Active Clutter Reduction Through Fusion with Passive Data

Jason M. Aughenbaugh, **Bryan A. Yocom**, and **Brian R. La Cour**
Applied Research Laboratories
The University of Texas at Austin
P.O. Box 8029, Austin, Texas 78713-8029 USA
{jason, byocom, blacour}@arlut.utexas.edu

Abstract – *Active and passive sonar systems often provide complementary information. Active systems generally provide good localization but are subject to high rates of clutter from bottom and volumetric scatterers. Passive systems generally have good resolution only in bearing and only can detect objects that emit acoustic energy. By fusing measurements from both types of systems, the passive data can be used to reduce the clutter in the active data. The result is improved target tracking, as shown in two examples, including the injection of a single target into a recorded active clutter dataset.*

Keywords: Bayesian tracking, passive sonar, active sonar, clutter reduction, data fusion

1 Introduction

Historically, passive sonar systems and active sonar systems have been used separately. Passive systems generally excel in long-range, low-frequency detection and classification, but are subject to false alarms from far away, loud interferers. Active systems tend to be used for shorter range detection and localization, but are subject to local clutter from bathymetric features and other scatterers. By fusing active and passive data in a single detection and localization system, one should realize overall gains in performance due to the complementarity of the information they provide.

For example, an undersea ridge will lead to significant clutter in an active system. However, an undersea ridge does not emit its own acoustic energy, and is therefore not detectable on a passive system. Consequently, the fusion of the active and passive data should decrease the probability of a false alarm on the ridge. Similarly, one may only be concerned with a target of interest (TOI) that emits acoustical energy, and therefore can be expected to be detectable on both active systems and passive systems. The result of the fused process should be a detection on the true TOI with fewer false alarms (or false tracks or errors in localization, depending on the tracker used) than the active system alone. Additionally, the incorporation of the active data into the passive data should greatly increase the localization of the target, which, for example, may only be known in bearing on the passive system.

Bayesian inference is recognized as a general framework for performing optimal target tracking. Fundamentally, it assumes that the uncertainty in our knowledge of the state of the target (or targets) may be represented well by probabilities. Bayes' theorem then provides the basic mechanism whereby measurements update these probabilities and, hence, our state of knowledge of the world. The probabilities serve as a common interface for different sensors, and thus the problem of data fusion becomes one of modeling the probabilities for different sensors, as well as a computational issue of how to implement the filter.

In this paper, we demonstrate that fusion of passive sonar data with active sonar data can reduce clutter and improve tracker performance. In Sec. 2, we provide an overview of the Bayesian fusion of active and passive sonar data in a variation of a track-before-detect system. In Sec. 3, we consider a simulated example using data generated with the Sonar Simulation Toolset (SST) [1, 2]. In Sec. 4, the example involves real world recorded active clutter data, into which we have injected targets with appropriate active echoes and radiated noise using SST. A brief summary is given in Sec. 5.

2 Background

There has been extensive work on representations of Bayesian filters, including variations of Kalman filters [3], grid-based methods [4], and particle filters [5]. In this paper, we adopt a grid-based implementation of the filter, based on details in an article that is under review [6] and summarized here.

2.1 Grid-based Bayesian tracker introduction

Grid-based Bayesian updating and tracking has a long history, including recent advances [4, 7, 8, 9, 10, 11, 12]. In a grid-based Bayesian filter, the state is discretized into a multi-dimensional grid. If the state space is inherently discrete, then a grid is an optimal representation. More commonly, the state space is continuous (or a mixture of continuous and discrete states). In this case, the grid represents a piece-wise constant approximation of the posterior density.

We consider a Bayesian tracking scheme for which the state space consists of the number of targets present (either zero or one in this paper) and the target's kinematic state, $\boldsymbol{s}$, given that it is present. Let p_n be the probability that a single target is present in a particular region of interest, and let $\rho_n(\boldsymbol{s})$ be the posterior probability density function after the n^{th} measurement is incorporated, where $n = 0$ corresponds to the prior distribution.

Two likelihood functions are relevant. The target likelihood function $L_n(\boldsymbol{y_n}|\boldsymbol{s})$ denotes the likelihood of receiving measurement $\boldsymbol{y_n}$ given that the target is in state $\boldsymbol{s}$. The no-target likelihood function $L_n(\boldsymbol{y_n}|\varnothing)$ denotes the likelihood of receiving measurement $\boldsymbol{y_n}$ when there is no target present. Measurement updates on the kinematic PDF are performed using Bayes' theorem, which for notational purposes we write as

$$\rho_n(\boldsymbol{s}) = L_n(\boldsymbol{y_n}|\boldsymbol{s})\rho_n^-(\boldsymbol{s})/E_n, \tag{1}$$

where ρ_n^- is the motion updated prior distribution and E_n is the partial Bayesian Evidence [13] for a target being present , where E_n is given as

$$E_n := \int L_n(\boldsymbol{y_n}|\boldsymbol{s})\rho_n^-(\boldsymbol{s})\mathrm{d}\boldsymbol{s}. \tag{2}$$

Similarly, the target probability is updated by

$$p_n = \frac{E_n p_n^-}{(1-p_n^-)L_n(\boldsymbol{y_n}|\varnothing) + p_n^- E_n}, \tag{3}$$

where p_n^- is the motion updated target probability [14].

