

1

Lower Bounds on Pose Estimation with High Range-Resolution Radar

by

Brian M. Smith

Submitted to the Department of Electrical Engineering and Computer Science

in Partial Fulfillment of the Requirements for the Degree of

Master of Engineering in Electrical Engineering and Computer Science

at the Massachusetts Institute of Technology

June 26, 2000

[September 2000]

Copyright 2000 MIT

Author _____

Brian M. Smith

Department of Electrical Engineering and Computer Science

June 26, 2000

Certified by _____


Jeffrey H. Shapiro, Julius A. Stratton Professor of Electrical Engineering
Thesis Supervisor

Accepted by _____

Arthur C. Smith

Chairman, Department Committee on Graduate Theses

BARKER



Lower Bounds on Pose Estimation with High Range-Resolution Radar

by
Brian M. Smith

Submitted to the
Department of Electrical Engineering and Computer Science

June 26, 2000

In Partial Fulfillment of the Requirements for the Degree of
Master of Engineering in Electrical Engineering and Computer Science

ABSTRACT

High range-resolution radar (HRR) is a useful tool for pose estimation: the problem of determining the orientation of a target object. Using a blocks-world tank as target, HRR models for return data for both rough and specular surface reflections are determined. The profiles are used to find Cramér-Rao type bounds on mean-squared error performance of any unbiased estimator which uses HRR data. The goal is to determine the effect on these bounds of the key parameters of the system: signal-to-noise ratio (SNR), target size, range bin width or range-resolution, and target reflectivity. Results are compared with minimum mean-squared error estimates obtained via simulation.

Thesis Supervisor: Jeffrey H. Shapiro

Table of Contents

	Page number
Abstract	2
Acknowledgements	4
I Introduction	5
II Pose Estimation with High Range-Resolution Radar	9
2.1 Pose Estimation	9
2.2 High Range-Resolution Radar	12
III High Range-Resolution Radar Data Models	16
3.1 Conditionally Gaussian Statistical Model	16
3.2 Lambertian/Rough-Surface Target Model	20
3.3 Specular Reflection Model	30
3.4 Composite Target Model	37
IV Cramér-Rao Bounds	39
4.1 Cramér-Rao Analysis	40
4.2 Analysis of Rough-Surface Target Model	42
4.3 Analysis of Specular Target Model	54
V Performance Comparisons	59
5.1 Rough Surface Performance Analysis	59
5.2 Specular Target Performance Analysis	64
VI Conclusions	66
References	68

Acknowledgements

This work was made possible by U.S. Army Research Office Grant #DAAH04-95-1-0494. In addition, the thesis could not have been done without the constant encouragement and aid of my advisor, Professor Jeffrey H. Shapiro.

Chapter I

Introduction

Radio Detection and Ranging, known by the acronym radar, has been used to detect the presence of and to determine the distance to and the velocity of remote objects since its development in the Massachusetts Institute of Technology's Radiation Lab during World War II. [1] However, recent advances in radar and other remote sensing technologies have permitted the collection of sophisticated images of remote objects. Synthetic aperture radar (SAR) systems can form high-resolution range and cross-range images of a target. [2] Similarly, passive forward-looking infrared (FLIR) systems can image the temperature on a target. In contrast, high range-resolution radar (HRR) systems use a high bandwidth pulse which provides more detailed information about the depth profile of an object, albeit at low resolution in the transverse spatial dimensions [3]. Automatic target recognition (ATR) systems attempt to classify the object under observation as one from a library of known possibilities.

In addition to identifying the class to which a target belongs, an ATR system may also give an estimate of its orientation, i.e., its pose. For a rigid-body target the pose parameter describes object orientation in three dimensions. When the target is on the ground at a known location, such as a tank on a flat plain, pose can be specified by a single rotation angle expressed on a 2π radian interval. Because radar signatures, even from the same object, are highly dependent on and vary significantly with changes in orientation, knowledge of pose increases the probability of correct identification. [4]

This thesis examines limits on the accuracy of orientation estimation for ground-based targets based on data from an HRR system. To achieve that goal, there must first be an explicit notion of orientation, including a metric that provides a sensible or natural measurement of the difference between two orientations. The Hilbert-Schmidt norm fulfills these requirements and is used to quantify the performance of the estimators. [5] Additionally, there must exist a statistical model to describe the information that is available from the radar system. A conditionally-Gaussian model is used to represent the data that will be generated by such a system. [6]

Prior work on HRR pose estimation has used Monte Carlo simulation to determine estimation performance for specific targets [2]. These simulations rely on detailed computer-aided-design (CAD) models to generate the parameters used in the return data statistical model. This approach has the benefit of employing accurate models for targets of interest, but it has several distinct disadvantages. A large number of Monte Carlo trials is required to obtain accurate results, and to determine the performance behavior as a function of quantities such as signal-to-noise ratio (SNR) or range-resolution requires an order of magnitude increase in the number of simulation trials. Furthermore, to study pose-estimation performance when data from several different remote sensors is optimally combined requires even more dramatic increases in simulation time. Finally, and perhaps most importantly, little physical intuition can be gained from simulation results: it is desirable to be able to show that, based on a certain theoretical model, performance is a specific, derived, function of the other system parameters.

The above shortcomings illustrate why analytically derived bounds on pose-estimation performance are an attractive alternative to Monte Carlo simulations. Cramér-Rao analysis provides a method of calculating these limits directly, without simulation. Given a probability distribution for some data, the Cramér-Rao method allows the determination of variance lower bounds on unbiased estimates of the unknown parameters of the distribution. [7] In the case of interest for this thesis, the HRR return profile from a known target is a random vector whose probability distribution is determined by the target's orientation angle. That random vector contains information about the orientation angle; the Cramér-Rao analysis allows us to determine how much information it can provide. If, for example, the distribution is highly sensitive to changes in orientation around some nominal pose angle, then it should be easy to accurately estimate the orientation angle when it is near that particular value. On the other hand, pose estimation will be difficult when the target is at an orientation for which the distribution is relatively constant over a wide range of poses.

This thesis consists of two main portions. In Chapters 2 and 3 there is background information on orientation estimation and the Hilbert-Schmidt metric and its measurement statistics, on the specifics of high range-resolution radar, on the statistical model for HRR return data, and on the Cramér-Rao bound. The second portion deals with the work done for this thesis. In Chapters 4 and 5 there is the derivation of the orientation dependence of the HRR statistics for a blocks-world target model, the analysis of Cramér-Rao type bounds based on these parameters, and comparisons between the derived bounds and the performance of a minimum-mean-squared error estimator. The

thesis concludes in Chapter 6 with a summary of accomplishments and some suggestions for future work.

Chapter II

Pose Estimation with High Range-Resolution Radar

An automatic target recognition system which estimates pose requires a method of stating its output, a formalization for expressing target pose. In addition, in order to evaluate the performance of the system, there must be a way to compare the pose estimate to the actual value and determine the magnitude of the error. A common formulation for measuring error is the Hilbert-Schmidt norm, which is described in the first section of this chapter. The second section of this chapter is devoted to the physics of high range-resolution radar. An ATR system needs data on which to base its decisions; this thesis investigates the accuracy of ATR using HRR data. It is therefore critical to understand the manner in which information about the remote target is derived, received, and processed by the ATR system.

2.1 Pose Estimation

Intuitively, the simplest way to describe the angular position of a point in 2-space, equivalent to the angular orientation of a ground-based rigid-body target, is the polar coordinate θ . The range of θ , without loss of generality, can be taken to be $-\pi < \theta \leq \pi$. However, this formulation has the disadvantage of lacking a simple effective metric, or method of describing the differences between two angles. Orientation is cyclic: given a range of $-\pi < \theta \leq \pi$, a pose estimate of -3.1 radians is an accurate measurement of an object placed at $+3.1$ radians, yet the standard metric $|\theta_1 - \theta_2|$ would give an error of 6.2

radians. This value is well above the intuitive error of approximately .08 radians. In fact, one would expect that the greatest possible error would occur when the estimate is directly opposite the actual value, for a difference of π ; 6.2 is even much larger than this intuitive maximum.

A method of stating orientation, which avoids this drawback, is the rotation matrix. If we represent a point in 2-space as the 2-vector $[x,y]^T$, then the action of the rotation matrix O ,

$$O = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

on the position vector is to rotate the point about the origin by the angle θ . Notice that the determinant of O is unity regardless of θ . The set of all possible O is known in group theory as SO_2 , the special orthogonal group of dimension 2. (Special refers to the determinant being equal to one; orthogonal is because the column vectors are perpendicular to each other.) SO_3 and even high order groups are defined in the same manner, and SO_3 performs all possible rotations in 3-space. These groups do not have the cyclic problem as above, for we can use the Hilbert-Schmidt norm to define differences between a pose O and an estimate \hat{O} based on $\hat{\theta}$ [5]:

$$\hat{O} = \begin{bmatrix} \cos \hat{\theta} & -\sin \hat{\theta} \\ \sin \hat{\theta} & \cos \hat{\theta} \end{bmatrix}$$

$$\|O - \hat{O}\|_{HS}^2 = \text{tr}((O - \hat{O})^T (O - \hat{O})) = 4 - 4(\cos(\theta - \hat{\theta}))$$

The Hilbert Schmidt Norm

Observe that, even if θ and $\hat{\theta}$ span an arbitrary boundary, the common sense of the metric is preserved due to the fact that $\cos(\theta) = \cos(2\pi - \theta)$. Also, as $\theta - \hat{\theta}$ approaches zero, $\|O - \hat{O}\|_{HS}$ approaches $\sqrt{2}|\theta - \hat{\theta}|$, which is a nice property – it is the intuitive metric that we first desired, scaled by a constant factor.

Let $p(\theta|D)$ be the conditional probability distribution for θ , conditional on knowledge of the sensor data D . The Hilbert-Schmidt estimator for θ is then defined in the natural way:

$$\hat{O}_{HS}(D) = \text{ArgMin}_{\hat{o}} \left[E \left(\|O - \hat{O}\|_{HS}^2 \right) \right]$$

and can be shown to satisfy $\hat{O}_{HS}(D) = \frac{\int_{-\pi}^{\pi} O_{\theta} p(\theta | D) d\theta}{\sqrt{\det \left(\int_{-\pi}^{\pi} O_{\theta} p(\theta | D) d\theta \right)}}$ where

$$O_{\theta} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

The Hilbert-Schmidt estimate minimizes the expectation of the squared Hilbert-Schmidt error on the distribution. An equivalent formulation of estimator is:

$$\hat{O}_{HS}(D) = \begin{bmatrix} \cos \hat{\theta}_{HS} & -\sin \hat{\theta}_{HS} \\ \sin \hat{\theta}_{HS} & \cos \hat{\theta}_{HS} \end{bmatrix}$$

where $\hat{\theta}_{HS} = \arg \left(\int_{-\pi}^{\pi} (\cos \theta) p(\theta | D) d\theta + i \int_{-\pi}^{\pi} (\sin \theta) p(\theta | D) d\theta \right)$, $-\pi < \hat{\theta}_{HS} \leq \pi$.

Given the Hilbert-Schmidt estimator, and an associated Hilbert-Schmidt error, we now have a convenient way to measure the error variance of any estimator. This error

can be compared with the Hilbert-Schmidt bound, by definition the minimum mean-squared error of any estimator.

2.2 High Range-Resolution Radar

A pulsed radar system works by emitting a short burst of electromagnetic energy, in the form of a carrier frequency modulated by an envelope. If there is a target object located nominally R_0 meters away, the electromagnetic wave may be reflected and then return to the antenna $2 R_0/c$ seconds after it was transmitted. The measured value of this delay can then be used to estimate the range. Many high range-resolution systems will use a chirp pulse of the form $A(t)\cos\left(\left(w_0 + \frac{\pi W}{T}t\right)t\right)$, where $A(t) > 0$ is the envelope of duration T and W is the chirp bandwidth. The frequency of this pulse increases linearly with time. An alternative approach varies the pulse's frequency in piecewise-constant discrete steps. Either method allows a pulse of the same width T in the time domain to occupy more bandwidth W in the frequency domain, yielding an autocorrelation between transmitted and received pulses that is more sharply peaked. That narrower autocorrelation peak translates into a more precise measurement of the time-of-flight of the pulse, and hence, of the distance to the target. Greater bandwidth, and therefore smaller effective pulse width, means improved range resolution for the HRR system. For a typical forward looking millimeter-wave radar system, the center carrier frequency might be at 85.5 GHz (corresponding to a wavelength of 3.51mm), with a chirp bandwidth of 300MHz. [3]

Given that the target is not simply a featureless plane oriented normal to the direction of the incoming radar pulse, then each of these features will also reflect the

electromagnetic wave. A portion of the target located at a distance R_0+r from the antenna will cause a reflection to appear $2r/c$ seconds later than the nominal range R_0 return. For every feature, in fact, there will be an appropriate reflection returning to the radar, and their sum creates the entire return for the radar pulse. The range resolution is the minimum distance between two features in the radar antenna's beam whose returns can be distinguished from their delay difference. For simplicity, this thesis assumes that the radar system transmits a single carrier frequency modulated by a transform-limited (unchirped) Gaussian pulse. If the pulse has width T , the distance $cT/2$ can be taken as approximately the range resolution. Alternatively stated, an arbitrary pulse of bandwidth W can resolve features which are located c/W meters apart.

