{T}}, \tag{12.1}$$

where $\theta_{m,k}$ is the action chosen from $\mathcal{A}$ for the mth task dwell of the kth frame.

Following the approach in [37, 11], the state vector has the form:

$$\boldsymbol{x}_k = \begin{bmatrix} \boldsymbol{x}_{1,k}^{\mathrm{T}}, & \boldsymbol{x}_{2,k}^{\mathrm{T}}, & \cdots, & \boldsymbol{x}_{M,k}^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}, \tag{12.2}$$

where each $\boldsymbol{x}_{m,k}, m \in \{1, \cdots, M\}$ is a single task state vector containing components for the particular task of interest. For tracking tasks, we use a tracking state vector which consists of kinematic variables (position, velocity, and acceleration) as well as the received signal-to-noise ratio (SNR). For classification tasks, the state vector is a scalar classification variable. The target class is a discrete random variable whose value is taken from a discrete, finite set of values, while the tracking state variables are continuous random variables. Therefore, the state of the surveillance region is described using a combination of a PDF for the continuous components and a PMF for the discrete components.

Given the action vector $\boldsymbol{\theta}_k$, the radar sensor transmits the selected waveforms, receives returns, and generates a measurement vector $\boldsymbol{z}_k$, which has the form:

$$\boldsymbol{z}_k = \begin{bmatrix} \boldsymbol{z}_{1,k}^{\mathrm{T}}, & \boldsymbol{z}_{2,k}^{\mathrm{T}}, & \cdots, & \boldsymbol{z}_{M,k}^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}, \tag{12.3}$$

where $\boldsymbol{z}_{m,k}$ is the measurement vector for the mth task during the kth frame. It may be a vector, a scalar, or empty.

The system state $\boldsymbol{x}_k$ evolves according to a state transition model, which we assume to be a first-order Markov model with initial target state PDF/PMF $q(\boldsymbol{x}_0)$ and a state transition PDF/PMF $q(\boldsymbol{x}_k|\boldsymbol{x}_{k-1})$. The measurement model is characterized by the conditional PDF/PMF, or likelihood function, $f(\boldsymbol{z}_k|\boldsymbol{x}_k;\boldsymbol{\theta}_k)$, which is a function of the system state vector $\boldsymbol{x}_k$ and the action vector $\boldsymbol{\theta}_k$.

The first element of the perception-action cycle is the perceptual processor. The goal of the perceptual processor is to combine known sensing actions and received measurements with the models of system state evolution and measurement statistics to construct a posterior PDF/PMF $f(\boldsymbol{x}_k|\boldsymbol{Z}_k;\boldsymbol{\Theta}_k)$ on the system state and a corresponding system state estimate $\hat{\boldsymbol{x}}_k(\boldsymbol{Z}_k)$.

For our first-order Markov motion model, the posterior PDF/PMF of $\boldsymbol{x}_k$ given $\boldsymbol{Z}_k$ can be obtained using the Bayes-Markov recursion:

$$\begin{aligned} f^+(\boldsymbol{x}_0) &= q(\boldsymbol{x}_0) \\ f^-(\boldsymbol{x}_k) \doteq f(\boldsymbol{x}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_k) &= \int q(\boldsymbol{x}_k|\boldsymbol{x}_{k-1}) f^+(\boldsymbol{x}_{k-1}) d\boldsymbol{x}_{k-1} \\ f^-(\boldsymbol{z}_k) \doteq f(\boldsymbol{z}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_k) &= \int f(\boldsymbol{z}_k|\boldsymbol{x}_k;\boldsymbol{\theta}_k) f^-(\boldsymbol{x}_k) d\boldsymbol{x}_k \\ f^+(\boldsymbol{x}_k) \doteq f(\boldsymbol{x}_k|\boldsymbol{Z}_k;\boldsymbol{\Theta}_k) &= \frac{f(\boldsymbol{z}_k|\boldsymbol{x}_k;\boldsymbol{\theta}_k) f^-(\boldsymbol{x}_k)}{f^-(\boldsymbol{z}_k)}, \end{aligned} \tag{12.4}$$

where $f^-(\boldsymbol{x}_k)$ is the predicted PDF/PMF obtained from the motion update step and $f^+(\boldsymbol{x}_k)$ is the posterior PDF/PMF obtained from the information update step.

The state estimation performance is characterized by the posterior Bayes risk, which is the expected value of a perceptual processor error function, $\epsilon(\hat{\boldsymbol{x}}(\boldsymbol{Z}_k), \boldsymbol{x}_k)$, with respect to the posterior PDF/PMF,

$$R^+(\boldsymbol{Z}_k;\boldsymbol{\Theta}_k) = \mathbb{E}_k^+ \left\{ \epsilon(\hat{\boldsymbol{x}}(\boldsymbol{Z}_k), \boldsymbol{x}_k) \right\}, \tag{12.5}$$

where $\mathbb{E}_k^+ \{\cdot\}$ denotes expectation with respect to $f^+(\boldsymbol{x}_k)$.[1] The state estimate is found by minimizing the posterior Bayes risk:

$$\hat{\boldsymbol{x}}_k(\boldsymbol{Z}_k) = \underset{\hat{\boldsymbol{x}}(\boldsymbol{Z}_k)}{\arg\min} \; R^+(\boldsymbol{Z}_k; \boldsymbol{\Theta}_k). \tag{12.6}$$

The second element of the perception-action cycle is the executive processor. The function of the executive processor is to select future sensing actions to optimize the performance of the state estimator that will include both the future and past observations. For single time-step scheduling, we define the joint conditional PDF/PMF of $\boldsymbol{x}_k$ and $\boldsymbol{z}_k$ conditioned on $\boldsymbol{Z}_{k-1}$ as:

$$f^{\uparrow}(\boldsymbol{x}_k, \boldsymbol{z}_k) \doteq f(\boldsymbol{x}_k, \boldsymbol{z}_k | \boldsymbol{Z}_{k-1}; \boldsymbol{\Theta}_k) = f(\boldsymbol{z}_k | \boldsymbol{x}_k; \boldsymbol{\theta}_k) f\left(\boldsymbol{x}_k | \boldsymbol{Z}_{k-1}; \boldsymbol{\Theta}_k\right). \tag{12.7}$$

From (12.4), the joint conditional PDF can also be expressed as:

$$f^{\uparrow}(\boldsymbol{x}_k, \boldsymbol{z}_k) = f^+(\boldsymbol{x}_k) f^-\left(\boldsymbol{z}_k\right). \tag{12.8}$$

If we take the expectation of the error function with respect to the joint conditional PDF, we obtain the predicted conditional (PC)-Bayes risk as

$$R^{\uparrow}(\boldsymbol{\theta}_k | \boldsymbol{Z}_{k-1}; \boldsymbol{\Theta}_{k-1}) = \mathbb{E}_k^{\uparrow}\{\epsilon(\hat{\boldsymbol{x}}(\boldsymbol{Z}_k), \boldsymbol{x}_k)\}, \tag{12.9}$$

where $\mathbb{E}_k^{\uparrow}\{\cdot\}$ denotes expectation with respect to $f^{\uparrow}(\boldsymbol{x}_k, \boldsymbol{z}_k)$. Using (12.5) and (12.8), the PC-Bayes risk can also be expressed as the expectation of the pos-

[1] The most commonly used error functions are the squared error, $(\hat{\boldsymbol{x}}(\boldsymbol{Z}_k) - \boldsymbol{x}_k)^2$, for continuous random variables, and the indicator function, $\mathbf{1}_{\boldsymbol{x}_k}(\hat{\boldsymbol{x}}(\boldsymbol{Z}_k))$, for discrete random variables. The resulting Bayes risks are the mean square error and probability of correct decision, respectively.

terior Bayes risk with respect to $f^-(\boldsymbol{z}_k)$, i.e.,

$$R^{\uparrow}(\boldsymbol{\theta}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1}) = \mathbb{E}^-_{\boldsymbol{z}_k}\{R^+(\boldsymbol{Z}_k;\boldsymbol{\Theta}_k)\}, \tag{12.10}$$

where $\mathbb{E}^-_{\boldsymbol{z}_k}\{\cdot\}$ denotes expectation with respect to $f^-(\boldsymbol{z}_k)$.

