

Sensor management: Past, Present, and Future

Alfred O. Hero III, *Fellow, IEEE* and Douglas Cochran, *Senior Member, IEEE*

Abstract—Sensor systems typically operate under resource constraints that prevent the simultaneous use of all resources all of the time. Sensor management becomes relevant when the sensing system has the capability of actively managing these resources; i.e., changing its operating configuration during deployment in reaction to previous measurements. Examples of systems in which sensor management is currently used or is likely to be used in the near future include autonomous robots, surveillance and reconnaissance networks, and waveform-agile radars. This paper provides an overview of the theory, algorithms, and applications of sensor management as it has developed over the past decades and as it stands today.

Index Terms—Active adaptive sensors, Plan-ahead sensing, Sequential decision processes, Stochastic control, Multi-armed bandits, Reinforcement learning, Optimal decision policies, Multi-stage planning, Myopic planning, Information-optimized planning, Policy approximation, Radar waveform scheduling

I. INTRODUCTION

Advances in sensor technologies in the last quarter of the 20th century led to the emergence of large numbers of controllable degrees of freedom in sensing devices. Large numbers of traditionally hard-wired characteristics, such as center frequency, bandwidth, beamform, sampling rate, and many other aspects of sensors' operating modes started to be addressable via software command. The same period brought remarkable advances in networked systems as well as deployable autonomous and semi-autonomous vehicles instrumented with wide ranges of sensors and interconnected by networks, leading to configurable networked sensing systems. These trends, which affect a broad range of sensor types, modalities, and application regimes, have continued to the present day and appear unlikely to abate: new sensing concepts are increasingly manifested with device technologies and system architectures that are well suited to providing agility in their operation.

A. O. Hero is with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109-2122, USA. D. Cochran is with the School of Mathematical and Statistical Sciences and the School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85287-5706, USA.

The term “sensor management,” as used in this paper, refers to control of the degrees of freedom in an agile sensor system to satisfy operational constraints and achieve operational objectives. To accomplish this, one typically seeks a policy for determining the optimal sensor configuration at each time, within constraints, as a function of information available from prior measurements and possibly other sources. With this perspective, the paper casts sensor management in terms of formulation and approximation of optimal planning policies. This point of view has led to a rich vein of research activity that extends and blends ideas from control, information theory, statistics, signal processing, and other areas of mathematical, statistical, and computational sciences and engineering. Our viewpoint is also slanted toward sensor management in large-scale surveillance and tracking systems for civilian and defense applications. The approaches discussed have much broader utility, but the specific objectives, constraints, sensing modalities, and dynamical models considered in most of the work summarized here have been drawn from this application arena.

Within its scope of attention, the intention of this paper is to provide a high-level overview; references are given to guide the reader to derivations of mathematical results, detailed descriptions of algorithms, and specifications of application scenarios and systems. The list of references, while extensive, is not exhaustive; rather it is representative of key contributions that have shaped the field and led to its current state. Moreover, there are several areas relevant or related to sensor management that are not within the scope of this survey. These include purely heuristic approaches to sensor management and scheduling as well as adaptive search methods, clinical treatment planning, human-in-the-loop systems such as relevance feedback learning, robotic vision and autonomous navigation (path planning), compressive and distilled sensing, and robust sensing based on non-adaptive approaches.

The most comprehensive recent survey on sensor management of which the authors are aware is the 2008 book [1]. This volume consists of chapters written collaboratively by numerous current contributors to the field specifically to form a perspicuous overview of the main methods and some noteworthy applications. The 1998 survey paper by A. Cassandra [2], while not

devoted to sensor management, describes a few applications of partially observed Markov decision process (POMDP) methods in the general area of sensor management and scheduling, thereby illustrating conceptual connections between sensor management and the many other POMDP applications summarized in the paper. The earlier 1982 survey paper by G. E. Monahan [3] does not consider sensor management applications, but gives an excellent overview of the base of theory and algorithms for POMDPs as they were understood a few years before sensor management was becoming established as an appreciable area of research. A 2000 paper by G. W. Ng and K. H. Ng [4] provides an overview of sensor management from the perspective of sensor fusion as it stood at that time. This point of view, although not emphasized in this paper or in [1], continues to be of interest in the research literature. Another brief survey from this period is given by X.-X. Liu *et al.* in [5], and a short survey of emerging sensor concepts amenable to active sensor management is given in [6].

Several doctoral dissertations on the topic of sensor management have been written in the past fifteen years. Most of these include summaries of the state of the art and relevant literature at the time they were composed. Among these are the dissertations of G. A. McIntyre (1998) [7], D. Sinno (2000) [8], C. M. Kreucher (2005) [9], R. Rangarajan (2006) [10], D. Blatt (2007) [11], J. L. Williams (2007) [12], M. Huber (2009) [13], and K. L. Jenkins (2010) [14].

The remainder of this paper is organized as follows. Section II describes the basic goals and defines the main components of a sensor management system. In Section III, the emergence of sensor management is recounted within a historical context that includes both the advancement of statistical methods for sequential definition, collection, and analysis of samples and the rise of sensor technologies and sensing applications enabling and calling for sensor management. Section IV gives an overview of some of the current state of the art and trends in sensor management and Section V describes some of the future challenges and opportunities faced by researchers in the field.

II. DESCRIPTION OF SENSOR MANAGEMENT

The defining function of sensor management is dynamic selection of a sensor, from among a set of available sensors, to use at each time during a measurement period in order to optimize some metric of performance. Time is usually partitioned into a sequence of epochs and one sensor is to be chosen in each epoch, thereby creating a discrete-time problem. The

term “sensor management” most often refers to closed-loop solutions to problems of this nature; i.e., the next sensor to employ is chosen while the sensor system is in operation and in view of the results obtained from prior sensor measurements. The term “sensor scheduling” is sometimes used to refer to feed-forward schemes for sensor selection, though this usage is not standardized and the two expressions are used interchangeably in some literature. In current applications of sensor management, and especially in envisioned future applications, the sensors available for selection in each time epoch are actually virtual sensors, each representing one choice of configuration parameters affecting the physical configurations and operating modes of a collection of sensors, sensor suites, sensor platforms, and the way data are processed and communicated among interconnected subsystems. With this perspective, selecting a sensor really means determining the values to which the available controllable degrees of freedom in a sensor system should be set.

Figure 1 illustrates the basic elements and operation of a closed-loop sensor management system. Once a sensor is selected and a measurement is made, information relevant to the sensing objective is distilled from the raw sensor data. This generally entails fusion of data representing disparate sensing modalities (e.g., optical and acoustic) and other properties, and further combining it with information gleaned from past measurements and possibly also side information from sources extrinsic to the sensor system. The fusion and signal processing components of the loop may produce ancillary information, such as target tracks or decisions about matters external to the sensor manager (e.g., direct an aircraft to take evasive action to avoid collision). For the purposes of sensor management, they must yield a state of information on the basis of which the merit of each possible sensor selection in the next time epoch may be quantified. Such quantification takes many forms in current approaches, from statistical (e.g., mean risk or information gain) to purely heuristic. From this point, the sensor manager must optimize its decision as to which sensor to select for the next measurement.

The notion of state is worthy of a few additional words. Heuristically, the state of information should represent all that is known about the scenario being sensed, or at least all that is relevant to the objective. Often this includes information about the physical state of the sensor system itself (e.g., the position and orientation of the air vehicle carrying one of the video sensors), which may constrain what actions are possible in the next step and thus the set of virtual sensors available to select in the upcoming epoch. Knowledge of the physical state




Fig. 1. Conceptual block diagram of a sensor management system. The sensor selector selects among sensor actions S1, S2, and S3 based on the output of the optimizer. The optimizer attempts to optimize a system performance metric, such as information gain or mean risk associated with decisions or estimates produced by signal processing algorithms that operate on fused sensor data.

frequently has utility extrinsic to the sensor manager, so some literature distinguishes physical and information states and their coupled dynamical models as depicted in Figure 2. This diagram evinces the similarity of sensor management and feedback control in many important respects, and indeed control theory is an important ingredient in current perspectives on sensor management. But sensor management entails certain aspects that give it a distinctive character. Chief among these is in the role of sensing. In traditional feedback control, sensors are used to ascertain information about the state of a dynamical plant. This information informs the control action through a control law or policy which in turn affects the state. In sensor management, the state of information is directly affected by the control action; i.e., rather than helping to decide what control action to invoke, the act of sensing is itself the control action.

