Effect of environmental prediction uncertainty on target detection and tracking

Lawrence D. Stone* and Bryan R. Osborn
Metron Inc, 11911 Freedom Drive, Reston, VA 20190-5602

ABSTRACT

Past attempts to use acoustic sensor performance predictions, typically probability of detection as a function of range, in naval undersea warfare tactical decision aids such as trackers and mission planning tools have met with great difficulty. These efforts have been hampered by the uncertainty often inherent in these predictions. In some cases, the use of incorrect predictions produced results or recommendations that were worse than not using the predictions at all. The goal of the work reported in this paper is to develop a Track–Before–Detect (TBD) system that accounts for this uncertainty and has the following features: (1) It produces results are at least as good as those obtained with no performance prediction information. (2) It produces a significant improvement in performance in some situations. In this paper we describe an extension of a TBD system called the Likelihood Ratio Tracker (LRT) that incorporates uncertainty in performance prediction. We have run LRT on data that simulate a multistatic active sonar detection system similar to one in use by the Navy. In these simulated cases, we have shown that using performance prediction improves LRT tracking and detection performance even in the presence of large prediction errors.

Keywords: Track-before-detect, Likelihood ratio, Bayesian

1. INTRODUCTION

In many undersea warfare situations it is desirable to predict the performance of acoustic sensors in detecting targets. These predictions can be used to develop efficient search plans and to improve localization resulting from underwater detection. The predictions rely on estimates of environmental variables that are used to calculate the components of the signal excess equation described below. This equation is used to predict probability of detection as a function of range and bearing of the target from a sensor.

The authors have developed a version of the Bayesian Track-Before-Detect (TBD) system called the Likelihood Ratio Tracker (LRT) that accounts for uncertainty in sensor performance prediction by extending the LRT state space to include uncertainty in predicted mean signal excess. (See Stone et al[1] for a description of likelihood ratio detection and tracking.) In this paper we compare the results of running LRT on simulated data with and without performance prediction. The results show that incorporating performance prediction along with its uncertainty improves the performance of the tracker in the situations considered below. The significance of this result is that it indicates that performance prediction can be used in naval tactical decision aids to improve their performance provided that the uncertainty in those predictions is properly accounted for.

Researchers working on the Uncertainty Directed Research Initiative of the Office of Naval Research have been exploring the process of characterizing and quantifying the uncertainty in acoustic environmental predictions. In addition they are developing methods for reflecting and displaying the effects of that uncertainty on tactical systems that rely on environmental inputs. Examples of such systems are those that predict sensor performance, recommend search plans, or track and detect targets. In this paper, we discuss how we characterize uncertainty in the environmental predictions for the components of the sonar equation for multistatic active detection, and how this characterization is quantified and incorporated into the Bayesian TBD system, LRT. We present an example showing the application of LRT to multistatic active detection and tracking that accounts for the uncertainty in the acoustic predictions. For this example, the explicit inclusion of environmental uncertainty in the LRT makes it more robust to errors in predicting signal excess and estimating detection probabilities. Furthermore, it produces improved detection and tracking performance compared to using no performance prediction estimates.

---

* stone@metsci.com; phone 703 787 8700 x440; fax 703 787 3518; metsci.com

2. DETECTION AND UNCERTAINTY MODEL

For multistatic active detection, the mean signal excess (in dB) is given by Urick[2] as

$$\overline{SE} = SL - TL_1 - TL_2 + TS - (RL \oplus AN) - DT \tag{1}$$

where

SL = source level of ping (dB)
TL_1 and TL_2 = transmission loss to the target and from the target to the receiver (dB)
TS = target strength (dB)
RL = reverberation (dB)
AN = ambient noise (dB)
DT = detection threshold (dB).

The detection threshold DT is the signal to noise ratio (in dB) at which detection occurs with probability 0.5. $\overline{SE}$ is the mean signal excess predicted for a single sensor in the multistatic active system. Urick[2] observes that about this mean there are short term fluctuations that that are approximately Gaussian in dB. Typically, the fluctuations have mean 0 and standard deviation of 8 or 9 dB. Let ξ be a Gaussian random variable with mean 0 and standard deviation σ. In Urick's model, detection occurs when $\overline{SE} + \xi > 0$. The variation represented by ξ is predictable only in a statistical sense and is already accounted for in many tactical decision aids used by the Navy.

The uncertainty that we are primarily concerned with comes from the possibility that we have mis-estimated the mean of one or more of the components of signal excess equation (1). This produces an uncertainty in $\overline{SE}$, the mean signal excess. This mis-estimation can be caused by a using a poor estimate of the sound speed profile, bottom type, or any environmental input required for computing $\overline{SE}$. In addition to environmental uncertainty, there may be uncertainty in the estimate of target strength TS. The resulting uncertainty in $\overline{SE}$ is represented by a probability distribution on $\overline{SE}$ as a function of target state. The initial distribution on prediction error is computed from the uncertainty distributions on each of the components in (1). These uncertainties are summed (their distributions are convolved) to produce the prior distribution on $\overline{SE}$ used by LRT.

3. DESCRIPTION OF LRT

In this section we provide a brief description of Bayes Markov single target tracking and then likelihood ratio detection and tracking. For a more detailed discussion see Stone et al[1].

3.1 Bayes Markov single target tracking

Let S be the state space of the target. Typically, the target state will be a vector of components including position, velocity, and possibly acceleration. There can be additional components that may be related to the identity or other features of the target. Let $X(t)$ be the (unknown) target state at time t. We start the problem at time 0 and are interested in estimating $X(t)$ for $t \geq 0$. The prior information about the target is represented by a Markov process $\{X(t); t \geq 0\}$. This process is specified by a prior probability distribution on the state of the target at time 0 and a Markov motion model that describes the target's motion through the state space S. There are one or more sensors that produce a set of K discrete observations or measurements $\mathbf{Y}(t) = (Y_1, \ldots, Y_K)$ obtained in the time interval [0, t]. The observations are received at the discrete (possibly random) times $(t_1, \ldots, t_K)$ where $t_0 = 0 \leq t_1 \ldots \leq t_K \leq t$.

