Simon Haykin

Cognitive Radar

A way of the future

This article discusses a new idea called cognitive radar. Three ingredients are basic to the constitution of cognitive radar: 1) intelligent signal processing, which builds on learning through interactions of the radar with the surrounding environment; 2) feedback from the receiver to the transmitter, which is a facilitator of intelligence; and 3) preservation of the information content of radar returns, which is realized by the Bayesian approach to target detection through tracking. All three of these ingredients feature in the echo-location system of a bat, which may be viewed as a physical realization (albeit in neurobiological terms) of cognitive radar.

Radar is a remote-sensing system that is widely used for surveillance, tracking, and imaging applications, for both civilian and military needs. In this article, we focus on future possibilities of radar with particular emphasis on the issue of cognition. As an illustrative case study along the way, we consider the problem of radar surveillance applied to an ocean environment.

According to the *Oxford English Dictionary*, cognition is "knowing, perceiving, or conceiving as an act." Given three distinct capabilities:

- the inherent ability of radar to sense its environment on a continuous basis
- the ability of phased-array antennas to electronically scan the environment in a fast manner
- the ever-increasing power of computers to digitally process signals

it is our conviction that it is indeed feasible to build a cognitive radar system using today's technology. Indeed, if ever there was a remote-sensing system well suited for cognition, radar is it.

From the moment a surveillance radar system is switched on, the system becomes electromagnetically linked to its surrounding environment in the sense that the environment has a strong and continuous influence on the radar returns (i.e., echoes). In so doing, the radar builds up its knowledge of the environment

 1053-5888/06/$20.00©2006IEEE

from one scan to the next and makes decisions of interest on possible targets at unknown locations in the environment; the locations are not known before the radar is switched on, but they become determined by the radar receiver once the targets under surveillance are declared. From signal processing and control theory, we know that it is not necessary for the radar to keep the entire record of past data. Rather, by adopting a state-space model of the environment and recursively updating the state vector representing an estimate of certain parameters pertaining to the environment, the need for storing the entire history of radar data on the environment is eliminated.

The requirement to update estimation of the environmental state is necessitated by the fact that the radar environment is nonstationary. A primary cause of nonstationarity is statistical variations in the weather and the presence of unknown targets at unknown locations. Recursive updating of a state is synonymous with adaptivity, which is the natural method for dealing with nonstationarity. In current designs of radar systems, however, adaptivity is usually confined to the receiver. For the radar to be cognitive, adaptivity has to be extended to the transmitter too. Moreover, the radar has to learn from experience on how to deal with different targets, large and small, and at widely varying ranges, all in an effective and robust manner.

COGNITIVE SIGNAL PROCESSING CYCLE

The Oxford dictionary definition of cognition includes conceiving, which might be taken to mean "the formulation of a hypothesis, and then testing that hypothesis for the likelihood of its correctness." This statement is in the spirit of the Bayesian approach to state estimation, with a probabilistic rating of alternatives. We are therefore emboldened to embrace the idea of Bayesian inference under the umbrella of cognitive radar.

This way of thinking leads us to the block diagram of Figure 1, which depicts the picture of a cognitive cycle performed by a cognitive radar system. The cycle begins with the transmitter illuminating the environment. The radar returns produced by the environment are fed into two functional blocks: radar-scene analyzer and Bayesian target-tracker. The tracker makes decisions on the possible presence of targets on a continuing time basis, in light of information on the environment provided to it by the radar-scene analyzer. The transmitter, in turn, illuminates the environment in light of the decisions made on possible targets, which are fed back to it by the receiver. The cycle is then repeated over and over again. Unlike a communication system, the feedback mechanism, which is a necessary requirement of a cognitive system, is easy to implement as the radar transmitter and receiver are usually co-located. Note also that although the process of target detection is not explicitly shown in the cognitive cycle of Figure 1, it is part and parcel of the Bayesian target-tracker, which performs detection through tracking as explained later.

Based on Figure 1, a cognitive radar distinguishes itself from an adaptive radar in three important respects.

- The radar continuously learns about the environment through experience gained from interactions with the environment and, in a corresponding way, continually updates the receiver with relevant information on the environment.
- The transmitter adjusts its illumination of the environment in an intelligent manner, taking into account such practical matters as the size of the target and its range, and consequently, making adjustments to the transmitted signal in an effective and robust manner.
- The whole radar system constitutes a dynamic closed feedback loop encompassing the transmitter, environment, and receiver.

It is well known that feedback is like a double-edged sword in that it can become harmful if it is used improperly; care must therefore be exercised in how the transmitter is designed in

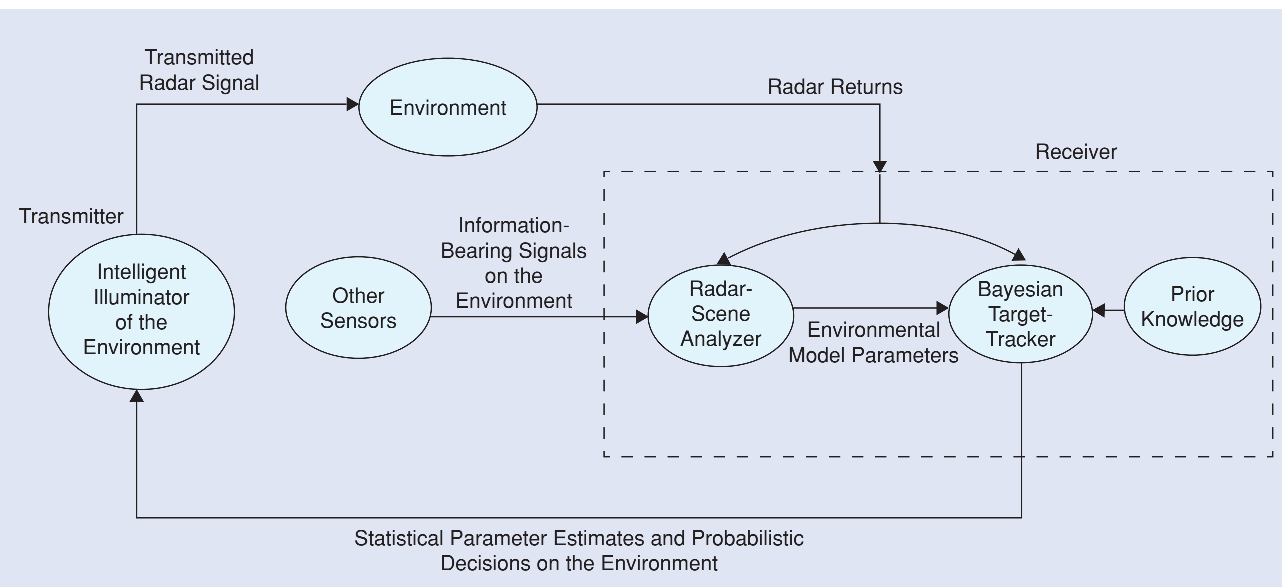

[FIG1] Block diagram of cognitive radar viewed as a dynamic closed-loop feedback system.

relation to the environment and receiver so as to maintain stable and reliable operation at all times.

One other important comment is in order. In reality, cognition is a two-way process, one being inside out and the other being outside in. These two parts of the cognitive process are so referred to, depending on whether the source of information leading to cognition resides inside or outside the receiver, respectively, as explained here.

