18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

Can This Target Be Tracked?

Steven Schoenecker
NUWC
Division Newport
Newport, RI
Email: steven.schoenecker@navy.mil

Peter Willett
Electrical Engineering
University of Connecticut
Storrs, CT
Email: willett@engr.uconn.edu

Yaakov Bar-Shalom
Electrical Engineering
University of Connecticut
Storrs, CT
Email: ybs@engr.uconn.edu

***Abstract*—In the field of target tracking, it is often assumed that as long as a target is present and detectable, it should be "trackable." Any failure to track a target that is generating detectable measurements is assumed to be due to a sub-optimal tracking algorithm or perhaps an algorithm that is not properly "tuned." There seems to exist the idea that "if this knob is turned slightly, or this parameter is adjusted a little, our tracker should be able to follow the target ..." This work shows that, as should really be expected, there are times when, even though there is a target present that is producing measurements, the output statistical distribution produced by these measurements cannot be differentiated from the output statistical distribution of the too-numerous clutter-generated measurements, and the target simply cannot be tracked.**

This degree of "trackability" is demonstrated by employing the Maximum Likelihood Probabilistic Multi-Hypothesis Tracker (ML-PMHT), which is a powerful non-Bayesian algorithm that uses a generalized likelihood ratio test (GLRT) to check for a target in the presence of clutter. For various combinations of measurement dimensionality, amplitude features, and classification features, we treat the ML-PMHT log-likelihood ratio (LLR) as a random variable (RV) transformation and then use extreme-value theory to calculate the probability density function (PDF) for the peak point in the LLR due to clutter as well as the PDF of the peak point in the LLR due to a target. In doing so, the tracking problem is reduced to a simple detection problem, making it possible to answer the question, "Can this target be tracked?"

Keywords: Tracking, trackability, ML-PMHT, maximum likelihood, extreme value theory

I. INTRODUCTION

The Maximum Likelihood Probabilistic Multi-Hypothesis Tracker (ML-PMHT) is a non-Bayesian algorithm that can be implemented as a powerful multitarget, multistatic active tracker. The ideas behind it were first introduced in [4], [23], [24], and [25]. ML-PMHT was first implemented as a multistatic tracker in [26] and [27], with more recent work done in [17].

At its core, ML-PMHT is a fundamentally simple algorithm. Assumptions are made about a target[1] and the environment in which the target is present. The assumptions for ML-PMHT are [5]

- A single target is present in each frame with known detection probability P_d. Detections are independent across frames.
- Any number of measurements per frame can originate from the target.
- The kinematics of the target are deterministic. The motion is usually parameterized as a straight line, although any other parameterization is valid.
- False detections (clutter) are uniformly distributed in the search volume.
- Target measurements are corrupted with zero-mean Gaussian noise with known variance.
- Measurements at different times, conditioned on the parameterized state, are independent.

With these assumptions, the ML-PMHT log-likelihood ratio can be constructed. It is written as [16]

$$\Lambda(\mathbf{x}, Z) = \sum_{i=1}^{N_w} \sum_{j=1}^{m_i} \ln \left\{ 1 + \frac{\pi_1}{\pi_0} V \rho_a \rho_c p[\mathbf{z}_j(i)|\mathbf{x}] \right\} \qquad (1)$$

Here, π_1 is the prior probability that a given measurement originates from the target, π_0 is the prior probability that a given measurement originates from clutter, V is the measurement search volume, $p[\mathbf{z}_j(i)|\mathbf{x}]$ is a target-centered Gaussian distribution, ρ_a is the amplitude likelihood ratio, ρ_c is the classification feature likelihood ratio, N_w is the number of scans processed by the tracker, m_i is the number of measurements in the i^{th} scan, Z is the entire batched set of measurements, and $\mathbf{x}$ is the target parameter vector (a target state vector sometimes augmented by other estimated parameters).

Running the tracker is done as follows: Once a batch of data is available, the expression in (1) is optimized over $\mathbf{x}$. If the LLR value for the optimal $\mathbf{x}$ is greater than a certain threshold, a target track is declared; if the LLR value at the optimal $\mathbf{x}$ is less than the threshold, the point is deemed to have been caused by clutter and is ignored.

Of the assumptions listed above, the second — that more than one measurement can originate from a target in a single scan — sometimes raises objections. The "typical" target-measurement generation model used in target tracking is that

Supported by NPS via ONR N00244-14-1-0033, N00014-13-1-0231, and ARO grants W991NF-06-1-0467 and W911NF-10-1-0369.
Proc. 18th Intn'l Conf. Info. Fusion. Washington, D.C., July 2015.

[1] ML-PMHT is easily extensible to multiple targets, but for this work on trackability, it will suffice to consider just a single target.

only zero or one measurement can originate from a target in a single scan. If this assumption is followed, the LLR in (1) would instead take the form of the Maximum Likelihood Probabilistic Data Association (ML-PDA) tracker [11], [12]. However, such concerns about the target-measurement generation model should be alleviated by previous work [16], [20] that shows that when there is at most one hit per scan, the ML-PMHT LLR converges to the ML-PDA LLR and there is virtually no performance difference between the two algorithms.