The received signal will differ in both phase and amplitude from the transmitted pulse. After bandpass/matched filtering it takes the form $r(t)\cos(\omega t + \phi(t))$, where $r(t)$ and $\phi(t)$, the magnitude and phase shift, are random processes whose statistics depend on the target characteristics and the random noise in the receiver. This sinusoid can be written in the more compact complex notation $\text{Re}(\underline{r}(t)e^{j\omega t})$, where $\underline{r}(t) = r(t)e^{j\phi(t)}$ is a complex envelope containing both the phase and amplitude information. With appropriate normalization, the instantaneous power received at the antenna in this formulation is then $|\underline{r}(t)|^2$. Samples of this power waveform at equal time spacings corresponding to the range resolution yield the HRR data to which the ATR algorithms have access and upon which they must base their orientation estimate. Each sample can be considered to be the average power received over the sampling period, equivalently, the energy received at the antenna over that time period. Also, since each different sample corresponds to a portion of the target object which is at a different, specified

distance from the radar system, each sample is termed a *range bin*. The typical HRR system described above has 512 range bins per profile, each corresponding to a range resolution, or range bin width, of 0.52 meters. [3]

The amplitude of the received electromagnetic wave, or equivalently the power detected by the antenna after processing, is dependent on several variables. From the standard monostatic radar equation [8]

$$P_R = P_T \frac{G_T}{4\pi R^2} \frac{\sigma_R}{4\pi R^2} e^{-2\alpha R} \epsilon$$

we see that the received power P_R is proportional to the transmitted power P_T , as well as to the antenna gain of the transmitter G_T , the overall power efficiency of radar ϵ , and most importantly to our analysis, the target cross section σ_R , at range R . In addition, the power is exponentially attenuated by scattering in the atmosphere with a decay coefficient α , and spreads out spherically from the transmitting and reflecting points, as seen in the two factors of $4\pi R^2$. The cross-sectional area σ_R is so important because the differences in range of any interesting features is so small compared to the total range to the target R , all of the other variables may be treated as constants. This simplifies the radar equation, in our area of interest, to a proportionality: $P_R = \kappa \sigma_R$

It is also important to note that while SAR, FLIR, and laser radar systems place many pixels across the face of a target object, the transverse spatial resolution of an HRR system is generally incapable of this imaging task. This work assumes that the target, while occupying several range bins, is confined to a single spatial pixel, i.e. that the target falls within a single beamwidth of the HRR radar. The availability of good depth information without transverse spatial information means that the return profile yields no

information about the horizontal or vertical placement of any feature within the beam.

As a result, there will be ambiguity in determining the orientation of objects with many sorts of symmetry. A tank with bilateral symmetry, for example, will produce the same HRR profile when viewed from its left side as it would when viewed from its right side.

Chapter III

High Range-Resolution Radar Data Models

A statistical model for a set of random data is needed in order to perform any sort of meaningful parameter estimation based on that data. High range-resolution radar data consists of a vector of range bin samples. The samples correspond to the return power reflected from the target object and received by the antenna within specified time intervals. Equivalently, because $R=ct/2$ shows that range is proportional to time delay, the samples also correspond to electromagnetic wave energy that was reflected off of the target at differing, specified ranges from the antenna.

Modeling the statistics of this random vector requires two stages. First, given a reflected electromagnetic wave with certain phase and amplitude properties, the distribution of the data that the radar receiver will yield must be modeled. Then, we need to know how a specific target, based on its dimensions, reflectivity, and orientation, will reflect the electromagnetic radar pulse and yield the above properties.

3.1 Conditionally Gaussian Statistical Model

As previously described, under proper normalization the instantaneous power in an electromagnetic wave is equal to $|r(t)|^2$. For the i^{th} range bin, the radar produces a value p_i , equal to the time-averaged power received over the time-delay interval corresponding to that range bin. It is convenient to work with the random variable $r_i = \sqrt{p_i}$, the square-root of the power. The vector \vec{r} comprised of these components is

the range profile, a real vector of HRR return data. We shall begin analysis, however, with the complex vector $\underline{\vec{r}}$, a vector of samples of the complex envelope.

Each complex random variable \underline{r}_i of the vector $\underline{\vec{r}}$ is the sum of the radar return due to the target, \underline{y}_i , and the receiver noise, \underline{n}_i :

$$\underline{r}_i = \underline{y}_i + \underline{n}_i$$

The receiver noise \underline{n}_i is typically receiver thermal noise, and hence can be taken to be statistically independent from bin to bin. In keeping with models of this kind of noise, \underline{n}_i is taken to be a zero-mean complex Gaussian random variable whose real and imaginary parts each have variance $\lambda_n/2$. The target reflection \underline{y}_i is itself a complex Gaussian random variable, whose mean \underline{m}_i and variance $\lambda_{gi}/2$ depend on the target's geometry and orientation. Explicitly, then, the reflection should be written $\underline{y}_i(Q, \theta)$, as a function of target Q and angle θ , but that dependence will be understood in the more compact notation. Finally, each \underline{y}_i , like receiver noise \underline{n}_i , in each bin, is statistically independent, conditioned on knowledge of the target and its pose. It follows that the statistics of a single range bin can be studied without loss of generality; further analysis will drop the identifying subscripts.

Physically, the distributions' \underline{m} and λ_g parameters correspond to properties of the reflected electromagnetic wave. The mean parameter \underline{m} represents the coherently reflected electromagnetic field, as if off of a smooth, mirror-like surface. The qualitative adjective 'smooth' can be more precisely defined to apply to surfaces whose variations and irregularities are much smaller than the wavelength of the carrier signal. These

coherent reflections are also described as *specular*. Herein, \underline{m} is determined from the geometry of the target using the far-field Fraunhofer diffraction approximation. The variance λ_g accounts for random minute surface variations of the target, manufacturing variations, mud and dirt, constructive and destructive interference of waves bouncing off of unmodeled features of the target – for any factors that are not taken into account in our simple geometric modeling and that vary even from target to target of the same type. To contrast with the mean, the surface variations which produce the interference for the variance parameter λ_g will have size on the order of the carrier signal's wavelength or larger. The model for computing λ_g as a function of the target geometry assumes that the parameter is proportional to the cross-section that the target presents to the radar system at a given range, multiplied by an orientation factor. The methods for deriving both \underline{m} and λ_g will be discussed in greater detail later.

Using the above Gaussian distributions for the real and complex parts of \underline{r} , the distribution for the real magnitude r can be shown to be: [9]

$$p_r(r | \theta) = \frac{2r}{\lambda} e^{-\frac{(r^2 + |\underline{m}|^2)}{\lambda}} I_0\left(\frac{2|\underline{m}|r}{\lambda}\right), r \geq 0;$$

where $I_0(x) = \frac{1}{2\pi} \int_0^{2\pi} e^{(x \cos \phi)} d\phi$, the zeroth order modified Bessel function and $\lambda = \lambda_n + \lambda_g$.

This probability distribution is named the Rice distribution. Similarly, the probability distribution for the power $p = r^2$ is found by the Leibniz rule to be:

$$p_p(p | \theta) = \frac{1}{\lambda} e^{-\frac{p + m^2}{\lambda}} I_0\left(\frac{2m\sqrt{p}}{\lambda}\right), p \geq 0.$$

Additionally, there are more specialized forms of these distributions, applicable when the mean \underline{m} is zero. The Rayleigh distribution describes the magnitude, and the exponential distribution applies to the power:

$$p_r(r|\theta) = \frac{2r}{\lambda} e^{-\frac{r^2}{\lambda}}, r > 0 \quad \text{and} \quad p_p(p|\theta) = \frac{1}{\lambda} e^{-\frac{p}{\lambda}}, p > 0$$

Examples of the Rayleigh and Rice distributions are plotted in Figure 3.1.

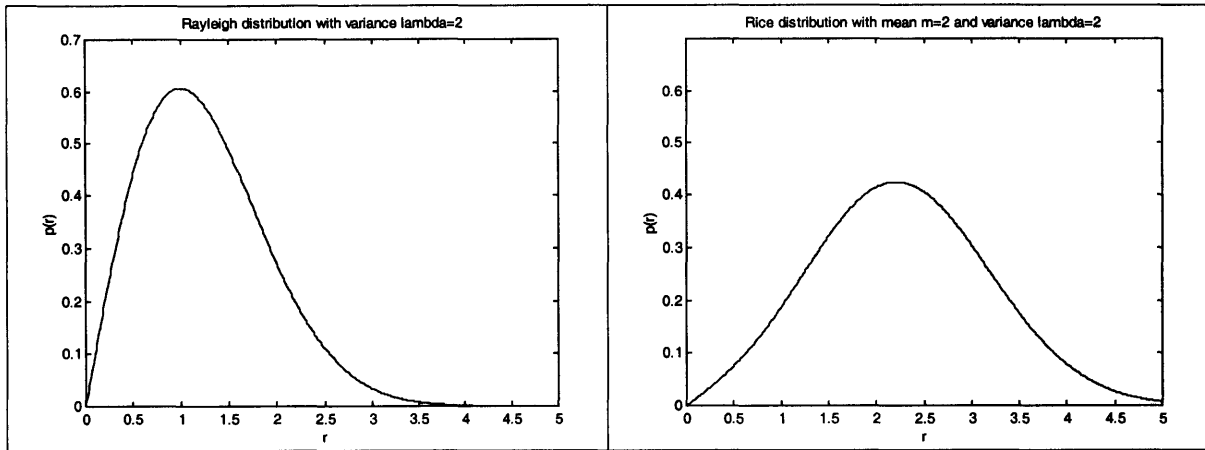


Figure 3.1: Rayleigh and Rice probability distributions

These conditionally-Gaussian distributions, once m and λ_g have been related to target geometry and orientation, will be the foundation for our Cramér-Rao analysis. Some prior HRR-based ATR work [6] has been performed with a deterministic model that lacks the reflection variance λ_g . In a train-and-test approach based on real data in which the performance of the deterministic model was compared to that of using the conditionally-Gaussian model, it was found that the conditionally-Gaussian model gave substantially better performance. [6] That is why we have selected it for use in our Cramér-Rao analysis.

3.2 Lambertian/Rough-Surface Target Model

The focus of this thesis lies in determining analytical bounds whose behavior can be linked, through underlying theory, to physical descriptions of the target and the radar system. To accomplish this goal, the work uses simplifying models of the target in order to draw out these scaling laws. The first of these is to employ a blocks-world model of the target of interest, here a generic three-block tank: body, turret, and gun barrel, as shown in Figure 3.2. The tank is viewed directly from ground level as it rotates around a vertical axis. Each block that comprises the tank has a height z_i , a width x_i , and a depth y_i .

Size Parameters for Standard Blocks-World Tank Model, meters								
Body			Turret			Gun Barrel		
x_1	y_1	z_1	x_1	y_1	z_1	x_1	y_1	z_1
3.8	10.28	1.5	3.48	8.0	1.4	0.3	6	0.3




Figure 3.2 Blocks-world tank model, seen from $\theta = 0$ (a) and $\theta = -\pi/2$ (b)

As stated above, the power received from a reflection, less thermal receiver noise, comes from two sources, coherent specular reflection and Lambertian fading interference. In another attempt at simplification, these two limiting reflectivity cases will be studied separately.

In the rough-surface model, the coherent return \underline{m} is taken to be zero. Under this condition, it was shown above that power takes an exponential distribution whose expected value in each range bin is the variance λ .

$$E[p | \theta] = \lambda$$

To determine λ , therefore, the amount of power that is reflected incoherently off of the target surface must be found. Given that the target is completely within one beam width, and that the target is sufficiently far away from the antenna (absolute range is much greater than range-resolution), we can assume that the power per unit area incident from the beam onto the target surface is a constant. Then, the reflected power from the target is proportional to the area presented by the target to the radar system. Finally, because this power does not return to the antenna all at the same time, but is distributed over time as the area is distributed over range, all of this becomes a function of distance from the antenna, as well.

The simplest case is a flat, but rough-surfaced, rectangular plate with width x_1 and height z_1 . Both dimensions are measured in coordinates normal to the propagation direction of the radar's incident plane wave, so that the surface area available to reflect the beam is the product $x_1 z_1$. Let R be the range from the radar system to the plane, with an appropriate offset such that the origin of R is at the target plane. $R > 0$ refers to a point between the plane and the radar; $R < 0$ for objects further from the radar receiver than the origin.

Allow the orientation $\theta = 0$ to be when the plane is normal to the incident plane wave; positive θ shall be counterclockwise rotation when the system is viewed from above (azimuthal angle of $\pi/2$). Recall, however, that the range profile, however, has no spatial resolution in the plane of the target, so for this target the profiles for poses $\theta = +\theta_0$ and $\theta = -\theta_0$ will have identical distributions.




Figure 3.3: Determining the Cross Section in a Simple Case

From Figure 3.3 above, the differential cross-section $\frac{d\sigma(R)}{dR}$ for a plane is seen to

be:

$$\frac{d\sigma(R)}{dR} = \begin{cases} z_i \cot \theta; R_{\min} \leq R \leq R_{\max} \\ 0; \text{elsewhere} \end{cases}$$

with the boundaries on R given by:

$$-R_{\min} = R_{\max} = \frac{1}{2} x_i \sin \theta$$

The differential cross section $\frac{d\sigma(R)}{dR}$ should intuitively have the property that

$$\int_{-\infty}^{+\infty} \frac{d\sigma(R)}{dR} dR = A_n, \text{ where } A_n \text{ is the projection of the surface on the normal plane. Indeed,}$$

observe that $\int_{\frac{1}{2}x_l \sin \theta}^{\frac{1}{2}x_l \sin \theta} (z_l \cot \theta) dR = (x_l \sin \theta)(z_l \cot \theta) = x_l z_l \cos \theta$ yields the expected result.

In addition to being proportional to the cross-section, however, the reflected power must be multiplied by an angular factor to account for the obliquity of the plane. For this case, the factor is $\cos \theta$. For a plate whose nominal $\theta = 0$ orientation differs by θ_0 degrees from the normal, for example the side face of a rectangular block, the factor would be $\cos(\theta - \theta_0)$. For the right or left side of the rectangular block, which are oriented at $\theta_0 = \pi/2$ when the target as a whole has pose $\theta = 0$, the factor is $\sin \theta$.

