An Algorithm for Automatic Target Recognition Using Passive Radar and an EKF for Estimating Aircraft Orientation

A Thesis
Presented to
The Academic Faculty

by

Lisa M. Ehrman

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in Electrical and Computer Engineering

School of Electrical and Computer Engineering
Georgia Institute of Technology
December 2005

An Algorithm for Automatic Target Recognition Using Passive Radar and an EKF for Estimating Aircraft Orientation

Approved by:

Professor Aaron Lanterman, Advisor
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Mark Richards
School of Electrical and Computer Engineering
Georgia Institute of Technology

Professor Steve McLaughlin
School of Electrical and Computer Engineering
Georgia Institute of Technology

Date Approved: October 31, 2005

To my dad, who taught me everything that really matters...

ACKNOWLEDGEMENTS

This work was funded by the NATO Consultation, Command, and Control Agency (NC3A), the U.S. Air Force Office of Scientific Research (grant F49620-03-1-0340), and start-up funds from the School of Electrical and Computer Engineering at the Georgia Institute of Technology. I would like to thank Dr. Paul Howland and Dr. Rene van der Heiden at NC3A for their support. I would also like to thank my advisor, Professor Aaron Lanterman, for his encouragement and assistance. I am also indebted to Major Larkin Hastriter and Lt. Col. Adam MacDonald for their assistance in obtaining aircraft flight paths. Special thanks also to my colleagues at GTRI, who have been exceedingly patient and supportive of my decision to complete my Ph.D. while working full-time. Finally, I would like to thank my husband, Brad, whose indefatiguable work ethic and support have kept me going when all else failed.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

Rather than emitting pulses, passive radar systems rely on "illuminators of opportunity," such as TV and FM radio, to illuminate potential targets. These systems are particularly attractive since they allow receivers to operate without emitting energy, rendering them covert. In addition, the fact that they do not require dedicated transmitters often translates into significant cost savings. Thanks to their potential for low-cost covert operation, passive radar has begun garnering increasing attention from both the media and the scientific community. Until recently, most of the research regarding passive radar has focused on detecting and tracking targets. Given the advances in these areas, the time is right for research that focuses on extending the capabilities of passive radar systems to include automatic target recognition. That is the goal of this dissertation.

The target recognition algorithm described in this dissertation uses the radar cross section (RCS) of potential targets, collected over a short period of time (for example, 60 seconds), as the key information for target recognition. To make the simulated RCS as accurate as possible, the received signal model accounts for aircraft position and orientation, propagation losses, and antenna gain patterns. An extended Kalman filter (EKF) estimates the target's orientation (and uncertainty in the estimate) from velocity measurements obtained from the passive radar tracker. Coupling the aircraft orientation and state with the known antenna locations permits computation of the incident and observed azimuth and elevation angles. The Fast Illinois Solver Code (FISC) simulates the RCS of potential target classes as a function of these angles. Thus, the approximated incident and observed angles allow the appropriate RCS to be extracted from a database of FISC results. Using this process, the RCS of each aircraft in the target class is simulated as though each is executing the same maneuver as the target detected by the system. Two additional scaling processes are required to transform the RCS into a power profile (magnitude only) simulating the signal in the receiver. First, the RCS is scaled by the Advanced Refractive Effects Prediction

System (AREPS) code to account for propagation losses that occur as functions of altitude and range. Then, the Numerical Electromagnetic Code (NEC2) computes the antenna gain pattern, further scaling the RCS. A Rician likelihood model compares the scaled RCS of the illuminated aircraft with those of the potential targets. To improve the robustness of the result, the algorithm jointly optimizes over feasible orientation profiles and target types via dynamic programming.

The primary contribution of this research is the development of a robust ATR procedure that can be used to covertly identify aircraft executing maneuvers ranging from straight-and-level trajectories to dog-fighting trajectories. Simulations indicate that the ATR algorithm performs extremely well at the highest anticipated noise levels when the aircraft being tracked executes predictable straight-and-level or banked turn trajectories. Even when the aircraft under track executes dog-fighting style trajectories, the algorithm's probability of error is still quite good. This is not meant to suggest that the algorithm would be perfect (or nearly so) in practice, as even the best models fail to account for all effects seen in the real world. However, the simulations suggest that the ATR algorithm has a great deal of potential for reliable performance if imported into a real system. Furthermore, the algorithm could be applied (with superficial changes) to low-frequency *active* radar systems.

The second major contribution of this research is the development of an EKF that estimates aircraft orientation from velocity measurements. This EKF plays a role in the ATR algorithm discussed in the previous paragraph, but also has the potential to be used in other applications. Although this dissertation applies the EKF to passive radar velocity measurements generated by state-of-the-art passive radar trackers, there is no reason it cannot be applied to measurements generated by other trackers, as well. The measurements and their error models would simply need to change to accommodate a different tracker and/or source.

The final contribution of this work is the derivation of closed-form approximations for the relative entropy and Chernoff information between two Rician distributions. Although a wide variety of problems are well-modeled by the Rician distribution, the Gaussian distribution is frequently used in its place because of mathematical convenience.

CHAPTER I

BACKGROUND

1.1 Automatic Target Recognition

Automatic target recognition (ATR) systems have been deployed to identify everything from land mines to aircraft, using data collected from infrared (IR), hyperspectral, laser radar (LADAR), and radar sensors, to name a few. Though the target classes and sensors may vary from one application to the next, the goal of this work remains the same; automatic target recognition systems attempt to automatically categorize targets into one of a set of classes.

For example, recent research [50] demonstrated that super-resolution techniques could greatly improve the performance of vehicle recognition schemes utilizing IR data. Another notable body of work [39] focused on extracting invariant features from the IR images to identify vehicles. Hyperspectral imagery is also frequently used to robustly detect and identify small targets, including mines and vehicles [10, 44]. Both synthetic aperture radar (SAR) and LADAR images have also been employed to identify vehicles, such as automobiles and tanks [45, 49, 46, 11, 6]. High-resolution range (HRR) profiles, which may be thought of as one-dimensional radar "images," have been explored for identifying vehicles and aircraft [30, 40].

Some sensors, such as infrared cameras, produce "images" that are readily interpreted by a human operator. Other sensors produce data that must go through some kind of signal processing to transform the data into an "image" form amenable to human interpretation. For instance, in synthetic aperture radar or magnetic resonance imaging, the raw data is nonsensical to human eyes; Fourier transform operations are needed to form images from such data for human analysis. When considering *automatic* target recognition, as performed by a computer, there is a question as to whether such an image-forming pre-processing step is necessary. Two schools of thought are prevalent in the literature. The first body of

literature assumes a traditional image-forming process and develops algorithms to work on those images. This approach has the advantage that the system designers may be better able to employ their visual intuition in creating algorithms. The second school of thought attempts ATR directly on the "unprocessed" data, in the hopes of avoiding a potential loss of information during the image formation stage. It also might sometimes be the case that a *computer* might have an easier time recognizing targets in the unprocessed data than it will in the processed data, although the opposite is true for humans.

This dissertation develops ATR algorithms for "passive radar" systems, which are described in detail in the next section. Some researchers [34, 33] have argued that low-frequency radar (i.e., using very high frequency (VHF) sources) is particularly well-suited for ATR, because the longer wavelengths are less susceptible to wild fluctuations as the target moves. Although there has been some work on exploring the possibility of forming images with passive radar data [12], the limited amount of data available from most passive radar systems, as well as the challenge of the underlying multidimensional autofocus problems, renders image formation from passive radar extremely difficult. For this reason, passive radar sources naturally lend themselves to the direct approach to ATR. Rather than basing recognition on images that are difficult to create, the recognition algorithms in this dissertation use the returned signal power (magnitude only) as the primary feature for identification.

1.2 Passive Radar

One potential ATR source that is the subject of burgeoning interest in both the media and scientific community [41, 26, 25, 52] is passive radar. Rather than emitting energy, passive radar sources rely on "illuminators of opportunity," such as TV and FM radio, to illuminate potential targets. While a traditional radar system contains both a transmitter and receiver, which may be co-located, a passive radar system requires only a receiver that is configured to exploit the desired illuminator(s) of opportunity. This enables passive radar systems to reap a number of benefits. Most notably, the fact passive radar systems do not emit energy renders them covert. This is a great strategic advantage if employed by the

military. Furthermore, because much of the cost of traditional radar systems is associated with the transmitter, passive radar systems are usually significantly less expensive to build and maintain.

Additionally, many illuminators of opportunity employ low frequencies. Although passive radar systems do not inherently require low frequencies, the allocation of bandwidth for commercial TV and FM radio implies that passive radar systems operate at lower frequencies than traditional microwave radar. As a fortunate though unintended consequence, the low-frequency signals exploited by passive radar are well-suited for ATR [34, 33, 13]. In addition to being less susceptible to inclement weather, the longer wavelengths result in target radar cross sections that vary "slowly" with small changes in the target state vector. Herman noted that the variation in radar cross section (RCS), as characterized by the number of nulls encountered as a target's aspect changes, is proportional to the electrical length of the target [28, 27]. At FM-band frequencies (100 MHz), a fighter-sized aircraft is approximately five wavelengths long. In contrast, at the X-band frequencies used by many traditional radars (10 GHz), the same aircraft would be 500 wavelengths long, making the ATR system very sensitive to small changes in target orientation. Although using a frequency in which the target is 500 wavelengths long would provide more variability amongst the target types, it may also require that the orientation be precisely known to achieve good results. Using the low-frequency source strikes a suitable balance; the RCS varies sufficiently from one aircraft class to another to allow for target discrimination, and it is not so sensitive to small changes in orientation that errors in the estimated orientation will break the ATR algorithm.

Despite its numerous benefits, passive radar was once deemed impractical. In the mid-1980s, Griffiths and Long [25] attempted to extract range information from backscattered television signals. Plagued by the low signal-to-noise ratio resulting from the available equipment and the range ambiguity inherent in the sync pulses of an analog TV signal, their results did not seem encouraging. A decade later, interest in passive radar was resurrected when Howland [29] successfully tracked targets by abandoning any attempt to directly measure range in favor of the velocity information contained in the Doppler-shifted TV

carrier and the angle-of-arrival information derived from a simple two-antenna array. Exotic track initialization algorithms, combined with an extended Kalman filter, fuse the Doppler and angle-of-arrival information into Cartesian coordinate tracks. Around that same time, Baltes and van Keuk [7] developed an algorithm for tracking multiple maneuvering targets with a network of passive radar sensors.

Interest in passive radar systems has since been spurred on by the successes of Lockheed Martin Mission Systems' Silent Sentry, which is capable of exploiting both analog TV and FM radio signals, and John Sahr's Manastash Ridge Radar [47, 48], which is used for upper atmospheric radio science. Recent analysis of the ambiguity functions of illuminators of opportunity [42, 23] has only bolstered the case for passive radar's feasibility. Ongoing research continues to demonstrate the potential of passive radar for tracking everything from aircraft to tornados [22].

In fact, now that target detection and tracking via passive radar have received so much attention, the focus of current research is beginning to shift to target imaging and identification. Some notable work regarding target imaging via passive radar has come from researchers at the University of Illinois, including Munson, Wu, Ye, Moulin, and Bresler [38, 51]. Herman and Moulin also demonstrated the potential for ATR using passive radar under the feature-based paradigm [27]. They proposed using a particle filter to simultaneously track and classify targets, in which the target's radar cross section (RCS) is one of the key features.

1.3 Research Contributions

The research presented in this dissertation is in the same vein as Herman's work, but attempts to solve the same problem with a much simpler algorithm. Rather than employing particle filters, which are difficult to implement and require a great deal of computational horsepower, the proposed algorithm is relatively easy to add to existing systems and implement in real time.