It is possible to create a framework that determines the trackability of a target with the ML-PMHT LLR. This is done by statistically quantifying the peak of the ML-PMHT LLR that is caused just by clutter, and then quantifying the peak of the ML-PMHT LLR that is caused just by a target. With this, the tracking problem is reduced to a simple detection problem — the peak in the ML-PMHT LLR due to clutter becomes the H_0 distribution, and the peak in the LLR due to the target becomes the H_1 distribution (here, the H_0 and the H_1 distributions are the distributions under the H_0 and H_1 hypotheses, respectively). By applying the Neyman-Pearson Lemma [14], a probability of false track acceptance (P_{FT}) is set. From this, a threshold can be computed; the H_1 distribution integrated to the right of this threshold determines the detection probability of target track (P_{DT}).

Finally, if the assumptions listed above about the target and the environment are correct, then ML-PMHT is an optimal algorithm in the sense that it uses all the available data, and it does not require any simplifying assumptions to work. Additionally, the ML-PMHT algorithm is a generalized likelihood ratio test (GLRT), which (at least asymptotically) is the optimal test according to classical detection theory [13], [14], [28]. Thus, trackability results obtained from an analysis of the ML-PMHT LLR can be generalized; if the framework here determines that ML-PMHT cannot track the target, then it is highly unlikely that *any* algorithm can track the target.

This work is meant to describe at a high level the entire trackability framework for all cases considered thus far, as well as introduce the latest addition to the framework. (Without this brief description of the entire framework, any discussion of new results would make little sense.) Previous works have examined different individual portions of this framework in much more detail. In [19], we derived the peak clutter ML-PMHT LLR PDF for all possible closed-form cases (these cases will be discussed shortly), and [22] derived the peak target LLR for the same closed-form cases. Next, [21] re-derived the peak clutter PDF when the clutter amplitudes were K-distributed. After this, [18] derived trackability for cases where the clutter amplitudes were K-distributed. Finally, here we present for the first time the initial theory and results from introducing a classification feature into the trackability framework.

II. ML-PMHT LLR Peak Due to Clutter

The PDF of the peak in the ML-PMHT LLR due to clutter is calculated by starting with the assumption that clutter measurements are uniformly distributed in the search volume. Now consider the ML-PMHT LLR for a single measurement, in a slightly different form from (1) for the case of two-dimensional measurements (i.e. azimuth and time-delay). This is written as

$$\Lambda_1(z) = \ln\left\{1 + K_d \rho_c \rho_a \exp^{-\frac{1}{2}\left[\frac{(z_1-\mu_1)^2}{\sigma_1^2} + \frac{(z_2-\mu_2)^2}{\sigma_2^2}\right]}\right\} \quad (2)$$

Here, z_1 and z_2 are the individual clutter-generated measurements (in the case of 2-D tracking, these would be in azimuth and time-delay space), μ_1 and μ_2 are the target locations in the two measurement space directions (what the actual values of μ_1 and μ_2 are turns out not to matter), and σ_1 and σ_2 are the individual standard deviation components of the measurement covariance matrix. Finally, the constant K_d for d measurement dimensions (in (2) $d = 2$) is

$$K_d = \frac{\pi_1}{\pi_0} \frac{V_d}{\sqrt{|2\pi \mathbf{R_d}|}} \quad (3)$$

where V_d is the d-dimensional search volume and $\mathbf{R_d}$ is the measurement covariance matrix.

The key first step in determining the PDF of the peak point in the ML-PMHT LLR due to clutter is treating (2) as simply a transformation of a uniform random variable, which from the ML-PMHT assumptions, is the distribution for clutter measurements. For certain limited cases, it is possible to compute this transformation analytically. However, it is usually necessary to compute some or all of the (multiple step) random variable transformation numerically.

A. Clutter cases with closed-form results

It is possible to analytically compute the RV transformation for one-dimensional (azimuth-only or time-delay-only), two-dimensional (azimuth and time-delay), and three-dimensional (azimuth, time-delay, and Doppler) measurements where no amplitude data or classification feature data is present (i.e. $\rho_a = \rho_c = 1$). The work for these analytical cases was performed in detail in [19]; the results are just presented here. The one-dimensional case is somewhat of a "toy" problem and is ignored. The PDF of the transformed RV in the case of two-dimensional clutter is given by

$$p_{w_2}(w_2) = \begin{cases} C\delta(w_2) & w_2 = 0 \\ 2\pi \frac{\sigma_1 \sigma_2}{V_1 V_2} \frac{\exp(w_2)}{\exp(w_2)-1} & 0 < w_2 \leq \ln(1 + K_2) \end{cases} \quad (4)$$

In this equation, C is a normalization constant to ensure the PDF integrates to one, and the "2" subscripts denote the result is for the two-dimensional measurement case.

The transformation result for three-dimensional clutter measurements, which would be time-delay, azimuth and range rate, is given by

$$p_{w_3}(w_3) = \begin{cases} C\delta(w_3) & w_3 = 0 \\ 4\pi\sqrt{2}\frac{\sigma_1\sigma_2\sigma_3}{V_1V_2V_3}\frac{\exp(w_3)}{\exp(w_3)-1}\sqrt{\ln\left(\frac{K_3}{\exp(w_3)-1}\right)} \\ & 0 < w_3 \leq \ln(1+K_3) \end{cases} \tag{5}$$

Finally, it is possible to analytically calculate the expression for 2-D measurements with Rayleigh distributed target and clutter amplitudes. In this case, the PDF is given by

$$p_{w_{2a}}(w_{2a}) = \begin{cases} C\delta(w_{2a}) & w_{2a} = 0 \\ \frac{2\pi\sigma_1\sigma_2}{V_1V_2}\left[1 - e^{\tau}\left(\frac{e^{w_{2a}}-1}{K_2'}\right)^{\frac{1}{K_\sigma}}\right]\frac{\exp(w_{2a})}{\exp(w_{2a})-1} \\ & 0 < w_{2a} \leq \ln\left(1 + K_2'e^{-K_\sigma\tau}\right) \end{cases} \tag{6}$$

where

$$K_2' = \frac{K_2}{\sigma^2}e^{K_\sigma\tau} \tag{7}$$

and

$$K_\sigma = \frac{1-\sigma^2}{\sigma^2} \tag{8}$$

where σ^2 is the expected target amplitude intensity and τ is the value of the detector threshold.