- The inside-out part of cognition is represented by prior knowledge on the environment, and which is an integral part of the receiver, as shown in Figure 1. The form of prior knowledge is naturally application dependent. For example, it may take the form of a geographic map, elevation model, or kinematics of noncooperative targets. The Bayesian target-tracker retrieves information from the prior-knowledge base and utilizes it for improved radar performance on a need-be basis. Prior knowledge may be viewed as long-term memory of the receiver.
- In contrast, the outside-in part of cognition may be viewed as short-term memory, which is developed by the receiver on the fly. It is initiated by the radar-scene analyzer in response to information-bearing signals gathered on the outside environment by the radar itself as well as other sensors working cooperatively with the radar.

It is noteworthy that the knowledge-based (KB) radar system described in [1] may be viewed as an inside-out cognitive system, embodying heuristics for determining how and when the signal-processing chain should be changed. The heuristics are developed through prior experimentation using a KB approach to target detection with human intervention; the human intervention is subsequently captured and then embedded into the receiver.

[FOR THE RADAR TO BE COGNITIVE, ADAPTIVITY HAS TO BE EXTENDED TO THE TRANSMITTER TOO.]

RADAR-SCENE ANALYSIS

The function of the radar-scene analyzer is to provide the receiver with information on the environment, which is of critical importance to the decisions made by the receiver on possible targets of interest. This function builds on two sources of information-bearing signals.

- radar returns, which are produced by the environment in response to the radar's own transmitted signal
- other relevant information on the environment (e.g., temperature, humidity, pressure, sea-state), which is gathered on the fly by sensors other than the radar itself.

These two sources of inputs constitute the stimuli for the outside-in part of radar cognition.

In a surveillance scenario, radar performance is affected significantly by the unavoidable presence of interference. Typically, the interference is dominated by clutter (i.e., radar returns produced by undesired targets). Accordingly, to design a target tracker which embodies target detection, we need two kinds of information, one pertaining to the clutter acting alone and the other pertaining to the target plus clutter.

STATISTICAL MODELING OF STATISTICAL REPRESENTATION OF CLUTTER- AND TARGET-RELATED INFORMATION

To describe how these two pieces of information can be addressed in specific terms, consider the case of a coherent radar dwelling on a particular patch of the ocean surface. With the radar being coherent, the radar returns contain amplitude as well as Doppler information on that patch. Correspondingly, the baseband version of the radar returns will be complex valued. Now, the dwelling process can be of a long-term nature, in which case the nonstationary character of the radar returns becomes quite noticeable. In situations of this kind, we may be forced to avoid modeling the actual Doppler spectrum (i.e., plot of average power versus frequency) of the radar returns. We do so by exploiting the following intuitively satisfying observations that the Doppler spectrum of clutter by itself is relatively smooth, whereas the spectral content of the radar echo from a target appears essentially as a line component.

However, when the target cross section is small and the target-to-clutter power ratio is therefore low, we need to enhance the line component due to the target. This enhancement may be achieved by performing the transformation of dividing the average power in each Doppler bin of the spectrum (pertaining to the range-azimuth resolution cell of interest) by the mean of its neighboring bins, say k in number [2], [3].

This transformation has the desired effect of accentuating the narrow peak of the line component due to the target and, at the same time, lowering the relatively wide peak of the clutter. Inspiration for the transformation, called a *peak filter*, is traced to the grouped periodogram test described in Priestly [4], which was itself inspired by prior work done by Tukey in 1949. The statistics of the peak filter output, in the absence of a target, may now be evaluated under three assumptions [2], [3].

- None of the k neighboring Doppler bins in the power spectrum contains a target.
- Inside a spectral window encompassing $(k+1)$ Doppler bins, the continuous clutter power spectrum is approximately constant.
- All $(k+1)$ ordinates of the power spectrum are sampled independently.

Under these three assumptions, the individual ordinates of the actual power spectrum have a $\mathcal{X}^2$ distribution with two degrees-of-freedom (DOF) [4]. Correspondingly, the peak-filter output, which divides each spectrum ordinate by k others, has a hypergeometric distribution, specifically an F-distribution with (2, $2k$) DOF [2], [3]. On this basis, the clutter statistics are described by the distribution $F_{2,2k}(\mathrm{z})$, where z is a random variable (i.e., average clutter power measurement). In

[5], a similar conclusion is made using stochastic differential equation theory.

Turning next to the target that is typically unknown, modeling its statistics is unfortunately not straightforward. For ease of implementation and due to lack of detailed knowledge about the target, it may be prudent to assume that the target has the same distribution that governs the clutter but with a difference. (This assumption may hold in the case of a small target moving on an ocean surface, in which case, the underlying dynamics of the clutter and the target are closely coupled.) Accordingly, if the clutter distribution is described by $F_{2,2k}(z)$, the target distribution is taken to be $(1/\gamma)F_{2,2k}(z/\gamma)$, where z is a power spectrum measurement and γ is the target-to-clutter power ratio [2], [3].

In addition to the target statistics, the receiver needs to have a model that accounts for the motion of the target. To this end, we may assume that the target has a Gaussian-distributed acceleration with variance σ^2, which characterizes the agility of the target. For a low standard deviation σ, the target is seen by the radar when it is not accelerating; on the other hand, for a high σ, the task of target detection may become difficult due to possible confusion of the target with small clutter peaks, hence the likelihood of the radar making a decision error.

In summary, for an ocean environment under surveillance by a coherent radar, information on radar returns processed by the radar scene analyzer for a particular range-azimuth cell may be modeled as follows:

- Clutter statistics, described by the F-distribution $F_{2,2k}(z)$, where z is a power spectrum measurement and k is the number of neighboring Doppler bins over which the measurement is averaged.
- Target-plus-clutter statistics, described by the scaled F-distribution $(1/\gamma)F_{2,2k}(z/\gamma)$, where γ is the target-to-clutter power ratio.
- Target motion, described by a Gaussian-distributed acceleration with a variance σ^2, which accounts for the target's agility.

It must be reemphasized, however, that this model is appropriate for the specific case of a target moving on an ocean surface. For other environmental scenarios, the radar designer is challenged to develop appropriate statistical models to describe the information content of radar returns on clutter and targets.

[THE FUNCTION OF THE RADAR-SCENE ANALYZER IS TO PROVIDE THE RECEIVER WITH INFORMATION ON THE ENVIRONMENT.]

BAYESIAN TARGET TRACKING

In [2] and [3], a Bayesian strategy is described for the coherent radar detection of small targets in the presence of sea clutter. Unlike conventional tracking algorithms that perform intermediate detections (i.e., hard decisions) on the radar returns, the new algorithm processes the radar returns directly. In [6], Bruno and Moura also describe a Bayesian approach to the tracking problem. Given a search space of R range-azimuth resolution cells and M possible targets, their algorithm is designed to track any of the targets. The algorithm does so by first computing the probability of each of the 2^M different target combinations. Specifically, the centroid of each target can be in any of the R resolution cells or else be absent. The Bayesian tracking approach described in [2] and [3], however, is different in that it is formulated in such a way that the algorithm can also operate in a smoothing mode, with the probability distribution of the smoothed output being conditional on both past and future observations.