The function $f(R, \theta) = \frac{d\sigma(R)}{dR} \cos \theta$ is a continuous function of range. What is needed, though, is to create a discrete function with a value for the variance at each range bin. To conform to the physical model of the radar system, we introduce a discretization in which the power sample for each range bin is the time average of the instantaneous power over the time-delay interval corresponding to that bin.

Let η be the range bin density, in units of range bins per meter. This quantity is defined so that $\frac{1}{\eta} = R_{res}$, the range resolution of the system. Let there be $2N+1$ total range bins, each identified by an index J such that $-N \leq J \leq N$. Center the $J=0$ range bin around $R=0$. The lower and upper edges of the J^{th} range bin are then $\frac{J-1/2}{\eta}$ and

$\frac{J+1/2}{\eta}$. The parameter $\lambda_g(\theta, J)$ for a single plane target can now be defined and is

illustrated in Figure 3.4 for a hypothetical, fixed orientation, target plate.

$$\lambda_g(\theta, J) = \kappa \int_{\frac{J-1/2}{\eta}}^{\frac{J+1/2}{\eta}} \left(\frac{d\sigma(R, \theta)}{dR} \cos \theta \right) dR = \kappa \int_{\frac{J-1/2}{\eta}}^{\frac{J+1/2}{\eta}} f(R, \theta) dR$$

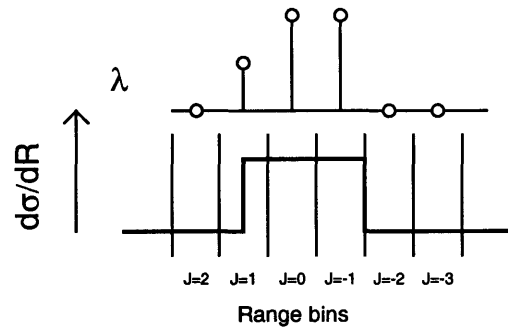


Figure 3.4 Hypothetical range profile at a fixed orientation.

In the slightly more complicated blocks-world model that we use to represent the tank, a rectangular prism viewed from zero azimuthal angle has four planar sides of interest; see Figure 3.5. Additionally, these planes do not rotate around their center points, as in the




Figure 3.5: A rectangular prism, or block.

above example, but around some other point in space. This axis of rotation is often, but not necessarily, through the center of the block. Suppose the block has dimensions x_1 by y_1 by z_1 , and let the block face that is oriented towards the radar when $\theta=0$ again have height z_1 and width x_1 . This will be known as the front of the target. The depth of the block is y_1 ; equivalently, the sides of the block which face the radar at orientations of $\theta = \pm\pi/2$ have dimensions of y_1 by z_1 . The axis of rotation will be located along the centerline running from the front face to the rear face, at a distance of p from the front of the target, as shown in Figure 3.6. When $p = \frac{y_1}{2}$, we have the special case in which the rotation axis is located at the center of the block.

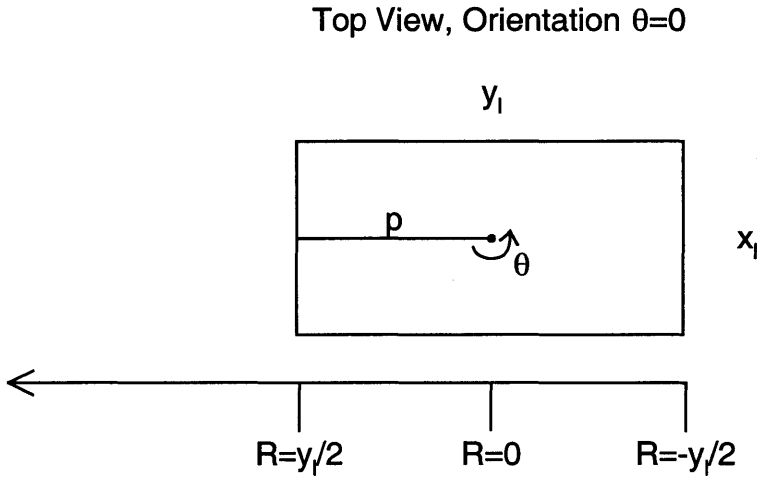


Figure 3.6: Dimensions of a block, top view, where $p = \frac{y_1}{2}$. The radar system is located far to the left of the block.

As the block rotates, the radar system will see various cross-sections of each of two faces. At a very small angle, for example, there will be a large cross-section and

return in the range bin corresponding to the front face of the target; also, there will a small cross-section for each of the bins along the side face of the target.

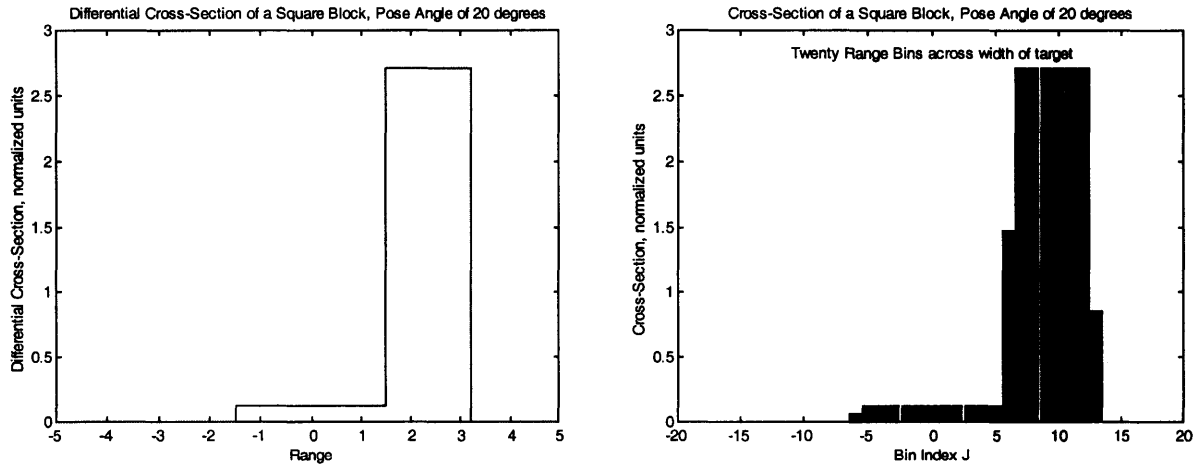


Figure 3.7: Continuous Differential Cross Section $f(R,\theta)$ as a function of range and the discretized $\lambda_g(\theta,J)$, both of the same square 5m x 5m target with a range resolution of 0.25m and pose angle of 20° .

Observe that for a rotating block at any given angle, the differential cross-section vs. range is a piecewise constant function with two non-zero portions, as shown in Figure 3.7. One portion represents the set of ranges over which two sides of the block appear; in the figure below, this corresponds to region II. The other piece contains the cross-section for only one face of the block, which appears in region III of Figure 3.8. There will be no reflection from the block in region I, where none of the target is located, nor in region IV, because the rear two sides of the block are hidden from the radar’s view by the two visible faces.




Figure 3.8: Corners of the block define the boundaries of different regions.

Also marked in Figure 3.8 are the three relevant corners of the block. Intuitively we can see that these boundaries move as the block rotates, and that the value of the cross-section in each of these regions changes. Let $d_1(\theta)$, $d_2(\theta)$, and $d_3(\theta)$ be the positions of these three marked corners in the range coordinates, with zero taken to be at the axis of rotation.

If the axis is at the center of the block, then at $\theta=0$, $d_1 = -\frac{y_l}{2}$ and

$d_2 = d_3 = \frac{y_l}{2}$. As θ increases from 0 to $\pi/2$, d_1 increases monotonically and d_2 decreases

monotonically. However, d_3 increases, until it reaches a maximum at $\theta_m = \tan^{-1}\left(\frac{x}{2p}\right)$

and then it decreases to $d_3 = \frac{x_l}{2}$ at $\theta = \frac{\pi}{2}$. Specifically, the values for these boundaries

are shown in Table 3.1, along with some convenient notation in Figure 3.9:




Figure 3.9: Graphical representation of notation.

Case $p = \frac{y_l}{2}$	General Case; p not necessarily equal to $\frac{y_l}{2}$
$\theta_m = \theta_f = \tan^{-1}\left(\frac{x}{y}\right)$	$\theta_f = \tan^{-1}\left(\frac{x}{2p}\right)$
$\theta_m = \theta_r = \tan^{-1}\left(\frac{x}{y}\right)$	$\theta_r = \tan^{-1}\left(\frac{x}{2(y-p)}\right)$
$d = d_f = \frac{1}{2}\sqrt{x^2 + y^2}$	$d_f = \sqrt{p^2 + (x/2)^2}$
$d = d_f = \frac{1}{2}\sqrt{x^2 + y^2}$	$d_f = \sqrt{(y-p)^2 + (x/2)^2}$
$d_1(\theta) = -d_2(\theta) = -\frac{1}{2}(y \cos \theta - x \sin \theta)$	$d_1(\theta) = -d_r \cos(\theta + \theta_r)$
$d_2(\theta) = \frac{1}{2}(y \cos \theta - x \sin \theta)$	$d_2(\theta) = d_f \cos(\theta + \theta_f)$
$d_3(\theta) = \frac{1}{2}(y \cos \theta + x \sin \theta) = d \cos(\theta - \theta_m)$	$d_3(\theta) = d_f \cos(\theta - \theta_f)$

Table 3.1: Positions of the corners of the rotating block.

Now computation of the differential cross section $\frac{d\sigma(R)}{dR}$ is straightforward. If θ

is such that $d_2 < R < d_3$ then there is a contribution from the front face with value

$z_l \cot(\theta)$. Also, if R lies between the limits d_1 and d_3 , then there is a contribution of

$z_i \tan(\theta)$ from the side face. Multiplying these by the correct obliquity factors ($\cos(\theta)$ for the front face, $\sin(\theta)$ for the side) and we have the continuous function of range $f(R, \theta)$. From this $\lambda_g(\theta, J)$ is computed by integrating over the depth of the range bin, as it was explicitly defined earlier. In figure 3.7 above, we can see that although the differential cross section is piecewise constant, the bins in which the boundaries, d_1 , d_2 , and d_3 , fall take on different values. In fact, it is these boundary bins in which λ_g changes the quickest with respect to θ .

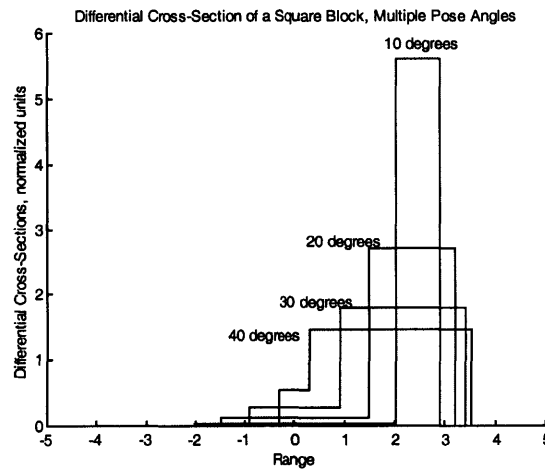


Figure 3.10: Cross-Section with obliquity factor at different orientations.

Figure 3.10 shows $f(R, \theta)$ versus range R for a rectangular block of dimensions 5m x 5m x 1m seen at four different poses. We can see that at the 10° pose, nearly all of the reflected power returns in the few bins corresponding to the front face. As pose angle increases, the number of bins receiving reflection from this face increases, however, the magnitude of the average power in each bin is decreasing.

Given the model for a rectangular prism, computing the λ_{gJ} for a blocks-world tank requires only the addition of three sets of λ_{gJ} , one each for the tank, body, and gun barrel of the tank. Each is computed with the dimensions x_1 , y_1 , and z_1 for the size tank of our choosing, and again the final value is just the sum of the three. To see this, remember that the λ 's are all proportional to the cross-sectional area, and that the high range-resolution radar has no way to distinguish where, in the transverse spatial direction, an object lies. By stacking three blocks on top of each other we are increasing the cross-sectional area seen by the radar system. Since all of the calculations from that point onwards are linear, addition of the λ 's is the proper course.

3.3 Specular Reflection Model

The above Lambertian model is valid when the depth of the surface irregularities on the target is large enough for complete incoherence – when no macroscopic Snell's law reflection off the target occurs. The limiting case in the opposite direction is a smooth, highly polished surface whose irregularities are of an order much smaller than the radar wavelength. Here, the surface acts as a mirror, and an electromagnetic wave reflects according to Snell's law. Intuitively this means a return will only be detected at the radar when a face is nearly normal to the beam, and that in this case the return will be of high intensity.

To develop a model for this specular-return intensity we will investigate the consequences of Fraunhofer diffraction. The Fraunhofer approximation is the far-field solution to a wave equation, which specifies the complex amplitude of an electromagnetic wave when the amplitude is known over some other plane. [10] It is a

simplification of the Huygens-Fresnel principle, derived from boundary conditions on the electromagnetic field, which is valid when the true range R from the source plane to the field point of interest satisfies $R \gg \frac{1}{4} kD^2$. Here, $k = \frac{\omega}{c} = \frac{2\pi}{\lambda}$ is the wave number of the radar's transmitted electromagnetic wave and D is the effective diameter of the radiating region in the source plane.

