Adaptive Waveform Design and Sequential Hypothesis Testing for Target Recognition With Active Sensors

Nathan A. Goodman, *Member, IEEE*, Phaneendra R. Venkata, and Mark A. Neifeld, *Member, IEEE*

***Abstract*—Cognitive radar is a recently proposed approach in which a radar system may adaptively and intelligently interrogate a propagation channel using all available knowledge including previous measurements, task priorities, and external databases. A distinguishing characteristic of cognitive radar is that it operates in a closed loop, which enables constant optimization in response to its changing *understanding* of the channel. In this paper, we compare two different waveform design techniques for use with active sensors operating in a target recognition application. We also propose the integration of waveform design with a sequential-hypothesis-testing framework that controls when hard decisions may be made with adequate confidence. The result is a system that updates multiple target hypotheses/classes based on measured data, customizes waveforms as the class probabilities change, and draws conclusions when sufficient understanding of the propagation channel is achieved.**

***Index Terms*—Cognitive radar, matched illumination, sequential detection.**

I. Introduction

ACTIVE sensors are increasingly being deployed in complex electromagnetic environments and as parts of low-power distributed sensor networks. In these environments, interference can be strong and nonstationary; meanwhile, transmit power and energy may be limited. Unfortunately, traditional sensors usually operate within a fixed frequency band, use a predefined suite of waveforms, and/or adapt to the propagation environment only through post-measurement signal processing [1]. Thus, in many cases it has been observed that traditional sensing modalities lack the flexibility necessary to provide adequate detection, tracking, and recognition performance in difficult/complex propagation and interference environments.

Effective sensor operation within complex environments requires adaptation via constant monitoring of interference, cooperation with other sensors, and optimized illumination waveforms. A newly proposed concept for optimizing the performance of active sensors within resource-constrained and interference-limited environments is cognitive radar [1]. The goal of cognitive radar (CR) is to provide the overall sensing system with the capability to integrate 1) information about the propagation environment, 2) the capabilities and positions of other sensors, 3) input from external knowledge bases, and 4) system objectives and priorities. This integration will enable a cognitive radar system to decide on its next course of action including such tasks as repositioning itself, requesting assistance from other sensors, or customizing an active waveform.

CR operates in a closed loop. Within this loop, active interrogations of the channel are optimized based on prioritized system objectives, understanding of the propagation channel, and other forms of prior knowledge. With each illumination, the system's understanding of the channel improves in response to collected data and other information. Haykin [1] suggests that such a CR system can be represented using a Bayesian formulation whereby many different channel hypotheses are given a probabilistic rating. As more information is collected, the parameters of the channel hypotheses and their relative likelihoods are updated. The goal of an illumination, therefore, is to efficiently reduce the uncertainty attributed to each channel hypothesis. Hard decisions are only made when confidence is sufficient or when necessity mandates an immediate action.

The two primary technologies applied to CR in this paper are matched illumination waveforms and sequential hypothesis testing. The first of these is related to the optimum design of illumination waveforms (i.e., probing signals). Probing signal design was considered in [2]–[5] where the goal was to identify the correct channel from among two known alternatives. When the number of hypotheses is greater than two, however, it is not possible to derive a closed-form solution for the optimum waveform. Useful alternative solutions have focused on maximizing the average or minimum distance between echoes from different hypotheses or on maximizing the average divergence between hypotheses [6]. Bell [3] considered the problem of extracting information from Gaussian ensembles of impulse responses in additive white Gaussian noise (AWGN). In [4], optimum waveforms for binary hypothesis testing were extended to include clutter echoes that depend on the transmitted signal. This formulation led to a nonlinear problem that required an iterative solution. The solution for the binary case has also been heuristically extended to the multihypothesis case in [7].

Other work has considered polarimetry [8], uncertainty in the assumed impulse responses [9], the Kullback-Leibler information criterion between target classes [10], and the theory of matched illumination waveforms [11], [12]. Signal design for improved target detection in the presence of range-Doppler-

Manuscript received September 1, 2006; revised January 25, 2007. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Arye Nehorai.

The authors are with the Department of Electrical and Computer Engineering, The University of Arizona, Tucson, AZ 85721 USA (e-mail: goodman@ece.arizona.edu; neifeld@ece.arizona.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSTSP.2007.897053

spread interference is considered in [13], [14], and several authors have also studied the waveforms used by echolocating bats [15], [16]. Echolocating bats use combinations of constant-frequency and modulated waveforms. Moreover, bats change the parameters of these waveforms as they get closer to their prey [15].

It is impossible for a CR system to decide *a priori* on the number of illumination waveforms that will be required to achieve a given level of confidence in the desired recognition decision. In order to enable the CR system to update its understanding of the channel without necessarily making hard target decisions after each illumination, we combine Bayesian channel representation and adaptive waveform design with the formalism of sequential hypothesis testing (SHT). In SHT, a decision is made after each observation to either select a hypothesis or, if a hypothesis cannot be selected with sufficient confidence, to make another observation. SHT was first motivated, formalized, and analyzed by Wald in [17], [18] where it was shown that a SHT procedure requires, on average, fewer observations than an equal-strength test with a fixed number of observations.

Multihypothesis sequential testing procedures have been provided in [19]–[22]. The test in [19] is often called the matrix sequential probability ratio test (SPRT) because it consists of a matrix of binary sequential tests. The procedure in [22] applies to multiple composite hypotheses. In this paper, we apply the relatively simple matrix multihypothesis test of [19]. The average number of observations for a sequential test has been extensively studied in [23]–[27]. While the original results presented in [17], [18] apply to binary problems where the observations are independent and identically distributed (i.i.d.), [23]–[27] extend the analysis of sequential testing procedures to multihypothesis tests such as the one in [19] and to non-i.i.d. scenarios. Of particular interest to us are the results in [23] and [26] because they include asymptotic lower bounds and approximate expressions for the average number of observations required by the matrix SPRT under non-i.i.d. observations. The non-i.i.d. case applies to our work because every observation is made with a different transmission waveform based on the current class probabilities. By *adapting* the SHT through updated waveforms, we desire to achieve an additional reduction in the average number of required observations (and, therefore, the number of transmissions).

In other related work, it should be mentioned that a sequential Bayesian approach to radar detection was presented in [28], [29]. In [28], the author demonstrated efficient sequential allocation of available energy to a set of detection problems in order to rapidly make decisions on the entire set. Although the work in [28], [29] used a fixed number of observations and did not explicitly consider waveform design, it is the earliest published work that the authors can find in which a radar system is treated as an adaptive interrogation system where the nature of the interrogations (i.e., the waveforms) are intelligently modified according to the results of previous interrogations.