3.1.1 Likelihood functions

Let y_k denote the value of the random variable Y_k which represents the kth measurement. Note Y_k can represent multiple measurements received at time t_k. We assume that we can compute the likelihood function

$$L_k(y_k \mid s) = \mathbf{Pr}\left\{Y_k = y_k \mid X(t_k) = s\right\} \text{for } s \in S. \tag{2}$$

Note, we use $\mathbf{Pr}$ to mean either probability or probability density as appropriate. The computation in (2) can account for correlation among sensor responses if that is required. Let $\mathbf{Y}(t) = (Y_1, Y_2, \ldots, Y_K)$ and $\mathbf{y} = (y_1, \ldots, y_K)$. Define

$$L(\mathbf{y} \mid s_1, \ldots, s_K) = \mathbf{Pr}\left\{\mathbf{Y}(t) = \mathbf{y} \mid X(t_1) = s_1, \ldots, X(t_K) = s_K\right\}$$

We assume that

$$\mathbf{Pr}\left\{\mathbf{Y}(t)=\mathbf{y} \mid X(u)=s(u), 0 \leq u \leq t\right\}=L\left(\mathbf{y} \mid s(t_1),\ldots,s(t_K)\right)=\prod_{k=1}^{K} L_k(y_k \mid s_k) \tag{3}$$

Equation (3) means that the likelihood of the data $\mathbf{Y}(t)$ received through time t depends only on the target states at the times $\{t_1,\ldots,t_K\}$ and that the likelihood function L_k is independent of L_j for $j \neq k$ given $X(t_k)=s_k$.

3.1.2 Bayes Markov recursion

Let

$$p_0(s)=\mathbf{Pr}\left\{X(0)=s\right\} \tag{4}$$

be the initial distribution on target state at time 0. The Markov motion model for the target has the transition function

$$q_k(s_k \mid s_{k-1})=\mathbf{Pr}\left\{X(t_k)=s_k \,\middle|\, X(t_{k-1})=s_{k-1}\right\} \text{ for } k \geq 1.$$

Define

$$p\left(t_k, s \mid \mathbf{Y}(t_k)\right)=\mathbf{Pr}\left\{X(t_k)=s \mid \mathbf{Y}(t_k)=(y_1,\ldots,y_k)\right\} \text{ for } s \in S \tag{5}$$

to be the posterior distribution on target state s at time t_k given the observations (measurements) $\mathbf{Y}(t_k)$ received, through time t_k. Under the above assumptions this posterior may be computed in the recursive fashion given below.

Bayes Markov Recursion for Single Target Tracking

Initial Distribution $\quad p\left(0,s \mid \mathbf{Y}(0)\right)=p_0(s) \text{ for } s \in S \quad (6)$

For $k \geq 1$ and $s \in S$,

Motion Update $\quad p^-\left(t_k,s \mid \mathbf{Y}(t_{k-1})\right)=\int q_k(s \mid s_{k-1})\, p\left(t_{k-1},s_{k-1} \mid \mathbf{Y}(t_{k-1})\right) ds_{k-1} \quad (7)$

Calculate Likelihood $\quad L_k(y_k \mid s)=\mathbf{Pr}\left\{Y_k=y_k \mid X(t_k)=s\right\} \text{ for } s \in S \quad (8)$

Information Update $\quad p\left(t_k,s \mid \mathbf{Y}(t_k)\right)=\frac{1}{C} L_k(y_k \mid s)\, p^-\left(t_k,s \mid \mathbf{Y}(t_{k-1})\right) \quad (9)$

The constant C in (9) normalizes the right-hand side to a probability distribution (density).

3.2 Likelihood ratio detection and tracking

In the Bayes recursion given above, we have assumed that there is one target present and that all the observations (measurements) are associated to that target. Likelihood ratio detection and tracking is based on an extension of the single target tracking methodology presented above to the case where there is either one or no target present.

For LRT we identify a region $\mathcal{R}$ of interest. As above we let S be the single target state space for a target in the region $\mathcal{R}$. If no target exists in $\mathcal{R}$, we shall formally refer to this state as the *null state* and designate it by the symbol ϕ. We augment the target state space S with this null state to make $S^+ = S \cup \phi$. The augmented state space S^+ includes not only all of the states within $\mathcal{R}$ but the discrete null state ϕ as well. We shall assume there is a probability (density) function p defined on S^+. Since we are assuming that there is at most one target in the region $\mathcal{R}$, we may write

$$p_0(\phi)+\int_{s \in S} p_0(s)\,ds=1.$$

We define the ratio of the posterior state probability (density) $p\left(t,s \mid \mathbf{Y}(t)\right)$ to the posterior null state probability $p\left(t,\phi \mid \mathbf{Y}(t)\right)$ as the *target likelihood ratio (density)* $\Lambda\left(t,s \mid \mathbf{Y}(t)\right)$; that is,

$$\Lambda\left(t,s \mid \mathbf{Y}(t)\right)=\frac{p\left(t,s \mid \mathbf{Y}(t)\right)}{p\left(t,\phi \mid \mathbf{Y}(t)\right)} \text{ for } s \in S. \tag{10}$$

Following the notation in the Bayes Markov recursion, we let

$$\Lambda^-(t_k, s \mid \mathbf{Y}(t_{k-1})) = \frac{p^-(t_k, s \mid \mathbf{Y}(t_{k-1}))}{p^-(t_k, \phi \mid \mathbf{Y}(t_{k-1}))} \text{ for } s \in S . \tag{11}$$

The likelihood ratio density has the same dimensions as the state probability density. Furthermore, from the likelihood ratio density one may easily recover the state probability density as well as the probability of the null state.