Ideally, the executive processor would choose the next action vector to minimize the PC-Bayes risk. However, the PC-Bayes risk does not generally have a closed-form analytical expression. As a substitute, we use an information-theoretic measure of utility that is analytically tractable and reflects the quality of the target state estimate through the amount of information it contains, which we denote as $I^{\uparrow}(\boldsymbol{\theta}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1})$. The executive processor optimization problem is then to select the next action vector to maximize the information,

$$\boldsymbol{\theta}_k = \arg\max_{\boldsymbol{\theta}} \; I^{\uparrow}(\boldsymbol{\theta}|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1}). \tag{12.11}$$

12.2.2. Independent Tasks

We assume that the task transition models are independent across tasks and that the measurements in each dwell are independent. Under these assumptions, the prior and transition PDF/PMFs and the likelihood function factor into a product of individual task terms, expressed as:

$$\begin{aligned} q(\boldsymbol{x}_0) &= \prod_{m=1}^{M} q(\boldsymbol{x}_{m,0}), \\ q(\boldsymbol{x}_k|\boldsymbol{x}_{k-1}) &= \prod_{m=1}^{M} q(\boldsymbol{x}_{m,k}|\boldsymbol{x}_{m,k-1}), \\ f(\boldsymbol{z}_k|\boldsymbol{x}_k;\boldsymbol{\theta}_k) &= \prod_{m=1}^{M} f(\boldsymbol{z}_{m,k}|\boldsymbol{x}_{m,k};\theta_{m,k}). \end{aligned} \tag{12.12}$$

With this factorization, the joint Bayes-Markov recursion in (12.4) decouples into M individual recursions, and the posterior PDF/PMF as well as the joint conditional PDF/PMF have the factored forms:

$$\begin{aligned} f\left(\boldsymbol{x}_k|\boldsymbol{Z}_k;\boldsymbol{\Theta}_k\right) &= \prod_{m=1}^{M} f(\boldsymbol{x}_{m,k}|\boldsymbol{Z}_{m,k};\boldsymbol{\Theta}_{m,k}) \\ f(\boldsymbol{x}_k,\boldsymbol{z}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_k) &= \prod_{m=1}^{M} f(\boldsymbol{x}_{m,k},\boldsymbol{z}_{m,k}|\boldsymbol{Z}_{m,k-1};\boldsymbol{\Theta}_{m,k}), \end{aligned} \tag{12.13}$$

where $\boldsymbol{Z}_{m,k} \doteq \{\boldsymbol{z}_{m,1},\cdots,\boldsymbol{z}_{m,k}\}$ and $\boldsymbol{\Theta}_{m,k} \doteq \{\theta_{m,1},\theta_{m,2},\cdots,\theta_{m,k}\}$ are the history of measurements and actions, respectively, for the mth task.

We further assume that the error function $\epsilon\left(\hat{\boldsymbol{x}}(\boldsymbol{Z}_k),\boldsymbol{x}_k\right)$ is the sum of task error functions,[2]

$$\epsilon\left(\hat{\boldsymbol{x}}(\boldsymbol{Z}_k),\boldsymbol{x}_k\right) = \sum_{m=1}^{M} \epsilon\left(\hat{\boldsymbol{x}}_m(\boldsymbol{Z}_k),\boldsymbol{x}_{m,k}\right). \tag{12.14}$$

Then the Bayes risk is the sum of task Bayes risks,

$$R^+(\boldsymbol{Z}_k;\boldsymbol{\Theta}_k) = \sum_{m=1}^{M} R^+(\boldsymbol{Z}_{m,k};\boldsymbol{\Theta}_{m,k}), \tag{12.15}$$

and the task state estimates can be obtained separately by minimizing the task Bayes risk:

$$\hat{\boldsymbol{x}}_{m,k}(\boldsymbol{Z}_{m,k}) = \underset{\hat{\boldsymbol{x}}_m(\boldsymbol{Z}_{m,k})}{\arg\min}\ R^+(\boldsymbol{Z}_{m,k};\boldsymbol{\Theta}_{m,k}). \tag{12.16}$$

[2]The task error functions may be weighted to account for different priorities among the tasks.

The PC-Bayes risk is also the sum of task PC-Bayes risks,

$$R^{\uparrow}(\boldsymbol{\theta}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1}) = \sum_{m=1}^{M} R^{\uparrow}(\theta_{m,k}|\boldsymbol{Z}_{m,k-1};\boldsymbol{\Theta}_{m,k-1}), \tag{12.17}$$

and the information-theoretic utility function (if chosen as described in the next section) is the sum of task information functions,

$$I^{\uparrow}(\boldsymbol{\theta}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1}) = \sum_{m=1}^{M} I_m^{\uparrow}(\theta_{m,k}|\boldsymbol{Z}_{m,k-1};\boldsymbol{\Theta}_{m,k-1}). \tag{12.18}$$

The executive processor still optimizes the joint information-theoretic utility function to determine the next set of actions. It should be emphasized that the optimization is in a global sense and the next action may not be the optimal solution for a particular radar task in isolation.

12.3. Information-theoretic Utility Function

We now develop the executive processor information-theoretic utility function.

12.3.1. Definitions

We begin by defining some basic continuous and discrete random variables and their PDFs and PMFs, then provide a review of the definitions of various information-theoretic quantities and some relationships between them. The material is taken from [25] and [26].

Let ξ denote a discrete random variable that takes on one of N_ξ values

in the set $\mathcal{C} = \{\xi_1, \xi_2, \cdots, \xi_{N_\xi}\}$ and let $f(\xi)$ denote its PMF. We denote the components of $f(\xi)$ as $f_i(\xi) = P(\xi = \xi_i); i = 1, \cdots, N_\xi$. Let d denote a discrete random variable that takes on one of N_d values in the set $\mathcal{D} = \{d_1, d_2, \cdots, d_{N_d}\}$ and let $f(d)$ denote its PMF. Let $f(\xi|d)$ denote the conditional PMF of ξ given d. We denote the components of $f(\xi|d)$ as $f_{i|j}(\xi|d) = P(\xi = \xi_i|d = d_j); i = 1, \cdots, N_\xi; j = 1, \cdots, N_d$.

The KLD between two PMFs for the discrete random variable ξ is defined as:

$$D\left(f(\xi)||g(\xi)\right) \doteq \sum_{i=1}^{N_\xi} f_i(\xi) \ln \frac{f_i(\xi)}{g_i(\xi)}, \tag{12.19}$$

For a discrete random variable, the entropy characterizes the level of uncertainty or information. The entropy of ξ is defined as:

$$H_\xi \doteq -\sum_{i=1}^{N_\xi} f_i(\xi) \ln f_i(\xi). \tag{12.20}$$

The entropy has the property $0 \leq H_\xi \leq \ln(N_\xi)$. The entropy is low when the PMF is concentrated on one of the classes and high when the PMF is distributed across the classes. The maximum value is obtained when all classes have the same probability. The conditional entropy of ξ given d is defined as:

$$H_{\xi|d} \doteq -\sum_{j=1}^{N_d} f_j(d) \sum_{i=1}^{N_\xi} f_{i|j}(\xi|d) \ln f_{i|j}(\xi|d). \tag{12.21}$$

The MI between ξ and d is defined as:

$$I_{\xi d} \doteq \mathbb{E}_d\left\{D\left(f(\xi|d)||f(\xi)\right)\right\} = \sum_{j=1}^{N_d} f_j(d) \left\{ \sum_{i=1}^{N_\xi} f_{i|j}(\xi|d) \ln \frac{f_{i|j}(\xi|d)}{f_i(\xi)} \right\}, \tag{12.22}$$

thus the MI is the expected KLD between the conditional PMF of ξ given d

and the marginal PMF of ξ, where the expectation is with respect to d. There are many relationships between these quantities, including:

$$H_{\xi|d} = H_{\xi} - I_{\xi d}. \tag{12.23}$$

Let $\boldsymbol{x}$ and $\boldsymbol{z}$ denote vectors of continuous random variables with PDFs $f(\boldsymbol{x})$ and $f(\boldsymbol{z})$, respectively. Let $f(\boldsymbol{x}|\boldsymbol{z})$ denote the conditional PDF of $\boldsymbol{x}$ given $\boldsymbol{z}$. The KLD between two PDFs for the continuous random variable $\boldsymbol{x}$ is defined as:

$$D\left(f(\boldsymbol{x})||g(\boldsymbol{x})\right) \doteq \int f(\boldsymbol{x}) \ln \frac{f(\boldsymbol{x})}{g(\boldsymbol{x})} d\boldsymbol{x}. \tag{12.24}$$

For continuous random variables, the entropy of $\boldsymbol{x}$ is defined as:[3]

$$H_{\boldsymbol{x}} \doteq -\int f(\boldsymbol{x}) \ln f(\boldsymbol{x}) d\boldsymbol{x}, \tag{12.25}$$

and the conditional entropy of $\boldsymbol{x}$ given $\boldsymbol{z}$ is defined as:

$$H_{\boldsymbol{x}|\boldsymbol{z}} \doteq -\int f(\boldsymbol{z}) \left\{ \int f(\boldsymbol{x}|\boldsymbol{z}) \ln f(\boldsymbol{x}|\boldsymbol{z}) d\boldsymbol{x} \right\} d\boldsymbol{z}. \tag{12.26}$$

The MI between $\boldsymbol{x}$ and $\boldsymbol{z}$ is defined as:

$$I_{\boldsymbol{x}\boldsymbol{z}} \doteq \mathbb{E}_{\boldsymbol{z}} \left\{ D\left(f(\boldsymbol{x}|\boldsymbol{z})||f(\boldsymbol{x})\right) \right\} = \int f(\boldsymbol{z}) \left\{ \int f(\boldsymbol{x}|\boldsymbol{z}) \ln \frac{f(\boldsymbol{x}|\boldsymbol{z})}{f(\boldsymbol{x})} d\boldsymbol{x} \right\} d\boldsymbol{z} \tag{12.27}$$

[3]For continuous random variables, the standard terminology for the definitions in (12.25) and (12.26) are the differential entropy and conditional differential entropy, and the standard notations are $h_{\boldsymbol{x}}$ and $h_{\boldsymbol{x}|\boldsymbol{z}}$, respectively. However, since we are using a combination of continuous and discrete random variables, we follow the convention in [26] and use the same terminology (entropy) and notation ($H_{\boldsymbol{x}}$) for both continuous and discrete random variables.

and we have the following relationship:

$$H_{\boldsymbol{x}|\boldsymbol{z}} = H_{\boldsymbol{x}} - I_{\boldsymbol{x}\boldsymbol{z}}. \tag{12.28}$$

12.3.2. Objective Function

In Section 12.2.1, we showed that the expected task performance is characterized by the PC-Bayes risk, which is the expected value of the posterior Bayes risk over the predicted PDF of the next measurement $f^-(\boldsymbol{z}_k)$. Since the PC-Bayes risk is difficult to compute analytically, we seek an information-theoretic measure of utility that is easier to analyze and results in an executive processor optimization problem whose solution is the same as, or close to, the solution that would be obtained if the PC-Bayes risk was used.

