CONTRIBUTED
P A P E R

Cognitive Radar: Step Toward Bridging the Gap Between Neuroscience and Engineering

This paper describes a radar system that emulates the way the visual brain observes its environment; it outperforms traditional radar and achieves very smooth waveform switching.

By Simon Haykin, *Fellow IEEE*, Yanbo Xue, *Member IEEE*, and Peyman Setoodeh, *Member IEEE*

ABSTRACT | In this paper, we describe a cognitive radar (CR) that mimics the visual brain. Although the visual brain and radar are different in that the visual brain does not transmit a probing signal to the environment while the active radar greatly relies on the probing signal it transmits to the environment, both of them are observers of the surrounding environment. As such, there is much that we can learn from the visual brain in building a new generation of CRs that outperform traditional radars. In this paper, we confine the discussion, in both analytic and experimental terms, to CR aimed at target tracking. From a theoretical perspective, using the posterior Cramér–Rao lower bound (PCRLB), it is shown that a cognitive tracking radar has the potential to improve tracking performance significantly. In particular, computer experiments are presented, which demonstrate that CR can indeed go beyond the theoretical limits of traditional active radars (TARs) as well as fore-active radars (FARs); the latter are radars equipped with feedback from the receiver to the transmitter. Moreover, computer experiments are presented to demonstrate another practical benefit resulting from the combined use of memory and executive attention in CR for a target-tracking application. Specifically, it is shown that with the provision of these two cognitive processes, the transition in switching from one transmit waveform to another goes forward in a smooth manner. Such a capability is beyond that of TAR or FAR.

KEYWORDS | Artificial intelligence; attention; cognitive radar (CR); fore-active radar (FAR); memory; perception–action cycle; posterior Cramér–Rao lower bound (PCRLB); tracking; traditional active radar (TAR)

I. INTRODUCTION

The visual brain is a powerful parallelized information-processing machine with a built-in capability to perform certain tasks such as target detection (recognition) reliably and target tracking accurately, at speeds far beyond the capability of any traditional radar system in existence today. Despite the many differences that exist between the visual brain and traditional radar, they do share a common feature: they are both observers of the outside world (environment). Radar is an active sensor in the sense that it transmits a probing signal to the environment and then builds a picture of the environment via a decision-making process based on the resulting radar returns. As an active sensor, radar is engaged in both perceiving radar returns and actuating the phenomenon that gives rise to the radar returns [1]. The key question is: what can we do to learn from the visual brain to significantly enhance the information-processing power of a traditional radar? The answer to this fundamental question lies in *cognition*, hence the innovative idea of the *cognitive radar (CR)*, which was inspired by the echo-locating system of bat in

Manuscript received March 30, 2012; accepted May 9, 2012. Date of publication July 31, 2012; date of current version October 16, 2012. This paper was supported by the National Science and Engineering Research Council (NSERC) of Canada.
The authors are with the Cognitive Systems Laboratory, McMaster University, Hamilton, ON L8S 4K1 Canada (e-mail: haykin@mcmaster.ca; yxue@ieee.org; setoodeh@ieee.org).
Digital Object Identifier: 10.1109/JPROC.2012.2203089

[2] and, then, the visual brain in [3]. The reason for the focused attention on the visual brain is the enormous literature in neuroscience on the visual system.

To proceed from a traditional radar to CR, the first step in this evolutionary process described in the literature may be credited to Kershaw and Evans [4]. In that paper, a link from the receiver to the transmitter was added, albeit in an offline manner, thereby modifying the radar into a closed-loop feedback control system. Hereafter, this kind of radar is referred to as a fore-active radar (FAR).

We may thus distinguish three classes of radar as follows.

1) *Traditional active radar (TAR)*, which operates in a feedforward manner; TAR may involve the use of adaptive filtering (e.g., Kalman filtering) for iterative state estimation in the receiver [5], as well as adaptive beamforming for matching the transmitted waveform to the environment [6].
2) *Fore-active radar (FAR)*, which distinguishes itself in using a feedback link connecting the receiver to the transmitter; in the radar literature, such a radar is said to be *fully adaptive*, in that it has global feedback that includes the environment within the feedback loop; and it could also involve adaptive filtering in the receiver and adaptive beamforming in the transmitter [7]. In control theory, it is known that feedback facilitates intelligence; we may therefore say that a fully adaptive radar has limited intelligence, which represents the first step toward cognition that is put to practical use in the transmitter for transmit-waveform selection [4], [7], [8] or resource allocation [9].
3) *Cognitive radar (CR)*, which differs from TAR as well as FAR by virtue of the following capability: the development of rules of behavior in a self-organized manner through a process called *learning from experience* that results from continued interactions with the environment [3].

A. Organization of the Paper

In this paper, we present a detailed study of CR that *mimics* the visual brain. The paper is organized as follows. Section II provides historical notes on human cognition, followed by Section III, which builds on the material covered in Section II and sets the stage for the formulation of CR. Section IV describes Fuster's paradigm of cognition. With Fuster's paradigm as the framework on which CR is based, Section V describes the engineering perspective of cognition. Then, in Section VI, we describe the perception–action cycle, followed by Section VII devoted to memory. In physical terms, the perception–action cycle and memory occupy distinct locations of their own within CR. Sections VIII and IX are devoted to two algorithmic mechanisms, namely attention and intelligence, respectively. Section X describes a cardinal characteristic of CR, namely the cyclic directed information-flow diagram. Sections XI, XII, and XIII are devoted to computational studies: the first two are aimed at demonstrating the information-processing power of CR as a target tracker, and the third one is aimed at the ways in which the underlying physical characteristic of transmit waveforms evolves as we move through different cycles of computation under different composition of CR. The paper concludes with summary and discussion in Section XIV.

II. HISTORICAL NOTES ON HUMAN COGNITION

A. Large-Scale Brain Networks in Cognition

Much of what we know about brain functions has relied on the following approach: simplistic mapping of cognitive constructs onto brain areas. The traditional approach to human cognition may be viewed as the *modular paradigm*, in which brain areas are postulated to act as independent processors to carry out specific functions.

However, a new paradigm is emerging in cognitive neuroscience, which goes beyond the modular paradigm. In this new paradigm, emphasis is placed on the following postulate [10]:

> Cojoint functions of brain areas, which work together as a large-scale network.

In this section, we highlight some important contributions that have been made to this new way of thinking about human cognition.

A large percentage of the information processing in the brain is performed in the cerebral cortex and it plays a key role in processes attributed to cognition by different researchers. Since the 1950s, Mountcastle's work on characterizing the columnar organization of the cerebral cortex has influenced research carried out in this field; he pointed out that cortical minicolumns are the basic functional units of cortex [11]. In 1978, based on the uniform appearance of the cortex, he proposed that all regions of the cortex may use a basic information-processing algorithm to accomplish their tasks [12]. This algorithm must be independent of the nature of the information-bearing sensory input. In other words, all kinds of sensory inputs (i.e., visual, audio, etc.) are coded in a standard form and fed to this basic processing algorithm [13].

As a pioneer in computational neuroscience, Marr followed a similar way of thinking that very few fundamental techniques are used by the cerebral cortex to process information for different tasks [14]. He was interested in developing a general computational theory for the brain based on biological evidence. However, in later years, he just focused on vision.

Following the same way of thinking and inspired by pioneering neuroscientists who had put a great deal of

thoughts and efforts into exploring what the essence of cognition is, Fuster proposed the concept of "cognit" for knowledge representation in the cerebral cortex; moreover, he proposed an abstract model for cognition based on five fundamental building blocks, namely perception, memory, attention, intelligence, and language [15].

Last, Sporns *et al.* [16] have broadened their view of the new paradigm for the conjoint functions of brain areas by considering the underpinnings of complex networks that cover a wide range of disciplines, which extend from biology to physics, and social sciences to informatics. In particular, they have addressed the issue of how the investigation of complex-network structures and dynamics could contribute to our understanding of brain and cognitive functions. To this end, they highlight a series of recent studies of complex brain networks, and attempt to identify information processing areas for future experimental and theoretical inquiry.

B. Cybernetics

Cybernetics is the study of the underlying structure of a *regulatory system* of interest, building primarily on information theory, control theory, and systems theory. For cybernetics to be applicable to a regulatory system, the system must reside inside a *closed feedback loop*, described as follows:

> Whenever the system acts on the environment in which it operates, the action taken by the system produces some change in the environment and that environmental change is fed back to the system in the form of information, whereafter the system is enabled to change its own behavior in a corresponding manner, and the cycle goes on.

This causal cyclic relationship between a regulatory system and its environment is a cardinal characteristic of cybernetics.

In the classic book entitled *Cybernetics* that was published in 1948, Wiener defined cybernetics as the study of control and communication in the animal and the machine [17]. Moreover, Wiener was an influential advocate of *automation* [18], the essence of which featured prominently in his second book entitled *The Human Use of Human Beings*, which was published in 1950. Needless to say, Wiener was a towering figure, widely considered to be the founding thinker of cybernetics, developed after the World War II.

To go on, McCulloch, neuropsychologist, and Pitts, mathematician, also contributed to the early development of cybernetics in their own pioneering ways. In their classic 1943 paper entitled "A logical calculus of the ideas immanent in nervous activity," they made five assumptions on the operation of neurons in the brain: the assumptions led to the formulation of a neural model widely known as the *McCulloch–Pitts neuron* [19], which is a binary device with a fixed threshold. As such, the neuron can only reside in one of two possible states. In particular, depending on the activities of its synapses, the neuron is either active or inactive; a synapse is a neurobiological link that connects one neuron to another. It can be justifiably said that the era of artificial neural networks started with the invention of the McCulloch–Pitts neuron.

Moving on to the last pioneering contribution that was made to the early development of cybernetics, we look to Walter for the construction of some of the *first autonomous, adaptive robots*, which are vividly described in his 1953 book on *The Living Brain* [20]. Walter had a passionate interest in brains as biological systems that have evolved through learning from the consequences of their own goal-directed actions, thereby resulting in complex behavior. In particular, he used this interest to build adaptive robots by incorporating into their design two cognitive functions: one being the formulation of a goal, and the other seeking the goal through a scanning mechanism. The end result of these endeavors was robots, namely tortoises, that would appear to simulate the most basic characteristics of animal and human behavior.

C. Artificial Neural Networks

The ability to learn from continuous interactions with the environment (world) is a distinctive characteristic of the brain. To model how the learning process can be mimicked artificially, a commonly used procedure is to train an *artificial neural network*,[1] made up of a set of computational units called *neurons*.

The input–output behavior of a neuron, denoted by k, in the network is mathematically described as follows:

$$u_k = \sum_{i=1}^{m} w_{ki} x_i$$

$$y_k = \phi(u_k + b_k)$$

where w_{ki} is an adjustable weight that connects node i to node k in the neural network, m is the number of weights, x_i is the input, u_k is an intermediate output, b_k is an adjustable bias, $\phi(\cdot)$ is a prescribed nonlinear function (e.g., sigmoid or hyperbolic tangent), and y_k is the overall output of the neuron.

Given a neural network, we may model the learning process algorithmically by using a supervised or unsupervised procedure as follows [22].

1) Supervised Learning or Learning With a Teacher: In conceptual terms, we may think of the teacher as having *knowledge* of the environment, with that knowledge being represented by a set of *input–output examples*. The environment is, however, unknown to the neural network.

[1]The handbook, edited by Arbib [21], provides the place occupied by neural networks, including their artificial counterparts, in the context of brain theory.

Suppose now that the teacher and the neural network are both exposed to a training vector, i.e., examples, drawn from the environment. By virtue of built-in knowledge, the teacher is able to provide the neural network with a desired response for that training vector; the desired response represents the *optimum* response to the training vector. The network parameters (i.e., weights of neurons in the network) are adjusted under the combined influence of the training vector, and an *error signal* defined as the difference between the desired response and the actual response of the network. Using an algorithm (e.g., the backpropagation algorithm), the network parameters are adjusted iteratively on an example-by-example basis until the neural network emulates the teacher accurately enough, at which point the training process is stopped. In effect, the environmental knowledge is transferred to the neural network through the training process and stored in the network's weights (parameters), representing *long-term memory*. When this condition is reached, we may then dispense with the teacher, thereby permitting the neural network to represent the environment completely by itself.

2) Self-Organized Learning: In this second form of learning, the requirement is for a neural network to discover significant features (patterns) of input data through the use of *unlabeled examples*. In other words, the neural network learns from the unlabeled examples without a teacher. To this end, we look to an *unsupervised* learning algorithm that is supplied with a set of *rules of local behavior*, where the term local means that the adjustments applied to the network's weights (parameters) are confined to the immediate local neighborhood of a particular neuron, and so it goes on. In such a scenario, the learning process tends to follow a neurobiological structure, recognizing that *network organization* is fundamental to the brain. It is here, we look to *Hebb's postulate of learning*, so called in honor of the psychologist Hebb [23]. In signal-processing terms, the postulate may be stated as follows.

1) If two neurons on either side of a synapse (connecting link) are activated simultaneously (i.e., synchronously), then the strength of that synapse is selectively increased.
2) If, on the other hand, the two neurons on either side of a synapse are activated asynchronously, then the synapse is selectively weakened or eliminated altogether.

Such a synapse is commonly called a *Hebbian synapse*.

What is remarkable is the fact that through synaptic amplification in accordance with Hebb's postulate of learning, together with two other principles of self-organization:

- cooperation between the neurons of a neural network, and
- structural information contained in the input data that is prerequisite to self-organized learning,

the neural network is enabled to extract the important features that characterize the input data, and do so without any form of supervision.

The important point to take from this brief discussion on learning using artificial neural networks is summed up as follows.

- For a neural network to learn in a supervised manner, the environment, from which the input–output examples are drawn, should be essentially *stationary*.
- On the other hand, a neural network, designed in a self-organized manner by exploiting Hebb's postulate of learning, has the built-in capability to *adapt* to statistical variations in the environment, provided that the rate of those statistical variations is slow compared to the rate at which the self-organized learning process is accomplished.

The brief *exposé* of artificial neural networks, just presented, is an integral part of *cognitive science*, which is an interdisciplinary subject devoted to the study of the mind and the many processes involved in it.

To summarize this section, particularly, large-scale brain networks, and cybernetics, there are many divergent views on how human cognition is described in the literature, which should not be surprising. This divergence makes the task of picking a basis rooted in human cognition for the study of CR that much more challenging.

III. WHERE DO WE GO FROM HERE?

At this juncture in the paper, it is instructive that we reflect back on the many points made in Section I and the lessons learned from Section II.

In Section I, we stressed the somewhat loose analogy between the visual brain and radar, hence the emphasis on learning from the visual brain to establish the basis for a new generation of radars, namely the *cognitive radar*, that goes beyond their traditional counterparts. With cognition as the driving force, we may look to neuroscience or cybernetics for the way in which the idea of cognition can be implemented. The logical source is neuroscience for the simple reason that CR so based can also account for relevant aspects in cybernetics, which is rationalized as follows. A cardinal characteristic of cognitive neuroscience is the perception–action cycle, which is, in fact, another way of describing the causal cyclic relationship that exists between a regulatory system and its environment in cybernetics.

Before going any further, it is equally instructive to describe the kind of CR that we have in mind. In this context, it can be argued that an adaptive control system, as we know it today, may well embody perception, attention, and intelligence. However, such a system lacks memory, which is known in neuroscience to play a critical function in cognition. It follows, therefore, for a radar system to be cognitive in the true sense of the word, it must embody all

the four functional blocks of cognition: perception–action cycle, memory, attention, and intelligence.

Moreover, the kind of CR we have in mind is one that goes along with the new paradigm in cognitive neuroscience, where emphasis is placed on cojoint functions working together under a single umbrella. To satisfy this requirement, we have opted for Fuster's vision on cognition as described in his book, entitled *Cortex and Mind: Unifying Cognition* [15], as the framework for reasons that will be discussed in Sections IV and V. At this stage in the paper, it suffices to say that the formulation of the CR theory, based on *Fuster's paradigm*, is validated by solid, ground-breaking results presented in the latter part of the paper. Accordingly, when we speak of CR, we have Fuster's paradigm in mind. Naturally, with Fuster's paradigm being rooted in neuroscience and CR being aimed at the engineering communities, we may describe CR, presented herein, as a brain-inspired engineering system.

IV. FUSTER'S PARADIGM OF COGNITION

First and foremost, this paper on CR is aimed at the engineering community, with the objective being that of describing an *orderly way*, according to which CR can be practically built. Moreover, with the objective also being that of mimicking the visual brain, it is important that we have an orderly structured framework, rooted in neuroscience, which satisfies the engineering need just mentioned.

In this paper, aimed for a single tracking radar, the current design focuses on the four most relevant underlying fundamental building blocks described by Fuster [15]:

1) *perception–action cycle*, which involves the propagation of feedback information about the environment from the receiver to the transmitter on a cyclic basis;
2) *memory*, part of which resides in the receiver and the other part resides in the transmitter, with these two parts being reciprocally coupled together; the *learning* process in CR is largely carried out in the memory;
3) *attention*, which is algorithmic in nature, with one part operating in the receiver and the other part operating in the transmitter;
4) *intelligence*, which is also algorithmic in nature but distributed throughout the radar system.