This paper evaluates the benefits and potential of CR using target recognition as a sample application. The target recognition application is chosen because of its previous use in the literature for evaluating matched illumination techniques. The main contribution of this paper is to propose and analyze the coupling of sequential hypothesis testing with adaptive, matched waveforms in order to implement a closed-loop active interrogation system. In doing so, we quantify the performance benefit of using adaptive, matched waveforms compared to other waveforms. For example, we compare our adaptive waveform approach to waveforms that are matched to the original ensemble of target hypotheses but do not adapt throughout the experiment. These comparisons show the benefit of cognitive, closed-loop operation. We also contribute by applying the matched illumination technique of [4], [5] to a multihypothesis scenario, by implementing the information-based matched illumination technique of [3] in a practical and concrete example, and by comparing the performance of the two techniques. Finally, despite the fact that we cannot predict the particular path a given sequential experiment will take (such as what particular waveforms will be transmitted), we demonstrate that prior results on the average number of observations for non-i.i.d. sequential tests can be used to approximate the number of illuminations required by our closed-loop system.

In the next section, we define the problem statement and assumed signal model. The SHT procedure and Bayesian update equations are presented in Section III. In Section IV, two waveform design techniques are summarized. Simulations results are presented in Section V, and conclusions are made in Section VI.

II. Problem Statement and Signal Model

We consider the target identification problem in which one of M possible targets is known to be present. Note that one of these could also be the null hypothesis. Each target hypothesis is characterized by a known impulse response $h_i(t)$, $i = 1, 2, \ldots, M$. The objective of the CR system is to identify which target is present as efficiently as possible in terms of transmitted energy. Each time the radar transmits a waveform, a noise-corrupted version of the reflected target echo is received. The CR system uses the noisy data to improve its understanding of the channel and, if possible, to draw conclusions related to its objectives.

Let the waveform transmitted by the active system be denoted as $s(t)$. We assume the waveform is energy-limited since otherwise no waveform or system optimization would be necessary. Practical waveforms are also time-limited; thus, we assume the waveform is nonzero only in the time interval $-T/2 < t < T/2$. These restrictions require that

$$\int_{-\frac{T}{2}}^{\frac{T}{2}} s^2(t)dt = E \tag{1}$$

where E is the energy allocated to a single transmission. When a signal is transmitted, one of M possible waveforms is received according to

$$y(t) = s(t) * h_i(t) + n(t), \quad i \in \{1, 2, \ldots, M\} \tag{2}$$

where $*$ denotes the convolution operator and $n(t)$ is AWGN at the receiver with average power normalized to $\sigma_n^2 = 1$. The target impulse responses are assumed to be nonnegligible for a finite duration of time T_h; therefore, the received waveform can be observed over a finite duration of length $T_y = T + T_h$.

To facilitate simulation on a computer, we use a discrete-time formulation of the signal model. In the discrete-time formulation, the transmit waveform is sampled with a sampling interval of T_s. Hence, the transmit waveform is represented by a length-L_s vector $\mathbf{s}$ where $L_s T_s = T$. Using a normalized sampling interval of $T_s = 1$, the transmit energy constraint is

$$\sum_{l=1}^{L_s} s^2(l) = \mathbf{s}^{\mathrm{T}}\mathbf{s} = E \tag{3}$$

where $(\cdot)^{\mathrm{T}}$ is the transpose operator and E is defined in arbitrary units of energy. Likewise, the target impulse responses are sampled at the same rate to produce length-L_h impulse response vectors $\mathbf{h}_i$. It is convenient to define a $L_y \times L_s$ target convolution matrix $\mathbf{Q}_i$ (shown here for $L_h < L_s$) according to [5]

$$\mathbf{Q}_i = \begin{bmatrix} h_i(1) & 0 & \cdots & \cdots & 0 \\ h_i(2) & h_i(1) & \ddots & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ h_i(L_h) & h_i(L_h-1) & \cdots & h_i(1) & 0 \\ 0 & h_i(L_h) & h_i(L_h-1) & \cdots & h_i(1) \\ \vdots & 0 & h_i(L_h) & \cdots & h_i(2) \\ \vdots & \vdots & 0 & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & h_i(L_h) \end{bmatrix} \tag{4}$$

where the length of the receive signal vector is $L_y = L_s + L_h - 1$. Using the convolution matrix, the received data are

$$\mathbf{y} = \mathbf{Q}_i\mathbf{s} + \mathbf{n}. \tag{5}$$

III. Sequential Probability Ratio Testing and Bayesian Updates

In this paper, we focus on the target recognition application. A finite number of target alternatives are possible, each characterized by a prior probability of being true P_i. Together, the target impulse responses and their probabilities form our Bayesian representation of the channel.

While the end user objective usually involves hard decisions concerning target identification, a CR system interrogates the radar channel in order to update its probabilistic understanding of the channel. In other words, the closed-loop operation of CR is not necessarily concerned with making hard decisions—it is the goal of the system designer to draw actionable conclusions from the cognitive system's Bayesian representation of the channel. Instead, CR updates the probabilistic channel description after each illumination, and only after one channel alternative is clearly favored over the others are hard decisions made. This operational concept is an ideal application for SHT.

SHT is a procedure for statistical decision-making based on a sequence of observations. After each observation, a decision is made to either accept one of the hypotheses (thereby ending the test) or to continue the test by making another observation. Therefore, the number of observations required to complete the test is a random variable. In particular, we will consider the sequential probability ratio test (SPRT) [18] where the decision at each stage (or iteration) of the test is based on the likelihood ratio test. For a given error rate, the SPRT requires, on average, fewer observations than decisions based on the most powerful test for a fixed number of observations. Detailed descriptions of the SPRT and its extensions can be found in [18]. In this paper, we apply the multihypothesis sequential test of [19], which is summarized below.