The target is assumed to follow an Integrated Ornstein-Uhlenbeck (IOU) process [4]. The kinematic update of the IOU process is implemented using a particle-filter like model in which each grid cell is sampled with particles, and then the IOU model is used as in [5] to update the particle cloud for each cell. Then, the relocated particles and their associated probability weights are aggregated back into the grid representation. Additional details are provided in [6]. Additionally, a birth-death process in which targets can appear inside the state space or exit the state space is implemented as described in [15].

2.2 Active sonar modeling

The crucial link between the actual sonar systems and the Bayesian filter is the definition of the likelihood functions. We choose a model that is based closely on the actual signal propagation process, as described in more detail in [11]. We define the likelihood model in terms of measured signal-to-noise ratio (SNR) values from the normalized matched filter output of a standard active signal processing chain. Both frequency modulated (FM) and continuous wave (CW) transmit waveforms can be considered, but for this paper we limit ourselves to FM data. In general, each SNR value (in units of squared amplitude) z_k is associated with a particular echo time of arrival (TOA) τ_k and angle of arrival (AOA) ϕ_k. Each measurement $\boldsymbol{y}$ consists of K returns y_k such that $\boldsymbol{y} = (y_1, \ldots, y_K)$, where for an FM waveform $y_k = (z_k, \tau_k, \phi_k)$.

2.2.1 Basic active likelihood model

The SNR values are assumed to be independent and follow a generalized Pareto distribution (GPD), so the likelihood function for the overall measurement $\boldsymbol{y}$ is

$$L(\boldsymbol{y}|\boldsymbol{s}) = \prod_{k=1}^{K} \frac{1}{\mu_k(\boldsymbol{s})}\left(1 + \xi\frac{z_k}{\mu_k(\boldsymbol{s})}\right)^{-(1+1/\xi)}, \tag{4}$$

where ξ is the GPD shape parameter, $\mu_k(\boldsymbol{s})$ represents the mean SNR (see Sec. 2.2.2) that one would expect to receive from measurement element y_k, given that the target state is $\boldsymbol{s}$. The GPD generalizes the exponential distribution presented in [11] and allows one to account for fatter tails in the observed normalizer output.

We choose to use thresholded measurements in order to reduce the false alarm rate and computational cost. Accordingly, only individual elements with SNRs that exceed a set threshold η are incorporated in to the measurement. Let $\boldsymbol{k} = (k_1, \ldots, k_I)$ denote an ordered sequence of indices corresponding to these threshold crossings. If there are no threshold crossings, then $\boldsymbol{k} = \varnothing$. For such measurements the likelihood is (see [11] for derivation)

$$L(\boldsymbol{k}, \boldsymbol{y}|\boldsymbol{s}) = \prod_{k=1}^{K} F_k(\eta|\boldsymbol{s}) \prod_{i=1}^{I} \frac{1 - F_i(\eta|\boldsymbol{s})}{F_i(\eta|\boldsymbol{s})} f_i(z_i|\boldsymbol{s}, \eta), \tag{5}$$

where $f_i(z_i|\boldsymbol{s}, \eta)$ and $F_i(\eta|\boldsymbol{s})$ are respectively the conditional probability density and the cumulative probability functions of the appropriate GPD.

2.2.2 Active SNR predictive modeling

The mean received SNR in units of amplitude squared at a point k is a sum of the uniform background clutter σ_0^2 and the energy provided by a target with nominal target strength of σ_{T}^2 in state $\boldsymbol{s}$, such that

$$\mu_k(\boldsymbol{s}) = \mu_k(\varnothing) + \sigma_{\mathrm{T}}^2 h_k(\boldsymbol{s}), \tag{6}$$

where $h_k(\boldsymbol{s}) \in [0, 1]$ is the response function, specifically the product of the array beam response $b(\phi)$ and waveform ambiguity function $|\chi(\tau, \nu)|^2$ for a target in state $\boldsymbol{s}$. The beam response, as approximated here, is given by

$$b(\phi) = \mathrm{sinc}^2(\phi/\Delta\phi), \tag{7}$$

where $\phi = \phi(\boldsymbol{s}) - \phi_k$, $\phi(\boldsymbol{s})$ is the bearing to the hypothesized target state, and $\Delta\phi$ is the nominal beam width. One could use a more detailed model when necessary.

For a linear FM waveform, the ambiguity function for $\tau \leq T$ is

$$\chi(\tau, \nu) = \left(1 - \frac{|\tau|}{T}\right) \mathrm{sinc}\left[(\nu - \tau B/T)\, T\left(1 - \frac{|\tau|}{T}\right)\right], \tag{8}$$

where T is the pulse length, B is the bandwidth, and $\tau = \tau(\boldsymbol{s}) - \tau_k$, $\nu = \nu(\boldsymbol{s}) - \nu_k$, and $\tau(\boldsymbol{s})$ and $\nu(\boldsymbol{s})$ are the TOA and

Doppler shift corresponding to the hypothesized target state $\boldsymbol{s}$ [16]. For $\tau > T$, $\chi(\tau, \nu) = 0$. The predictive SNR modeling also accounts for detection ranges and blanking regions.