According to Fuster, memory is driven by the perception–action cycle; attention is memory driven; and intelligence builds on all three of them. It will be explained later that different building blocks of cognition are distributed in the system. Some of them such as perception–action cycle and memory involve physical elements. Others such as attention and intelligence do not occupy distinct physical spaces of their own. Rather, attention and intelligence manifest themselves algorithmically through interactions with memory and perception in a distributed manner. This is in agreement with the distributed processing in the brain [24], [25]. As mentioned before, in the neuroscience community, recently, there has been a shift of opinion in favor of associating cognition with dynamic interactions of distributed brain areas operating in large-scale networks rather than associating cognitive functions to particular brain areas [10].

Also, the concept of neural reuse has been introduced, in which it is justifiably argued that different brain areas can be redeployed and reused by different networks for different purposes [26].

V. ENGINEERING PERSPECTIVE OF COGNITION

For its formulation, Fuster's paradigm, rooted in cognitive neuroscience, has relied on innovative conceptualization of ideas, supported by detailed experimentations. With this paradigm as the frame of reference for CR, it is therefore not surprising to find that certain functional blocks of CR, namely memory and attention tend to be largely conceptualized at this early stage of the development of CR.

By building on the following concepts that are well developed in the engineering literature, we now have a strong basis for describing the perception–action cycle, which represents the first cardinal characteristic of CR:

- first, the Bayesian paradigm for probabilistic modeling of perceiving the radar environment in the receiver;
- second, Bellman's dynamic programming for the transmitter to optimally control the receiver via the radar environment;
- third, information theory for mathematical description of the link connecting the receiver to the transmitter.

Nevertheless, as already mentioned, research into cognition is in its early stages of development, be that in mathematical terms or practical applications. As such, there is much to be done on the following two topics:

- first, the important role that memory, and therefore learning, plays in the mathematical modeling and design of CR;
- second, a rigorous mathematical theory of CR, viewed as a complex system of systems.

Needless to say, both of theses two topics are beyond the scope of the paper.

Elaborating further on the Bayesian paradigm, there are two aspects of the paradigm that impact radar theory and design in a generic sense [27], [28]:

- hypothesis testing, which is basic to target detection;
- estimation theory, which is basic to target tracking.

Both of these processes are performed in the receiver. Owing to space limitation, the scope of the paper is confined to target tracking. Here again, when we speak of target tracking

in the context of the radar literature, we usually think of multiple-target tracking. However, the primary objective of this paper being that of demonstrating how cognition can be explained in improving radar performance and given space limitation as another constraint, the study of CR will be restricted to single-target tracking.

The following four sections describe the fundamental building blocks of cognition.

VI. PERCEPTION–ACTION CYCLE

Fig. 1 shows the block diagram of CR, with the perception–action cycle and memory occupying their distinctive places within the radar system [3]. As mentioned previously, both attention and intelligence are algorithmic mechanizations that are driven by the perception–action cycle and memory. Through waveform adaptation, which is performed in the perception–action cycle, CR gains control over certain aspects of the sensing process. Therefore, the perception–action cycle ties together estimation (through sensing) and control. Hence, cognition equips a radar with *controlled-sensing* ability to counteract the effect of nuisances (e.g., clutter) [29].

Perception, performed in the receiver, operates on measurements (i.e., radar returns) from unknown target with a dual objective: reliable detection of the target and tracking its behavior across time. As mentioned previously, herein we focus on target tracking, which is treated as a state-estimation problem under a Bayesian framework [2]. When the radar environment is nonlinear, we resort to an approximate form of Bayesian filtering [27].

To address the tracking problem, the traditional approach is to start with a state–space model that consists of a pair of equations: the system equation describes evolution of the state across time with system noise as the driving force, and the measurement equation describes dependence of the measurements on the state corrupted by measurement noise. Typically, but not always, the state–space model is nonlinear, which requires approximating the optimal Bayesian filter in some sense. Hereafter, we look to an *approximate Bayesian filter* to perform perception of the environment, which is naturally performed in the receiver.

The perception–action cycle requires that the receiver be linked to the transmitter. In a monostatic radar, where the receiver and the transmitter are collocated, such a requirement is relatively straightforward to handle. Accordingly, the receiver is also responsible for computing *feedback information* about the radar environment and then supplying it to the transmitter for action in the environment. With state estimation playing the role of perception in mathematical terms, it is logical to formulate the feedback information using the state-estimation error vector.

Next, turning to the transmitter, its function is to *control* the receiver indirectly through illumination of the environment. Here again, with optimal control in mind, we may look to *Bellman's dynamic programming* as the method of choice [2]. However, when dimensionality of the state space, action space, measurement space, or combination thereof, is high, which is typically the case in difficult tracking problems, we have to resort to *approximate dynamic programming* for mitigating the curse-of-dimensionality problem [8].

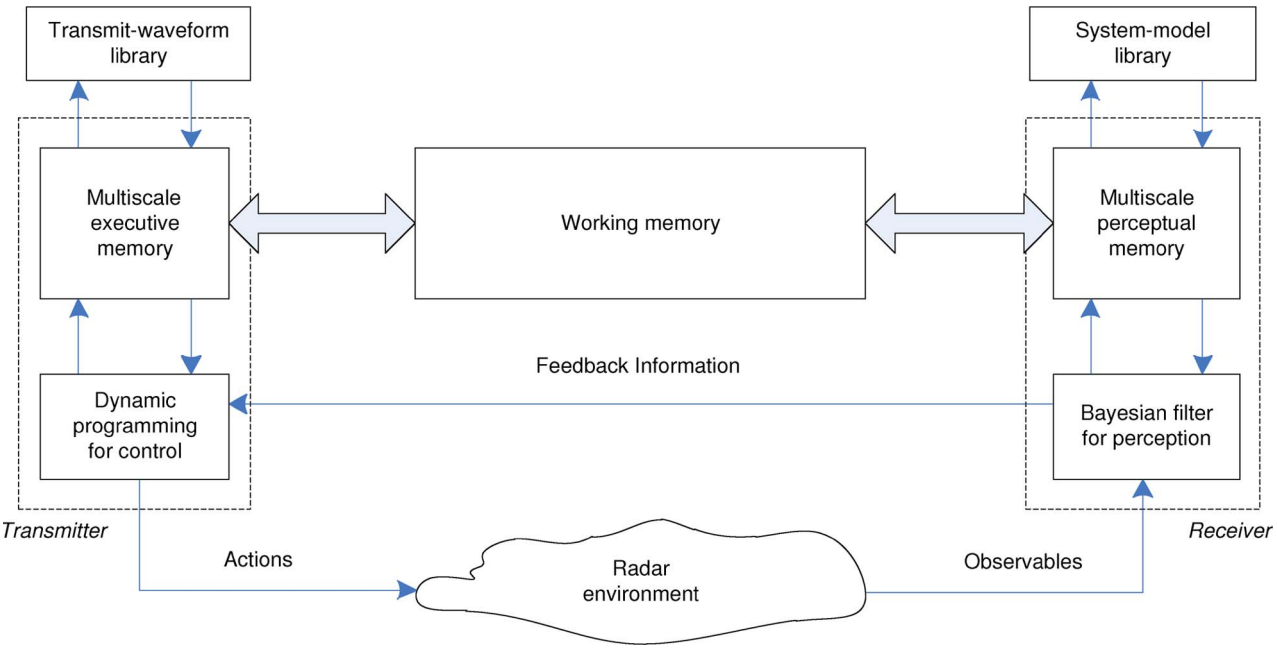

Fig. 1. ***Block diagram of CR with memory.***

Thus, in light of what we have just described, the perception–action cycle embodies four functional blocks: approximate Bayesian filter for environmental perception in the receiver, linkage for feedback information from the receiver to the transmitter, approximate dynamic programming for receiver-control performed in the transmitter and, finally, the state–space model of the radar environment. The perception–action cycle may, therefore, be viewed as a *closed-loop feedback control system*, which, in physical terms, is clearly visible in Fig. 1. In computational terms, we thus say:

> The function of the perception–action cycle is to produce information gain about the radar environment by processing the received signal (radar returns), with magnitude of the gain progressively increasing from one cycle to the next.

A. Bayesian Filtering for Optimal Perception in the Receiver

Following the traditional procedure used in the analytic study of a radar system, we proceed on the basis of a baseband model developed for describing the transmission of a radar signal across the environment [4], [8]. The idea behind baseband modeling is that the transmitted and received signals are represented by the complex envelopes of their actual radio-frequency (RF) forms. Most importantly, baseband modeling dispenses with the carrier frequency without incurring loss of information; in so doing, mathematical analysis of radar signal is simplified considerably.

In the rest of this section, we denote the cycle at the receiver in discrete time k, corresponding to the waveform $\boldsymbol{\theta}_{k-1}$, which was transmitted at cycle $k-1$.

Recognizing that in a real-life situation the radar environment operates in continuous time, the *system equation* of the state–space model is described by [30]

$$d\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t), t)\, dt + \sqrt{\mathbf{Q}}\, d\boldsymbol{\beta}(t) \tag{1}$$

where $\mathbf{x}(t) \in \mathbb{R}^{N_x}$ denotes state of the system at continuous time t; $\boldsymbol{\beta}(t) \in \mathbb{R}^{N_x}$ denotes the standard Brownian motion with increment $d\boldsymbol{\beta}(t)$ that is independent of $\mathbf{x}(t)$ and it plays the same role as system noise; $\mathbf{f} : \mathbb{R}^{N_x} \times \mathbb{R} \to \mathbb{R}^{N_x}$ is a known nonlinear function with appropriate regularity properties; and $\mathbf{Q} \in \mathbb{R}^{N_x \times N_x}$ is called the diffusion matrix. It is also assumed that the initial conditions and noise processes are statistically independent.

To discretize the stochastic differential equation (1) in an efficient manner, we may use the Itô–Taylor expansion. This expansion of higher order is theoretically more accurate than that of lower order [31]. Hereafter, the Itô–Taylor expansion of order 1.5 is chosen because of its most effective approximation capability.

Turning next to the *measurement equation*, it is generally described in discrete time k. Thus, dependence of the measurement vector $\mathbf{z}_k$ on the state $\mathbf{x}_k$ is expressed in the discrete form

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{w}_k(\boldsymbol{\theta}_{k-1}) \tag{2}$$

where the vector $\mathbf{w}_k(\boldsymbol{\theta}_{k-1})$ denotes the measurement noise that acts as the *driving force*. It is in the dependence of this noise on the waveform-parameter vector $\boldsymbol{\theta}_{k-1}$ generated at time $k-1$ that the transmitter controls accuracy of the state estimation in the receiver at time k, yielding the measurement $\mathbf{z}_k$.[2]

Examination of the system noise in (1) and measurement noise in (2) reveals an important physical fact in modeling the radar environment:

> Evolution of the target's state across time is governed by the target's dynamics itself, independent of the radar system; the radar interacts with the target only when the measurement system is switched on.

Application of the state–space model described in (1) and (2) hinges on three basic assumptions:

1) the nonlinear vectorial functions $\mathbf{f}(\cdot)$ and $\mathbf{h}(\cdot)$ in (1) and (2) are both smooth and otherwise arbitrary;
2) the system noise and measurement noise are zero-mean Gaussian distributed and statistically independent of each other;
3) the covariance matrix of system noise is known.

Examining (1) and (2), we immediately see that the state $\mathbf{x}_k$ is *hidden* from the receiver, and the challenge for the receiver is therefore to exploit dependence of the measurement vector on the state to compute an estimate of the state and do so in a sequential and online manner.

Before proceeding further, two other points deserve particular attention.

1) Mathematical description of the perception–action cycle proceeds with perception of the environment by the receiver *before* action is taken by the transmitter; in other words, we have the

[2]Building on the model proposed in [4], the measurement noise in (2) was considered as a function of the transmitted waveform parameters. However, in practice, the measurement noise is also affected by other factors such as calibration residues and system nonlinearities, which are not directly dependent on the transmitted waveform. In other words, the measurement noise consists of two terms; one is controllable by the transmitted waveform and the other one is not. For a concise treatment of the concept of CR, part of the measurement noise that is independent of the transmitted waveform has been neglected in this paper. Analysis based on (2) provides the ideal situation which can represent a benchmark for more practical situations that can be elaborated on in future studies.

pair $(\boldsymbol{\theta}_{k-1}, \mathbf{z}_k)$, where $\mathbf{z}_k$ is the measurement made at the receiver at time k in response to the transmit-waveform parameter vector $\boldsymbol{\theta}_{k-1}$.

2) The transmit waveform, generated at time $k-1$, is obtained by combining linear frequency-modulated (LFM) waveform pulse with Gaussian amplitude modulation [4], [32]. The transmit-waveform parameter vector is defined as $\boldsymbol{\theta}_{k-1} = [\lambda_{k-1}\ b_{k-1}]^T$, where λ and b denote the duration of the Gaussian envelope for the LFM chirp transmit signal and the chirp rate of the LFM pulse, respectively; the superscript T denotes matrix transposition.

The Bayesian filter is known to be optimal, at least in conceptual terms [33]. Unfortunately, when the state–space model is nonlinear as it is in (1) and (2), the Bayesian filter is no longer computationally feasible, hence the practical need for its approximation. To do this approximation under the Gaussian assumption, there are three possible continuous-discrete approaches, based on the following:

- extended Kalman filter;
- unscented Kalman filter;
- cubature Kalman filter.

In [34], these three approaches have been compared for a relatively difficult tracking problem that involves a *corner turn*, and therein it is demonstrated that the continuous-discrete cubature Kalman filter (CD-CKF) outperforms the other two approaches. Accordingly, we have chosen to work with the CD-CKF as the tool for perception of the environment in the receiver.

Under the Gaussian assumption, the Bayesian filter reduces to the problem of how to compute moment integrals whose integrands have the following form:

$$\text{nonlinear function} \times \text{Gaussian}. \tag{3}$$

To numerically compute such integrals, we consider the *cubature rule of third degree*. Following the analysis described in [34], based on this cubature rule, two steps involved in each recursion of the CD-CKF are as follows.

1) *Time update*, which focuses on computing the estimate of the state and associated covariance for consecutive steps, until we end up with the time update ($\hat{\mathbf{x}}^m_{k|k}$ and $\mathbf{P}^m_{k|k}$), respectively, where m is the number of time-update iterations per sampling interval. These estimates are then used to compute the predicted state and associated covariance, denoted by $\hat{\mathbf{x}}^m_{k|k-1}$ and $\mathbf{P}^m_{k|k-1}$, respectively.
2) *Measurement update*, which, for every new measurement $\mathbf{z}_k$, propagates the cubature points to predict the one-step prediction of the measurement and associated covariance, denoted by $\hat{\mathbf{z}}_{k|k-1}$ and $\mathbf{P}_{\mathbf{zz},k|k-1}$, respectively. Also, the cross covariance between the state and the measurement, denoted by $\mathbf{P}_{\mathbf{xz},k|k-1}$, is computed. Next, the new measurement $\mathbf{z}_k$ is incorporated into the measurement update of the state and its covariance, yielding the filtered estimate of the state $\hat{\mathbf{x}}_{k|k}$ and its covariance $\mathbf{P}_{k|k}$.

Finally, the *Kalman gain* $\mathbf{G}_k$ is computed, which is defined by

$$\mathbf{G}_k = \mathbf{P}_{\mathbf{xz},k|k-1}\mathbf{P}^{-1}_{\mathbf{zz},k|k-1} \tag{4}$$

which has been singled out because of its relevance to what follows. In particular, examining the formulas derived in [34] for the CD-CKF algorithm, we find that with the exception of $\mathbf{P}_{\mathbf{zz},k|k-1}$, all the estimates and covariances, described under the time update and measurement update, are independent of the transmit-waveform vector $\boldsymbol{\theta}_{k-1}$. To be more specific, we find that the expression for $\mathbf{P}_{\mathbf{zz},k|k-1}$ has the following format, expressed in words:

$$\mathbf{P}_{\mathbf{zz},k|k-1} = \begin{pmatrix}\text{sum of two terms}\\ \text{independent of } \boldsymbol{\theta}_{k-1}\end{pmatrix} + \begin{pmatrix}\text{covariance matrix of}\\ \text{measurement noise } \mathbf{R}_k(\boldsymbol{\theta}_{k-1})\end{pmatrix}. \tag{5}$$

For reasons that will become apparent in Section VI-D on dynamic programming, (5) plays a key role in the perception–action cycle.

B. Shannon's Entropy Versus Fisher Information

Discussion of perception, involving state estimation, would be incomplete without incorporating a theoretical limit on its estimate. To pave the way for this theoretical limit in Section VI-C, it is informative that we briefly review the two notions: *Shannon's entropy* and *Fisher information*.