The foundation of our approach is the observation that the entropy of the posterior distribution is a suitable information-theoretic surrogate for the posterior Bayes risk. It is low when the posterior distribution contains a lot of information about the state and good state estimates with low Bayes risk can be obtained. If we take the expected value of the posterior entropy with respect to $f^-(\boldsymbol{z}_k)$, we obtain the corresponding surrogate for the PC-Bayes risk, which turns out to be the conditional entropy, defined in (12.21) and (12.26).

The entropy of the predicted and posterior PDF/PMFs are defined as:

$$\begin{aligned} H_k^-(\boldsymbol{Z}_{k-1}; \boldsymbol{\Theta}_{k-1}) &= -\int f^-(\boldsymbol{x}_k) \ln f^-(\boldsymbol{x}_k) d\boldsymbol{x}_k, \\ H_k^+(\boldsymbol{z}_k | \boldsymbol{Z}_{k-1}; \boldsymbol{\theta}_k, \boldsymbol{\Theta}_{k-1}) &= -\int f^+(\boldsymbol{x}_k) \ln f^+(\boldsymbol{x}_k) d\boldsymbol{x}_k, \end{aligned} \tag{12.29}$$

and the conditional entropy is

$$H^{\uparrow}\left(\boldsymbol{\theta}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1}\right) = -\int f^{-}(\boldsymbol{z}_k) H_k^{+}\left(\boldsymbol{z}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\theta}_k,\boldsymbol{\Theta}_{k-1}\right) d\boldsymbol{z}_k. \tag{12.30}$$

Using the property in (12.28), the conditional entropy can be written as

$$H^{\uparrow}\left(\boldsymbol{\theta}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1}\right) = H_k^{-} - I_{\boldsymbol{xz}}(\boldsymbol{\theta}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1}). \tag{12.31}$$

where

$$I_{\boldsymbol{xz}}(\boldsymbol{\theta}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1}) \doteq \mathbb{E}_{\boldsymbol{z}_k}^{-}\left\{\int f^{+}(\boldsymbol{x}_k)\ln\frac{f^{+}(\boldsymbol{x}_k)}{f^{-}(\boldsymbol{x}_k)}d\boldsymbol{x}_k\right\} \tag{12.32}$$

is the MI (or expected KLD) between the target state and the next measurement, defined in (12.22) and (12.27).

Since H_k^{-} in (12.31) is not a function of the action vector $\boldsymbol{\theta}_k$, the MI is an equivalent information-theoretic utility function to the conditional entropy, which can be *maximized* to obtain the next action vector.

The MI is a measure of the expected *gain* in information [8, 10, 13, 16, 29]. For our independent task model, the global MI decomposes into the sum of the task MIs as:

$$I_{\boldsymbol{xz}}(\boldsymbol{\theta}_k|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1}) = \sum_{m=1}^{M} I_{\boldsymbol{xz};m}(\theta_{m,k}|\boldsymbol{Z}_{m,k-1};\boldsymbol{\Theta}_{m,k-1}). \tag{12.33}$$

To summarize, our information-theoretic method for selecting the next action vector is to find the action vector that maximizes the MI,

$$\boldsymbol{\theta}_k = \arg\max_{\boldsymbol{\theta}}\ I_{\boldsymbol{xz}}(\boldsymbol{\theta}|\boldsymbol{Z}_{k-1};\boldsymbol{\Theta}_{k-1}). \tag{12.34}$$

12.4. Tracking Task

For a tracking task, the tracking state vector consists of kinematic variables (position, velocity, and possibly acceleration) and the received SNR. For three-dimensional (3D) tracking, the state vector has the form:

$$\boldsymbol{x}_{m,k} = \left[x_{m,k}, \dot{x}_{m,k}, \ddot{x}_{m,k}, y_{m,k}, \dot{y}_{m,k}, \ddot{y}_{m,k}, z, \dot{z}_{m,k}, \ddot{z}_{m,k}, s_{m,k}\right]^{\mathrm{T}}. \tag{12.35}$$

We assume an initial tracking state distribution that is multivariate Gaussian with mean $\boldsymbol{\mu}_{m,0}$ and covariance matrix $\boldsymbol{\Sigma}_{m,0}$, which we denote as $q(\boldsymbol{x}_{m,0}) = \mathcal{N}\left(\boldsymbol{x}_{m,0}; \boldsymbol{\mu}_{m,0}, \boldsymbol{\Sigma}_{m,0}\right)$. We assume a linear motion model of the form:

$$\boldsymbol{x}_{m,k} = \boldsymbol{F}_m \boldsymbol{x}_{m,k-1} + \boldsymbol{e}_{m,k}, \tag{12.36}$$

where $\boldsymbol{F}_m$ is the state transition matrix and $\boldsymbol{e}_{m,k}$ is zero-mean additive white Gaussian noise (AWGN) with covariance matrix $\boldsymbol{Q}_m$. The transition PDF is then $q(\boldsymbol{x}_{m,k}|\boldsymbol{x}_{m,k-1}) = \mathcal{N}\left(\boldsymbol{x}_{m,k}; \boldsymbol{F}_m \boldsymbol{x}_{m,k-1}, \boldsymbol{Q}_m\right)$.

We assume that when new targets enter the surveillance region, they are detected by a separate surveillance process and provided to the system with an initial tracking state distribution. Targets may arrive at the onset of the simulation or part way through.

We assume measurements are received with detection probability $P_D\left(\boldsymbol{x}_{m,k}; \theta_{m,k}\right)$ and false alarm probability P_F. The detection probability is determined by the detection threshold, which is set to achieve a desired P_F, and the received SNR, which is a function of the sensor parameters and target state. When measurements are received, we assume they follow a nonlinear, AWGN measurement model of the form:

$$\boldsymbol{z}_{m,k} = \boldsymbol{h}_{m,k}\left(\boldsymbol{x}_{m,k}\right) + \boldsymbol{n}_{m,k}, \tag{12.37}$$

where $\boldsymbol{h}_{m,k}(\boldsymbol{x}_{m,k})$ is a nonlinear transformation from the target state space to the radar measurement space and $\boldsymbol{n}_{m,k}$ is the measurement error, which is modeled as a zero-mean Gaussian random vector with covariance matrix $\boldsymbol{R}_{m,k}(\theta_{m,k})$. The single target likelihood function is then

$$f(\boldsymbol{z}_{m,k}|\boldsymbol{x}_{m,k};\theta_{m,k}) = P_D(\boldsymbol{x}_{m,k};\theta_{m,k})\,\mathcal{N}(\boldsymbol{z}_{m,k};\boldsymbol{h}_{m,k}(\boldsymbol{x}_{m,k}),\boldsymbol{R}_{m,k}(\theta_{m,k})). \tag{12.38}$$

In the perceptual processor, we implement the Bayes-Markov recursion expressions in (12.4) using the extended Kalman filter (EKF). It requires evaluating the Jacobian matrix, which is defined as:

$$\tilde{\boldsymbol{H}}_{m,k}(\boldsymbol{x}) = \left[\nabla_{\boldsymbol{x}}\boldsymbol{h}_{m,k}^{\mathrm{T}}(\boldsymbol{x})\right]^{\mathrm{T}}. \tag{12.39}$$

The EKF recursion is as follows:

Initialization:

$$\begin{aligned}\boldsymbol{\mu}_{m,0}^{+} &= \boldsymbol{\mu}_{m,0}\\ \boldsymbol{P}_{m,0}^{+} &= \boldsymbol{\Sigma}_{m,0}\end{aligned} \tag{12.40}$$

Motion update:

$$\begin{aligned}\boldsymbol{\mu}_{m,k}^{-} &= \boldsymbol{F}_m\boldsymbol{\mu}_{m,k-1}^{+}\\ \boldsymbol{P}_{m,k}^{-} &= \boldsymbol{F}_m\boldsymbol{P}_{m,k-1}^{+}\boldsymbol{F}_m^{\mathrm{T}} + \boldsymbol{Q}_m\end{aligned} \tag{12.41}$$

Information update, when a measurement is received:

$$\begin{aligned}
\boldsymbol{H}_{m,k} &= \tilde{\boldsymbol{H}}_{m,k}(\boldsymbol{\mu}_{m,k}^{-}) \\
\boldsymbol{K}_{m,k} &= \boldsymbol{P}_{m,k}^{-}\boldsymbol{H}_{m,k}^{\mathrm{T}}\left[\boldsymbol{H}_{m,k}\boldsymbol{P}_{m,k}^{-}\boldsymbol{H}_{m,k}^{\mathrm{T}} + \boldsymbol{R}_{m,k}(\theta_{m,k})\right]^{-1} \\
\boldsymbol{\mu}_{m,k}^{+} &= \boldsymbol{\mu}_{m,k}^{-} + \boldsymbol{K}_{m,k}\left[\boldsymbol{z}_{m,k} - \boldsymbol{h}_{m,k}(\boldsymbol{\mu}_{m,k}^{-})\right] \\
\boldsymbol{P}_{m,k}^{+} &= \boldsymbol{P}_{m,k}^{-} - \boldsymbol{K}_{m,k}\boldsymbol{H}_{m,k}\boldsymbol{P}_{m,k}^{-}
\end{aligned} \tag{12.42}$$

