Distribution Statement A.
Approved for public release;
distribution unlimited.

Adaptive prediction of battlefield signal propagation and detection

D. Keith Wilson[a], Carl R. Hart,[a] Chris L. Pettit,[b]
Daniel J. Breton,[a] and Edward T. Nykaza[c]

[a]U.S. Army Engineer Research and Development Center,
72 Lyme Rd., Hanover, NH, USA 03755

[b]U.S. Naval Academy, 590 Holloway Rd., Annapolis, MD, USA 21402

[c]U.S. Army Engineer Research and Development Center,
2902 Newmark Dr., Champaign, IL, USA 61822

ABSTRACT

Predictions of signal propagation and sensor system performance aid in understanding the vulnerabilities of friendly forces to detection and tracking by hostile forces, and vice versa. In practice, the accuracy of such predictions is limited by simplifications in the modeling physics and by uncertainties in the environmental inputs to the models. To address this problem, we consider the possibility of rapidly updating initial model predictions as more information on the actual sensor performance becomes available during an operation. Specifically, we propose a Bayesian sequential updating approach that incorporates realistic signal distributions for randomly scattered signals transmitted along a single path or multiple paths. We discuss how, in the Bayesian context, these various distributions represent likelihood functions, which are conveniently paired with their conjugate priors to provide efficient updates to the uncertain signal parameters. We also discuss how the Bayesian updating problem is closely connected to the problem of describing the impact of parametric uncertainties on the signal and noise distributions, and hence upon the receiver operating characteristic (ROC curve, a plot of the probability of detection vs. the probability of false alarm). In particular, we show that parametric uncertainties in quantities such as the mean signal and noise power substantially raise the tails of the distributions, which leads to severely degraded ROC curves. Formulation of automatic target recognition (ATR) algorithms, that use Bayesian updating to help compensate the signature for environmental effects, is also discussed.

Keywords: signal fading, environmental effects, uncertainty modeling, Bayesian methods, sensor performance, receiver operating characteristic

1. Introduction

Predictions of signal propagation and sensor system performance aid in understanding of the vulnerabilities of friendly forces to detection and tracking by hostile forces, and vice versa. Such predictions would typically be produced and analyzed as part of the mission planning process. The accuracy of such analyses is limited, however, by simplifications in the modeling physics and by uncertainties in the environmental inputs to the models. For example, the predictions often depend on atmospheric forecast or climatological data of limited resolution. Or, an operation may be undertaken in a data-denied area for which limited terrain elevation and landcover data are available. To address this problem, we consider dynamically updating the model predictions as more information on the actual sensor performance becomes available during mission execution. That is, we start with an initial imperfect prediction based on limited data, and then "nudge" the prediction in real time, as we gain a better understanding of how the systems are actually performing. With improvements in sensor networks and communication bandwidth, such an adaptive prediction process is becoming increasingly feasible.

Specifically, we propose here adaptive prediction based on Bayesian sequential updating. This approach can be done very rapidly, can incorporate realistic signal distributions and uncertainties in the parameters of these distributions, is rigorous, and applies regardless of the specific sources of the discrepancies between initial predictions and observations. A Bayesian approach is warranted by the complexity of real-world propagation and the limitations of computational methods to describe these processes.

The signal distributions we consider here are specifically intended for situations in which the signal is randomly scattered as it propagates along a fixed path between and a source and receiver. Although other types of random signal behavior are practically important, it is reasonable to focus on scattering initially, because this phenomenon is important in many problems. In this paper, we link the problems of modeling parametric uncertainties in scattering processes (e.g., uncertainties in the strength of the scattering process) to the problem of Bayesian updating. As far as we know, this connection has not been made previously.

After providing some background on the problem of parametric uncertainties in signal scattering (Sec. 2), we consider two realistic distributions for random signals propagating along a single transmission path (communication channel), namely: (1) exponential, as appropriate to strong signal fading, and (2) log-normal, as appropriate to weak fading. These are the topics of Secs. 3 and 5, respectively. In between, in Sec. 4, we make the connection between parametric uncertainties and Bayesian adaptation. We discuss how, in the Bayesian context, these various distributions represent likelihood functions, which are conveniently paired with their conjugate priors to provide computationally efficient updates to the uncertain signal parameters. The technique is demonstrated using simulations in which signal observations are collected at single or multiple receivers, and used to sequentially refine the prior distributions for the signal parameters.

In Sec. 6, we consider other statistical models for random signals, besides exponential and log-normal. In particular, we consider single-path models appropriate for both weak and strong fading, namely the gamma and Ricean distributions. For transmissions involving weak scattering along multiple paths (i.e., multiple emitter and sensor locations), we consider a jointly log-normal distribution, which includes correlations between the transmitted signals.

Sec. 7 considers some additional applications of the basic framework developed here. In particular, we discuss the impact of parametric uncertainties on the signal and noise distributions, and hence upon the receiver operating characteristic (ROC curve, a plot of the probability of detection vs. the probability of false alarm). The parametric uncertainties in quantities such as the mean signal and noise power

substantially raise the tails of the distributions, which leads to severely degraded ROC curves. Sec. 7 also includes a preliminary discussion of how Bayesian adaptation might be used to improve the performance of automatic target recognition (ATR) algorithms in the presence of random signal scattering and uncertainties.

2. Randomly scattered signals with parametric uncertainties

A conceptual diagram of the general problem addressed in this paper is shown in Fig. 1. The signal at a receiver consists of the signal of interest, which propagates along a direct path and multiple randomly scattered paths. The resulting random signal variations are often called signal *fading* or *scintillations*. Noise (when explicitly considered) originates from multiple, random sources. The random scattering and random noise mechanisms lead to a probabilistic distribution for the signal and noise. Based on these distributions, we wish to calculate the probabilities of detection (for the signal of interest) and false alarm. In the discussion to follow, *weak* scattering means that the signal power propagating along the direct path dominates the received signal, whereas *strong* scattering means that the randomly scattered contributions dominate.

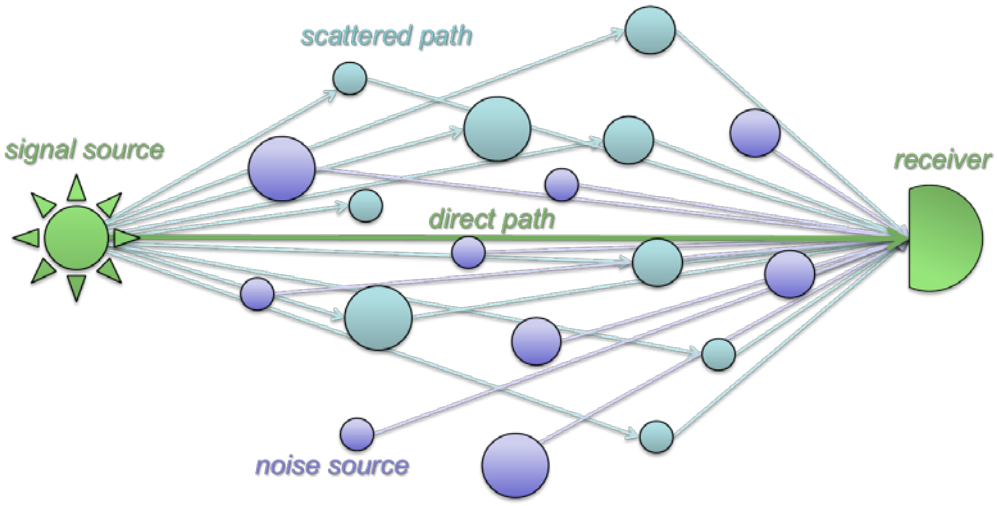

Figure 1. Conceptual diagram of the problem considered in this paper. The received signal consists of the signal of interest, which propagates along a direct path and multiple randomly scattered paths, plus noise, which originates from multiple, random sources.

