Standard Bayesian Approach to Quantized Measurements and Imprecise Likelihoods

LAWRENCE D. STONE
STEPHEN L. ANDERSON

In this paper we show that the standard definition of likelihood function used in Bayesian inference simply and correctly handles quantized measurements and imprecise likelihood functions. Some recent papers have stated or implied that methods involving random sets, fuzzy membership functions, generalized likelihood functions, or Dempster-Shafer concepts are required to handle imprecise likelihood functions and quantized measurements. While it is true that one can use these methods, in the spirit of Occam's razor, we feel the simplest correct solution is the best.

Manuscript received on December 19, 2013; revised April 10, 2014 and April 22, 2014; released for publication May 11, 2014.

Refereeing of this contribution was handled by Jean Dezert.

This is an extended version of the paper "Standard Bayesian approach to quantized measurements and imprecise likelihoods" published at the 16th International Conference on Information Fusion (Fusion 2013).

Authors' address: Metron Inc., 1818 Library Street, Reston Virginia, 20190 (e-mail: stone@metsci.com and Anderson@metsci.com).

1 INTRODUCTION

Some recent papers have stated or implied that methods involving random sets, fuzzy membership functions, generalized likelihood functions, or Dempster-Shafer concepts are required to handle quantized measurements and imprecise likelihood functions. In particular, reference [3] considers the problem of constructing likelihood functions for quantized measurements and proposes that these types of measurements require a generalization of the standard notion of likelihood function that involves the use of random sets, concepts from fuzzy logic and Dempster-Shafer theory, as well as generalized or imprecise likelihood functions. Similarly, reference [4] presents examples of problems where the author claims that imprecise likelihood functions (a generalization of standard likelihood functions) are required.

The purpose of this paper is to show that the standard concept of likelihood function as defined in [1] or [6] is sufficient to solve the problems presented in [3] and [4] in an easy and straightforward manner. This is an important point because in the spirit of Occam's razor, we believe the simplest correct solution to a problem is the best one. Simplicity allows readers to clearly understand the nature of the problem and its solution. It facilitates the use of a concept in applications and makes it easier to extend it to more challenging problems. It enables progress.

This suggests the following question which we pose but do not presume to answer here: What situations involving quantized measurements or imprecise likelihood functions *require* the use of alternate, non-Bayesian models of uncertainty?

Section 2 of this paper presents the standard Bayesian inference formulation. Section 3 shows how this formulation can be used to handle the quantized measurement examples presented in [3]. Section 4 shows that the examples given in [4] can be readily handled using standard Bayesian likelihood functions and that a generalization to imprecise likelihood functions is not required for these examples.

Reference [2] investigates the problem of tracking a target with quantized measurements. The authors assume a Gaussian motion model for the target and develop an approximate Minimum Mean Squared Error (MMSE) solution. They provide a numerical algorithm for obtaining this solution. This is very impressive work, and one must admire the authors for the cleverness of their solution. However, the solution is complex and does require many special assumptions. In contrast to this we present in Section 5 a particle filter approach to solving this problem with standard Bayesian likelihood functions that is very simple and general. One is not constrained to Gaussian motion models or measurement errors. It is straight-forward to incorporate a wide variety of types of measurements and motion models. We present an example to illustrate this approach.

2 BAYESIAN INFERENCE FORMULATION

In order to clarify what we mean by a standard Bayesian approach, we give the formulation of the basic Bayesian inference problem that is presented in [6] and is consistent with that in [1].

There is an unknown parameter Θ that we wish to estimate. There is a prior distribution p_0 on Θ such that

$$p_0(\theta) = \Pr\{\Theta = \theta\} \tag{1}$$

where Pr indicates either probability or probability density as appropriate. We obtain a measurement Z from a sensor. The measurement is viewed as a random variable whose distribution depends on θ. We define the likelihood function

$$l(z \mid \theta) = \Pr\{Z = z \mid \Theta = \theta\}. \tag{2}$$

If we receive a measurement $Z = z$, we compute the posterior distribution

$$p_1(\theta \mid Z = z) = \frac{l(z \mid \theta)p_0(\theta)}{\int l(z \mid \theta')p_0(\theta')d\theta'} \tag{3}$$

where integration is replaced by summation if the distribution on Θ is discrete.

Note, when we use the term likelihood function to describe (2), we mean the function $l(z \mid \cdot)$ obtained by holding the measurement z fixed and letting the parameter θ vary. $l(z \mid \cdot)$ need not be a probability (density) function. It may integrate to a number different than 1. We use the notation l in place of the more usual p to emphasize this point.

3 QUANTIZED MEASUREMENTS

Reference [3] sought to illustrated the necessity of its approach by presenting examples of performing inference using quantized measurements. In this section, we use the same examples to show that standard likelihood functions and the Bayesian inference process as given in (1)–(3) provide a straight-forward and correct way of incorporating quantized measurements into Bayesian inference. No generalization is required, and no extensions of the standard Bayesian probability concepts are needed.

In the digital voltmeter example given in [3], measurements are taken by a digital voltmeter that provides voltage readings to two decimal places. From the digital voltmeter measurement, we wish to estimate the actual voltage Θ. Let p_0 be the prior on Θ. We consider three cases, measurements without noise, measurements with noise, and measurements where the quantization is unknown.

3.1 Quantized Measurements without noise—known quantization

If there is no noise added to the actual voltage, then any voltage in the interval $(199.975, 199.985]$ will produce a measurement of 199.98. The measurement space is a discrete set of points on the real line of the form $j \times 0.0$ where j is an integer such that $-\infty < j < \infty$. Any voltage in the set

$$S_j = (j \times 0.01 - 0.005, j \times 0.01 + 0.005] \tag{4}$$

will produce a measurement $Z = j \times 0.01$. From the definition of likelihood function in (2), we have

$$\begin{aligned} l(j \times 0.01 \mid \theta) &= \Pr\{Z = j \times 0.01 \mid \theta\} \\ &= \begin{cases} 1 & \text{if } \theta \in S_j \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \tag{5}$$

For notational convenience, we shall use $Z = j$ for the measurement and $l(j \mid \theta)$ for the likelihood function in (5).