As described in prior publications [18, 17, 21, 19, 20], this approach uses the radar cross section (RCS) of potential targets as the key information for target recognition. The passive

radar measurements are collected for a short period of time (for example, 60 seconds), and processed in a batch. To make the simulated RCS as accurate as possible, the received signal model accounts for aircraft position and orientation, propagation losses, and antenna gain patterns. An EKF estimates the target's orientation (and uncertainty in the estimate) from velocity measurements obtained from the passive radar tracker during the time range of interest. Coupling the aircraft orientation and state with the known antenna locations permits computation of the incident and observed azimuth and elevation angles during this same time span. This is depicted in Figure 1.

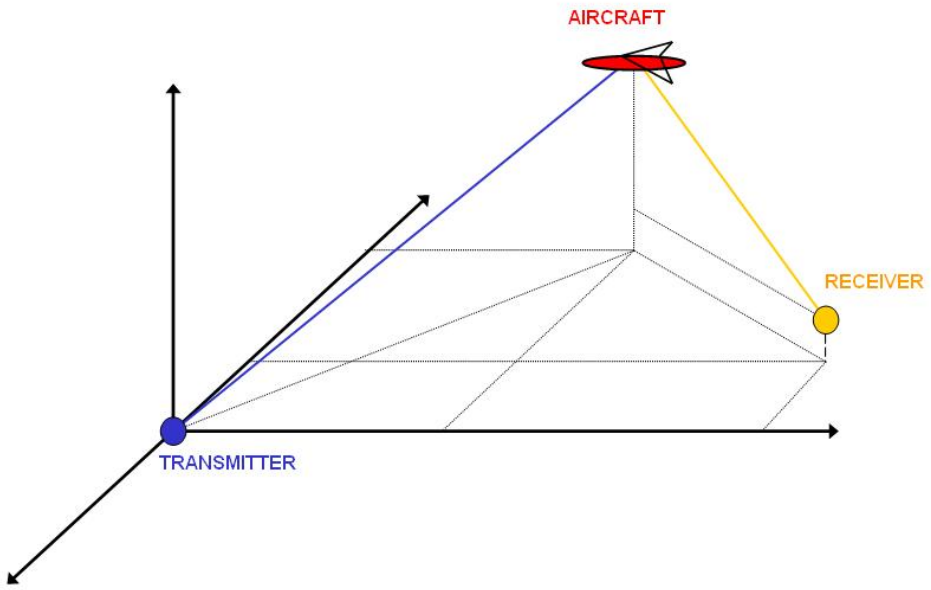

Figure 1: Computing the incident and observed azimuths and elevations on the aircraft

The Fast Illinois Solver Code (FISC) [4] simulates the RCS of potential target classes as a function of these angles. Thus, the approximated incident and observed angles allow the appropriate RCS to be extracted from a database of FISC results. Using this process, the RCS profile of each aircraft in the target class is simulated as though each is executing the same maneuver as the target detected by the system. Two additional scaling processes are required to transform the RCS profile into a power profile (magnitude only) simulating the signal in the receiver. First, the RCS profile is scaled by the Advanced Refractive Effects Prediction System (AREPS) [1] code to account for propagation losses that occur

as functions of altitude and range. Then, the Numerical Electromagnetic Code (NEC2) [2] computes the antenna gain pattern, further scaling the RCS. A Rician likelihood model compares the scaled RCS profile of the illuminated aircraft (spanning the entire time range of interest) with those of the potential targets. To improve the robustness of the result, the algorithm jointly optimizes over feasible orientation profiles and target types. The ATR algorithms also permit exploitation of multiple transmitters, possibly employing different polarizations.

The contributions of this research are numerous. The primary contribution is the development of an ATR procedure for quickly and covertly identifying aircraft that could operate in real time and be easily incorporated into existing passive radar systems. Secondary contributions include the development of an EKF for estimating aircraft orientation from velocity measurements, and closed-form approximations of the relative entropy and Chernoff information between two Rician distributions.

1.4 Organization of Dissertation

The body of the dissertation begins with the derivation of a coordinated flight model and its development into an EKF in Chapter 2. Chapter 3 then describes a simple ATR algorithm that uses the basic coordinated flight model (rather than the orientation-estimating EKF). Results using this simple ATR algorithm are given in Chapter 4. Since this simple algorithm is revealed to lack robustness, further development is required. Chapter 5 estimates the performance of the algorithm from Chapter 3 under both Neyman-Pearson and Bayesian frameworks. It concludes with a demonstration of how these performance estimates can be utilized to estimate the length of time a target must be tracked to identify it with a desired probability of error. Since the ATR algorithm processes a segment of data collected during a time span of interest, this application of the Chernoff information is particularly relevant. A more robust ATR algorithm is developed in Chapter 6. This algorithm jointly estimates the aircraft orientation and type. Engineering decisions in the deployment of this system, such as the length of time to consider when estimating the orientation, are based on results from Chapter 5. The more robust algorithm does indeed perform well, as is demonstrated

in Chapter 6. Its robustness to modeling errors (both in the RCS database and atmospheric attenuation database) is explored in Chapter 7. Conclusions are then provided in Chapter 8.

CHAPTER II

ESTIMATING AIRCRAFT ORIENTATION FROM VELOCITY ESTIMATES

Aircraft RCS is highly aspect dependent. It follows that for an RCS-based recognition scheme to be successful, the incident and observed aspect angles on the aircraft must be well characterized. This poses a challenge. The aspect angles are functions of the target location, relative to the antenna locations, and the target orientation. Although passive radar tracking algorithms, such as the one developed by Howland [29], are capable of providing position and velocity estimates, nothing has yet been published that estimates target orientation. This chapter attempts to fill this void by developing a coordinated flight model to estimate aircraft orientation from velocity estimates. The limitations of the model are discussed in Section 2.2. In response to these limitations, the coordinated flight model from Section 2.1 is developed into an EKF in Section 2.3.

2.1 A Coordinated Flight Model

The conventions used for orientation are depicted in Figure 2. As is fairly standard, yaw, the first element of the orientation vector, is defined to be the rotation about the vector coming out the top of the aircraft, pitch, the second element, is the rotation about the vector extending out the wing of the aircraft, and the third element, roll, is the rotation about the vector in the direction of the aircraft's nose [16].

2.1.1 Estimating Yaw

Using the coordinated flight model, the yaw is modeled as a function of the aircraft motion in the p_x-p_y plane[1], and is expressed by

[1]In this chapter, the cartesian position coordinates are referred to as p_x, p_y, and p_z. This is done to avoid confusion between the p_x axis and state of the extended Kalman filter in Section 2.3, x.

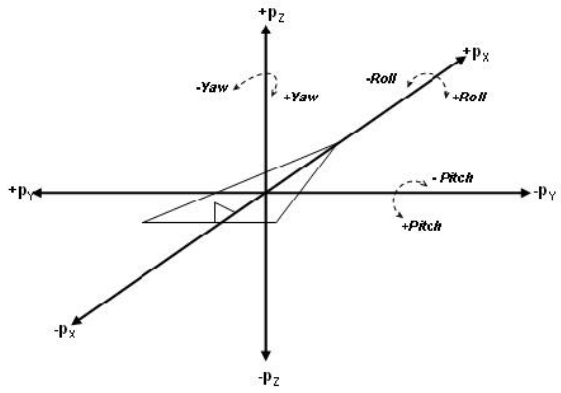

Figure 2: Coordinate convention used to compute aircraft yaw, pitch, and roll

$$\psi = \arctan\left(\frac{v_Y}{v_X}\right), \tag{1}$$

where v_X and v_Y are the p_x and p_y components of the velocity vector. Under this convention, yaw is defined relative to a cartesian coordinate system on the earth's surface. (Because of the earth's curvature, this cartesian system is an approximation.) Thus, an aircraft flying down the x-axis of the coordinate system is said to have a yaw of zero degrees, while one flying down the y-axis is said to have a yaw of 90 degrees.

2.1.2 Estimating Pitch

As should be clear from Figure 2, the aircraft's pitch describes the angle between the total velocity vector and the velocity vector in the p_x-p_y plane, and can be thought of as the angle at which the aircraft is changing altitude. Let (v_X, v_Y, v_Z) denote the aircraft's velocity. The pitch is then approximated under the coordinated flight model as

$$\theta = \arctan\left(\frac{v_Z}{\sqrt{v_X^2 + v_Y^2}}\right). \tag{2}$$

As long as the sampling rate of the aircraft's position is on the order of seconds, this linear approximation of the pitch is sufficiently accurate for these purposes.

2.1.3 Estimating Roll

Estimating the aircraft's roll is more difficult, as rotation about the vector extending through the aircraft's nose in the direction of motion can be performed by high-performance aircraft without any deviation in the flight path or velocity profile. Thus, the uncertainty associated with the estimate of the aircraft's roll is much higher than for the aircraft's yaw and pitch. Even so, the coordinated flight model assumes that the targets execute coordinated maneuvers, and that non-zero roll is only likely when the aircraft is executing a banked turn. Otherwise, the model dictates that the aircraft's roll is approximated as zero.

Since the coordinated flight model only estimates non-zero roll angles during banked turns, it is natural to begin the derivation of the roll angle by examining the free body diagram of an aircraft undergoing a banked turn with radius R. The three main forces acting on the aircraft in this plane are depicted in the free body diagram shown in Figure 3. For the sake of this example, the aircraft is flying directly out of the page along the +x-axis. The lift, L, acts orthogonally to the top of the aircraft, the weight, W, pulls directly towards the ground, and the centrifugal force, F_C, pushes out along the vector originating at the center of the turn.

3.

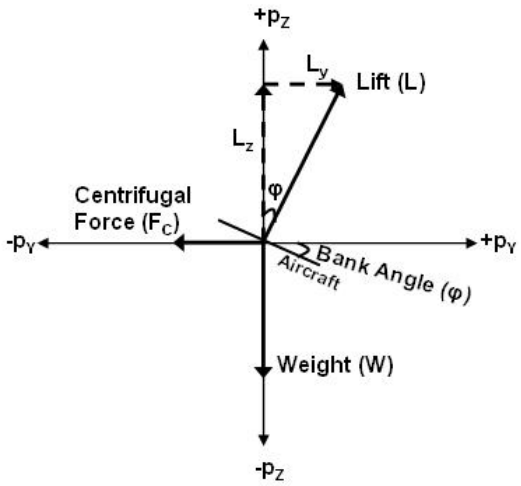

Figure 3: Free body diagram of an aircraft executing a constant-altitude banked turn

The coordinated flight model, by definition, assumes that the aircraft isn't slipping at

all in the p_x-p_y plane while turning. Thus, it follows that

$$\sum F_y = 0. \tag{3}$$

Since the only forces acting in the p_y-dimension are the lift and centrifugal force, it follows that

$$L_Y = F_C. \tag{4}$$

The centrifugal force is derived from Newton's Second Law, in which $\frac{W}{g}$ is substituted for mass and $\frac{|v|^2}{R}$ is substituted for acceleration. Here W is the weight and g is the gravitational constant. The y-component of lift is related to the total lift through a simple trigonometric relation. Substitution of these expressions into (4) results in

$$L\sin(\phi) = \frac{W|v|^2}{Rg}. \tag{5}$$

Solving (5) for the roll angle yields

$$\phi = \arcsin\left(\frac{W|v|^2}{LRg}\right). \tag{6}$$

Since the aircraft is also assumed to be executing a constant-altitude maneuver in this example, it follows that

$$\sum F_z = 0. \tag{7}$$

Since the lift and weight are the only two forces acting in the p_z-dimension,

$$L_Z = W. \tag{8}$$

L_Z can be written as a function of lift and roll angle. This relation is substituted into (8), resulting in

$$L\cos(\phi) = W, \tag{9}$$

which can be solved for the roll angle, giving

$$\phi = \arccos\left(\frac{W}{L}\right). \tag{10}$$

Both weight and lift vary with aircraft type, which is unknown [31]. For this reason, it is advantageous to derive an equation for the roll angle that is independent of weight and lift. Fortunately, this can be accomplished by combining (6) and (10) to give

$$\phi = \arctan\left(\frac{|v|^2}{Rg}\right). \tag{11}$$

Aircraft velocity is easily computed from the estimated flight path, and the force of gravity is a known constant. Thus, the only remaining unknown in (11) is the radius of the turn. Computing the radius of turn would be relatively simple if aircraft always flew in perfect circles and began the turns instantaneously. However, this is not the case. A more realistic model dictates gradual transition into and out of banked turns. This gradual transition can be thought of as moving from a turn with a radius of curvature approaching infinity to a turn with a particular finite radius, or similarly, moving from a turn with a finite radius to one with a radius approaching infinity. Thus, the transition made by an aircraft into or out of a banked turn is essentially just a curve with a uniformly changing degree of curvature. This type of curve is known as a spiral curve.