B. Numerical clutter cases

Adding amplitude (other than for the single Rayleigh case described above) or classification feature data to the tracker measurements makes the single-variable transformation described by (2) more difficult. For amplitude data, we will consider cases where the clutter amplitudes are Rayleigh distributed as well as K-distributed. The former is the "traditional" treatment in target tracking [8], [12], while the latter has recently been posited to more accurately describe clutter amplitudes [1], [2], and [3]. (The target amplitudes will still have a Rayleigh distribution in both cases.) In general, when using these amplitude distributions, the transformation of a uniform random variable with the single-measurement ML-PMHT LLR (2) does not have a closed-form solution; the transformation must be done numerically [21].

After this, the effect of classification feature data will be examined (for simplicity, the amplitude LR ρ_a will again be set to one, although this does not have to be the case). Here, there is some "generic" classification feature data present; it is assumed that both the clutter classification feature PDF and the target classification feature PDF have Gaussian distributions. (Future work will examine distributions that are more heavy-tailed than a Gaussian.) Again, as in the case of K-distributed clutter amplitudes, the single-measurement ML-PMHT transformation described by (2) does not have a closed form solution but rather must be done numerically.

In general, the numerical single-measurement transformation (for both amplitude and classification feature data) is done in the following manner. The single-measurement LLR has the form

$$W = \ln(1 + XY) \tag{9}$$

where W, X and Y are all random variables. The RV Y represents the 2-D or 3-D target-centered Gaussian distribution; the RV X represents the amplitude LR or the classification LR. In general, it is preferable, if possible, to calculate the PDFs of X and Y analytically; this will lead to more accurate results. However, the product of X and Y, and by extension, the expression for W (with the exception of the cases listed above) must be done numerically.

C. Single-measurement PDF to batch and peak clutter PDFs

Examples of four different single-measurement transformations are shown in Figure 1. This figure has 2-D and 3-D plots with no amplitude, a 2-D case with K-distributed clutter amplitudes, and a 2-D case with classification features. For the 2-D case with the amplitude LR, the clutter amplitude is K-distributed; the K-distribution parameter α that describes the heaviness of the tail of the distribution is set to $\alpha = 1$. For the 2-D case with the classification feature, the classification feature LR is calculated with a target feature PDF $\sim \mathcal{N}(2, 1)$, and a clutter feature PDF $\sim \mathcal{N}(1, 1)$. Notice that for both the K-distributed case and the classification feature case, some of the mass of the PDF is shifted to the right (as compared to the no-amplitude 2-D PDF); the support for the single transformed clutter measurement has increased.

The ML-PMHT tracker is a batch algorithm; once the PDF for the single-measurement transformation is obtained, it is necessary to calculate the PDF of the batch of data that is processed by the tracker. Fortunately, one of the other assumptions of ML-PMHT is that the individual measurements are independent, so the transformed single-measurement PDFs will be independent as well. Thus, the PDF of a batch of N_c measurements (where the 'c' subscript denotes clutter measurements) is obtained by convolving the single-measurement PDF with itself $N_c - 1$ times. (It is actually more accurate and efficient to calculate the characteristic function of the single-measurement PDF, exponentiate it by N_c, and then take the inverse characteristic function of the product.) The batch PDF that results represents the distribution of LLR values that would be output if the ML-PMHT algorithm was randomly sampled throughout parameter space. This is exactly what is seen in Figure 2; the theoretical plot is the result of the theoretical single-measurement PDF convolved with itself $N_c - 1$ times, whereas the empirical plot is the result from sampling the output of the ML-PMHT tracker with only clutter measurements fed to it.

It necessary to characterize the statistics of the *peak* point of the clutter-only LLR; the batch PDF shown in Figure 2 describes the statistics of the LLR everywhere in target parameter vector space. To go from the batch PDF to the

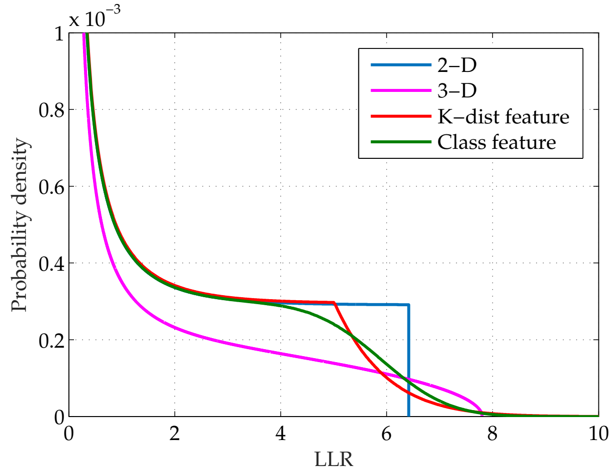

Figure 1. Single clutter measurement transformed to LLR for 2-D, 3-D, 2-D with K-distributed amplitudes, and 2-D with classification feature.