Sensor management is motivated and enabled by a small number of essential elements. The following para-




Fig. 2. A control-theoretic view of sensor management casts the problem as that of optimally controlling a state, sometimes regarded as consisting of separate information and physical components, through the selection of measurement actions.

graphs describe these and explain the roles they play in the current state of the subject. First, a summary of waveform-agile radar is given to provide the context of a current application for the more general descriptions that follow.

A. Sensor management application – Waveform-agile radar

Among the most well developed focus applications of sensor management is real-time closed-loop scheduling of radar resources. The primary feature of radar systems that makes them well suited for sensor management is that they offer several controllable degrees of freedom. Most modern radars employ antenna arrays for both the transmitter and receiver, which often share the same antenna. This allows the illumination pattern on transmit as well as the beam pattern on receive to be adjusted simply by changing parameters in a combining algorithm. This ability has been capitalized upon, for example, by adaptive signal processing techniques such as adaptive beamforming on both transmit and receive and more recently by space-time adaptive processing (STAP). The ability for the transmitter to change waveforms in a limited way, such as switching between a few pre-defined waveforms in a library, has existed in a few radar systems for decades. Current radar concepts allow transmission of essentially arbitrary waveforms, with constraints coming principally from hardware limitations such as bandwidth and amplifier power. They also remove traditional restrictions that force the set of transmit antenna elements to be treated as a phased array (i.e., all emitting the same waveform except for phase factors that steer the beam pattern), thereby engendering the possibility of the

transmit antennas simultaneously emitting completely different waveforms. This forms the basis of one form of so-called multi-input multi-output (MIMO) radar.

Two more aspects of the radar application stand out in making it a good candidate for sensor management. One is that pulse-Doppler radars have discrete time epochs intrinsically defined by their pulse repetition intervals and often also by their revisit intervals [15]. Also, in radar target tracking applications there are usually well defined performance metrics and well developed dynamical models for the evolution of the targets' positions, velocities, and other state variables. These metrics and models directly enhance the sensor manager's ability to quantitatively predict the value of candidate measurements before they are taken.

In view of these appealing features, it is no surprise that radar applications have received a large amount of attention as sensor management has developed. The idea of changing the transmitted waveform in a radar system in an automated fashion in consideration of the echo returns from previously transmitted waveforms dates to at least the 1960s, though most evidence of this is anecdotal rather than being documented in the research literature. The current generation of literature on closed-loop waveform management as a sensor management application began with papers of D. J. Kershaw and R. J. Evans [16], [17] and S. M. Sowelam and A. H. Tewfik [18], [19] in the mid-1990s, roughly corresponding to the ascension of sensor management literature in broader contexts. Among the early sensor management papers that focused on closed-loop beam pattern management were those of V. Krishnamurthy and Evans in the early 2000s [20], [21]. Several contributions by numerous authors on these and related radar sensor management applications have appeared in the past decade. Among the topics addressed in this recent literature are radar waveform scheduling for target identification [22], target tracking [23], clutter and interference mitigation [24], [25], and simultaneously estimating and tracking parameters associated with multiple extended targets [26]. There has also been recent interest in drawing insights for active radar and sonar sensor management from biological echolocation systems [27] and in designing optimal libraries of waveforms for use with radar systems that support closed-loop waveform scheduling [28].

B. Controllable Degrees of Freedom

Degrees of freedom in a sensor system over which control can be exercised with the system in operation provide the mechanism through which sensors can be managed. In envisioned applications, they include di-

verse sets of parameters, including physical configuration of the sensor suite, signal transmission characteristics such as waveform or modulation type, signal reception descriptors ranging from simple on/off state to sophisticated properties like beamform. They also include algorithmic parameters that affect local versus centralized processing trade-offs, data sharing protocols and communication schemes, and typically numerous signal processing choices.

Many characteristics of current and anticipated sensor systems that are controllable during real-time operation were traditionally associated with subsystems that were designed independently. Until relatively recently, transduction of physical phenomena into electrical signals, analog processing, conversion to digital format, and digital processing at various levels of information abstraction were optimized according to performance criteria that were often only loosely connected with the performance of the integrated system in its intended function. Further integrated operation of such subsystems generally consisted of passing data downstream from one to the next in a feed-forward fashion. Integrated real-time authority over controllable degrees of freedom spanning all of this functionality not only allows joint optimization of systemic performance metrics but also accommodates adaptation to changing objectives.

In the radar sensor management example, the ease and immediacy of access (i.e., via software command) to crucial operating parameters such as antenna patterns and waveforms provides the means by which a well conceived algorithm can manage the radar in each time epoch.

C. Constraints

The utility of sensor management emerges when it is not possible to process, or even collect, all the data all the time. Operating configurations of individual sensors or entire sensor systems may be intrinsically mutually exclusive; e.g., the transmitter platform can be in position A or in position B at the time the next waveform is emitted, but not both. One point of view on configurable sensors, discussed in [29], imagines an immense suite of virtual sensor systems, each defined by a particular operating configuration of the set of physical sensors that comprises the suite. Limitations preventing an individual sensor from being in multiple configurations at the same time are seen as constraints to be respected in optimizing the configuration of the virtual sensor suite. This is exactly the case in the waveform-agile radar example, where only one waveform can be transmitted on each antenna element at any given time.

Restrictions on communications and processing resources almost always constrain what signal processing is possible in networked sensor applications. Collecting all raw data at a single fusion center is seldom possible due to bandwidth limitations, and often to constraints imposed by the life and current production of batteries as well. So it is desirable to compress raw data before transmission. But reducing the data at the nodes requires on-board processing, which is typically also a limited resource.

D. Objective Quantification

When controllable degrees of freedom and constraints are present, sensor management is possible and warranted. In such a situation, one would hope to treat the selection of which sensing action to invoke as an optimization problem. But doing so requires the merit of each possible selection to be represented in such a way that comparison is possible; e.g., by the value of a cost or objective functional.

The value of a specified set of data collection and processing choices generally depends on what is to be achieved. For example, one set of measurements by a configurable chemical sensor suite may be of great value in determining whether or not an analyte is an explosive, but the best data to collect to determine the species of a specimen already known to be an explosive may be quite different. Moreover, the objective may vary with time or state of knowledge: once a substance is determined to be an explosive, the goal shifts to determining what kind of explosive it is, then how much is present, then precisely where it is located, etc. Consequently, predictively quantifying the value of the information that will be obtained by the selection of a particular sensing action is usually difficult and, at least in principle, requires a separate metric for each sensing objective that the system may be used to address. The use of surrogate metrics, such as information gain discussed in Section IV, has proven effective in some applications. With this approach, the role of a metric designed specifically for a particular sensing objective is undertaken by a proxy, usually based on information theoretic measures, that is suited to a broader class of objectives. This approach sacrifices specificity in exchange for relative simplicity and robustness, especially to model mismatch.

Management of radar beamforms and waveforms for target tracking, though not trivial, is one of the most tractable settings for objective quantification. The parameters can be chosen to optimize some function of the track error covariance, such as its expected trace or determinant, at one or more future times; e.g., after

the next measurement, after five measurement epochs, or averaged over the next ten epochs. Computation or approximation of such functions is assisted by the tracker's underlying model for the dynamical evolution of the target states. The use and effectiveness of waveform management in such applications is discussed in [1, Ch. 10], which also cites numerous references.

III. HISTORICAL ROOTS OF SENSOR MANAGEMENT

It has long been recognized that appropriate collection of data is essential in the design of experiments to test hypotheses and estimate quantities of interest. R. A. Fisher's classical work [30], which encapsulated most of the ideas on statistical design of experiments developed through the first part of the 20th century, primarily addressed the situation in which the composition of the sample to be collected is to be determined in advance of the experiment. In the early 1950s, the idea of using closed-loop strategies in experiment design emerged in connection with sequential design of experiments. In his 1951 address to the Meeting of the American Mathematical Society [31], H. Robbins observed:

A major advance now appears to be in the making with the creation of a theory of the sequential design of experiments, in which the size and composition of the samples are not fixed in advance but are functions of the observations themselves.