The radius of curvature at any point along a spiral curve can be found using the method of osculating circles. In this method, a "circular arc drawn tangent to the spiral at point P has a radius r equal to the radius of curvature of the spiral at the point of tangency" [37]. The circular arc is known as an osculating circle, whose radius of curvature is given by

$$R(t) = \frac{[\dot{x}(t)^2 + \dot{y}(t)^2]^{3/2}}{\dot{x}(t)\ddot{y}(t) - \ddot{x}(t)\dot{y}(t)}, \tag{12}$$

where $x(t)$ and $y(t)$ are the p_x and p_y positions as functions of time. Using this method, the radius of curvature is found for every position in the flight path. The resulting radius function is then substituted into (11) to compute roll as a function of time.

The sign of the roll angle must also be considered. Because the expression for roll given by (11) is a function of the magnitude of the velocity, the radius of acceleration, and the

pull of gravity, it always results in a positive roll angle. From Figure 2, it is clear that positive roll angles are associated with rolls in which the aircraft's right wing is down. If the aircraft is flying in a clockwise manner, then the roll angle should be positive, given the current convention. Conversely, counter-clockwise flight should result in negative roll angles.

A minor change is required to extend the expression in (11) to the general case in which the aircraft's pitch may be non-zero. In particular, the assumption that the sum of forces in the z-direction is equal to zero no longer applies.

Equations 4 through 6 provide a method for computing roll angles that is not dependent upon maintaining a constant altitude. However, the equations are dependent upon aircraft weight and lift, parameters which vary from one aircraft to another. Since aircraft type is not known, a method for computing the roll angle that is not dependent upon aircraft weight and lift is still desirable.

One solution to this dilemma is to select a new coordinate system in which the aircraft is not changing altitude. This is accomplished by multiplying each state by a coordinate rotation matrix that puts all of the aircraft motion at that instant along the new p_x-axis. Equation 11 then applies, with two minor changes. First, the radius of curvature must now reflect the curvature in the new p_x-p_y plane, denoted by R_{XY}. Then, the norm of the velocity must be scaled by the cosine of the pitch to produce the norm of the velocity in the new p_x-p_y plane. Thus, the general equation for the aircraft roll is

$$\phi = \arctan\left(\frac{|v|^2\cos(\theta)}{R_{XY}g}\right). \tag{13}$$

2.2 Limitations of the Coordinated Flight Model

The coordinated flight model developed in Section 2.1 has two major limitations. First, the coordinated flight model is bound by the assumption that the aircraft nose points in the direction of motion. In reality, aircraft engaged in ascent or descent maneuvers are not likely to have their noses pointed directly along the velocity vector. Furthermore, aircraft are also affected by "crabbing," a condition commonly occurring in conjunction with high wind or

severe weather in which there is a slight difference between the direction of the aircraft's nose and the direction of the aircraft's motion. In this case, the aircraft yaw predicted by the coordinated flight model is also biased. The more serious limitation of the coordinated flight model is that it only provides a state estimate, and makes no attempt to model the uncertainty associated with each element of the state.

2.3 An EKF for Estimating Aircraft Orientation from Velocity Estimates

This section combats the limitations of the coordinated flight model by developing it into an EKF for estimating aircraft orientation from velocity estimates. Consider the two distinct types of error present in this problem. Measurement error, as the name implies, is the error on the velocity measurements received from the passive radar tracker. Modeling error is also present, and can be thought of as the error that would occur even if the measurements were perfectly known. The coordinated flight model treats the measurements as though they are perfectly known and provides no estimate of the uncertainty resulting from modeling error. The EKF, in contrast, blends both types of error to create its estimates of the target's state and covariance. Furthermore, the EKF accounts for "crabbing" and other modeling errors with the process noise. In doing so, the EKF sidesteps both limitations of the coordinated flight model.

2.3.1 The EKF Model

This development of an EKF that estimates the aircraft orientation from velocity estimates begins with the assumption that the best state model is a straight-and-level constant-velocity trajectory, described by

$$x_k = x_{k-1} + q_k, \tag{14}$$

where q_k is the process noise term, modeled as a zero-mean Gaussian random variable with covariance Q_k. The state at time t_k, denoted as x_k, is a vector consisting of the target's yaw, pitch and roll.

The measurement model is then described by

$$z_k = h_k(x_k) + s_k, \tag{15}$$

where z_k is the measurement vector at time k consisting of the x, y, and z components of the target's velocity, $h_k(x_k)$ is a nonlinear mapping of the state into the measurement space, and s_k is the measurement noise term, modeled as a zero-mean Gaussian random variable with covariance S_k.

The standard EKF equations are given next, to clarify notation. The state, $x_{k-1|k-1}$ and its covariance, $P_{k-1|k-1}$, are extrapolated from time t_{k-1} to time t_k with

$$x_{k|k-1} = F_k x_{k-1|k-1} \tag{16}$$

and

$$P_{k|k-1} = F_k P_{k-1|k-1} F_k^T + Q_k. \tag{17}$$

Given the state model, it follows that $F_k = I$. Thus, the time-update equations for this EKF reduce to

$$x_{k|k-1} = x_{k-1|k-1} \tag{18}$$

and

$$P_{k|k-1} = P_{k-1|k-1} + Q_k. \tag{19}$$

As is standard practice [5], the process noise, Q_k, is selected to tune the filter to the application.

The measurement update is performed using

$$K_k = P_{k|k-1} H_k \left(H_k P_{k|k-1} H_k^T + S_k \right)^{-1}, \tag{20}$$

$$x_{k|k} = x_{k|k-1} + K_k \left(z_k - h_k(x_{k|k-1}) \right), \tag{21}$$

and

$$P_{k|k} = (I - K_k H_k)\, P_{k|k-1}, \tag{22}$$

where S_k is the measurement covariance and H_k is the linearization of the nonlinear mapping $h_k(x_{k|k-1})$.[2] This completes one update cycle in the EKF.

Although EKFs are quite common, initializing them is often something of an art. One advantage of this EKF is that the initialization procedure flows very naturally from its development. No exotic algorithms are required. Rather, the state vector of this EKF is simply initialized using the coordinated flight model.

2.3.2 Solving for h and H

Solving for $h_k(x_{k|k-1})$ and H_k is far from trivial. First, it is necessary to solve for the x, y, and z components of the target's velocity, which comprise the measurement, in terms of the yaw, pitch, and roll, which comprise the state. To make this problem mathematically tractable, the radius of curvature is computed from the measurements off-line via (12), and is treated as noiseless in this derivation. Note that it is not necessary to wait until the entire trajectory has been collected to compute R_{XY}. Rather, it is only necessary to wait for two measurements, thus allowing for updated velocity and acceleration measurements, to compute R_{XY}. For example, $R_{XY}(t)$ can be computed using measurements through time $t + 2\Delta_t$, where Δ_t is the time between measurements. Note that hereafter we will drop the subscript XY from R_{XY} for notational compactness.

Solving (1) for v_Y results in

$$v_Y = v_X \tan(\psi). \tag{23}$$

Similarly, solving (2) for v_Z results in

[2] As is often the practice when implementing EKFs, the covariance is inflated following the measurement update step to account for imperfections resulting from the linear approximations underlying the EKF derivation. Determining whether this is necessary and how much inflation is required are questions typically addressed through analysis of multiple sets of training data. Our training data is admittedly limited. One path for future work is to re-tune the filter using a more extensive set of training data.

$$v_Z = \sqrt{v_X^2 + v_Y^2}\tan(\theta). \tag{24}$$

From (23), it follows that

$$v_X^2 + v_Y^2 = v_X^2 + v_X^2\tan^2(\psi), \tag{25}$$

which reduces to

$$v_X^2 + v_Y^2 = v_X^2\sec^2(\psi), \tag{26}$$

Using this result reduces (24) to

$$v_Z = v_X\sec(\psi)\tan(\theta). \tag{27}$$

Equation 13 is then rewritten for $|v|^2$, resulting in

$$|v|^2 = Rg\sec(\theta)\tan(|\phi|). \tag{28}$$

From (26) and (27), $|v|^2$ is also given by

$$|v|^2 = v_X^2\sec^2(\psi) + \left(v_X\sec(\psi)\tan(\theta)\right)^2, \tag{29}$$

which reduces to

$$|v|^2 = v_X^2\sec^2(\psi)\sec^2(\theta). \tag{30}$$

Substituting (30) into (28) and solving for v_X results in

$$v_X = \cos(\psi)\sqrt{Rg\cos(\theta)\tan(|\phi|)}. \tag{31}$$

This is then substituted back into (23) and (24) to solve for v_Y and v_Z in terms of ψ, θ, ϕ, and R. Thus, the set of equations that map the state vector into the measurement space is summarized by

$$h(x) = \begin{bmatrix} v_X \\ v_Y \\ v_Z \end{bmatrix} = \begin{bmatrix} \cos(\psi)\sqrt{Rg\cos(\theta)\tan(|\phi|)} \\ \sin(\psi)\sqrt{Rg\cos(\theta)\tan(|\phi|)} \\ \tan(\theta)\sqrt{Rg\cos(\theta)\tan(|\phi|)} \end{bmatrix}. \tag{32}$$

The linearization of this nonlinear mapping, used in (20) and (22), is described by

$$H = \begin{bmatrix} \frac{\partial v_X}{\partial \psi} & \frac{\partial v_X}{\partial \theta} & \frac{\partial v_X}{\partial \phi} \\ \frac{\partial v_Y}{\partial \psi} & \frac{\partial v_Y}{\partial \theta} & \frac{\partial v_Y}{\partial \phi} \\ \frac{\partial v_Z}{\partial \psi} & \frac{\partial v_Z}{\partial \theta} & \frac{\partial v_Z}{\partial \phi} \end{bmatrix}, \tag{33}$$

where the partial derivatives are given by

$$\frac{\partial v_X}{\partial \psi} = -\sin(\psi)\sqrt{Rg\cos(\theta)\tan(|\phi|)}, \tag{34}$$

$$\frac{\partial v_Y}{\partial \psi} = \cos(\psi)\sqrt{Rg\cos(\theta)\tan(|\phi|)}, \tag{35}$$

$$\frac{\partial v_Z}{\partial \psi} = 0, \tag{36}$$

$$\frac{\partial v_X}{\partial \theta} = \frac{-1}{2}\cos(\psi)\sin(\theta)\sqrt{Rg\sec(\theta)\tan(|\phi|)}, \tag{37}$$

$$\frac{\partial v_Y}{\partial \theta} = \frac{-1}{2}\sin(\psi)\sin(\theta)\sqrt{Rg\sec(\theta)\tan(|\phi|)}, \tag{38}$$

$$\frac{\partial v_Z}{\partial \theta} = \frac{-1}{2}\tan(\theta)\sin(\theta)\sqrt{Rg\sec(\theta)\tan(|\phi|)} + \sec^2(\theta)\sqrt{Rg\cos(\theta)\tan(|\phi|)}, \tag{39}$$

$$\frac{\partial v_X}{\partial \phi} = \frac{\cos(\psi)}{2\cos^2(|\phi|)}\sqrt{\frac{Rg\cos(\theta)}{\tan(|\phi|)}}, \tag{40}$$

$$\frac{\partial v_Y}{\partial \phi} = \frac{\sin(\psi)}{2\cos^2(|\phi|)}\sqrt{\frac{Rg\cos(\theta)}{\tan(|\phi|)}}, \tag{41}$$

and

$$\frac{\partial v_Z}{\partial \phi} = \frac{\tan(\theta)}{2\cos^2(|\phi|)}\sqrt{\frac{Rg\cos(\theta)}{\tan(|\phi|)}}. \tag{42}$$

2.3.3 Numerical Issues

Although the H matrix computed in Section 2.3.2 applies, in theory, to any target trajectory, it presents a numerical problem when the target follows a straight, constant-velocity trajectory in the p_x-p_y plane. In such a case, the radius of curvature goes to infinity while the roll angle goes to zero. Thus, the product of R and $\tan(|\phi|)$ found in eight of the nine elements of H is numerically unstable.