The posterior on the actual voltage Θ is computed by,

$$\begin{aligned} p_1(\theta \mid Z = j) &= \frac{l(j \mid \theta)p_0(\theta)}{\int l(j \mid \theta')p_0(\theta')d\theta'} \\ &= \begin{cases} \dfrac{p_0(\theta)}{\int_{S_j} p_0(\theta')d\theta'} & \text{if } \theta \in S_j \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \tag{6}$$

3.2 Quantized Measurements with Noise—Known Quantization

In this example we suppose the received voltage r at the digital voltmeter is the true voltage θ plus noise ε. Specifically, the true voltage is $\theta + \varepsilon$ where ε has the density function

$$f(y) = \Pr\{\varepsilon = y\} \quad \text{for } -\infty < y < \infty.$$

The digital voltmeter produces measurements to two decimal places as above. In this case the likelihood function becomes

$$\begin{aligned} l(j \mid \theta) &= Pr\{Z = j \mid \theta\} \\ &= \Pr\{j \times 0.01 - 0.005 < \theta + \varepsilon \\ &\qquad \le j \times 0.01 + 0.005\} \\ &= \Pr\{j \times 0.01 - \theta - 0.005 < \varepsilon \\ &\qquad \le j \times 0.01 - \theta + 0.005\} \\ &= \int_{j \times 0.01 - \theta - 0.005}^{j \times 0.01 - \theta + 0.005} f(y)dy. \end{aligned} \tag{7}$$

and the posterior distribution on Θ given the measurement $Z = j$ is

$$\begin{aligned} p_1(\theta \mid Z = j) &= \frac{p_0(\theta)l(j \mid \theta)}{\int p_0(\theta')l(j \mid \theta')d\theta'} \\ &= \frac{p_0(\theta)\int_{j \times 0.01 - \theta + 0.005}^{j \times 0.01 - \theta + 0.005} f(y)dy}{\int p_0(\theta')\left(\int_{j \times 0.01 - \theta' - 0.005}^{j \times 0.01 - \theta' + 0.005} f(y)dy\right)d\theta'} \end{aligned}$$

As an example, let us consider the situation where

$$f(y) = \eta(y, 0, \sigma^2)$$

where $\eta(\cdot,0,\sigma^2)$ is the probability density function for the normal distribution with mean 0 and variance σ^2. The notation $\eta(\cdot,0,\sigma^2)$ is used to indicate the function of one variable obtained by fixing the values of the 2nd and 3rd variables at 0 and σ^2. We use a similar notation for a function of two variables when we wish to fix the value of one of the variables. Let

$$\Phi(z,\sigma^2) = \int_{-\infty}^{z} \eta(y,0,\sigma^2)dy \quad \text{for } -\infty < z < \infty. \quad (8)$$

Then the likelihood function $l(j \mid \theta)$ in (7) becomes

$$l(j \mid \theta) = \Phi(j \times 0.01 - \theta + 0.005, \sigma^2) - \Phi(j \times 0.01 - \theta - 0.005, \sigma^2). \quad (9)$$

Figure 1 shows plots of the likelihood function in (9) for $Z = 10$ and $\sigma^2 = 0.0001$, 0.0004, and 0.0016.

Reference [4] also presents a quantized measurement example that the author claims requires the use of imprecise likelihood functions. In fact one can produce the example in [4] and obtain Figure 1 in [4] if he considers the case where the quantization has bin-size 15 and the measurement $Z = j$ indicates the interval $(15j, 15(j+1)]$. In this case the sets S_j in (4) become

$$S_j = (15j, 15(j+1)],$$

and the likelihood function in (9) becomes

$$l(j \mid \theta) = \Phi(j \times (15+1) - \theta, \sigma^2) - \Phi(j \times 15 - \theta, \sigma^2)$$

where θ plays the role of the variable z in [4].

3.3 Quantized Measurements when Quantization is Unknown

In the case where the quantization is unknown, we expand the state space on which we perform inference to simultaneously estimate the voltage and the quantization. To illustrate how this is done within the conventional Bayesian inference formalism, we consider the case where there is no noise added to the voltage and the quantized bins have a known and equal size. However, we do not know the anchor point for the bins.

Following Example 1, we take the bin size to be 0.01. However we do not know the anchor point of the bins. Specifically there is a unknown parameter Δ such that $-0.005 \le \Delta \le 0.005$ and

$$S_j(\Delta) = (j \times 0.01 - 0.005 + \Delta, j \times 0.01 + 0.005 + \Delta].$$

The inference problem is to estimate both θ and Δ.

In classic Bayesian fashion, we impose a prior distribution on (θ,Δ) which represents our prior knowledge (or uncertainty) about (θ,Δ). As an example we suppose that the priors on the two parameters are independent and the joint density on (θ,Δ) is given by

$$g_0(\theta,\delta) = p_0(\theta)q_0(\delta) \quad \text{for } -\infty < \theta < \infty \ -0.005 < \delta \le 0.005.$$

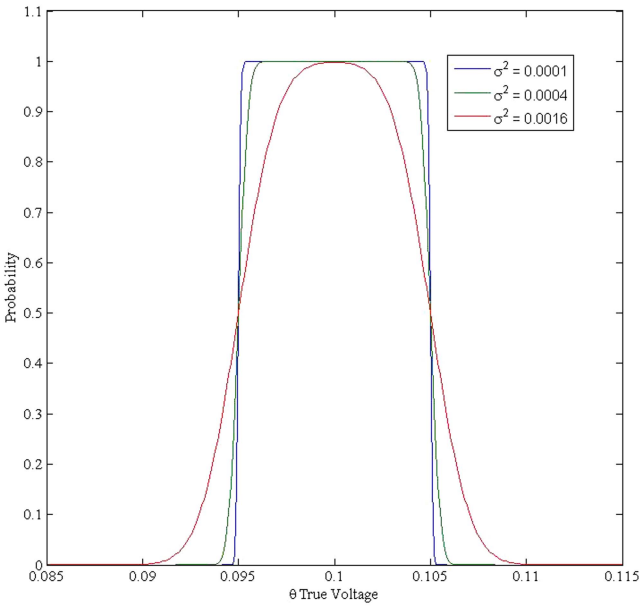

Fig. 1. Likelihood functions for $j = 10$ (0.1 volt reading on voltmeter) when $\sigma^2 = 0.0001$, 0.0004, and 0.0016.