Let there be M target hypotheses denoted as $H_1, H_2, \ldots, H_M$. Further, let $\alpha_{i,j}$ for $i \neq j$ be the desired probability of incorrectly selecting H_j given that H_i is true. The likelihood ratio including prior probabilities for a pair of hypotheses i and j after the k^{th} data observation is

$$\Lambda_{i,j}^k = \frac{p_{i1}(\mathbf{y}_1)p_{i2}(\mathbf{y}_2)\cdots p_{ik}(\mathbf{y}_k)}{p_{j1}(\mathbf{y}_1)p_{j2}(\mathbf{y}_2)\cdots p_{jk}(\mathbf{y}_k)}\frac{P_i}{P_j} \tag{6}$$

where $p_{ik}(\mathbf{y}_k)$ is the pdf of the k^{th} observation under the i^{th} hypothesis and $\mathbf{y}_k$ is the received data due to the k^{th} illumination waveform. The experiment is terminated and H_m is selected when the condition

$$\Lambda_{m,j}^k > \frac{1-\alpha_{m,j}}{\alpha_{m,j}} \quad \text{for all } j \neq m \tag{7}$$

is met for some m. After a given illumination and data collection, if the condition in (7) is not met, then another illumination cycle is made. If the condition is met, then the number of iterations/illuminations required to make the decision is $K = k$. The threshold in (7) to which the likelihood ratio is compared is taken from [25], [26]. It is the threshold necessary to prevent the average rate of making an error in favor of H_j when H_i is true from exceeding $\alpha_{i,j}$. Note that since the SPRT is based on a Bayesian approach, obtaining this error rate requires accurate knowledge of the target prior probabilities. Since these probabilities can be difficult to define for applications such as radar target recognition, this could be a limiting factor in achieving desired error performance levels.

For many investigations of the SPRT in the literature, the pdf of the observations does not change with iteration number. That is, $p_{i1}(\mathbf{y}) = p_{i2}(\mathbf{y}) = \cdots = p_{ik}(\mathbf{y})$, and the iteration index can be dropped from the pdf notation. In our case, however, we are updating the illumination waveform at each iteration. The pdf of the observations under AWGN is

$$p_{ik}(\mathbf{y}_k) = \frac{1}{\left(\sqrt{2\pi\sigma_n^2}\right)^{L_y}} \times \exp\left[-\frac{1}{2\sigma_n^2}(\mathbf{y}_k - \mathbf{Q}_i\mathbf{s}_k)^{\mathrm{T}}(\mathbf{y}_k - \mathbf{Q}_i\mathbf{s}_k)\right]. \tag{8}$$

Since the mean of the pdf depends on the transmit signal, a different pdf applies to each observation. Fortunately, the thresholds used by sequential testing to terminate the experiment do not depend on the actual distribution of the data, which means that the same likelihood threshold can be used regardless of the shape or energy of the transmit signal.

When the distribution of the likelihood ratio can be obtained, it can be used to approximate the statistical moments of the number of iterations, K. This is most easily done when the observations of the sequential test are i.i.d.; however, useful asymptotic results do exist for non-i.i.d. cases. In particular, we use the results of [23], [26] for the matrix SPRT to obtain asymptotic approximations for the average value of K, called the average sample number (ASN). From [23], [26], we know that the average number of illuminations, or ASN, can be approximated for small probabilities of error if the following convergence holds almost surely:

$$\frac{1}{f(k)} \log \Lambda_{i,j}^{k} \xrightarrow[k\to\infty]{} q_{ij} \tag{9}$$

where $f(k)$ is an increasing function and the q_{ij}'s are positive finite constants. More specifically, [23] proves that if (9) holds for $f(k)=k$, then the ASN of a sequential test when H_i is the true hypothesis can be lower bounded by

$$K_{lb} = \max_{j\neq i} \left(\frac{|\log \alpha_{i,j}|}{q_{ij}} \right). \tag{10}$$

Furthermore, as the desired error rates approach zero, the bound becomes tighter and becomes an approximation to the ASN.

When the test's observations are i.i.d., the q_{ij}'s in (9) and (10) are the Kullback–Leibler (KL) distances between the i^{th} and j^{th} hypotheses. In our application, however, the KL distances vary with the illumination waveform. Hence, in the adaptive waveform case there is not a unique KL distance to be used and it is not immediately clear what the q_{ij} values should be. In fact, it is not certain that the convergence in (9) will hold. Intuitively, however, we reason that as $k \to \infty$, the probability of the true hypothesis will approach one while the probabilities of the other hypotheses will approach zero. In this limit, the transmitted waveform will stabilize, which results in i.i.d. observations for large k. Furthermore, the observations for large k should always converge to the same distribution regardless of the particular waveforms used for small k. Hence, we conclude that (9) should still hold for our closed-loop observation system, and based on this conclusion, it should be possible to numerically estimate the q_{ij}'s through Monte Carlo simulation. This estimation has been performed with good results in Section V. For rigorous derivation of (9) and (10), please see [23], [26].

IV. Waveform Design Techniques

The two matched-illumination waveform design techniques that we will investigate are summarized in this section. The first technique is based on the work in [4], [5], [7]–[9] in which waveform pulse shaping was optimized to yield maximum SNR at the output of the receiver matched filter. This technique provides a provably optimal transmission waveform for the $M=2$ case and is heuristically extended for $M>2$. For $M=2$, define the target autocorrelation matrix as

$$\mathbf{\Omega} = (\mathbf{Q}_1 - \mathbf{Q}_2)^{\mathrm{T}} \mathbf{R}_n^{-1} (\mathbf{Q}_1 - \mathbf{Q}_2). \tag{11}$$

The optimum transmit signal is the one that separates the two signal echoes as far as possible in the whitened received signal space. The solution is the transmission waveform vector that maximizes the quantity $\mathbf{s}^{\mathrm{T}}\mathbf{\Omega}\mathbf{s}$, which is the eigenvector corresponding to the largest eigenvalue of $\mathbf{\Omega}$. For the present assumption of white noise, the noise covariance matrix is $\mathbf{R}_n = \mathrm{E}[\mathbf{n}\mathbf{n}^{\mathrm{T}}] = \sigma_n^2 \mathbf{I}$ where $\mathbf{I}$ is the identity matrix. Thus, the optimum waveform for the binary-hypothesis case is the eigenvector corresponding to the largest eigenvalue of $\mathbf{\Omega} = (\mathbf{Q}_1 - \mathbf{Q}_2)^T (\mathbf{Q}_1 - \mathbf{Q}_2)$. Once the eigenvector is found, it is normalized to have the proper energy as defined by (3).