Robbins attributes the first application of this idea to Dodge and Romig in 1929 [32] in the context of industrial quality control. They proposed a double sampling scheme in which an initial sample is collected and analyzed, then a determination about whether to collect a second sample is based on analysis of the first sample. This insight was an early precursor to the development of sequential analysis by Wald and others during the 1940s [33], and ultimately to modern methods in statistical signal processing such as sequential detection [34]. In the interim, H. Chernoff made substantial advances in the statistical study of optimal design of sequences of experiments, particularly for hypothesis testing and parameter estimation [35], [36]. Many results in this vein are included in his 1972 book [37]. Also in 1972, V. V. Fedorov's book [38] presented an overview of key results, many from his own research, in optimal experimental design up to that time. The relevance of a portion of Fedorov's work to the current state of sensor management is noted in Section IV.

One view of the *raison d'être* for sensors, particularly among practitioners of sensor signal processing, is to

collect samples to which statistical tests and estimators may be applied. From this perspective, the advancement of sensor signal processing over the latter half of the 20th century paralleled that of experimental design. By the early 1990s, a rich literature on detection, estimation, classification, target tracking and related problems had been compiled. Nearly all of this work was predicated on the assumption that the data were given and the goal was to process it in ways that are optimally informative in the context of a given application. There were a few notable cases in which it was assumed the process of data collection could be affected in a closed-loop fashion based on data already collected. In sequential detection theory, for example, the data collection is continued or terminated at a given time instant (i.e., binary feedback) depending on whether a desired level of confidence about the fidelity of the detection decision is supported by data already collected. An early example of closed-loop data collection involving a dynamic state was the “measurement adaptive problem” treated by L. Meier *et al.* in 1967 [39]. This work sought to simultaneously optimize control of a dynamic plant and the process of collecting measurements for use in feedback. Another is given in a 1972 paper of M. Athans [40] that considers optimal closed-loop selection of the linear measurement map in a Kalman filtering problem.

One of the first contexts in which the term “sensor management” was used in the sense of this discussion¹ was in automating control of the sensor systems in military aircraft (see, e.g., [42]). In this application, the constrained resource is the attention of the pilot, particularly during hostile engagement with multiple adversaries, and the objective of sensor management is to control sensor resources in such a way that the most important information (e.g., the most urgent threats) are emphasized in presentation to the pilot. Applications associated with situational awareness for military aircraft continue to be of interest, and this early vein of application impetus expanded throughout the 1990s to include scheduling and management of aircraft-based sensor assets for surveillance and reconnaissance missions (see, e.g., [43], [44] and [1, ch. 11]).

Also beginning in the 1980s, sensor management was actively pursued under the label of “active vision” for applications in robotics [45]. This work sought to exercise feedback control over camera direction and sometimes other basic parameters (e.g., zoom or focal distance) to improve the ability of robotic vision systems

to contribute to navigation, manipulation, and other tasks entailed in the robot’s intended functionality.

The rapid growth of interest in sensor management beginning in the 1990s can be attributed in large part to developments in sensor and communications technologies. New generations of sensors, encompassing numerous sensing modalities, are increasingly agile. Key operating parameters, once hard-wired, can be almost instantly changed by software command. Further, transducers can be packaged with A/D converters and microprocessors in energy efficient configurations, in some cases on a single chip, creating sensors that permit on-board adaptive processing involving dynamic orchestration of all these components. At the same time, the growth of networks of sensors and mobile sensor platforms is contributing even more controllable degrees of freedom that can be managed across entire sensor systems. From a purely mathematical point of view, it is almost always advantageous to collect all available data in one location (i.e., a “fusion center”) for signal processing. In today’s sensor systems, this is seldom possible because of constraints on computational resources, communication bandwidth, energy, deployment pattern, platform motion, and many other aspects of the system configuration. Even highly agile sensor devices are constrained to choose only one configuration from among a large collection of possibilities at any given time.

These years spawned sensor management approaches based on the modeling sensor management as a decision process, a perspective that underpins most current methods as noted in Section IV. Viewing sensor management in this way enabled tapping into a corpus of knowledge on control of decision processes, Markov decision processes in particular, that was already well established at the time [46]. Initial treatments of sensor management problems via POMDPs, beginning with D. Castañón’s 1997 paper [47], were followed shortly by other POMDP-based ideas such as the work of J. S. Evans and Krishnamurthy published in 2001–2002 [48], [49]. These were the early constituents of a steady stream of contributions to the state of the art summarized in Section IV-B. A formidable obstacle to the practicality of the POMDP approach is the computational complexity entailed in its implementation, particularly for methods that look more than one step ahead. Consequently, the need for approximation schemes and the potential merit of heuristics to provide computational tractability was recognized from the earliest work in this vein.

The multi-armed bandit (MAB) problem is an important exemplar of a class of multi-stage decision problems where actions yielding large immediate rewards must be balanced with others whose immediate rewards are

¹The phrase comes up in various literature in ways that are related to varying degrees to our use in this paper. To maintain focus, we have omitted loosely related uses of the term, such as in clinical patient screening applications [41].

smaller, but which hold the potential for greater long-term payoff. While two-armed and MAB problems had been studied in previous literature, the origin of index policy solutions to MAB problems dates to J. C. Gittins in 1979 [50]. As discussed in Section IV-C, under certain assumptions, an index solution assigns a numerical index to each possible action at the current stage of an infinitely long sequence of plays of a MAB. The indices can be computed by solving a set of simpler one-armed bandit problems and their availability reduces the decision at each stage to choosing the action with the largest index. The optimality of Gittins' index scheme was addressed by P. Whittle in 1980 [51].

As with POMDPs, the MAB perspective on sensor management started receiving considerable research attention around 2000. Early applications of MAB methodology to sensor management include the work of Krishnamurthy and R. J. Evans [20], [21] who considered a multi-armed bandit model with Markov dynamics for radar beam scheduling. The 2002 work of R. Washburn *et al.* [52], although written in the context of more general dynamic resource management problems, was influential in the early develop of MAB approaches to sensor management.

A theory of information based on entropy concepts was introduced by C. E. Shannon in his classic 1948 paper [53] and was subsequently extended and applied by many others, mostly in connection with communication engineering. Although Shannon's theory is quite different than that of Fisher, sensor management has leveraged both in various developments of information-optimized methods. These were introduced specifically to sensor management in the early 1990s by J. Manyika and H. Durrant-Whyte [54] and by W. W. Schmaedeke [55]. As remarked in Section IV, information-based ideas were applied to particular problems related to sensor management even earlier. Fisher's information theory was instrumental in the development of the theory of optimal design of experiments, and numerous examples of applications of this methodology have appeared since 2000; e.g., [56], [57]. Measures of information led to sensor management schemes based on information gain, which developed into one of the central thrusts of sensor management research over the past decade. Some of this work is summarized in Section IV-D, and a more complete overview of these methods is provided in [1, Ch. 3].

From foundations drawing on several more classical fields of study, sensor management has developed into a well-defined area of research that stands today at the crossroads of the disciplines upon which it has been built. Key approaches that are generally known to researchers

in the area are discussed in the following section of this paper. But sensor management is an active discipline, with new work and new ideas appearing regularly in the literature. Some noteworthy recent developments include work by V. Gupta *et al.* which introduces random scheduling algorithms that seek optimal mean steady state performance in the presence of probabilistically modeled effects [58], [59], [60]. K. L. Jenkins *et al.* very recently proposed the use of random set ideas, similar to those applied in some approaches to multi-target tracking, in sensor management [61], [62]. These preliminary investigations have resulted in highly efficient algorithms for certain object classification problems. Also very recently, D. Hitchings *et al.* introduced new stochastic control approximation schemes to obtain tractable algorithms for sensor management based on receding horizon control formulations [63]. They also proposed a stochastic control approach for sensor management problems with large, continuous-valued state and decision spaces [64].

Despite ongoing progress, sensor management still holds many unresolved challenges. Some of these are discussed in Section V.