The basic problem of detecting a signal in noise is discussed in many textbooks (e.g., Burdic 1991, Scharf 1991). The general starting point is the joint probability density function (pdf) for the signal s and noise n, which depends on some parameter set θ. We indicate this dependence as $p(s,n|\theta)$, where the vertical bar indicates that the random variables on the left, in this case s and n, are conditioned upon the values on the right, in this case θ. Here, θ may include one or more parameters. If the signal and noise are independent, $p(s,n|\theta) = p(s|\theta_s)p(n|\theta_n)$, where θ_s represents the subset of parameters from θ impacting the signal pdf, and θ_n is likewise the subset of parameters impacting the noise pdf.

The preceding framework can be generalized to include parametric uncertainties by introducing a compound probability density function, which consists of a pdf representing the solution for fixed values of the parameters θ, modulated by a higher-level distribution representing the variability of those parameters in space and/or time. Mathematically, the compound formulation for the joint pdf of the signal and noise is given by

$$p(s,n|\chi) = \int p(s,n|\theta)p(\theta|\chi)d\theta, \quad (1)$$

where χ is the set of so-called *hyperparameters*, which describe the distribution of θ. The dimensionality of the integral equals the number of uncertain parameters in θ.

If the signal and noise are independent, we have

$$p(s|\chi_s) = \int p(s|\theta_s)p(\theta_s|\chi_s)d\theta_s \tag{2}$$

and

$$p(n|\chi_n) = \int p(n|\theta_n)p(\theta_n|\chi_n)d\theta_n, \tag{3}$$

where χ has similarly been partitioned into hyperparameters impacting the signal and noise parameters.

3. Strong scattering along a single transmission path

In conditions of strong scattering, the received signal consists of many independent, randomized contributions. The real and imaginary parts of the complex signal are then zero-mean and normally distributed. This leads to a simple exponential pdf for the signal power; or, equivalently, a Rayleigh distribution for the amplitude (e.g., Burdic 1991, Dashen 2010). In this section, we consider how parametric uncertainties impact the exponential pdf.

Referring back to Eq. (2), in this case θ_s consists of a single parameter, representing the mean strength (power) of the scattered signal. Specifically, we take this parameter to be the *inverse* of the scattered signal strength, which we designate as λ. The reason for using the inverse is that this choice will lead to analytic solutions, as will be demonstrated shortly. We thus have

$$p(s|\chi_s) = \int_0^\infty p\,(s|\lambda)p(\lambda|\chi_s)\,d\lambda, \tag{4}$$

where, for strong scattering, $p(s|\lambda)$ is given by an exponential distribution:

$$p(s|\lambda) = \mathrm{Exp}(s|\lambda),$$

in which

$$\mathrm{Exp}(x|\lambda) = \lambda\exp(-\lambda x). \tag{5}$$

The mean of the exponential pdf is $m = \lambda^{-1}$ and the variance is $\sigma^2 = \lambda^{-2}$. Hence the normalized variance, σ^2/m^2, is 1 for the exponential pdf. This so-called *saturation* of the signal variations is a defining characteristic of strong scattering.

A variety of forms may be appropriate for $p(\lambda|\chi_s)$, depending on the nature of the parametric uncertainties being modeled. We adopt here a gamma pdf, mainly because it is appropriate for modeling positive definite quantities such as (inverse) mean signal power, and has the flexibility to address a variety situations, from those in which a random variable varies little from the mean, to those in which it is highly variable with a long-tailed distribution. The gamma pdf is given by

$$\mathrm{Gamma}(x|\alpha,\beta) = \frac{\beta^\alpha x^{\alpha-1}}{\Gamma(\alpha)}e^{-x\beta} \quad \text{for } x \geq 0 \text{ and } \alpha,\beta > 0, \tag{6}$$

where α is the shape parameter and β is the scale parameter. The mean of the gamma pdf is $m = \alpha/\beta$ and the variance is $\sigma^2 = \alpha/\beta^2$. Hence the parameters α and β can be determined by setting $\alpha = m^2/\sigma^2$ and $\beta = m/\sigma^2$. Note that the variance as normalized by the squared mean is $\sigma^2/m^2 = 1/\alpha$. Hence, when the shape factor α is large, the gamma pdf becomes relatively more peaked around the mean value.

Several cases of the gamma pdf are plotted in Fig. 2. In each case, the mean has been set to 1, which implies that $\alpha = \beta$. The value of β (and hence α) is varied from 0.5 through 8 is shown in the figure legend. The case $\alpha = \beta = 1$ corresponds to the exponential pdf.

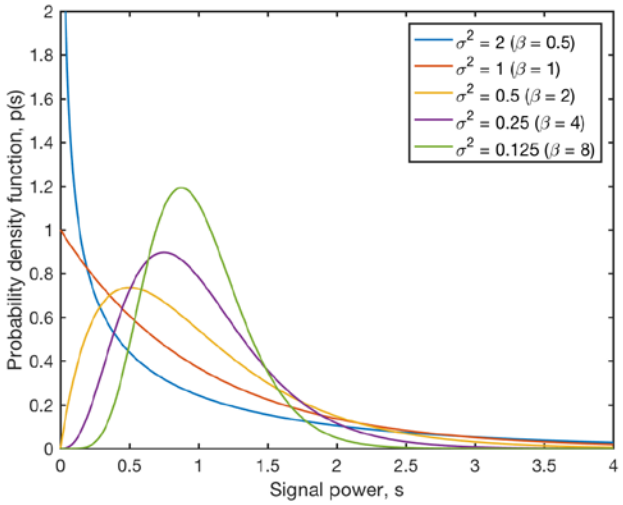

Figure 2. Gamma distribution for several values of the parameter β as indicated in the legend. Also shown in the legend are the corresponding values of the variance. The parameter α is set to β, so that the mean of the distribution is 1 for each case.