When $M>2$, the strategy of creating maximum signal separation is less clear. There are several possible distances to maximize; hence, it is unclear if the average distance between receive echoes should be maximized, the minimum distance, or some other criterion. Moreover, even if the optimum signal criterion was known, it is not clear if a unique solution can be found. This is because we lack direct control of the receive signals. Instead, we have indirect control via a transmission waveform that interacts with target impulse responses. A logical approach suggested in [7] is to form an overall target autocorrelation matrix via a weighted sum of the individual matrices for each binary pair. In [7], the autocorrelation matrix for the M-ary case is suggested to be in the form

$$\mathbf{\Omega} = \sum_{i=1}^{M-1} \sum_{j=i+1}^{M} w_{i,j} (\mathbf{Q}_i - \mathbf{Q}_j)^{\mathrm{T}} (\mathbf{Q}_i - \mathbf{Q}_j) \tag{12}$$

where $w_{i,j}$ is a weight factor that accounts for the relative importance of discriminating between hypotheses i and j. Unfortunately, a specific suggestion for the weight factor is not provided in [7]. Using the target probabilities, we studied two options: $w_{i,j} = P_i + P_j$ and $w_{i,j} = P_i P_j$. The first option is an *ad hoc* selection that seems to make intuitive sense, but one can see that even if the probability assigned to one of the targets is zero, that target will still factor into the overall waveform. The second option is known to maximize the average divergence between output echoes [6], and we found that this weighting provided better results. After each transmission and reception by the radar system, if a decision cannot be made, the hypothesis probabilities are updated. Since the weights are based on these probabilities, the target autocorrelation matrix is modified after each data collection, which leads to a new waveform on the next transmission. Waveforms obtained from this approach will be referred to as eigen-based waveforms or the *eigensolution*.

The second matched-illumination waveform design technique is based on mutual information as described in [3]. Let $\mathbf{h}(t)$ be a random process that can be thought of as an ensemble of target impulse responses. We will assume that all of the sample functions of $\mathbf{h}(t)$ have finite energy and are causal impulse responses. If we further assume that $\mathbf{h}(t)$ is a Gaussian random process, then we can find the waveform that maximizes the mutual information between the ensemble of impulse responses and the received waveform. Let the waveform have finite energy E, be confined to the time interval $-T/2 < t < T/2$, and be essentially bandlimited such that most of its energy is contained within the frequency band $|f| \leq 1/2T_s$. The information-maximizing waveform under these constraints has the magnitude-squared spectrum defined by [3]

$$|S(f)|^2 = \begin{cases} \max\left[0, A - \frac{\sigma_n^2 T_y}{2\sigma_H^2(f)}\right] & |f| \leq \frac{1}{2T_s} \\ 0 & |f| > \frac{1}{2T_s} \end{cases}. \tag{13}$$

where T_y is the receive observation interval defined in Section II and the quantity $\sigma_H^2(f)$ is called the spectral variance and is defined by

$$\sigma_H^2(f) = \mathrm{E}\left\{ |\mathbf{H}(f) - \mathrm{E}\{\mathbf{H}(f)\}|^2 \right\}. \tag{14}$$

The constant A in (13) enforces the finite-energy constraint by satisfying

$$E = \int_{-\frac{1}{2T_s}}^{\frac{1}{2T_s}} \max\left[0, A - \frac{\sigma_n^2 T_y}{2\sigma_H^2(f)}\right] df. \tag{15}$$

In the present application, the objective is to decide between a finite number of target hypotheses that do not form a Gaussian ensemble. Hence, the waveform based on (13)–(15) is no longer optimum in terms of mutual information. Nevertheless, the approach is still intuitively satisfying and can be integrated with our Bayesian representation of the target hypotheses by defining the ensemble variance according to

$$\sigma_H^2(f) = \sum_{i=1}^{M} P_i \left|H_i(f)\right|^2 - \left|\sum_{i=1}^{M} P_i H_i(f)\right|^2. \tag{16}$$

Finally, we point out that the solution in (13) is obtained by performing the waterfilling action [30] on the function $\sigma_n^2 T_y / 2\sigma_H^2(f)$. Therefore, this waveform technique will be called the waterfilling solution.

In the following section, we apply these two matched illumination techniques to a multihypothesis target identification scenario. Using Monte Carlo simulation, we compare the two techniques and quantify the improvement obtained by adapting the matched waveforms in response to previous observations. In addition to the adaptive matching, the following results are interesting because the two matched illumination techniques have rarely been applied and compared in practical scenarios.

V. Results

In this section, we demonstrate the benefits of matched illumination for M-ary target identification and the improved convergence obtained through the adaptive CR approach. We also compare the performance of different waveform design approaches. To estimate the ASN for the closed-loop system, 1500 different sets of M target impulse responses are created by generating sample functions of a Gaussian random process with specified power spectral density (PSD). Once the impulse responses are generated, it is assumed that they are known perfectly. For each set of impulse responses, the adaptive sequential testing procedure is performed. The specified error rate is $\alpha_{i,j} = 0.01$ for all i and j, and the target probabilities prior to any observations are all equal. The length of all impulse responses and the waveform vector is $L_h = L_s = 31$ samples. For each trial we note the number of illuminations required for making a decision, and the final metric is the number of illuminations averaged over all trials. In these experiments, the 90% confidence intervals for 1500 trials are less than 10% of the absolute value of the estimated value.

After obtaining the ASN for each data point, we perform another Monte Carlo simulation to obtain a numerical estimate of the q_{ij}'s necessary for the asymptotic approximation. First, for each transmitted energy level and adaptive waveform technique, we define a fixed-length experiment equal to approximately twice the ASN of the variable-length sequential test. That is, the fixed length is $L \approx 2 \cdot \mathrm{ASN}$. Then, for each of the 1500 sets of impulse responses, we repeat the fixed-length test 50 times and store $\Lambda_{m,j}^k$ (m is the true hypothesis) for each iteration of each trial. After 50 trials, we average the $\Lambda_{m,j}^k$'s for each value of j to obtain an average path, $\overline{\Lambda}_{m,j}^k$, taken by each of the binary SPRT's over L observations. The value of L was chosen experimentally such that a plot of $\overline{\Lambda}_{m,j}^k/k$ versus k approached a constant value prior to termination of the fixed-length test. This constant value was taken to be q_{mj}. Finally, the lower bound for that particular set of hypotheses was computed as

$$K_{lb} = \max_{j \neq m}\left(\frac{|\log 0.01|}{q_{mj}}\right), \tag{17}$$

and the final data point for the lower bound curve was obtained by averaging the lower bound values over all impulse response sets.

For Fig. 1, we generated 1500 sets of $M = 4$ impulse responses from a flat PSD. Therefore, the targets are all sample functions of the same random process. Fig. 1 shows the average number of iterations required for each waveform design approach as a function of energy units allocated to a single illumination. The nonadaptive curves refer to the situation where the waterfilling or eigensolution waveforms were initially matched to the target ensemble with equal prior probabilities, but were not adapted as the hypothesis probabilities changed. The best-performing waveform is the eigensolution with adaptivity because it requires, on average, the fewest illuminations to reach a decision. The waterfilling waveform performs slightly poorer than the eigensolution while the impulse waveform (defined as $\mathbf{s} = [\sqrt{E}\ 0\ \cdots\ 0]^{\mathrm{T}}$) and nonadaptive approaches perform significantly worse. The lower bound (LB) curves, shown only for the two adaptive waveform techniques, are good approximations of the ASN. Furthermore, this approximation would improve if the error rate were reduced.

