An Estimation Error Bound for Pixelated Sensing

Chris Kreucher[a] and Kristine Bell[b]
[a]Integrity Applications Incorporated, Ann Arbor MI, 48108
[b]Metron, Inc., 1818 Library St., Suite 600, Reston, VA 20190

ABSTRACT

This paper considers the ubiquitous problem of estimating the state (e.g., position) of an object based on a series of noisy measurements. The standard approach is to formulate this problem as one of measuring the state (or a function of the state) corrupted by additive Gaussian noise. This model assumes both (i) the sensor provides a measurement of the true target (or, alternatively, a separate signal processing step has eliminated false alarms), and (ii) The error source in the measurement is accurately described by a Gaussian model. In reality, however, sensor measurement are often formed on a grid of pixels – e.g., Ground Moving Target Indication (GMTI) measurements are formed for a discrete set of (angle, range, velocity) voxels, and EO imagery is made on (x, y) grids. When a target is present in a pixel, therefore, uncertainty is not Gaussian (instead it is a boxcar function) and unbisased estimation is not generally possible as the location of the target within the pixel defines the bias of the estimator. It turns out that this small modification to the measurement model makes traditional bounding approaches not applicable. This paper discusses pixelated sensing in more detail and derives the minimum mean squared error (MMSE) bound for estimation in the pixelated scenario. We then use this error calculation to investigate the utility of using non-threshoded measurements.

Keywords: Estimation, Detection, Error bounds, Pixelated Sensing.

1. INTRODUCTION

This paper discusses the problem of estimating the state of an object using sensor measurements with special attention to “pixelization”. Pixelization refers to the fact that processed sensor data is typically provided to an estimation algorithm on a discrete grid of pixels, rather than as a set of point measurements of targets. For example, Electro-Optical (EO) data is typically made on a 2D spatial grid, with resolution defined by the wavelength and lens diameter, among other things. Likewise, Ground Moving Target Indicator (GMTI) data is typically processed to set of range/angle/radial velocity voxels where the resolution is defined by transmit bandwidth, angle subtended during the collect, pulse repetition interval and so on.

Pixelization is not typically dealt with explicitly in detection, tracking, and estimation algorithms. Instead, the effect of the pixelated nature of the sensor is typically modeled as an additive error source. However, pixelization imparts two interesting properties that this type of approximation does not capture: First, the measurement likelihood is uniform over the pixel, rather than Gaussian. In other words, when a pixel is detected, the implication is that we have learned only that the state of the target maps into that pixel. Second, point estimators are not unbiased unless the target happens to be in the center of a cell.

This paper includes two contributions. First, we derive the minimum mean square error estimator (MMSE) for a simply modeled pixelated sensor. We then discuss why conventional bounding approaches face problems with this model. Next, we compare the difference between this bound for the case where we elect to threshold (“detect”) the data before performing estimation (for the purposes of minimizing communication bandwidth, for example) to the case where the “raw” (i.e., non-thresholded) measurements can be used.

2. THE PIXELATED SENSOR MODEL

Assume we desire an estimate of the state (e.g., position) of a target, x, from a series of measurements made by a pixelated sensor with N cells. For the purposes of exposition, it is useful to describe the case where the target state is 1D and the sensor pixels are in the same units of the desired unknown x. However, the approach extends in a straightforward manner to the more general case.

Signal Processing, Sensor/Information Fusion, and Target Recognition XXV, edited by Ivan Kadar, Proc. of SPIE
Vol. 9842, 98421E · © 2016 SPIE · CCC code: 0277-786X/16/$18 · doi: 10.1117/12.2224059

The measurement in cell i will be denoted z_i, for cells $i = 1 \cdots N$. In the standard "detection/no-detection" case, $z_i \in [0,1]$, whereas in the more general "raw" case $z_i \in R$. We will use the notation i_x to denote the cell that x maps into, and z_x will denote the measurement in that cell.

With that as background, the generic sensing likelihood model is

$$p_{generic}(z_i|x) = \begin{cases} p_1(z_i) & i = i_x \\ p_0(z_i) & i \neq i_x \end{cases}, \tag{1}$$

where p_1 is the target-present distribution and p_0 is the background-only distribution.