In the present development, the parameters α and β together comprise the set χ_s. By substituting Eq. (5) for $p(s|\lambda)$ and Eq. (6) for $p(\lambda|\alpha,\beta)$ into Eq. (4), we have

$$p(s|\alpha,\beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \int_0^\infty \lambda^\alpha \, e^{-(\beta+s)\lambda} d\lambda.$$

The integral can be found in standard tables, with result

$$p(s|\alpha,\beta) = \text{Lomax}(s|\alpha,\beta) = \frac{\alpha\beta^\alpha}{(s+\beta)^{\alpha+1}}. \tag{7}$$

The preceding pdf is called a *Lomax* or *Pareto Type II* distribution (Lomax 1954). This equation represents the pdf for the signal *after marginalizing for the uncertain parameters*. Put another way, $p(s|\lambda)$ is the signal pdf for a particular value of λ. The Lomax distribution results when λ is uncertain.

Figure 3 shows the Lomax pdf for various values of the parameter β. For this plot, we have fixed the signal mean to 1, which constrains α to the value $\beta + 1$. Also shown on the plot is the exponential distribution for a mean of 1 ($\lambda = 1$). The Lomax distribution converges to the exponential as β (and hence α) is increased. When β is small, the tail of the distribution (corresponding to large values of s) is substantially elevated in comparison to the exponential distribution. This behavior is particularly evident when the pdf is plotted on a logarithmic axis.

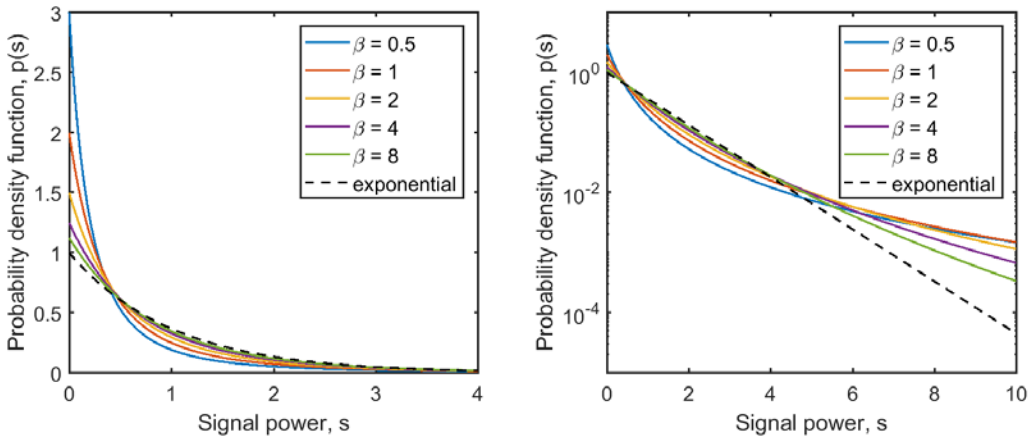

Figure 3. Lomax distribution. The parameter α is set to $\beta + 1$, so that the mean of the distribution is always 1. The dashed curve is the exponential distribution for a mean of 1. Left: signal power values up to 4, with a linear scale for the pdf. Right: signal power values up to 10, with a logarithmic scale for pdf.

The preceding formulation for strongly scattered signals is very similar to one originally proposed by Tatarskii and Zavorotnyi (1985) and Gurvich and Kukharets (1986). Those authors modeled wave scattering (RF and acoustic) by atmospheric turbulence using an exponential pdf, in which the strength of the turbulence varies in time and space. This phenomenon is known as turbulent *intermittency*. Wilson et al. (1996) demonstrated agreement of this formulation with data for near-ground sound propagation in open terrain. The primary difference between those papers and the present one is that a log-normal pdf was used for $p(\theta_s|\chi_s)$, which is justifiable by turbulence theory. However, the gamma pdf used here leads to similar signal behavior and (unlike the log-normal pdf for $p(\theta_s|\chi_s)$) enables analytical solutions.

4. Connection to Bayesian statistics

The previous discussion of compound distributions has a useful connection to Bayesian inference and adaptation, which is explored in this section. Readers may refer to Gelman et al. (2014) or other suitable textbooks for an introduction to Bayesian statistics.

Suppose we regard $p(\theta_s|\chi_s)$ in Eq. (2) as the *prior* distribution for the parameters θ_s; that is, $p(\theta_s|\chi_s)$ is the assumed distribution for θ_s before any observations of s are made. Next, we perform an experiment to collect a sample of s. We would then like to know the *posterior* distribution, $p(\theta_s|s,\chi_s)$, which indicates the improvement in our knowledge of θ_s resulting from the sample. By Bayes' theorem,

$$p(\theta_s|s,\chi_s) = \frac{p(s|\theta_s,\chi_s)\,p(\theta_s|\chi_s)}{p(s|\chi_s)}. \tag{8}$$

Here, $p(s|\theta_s,\chi_s)$ is the *likelihood* function for the signal observation; that is, the signal distribution conditioned upon a particular set of parameters. Using Eq. (2) to rewrite the denominator, we have

$$p(\theta_s|s,\chi_s) = \frac{p(s|\theta_s,\chi_s)\,p(\theta_s|\chi_s)}{\int p(s|\theta_s{}',\chi_s)\,p(\theta_s{}'|\chi_s)\,d\theta_s{}'}. \tag{9}$$

Although the denominator in Bayes' theorem is often not regarded as intrinsically interesting (being referred to as the "normalization"), in the context of the problem described Sec. 2, it does have an important interpretation. Namely, it is the probability of observing a particular value of the signal after marginalizing over the uncertain parameters. Put another way, whereas $p(s|\theta_s, \chi_s)$ is the pdf of the signal for the signal given the parameters θ_s and hyperparameters χ_s, $p(s|\chi_s)$ is the pdf of the signal after accounting for the uncertainty in θ_s.

For the problem described in the preceding section, in which the parameter set θ_s consists of λ, and χ_s consists of α and β, we have all the information needed to calculate $p(\theta_s|s, \chi_s) = p(\lambda|s, \alpha, \beta)$. Namely, the likelihood function $p(s|\theta_s, \chi_s) = p(s|\lambda, \alpha, \beta)$ in the numerator is given by the exponential distribution, and the prior $p(\theta_s|\chi_s) = p(\lambda|\alpha, \beta)$ is given by the gamma distribution. We have already determined that the denominator $p(s|\chi_s) = p(s|\alpha, \beta)$ is given by the Lomax distribution, Eq. (7). Substituting these three results into Eq. (8), we find after some algebra

$$p(\lambda|s, \alpha, \beta) = \frac{(\beta + s)^{\alpha+1}\lambda^{\alpha}}{\Gamma(\alpha + 1)} e^{-(\beta+s)\lambda}.$$

Comparing this result to Eq. (6), we see that $p(\lambda|s, \alpha, \beta)$ is a gamma distribution, namely

$$p(\lambda|s, \alpha, \beta) = \text{Gamma}(\lambda|\alpha + 1, \beta + s).$$

Thus, the posterior distribution, like the prior, is a gamma distribution, although the parameters are updated from α and β to $\alpha' = \alpha + 1$ and $\beta' = \beta + s$. When the family of the prior and posterior distributions is preserved in this manner, the prior distribution is said to be the *conjugate prior* of the likelihood function. We have thus shown that the gamma distribution is the conjugate prior of an exponential distribution likelihood function.