Specifically, the algorithm, referred to as a direct tracking algorithm in [2] and [3], consists of three basic steps:

1) For a given search area, radar returns are collected over a certain period of time.
2) For each range-azimuth resolution cell in the search space, the probability that the cell contains a target is computed.
3) With the evolution of target probability distribution resulting from the recursive computation of step 2 over time, target tracks are detected, and corresponding hard decisions on possible targets are subsequently made.

In effect, the algorithm (formulated in probabilistic terms) may be viewed as a soft-decision detection procedure.

To set the stage for the Bayesian framework, let there be a total of R range-azimuth resolution cells in the search space S, and let $r \in S$ denote a resolution cell in question. Let ε_t^r denote the event of a single target occurring in resolution cell r at discrete time t. Let the vector $\mathbf{z}_t$ denote the frame that is made up of the spectral measurements for all R resolution cells at time t. The matrix

$$\mathbf{Z}_t = [\mathbf{z}_t, \mathbf{z}_{t-1}, \ldots, \mathbf{z}_2, \mathbf{z}_1] = [\mathbf{z}_t, \mathbf{Z}_{t-1}]$$

denotes the full set of all the available frames extending up to and including time t. Then, according to this notation, the vector $\mathbf{z}_t$ denotes the current frame and the remaining matrix $\mathbf{Z}_{t-1}$ denotes the combined set of all past frames. By the same token $\mathbf{Z}_{t+1}$ denotes the combination of a future frame $\mathbf{z}_{t+1}$, the current frame $\mathbf{z}_t$, and all past frames $\mathbf{Z}_{t-1}$.

Following the traditional approach to state estimation, we may now identify three different forms of the Bayesian target-tracker:

- one-step predictor, whose output is described by the conditional probability $\mathrm{P}(\varepsilon_t^r|\mathbf{Z}_{t-1})$
- filter, whose output is described by the conditional probability $\mathrm{P}(\varepsilon_t^r|\mathbf{Z}_t)$
- smoother, whose output is described by the expanded conditional probability $\mathrm{P}(\varepsilon_t^r|\mathbf{Z}_{t+1})$.

Smoothing uses more information than both prediction and filtering and may therefore be more accurate than both of them in a statistical sense. On the other hand, however, only prediction and filtering can be implemented in real time.

ONE-STEP TRACKING PREDICTION

Consider the joint event $(\varepsilon_t^r, \varepsilon_{t-1}^q)$, which describes a target occurring in resolution cell q at time $t-1$ and then moving into resolution cell r at time t. From probability theory, we may express the output of the tracking predictor at time t as

$$\begin{aligned} \mathrm{P}(\varepsilon_t^r|\mathbf{Z}_{t-1}) &= \sum_{q=1}^{R} \mathrm{P}(\varepsilon_t^r, \varepsilon_{t-1}^q|\mathbf{Z}_{t-1}) \\ &= \sum_{q=1}^{R} \mathrm{P}(\varepsilon_t^r|\varepsilon_{t-1}^q, \mathbf{Z}_{t-1})\, \mathrm{P}(\varepsilon_{t-1}^q|\mathbf{Z}_{t-1}). \end{aligned} \quad (1)$$

However, given the fact that the event ε_{t-1}^q has occurred at time $t-1$, it makes the previous measurements matrix $\mathbf{Z}_{t-1}$ irrelevant. In other words, occurrence of the event ε_{t-1}^q conveys exactly the same amount of information as the joint event $(\varepsilon_{t-1}^q, \mathbf{Z}_{t-1})$. Accordingly, (1) reduces to the simpler form

$$\mathrm{P}(\varepsilon_t^r|\mathbf{Z}_{t-1}) = \sum_{q=1}^{R} \mathrm{P}(\varepsilon_t^r|\varepsilon_{t-1}^q)\, \mathrm{P}(\varepsilon_{t-1}^q|\mathbf{Z}_{t-1}). \quad (2)$$

The conditional probability $\mathrm{P}(\varepsilon_{t-1}^q|\mathbf{Z}_{t-1})$ is the output of the tracking filter working on resolution cell q at time $t-1$. We also recognize $\{\mathrm{P}(\varepsilon_{t-1}^r|\varepsilon_{t-1}^q)\}_{q,r}$ as the set of probabilities that event ε_t^r follows event ε_{t-1}^q. This set of probabilities is referred to as the *transition matrix* of the tracker, the formulation of which exploits the statistical model of target motion as supplied by the radar-scene analyzer. It is noteworthy that the less agile the target is, the smaller the jumps the target is expected to make in the search space S, thereby causing the transition matrix to be more sparse. In any event, given the tracking filter output at time $t-1$ and the transition matrix, we may use (2) to compute the output of the corresponding tracking predictor at time t.

TRACKING FILTER

Consider next the issue of computing the output of the tracking filter at time t, which is defined by the posterior probability $\mathrm{P}(\varepsilon_t^r|\mathbf{Z}_t)$. Applying Bayes' rule to this probability yields

$$\mathrm{P}(\varepsilon_t^r|\mathbf{Z}_t) = \frac{p(\mathbf{Z}_t|\varepsilon_t^r)\,\mathrm{P}(\varepsilon_t^r)}{p(\mathbf{Z}_t)}, \quad (3)$$

where $p(\mathbf{Z}_t|\varepsilon_t^r)$ is the conditional probability density function of the current measurements matrix $\mathbf{Z}_t$ given the occurrence of event ε_t^r, and $\mathrm{P}(\varepsilon_t^r)$ is the prior probability of that event. The probability density function $p(\mathbf{Z}_t)$ in the denominator is the evidence, which acts merely as a normalizing function. Since, by definition, $\mathbf{Z}_t = (\mathbf{z}_t, \mathbf{Z}_{t-1})$, we may rewrite (3) by expanding the numerator:

$$\begin{aligned} \mathrm{P}(\varepsilon_t^r|\mathbf{Z}_t) &= \frac{p(\mathbf{z}_t, \mathbf{Z}_{t-1}|\varepsilon_t^r)\,\mathrm{P}(\varepsilon_t^r)}{p(\mathbf{Z}_t)} \\ &= \frac{p(\mathbf{z}_t|\varepsilon_t^r, \mathbf{Z}_{t-1})\,p(\mathbf{Z}_{t-1}|\varepsilon_t^r)\,\mathrm{P}(\varepsilon_t^r)}{p(\mathbf{Z}_t)}. \end{aligned} \quad (4)$$

Recognizing that the occurrence of event ε_t^r makes past measurements $\mathbf{Z}_{t-1}$ irrelevant, we may simplify (4) as

$$\mathrm{P}(\varepsilon_t^r|\mathbf{Z}_t) = \frac{p(\mathbf{z}_t|\varepsilon_t^r)\,p(\mathbf{Z}_{t-1}|\varepsilon_t^r)\,\mathrm{P}(\varepsilon_t^r)}{p(\mathbf{Z}_t)}. \quad (5)$$