IV. STATE OF THE ART IN SENSOR MANAGEMENT

The theory of decision processes provides a unifying perspective for the state of the art in sensor management research today. A decision process, described in more detail below, is a time sequence of measurements and control actions in which each action in the sequence is followed by a measurement acquired as a result of the previous action. With this perspective, the design of a sensor manager is formulated as the specification of a decision rule, often called a policy, that generates realizations of the decision process. An optimal policy will generate decision processes that, on the average, will maximize an expected reward; e.g., the negative mean-squared tracking error or the probability of detection. A sound approach to sensor management will either approximate an optimal policy in some way or else attempt to analyze the performance of a proposed heuristic policy. In this section we will describe some current approaches to design of sensor management policies. The starting point is a formal definition of a decision process.

A. Sensor management as a decision process

Assume that a sensor collects a data sample y_{t+1} at time t after taking a sensing action a_t . It is typically assumed that the possible actions are selected from a finite action space \mathcal{A} , that may change

over time. The selected action a_k depends only on past samples $\{y_k, y_{k-1}, \dots, y_1\}$ and past actions $\{a_{k-1}, a_{k-2}, \dots, a_0\}$, and the initial action a_0 is determined offline. The function that maps previous data samples and actions to current actions is called a policy. That is, at any time t , a policy specifies a mapping γ_t and, for a specific set of samples, an action $a_t = \gamma_t(\{a_k\}_{k < t}, \{y_k\}_{k \leq t})$. A decision process is a sequence $\{a_k, y_{k+1}\}_{k \geq 0} = \{a_0, y_1, a_1, y_2, a_2, y_3, \dots\}$, which is typically random and can be viewed as a realization from some generative model specified by the policy and the sensor measurement statistics.

A well designed sensor manager will formulate the policy with the objective of maximizing an average reward. The reward at time t is a function $R_t(\{a_k\}_{k < t}, \{s_k\}_{k \leq t})$ of the action sequence $\{a_k\}_{k > 0}$ and a state sequence $\{s_k\}_{k > 0}$, describing the environment or a target in the environment. The state s_k might be continuous (e.g., the position of a moving target) or discrete (e.g., $s_k = 1$ when the target is moving and $s_k = 0$ when it is not moving). It is customary to model the state as random and the data sample y_k as having been generated by the state s_k in some random manner. In this case, there exists a conditional distribution of the state sequence given the data sequence and the average reward at time t can be defined through the statistical expectation $E[R_t(\{a_k\}_{k < t}, \{s_k\}_{k \leq t})]$.

An optimal action policy will maximize the average reward at each time t during the sensor deployment time period. The associated optimization must be performed over the set of mappings γ_t defined on the cartesian product spaces $\times_{k=1}^t \{\mathcal{A}_{k-1} \times \mathcal{Y}\}$ and mapping to \mathcal{A}_t for $t = 0, 1, \dots$. Due to the high dimensionality of the cartesian product spaces, no tractable methods exist for determining optimal action policies under this degree of generality. Additional assumptions on the statistical distributions of the decision process and state process are needed to reduce the dimensionality of the optimization spaces.

When the unknown state s_k is not recoverable from y_k then the decision process is called a partially observable decision process. The partially observable case is common in actual sensing systems where the measurements y_k are typically contaminated by noise or clutter. However, policy optimization generally presents more mathematical difficulties in the partially observable case than in the perfectly observable case.

B. Markov decision processes

A natural way to simplify the task of policy optimization is to assume that the general decision process de-

scribed in Section IV-A satisfies some additional Markovian properties. To make the general decision process Markovian one imposes the assumption that the state sequence is dependent only on the most recent state and action given the entire past. Specifically, we assume that $P(s_{t+1} | \{s_k, a_k\}_{k \leq t}) = P(s_{t+1} | s_t, a_t)$, the conditional state transition probability, and $P(y_t | \{s_k, a_k\}_{k \leq t}) = P(y_t | s_t, a_t)$, the measurement likelihood function given action a_t .

We additionally restrict the reward to be additive over time and only consider policies that depend on the most recent measurement, i.e., $R_t(\{a_k\}_{k < t}, \{s_k\}_{k \leq t}) = \sum_{k=0}^t R_t(a_k, s_k)$ and the associated mapping γ_t is restricted to be from $\mathcal{Y} \times \mathcal{A}_{t-1}$ to \mathcal{A}_t . When the state can be recovered from the measurements the resultant process is called a Markov decision process (MDP). When the state is not recoverable from the measurements the resultant process is called a partially observable Markov decision process (POMDP).

For MDP or POMDP models the optimal restricted policy can be determined by backwards induction over time. In particular, there is a compact recursive formula, known as Bellman's equation, for determining the mapping γ_{t-1} from the mapping γ_t . In special cases where the state and the measurements obey standard dynamical stochastic state models (e.g., the linear-Gaussian model assumed in the Kalman filter), this optimal restricted policy is in fact the overall optimal policy. That is, the overall optimal policy only depends on the most recent measurements. Furthermore, as shown by E. J. Sondik [65], the optimal policy can be found by linear programming. For more details on MDPs, POMDPs, and Bellman's equation and solutions, the reader is referred to [1, Ch. 2].

As noted in Section III, the application of MDP and POMDP methods to sensor management problems can be traced back to the mid 1990's. In their 1994 overview of the field of sensor management [43], S. Musick and R. Malhotra suggested that a comprehensive mathematical framework was needed to assess and optimize scheduling over sensor and inter-sensor actions. Anticipating the future application of POMDP and reinforcement learning approaches, went on to suggest adaptive control, state space representations, and mathematical programming as the components of a promising framework. However, to be applied to practical large scale sensor management problems approximate solutions to the POMDP would be necessary. Castañón's 1997 policy rollout approximation [47] was the earliest successful application of the POMDP to sensor management.

Several types of approximations to the optimal POMDP sensor management solution are discussed in

[1, Ch. 2] under the heading of approximate dynamic programming. These include: offline learning, rollout, and problem approximation techniques. Offline learning techniques use offline simulation to explore the space of policies and include the large class of reinforcement learning methods [44], [66], [67]. Rollout uses real-time simulation to approximate the rewards of a sub-optimal policy [47], [68]. Problem approximation uses a simpler approximate model or reward function for the POMDP as a proxy for the original problem and includes bandit and information gain approaches, discussed in the following subsections.

POMDP approaches have been applied to many different sensing systems. One of the most active areas of application has been distributed multiple target tracking, see for example [69] and references therein. When target dynamics are non-linear and environments are dynamically changing, the states of targets can be tracked by a particle filter [70]. This filter produces an estimate of the posterior density of the target tracks that is used by the scheduler to predict the value of different sensing actions [67]. Managers for many other sensing systems have been implemented using POMDPs and reinforcement learning, for example, multifunction radar [71], underwater sensing applications [72], passive radar [73], and air traffic management [74].

C. Multi-armed bandit decision processes

A multi-armed bandit (MAB) is a model for sequential resource allocation in which multiple resources (the arms of the bandit) are allocated to multiple tasks by a controller (also called a processor). When a particular arm a_t of the bandit is pulled (a control action called a “play”) at time t the MAB transitions to a random state x_t and pays out a reward depending on the state. As in a MDP, successive MAB control actions produce a sequence of actions and states. When the MAB action-state sequence is Markovian it is a special case of a MDP or POMDP process.

In some cases, the optimal policy for a k -arm MAB problem can be shown to reduce to a so-called index policy. An index policy is a simpler mapping that assigns a score (or index) to each arm of the MAB and pulls only the arm having maximum score at a given time. The key to the simplification is that these scores, the Gittins indices mentioned in Section III, can be determined by solving a much simpler set of k different single-armed bandit problems. Gittins index policies exist when the actions are not irrevocable; meaning that any available actions not taken at the present time can be deferred to the future, producing the same sequence of future

rewards, except for a discount factor. The significance of Gittins index policies is that they are frequently much simpler to compute than backwards induction solutions to optimal policies for MDPs and POMDPs. Thus they are sometimes used to approximate these optimal policies; e.g., using rollout with MAB index-rules as the base policy [75]. See [1, Ch. 6] for further discussion of index policies and their variants.