Note that we could carry on in this manner and incorporate additional samples of s, say s_i, where $i = 1, 2, \ldots, M$. The posterior distribution at each step becomes the prior distribution for the next step. Thus we have

$$\alpha' = \alpha + M,$$

and

$$\beta' = \beta + \sum_{i=1}^{M} s_i.$$

This useful result shows how an initial estimate for α and β (e.g., one based on a prediction from a propagation model) can be updated as new observations of the signal power become available.

Let us consider a numerical example illustrating the refinement of the pdf for the inverse mean signal power (λ) as more data are collected. Without loss of generality, we will assume that the actual value of λ is 1. Initially, we start with a gamma pdf with $\alpha = 2.5$ and $\beta = 1.6$ to represent the uncertainty in our knowledge of λ (the prior). These values imply an initial distribution for λ with a mean of 1.5625. The chosen values of α and β produce a very broad distribution, as indicative of a situation where our knowledge of the actual value of λ is limited. (In Bayesian terminology, the prior is rather *uninformative*.) We then begin to collect signal data s_i, each of which is a sample from the exponential pdf, Eq. (5). Note

that, because of the random scattering, each of the observations on s_i provides additional information on the true value of λ, although we never know this value exactly.

Figure 4 shows an example simulation with updates to the gamma pdf for λ after 1, 4, 16, 64, and 256 random trials (samples of s_i). After a single observation, the distribution remains close to the prior. After 256 observations, the distribution has a fairly sharp peak near the true value of λ.

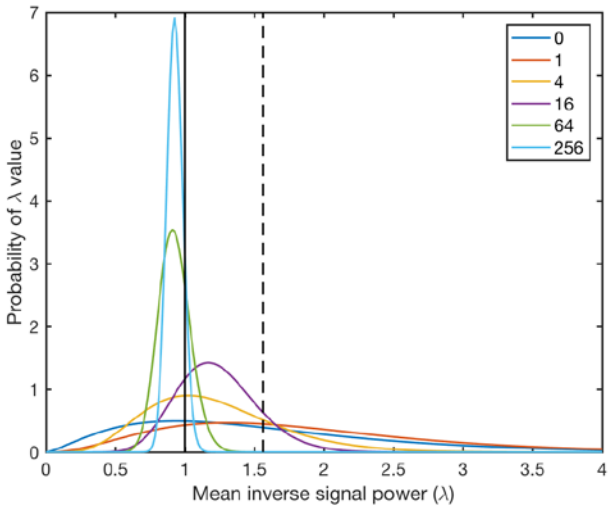

Figure 4. Refinement of the pdf for the signal distribution parameter λ as more signal samples are collected. The curve labelled 0 is the initial assumed distribution (prior). The subsequent curves show updated (posterior) distributions after 1, 4, 16, 64, and 256 random trials. The dashed vertical line indicates the mean of the prior (initial assumed mean); the solid vertical line indicates the correct value of λ.

5. Weak scattering (Rytov approximation)

In Sec. 3, we considered the case of strong scattering, as described by an exponential pdf for the signal power. A log-normal distribution (meaning that the logarithm is normally distributed) is widely employed for weak scattering; this signal behavior is predicted by the well known Rytov approximation (e.g., Andrews and Phillips 2005, Dashen et al. 2010). In both RF and acoustical modeling, the signal is often measured in decibels (dB), which is a logarithm of the signal power or amplitude. Hence adoption of a log-normal pdf would amount to assuming that signal in dB is normally distributed.

For the log-normal model, we represent the natural logarithm of the signal power, $\eta = \ln s$, as:

$$\mathcal{N}(\eta|\mu,\phi) = \frac{1}{\phi\sqrt{2\pi}}\exp\left[-\frac{(\eta-\mu)^2}{2\phi^2}\right]. \tag{10}$$

Here μ and ϕ^2 are the mean and variance of η. We call these parameters the log-mean and log-variance, respectively. By a transformation of variables, we find that the pdf for the signal s is

$$p(s|\chi) = \text{Lognorm}(s|\mu,\phi),$$

where $\text{Lognorm}(\cdot)$ is

$$\text{Lognorm}(s|\mu,\phi) = \frac{1}{s\phi\sqrt{2\pi}} \exp\left[-\frac{(\ln s - \mu)^2}{2\phi^2}\right]. \tag{11}$$

The mean of the signal s can be shown to be

$$m = e^{\mu+\phi^2/2},$$

whereas the variance is

$$\sigma^2 = \left(e^{\phi^2} - 1\right)e^{2\mu+\phi^2} = \left(e^{\phi^2} - 1\right)m^2.$$

We thus find for the log-mean,

$$\mu = \ln m - \frac{1}{2}\ln\left(1 + \frac{\sigma^2}{m^2}\right)$$

and for the log-variance,

$$\phi^2 = \ln\left(1 + \frac{\sigma^2}{m^2}\right).$$

Several cases of the log-normal pdf are plotted in Fig. 5. In each case, the mean has been set to 1, which implies that $\mu = -\phi^2/2$. The value of σ^2 is set to five different values as indicated in the figure legend.

Compound models for the log-normal distribution are most naturally formulated using the random variable $\eta = \ln s$, in which case the baseline distribution is normal. For Eq. (2), we have

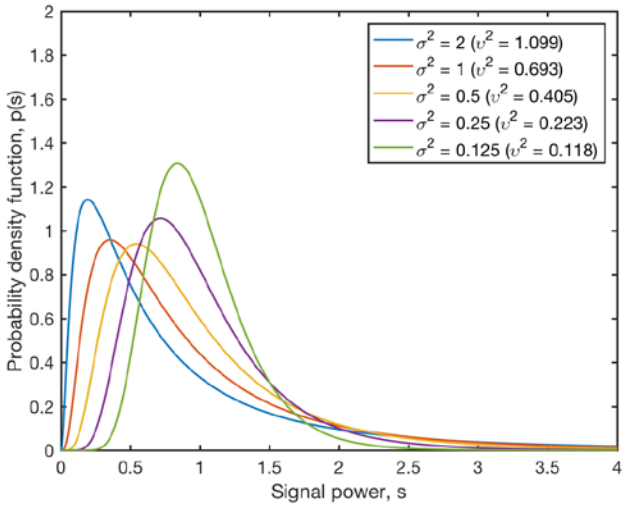

Figure 5. Log-normal distribution for several values of the variance, as indicated in the legend. Also shown in the legend are the corresponding values of the parameter ϕ^2. The mean of the distribution is 1 for each case.

$$p(\eta|\chi_s) = \int p(\eta|\theta_\eta)p(\theta_\eta|\chi_\eta)d\theta_\eta,$$

where θ_η are the parameters of the distribution for η, and χ_η are the hyperparameters for the distribution of θ_η. For the log-normal distribution, θ_η may consist of μ or ϕ, or both.