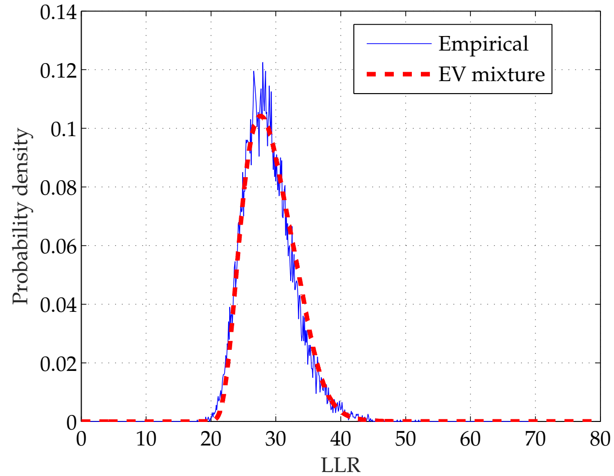

Figure 3. Empirical and theoretical EV mixtures for peak clutter PDF.

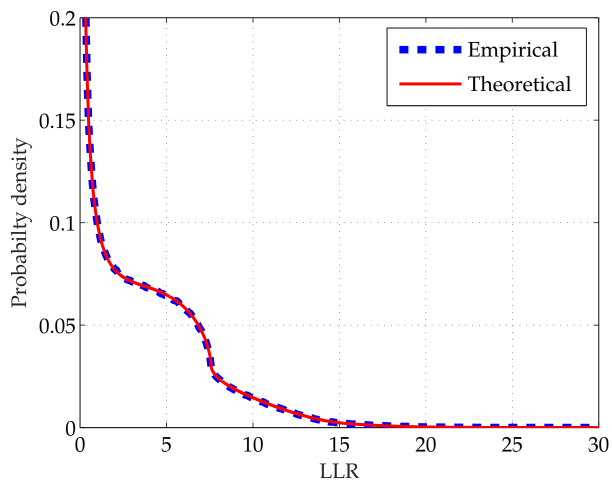

Figure 2. Batch clutter PDF for 2-D measurement with classification feature.

"peak" PDF, we turn to extreme-value theory [7], [9], [10]. In an "actual" ML-PMHT implementation, the batch of ML-PMHT data is optimized with some numerical optimization scheme that hopefully finds the global maximum of the LLR. If we know how accurately the optimizer finds this global maximum, then we can calculate how many times we would have to sample the LLR so that the maximum sample has the same statistics as the actual peak point found by the optimizer. From extreme value theory, this peak point is well-approximated by a Gumbel distribution [6], which has the form of

$$f(x) = \frac{1}{\beta} \exp\left[-\left(\frac{x-\nu}{\beta}\right) - \exp\left(-\frac{x-\nu}{\beta}\right)\right] \qquad (10)$$

Given the number of samples N_s required to get the same accuracy as the optimizer output, as well as the batch PDF, it is possible to calculate the Gumbel distribution that represents the peak PDF (for a much more detailed description of this process, see [19]).

It turns out that the actual optimizer used for ML-PMHT (in this case taken from Matlab's optimization toolbox) outputs with a range of accuracies, so the resultant peak clutter PDF is a mixture of Gumbel distributions. An example of an empirical peak PDF for 2-D measurements with K-distributed clutter ($\alpha = 0.1$), as well as the theoretical EV mixture PDF, is shown in Figure 3. It should be noted that the empirical PDF was obtained via repeatedly simulating a 2-D clutter-only set of measurements, and then running the tracker on it to find the global peak. This took on the order of 10 hours to complete. In contrast, calculating the theoretical peak clutter PDF took only on the order of seconds.

III. ML-PMHT LLR Peak Due to the Target

It is now necessary to calculate the PDF of the peak in the ML-PMHT LLR due to target measurements. This process is very similar to the procedure in the clutter case, but even simpler. Whereas in the case of clutter, a measurement was assumed to have a uniform distribution in measurement space, a target measurement is assumed to have a normal distribution in measurement space, centered around the actual target location. As a result, the transformed single-measurement LLR is just a Gaussian RV transformed by (2).

A. Closed-form and numerical single-measurement target cases

The resultant PDF for 1-D, 2-D and 3-D target measurements with no amplitude can be compactly expressed as [22]

$$p_{t_d}(t_d) = \frac{\exp(t)}{K_d \Gamma(d/2)} \left\{ \ln \left[\frac{K_d}{\exp(t) - 1} \right] \right\}^{d/2-1}$$
$$0 \leq t_d \leq \ln(1 + K_d) \qquad (11)$$

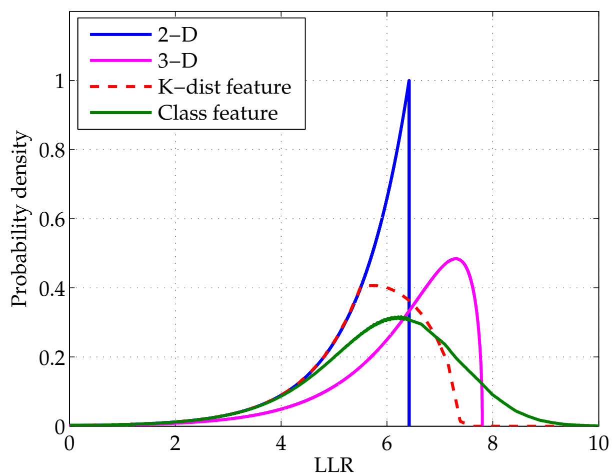

Figure 4. Single target measurement transformed to LLR for 2-D, 3-D, 2-D with K-distribution, and 2-D with classification feature.