The first term $p(\mathbf{z}_t|\varepsilon_t^r)$ in the numerator of (5) is the probability density function of measurement $\mathbf{z}_t$, given that there is a target in cell r at time t. The second term $p(\mathbf{Z}_{t-1}|\varepsilon_t^r)$ is computed by using the recursive formula

$$\begin{aligned} p(\mathbf{Z}_{t-1}|\varepsilon_t^r) &= p(\mathbf{z}_{t-1}, \mathbf{Z}_{t-2}|\varepsilon_t^r) \\ &= \sum_{q=1}^{R} p(\mathbf{z}_{t-1}|\varepsilon_{t-1}^q)\,p(\mathbf{Z}_{t-2}|\varepsilon_{t-2}^q)\,\mathrm{P}(\varepsilon_{t-1}^q|\varepsilon_t^r) \end{aligned} \quad (6)$$

where, as before, $p(\mathbf{z}_{t-1}|\varepsilon_{t-1}^q)$ is input from the radar-scene analyzer and $p(\mathbf{Z}_{t-2}|\varepsilon_{t-1}^q)$ is the one-step delayed version of $p(\mathbf{Z}_{t-1}|\varepsilon_t^r)$; hence the reference to (6) as a recursive formula. The matrix of probabilities $\{\mathrm{P}(\varepsilon_{t-1}^q|\varepsilon_t^r)\}_{q,r}$ is the inverse transition matrix, which is defined by the probabilities that event ε_q^{t-1} preceded event ε_t^r; the term inverse is used here merely to imply the role reversal of these two events with respect to the transition matrix under (2). The following two points are noteworthy:

- The recursive formula of (6) is identical to the hidden Markov model (HMM) filter for a Markov chain $\{\varepsilon_t\}$ with transition probabilities $\{\mathrm{P}(\varepsilon_{t-1}^q|\varepsilon_t^r)\}$.
- Given the posterior probability distribution of (4), the conditional mean estimate (i.e., minimum mean-square estimate) of the event ε_t^r over the entire search space S can be computed as the summation $\sum_{r=0}^{R} \varepsilon_t^r \mathrm{P}(\varepsilon_t^r|\mathbf{Z}_t)$.

We may also compute the conditional probability density function $p(\mathbf{Z}_{t-1}|\varepsilon_t^r)$ in another way by recasting the recursive formula of (6) as follows:

$$\begin{aligned} p(\mathbf{Z}_{t-1}|\varepsilon_t^r) &= \frac{p(\mathbf{Z}_{t-2})}{\mathrm{P}(\varepsilon_t^r)} \sum_{q=1}^{R} p(\mathbf{z}_{t-1}|\varepsilon_{t-1}^q) \\ &\quad \times \mathrm{P}(\varepsilon_{t-1}^q|\mathbf{Z}_{t-2})\,\mathrm{P}(\varepsilon_t^r|\varepsilon_{t-1}^q). \end{aligned} \quad (7)$$

Then, substituting (7) into (5), we get the new formula for computing the posterior probability at the output of the tracking filter

$$\begin{aligned} \mathrm{P}(\varepsilon_t^r|\mathbf{Z}_t) &= \frac{p(\mathbf{z}_t|\varepsilon_t^r)}{p(\mathbf{z}_1, \mathbf{z}_2|\mathbf{Z}_{t-2})} \sum_{q=1}^{R} p(\mathbf{z}_{t-1}|\varepsilon_{t-1}^q) \\ &\quad \times \mathrm{P}(\varepsilon_{t-1}^q|\mathbf{Z}_{t-2})\,\mathrm{P}(\varepsilon_t^r|\varepsilon_{t-1}^q), \end{aligned} \quad (8)$$

where the probability $\mathrm{P}(\varepsilon_{t-1}^q|\mathbf{Z}_{t-2})$ is the delayed version of the tracking predictor output, and the probabilities $\{\mathrm{P}(\varepsilon_t^r|\varepsilon_{t-1}^q)\}_{q,r}$ are elements of the transition matrix.

On the basis of (2) and (8), we may now construct the block diagram of Figure 2 for the Bayesian direct filtering system. The diagram is in the form of a closed-loop feedback system that operates by propagating a state vector of probabilities from one iteration to the next. Most important, the right relationship must be established between the radar parameters and statistical characteristics of clutter and target-plus-clutter for the tracker to maintain a stable operation.

TRACKING SMOOTHER

An attractive feature of the Bayesian tracker as described herein is the fact that it is straightforward to make its operation conditional on both past and future spectral measurements. The result of this expansion is a target tracking smoother, for which the output is expressed as

$$P\left(\varepsilon_t^r|\mathrm{z}_t, \mathrm{Z}_{t-1}, \mathrm{z}_{t+1}\right) = \frac{p\left(\mathrm{z}_t|\varepsilon_t^r\right) p\left(\mathrm{Z}_{t-1}|\varepsilon_t^r\right) p\left(\mathrm{z}_{t+1}|\varepsilon_t^r\right) P\left(\varepsilon_t^r\right)}{p\left(\mathrm{z}_t, \mathrm{Z}_{t-1}, \mathrm{z}_{t+1}\right)}. \quad (9)$$

The factorization of terms in the numerator of (9) assumes that the radar is treated as a first-order Markov model, in which case the conditional dependence of the distribution of past measurements Z_{t-1} on the future measurements z_{t+1} may be ignored; that is, we may set $p(\mathrm{Z}_{t-1}|\varepsilon_t^r, \mathrm{z}_{t+1})$ equal to $p(\mathrm{Z}_{t-1}|\varepsilon_t^r)$.

The additional factor $p\left(\mathrm{z}_{t+1}|\varepsilon_t^r\right)$ in the numerator of (9) is computed by running the right-hand side of the recursive equation (6) backwards in time [2], [3]. Thus, whereas the target-tracking filter operates in the forward direction only, the target-tracking smoother operates in the forward as well as backward direction. Accordingly, decisions made on possible targets using the tracking smoother contain more information than the corresponding tracking filter and may therefore be more reliable. However, this improvement in performance is gained at the expense of two factors: increased computational complexity and nonreal time operation.

EXPERIMENTAL RESULTS: CASE STUDY OF SMALL TARGET IN SEA CLUTTER

In [2] and [3], the performance of the Bayesian target detector is evaluated using real-life radar data under varying conditions. The data were collected by means of the McMaster IPIX radar, which is a highly configurable coherent multifunction X-band radar built specifically for research purposes. For a subset of the database collected at a site in Dartmouth (Nova Scotia), the radar was operated in the dwell mode with a l° pencil beam and fixed radio frequency of 9.39 GHz. The radar was mounted about 30 m above sea level, with the target of interest being located at about 2.5 km offshore. The target was a sphere (1 m in diameter) made up of wire covered in foam. Radar range was sampled at 15 m intervals, obtained by using a 200 ns rectangular pulse. (The actual range resolution of the radar was 30 m.) The pulse-repetition frequency (PRF) was 2,000 Hz, but the pulse alternated between horizontal (H) and vertical (V) polarization, so the effective single-polarization PRF was 1,000 Hz. For each pulse, both H and V polarizations were recorded simultaneously, resulting in a matrix of four possible transmit/receive polarizations: HH, HV, VH, and VV. For each combination in the matrix, the amplitude and phase of the radar returns were stored in the form of in-phase (I) and quadrature (Q) components. (To promote further research in the radar area, we have created a comprehensive Web site has been created from which IPIX radar data sets are available. For details of the site, see http://soma.mcmaster.ca/ipix.)

In [3], three data sets from the Dartmouth database were used to test the Bayesian target detector. The results pertaining to one of those datasets is reproduced in Figure 3. The upper part of the figure displays the Doppler-time image of the raw radar dataset, using a 64-sample sliding window. The lower part of the figure displays the resulting output of the Bayesian direct tracking smoother. Each pixel in the image represents the probability of a target being present in the corresponding resolution cell: the darker the pixel, the higher the probability of target occurrence. Note also that the dark traces included along the 500 Hz-line indicate the points in time where the target was invisible to the radar or when the radar failed to detect the target.