This model is often referred to as the "point-target" model. It is only explicitly true when the target is small relative to the sensor pixel size and the target does not produce artifacts (e.g., sidelobes) in adjacent cells. We will assume these conditions are met in this paper.

Radar tracking is an interesting motivating example for this model. In this setting, a series of radar pulses are aggregated together into a coherent pulse interval (CPI). They are then processed to construct a range-Doppler map (RDM), which is used to provide input to a tracker. Each cell in the RDM corresponds to a hypothesized target range and radial velocity. A common model for the target and non-target distributions in Radar is to use Rayleigh statistics [1], i.e.,

$$p_{raw}(z_i|x) = \begin{cases} \dfrac{z_i}{\lambda_t^2} e^{-z_i^2/\lambda_t^2} & i = i_x \\ \dfrac{z_i}{\lambda_b^2} e^{-z_i^2/\lambda_b^2} & i \neq i_x \end{cases}. \tag{2}$$

Often times, the RDM is thresholded (i.e., only those pixels with amplitude over a threshold are preserved). This results in a probability of correctly detecting the target pixel, called p_d, and a probability of falsely detecting a non-target pixel, called p_f. For the thresholded case, the model is

$$p_{thr}(z_i|x) = \begin{cases} p_d & i = i_x, z_i = 1 \\ (1 - p_d) & i = i_x, z_i = 0 \\ p_f & i \neq i_x, z_i = 1 \\ (1 - p_f) & i \neq i_x, z_i = 0 \end{cases}. \tag{3}$$

It can be shown that the thresholded and non-thresholded models are linked under the Rayleigh assumption. If we define $\text{SNR} = \lambda_t^2/\lambda_b^2 - 1$, and use constant false alarm rate (CFAR) thresholding, we have $p_d = p_f^{\frac{1}{SNR+1}}$.

Furthermore, if we define the indicator function

$$I_{i,x} = \begin{cases} 1 & i = i_x \\ 0 & i \neq i_x \end{cases}, \tag{4}$$

then these two models can be written compactly as

$$p_{raw}(z_i|x) = \frac{z_i}{\lambda_t^2} e^{-z_i^2/\lambda_t^2} I_{i,x} + \frac{z_i}{\lambda_b^2} e^{-\frac{z_i^2}{\lambda_b^2}} (1 - I_{i,x}), \tag{5}$$

and

$$p_{thr}(z_i|x) = \begin{cases} p_d I_{x,i} + p_f(1 - I_{x,i}) & z_i = 1 \\ (1 - p_d) I_{x,i} + (1 - p_f)(1 - I_{x,i}) & z_i = 0 \end{cases}. \tag{6}$$

At this point it is useful to highlight the distinction between the pixelated model just presented and the usual method for modeling sensor measurements. The typical detection-based model is that a sensor receives a set of data and first performs a detection process to determine which of the cells have threshold exceendances. The threshold exceendances are then gated, associated, and subjected to other sophisticated processing to (in the ideal case) reject all of the false alarms and be left with only the correct detection $z_{correct}$. This remaining single (assumed correct) detection is then treated as originating from the center of the cell. Then the model is

$$p(z_{correct}) \sim N(x, \sigma^2). \tag{7}$$

This approach approximates sensor pixelization using measurement variance σ^2, typically matched to a function of the cell size. Additionally, it assumes that the uncertainty is centered around the true value of the unknown. There is therefore a mismatch between the assumed model and the actual measurements. Nevertheless, filtering and estimation approaches based on this model often work well in practice if the variance is tuned well.

Of interest here are fundamental bounds on performance that explicitly incorporates the pixelated nature of the sensor. While quite simple, the boxcar models presented above have properties which prevent the direct application of the usual estimation theoretic techniques.