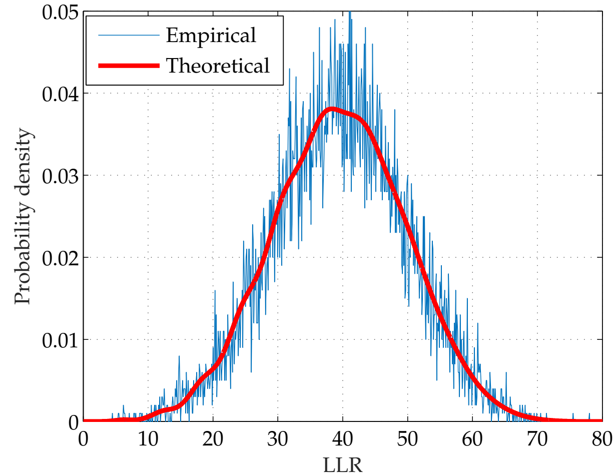

Figure 5. Empirical and theoretical batch/peak PDF for 2-D target measurements with classification feature data.

Again, $d \in \{1, 2, 3\}$ is the number of measurement dimensions. As with the clutter case, it is also possible to obtain an analytic expression for the 2-D single measurement PDF when the clutter and target amplitudes have a Rayleigh distribution. In this case, the expression is

$$p_{t_{2a}}(t_{2a}) = \begin{cases} \frac{1}{|K_\sigma|\sigma^2+1} e^{-|K_\sigma|\tau} \frac{\exp(t_{2a})}{K_2'} & t_{2a} < \ln(1 + K_2' e^{|K_\sigma|\tau}) \\ e^{\frac{\tau}{\sigma^2}} \frac{1}{|K_\sigma|\sigma^2+1} \left(\frac{K_2'}{\exp(t_{2a})-1} \right)^{\frac{1}{|K_\sigma|\sigma_2}} \frac{\exp(t_{2a})}{\exp(t_{2a})-1} & \\ & t_{2a} > \ln(1 + K_2' e^{|K_\sigma|\tau}) \end{cases} \tag{12}$$

However, when other amplitude and/or classification feature data is added, in general, the single-measurement transformed target PDF must be calculated numerically. The process is the same as described in section II-B. Examples of this transformation are shown in Figure 4, again for 2-D, 3-D, 2-D with K-distributed amplitudes, and 2-D measurements with a classification feature.

B. Target batch/peak PDFs

As with the clutter case, now that the single-measurement transformed PDF is available, it is necessary to calculate the batch PDF. This is done by again using the assumption that the individual measurements are independent; if N_t is the total number of target measurements, then the single measurement transformed PDF is convolved with itself $N_t - 1$ times, which will produce the target batch PDF.

In reality, the number of target measurements N_t is not fixed, but is actually a random variable itself. If the probability of target detection in a scan is P_d, then the number of target detections in a batch will be a binomial random variable with mean $P_d N_w$ (which completely describes the distribution). Thus, if $N_w = 11$ (the batch size typically used in the ML-PMHT tracker implementation), we can use the binomial distribution to calculate $P\{N_t = 0\}$, $P\{N_t = 1\}$, $P\{N_t = 2\}$, $\ldots$, $P\{N_t = 11\}$ (where P denotes probability). Then, each individual probability is used to weight the individual target batch PDFs (i.e. the batch PDF for $N_t = 1, 2, \ldots$, etc.) to create a mixture batch target PDF.

At this point, everything done for calculating the target batch PDF was the same as was done for the clutter batch PDF, with the exception that the initial RV was a uniform RV in the clutter case, and a Gaussian RV in the target case. In the clutter case, it was necessary to use extreme value theory to calculate the peak PDF from the batch. This is because the clutter batch PDF represents the statistical distribution of ML-PMHT LLR values caused by *all* the points in parameter space. In contrast, the target batch PDF represents the statistical distribution of the ML-PMHT LLR caused by target measurements *at just one point — the true target location*. Thus, in the target case, *the batch PDF is the peak PDF*. There is no reason to use extreme-value theory for target-generated measurements.

An example of a target batch/peak PDF is shown in Figure 5. Again, the empirically-determined curve matches extremely well with the theoretically determined curve.

IV. TRACKABILITY RESULTS

At this point, we have the PDF that statistically describes the maximum point in the ML-PMHT LLR caused by clutter, and we have the PDF that statistically describes the maximum point in the ML-PMHT LLR caused by the target. This reduces the tracking problem to a simple detection problem. The peak clutter distribution is the H_0 distribution, and the peak target distribution is the H_1 distribution. By using the Neyman-Pearson lemma [15], we can pick an acceptable probability of false track P_{FT} (the typical value selected was $P_{FT} = 0.01$)

and use the H_0 distribution to calculate a tracking threshold. Any optimized ML-PMHT LLR outputs above this value will be deemed to be from the target; any optimized ML-PMHT LLR outputs below this will be ignored as they are assumed to be from clutter.