2.3 Passive sonar modeling

Assume there is one narrow-band plane-wave signal impinging on an N-element uniformly-spaced horizontal line array (HLA). The AOA of the signal, denoted by ϕ, is measured relative to the line of bearing of the array (i.e. $0°$ is forward endfire and $90°$ is broadside). Note that when a linear array is used the AOAs from $0°$ to $180°$ are indistinguishable from those at $180°$ to $360°$. This is often referred to as left-right ambiguity. The frequency domain model used for the sampled array output is then given by the $N \times 1$ vector,

$$\boldsymbol{y}(f, t_j) = s(f, t_j)\,\boldsymbol{v}(f, \phi) + \boldsymbol{n}(t_j), \qquad (9)$$

where f represents frequency, j are the times of the signal snapshots, $s(f, t_j)$ is the signal amplitude, $\boldsymbol{v}(f, \phi)$ represents the so-called steering vector to a signal in direction ϕ [17], and $\boldsymbol{n}(t_j)$ is additive (spatially and temporally) white noise. It is assumed that the signal and noise are mutually independent, zero-mean Gaussian random processes with variance σ_{s}^2 and σ_{n}^2, respectively. The model ignores ambient noise that may be colored. Since it is also assumed that the frequency of interest is known, we will suppress f dependence from here forward.

Based on the model of (9) and assuming that the signals are uncorrelated, then the observation data covariance matrix has the form

$$\mathbf{R}_{\boldsymbol{y}} = E\left\{\boldsymbol{y}(t_j)\,\boldsymbol{y}(t_j)^H\right\} \qquad (10)$$

$$= \sigma_{\mathrm{s}}^2 \boldsymbol{v}(\phi)\,\boldsymbol{v}(\phi)^H + \sigma_{\mathrm{n}}^2 \mathbf{I} \qquad (11)$$

where σ_{s}^2 and σ_{n}^2 represent the signal and noise powers, respectively. Each measurement, $\boldsymbol{y}$, will consist of a collection of snapshots $\boldsymbol{y}(t_j)$, $j = 1, 2, ..., J$. The sample covariance matrix is then defined to be

$$\mathbf{C}_{\boldsymbol{y}} = \frac{1}{J}\sum_{j=1}^{J} \boldsymbol{y}(t_j)\,\boldsymbol{y}(t_j)^H. \qquad (12)$$

Given that the signal and noise are zero-mean Gaussian random processes, the log-likelihood function (LLF) is given by [17]

$$\ln L(\mathbf{y}|\mathbf{s}) = -NJ\ln(\pi) - J\ln\left(|\mathbf{R}_{\boldsymbol{y}}(\boldsymbol{s})|\right) - J\mathrm{tr}\left(\mathbf{R}_{\boldsymbol{y}}^{-1}(\boldsymbol{s})\,\mathbf{C}_{\boldsymbol{y}}\right). \qquad (13)$$

Note that there is a non-linear relationship between the state, $\boldsymbol{s}$, and AOA, ϕ. Substituting the form for $\mathbf{R}_{\boldsymbol{y}}$, given by (11), into (13) and removing constants that are independent of ϕ we obtain the simplified expression

$$\ln L(\boldsymbol{y}|\boldsymbol{s}) = \ln L(\boldsymbol{y}|\varnothing) - J\ln\left(1 + N\frac{\sigma_{\mathrm{s}}^2}{\sigma_{\mathrm{n}}^2}\right) + J\left(\frac{N}{\sigma_{\mathrm{n}}^2}\cdot\frac{N\frac{\sigma_{\mathrm{s}}^2}{\sigma_{\mathrm{n}}^2}}{1 + N\frac{\sigma_{\mathrm{s}}^2}{\sigma_{\mathrm{n}}^2}}\right) P_{\mathrm{CBF}}(\phi) \qquad (14)$$

where $P_{\mathrm{CBF}}(\phi) = \boldsymbol{v}(\phi)^H \mathbf{C}_{\boldsymbol{y}} \boldsymbol{v}(\phi)/N^2$ is the output power of a conventional beamformer (CBF), and $\ln L(\boldsymbol{y}|\varnothing)$ is the no-target ($\sigma_{\mathrm{s}}^2 = 0$) log-likelihood, given as

$$\ln L(\boldsymbol{y}|\varnothing) = -K\left[N\ln(\sigma_n^2) + \frac{1}{\sigma_n^2}\mathrm{tr}(\mathbf{C}_{\boldsymbol{y}})\right]. \qquad (15)$$

Unfortunately, the parameters σ_{s}^2 and σ_{n}^2 are usually not known *a priori*. To that end, the parameters are replaced with their maximum likelihood estimates (MLEs) when evaluating (14). A method for computing these MLEs can be found in [18]. Although this model assumes only one signal is present, it may be generalized to account for interferers as discussed in [19].

3 Example with Simulation data

The first example involves a purely simulated data set and explores the general concepts and feasibility of the approach for fusing active and passive sonar data. It also reveals the value of combining two complementary information sources for improving target localization.

3.1 Example description

The raw acoustic data is generated using SST [1, 2] as mentioned previously. The scenario layout is summarized in Fig. 1. The example consists of a single point target (the TOI) with initial position of [-12, 28] km and a constant velocity of [0, -4.5] m/s. The simulation is run for 225 minutes. The target is modeled to emit a narrowband signal with a center frequency of 500 Hz and a source level of 141 dB re 1 μPa in the frequency bins processed. In addition, it has an active target strength of 27 dB. The ambient noise is 62 dB re 1 μPa in the frequency bins processed.

In addition to the TOI, two discrete clutter objects are modeled. The motivation for these objects is to approximate the effects of fixed scatterers such as bathymetric features or wrecks on the ocean floor. The first fixed clutter object (labeled C1) is a point scatterer located at [7.05, -13] km having a target strength of 28 dB. The second (labeled C2) is a point scatterer with a target strength of 28 dB located at [9.15, 7] km. Neither radiates any acoustic energy.

