1 Introduction

I see the emergence of a new discipline, called Cognitive Dynamic Systems, which builds on ideas in statistical signal processing, stochastic control, and information theory, and weaves those well-developed ideas into new ones drawn from neuroscience, statistical learning theory, and game theory. The discipline will provide principled tools for the design and development of a new generation of wireless dynamic systems exemplified by cognitive radio and cognitive radar with efficiency, effectiveness, and robustness as the hallmarks of performance.

This quotation[1] is taken from a point-of-view article that appeared in the *Proceedings of the Institute of Electrical and Electronics Engineers* (Haykin, 2006a). In retrospect, it is perhaps more appropriate to refer to Cognitive Dynamic Systems as an "integrative field" rather than a "discipline."

By speaking of cognitive dynamic systems as an integrative field, we mean this in the sense that its study integrates many fields that are rooted in neuroscience, cognitive science, computer science, mathematics, physics, and engineering, just to name a few. Clearly, the mixture of fields adopted in the study depends on the application of interest. However, irrespective of the application, the key question is

What is the frame of reference for justifying that a so-called cognitive dynamic system is indeed cognitive?

In this book, we adopt *human cognition* as *the frame of reference*. As for applications, the book focuses on cognitive radar and cognitive radio.

With these introductory remarks, the study begins with the next section.

1.1 Cognitive dynamic systems

A system, be it linear or nonlinear, is said to be *dynamic* if *time* plays a key role in its input–output behavior. In this book, we are interested in a new class of dynamic systems called *cognitive dynamic systems*, the study of which is inspired by the unique neural computational capability of the human brain[2] and the viewpoint that human cognition is a form of computation.[3]

To be specific, we say that a dynamic system, operating in an environment to be explored, is *cognitive* if it is capable of four fundamental functions (tasks) that are basic to *human cognition*:

(1) the perception–action cycle;
(2) memory;

(3) attention; and
(4) intelligence.

The perception–action cycle implies that a cognitive dynamic system has two functional parts; namely, *perceptor* and *actuator*. The cycle begins with the perceptor perceiving the environment (world) by processing the incoming stimuli, called *observables* or *measurements*. In response to *feedback information* from the perceptor about the environment, the actuator acts so as to control the perceptor via the environment, and the cycle goes on. In effect, the perceptor "guides" the actuator by virtue of what it has *learned* about the environment, and the actuator "controls" the perceptor by acting in the environment. The benefit resulting from the perception–action cycle is that of *maximizing information gained* from the environment.

Typically, the environment is *nonstationary*, which means that the underlying behavior of the environment continually changes with time. Given such an environment to deal with, we may now go on to say that a cognitive dynamic system *must also have memory*, desirably of a multiscale variety. This requirement is needed for the system to do the following:

- *learn* from the environment and store the knowledge so acquired;
- continually *update* the stored knowledge in light of environmental changes; and
- *predict* the consequences of actions taken and/or selections made by the system as a whole.

As for attention, a cognitive dynamic system must be equipped with the capability to focus its information-processing power on a target or footprint in the environment that is considered to be of special interest or strategic importance; this is done by *prioritizing* the allocation of available resources.

Turning to intelligence, among the above-mentioned four functions, it is by far the most difficult one to describe. Nevertheless, intelligence is the single most important function in a cognitive dynamic system. For the present, it suffices to say that intelligence is based on the perception–action cycle, memory, and attention for its functionality. Most importantly, it is the presence of feedback at multiple levels in the system that facilitates intelligence, which, in turn, makes it possible for the system to make *intelligent decisions* in the face of inevitable uncertainties in the environment. The feedback can itself take one of two forms:

- *global feedback*, which embodies the environment, and
- *local feedback*, which does not.

Typically, the local and global feedback loops are distributed throughout a cognitive dynamic system. The extent of feedback loops naturally varies from one cognitive dynamic system to another, depending on the application of interest.