The Fraunhofer approximation says that if the complex amplitude of an electromagnetic wave is known over some aperture, some portion of a plane, and is assumed to be zero elsewhere on the plane, then the complex amplitude of the electromagnetic field at a point a distance z and at coordinates x_0 and y_0 can be determined. Given a coordinate system x_1, y_1 in the known plane, we can specify this complex wave as $U(x_1, y_1)e^{-j\omega t}$. The Fraunhofer approximation for a monochromatic source states that the complex amplitude at the remote point can be calculated via:

$$U(x_0, y_0) = \frac{e^{jkz} e^{j\frac{k}{2z}(x_0^2 + y_0^2)}}{j\lambda z} \iint_{\infty} U(x_1, y_1) e^{-j\frac{2\pi}{\lambda z}(x_0 x_1 + y_0 y_1)} dx_1 dy_1$$

Since we are interested in only the center point ($x_0=y_0=0$) of the distribution, this formula simplifies greatly.

$$U = \frac{e^{jkz}}{j\lambda z} \iint_{\infty} U(x_1, y_1) dx_1 dy_1 \quad [10]$$

To derive the reflection model, we use two applications of the Fraunhofer far-field approximation. The first of these is a demonstration that, at a plane normal to the beam at the surface of the target, the incoming radar beam can be taken to be a plane wave, i.e. to have the form $Ae^{j(kz - \omega t)}$. We then look at how Snell's Law reflection off of a surface will alter the electromagnetic field due to the radar. The complex amplitude

$U(x_1, y_1)$ resulting from this calculation is then used as the known data in the Fraunhofer equation to determine the reflection strength seen by the antenna.

These Fraunhofer calculations are valid only in the case of a monochromatic wave, i.e., a continuous signal at a constant frequency. The electromagnetic pulses used in radar, however, comprise a range of different frequency components. To account for this, we first perform the Fraunhofer approximation for the general monochromatic case. This process creates a linear transfer function; for any single-frequency input to the reflection system, the output is a wave of the same frequency with a multiplicative change in its amplitude and phase. Then, for whatever pulsed time-domain signal is desired, we can perform the Fourier transform to determine the weighting amounts for each frequency component of the waveform. We then multiply these frequency-domain weights by the frequency-dependant transfer function to determine the effect of reflection and transmission on the waveform as an entity. An inverse Fourier transform then returns the time domain waveform received by the antenna.

Let D_R be the diameter of the circular radar transmitting antenna and P_T be the transmitted power. Given that $P = \int_{surf} |E|^2 dA$ and that the power is evenly distributed over

the antenna surface, the electromagnetic field strength of the transmitted wave E is

$\sqrt{P_T \frac{4}{\pi D_R^2}}$. To verify that the Fraunhofer approximation will be valid in this case, the

requirement for operation in the far-field must be fulfilled: $R \gg \frac{1}{4} k D_R^2$. Using a

nominal wavelength of 3.51mm and antenna diameter of 30.5 cm [3] the limiting distance

is approximately 50m, much less than a typical application where the range will be on the order of 1km.

Applying the approximation to determine the complex amplitude of the electromagnetic field $U(x_0, y_0)$ at a plane a distance z from the antenna yields

$$U(x_0, y_0) = \frac{e^{jkz}}{j\lambda z} \sqrt{P_T \frac{\pi D_R^2}{4}}. \text{ This implies that } E(z, t) = \sqrt{P_T \frac{1}{(\lambda z)^2} \frac{\pi D_R^2}{4}} e^{j(kz - \omega t)} \text{ which is, as}$$

we desired in the far-field region where the target will be located, a plane wave. Define,

$$\text{for later convenience, } A = \sqrt{P_T \frac{1}{(\lambda z)^2} \frac{\pi D_R^2}{4}} \text{ to be the real amplitude of this wave.}$$

Now we examine the target geometry to determine what effect interaction with the target will have on the waveform. For reflection off of a flat mirror-like plate, let the known Fraunhofer plane be normal to the incoming radar pulse, through the center of the plate. The aperture will have dimensions of height z_1 and width $x_1 \cos\theta$ or $y_1 \cos\theta$, depending on whether the pose is near $\theta=0$ or $\theta=\pi/2$, i.e. upon which face of the block is exposed to the radar. At the small angles off normal incidence, angles for which Snell's law behavior will be appropriate, we can approximate $\cos\theta$ as unity. For simplicity, we then let the width be w . Outside of this aperture ($|x_1| > \frac{w}{2}$ or $|y_1| > \frac{z_1}{2}$) assume that $U(x_1, y_1)=0$. Within the boundaries of this rectangular region, the wave will have to travel a slightly longer or shorter distance to reach the imaginary $z=0$ reference plane, as shown in Figure 3.11.

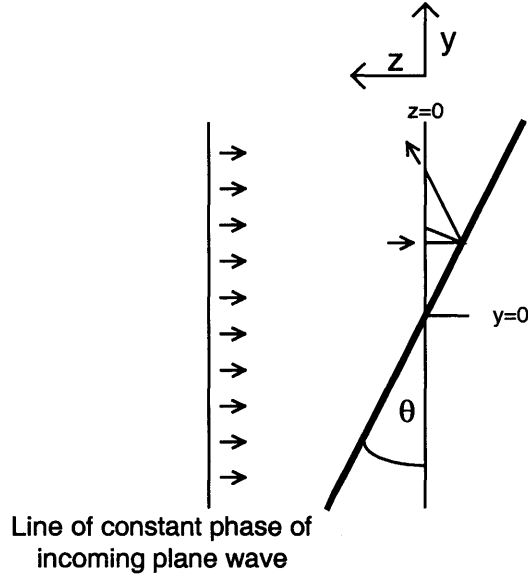


Figure 3.11: Extra path length traveled by a wave reflecting off of a tilted plate

A wave that enters the plane $z=0$ at a position y must travel an additional distance of $y \tan \theta$ before being reflected off of the plate. Then, to return to the imaginary plane, it must travel an additional distance $y \tan \theta / \cos 2\theta$, for a total extra length of

$y \tan \theta \left(1 + \frac{1}{\cos 2\theta} \right)$. However, that wave now has a y -coordinate of $y(1 + \tan \theta \tan 2\theta)$.

Taking the ratio of these quantities determines the excess distance traveled by a wave which leaves the $z=0$ plane with a normalized y -coordinate. Substitution with standard trigonometric identities reveals that a wave leaving the plane at y_1 will have traveled $y_1 \sin 2\theta$ greater path length than a wave leaving at the center of the plate, where $y_1 = 0$.

Accordingly, the phase will be shifted along that plane by a factor of $ky_1 \sin 2\theta$.

However, the real amplitude of the wave A will remain unchanged, so that

$$U(x_1, y_1) = Ae^{jky_1 \sin 2\theta}.$$

The Fraunhofer approximation can now be applied a second time, to determine the effect that this reflection from the target will have on the waveform received back at the radar antenna. In this far-field regime the specular return (in units $\sqrt{\text{Watts}}$) is

$\sqrt{\frac{\pi D_R^2}{4}} U(x_0, y_0)$, where

$$U(x_0, y_0) = \frac{e^{jkz}}{j\lambda z} A \int_{-\frac{z_l}{2}}^{\frac{z_l}{2}} dx_1 \int_{-\frac{w}{2}}^{\frac{w}{2}} dy_1 e^{jk \sin 2\theta y_1}$$

$$= \frac{e^{jkz}}{j\lambda z} A z_l \frac{1}{\frac{k}{2} \sin 2\theta} \text{sinc}\left(w \frac{k}{2} \sin 2\theta\right).$$

Using a small angle approximation of $\sin \theta = \theta$, we see that the signal detected at the receiver is again a plane wave, but one whose amplitude is a sinc function with respect to θ . In fact, the sinc resulting from reflection off of this perfectly flat surface is extremely narrow: the first zero-crossing will occur when the plate is tilted only $\theta = \frac{\pi}{wk} = \frac{\lambda}{2w}$

radians off from normal. For even a very small target with a width of 1m, that becomes an offset of 1/10 of a degree.

The preceding analysis is for a monochromatic source. The HRR system we will consider uses a carrier frequency ω_0 modulated by a Gaussian pulse envelope of width τ , such that the range resolution of the radar $R_R = \frac{c\tau}{2}$. In the frequency domain, this pulse transforms into another Gaussian

$$\frac{A}{\sqrt{2\pi\tau^2}} e^{-\frac{t^2}{2\tau^2}} \longleftrightarrow A e^{-\frac{\omega^2\tau^2}{2}}$$

so that our passband pulse transforms into a Gaussian centered around the carrier frequency ω_0 :

$$\frac{A}{\sqrt{2\pi\tau^2}} e^{-\frac{t^2}{2\tau^2}} e^{j\omega_0 t} \longleftrightarrow A e^{-\frac{(\omega-\omega_0)^2 \tau^2}{2}}$$

For all frequencies, weight this by the amplitudes determined through the Fraunhofer approximation, remembering that the wavenumber k is a function of frequency and must be expressed as such in order to perform the integration in the inverse Fourier transform. The results of this calculation show that two Gaussian pulses are received at the antenna. Each is shifted in phase and in delay as if it had been reflected from a single reflector, normal to the plane wave, located at the two corners of the plate. However, because the width of the pulse is so great compared to the period of the carrier, very little separation of the pulses can occur as the target rotates before the sinc function decays to an insignificant value. In light of this, and of the fact that almost no significant information may be gained if reflection only occurs at extremely small angles, for purposes of analysis we shall arbitrarily broaden the width of the sinc. It shall be approximated as a Gaussian with width σ , so that the mean value of the return in a single range bin will be

$$K e^{-\frac{\theta^2}{2\sigma^2}}, \text{ where } K = \frac{1}{\lambda R} \sqrt{\frac{\pi D_R^2}{4}} A z_l w = \sqrt{P_T \frac{1}{(\lambda R)^2} \frac{\pi D_R^2}{4} \frac{z_l w}{(\lambda R)}}.$$

The angular width over which the reflection is seen σ shall be somewhat broadened from the theoretical value computed from the model. This is done so that detection over a wider range of orientations is possible; however, to keep the analysis under control, the value of σ shall still be small enough such that the return from a surface remains in a single range bin.

These amplitudes can be found for each surface of the tank, so reflection return will appear in our specular model, with different amplitudes, in three range bins.

The preceding, seemingly arbitrary angular broadening of the specular reflection from a single plate is actually quite appropriate. A tank does not consist of three large reflectors, but rather of many much smaller facets. A smaller facet implies a wider sinc function, which we can approximate fairly well with a Gaussian. The Fraunhofer analysis method can be extended to returns occurring from multiple facets in as many range bins as desired, although additional work would be required to properly account for multiple-facet reflections.

The statistical model for the case of pure specular reflection, therefore, involves a distribution with both variance λ and mean m parameters, as opposed to the purely Lambertian model for which there was no coherent reflection parameter m . We shall model $m(\theta)$ as the artificially broadened Gaussian described above which will appear, for a single reflective surface, in a single range bin. It is taken to be zero in all other range bins. The variance parameter λ in this case is due solely to thermal receiver noise, and is therefore not a function of the orientation and will be constant across all range bins.

3.4 Composite Target Model

The Lambertian and specular models are two limiting cases, appropriate when the all of the power reflected off of the target surface is found in either coherent or completely incoherent radiation. Most real targets, however, reflect a combination of the two types of radiation. Although the analysis of such a model is not studied in this thesis,

the Rician parameters m and λ can be straightforwardly determined for an appropriate statistical probability distribution. To do so, another descriptive quantity of the physical system must be defined. Let a fraction α of the total reflected power P reflected by the target's physical cross section due to specular reflection, and then the remaining fraction $(1 - \alpha)$ will be due to incoherent reflection. The mean and variance parameters can be computed independently, as described above in this chapter. Since the total reflected power is the sum from the two types of reflection, then $m_{composite}(\theta) = m_{specular}(\theta)\sqrt{\alpha}$ and $\lambda(\theta) = \lambda_n + \lambda_{g,Lambertian}(\theta)(1 - \alpha)$ are the appropriate mean and variance parameters for case of composite reflection, where $0 \leq \alpha \leq 1$.

Chapter IV

Cramér-Rao Bounds

The general problem of *estimation* is to provide a choice for the value, i.e. an estimate, of some quantity, taken from a continuous set of possibilities. Depending on the quantity we desire to estimate, there are two methods of thinking about estimation. In the first, the quantity is a random variable with a known prior density; in the second, it is a non-random parameter of the probability density of the data to which we are allowed access. For most of this analysis, the target's pose is treated as an unknown parameter θ . It is a set quantity, not a random variable, which we want to determine from the statistical radar-return data. The analysis proceeds by determining, for each possible true value of θ , how well an estimator of that orientation could perform. Performance is measured, as described earlier, by the Hilbert-Schmidt metric, which at high signal-to-noise ratios is approximately equal to the minimum-mean-squared error on pose angle estimation. The Cramér-Rao bound, which is a lower bound on the variance of any unbiased estimator, is used to put a best-case limit on that performance. A drawback of the Cramér-Rao bound is that it is only valid for unbiased estimates of the unknown parameter; the implications of that problem are addressed in the cases in which it becomes an issue. Finally, an estimation of the overall performance of pose estimators, when the pose is viewed as a random variable with a uniform distribution, is found by averaging the error variance over the set of all possible orientations.

4.1 Cramér-Rao Analysis

In general, the mean-squared error of an estimator $\hat{\theta}$ of an unknown non-random parameter θ is the sum of two components, the variance of the estimator and the bias.