This rapid and accurate determination of the ML-PMHT tracking threshold is a powerful result by itself. Prior to this, the best way to calculate this threshold was to empirically estimate the Gumbel distribution of the clutter peak in the manner described by [6]. While this method is accurate, it is slow, making real-time implementation of the ML-PMHT tracker with the most accurate threshold impossible. (The threshold will change as environmental conditions change, and it is necessary to have an algorithm that can update the threshold in real-time to account for this.)

More importantly, beyond the tracking threshold determination, it is now possible to make statements about trackability for a given target in a given environment. From classical detection theory, we can compute receiver operating characteristic curves (ROC), or perhaps more appropriately, *tracker operating characteristic curves* (TOC). Or, we can "measure" in a binary sense, if a given target is trackable in a given environment. As described above, the P_{FT} value applied to the H_0 curve determines the tracking threshold. When the right side of the H_1 curve is integrated starting at this threshold, it produces a value for P_{DT}. If the resultant P_{DT} value is greater than 0.5, we declare the target to be trackable. Of course, the selected values of P_{FT} and P_{DT} to determine trackability are arbitrary — different ones can certainly be selected, but the methodology to determine trackability is the same.

With this in place, trackability for a given target in a given environment can be computed. Results are presented for three different cases in Figures 6, 7, and 8. The inputs for these runs are given in Table II. For the figures, trackability boundaries are shown as a function of probability of target detect in a scan (P_d) vs. clutter density (λ). As is noted in each figure, above the curves is the "trackable" region; below the curves is the "non-trackable" region.

Figure 6 shows the effect of processing Doppler data. It shows trackability for 2-D measurements (azimuth and time-delay) versus several 3-D measurement cases (azimuth, time-delay and Doppler/range-rate). In general, processing Doppler increases trackability; however, an interesting effect is seen as the Doppler error increases. The 3-D curves approach and then actually cross over the 2-D curves — for $\sigma_d = 5$ units/sec, the 2-D case is better than the 3-D case — i.e. *it is better to ignore the Doppler entirely*. This is counterintuitive at first — it would seem that if more information is processed by any algorithm, it should perform better. However, what is happening here is driven by the ratio of the Doppler error to the Doppler volume (from Table II, the Doppler volume is 10 units/sec). With a (normally distributed) Doppler error of 5 units/sec, there is a non-insignificant chance that valid target measurements will have a Doppler component that is outside of the window being processed and thus will be effectively thrown away. (Measurements do not actually get thrown away by the trackability framework, but mathematically, this is exactly the effect.) However, if the Doppler were ignored, these valid target-generated measurements would be used, which increases trackability.

Figure 7 shows the effect of processing an amplitude feature. (Note that all the curves for this plot are 2-D measurements, and $\rho_c = 1$.) First, there is the 2-D curve with no amplitude that serves as a baseline. Now consider the case of both the target and the clutter amplitudes having a Rayleigh distribution. This is the "typical" amplitude likelihood ratio that is present in much of the literature [8], [12]. The amplitudes are normalized so that the clutter measurements have an expected intensity (amplitude squared) of one, and the target measurements have some expected intensity σ^2. An issue with this model is that the Rayleigh distribution is very light-tailed, which causes the amplitude LR ρ_a to get very large as received amplitude increases — essentially the algorithm is "deciding" that received high amplitude measurements are almost certainly from the target and cannot be from clutter. This has the result of increasing predicted trackability. This is why the Rayleigh curve in Figure 7 shows the best trackability performance.

Recent work on received acoustic clutter has posited that the clutter amplitude distribution, instead of being described by a Rayleigh distribution, is better described by a K-distribution [1], [2], and [3]. This is a heavier tailed distribution; its PDF has a parameter α that controls the heaviness of the PDF's tail. Typical values for α range between 0.1 and 50; the smaller the value of α, the heavier the tail of the distribution.

This heavy tail for the clutter PDF has the effect of (probably realistically) deweighting high-amplitude measurements, and thus reducing their effect on the ML-PMHT LLR. This is exactly what is seen in Figure 7. For light-tailed K-distributions — large values of α (i.e. $\alpha \sim 50$) — the K-distributed clutter curve is close to that of the Rayleigh distributed clutter curve. As the value of α gets smaller and the clutter amplitude tail gets heavier, trackability decreases. At the heaviest-tailed level analyzed, $\alpha = 0.1$, the K-distributed curve is practically on top of the 2-D curve using no amplitude information at all. This shows that processing the amplitude likelihood ratio when the clutter amplitude has a very heavy-tailed distribution does not add much to a tracker's performance.

Finally, the effect of a classification feature is shown in Figure 8. Three 2-D cases are examined and compared to the "clean" 2-D case: Case 1, where the clutter classification distribution has a higher variance than the target classification distribution; Case 2, where the clutter and the target classification distributions have equal variances; and finally, Case 3, where the target classification distribution has a larger variance than the clutter classification distribution. The exact mean and variance values of the classification feature distributions for

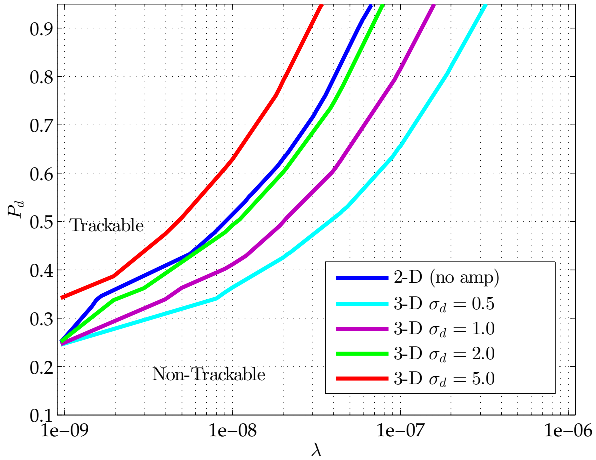

Figure 6. 2-D vs. 3-D trackability as a function of λ vs. P_d.