The likelihood function for the observation $Z = j \times 0.01$ is

$$l(j \mid (\theta,\delta)) = \Pr\{Z = j \times 0.01 \mid (\Theta,\Delta) = (\theta,\delta)\} = \begin{cases} 1 & \text{if } \theta \in S_j(\delta) \\ 0 & \text{otherwise.} \end{cases}$$

Let g_1 be the posterior joint density on (θ,Δ) given $Z = j$. Then

$$g_1((\theta,\delta) \mid Z = j) = \frac{l(j \mid (\theta,\delta))p_0(\theta)q_0(\delta)}{\int_{-0.005}^{0.005}\int l(j \mid (\theta',\delta'))p_0(\theta')q_0(\delta')d\theta'd\delta'} = \begin{cases} \dfrac{p_0(\theta)q_0(\delta)}{\int_{-0.005}^{0.005}\int_{S_j(\delta')} p_0(\theta')q_0(\delta')d\theta'd\delta'} & \text{for } \theta \in S_j(\delta) \\ 0 & \text{otherwise.} \end{cases}$$

3.4 Alternate Quantization Models

Digital signal processing involves quantized measurements. The effect of this quantization is sometimes modeled as adding random and independent noise to the measurements. This is discussed in [7] which notes "It has been shown to be a valid model in cases of high resolution quantization (small Δrelative to the signal strength) with smooth probability density functions. However, additive noise behaviour is not always a valid assumption, and care should be taken to avoid assuming that this model always applies. In actuality, the quantization error…is deterministically related to the signal rather than being independent of it."

4 IMPRECISE LIKELIHOOD FUNCTIONS

In this section we show that the localization example given in IV of [4] can be computed in a straight-

forward way by standard Bayesian likelihood functions without resorting to imprecise likelihood functions or other generalizations of Bayesian inference.

For the convenience of the reader we reproduce the description of the example given in [4].

4.1 Localization using RSS

Received signal strength (RSS) is often used for localizing an emitting energy source although the source level of the emitter is unknown. As an example, consider the two-dimensional situation shown in Figure 2 where the unknown source position $X = (X_1, X_2)$ is located inside the square defined by

$$5 \le X_1 \le 95; \quad 5 \le X_2 \le 95, \tag{10}$$

and the prior on distribution on (X_1, X_2) is uniform over this square. The unknown source level A (in dB) of the emitter has a uniform distribution on $[25, 65]$. There are 12 receivers uniformly spaced on a circle of radius 50 centered at $(50, 50)$ as shown by the squares in Figure 2. The receivers are numbered in counter-clockwise fashion starting with receiver 1 at the 3 o'clock position. Let (x_1^i, x_2^i) be the location of the ith receiver.

For $i = 1, \ldots, 12$, the measurement Z_i of RSS at the ith receiver satisfies the following equation

$$Z_i = A - 10\theta_i \log(d_i(x)/d_0) + v_i \tag{11}$$

where

$$d_i(x) = \sqrt{(x_1^i - x_1)^2 + (x_2^i - x_2)^2}$$

is the distance from the location of the ith receiver to the source given the source is located at $x = (x_1, x_2)$,

θ_i is an unknown propagation loss factor where $2 \le \theta_i \le 4$,

$d_0 = 10$ is a reference distance,

v_i has a Gaussian distribution with mean 0 and variance 4,

v_i is independent of v_j for $i \ne j$.

4.1.1 Likelihood Function

Reference [4] does not provide an explicit formula for the likelihood function employed in this example. For our computations, we assume that θ_i is uniformly distributed over $[2.0, 4.0]$ for each receiver and that θ_i is independent of θ_j for $i \ne j$ and obtain an explicit likelihood function as follows.

Let η_{04} be the density function for a Gaussian distribution with mean 0 and variance 4. Then

$$\begin{aligned} &\Pr\{Z_i = z \mid X = x, A = a, \text{ and } \theta_i = \theta\} \\ &\quad = \eta_{04}(z - a + 10\theta \log(d_i(x)/d_0)) \end{aligned}$$

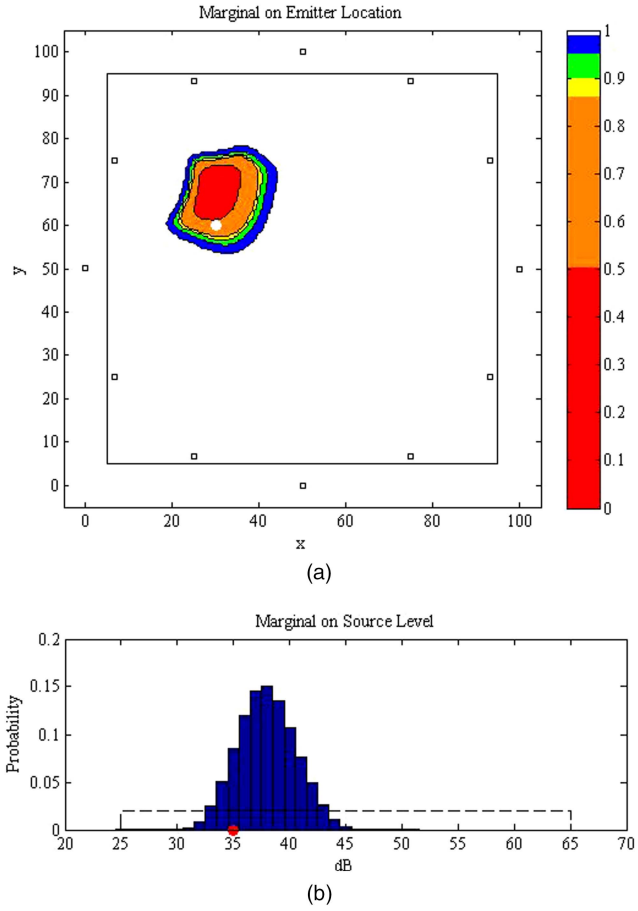

Fig. 2. Marginal distributions on position and source level.