Entropy is a measure of our uncertainty about an event in Shannon's information theory. Specifically, the entropy of a discrete random vector $\mathbf{X}$ with alphabet $\mathcal{X}$ is defined as follows [35]:

$$H(\mathbf{X}) = -\sum_{\mathbf{x}\in\mathcal{X}} p(\mathbf{x}) \log p(\mathbf{x}). \tag{6}$$

It can also be interpreted as the expected value of the term $1/\log p(\mathbf{x})$

$$H(\mathbf{X}) = \mathbb{E}\left[\frac{1}{\log p(\mathbf{x})}\right] \tag{7}$$

where $\mathbb{E}$ is the expectation operator and $p(\mathbf{x})$ is the probability density function (pdf) of $\mathbf{X}$.

As mentioned previously, relevant portion of the measurement can be interpreted as information. In this line of thinking, a summary of the amount of information with regard to the variables of interest is provided by the *Fisher information matrix* [36]. To be more specific, the Fisher information plays two basic roles:

1) it is a measure of the ability to estimate a quantity of interest;
2) it is a measure of the state of disorder in a system or a phenomenon of interest.

The first role implies that the Fisher information matrix has a close connection to the estimation-error covariance matrix and can be used to calculate the confidence region of estimates. The second role implies that the Fisher information has a close connection to Shannon's entropy.

Elements of the Fisher information matrix at time k are defined by [36]

$$\mathbf{J}_k^{ij} = \mathbb{E}\left[\nabla_{\mathbf{x}_i} \log p(\mathbf{X}_k, \mathbf{Z}_k) \nabla_{\mathbf{x}_j}^T \log p(\mathbf{X}_k, \mathbf{Z}_k)\right] \quad (8)$$

where $\mathbf{X}_k = \{\mathbf{x}_0, \ldots, \mathbf{x}_k\}$ and $\mathbf{Z}_k = \{\mathbf{z}_0, \ldots, \mathbf{z}_k\}$; the $\mathbf{X}_k$, denoting a set here, should not be confused with the random variable $\mathbf{X}$. The definition in (8) is based on the outer product of the gradient of $\log p$ with itself, where the gradient is denoted by $\nabla_{\mathbf{x}}$ in the equation. There is also an equivalent definition based on the second derivative of the logarithm on the right-hand side of (8), yielding

$$\mathbf{J}_k^{ij} = -\mathbb{E}\left[\frac{\partial^2 \log p(\mathbf{X}_k, \mathbf{Z}_k)}{\partial \mathbf{x}_i \partial \mathbf{x}_j}\right]. \quad (9)$$

The definitions of Shannon's entropy H and Fisher information $\mathbf{J}$ show that both are functions of the corresponding pdf. The connection between these two information measures is revealed by examining the way that they are affected by the shape of the pdf. A relatively broad and flat pdf, which is associated with lack of predictability and high entropy, has small gradient contents and, in effect therefore, low Fisher information. On the other hand, if the pdf is relatively narrow and has sharp slopes around a specific value of $\mathbf{x}$, which is associated with bias toward that particular value of $\mathbf{x}$ and low entropy, it has large gradient contents and therefore high Fisher information. In summary, there is a duality between Shannon's entropy and Fisher information. However, a closer look at their mathematical definitions reveals an important difference [1].

- A rearrangement of the tuples $\{\mathbf{x}_i, p(\mathbf{x}_i)\}$ may change the shape of the pdf curve significantly but it does not affect the value of the summation in (6) because the summation can be calculated in any order. Since H is not affected by local changes in the pdf curve, it can be considered as a global measure of the behavior of the corresponding pdf.
- On the other hand, such a rearrangement of points changes the slope and therefore gradient of the pdf curve, which, in turn, changes the Fisher information significantly. Hence, the Fisher information is sensitive to local rearrangement of points and can be considered as a local measure of the behavior of the corresponding pdf.

Both entropy (as a global measure of smoothness in the pdf) and Fisher information (as a local measure of smoothness in the pdf) can be used in a variational principle to infer about the pdf that describes the phenomenon under consideration. However, the local measure may be preferred in general [1]. This leads to another performance metric that is used in simulations to compare different radar configurations, which is discussed in Section VI-C.

C. Posterior Cramér–Rao Lower Bound (PCRLB)

To assess the performance of an estimator, a lower bound is always desirable. Such a bound is a measure of performance limitation that determines if the design criterion is realistic and implementable or not. The Cramér–Rao lower bound (CRLB) is a lower bound that represents the lowest possible mean-square error (MSE) in the estimation of deterministic parameters for all unbiased estimators. It can be calculated as the inverse of the Fisher information matrix. For random variables, a similar version of the CRLB, namely the PCRLB, was derived in [37], as shown by

$$\mathbf{P}_{k|k} = \mathbb{E}\left[(\mathbf{x}_k - \hat{\mathbf{x}}_k)(\mathbf{x}_k - \hat{\mathbf{x}}_k)^T\right] \geq \mathbf{J}_k^{-1} \quad (10)$$

which is also referred to as the Bayesian CRLB [38], [39]. To calculate it in an online manner, an iterative version of the PCRLB for nonlinear filtering using state–space models was proposed in [40], where the posterior information matrix of the hidden state is decomposed for each discrete-time instant by virtue of the factorization of the joint pdf of the state variables. In this way, an iterative structure is obtained for evolution of the information matrix. For a nonlinear system with the following state–space model:

$$\mathbf{x}_{k+1} = \mathbf{f}_k(\mathbf{x}_k) + \boldsymbol{\nu}_k \quad (11)$$

$$\mathbf{z}_k = \mathbf{h}_k(\mathbf{x}_k) + \mathbf{w}_k \quad (12)$$

the sequence of posterior information matrices $\{\mathbf{J}_k\}$ for estimating state vectors $\{\mathbf{x}_k\}$ was calculated in [40] as

follows:

$$\mathbf{J}_{k+1} = \mathbf{D}_k^{22} - \mathbf{D}_k^{21}\left(\mathbf{J}_k + \mathbf{D}_k^{11}\right)^{-1}\mathbf{D}_k^{12} \tag{13}$$

where

$$\mathbf{D}_k^{11} = \mathbb{E}\left[\nabla_{x_k} f_k(\mathbf{x}_k)\mathbf{Q}_k^{-1}\nabla_{x_k}^T f_k(\mathbf{x}_k)\right] \tag{14}$$

$$\mathbf{D}_k^{12} = -\mathbb{E}[\nabla_{x_k} f_k(\mathbf{x}_k)]\mathbf{Q}_k^{-1} \tag{15}$$

$$\mathbf{D}_k^{21} = \left[\mathbf{D}_k^{12}\right]^T \tag{16}$$

$$\mathbf{D}_k^{22} = \mathbb{E}\left[\nabla_{x_{k+1}} h_{k+1}(\mathbf{x}_{k+1})\mathbf{R}_{k+1}^{-1}\nabla_{x_{k+1}}^T h_{k+1}(\mathbf{x}_{k+1})\right] + \mathbf{Q}_k^{-1}. \tag{17}$$

The $\mathbf{Q}_k$ and $\mathbf{R}_k$ are the process and measurement noise covariance matrices, respectively.

As will be described in Section VI-D, CR manipulates the measurement noise covariance matrix $\mathbf{R}_{k+1}$ by adapting the waveform parameters in a way that the covariance is reduced iteratively. In the above formulas, the measurement noise covariance matrix only appears in (17). For the linearized measurement equation

$$\mathbf{z}_k = \mathbf{H}_k\mathbf{x}_k + \mathbf{w}_k \tag{18}$$

which is beneficial for illustrative purposes, (17) is reduced to

$$\mathbf{D}_k^{22} = \mathbf{Q}_k^{-1} + \mathbf{H}_{k+1}^T\mathbf{R}_{k+1}^{-1}\mathbf{H}_{k+1}. \tag{19}$$

If the elements of $\mathbf{R}_{k+1}$ reduce to small values, elements of $\mathbf{R}_{k+1}^{-1}$ will have large values, and elements of $\mathbf{D}_k^{22}$, therefore, will also have large values. Then, according to (13), elements of $\mathbf{J}_{k+1}$ and $\mathbf{J}_{k+1}^{-1}$ will, respectively, have large and small values. CR iteratively reduces the lower bound by manipulating $\mathbf{R}_{k+1}$. Theoretically speaking, as the elements of $\mathbf{R}_{k+1}$ reduce iteratively and go toward zero, the measurement noise is gradually eliminated from the measurement equation and the measurements will contain only useful information about the hidden state. Then, the radar will be able to provide accurate estimates of the state. In other words, in theory, CR would be capable of pushing the PCRLB down to zero, which will be confirmed by computer experiments later on. However, implementing such a system in practice would be very challenging.

Section VI-D discusses the sensitivity analysis.

D. Sensitivity Analysis

The derivation of the approximate Bayesian filter is based on the assumption that the system model and noise statistics are known. In practice, it is quite likely that the true distributions deviate from the assumed nominal ones. Therefore, it would be desirable to modify the filter in a way that it is to some extent desensitized with respect to modeling errors and implementation approximations. Since there is no free lunch, this *robustification* may deteriorate the performance of the filter [5], [41]. According to Huber and Ronchetti [42], an algorithm is said to be *robust* when its performance is not impaired by small deviations of the actual pdf from the pdf of the assumed model. In other words, a robust algorithm's performance must be acceptable for a range of possible deviations from the nominal model [43].

There are two approaches to design filters that are able to deal with the problem of not knowing the true distribution.

- *Stochastic approach*, in which a *risk-sensitive* formulation of the nonlinear filter is considered. Rather than minimizing the MSE, the risk-sensitive filter minimizes the mean-exponential-square error. An error bound can be obtained for the risk-sensitive filter, which consists of two terms that are respectively associated with the good performance under nominal conditions and acceptable performance under perturbed conditions. The second term determines the range of permissible deviations from the nominal model. Compared to the optimal Bayesian filter, the risk-sensitive filter will have a performance which may be slightly worse under nominal conditions but deteriorates slower under perturbed conditions [44].
- *Deterministic approach*, in which the *min–max* technique is used to achieve optimal performance under the worst case condition. In this case, an error bound can be obtained in terms of disturbance energy [45].

Another comment is in order, regarding the heavy-tail distributions. In Kalman-filter-based estimators, due to the Gaussian assumption, neither prior nor likelihood dominates the other one because they have similar tails. Therefore, the posterior will be a compromise between the prior and the likelihood. If one of them has a heavier tail (i.e., it is weaker), then the posterior converges to the stronger distribution that has a lighter tail, be it the prior or the likelihood. In this context, when the innovation term is large or there is a large difference between the actual measurement and its predicted value, convergence of the posterior to the prior or the likelihood will have two different meanings. Convergence to the prior means that the filter somehow considers the new measurement as an outlier and has the tendency to ignore it. On the other hand, convergence to the likelihood can be interpreted as acknowledgement of a shift in structure [46]. In the formulation of the risk-sensitive

filter, cost function is a weighted summation of two terms: an accumulated error cost up to instant $k-1$ and the error cost at time k, where each one of these terms is multiplied by a risk parameter. Therefore, the risk-sensitive filter will be able to handle the situation of occurrence of a large innovation delicately. Moreover, in such a situation, entropy as a global measure of the behavior of the corresponding pdf will help the perceptual associative memory to decide whether a change of the model structure is necessary. Modifying the perceptor according to the risk-sensitive formulation is a topic for future research.

Now that the receiver is covered thoroughly, the controller structure in the transmitter is investigated next.

E. Dynamic Programming for Control in the Transmitter

Previously, we remarked that the measurement noise covariance in the measurement equation (2) depends only on the transmitted waveform parameter vector $\boldsymbol{\theta}=[\lambda, b]$, which applies to LFM waveform pulse combined with Gaussian amplitude modulation. Hence, if the waveform parameters are selected optimally, any *action* taken by the transmitter will result in an optimal control of the receiver via the environment; in other words, the receiver is under control of the transmitter. With this point in mind, we may now address algorithmic formulation for waveform selection in the transmitter.

Before proceeding further, however, one important remark deserves particular attention: the state of the target is *hidden* from the receiver. This fact, in turn, poses a practical problem in the following sense: the formulation of Bellman's dynamic programming demands not only that the environment be Markovian but also that the controller have perfect knowledge of the state. In reality, however, the transmitter of a radar tracker has an *imperfect* estimate of the state reported to it by the receiver. Accordingly, we are faced with an *imperfect state-information problem*. To resolve this problem, we follow [47] by introducing a new information state vector defined by [8]

$$\mathbf{I}_k \triangleq (\mathbf{Z}_k, \boldsymbol{\Theta}_{k-1}), \qquad \text{with } \mathbf{I}_0 = \mathbf{z}_0 \tag{20}$$

where

$$\mathbf{Z}_k = [\mathbf{z}_0\ \mathbf{z}_1\ \cdots\ \mathbf{z}_k] \tag{21}$$

$$\boldsymbol{\Theta}_{k-1} = [\boldsymbol{\theta}_0\ \cdots\ \boldsymbol{\theta}_{k-1}]. \tag{22}$$

From these three equations, we readily obtain the recursion

$$\mathbf{I}_k = (\mathbf{I}_{k-1}, \mathbf{z}_k, \boldsymbol{\theta}_{k-1}) \tag{23}$$

which may be viewed as the state evolution of a new dynamic system with perfect state information, and therefore amenable to dynamic programming. According to (23), we may now say:

- $\mathbf{I}_{k-1}$ is the old value of the state;
- $\boldsymbol{\theta}_{k-1}$ is the waveform parameter vector computed at time $k-1$, and on which the receiver operates at time k;
- the current measurement $\mathbf{z}_k$ is viewed as a *random disturbance* resulting from the control decision $\boldsymbol{\theta}_{k-1}$.

At any cycle time k, the waveform-selection algorithm seeks to find the set of best waveform parameters by minimizing a *cost-to-go function* for a rolling horizon of L steps, that is, to minimize the cost incurred in steps $k : k+L-1$. Denoting the control policy for the next L steps by $\pi_k = \{\mu_k, \ldots, \mu_{k+L-1}\}$ with the policy function $\mu(\mathbf{I}_k) = \boldsymbol{\theta}_k \in \mathcal{P}_{k+1}$ mapping the information vector into a selection in the waveform library $\mathcal{P}_{k+1}$, we wish to find a control policy π_k at time k corresponding to the solution of the following minimization [4], [8]:

$$\min_{\pi_k} \mathbb{E}\left[\sum_{i=k}^{k+L-1} g(\boldsymbol{\theta}_k)\right] \tag{24}$$

where L is the number of time steps looking ahead, and the cost-to-go function $g(\boldsymbol{\theta}_k)$ is defined in terms of the posterior expected state-estimation error $\varepsilon_{k+1|k+1}$ that is dependent on $\boldsymbol{\theta}_k$. The posterior covariance of the state (the updated filtered state-error covariance matrix), namely $\mathbf{P}_{k+1|k+1}(\boldsymbol{\theta}_k)$, is implicitly included under the summation in formulating the cost-to-go function; the cost function $g(\cdot)$ inside the summation in (24) is the Fisher information metric defined by the trace of $\mathbf{P}_{k+1|k+1}(\boldsymbol{\theta}_k)$. The $\boldsymbol{\theta}_k$ is the "unknown" to be computed for the perception–action cycle to continue on.

With optimal control being the essence of dynamic programming, we have to address the following question:

> How do we optimize the cost function $g(\boldsymbol{\theta}_k)$ with respect to the "unknown" transmit-waveform vector $\boldsymbol{\theta}_k$ in a computationally feasible manner?

To answer this question, we first note that the feedback information computed in the receiver, to be passed onto the transmitter, is defined by the one-step prediction covariance $\mathbf{P}_{k+1|k}$. The desired posterior (smoothed) covariance $\mathbf{P}_{k+1|k+1}$ is related to $\mathbf{P}_{k+1|k}$ by the *Riccati equation*, as follows [8], [34]:

$$\mathbf{P}_{k+1|k+1} = \mathbf{P}_{k+1|k} - \mathbf{G}_{k+1}\mathbf{P}_{\mathbf{zz},k+1|k}\mathbf{G}_{k+1}^T. \tag{25}$$

The product term $\mathbf{G}_{k+1}\mathbf{P}_{\mathbf{zz},k+1|k}\mathbf{G}_{k+1}^{T}$ may be expressed as follows, using (4):

$$\mathbf{G}_{k+1}\mathbf{P}_{\mathbf{zz},k+1|k}\mathbf{G}_{k+1}^{T} = \mathbf{P}_{\mathbf{xz},k+1|k}\mathbf{P}_{\mathbf{zz},k+1|k}^{-1}\mathbf{P}_{\mathbf{xz},k+1|k} \quad (26)$$

where we have made use of the symmetric property of the covariance matrices on the right-hand side of (26). Now, we make two important observations.

- The cross covariance $\mathbf{P}_{\mathbf{xz},k+1|k}$ is independent of $\boldsymbol{\theta}_k$.
- On the other hand, the measurement covariance $\mathbf{P}_{\mathbf{zz},k+1|k}$ is dependent on $\boldsymbol{\theta}_k$, with the dependence being solely limited to the measurement noise covariance $\mathbf{R}(\boldsymbol{\theta}_k)$.

We are, therefore, fortunate to find that the optimization is feasible in an online manner because $\boldsymbol{\theta}_k$ shows up only in $\mathbf{R}(\boldsymbol{\theta}_k)$; this idea was first described in [4].