One solution is to switch to an entirely new state vector when this condition occurs. Consider the equations mapping the state vector into the measurement space, summarized in (32). The $\sqrt{Rg\cos(\theta)\tan(|\phi|)}$ term found in each of the mappings is equivalent, through (13), to $|v_{XY}|$. Since the velocities are treated as "measurements" by the EKF that estimates orientation, $|v_{XY}|$ is ideally suited to join the yaw and pitch in comprising the new state vector. Note that v_{XY} is in no way approximating the roll. A new set of mapping equations are required to transform this new state vector into the measurement space.

This solution poses two challenges. First, criteria for selecting the appropriate state model must be selected. Once this has been done, it is necessary to chose a method for initializing the new state after the state model has changed. Neither task is insurmountable. For example, Howland [29] tackles a similar problem through the creation of a maneuver threshold. Maneuvers, or flight that is not straight-and-level, are only assumed to occur when a chosen maneuver parameter exceeds this maneuver threshold. When this occurs, the maneuver state model described in Section 2.3.2 is used. If the parameter is below the threshold, constant-velocity flight is assumed and the alternate state model, incorporating $|v_{XY}|$ instead of ϕ, is used. Furthermore, the relation between $|v_{XY}|$ and ϕ leads to natural initialization procedures when a model switch occurs. When switching from the non-maneuver model to the maneuver model, ϕ is assumed to be transitioning away from zero, so zero is a natural initial value. When switching from the maneuver model to the non-maneuver model, the $|v_{XY}|$ state is initialized using the most recent measurement.

2.3.4 Solving for h and H under the Non-Maneuvering Model

If the non-maneuvering model is selected, the nonlinear mapping of the state vector into the measurement space must change to reflect the new state vector. The new nonlinear mapping is given by

$$h(x) = \begin{bmatrix} v_X \\ v_Y \\ v_Z \end{bmatrix} = \begin{bmatrix} |v_{XY}| \cos(\psi) \\ |v_{XY}| \sin(\psi) \\ |v_{XY}| \tan(\theta) \end{bmatrix}, \tag{43}$$

and its linearization matrix, H, is given by

$$H = \begin{bmatrix} \frac{\partial v_X}{\partial \psi} & \frac{\partial v_X}{\partial \theta} & \frac{\partial v_X}{\partial |v_{XY}|} \\ \frac{\partial v_Y}{\partial \psi} & \frac{\partial v_Y}{\partial \theta} & \frac{\partial v_Y}{\partial |v_{XY}|} \\ \frac{\partial v_Z}{\partial \psi} & \frac{\partial v_Z}{\partial \theta} & \frac{\partial v_Z}{\partial |v_{XY}|} \end{bmatrix}, \tag{44}$$

or,

$$H = \begin{bmatrix} -\sin(\psi) & 0 & \cos(\psi) \\ \cos(\psi) & 0 & \sin(\psi) \\ 0 & \sec^2(\theta) & \tan(\theta) \end{bmatrix}. \tag{45}$$

2.3.5 Results

Four trajectories are used to demonstrate the EKF's performance. The first is a straight-and-level maneuver in which the aircraft flies directly away from the sensor with a yaw of 150^o. This flight path is then rotated by ninety degrees to create the second straight-and-level maneuver. In this case, the aircraft flies broadside to the sensor with a yaw of 240^o. The third trajectory consists of a constant-altitude circular banked turn, and the fourth was recorded on-board a maneuvering F-15 at Edwards Air Force Base.[3] (A 3-D view of the Edwards trajectory appears in Figure 4 with times labeled.) In all cases, Howland's EKF that estimates positions and velocities feeds "measurements" to the EKF that estimates

[3]The F-15C trajectory was obtained from the Joint Helmet Cuing System, Mission JH-16, conducted by the 445th Flight Test Squadron at Edwards Air Force Base in May 2000. Thanks to Major Larkin Hastriter and Lt. Col. Adam MacDonald for their assistance in obtaining this aircraft flight path.

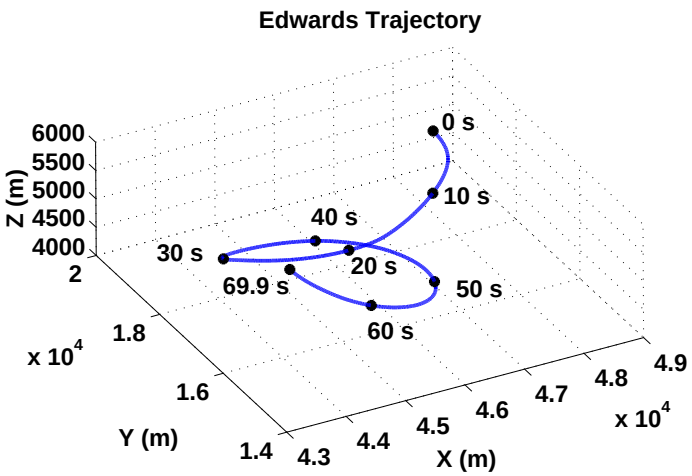

Figure 4: 3-D view of Edwards trajectory

orientation. Using Howland's EKF to generate measurements is appropriate for the passive radar application described here, but could be replaced by some other means of generating velocity measurements to apply the technique to a different application.

Figure 5a shows the yaw of the aircraft executing the straight-and-level maneuver, as well as $1-\sigma$ error bars obtained from the covariance matrix. Similar results involving aircraft pitch are shown in Figure 5b. Since the roll is assumed to be zero and is not estimated by the model when no acceleration is detected, it is neglected here. Note that the uncertainties in yaw and pitch are fairly small and stem almost entirely from the measurement noise. This is also the case in the second straight-and-level maneuver. The estimated pitch and roll for this trajectory are shown in Figure 6.

The uncertainties associated with the orientation estimates become much larger when the aircraft accelerates due to limitations in the coordinated flight model reflected in the increased process noise. This is the case in the last two trajectories. The estimated yaw, pitch, and roll, along with the error bounds, are shown in Figures 7 through 9 for the banked turn trajectory. Since the acceleration is even larger in the Edwards trajectory, the error bars also grow. This is reflected in Figures 10 through 12. Although the estimated orientations are not always near the true values, the true values are usually within a standard deviation of the estimates.

This EKF will be used in Chapter 6 to develop a more robust recognition algorithm. In particular, the algorithm will jointly estimate aircraft type and orientation, from within a feasible set. The feasible set is determined by adding $1-\sigma$ error bars from the EKF's estimated covariance to the estimated states.

Before moving on, two issues should be noted. The practices of tuning the process noise and inflating the measurement covariance frequently arise when designing EKFs, yet they are rarely addressed in EKF literature. As is typically the case, this EKF has been tuned for a particular application using the training set of aircraft data at our disposal. One avenue for future work is to obtain more extensive databases of aircraft data and re-tune the EKF.

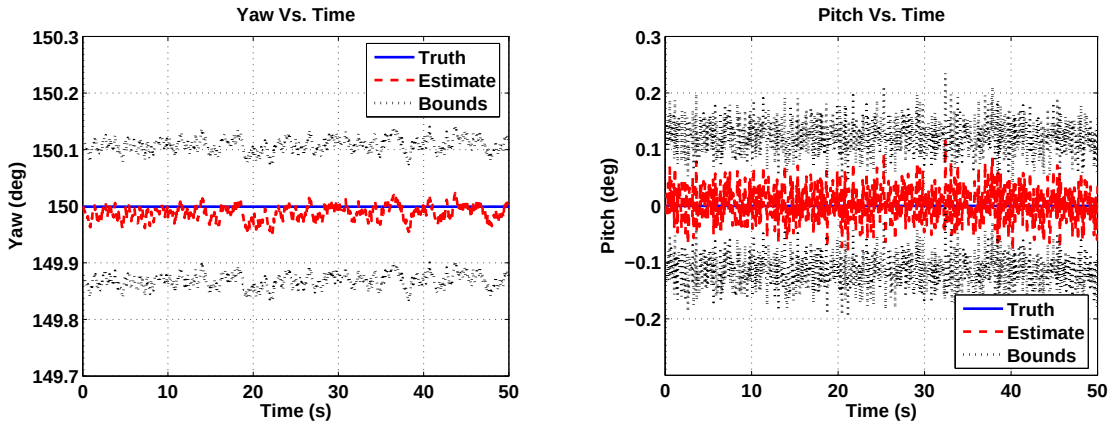

Figure 5: Orientation predicted by the EKF for straight-and-level trajectory #1: a) Yaw (left), b) Pitch (right)

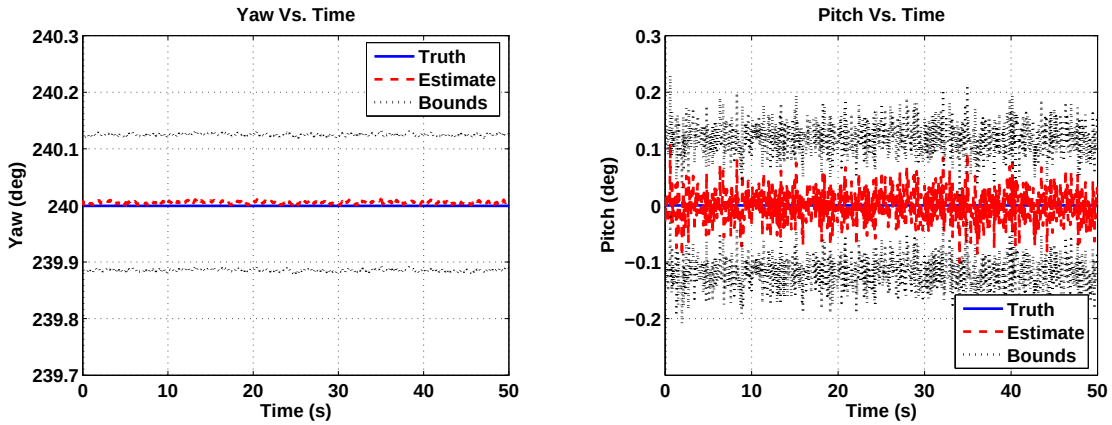

Figure 6: Orientation predicted by the EKF for straight-and-level trajectory #2: a) Yaw (left), b) Pitch (right)