and

$$\begin{aligned} l(z \mid x, a) &= \Pr\{Z_i = z \mid X = x, A = a\} \\ &= \frac{1}{2}\int_2^4 \eta_{04}(z - a + 10\theta \log(d_i(x)/d_0))d\theta \\ &= \frac{\int_{20\log(d_i(x)/d_0)}^{40\log(d_i(x)/d_0)} \eta_{04}(z - a + y)dy}{20\log(d_i(x)/d_0)} \\ &= \frac{\Phi(z - a + 40\log(d_i(x)/d_0), 4)}{20\log(d_i(x)/d_0)} \\ &\quad - \frac{\Phi(z - a + 20\log(d_i(x)/d_0), 4)}{20\log(d_i(x)/d_0)} \end{aligned} \tag{12}$$

is the likelihood for the measurement $Z_i = z$ given $X = x$ and $A = a$.

4.1.2 Results

Following [4] we simulated measurements at the 12 receivers using the model in (11) for a source located at $X = (30, 60)$ with source level $A = 35$. We set $\theta_i = 2.3$ for $i = 1, \ldots, 6$ and $\theta_i = 3.5$ for $i = 7, \ldots, 12$. Using the resulting measurements and the likelihood function in (12), we computed the joint posterior distribution on emitter position X and source level A. Figure 2 shows the resulting marginals on X and A. For numerical convenience we computed the distributions on a grid that

has 200 by 200 cells in position and 41 cells in source level. We calculated the posterior probability in each of these cells given the measurements from the 12 sensors. The color bar next to the position marginal indicates containment. The red region is the 50% containment region (i.e., the region with the smallest number of cells that contains 50% probability). The red plus orange region is the 86% containment region, and so on for the yellow, green, and blue regions. Observe that the emitter's location (shown as a white dot in Figure 2) is close to edge of the 50% containment region. The marginal on source level is represented by a bar graph with the height of the bar being equal to the posterior probability in the cell containing the bar. The cells are 1 dB in width. The actual source level, 35 dB, is in a reasonable location in this distribution.

In order to test whether the above method produces a good representation of the uncertainty in the posterior estimate of the position X of the emitter, we followed the approach of Section V in [4] and simulated 1000 sets of measurements at the 12 receivers. For each replication of the simulation, we made an independent draw for the value of θ_i from a uniform distribution over $[2.0, 4.0]$ and computed the resulting measurement Z_i from (11) for $i = 1, \ldots, 12$. The location of the emitter and source level remained fixed at $X = (30, 60)$ and $A = 35$ for all replications. For each replication, we computed the marginal on emitter position as in Figure 2. To test whether the resulting distributions correctly represent the uncertainty in the location of the emitter, we computed the following Kolmogorov-Smirnov (KS) graph.

For the nth replication, we started with the cell containing the emitter (the one containing the point $(30, 60)$) and summed the probability in all cells having probability greater than or equal to the probability in the emitter cell. This produced a containment region and containment probability c_n for $n = 1, \ldots, 1000$. Next we ordered the containment probabilities into a set $\{\hat{c}_n; n = 1, \ldots, 1000\}$ such that $0 \le \hat{c}_1 \le \hat{c}_2 \ldots \le \hat{c}_{1000} \le 1$. We then plotted the points $(n/1000, \hat{c}_n)$ for $n = 1, \ldots, 1000$. This plot is shown in blue in Figure 3 which is the empirical distribution for the containment probability produced by the marginal distribution on position. If the marginal distributions accurately represent the uncertainty in emitter location, this distribution should converge to the red straight line in Figure 3 as the number of replication increases to infinity. That is, the percentage of replications in which the target is inside the p percent containment region should be p percent. One can see that the empirical distribution is indeed a close fit to the straight line indicating an accurate representation of the uncertainty in the marginal. One could perform a KS test to test the hypothesis that the empirical distribution is the same except for sampling noise as the straight line. However, it is clear from Figure 3 that the fit is very good.

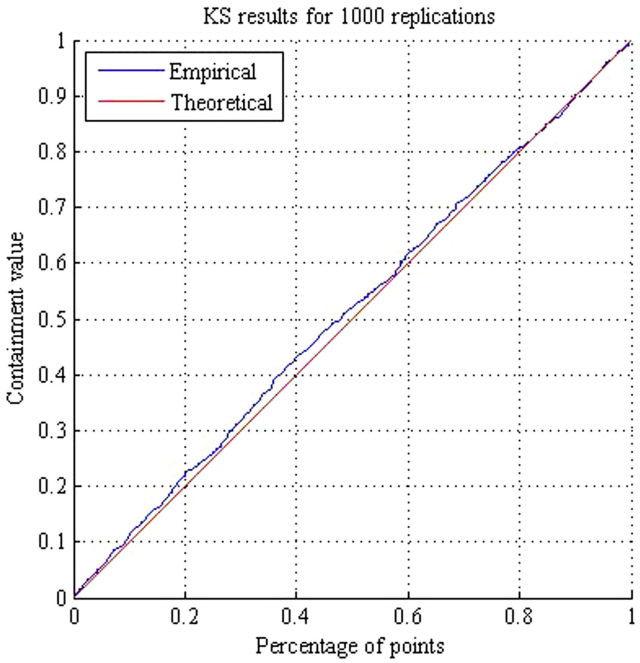

Fig. 3. KS graph for marginal distribution on emitter location.