Table I
VALUES USED FOR THREE CLASSIFICATION CASES

Case	μ_c	σ_c	μ_t	σ_t
1	1	3	2	1
2	1	1	2	1
3	1	1	3	2

the three cases are provided in Table I.

It does not appear as if there is a single factor that drives trackability as a function of classification feature parameters. Intuitively, it would seem that the normalized distance between the clutter and target clutter features (the difference in means divided by the sum total of the variances) would be a large factor in trackability performance. However, Case 1 and Case 2 have a fairly large difference in this metric but have similar trackability curves. Interestingly, what drives the big increase in trackability for Case 3 is the *target* classification feature variance. As this target classification feature variance increases, the support for the transformed single target measurement PDF increases, which in turn spreads the mass of the target batch/peak PDF to the right as well. In contrast, for all three cases, the peak clutter PDFs were fairly similar. Further work will be done to see if the target classification feature variance always has the same effect. Additionally, note that to this point, the classification feature was limited to Gaussian distributions for the target and clutter with *known* variances. Upcoming work will examine trackability in cases when some or all of the parameters in those distributions are not exactly known.

V. CONCLUSIONS

In this work we have presented a framework that is used to determine the trackability of a given target in a given environment for various combinations of measurement dimensionality, amplitude features, and classification features. It uses

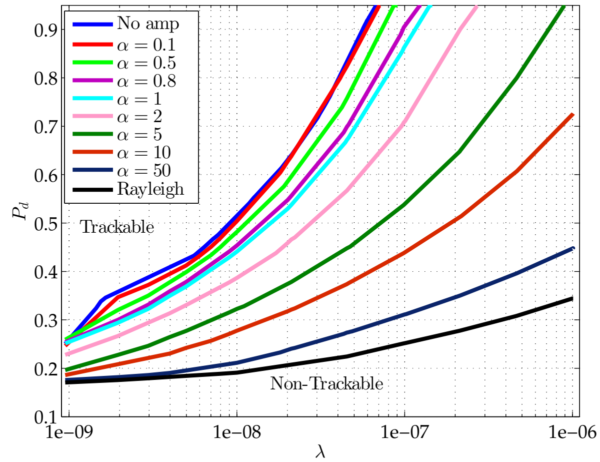

Figure 7. Trackability plot for various amplitude features as a function of λ vs. P_d.

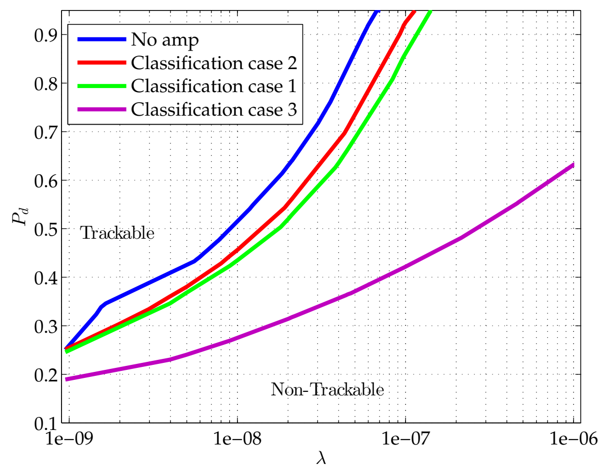

Figure 8. Trackability plot for various classification features as a function of λ vs. P_d. Cases defined in Table I.

Table II
VALUES USED FOR FIGURES 6-8.

2-D and 3-D tracking parameters	
Angular volume	360°
Angular error	5°
Time delay volume	60 sec
Time delay error	0.1 sec
Range-rate volume	10 units per sec
Range-rate error	1 units per sec
Amplitude threshold (K-dist cases)	7 dB
Expected target amplitude (K-dist cases)	10 dB
N_w	11
π_1	0.1512
π_0	0.8488

the assumptions of the ML-PMHT algorithm and extreme-value theory to calculate a PDF that describes the statistics of the peak point in the tracker LLR caused by clutter as well as the PDF that describes the peak point in the tracker LLR caused by the target. With the clutter peak distribution, it is possible to rapidly and accurately determine tracking thresholds for the ML-PMHT algorithm. When the target-due LLR peak distribution is added, it becomes possible, for some desired P_{FT} value, to determine if the target has a P_{DT} level that is greater than some desired level and is thus able to be tracked.

The ML-PMHT algorithm is an optimal algorithm under the assumptions given: the measurement-to-target association model and the parametric trajectory sought. Empirically, the former makes little difference when an ML method is used, and the latter is a good approximation to any trajectory over data batches of reasonable size. Given this, and in light of the fact that it uses all the data available (unlike a "hard association" algorithm such as the MHT) and it needs no simplifying assumptions to work (as opposed to a scan-by-scan tracker such as the JPDAF), it seems reasonable to claim that statements made for the limits of ML-PMHT performance apply to all trackers. That is, it places a rough upper bound on any other tracking algorithm; the framework presented here provides the answer to the question, "Can this target be tracked?"