In LRT, the measurement likelihood ratio takes the place of the measurement likelihood function in Bayesian single target tracking. It is the ratio of the likelihood of receiving the measurement given a target is present to the likelihood of receiving the measurement given no target is present. Specifically, we define the measurement likelihood ratio for the measurement $Y_k = y_k$ to be

$$\mathcal{L}_k(y_k \mid s) = \frac{L_k(y_k \mid s)}{L_k(y_k \mid \phi)} \text{ for } s \in S .$$

3.2.1 Likelihood ratio detection and tracking recursion

One can extend the Bayes Markov recursion for a single target given above to likelihood ratio detection and tracking by adding the state variable ϕ to the state space S. Having calculated the posterior distribution (density) on S^+, one then uses (10) to compute the posterior target likelihood ratio. Instead of doing this, we generally make the assumption below which allows us to use a simplified version of the recursion. Specifically, we assume that

$$\begin{aligned} p^-(t_k, \phi \mid \mathbf{Y}(t_{k-1})) &= q_k(\phi \mid \phi) p(t_{k-1}, \phi \mid \mathbf{Y}(t_{k-1})) + \int_S q_k(\phi \mid s) p(t_{k-1}, s \mid \mathbf{Y}(t_{k-1})) ds \\ &= p(t_{k-1}, \phi \mid \mathbf{Y}(t_{k-1})). \end{aligned} \tag{12}$$

This assumption says that under the motion model, the amount of probability mass that moves out of the region $\mathcal{R}$ is equal to the amount that moves in for a given time period. Under this assumption we obtain the following simplified likelihood ratio detection and tracking recursion.

Simplified Likelihood Ratio Detection and Tracking Recursion

Initialize Likelihood Ratio $$\Lambda(0, s \mid \mathbf{Y}(0)) = \frac{p_0(s)}{p_0(\phi)} \text{ for } s \in S \tag{13}$$

For $k \geq 1$ and $s \in S$

Motion Update $$\Lambda^-(t_k, s \mid \mathbf{Y}(t_{k-1})) = q_k(s \mid \phi) + \int_S q_k(s \mid s_{k-1}) \Lambda(t_{k-1}, s_{k-1} \mid \mathbf{Y}(t_{k-1})) ds_{k-1} \tag{14}$$

Measurement Likelihood Ratio $$\mathcal{L}_k(y_k \mid s) = \frac{L_k(y_k \mid s)}{L_k(y_k \mid \phi)} \text{ for } s \in S \tag{15}$$

Information Update $$\Lambda(t_k, s \mid \mathbf{Y}(t_k)) = \mathcal{L}_k(y_k \mid s) \Lambda^-(t_k, s \mid \mathbf{Y}(t_{k-1})) \text{ for } s \in S . \tag{16}$$

3.2.2 Extension of LRT state space

Ordinarily LRT uses the target's kinematic variables for the tracker state space. For example, the state is often taken to be the target's 2 (or 3) dimensional position and velocity. For this analysis, we have extended the kinematic state space (2 dimensional position and velocity) to include mean signal excess prediction error. As sensor responses are obtained, LRT produces a joint estimate of target kinematic state and $\overline{SE}$ prediction error.

3.2.3 Numerical implementation of LRT

The implementation of LRT used to produce the results in the examples below employs a discrete grid in position, velocity, and mean signal excess prediction error. The methodology is similar to the gridded version of Nodestar described in Chapter 3 of Stone et al.[1]

4. MSA MEASUREMENTS AND LIKELIHOOD RATIO FUNCTIONS

The examples presented in this paper involve multi-static active (MSA) sonar. There are a set of 10 to 30 buoy pairs. Each pair consists of a receiver buoy and a source buoy with two explosive charges. These buoy pairs are distributed over an area where a target submarine may be present. The charges on the source buoys are detonated sequentially, one every few minutes. Between detonations, the hydrophones in the receiver buoys listen for echoes of the shockwave as it scatters off objects, potentially including a moving target submarine. The time between the reception of the direct blast and an echo produces an ellipse of possible locations for the echo producer. By accumulating a number of these echoes from the target, it is possible to call a detection and develop a track on the target, i.e., a distribution on the target's position and velocity.

The signal processing algorithms associated with these buoys are presumed to process the time-series at the receiver hydrophones and call detections. These algorithms produce a set of time values for each hydrophone where the signal or matched-filter output exceeds some threshold. Each of these "detections" comes from one of three things: random stochastic fluctuations in the noise signal (false alarms), clutter echo, or a target echo. These are the measurements used by LRT. The algorithm begins with a prior over a gridded state space with position and velocity and $\overline{SE}$ error as dimensions. As sensors report information — in this case the series of detection times from the signal processor — their information content is incorporated into the posterior likelihood ratio on state space by the use of a measurement likelihood ratio function. Between sensor measurements, the target likelihood ratio surface evolves according to a motion model prescribed for the target. If a target is present, then we expect to see a peak form in its vicinity as time progresses and the measurement likelihood ratio peaks in the vicinity of the target reinforce one another. In order to process these measurements, we construct a measurement likelihood ratio function for use by LRT.

4.1 Measurement likelihood ratio function – known probability of detection

To simplify our presentation, we first introduce the measurement likelihood ratio function for MSA sonar in the case where the detection probability is known. For a single ping, let $\boldsymbol{\tau}_l = \{\tau_1, \ldots, \tau_n\}$ be the set of echo times detected at buoy l. The measurement likelihood ratio for $\boldsymbol{\tau}_l$ is

$$\begin{aligned} \mathcal{L}_l(\boldsymbol{\tau}_l \mid s) &\equiv \frac{\mathbf{Pr}\{\text{Obtaining the set of echo times } \boldsymbol{\tau}_l \text{ at buoy } l \mid \text{target in state } s\}}{\mathbf{Pr}\{\text{Obtaining the set of echo times } \boldsymbol{\tau}_l \text{ at buoy } l \mid \text{no target present}\}} \\ &= P_d^{\,l}(s) \sum_{i=1}^{n} \frac{1}{\lambda} \left(\frac{f_l(\tau_i \mid s)}{w(\tau_i)} \right) + \left(1 - P_d^{\,l}(s)\right) \end{aligned} \tag{17}$$

where

$s =$ target state (x, v), position and velocity

$f_l(\tau \mid s) = \mathbf{Pr}\{\text{receiving an echo at buoy } l \text{ at time } \tau \mid \text{target detected in state } s\}$

$w(\tau) = \mathbf{Pr}\{\text{receiving an echo at time } \tau \mid \text{ echo generated by false alarm}\}$