To further test the accuracy of the marginal posterior on location, we performed another set of 1000 replications similar to the ones above. For each of these replications, we made an independent draw from a uniform distribution over the square defined by (10) for the location of the emitter and an independent draw from a uniform distribution over $[25, 65]$ for the source level. As above, we made independent draws for the values of $\theta_i$for $i = 1, \ldots 12$. We then computed the KS graph for the posterior marginal distribution on emitter location. The result looked very similar to Figure 3 further confirming the accuracy of these posterior distributions.

5 TRACKING MOVING TARGETS WITH QUANTIZED MEASUREMENTS

The target considered in Section 4 is stationary. One can also track moving targets using sensors with quantized measurements. Since the measurements do not satisfy the linear-Gaussian assumptions required for a Kalman filter, we perform the tracking using a particle filter as described in [5] or Chapter 3 of [6].

The general procedure is straightforward. The particles are motion updated to the time of a measurement. The likelihood function for a quantized measurement is applied to the weight of each particle to produce the posterior distribution on target state. The particles are resampled and then motion updated to the time of the next measurement.

In particular, suppose that the target state distribution at time t is represented by the set of particles

$$\{(x_n(t), w_n(t)\} \quad \text{for } n = 1, \ldots, N$$

where $x_n(t)$ is the state of the nth particle at time t and w_n is its probability. This set of particles represents a discrete probability approximation to the distribution on target state at the time t. Suppose we obtain a quantized measurement $Z = z$. Let $l(z \mid x)$ be the likelihood function for this measurement. The posterior distribution on target state at time t is given by

$$\{(x_n(t), \tilde{w}_n(t)\} \quad \text{for } n = 1, \ldots, N \tag{13}$$

where

$$\tilde{w}_n(t) = \frac{l(z \mid x_n(t))w_n(t)}{\sum_{n'=1}^{N} l(z \mid x_{n'}(t))w_{n'}(t)}.$$

If the next measurement is received at time $t' > t$, the posterior particle filter representation in (13) can be motion-updated to the time t' to act as a proposal distribution for the incorporation of the measurement at time t'. The posterior in (13) is typically resampled before the motion update is performed. If desired, other proposal distributions can used to improve particle filter performance as discussed in [5].

In the following example we consider an underwater acoustic detection situation that goes beyond the quantized measurement examples considered in [3]. In particular, the bins are unions of disjoint intervals. In this example we employ a particle filter to track a moving target.

5.1 Likelihood Function for Acoustic Detection

When a passive acoustic sensor is located in a deep water region of the ocean, the sound propagation conditions often produce detection areas that are disjoint. For example, there may be good detection conditions from the sensor's location out to range 5 nm. This is typically called the direct path region. In addition there are often convergence zone regions at ranges of roughly 30 nm, 60 nm, and even farther out. A convergence zone is a region where the acoustic rays converge and produce low propagation loss and increased detection probability for the sensor. Suppose the convergence zones are 5 nm wide. It is often the case that the uncertainty about the source level of a potential target means that although one cannot calculate the detection probability as a function of range, one does know that if a target has been detected, it is in one of these zones. In this case, a detection means that the target is in one of the above range intervals, i.e., its range is in the union of disjoint intervals

$$S = [0,5] \cup [27.5, 32.5] \cup [57.5, 62.5]. \tag{14}$$

Generally one does not know the edges of the intervals in S exactly. Depending on the source level of the target and the ambient noise in the ocean, these areas can be a bit larger or smaller than the nominal numbers in (14). We will model this uncertainty with a likelihood function that is similar to the one given in the example in Section 3.2 with the exception that there is only one bin corresponding to a detection, which we denote by $Z = 1$. Specifically we let r denote the range of the target and ε_i be mutually independent normally distributed random variables with mean 0 and variance σ_i^2 for $i = 1,2,3$. Then the likelihood function l_d for a detection becomes

$$\begin{aligned} l_d(1 \mid r) &= \Pr\{Z = 1 \mid \text{target at range } r\} \\ &= \Pr\left\{ \begin{array}{c} r \le 5 + \varepsilon_1 \text{ or } 27.5 - \varepsilon_2 \le r \le 32.5 + \varepsilon_2 \\ \text{or } 57.5 - \varepsilon_3 \le r \le 62.5 + \varepsilon_3 \end{array} \right\} \\ &\approx \begin{cases} \Pr\{r \le 5 + \varepsilon_1\} & \text{for } 0 \le r \le 20 \\ \Pr\{27.5 - \varepsilon_2 \le r \le 32.5 + \varepsilon_2\} & \text{for } 20 < r \le 45 \\ \Pr\{57.5 - \varepsilon_3 \le r \le 62.5 + \varepsilon_3\} & \text{for } 45 < r < \infty. \end{cases} \end{aligned} \tag{15}$$

The approximation in the last line is essentially an equality if $\sigma_i^2 < 4$ for $i = 1,2,3$. In terms of Φ defined in (8), the likelihood function in (15) becomes

$$l_d(1 \mid r) = \begin{cases} 1 - \Phi(r - 5, \sigma_1^2) & \text{for } 0 \le r \le 20 \\ \min\{1 - \Phi(27.5 - r, \sigma_2^2), 1 - \Phi(r - 32.5, \sigma_2^2)\} & \\ & \text{for } 20 < r \le 45 \\ \min\{1 - \Phi(57.5 - r, \sigma_3^2), 1 - \Phi(r - 62.5, \sigma_3^2)\} & \\ & \text{for } 45 < r < \infty. \end{cases} \tag{16}$$

Figure 4 shows the likelihood function $l_d(1 \mid \cdot)$ when $\sigma_i^2 = 0.5$ for $i = 1,2,3$.