The estimator $\hat{\theta}$ is a function of the data provided through some statistical distribution based on the parameter θ , which means that $\hat{\theta}$ is itself a random variable with mean

$m_{\hat{\theta}} = E[\hat{\theta}]$ and variance $\sigma_{\hat{\theta}}^2 = E[\hat{\theta}^2] - m_{\hat{\theta}}^2$. The mean-squared error $E\left[(\hat{\theta} - \theta)^2\right]$ is then

$\sigma_{\hat{\theta}}^2 + (m_{\hat{\theta}} - \theta)^2$, the variance added to the square of the bias. Most often, we are

interested in finding the unbiased estimator with the minimum possible variance to

minimize the total mean-squared error. Cramér-Rao analysis is a tool which, by

linearizing the behavior of the probability distribution for the data around an operating

point, allows us to do precisely that. There do exist cases, however, especially with

distributions which are highly non-linear functions of the parameter which we desire to

estimate, for which no unbiased estimator exists, or for which a small bias may be a good

tradeoff for a large reduction in the variance of the estimator. When this case applies to

the problem of pose estimation with HRR data, some alternate methods of approximating

the minimum-mean squared error will be discussed.

For the most part, however, and especially in the good performance region of high

signal-to-noise ratios, an unbiased estimate is the most desirable. The Cramér-Rao bound

is therefore the primary tool used to place performance bounds on the problem of

orientation estimation. It achieves this goal by determining how sensitive the statistical

distribution of the data is to changes in the unknown non-random parameter, here the

orientation angle θ . The sensitivity measurement that is used is the Fisher Information [7]:

$$I_y(\theta) = E \left[\left(\frac{\partial}{\partial \theta} \ln p_y(y | \theta) \right)^2 \right],$$

where y is the data to which the estimator has access. Also, an alternate, equivalent formulation of the Fisher Information can be shown to be:

$$I_y(\theta) = -E \left[\frac{\partial^2}{\partial \theta^2} \ln p_y(y | \theta) \right].$$

Additionally, given access to a random vector \mathbf{y} composed of statistically independent components, the total Fisher information provided by this vector equals the sum of the Fisher informations provided by the individual components.

The Cramér-Rao bound is the inverse of the Fisher Information:

$$\sigma_{\hat{\theta}}^2(\theta) \geq \frac{1}{I_y(\theta)} \quad [7]$$

As demonstrated earlier, the return data from an HRR radar in each bin is statistically modeled to have a Rician distribution, with parameters $\lambda(\theta) = \lambda_g(\theta) + \lambda_n$ and $m(\theta)$ determined by the target's geometry, orientation, and the radar receiver noise:

$$p_r(r | \theta) = \frac{2r}{\lambda} e^{-\frac{(r^2+m^2)}{\lambda}} I_0\left(\frac{2mr}{\lambda}\right), r \geq 0.$$

The Fisher Information of each bin is then found to be very difficult to work with analytically, involving several higher-order modified Bessel functions. To avoid that problem, only the two limiting reflectivity cases, completely Lambertian and completely specular, are discussed in this thesis.

4.2 Analysis of Rough-Surface Target Model

In the Lambertian reflection model, the mean parameter m of the Rician distribution is zero. The probability distribution for each bin then become Rayleigh:

$$p_r(r|\theta) = \frac{2r}{\lambda} e^{-\frac{r^2}{\lambda}}, r > 0$$

where $\lambda = \lambda(\theta, J) = \lambda_g(\theta, J) + \lambda_n$. The Fisher Information about θ in the data r is found to be:

$$I_r(\theta) = \frac{1}{\lambda^2} \left(\frac{\partial \lambda}{\partial \theta} \right)^2$$

For the reflection pattern of a single rectangular block, there are three main effects which contribute to the change of variance with respect to pose in each bin. The first of these is the change in the reflection of the front face of the block. When this face spans the entire bin width, its contribution to variance is $\lambda_{gJ, \text{frontface}}(\theta) = \kappa z_l \cot \theta \cos \theta$. We then find

$\frac{\partial \lambda}{\partial \theta} = -\kappa z_l \cos \theta (\csc^2 \theta + 1)$. Similarly, there is a contribution from the side face, if that

face also spans the entire bin, of $\lambda_{gJ, \text{sideface}}(\theta) = \kappa z_l \tan \theta \sin \theta$. This term has derivative

$\frac{\partial \lambda}{\partial \theta} = \kappa z_l \sin \theta (\sec^2 \theta + 1)$. The third, slightly more complicated, effect occurs when the

bin under investigation contains one of the block's edges, i.e. one of the boundaries (d_1 , d_2 , or d_3 defined in Section 3.2) of the spatial regions for which the above two faces may contribute to the total reflection strength. In this case, both the contributions of the two faces change, as do the fractions of the bin in which the reflection is received.

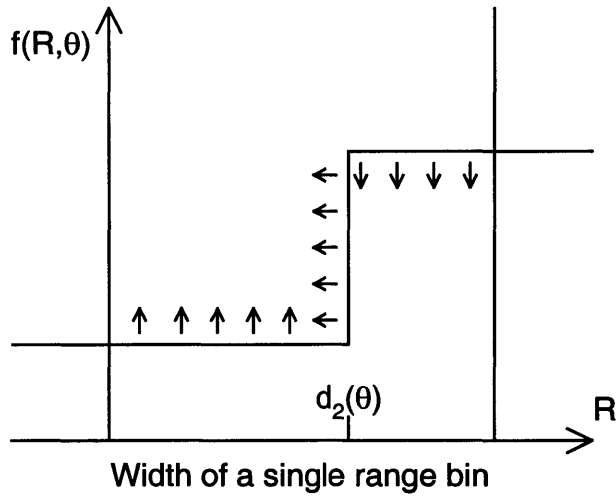
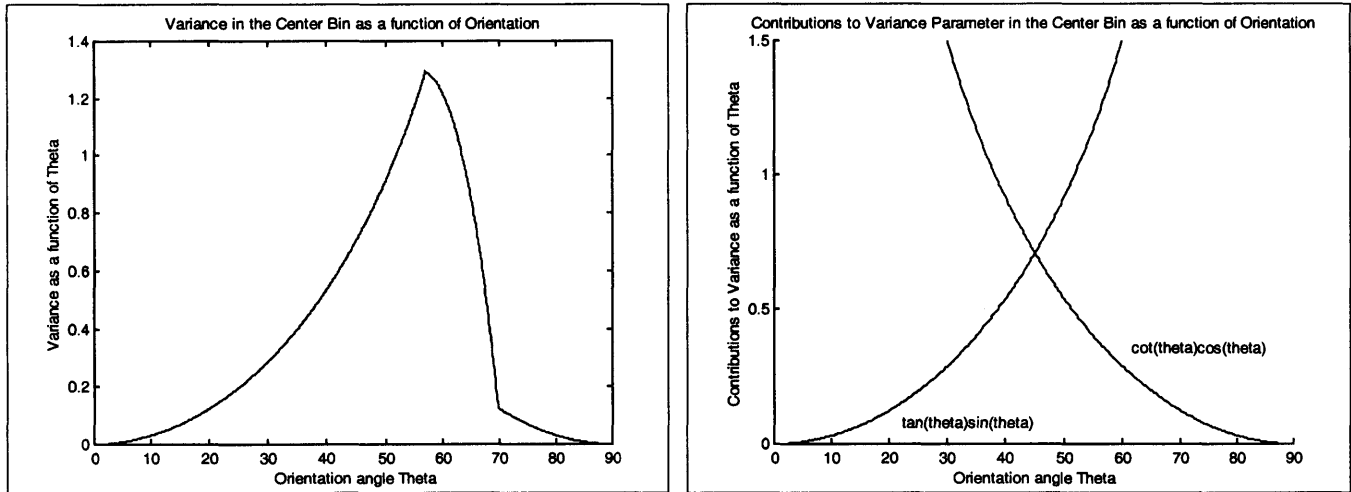


Figure 4.1: Three effects contribute to the change in variance parameter λ_g as orientation changes and a block's edge moves through a range bin.

Recall from Section 3.2 that the λ_g parameter in each bin is the average, over the range, of the function $f(R, \theta)$ shown in Figure 4.1 above. The change in $f(R, \theta)$ as the pose angle θ increases is indicated by the arrows; the sum of these changes integrated over the width of the entire range bin becomes the change in the variance of the range bin, $\frac{\partial \lambda}{\partial \theta}$. As the edge d_2 enters or leaves this particular range bin, the derivative $\frac{\partial \lambda}{\partial \theta}$ as a function of θ will be discontinuous, even though the variance $\lambda(\theta)$ remains a continuous function. As will be shown, this implies that the Fisher information and therefore the Cramér-Rao bound are discontinuous functions of the orientation θ .

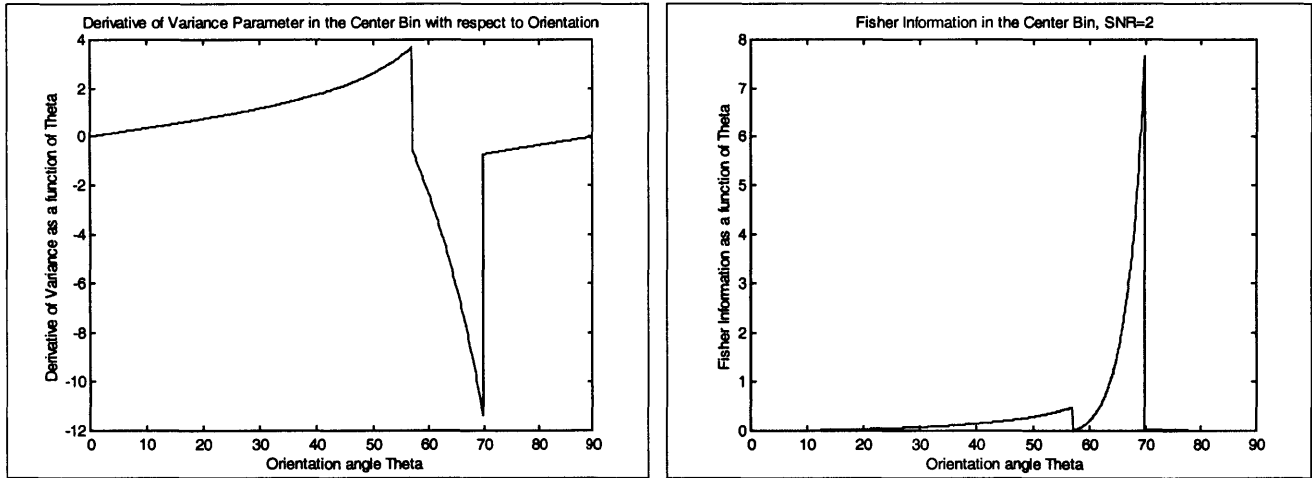
Following, we have plotted the variance parameter (Figure 4.2a), its derivative (Figure 4.3a), and the Fisher information for a single range bin (Figure 4.3b), for a rectangular block with unit width whose depth is twice that large, and whose height equal its width. The range resolution is four range bins per unit width .



(a) (b)

Figure 4.2: The variance parameter λ is composed of contributions from the front face, proportional to $\cot(\theta)\cos(\theta)$, and from the side face, proportional to $\tan(\theta)\sin(\theta)$.

When the block is at the $\theta=0$ orientation, all of the front face is found much closer to the radar than the center of the block (at the distance $y/2$, precisely), so there is no reflection from that face received at the radar in this bin. Also, the side face of the block is parallel to the radar's line of sight, so it has no visible cross section at this orientation. Thus, there is no power received by the radar at this orientation, in this bin, due to reflection from the side of the target, either. However, as θ increases, a small cross section from the side face appears and grows in magnitude as the face becomes closer to head-on with the radar. At some point, then, the back edge of the side face begins to leave the bin, and the variance of the bin decreases as less cross section is available for reflection. Around the same angle, the back edge of the front face then begins to enter the bin as well, and the contribution from this face can be seen clearly at the far right of Figure 4.2a. In, Figure 4.2b, we show the two trigonometric functions, whose contributions represent those of the front and side faces, respectively.



(a) (b)
Figure 4.3: The derivative of the variance and the Fisher Information.

The discontinuities in the derivative due to the back corner of the side face entering and leaving this particular bin are clearly visible in Figure 4.3a. The two plots in Figure 4.3 also vividly demonstrate that the Fisher information provided at the angles at which the corner is moving through the bin is significantly greater than that available at other orientations. In the more complete case, when the data from multiple bins is available to aid in pose estimation, it is therefore the bins that contain edges that are most helpful in refining the estimate.

Signal-to-Noise Ratio (SNR) is defined as for a rectangular block to be the ratio

$$\frac{\lambda_{G,J}(\theta)}{\lambda_n}$$

where the denominator is variance parameter due to thermal noise. The

numerator is the parameter due to geometry, measured in the bin at which the larger of the front and side faces appears, at the orientation such that this face is normal to the incoming radar beam.

To obtain the Fisher Information contained in the entire vector of range information, the individual Fisher Informations of each bin are summed. Figure 4.4 is a

plot of the total Fisher information versus pose angle θ from a 1m x 2m x 1m sized rectangular block, with a range resolution of 0.25m, at a SNR of 2.