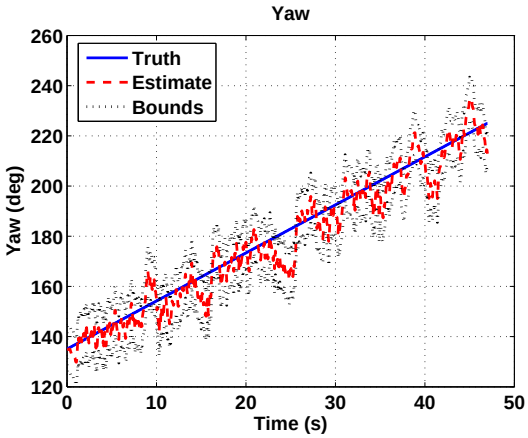

Figure 7: Yaw predicted by the EKF: banked turn trajectory

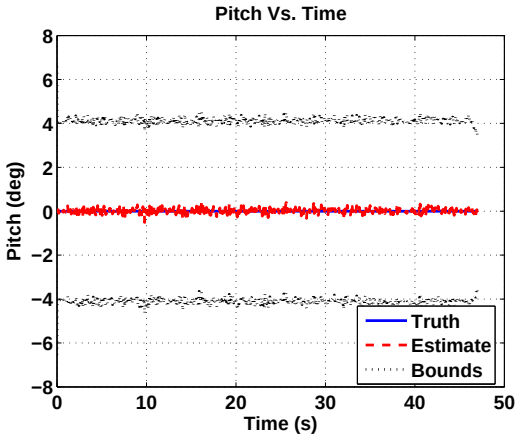

Figure 8: Pitch predicted by the EKF: banked turn trajectory

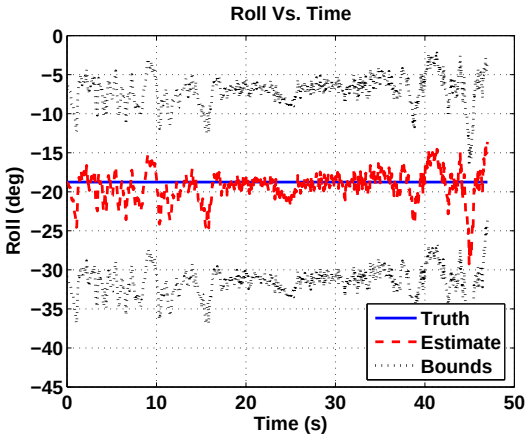

Figure 9: Roll predicted by the EKF: banked turn trajectory

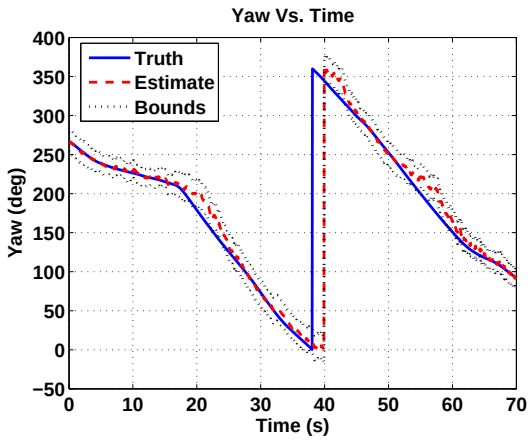

Figure 10: Yaw predicted by the EKF: Edwards trajectory

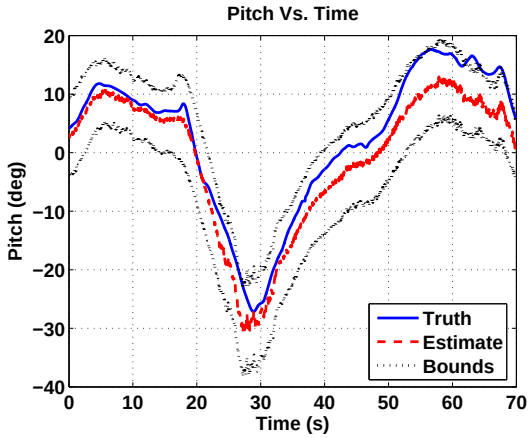

Figure 11: Pitch predicted by the EKF: Edwards trajectory

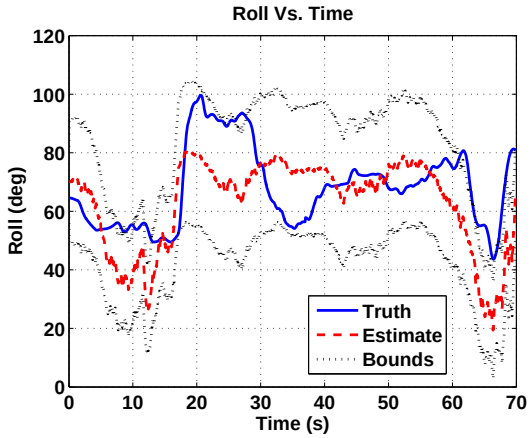

Figure 12: Roll predicted by the EKF: Edwards trajectory

CHAPTER III

A SIMPLE PASSIVE RADAR ATR ALGORITHM BASED ON A COORDINATED FLIGHT MODEL

Sections 2.1 through 2.3 describe two methods for estimating the aircraft orientation from velocity measurements. Once the aircraft orientation has been estimated, it is appended to the existing trajectory, creating a supplemented flight profile that covers the aircraft's maneuver during the entire time span of interest. Together with the system geometry, this supplemented flight profile provides all the information necessary to extract the appropriate RCS for each aircraft in the target library from the database.

3.1 Modeling RCS from a Low-Frequency, Passive Source

Passive radar is garnering increasing attention from the scientific community, but even so, very few systems are currently operational. As such, there is no "typical" set of specifications for passive radar systems. To ensure that this research is not based on unrealistic specifications, a passive radar demonstration system currently being developed by NATO is modeled. The relevant parameters for the potential transmitters and receiver are shown in Table 1. This in no way implies that the work only applies to the NATO system. It merely reflects the difficulty of obtaining detailed specifications for more well-known systems, such as Lockheed Martin's Silent Sentry.

Modeling the power profile arriving at the receiving antenna is a multi-step process. The supplemented flight profile, containing aircraft position and orientation for the entire time range of interest, is used to determine the incident and observed angles. The incident and observed angles are then used to access a database of FISC results,[1] which are available

[1] In the interest of limiting the number of necessary FISC runs, both transmitters are modeled as operating at 100 MHz. The error induced through this assumption is much smaller than the difference between the power profiles (magnitude only) that are used to identify aircraft, so it is deemed negligible. More anaylsis on this topic is provided in Section 7.2.

Table 1: Transmitter and receiver parameters.

Parameter	Vertically Polarized Transmitter	Horizontally Polarized Transmitter	Receiver
Latitude (N)	52^o 01' 00"	51^o 31' 00"	52^o 06' 36"
Longitude (E)	05^o 03' 00"	03^o 53' 00"	04^o 19' 26"
Altitude (m ASL)	375	103	100
Frequency (MHz)	104	96	–
Peak Power (kW)	100	100	–
Direction	Omni-Directional	Omni-Directional	320^o
Polarization	Vertical	Horizontal	Vertical

for each aircraft model in the target library. Note that the type and number of targets in the target library has been limited by the availability (and price) of FISC-compatible target CAD models. The models used in this study were obtained from Digimation [3], a subsidiary of Viewpoint, and, when necessary, were made compatible with FISC using GeomFix by MATIS, Inc. A set of signal profiles (magnitude only) are created when data is extracted from the RCS database. Additional scaling is required to make these signal profiles represent the power signals arriving at the receiver due to the illuminated targets. Some significant factors that must be considered are propagation losses between the aircraft and antennas, and antenna gain. The propagation losses, which include effects due to multipath, are modeled using AREPS. As is likely to be the case for most passive radar applications, the transmitting antenna exploited by the NATO system is modeled as being omni-directional in azimuth. Thus, the only antenna gain pattern modeled here corresponds to the receiver and is accomplished with NEC2. To cut down on the length of time required to execute the simulation, databases are also created for AREPS and NEC2. The overall result of this process is a power profile that is scaled to account for propagation losses and antenna gain.

The necessity of the FISC database, as well as its limitations, are worth mentioning. Ideally, the simulation process described here would run FISC for every new set of incident and observed angles. However, the lengthy run-time and massive memory requirements of FISC render this option unfeasible. A more attractive option is the creation of a database of FISC results, in which aircraft RCS is sampled sufficiently to meet the Nyquist sampling criterion. In particular, the angular sampling of the RCS should satisfy,

$$\Delta\theta \leq \frac{c}{(2f_0)(size)}, \tag{46}$$

where c is the speed of light, f_0 is the radar frequency, and $size$ is the longest dimension of the aircraft [36]. The minimum angular sampling required for each aircraft when the frequency is 100 MHz is shown in Table 2.

Table 2: Minimum angular sampling required for each aircraft in target class.

Aircraft	Longest Dimension (m)	Minimum Spacing (deg)
F-15	19.3	4.5
Falcon-20	17.1	5.0
Falcon-100	13.7	6.2
T-38	14.0	6.2

Thus, using an angular spacing of 4^o, a database of FISC runs can be created that is sufficiently full to represent the RCS of each aircraft type in the study, without aliasing. This database can be quickly accessed and allows for the creation of a power profile for virtually any desired flight path. Sampling every 4^o in both azimuth and elevation around the aircraft is admittedly overkill, since the points become very closely spaced near $\pm 90^o$. However, this route was chosen because it was simple to model and implement. Future work could explore more intelligent methods for sampling the azimuths and elevations, thereby reducing the size of the RCS database.

3.2 Modeling the Power Arriving at the Receiver

Two changes differentiate this process from that used to find the precomputed RCS of known targets in the target class. First, when simulating the RCS of detected targets, the simulation is run with the real aircraft orientations in place of the estimated ones. These power profiles are then corrupted with additive white Gaussian noise, which acts independently on the real and imaginary parts of the signal. This noise model is quite common in modeling passive radar systems [27]. Along these lines, the simulated received profile is expressed as

$$P_{SIM} = (\sqrt{P_R} + w_R)^2 + w_I^2, \tag{47}$$

where P_R is the real component of the power profile prior to being corrupted by noise, and w is zero-mean additive white Gaussian noise, which has real and imaginary components, w_R and w_I [27]. The noise power is computed (in Watts) using,

$$P_N = \frac{kT_0N_F}{CPI}, \tag{48}$$

where k is Boltzmann's constant, T_0 is temperature in Kelvin, N_F is the unitless noise figure, and CPI is the coherent processing interval of the system [8]. To match the NATO system, the CPI is set equal to 0.5 seconds, and T_0 is set equal to 290 K.

Selection of the noise figure, N_F, is more difficult. If the noise figure is only expected to account for thermal noise and out-of-band interference, then a conservative estimate of the noise figure in a city environment might be 30 dB [24]. Because this number includes out-of-band interference, which is typically quite high in an urban enviroment due to the abundance of other emitters, this noise figure is significantly larger than the 2-4 dB normally witnessed with active radar. In the absence of data regarding the noise power of the NATO system, it is reasonable to use 30 dB as an approximation. The noise figures used in this study are given in Table 3, along with the corresponding noise power. Note that some of the larger noise figures shown the Table 3 are not expected to occur in a real setting; they are merely included to demonstrate the breaking point of the algorithm.

3.3 Noise Power Due to Direct Path Interference

The noise figure accounts for thermal noise and out-of-band interference, but until now the issue of transmitter interference has not been addressed. Typically, this direct path interference manifests itself as a spike in the cross-ambiguity function. Since the transmitter's power and location are known, and since the direct path interference spike occurs along the axis with zero velocity, this spike can usually be identified and removed. The more treacherous effect of transmitter interference is that is can raise the "thumbtack" noise floor of

Table 3: Noise figure and noise power.