1.2 The perception–action cycle

In diagrammatic form, much of what we have just described is captured so illustratively in Figure 1.1. On the right-hand side of the figure, we have the perceptor of a cognitive dynamic system that is responsible for perception of the environment. On the left-hand side of the

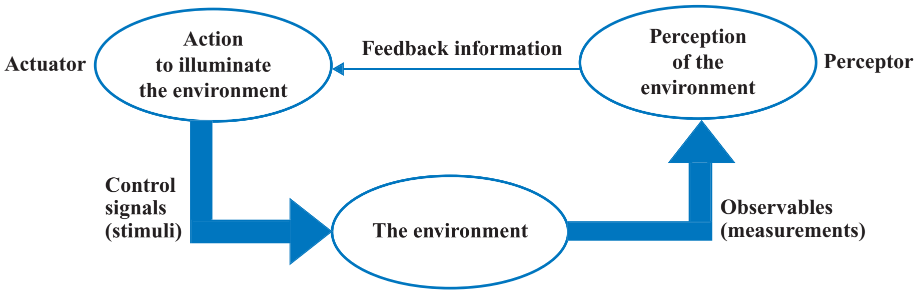

Figure 1.1. The perception–action cycle of a cognitive dynamic system in its most generic sense.

figure, we have the actuator of the system that is responsible for action in the environment. And, above all, we have a feedback link connecting the perceptor to the actuator.

Figure 1.1 is commonly referred to as the *perception–action cycle* of a cognitive dynamic system. This remarkable operation, in its generic form, stands out in *human cognition*, which, as mentioned at the beginning of this introductory chapter, embodies the following (Sejnowski, 2010):

- a global feedback loop, embodying perception and action so as to maximize information gain about the environment;
- multiscale memory that is organized to predict the consequences of actions;
- a memory-based attentional mechanism that prioritizes the allocation of resources; and
- a feedback-based decision-making mechanism that identifies intelligent choices in uncertain environments.

We are repeating what we described earlier in order to emphasize that these four distinctive properties of human cognition constitute the *ideal framework*, against which a dynamic system should be assessed for it to be cognitive.

1.3 Cognitive dynamic wireless systems: radar and radio

A new generation of engineering systems is being inspired by human cognition, the structures of which vary in details from one system to another, depending on the application of interest. In this context, two dynamic wireless systems stand out:

- *cognitive radar*, for improved performance in remote-sensing applications for system accuracy and reliability; and
- *cognitive radio*, for solving the underutilized electromagnetic spectrum problem.

1.3.1 Cognitive radar

There is a remarkable analogy between the *visual brain* and radar. Indeed, the perception–action cycle of Figure 1.1 applies equally well to cognitive radar by merely changing the ways in which the transmitter (actuator) and receiver (perceptor) are actually implemented.

Specifically, the function of the receiver in a radar system is to produce an *estimate of the state* of an unknown target located somewhere in the environment by processing a *sequence of observables* dependent on the target state. In effect, perception of the environment takes the form of *state estimation*. As for the transmitter in the system, its function is to adaptively select a transmitted waveform that illuminates the environment in the best manner possible. In target detection, the issue of interest is to decide as reliably as possible whether a target is present or not in the observables. In target tracking, on the other hand, the issue of interest is to estimate the target parameters (e.g. range and velocity) as accurately as possible.

With radar intended for remote-sensing applications and with its transmitter and receiver being typically collocated, much can be learned from the human brain to make a radar system cognitive.

1.3.2 Cognitive radio

The practical use of cognitive radio is motivated by the desire to address the *electromagnetic spectrum underutilization problem*. In today's wireless communications world, we typically find that only a small fraction of the radio spectrum assigned to legacy operators by government agencies is actually employed by primary (licensed) users. The underutilized subbands of the spectrum are commonly referred to as *spectrum holes*. The function of a cognitive radio may then be summarized as follows:

(1) The radio receiver is equipped with a *radio scene analyzer*, the purpose of which is to identify where the spectrum holes are located at a particular point in time and space.
(2) Through an external feedback link from the receiver to the transmitter, the information on spectrum holes is then passed to the radio transmitter, which is equipped with a *dynamic spectrum manager* and *transmit-power controller*. The function of the transmitter is to allocate the spectrum holes among multiple secondary (cognitive radio) users in accordance with prioritized needs.