The passive system consists of one 34.5 m long horizontal line array (HLA) consisting of 25 uniformly spaced elements centered at [-5, 0] km and oriented east-west. The data is sampled at 2400 Hz and processed over a frequency band from 498.6 to 502.1 Hz in 90 second scans.

The active systems consists of a source (labeled S in Fig. 1) at [5, -1] km and a receiver (labeled R) at [0,-5] km. The LFM sweep has a center frequency of 825 Hz, a bandwidth of 50 Hz, a source level of 215 dB, and a pulse length of 1 s. Each scan lasts for 30 seconds, leading to a maximum bistatic range of about 45 km, or loosely speaking a range of about 20-25 km from the receiver, and scans occur every five minutes. The receiver consists of a horizontal plane array with an azimuthal beam width of 13 degrees. The data is simulated

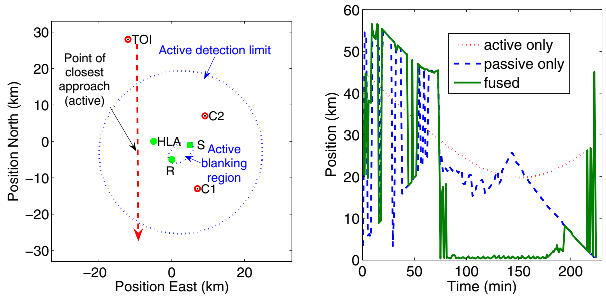

Figure 1: Scenario layout and root mean squared error (RMSE) in position estimates for Example 1

at the element level in SST and is then passed through a standard processing chain of a beamformer, a matched filter, and a normalizer.

For this active system, the TOI appears to have an observed SNR in the normalizer output of around 13 dB just as the target enters the detection region around time 51 minutes. These numbers were found by analyzing the normalizer output and comparing the peak returns in the vicinity of the target's known true position. For the four scans surrounding the closest point of approach, the observed peak SNRs in the normalizer output near the target's true position were about 16-19 dB. For clutter object C1, the observed SNRs are around 17-19 dB, while for clutter object C2 the observed SNRs are around 14-15 dB. Note that the TOI has observed SNRs that are equal to or small than those for clutter object C1.

For the Bayesian tracker settings, we assume that the TOI has a received SNR of 12 dB (input σ_{T}^2 in (6)), the threshold η is set to 10 dB, and the GPD shape parameter ξ is 0.30. The fixed grid covers a region from -33 km to 33 km north and -33 km to 33 km east, as well as velocities in the region -12 m/s to 12 m/s north and -12 m/s to 12 m/s east. There are 51 grid cells in each position dimension, and 7 grid cells in each velocity dimension.

3.2 Results

We compare three different scenarios for this example. The first uses the active system only, while the second considers only the passive system. The third scenario fuses the measurements from both the active and the passive systems. The first main criterion that will be used in comparing the results are the root mean squared error (RMSE) of the position estimates (i.e. localization error). In general, it is also important to look at the target probabilities, but in this example, the target probabilities rise to near 1 very quickly for all methods due to the presence of the clutter objects (for active) and the relatively loud received level from the TOI on the passive system. The RMSEs for all three scenarios are shown in Fig. 1. There are obvious differences in the results, with the fused scenario performing much better than the other two, although even the fused does not localize the target until about 30 minutes after it enters the active detection region. While this is interesting, the real value of the example is derived from examining each scenario independently before comparing across methods. In this way, the actual complementary between the active and passive data can be explored.

3.2.1 Active only scenario

The active only scenario fails to track the target for primarily one reason: the persistent clutter objects are tracked instead. Specifically, the tracker locks on to object C1. At the start of the simulation, C1 and C2 are inside the detection region. With received SNRs of around 18 dB and 15 dB respectively, they are easily identified with the tracker in a single scan. C1 is louder than C2 for two reasons. First, C1 is slightly closer (about 1.5 km closer in bistatic distance). Second, it is nearly aligned with a beam, whereas C2 falls nearly midway between beams, resulting in scalloping losses.

As shown in the example posterior from the second active scan (provided in the top left of Fig. 2), the tracker has already concluded that the probability of there being a target in the state space is high, and it has localized the maximum *a posteriori* (MAP) estimate (labeled with the arrow), to C1. There is also some probability collecting near C2, which reflects the consistent echos in that region. One can also see areas of high probability scattered around the detection annulus. These correspond to random clutter returns.

By scan six (top right of Fig. 2), the tracker is doing a very good job of localizing the two clutter objects. This is evidenced by the probability mass clustering near each object and being cleared out from the regions with no consistent returns. The shading is subtle in the figure, but it is important to note that the regions not contained in the detection annulus have a higher probability than regions inside the detection region where there have been no consistent returns. Because the likelihood models (via the predicted SNR) include information about detection ranges, the tracker maintains the possibility that there is a target in regions it cannot "hear."

Over time, the tracker isolates the stronger of the two clutter objects, as shown in the bottom left of Fig. 2 for scan 31. At this point, nearly all of the mass is focused near C1.

Occasionally, the returns from C1 are relatively weak, while simultaneously relatively loud returns from the TOI are present. However, the prior mass on C1 keeps the tracker focused there, even though some mass moves to the TOI and C2, as shown in the lower right of Fig. 2 for scan 37. The bulk of the probability (including the MAP estimate) remains with C1, but there is some mass now visible around the TOI and C2.