In a physical context, the state–space model of CR is an infinite-dimensional continuous-valued space. Moreover, the dimension of the model grows exponentially with depth of the optimization horizon L. Specifically, at each step of the dynamic programming algorithm, we need to examine an infinite number of possibilities such that the perfect state-information vector $\mathbf{I}_k$ can evolve to a new value on the next time step. To simplify this very burdensome computation, we apply an approximation technique, namely the cubature rule of third degree, to the expectation involved in defining the cost function $g(\boldsymbol{\theta}_k)$; see [8] for details. The net result of doing all this is an approximation form of dynamic programming, which is computationally feasible so long as the optimization horizon L is kept not too large.

VII. MEMORY

Next, we turn to the requirement that CR has to *learn* from the experience gained through continued interactions with the radar environment [2], [3]. This requirement is satisfied by equipping the radar with a memory system, one part of which resides in the receiver, another part resides in the transmitter, and the two of them are reciprocally coupled in the manner described in Fig. 1.

These three kinds of memory are discussed in what follows, one by one.

A. Perceptual Memory

The part of memory that resides in the receiver is called *perceptual memory*. It would be desirable for the perceptual memory to have a multiscale structure. The idea behind this kind of structure is referred to in the neural network literature as *features of features* [48]. Basically, through a learning process, the first layer of the perceptual memory extracts the important features that characterize the incoming measurement vector $\mathbf{z}_k$ received at cycle k. Naturally, these features act as the input to the second layer of the perceptual memory, which, in turn, goes on to extract the features of features that characterize the original measurement $\mathbf{z}_k$, and so on for the third layer of perceptual memory.

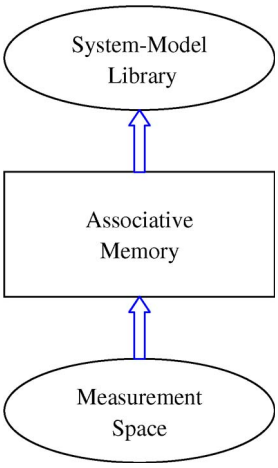

Fig. 2. *Perceptual memory as a heterogeneous associative memory.*

The idea of features of features, as just described, is simply a learning process that goes on in an *associative memory* of the *heterogeneous* kind [22], [49]–[51]. To elaborate, consider the diagram depicted in Fig. 2, where we have a measurement space at the input end of an associative memory and a system-model library at the output end; this library consists of a grid of points, each one of which represents a different set of values of the nonlinearity (responsible for transition from one state to another) and covariance of the system noise. Basically, the associative memory acts as a *correlator* between the measurement space at the one end and the system-model space at the other end for the purpose of matching each point in the system-model space to a corresponding point in the measurement space.

B. Executive Memory

The part of the memory that resides in the transmitter is called the *executive memory*, whose structural composition follows a similar format to that of the perceptual memory, with the following basic difference:

> Whereas the perceptual memory sees the radar environment through the measurement vector directly, the executive memory sees the radar environment indirectly through the feedback information about the environment supplied to the transmitter by the receiver.

Note also that the executive memory is reciprocally coupled to another library called the *transmit-waveform library*. Each grid point in this second library represents a different combination of chirp rate of the LFM signal and duration of its Gaussian envelope: two parameters that define the transmit-waveform vector $\boldsymbol{\theta}$. The executive

memory operates in a manner similar to the perceptual memory except for the following difference: the matching (correlation) process involves the feedback information rather than the measurement vector. Equally well, except for differences in terminology, Fig. 2 also applies to the executive memory, representing another example of associative memory.

C. Working Memory

In order to exploit the full capability available in having the perceptual memory and executive memory, it is desirable to have them reciprocally coupled together. This reciprocal coupling is achieved by means of the *working memory*, as shown in Fig. 1. The working memory is a dedicated memory with limited capacity that provides an interface between perception and action by linking the perceptual and executive memories together. Unlike the perceptual and executive memories that are long-term memories, working memory has a short-term nature, and it is therefore used for temporary information storage [52].

In the memory viewed as a whole, we thus have an integrated system that enables all the information-processing steps performed in each cycle of the perception–action cycle mechanism to proceed in a *synchronized (coherent) fashion* across time. This *self-organized temporal behavior* is another cardinal characteristic of CR.

Having described the composition of memory, we may now go on to describe its function as follows:

> The function of working memory in CR is to predict the consequences of actions taken by the radar's receiver, transmitter or both.

In CR, the working memory consists of a neural network of the *heterogeneous associative kind* [22], [49]–[51]. For example, suppose that in a particular cycle in the perception–action cycle, the perceptual memory picked a grid point in the system-model library that is incorrectly matched to the measurement vector at the receiver input. In such a situation, the working memory corrects that action on the next cycle. A similar statement applies when a point in the transmit-waveform library is incorrectly picked in the transmitter.

D. Guidelines for Designing the Memory

As pointed out previously in this section, the use of multiscale memory is a highly desirable design requirement for the CR tracker to perform its assigned task exceptionally well. In other words, *hierarchical depth* should be built into the design of memory [53].

There are various ways in which this requirement can be satisfied; the most elegant of them all is based on the *encoding–decoding principle* [27]. According to this principle, the memory is composed of two components:

- *multiscale encoder*, the function of which is to *map* the incoming measurement vector onto an abstract feature that is easy to understand;
- *multiscale decoder*, which operates on the features produced by the encoder to reconstruct the original measurement vector as closely as possible.

The simplest way of constructing the encoder–decoder is to use a multilayer perceptron, called the *replicator* or *auto-encoder* [22].

1) Design of the Perceptual and Executive Memories: A block diagram of the complete memory system is represented by Fig. 3, which embodies the perceptual memory (on the right), the executive memory (on the left), and the working memory (in the middle). Both the perceptual memory and the executive memory have multiple hidden layers. In the following, we will start from the design of the multiscale perceptual memory.

As already mentioned, a replicator neural network is a simple choice of implementing the multiscale memory [54]. The replicator neural network described in [54] is a special kind of multilayer perceptron network with three hidden layers, each of which follows a strict structural configuration. Therefore, the *backpropagation algorithm* is applicable as the training procedure [22]. The input layer of source nodes and the output layer of neurons have the same size. However, in going up through the encoder, the size of the hidden layers gets progressively smaller until we get to the "bottleneck" at hidden layer 3. Then, with hidden layer 3 common to the encoder and the decoder, we see that the two of them form a symmetric pair with respect to each other. Likewise, the activation functions of neurons differ from layer to layer, summarized as follows [54].

- For the second and fourth hidden layers, the activation functions are defined by the hyperbolic tangent function

$$\varphi^{(2)}(\nu) = \varphi^{(4)}(\nu) = \tanh(\nu) \tag{27}$$

  with ν as the input to a neuron in a particular layer.
- For the output (fifth) layer, we use a linear activation function defined by

$$\varphi^{(5)}(\nu) = \nu. \tag{28}$$

- The activation function for neurons in the middle (third) layer is defined as follows:

$$\varphi^{(3)}(\nu) = \frac{1}{2} + \frac{1}{2(N-1)} \sum_{j=1}^{N} \tanh\left(a\left(\nu - \frac{j}{N}\right)\right) \tag{29}$$

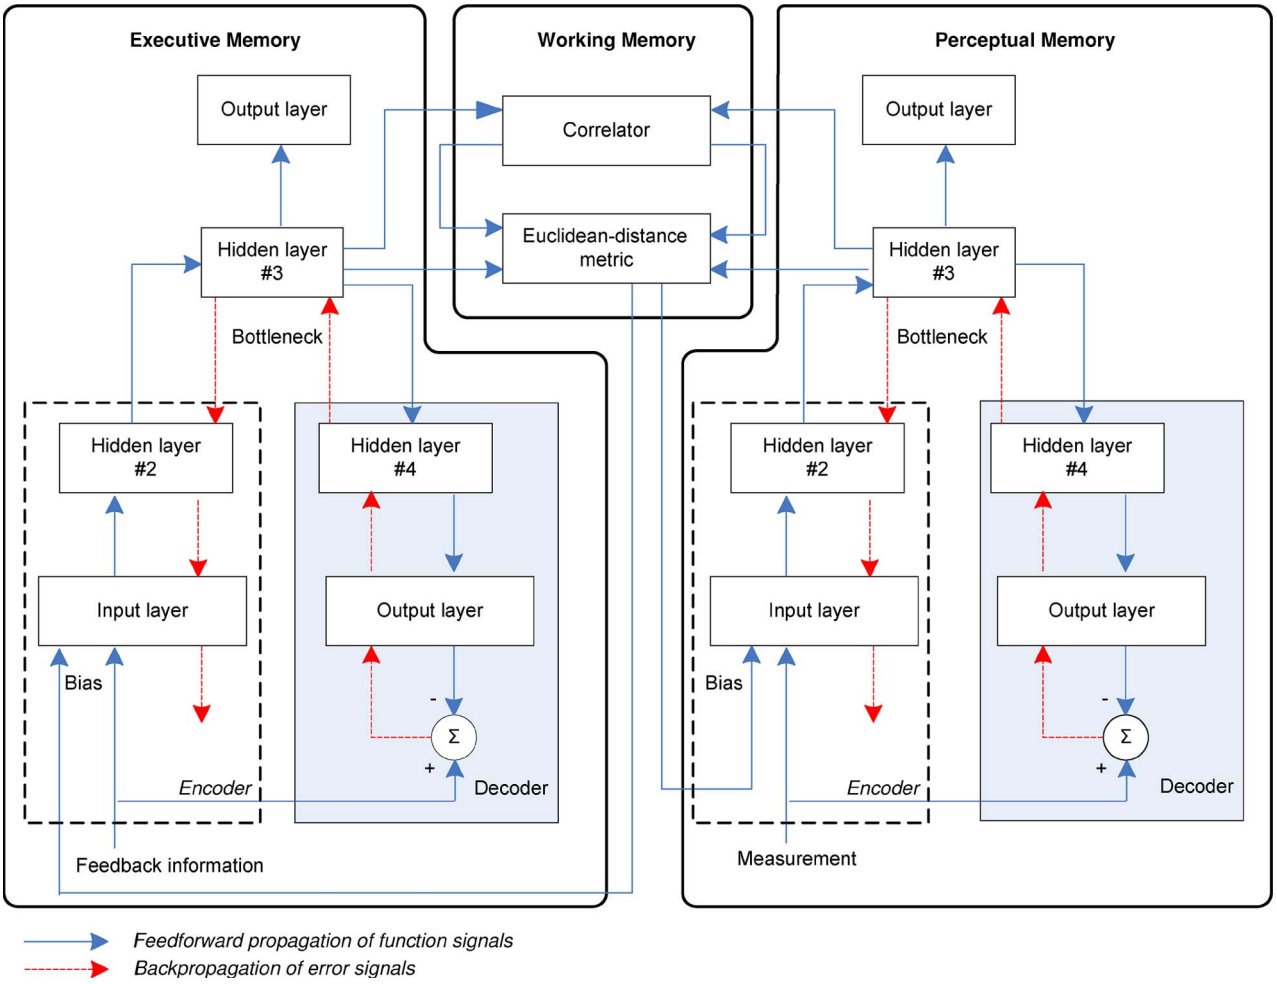

Fig. 3. ***Block diagram of the memory system, embodying perceptual memory, executive memory, and working memory.***

which describes a smooth staircase activation function with N treadles, thereby quantizing the vector of the third-layer outputs into $K = N^n$ grid points, with n as the number of neurons of this third layer.

Note that in Fig. 3, aggregation of the encoder, the decoder, and the bottleneck layer is a complete representation of the replicator network. Training of the replicator network is carried out in an *unsupervised* manner, because the measurement vector is applied to both the input layer and the output layer, where, in the latter case, the measurement vector takes on the role of *desired response*. The overall effect of the unsupervised training procedure can be rationalized in two ways. On the one hand, the signals that flow from the input layer to the middle (bottleneck) layer can be regarded as a process of *data compression*, i.e., data encoding. On the other hand, the signals that flow from the middle layer to the output layer represent a process of *data decompression*, i.e., data decoding.

When the training of the replicator network is finished, the decoder part of the replicator network, shown shaded in Fig. 3, is no longer needed, as there is no longer a requirement for decoding. To this end, we purposely isolate the decoder from the network and connect the bottleneck layer (i.e., the third layer) to a new output layer which has the same dimensionality of the system model.

We thus have a *hybrid neural network* with the first part being the encoder of the replicator network and the second part being only a single output layer. To train this hybrid neural network as a whole, we fix the weights in the encoder and train the output weights by using the system model extracted from the library as the teacher signal. Algorithmically, we use supervised training based on the least-mean-square (LMS) algorithm, described in [56]

$$\hat{\mathbf{w}}_o(k+1) = \hat{\mathbf{w}}_o(k) + \eta \mathbf{x}(k) e(k) \tag{30}$$

where η is the learning-rate parameter and $e(k)$ is the error signal, defined by the difference between an ideal (teacher)

signal $d(k)$ in the system-model library and output of the encoder in the replicator network, i.e., $e(k) = d(k) - \hat{d}(k)$. The net result of the computations described herein is extraction of those important features that characterize the radar environment.

At this point, we have finished training the perceptual memory and it can be switched to running mode with all the weights fixed. One last note we need to mention is that the input layer of the perceptual memory has a bias input. It is of no practical use at the training stage. As such, it is set as 1 temporarily. However, it will be connected to the working memory at the running stage.

Design of the executive memory follows a strategy similar to that described above for the perceptual memory. It is worthwhile here to emphasize the following points:

- another replicator network is first constructed and its decoder is isolated at a later stage;
- a new output layer is connected to the encoder to construct another hybrid neural network;
- the desired transmit waveform obtained from the transmit-waveform library is selected as the teacher signal to train this second hybrid neural network;
- the bias input of the executive memory is also set as 1 at the training stage and connected to working memory at the running stage.

2) Design of the Working Memory: As described previously in this section, the function of working memory is to reciprocally couple the transmitter and the receiver in order to synchronize their operations. As such, the design of working memory should take into consideration the information extracted by both the perceptual memory and the executive memory, as depicted by the middle part of Fig. 3.

Let the output vectors of hidden layer #3, i.e., the bottleneck layer, of the perceptual memory and the executive memory be $\mathbf{o}_p(k) \in \mathbb{R}^{L_p \times 1}$ and $\mathbf{o}_e(k) \in \mathbb{R}^{L_e \times 1}$, respectively. To compute the output weights, we collect the features extracted by the bottleneck layer into a matrix

$$\mathbf{O}_p = [\mathbf{o}_p(1), \mathbf{o}_p(2), \ldots, \mathbf{o}_p(M)] \tag{31}$$

$$\mathbf{O}_e = [\mathbf{o}_e(1), \mathbf{o}_e(2), \ldots, \mathbf{o}_e(M)] \tag{32}$$

where M is the number of training data sets.

If we can obtain the correlation matrix of $\mathbf{O}_p$, denoted by $\mathbf{r}_{pp} \in \mathbb{R}^{L_p \times L_p}$ and the cross-correlation vector between $\mathbf{O}_p$ and $\mathbf{O}_e$ denoted by $\mathbf{r}_{pe} \in \mathbb{R}^{L_p \times L_e}$, we may then use the least squares solution of the weights from perceptual memory to executive memory, as shown by

$$\mathbf{w}_{pe} = \mathbf{r}_{pp}^{-1}\mathbf{r}_{pe} \tag{33}$$

which is the solution to the Wiener–Hopf equation [56].

Similarly, we can obtain the Wiener solution of weights from the executive memory to the perceptual memory, expressed by

$$\mathbf{w}_{ep} = \mathbf{r}_{ee}^{-1}\mathbf{r}_{ep}. \tag{34}$$

When the training of the working memory is finished, the two sets of weights, i.e., $\mathbf{w}_{pe}$ and $\mathbf{w}_{ep}$, are fixed, respectively. Then, the system is switched to the running mode. For any unknown output produced by the perceptual memory at hidden layer #3, it will correspondingly induce an output at the same hidden layer of the executive memory by going through the working memory from the receiver to the transmitter. In this scenario, the dimension of the output from the working memory is in accordance with the bottleneck layer of the executive memory. The Euclidian distance between them represents the mismatch, if any, between the perceptual memory and the executive memory, to which we refer as the *perceptual consequence*. To synchronize the executive memory to the perceptual memory, this Euclidian distance is applied to the executive memory as a bias input.

In the same manner, we can go through the working memory from the transmitter to the receiver. Here, the Euclidian distance between the output of the working memory and the bottleneck layer from the perceptual memory is calculated to represent the *executive consequence*, which is then applied to the perceptual memory as a bias input.

The two consequences, perceptual and executive, are used to account for the consequences of action taken in the receiver and the transmitter, if and when they are required in the operating mode.

Remark: The current design of the memory elements is based on artificial neural networks, which heavily rely on heuristics. For instance, choosing the training data set is a critical design factor. On the one hand, training data set must be rich enough to excite all dynamic modes of the system whose behavior is going to be learned by the neural network; on the other hand, there must be a tradeoff between training and generalization ability of the network. In other words, neural networks should not be overtrained in order to be able to handle slightly different situations. Cross validation is one of the techniques that can be used for this purpose [22], [57]. Since some system designers may be bothered by heuristics, the systematic design of memory elements as part of a rigorous mathematical theory of CR is suggested as a future research topic.