$\lambda =$ mean number of false alarms per ping per buoy

$P_d^{\,l}(s) = \mathbf{Pr}\{\text{buoy } l \text{ detects a target in state } s \text{ on a single ping}\}.$

The derivation of (17) is given in the appendix. To form the composite measurement likelihood ratio function for a single ping, we multiply the likelihood ratio for each of the buoys as follows. Let $\boldsymbol{\tau} = (\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_{N_r})$. Then

$$\mathcal{L}_{\mathbf{C}}(\boldsymbol{\tau} \mid s) = \prod_{l=1}^{N_r} \mathcal{L}_l(\boldsymbol{\tau}_l \mid s) \tag{18}$$

where N_r is the number of receiver buoys. This measurement likelihood ratio is multiplied into the motion updated cumulative likelihood ratio to compute the posterior cumulative likelihood ratio in the information update step in (16). This process is repeated for each ping.

The measurement density function $f_l(\tau \mid s)$ represents the "ellipse" information from the time of detection. Let x_s^m be the position of source m and x_r^l be the position of the receiver buoy l. Then the measurement error density function for the arrival time τ from source m is

$$f_l\left(\tau \mid s=(x,v)\right)=\frac{1}{\sqrt{2\pi}\sigma_t}\exp\left\{\frac{\left(c_s\tau-d(x_s^m,x)-d(x,x_r^l)\right)^2}{c_s^2 2\sigma_t^2}\right\}$$

where $d(x_1,x_2)$ is the distance between positions x_1 and x_2, and c_s is the sound speed. We use σ_t = 1 second. This form of the measurement error density function can also be used to account for uncertainty in buoy position to first order, as we have shown elsewhere. The false alarm density function w is uniform over a 90 second interval.

The $P_d^{\,l}(s)$ factor is where we require a performance prediction model for the sensors, and this is where environmental information is incorporated into the LRT algorithm. To compute $P_d^{\,l}(s)$ for a given source m, we begin by computing the mean signal excess $\overline{SE}$ in dB at the receiver buoy l as follows.

$$\overline{SE}_{lm}(s)=SL_m-TL_m(s)+TS_{lm}(s)-TL_l(s)-(RL_{lm}(s)+AN)-DT \tag{19}$$

where

- $SL_m =$ source level at source m
- $TL_m(s) =$ transmission loss from source m to state s
- $TS_{lm}(s) =$ target strength for a signal from source m reflected from a target in state s to sensor l
- $TL_l(s) =$ transmission loss from state s to sensor l
- $RL_{lm}(s) =$ reverberation at sensor l from source m at the time an echo arrives from state s
- $AN =$ ambient noise level
- $DT =$ detection threshold.

The value of $\overline{SE}_{lm}$ given by equation (19) represents the mean signal excess level. The actual value of signal excess $\widetilde{SE}$ is a random variable with fluctuations about this mean that are normally distributed with mean 0 and standard deviation σ_{SE} generally taken to be between 5 to 10 dB. Detection occurs when $\widetilde{SE}>0$. Having calculated $\overline{SE}_{lm}(s)$ from (19), we can compute $P_d^{\,k}(s)$ as follows

$$P_d^{\,k}(s)=\int_0^\infty N(x,\overline{SE}(s),\sigma_{SE}^2)\,dx$$

where $N(\cdot,\mu,\sigma^2)$ is the density function for a normal distribution with mean μ and variance σ^2.

4.2 Incorporating prediction uncertainty

The procedure for computing the likelihood function described in the previous section assumes we know the terms in equation (19) with certainty. When there is uncertainty in these terms, we must modify our procedure. One method for doing this is to extend the LRT state space to include an "environmental-uncertainty" dimension where each value in this dimension represents one possible environment, E. The likelihood function then becomes

$$\mathcal{L}_l\left(\boldsymbol{\tau} \mid s,E\right)=P_d^{\,l}(s,E)\sum_{i=1}^{n}\frac{1}{\lambda}\left(\frac{f_l(\tau_i \mid s)}{w(\tau_i)}\right)+\left(1-P_d^{\,l}(s,E)\right),$$

where we now have a likelihood function that is defined not only over the target's kinematic state space, but also over the extra environmental dimension. The values of $E\in\{E_1,E_2,\ldots\}$ represent all possible environments. For the examples that we consider below, each environment consists essentially of a reverberation and transmission loss characterization. Furthermore, as we are using range independent environments for this work, we have $RL\equiv RL(d,t)$ and $TL\equiv TL(d)$: reverberation is a function of the distance between source and receiver and time since blast while transmission loss is a function of distance only. Thus, $E_i\equiv\{RL^i,TL^i\}$. The Applied Research Laboratory, University of Texas, and the Applied Physics Laboratory at University of Washington have computed these functions for us from more basic environmental parameters such as bottom type and sound speed profiles appropriate for a specific ocean area.

The number of environments needed to capture the full range of uncertainty is infinite because there is a continuum of uncertainties. For LRT, we approximate this continuous distribution of uncertainties by choosing a discrete set of environments $\{\hat{E}_1, \hat{E}_2, ..., \hat{E}_N\}$ to include in the LRT state space. Our first approach has been to form the mean over all environments provided to us, and let each possible environment be a constant offset from this mean. Thus, we compute:

$$\overline{RL}(D,t) = \frac{1}{N}\sum_{i=1}^{N} RL^i(D,t)$$

and

$$\overline{TL}(d) = \frac{1}{N}\sum_{i=1}^{N} TL^i(d).$$

Each environment is then characterized by these mean curves and a constant offset δ_i: $\hat{E}_i \equiv \{\overline{RL}, \overline{TL}, \delta_i\}$. For the i^{th} environment we compute mean signal excess by

$$\overline{SE}_{lm}(s) - \delta_i$$

where $\overline{SE}_{lm}(s)$ is computed by equation (19). We can think of δ as a random variable that equals the (unknown) constant error in the mean signal excess computation in equation (19). In this case the measurement likelihood ratio function in (17) becomes

$$\mathcal{L}_l(\boldsymbol{\tau} \mid s, \delta) = P_d^{\,l}(s,\delta)\sum_{i=1}^{n}\frac{1}{\lambda}\left(\frac{f_l(\tau_i \mid s)}{w(\tau_i)}\right) + \left(1 - P_d^{\,l}(s,\delta)\right) \tag{20}$$

where

$$P_d^{\,l}(s,\delta) = \int_0^{\infty} N(x, \overline{SE}_{lm}(s) - \delta, \sigma_{SE}^2)\, dx\,.$$

The effect of adding the variable E (or δ) to the state space is that LRT is calculating a joint distribution on kinematic state and δ. We expect to see a peak in the cumulative likelihood ratio surface for the environment with δ_i closest to the actual error. The astute reader will note that appropriate value of δ_i may change as the target moves through the state space. The hope is that the target moves slowly enough that the appropriate value of δ_i changes slowly. We will address this issue in future studies.