5.2 Acoustic Tracking Example

For this example we consider a target moving at 14 kn in a 70 nm by 70 nm square as shown in Figure 5. There are six sensors located at $(0,0)$, $(0,35)$, $(0,70)$, $(70,0)$, $(70,35)$, and $(70,70)$. These are shown as white dots in Figure 5. Each of these sensors has the detection characteristics described in Section 5.1 and in particular has the detection likelihood function given by and plotted in Figure 4. There is an independent detection opportunity for each of these sensors every 5 minutes. The detection opportunities occur simultaneously for all sensors. Note that the detections produce range information but no bearing information.

In order to show the coverage of the sensor field over the 70 nm $\times$ 70 nm square, we computed the sum of the likelihood functions for the six sensors evaluated at each point in the square. This sum yields the expected number of sensors that can detect a target at each point in the square. The results are color-coded and plotted on the square in Figure 5. The color code is given by the color bar on the right. Note, this is not a calculation of the likelihood function resulting from the detections at an opportunity time. That likelihood function is obtained by pointwise *multiplying* the likelihood functions for the sensors obtaining a detection at that time.

For the example, the target follows the track shown in white in Figure 5 moving at a constant 14 kn from the

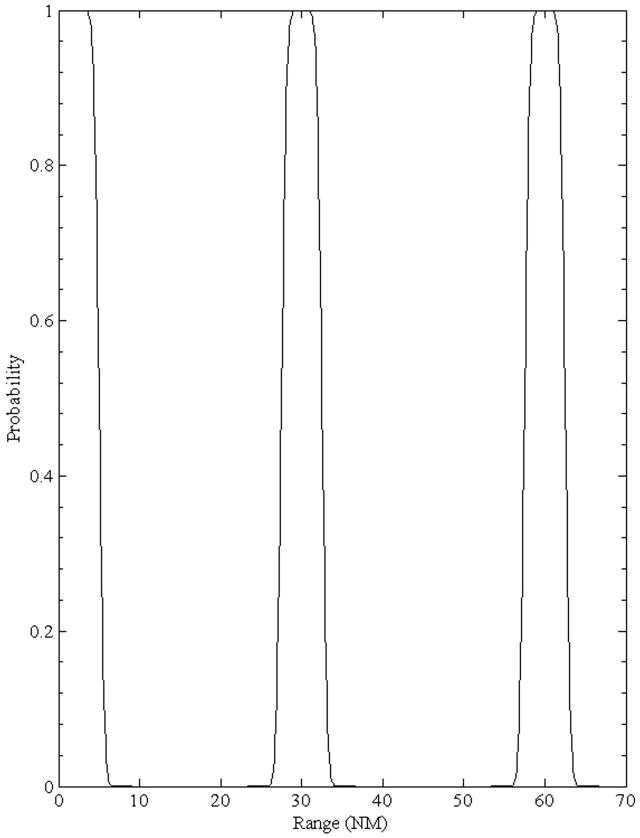

Fig. 4. Likelihood function $l_d(1 \mid \cdot)$ for accoustic detection in a convergence zone environment.

bottom to the top of the figure. The maneuver occurs at 2.5 hours.

Using the likelihood function in (16) and a particle filter that is a minor modification of the one described in Section 1.3 of [6], we estimated the track of this moving target. We simulated detections as follows. At each opportunity time, the simulation calculated the range of the target from each sensor. A detection was called with probability equal to the likelihood function value at that range. Detections are independent from sensor to sensor.

We used 25,000 particles. The particle paths were initialized from the first sensor detection as follows. For each particle, the range was randomly drawn from a uniform distribution on the convergence zone intervals with a Gaussian component having mean 0 and variance 0.5 nm added. The bearing from the sensor was chosen uniformly over the interval 0 to 360 degrees. By doing this we obtained 25,000 equally weighted independent points from the posterior distribution on target position given the first detection.

The initial speed for each particle was drawn from a uniform distribution on the interval 2 to 30 kn. The initial course was drawn from a uniform distribution over the interval 0 to 360 degrees. The particles change velocity according to an exponential distribution with mean 0.5 hours. When a velocity change takes place, a new velocity is chosen from a distribution that produces

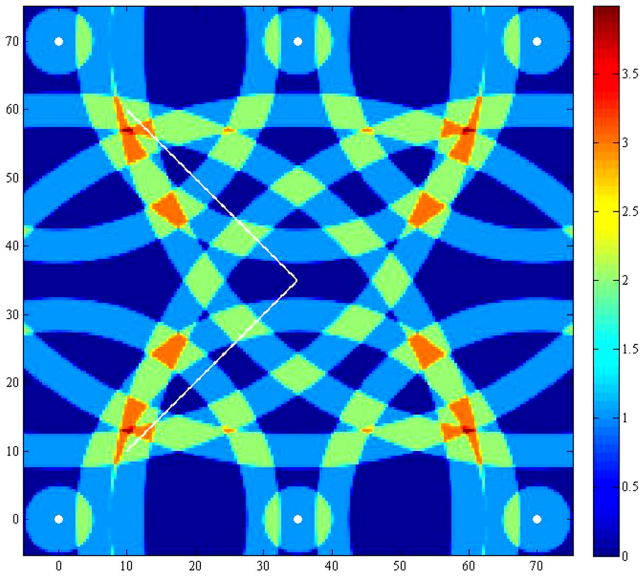

Fig. 5. Sensor field and target track.

a mean change of 30 degrees in course and 2 kn in speed. See Section 1.3.3 of [6] for the details of this motion model.

Figures 6 and 7 show the filter output at 0.5 and 2.0 hours. In the figures, we show only a 500 point sample of the particles in the figure. At 0.5 hours, the distribution has two modes. By 2.0 hours the second mode has disappeared, and the remaining particle cloud is centered at the target's location.

Figures 8 and 9 show the filter output at 3 hours and 5 hours. At 3 hours, just after the maneuver, the particle distribution has spread out. By 5 hours it has condensed around the position of the target.