The situation shown in the top right of Fig. 2 is a candidate for separating out two targets in some manner as part of the tracker. Alternatively, if the fusion operator becomes convinced that these are bathymetric features or wrecks, they can be accounted for as known clutter objects using the techniques outlined in [11].

The example posteriors shown in Fig. 2 help to explain the RMSE result shown in Fig. 1. The tracker, using the MAP point estimate, does not track the TOI appropriately because it instead tracks the louder object C1. However, results such

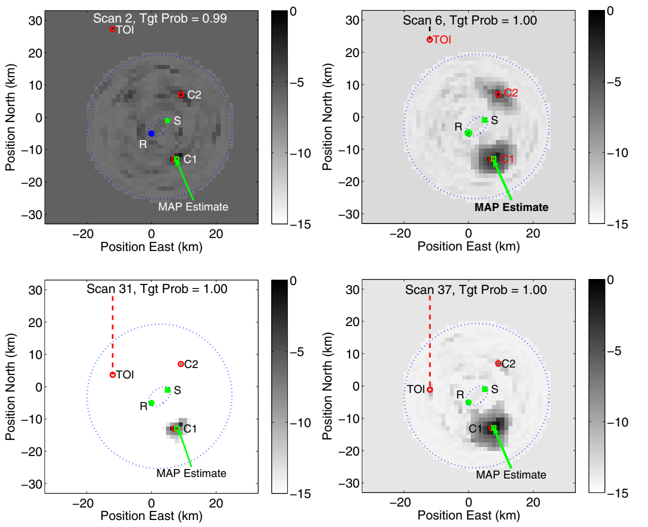

Figure 2: Posteriors for active only scenario of example 1. Gray scale reflects log of probability.

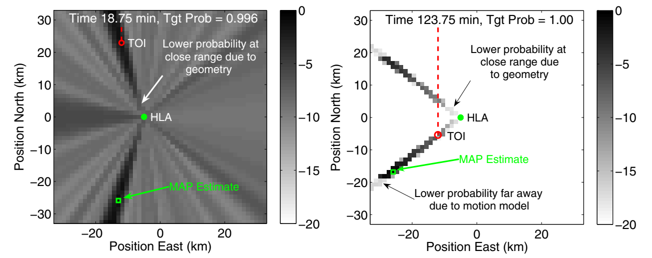

Figure 3: Posteriors for passive only scenario of example 1. Gray scale reflects log of probability.

as the example in the lower right of Fig. 2, as well as examination of the likelihoods (not shown), reveal that there are detectable echoes arriving from the target. Fusion with a different sensor that can also detect energy from the target but not the clutter objects may provide complementary information that reinforces TOI and downplays C1. Specifically, a passive sensor can provide this information, as shown in the following.

3.2.2 Passive only scenario

The inability of the passive receiver alone to localize the target (shown in Fig. 1) is not surprising, because the passive receiver only gives an estimate of the bearing to the target. However, the details of the passive tracking behavior are important to consider in order to better understand the fused result.

Due to the linear nature of the array, there is a left-right ambiguity in the bearing estimate. The ambiguity is clearly visible on the right side of Fig. 3. Due to numerics, the MAP estimate will fall on one of the two lines of bearing corresponding to the ambiguity, so it may flip-flop between the two correct ridges of high probability in bearing.

The MAP estimate is defined as the grid cell with the highest posterior probability, which leads to an interesting effect in the tracker. Because beams spread out in Cartesian space as range increases, a particular grid cell far from the receiver spans a smaller angle than a grid cell close to the receiver spans.

Consider a sequence of grid cells with centers falling exactly on the line of bearing with the highest marginal probability. The cell that is farther away will have the highest posterior probability due to simple geometry. The probability assigned to a grid cell approximates the integral of the true PDF over the states in that cell, and the PDF falls off as one moves off beam. Therefore, cells at close range, which cover a wider span of bearings, will have lower total probability. This decrease in probability at close range is apparent in the left side of Fig. 3.

On the right side of Fig. 3, a posterior from later in the simulation is shown. While the same effect is seen near the receiver, a competing effect is apparent far from the receiver where the probability falls off again. This is due to the modeling of of the velocity state space. The limit of ± 12 m/s in each velocity dimensions bounds the possible motion of the target in the motion update. For the target to maintain a long range from the receiver and still achieve the observed bearing rate, it would have to be traveling with a greater velocity than modeled. This is a positive effect of the tracker, assuming the bounds are properly chosen for the type of target of interest.

Another important aspect of the posterior on the left side of Fig. 3 is the presence of areas of high probability that do not correspond to the target. These are due to random ambient noise. Even after 13 scans, there is still some probability that there is a TOI at various bearings. At the beginning of the simulation, the received level from the target is around 52 dB re 1 μPa, whereas the ambient noise is around 62 dB re 1 μPa. Later in the simulation, the effects from the ambient noise die out because the random noise is not consistent over a time and because the received level of the target increases as it gets closer to the array.

3.2.3 Fused scenario

The ability of the passive sensor to localize the bearing of the TOI in directions that do not correspond to active clutter objects suggests that it may be used to reduce the clutter in the active tracking problem. As noted previously, the Bayesian formulation is a natural mechanism by which to fuse the active and passive data. The same target simulation and active and passive data from the previous two sections is used again but in a fused manner, with select posteriors shown in Fig. 4.