Information update, when no measurement is received:

$$\begin{aligned}
\boldsymbol{\mu}_{m,k}^{+} &= \boldsymbol{\mu}_{m,k}^{-} \\
\boldsymbol{P}_{m,k}^{+} &= \boldsymbol{P}_{m,k}^{-}.
\end{aligned} \tag{12.43}$$

The posterior Bayes risk for the tracking state vector is the trace of the posterior mean square error (MSE) matrix. The solution to (12.16) is the mean of the posterior PDF:

$$\hat{\boldsymbol{x}}_{m,k}(\boldsymbol{Z}_{m,k}) = \mathbb{E}_{k}^{+}\left\{\boldsymbol{x}_{m,k}\right\} = \boldsymbol{\mu}_{m,k}^{+}. \tag{12.44}$$

In the executive processor, the MI has the form:

$$\begin{aligned}
&I_{\boldsymbol{xz};m}(\theta_{m,k}|\boldsymbol{Z}_{m,k-1};\boldsymbol{\Theta}_{m,k-1}) = \\
&\frac{P_D\left(\boldsymbol{x}_{m,k};\theta_{m,k}\right)}{2}\left(\ln|\boldsymbol{P}_{m,k}^{-}| + \ln\left|\left[\boldsymbol{P}_{m,k}^{-}\right]^{-1} + \boldsymbol{H}_{m,k}^{\mathrm{T}}\boldsymbol{R}_{m,k}(\theta_{m,k})^{-1}\boldsymbol{H}_{m,k}\right|\right),
\end{aligned} \tag{12.45}$$

where $|\cdot|$ denotes the determinant of a matrix.

The probability of detection $P_D\left(\boldsymbol{x}_{m,k};\theta_{m,k}\right)$ and the estimation covariance matrix $\boldsymbol{R}_{m,k}(\theta_{m,k})$ depend on the characteristics of the radar waveform, and generally have closed form analytical expressions. The probability of detection

also depends on the current state of the target $\boldsymbol{x}_{m,k}$. To evaluate the MI, we substitute the probability of detection evaluated at the predicted state estimate, $P_D\left(\boldsymbol{\mu}_{m,k}^{-};\theta_{m,k}\right)$. The remaining terms in the MI are available from the EKF recursion, so they are straightforward to evaluate.

12.5. Classification Task

For a classification task, the state vector is a scalar classification variable, $x_{m,k}$, which takes on one of a discrete set of N_c values in the set $\mathcal{C}$,

$$x_{m,k} \in \mathcal{C} = \{1, 2, \cdots, N_c\}. \tag{12.46}$$

The prior PMF $q(x_{m,0})$ is represented by the $N_c \times 1$ vector $\boldsymbol{q}_m$, which consists of the N_c probabilities

$$[\boldsymbol{q}_m]_i = P(x_{m,0} = i); \qquad i = 1, \cdots, N_c. \tag{12.47}$$

The transition model $q(x_{m,k}|x_{m,k-1})$ is represented by the $N_c \times N_c$ transition matrix $\boldsymbol{\Upsilon}_m$, where

$$[\boldsymbol{\Upsilon}_m]_{ij} = P(x_{m,k} = i|x_{m,k-1} = j); \qquad i, j = 1, \cdots, N_c. \tag{12.48}$$

Depending on the application, switching between classes may or may not be possible. For example, if the class represents a behavior of the target, then it is possible to change classes, however if the class represents a fixed property of the target such as the type of vehicle or aircraft, then the class cannot change and $\boldsymbol{\Upsilon}_m$ is equal to the identity matrix.

We assume that when new targets enter the surveillance region, they are detected by a separate surveillance process and provided to the system with an initial classification state distribution.

We assume that the radar data is processed to produce a discrete valued measurement of target class, i.e., $z_{m,k} \in \mathcal{C} = \{1, 2, \cdots, N_c\}$ and the likelihood function $f(z_{m,k}|x_{m,k}; \theta_{m,k})$ is represented by the $N_c \times N_c$ likelihood matrix $\boldsymbol{L}_m(\theta_{m,k})$, where

$$[\boldsymbol{L}_m]_{ij}(\theta_{m,k}) = P(z_{m,k} = i|x_{m,k} = j; \theta_{m,k}); \qquad i, j = 1, \cdots, N_c. \tag{12.49}$$

In the perceptual processor, the Bayes-Markov recursion can be computed exactly. Let $f_i^-(x_{m,k}) \doteq P(x_{m,k} = i|\boldsymbol{Z}_{m,k-1}; \boldsymbol{\Theta}_{m,k-1})$ denote the predicted PMF and $f_i^+(x_{m,k}) \doteq P(x_{m,k} = i|\boldsymbol{Z}_{m,k}; \boldsymbol{\Theta}_{m,k})$ denote the posterior PMF. The recursion is as follows:

$$\begin{aligned}
f_i^+(x_{m,0}) &= [\boldsymbol{q}_m]_i; \qquad i = 1, \cdots, N_c, \\
f_i^-(x_{m,k}) &= \sum_{j=1}^{N_c} [\boldsymbol{\Upsilon}_m]_{ij} f_j^+(x_{m,k-1}); \qquad i = 1, \cdots, N_c \\
f^-(z_{m,k}) &= \sum_{j=1}^{N_c} [\boldsymbol{L}_m]_{z_{m,k},j}(\theta_{m,k}) f_j^-(x_{m,k}) \\
f_i^+(x_{m,k}) &= \frac{[\boldsymbol{L}_m]_{z_{m,k},i}(\theta_{m,k}) f_i^-(x_{m,k})}{f^-(z_{m,k})}; \quad i = 1, \cdots, N_c.
\end{aligned} \tag{12.50}$$

The posterior Bayes risk for the classification state vector is the posterior probability of incorrect classification. The solution to (12.16) is the maximum of the posterior PMF:

$$\hat{x}_{m,k}(\boldsymbol{Z}_{m,k}) = \arg\max_{i \in \mathcal{C}} \; f_i^+(x_{m,k}). \tag{12.51}$$

In the executive processor, to compute the classification MI, we first define the prior probability of receiving classification measurement $z_{m,k}$ using the notation

$$f_j^-(z_{m,k}) \doteq P(z_{m,k} = j | \boldsymbol{Z}_{m,k-1}; \boldsymbol{\Theta}_{m,k}) \tag{12.52}$$

and the posterior probability of class $x_{m,k}$ after receiving observation $z_{m,k}$ using the notation

$$f_{i|j}^+(x_{m,k}|z_{m,k}) \doteq P(x_{m,k} = i | z_{m,k} = j, \boldsymbol{Z}_{m,k-1}; \boldsymbol{\Theta}_{m,k}). \tag{12.53}$$

These quantities are computed for every $z_{m,k}$ using the expressions in (12.50). Then the MI is computed directly from the definition in (12.22):

$$I_{xz;m}(\theta_{m,k}|\boldsymbol{Z}_{m,k-1}; \boldsymbol{\Theta}_{m,k-1}) = \sum_{j=1}^{N_c} f_j^-(z_{m,k}) \left\{ \sum_{i=1}^{N_c} f_{i|j}^+(x_{m,k}|z_{m,k}) \ln \frac{f_{i|j}^+(x_{m,k}|z_{m,k})}{f_j^-(z_{m,k})} \right\}. \tag{12.54}$$

12.6. Simulation Example

This section demonstrates our information-theoretic approach to fully adaptive radar resource management in a multitarget tracking, classification, and track initialization simulation.

12.6.1. Simulation Scenario

Our model problem consists of an agile multimode radar platform and three airborne targets as illustrated in Figure 12.2. The radar platform is charged with performing tracking and classification of the targets in the surveillance region,

which is indicated by the box. Targets 1 and 3 are always in the surveillance region, while Target 2 enters about one-third of the way into the 250-second simulation.

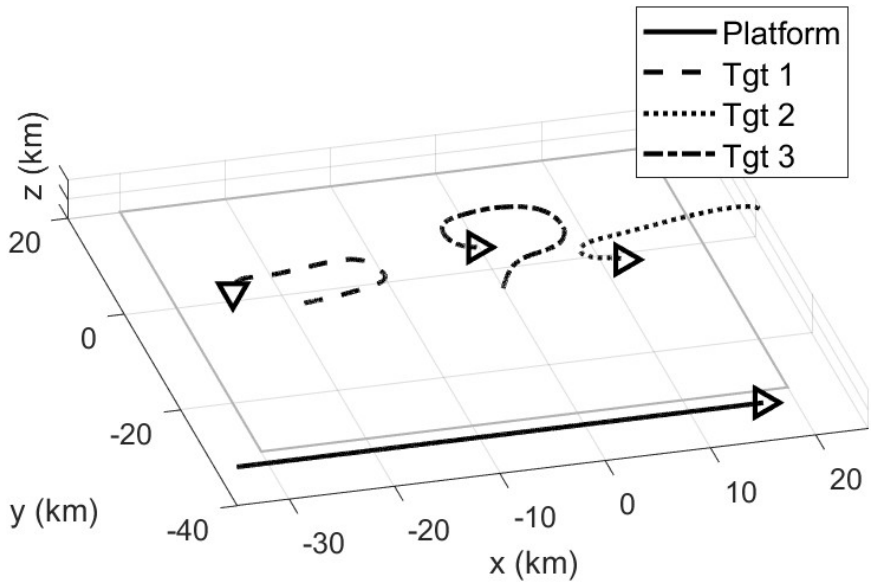

Figure 12.2: Our simulation has an airborne radar platform and three airborne targets. Target trajectories are derived from the TrajAir [38] collected aircraft dataset. The surveillance region is indicated by the solid box.