This example shows two things. First that even complicated quantized measurements can be represented by standard likelihood functions, and second these likelihood functions can be easily used in a particle filter to track moving targets.

6 CONCLUSIONS

The examples have shown how to construct likelihood functions for quantized measurements using the standard Bayesian approach with standard likelihood functions.

In Section 3, we showed that the quantized measurement examples presented in [3] can be treated without employing the notions of generalized likelihood functions.

In Section 4 we have shown that the examples presented in [4] can be handled with standard Bayesian likelihood functions without the extra complexity required for imprecise likelihood functions. If anything, the results presented in Section 4.1.2 are somewhat better than the ones obtained by the use of imprecise likelihood functions in [4].

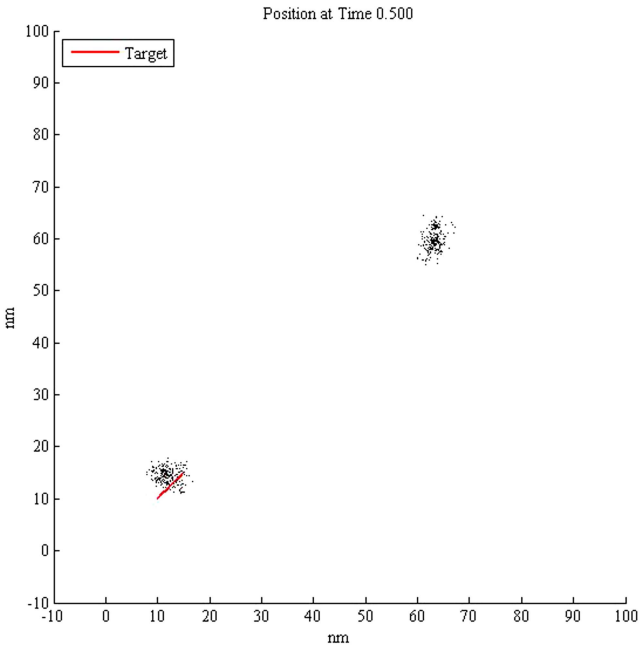

Fig. 6. Position marginal at 0.5 hr.

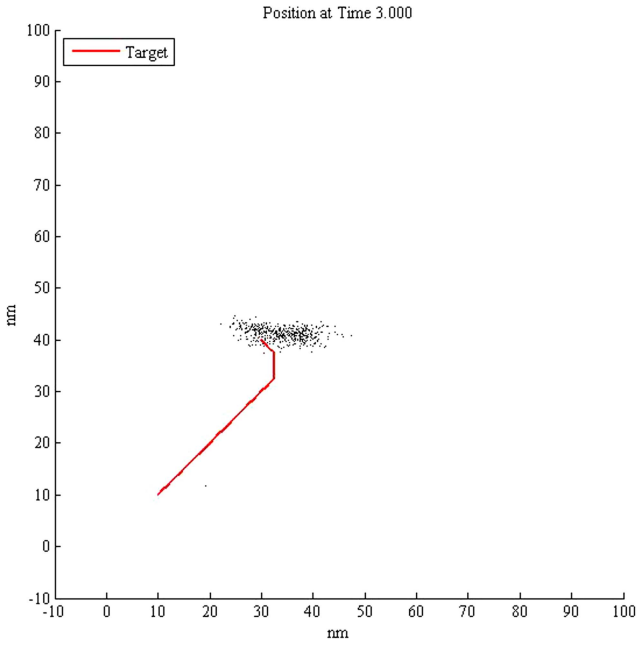

Fig. 8. Postion marginal at 3 hr.

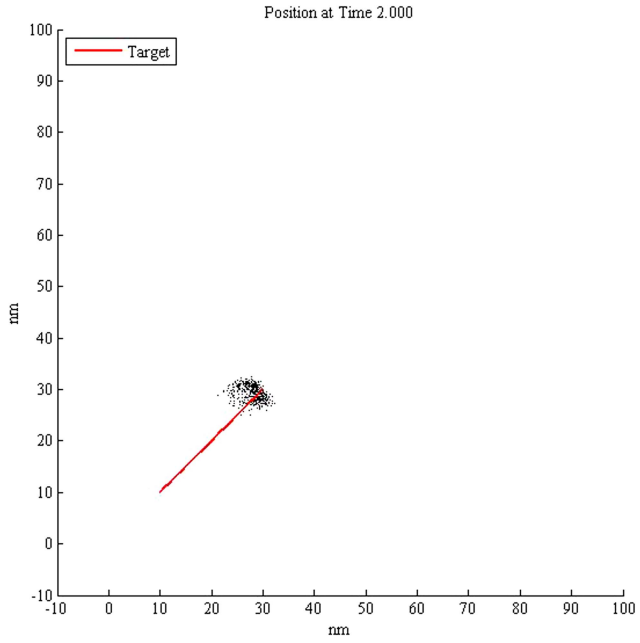

Fig. 7. Position marginal at 2.0 hr.

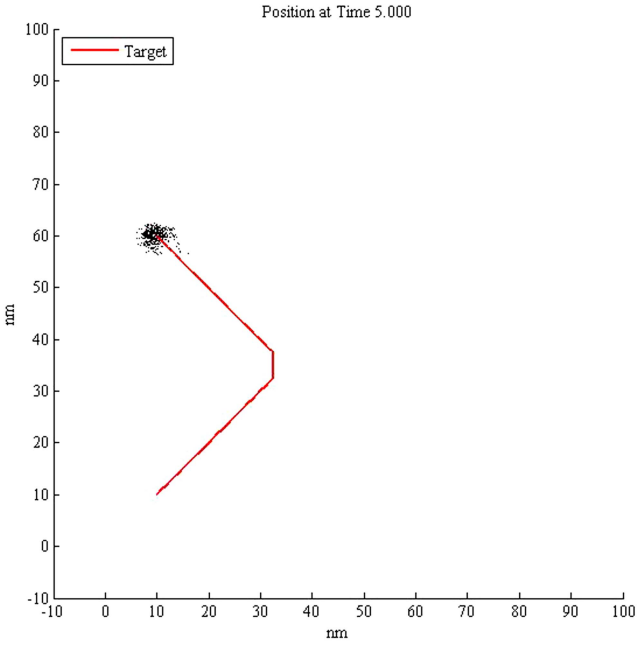

Fig. 9. Postion marginal at 5 hr.