VIII. ATTENTION

Examining the block diagram of Fig. 1, we see that both the perception–action cycle and the memory occupy distinct *physical places* of their own within CR. However, this is not

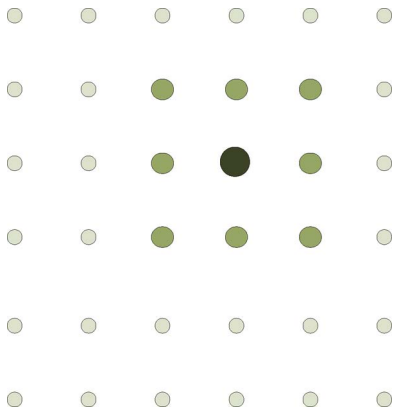

Fig. 4. *Illustrating the explore–exploit strategy performed by the attentional mechanism: the darkened grid point was selected on the preceding cycle; together with the surrounding eight grid points (picked at the current cycle) they constitute the localized cluster to be passed onto the controller.*

so when it comes to attention. Rather, memory-based attention manifests itself in the receiver and the transmitter of CR through algorithmic mechanisms. We thus have *perceptual attention* in the receiver and *executive attention* in the transmitter.

In a tracking problem that is the issue of interest in this paper, the so-called *explore–exploit strategy* plays a key role in formulating an algorithm for the executive attention. From the discussion presented on the executive memory, we recall that a particular grid point in the transmit-waveform library is selected to closely match the feedback information on the preceding cycle. Then, for the current cycle in the perception–action cycle, we may therefore look to that particular point as a *center point* to expand on. Specifically, we look to the subset of grid points in the transmit-waveform library that lie in the immediate neighborhood of the center point, as illustrated in Fig. 4. In other words, the computational effort involved in a *global search* of the entire library is replaced with a *local search* based on a small cluster of grid points. Typically, with changes in the execution of explore–exploit strategy being relatively slow from one cycle to the next, we expect that evolution of the local search to be correspondingly smooth.[3]

An important point to note here is that it is the explore–exploit strategy that makes the inclusion of attention in CR a much needed process. Simply put, if we are to adopt a global search in designing CR, there is no need for attention; the end result of this adoption could be a highly demanding computational complexity.

Attention has yet another key role to play in CR. Adding executive attention to the radar imposes temporal stability on the time rate of change in the transmit waveform. In other words, the executive attention implements a version of the *principle of minimum disturbance* [55], [56]. In a sense, by adding memory and attention to radar, the system more closely mimics the visual brain. Evidence has been found by neuroscientists that in the visual system, each cell's response is normalized by the integrated activity of its neighboring neurons. In a related context in [57], it was shown that the entropy of the response distribution can be decreased by such a gain-control mechanism through reducing neuron-response variability both within and between scenes. In Section XIII, it will be demonstrated by computer experiments that the principle of minimum disturbance, manifested by executive attention in CR, will lead to a smoother waveform evolution through time, which is of practical benefit to the transmitter.[4]

IX. INTELLIGENCE

As with attention, intelligence does not occupy a distinct physical space of its own in CR. Rather, intelligence also manifests itself through an algorithmic mechanism that is driven by the combination of attention, memory, and the perception–action cycle. With this kind of a structure, it is therefore not surprising to find that intelligence is the most powerful of all the four principles of cognition under Fuster's paradigm of cognition.

With the localized cluster of grid points in the transmit-waveform library provided by the executive attentional mechanism, as pictured in Fig. 4, it is now the function of the controller in the transmitter to do the following:

> Find the "best" waveform $\boldsymbol{\theta}_k$ among the localized cluster of transmit waveforms for which the cost function $g(\boldsymbol{\theta}_k)$ is minimized.

For the special case of $L = 1$, solving this optimization problem can be computationally much simpler than that of (24), depending on how large the transmit waveform library is. The above statement can be expanded to account for $L > 1$, if so desired.

We thus have two contrasting strategies for building the controller:

1) *global search* of the grid points in the transmit-waveform library, whereby executive attention is bypassed but global optimization for waveform

[3]In a real-world application of radar, we have to consider the issue of target detection. It is here where perceptual attention, implemented using the explore–exploit strategy comes into play. In this paper, however, target tracking has been the issue of interest, with the premise that a target has been detected, ready for tracking. In this limited objective, there is no need for perceptual attention.

[4]The practical importance of a smooth transition of the transmitted waveform from one cycle to the next is that it prolongs the life of the microwave device (i.e., magnetron, klystron, or traveling-wave tube) in the radar transmitter.

selection is guaranteed at the expense of increased computational complexity;

2) *local search* of the grid points using the explore–exploit strategy, whereby executive attention comes into play, resulting in a suboptimal selection of the transmit waveform and therefore a somewhat degraded tracking accuracy. But the practical benefit gained in using the local search is the conservation of computational resources.

What we have just described here is a tradeoff between conservation of computational resources and target tracking accuracy, which is a perfect example of the *no-free-lunch theorem* [59].

Moreover, examining the block diagram of Fig. 1, there may be an abundance of local and global feedback loops, depending on how many layers are built into the memory. Feedback, if used properly, may enhance the information-processing power of the controller by reducing the effect of uncertainties and disturbances. In other words, feedback is the facilitator of intelligence, improving performance of the controller.

For functionality, we may therefore make the following statement:

> The function of intelligence, distributed throughout CR, is to enable the controller in the transmitter to pick a transmit-waveform vector $\boldsymbol{\theta}$, so as to exercise control over the receiver in a robust manner in the face of environmental uncertainties and disturbances.

The more that such a choice reduces the feedback information on a cycle-by-cycle basis, the closer the state estimator in the receiver assumes a *deterministic* form, which, in turn, means that CR, as a whole, is robust.

Table 1 presents a summary of the cognitive processes used in CR.

Table 1 Summary of Cognitive Processes

1	**Perception-action cycle**: form of global feedback made up of state estimator implemented using Bayesian filter in the receiver, feedback information, controller implemented using Bellman's dynamic programming in the transmitter, and state-space model of the radar environment.
2	**Memory**: learning machine based on inputs directly from perception-action cycle; it is made up of associative perceptual memory in the receiver, associative executive memory in the transmitter, the two of which are reciprocally coupled together through associative working memory.
3	**Attention**: memory-driven algorithmic mechanism that indirectly exploits inputs from perception-action cycle; it is made up of perceptual attention in the receiver and executive attention in the transmitter.
4	**Intelligence**: attention-driven algorithmic mechanism that is distributed in an abundant use of feedback loops throughout the radar; it exploits inputs from attention directly, and indirectly from memory and perception-action cycle.

X. CYCLIC DIRECTED INFORMATION FLOW

At this point in the discussion, in light of what we have already learned about perception of the radar environment in the receiver and action in the transmitter, it is instructive that we formulate a *cyclic directed information flow graph*, as depicted in Fig. 5 for CR of Fig. 1 in its most generic form. Examination of this figure reveals two fundamental transmission paths, one being *perceptual pathway* (i.e., feedforward) and the other *executive pathway* (i.e., feedback).

A. Perceptual Pathway

The perceptual pathway in Fig. 5 embodies the dynamics of the perception–action cycle in the receiver influenced by the perceptual memory. It begins with the measurement $\mathbf{z}_k$ at the receiver input, initiating the following sequence of computations at time k:

- filtered estimation of the state $\mathbf{x}_k$, denoted by $\hat{\mathbf{x}}_{k|k}$, and associated error covariance $\mathbf{P}_{k|k}$, followed by
- predicted estimation of the future state $\mathbf{x}_{k+1}$, denoted by $\hat{\mathbf{x}}_{k+1|k}$, and finally
- covariance of the prediction error vector $\boldsymbol{\varepsilon}_{k+1|k}$, denoted by $\mathbf{P}_{k+1|k}$.

The perceptual pathway, based on an approximate Bayesian filter (e.g., the CD-CKF) in the receiver, is depicted on the right-hand side of Fig. 5. Basically, this pathway of directed information flow describes the *perceptual dynamics* of the receiver.

B. Executive Pathway

The executive pathway in Fig. 5 embodies the action part of the perception–action cycle in the transmitter, influenced by the executive memory. It begins with the feedback information $\mathbf{P}_{k+1|k}$ at the transmitter input delivered by the receiver, which initiates the following sequence of computations that look into the future by one cycle:

- formulation of the cost-to-go function $g(\boldsymbol{\theta}_k)$, defined in terms of the posterior expected state-estimation error $\boldsymbol{\varepsilon}_{k+1|k+1}$ that depends on the "to be updated" parameter vector $\boldsymbol{\theta}_k$, followed by
- optimization of the cost-to-go function, yielding the unknown $\boldsymbol{\theta}_k$, and finally
- setting the stage for a repeat of the next perception–action cycle.

The cost-to-go function provides the transmitter with a measure of how well the receiver is doing its job in extracting information about the environment from the radar returns.

Recognizing that dependence of the covariance $\mathbf{P}_{\mathbf{zz},k|k-1}$ on $\boldsymbol{\theta}_{k-1}$ in (5) shows up in one place only, namely the measurement noise covariance $\mathbf{R}_k(\boldsymbol{\theta}_{k-1})$, in a corresponding way, we find that dependence of the cost function $g(\boldsymbol{\theta}_k)$ on the unknown $\boldsymbol{\theta}_k$ shows up only in the

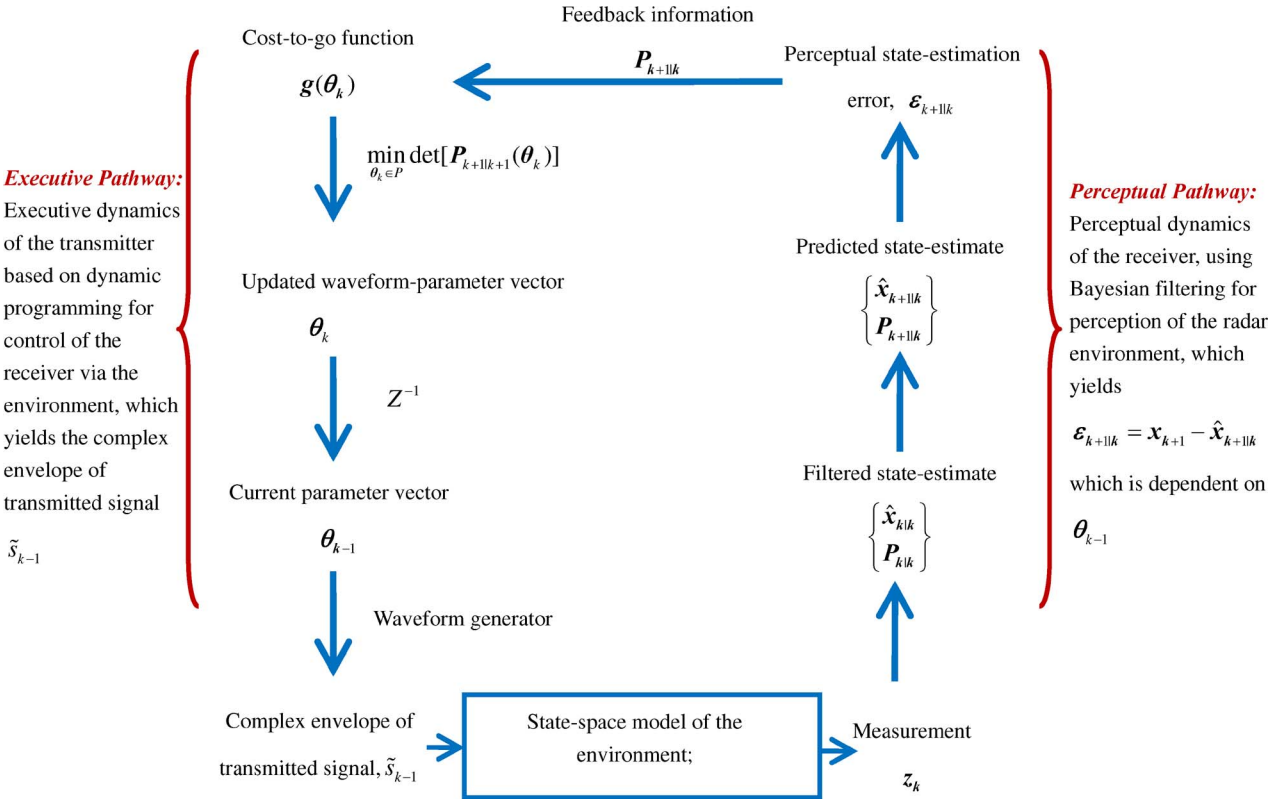

Fig. 5. *The basic cycle of directed information flow graph in its most generic form. The symbol Z^{-1} represents bank of unit-time delays.*

updated noise covariance $\mathbf{R}_k(\boldsymbol{\theta}_k)$. Accordingly, the optimization described in (24) is perfectly feasible in an online manner.

The executive pathway, centered on dynamic programming in the transmitter, is depicted on the left-hand side of Fig. 5. This path describes the *executive dynamics* of the transmitter, a primary function of which is to compute the "new" waveform-parameter vector $\boldsymbol{\theta}_k$ for use at the next cycle.

C. How Can We Build on the Directed Information Flow Graph to Better Understand the Role of Memory in Cognition?

First and foremost, Fig. 5 teaches us how the perception–action cycle in CR progresses from one cycle to the next, highlighting the mathematical roles played by the Bayesian filter for perception of the environment carried out in the receiver, feedback information from the receiver to the transmitter, and Bellman's dynamic programming aimed at actuation in the radar environment. From control theory, we may go on to make the statement:

> The amount of information contained in the posterior expected estimation error is progressively reduced from one cycle to the next by virtue of the feedback embodied in perception–action cycle acting alone.

The next not-so-obvious issue to explain is: how does perceptual memory come into play to further improve the radar performance? To answer this question, we remind ourselves of the role of perceptual memory in cognition. To this end, we say that this role is trying to find the particular element in the system-model library that is the closest match to every new measurement vector $\mathbf{z}_k$, for all k. In so doing, the Bayesian filter in the receiver becomes equipped with the right description of the state–space model for every measurement vector; consequently, a further reduction in the estimation error is brought into play on each cycle and target-tracking accuracy is thereby improved further.

Turning next to the transmitter: how does the executive memory help us to improve radar performance? Here again, the improvement is attributed to a matching process, albeit differently. To be specific, the function of the executive memory is to select that particular element in the transmit-waveform library that is the closest match to the feedback information. This feedback information provides an indirect means for the transmitter to sense the environment via the receiver. Hence, by selecting the right

transmit waveform to match the feedback information for every measurement vector $\mathbf{z}_k$, the transmitter is matched adaptively to the radar environment.

This process, matching the transmitted waveform to the radar environment on a cycle-by-cycle basis in CR, is the very essence of what Gjessing described in his classic book, entitled *Target Adaptive Matched Illumination Radar* over 25 years ago [6]. Gjessing stated that there is usually enough information available about most of the targets that a radar will face regarding their size, shape, rigidity, surface material composition, and motion pattern. If the radar exploits this prior information about the target, it can detect and identify the target faster, which is what CR provides.

Finally, turning to working memory: what is its role in the cognitive process? Referring back to Fig. 1, we see that the working memory is reciprocally coupled to the executive memory on the left and reciprocally coupled to the perceptual memory on the right. The net result of this reciprocal coupling is that CR assumes the role of a synchronous information-processing machine, which is beyond the capability of TAR or FAR, with all the practical benefits that resulted from this synchrony.

XI. EXPERIMENTAL GROUNDWORK

The stage is now set for an experimental study of CR and comparison of it to other related radar configurations. To this end, we consider an air-traffic-control problem, the objective of which is to track the trajectory of an aircraft that executes a maneuver at (nearly) constant speed and turn rate in free space. We are interested in this problem because:

1) air traffic control is important for both military and civilian applications;
2) CR performs well for a tracking problem with continuous system equation and discrete measurement equation.

In [34], this tracking scenario was used to compare the performance of three different estimators: the continuous-discrete versions of the extended Kalman filter, the unscented Kalman filter, and the cubature Kalman filter. Therein, it was shown that the continuous-discrete cubature Kalman filter (CD-CKF) outperforms the other two by wide margins, hence its adoption as the candidate for state estimation in additive white Gaussian noise for our experimental study. The objective of the study is to show how the tracking accuracy is improved by embedding cognition of varying complexity into the FAR.

The Appendix summarizes the important parameters and other related matters pertaining to the simulation.

A. State–Space Model

In the aviation language, the scenario considered in this study is commonly referred to as tracking of a target with (nearly) *coordinated turn* [5]. In this scenario, motion in the horizontal plane and motion in the vertical plane are assumed to be decoupled. Hence, we can write the coordinated turn in a 3-D space, subject to fairly small noise modeled by independent Brownian motions, as shown in (1).