The radar platform flies at a constant heading with constant velocity of 200 m/s. The three targets are taken from the TrajAir [38] data collection, which is a set of recorded trajectories over the Pittsburgh-Butler Regional Airport in Pennsylvania. Since the TrajAir targets are general aviation aircraft operating over a narrow region, we spatially shifted and scaled the recorded flight data to construct speeds and locations consistent with maneuvering targets. Target 1 has a range between 40 km and 60 km to the radar platform. Target 2 has a range between 50 km (at entry into the surveillance region) and 35 km to the radar platform (at the end of the simulation). Target 3 has a relatively

constant range of approximately 45 km to the radar platform. All targets vary their speeds between 100 m/s and 300 m/s during the 250-second simulation in accordance with the (scaled) recorded data. The simulation has four tasks: track the three targets and identify (classify) Target 1.

12.6.2. Sensor and Task Models

The tracking state vector $\boldsymbol{x}_{m,k}$ consisting of the three-dimensional position, velocity, acceleration, and SNR as specified in (12.35). We define the SNR in decibels as $s_{m,k} = 10\log_{10}\zeta_{m,k}$, where $\zeta_{m,k}$ is the SNR in linear scale. The classification state variable $x_{m,k}$ is assumed to be one of $N_c = 5$ classes. For tracking, we use the Singer model [39] for target motion as the tracking transition model and for classification, we assume a transition matrix with diagonal entries $[\boldsymbol{\Upsilon}_m]_{ii} = 0.95$ and off-diagonal entries $[\boldsymbol{\Upsilon}_m]_{ij} = 0.0125$. Although the target class should not change, providing some probability for transition forces the radar to make periodic classification measurements and prevents the state estimate from getting stuck in an incorrect class due to an incorrect measurement.

We assume the radar transmits a waveform and receives returns through an antenna with a fixed azimuth beamwidth ($\Delta\phi$) and elevation beamwidth ($\Delta\theta$). The transmitted waveform is characterized by its center frequency (f_c), pulse bandwidth (B_p), PRF (f_p), and number of pulses (N_p). We model the tracking measurement process as providing detections which are estimates of target range (R), range-rate ($\dot{R}$), azimuth angle (ϕ), elevation angle (θ), and SNR in decibels ($s = 10\log_{10}\zeta$). For this simulation, we fix f_c, $\Delta\phi$, $\Delta\theta$ and assume that the radar can select B_p, f_p, and N_p from a collection of options. We

also assume that the detection threshold and the false alarm rate P_F are fixed. Let $B_{p;m,k}$, $f_{p;m,k}$, and $N_{p;m,k}$ denote the parameters of the selected waveform for the mth tracking task. The detection probability is given by [40]:

$$P_D(\zeta_{m,k};\theta_{m,k}) = Q_{MAR}\left(\sqrt{2N_{p;m,k}\zeta_{m,k}}, \sqrt{-2\ln P_F}\right), \tag{12.55}$$

where $Q_{MAR}(a,b)$ is the Marcum Q-function.

The estimation covariance matrix is a diagonal matrix whose components are [41, 42]:

$$\begin{aligned}
\left[\boldsymbol{R}_{m,k}(\theta_{m,k})\right]_R &= \left[2N_{p;m,k}\zeta_{m,k}\left(\frac{2}{c}\right)^2(3B_{p;m,k})^2\right]^{-1} \\
\left[\boldsymbol{R}_{m,k}(\theta_{m,k})\right]_{\dot{R}} &= \left[2N_{p;m,k}\zeta_{m,k}\left(\frac{4\pi f_c}{c}\right)^2\left(\frac{1}{12B_{p;m,k}^2}+\frac{(N_{p;m,k}^2-1)}{12f_{p;m,k}^2}\right)\right]^{-1} \\
\left[\boldsymbol{R}_{m,k}(\theta_{m,k})\right]_{\phi} &= \left[2N_{p;m,k}\zeta_{m,k}\left(\frac{1.782\pi}{\Delta\phi}\right)^2\right]^{-1} \\
\left[\boldsymbol{R}_{m,k}(\theta_{m,k})\right]_{\theta} &= \left[2N_{p;m,k}\zeta_{m,k}\left(\frac{1.782\pi}{\Delta\theta}\right)^2\right]^{-1} \\
\left[\boldsymbol{R}_{m,k}(\theta_{m,k})\right]_{s} &= \left(\frac{10}{\ln(10)}\right)^2,
\end{aligned} \tag{12.56}$$

where $c = 3\times 10^8$ m/s is the speed of light.

Classification measurements are characterized similarly. The radar is assumed to transmit and receive a CPI, process the data through a target recognition algorithm and return a discrete classification call. We characterize the performance of a classification waveform by the probability that the discrete classification call is correct, denoted by p_{dc}. Let $p_{dc;m,k}$ denote the value corresponding to the selected waveform for the mth task. The classification likelihood matrix has the form:

$$[\boldsymbol{L}_m]_{ij}(\theta_{m,k}) = \begin{cases} p_{dc;m,k} & i = j \\ \dfrac{1 - p_{dc;m,k}}{N_c - 1} & i \neq j. \end{cases} \tag{12.57}$$

12.6.3. Waveforms

We assume that the radar resource management frame length is the same as the track update interval, which is $T_F = 100$ ms. The 100-ms frame is used to perform both surveillance (which detects new targets) and service of existing targets by performing tracking and classification tasks. Our focus in this study is on scheduling the agile radar to best perform the latter of those two functions, tracking and classification of known targets.

As such, we assume that 90 ms of the 100-ms frame time is reserved for performing the surveillance function, leaving 10 ms to perform tracking and classification of the known targets. The surveillance function is assumed to be performed using some standard method (e.g., periodic scan with detection revisits and coarse initialization) and provides coarsely initialized targets with a large covariance which then become part of the tracking and classification tasks.

The resource management algorithm is free to perform a measurement for each task during the tracking and classification portion of the timeline or elect to measure any subset of the tasks as long as the total measurement time is less than or equal to the 10-ms time budget. The waveforms available for tracking and classification, their parameters, and dwell times are given in Table 12.1. Also included is the "no observation" waveform #0. This allows the scheduler to choose not to measure some of the tasks at the current frame, thereby leaving extra time for other tasks. The fixed tracking waveform parameters are $f_c = 3$

Table 12.1: Our airborne radar scheduler selects between classification and tracking waveforms. When selecting a tracking waveform, the scheduler chooses bandwidth, pulse count, and PRF. When selecting classification waveform, the scheduler chooses the classification accuracy.

Waveform	Parameter(s)			Dwell Time
0	N/A			0.0
#	B_p (MHz)	f_p (kHz)	N_p	T (ms)
1,2,3	2.5, 5, 7.5	20	25	1.25
4,5,6	2.5, 5, 7.5	10	25	2.50
7,8,9	2.5, 5, 7.5	10	50	5.00
#	p_{dc}			T (ms)
10	0.6			1.0
11	0.75			2.5

GHz, $\Delta\phi = 2°$, $\Delta\theta = 6°$, and $P_F = 10^{-6}$.

The set of waveforms allows the scheduler to perform trade-offs between receiving information on different tasks as well as receiving different types of information. For example, it may select a longer CPI (yielding higher P_D or p_{dc}) for one task at the expense of a shorter CPI (yielding lower P_D or p_{dc}) or no CPI for other tasks.

12.6.4. Waveform Optimization

The predicted utility of a sensing action is scored using the information-theoretic utility function in (12.33). The objective is to maximize the total information gain over the four tasks subject to the timing constraint. This is accomplished by selecting an information-optimal set of waveforms to transmit

for the collection of tasks at each time step.

The numerical waveform optimization proceeds as follows. We first compute the information-theoretic utility of each waveform for each task. The computational complexity is $O(LM)$, where L is the number of waveforms and M is the number of tasks. Next, we record the maximum utility waveform from among all waveforms having the same CPI for each task. This reduces the number of permutations under consideration by discarding any equal-CPI waveforms that have lower utility than another choice for the same task. As such, we do not explicitly enumerate the L^M combinations of waveforms that could be transmitted together. Finally, we apply a discrete double auction parameter selection (DDAPS) [43, 44] approach to find the set of M waveforms that maximize the total MI subject to the timing constraint.

12.6.5. Simulation Results

Simulation example results for a single trial are shown in Figures 12.3 and 12.4. Figure 12.3 shows the coarsest level description of results, indicating how much time (of the 10-ms budget) is allocated to each task at each of the time steps. Since each time step is 100 ms, there are $2,500$ time steps in the 250-second simulation.