Let us consider μ to be an unknown parameter, and assume that μ is normally distributed with mean m_μ and variance σ_μ^2. Hence we have

$$p(\eta|m_\mu, \sigma_\mu^2) = \int p(\eta|\mu)p(\mu|m_\mu, \sigma_\mu^2)d\mu,$$

where

$$p(\mu|m_\mu, \sigma_\mu^2) = \mathcal{N}(\mu|m_\mu, \sigma_\mu^2) = \frac{1}{\sigma_\mu\sqrt{2\pi}}\exp\left[-\frac{(\mu - m_\mu)^2}{2\sigma_\mu^2}\right].$$

Utilizing Bayes' theorem, Eq. (8), we have the following equation for updating the estimate of μ after a new observation of the log-signal, η:

$$p(\mu|\eta, m_\mu, \sigma_\mu^2) = \frac{p(\eta|\mu, m_\mu, \sigma_\mu^2)p(\mu|m_\mu, \sigma_\mu^2)}{p(\eta|m_\mu, \sigma_\mu^2)}.$$

After some algebra and use of integral tables, we find

$$p(\mu|\eta, m_\mu, \sigma_\mu^2) = \mathcal{N}(\mu|m'_\mu, \sigma'^2_\mu) = \frac{1}{\sigma'_\mu\sqrt{2\pi}}\exp\left[\frac{(\mu - m'_\mu)^2}{2\sigma'^2_\mu}\right],$$

where

$$m'_\mu = \frac{\phi^2 m_\mu + \sigma_\mu^2\eta}{\phi^2 + \sigma_\mu^2} = \left(\sigma_\mu^{-2} + \phi^{-2}\right)^{-1}\left(\sigma_\mu^{-2}m_\mu + \phi^{-2}\eta\right)$$

and

$$\sigma'^2_\mu = \frac{\phi^2\sigma_\mu^2}{\phi^2 + \sigma_\mu^2} = \left(\sigma_\mu^{-2} + \phi^{-2}\right)^{-1}.$$

Figure 6 is similar to Fig. 4, except that a simulation of Bayesian updating for the log-mean (as based on the just-described log-normal signal model) is shown. In this simulation, the mean m and variance σ^2 were both set to 1, as expected in strong scattering. The resulting values of the log-signal parameters are $\mu = -0.3466$ and $\phi^2 = 0.6931$. For the prior distribution of μ, we somewhat arbitrarily set $m_\mu = -0.1466$ and $\sigma_\mu^2 = 2.773$, to represent an uninformative prior. The figure clearly demonstrates the convergence to a peaked distribution around the correct value of μ as more data samples become available.

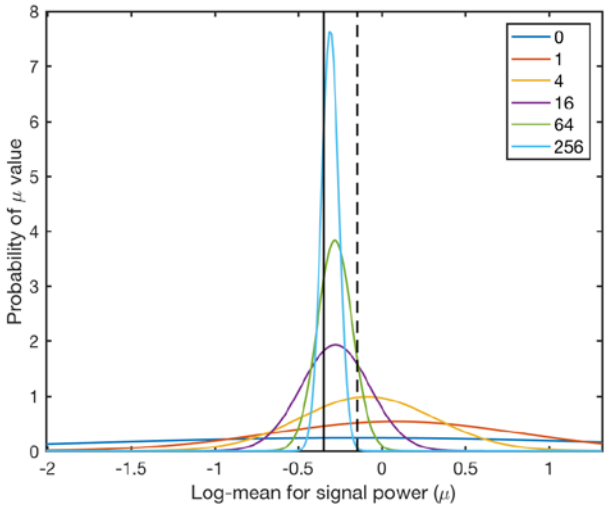

Figure 6. Refinement of the pdf for the log-mean signal distribution parameter, μ, as more log-signal samples are collected. The curve labelled 0 is the initial assumed distribution (prior). The subsequent curves show updated (posterior) distributions after 1, 4, 16, 64, and 256 random trials. The vertical dashed line indicates the initial mean of the prior; the solid line indicates the correct value of μ.

Up until this point, we have considered signal transmissions along a single path. Let us now suppose there are multiple paths, and that the transmissions along these paths may be correlated to some degree. In particular, we have in mind a situation where there are multiple source and receiver (monitor) locations. Each source/receiver pair provides a different transmission path, so that there are a total of $K = N_s \times N_r$ paths, where N_s is the number of sources and N_r the number of receivers. The multivariate distribution represents the cross correlations of the received signal power along each of these paths. These cross correlations on the proximity and overlap of the paths. This case is fully analyzed in Wilson et al. (2017). For present purposes, we will simply mention that the updating equations are very similar to those for a single path, except that the mean generalizes to *K*-element vector, and the variance to a $K \times K$ covariance matrix.

Figure 7 is similar to Fig. 6, except that a bivariate case ($K = 2$) is shown; that is, there are two propagation paths. In this simulation, the mean $\mathbf{m}$ was set to $[1; 2]$ and the variances $\boldsymbol{\sigma}^2$ to $[1; 4]$, in accordance with strong scattering (for which the ratio of the variance to the squared mean is one). The resulting mean for the log-signal parameters is $\boldsymbol{\mu} = 0.3466[-1; 1]$. (Note that the means of the two signals differ, in order to add realism to the problem.) The matrix $\boldsymbol{\Phi}$ is $\phi^2[1\ \rho; \rho\ 1]$, where $\phi^2 = 0.6931$ and $\rho = 0.5$. For the prior distribution of $\boldsymbol{\mu}$, we set $\boldsymbol{\mu}_\mu = -0.1466[1; 1]$ and $\Sigma_\mu = 2.773[1\ \rho; \rho\ 1]$, to represent an uninformative prior. As with the previous single variate example, the bivariate update converges well to distributions that are sharply peaked around the correct values of the means.

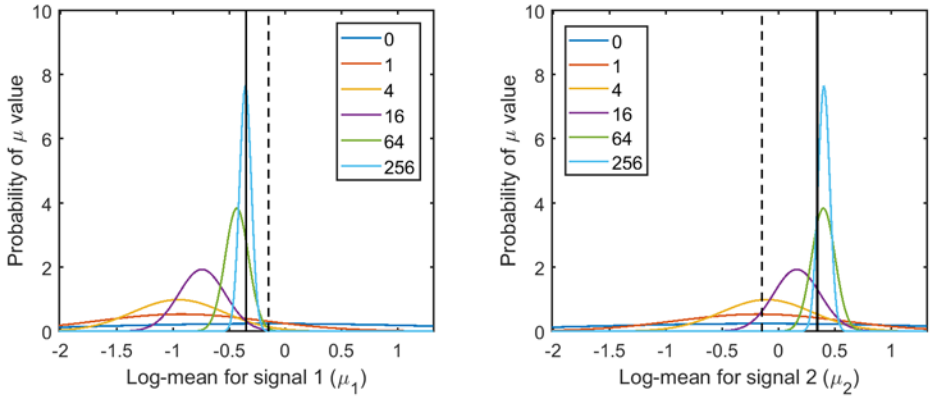

Figure 7. Same as Fig. 6, except for a bivariate normal distribution. Left is the prior distribution and its updates for the first transmission path; right is for the second transmission path.