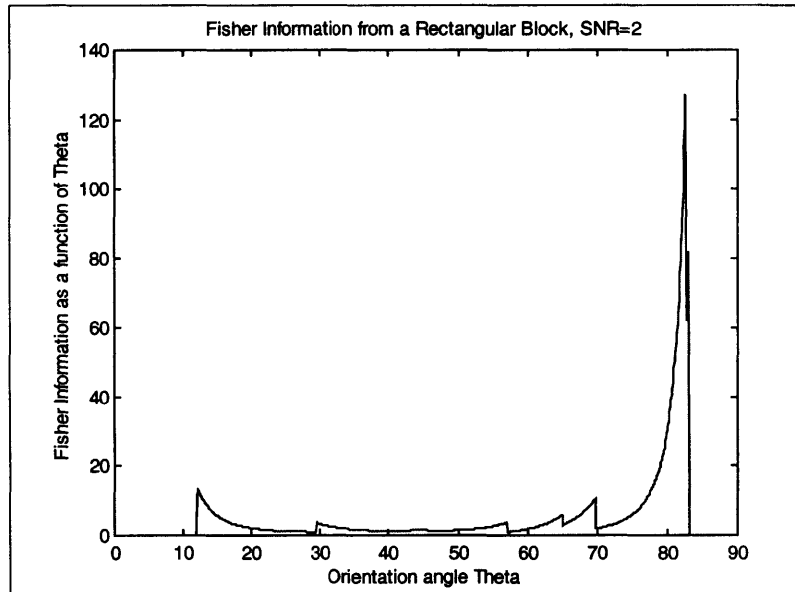


Figure 4.4 Fisher Information as function of Orientation

The Cramér-Rao bound, the inverse of the Fisher Information and the focus of this thesis, for this block is plotted versus pose θ in Figure 4.5 below.

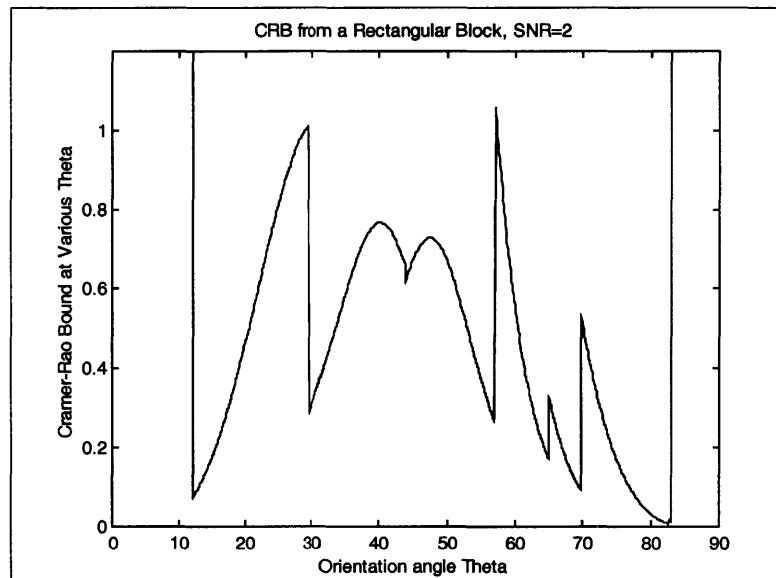


Figure 4.5: The Cramér-Rao bound on Pose Estimation of a Rectangular Block

There are several features of this bound which deserve further investigation. The first of these is the striking jaggedness of the bound, i.e. the large discontinuities due to the corners of the block entering and leaving range bins. At first glance, it seems plausible that, because any particular corner is always within one bin or another, there is always at least one bin with a large Fisher Information. However, that argument is not exactly correct. There is, at most angles, at least one bin with a large derivative $\frac{\partial \lambda}{\partial \theta}$.

However, consider the case illustrated below in Figure 4.6. Because the largest component of the derivative in either bin is due to the movement of the boundary




Figure 4.6: The function $f(R, \theta)$ as it changes with increasing angle θ

$d_2(\theta)$, the derivative in the first bin after this corner switches bins will be approximately equal to the value of the derivative in the bin on the right just before that occurrence takes place. However, the value of the variance λ itself will be much smaller in the left bin at that instant. Since the Fisher Information is the square of the ratio of these quantities, the

total Fisher Information will be dramatically larger at the angles slightly after the bin crossing than slightly before. Another consequence is that the information will decrease as the edge travels further through the bin, an effect that can be seen in the plot of the Cramér-Rao bound above.

The ultimate cause of these discontinuities is the assumption that range to the center of the block is known exactly. In practice, this distance will not usually be known to any greater accuracy than the range resolution, making the exact angle at which one of the block's corners traverses from one range bin to another a moot point. The cross section/obliquity function $f(R,\theta)$ could be smoothed to represent this uncertainty in range knowledge; for example, we could assume that the actual absolute range to the target is a random variable with a normal distribution and modify the function accordingly. However, this will likely result simply in a leveling out of the Fisher Information and therefore the Cramér-Rao bound (CRB). Since the ultimate goal is to integrate the CRB over a uniform prior density, i.e. average its value over all possible orientations, the final effect of this smoothing would be minimal.

A second observation is that the Cramér-Rao bound predicts, on the average, a greater ability to estimate pose for orientation angles $\frac{\pi}{4} \leq \theta \leq \frac{\pi}{2}$ than it does for the angles $0 \leq \theta \leq \frac{\pi}{4}$. This agrees with intuition. At the greater angles, the larger face is presented towards the radar. Accordingly, there is a greater cross section available to reflect the radar signal, and better information available to aid in estimation.

Finally, at orientations near $\theta = 0$ and $\theta = \frac{\pi}{2}$, there is a dramatic increase in the Cramér-Rao bound; the bound very quickly heads towards infinity and there is a

singularity located at each of those angles. It can be shown that these increases occur when a block face is completely contained within one range bin. This effect is highly counter-intuitive; one would expect that at these angles, the radar return profile will be very distinctive – almost all of the return is located in one single bin. It is precisely that fact, that nearly all of the return is to be found in one bin, that leads to such a large bound. Over the range of angles for which a face does appear in a single bin, the return is distinctive, but changes very little from one orientation to another. Since Cramér-Rao bound analysis is completely based on the determination of how quickly a random vector or variable changes with respect to changes in the parameter which is to be estimated, an invariant statistical distribution for the random variable provides virtually no information about what, within that particular range, the pose might be. In fact, the Cramér-Rao bound is derived from a linearization of the probability distribution at a certain point. Thus, because HRR can make no distinction between a block oriented at $+\phi$ and one at $-\phi$, there is necessarily an inflection point at $\theta = 0$ where there is no change in the statistical distribution. At this point, the Fisher Information is necessarily zero and the CRB is infinite.

This singularity occurrence reflects the difficulty of creating an unbiased pose estimator for these special orientations. However, to show that it is possible to make a somewhat good estimate of the pose at these orientations, a method other than Cramér-Rao analysis may be used. If the question “Is the pose angle such that the entire front face of the target block is to be found within one range bin?” is asked, the overall problem becomes one of decision between two alternatives rather than an estimation, with a continuum of possible results.

Specifically, the set of possible angles $0 \leq \theta \leq \frac{\pi}{2}$ can be divided into M different bins, each of width $\frac{\pi}{2M}$. For each angular bin, compute the variance $\lambda_{g,j}(\theta_m)$, where θ_m is the angle at the center of each bin, i.e. $\theta_m = \frac{\pi}{2M}(2m-1)$. Using a maximum likelihood binary decision rule, we can determine the probability that, in a choice between the true orientation θ_i and an alternative θ_j , the correct orientation is identified. The conditional probability of correct decision in this M-ary decision problem can then be bounded by the sum of the errors in the M-1 binary decision problems:

$$P[\text{declaring } \theta_i | \theta_i, \text{true}] \geq 1 - \sum_{j \neq i} P[\text{declaring } \theta_j | \theta_i, \text{true}],$$

where the probability on the left side of the equation refers to the M-ary problem and the right side refers to a binary decision between θ_i and θ_j only.

Further, the set of probabilities $P[\text{declaring } \theta_j | \theta_i, \text{true}]$ can be used to approximate the minimum mean-squared error of a maximum likelihood estimator. The error is the sum of the error from choosing the wrong angular bin, in addition to an error from declaring the orientation to be the particular angle at the center of the bin, when it could be uniformly distributed within it:

$$E\left[(\hat{\theta} - \theta)^2\right] \approx \sum_{j \neq i} (\theta_i - \theta_j)^2 P[\text{declaring } \theta_j | \theta_i, \text{true}] + \frac{1}{12} \left(\frac{\pi}{2M}\right)^2$$

This formula would provide a more precise value for the mean-squared error if the probabilities used were the error probabilities in the M-ary decision problem, rather than the binary error probabilities. However, especially for a vector of random variables, which the range profile provides in this case, finding the binary probabilities is a solvable

problem, whereas the M-ary maximum-likelihood error probability is much more difficult to find. In addition, if the true orientation is θ_1 , then a minimum on the estimation error can be found by assuming that, in a binary decision between θ_2 and $\theta_{j>2}$, θ_2 will always be chosen. Any cases in which this choice does not occur will increase the mean-square error of the result, so that

$$\left(\frac{\pi}{2M}\right)^2 P[\text{declaring } \theta_2 | \theta_1 \text{ true}] + \frac{1}{12} \left(\frac{\pi}{2M}\right)^2 \text{ becomes a minimum bound.}$$

To compute the binary probabilities, we use the Van Trees approximation for Receiver Operating Characteristics (ROC). [11] In a likelihood ratio test with data \mathbf{X} and the two hypotheses θ_i and θ_j , the log-likelihood

$$\ell(\mathbf{X}) \equiv \ln \left[\frac{P_{\mathbf{x}|\theta_i}(\mathbf{X}|\theta_i)}{P_{\mathbf{x}|\theta_j}(\mathbf{X}|\theta_j)} \right]$$

is compared to some threshold γ in order to determine whether to declare hypothesis θ_i or to declare θ_j . To choose the hypothesis with greatest probability of being true, the threshold γ is set to zero.

The Van Trees approximation transforms the coordinate system such that

$$\mu(s) \equiv \ln(E[e^{s\ell} | \theta_j]), \quad 0 \leq s \leq 1$$

and concludes that

$$P[\text{declaring } \theta_j | \theta_i \text{ true}] \approx Q[(1-s)\sqrt{\dot{\mu}(s)}] \exp \left[\frac{(1-s)^2 \ddot{\mu}(s)}{2} + \mu(s) + (1-s)\dot{\mu}(s) \right]$$

where $Q(x) \equiv \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}}$ and $\dot{\mu}(s) = \gamma$.

To apply these results to the problem at hand, i.e. the rough-surface model of radar return, we examine the square of the data that we considered earlier. These values then take an exponential, rather than Rayleigh, statistical distribution in each bin; however the variance parameters λ remain unchanged. For vectors \mathbf{P} of independent exponentially distributed random variables the log-likelihood ℓ is:

$$\ell(\mathbf{P}) = \ell_B + \sum_{allJ} p_J \left(\frac{1}{\lambda_J(\theta_j)} - \frac{1}{\lambda_J(\theta_i)} \right)$$

where $\ell_B \equiv \sum_{allJ} \ln \left(\frac{\lambda_J(\theta_j)}{\lambda_J(\theta_i)} \right)$. It is also verified that $\mu(s)$ converges for all $0 \leq s \leq 1$ for an exponentially distributed random vector. The semi-invariant moment-generating function, $\mu(s)$, is then given by:

$$\mu(s) = s\ell_B - \sum_{allJ} \ln \left[1 + \left(\frac{\lambda_J(\theta_j)}{\lambda_J(\theta_i)} - 1 \right) s \right]$$

The derivatives $\dot{\mu}(s)$ and $\ddot{\mu}(s)$ follow straightforwardly.

Next, we are ultimately interested in finding the probability of error using a decision rule for which the log-likelihood threshold γ is zero, i.e. the maximum likelihood decision region for any data vector \mathbf{P} . Since $\gamma = \mu(s)$ is a monotonically increasing function of s , we can use Newton's approximation method to find the s for which $\dot{\mu}(s) = 0$ without fear of the algorithm looping within a local extrema rather than converging to the correct result. With knowledge of the appropriate s , $\mu(s)$ and $\ddot{\mu}(s)$ can be calculated, and therefore the error probabilities for each of the M possible are computed.

For the same rectangular block analyzed for the Cramér-Rao bound earlier in this section, we can compute the probabilities of declaring θ_j when the true orientation is θ_1 , the center orientation of the angular bin in which the front face is located completely within a single range bin.

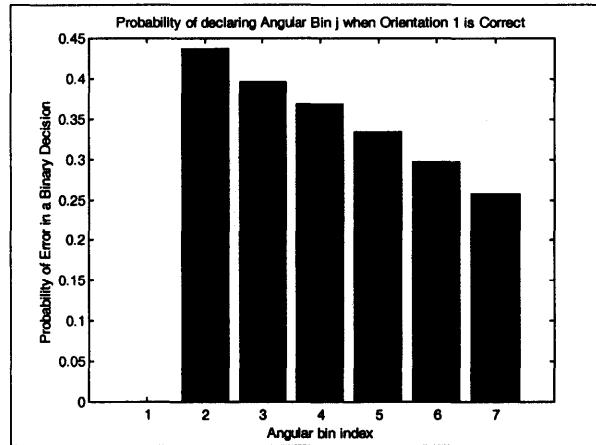


Figure 4.7 Binary decision probabilities of error.

In figure 4.7 above, the probabilities of declaring θ_j when θ_1 is the true orientation are shown for the case of SNR=2. These probabilities are so large that the SNR used for these calculations clearly falls outside the good-performance region that we desire to work in. Nonetheless, the bound on MMSE can be obtained from this data, a similar calculation is done when the true orientation is near $\frac{\pi}{2}$, and the results, once grafted onto the plot of the CRB obtained earlier, provide the composite bound is shown in Figure 4.8.