Fig. 1 clearly demonstrates the benefit of the SHT implementation of CR. Updating the system's understanding of the channel after each observation and updating a waveform to match that understanding reduces the average number of illumination cycles. Since each illumination has the same energy, reducing the number of illuminations also reduces the total amount of energy expended, which is useful in energy-limited applications such as distributed sensor networks. The impulse waveform performs approximately as well as the nonadaptive matched illumination approaches because it is matched to the PSD from which the impulse responses were created. If the bandwidth of the impulse signal were not matched to the bandwidth of the impulse responses, the performance would be much worse.

Fig. 2 shows results for a scenario in which the impulse responses are generated from different PSDs. For Fig. 2, each impulse response was generated from a PSD shaped like a Hanning window, but the windows were shifted by different amounts for each of the four hypotheses. Therefore, each hypothesis had

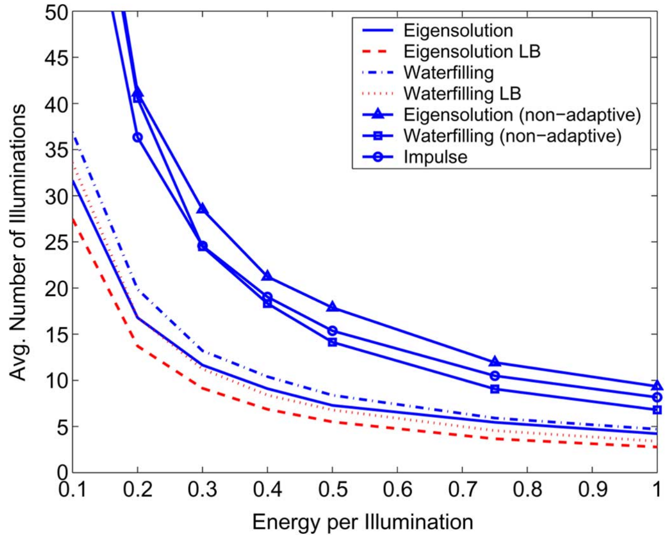

Fig. 1. Average illuminations to reach a decision versus energy per illumination.

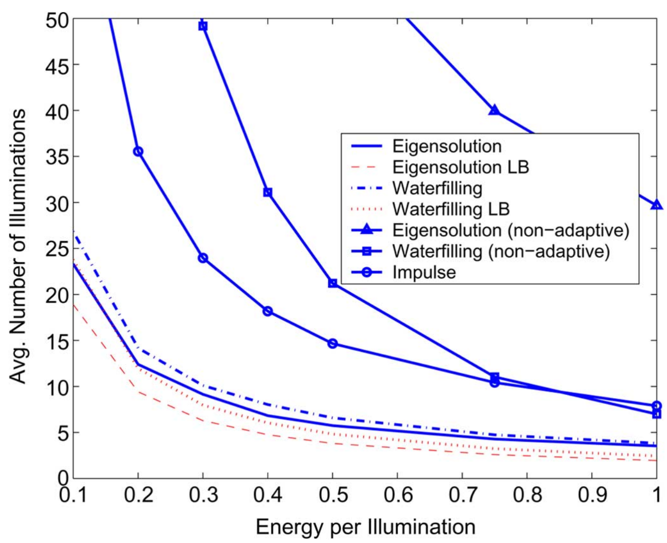

Fig. 2. Average illuminations to reach a decision versus energy per illumination.

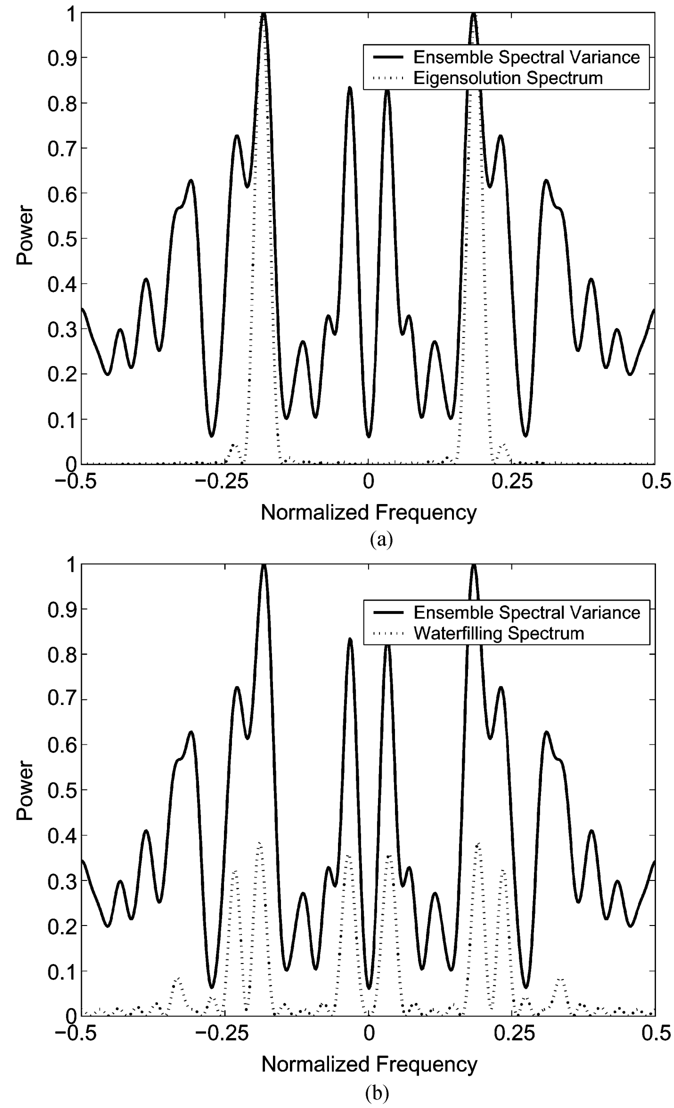

Fig. 3. Waveform spectra compared to ensemble variance. Eigensolution spectrum is shown in (a) and the waterfilling spectrum is shown in (b).

its energy concentrated (on average) into a different frequency band. Since the PSDs underlying the hypotheses were different, the best performance of the adaptive eigensolution and waterfilling waveforms is improved over the results of Fig. 1. It is interesting, however, that performance of the nonadaptive eigensolution and waterfilling waveform is actually worse than in Fig. 1. The performance of the impulsive waveform is relatively unchanged.