Since Target 2 does not enter the surveillance region until $t = 100$ s, no resources are allocated to it until that point. At the beginning of the simulation ($t = 0$ s to $t = 5$ s), the scheduler elects to split its time between tracking dwells on Targets 1 and 3 and classification dwells on Target 1. Once Target 1 is identified (at approximately $t = 5$ s), the scheduler divides its time between tracking dwells on Targets 2 and 3, only occasionally revisiting with a classi-

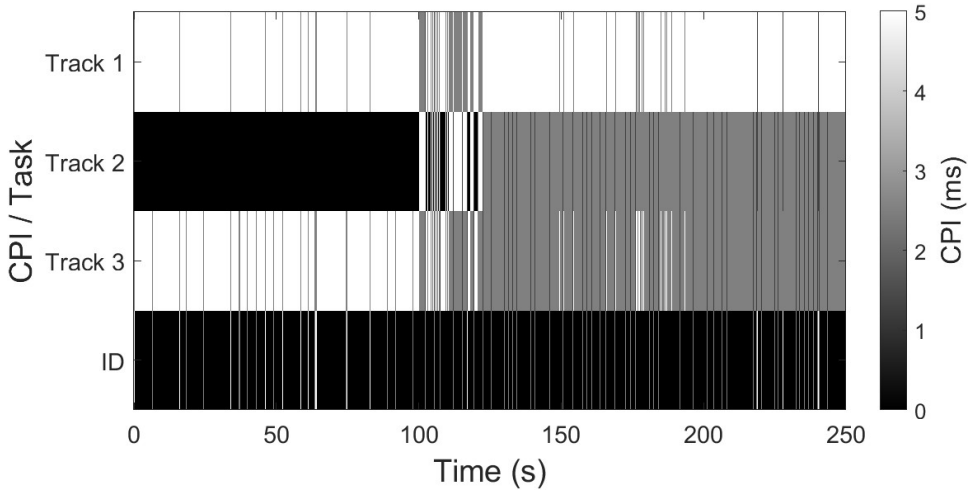

Figure 12.3: The time (out of 10 ms) the scheduler gives to each task for a single simulation. At the beginning, the scheduler performs some classification dwells and then rarely thereafter. Track dwells are initially divided between Tracks 1 and 3. Once Target 2 enters, it is preferentially measured. Later in the simulation, the scheduler gives extra resources to Track 1 due to its distance.

fication dwell on Target 1 since our model includes some time-induced class confusion for the reasons described in Section 12.6.2. When Target 2 enters at $t = 100$ s, the surveillance process provides a very coarse estimate of target state. This large track covariance leads the scheduler to preferentially use high-CPI dwells on Target 2 temporarily. Shortly thereafter, the range from the radar platform to Target 1 has increased enough that the scheduler now elects to use high CPI dwells on Track 1 to receive larger P_D at the expense of reduced P_D on Tracks 2 and 3.

A richer depiction of the scheduling choices for the three targets is given in Figure 12.4, which shows the bandwidth, PRF, and pulse count selected for each track at each scheduling time. As described in Section 12.6.2, these parameter selections influence the statistics of the measurements as well as SNR and detection probability.

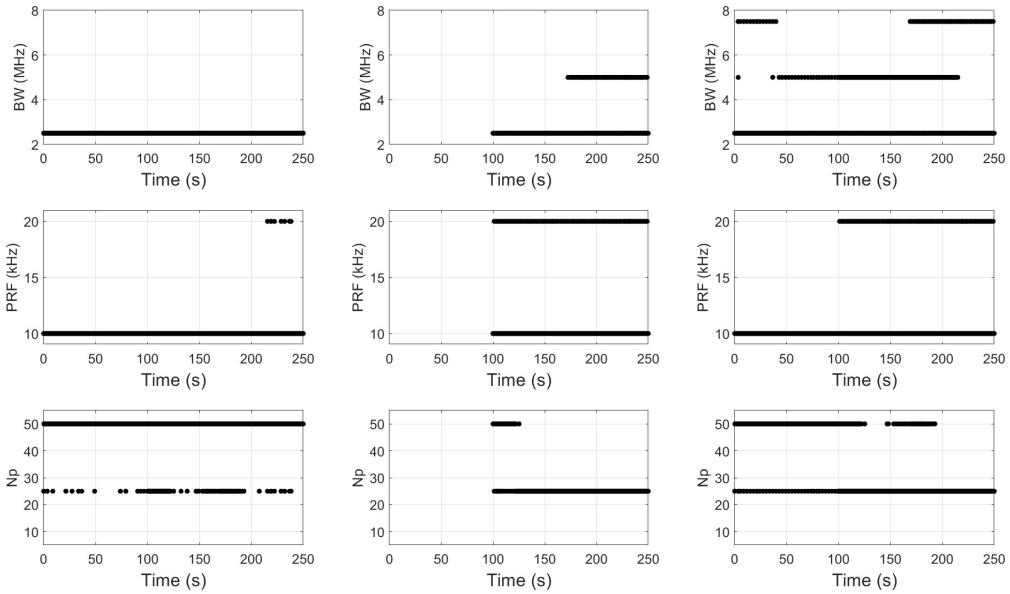

Figure 12.4: The selected waveforms for tracking Target 1 (left), Target 2 (middle), and Target 3 (right) for a single simulation. The scheduler varies the bandwidth, PRF, and pulse count to trade between detectability and range/range-rate resolution subject to the total time constraint.

The parameter selections correspond to the trade-offs the scheduler must make when selecting the waveform parameters. First, the limited time budget means that if the scheduler elects to use a longer CPI on one task to improve detectability that necessarily reduces the amount of time available for other tasks. Additionally, the scheduler may choose to get improved range resolution through higher bandwidth but that leads to reduced detection probability due to the correspondingly larger noise bandwidth entering the radar. This can be compensated for by increasing CPI at the expense of other targets. We see from the figure that the scheduler opts to choose different bandwidths, PRFs, and CPIs (via selection of PRF and pulse count) throughout the simulation as the received detections and track estimate warrant.

We now turn to a Monte Carlo analysis of the simulation. Figures 12.5 to

12.9 show the result of 100 Monte Carlo realizations of the simulation scenario. All trials share the same sensor and target trajectories, but the measurement realizations are drawn randomly each time. This leads to variations in the scheduling decisions and the MI, RMSE, and probability of correct class for each trial.

First, Figure 12.5 shows the average (over 100 trials) amount of time (of the 10 ms) that is allocated to each task at each time step in the simulation. This result is similar to that of the corresponding single simulation figure, showing how the scheduler elects to move its attention between the four tasks as the situation evolves. The prominent features include dividing attention equally between Tracks 1 and 3 early; focusing (briefly) on Track 2 at $t = 100$ s; and giving extra time to Track 1 near the end when its range to the platform is largest.

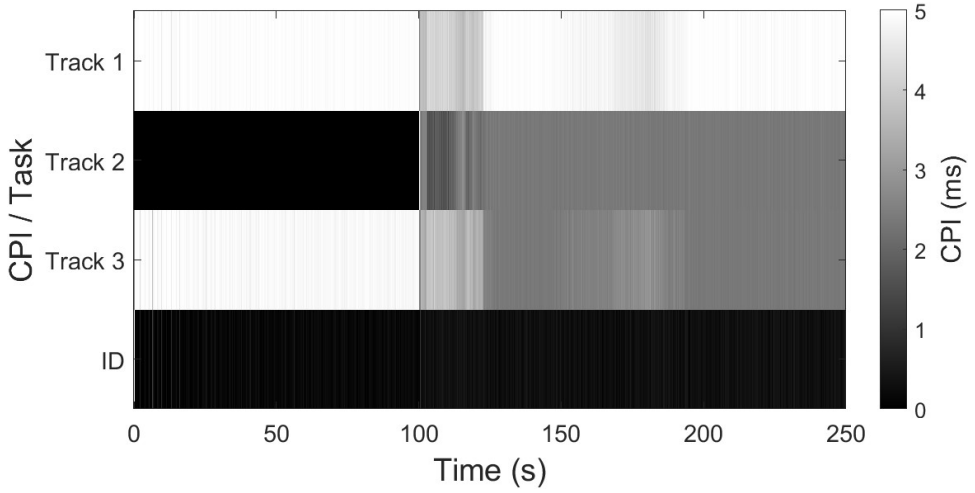

Figure 12.5: The average (over 100 trials) amount of time (out of 10 ms) the scheduler dedicates to each task during the simulation.

Next, Figure 12.6 illustrates the average MI for each of the four tasks (three

tracking task and one classification task) for the waveforms selected as well as the total MI as a function of time. The MI for each waveform is computed as described in (12.46), where the information flow the observation is expected to provide is weighted by the detection probability. The set of waveforms to transmit at a particular time is chosen to maximize the total MI subject to the total time constraint. Therefore, while the total MI represents the maximum total MI possible given the time constraints, the MI of the individual tasks is not necessarily the largest MI that any waveform could produce on that task.

The total MI is approximately constant over the simulation. Notable exceptions happen at $t \approx 15$ s, where Track 3 performs a maneuver and at $t = 100$ s, when Track 2 enters the region and has a large covariance provided by the detection engine. An additional feature this analysis reveals is that the MI of Track 1 declines near the end of the simulation as its range increases and detection probability is correspondingly decreased.