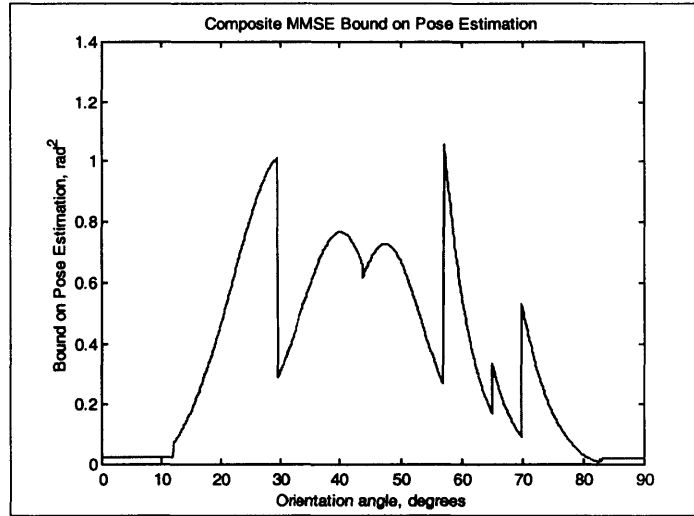


Figure 4.8 Bound obtained using both CRB and Van Trees error estimation

All of this analysis was performed using a single rectangular block as a target; the extension to a three-block tank is straightforward and will be discussed in Chapter V.

4.3 Analysis of Specular Target Model

It can be shown that the Rice distribution,

$$p_r(r|\theta) = \frac{2r}{\lambda} e^{-\frac{(r^2+m^2)}{\lambda}} I_0\left(\frac{2|m|r}{\lambda}\right), r > 0$$

can be approximated, when $m^2 \gg \lambda$, by a Gaussian probability distribution with the same mean and whose variance is one-half that of the Rician's parameter:

$$p_r(r|\theta) = \frac{1}{\sqrt{2\pi \frac{\lambda}{2}}} e^{-\frac{(r-m)^2}{2\left(\frac{\lambda}{2}\right)}}$$

As a demonstration, Figure 4.9 shows a concurrent plot of the two probability distributions.

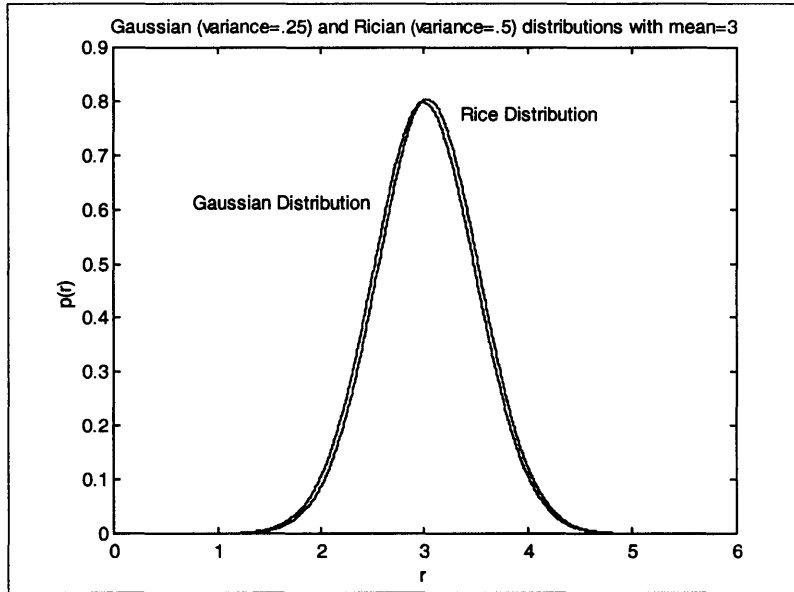


Figure 4.9: Comparison of Gaussian and Rician Distributions

Under the model described in Chapter III, the radar return profile for a rectangular block will consist of only noise in every bin unless face of the target is close to being normal to the radar beam. In this case, that range bin will have a Gaussian distribution,

with mean $m = Ke \frac{\theta^2}{2\sigma^2}$ and accordingly, full statistical description

$$p_r(r|\theta) = \frac{1}{\sqrt{2\pi \frac{\lambda}{2}}} e^{-\frac{\left(r - Ke \frac{\theta^2}{2\sigma^2}\right)^2}{2\left(\frac{\lambda}{2}\right)}} .$$

Applying Cramér-Rao analysis to a Gaussian distribution with mean $m(\theta)$ and constant

variance $\frac{\lambda}{2}$ yields the Fisher Information $\frac{2}{\lambda} \left(\frac{\partial m}{\partial \theta}\right)^2$. For the mean of this model, the

Cramér-Rao bound becomes $\frac{\lambda}{2} \frac{\sigma^4}{K^2 \theta^2 e^{\frac{\theta^2}{\sigma^2}}}$. This bound is plotted in Figure 4.10a, below

for a $\frac{K^2}{\lambda}$ Signal-to-Noise ratio of 25 and a reflection-decay standard-deviation parameter σ of 4 degrees.

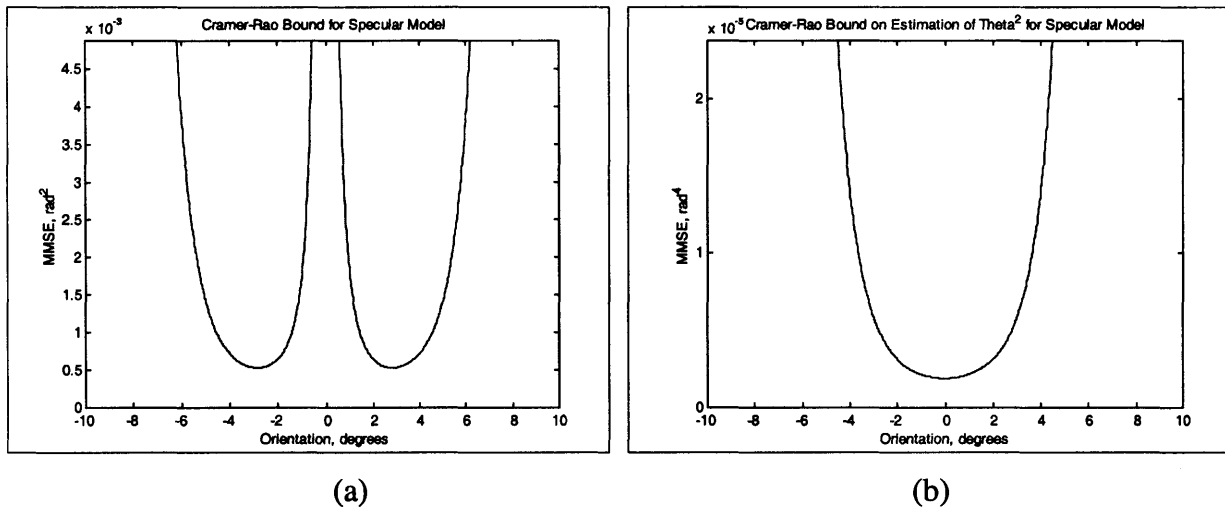


Figure 4.10 The Cramér-Rao bounds on the estimation of θ and θ^2 .

The Cramér-Rao bound on the estimation of θ in the specular case, just as with the rough-surface model, has a singularity at $\theta = 0$, when the face is perpendicular to the radar beam. Again, intuition would expect that at these smallest angles it should be easier to make a good estimate of pose, because at this orientation the mean takes its greatest value in comparison to the noise. Greater signal implies that estimation should be more accurate.

Remembering that HRR data can give no indication on whether a symmetric target is oriented at $\theta = +\delta$ or at $\theta = -\delta$, it is useful to determine the Cramér-Rao bound on the estimation of θ^2 . Normally, in an estimation of the square of a parameter, any

information on the sign of the value would be lost, but since in this case there is no sign data in the profile it makes a logical choice. Computing the Fisher Information

$\frac{2}{\lambda} \left(\frac{\partial m}{\partial \theta^2} \right)^2$ produces a Cramér-Rao bound of $\frac{4 \frac{\lambda}{2} \sigma^4}{K^2 e^{\frac{\theta^2}{\sigma^2}}}$. This is also plotted in Figure

4.10b. above, on an equivalent scale. (The vertical axis for the plot on the left ranges from 0 to σ^2 , while on the right plot it varies from 0 to σ^4 .)

To attempt to compare the Cramér-Rao bound on the estimation of θ^2 with the mean-squared error of a specific estimator, the mean-square error $E[(\theta^2 - \hat{\theta}^2)^2]$ of the maximum-likelihood estimator $\hat{\theta}^2 = -2\sigma^2 \ln\left(\frac{r}{K}\right)$ was computed. Using the approximation

$$\ln(1+z) \approx z - \frac{1}{2}z^2 \text{ for } |z| \ll 1$$

it can be shown that, for a Gaussian distribution with mean m and variance $\frac{\lambda}{2}$,

$$E[\ln x] \approx \ln m - \frac{1}{2} \frac{\lambda/2}{m^2} \text{ when } m^2 \gg \lambda. \text{ Strictly speaking, the far tail of the}$$

Gaussian distribution must be ignored, because the logarithm of a negative number is an imaginary number and the integral will not converge in that region. However, because the Gaussian is only approximating the Rice distribution this is not a fundamental

problem. Further, approximating $E[(\ln x)^2] \approx (\ln m)^2 + (1 - \ln m) \frac{\lambda/2}{m^2}$ and substituting

these into the expansion of $E\left[(\theta^2 - \hat{\theta}^2)^2\right]$ produces a mean-square error of $\frac{4\lambda\sigma^4}{K^2 e^{\frac{\theta^2}{\sigma^2}}}$,

exactly the bound predicted through Cramér-Rao analysis.

Chapter V

Performance Comparisons

Up to this point, this thesis has described methods for computing minimum mean-squared error bounds for pose estimation of rectangular and blocks-world target models. A major goal of the work, however, is to determine the effects that variations in the physical parameters of the radar-target system have on these bounds. Much of this behavior is found to be reducible to the effects of varying the signal-to-noise ratio. For example, increasing the cross-section of a target may have exactly the same effect as reducing the value of the receiver noise by the same proportion. SNR has been previously defined in both specular and rough-surface cases as the ratio of maximum power received in a bin, due to reflection from the target, to the power resulting from the thermal receiver noise, allowing changes in many parameters to be absorbed into this one variable. Finally, the performance results obtained through the models are compared with Monte-Carlo simulations that find the Hilbert-Schmidt estimate and associated error for many trials of the statistical data.

5.1 Rough-Surface Performance Analysis

An examination of the form of the Fisher Information for each individual range bin in the rough-surface model produces a few observations. The first of these is that, even as signal-to-noise ratio grows without bound, there may be a finite limit on the fidelity of the pose estimate. In other words, the minimum mean-squared error cannot necessarily be reduced to an arbitrarily small value by increasing the SNR. Recall the

form of the Fisher Information for a single bin, based on a Rayleigh probability distribution:

$$\left(\frac{\partial \lambda_j / \partial \theta}{\lambda_{g,j} + \lambda_n} \right)^2$$

Defining $\tilde{\lambda}_{g,j} \equiv \frac{\lambda_{g,j}}{\max_{J,\theta} \lambda_{g,j}}$ as SNR-normalized variance for each bin, the total Fisher

Information available from the radar-return range profile becomes

$$I_R = \sum_J \left(\frac{\partial \tilde{\lambda}_{g,j} / \partial \theta}{\tilde{\lambda}_{g,j} + \frac{1}{SNR}} \right)^2.$$

The values of $\lambda_{g,j}$ and $\partial \lambda_j / \partial \theta$ are both proportional to the target's height z_1 , so increasing this parameter by a constant factor is equivalent to reducing the receiver noise by the same factor, i.e. the SNR is constant. However, even if we could increase the signal-to-noise ratio to infinity, perhaps by setting receiver noise to zero, there is still a maximum Fisher Information in each range bin. This implies that the CRB cannot be reduced to zero in this manner. While at very low signal-to-noise ratios, the Cramér-Rao bound will improve as $\frac{1}{SNR^2}$, the performance will eventually plateau. The only way to further increase the Fisher Information is to increase the number of bins in which the target appears; i.e. increase the length x_1 and the width y_1 or increase the range bin density $n = \frac{1}{R_{res}}$. For a given target, this means that increased range-bin density, which requires increased radar-pulse bandwidth, is necessary. Figure 5.1a below displays the variation

in the CRB, integrated over a uniform prior distribution, as SNR is increased via decreasing the variance contribution of receiver noise. In Figure 5.1b, the width and length of the tank are increased proportionally over a range of 0.4 to 2.5 times the size of the basic tank model. The tank used is the blocks-world model described in section 3.2, with a standard range-resolution [3] of 0.5 meters.

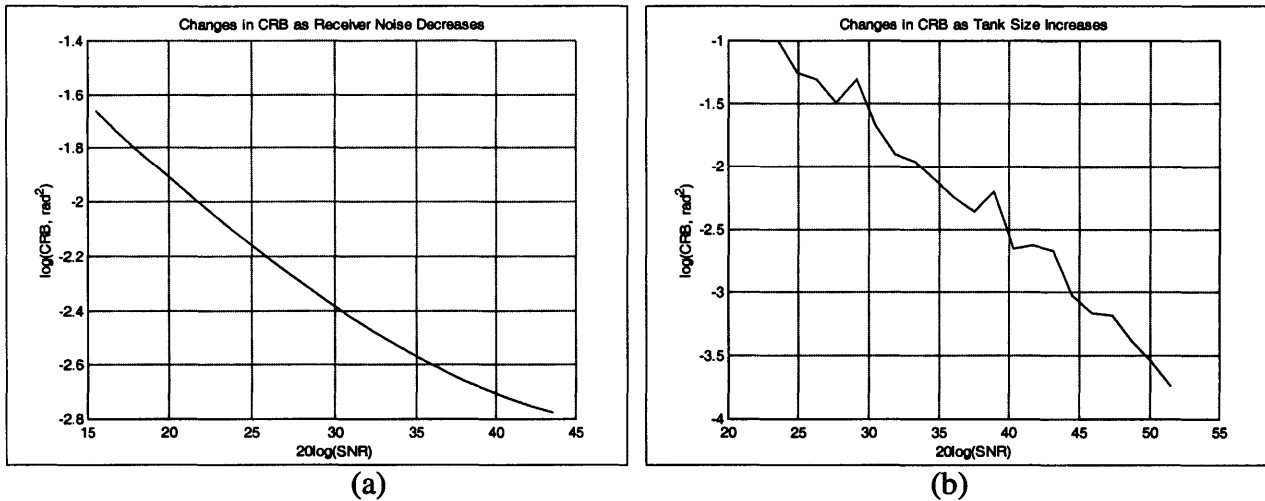


Figure 5.1: Plots comparing the change in CRB as SNR increases. In plot (a), receiver noise is decreased to provide better SNR. In plot (b), the sizes x_1 and y_1 are increased while keeping the receiver noise unchanged.