Figure 3 and several other results reported in [3] attest to the effectiveness of the Bayesian direct tracker. In particular, even for a dataset with an average target-to-clutter power ratio as low as −7 dB, Figure 3 clearly demonstrates the visibility of the target most of the time.

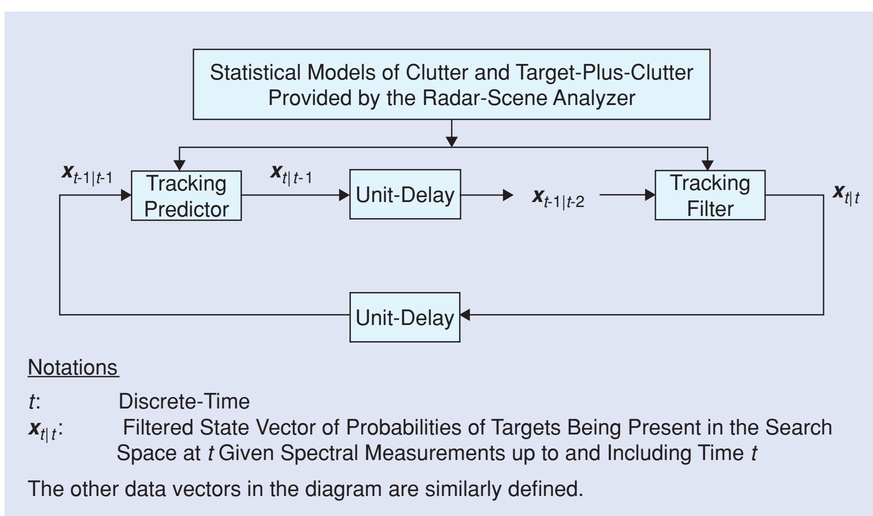

[FIG2] Block diagram of the Bayesian direct filtering system.

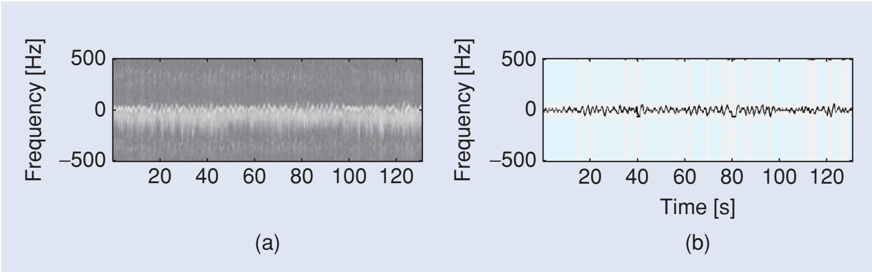

[FIG3] (a) A 64 sample sliding window time-Coppler image of raw radar data set 3 in [3]. (b) The output of the Bayesian tracking smoother; each pixel represents the pobability of a target in the corresponding resolution cell.

Using two different real-life radar data sets and computer-simulated data, a comparative evaluation of the Bayesian approach to target-detection-through-tracking has been made against a new detection strategy called the *correlation anomaly receiver* that follows from the theory of stochastic differential equations [5]. The results of this evaluation, reported in [9], show that the Bayesian receiver's performance is superior to that of the correlation anomaly receiver.

PRACTICAL IMPLICATIONS OF THE BAYESIAN TARGET TRACKER

To the best of my knowledge, the Bayesian target tracker described in detail in [2] and [3] and highlighted herein is the first study on the feasibility of direct target-tracking without intermediate detections. The use of a Bayesian approach to direct tracking, combined with complete reliance on soft decisions (i.e., avoiding hard decisions through intermediate detections), has some important practical implications.

- Unlike hard decisions, the soft decisions made by the Bayesian target-tracker preserve the information content of the radar returns; this approach follows the principle of information preservation learned from Shannon's information theory [7].
- The Doppler-time image produced by the Bayesian direct target tracker makes it possible for the radar to see the motion of the target in a manner comparable to the human eye. Indeed, we conjecture that an experienced human operator could not do a better job of following the target than the Bayesian tracker, especially so when the tracker is operated in the smoothing mode; in this mode, the tracker exploits the combined benefit of forward and backward computations.
- The basic idea behind the Bayesian approach is to view the information contained in the radar image as a probability distribution that characterizes the likelihood of a particular resolution cell containing a target. The distribution, in the form of a posterior probability density function, is determined in part by the statistical structure of the radar scene (i.e., the outside world) and in part by the way in which echoes from the world are actually encoded by the radar itself. Accordingly, the Bayesian approach distinguishes itself from other approaches by invoking an explicit statistical structure of the world which, in reality, is a fundamental necessity.
- The edited book [8] presents a number of theoretical frameworks for studying visual perception which, in varying degrees, are all founded on Bayesian principles. In a way, this book lends further support to the Bayesian radar-target tracker, the theory of which is embodied in (1)–(9), depending on the mode of operation.

ADAPTIVE RADAR ILLUMINATION

As it stands, there is no optimization being performed on the posterior probability distribution $\mathrm{P}(\varepsilon_t^r|\mathbf{Z}_t)$ computed by the Bayesian target tracker [10]. The practical issue with adaptive radar illumination (transmission) is how to observe past radar returns and extract useful information in order to decide or select the radar waveform for the next transmission in some optimal fashion.

In an implicit sense, the present spectral measurements at time t, denoted by $\mathbf{z}_t$ and the past measurements denoted by $\mathbf{Z}_{t-1}$, are all dependent on the transmitted signal. This dependence suggests that the whole radar system can be made adaptive by adjusting certain parameters in the transmitted signal in response to the probabilistic decisions made by the Bayesian tracker on the environment under surveillance. Note however that by doing so, the radar system assumes the form of a stochastic control system involving a state-space model governed by the posterior distribution of (4); the optimal solution to such partially observable stochastic-control problems is NP hard. Fortunately, there are suboptimal procedures such as reinforcement learning that can yield acceptable solutions; this issue is discussed later.

There are many ways in which parameters of the transmitted signal can be adjusted [11]. One practical way is to use burst waveforms, with each burst made up of a sequence of uniformly spaced, nonoverlapping subpulses of fixed duration. The pulse amplitudes are held constant for two reasons: unforeseen difficulties with dynamic range requirements are avoided and the target-to-clutter power ratio may not be sensitive enough to pulse-amplitude adjustments.

The logical strategy is then to adjust the phase of each transmitted RF pulse in accordance with feedback sent to the transmitter from the receiver. Here we have the choice of a phase response that varies with time according to a square law that results in linear frequency modulation (FM) or a cubic law that results in nonlinear FM. Both of these configurations are well known for their pulse-compression characteristics, with the nonlinear FM being more effective than the linear FM [12].