Unlike radar, where the transmitter and receiver are ordinarily collocated, in a radio (wireless) communication system the transmitter and receiver are located in different places. Accordingly, for the receiver to send the transmitter information about the spectrum holes, we require the use of a *low-bandwidth feedback link* connecting the receiver to the transmitter.

With radio intended for wireless communications and with its transmitter and receiver being separately located, there is still a great deal we can learn from the human brain; but to make the radio cognitive, we have to use *engineering ingenuity*.

1.4 Illustrative cognitive radar experiment

We will now motivate the power of cognitive information-processing by considering a simple cognitive radar tracker.

The function of the receiver is to estimate the state of a target in space and thereby track its motion across time, given a set of observables (measurements) obtained on the target. With this objective in mind, a sequential state estimator suffices to take care of the perceptive needs of the receiver.

To elaborate on the sequential state estimator, for the sake of simplicity we assume that the environment is described by a *linear state-space model*, comprised of the following pair of equations:

$$\mathbf{x}(n) = \mathbf{A}\mathbf{x}(n-1) + \boldsymbol{\omega}(n), \tag{1.1}$$

$$\mathbf{y}(n) = \mathbf{B}\mathbf{x}(n) + \mathbf{v}(\boldsymbol{\theta}_{n-1}), \tag{1.2}$$

where n denotes *discrete time*. The vector $\mathbf{x}(n)$ denotes the *state* of the environment at time n; the state evolution across time in (1.1) is called the *system equation*. The vector $\mathbf{y}(n)$ denotes the *measurements* recorded digitally by the receiver as input at time n, hence the reference to (1.2) as the *measurement equation*. Transition of the state at time $n-1$ to that at time n is described by the matrix $\mathbf{A}$. Correspondingly, dependence of the measurements (observables) on the state at time n is described by the *measurement matrix* $\mathbf{B}$.

Naturally, the imposition of a mathematical model on the environment, as described in (1.1) and (1.2), gives rise to *uncertainties* about the physical behavior of the environment. These uncertainties are accounted for by introducing the *system noise* $\boldsymbol{\omega}(n)$ in (1.1) and *measurement noise* $\mathbf{v}(\theta_{n-1})$ in (1.2). The measurement noise is denoted by $\mathbf{v}$, the composition of which is dependent on the action of the transmitter. That action is controlled by a transmit-waveform parameter vector $\boldsymbol{\theta}_{n-1}$; the reason for (partially or in full) assigning time $n-1$ to this vector is to account for the propagation delay between the transmitter and receiver.

In what follows, we assume that the process noise $\omega(n)$ and measurement noise $\mathbf{v}(\boldsymbol{\theta}_{n-1})$ are both *stationary Gaussian processes* of zero mean; their respective covariance matrices are denoted by $\mathbf{Q}$ and $\mathbf{R}_\theta$. Invoking these assumptions and recognizing that the state-space model described in (1.1) and (1.2) is linear, it follows that the solution to the sequential state-estimation problem – that is, estimating the state $\mathbf{x}(n)$ given the sequence of observables $\{\mathbf{y}(i)\}_{i=1}^{n}$ – is to be found in the classic *Kalman filter*. The issue of sequential state estimation is discussed in Chapter 4. As such, in this introductory chapter, we will proceed on the premise that we know how to formulate the Kalman filtering algorithm. We may, therefore, go on to say that, with tracking as the issue of interest, the Kalman filter, formulated on the basis of the state-space model of (1.1) and (1.2), adequately fulfills the perceptive needs of the receiver.