6. Other signal models

Characterization of signal and noise power distributions is important for many applications involving signal detection and the robustness of communication channels. Up until this point we have considered just two particular limiting cases, namely strong and weak (Rytov approximation) scattering. Other formulations are available for signal distributions that bridge the weak and strong regimes, and represent other phenomena such as multipath propagation, interference, and superposition of multiple sources (e.g., Burdic 1991, Dashen et al. 2010, Wepman and Sanders 2011). The amplitude of a signal that is weakly scattered or embedded in random Gaussian noise may be described with a Ricean distribution (Strohbehn et al. 1975, Suzuki 1977). The Ricean distribution also interpolates between strong scattering and weak scattering in the Born approximation (Andrews and Phillips 2005). Multiple sources of equal amplitude and random phase lead to a chi-square distribution for power (Burdic 1991). Beyond these idealized cases, however, analytical solutions are often unavailable and it becomes necessary to fit an appropriate, general distribution to empirical data or simulations. Some of the distributions which have been commonly used in such situations include log-normal (Turin et al. 1972, Strohbehn et al. 1975), gamma, generalized gamma (Ewart 1989), Nakagami (1993), and Weibull (Tzeremes and Christodoulou 2002). Various arguments can be made in favor of these distributions, depending on the nature of the multipath and/or source distribution that is present, as well as empirical observations.

In Table 1, we list a number of important signal distributions, and indicate connections to Bayesian statistics. Some additional background underlying the results shown in the table can be found in Wilson et al. (2017). The table deals explicitly with signal power, as opposed to signal amplitude (proportional to the square root of power.) Note that some commonly encountered distributions, such as the Rayleigh and Nakagami, are the amplitude counterparts of the power distributions shown in the table (exponential and gamma, respectively). As discussed earlier, the basic (unconditioned) pdf for the random signal corresponds to the likelihood function. Bayesian conjugate priors, where known,[*] are available for the cases shown in red (although this relationship has apparently not been previously exploited). One noteworthy gap, evident from Table 1, is that there is no *multivariate* (multiple path) solution available

[*]See, for example, https://en.wikipedia.org/wiki/Conjugate_prior#Table_of_conjugate_distributions, which lists significant cases of likelihood functions with known conjugate priors. (Accessed: 2017-10-12.)

for strong scattering. (Recall that in Sec. 5 we discussed the multiple path formulation for weak scattering.) This gap is unfortunate, because strong scattering is commonly encountered in practice.

Table 1. Models for scattered signal power distributions with connections to Bayesian statistics.

Likelihood function (random signal model)	Physical interpretation in terms of signal scattering	Prior distribution (for parameters of signal model)	Posterior distribution (for parameters of signal model)	Posterior predictive (marginalized signal model)
Exponential	Strong scattering, single path	Gamma	(non-analytic)	K-distribution
Exponential	Strong scattering, single path	Inverse gamma	Inverse gamma	Lomax
Exponential	Strong scattering, single path	Log-normal (turbulent intermittency)	(non-analytic)	(non-analytic)
Gamma	Weak (approx.) or strong, single path	Gamma	(non-analytic)	Generalized K-distribution
Gamma	Weak (approx.) or strong, single path	Inverse gamma	Inverse gamma	Generalized beta--prime
Log-normal	Weak scattering, single path	Normal	Normal	Student T
Log-normal, multivariate	Weak scattering, multiple paths	Multivariate normal	Multivariate normal	Student T, multivariate
Ricean	Weak (Born) to strong scattering	(?)	(?)	(?)

7. Applications

7.1 Receiver operating characteristics

In this section, we consider the impact of parametric uncertainties on detection of signals in noise. The basic problem and notation were described in Sec. 2. Calculations of detection and false-alarm probabilities require knowledge of the pdf for two cases: (1) when the observed signal x consists only of noise ($x_0 = n$), and (2) when x consists of signal plus noise ($x_1 = s + n$). The pdf for the first case, which we designate as $p_0(x)$, can be found by integrating (marginalizing) the joint pdf for the signal and noise over the signal:

$$p_0(x|\theta_n) = p(n|\theta_n) = \int p(s, n|\theta)\, ds. \tag{12}$$

The pdf for the second case, $p_1(x)$, can be found by setting $n = x - s$, and integrating over s:

$$p_1(x|\theta) = \int p(s, x - s|\theta)\, ds. \tag{13}$$

Given the pdfs $p_0(x)$ and $p_1(x)$, we can determine the probabilities of false alarm and detection using the equations

$$P_{\mathrm{fa}}(\theta_n) = \int_{\gamma}^{\infty} p_0\,(x|\theta_n)\,dx \tag{14}$$

and

$$P_{\mathrm{d}}(\theta) = \int_{\gamma}^{\infty} p_1\,(x|\theta)\,dx. \tag{15}$$

Here, γ is the detection threshold, a quantity set by the detection algorithm. The threshold determines the trade-off between the detection and false-alarm probabilities. A plot of P_{d} as a function of P_{fa} is called the receiver operating characteristic, or ROC (curve). Of course, we would normally desire a high probability of detection and low probability of false alarm, but this is not always possible. The constant false-alarm rate (CFAR) detector provides a simple and useful approach to setting the threshold; namely we set the probability of false alarm, and then use Eq. (15) to infer γ. (See, for example, Burdic (1991).)

Let us consider a numerical example. We assume that the signal and noise both obey gamma distributions, $p(s|\lambda_s; k_s) = \mathrm{Gamma}(s|k_s, \lambda_s)$ and $p(n|\lambda_n; k_n) = \mathrm{Gamma}(n|k_n, \lambda_n)$, respectively. The parameters λ_s and λ_n are uncertain and also modeled by gamma distributions, with hyperparameters (α_s, β_s) and (α_n, β_n), respectively. The parameters k_s and k_n are considered fixed.

For the gamma distribution, the mean value of s is $\langle s\rangle = \beta_s k_s/(\alpha_s - 1)$, whereas the mean value of n is $\langle n\rangle = \beta_n k_n/(\alpha_n - 1)$. In the following calculations, we will set $\langle s\rangle = 1$, without loss of generality. Defining the mean signal-to-noise ratio as $\mathrm{SNR} = \langle s\rangle/\langle n\rangle$, we have $\langle n\rangle = 1/\mathrm{SNR}$. The values for the α parameters, $\alpha_s = \beta_s k_s + 1$ and $\alpha_n = \beta_n k_n(\mathrm{SNR}) + 1$, then follow. Hence there are four free parameters in the problem as defined, namely β_s, k_s, β_n, and k_n.