SIMULATION EXPERIMENTS IN SUPPORT OF ADJUSTABLE FM

In [13] and [14], DeLong and Hofstetter describe an adaptive transmission strategy that uses the above-mentioned scheme of burst waveforms. Their detailed theoretical study was in two parts: the first one published in 1967 investigated the use of adjustable pulse amplitudes and the second one published later in 1969 focused on the use of adjustable pulse phases with limited dynamic range. For both studies, they used the signal-to-interference ratio as the index of performance, with interference being composed of clutter and receiver noise. For performance optimization, they used a procedure based on the Karush-Kuhn-Tucker theorem [15]. The conclusions reported by DeLong and Hofstetter in those two early, but still very much valid, papers may be summarized as follows:

- Adjustment of phase is a more practical approach than amplitude for the design of adaptive radar transmitters.
- The use of constant amplitude, quadratic-phase burst waveforms provides a significantly better ambiguity pattern than the corresponding constant-amplitude, zero-phase burst waveforms.

These findings have been confirmed in simulation experiments reported in [16], using the following system parameters:

Burst waveform:	32 subpulses
Single-pulse signal-to-interference ratio:	5 dB
Clutter-to-target cross-section ratio:	100

In [16]–[18], the Delong-Hofstetter algorithm using constant-amplitude, square-phase burst waveforms was found to reach a peak signal-to-interference ratio of 16.9 dB after 30 iterations of the algorithm; the performance of this system is almost equivalent to perfect clutter suppression in an environment highly dominated by clutter.

THE SELECTION OF WAVEFORMS TO BE USED FOR ADAPTIVE RADAR TRANSMISSION IS APPLICATION DEPENDENT.

ECHO-LOCATION IN BATS

From the introductory section, we recall that a cognitive radar system embodies three fundamental ingredients:

- learning from the environment through experience
- adjustment of the transmitted signal in an intelligent manner
- feedback from the receiver to the transmitter to make this adjustment possible.

All these three features are part and parcel of the echo-location system of a bat. Accordingly, there is much that we can learn from the echo-location system of a bat [19]–[22]. Most echo-locating bats are blind. (We say "most" because not all species of bats are blind; note also that not all bats use echo-location.) To see the world around it, the bat uses sonar, which is an active echo-location system. In addition to providing information about how far away a target (i.e., flying insect) is, the bat's sonar conveys information about the relative velocity of the target, the size of various features of the target, and azimuth and elevation of the target [21]–[23]. The complex neural computations needed to extract all this information from the target echo occur within a brain the size of a plum. Indeed, an echo-locating bat can pursue and capture its target with a facility and success rate that would be the envy of a radar engineer. How then does the bat perform all these remarkable tasks? The answer to this fundamental question lies in the fact that soon after birth, the bat uses its innate hard-wired brain to build up rules of behavior through what we usually refer to as experience, hence the remarkable ability of the bat for echo-location.

The bat uses its mouth (or nose) to broadcast echo-location sounds and its auditory system as the sonar receiver. The emitted sounds consist of burst waveforms whose characteristics are highly diverse, varying with both species and being situation specific. The transmitted sound characteristics are summarized here:

Duration:	0.3 to 300 ms
Frequency:	12 to 200 kHz
Structure:	FM component,
	or constant-frequency (CF) component followed by FM component.

The CF component can be single or multiple harmonic. The FM component can be of a downward or upward kind, with the FM sweep varying linearly or nonlinearly with time. The use of FM is intended to improve the echo-location system's resolution capability of the bat. (It is noteworthy that an echo-location bat's emitted sounds consist of burst waveforms just as the adaptive transmission strategy used in the DeLong-Hofstetter algorithm consists of burst waveforms.)

Broadly speaking, the adaptive behavior of bats may be categorized as follows [21]:

- *Velocity-dependent adaptation*, which involves adjustment of the transmitted sound frequency; this form of adaptation is most salient in species of CF-FM bats. These CF-FM bats also appear to make adjustments in temporal patterning as they close in on their targets.
- *Range-dependent adaptation,* which involves adjustment of the emitted-sound duration, bandwidth, and repetition rate; this second form of adaptation is most salient in bats using only FM. These bats also appear to make adjustments in the transmitted sound bursts during target approach.

Echoes from targets (i.e., insects) are represented in the auditory system by neuronal activities that are sensitive to different combinations of acoustic inputs produced in response to the transmitted sound bursts. In particular, three principal dimensions of the bat's auditory representation have been identified [19]:

- echo frequency, which is initially encoded in the auditory periphery cochlea by place in the cochlear
- echo amplitude, which is encoded by the neuronal responses under the previous and other neurons tuned to different dynamic ranges in the central nervous system

- echo delay, which is encoded through neuronal computations that produce target-range tuning responses.

There are two principal (neuronal) computations that are performed by the bat's brain for image-forming purposes. One is the spectrum of the incoming echo, which is intended for the extraction of target shape. The other is delay in the received echo with respect to the transmitted sound bursts, which is intended for the extraction of target range. To carry out these computations, frequency-based information contained in the incoming echo spectrum is converted into estimates of the spatial (time) structure of the target.

In short, the echo-location system of a bat is very plastic, in that the parameters of the transmitted sound bursts can be changed considerably during the different phases of the target-pursuit sequence. We are therefore justified to view the echo-location system of a bat as physical proof (albeit in neurobiological terms) of cognitive radar.

THERE IS MUCH THAT WE CAN LEARN FROM THE ECHO-LOCATION SYSTEM OF A BAT.

DISCUSSION

Three important conclusions can be drawn from the presentations made in this article.

1) Intelligence is a necessary requirement for the radar to be cognitive. A striking difference is discernible between the presentations we have made on adaptive radar illumination and echo-location in bats. Simply put, in signal processing terms, the echo-location systems of bats are far more plastic than the adaptive radar systems that are currently in use or being contemplated. This important point is best illustrated by the spectograms shown in Figure 4, which were produced by four different bat species in their respective target (insect)-pursuit sequences. The significant characteristic that is immediately apparent from this figure is that the transmitted signal duration decreases and the burst repetition rate increases as the bat gets closer to its target. In doing this, the bat is using acquired knowledge of the distance from its target to adjust the parameters of its transmitted sound bursts. For a radar system to be cognitive, therefore, it is a fundamental necessity for the radar transmitter to learn from continuing interactions with the environment and intelligently use the information extracted by the receiver on targets under surveillance, all of this being done on the fly during the different phases of the target-track sequence.

2) Feedback from the receiver to the transmitter is the facilitator of intelligent signal processing. We say feedback is the facilitator of intelligence, because it is through feedback from the receiver to the transmitter that cognitive radar is enabled to learn from interactions with the environment. More is said on this issue later in "Learning."

3) The preservation of information in radar returns is of crucial importance to receiver performance. The results presented earlier on the Bayesian target-tracker emphasize the signal processing power of the Bayesian approach. This approach is the only statistical approach, in which a model of the received signal accounts for two factors contributing to the specification of information:

- statistical nature of interference (i.e., radar clutter and noise)
- explicit statistical structure of the radar environment (i.e., outside world), including targets.

In the past, the Bayesian approach has been criticized for requiring a model that includes a statistical structure of the radar environment. In response to such criticism, we merely have to emphasize that if we are to account for the physical realities that are responsible for the generation of radar returns, then the inclusion of a statistical structure of the radar environment is a fundamental requirement for preserving the information content of the received signal.