In particular, the support of the distribution is a function of the parameter. This is a well-known situation in which the ordinary Cramer-Rao lower bound (CRLB) does not apply [3][5][7] because the exchange of integration and differentiation in the CRB derivation is no longer valid. Recently, [7] has developed and extension called the Cramer-Rao-Leibniz lower bound (CRLLB) for parameter-dependent support. However, our case has support endpoints that are not differentiable with respect to the parameter. More generically applicable bounds, such as the Barankin family of bounds (including the Chapman-Robbins bound [4], the Bobrovsky-Zakai bound [5], and the Keifer bound [1]) are also not applicable with this model because it is not true under the boxcar pixelated model that $p(z|x) = 0$ implies $p(z|x + h) = 0$.

3. THE MMSE ESTIMATION BOUND

Despite the inability to use the conventional bound techniques, we can derive the minimum mean squared error estimator (MMSE) as follows. Let the prior $p(x)$ be uniform over $x \in (x_L, x_R)$. Then for measurement $z = (z_1, z_2, \cdots, z_N)$, define the per-cell likelihood ratio

$$L(z_i) = p_1(z_i)/p_0(z_i). \tag{8}$$

Assuming conditional independence across the cells,

$$p(z|x) = \prod p(z_i|x) = L(z_x)\prod p_0(z_i). \tag{9}$$

Using Bayes rule, we can compute the posterior given the set of measurements as

$$p(x|z) \propto L(z_x)/\sum L(z_i). \tag{10}$$

Finally, the MMSE, $\hat{x}(z)$, is given by

$$\hat{x}(z) = \frac{\int xL(z_x)dx}{\sum L(z_i)} = \frac{\sum x_i L(z_i)}{\sum L(z_i)}, \tag{11}$$

where x_i is the center of measurement cell *i*.

The actual mean squared error (MSE) is

$$\mathbb{E}_z[(\hat{x} - x_{true})^2] = \int p(z|x_{true})(\hat{x}(z) - x_{true})^2 dz = \int L(z_{xtrue})\prod p_0(z_j)\left(\frac{\sum x_i L(z_i)}{\sum L(z_i)} - x_{true}\right)^2 dz, \tag{12}$$

where the integral is to be interpreted as the N-D integration over the N sensor cells. As expected the MSE is a function of the target location x_{true}.

We now specify the form of the likelihood ratio for the two radar cases we considered before. These come directly from the earlier model definitions. For the thresholded case, we have

$$L_{thr}(z_i) = \begin{cases} p_d/p_f & z_i = 1 \\ (1 - p_d)/(1 - p_f) & z_i = 0 \end{cases} \tag{13}$$

and

$$\prod p_0(z_j) = (1 - p_f)^M p_f^{N-M} \tag{14}$$

and for the non-thresholded case we have

$$L_{raw}(z_i) = \frac{\lambda_b^2}{\lambda_t^2}\ exp\left[z_i^2\,\frac{\lambda_t^2 - \lambda_b^2}{2\lambda_t^2\lambda_b^2}\right] \tag{15}$$

and

$$\prod p_0(z_j) = \prod \frac{z_j}{\lambda_b^2} e^{-\frac{z_j^2}{2\lambda_b^2}} \tag{16}$$

Neither of these likelihoods admit an analytic solution to the MSE. However, both integrals can be evaluated using the principle of importance sampling. When the number of sensor cells is small, we can even enumerate the 2^N cases in the thresholded setting.

In the thresholded case, the MSE depends on the false alarm setting, which is a free parameter. Figure 1 shows the MSE as a function of SNR for different values of p_f. We find that one should elect to choose small p_f so the asymptotic MSE is small.

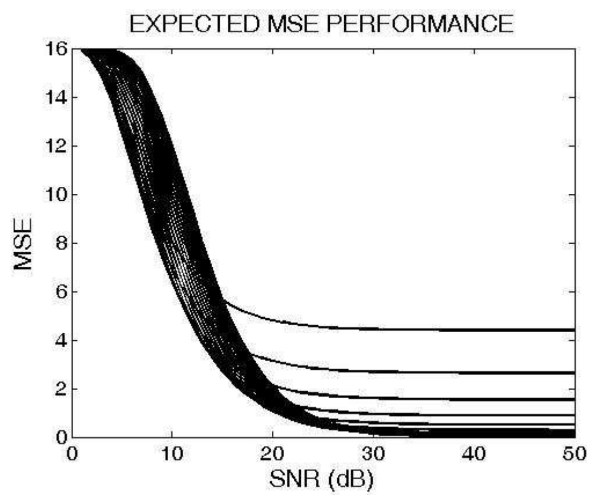

Figure 1. False alarm rate affects the MSE.