1.4.1 The experiment

The target parameters to be estimated are the *delay* and *Doppler shift*. The delay τ is defined in terms of the range ρ (i.e. distance of the target from the radar) by

$$\tau = \frac{2\rho}{c}, \tag{1.3}$$

where c is the speed of electromagnetic wave propagation (i.e. the speed of light). The Doppler shift f_D is defined in terms of the range rate $\dot{\rho}$ (i.e. velocity of the target) by

$$f_\mathrm{D} = \frac{2 f_\mathrm{c}\dot{\rho}}{c}, \tag{1.4}$$

where f_c is the transmitted carrier frequency and the dot in $\dot{\rho}$ denotes differentiation with respect to time.

1.4.2 The environment

The unknown target is located at a distance of 3 km from the radar and it is moving at a speed of 200 m/s.

1.4.3 The radar

The radar is an X-band radar operating at the frequency f_c = 10.4 GHz; it is located at the origin. The 0 dB *signal-to-noise ratio* (SNR) at the receiver input is defined at 80 km.

1.4.4 State-space model

The state-space model of (1.1) and (1.2) is parameterized as follows:

$$\mathbf{A} = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix},$$

$$\mathbf{B} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

where T is the sampling period. The system noise is modeled as the target-acceleration noise, with its zero-mean covariance defined by (Bar-Shalom *et al*., 2001)

$$\mathbf{Q} = \sigma_v^2 \begin{bmatrix} \frac{1}{4}T^4 & \frac{1}{2}T^3 \\ \frac{1}{2}T^3 & T^2 \end{bmatrix},$$

where the variance $\sigma_v^2 = 0.49$.

1.4.5 Simulation results

Figure 1.2 presents the results of Monte Carlo simulations performed to evaluate the tracking performance of three different radar configurations:[4]

(1) *Traditional active radar* with a fixed transmit waveform, in which case the measurement equation (1.2) simplifies to

$$\mathbf{y}(n) = \mathbf{B}\mathbf{x}(n) + v(n). \tag{1.5}$$

It is assumed that the matrices **A** and **B** and the covariance matrices **Q** and **R** are all known.