As a simple example, consider the aforementioned wide area search problem for the case of a single non-moving target that could be located in one of k locations with equal probability. Assume that in each time epoch a sensor can look at a single location with specified probabilities of correct detection and false alarm. Further assume that the reward is decreasing in the amount of time required by the sensor to correctly find the target. Identify each sensing action (location) as an arm of the MAB and the un-normalized posterior probability of the true target location as the state of the MAB. Under these assumptions, the optimal MAB policy for selecting arms is an index policy and specifies the optimal wide area search scheduler. For further details on this application of MAB to sensor management see [1, Ch. 7].

Bandit models were proposed for search problems like the above several decades ago [76], but their application to sensor management is relatively recent. Early applications of the multi-armed bandit model to sensor management were Krishnamurthy’s treatment of the radar beam scheduling for multiple target tracking problem [20], [21] and Washburn *et al.*’s application to general problems of sensor resource management [52]. As another example, arm acquiring bandits have been proposed by Washburn [1, Ch. 7] for tracking targets that can appear or disappear from the scene. Also discussed in [1, Ch. 7] are restless bandits, multi-armed bandits in which the states of the arms not played can evolve in time. Sensor management application of restless bandits include radar sensor management for multi-target tracking (see, e.g., [77]).

D. Information-optimized decision processes

The MDP/POMDP and MAB approaches to sensor management involve searching over multi-stage look-ahead policies. Designing a multi-stage policy requires evaluating each available action in terms of its impact on the potential rewards for all future actions. Myopic sensor management policies have been investigated as low complexity alternatives to multi-stage policies. Myopic policies only look ahead to the next stage; i.e., they compute the expected reward in the immediate future to determine the best current action. Such greedy policies

benefit from computational simplicity, but at the expense of performance loss compared to multi-stage optimal policies. Often this loss is significant. However, there are cases where the myopic loss approach gives acceptable performance, and indeed is almost optimal in special cases.

The most obvious way to obtain myopic sensor scheduling policies is to only consider the effect of the control action on the immediate reward; i.e., to truncate the future reward sequence in the multi-stage POMDP scheduling problem. This approach is called the optimal one-step look-ahead policy. However, it has often been observed that a myopic policy can achieve better overall performance by maximizing a surrogate reward, such as the mutual information between the data and the target. The information gain, discussed in more detail below, has the advantage that it is a more fundamental quantity than a task-specific reward function. For example, unlike many reward functions associated with estimation or detection algorithms, the mutual information is invariant to invertible transformations of the data. This and other properties lead to myopic policies that are more robust to factors such as model mismatch and dynamically changing system objectives (e.g., detection versus tracking), while ensuring a minimal level of system performance. For further motivation and properties of information theoretic measures for sensor management the reader may wish to consult [1, Ch. 3].

Information theoretic measures have a long history in sensor management. Optimization of Fisher information was applied to the related problem of optimal design of experiments (DOE) by Fisher, discussed in Section III, in the early part of the twentieth century [30]. Various functions of the Fisher information matrix, including its determinant and trace, have been used as reward functions for optimal DOE [38]. More recently, sensor management applications of optimal DOE have been proposed; e.g., in managed sensor fusion [78], in sensor managed unexploded ordnance (UXO) detection [56], in multi-sensor scheduling [79], and in sensor management for robotic vision and navigation [57]. However, Fisher information approaches to sensor management have several drawbacks. Notable among these are that the Fisher information requires specification of a parametric model for the observations. It is also a local measure of information that does not apply to discrete targets or mixtures of discrete and continuous valued targets. Model mismatch and/or discrete valued quantities frequently arise in sensor management applications. For example, discrete values arise when there is categorical side information about the target or clutter, or a target that transitions between two states like stopping and moving.

These are principal reasons that non-local information measures such as entropy and mutual information have become more common in sensor management.

In his 1998 PhD thesis [7], McIntyre cites the work of Barker [80] and Hintz and McVey [81] as the first to apply entropy to sensor management problems in 1977 and 1991, respectively. However, while the problems they treated are special cases of sensor management, they did not treat the general sensor management problem nor did they use the term in their papers. The first papers we know of that applied entropy measures explicitly to sensor management were Manyika and Durrant-Whyte [78] in 1992 and Schmaedeke [55] in 1993. The information measure used in these papers was the expected update in posterior entropy, called the information gain, that is associated with a given candidate sensor action.

These early information theoretic sensor management papers assumed Gaussian observations and linear dynamics, in which case the entropy and information gain have closed form mathematical expressions. Subsequently, the linear-Gaussian assumptions have been relaxed by using non-parametric estimation of entropy and information gain. Other information gain measures have also been introduced, such as the Kullback-Leibler (KL) divergence, the KL discrimination, and the Rényi entropy. The reader can consult the book [1] and, in particular, early papers by Schmaedeke and K. Kastella [82], [83], R. Mahler [84], Hintz and McIntyre [85], and Kreucher *et al.* [86], [87].

At first information gain sensor management methods were focused on single modality tracking of simple passive targets. In recent years, information gain has been applied to increasingly general models and sensor management tasks. For example, information driven methods have been applied to dynamic collaborative sensing with communication costs [88], multi-sensor information fusion [89], target tracking with uncertain sensor responses [90], multi-target tracking in large dynamic sensor networks [91], multi-modality multi-target tracking with time varying attenuation and obscuration [92], [67], robot path planning [93], and active camera control for object recognition and tracking using mutual information [94].

A striking mathematical result on the capabilities of information driven sensor management was obtained by J. L. Williams *et al.* [95] in connection with the general problem of information gathering in the context of graphical models under the assumption of conditionally independent measurements. In 2005 Guestrin *et al.* [96] showed that the conditional mutual information is submodular in the context of general machine learning problems. The significance of this result for sensor

management is that the maximizer of a submodular objective function can be well approximated using greedy optimization algorithms. Using this insight, Williams *et al.* established in [95] that greedy sequential methods for measurement planning are guaranteed to perform within a factor of $1/2$ of the optimal multi-stage selection method. Furthermore, this bound is independent of the length of the planning horizon and is sharp. The remarkable results of [95] are significant in that they provide theoretical justification for the computationally simpler myopic strategy and provide the designer with a tool to gauge the expected loss with respect to the optimal, but intractable, multi-stage policy. The bound was used to design resource constrained, information driven sensor management algorithms that exploit the submodularity property. The algorithm monotonically reduces an upper bound on the optimal solution that permits the system designer to terminate computation early with a near-optimal solution. These results are further elaborated in [12], [97], [98], [95].

V. OPPORTUNITIES ON THE HORIZON

Despite intensive research activity over the past fifteen years, and particularly in the past decade, formidable challenges remain to be addressed in order for sensor management to be genuinely viable in large-scale sensing systems. A central issue is computational feasibility of even approximate methods when scaled to problems that involve large numbers of controllable parameters, pose acute time constraints, or can only be adequately addressed by methods that look multiple steps ahead.

One arena of current investigation seeking to address the complexity issue involves sparse convex optimization approaches. The selection of an action sequence among a large number of possible sequences is similar to variable selection in sparse (lasso) regression [99] and compressive sensing [100], among other areas. This insight led R. Rangarajan *et al.* [101] to apply convex relaxation to optimal waveform design. A similar approach was later applied by S. Joshi and S. Boyd [102] to sensor selection. The use of such convex relaxation principles to develop tractable approximations to more complex sensor management combinatorial optimization problems, such as multi-stage planning, may lead to computational breakthroughs.

Another circle of current research offering some promise with regard to mitigating complexity involves the use of statistical machine learning tools. Often difficult problems in one domain can be reduced to equivalent problems in another domain for which different and effective solution tools have been developed. For example, the celebrated boosting method of Y. Freund

and R. Schapire for learning optimal classifiers [103] was directly motivated by optimal multi-armed bandit strategies. Conversely, by casting offline learning of optimal POMDP policies as an equivalent problem of learning optimal classifiers [104], [105], D. Blatt *et al.* [106] developed a boosting approach to learning optimal sensor management policies for UXO and radar sensing applications. It is likely that other advances in statistical machine learning can have positive impact on sensor management.

The authors are aware of ongoing research involving new approximation schemes that adaptively partition the information state space in an MDP problem in a way that allows controllable tradeoff of computational efficiency and approximation fidelity. This work, as yet unpublished, casts fidelity in terms of preserving the ranking of possible actions in terms of expected loss rather than preserving the actual values of expected loss.