The 7-D state vector of the aircraft is described as follows:

$$\mathbf{x} = [\epsilon, \dot{\epsilon}, \eta, \dot{\eta}, \zeta, \dot{\zeta}, \omega]^T \tag{35}$$

with ϵ, η, and ζ denoting positions and $\dot{\epsilon}$, $\dot{\eta}$, and $\dot{\zeta}$ denoting velocities in the x, y, and z Cartesian coordinates, respectively. The parameter ω denotes the turn rate; the time evolution of the state is governed by the nonlinear function

$$\mathbf{f}(\mathbf{x}) = [\dot{\epsilon}, -\omega\dot{\eta}, \dot{\eta}, \omega\dot{\epsilon}, \dot{\zeta}, 0, 0]^T.$$

The noise term is defined as

$$\boldsymbol{\beta}(t) = [\beta_1(t), \beta_2(t), \ldots, \beta_7(t)]^T \tag{36}$$

with $\beta_i(t)$, $i = 1, 2, \ldots, 7$, being all mutually independent standard Brownian motions, which account for unpredictable modeling errors; and finally, the diffusion matrix is

$$\mathbf{Q} = \mathrm{diag}\big([0, \sigma_1^2, 0, \sigma_1^2, 0, \sigma_1^2, \sigma_2^2]\big).$$

It is assumed that the radar is located at the origin and equipped to measure the range ρ and azimuth angle θ with a measurement sampling time T. Hence, the measurement equation can be written as

$$\begin{bmatrix} \rho_k \\ \theta_k \end{bmatrix} = \begin{bmatrix} \sqrt{\epsilon_k^2 + \eta_k^2 + \zeta_k^2} \\ \tan^{-1}\left(\dfrac{\eta_k}{\epsilon_k}\right) \end{bmatrix} + \mathbf{w}_k$$

where the measurement noise is Gaussian $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_k)$ with $\mathbf{R}_k = \mathrm{diag}([\sigma_\rho^2, \sigma_\theta^2])$.

To simulate the target trajectory, we use $\sigma_1 = \sqrt{0.5}$; $\sigma_2 = \sqrt{5e-7}$; $\sigma_\rho = 50$ m; $\sigma_\theta = 0.001°$; and the initial state is $\mathbf{x}_0 = [10$ km, 100 m/s, 10 km, 150 m/s, 5 km, 0 ms^{-1}, 30 °/s$]^T$. Fig. 6 shows the ideal trajectory of the target for these parameters.

We use the notation $\mathbf{x}_k^j$ to denote $\mathbf{x}(t)$ at time $t = t_k + j\Delta$, where $1 \le j \le m$ and $\Delta = T/m$ with m being

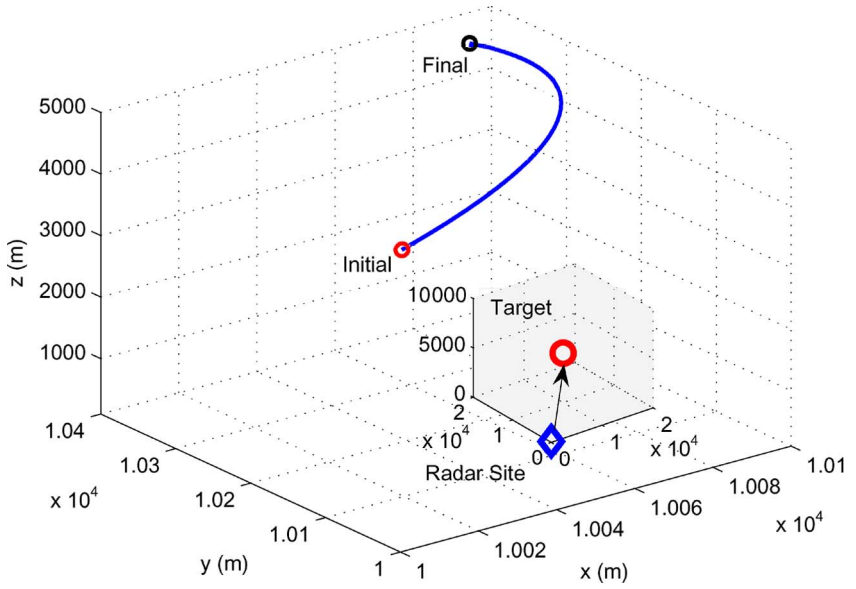

Fig. 6. *Target trajectory, experiencing a corner turn.*

an integer. Applying the Itô–Taylor expansion of order 1.5 to (1), we get the stochastic difference equation

$$\mathbf{x}_k^{(j+1)} = \mathbf{f}_d\left(\mathbf{x}_k^j\right) + \sqrt{\mathbf{Q}}\boldsymbol{\omega} + \left(\mathbb{L}\mathbf{f}\left(\mathbf{x}_k^j\right)\right)\psi \tag{37}$$

where

$$\mathbf{f}_d(\mathbf{x}) = \begin{bmatrix} \epsilon + \Delta\dot{\epsilon} - \Delta^2\omega\dot{\eta}/2 \\ \dot{\epsilon} - \Delta\omega\dot{\eta} - \Delta^2\omega^2\dot{\epsilon}/2 \\ \eta + \Delta\dot{\eta} + \Delta^2\omega\dot{\epsilon}/2 \\ \dot{\eta} + \Delta\omega\dot{\epsilon} - \Delta^2\omega^2\dot{\eta}/2 \\ \zeta + \Delta\dot{\zeta} \\ \dot{\zeta} \\ \omega \end{bmatrix}$$

and

$$\mathbb{L}\mathbf{f}(\mathbf{x}) = \begin{bmatrix} 0 & \sigma_1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -\sigma_1\omega & 0 & 0 & -\sigma_2\dot{\eta} \\ 0 & 0 & 0 & \sigma_1 & 0 & 0 & 0 \\ 0 & \sigma_1\omega & 0 & 0 & 0 & 0 & -\sigma_2\dot{\epsilon} \\ 0 & 0 & 0 & 0 & 0 & \sigma_1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

To generate independent trajectories, we use $m = 1000$ time-steps/sampling interval.

B. Construction of the Two Libraries

The experimental results presented in this section are based on simulating an *X*-band radar with carrier frequency $f_c = 10.4$ GHz. We use LFM with both upsweep and down-sweep chirps, which constitute the waveform library as follows:

$$\mathcal{P} = \{\lambda \in [10e-6, 100e-6], \quad b \in [-100e9, -10e9] \cup [10e9, 100e9]\} \tag{38}$$

with grid step sizes $\Delta_\lambda = 10e-6$ and $\Delta_b = 10e9$. The assumed waveform is the LFM pulse combined with Gaussian amplitude modulation. The sampling frequency is set to $f_s = 400$ Hz and the update frequency of the cognitive tracking algorithm is set to 20 Hz. This completes the design of the waveform library block in Fig. 1.

To design the system-model library block in Fig. 1, we gather together all the unknown parameters in (37). A careful examination of (37) reveals that we need to gather only three parameters σ_1, σ_2, and ω, where the first two represent the modeling error for velocity, and the last one represents the corner turn. The direct benefit we can gain from this operation is the reduced dimensionality of the system-model library from seven to three.

C. Performance Metric

We use the ensemble-averaged root-mean-square error (RMSE) as the metric to evaluate the performance of TAR, FAR, and CR with one layer of memory, as well as

multiscale memory. The ensemble-averaged RMSE for the range is defined by

$$\aleph_{\mathrm{p}}(k)=\sqrt{\frac{1}{N}\sum_{n=1}^{N}\left(\left(\epsilon_k^n-\hat{\epsilon}_k^n\right)^2+\left(\eta_k^n-\hat{\eta}_k^n\right)^2+\left(\zeta_k^n-\hat{\zeta}_k^n\right)^2\right)}$$

where $(\epsilon_k^n, \eta_k^n, \zeta_k^n)(\hat{\epsilon}_k^n, \hat{\eta}_k^n, \hat{\zeta}_k^n)$ are, respectively, the true and estimated positions at time index k in the nth Monte Carlo run. In a similar manner, we may also define the ensemble-averaged RMSE for velocity, denoted by $\aleph_{\mathrm{v}}(k)$.

XII. EXPERIMENTAL RESULTS: THEORETICAL CONSIDERATIONS

This section presents two sets of simulation results, namely the PCRLB and tracking accuracy. The simulations are performed for three different radar configurations: TAR, FAR, and CR, where the cognitive building blocks are added to the radar one after another in the following order: perception–action cycle, memory, and attention. Although the proposed successive addition of elements of cognition is a logical approach, trying other possible combinations would be insightful for understanding the relative value of each element. For instance, in the proposed structure, attention is memory driven but it can be tested in isolation without the memory elements. In that case, we will have a fully adaptive radar equipped with attention. However, in order to have a radar system, which is truly cognitive, all of the cognitive building blocks must be in use.

A. PCRLB

To have an appreciation of the lower bound, Fig. 7(a) and (b) plots the square root of the PCRLB and the standard deviation of the estimation error versus time for range and range rate estimates, respectively. It confirms the theoretical prediction that, by adding the cognitive building blocks one after another to the radar, the bound has been pushed lower and lower. The initial transients in the square root of the PCRLB are mainly due to the fact that the perceptual associative memory initially experiences *a moment of indecision* [60] about the system model and noise characteristics, and it takes a few iterations until it settles. The following is obvious from Fig. 7(a) and (b).

1) The standard deviation of the estimation error achieved by different members of the CR family are lower than the square root of the PCRLB of TAR. It means that all members of the CR family have been able to go beyond what a traditional radar can ever potentially achieve and provide a level of accuracy, which is better than the theoretical limit imposed on a traditional radar.
2) Increasing the degree of cognitivity of the radar improves its performance. As shown in Fig. 7(a) and (b), CR with multiscale memory (CRm) performs better than CR with one layer of memory (CR1), which, in turn, performs better than FAR that uses only feedback. The fully CR with memory and attention (CRa) performs better than CR with one layer of memory but slightly worse than CR with multiscale memory. Hence, it settles for suboptimality by slightly sacrificing performance for conserving computational resources.
3) These figures demonstrate that although the theoretical PCRLB is close to zero and the estimation error variance was quite improved by cognitizing the radar, the error variance does not meet the lower bound. This is reasonable because of issues related to design and implementation of some components such as memory. However, it leaves a wide margin for researchers interested in CR to improve on.

B. Tracking Accuracy

The performance of the cognitive tracker can be best illustrated by the RMSE and bias of the estimates. Fig. 8(a) and (b) depicts the ensemble-averaged RMSE curve of the range and range rate, i.e., $\aleph_{\mathrm{p}}$ and $\aleph_{\mathrm{v}}$, respectively. Biases of the range and range-rate estimates are shown in Fig. 9(a) and (b). Examination of the results presented in these two figures leads us to the following observations.

1) As the fundamental building blocks of cognition are added to the radar one by one, the radar becomes more and more sophisticated, starting from TAR, FAR, to CR with single memory and then a multiscale memory, thereby resulting in better and better system performance. To elaborate on our claim, inclusion of the perception–action cycle decreases the RMSE significantly. The inclusion of memory into CR further pushes the RMSE curve to a much lower level.
2) With increased memory scale, two comments have to be made. First, the accuracy of CR is also enhanced. The accumulative RMSE for range and range rate is summarized in Table 2. From the results presented in this table, we see that CR system with multiscale memory can achieve an accuracy of 0.47 m for range and 3.30 m/s for range rate, which are far beyond the reach of traditional radar systems. Second, the computation is dramatically reduced. In [8], the complexity of the waveform selection part is shown to be in the order of $\mathcal{O}((N_g)^L)$, where N_g is the size of the waveform library and L is the length of horizon of dynamic programming. It is obvious that when the *explore–exploit* strategy is used, searching of the waveform library is constrained to the local neighborhood of current waveform, which then greatly reduces the computational

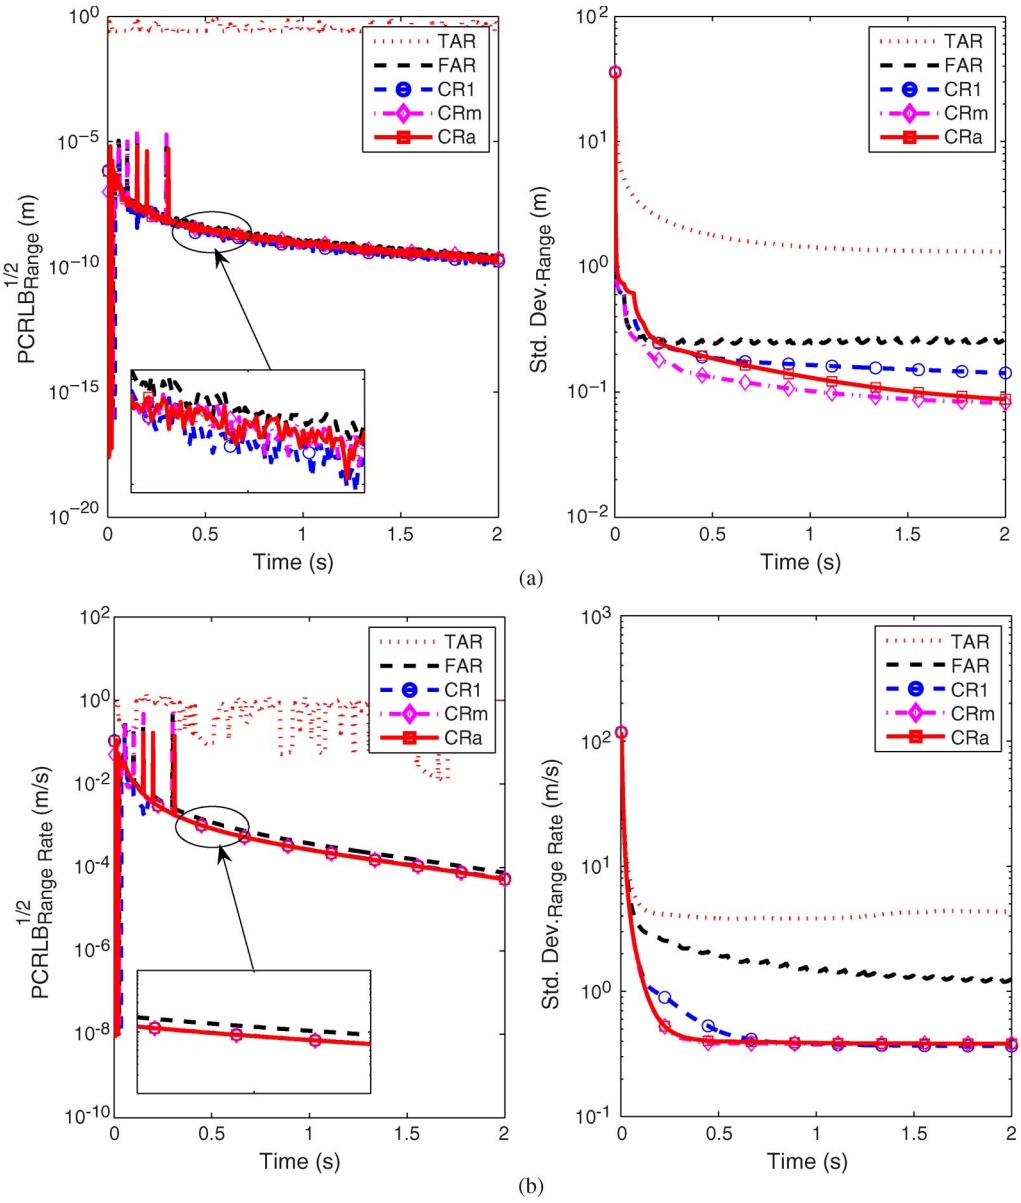

Fig. 7. ***Square root of PCRLB versus standard deviation for (a) range and (b) range-rate estimation errors.***

complexity at the cost of a reduction in tracking performance. In Table 2, it is not difficult to see that the accuracy of CR with multiscale memory and attention is reduced to 0.69 m for range and 3.53 m/s, validating the no-free-lunch theorem.

Based on the results presented in Table 2, we make two more observations.

1) For range estimation, in going from TAR to FAR, we have an order of magnitude improvement; then in going from FAR to CR with one level of memory, we have another order of magnitude improvement; and in going onto CR with multiscale memory, we have a further improvement of 18.97%. However, with the addition of attention to CR, the range accuracy is dropped by 0.22 m but computational complexity is reduced.
2) For range-rate estimation, the improvements achieved through the use of memory are much more profound. Specifically, in going from TAR to FAR, we have improved the accuracy by about 73.78%; then in going from FAR to CR with one level of memory, we have another improvement of about 68.56%; and in going onto CR with multiscale memory, we have a further improvement of almost 21.62%. Here again, with the addition of attention to CR, the range-rate accuracy is reduced by 0.23 m/s.