4.3 Measurement likelihood ratio without performance prediction

In previous Metron work using LRT for MSA sonar, we have used a simple empirical model. In this model, we set a maximum probability of detection at broadside aspect for the target to the sensor and let the probability fall off according to a Gaussian density in bistatic aspect angle where σ_φ is the standard deviation of the density. Let φ_0^l be the bistatic aspect angle with maximum response for sensor l and $\varphi^l(s)$ be the bistatic angle for a target in state s. This will depend on the location of the source buoy producing the ping as well as the receiver buoy. The functional form of $P_d^{\,l}(s)$ is

$$P_d^{\,l}(s) = P_d^{\text{MAX}}(s)\exp\left\{-\frac{\left(\varphi_0^l - \varphi^l(s)\right)^2}{2\sigma_\varphi^{\,2}}\right\}.$$

The likelihood function in this case becomes

$$\mathcal{L}_l(\boldsymbol{\tau} \mid s) = P_d^{\text{MAX}}(s)\exp\left\{-\frac{\left(\varphi_0^l - \varphi^l(s)\right)^2}{2\sigma_\varphi^{\,2}}\right\}\sum_{i=1}^{n}\frac{1}{\lambda}\left(\frac{f_l(\tau_i \mid s)}{w(\tau_i)}\right) + 1 - P_d^{\text{MAX}}(s)\exp\left\{-\frac{\left(\varphi_0^l - \varphi^l(s)\right)^2}{2\sigma_\varphi^{\,2}}\right\}. \tag{21}$$

Equation (21) assumes that we know $P_d^{\text{MAX}}(s)$. When there is substantial uncertainty in estimating this probability, we assume it is uniformly distributed on [0,1]. Integrating the likelihood function in (21) over this distribution on $P_d^{\text{MAX}}(s)$ yields

$$\mathcal{L}_l(\boldsymbol{\tau}\,|\,s)=\frac{1}{2}\exp\left\{-\frac{\left(\varphi_0^l-\varphi^l(s)\right)^2}{2\sigma_\varphi^{\ 2}}\right\}\sum_{i=1}^{n}\frac{1}{\lambda}\left(\frac{f_l(\tau_i\,|\,s)}{w(\tau_i)}\right)+1-\frac{1}{2}\exp\left\{-\frac{\left(\varphi_0^l-\varphi^l(s)\right)^2}{2\sigma_\varphi^{\ 2}}\right\}. \tag{22}$$

The value of σ_φ is set to approximately 15 degrees. Equation (22) gives the likelihood function used in the previous MSA sonar versions of LRT. We call this the *baseline* case. It has a couple of obvious drawbacks. The relationship between detection probability and aspect angle has been estimated in an empirical fashion. One does not know how well this prediction will hold up outside of the situations which produced the data used to form this prediction. Second the likelihood function clearly does not account for range. According to the above likelihood function, targets at very long ranges from the receiver are as likely to produce detections as targets that are close in range. Furthermore, by integrating over a uniform $[0,1]$ distribution on probability of detection, we have in effect assumed that $P_d^{\mathrm{MAX}}(s)$ = 0.5. This will reduce the likelihood ratio for contacts obtained from targets with an average $P_d^{\mathrm{MAX}}(s)$ < 0.5. One approach is to set $P_d^{\mathrm{MAX}}(s)$ equal to some minimal detectable level (i.e. the smallest P_d for which the system is expected to generate a detection). In this case we sacrifice some performance on stronger targets so that we can increase the chance of detecting weaker ones. This is the *minimum detectable level* likelihood ratio function. Section 6.3.4 of Stone et al[1] discusses this approach.

5. SIMULATION RESULTS

We present two comparisons of the performance of LRT with and without performance prediction. In both comparisons, the version of LRT with performance prediction performs better than the one without.

5.1 Baseline versus large environmental uncertainty

For this case, we compare the output of LRT using the expanded LRT state space which includes environmental uncertainty and the likelihood function in (20) with the *baseline* case using the likelihood function in (22) and the standard kinematic LRT state space. For the expanded state space, we have assumed that the prior distribution on δ, the error in mean signal excess prediction is uniform between -30 and 30 dB. This is a case of great uncertainty in the environmental and performance predictions. The results below show that using environmental information, even with great uncertainty, improves tracker performance in the simulated situation considered here.

Researchers at the Applied Physics Laboratory at the University of Washington provided environmental predictions for two bottom types: low-frequency bottom loss which they designated LFBL and Fulford. Both bottom types, and therefore the resulting RL/TL curves, are assumed to be equally likely. We have used the Fulford bottom type to produce the simulated detection data used by the tracker. Thus this becomes the true bottom type. This is not known by the tracker. The conditions of the environment are otherwise typical of the East China Sea. The bottom depth is 100m, and the sound speed profile is range independent.

The target in these simulations moves due east parallel to and 15 nm north of a buoy field. The probability of detection for the target for any given source-receiver pair for each blast is between 0 and 0.1. Figures 1 – 3 below compare the performance of these two methods at low and medium false alarm rates. The average number of false alarms per buoy per ping is 1 for the low rate and 20 for the medium rate.