Noise Figure (dB)	Noise Power (dBW)
30	-171
35	-166
40	-161
45	-156
50	-151
55	-146
60	-141
65	-136
70	-131
75	-126
80	-121
85	-116
90	-111
95	-106
100	-101

the ambiguity function, potentially masking the target spike. To be thorough, this should also be considered when computing the noise figure.

If the ambiguity function is normalized such that the direct path spike has unit height, then the average pedestal height, or sideband power, is given by

$$P_{pedestal} = \frac{1}{B \times CPI}, \tag{49}$$

where B is the signal bandwidth, and CPI is the coherent processing interval [43]. To match the NATO system, values of 45 kHz and 0.5 seconds are used for B and CPI, respectively. If propagation losses and antenna gain are neglected, the pedestal power is 44 dBW below the direct-path spike. Since the NATO transmitter power is 50 dBW, the sideband power is 6 dBW. Propagation losses and antenna gain play a significant role, lowering the pedestal power by 95 dBW. The electronics in the receiver of the NATO system also mitigate the problem by suppressing the direct path signal by 70 dBW, which reduces the sideband power to -159 dBW. More sophisticated filters could be implemented to further reduce the noise figure, but using the specifications of the system being modeled, the effective noise figure falls between 40 and 45 dB. Thus, the effects due to transmitter interference are far

more significant than those due to thermal noise and out-of-band interference.

3.4 Target Recognition

The simulation process is repeated for each member of the target class, resulting in a set of simulated power profiles corresponding to the detected targets. The automatic target recognizer compares these noisy received profiles to the library of precomputed power profiles. Equation 47 leads to a Rician likelihood model [27], whose probability density function is given by

$$p_x(x) = \frac{x}{\sigma_w^2} e^{-\left(\frac{x^2+s^2}{2\sigma_w^2}\right)} I_0\left[\frac{xs}{\sigma_w^2}\right], \tag{50}$$

where $I_0()$ is a modified Bessel function of the first kind.

To apply the Rician density to the simulated and truth profiles, associate x with the square root of the measured power of the detected target, and s with the square root of the precomputed power. The noise power, which equals the noise variance, is then equated with σ_w^2. Since each measurement is assumed to be an independent sample from a process, the data loglikelihood is

$$\ln(p_x(\bar{x})) = \sum_{i=1}^{n} \ln\left(\frac{x_i}{\sigma_w^2}\right) + \ln\left(I_0\left[\frac{x_i s_i}{\sigma_w^2}\right]\right) - \left(\frac{x_i^2 + s_i^2}{2\sigma_w^2}\right), \tag{51}$$

where n is the number of measurements collected during the time span of interest. (Section 5.2.5 provides a method for gauging how long the time span of interest should be to ensure that the probability of error is below a desired level.)

Since the first term in (51) does not vary with the precomputed power profiles, it can be neglected when computing the likelihood scores for each aircraft in the target class. Thus, each aircraft in the target class is scored with

$$L(\bar{x}) = \sum_{i=1}^{n} \ln\left(I_0\left[\frac{x_i s_i}{\sigma_w^2}\right]\right) - \left(\frac{x_i^2 + s_i^2}{2\sigma_w^2}\right). \tag{52}$$

A maximum-likelihood scheme determines the winner [17].

CHAPTER IV

RESULTS USING THE SIMPLE ATR ALGORITHM BASED ON A COORDINATED FLIGHT MODEL

The method described in Chapter 3 is applied to scenarios of increasing complexity. The locations of the aircraft maneuvers relative to the transmitters and receiver are shown in Figure 13. Note that the transmitters shown in Figure 13 correspond to the real transmitters modeled in Section 3.1, and will heretoafter be referred to as the horizontally polarized transmitter and vertically polarized transmitter. The first scenario involves a simple straight-and-level trajectory in which the aircraft fly directly away from the receiver. This trajectory is then rotated by ninety degrees so that the aircraft fly broadside to the receiver. The third trajectory is a constant-altitude circular banked turn. Although aircraft are not likely to exercise this trajectory in practice, it is included in the study as a trajectory of intermediate difficulty. Finally, a flight profile recorded on-board a maneuvering F-15 at Edwards Air Force Base[1] is used to provide a more realistic test of the algorithm. The Edwards trajectory came complete with measured aircraft orientation, allowing a unique opportunity to quantify the performance degradation induced from having to estimate aircraft orientation. First, the algorithm is executed with the real aircraft orientation angles used in place of the ones estimated by the coordinated flight model. This serves as a baseline for comparison. Next, the simulation is run using the estimated aircraft orientation angles. Finally, a test is conducted in which the aircraft position is errantly estimated to be 300 m north and 300 m west of its actual location [2]. This gauges performance degradation due to errors in the position estimates.

The target library contains four aircraft. The F-15 and T-38A are both fighter aircraft,

[1]The F-15C trajectory was obtained from the Joint Helmet Cuing System, Mission JH-16, conducted by the 445th Flight Test Squadron at Edwards Air Force Base in May 2000.

[2]These errors are consistent with those estimated by Howland's passive radar tracker [29] under similar settings.

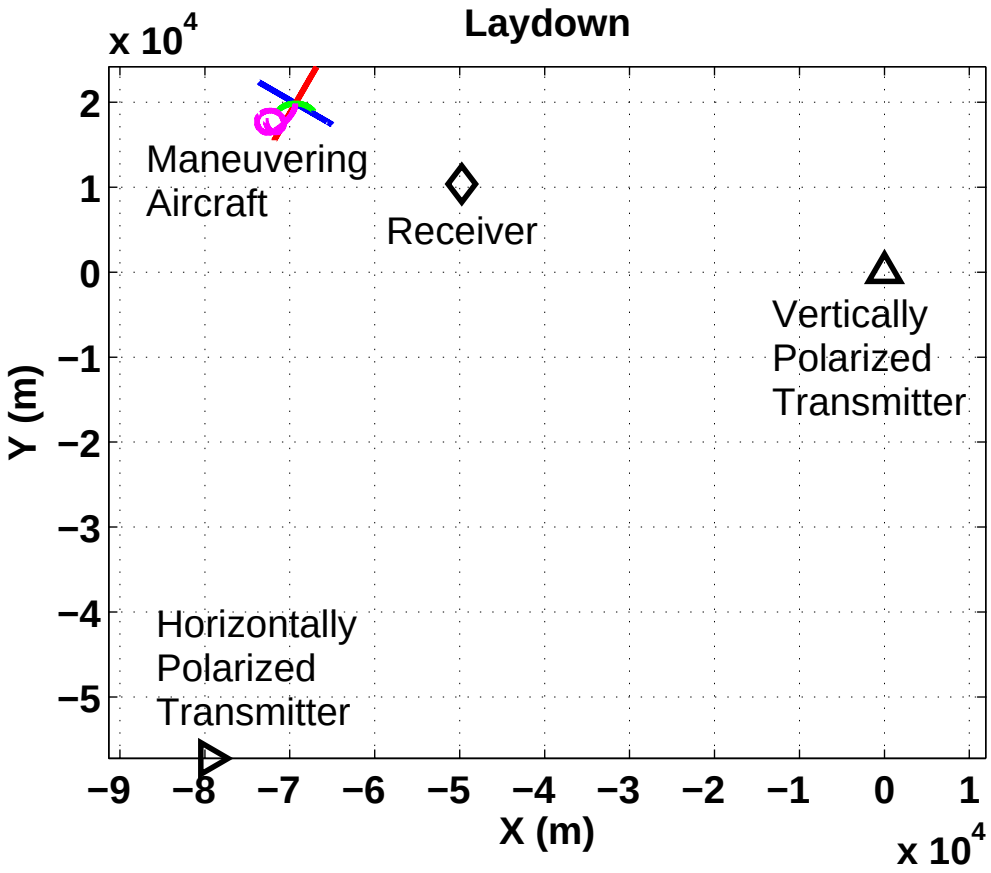

Figure 13: Trajectory laydowns

while the Falcon-20 and Falcon-100 are both commercial. The same library is used for all scenarios. To test each scenario, 400 Monte Carlo trials with different noisy "truth" profiles are given to the algorithm along with position and velocity measurements. A quarter of the profiles corresponds to each of the four aircraft being present. The trials are then repeated at noise figures ranging from 30 to 100, in increments of 5 dB. While the anticipated noise figures do not exceed 45 dB, the algorithm is tested under more stressing conditions to determine its breaking point.

4.1 Results Using a Vertically Polarized Transmitter

Simulations are conducted in which both horizontally and vertically polarized transmitters are exploited. This section focuses on the results obtained when the illuminator of opportunity is vertically polarized.

4.1.1 Straight-and-Level Trajectory #1

The aircraft in the first straight-and-level trajectory travel at a speeds of 200 m/s and altitudes of 8 km. The magnitudes of the power profiles resulting from this trajectory are shown in Figure 14a for all four targets. Since the targets are flying directly away from the receiver, exposing only a small range of aspect angles, these profiles are similar for all four targets. Thus, despite the simplicity of the trajectory, this is a reasonably difficult recognition problem. Even so, the algorithm performs extremely well at the anticipated noise levels. Figure 14b displays the probability of error as a function of noise figure for all four aircraft in the target library. Recognition errors are only noted once the noise figure hits 55 dB. This does not suggest that such an algorithm, if implemented in a real passive radar system, would perform without error at expected noise levels; real-world scenarios will always contain effects not modeled in even the best of simulations. It does suggest, however, that the algorithm has a great deal of potential for excellent performance.

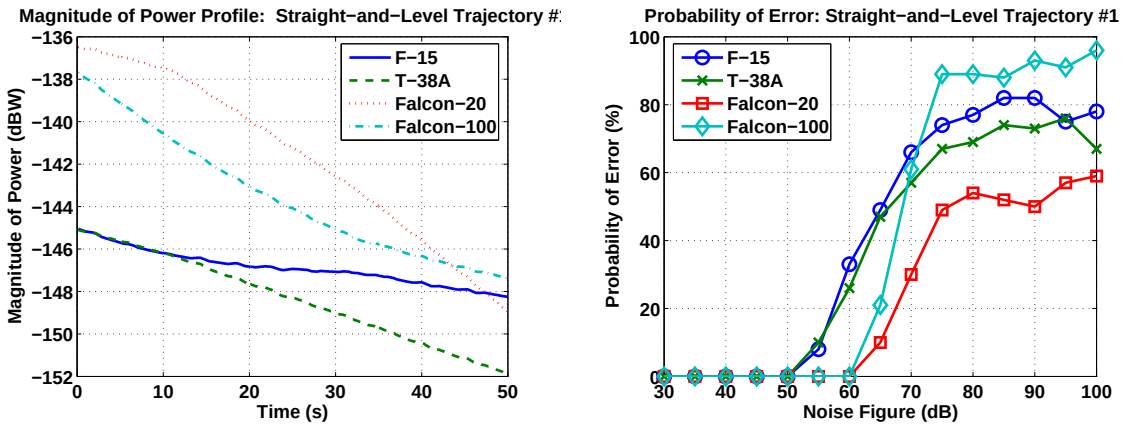

Figure 14: Straight-and-level trajectory #1: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

Also note that for noise figures greater than 70 dB, the probability of error averaged over all four aircraft types is roughly 75%. This result should be intuitive. Once the noise figure is that large, the noise swamps the target signal. The algorithm has little upon which to base its recognition and its cumulative odds of a correct match are reduced to one in four. The less intuitive result is that the probability of error for each *individual* aircraft is

something other than 75% and approaches a constant value that doesn't change significantly as the noise level increases. For example, the probability of error when the true target is a Falcon-20 remains between 50% and 60% for noise figures greater than 70 dB. However, this behavior is also attributed to noise swamping the underlying signal. When the noise figure is 100 dB, as shown in Figure 15, it is roughly 40 dB larger than the underlying signal. Under these conditions, the underlying signal has virtually no impact on the recognition outcome, and the ATR algorithm is always biased towards certain aircraft.