In Section 5, we have shown that quantized measurements can be applied to moving target tracking problems using particle filters and likelihood functions for the quantized measurements in a straightforward, standard Bayesian fashion.

The power of a likelihood function is that it converts measurements from (almost) any measurement space into a function on the target state space. This allows us to incorporate the information in these measurements into the posterior distribution on the target state space. The examples given above illustrate this process with quantized measurements, but the method is applicable to wide range of types of measurements and sensors. In particular, it is applicable to any measurement for which one can compute a likelihood function using the definition in (2). This is why likelihood functions are the common currency of information in Bayesian inference. The examples given above demonstrate this fact.

We have discussed above the virtues of using the simplest solution to a problem. We would be remiss if we did not also point out that there can be drawbacks to unnecessary complexity. For example, if we employ Dempster Shafer methods to handle quantized measurements, then we will be limited in applications to finite

discrete state spaces since there has been no satisfactory extension of Dempster-Shafer theory to continuous state spaces. Even if the state space is finite, the computations involved with Dempster-Shafer methods grow exponentially with the size of the state space which limits its applicability to real problems.

REFERENCES

[1] J. O. Berger
Statistical Decsion Theory and Bayesian Analysis, 2nd ed.
New York, NY: Springer-Verlag, 1985.

[2] Z. Daun, V. P. Jilkov, and X. R. Li
State estimation with quantized measurements: approximate MMSE approach.
In *Proceedings of 11th International Conference on Information Fusion*, Cologne, Germany, July 2008, 1067–1072.

[3] R. Mahler
General Bayes filtering of quantized measurements.
In *Proceedings of 14th International Conference on Information Fusion*, Chicago, USA, July 5–8, 2011, 346–352.

[4] B. Ristic
Bayesian estimation with imprecise likelihoods: Random set approach.
IEEE Signal Processing Letters, 18 (July 2011), 394–398.

[5] B. Ristic, S. Arulampalam, and N. Gordon
Beyond the Kalman Filter.
Boston, MA: Artech House, 2004.

[6] L. D. Stone, R. L. Streit, T. L. Corwin, and K. L. Bell
Bayesian Multiple Target Tracking 2nd ed.
Boston, MA: Artech House, 2014.

[7] http://en.wikipedia.org/wiki/Quantization.

Lawrence D. Stone obtained his B.S. in mathematics from Antioch College in 1964 and his M.S. and Ph.D. in mathematics from Purdue University in 1965 and 1967. He is Chief Scientist at Metron Inc. He is a member of the National Academy of Engineering and a fellow of the Institute for Operations Research and Management Science. In 1975, the Operations Research Society of America awarded the Lanchester Prize to Dr. Stone's text, *Theory of Optimal Search*. In 1986, he produced the probability maps used to locate the *S.S. Central America* which sank in 1857, taking millions of dollars of gold coins and bars to the ocean bottom one and one-half miles below. In 2010 he led the team that produced the probability distribution that guided the French to the location of the underwater wreckage of Air France Flight AF447. He is a coauthor of the 2014 book, *Bayesian Multiple Target Tracking*. He continues to work on a number of detection and tracking systems for the United States Navy and Coast Guard including the Search And Rescue Optimal Planning System used by the Coast Guard since 2007 to plan searches for people missing at sea.

Stephen L. Anderson obtained his B.A. in mathematics from the University of Utah in 1975 and his Ph.D. in mathematics from Brown University in 1980. He is a senior analyst at Metron Inc. His current work is primarily in the development of algorithms to assist in finding isolated personnel such as down pilots. The research focuses on developing motion models and applying search theory to optimize recovery efforts. From 1998 to 2010 Dr. Anderson was the principal developer of the Nodestar algorithm for use aboard US Navy fast attack submarines. The core of Nodestar is a non-linear data fusion methodology for multiple target tracking. The algorithm's capabilities have expanded to include all of the acoustic and non-acoustic organic sources available aboard the submarines. Dr. Anderson has also applied non-linear filtering theory to meteorological problems. He is co-author (with Jeffrey Anderson) of "A Monte Carlo Implementation of the Nonlinear Filtering Problem to Produce Ensemble Assimilations and Forecasts" which appeared in the Monthly Weather Review of the American Meteorological Society in December 1999. In 1995 Dr. Anderson participated in the analysis of flight information databases to develop delay models for the FAA. Additionally, in 1996 he helped in the development analytic models for the effects of unscheduled FAA outages.

Prior to joining Metron, Dr. Anderson was an Associate with Daniel H. Wagner, Associates. For two years, he was the coordinator of the NRL Ocean Surveillance Tracker/Correlator Testbed. In this capacity, he directed the evaluation of ocean surveillance algorithms in all stages of development. The evaluation included review of the algorithm design, verification that the algorithm as implemented reflected the design, and testing of the algorithm on simulated data. Additionally, he participated in the development of an ocean surveillance correlator/tracker used as a baseline in the evaluation process.

Dr. Anderson directed a similar effort for evaluation of Strategic Defense Initiative correlator/tracker algorithms. The basic modules which he designed and assisted in coding are still in use in modified forms, including a version which was used in parallelization studies at Sandia National Laboratory to simulate ten thousand targets. Also for missile defense, he assisted in the development of an algorithm for merging tracks from multiple independent tracking algorithms. Additionally, he was responsible for the testing to verify the basic soundness of the overall concept.