In Figures 1 – 3 below, the simulation has been run for about 2 hours using 72 pings spaced 100 seconds apart. The figures show LRT output at the end of the simulation period. Figure 1 shows the output of the version of LRT with performance prediction for the low false alarm case. LRT has detected the target and correctly identified the performance prediction error (which is between 10 and 15 dB). Figure 2 compares this to output of LRT with no performance prediction, in particular using the baseline likelihood function. One can see that LRT without performance prediction (baseline) fails to detect the target. When we have a medium false alarm rate (20 per receiver per ping), Figure 3 shows the baseline method again fails to detect the target, but adding an environmental uncertainty dimension to the state space and using performance prediction allows LRT to detect the target. We also ran the simulation using a high false alarm rate case (~100 per receiver per ping). In this case neither method detected the target.

5.2 Minimum detectable level versus small environmental uncertainty

For this case, we compare the output of LRT using the expanded state space with small environmental uncertainty and LRT with the standard kinematic state space using the likelihood function in (21) with $P_d^{\mathrm{MAX}}(s)$ = 0.05, a reasonable minimal detectable level and the actual mean Pd of the target in this exercise. For the expanded state space, we have assumed that the prior distribution on δ, the error in mean signal excess prediction, is uniform between -5 and 5 dB. This is a case of low uncertainty in the environmental and performance predictions.

We have used the Fulford bottom type environmental model to produce the simulated detection data used by the tracker. Thus this becomes the true bottom type. This is not known by the tracker. The tracker uses the Fulford model as the mean curve with a set of δ values that range from -5 dB to 5 dB. The conditions of the environment are otherwise typical of the East China Sea. The bottom depth is 100 m, and the sound speed profile is range independent.

The target in these simulations moves east parallel to and 10 nm north of a buoy field. The probability of detection for the target for any given source-receiver pair for each blast is between 0 and 0.2. Figures 4 – 5 show the cumulative LRT surfaces after 48 pings spaced 100 seconds apart. They compare the performance of the two methods.

In Figure 4 where the false alarm rate is low, the target shows a much sharper and higher peak in the cumulative likelihood ratio surface of the LRT with performance prediction and low uncertainty (right). In Figure 5, we see the same comparison, here with a medium false alarm rate. In this case, both methods would have tracked the target; however, the peak likelihood ratio with performance predication and low uncertainty (right) is higher than with the minimum detectable level method (left), and the target would have been detected sooner.

6. CONCLUSIONS

This work shows that using performance prediction in a manner that accounts for the uncertainty in the predications improves the performance of LRT compared to using likelihood functions that do not use performance prediction. Note that in order to compute a likelihood ratio we have to assume some probability of detection. In the IASW baseline case we have in effect used a probability of detection equal to 0.5 everywhere in the position component of the state space. While we are using a probability of detection, we are not using a performance prediction in a meaningful sense, only in the trivial sense of picking a single probability of detection to be used independent of target state.

When we do have a good characterization of the environment with less uncertainty, the tactical picture produced by the tracker should reflect this, as it does in the comparison shown in case 5.2 above. As an example, if we know that the signal excess falls off rapidly after some range, the LRT tracker will reject detections with large ranges. In the future, we would like to examine alternate representations of the environmental dimension. In some cases it may be possible to accurately represent the full range of environments with a manageable number of RL/TL curves, allowing us to reap the benefits of the frame-to-frame and sensor-to-sensor correlations in target strength demonstrated in these examples.

REFERENCES

1. L. D. Stone, C. A. Barlow, and T. L. Corwin, *Bayesian Multiple Target Tracking*, Artech, Boston, 1999.
2. R. J. Urick, *Principles of Underwater Sound*, 3rd Ed, McGraw-Hill, New York, 1983.

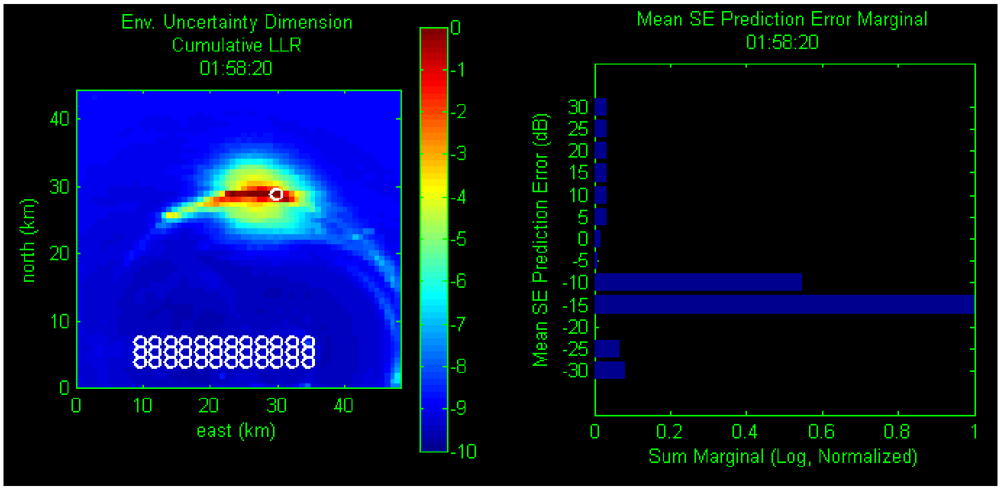

Figure 1. Low false alarm rate: On the left is the final log-likelihood ratio position marginal for the extended state space. On the right is the marginal distribution on δ . The peak has converged on the correct value of delta. The false alarm rate is 1\ping\receiver.