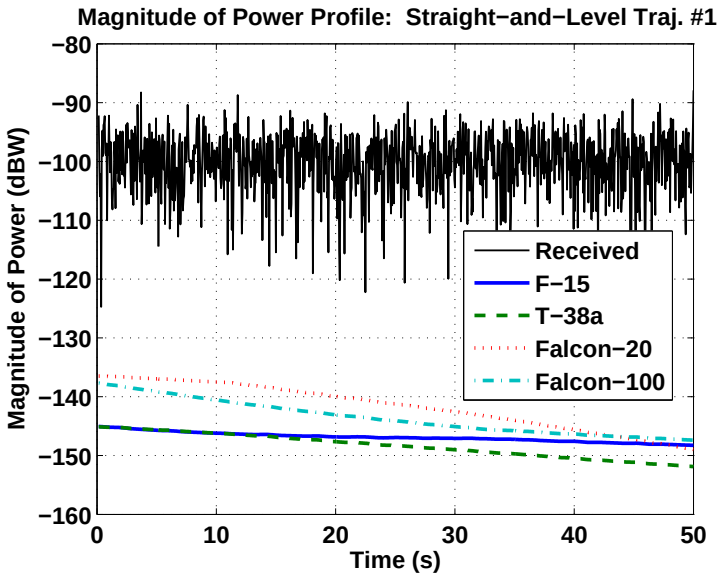

Figure 15: Straight-and-level trajectory #1, magnitude of power profile with a noise figure of 100 dB

To provide further insight into the recognition errors committed by the algorithm, a confusion matrix is shown in Table 4 for the case in which the noise figure is 65 dB. The aircraft listed across the top row correspond to the aircraft identified by the algorithm, while the aircraft listed in the leftmost column correspond to those that are actually present. For example, when the F-15 is actually present, it is correctly identified in 51 of 100 Monte Carlo trials, and is incorrectly identified as being a T-38 in 41 trials and a Falcon-100 in 8 trials. The confusion matrix reveals that in this scenario, the F-15 and T-38 are most likely to be swapped by the algorithm. Furthermore, if the algorithm incorrectly identifies either of the commercial aircraft, it is most likely to have swapped them.

Table 4: Confusion matrix for straight-and-level trajectory #1 with noise figure = 65 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	51	41	0	8
T-38A	40	53	0	7
Falcon-20	0	0	90	10
Falcon-100	6	3	12	79

4.1.2 Straight-and-Level Trajectory #2

The straight-and-level trajectory is then rotated ninety degrees so that the aircraft fly broadside to the receiver, exposing a much wider range of aspect angles. This angular diversity is evident when examining the power profiles, shown in Figure 16a. The probability of error for each aircraft is plotted in Figure 16b. Thanks to the increased diversity in the profile shapes, the algorithm is able to correctly identify the aircraft in all the Monte Carlo trials with noise figures less than 70 dB. When errors begin to be made, the F-15 is the least likely to be confused with another aircraft. The confusion matrix for a noise figure of 75 dB, shown in Table 5, makes this very clear.

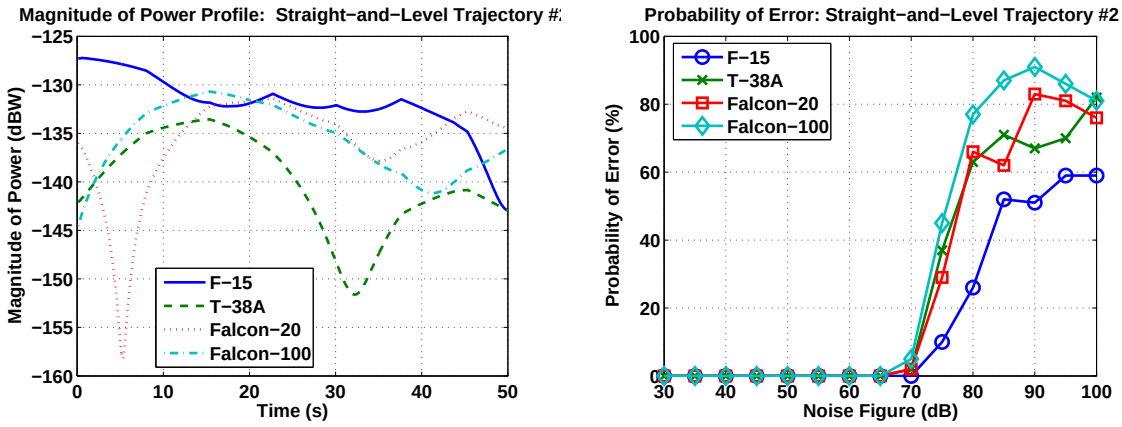

Figure 16: Straight-and-level trajectory #2: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

Table 5: Confusion matrix for straight-and-level trajectory #2 with noise figure = 65 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	90	0	1	9
T-38A	0	63	17	20
Falcon-20	0	16	71	13
Falcon-100	1	20	24	55

4.1.3 Banked Turn Trajectory

Aircraft executing the banked turn trajectory fly at speeds of 100 m/s and altitudes of 8 km. The power profiles resulting from this trajectory are presented in Figure 17a. Probability of error curves are shown in Figure 17b. Two factors conspire to make this a more challenging test of the recognition algorithm than the previous trajectory. First, a somewhat smaller range of angles is presented to the receiver. Additionally, the algorithm must estimate the aircraft's changing yaw, pitch, and roll. Even so, the results are still quite encouraging. The recognition algorithm makes no mistakes in the Monte Carlo trials until the noise figure reaches 65 dB, well above the expected maximum noise figure of 45 dB. At that point, the recognition errors mainly consist of swapping the two fighters. This is clear from Table 6.

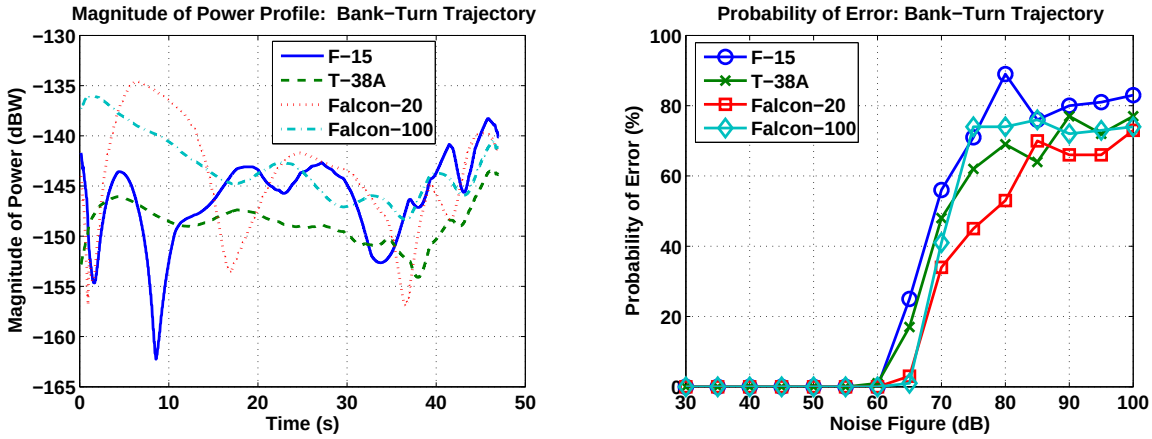

Figure 17: Banked turn trajectory: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

Table 6: Confusion matrix for banked turn trajectory with noise figure = 65 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	75	25	0	0
T-38A	14	83	0	3
Falcon-20	1	0	97	2
Falcon-100	1	0	0	99

4.1.4 Edwards Trajectory, with True Orientation Angles

The Edwards trajectory, recorded on-board a maneuvering F-15 at Edwards Air Force Base, provides a significantly more stressing test of the recognition algorithm. Unlike the first three trajectories, this trajectory is not idealized. Although the coordinated flight model is still a suitable model, nothing guarantees that the aircraft will actually maintain strictly coordinated flight throughout the entire maneuver. Since errors in the estimated orientation are more likely than before, this section establishes a performance baseline by providing the algorithm with the true aircraft orientation. Power profiles (magnitude only) and probability of error plots are shown in Figure 18. When provided with the true orientations, the recognition algorithm is able to correctly identify aircraft in all Monte Carlo trials with noise figures below 70 dB.

A confusion matrix corresponding to a noise figure of 75 dB is given in Table 7. This time, the recognition mistakes are more evenly distributed amongst the aircraft types. When an aircraft is misidentified, it is fairly likely to be misidentified as any of the other aircraft; swapping mistakes are less common than using some of the other trajectories.

Table 7: Confusion matrix for Edwards trajectory (using true orientation angles) with noise figure = 75 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	76	6	11	7
T-38A	13	52	17	18
Falcon-20	21	21	45	13
Falcon-100	19	21	18	42

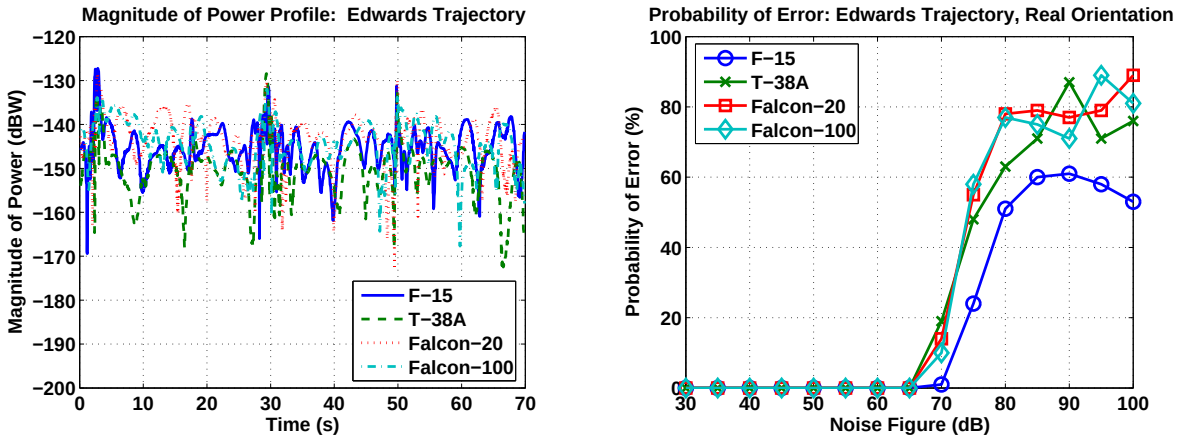

Figure 18: Edwards trajectory: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

4.1.5 Edwards Trajectory, with Approximated Orientation Angles

The simulations from Section 4.1.4 are repeated with the coordinated flight model back in the loop. Now that the algorithm must estimate the orientation, serious errors are made. Power profiles and probability of error plots are shown in Figure 19. The F-15 and Falcon-20 profiles, when computed from the approximated orientations, are more similar to the true Falcon-100 profile than to the F-15 and Falcon-20 profiles. Thus, these aircraft are incorrectly labeled as Falcon-100s. The confusion matrix in Table 8 demonstrates this problem. The trend only stops for noise figures above 50 dB, when the noise begins to swamp the signal and the probability of error of all aircraft tends towards 75 %. Clearly, further work is needed to make this algorithm more robust.