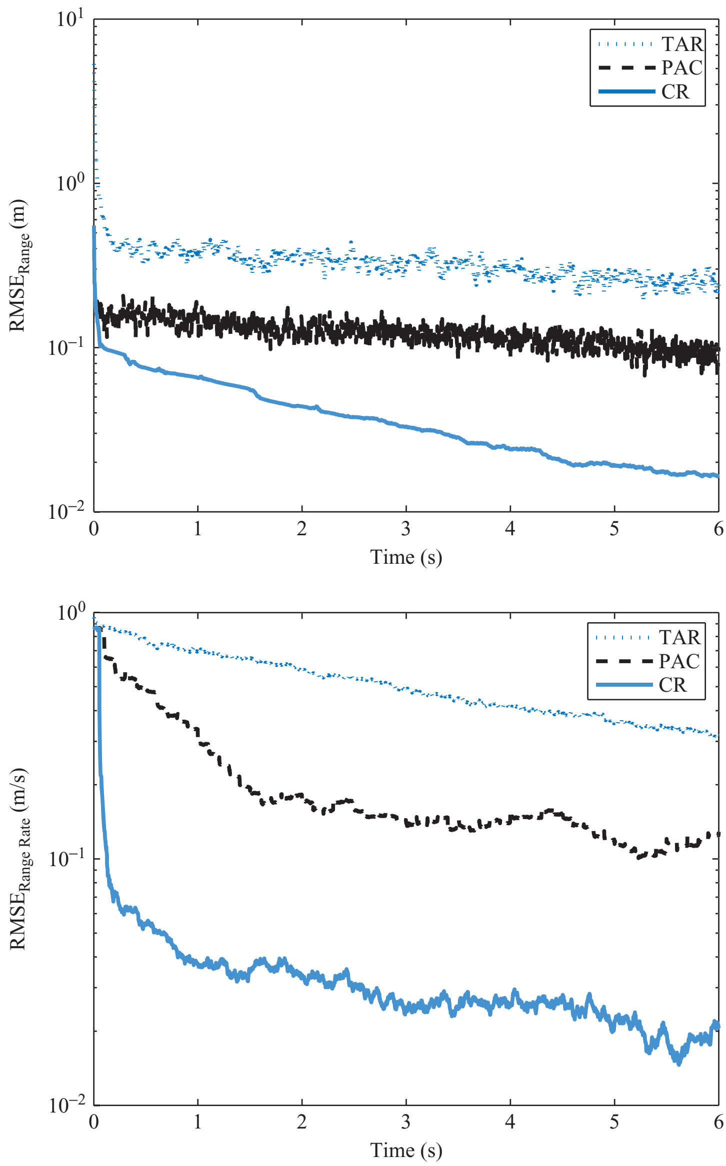

Figure 1.2. Demonstrating the information-processing power of global feedback and cognition in radar tracking. (a) Root mean-squared-error (RMSE) of target range, measured in meters. (b) RMSE of range rate, measured in meters per second. TAR: traditional active radar; PAC: the perception–action cycle, as in the first stage toward cognition in radar, or equivalently, the fore-active radar; CR: cognitive radar.

(2) *Perception–action cycle mechanism*, which is the first step toward radar cognition; the same mechanism is applicable to a second class referred to as the *fore-action radar*.[4] Accordingly, the measurement equation (1.2) holds.

(3) *Cognitive radar*, the transmitter of which is equipped not only with a transmit-waveform library but also another library in the receiver; the second library

makes it possible for the receiver to select appropriate values for the matrix **A** and covariance matrix **Q** in the system equation (1.1); that is, continually *model* the environment.

Figure 1.2a presents the RMSE for the target range and Figure 1.2b the RMSE for the target range-rate plotted versus time. In both parts of the figure, the top dashed curves refer to the traditional active radar, the middle dashed bold curves refer to the perception–action cycle mechanism acting alone (i.e. fore-active radar), and the bottom full curves refer to the cognitive radar.

The results presented in Figure 1.2 lead us to report two important findings:

(1) The use of global feedback in the perception–action mechanism acting alone makes a significant difference in the tracking accuracy of the radar, compared with the traditional active radar with no feedback.
(2) The addition of memory to the perception–action mechanism as in the cognitive radar brings in even more significant improvement to tracking accuracy of the radar.

1.5 Principle of information preservation

Having just reported findings (1) and (2) on how the fore-active radar and the cognitive radar compare with a traditional active radar in terms of tracking accuracy, we may now pose a related question:

Over and above the improvements in tracking accuracy, what else do the results of Figure 1.2 teach us?

1.5.1 Feedback information

Our first step in answering this fundamental question is to reiterate that the environmental state of the target consists of two parameters:

(1) *range* ρ, which defines how far away the target is from the radar;
(2) *range rate* $\dot{\rho}$, which defines the velocity of the target.

Since both the system and measurement equations, (1.1) and (1.2), are corrupted by additive noise processes, it follows that both the range ρ and range rate $\dot{\rho}$ are *random variables*. According to Shannon's information theory (Shannon, 1948), we may, therefore, say that *information* about the target's state is contained in the sequence of measurements $\{\mathbf{y}(i)\}_{i=0}^{n}$. Moreover, in view of this statement, we may go on to speak of *feedback information*, defined in terms of the *error* between the actual state of the target and its estimate computed by the receiver as the result of operating on the sequence of measurements. It is by virtue of this feedback information passed to the transmitter by the receiver that the feedback loop around the environment is closed.

1.5.2 Bayesian filtering of the measurements

With state estimation as a central issue of interest in cognitive radar, we look to the *Bayesian filter as the optimal recursive data-processing algorithm*. Basically, the Bayesian filter combines all the available measurements (data) plus *prior* knowledge about the system and measuring devices and uses them all to produce the *optimal estimate* of hidden target parameters. To perform this estimation, the Bayesian filter propagates the *posterior* (i.e. probability density function of the state, conditioned on the measurements) from one estimation recursion to the other. The rationale for focusing on the posterior is that it contains *all the information* about the state that is available in the measurements. Hence, the optimal estimate is obtained by *maximizing* the posterior, which is the "best" that can be achieved (Ho and Lee, 1964).