The area of adversarial sensor management, which deals with situations where an adversary can control some aspects of the scenario to deliberately confound the sensor manager's objectives, presents opportunities for new sensor management research directions involving game theory and other methods. Recent work on POMDP for smart targets, i.e., targets that can react when they sense that they are being probed, is a step in this direction [107],[108]. Adversarial multi-armed bandits [109] and game theoretic solutions to adversarial multimobile sensing have also been proposed [110]. However, there are presently very few fundamental results on performance in adversarial environments; e.g., generalizations of the non-adversarial bounds of Castañón [111] and Williams [12] for POMDPs or those of K. D. Glazebrook and R. Minty for multi-armed bandits [112].

VI. CONCLUDING REMARKS

In this overview article on sensor management we have described the primary models and methods around which recent research in the field has been centered. We have also attempted to expose the historical roots in classical work spanning sequential analysis, optimal design of experiments, information theory, and optimal control. In our discussion of current trends and future research opportunities, we point out formidable challenges to achieving the performance gains in real-world systems that we believe are potentially possible. The computational viability of scaling the methods described in this paper to large-scale problems involving sensing systems with many controllable parameters, applications with fast operating tempos, and scenarios calling for non-myopic optimization depends upon substantial advances

in efficient and certifiable approximation in all the main components depicted in Figure 1.

Nevertheless, there is much room for optimism. The past two decades have seen intense research activity that has legitimized sensor management as a field of study and established its mathematical foundations. These have drawn on, adapted, and blended ideas from several established areas, including Markov decision processes, multi-armed bandit scheduling, and information gain myopic planning. The applications and technological advances that spurred the profound growth of interest in sensor management during this period continue to provide more and more opportunities for sensor management, and in some cases demand it. In our own application regime of surveillance and reconnaissance for security and defense applications, future operational concepts envision increasingly versatile networked collections of sensor assets, a large fraction of which will be mounted on autonomous or semi-autonomous platforms, providing situational awareness at levels of abstraction considerably higher than target tracks and emission source localizations. We remain hopeful that a combination of significant incremental advances and *bona fide* breakthroughs will enable sensor management to rise to meet such visions.

In closing, we wish to acknowledge the role of numerous sponsored research programs that have enabled and shaped the development of sensor management over the past decade. Some such activities of which we are aware include DARPA's *Integrated Sensing and Processing* and *Waveforms for Active Sensing* programs, which ran from 2001 through 2006. The U.S. Department of Defense has also invested in academic research in sensor management through several Multidisciplinary University Research Initiatives (MURIs) since the early 2000s. These have been managed by DARPA, the U.S. Air Force Office of Scientific Research (AFOSR), and the U.S. Army Research Office (ARO). We are also aware of sponsored work through the U.S. Air Force Research Laboratory, the Australian Defence Science and Technology Organisation, a few other government agencies, and several industrial sources. This list is by no means comprehensive, but it illustrates the recognition of sensor management as a valuable emerging area of study by major research organizations. Further, this trend is ongoing. For example, two new MURI projects related to sensor management have recently been initiated, one by AFOSR in 2010 entitled *Control of Information Collection and Fusion* and the most recent by ARO in 2011 entitled *Value of Information for Distributed Data Fusion*.

REFERENCES

- [1] A. Hero, D. A. Castañón, D. Cochran, and K. Kastella, *Foundations and Applications of Sensor Management*. New York: Springer, 2008.
- [2] A. R. Cassandra, "A survey of POMDP applications," in *Working Notes of AAAI 1998 Fall Symposium on Planning with Partially Observable Markov Decision Processes*. Citeseer, 1998, pp. 17–24.
- [3] G. E. Monahan, "A survey of partially observable Markov decision processes: theory, models and algorithms," *Management Science*, vol. 28, no. 1, pp. 1–16, January 1982.
- [4] G. W. Ng and K. H. Ng, "Sensor management — What, why and how," *Information Fusion*, vol. 1, no. 2, pp. 67–75, 2000.
- [5] X.-X. Liu, S.-L. Shen, and Q. Pan, "A survey of sensor management and methods," *Acta Electronica Sinica*, vol. 30, no. 3, pp. 394–398, 2002.
- [6] D. Cochran, "Waveform-agile sensing: Opportunities and challenges," in *Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc.*, vol. V, March 1995, pp. 877–880.
- [7] G. A. McIntyre, "A comprehensive approach to sensor management and scheduling," Ph.D. dissertation, George Mason University, 1998.
- [8] D. Sinno, "Attentive management of configurable sensor systems," Ph.D. dissertation, Arizona State University, May 2000.
- [9] C. Kreucher, "An information-based approach to sensor resource allocation," Ph.D. dissertation, EECS Department, University of Michigan, 2005.
- [10] R. Rangarajan, "Resource constrained adaptive sensing," Ph.D. dissertation, EECS Department, University of Michigan, 2006.
- [11] D. Blatt, "Performance evaluation and optimization for inference systems: Model uncertainty, distributed implementation, and active sensing," Ph.D. dissertation, EECS Department, University of Michigan, 2007.
- [12] J. L. Williams, "Information theoretic sensor management," Ph.D. dissertation, Massachusetts Institute of Technology, 2007.
- [13] M. Huber, *Probabilistic Framework for Sensor Management*. KIT Scientific Publishing, 2009.
- [14] K. L. Jenkins, "Fast adaptive sensor management for feature-based classification," Ph.D. dissertation, Boston University, 2010.
- [15] M. A. Richards, *Fundamentals of Radar Signal Processing*. McGraw-Hill, 2005.
- [16] D. J. Kershaw and R. J. Evans, "Optimal waveform selection for tracking systems," *IEEE Trans. on Inform. Theory*, vol. 40, no. 5, pp. 1536–1550, September 1994.
- [17] —, "Waveform selective probabilistic data association," *IEEE Trans. on Aerosp. and Electron. Systems*, vol. 33, no. 4, pp. 1180–1188, October 1997.
- [18] S. Sowelam and A. H. Tewfik, "Optimal waveform selection in range-doppler imaging," in *IEEE Int. Conf. on Image Processing*, November 1994, pp. 441–445.
- [19] S. M. Sowelam and A. H. Tewfik, "Optimal waveform selection for radar target classification," in *IEEE Int. Conf. on Image Processing*, October 1997, pp. 476–479.
- [20] V. Krishnamurthy and R. J. Evans, "Hidden Markov model multiarm bandits: a methodology for beam scheduling in multitarget tracking," *IEEE Trans. on Signal Processing*, vol. 49, no. 12, pp. 2893–2908, 2001.
- [21] —, "Correction to 'Hidden Markov model multiarm bandits: a methodology for beam scheduling in multitarget tracking'," *IEEE Trans. on Signal Processing*, vol. 51, no. 6, pp. 1662–1663, 2003.