REFERENCES

[1] D. Abraham, "Detection-threshold approximation for non-Gaussian backgrounds," *IEEE Journal of Oceanic Engineering*, vol. 35, no. 2, pp. 355–365, 2010.

[2] D. Abraham and A. Lyons, "Novel physical interpretations of K-distributed reverberation," *IEEE Journal of Oceanic Engineering*, vol. 27, no. 4, pp. 800–813, 2002.

[3] ——, "Simulation of non-Rayleigh reverberation and clutter," *IEEE Journal of Oceanic Engineering*, vol. 29, no. 2, pp. 347–362, 2004.

[4] D. Avitzour, "A maximum likelihood approach to data association," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 28, no. 2, pp. 560–566, 1996.

[5] Y. Bar-Shalom, P. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. YBS Publishing, 2011.

[6] W. Blanding, P. Willett, and Y. Bar-Shalom, "Offline and real-time methods for ML-PDA track validation," *IEEE Transactions on Signal Processing*, vol. 55, no. 5, pp. 1994–2006, 2007.

[7] E. Castillo, *Extreme Value Theory in Engineering*. Boston: Harcourt Brace Jovanovich, 1988.

[8] M. R. Chummun, Y. Bar-Shalom, and T. Kirubarajan, "Adaptive early-detection ML-PDA estimator for LO targets with EO sensors," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 2, pp. 694–707, 2002.

[9] S. Coles, *An Introduction to Statistical Modeling of Extreme Values*. London: Springer-Verlag, 2001.

[10] E. Gumbel, *Statistics of Extremes*. New York: Columbia University Press, 1958.

[11] C. Jauffret and Y. Bar-Shalom, "Track formation with bearing and frequency measurements in clutter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 26, no. 6, pp. 999–1010, 1990.

[12] T. Kirubarajan and Y. Bar-Shalom, "Low observable target motion analysis using amplitude information," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 32, no. 4, pp. 1637–1382, 1996.

[13] G. Moustakides, G. Jajamovich, A. Tajer, and X. Wang, "Joint detection and estimation: Optimum tests and applications," *IEEE Transactions on Information Theory*, vol. 58, no. 7, pp. 4215–4229, 2012.

[14] N. Mukhopadhyay, *Probability and Statistical Inference*. New York: Marcel Dekker, 2000.

[15] J. Neyman and E. Pearson, "On the problem of the most efficient tests of statistical hypotheses," *Philosophical Transactions of the Royal Society of London*, vol. Series A, Containing Papers of a Mathematical and Physical Character, no. 232, pp. 289–337, 1933.

[16] S. Schoenecker, P. Willett, and Y. Bar-Shalom, "A comparison of the ML-PDA and the ML-PMHT algorithms," in *Proceedings of the 14th International Conference on Information Fusion*, Chicago, IL, 2011.

[17] ——, "The ML-PMHT multistatic tracker for sharply maneuvering targets," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 4, pp. 2235–2249, 2013.

[18] ——, "The effect of K-distributed clutter on trackability," *Submitted to IEEE Transactions on Signal Processing*, 2014.

[19] ——, "Extreme-value analysis for ML-PMHT, part 1: Threshold determination for false track probability," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 4, pp. 2500–2514, 2014.

[20] ——, "ML-PDA and ML-PMHT: Comparing multistatic sonar trackers for VLO targets using a new multitarget implementation," *Journal of Oceanic Engineering*, vol. 39, no. 2, pp. 303–317, 2014.

[21] ——, "ML-PMHT track detection threshold determination for K-distributed clutter," in *Proceedings of the SPIE Conference on Signal and Data Processing of Small Targets, #9092-23*, Baltimore, MD, 2014.

[22] S. Schoenecker, P. Willett, Y. Bar-Shalom, and T. Luginbuhl, "Extreme-value analysis for ML-PMHT, part 2: Probability of target track detection and the fundamental trackability of targets," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 4, pp. 2515–2527, 2014.

[23] R. Streit and T. Luginbuhl, "A probabilistic multi-hypothesis tracking algorithm without enumeration," in *Proceedings of the 6th Joint Data Fusion Symposium*, Laurel, MD, June 1993.

[24] ——, "Maximum likelihood method for probabilistic multi-hypothesis tracking," in *Proceedings of the SPIE Conference on Signal and Data Processing of Small Targets, #2235*, Orlando, FL, 1994.

[25] ——, "Probabilistic multi-hypothesis tracking," Naval Undersea Warfare Center, Tech. Rep. TR 10428, 1995.

[26] P. Willett and S. Coraluppi, "MLPDA and MLPMHT applied to some MSTWG data," in *Proceedings of the 9th International Conference on Information Fusion*, Florence, Italy, July 2006.

[27] P. Willett, S. Coraluppi, and W. Blanding, "Comparison of soft and hard assignment ML trackers on multistatic data," in *Proceedings of the IEEE Aerospace Conference*, Big Sky, March 2006.

[28] O. Zeitouni, J. Ziv, and N. Merhav, "When is the generalized likelihood ratio test optimal?" *IEEE Transactions on Information Theory*, vol. 48, no. 5, pp. 1597–1062, 1992.