Under the combined assumption of linearity and Gaussianity, the Bayesian filter reduces to the *Kalman filter*, hence its choice as the functional block for processing the measurements in the receiver in the motivational experiment of Section 1.4.

1.5.3 Information preservation through cognition

Moving on to the transmitter for action in the environment, the feedback information highlighted earlier in this section provides a basis of a *cost-to-go function* that looks to the future by one time step. Recognizing that this function is also dependent on the parameters that define the transmitted waveform, a primary function of the transmitter, therefore, is to select a set of transmit-waveform parameters that minimizes the cost-to-go function. Thus, for every cycle of the radar's perception–action cycle, the transmit-waveform parameters are selected such that perception of the radar environment in the receiver and action in the environment performed in the transmitter are optimized in a sequential manner. This process is repeated on a cycle-by-cycle basis.

To assess the overall radar performance, we need a *metric* that provides an assessment of how close the estimated state of the target is to its actual value. For the experiment, following the traditional approach in statistical signal processing, this metric was defined simply as the RMSE between the actual state and its estimated value using the Kalman filter, with both of them defined for one step into the future. What we have just described here provides the mathematical justification for improved tracking accuracy through the use of global feedback from the receiver to transmitter, reported previously under point (i) at the end of Section 1.4.

Next, through the following combination of system additions:

- *perceptual memory*, reciprocally coupled with the Kalman filter in the receiver;
- *executive memory*, reciprocally coupled with the transmit-waveform selector in the transmitter; and
- *working memory*, reciprocally coupling the executive and perceptual memories,

the transmitter and receiver are continuously matched together in their respective operations in an adaptive manner. It is through this adaptive process distributed in different parts of the radar receiver and transmitter that we are able to explain the additional significant accuracy improvement reported under point (2) at the end of Section 1.4.

Now, we are ready to answer the question that was posed at the beginning of this section. In using the Kalman filter by itself in the receiver, information about the state of the target contained in the measurements is *preserved to some extent* (Kalman, 1960). In adaptively *matching illumination* of the environment with the target through the use of feedback from the receiver to the transmitter (Gjessing, 1986), information about the state contained in the measurements is *preserved even more*. Finally, in making the radar cognitive through the provision of distributed memory, *further improvement in information preservation is achieved* by the cognitive radar system.

What we have just described here is now summed up as follows:

> Cognitive information processing provides a powerful tool for fulfilling the principle of information preservation, which is aimed at preserving information about a hidden target state that is contained in the measurements.

In the statement just made on the principle of information preservation, we have emphasized the role of cognition. Information preservation may also be viewed as *information gain* in the following sense: the more we preserve information contained in the measurements about the target's state, the closer the estimated state is to its actual value. This is another way of saying that we are progressively "gaining" information about the target's state from one cycle to the next in following the perception–action cycle. By saying so, we have justified the previous use of "information gain" in describing the role of a global feedback loop in the perception–action cycle.

1.5.4 Concluding remarks

To conclude this discussion, it should be noted that successive improvements in information preservation, exemplified by corresponding improvements in state-estimation accuracy, are achieved at the expense of increased computational complexity. It is *apropos*, therefore, that we complement our previous statement by saying:

> There is "no free lunch," in that for every gain we make in practice there is a price to be paid.

1.6 Organization of the book

The rest of the book is organized in seven chapters, as summarized here.

Chapter 2 is devoted to a detailed discussion of the *perception–action cycle*, which is the *baseline for the operation of every cognitive dynamic system*. This chapter also identifies three kinds of memory:

- *perceptual memory*, which is an integral part of the receiver;
- *executive memory*, which is an integral part of the transmitter; and
- *working memory*, which reciprocally couples the executive memory to perceptual memory.