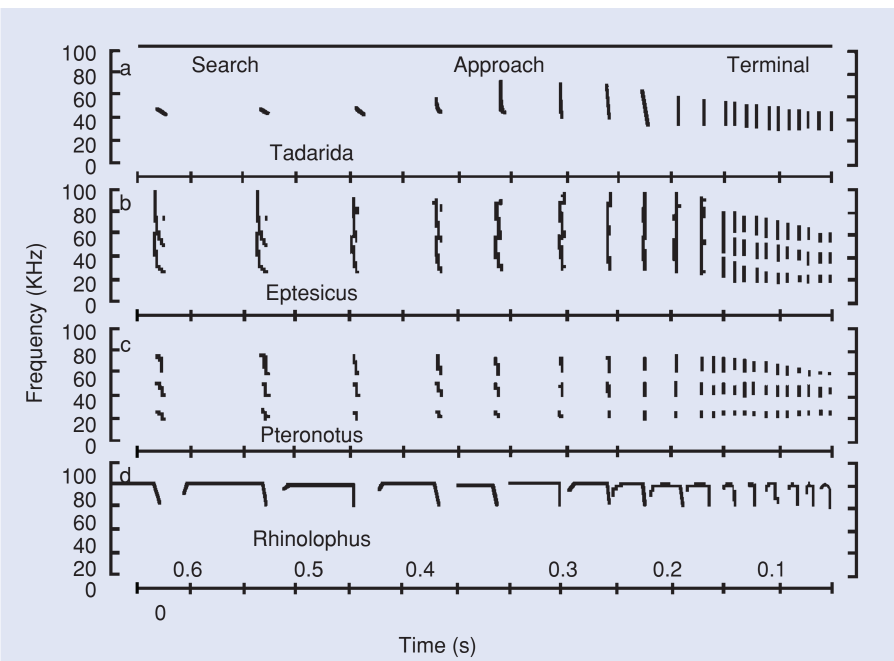

[FIG4] Spectrograms of sonar signals produced by four different species of bats as they advance from the search to approach and finally to the terminal phase of insect pursuit. (Reproduced from [20] with permission of the University of Chicago Press.)

LEARNING

Throughout this article, we have emphasized that learning is a basic ingredient of cognitive radar. In a generic sense, the learning process can take two different forms: offline, and online.

Through offline learning, knowledge is acquired about the environment and then embedded in the receiver. In the radar context, an established way of accomplishing this acquisition is to collect real-life data by conducting ground-truthed experiments on the environment under varying conditions. Then, by performing statistical analysis on the radar data and formulating models on clutter and targets, the acquisition of knowledge of the environment is accomplished; see, for example [23] and [24]. In any event, the offline learning takes place through the intervention of the experimenter.

[AN ECHO-LOCATING BAT CAN PURSUE AND CAPTURE ITS TARGET WITH A FACILITY AND SUCCESS RATE THAT WOULD BE THE ENVY OF A RADAR ENGINEER.]

Among the many different online learning procedures, reinforcement learning [25] stands out as the procedure best suited for cognitive radar. In the modern approach to reinforcement learning, also referred to as neurodynamic programming [26], Bellman's dynamic programming (rooted in control theory) provides the theoretical foundation of the procedure. However, Bellman's dynamic programming suffers from the curse of dimensionality, which limits its practical utility. Neurodynamic programming overcomes this limitation by using a neural network for the approximation of dynamic programming in a physically realizable manner. Stated in simple terms, neurodynamic programming enables a learning system to do two things [26]:

- make good decisions by observing the system's own behavior, which is achieved by using Monte Carlo simulations in an offline manner
- improve the system's actions through a built-in reinforcement mechanism, which is achieved by using iterative optimization in an online manner.

The net result is a learning procedure that permits cognitive radar to learn through interactions with the environment on a continuing basis in a way which, loosely speaking, mimics the way in which the echo-location bat learns from its environment.

APPLICATIONS

A discussion of cognitive radar would be incomplete without some applications where it has the potential to make a difference. In what follows, we address two applications of cognitive radar, one dealing with multifunction radars that are expensive and the other dealing with noncoherent radars that are inexpensive.

MULTIFUNCTION RADARS

Thanks to continuing advances and improvements on two fronts, namely, phased-array antennas and computers, multifunction radars are fast becoming, if not already, the norm in building sophisticated radar systems. For example, the radar may have to deal with a fading target due to the presence of multipath produced by close proximity of the target (e.g., sea-skimming missile) to the sea surface in a hostile marine environment. One way of mitigating the fading problem is to increase the dwell time in order to track the target with adequate accuracy. In such an environment, we may identify two problems that require serious attention:

- *agility,* which mandates the use of phased-array antennas oriented to provide 360° coverage (e.g., four arrays at 90° with respect to each other)
- *fast response,* which is attained by using powerful computers that enable the radar to adapt its transmission waveforms so as to detect, track, and paint the target rapidly enough for the engagement to occupy no more than 30 s to couple of minutes.

Typically, while attending to the fading target, the radar is also required to handle other threatening targets. The radar is therefore faced with a new problem, namely, resource management [27], [28]. Neurodynamic programming provides a principled approach for a solution (albeit suboptimal but perhaps adequate) to the resource management problem.

NONCOHERENT RADAR NETWORK

For an entirely different application that could benefit from the use of cognitive radar, consider the international border-security problem. To be more specific, consider the Great Lakes St. Lawrence Seaway; there are two challenging problems with this large open border between the United States of America and Canada [29]:

- the protection of assets and populations of people from terrorism
- the prevention of illegal crossings across the border.

A cost-effective, all-weather, and all-day solution to both of these challenges is a cognitive noncoherent radar network. The network would be made up of inexpensive commercial off-the-shelf marine radars, which are distributed across the border. The only discriminant available for surveillance with such simple radars is amplitude, which severely limits the capability of the radar to detect noncooperative targets with small radar cross-section in the presence of lake clutter. To mitigate this serious problem, Weber et al. [29] depart from conventional radar signal processing by purposely setting low detection thresholds. Naturally, the false-alarm rates are raised to levels higher than a conventional processor. But, most importantly, the noncooperative small targets are now detectable. Then through the use of a sophisticated tracking algorithm, the real targets are extracted and the false-alarm rates are reduced to an acceptable level.

Given a network of such noncoherent radars, which also incorporates a central base station, the real-target tracks computed by the component radars are transmitted by a communication channel (wireline or wireless) to that station.

Consequently, we have yet another new problem, namely, multi-sensor fusion. Given the limited computing resources at the base station, the challenge here is how to design a cognitive radar network that produces a map in real-time for the entire Great Lakes St. Lawrence Seaway, which identifies the tracks of all noncooperative targets operating therein and does so in the most reliable manner possible.

In both of the applications addressed herein, another extremely challenging issue is that of knowing how to define a metric by means of which it can be said that the task in question has been accomplished. Stated in another way, what is the essence of the description of the environmental scene that is under surveillance? The traditional radar specifications, based on the probability of detection and the problem of false alarm (which are never measured anyway in a real-time setting) are unsuitable. Rather, we need a new metric that addresses specifically what the end user needs to see. The formulation of this metric is further exasperated when the application at hand involves several tasks and the tasks have to be prioritized. Here again, a cognitive approach that learns over time may well provide an answer, as it is often the case with humans [30].