Despite having somewhat similar performance, the waveforms produced by the eigensolution and the waterfilling solution have quite different characteristics. In Fig. 3, we show sample power spectra of signals produced by the two techniques and compare them to the ensemble variance defined in (16). Both power spectra are for the first illumination in a sequential test. Fig. 3(a) compares the power spectrum of the eigensolution and the ensemble variance. The eigensolution focuses most of its energy into one or two narrow frequency bands that are most useful. This effect is clearly seen in Fig. 3(a) and has been noted in [5]. The waterfilling waveform, however, spreads its energy into several narrow bands as seen in Fig. 3(b). Both the ensemble variance and the eigensolution spectrum are normalized to their peak power. In order to show that the waterfilling solution allocates less energy to more bands, the waterfilling spectrum in Fig. 3(b) is normalized to the eigensolution spectrum's peak power. In this case, the waterfilling spectral peak is approximately 40% of the eigensolution's spectral peak.

The differing spectral characteristics of these two waveform approaches translate into differences in the distances between target echoes. In the same way that larger distance between symbols in a communication system reduces errors, larger distance between target echoes makes it easier for the system to make an accurate identification. Defining the distance between (noise-free) echoes as

$$d_{ij} = \left[(\mathbf{Q}_i\mathbf{s} - \mathbf{Q}_j\mathbf{s})^{\mathrm{T}}(\mathbf{Q}_i\mathbf{s} - \mathbf{Q}_j\mathbf{s})\right]^{\frac{1}{2}} \tag{18}$$

we have computed the average and minimum distance between target echoes for both waveform design approaches and for 100

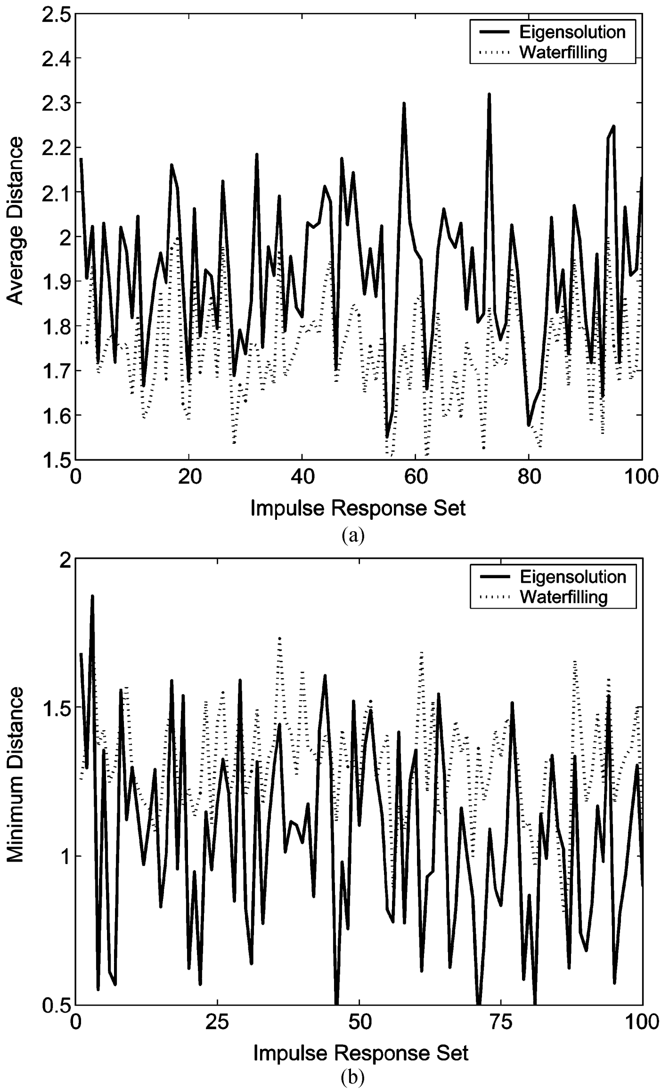

Fig. 4. Comparison of the (a) average distance and (b) minimum distance between receive echoes for both waveform techniques.

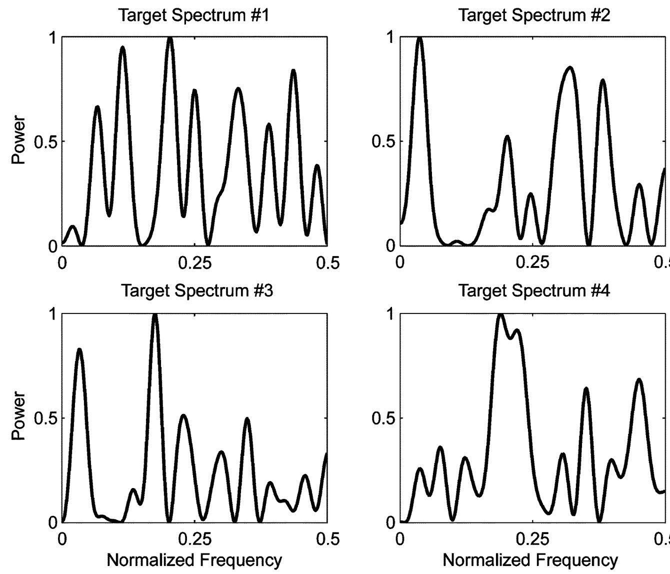

Fig. 5. Positive-frequency spectra of four sample target hypotheses.

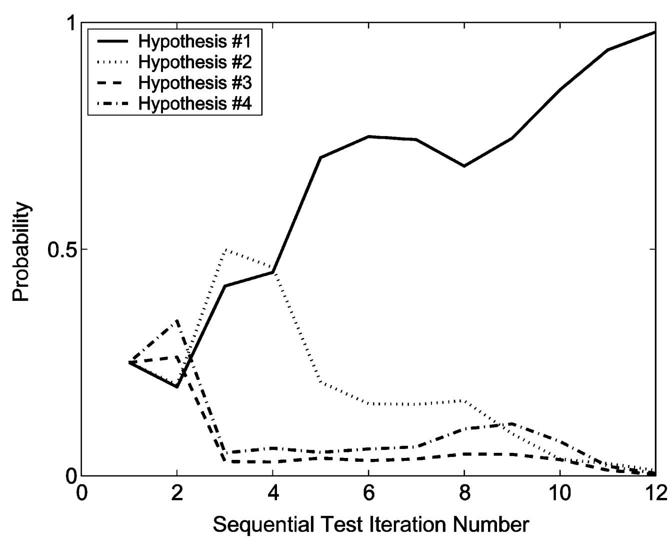

Fig. 6. Progression of hypothesis probabilities over a realization of a sequential test.

different impulse response sets. In Fig. 4(a), we see the average distance between received target echoes. The eigensolution waveforms clearly produce echoes with larger average distances than the waterfilling waveforms. However, the eigensolution focuses energy into only one or two narrow frequency bands, which may not be good frequency bands for all hypotheses. In other words, a few well-separated hypotheses can increase the average distance between echoes at the expense of a few poorly separated echoes. The waterfilling solution, on the other hand, spreads energy into several more frequency bands. As seen in Fig. 4(b), this characteristic means that the waterfilling solution usually produces the largest minimum distance between echoes. For this application, average distance seems to be more important since the eigensolution outperforms the waterfilling solution.