In the upper left posterior in Fig. 4, the MAP estimate remains on C1. The passive data retains some ambiguity in the bearing up to this point. As the received level from the target increases (by virtue of getting closer to the HLA) and random noise evens out, the bearing becomes more certain and accurate. The upper right posterior in Fig. 4 is a few scans later, and the passive has started to clear out the probability in the directions corresponding to the clutter objects (which are not radiating an acoustic energy).

After 18 minutes (bottom left of Fig. 4), the results of the fusion become more obvious. The passive sensor has isolated the target bearing, and the active system has cleared out some of the mass inside its detection annulus where it is not consistently receiving echoes. There is still some probability mass associated with each clutter object due to the active returns, but the MAP estimate has now moved outside the active detection region. The MAP estimate is at the boundary due to the geometric effect discussed in Sec. 3.2.2, and it is again subject to flipping between the two lines of bearing.

At this point, the active and passive systems are essentially battling each other. The passive system senses a target in directions that contradicts what the active system is sensing, and vice versa. The likelihood functions in Sec. 2 provide a natural scaling between the two systems that is data dependent and guides this battle.

Once the target enters the active detection annulus, the active and passive systems can both sense it. For measurements near the TOI, the two systems reinforce the probability of a target in that region. The disagreement in the directions of the clutter objects persists. Consequently, the agreement in the area of the TOI leads to an accurate MAP estimate near the TOI, as shown in the bottom right of Fig. 4. Several important phenomena are visible in this posterior.

First, the probability along the passive system's ambiguity lobes is greater outside the active detective region. The active likelihood model accounts for the detection annulus, so no matter what the active systems actually receives, there is always some probability that there is a target outside the detection annulus. Inside the detection annulus, probability increases where there are consistent returns and decreases where there are either no returns or inconsistent returns over several scans. Near the TOI, there are active returns consistent with the moving target, and the passive system reinforces this region as being a likely location for a target. While some probability mass remains on each clutter object (barely visible in the figure), the contradictory active and passive data on those objects is overwhelmed by the consistent data on the actual TOI. This leads to accurate localization once the target has entered the active detection annulus.

4 Example with DEMUS Active Clutter Measurements

This example uses data that was collected during the DEMUS'04 sea-trial, which was conducted in the Malta Plateau region during September of 2004. DEMUS'04 was performed as part of the Deployable Multistatic Sonar Joint Research Project formed between the NATO Undersea Research Centre, the U.S. Office of Naval Research , and the U.K. Ministry of Defense. The sea-trial involved the use of, among other assets, the Deployable Multistatic Undersea Surveillance (DEMUS) system, a set of moored active sonar sources and receivers.

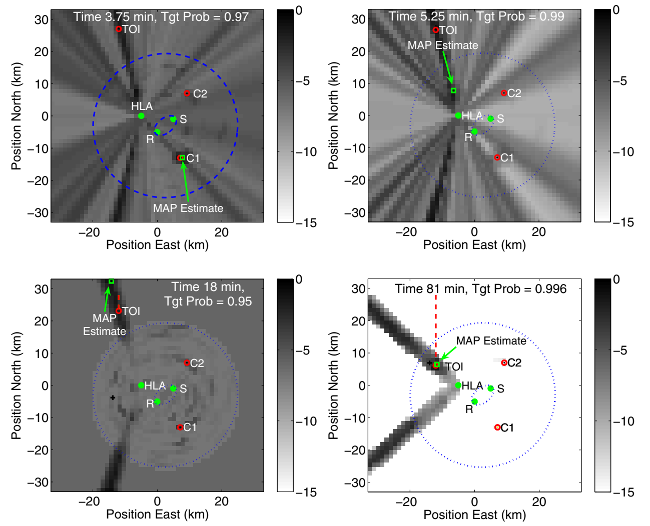

Figure 4: Posteriors for fused scenario of example 1. Gray scale reflects log of probability.

4.1 Data description

The configuration is summarized in Fig. 5, which also includes details about five regions where active clutter is persistently found throughout the recorded data. For the purpose of this example, three components are of interest: one DEMUS source, denoted S located at the origin, and two DEMUS receivers, denoted R1 (located at [5.3814, -8.0006] km) and R2 (located at [-1.8352, -6.1560] km. Each was moored to the bottom at a depth of about 50 m in an area with some 100 m of water depth.

For this particular data set, denoted E06, the DEMUS source pinged once every two minutes, including an LFM pulse centered at 2350 Hz with a bandwidth of 400 Hz and 1-sec duration. A total of 60 scans, each about 34 sec long, were performed every two minutes. No targets were present in the original data.

Using SST, a single point target was injected into the data set using the same basic source and receiver parameters as were used in the true DEMUS test. The target has an initial position of [15, 21] km and a constant velocity of [0, -4] m/s. The target strength was configured such that the target has an SNR of around 14 dB at R1 when it enters the detection annulus, and has a peak SNR of around 18 dB on R1 at closest approach. A radiated noise profile was also generated for the target using SST. The passive HLA and target are given the same parameters as those described in Sec. 3.1, except the array is now centered at [0, 5] km and scans are 60 seconds long.

The tracker properties are the same as in the previous example. Also shown is the RMSE error performance of each method. Clearly the fused approach provides an advantage over the other approaches, as we will explore in more detail in the following.