A more detailed look at the MI and its optimization is given in Figure 12.7. For each of the four tasks, we compare the MI (again, averaged over 100 trials) in the four-task simulation with the MI in a simulation that has only a single task. Broadly speaking, we find that there is a small loss in MI when all four tasks are present, resulting from the time constraint which prevents the scheduler from selecting the best possible waveform for each particular task at all times. A second feature of the comparison is that occasionally the four-task approach has more MI than the single-task MI (e.g., Track 3 at $t = 15$ s). This is because the MI is a measure of the information *gain* the measurement provides. In the single-task case, the scheduler is able to use the full CPI to follow Track 3, meaning that it is well followed even as it maneuvers. In contrast, when the time is to be divided among many tasks, Track 3 is not as

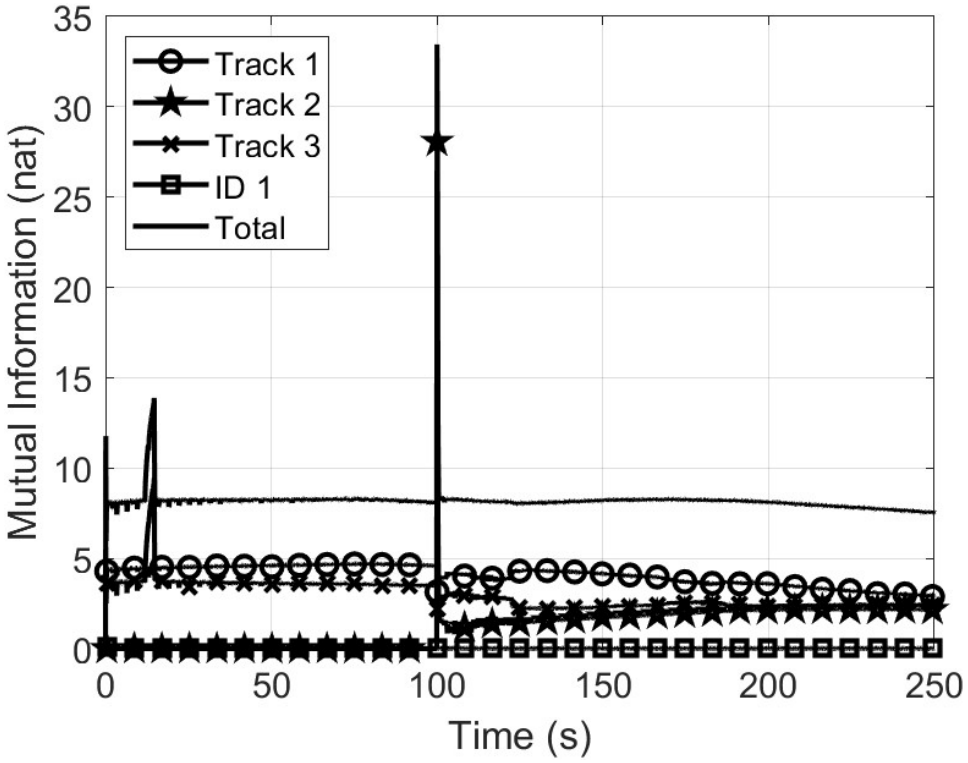

Figure 12.6: The average (over 100 trials) of the MI for each task, and the total MI.

well followed and is temporarily out of track as it maneuvers. This leads to a large expected information flow from making an observation immediately after the maneuver.

Figure 12.8 shows the average position and velocity RMSE as a function of time for the four-task simulation. Track 1 and Track 3 are detected at onset, while Track 2 is detected after $t = 100$ s. Track 2 is initialized with a large covariance and has a short period of larger RMSE while the track is refined. After that period, the RMSEs of both position and velocity are similar for the three targets until the end of the simulation where Target 1's distance makes it significantly more difficult to detect.

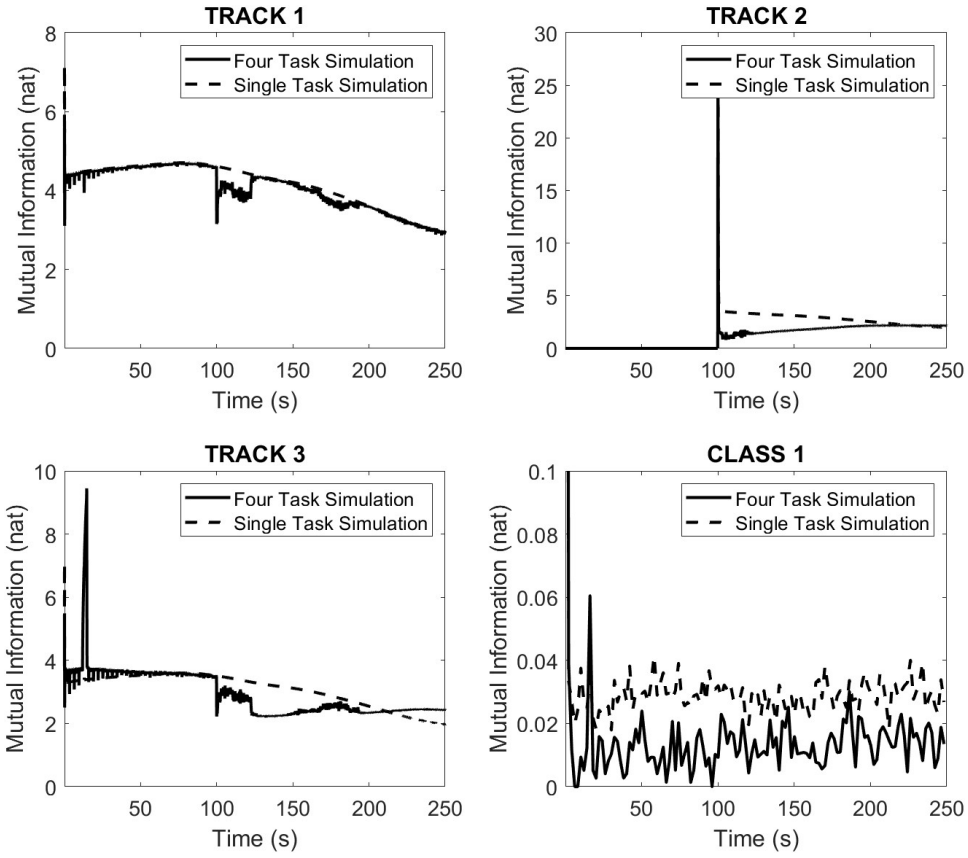

Figure 12.7: A comparison of the average (over 100 trials) MI in a four-task simulation with the MI in a simulation that has only a single task.

Finally, Figure 12.9 shows the posterior probability of correct class for Target 1 averaged over 100 trials. As mentioned earlier, our classification state transition model diffuses the class probability over time, requiring occasional revisits with a classification waveform to update the probability. We find empirically that the system maintains a high probability in the correct class.

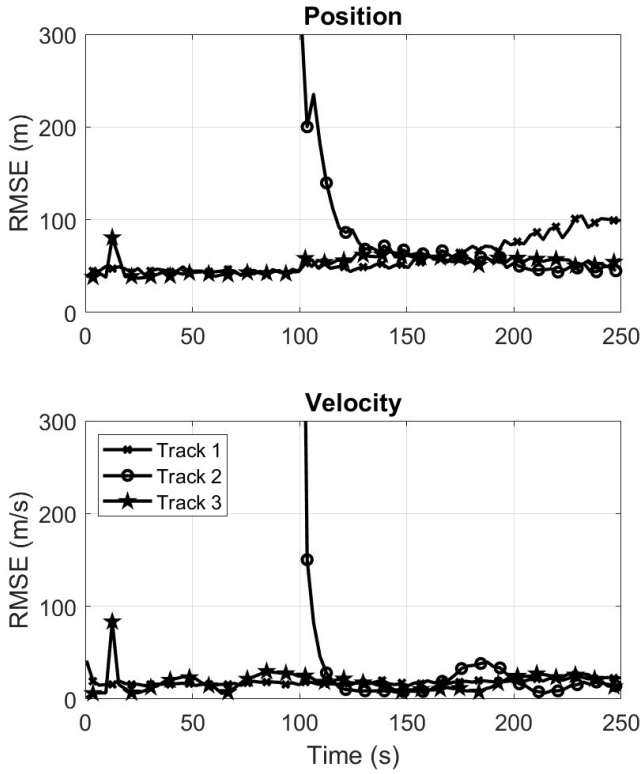

Figure 12.8: The average (over 100 trials) position and velocity RMSE tracking error for the three targets as a function of time.

12.7. Conclusion

New capabilities in modern multi-mode radars enable techniques to optimize waveform selection for multi-target detection, tracking, and classification. This chapter presented a rigorous model-based mathematical approach to fully adaptive radar resource management. The approach has two main components, which are based on the perception-action cycle of cognition. First, the perceptual processor constructs an probabilistic estimate of the state of the surveillance region. Next, the executive processor uses this estimate to determine the sensing actions expected to yield the most utility. In this chapter, we characterized utility using an information-theoretic measure, the MI.

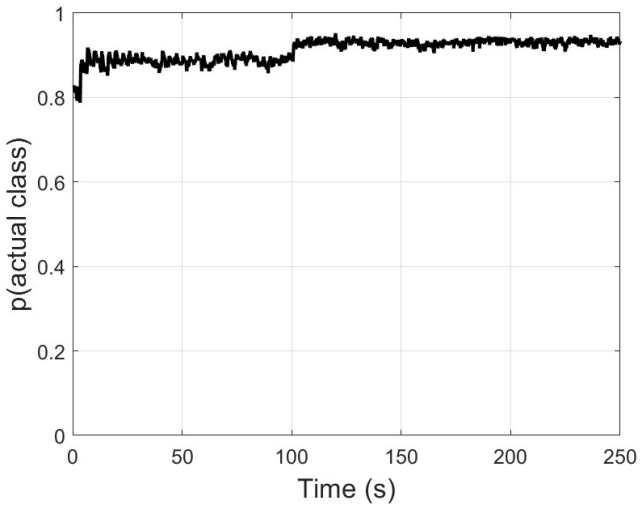

Figure 12.9: The average (over 100 trials) posterior probability in the correct class for the identification task.