In the raw (non-thresholded) case, there is no free parameter. Figure 2 compares the non-threshoded case to the thresholded case when $p_f = 1e - 5$. It shows the benefit of non-thresholded measurements in an example problem can be as much as $4dB$.

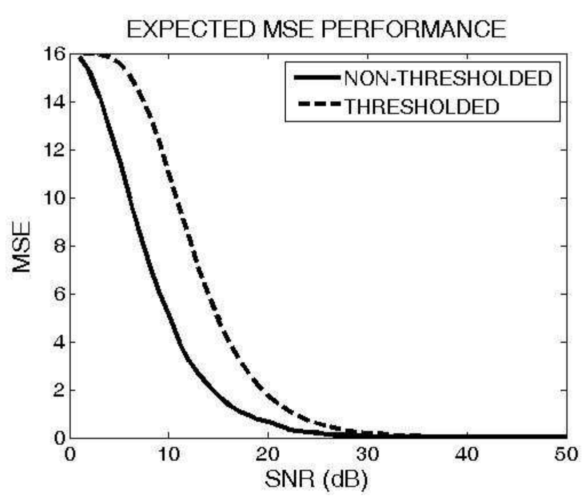

Figure 2. The benefit of not thresholding.

Since the MSE is a function of the target location in a cell, we also derive the average MSE. The average MSE is found by averaging over the (assumed uniform) prior locations of x in a cell as

$$\mathbb{E}_{z,x}[(\hat{x}(z)-x)^2] = \frac{1}{x_R - x_L}\int\int L(z_x)\prod p_0(z_j)\left(\frac{\sum x_i L(z_i)}{\sum L(z_i)} - x\right)^2 dzdx, \tag{17}$$

i.e., the average MSE is found by averaging the MSE over the possible target locations.

It is informative to look at the average MSE in the special case where SNR is very high. In the case of thresholded measurements, this means $p_d \sim 1$ and $p_f \sim 0$ – that is, we can identify the target cell precisely with no false alarms or missed detections. In this case, we find that the average MSE as given by Eq. (17) is $\delta^2/12$, where δ is the cell size. This is identical to the variance bound derived in [7] (using the substitution $\delta = 2a$), which analyzed the special case where $p_d = 1$, $p_f = 0$, and the likelihood was uniform and centered on the parameter.

4. CONCLUSION

This paper has discussed the ubiquitous problem of estimating the state (e.g., position) of an object based on a series of noisy measurements. The standard approach in the literature is to formulate this problem as one of measuring the state (or a function of the state) in Gaussian noise. In many applications, however, measurements are "pixelated" – i.e., they are presented on an sensor array. When a target is present in a pixel, therefore, the uncertainty is not Gaussian (instead it is a boxcar function) and the estimator is not unbiased (the location of the target within the pixel defines the bias of the estimator). This paper has looked at this in more detail and derived an MSE bound for estimation. Particular attention has been given to the distinction between thresholded and non-thresholded measurements.

REFERENCES

[1] Skolnik, M. Introduction to Radar Systems. McGraw-Hill, 2001.

[2] Kiefer, J., "On minimum variance estimators". Annals of Math. Statistics, 23, 627-629, 1952.

[3] Rao, C. R. "Information and accuracy attainable in the estimation of statistical parameters". Bull. of Calcutta Maths. 37,81-91, 1945.

[4] Chapman and Robbins, "Minimum variance estimation without regularity assumptions", Annals Math. Stat. 22: 581–586, 1951.

[5] Van Trees, H and Bell, K., "Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking", Wiley, 2007.

[6] Morelande, M., Kreucher, C., and Kastella, K., "A Study of Factors Affecting Mutiple Target Tracking with a Pixelized Sensor", Proc. SPIE V. 5913, 2005.

[7] Bar-Shalom, Y., Osborne, R., Willet, P., and Daum, F., "CRLB for Likelihood Functions with Parameter-Dependent Support", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 3, July 2014.