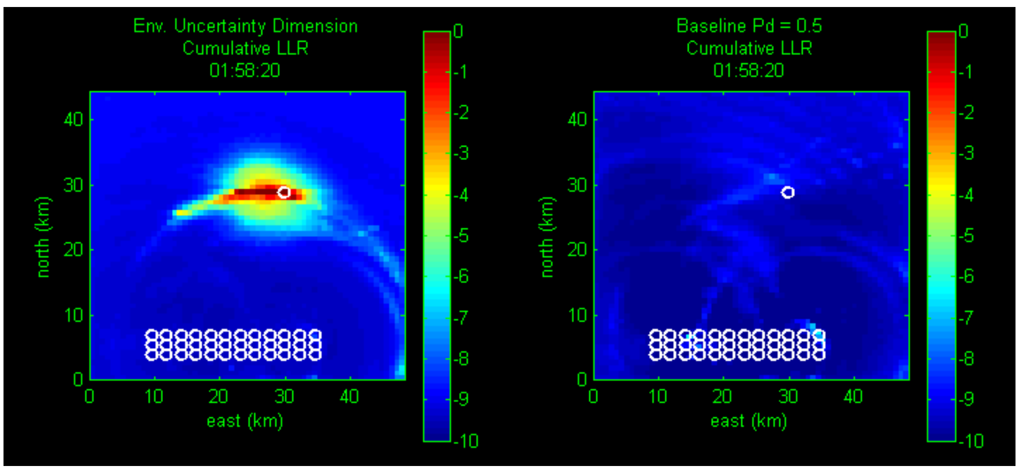

Figure 2. Low false alarm rate: Final log-likelihood ratio position marginal (false alarm rate is 1\ping\receiver). The version with performance prediction detected the target (left) while the baseline version (right) did not.

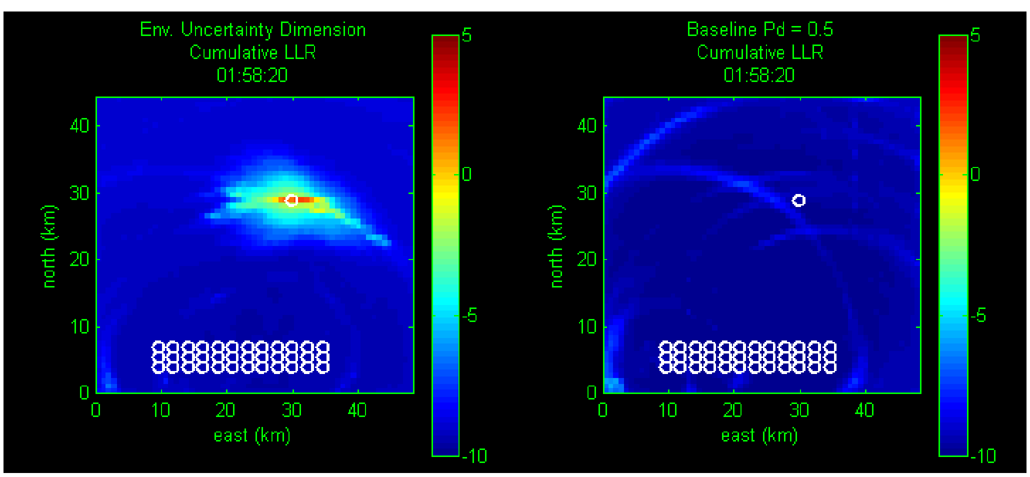

Figure 3. Medium false alarm rate: Final log-likelihood ratio position marginal for the two methods where the false alarm rate is 20 per ping per receiver. The version with performance prediction (left) detects the target while the baseline version (right) does not**.**

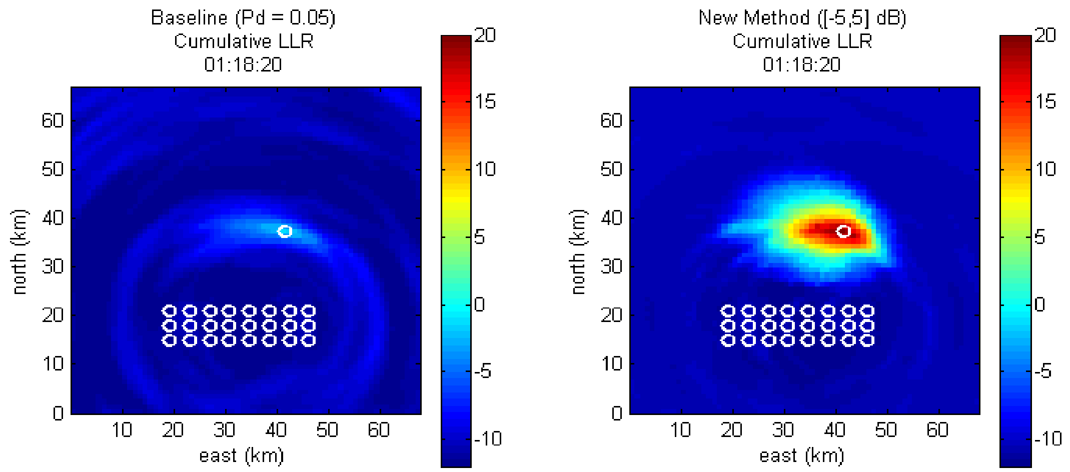

Figure 4: Low false alarm rate. Cumulative log-likelihood ratio surface for low false alarm rate case (15 per buoy per ping). On the left is the position marginal of the minimum detectable level case with Pd = 0.05. On the right is the surface for the new method which uses performance prediction and an expanded state space.

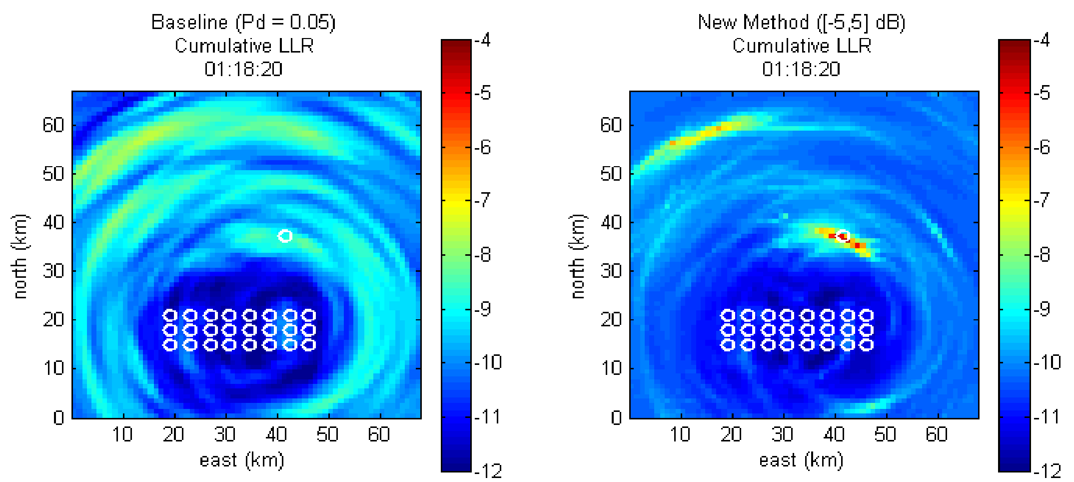

Figure 5: Medium false alarm rate. Cumulative log-likelihood ratio surface for medium false alarm rate case (25 per buoy per ping). On the left is the position marginal of the minimum detectable level case with Pd = 0.05. On the right is the surface for the new method which uses performance prediction and an expanded state space.