ACKNOWLEDGMENTS

The author is grateful to Prof. C.J. Baker, University College London, U.K.; Dr. Vikram Krishnamurthy, University of British Columbia, Vancouver; Dr. A. Farina, Selex-Sistemi Integrati, Italy; Mr. Brian Currie, formerly of McMaster University; Dr. T. Nohara, Sicom Systems Ltd., Fonthill, Ontario; Dr. T. Field, McMaster University; and Dr. M. Bell, Purdue University, for many useful comments on an early draft of this article. Helpful input from Dr. J. McWhirter, QinteQ, U.K., and Dr. M. McDonald, DRDC Ottawa, are also appreciated. Comments made by Dr. Paul Faure, McMaster University, particularly on the echo-location system of a bat have been most helpful. Helpful inputs from my two graduate students, Kenny Szeto and Ulas Gunturkun, as well as constructive comments from three reviewers of the article are acknowledged. Last but by no means least, permission of the University of Chicago Press to reproduce Figure 4 is deeply appreciated.

AUTHOR

Simon Haykin received his B.Sc. (first-class honors), Ph.D., and D.Sc., all in electrical engineering from the University of Birmingham, England. He is the author of numerous books, including the most widely used books: *Communication Systems* (Wiley, fourth edition), *Adaptive Filter Theory* (Prentice-Hall, fourth edition), *Neural Networks: A Comprehensive Foundation* (Prentice-Hall, second edition), and a new book on adaptive radar: *Toward the Development of Cognitive Radar,* which is about to be published (Wiley), as well as numerous refereed journal papers. He is a Fellow of the Royal Society of Canada, recipient of the honorary degree of Doctor of Technical Sciences from ETH, Zurich, Switzerland, and the Henry Booker Gold Medal from URSI, as well as other prizes and awards.

REFERENCES

[1] G. Capraro, A. Farina, H. Griffiths, and M. Wicks, "Knowledge-based radar signal and data processing: A tutorial overview," *IEEE Signal Processing Mag.*, vol. 23, no. 1, pp. 18–29, 2006.

[2] R. Bakker, T. Kirubarajan, B. Currie, and S. Haykin, "Adaptive radar detection: A bayesian approach," in *Proc. EPSRC IEE Workshop Nonlinear Non-Gaussian Signal Processing*, Peebles, Scotland, 8–12 July 2002.

[3] R. Bakker, G. Lopez, and S. Haykin, "Bayesian approach to the direct filtering of radar targets in clutter," Adaptive Systems Laboratory, McMaster Univ., Canada, Report 02-02, Aug. 2002.

[4] M.B. Priestly, *Spectral Analysis and Time Series*. New York: Academic, 1981.

[5] T.R. Field and R.J.A. Tough, "Stochastic dynamics of the scattering amplitude generating K-distributed noise," *J. Math. Phys.*, vol. 44, pp. 5212–5223, 2003.

[6] M.G.S. Bruno and J.M.F. Moura, "Multiframe detector/tracker: Optimal performance," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 37, pp. 925–944, July 2001.

[7] A. Viterbi, "Wireless digital communication: A view based on three lessons learned," *IEEE Commun. Mag.*, vol. 29, no. 9, pp. 33–36, Sept. 1991.

[8] *Perception as Bayesian Inference*, D.C. Knill and W. Richards, Eds. Cambridge, U.K.: Cambridge Univ. Press, 1996.

[9] B. Currie and S. Haykin, "Bayesian detector evaluation and comparison," Defense Research Development-Ottawa (DRDC), Final Rep. under Contract W7714-020683/001/SV, Apr. 2004.

[10] D.T. Gjessing, *Target Adaptive Matched Illuminate Radar: Principles and Applications* (IEEE Electromagnetic Waves Series 22). Stevenage, U.K.: Peregrinus, 1986.

[11] M.R. Bell and S. Monroq, "Diversity waveform signal processing for delay-Doppler measurement and imaging," *Dig. Signal Process*, vol. 12, pp. 329–346, Apr./July 2002.

[12] P. Weber and S. Haykin, "Non-linear FM pulse compression radar systems," Communications Research Laboratory, McMaster University, Canada, Rep. 120, 1983.

[13] D.F. DeLong and E.M. Hofstetter, "On the design of optimum radar waveforms for clutter rejection," *IEEE Trans. Inform. Theory*, vol. IT-13, pp. 454–463, 1967.

[14] D.F. DeLong and E.M. Hofstetter, "The design of clutter-resistant radar waveforms with limited dynamic range," *IEEE Trans. Inform. Theory*, vol. IT-15, pp. 376–385, 1969.

[15] D.P. Bertsekas, *Nonlinear Programming*. Athenas Scientific, 1995.

[16] S. Haykin, T. Kirubarajan, and B. Currie, "Adaptive radar for improved small target detection in a maritime environment," ASL Report 03-01, Adaptive Systems Laboratory, McMaster Univ., Canada, 2003.

[17] S. Haykin, T. Kirubarajan, and B. Currie, "Simulation Setup for Optimum Energy-Constrained and Amplitude-Constrained Waveform Design," Adaptive Systems Laboratory, McMaster University, Canada. ASL Report 03-02, 2003.

[18] S. Haykin, T. Kirubarajan, and B. Currie, "Literature Search on Adaptive Radar Transmit Waveforms," ASL Report 02-01, 2002, Adaptive Systems Laboratory, McMaster University, Canada.

[19] J.A. Simons, "The resolution of target range by echo-locating bats," *J. Acoust. Soc. Amer.*, vol. 54, pp. 157–173, 1973.

[20] N. Suga, "Cortical computation maps for auditory imaging," *Neural Netw.*, vol. 3, pp. 3–21, 1990.

[21] J.A. Thomas, C.F. Moss, and M. Vater, *Echolocation in Bats and Dolphins*. Chicago, IL: Univ. of Chicago Press, 2004.

[22] G. Schuller and C.F. Moss, "Vocal control and acoustically guided behavior in bats," In J.A. Thomas, C.F. Moss, and M. Vater, *Echolocation in Bats and Dolphins*. Chicago, IL: Univ. of Chicago Press, 2004, pp. 3–16.

[23] S. Haykin, R. Bakker, and B.W. Currie, "Uncovering nonlinear dynamics: The case study of sea clutter," *Proc. IEEE,* vol. 90, no. 5, pp. 860–881, May 2002.

[24] M. Greco and F. Gini, "X-band sea clutter non-stationarity: The influence of long waves," in S. Haykin, Ed., *Adaptive Radar: Toward the Development of Cognitive Radar*. Hoboken, NJ: Wiley 2006, ch. 5.

[25] S. Haykin, *Neural Networks: A Comprehensive Foundation*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1999.

[26] D.P. Bertsekas and J.N. Tsitsiklis, *Neuro-Dynamic Programming*. Athenas Scientific, 1996.

[27] V. Krishnamurthy and R.J. Evans, "Beam scheduling for electronically scanned array tracking systems using multi-arm bandits," *IEEE Trans. Signal Processing*, vol. 49, pp. 2893–2908, 2001.

[28] S.L.C. Miranda, C.J. Baker, K.D. Woodbridge, and H.D. Griffiths, "Knowledge-based resource management for multifunction radar," *IEEE Signal Processing Mag.*, vol. 23, no. 1, pp. 66–76, 2006.

[29] P. Weber, A. Premji, T.J. Nohara, and C. Krasnor, "Low-cost radar surveillance of inland waterways for homeland security applications," in *Proc. 2004 IEEE Int. Radar Conf.*, Philadelphia, PA, 26–29 Apr. 2004.

[30] C.J. Baker, private communication, 17 Jan. 2005.

SP