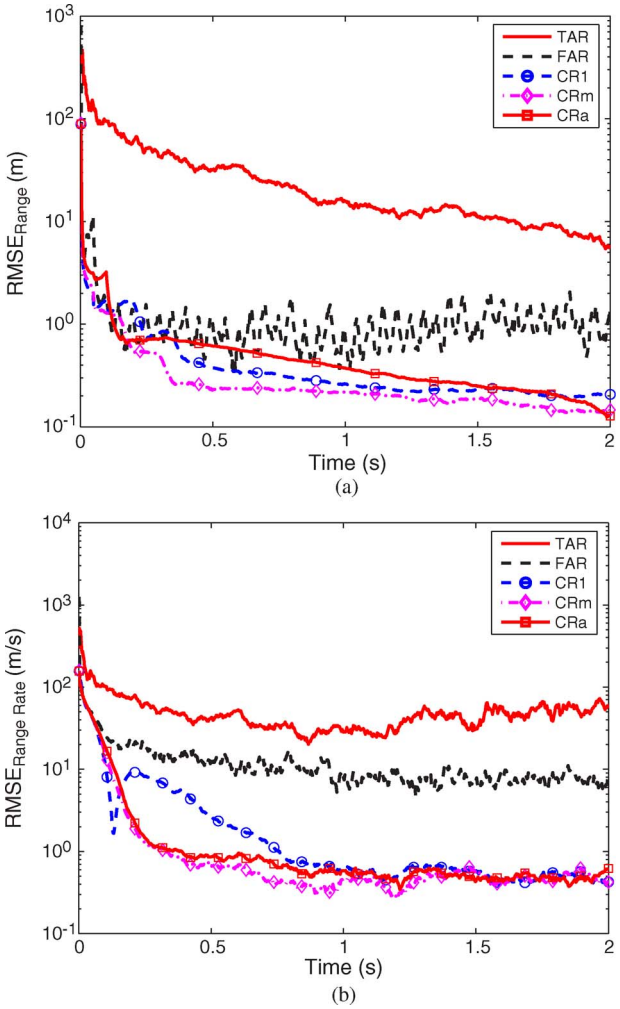

Fig. 8. *RMSE of (a) range and (b) range-rate estimates for TAR, FAR, CR with one layer of memory (CR1), CR with multiscale memory (CRm), and CRm with executive attention (CRa).*

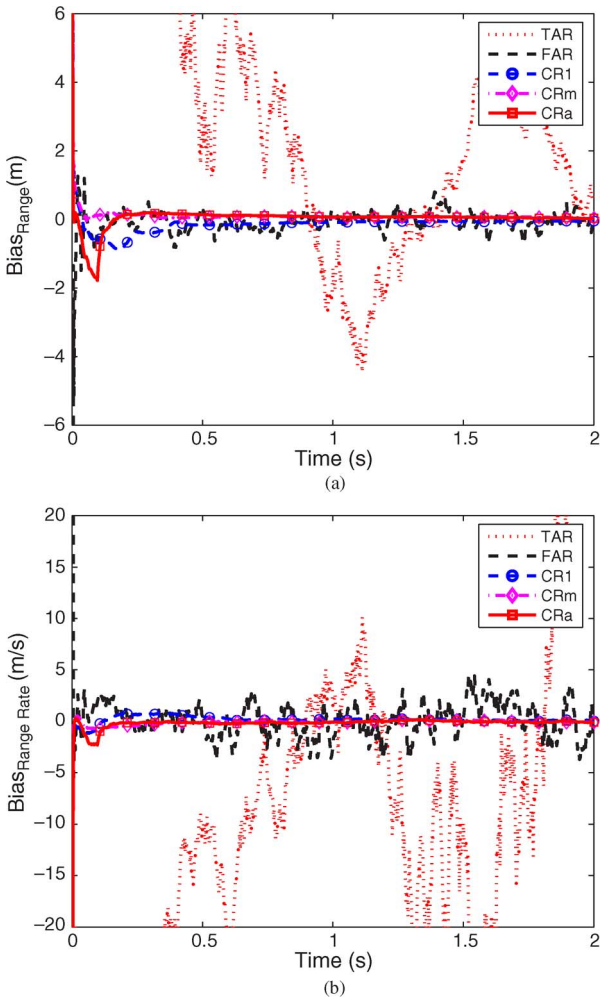

Fig. 9. *Bias of (a) range and (b) range-rate estimates for TAR, FAR, CR with one layer of memory (CR1), CR with multiscale memory (CRm), and CRm with executive attention (CRa).*

A question of interest is: how, in physical terms, do we explain the ground-breaking results reported on CR in Fig. 8(a) and (b)? The answer to this important question is partly contributed to the provision of two libraries in CR, depicted in Fig. 1, namely the system-model library in the receiver and the transmit-waveform library in the transmitter. These two libraries may be viewed as *prior information*, or knowledge to be more precise. As explained in Section X, availability of the system-model library enables the receiver to select the system model that best matches every new measurement on a cycle-by-cycle basis, via the perceptual memory. Similarly, availability of the transmit-waveform library enables the transmitter to select the particular transmit waveform that adaptively matches the environment in the best manner possible.

In a way, by virtue of the feedback link from the receiver to the transmitter, FAR can avail itself of a transmit-waveform library, and thereby adaptively match the environment for every new measurement on a cycle-by-cycle basis. However, FAR does not have memory and therefore lacks the capability available to CR over FAR, as reported in Fig. 8(a) and (b).

Finally, turning to TAR, it has no prior information about the environment, be it in the receiver or the transmitter, hence its inferior tracking performance compared to FAR, again as reported in Fig. 8(a) and (b).

Table 2 Accumulative RMSE for Range and Range Rate

accumulative RMSE	TAR	FAR	CR1	CRm	CR
range(m)	29.24	2.12	0.58	0.47	0.69
range-rate(m/s)	51.07	13.39	4.21	3.30	3.53

XIII. EXPERIMENTAL RESULTS: PRACTICAL CONSIDERATIONS

Previously, in Section VIII, we have pointed out that attention improves temporal stability of the time rate of change in the transmit waveform. With this point in mind, we now try to provide insight into the behavior of different radar configurations. Specifically, in Fig. 10, the waveform transitions for chirp rates are plotted for different scenarios. In order to study the waveform transition from an optimization point of view, we randomly choose reference time steps and plot their corresponding error surfaces, with Figs. 11–14 depicting the error surfaces corresponding to FAR, CR with one layer of memory, CR with multiscale memory, and CR with multiscale memory and attention (CRa). The highly revealing results depicted in these figures lead to the following conclusions.

1) In FAR, the controller's decision for choosing the optimal transmit waveform switches between two sets of waveform parameters. As time passes, the controller is not able to settle on one waveform and keeps switching between the two, and therefore, shows a bistable behavioral pattern (Fig. 10). As shown in Fig. 11, there are two local minima on the error-performance surface, which correspond to these two sets of waveform parameters. At different instants, the controller, seeking the minimum of the error surface, will end up in one of these two points. It will be easier to locate the minimum of the error-performance surface by looking at the contours on the waveform-parameter plane (horizontal plane) rather than the 3-D surface.
2) In CR with memory (both single layer and multiscale), the controller switches between two sets of waveform parameters for a while and then settles down on only a single one. In other words,

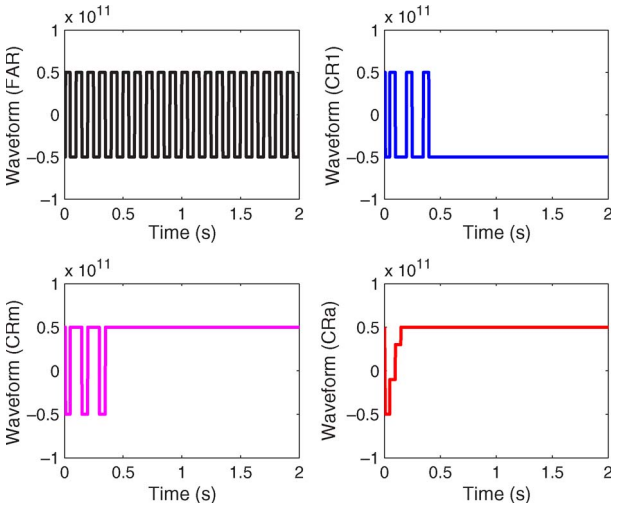

Fig. 10. *Chirp-rate selection for FAR, CR with one layer of memory (CR1), CR with multiscale memory (CRm), and CRm with executive attention (CRa).*

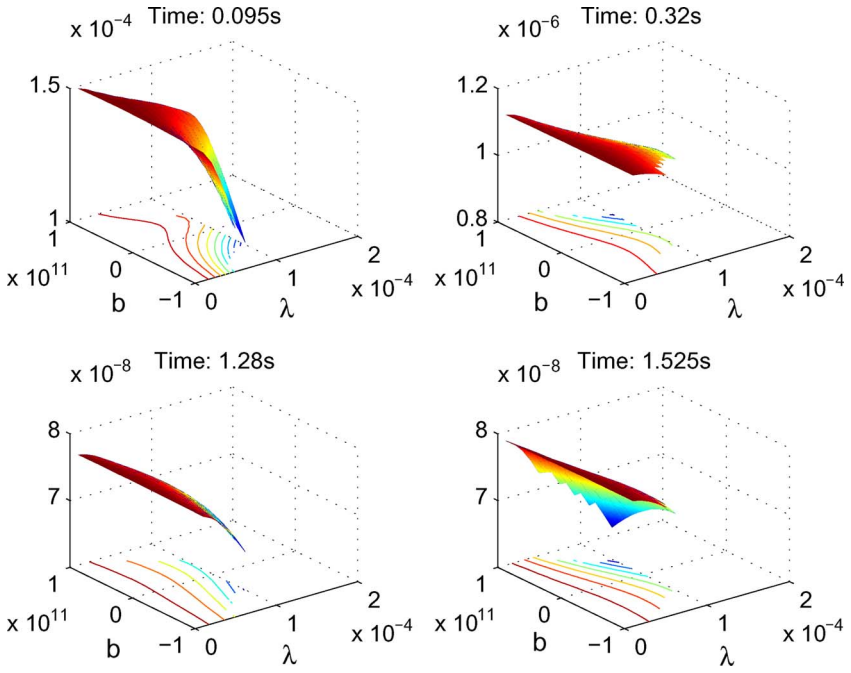

Fig. 11. *Error surface at different time steps for FAR.*

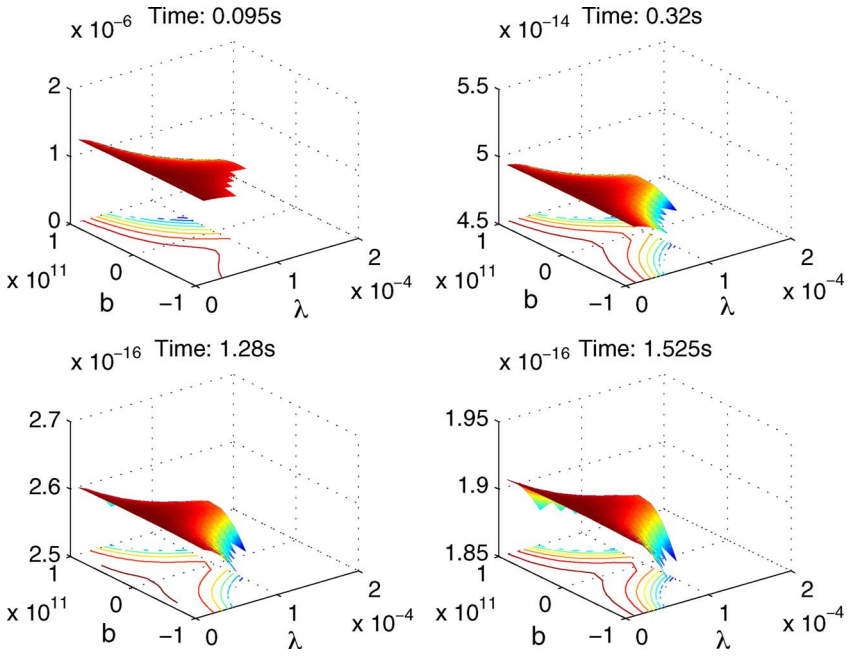

Fig. 12. *Error surface at different time steps for CR with one-layer of memory.*

it shows a monostable behavioral pattern (Fig. 10). The corresponding error-performance surfaces are shown in Figs. 12 and 13. The location of the minimum point of the error-performance surface can be easily recognized from error contours on the waveform-parameter plane.

3) In CR with memory and attention, we observe a dramatic result: although the performance of CR with attention degrades slightly, the waveform has a much smoother transition, which in practical cases is an advantageous factor for the radar transmitter (Fig. 10). Furthermore, when CR is equipped with

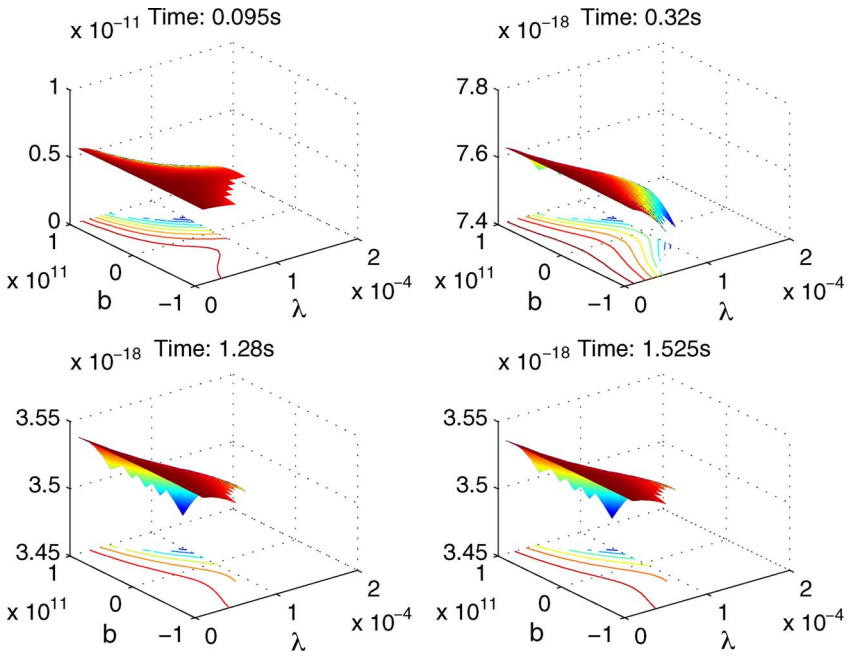

Fig. 13. *Error surface at different time steps for CR with multiscale memory.*

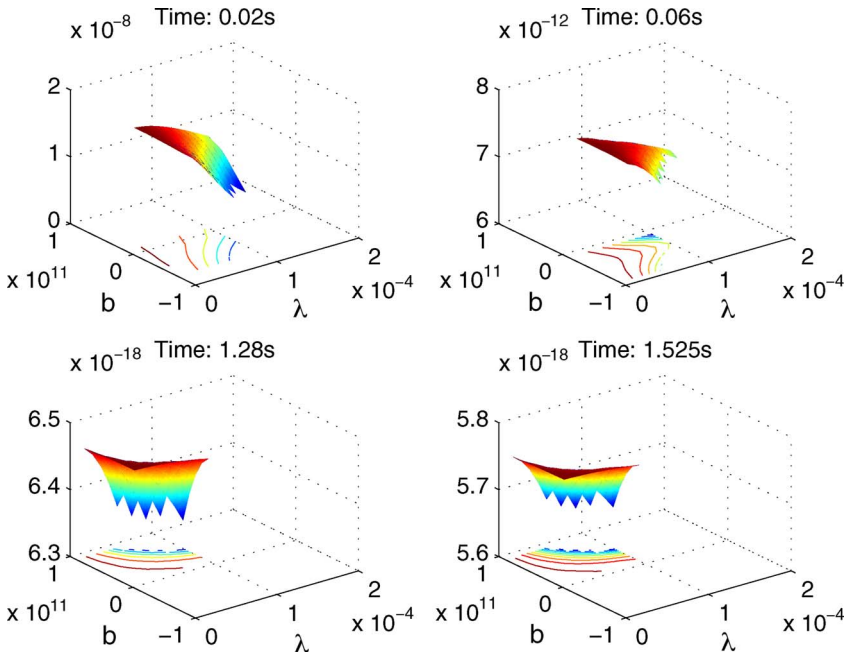

Fig. 14. ***Error surface at different time steps for CR with executive attention.***

executive attention, the (sub)optimal waveform has been selected locally, which means the saving of computational resources. Here again, the combination of these two factors validates the no-free-lunch theorem. Since the optimal waveform is selected based on a local search among the neighbors of the current grid point, the 3-D error-performance surface in Fig. 14 shows a pattern with multiple-local minima, prompting us to make the following statement:

> Each one of these local minima coincides with one of the neighbors of the current grid point and the best one among them is chosen as the next grid point.

Although intelligence in CR, empowered by the perception–action cycle, memory, and attention, makes a difference in tracking accuracy, it is in the attainment of a smooth transition of the transmit waveform from one cycle to the next where CR distinguishes itself from TARs and FARs.