APPENDIX

This appendix provides the derivation of the measurement likelihood ratio function in (17).

Let $\mathbf{y}_l = (y_1, \dots y_n)$ be the set measurements received at a single buoy from ping k corresponding to time t_k. These measurements can be arrival times of echoes or pairs of arrival times and estimates of bearing to the target. We wish to compute the measurement likelihood ratio

$$\mathcal{L}_l\left(\mathbf{y}_l \middle| s\right) = \frac{L_l\left(\mathbf{y}_l \middle| s\right)}{L_l\left(\mathbf{y}_l \middle| \phi\right)} \text{ for } s \in S$$

which is the ratio of the likelihood of receiving the observation $\mathbf{y}_l = (y_1, \dots y_n)$ given the target is in state s to likelihood of receiving these measurements given there is no target present. The measurement likelihood ratios from different buoys are be multiplied together to form the composite measurement likelihood ratio for ping k.

As above, we define the probability of detection by $P_d^l(s) = \mathbf{Pr}\{\text{target is detected by buoy } l \mid X(t_k) = s\}$. This function captures the bistatic scattering geometry. It will depend on the positions of the source and receiver as well as the heading and position of the target. We define the probability density function of an observation to be

$$f(y|s) = \mathbf{Pr}\left\{\text{measurement from target} = y \middle| X(t_k) = s\right\}.$$

For the false-alarm model, let Φ_k be the random variable which equals the number of false alarms received at the buoy from the k th ping. In addition, let

$$w_n(y_1,\ldots,y_n) = \mathbf{Pr}\left\{\text{obtaining false alarm measurements } (y_1,\ldots,y_n) \middle| \Phi_k = n\right\}.$$

Using these definitions, we compute

$$\begin{aligned}
L_l(\mathbf{y}_l|s) &= P_d^l(s)\mathbf{Pr}\left\{\mathbf{y}_l \mid X(t_k) = s \,\&\, \text{ target detected}\right\} + \left(1 - P_d^l(s)\right)\mathbf{Pr}\left\{\mathbf{y}_l \middle| X(t_k) = s \,\&\, \text{ target not detected}\right\} \\
&= P_d^l(s)\sum_{i=1}^{n}\mathbf{Pr}\left\{y_i \middle| X(t_k) = s \,\&\, \text{ target produces measurement } i\right\} \times \mathbf{Pr}\left\{y_i \text{ is produced by the target}\right\} \\
&\quad + \left(1 - P_d(s)\right)\mathbf{Pr}\left\{\mathbf{y}_l \middle| X(t_k) = s \,\&\, \text{ target not detected}\right\} \\
&= P_d^l(s)\sum_{i=1}^{n}\frac{1}{n} f(y_i|s)\mathbf{Pr}\left\{\Phi_k = n-1\right\} w_{n-1}(y_1,\ldots,\hat{y}_i,\ldots,y_n) + \left(1 - P_d^l(s)\right)\mathbf{Pr}\left\{\Phi_k = n\right\} w_n(y_1,\ldots,y_n).
\end{aligned}$$

Note $(y_1,\ldots,\hat{y}_i,\ldots,y_n)$ denotes the $n-1$ vector obtained from $(y_1,\ldots,y_n)$ by deleting the i th component y_i. We also have from the definitions that

$$L_l(\mathbf{y}_l|\phi) = \mathbf{Pr}\left\{\Phi_k = n\right\} w_n(y_1,\ldots,y_n).$$

Dividing these two equations we obtain

$$\mathcal{L}_k(\mathbf{y}_k \mid s) = P_d^l(s)\sum_{i=1}^{n}\frac{1}{n}\left(\frac{\mathbf{Pr}\left\{\Phi_k = n-1\right\}}{\mathbf{Pr}\left\{\Phi_k = n\right\}}\right)\left(\frac{f(y_i|s)\, w_{n-1}(y_1,\ldots,\hat{y}_i,\ldots,y_n)}{w_n(y_1,\ldots,y_n)}\right) + \left(1 - P_d^l(s)\right). \tag{23}$$

Assuming that the false measurements are independent and identically distributed, i.e.,

$$w_n(y_1,\ldots,y_n) = \prod_{i=1}^{n} w(y_i),$$

we have

$$\mathcal{L}_l(\mathbf{y}_l|s) = P_d^l(s)\sum_{i=1}^{n}\frac{1}{n}\left(\frac{\mathbf{Pr}\left\{\Phi_k = n-1\right\}}{\mathbf{Pr}\left\{\Phi_k = n\right\}}\right)\left(\frac{f(y_i|s)}{w(y_i)}\right) + \left(1 - P_d^l(s)\right).$$

Now assume that the number of false alarms at a buoy for a single ping is Poisson distributed with mean λ, that is

$$\mathbf{Pr}\left\{\Phi_k = j\right\} = \frac{\lambda^j}{j!}\exp(-\lambda).$$

Then we have

$$\mathcal{L}_l(\mathbf{y}_l|s) = P_d^l(s)\sum_{i=1}^{n}\frac{1}{\lambda}\left(\frac{f(y_i|s)}{w(y_i)}\right) + \left(1 - P_d^l(s)\right). \tag{24}$$

When $y_i = \tau_i$, (24) specializes to (17).