Analytical solutions are not available for all the distributions necessary to calculate the ROC curves, particularly $p_1(x)$. However, Monte Carlo methods provide a viable and relatively simple way to study the behavior of these curves. The basic idea is that, for specified values of k_s and β_s, we draw a large number M of random samples for λ_s from a gamma distribution with the parameters $(\alpha_s = \beta_s k_s + 1\,, \beta_s)$. Using these values for λ_s and the specified value for k_s as parameters in a gamma distribution, we then draw M random values for s. A set of M random values for n are generated similarly. Using the random noise samples to approximate $p_0(x)$, we can determine the value of γ corresponding to a target value of P_{fa}.

ROC curves based on such Monte Carlo simulations are shown in Figs. 8 and 9. A CFAR detector is employed. For both figures, the SNR is set to 2. The signal and noise have gamma distributions as previously described, with shape factors $k_s = 1$ and $k_n = 4$ for Fig. 8, and $k_s = k_n = 4$ for Fig. 9. The scaling parameters λ_s and λ_n also follow gamma distributions, with β_s and β_n being varied together from $\beta = 0.5$, 1, 2, 4, and 8, as shown in the legends of the figures. The parameters α_s and α_n are determined as described above. The value of β is seen to have a very substantial impact on the ROC curves. For large β (cases where the mean is much larger than the variance), the calculations converge on the conventional prediction, which does not include variations in the value of λ (i.e., there is no intermittency in the scattering). Smaller values of β lead to the occurrence of more extreme values for λ, which unfavorably impacts detection, particularly by increasing the occurrence of false alarms.

The primary conclusion to be drawn is that parametric uncertainties have a first-order impact on the ROC curves. In practice, the ROC curve is substantially worse than we would expect if the signal parameters are perfectly known (which is typically the case in practice). To help overcome this problem, we could use the Bayesian sequential updating to reduce the parametric uncertainties as additional observations of the signal and noise become available.

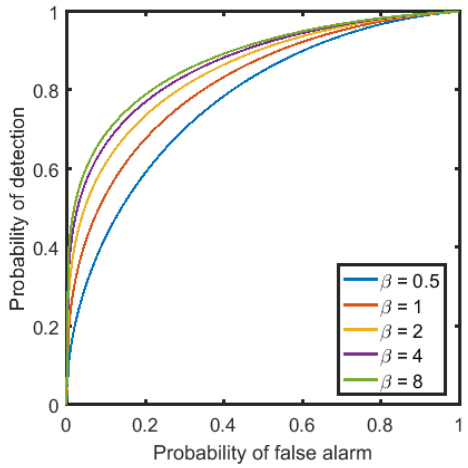

Figure 8. Receiver operating characteristic (ROC) curves corresponding to a gamma-distributed signal with shape factor $k_s = 1$ and gamma-distributed noise with $k_n = 4$. The signal-to-noise ratio (SNR) is 2, and the hyperparameter β is varied as shown in the legend.

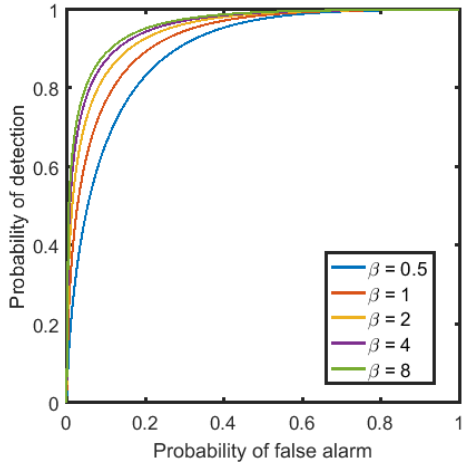

Figure 9. Same as Fig. 8, except that $k_s = 4$ and $k_n = 4$.

7.2 Automatic target recognition

Due to random processes impacting the generation and transmission of acoustic, RF, optical, and other types of signals, the apparent signature of a target, as observed at a receiver, varies randomly. The random variations in the signature degrade the performance of automated target recognition (ATR) algorithms. As discussed in this paper, one of the strengths of the Bayesian formulation is that it provides a way to

update our knowledge of the signal parameters as more observations are made of the received signal. Thus we can more accurately estimate the true signature. We provide here a preliminary analysis for how the Bayesian updating of signal parameters might be used as part of an ATR algorithm. Application of this analysis has not yet been attempted.

The formulation of the problem is illustrated in Fig. 11. The source signal propagates randomly along multiple transmission paths. Each sensor receives a version of the feature set (vector) which is distorted during transmission, by a combination of deterministic and random effects. The variations in the signature are modeled by the hyperparameters, namely the *statistical parameters for the received signal feature set*, at each receiver location.

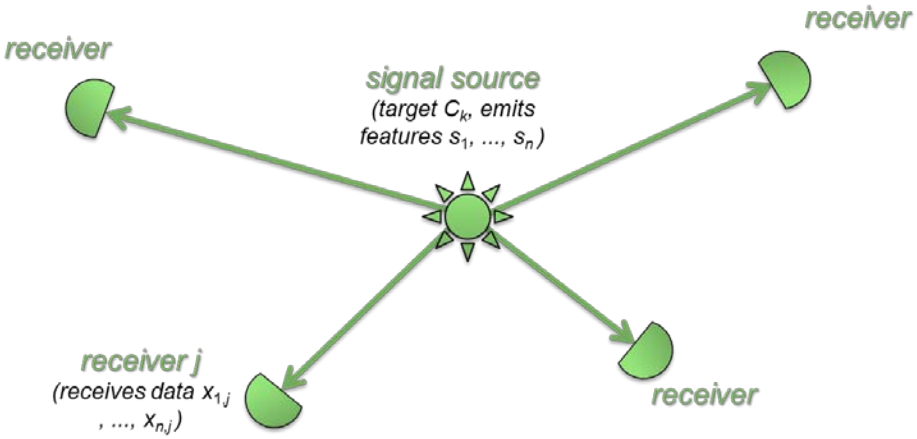

Figure 10. Transmission problem involving a single source and multiple receivers. Each sensor receives a randomized set of data, or features, which together represent the source signature.

Let us define C_k as the target class, indexed by k, $\mathbf{x}$ as the vector of signal features (all features, all receivers), $\boldsymbol{\theta}$ as the vector of pdf parameters for signal features (all features, all receivers). Using various probabilistic identities, we can then formulate the posterior probability for a given target as follows:

$$p(C_k|\mathbf{x}) = \frac{p(C_k)p(\mathbf{x}|C_k)}{p(\mathbf{x})} = \frac{p(C_k)}{p(\mathbf{x})}\int p(\mathbf{x}|\boldsymbol{\theta})p(\boldsymbol{\theta}|C_k)d\boldsymbol{\theta}. \tag{16}$$

Here, $p(C_k)$ is the prior probability, $p(\mathbf{x}|C_k)$ is the likelihood of the observed feature set given target C_k, $p(\mathbf{x})$ is the unconditioned probability for the observed feature set (a normalizing factor), $p(\mathbf{x}|\boldsymbol{\theta})$ is the probability for observed feature set given modeled propagation from target C_k, and $p(\boldsymbol{\theta}|C_k)$ is the probability for modeled signal parameters (hyperparameters) for target C_k.