XIV. CONCLUDING REMARKS

In this paper, inspired by the visual brain, we have proposed a novel way of building cognitively controlled sensing into radar systems. CR is built around a perception–action cycle, which is the first stage to cognition. While the perception part is based on a Bayesian filter and therefore seeks Bayesian optimality, the action part is based on an optimal controller in the sense of Bellman's dynamic programming. By adding memory and attention to the perception–action cycle, the radar system is enabled to reach a level of performance under a wide range of environmental conditions that can be interpreted as showing intelligent behavior in the true sense of the word. Different building blocks of such a radar system were discussed in detail and guidelines for designing each building block were proposed. Theoretical results suggest that the proposed structure for CR has the potential to significantly improve the tracking performance of current radar systems. Computer experiments show that CR can indeed go beyond the theoretical limits of traditional radars.

Moreover, the experimentally demonstrated benefit, gained from the combined use of memory and executive attention in CR for target tracking, is the smooth manner in which the controller selects the transmit waveform from the transmitter's library as the radar progresses in moving forward from one cycle to the next, mimicking what goes on in the visual brain. This new capability of practical importance, enabled by cognition, is beyond that of TAR or FAR.

In dynamic terms, we may describe the attributes of memory in CR as follows:

- memory is predictive, predicting the consequence of action;
- executive attention, driven by memory, smooths out the transition of transmit waveform from one cycle to the next.

In mathematical terms, it would be fair to say that we now have a strong basis for the perception–action cycle. However, recognizing that in an overall sense, CR is in its

early stages of development, there is much to be done in other areas that require further research. To be specific, we may mention three specific areas:

1) the important role that memory plays in modeling and design of CR;
2) an equally important role that learning through interaction with the environment plays in improving the performance of CR on the fly;
3) a rigorous mathematical theory of CR, viewed as a complex system of systems.

All three areas were beyond the scope of this paper.

Among the four functional building blocks of cognition under Fuster's paradigm, we have singled out memory for more detailed research because it is the block that brings CR that much closer to the visual brain.

Also, there is room for extending the proposed framework to consider scenarios, which include the presence of clutter in the radar returns, the probability of target detection, and the presence of more than one target in the area of interest as in a multitarget tracking case.

One final comment is in order. A bat can pursue and capture its target with a facility and success rate that would be the envy of a radar engineer [2]. In [61], the time–frequency representation of the bat sonar signal is provided in the four phases of action, which are hunting, approach, pursuit, and capture of a prey. The ability of bat's small brain, about the size of a plum, to adaptively change the transmitted waveform parameters to its actual need is fascinating. In the design of radar systems, it would be of interest to take further steps toward mimicking mother nature by refining Fuster's paradigm of cognition to account for bat's four phases of action. ■

APPENDIX
SUMMARY OF IMPORTANT PARAMETERS AND OTHER RELATED MATTERS OF THE SIMULATION

It is assumed that the radar is located 100 m above the ground and simulation results are based on 50 Monte Carlo runs of the experiment.

A. Track Initialization

In CD-CKF, the dimensionality of the state vector was chosen to be 7. To initialize the CD-CKF algorithm, the initial state density was assumed to be Gaussian. Therefore, the two-point differencing method was adopted [5], where the first two measurements were used to estimate the states' statistics. PCRLB was initialized, based on an initial guess made by the perceptual associative memory about a model and the initial state estimate.

B. Memory

The multiscale memory is designed using a replicator network. For the perceptual memory, the replicator network is made up of an input layer of source nodes, three hidden layers, and an output layer, whose sizes are 20, 10, 5, 10, and 20, respectively.

The weights of the replicator network were initialized as uniformly distributed small numbers on the interval of $(-0.125, 0.125)$. The learning rate was linearly annealed on the region of $[10^{-1}, 10^{-5}]$. To train the replicator network, a number of 600 training data sets were selected and the number of epochs was set as 100.

When training of the replicator network is done, the decoder is no longer needed and a new output layer of size 2 was connected to the encoder. The LMS algorithm was adopted to train the weights between the encoder and this new output layer. The learning rate was defined as $\eta = 0.1$. Furthermore, the number of epochs was also set as 100. However, for every epoch, if the MSE was found to be lower than the predefined threshold of 10^{-10}, the algorithm was terminated.

Likewise, the executive memory has exactly the same configuration and follows the same training procedure. As for the bias input used in the replicator networks for both the perceptual memory and the executive memory, it was fixed at 1 during the training stage. The working memory is responsible for the bias input at the running stage.

Furthermore, the working memory reciprocally couples the perceptual memory and the executive memory by connecting their bottleneck layers together. To train the working memory, 2000 data sets were collected from bottleneck layers for both the perceptual memory and the executive memory.

Acknowledgment

The detailed feedback notes on different aspects of the paper by four reviewers have not only reshaped the paper into its present form but also made the paper the best it could be.

REFERENCES

[1] B. R. Frieden, *Science From Fisher Information: A Unification*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[2] S. Haykin, "Cognitive radar: A way of the future," *IEEE Signal Process. Mag.*, vol. 23, no. 1, pp. 30–40, Jan. 2006.

[3] S. Haykin, "New generation of radar systems enabled with cognition," presented at the IEEE Radar Conf., Arlington, VA, May 2010, Keynote Lecture.

[4] D. J. Kershaw and R. J. Evans, "Optimal waveform selection for tracking systems," *IEEE Trans. Inf. Theory*, vol. 40, no. 5, pp. 1536–1550, Sep. 1994.

[5] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation With Applications to Tracking and Navigation*. New York: Wiley-Interscience, 2001.

[6] D. T. Gjessing, *Target Adaptive Matched Illumination Radar: Principles and Applications*, London, U.K.: Peter Peregnins Ltd. on behalf of the Institution of Electrical Engineers, 1986.

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

[8] S. Haykin, A. Zia, Y. Xue, and I. Arasaratnam, "Control-theoretic approach to tracking radar: First step towards cognition," *Digital Signal Process.*, vol. 21, no. 5, pp. 576–585, 2011.

[9] V. Krishnamurthy and D. V. Djonin, "Optimal threshold policies for multivariate POMDPs in radar resource management," *IEEE Trans.*

Signal Process., vol. 57, no. 10, pp. 3954–3969, Oct. 2009.

[10] S. L. Bressler and V. Menon, "Large-scale brain networks in cognition: Emerging methods and principles," *Trends Cogn. Sci.*, vol. 14, pp. 277–290, 2010.

[11] V. B. Mountcastle, "The columnar organization of the neocortex," *Brain*, pp. 701–722, 1997.

[12] V. B. Mountcastle, "An organizing principle for cerebral function: The unit model and the distributed system," in *The Mindful Brain*, G. M. Edelman and V. B. Mountcastle, Eds. Cambridge, MA: MIT Press, 1978.

[13] V. B. Mountcastle, *Perceptual Neuroscience: The Cerebral Cortex*. Cambridge, MA: Harvard Univ. Press, 1998.

[14] D. Marr, "A theory for cerebral neocortex," *Proc. Roy. Soc. Lond. B, Biol. Sci.*, vol. 176, pp. 161–234, 1972.

[15] J. M. Fuster, *Cortex and Mind: Unifying Cognition*. Oxford, U.K.: Oxford Univ. Press, 2003.

[16] O. Sporns, D. R. Chialvo, M. Kaiser, and C. C. Hilgetag, "Organization, development and function of complex brain networks," *Trends Cogn. Sci.*, vol. 8, no. 9, pp. 418–425, 2004.

[17] N. Wiener, *Cybernetics: Or the Control and Communication in the Animal and the Machine*, 2nd ed. Cambridge, MA: MIT Press, 1965.

[18] N. Wiener, *The Human Use of Human Beings: Cybernetics and Society*. New York: Houghton Mifflin, 1950.

[19] W. S. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bull. Math. Biol.*, vol. 5, no. 4, pp. 115–133, 1943.

[20] W. G. Walter, *The Living Brain*. New York: Norton, 1953.

[21] M. A. Arbib, *The Handbook of Brain Theory and Neural Networks*, 2nd ed. Cambridge, MA: MIT Press, 2003.

[22] S. Haykin, *Neural Networks and Learning Machines*, 3rd ed. Englewood Cliffs, NJ: Prentice-Hall, 2009.

[23] D. O. Hebb, *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley, 2003.

[24] M. M. Mesulam, "Large scale neurocognitive networks and distributed processing for attention, language, and memory," *Neurobiol. Progr.*, pp. 597–613, 1990.

[25] P. T. Fox and K. J. Friston, "Distributed processing; distributed functions?" *NeuroImage*, vol. 61, pp. 407–426, 2012.

[26] M. L. Anderson, "Neural reuse: A fundamental organizational principle of the brain," *Behav. Brain Sci.*, vol. 33, pp. 245–313, 2010.

[27] S. Haykin, *Cognitive Dynamic Systems: Perception-Action Cycle, Radar and Radio*. Cambridge, U.K.: Cambridge Univ. Press, 2012.

[28] B.-T. Vo, B.-N. Vo, and A. Cantoni, "Bayesian filtering with random finite set observations," *IEEE Trans. Signal Process.*, vol. 56, no. 4, pp. 1313–1326, Apr. 2008.

[29] S. Soatto, "Actionable information in vision," in *Proc. Int. Conf. Comput. Vis.*, 2009, pp. 2138–2145.

[30] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*. New York: Academic, 1970.

[31] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*. Berlin, Germany: Springer-Verlag, 1999.

[32] D. J. Kershaw and R. J. Evans, "Waveform selective probabilistic data association," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 33, no. 4, pp. 1180–1188, Oct. 1997.

[33] Y. C. Ho and R. C. K. Lee, "A Bayesian approach to problems in stochastic estimation and control," *IEEE Trans. Autom. Control*, vol. AC-9, no. 4, pp. 333–339, Oct. 1964.

[34] I. Arasaratnam, S. Haykin, and T. R. Hurd, "Cubature Kalman filter for continuous-discrete systems: Theory and simulations," *IEEE Trans. Signal Process.*, vol. 58, no. 10, pp. 4977–4993, Oct. 2010.

[35] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. New York: Wiley-Interscience, 2006.

[36] J. C. Spall, *Introduction to Stochastic Search and Optimization*. New York: Wiley-Interscience, 2003.

[37] H. L. Van-Trees, *Detection, Estimation and Modulation Theory, Part I*. New York: Wiley, 1968.

[38] J. Dauwels, "Computing Bayesian Cramér-Rao bounds," in *Proc. Int. Symp. Inf. Theory*, 2005, pp. 425–429.

[39] H. L. Van-Trees and K. L. Bell, *Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking*. New York: Wiley/IEEE Press, 2007.

[40] P. Tichavský, C. H. Muravchik, and A. Nehorai, "Posterior Cramér-Rao bounds for discrete-time nonlinear filtering," *IEEE Trans. Signal Process.*, vol. 46, no. 5, pp. 1386–1396, May 1998.

[41] C. J. Masreliez and R. D. Martin, "Robust Bayesian estimation for the linear model and robustifying the Kalman filter," *IEEE Trans. Autom. Control*, vol. AC-22, no. 3, pp. 361–371, Jun. 1977.

[42] P. J. Huber and E. M. Ronchetti, *Robust Statistics*, 2nd ed. New York: Wiley, 2009.

[43] S. A. Kassam and H. V. Poor, "Robust techniques for signal processing: A survey," *Proc. IEEE*, vol. 73, no. 3, pp. 433–481, Mar. 1985.

[44] R. K. Boel, M. R. James, and I. R. Petersen, "Robustness and risk-sensitive filtering," *IEEE Trans. Autom. Control*, vol. 47, no. 3, pp. 451–461, Mar. 2002.

[45] W. H. Fleming and W. M. Mceneaney, "Risk sensitive and robust nonlinear filtering," in *Proc. 36th IEEE Conf. Decision Control*, San Diego, CA, 1997, pp. 1088–1093.

[46] R. J. Meinhold and N. D. Singpurwalla, "Robustification of Kalman filter models," *J. Amer. Stat. Assoc.*, vol. 84, pp. 479–486, 1989.

[47] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Nashua, NH: Athena Scientific, 2005.

[48] O. G. Selfridge, "Pandemonium: A paradigm for learning," in *Proc. Symp. Mech. Thought Processed*, Nov. 1958, pp. 513–526.

[49] T. Kohonen, *Associative Memory: A System-Theoretical Approach*. New York: Springer-Verlag, 1977.

[50] G. E. Hinton and J. A. Anderson, *Parallel Models of Associative Memory*. London, U.K.: Lawrence Erlbaum, 1989.

[51] M. R. W. Dawson, *Minds and Machines: Connectionism and Psychological Modeling*. Oxford, U.K.: Blackwell, 2004.

[52] A. Baddeley, "Working memory: Looking back and looking forward," *Nature Rev. Neurosci.*, vol. 4, pp. 829–839, 2003.

[53] Y. Bengio and Y. LeCun, "Scaling learning algorithms towards AI," in *Large-Scale Kernel Machines*, L. Bottou, O. Chapelle, D. DeCoste, and J. Weston, Eds. Cambridge, MA: MIT Press, 2007.

[54] R. Hecht-Nielsen, "Replicator neural networks for universal optimal source coding," *Science*, vol. 269, pp. 1860–1863, 1995.

[55] B. Widrow and M. A. Lehr, "30 years of adaptive neural networks: Perceptron, madaline, and backpropagation," *Proc. IEEE*, vol. 78, no. 9, pp. 1415–1442, Sep. 1990.

[56] S. Haykin, *Adaptive Filter Theory*, 4th ed. Englewood Cliffs, NJ: Prentice-Hall, 2002.

[57] S.-I. Amari, N. Murata, K. R. Müller, M. Finke, and H. H. Yang, "Asymptotic statistical theory of overtraining and cross-validation," *IEEE Trans. Neural Netw.*, vol. 8, no. 5, pp. 985–996, Sep. 1997.

[58] N. Brady and D. J. Field, "Local contrast in natural images: Normalisation and coding efficiency," *Perception*, vol. 55, no. 8, pp. 1041–1055, 2000.

[59] D. H. Wolpert and W. G. Macready, "No free lunch theorems for optimization," *IEEE Trans. Evol. Comput.*, vol. 1, no. 1, pp. 67–82, Apr. 1997.

[60] A. O'Hagan, "A moment of indecision," *Biometrika*, vol. 68, no. 1, pp. 329–330, 1981.

[61] F. Hlawatsch and G. F. Boudreaux-Bartels, "Linear and quadratic time-frequency signal representations," *IEEE Signal Process. Mag.*, vol. 9, no. 2, pp. 21–67, Apr. 1992.

ABOUT THE AUTHORS

Simon Haykin (Fellow, IEEE) received the B.Sc. (first class honors), Ph.D., and D.Sc. degrees in electrical engineering from the University of Birmingham, Birmingham, U.K.

Having worked in signal processing applied to radar and communications for a good part of his professional life, all along in the past 15 years, he had the vision of revisiting the fields of radar and communications from a brand new perspective. That vision became a reality in the early years of this century with the publication of two seminal journal papers: 1) "Cognitive radio: Brain-empowered wireless communications," IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (February 2005); and 2) "Cognitive radar: A way of the future," IEEE SIGNAL PROCESSING MAGAZINE (January 2006). Cognitive radio and cognitive radar are two important parts of cognitive dynamic systems, on which his new book was published in March 2012. He is the author/coauthor of close to 50 books.

Prof. Haykin is a Fellow of the Royal Society of Canada. In 1999, he was awarded the Honorary Degree of Doctor of Technical Sciences by ETH, Zurich, Switzerland. In 2002, he was the first recipient of the Booker Gold Medal, which was awarded by the International Scientific Radio Union (URSI).

Yanbo Xue (Member, IEEE) received the B.Sc. degree in automation engineering and the M.A.Sc. degree in control theory and engineering from Northeastern University, Shenyang, China, in 2001 and 2004, respectively, and the Ph.D. degree in electrical and computer engineering from McMaster University, Hamilton, ON, Canada, in 2010. His Ph.D. dissertation titled "Cognitive radar: Theory and simulations" was the first doctoral thesis on cognitive radar.

He is currently working as a Postdoctoral Fellow at the McMaster Institute for Automotive Research and Technology (MacAuto). His major research interests are on two aspects: 1) smart grid and hybrid plug-in electric vehicle (HPEV), which includes modeling and optimization of microgrid and smart grid, energy management systems, and vehicle-to-grid (V2G) technology; and 2) cognitive dynamic systems with application to cognitive radar, target detection and tracking, neural networks and learning machines, control theory, and array signal processing.

Dr. Xue was the recipient of Young Scientist Travel Grant (YSTG) of the 2004 International Symposium on Antennas and Propagation (ISAP) held in Japan.

Peyman Setoodeh (Member, IEEE) received the B.Sc. and M.Sc. (first class honor) degrees in electrical engineering from Shiraz University, Shiraz, Iran, and the Ph.D. degree in computational engineering and science from McMaster University, Hamilton, ON, Canada.

He is currently a Postdoctoral Fellow at McMaster University, where he collaborates with both the Cognitive Systems Laboratory (CSL) and the Centre for Mechatronics and Hybrid Technology (CMHT) Research. Also, he is a Lecturer in the Department of Electrical and Computer Engineering. His research interests include cognitive machines, complex adaptive systems, nonlinear estimation and control, game theory, and optimization.

Dr. Setoodeh is a recipient of the Monbukagakusho Scholarship from the Ministry of Education, Culture, Sports, Science and Technology in Japan.