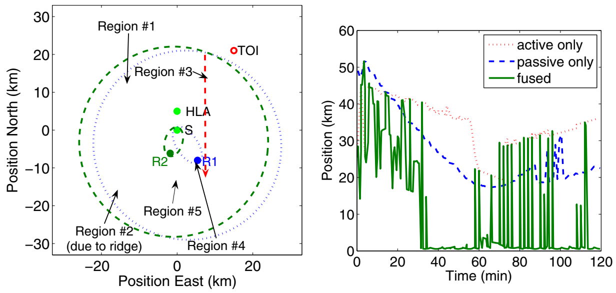

Figure 5: Example 2 layout and RMSE

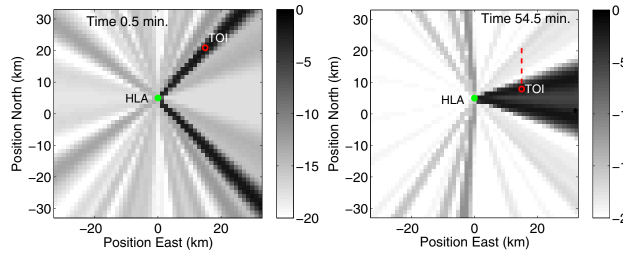

Figure 7: Likelihoods for passive only scenario of example 2. Gray scale reflects log of likelihood.

4.2 Active results with real clutter data

Having considered the posterior distributions in the previous example, we here focus on the likelihood functions because they better reveal the information presented by each individual scan and receiver. For example, clutter regions #2, #3, and #4 (labeled in Fig. 5) are apparent on the sample scan from R1 shown in the top left of Fig. 6. On R2, only areas #1 and #2 are apparent, as shown in the sample scan in the top right of Fig. 6.

Once the TOI enters the detection annulus, some likelihood exists around it. However, the SNR of the TOI is still not as high as the SNR of some clutter regions, as shown in the lower left and lower right images of Fig. 6. In the lower left, the peak is at some transient clutter object. In the lower right example, the peak of the likelihood function in on clutter region #1. Given these differences in scale of the likelihoods (which reflect an underlying difference in SNR), it is not surprising that the active system alone fails to detect the target—it is much quieter than the clutter. However, the hope is that the passive system can complement this information while ruling out the clutter regions.

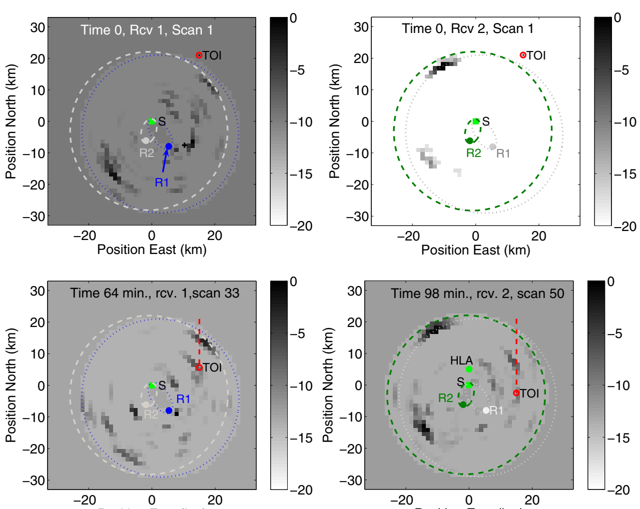

Figure 6: Likelihood functions for active only scenario of example 2. Gray scale reflects log of likelihood.

4.3 Passive results

The passive results are analogous to those in Sec. 3.2.2, although the layout of the HLA and TOI differ. Two example likelihood functions are shown in Fig. 7. At the start of the simulation, the received signal levels from the TOI are sufficient to isolate the bearing, as the received signal level is initially around 54 dB re 1 μPa with 62 dB re 1 μPa ambient noise. Later in the simulation, the TOI is at endfire, so the localization in bearing is relatively poor, as shown on the right side of Fig. 7, but it may still be enough to disambiguate the TOI from all but perhaps clutter region #4.

4.4 Fused results with real active clutter

The active and passive data from the previous two sections is now fused into one scenario. The RMSE performance shown in Fig. 5 reveals mostly small errors in the fused result once the target enters the active detection annulus. However, there are some large errors even after it is first localized well. Four posteriors are shown in Fig. 8 in order to examine the behavior in more detail.

The upper left of Fig. 8 shows the posterior (not the likelihoods, as were shown in the previous two sections) after the tenth active scan from receiver 1, which occurs just as the TOI enters the detection annulus for receiver 1. Clutter region #2 is obvious, and regions #4 and #5 also appear faintly. The TOI is very near clutter region #3, which is also quite apparent. However, the passive lines of bearing corresponding to the TOI are also apparent.

The upper right of Fig. 8 shows the posterior 16.5 minutes later and after a passive scan. Here the MAP estimate is right on the TOI, and a large mound of probability mass can be seen surrounding it. The passive lines of bearing are also visible. There is still some probability mass around clutter regions #2 and #5, and there is a concentration of mass to the east of the source. The addition of one more active scan leads to the posterior shown in lower left of Fig. 8, in which nearly all of the probability is concentrated on the TOI.

The lower right posterior in Fig. 8 shows one example of when the MAP estimate jumps off of the target. Although there is a strong concentration of probability still on the TOI, the MAP estimate jumps to clutter region #5 due to a very loud active return from that region. Because this clutter is not pronounced on every scan from receiver 2, is almost non-existent on receiver 1, and is non-existent on the passive, the fused MAP estimate quickly returns to the TOI. This reveals the potential sensitivity of the fused system to a very strong return from one sensor, while also revealing the robustness of the long term behavior.