Our model problem included multiple tracking and classification tasks, a bank of variable-length and variable-utility waveforms with a fixed time budget. Our technique selected the collection of waveforms to transmit during a radar frame subject to the total time constraint. Our techniques were illustrated with a multi-target simulation using collected aircraft trajectories and first-principles models of radar waveform utility.

Recent trends in cognitive radar include using data-driven and learning-based methods in place of or in conjunction with model-based methods. Future work might include comparison to these emerging techniques..

Acknowledgment

This material is based upon work supported by the Air Force Research Laboratory (AFRL) under Contract No. FA8649-20-P-0940. Any opinions, findings

and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of AFRL.

Bibliography

[1] S. Haykin, "Cognitive radar: a way of the future," *IEEE Signal Processing Magazine*, vol. 23, no. 1, pp. 30–40, January 2006.

[2] ——, *Cognitive Dynamic Systems: Perception-Action Cycle, Radar and Radio*. Cambridge University Press, 2012.

[3] J. R. Guerci, *Cognitive Radar: The Knowledge-Aided Fully Adaptive Approach*. Artech House, 2010.

[4] M. Greco, F. Gini, P. Stinco, and K. Bell, "Cognitive radars: on the road to reality: progress thus far and possibilities for the future," *IEEE Signal Processing Magazine*, vol. 35, no. 4, pp. 112–125, July 2018.

[5] S. Gurbuz, H. Griffiths, A. Charlish, M. Rangaswamy, M. Greco, and K. Bell, "An overview of cognitive radar: past, present, and future," *IEEE Aerospace and Electronic Systems Magazine*, vol. 34, no. 12, pp. 6–18, December 2019.

[6] J. M. Fuster, *Cortex and Mind: Unifying Cognition*. Oxford University Press, 2010.

[7] M. Hernandez, T. Kirubarajan, and Y. Bar-Shalom, "Multisensor resource deployment using posterior Cramér-Rao bounds," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 40, no. 2, pp. 399–416, April 2004.

[8] C. Kreucher, K. Kastella, and A. Hero, "Information-based sensor management for multitarget tracking," in *Proc. SPIE 5204*, 2003, pp. 480–489.

[9] ——, "A Bayesian method for integrated multitarget tracking and sensor management," in *Proc. 6th International Conference on Information Fusion*, 2003, pp. 704–711.

[10] C. Kreucher, A. Hero, and K. Kastella, "A comparison of task driven and information driven sensor management for target tracking," in *Proc. 44th IEEE Conference on Decision and Control*, 2005, pp. 4004–4009.

[11] C. Kreucher, K. Kastella, and A. Hero, "Sensor management using an active sensing approach," *Signal Processing*, vol. 85, no. 3, pp. 607–624, March 2005.

[12] C. Kreucher, D. Blatt, A. Hero, and K. Kastella, "Adaptive multimodality sensor scheduling for detection and tracking of smart targets," *Digital Signal Processing*, vol. 16, no. 5, pp. 546–567, September 2006.

[13] C. Kreucher, K. Kastella, and A. Hero, "Multi-platform information-based sensor management," in *Proc. SPIE 5820*, 2005, pp. 141–151.

[14] C. Kreucher, A. Hero, K. Kastella, and M. Morelande, "An information-based approach to sensor management in large dynamic networks," *Proceedings of the IEEE*, vol. 95, no. 5, pp. 978–999, May 2007.

[15] P. Chavali and A. Nehorai, "Scheduling and power allocation in a cognitive radar network for multiple-target tracking," *IEEE Transactions on Signal Processing*, vol. 60, no. 2, pp. 715–729, February 2012.

[16] R. Romero and N. Goodman, "Cognitive radar network: cooperative adaptive beamsteering for integrated search-and-track application," *IEEE*

Transactions on Aerospace and Electronic Systems, vol. 49, no. 2, pp. 915–931, April 2013.

[17] K. Bell, C. Baker, G. Smith, J. Johnson, and M. Rangaswamy, "Cognitive radar framework for target detection and tracking," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 8, pp. 1427–1439, December 2015.

[18] A. Charlish and F. Hoffmann, "Cognitive radar management," in *Novel Radar Techniques and Applications Volume 2: Waveform Diversity and Cognitive Radar, and Target Tracking and Data Fusion*. Institution of Engineering and Technology, 2017, pp. 157–193.

[19] W. Schmaedeke, "Information-based sensor management," in *Proc. SPIE 1955*, 1993.

[20] J. Manyika and H. F. Durrant-Whyte, "Information-theoretic approach to management in decentralized data fusion," in *Proc. SPIE 1828*, 1992.

[21] K. Bell, G. E. Smith, A. E. Mitchell, and M. Rangaswamy, "Multiple task fully adaptive radar," in *Proc. 52nd Asilomar Conference on Signals, Systems, and Computers*, 2018, pp. 1344–1348.

[22] K. Bell, C. Kreucher, and M. Rangaswamy, "An evaluation of task and information driven approaches for radar resource allocation," in *Proc. 2021 IEEE Radar Conference*, 2021.

[23] A. E. Mitchell, G. E. Smith, K. L. Bell, A. J. Duly, and M. Rangaswamy, "Cost function design for the fully adaptive radar framework," *IET Radar, Sonar, and Navigation*, vol. 12, no. 12, pp. 1380–1389, December 2018.

[24] R. Nadjiasngar and A. Charlish, "Quality of service resource management for a radar network," in *Proc. 2015 IEEE Radar Conference*, 2015, pp. 344–349.

[25] T. Cover and J. Thomas, *Elements of Information Theory*. Wiley, 1991.

[26] S. M. Kay, *Information Theoretic Signal Processing and its Applications*. Sachuest Point Publishers, 2020.

[27] F. Liese and I. Vajda, "On divergences and informations in statistics and information theory," *IEEE Transactions on Information Theory*, vol. 52, no. 10, pp. 4394–4412, November 2006.

[28] J. M. Aughenbaugh and B. R. LaCour, "Metric selection for information theoretic sensor management," in *Proc. 11th International Conference on Information Fusion*, 2008.

[29] C. Yang, I. Kadar, E. Blasch, and M. Bakich, "Comparison of information theoretic divergences for sensor management," in *Proc. SPIE 8050*, 2011.

[30] D. A. Castañón, R. Mahler, K. J. Hintz, J. Reich, I. Kadar, M. Farooq, T. Kirubarajan, R. Tharmarasa, T. Sathyan, and A. Sinha, "Issues in resource management with applications to real-world problems," in *Proc. SPIE 6235*, 2006.

[31] C. Kreucher, "An information-based approach to sensor resource allocation," PhD thesis, University of Michigan, February 2005.

[32] A. E. Mitchell, G. E. Smith, K. L. Bell, A. J. Duly, and M. Rangaswamy, "Hierarchical fully adaptive radar," *IET Radar, Sonar, and Navigation*, vol. 12, no. 12, pp. 1371–1379, December 2018.

[33] G. E. Smith, Z. Cammenga, A. E. Mitchell, K. L. Bell, J. T. Johnson, M. Rangaswamy, and C. J. Baker, "Experiments with cognitive radar," *IEEE Aerospace and Electronic Systems Magazine*, vol. 31, no. 12, pp. 34–46, December 2016.

[34] K. Bell, G. Smith, A. Mitchell, and M. Rangaswamy, "Fully adaptive radar for target classification," in *Proc. 2019 IEEE Radar Conference*, 2019.

[35] K. Bell, C. Kreucher, A. Brandewie, and J. Johnson, "Fully adaptive radar resource allocation for tracking and classification," in *Next Generation Cognitive Radar*, K. V. Mishra, B. M. Shankar, and M. Rangaswamy, Eds. IET Press (Radar, Electromagnetics and Signal Processing Technologies Series), 2024.

[36] A. Charlish, K. Bell, and C. Kreucher, "Implementing perception-action cycles using stochastic optimization," in *Proc. 2020 IEEE Radar Conference*, 2020.

[37] C. Kreucher, K. Kastella, and A. Hero, "Multitarget tracking using the joint multitarget probability density," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1396–1414, October 2005.

[38] J. Patrikar, B. Moon, J. Oh, and S. Scherer, "Predicting like a pilot: dataset and method to predict socially-aware aircraft trajectories in non-towered terminal airspace," in *Proc. 2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 2525–2531.

[39] R. A. Singer, "Estimating optimal tracking filter performance for manned maneuvering targets," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 6, no. 4, pp. 473–483, July 1970.

[40] M. Richards, *Fundamentals of Radar Signal Processing.* McGraw-Hill, 2005.

[41] H. L. Van Trees, *Optimum Array Processing.* Wiley, 2002.

[42] A. Dogandzic and A. Nehorai, "Cramér-Rao bounds for estimating range, velocity, and direction with an active array," *IEEE Transactions on Signal Processing*, vol. 49, no. 6, pp. 1122–1137, June 2001.

[43] V. L. Smith, "An experimental study of competitive market behavior," *Journal of Political Economy*, vol. 70, no. 2, p. 111–137, April 1962.

[44] D. Friedman, "The double auction market institution: a survey," in *The Double Auction Market: Institutions, Theories and Evidence*, D. Friedman and J. Rust, Eds. Perseus Publishing, 1993, ch. 1, pp. 3–25.