We may view the update process as consisting of two steps. First, is the classifier update, which provides the probability that a particular target is being observed, given the observed feature values:

$$p(C_k|\mathbf{x}) = \frac{p(\mathbf{x}|C_k)p(C_k)}{p(\mathbf{x})}. \tag{17}$$

Next is the hyperparameter update, which updates the propagation prediction for a particular target, given the observed feature values:

$$p(\boldsymbol{\theta}|\mathbf{x}, C_k) = \frac{p(\mathbf{x}|\boldsymbol{\theta})p(\boldsymbol{\theta}|C_k)}{p(\mathbf{x}|C_k)}. \tag{18}$$

In a sequential updating process, the posterior at the current time step becomes the prior at the next time step; that is, $p(C_k|\mathbf{x})$ becomes the new $p(C_k)$, and $p(\boldsymbol{\theta}|\mathbf{x}, C_k)$ becomes the new $p(\boldsymbol{\theta}|C_k)$. Hence our knowledge of the target probability and the true signal parameters improves as more observations are collected.

8. Conclusions and future work

There is an extensive literature on how wave scattering impacts the statistical distributions of signals; many pertinent references were provided here. However, signal processing methods rarely utilize this understanding of signal behavior directly. In this paper, we have endeavored to more fully exploit the connection between signal processing and models for wave scattering.

In particular, we explored a connection between modeling of parametric uncertainties in the signal statistics, as based on a compound pdf formulation, and Bayesian inference of the signal parameters. The compound formulation pairs a physics-based model for the wave scattering with a higher-level distribution describing uncertainties in the wave scattering parameters, using hyperparameters. In the Bayesian context, the physics-based scattering model corresponds to the likelihood function. The likelihood function can be conveniently paired with its conjugate prior, if available, to efficiently update the uncertain signal parameters. The original (prior) distributions, as predicted using an initial forecast based on available weather and terrain data, can then be refined as additional observations are collected.

A solution was presented to the problem when the signal parameters are predicted through an initial estimate and then updated by repeated sound transmissions made along the same path, and used to refine the parameter distribution. The approach is thus adaptive in the sense that it learns the characteristics of the transmission channel. Some preliminary simulation results were provided which support the feasibility of this general approach. With improvements in sensor networks and communication bandwidth, such adaptive prediction and classification methods are becoming increasingly feasible.

We also explored two practical applications of this formulation, namely calculation of receiver operating characteristics (ROC curves) and automatic target recognition (ATR). Regarding the former application, we showed how the probabilities of false alarm and detection are significantly altered when parametric uncertainties in the signal and noise distributions are included in the ROC curves. One of the primary predictions is that the tails of the pdfs for the signal and noise are significantly elevated when uncertainties are included. Since detector thresholds are normally set to make false alarms rare, and we normally wish to operate with a probability of detection close to 1, the tails of the signal and noise distributions are very important. However, these parts of the distributions are also the least predictable. It is plausible that, in practice, extreme values of the signal and noise are more frequent than expected, due to our limited knowledge of the many processes impacting the signal and noise. This effect significantly degrades system performance.

Regarding the ATR problem, we described a Bayesian approach to target classification that incorporates physics-based signal propagation predictions and wave scattering models. The approach provides an optimal posterior distribution, as appropriate for use in a Bayesian classifier. By basing ATR algorithms on realistic statistics for the signal propagation, we are hopeful that this approach will help to overcome the “brick wall” in ATR that results from environmental impacts on signatures.

Acknowledgements

This research was funded by the U.S. Army Engineer Research and Development Center (ERDC) Environmental Quality/Installations business area. Permission to publish was granted by Director, Cold Regions Research and Engineering Laboratory. Any opinions expressed in this paper are those of the authors, and are not to be construed as official positions of the funding agency or the Department of the Army unless so designated by other authorized documents.

References

Andrews, L. C., and Phillips, R. L. (2005): *Laser Beam Propagation through Random Media* (SPIE Press, Bellingham, WA).

Burdic, W. S. (1991): *Underwater Acoustic System Analysis* (Prentice Hall, Englewood Cliffs, NJ).

Dashen, R., Flatté, S. M., Munk, W. H., Watson, K. M., and Zachariasen, F. (2010): *Sound Transmission through a Fluctuating Ocean* (Cambridge University Press, Cambridge, 2010).

Ewart, T. E. (1989): "A model of the intensity probability distribution for wave propagation in random media," *J. Acoust. Soc. Am.* **86**(4), 1490–1498 (1989).

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2014): *Bayesian Data Analysis* (Chapman & Hall/CRC, Boca Raton).

Gurvich, A. S. and Kukharets, V. P. (1986): "The influence of intermittence of atmospheric turbulence on the scattering of radio waves," *Sov. J. Commun. Technol. Electron.* **30**, 52–58.

Hashemi, H. (1993): "The indoor radio propagation channel," *Proc. IEEE* **81**(7), 943–968.

Lomax, K. (1954): "Business failures: Another example of the analysis of failure data," *Journal of the American Statistical Association* **49**(268), 847–852.

Scharf, L. L. (1991): *Statistical Signal Processing* (Addison-Wesley, Reading, MA, 1991).

Strohbehn, J. W., Wang, T.-i., and Speck, J. P. (1975): "On the probability distribution of line-of-sight fluctuations of optical signals," *Radio Sci.* **10**(1), 59–70.

Suzuki, H., "A statistical model for urban radio propagation" (1977): *IEEE Transactions on Communications* **25**(7), 673–680.

Turin, G. L., Clapp, F. D., Johnston, T. L., Fine, S. B., and Lavry, D. (1972): "A statistical model of urban multipath propagation," *IEEE Trans. Vehicular Technology* **21**(1), 1–9.

Wepman, J. A. and Sanders, G. A. (2011): *Wideband Man-Made Radio Noise Measurement VHF and Low UHF Bands* (US Department of Commerce, National Telecommunications and Information Administration).

Tatarskii, V. I. and Zavorotnyi, V. U. (1985): "Wave propagation in random media with fluctuating turbulent parameters," *J. Opt. Soc. Am. A* **2**, 2069–2076.

Tzeremes, G. and Christodoulou, C. (2002): "Use of Weibull distribution for describing outdoor multipath fading," in *Antennas and Propagation Society International Symposium, 2002* (IEEE) 232–235.

Wilson, D. K., Wyngaard, J., and Havelock, D. (1996): "The effect of turbulent intermittency on scattering into an acoustic shadow zone," *J. Acoust. Soc. Am.* **99**(6), 3393–3400.

Wilson, D. K., Breton, D. J., Hart, C. R., and Pettit, C. L. (2017): "Impact of signal scattering and parametric scattering on receiver operating characteristics," SPIE Defense and Commercial Sensing (International Society for Optics and Photonics, Anaheim).