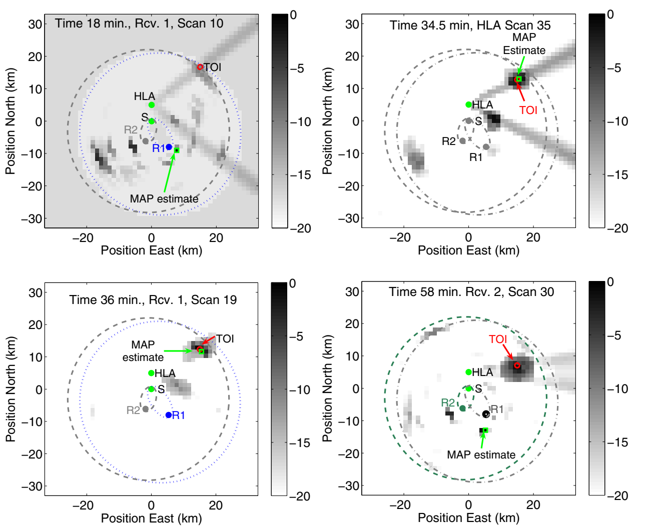

Figure 8: Example posterior distributions for fused scenario of example 1. Color map reflects log of posterior probability.

5 Summary

Active and passive sonar systems naturally provide complementary information, and Bayesian inference provides a natural means for fusing different types of information that can be summarized in a probabilistic manner. Building off of previous work on separate active and passive modeling and Bayesian tracking, we have shown the benefit that fusing active and passive sonar information provides for target localization, and in particular the mitigation of active clutter. The first, simulated example has illustrated the basic complementarity of the active and passive systems. The second example, which is based on real world, recorded active sonar data containing clutter returns, is a more realistic example of the type of challenges faced in sonar tracking. The benefit to fusing active and passive information that was shown in this example strongly suggests that the benefit will extend to real world applications. Future work will continue to explore more realistic scenarios, with the eventual goal of using an entire real-world, active-passive data set.

Acknowledgments

This work was supported by the Office of Naval Research under contract no. N00014-06-G-0218-01. DEMUS'04 data was made possible by the Deployable Multistatic Sonar Joint Research Project (JRP), a collaboration between NATO Undersea Research Centre (NURC), MOD (U.K.), and the Office of Naval Research (U.S.).

References

[1] R. P. Goddard, "The Sonar Simulation Toolset," in *OCEANS'89*, vol. 4, September 1989, pp. 1217–1222.

[2] ——, "The Sonar Simulation Toolset, release 4.6: Science, mathematics, and algorithms," Applied Physics Laboratory, University of Washington, Seattle, WA, Tech. Rep. Technical Report APL-UW 0702, March 2008.

[3] Y. Bar-Shalom and T. E. Fortmann, *Tracking and data association*. San Diego: Academic Press, 1988.

[4] L. D. Stone, C. A. Barlow, and T. L. Corwin, *Bayesian multiple target tracking*. Boston: Artech House, 1999.

[5] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, Feb. 2002.

[6] J. Aughenbaugh and B. R. La Cour, "Measurement-guided likelihood sampling for grid-based Bayesian tracking," *Journal of Advances in Information Fusion*, 2009 under review.

[7] G. Kitagawa, "Non-gaussian state-space modeling of nonstationary time series," *Journal of the American Statistical Association*, vol. 82, no. 400, pp. 1032–1041, 1987.

[8] S. C. Kramer and H. W. Sorensen, "Recursive bayesian estimation using piece-wise constant approximations," *Automatica*, vol. 24, no. 6, pp. 789–801, 1988.

[9] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA: The MIT Press, 2005.

[10] C. Jemmott, R. L. Culver, and N. Bose, "Passive sonar target localization using a histogram filter with model derived priors," in *IEEE Asilomar Conference on Signals, Systems, and Computers.*, Pacific Grove, CA, 2008.

[11] J. M. Aughenbaugh and B. R. La Cour, "Use of prior information in active sonar tracking," in *Proc. 12th Int. Conf. on Information Fusion*, Seattle, WA, Jul. 2009.

[12] C. M. Kreucher, B. Shapo, and R. Bethel, "Multitarget detection and tracking using multi-sensor passive acoustic data," in *Proc. IEEE Aerospace Conference*, Mar. 2009, pp. 1–16.

[13] D. J. C. MacKay, "Bayesian interpolation," in *Neural Computation*, vol. 4, no. 3, May 1992, pp. 415–447.

[14] B. R. La Cour, "Ensemble-based Bayesian detection and tracking." in *Proceedings of Meetings on Acoustics*, vol. 1, 2008.

[15] B. La Cour, "Stationary priors for Bayesian target tracking," in *Proceedings of the 11th International Conference on Information Fusion (Fusion 2008)*, Cologne, Germany, 30 June – 3 July 2008.

[16] P. M. Woodward, *Probability and information theory, with applications to radar*. New York: McGraw Hill, 1953.

[17] H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part IV, Optimum Array Processing*. New York: Wiley, 2002.

[18] R. E. Bethel, "Joint detection and estimation in a multiple signal array processing environment." Ph.D. dissertation, George Mason University, Fairfax, VA, Aug. 2002, available at (http://gunston.gmu.edu/kbell/pubs/BethelDissert.pdf).

[19] B. A. Yocom, T. W. Yudichak, and B. R. La Cour, "Bayesian passive sonar tracking in the presence of known interferers," in *Proc. 12th Int. Conf. on Information Fusion*, Seattle, WA, Jul. 2009.