Finally, we present an example of how the waveforms and hypothesis probabilities evolve over a single sequential test trial. In Fig. 5, we show the positive half of the normalized power spectrum for four different target hypotheses generated from flat PSD's. These are the target hypotheses that we are attempting to discriminate in this experiment. The first hypothesis is the true hypothesis in this case. Fig. 6 shows the evolution of the hypothesis probabilities, and Fig. 7 shows four different waveforms used during the test. Prior to observing any data, the hypotheses are equally probable. The customized waveform for this equal probability case is shown in the upper-left panel of Fig. 7. After the first transmission, the probabilities do not change significantly, but after the second transmission, the first and second hypotheses become the clear favorites. This change in likelihood causes the waveform to change—careful comparison of waveforms #1 and #3 in Fig. 7 shows that the narrow band of signal energy has shifted in response to the updated probabilities.

For the next several iterations, the likelihood of hypothesis #1 becomes stronger while hypothesis #2 remains the next likely candidate. Since the same two hypotheses are clearly favored over this interval, the waveform spectra do not change considerably. Finally, after eight transmissions, the second and fourth hypotheses become approximately equally likely. This

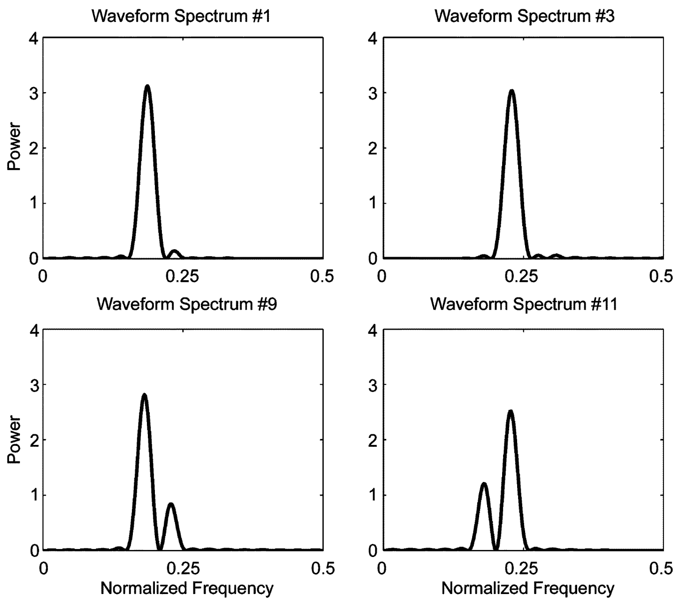

Fig. 7. Eigensolution waveform spectra at several key iterations of the sequential test.

results in waveform #9, where transmit energy is somewhat divided between two frequency bands. These are the bands that are useful for discriminating hypothesis #1 from hypotheses #2 and #4. When the likelihood of the first hypothesis becomes even stronger after the tenth transmission, the relative importance of the two frequency bands shifts a final time in waveform #11. After the twelfth transmission, the system is sufficiently confident to decide in favor of hypothesis #1 and the experiment is terminated.

VI. Conclusions

We have proposed and simulated a closed-loop active sensor by updating the probabilities on an ensemble of target hypotheses while adapting customized waveforms in response to prior measurements. We have compared the performance of two different waveform design techniques—one based on the eigensolution of [4], [5] and one based on information theory [3]. For both waveforms, the advantage of adapting waveforms on the fly in response to previously received data was substantial, and the eigensolution slightly outperformed the waterfilling-based approach. Moreover, we combined closed-loop radar operation with a sequential testing procedure that allows the system to update its understanding of the propagation channel until a hard decision can be made with confidence. These supporting technologies, acting together, yield a radar system that uses measured data and other information to continually update its understanding of the radar channel. In response to that understanding, the radar system transmits waveforms customized to the environment and objective at hand while making conclusions when sufficient confidence is achieved.

References

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

[2] R. Esposito and M. Schumer, "Probing linear filters—Signal design for the detection problem," *IEEE Trans. Inform. Theory*, vol. 16, no. 2, pp. 167–171, Mar. 1970.

[3] M. R. Bell, "Information theory and radar waveform design," *IEEE Trans. Inform. Theory*, vol. 39, no. 5, pp. 1578–1597, Sep. 1993.

[4] S. U. Pillai, H. S. Oh, D. C. Youla, and J. R. Guerci, "Optimum transmit-receiver design in the presence of signal-dependent interference and channel noise," *IEEE Trans. Inform. Theory*, vol. 46, no. 2, pp. 577–584, Mar. 2000.

[5] D. A. Garren, M. K. Osborn, A. C. Odom, J. S. Goldstein, S. U. Pillai, and J. R. Guerci, "Enhanced target detection and identification via optimised radar transmission pulse shape," *Proc. IEEE*, vol. 148, no. 3, pp. 130–138, Jun. 2001.

[6] E. Mosca, "Probing signal design for linear channel identification," *IEEE Trans. Inform. Theory*, vol. 18, no. 4, pp. 481–487, Jul. 1972.

[7] J. R. Guerci and S. U. Pillai, "Theory and application of optimum transmit-receive radar," in *Proc. IEEE 2000 Int. Radar Conf.*, Washington, DC, May 8–12, 2000, pp. 705–710.

[8] D. A. Garren, A. C. Odom, M. K. Osborn, J. S. Goldstein, S. U. Pillai, and J. R. Guerci, "Full-polarization matched-illumination for target detection and identification," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 38, no. 3, pp. 824–837, Jul. 2002.

[9] D. A. Garren, M. K. Osborn, A. C. Odom, J. S. Goldstein, S. U. Pillai, and J. R. Guerci, "Optimal transmission pulse shape for detection and identification with uncertain target aspect," in *Proc. IEEE 2001 Radar Conf.*, Atlanta, GA, May 1–3, 2001, pp. 123–128.

[10] S. M. Sowelam and A. H. Tewfik, "Waveform selection in radar target classification," *IEEE Trans. Inform. Theory*, vol. 46, no. 3, pp. 1014–1029, May 2000.