- [22] N. A. Goodman, P. R. Venkata, and M. A. Neifeld, "Adaptive waveform design and sequential hypothesis testing for target recognition with active sensors," *IEEE Journ. Selected Topics in Signal Processing*, vol. 1, no. 1, pp. 105–113, June 2007.
- [23] S. P. Sira, Y. Li, A. Papandreou-Suppappola, D. Morrell, D. Cochran, and M. Rangaswamy, "Waveform-agile sensing for tracking," *IEEE Signal Processing Magazine*, vol. 26, no. 1, pp. 53–64, 2009.
- [24] M. Greco, F. Gini, P. Stinco, A. Farina, and L. Verrazzani, "Adaptive waveform diversity for cross-channel interference mitigation," in *Proc. IEEE Radar Conf.*, May 2008, pp. 1–6.
- [25] S. P. Sira, D. Cochran, A. Papandreou-Suppappola, D. Morrell, W. Moran, S. D. Howard, and R. Calderbank, "Adaptive waveform design for improved detection of low-RCS targets in heavy sea clutter," *IEEE Journ. Selected Topics in Signal Processing*, vol. 1, no. 1, pp. 56–66, June 2007.
- [26] A. Leshem, O. Naparstek, and A. Nehorai, "Information theoretic adaptive radar waveform design for multiple extended targets," *IEEE Journ. Selected Topics in Signal Processing*, vol. 1, no. 1, pp. 42–55, 2007.
- [27] M. Vespe, G. Jones, and C. J. Baker, "Lessons for radar," *IEEE Signal Processing Magazine*, vol. 26, no. 1, pp. 65–75, 2009.
- [28] D. Cochran, S. Suvorova, S. D. Howard, and B. Moran, "Waveform libraries," *IEEE Signal Processing Magazine*, vol. 26, no. 1, pp. 12–21, 2009.
- [29] D. Sinno and D. Cochran, "Estimation with configurable and constrained sensor systems," in *Defence Applications of Signal Processing*, D. Cochran, W. Moran, and L. B. White, Eds. Elsevier, 1998.
- [30] R. A. Fisher, *The Design of Experiments*. Oliver and Boyd, Edinburgh, 1935.
- [31] H. Robbins, "Some aspects of the sequential design of experiments," *Bulletin of the American Mathematical Society*, 1952.
- [32] H. F. Dodge and H. G. Romig, "A method of sampling inspection," *Bell System Technical Journal*, vol. 8, pp. 613–631, 1929.
- [33] A. Wald, *Sequential Analysis*. New York: Wiley, 1947.
- [34] C. W. Helstrom, *Elements of Signal Detection and Estimation*. Prentice Hall, 1994.
- [35] H. Chernoff, "Locally optimal designs for estimating parameters," *Ann. Math. Stat.*, vol. 24, no. 4, pp. 586–602, December 1953.
- [36] —, "Sequential design of experiments," *Ann. Math. Stat.*, vol. 30, no. 3, pp. 755–770, September 1959.
- [37] —, *Sequential analysis and optimal design*. SIAM, 1972.
- [38] V. V. Fedorov, *Theory of Optimal Experiments*. Academic Press, 1972.
- [39] L. Meier III, J. Perschon, and R. M. Dressier, "Optimal control of measurement subsystems," *IEEE Trans. Automatic Control*, vol. 12, no. 5, pp. 528–536, October 1967.
- [40] M. Athans, "On the determination of optimal costly measurement strategies for linear stochastic systems," *Automatica*, vol. 8, no. 4, pp. 397–412, July 1972.
- [41] J. R. Landoll and C. A. Caceres, "Automation of data acquisition in patient testing," *Proceedings of the IEEE*, vol. 57, no. 11, pp. 1941–1953, 1969.
- [42] S. G. Bier, P. L. Rothman, and R. A. Manske, "Intelligent sensor management for beyond visual range air-to-air combat," in *Proc. IEEE National Aerosp. and Electron. Conf. (NAECON)*, vol. 1, May 1988, pp. 264 – 269.
- [43] S. Musick and R. Malhotra, "Chasing the elusive sensor manager," in *Proceedings of the IEEE 1994 National Aerospace and Electronics Conference (NAECON)*. IEEE, 1994, pp. 606–613.
- [44] R. Malhotra, E. P. Blasch, and J. D. Johnson, "Learning sensor-detection policies," in *Proc. IEEE National Aerosp. and Electron. Conf. (NAECON)*, vol. 2. IEEE, 1997, pp. 769–776.
- [45] J. Aloimonos, I. Weiss, and A. Bandyopadhyay, "Active vision," *International Journal of Computer Vision*, vol. 1, no. 4, pp. 333–356, 1988.
- [46] A. Segall, "Optimal control of measurement subsystems," *IEEE Trans. Automatic Control*, vol. 22, no. 2, pp. 179–186, April 1977.
- [47] D. A. Castañón, "Approximate dynamic programming for sensor management," in *Proc. of IEEE Conf. Decision and Control*, vol. 2, 1997, pp. 1202–1207.
- [48] J. S. Evans and V. Krishnamurthy, "Optimal sensor scheduling for hidden Markov model state estimation," *Int. J. Contr.*, vol. 74, no. 18, pp. 1737–1742, December 2001.
- [49] V. Krishnamurthy, "Algorithms for optimal scheduling and management of hidden Markov model sensors," *IEEE Trans. on Signal Processing*, vol. 50, no. 6, pp. 1382–1397, 2002.
- [50] J. C. Gittins, "Bandit processes and dynamic allocation indices," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 41, no. 2, pp. 148–177, 1979.
- [51] P. Whittle, "Multi-armed bandits and the Gittins index," *Journ. Royal Statist. Soc. B*, vol. 42, no. 2, pp. 143–149, 1980.
- [52] R. B. Washburn, M. K. Schneider, and J. J. Fox, "Stochastic dynamic programming based approaches to sensor resource management," in *Proc. Int. Conf. on Inf. Fusion*, vol. 1, 2002, pp. 608–615.
- [53] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, July 1948.
- [54] J. Manyika and H. Durrant-Whyte, *Data fusion and sensor management: a decentralized information-theoretic approach*. Prentice Hall PTR Upper Saddle River, NJ, USA, 1995.
- [55] W. W. Schmaedeke, "Information-based sensor management," in *Proc. Meeting of Intl. Soc. for Optical Engin. (SPIE)*, vol. 156, 1993.
- [56] Y. Zhang, X. Liao, and L. Carin, "Detection of buried targets via active selection of labeled data: Application to sensing subsurface UXO," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 11, pp. 2535–2543, 2004.
- [57] K. Tian and G. Zhu, "Sensor management based on Fisher information gain," *Journal of Systems Engineering and Electronics*, vol. 17, no. 3, pp. 531–534, 2006.
- [58] V. Gupta, T. H. Chung, B. Hassibi, and R. M. Murray, "On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage," *Automatica*, vol. 42, no. 2, pp. 251–260, February 2006.
- [59] —, "Sensor scheduling algorithms requiring limited computation," in *Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc.*, vol. 3, May 2004, pp. 825–828.
- [60] T. H. Chung, V. Gupta, B. Hassibi, J. Burdick, and R. M. Murray, "Scheduling for distributed sensor networks with single sensor measurement per time step," in *Proc. IEEE Int. Conf. on Robotics and Automation*, vol. 1, April 2004, pp. 187–192.
- [61] K. L. Jenkins and D. A. Castañón, "Information-based adaptive sensor management for sensor networks," in *Proc. American Control Conf.*, June 2010, pp. 4934–4940.
- [62] —, "Adaptive sensor management for feature-based classification," in *Proc. of IEEE Conf. Decision and Control*, December 2010, pp. 522–527.
- [63] D. Hitchings and D. A. Castañón, "Receding horizon stochastic control algorithms for sensor management," in *Proc. American Control Conf.*, June 2010, pp. 6809–6815.
- [64] —, "Adaptive sensing for search with continuous actions and observations," in *Proc. of IEEE Conf. Decision and Control*, December 2010, pp. 7443–7448.