Figure 5.1b shows that, as the size of the target grows, the MMSE appears to decrease as $\frac{1}{\text{SNR}^2}$, possibly because there are two effects influencing an algorithm's pose estimation ability. First, there is the increase in the number of bins that provide a radar return, and therefore can provide applicable Fisher Information. Secondly, a longer moment arm for rotation means that, over the same range of orientations, the corners of the block will be sweeping each range bin at a faster pace, increasing the value of $\frac{\partial \lambda}{\partial \theta}$ in each bin by a factor while leaving λ_g invariant.

Increasing SNR by changing the target height z_1 or the receiver noise λ_n , however, only seems to produce a $\left(\frac{1}{SNR}\right)$ change in the MMSE in this regime, and we can see that as the SNR increases, the effect becomes even less pronounced. The limit on pose estimation ability, discussed above, is being reached and can be seen at the right extreme of Figure 5.2a as the CRB begins to flatten out.

The predictions of the Cramér-Rao bound can be compared to the mean-square error of the Hilbert-Schmidt estimator. The first step in determining the HSE is to generate a random range profile \mathbf{R} according to the statistical model for a particular orientation angle θ . Then the Hilbert-Schmidt estimate is found: for each of a set of N possible orientation angles, the conditional probability density

$p_{R|\theta_i}(R|\theta_i) = \prod_j p_r(R_j | \theta_i)$ is computed. Because we assume a uniform prior density for pose, we obtain the posterior distribution for pose θ by normalizing the conditional density such that $\alpha \sum_N p_{R|\theta_i}(R|\theta_i) = 1$. In a similar manner we compute

$E[\cos \theta] = \frac{1}{\alpha} \sum_{i=1}^N p_{R|\theta_i}(R|\theta_i) \cos \theta_i$ and $E[\sin \theta] = \frac{1}{\alpha} \sum_{i=1}^N p_{R|\theta_i}(R|\theta_i) \sin \theta_i$. The Hilbert-Schmidt estimate is then $\hat{\theta}_{HS} = \tan^{-1}\left(\frac{E[\sin \theta]}{E[\cos \theta]}\right)$. Finally, the squared-error $(\theta - \hat{\theta}_{HS})^2$ is

calculated for this trial. Over a set of many trials, the mean of these squared-errors will converge to the mean-squared error, which in the high-performance regime is approximately one-half of the Hilbert-Schmidt error. Figure 5.2 shows a comparison of the empirically calculated MSE (600 trials at each of 45 angles, SNR=46dB, for the

standard tank model) with the theoretically derived Cramér-Rao/Van Trees composite bound.

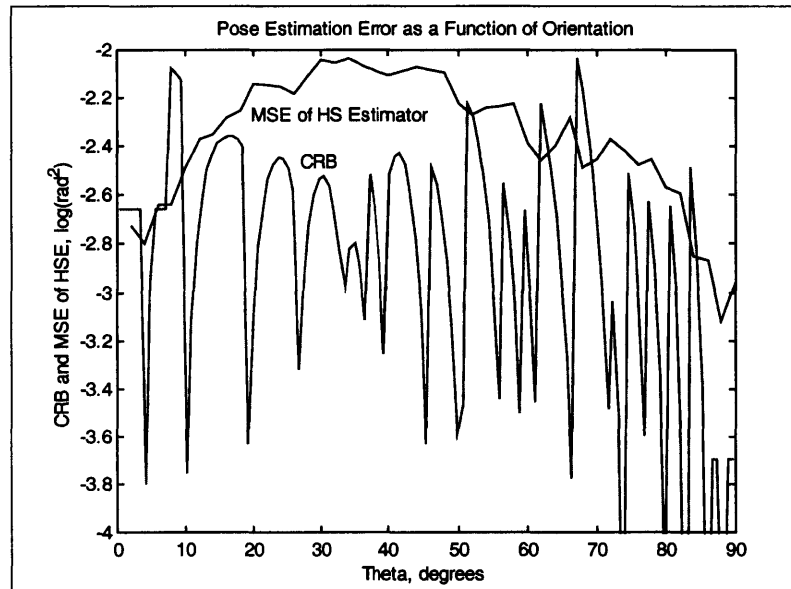


Figure 5.2: CRB and Mean Square Error of the Hilbert-Schmidt Estimator

At some points in the figure above, the minimum mean-squared error obtained by simulation becomes smaller than the Cramér-Rao bound, which is an apparent inconsistency – the CRB is by definition the minimum mean-squared error of any unbiased estimator. However, this is a manifestation of the same behavior which is seen at the two orientation extrema, $\theta = 0$ and $\theta = \frac{\pi}{2}$. Although estimators may exist which have better than CRB mean-squared errors at certain values of the unknown parameter, a deeper analysis of the estimator will show that it cannot be unbiased. It is expected that, as signal-to-noise ratio increases that an unbiased estimator, who best possible mean-squared error is predicted by the Cramér-Rao bound, will outperform all other valid estimators.

The variation of the bound with respect to changes in parameters can also be checked against the simulation results. Below, in Figure 5.3, is the same behavior captured in Figure 5.1a above, the changes in error variance with receiver noise. These plots also include the MSE calculated by the Monte-Carlo method for a comparison, using 250 trials at each of 90 poses for each value of SNR.

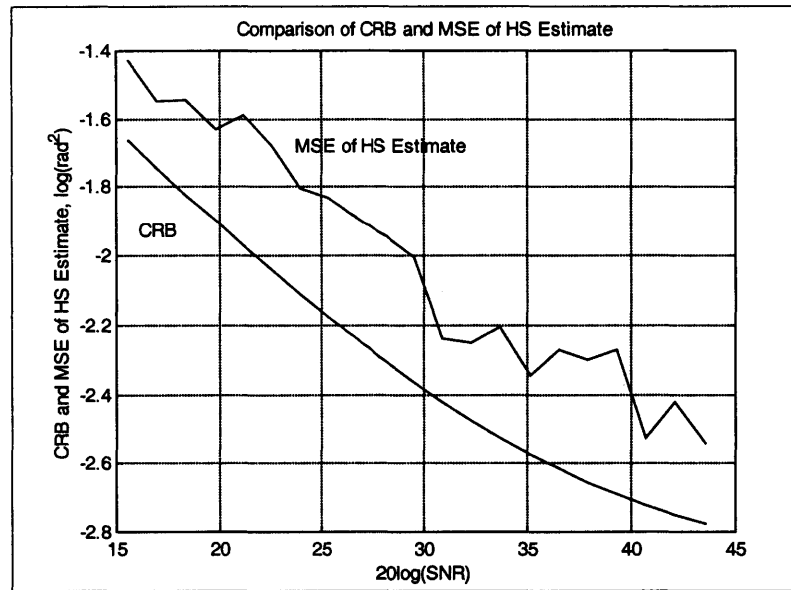


Figure 5.3: Comparison of CRB and MSE of Hilbert-Schmidt Estimate as SNR increases

5.2 Specular Results

By inspection of the Cramér-Rao bounds derived in Chapter IV for a single block,

$\frac{\lambda}{2} \sigma^4$, the error variance of an estimate of orientation angle θ , we see that the error $K^2 \theta^2 e^{-\frac{\theta^2}{\sigma^2}}$

decreases as $\frac{1}{SNR}$ for a signal-to-noise ratio defined as $\frac{K^2}{\lambda}$, maximum reflected power to receiver noise power. This ratio is a function of the peak power transmitted, as well as of the cross-section of the reflecting surface. Also, the angular range over which an

estimate can be made is a strong function, obviously, of the angular width over which the

reflection occurs. The CRB can be rewritten $\frac{\lambda}{2} e^{\frac{4 \ln(\sigma) \theta^2}{\sigma^2}}$ to show the dependence more

explicitly. From the physical model, this σ parameter will increase as the size of the reflecting surface decreases. A decrease in the surface's area then has two opposing effects, the decrease in SNR acting to cancel out some of the increase in angular width for reflection.

For multiple blocks whose reflections appear in different range bins and therefore cause no interference, the Fisher Information from each face can be added in order to compute a complete Cramér-Rao bound. This specular model can even be extended to a target with many small reflecting facets. Each plate may have a different placement on the target and therefore cause specular reflection at different orientation angles. To determine the radar return, however, the effects of interference, shadowing, and reflection off of multiple surfaces must be taken into account and the Fisher Information will take a more complicated form.

Chapter VI

Conclusions

In this thesis, we have created statistical models, based on physical principles, of High Range-Resolution Radar return profiles. Two limiting cases of reflectivity were studied, the rough-surface Lambertian reflection model and the smooth-surface, specular reflection model. We also determined how to compute the parameters of each statistical model for an arbitrary target composed of rectangular blocks, and then did so for a standard blocks-world model of a tank. Since these parameters vary as the orientation of the tank changes, the data from the statistically-generated range profiles can be used to make an estimate of that orientation angle.

The thesis then continues by determining performance bounds on the accuracy of algorithms intended to estimate that orientation. First, we determined the process by which to compute these bounds for the particular models used, and then evaluated the bounds for the tank target. We found the dependence of the bounds on the physical parameters of the system and showed how the error variance decreases inversely with increasing signal-to-noise ratio. The lower bounds were also compared to empirical estimates obtained, using the statistical models, via Monte-Carlo simulation.

There are, however, some remaining issues that would be helped by more elegant solutions. The most glaring of these is the specular return model, in which we approximated a sinc function as a Gaussian, and artificially and arbitrarily determined the width of the function. This model may actually have behavior conforming to reality, but without a more detailed analysis of reflection from a multi-faceted target there is little by

which to judge its accuracy. On a similar note, there is no Cramér-Rao analysis of a composite reflection surface, which reflects some portion of the radar beam coherently and the rest incoherently. The lack of a tractable approximation to the Rice distribution when the m^2 and λ parameters are of comparable magnitudes makes that particular project difficult to approach. Analysis of the pure specular case would also be aided by such an approximation, so that the Cramér-Rao bound not be restricted to the regime in which the reflected power is an order of magnitude or more greater than the power inherent in the receiver noise.

Additionally, we identified problems with the Cramér-Rao bound itself because it acts on the linearization of the distribution around a specific value of the unknown parameter. Often, as we showed at the extreme orientations $\theta = 0$ and $\theta = \frac{\pi}{2}$, the ability to make an estimation of the parameter is not limited by the curvature of the distribution at that point. A biased estimate, which this version of the CRB cannot consider, could in fact provide a much smaller mean-squared error than the minimum variance unbiased estimator.

Finally, sensor fusion could be attempted, in which the information from a High Range-Resolution Radar profile is optimally combined with the data from a Laser Radar or an Infrared Thermal sensor image. That analysis could identify performance trade-offs and highlight any particular benefits of using multiple sensors to obtain images of a particular target. For example, the inability of HRR to judge the transverse positioning of a reflector could be compensated for by adding another kind of sensor.

References

- [1] L.N. Ridenour, Ed., *Radar Systems Engineering*, MIT Radiation Laboratory Series, 1st Edition, McGraw Hill Book Company, New York, 1947
- [2] V. Kedia, "Performance Analysis of Automatic Target Recognition Using High Resolution Radar," MS thesis, Washington University, August 1998
- [3] J.K. Bounds, "The Infrared Airborne Radar Sensor Suite," The Research Laboratory of Electronics at the Massachusetts Institute of Technology Technical Report No. 610, December 1996
- [4] U. Grenander, A. Srivastava, M. Miller, "Asymptotic Performance Analysis of Bayesian Object Recognition", *Submitted to IEEE Transactions on Information Theory*, February 1998
- [5] U. Grenander, M. Miller, and A. Srivastava, "Hilbert-Schmidt Lower Bounds for Estimators on Matrix Lie Groups," *IEEE Trans. Pattern Analysis and Machine Intell.*, **20**(8), 1-13, August 1998
- [6] J.A. O'Sullivan, S.P. Jacobs, and V. Kedia, "Stochastic models and performance bounds for pose estimation using high resolution radar data," *SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery V*, Vol. 3370, 576-587, April 1998
- [7] A.S. Willsky, G.W. Wornell, J.H. Shapiro, "Stochastic Processes, Estimation, and Detection," Massachusetts Institute of Technology course notes, Spring 1999
- [8] M.I. Skolnik, *Introduction to Radar Systems*, 2nd Edition, McGraw-Hill Book Company, New York, 1980
- [9] J. Kostakis, M. Cooper, T.J. Green, Jr., M.I. Miller, J.A. O'Sullivan, J.H. Shapiro, D.L. Snyder, "Multispectral Sensor Fusion for Ground-Based Target Orientation Estimation: FLIR, LADAR, HRR," *Proc. SPIE*, Vol. 3718, pp14-24, 1999
- [10] J.W. Goodman, *Introduction to Fourier Optics*, McGraw-Hill Book Company, San Francisco, 1968
- [11] J.H. Shapiro, "Extended Version of Van Trees's Receiver Operating Characteristic Approximation," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 35, No. 2, April 1999