[11] D. T. Gjessing, *Target Adaptive Matched Illumination Radar Principles and Applications*, ser. IEE Electromagnetic Wave Series (22). London, U.K.: Peter Peregrinus, 1986.

[12] D. T. Gjessing, "Characterization of low observable targets and wakes by a matched illumination multifrequency, multistatic radar system," in *Proc. 1997 Radar Edinburgh Int. Conf.*, Oct. 1997, pp. 548–552.

[13] D. F. DeLong, Jr. and E. M. Hofstetter, "On the design of optimum radar waveforms for clutter rejection," *IEEE Trans. Inform. Theory*, vol. IT-13, no. 3, pp. 454–463, Jul. 1967.

[14] D. F. DeLong, Jr. and E. M. Hofstetter, "The design of clutter-resistant radar waveforms with limited dynamic range," *IEEE Trans. Inform. Theory*, vol. IT-15, no. 3, pp. 376–385, May 1969.

[15] J. A. Simmons, "The resolution of target range by echolocating bats," *J. Acoust. Soc. Amer.*, vol. 54, no. 1, pp. 157–173, 1973.

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

[17] A. Wald, *Sequential Analysis*. New York: Wiley, 1947.

[18] A. Wald, "Sequential tests of statistical hypotheses," *Ann. Math. Statist.*, vol. 16, no. 2, pp. 117–186, Jun. 1945.

[19] P. Armitage, "Sequential analysis with more than two alternative hypotheses and its relation to discriminant function analysis," *J. R. Statist. Soc.*, ser. B, vol. 12, no. 1, pp. 137–144, 1950.

[20] C. W. Baum and V. V. Veeravalli, "A sequential procedure for multihypothesis testing," *IEEE Trans. Inform. Theory*, vol. 40, no. 6, pp. 1994–2007, Nov. 1994.

[21] G. Lorden, "Nearly-optimal sequential tests for finitely many parameter values," *Ann. Statist.*, vol. 5, no. 1, pp. 1–21, 1977.

[22] I. V. Pavlov, "Sequential procedure of testing composite hypotheses with applications to the Kiefer-Weiss problem," *Theory Prob. Applicat.*, vol. 35, no. 2, pp. 280–292, 1990.

[23] V. P. Dragalin, A. G. Tartakovsky, and V. Veeravalli, "Multihypothesis sequential probability ratio tests—Part I: Asymptotic optimality," *IEEE Trans. Inform. Theory*, vol. 45, no. 7, pp. 2448–2461, Nov. 1999.

[24] V. P. Dragalin, A. G. Tartakovsky, and V. Veeravalli, "Multihypothesis sequential probability ratio tests—Part II: Accurate asymptotic expansions for the expected sample size," *IEEE Trans. Inform. Theory*, vol. 46, no. 4, pp. 1366–1383, Jul. 2000.

[25] A. G. Tartakovsky, "Asymptotically optimal sequential tests for nonhomogeneous processes," *Sequen. Anal.*, vol. 17, no. 1, pp. 33–61, 1998.

[26] A. G. Tartakovsky, "Asymptotic optimality of certain multihypothesis sequential tests: Non-i.i.d. case," *Statist. Infer. Stochast. Process.*, vol. 1, no. 3, pp. 265–295, 1998.

[27] N. V. Verdenskaya and A. G. Tartakovskii, "Asymptotically optimal sequential testing of multiple hypotheses for nonhomogeneous Gaussian processes in an asymmetric situation," *Theory Prob. Applicat.*, vol. 36, no. 3, pp. 536–547, 1991.

[28] D. R. Fuhrmann, "Active-testing surveillance systems, or, playing twenty questions with a radar," in *Proc. 11th Annu. Adaptive Sensor and Array Processing (ASAP) Workshop*, Lexington, MA, Mar. 11–13, 2003, MIT Lincoln Lab..

[29] D. Fuhrmann and L. Boggio, “Active-testing surveillance for multiple target detection with composite hypotheses,” in *Proc. 2003 IEEE Workshop on Statistical Signal Processing*, St. Louis, MO, Sep. 2003, pp. 641–644.

[30] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.

Nathan A. Goodman (S’98-M’02) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Kansas, Lawrence, in 1995, 1997, and 2002, respectively.

From 1996 to 1998, he was an RF Systems Engineer for Texas Instruments, Dallas, TX. From 1998 to 2002, he was a Graduate Research Assistant in the Radar Systems and Remote Sensing Laboratory, University of Kansas. He is currently an Assistant Professor in the Department of Electrical and Computer Engineering, University of Arizona, Tucson. Within the department, he directs the Laboratory for Sensor and Array Processing. His research interests are in radar and array signal processing.

Dr. Goodman was awarded the Madison A. and Lila Self Graduate Fellowship from the University of Kansas in 1998. He was also awarded the IEEE 2001 International Geoscience and Remote Sensing Symposium Interactive Session Prize Paper Award.

Phaneendra R. Venkata received the B.S. degree from University of Madras, India, in 2002 and the M.S. degree in electrical and computer engineering from the University of Arizona, Tucson, in 2006.

From 2003 to 2005, he was a Graduate Research Assistant in both the Laboratory for Sensor and Array Processing and the Intelligent Large Scale Sensing Systems (ILS^3) Laboratory, University of Arizona. Since February 2006, he has been with Qualcomm as a DSP Applications Engineer.

Mark A. Neifeld (M’06) received the B.S.E.E. degree from the Georgia Institute of Technology, Atlanta, in 1985 and the M.S. and Ph.D. degrees from the California Institute of Technology, Pasadena, in 1987 and 1991, respectively.

From 1985 to 1986 he was also a Member of Technical Staff at TRW Systems Engineering and Analysis Laboratory, Redondo Beach, CA. Following completion of his thesis work at Caltech he accepted a one-year post-doctoral position at the NASA Jet Propulsion Laboratory, Pasadena, where he studied the application of parallel image processing techniques to problems in pattern recognition. In 1991, he joined the faculty of the Department of Electrical and Computer Engineering and the Optical Sciences Center at the University of Arizona, Tucson. He has coauthored more than 70 journal articles and more than 200 conference papers in the areas of optical communications and storage, coding and signal processing, and optical imaging and processing systems. His current interests include information- and communication-theoretic methods in image processing, nontraditional imaging techniques that exploit the joint optimization of optical and post-processing degrees of freedom, coding and modulation for fiber and free-space optical communications, and applications of slow and fast light.

Dr. Neifeld is a Fellow of the Optical Society of America and a member of SPIE and APS. He has served on the organizing committees of numerous conferences and symposia. He has also been a two-term topical editor for *Applied Optics* and a three-time Guest Editor of special issues of *Applied Optics*.