- [65] E. J. Sondik, "The optimal control of partially observable Markov processes," Ph.D. dissertation, Stanford University, 1971.
- [66] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, vol. 101, pp. 99–134, 1998.
- [67] E. K. P. Chong, C. Kreucher, and A. Hero, "Partially observable Markov decision process approximations for adaptive sensing," *Discrete Event Systems*, vol. 19, no. 3, pp. 377–422, September 2009.
- [68] Y. Li, L. W. Krakow, E. K. P. Chong, and K. N. Groom, "Approximate stochastic dynamic programming for sensor scheduling to track multiple targets," *Digital Signal Processing*, vol. 19, no. 6, pp. 978–989, 2009.
- [69] J. Liu, M. Chu, and J. E. Reich, "Multitarget tracking in distributed sensor networks," *IEEE Signal Processing Magazine*, vol. 24, no. 3, pp. 36–46, 2007.
- [70] C. Kreucher, K. Kastella, and A. Hero, "Multitarget tracking using a particle filter representation of the joint multitarget probability density," *IEEE Trans. on Aerosp. and Electron. Systems*, vol. AES-39, no. 4, pp. 1396–1414, 2005.
- [71] V. Krishnamurthy and D. V. Djonin, "Optimal threshold policies for multivariate POMDPs in radar resource management," *IEEE Trans. on Signal Processing*, vol. 57, no. 10, pp. 3954–3969, 2009.
- [72] S. Ji, R. Parr, and L. Carin, "Nonmyopic multiaspect sensing with partially observable Markov decision processes," *IEEE Trans. on Signal Processing*, vol. 55, no. 6, pp. 2720–2730, 2007.
- [73] T. Hanselmann, M. Morelande, B. Moran, and P. Sarunic, "Sensor scheduling for multiple target tracking and detection using passive measurements," in *Proc. Int. Conf. on Information Fusion*, 2008, pp. 1–8.
- [74] M. J. Kochenderfer, L. P. Espindle, J. K. Kuchar, and J. D. Griffith, "A comprehensive aircraft encounter model of the national airspace system," *Lincoln Laboratory Journal*, vol. 17, no. 2, pp. 41–53, 2008.
- [75] D. P. Bertsekas and D. A. Castañón, "Rollout algorithms for stochastic scheduling problems," *Journal of Heuristics*, vol. 5, no. 1, pp. 89–108, 1999.
- [76] S. J. Benkoski, M. G. Monticino, and J. R. Weisinger, "A survey of the search theory literature," *Naval Research Logistics*, vol. 38, no. 4, pp. 469–494, 1991.
- [77] B. F. La Scala and B. Moran, "Optimal target tracking with restless bandits," *Digital Signal Processing*, vol. 16, no. 5, pp. 479–487, 2006.
- [78] J. M. Manyika and H. F. Durrant-Whyte, "Information-theoretic approach to management in decentralized data fusion," in *Proc. Meeting of Intl. Soc. for Optical Engin. (SPIE)*, vol. 202, 1992.
- [79] M. L. Hernandez, T. Kirubarajan, and Y. Bar-Shalom, "Multisensor resource deployment using posterior Cramér-Rao bounds," *IEEE Trans. on Aerosp. and Electron. Systems*, vol. 40, no. 2, pp. 399–416, 2004.
- [80] W. H. Barker, "Information theory and the optimal detection search," *Operations Research*, vol. 25, no. 2, pp. 304–314, 1977.
- [81] K. J. Hintz and E. S. McVey, "Multi-process constrained estimation," *IEEE Trans. on Systems, Man, and Cybernetics*, vol. 21, no. 1, pp. 434–442, January/February 1991.
- [82] W. Schmaedeke and K. Kastella, "Event-averaged maximum likelihood estimation and information-based sensor management," in *Proc. Meeting of Intl. Soc. for Optical Engin. (SPIE)*, vol. 2232, Orlando, June 1994, pp. 91–96.
- [83] K. Kastella, "Discrimination gain to optimize detection and classification," *IEEE Trans. on Systems, Man and Cybernetics, Part A: Systems and Humans*, vol. 27, no. 1, pp. 112–116, 1997.
- [84] R. Mahler, "Global optimal sensor allocation," in *Proceedings of the Ninth National Symposium on Sensor Fusion*, vol. 1, 1996, pp. 167–172.
- [85] K. J. Hintz and G. A. McIntyre, "Information instantiation in sensor management," in *Proceedings of the SPIE International Symposium on Aerospace/Defense Sensing & Control, Orlando, FL*, vol. 3374, 1998, pp. 38–47.
- [86] C. Kreucher, K. Kastella, and A. O. Hero, "Information based sensor management for multitarget tracking," in *Proc. Meeting of Intl. Soc. for Optical Engin. (SPIE)*, San Diego, 2003.
- [87] —, "Multi-target sensor management using alpha-divergence measures," in *3rd Workshop on Information Processing for Sensor Networks*, Palo Alto, CA, 2003.
- [88] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration for tracking applications," *IEEE Signal Processing Magazine*, vol. 19, no. 2, pp. 61–72, 2002.
- [89] N. Xiong and P. Svensson, "Multi-sensor management for information fusion: Issues and approaches," *Information Fusion*, vol. 3, no. 2, pp. 163–186, 2002.
- [90] M. P. Kolba and L. M. Collins, "Information-based sensor management in the presence of uncertainty," *IEEE Trans. on Signal Processing*, vol. 55, no. 6, pp. 2731–2735, 2007.
- [91] C. M. Kreucher, A. O. Hero, K. D. Kastella, and M. R. 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.
- [92] C. Kreucher, K. Kastella, and A. O. Hero, "Sensor management using an active sensing approach," *Signal Processing*, vol. 85, no. 3, pp. 607–624, 2005.
- [93] G. Zhang, S. Ferrari, and M. Qian, "An information roadmap method for robotic sensor path planning," *Journal of Intelligent and Robotic Systems*, vol. 56, no. 1, pp. 69–98, 2009.
- [94] J. Denzler and C. M. Brown, "Information theoretic sensor data selection for active object recognition and state estimation," *IEEE Trans. on Pattern Anal. and Machine Intell.*, vol. 24, no. 2, pp. 145–157, 2002.
- [95] J. L. Williams, J. W. Fisher III, and A. S. Willsky, "Performance guarantees for information theoretic active inference," in *Proc. Int. Conf. on Artificial Intelligence and Statistics*, M. Meila and X. Shen, Eds., March 2007, pp. 616–623.
- [96] A. Krause and C. Guestrin, "Near-optimal nonmyopic value of information in graphical models," in *Uncertainty in Artificial Intelligence*, July 2005.
- [97] J. L. Williams, J. W. Fisher III, and A. S. Willsky, "Approximate dynamic programming for communication-constrained sensor network management," *IEEE Transactions on Signal Processing*, vol. 55, no. 8, pp. 3995–4003, August 2007.
- [98] —, "A constraint generation integer programming approach to information theoretic sensor resource management," in *Proc. IEEE Workshop on Statistical Signal Processing*, August 2007.
- [99] T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York: Springer, 2001.
- [100] E. Candés, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Comm. Pure Appl. Math.*, vol. 59, pp. 1207–1223, 2005.
- [101] R. Rangarajan, R. Raich, and A. O. Hero, "Single-stage waveform selection for adaptive resource constrained state estimation," in *Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc.*, May 2006.
- [102] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Trans. on Signal Processing*, vol. 57, no. 2, pp. 451–462, 2009.

- [103] Y. Freund and R. Schapire, “A decision theoretic generalization of online learning and an application to boosting,” *Journ. of Computer and System Sciences*, vol. 55, pp. 119–139, 1997.
- [104] J. Langford and B. Zadrozny, “Reducing T-step reinforcement learning to classification,” in *Proceedings of the Machine Learning Reductions Workshop*, 2003.
- [105] D. Blatt and A. Hero, “From weighted classification to policy search,” in *Advances in Neural Information Processing Systems 18*, Y. Weiss, B. Schölkopf, and J. Platt, Eds. Cambridge, MA: MIT Press, 2006, pp. 139–146.
- [106] D. Blatt and A. O. Hero, “Optimal sensor scheduling via classification reduction of policy search (CROPS),” in *International Conference on Automated Planning and Scheduling*, 2006.
- [107] C. Kreucher, A. O. Hero, D. Blatt, and K. Kastella, “Adaptive multi-modality sensor scheduling for detection and tracking of smart targets,” in *Workshop on Defense Applications of Signal Processing*, 2004.
- [108] B. Liu, C. Ji, Y. Zhang, and C. Hao, “Blending sensor scheduling strategy with particle filter to track a smart target,” *Wireless Sensor Networks*, no. 4, November 2009.
- [109] T. Uchiya, A. Nakamura, and M. Kudo, “Algorithms for adversarial bandit problems with multiple plays,” in *Algorithmic Learning Theory*. Springer, 2010, pp. 375–389.
- [110] M. Wei, G. Chen, E. Blasch, H. Chen, and J. B. Cruz, “Game theoretic multiple mobile sensor management under adversarial environments,” in *Proc. Int. Conf. on Inf. Fusion*, 2008, pp. 1–8.
- [111] D. Castañón, “Stochastic control bounds on sensor network performance,” in *Proc. of IEEE Conf. Decision and Control and European Control Conf.* IEEE, 2005, pp. 4939–4944.
- [112] K. D. Glazebrook and R. Minty, “A generalized Gittins index for a class of multiarmed bandits with general resource requirements,” *Mathematics of Operations Research*, vol. 34, no. 1, pp. 26–44, 2009.