Table 8: Confusion matrix for Edwards trajectory (using approximated orientation angles) with noise figure = 45 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	0	0	0	100
T-38A	0	100	0	0
Falcon-20	0	0	0	100
Falcon-100	0	0	0	100

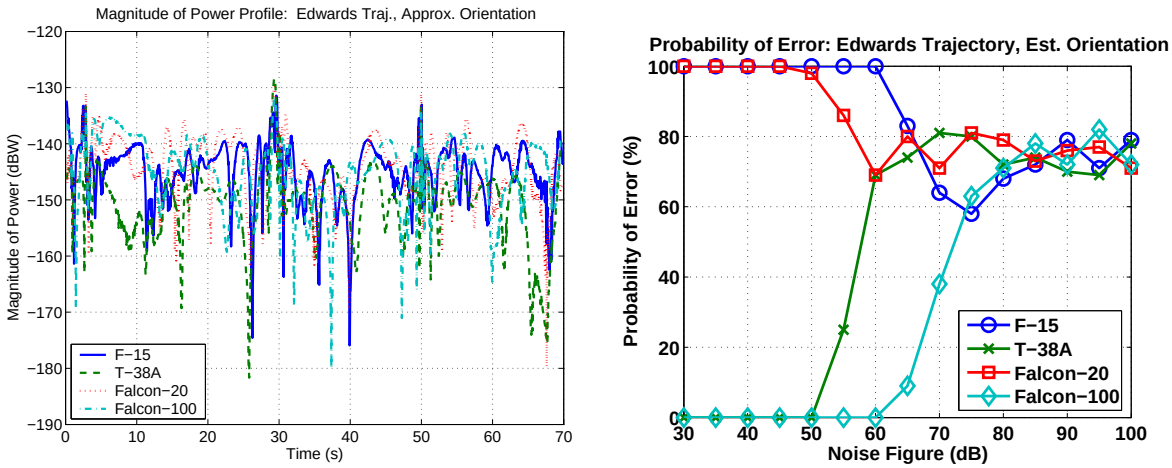

Figure 19: Edwards trajectory with approximated orientation: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

4.1.6 Edwards Trajectory, with Approximated Orientation Angles and Errors in the Position Estimates

The tests in Section 4.1.5 are then repeated with biased position measurements provided to the algorithm. Since the coordinated flight model operates on velocity, not position, estimates, the estimated orientations are unchanged. However, the incident and observed angles used to compile the expected magnitudes of the power profiles, shown in Figure 20a, are affected by the biased position measurements. Despite this change, the probabilities of error shown in Figure 20b are similar to those in Section 4.1.5. In fact, the confusion matrices are identical to those in Section 4.1.5 for noise figures below 50 dB. The incorrect orientation estimates appear to have a much more dramatic effect on the results than somewhat biased position estimates. Note that this trend may change if the position estimates from the passive radar tracker were substantially worse.

4.2 Results Using a Horizontally Polarized Transmitter

The Monte Carlo trials from Section 4.1 are then repeated using the horizontally polarized transmitter shown in Figure 13. This section is not intended to provide insight into the algorithm's performance as a function of antenna polarization. Rather, it lays the groundwork for Chapter 6, which exploits echoes from both transmitters when making a

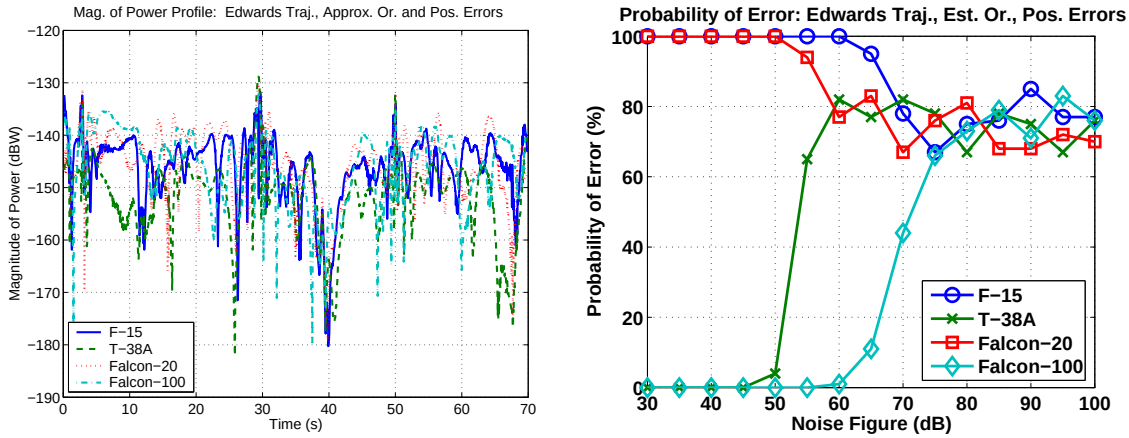

Figure 20: Edwards trajectory with approximated orientation: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

recognition decision. Transmitters of both horizontal and vertical polarizations are used because such circumstances will be quite common when implementing such algorithms in real-world settings.

4.2.1 Straight-and-Level Trajectory #1

The trajectory described in Section 4.1.1 is repeated using the horizontally polarized transmitter. The power profiles appear in Figure 21a, while the probability of error curves appear in Figure 21b. In this case, the algorithm begins making recognition mistakes when the noise figure reaches 50 dB. The confusion matrix shown in Table 9 reveals that mistakes at this noise level are limited to swapping the T-38A and Falcon-20. Again, a noise figure of 50 dB is above the anticipated noise level of a real system, indicating that the algorithm has great potential.

Table 9: Confusion matrix for straight-and-level trajectory #1 with noise figure = 50 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	100	0	0	0
T-38A	0	91	9	0
Falcon-20	0	9	91	0
Falcon-100	0	0	0	100

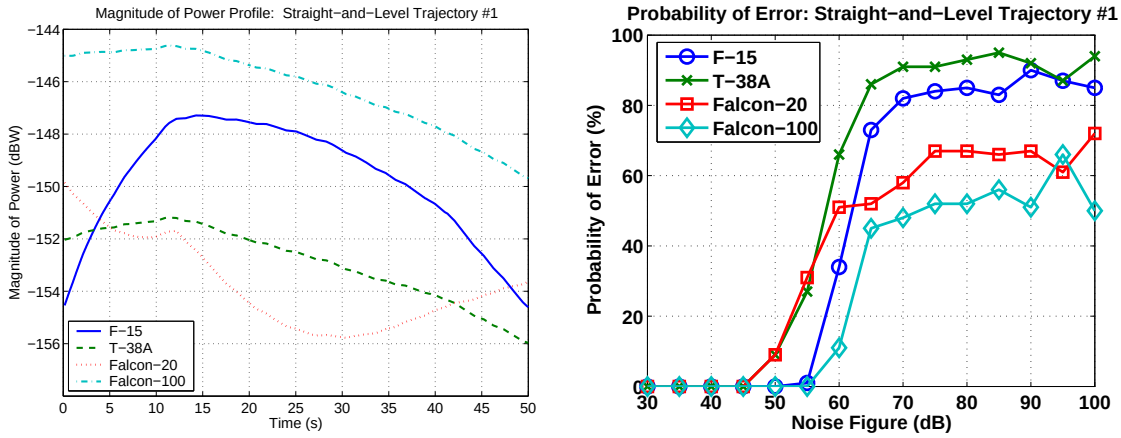

Figure 21: Straight-and-level trajectory #1: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

4.2.2 Straight-and-Level Trajectory #2

The power profiles for the second straight-and-level maneuver are shown in Figure 22a, and the power profiles are shown in Figure 22b. As before, the recognition algorithm's performance improves when a wider range of aspect angles are presented to the receiver. The results are still excellent at the anticipated noise figure of 45 dB. This time, the algorithm's mistakes are more evenly distributed amongst all aircraft classes, as is clear from the confusion matrix in Table 10.

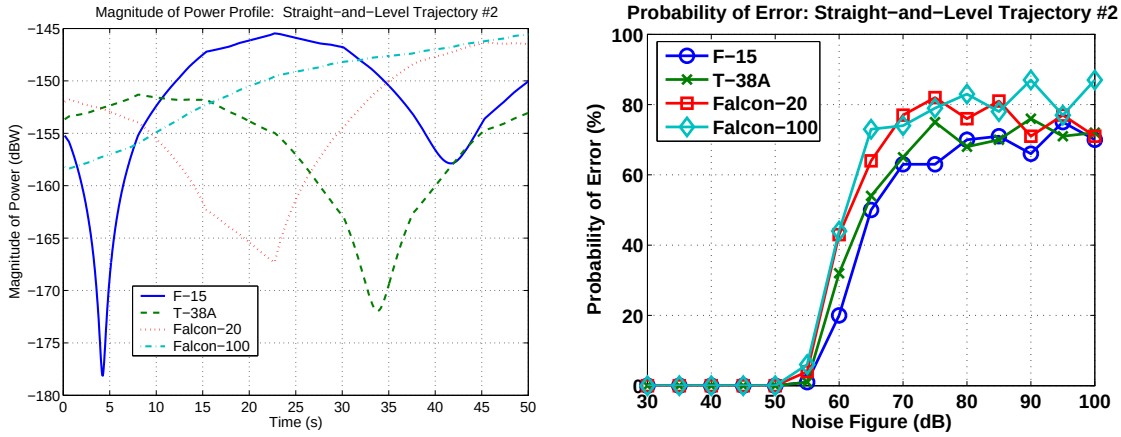

Figure 22: Straight-and-level trajectory #2: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

Table 10: Confusion matrix for straight-and-level trajectory #2 with noise figure = 60 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	80	8	1	11
T-38A	12	68	12	8
Falcon-20	3	17	57	23
Falcon-100	11	12	21	56

4.2.3 Banked Turn Trajectory

The algorithm performance is also quite good when the aircraft execute the banked turn maneuver. The power profiles appear in Figure 23a, while the probability of error curves are shown in Figure 23b. Once again, the algorithm's performance is very encouraging at the anticipated noise levels. As in the second straight-and-level trajectory, the classification errors are fairly evenly distributed over the four aircraft classes. This is clear from the confusion matrix in Table 11.

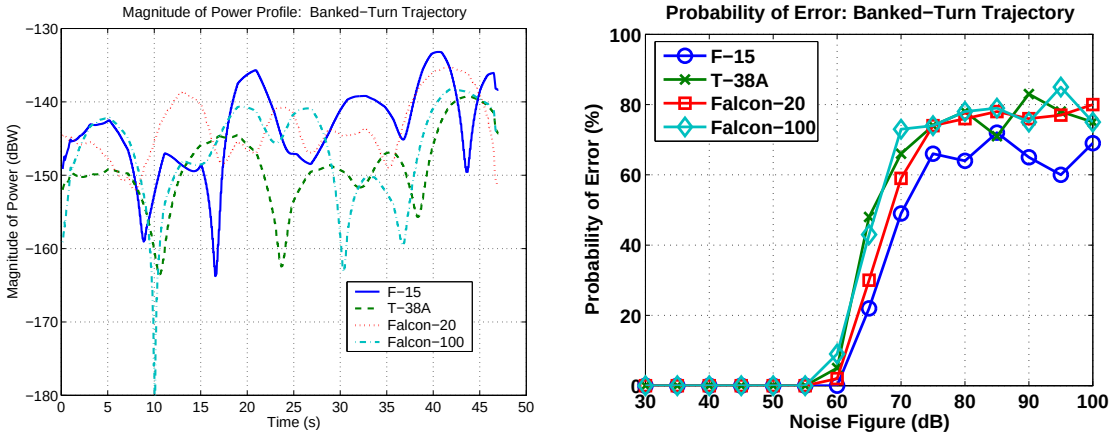

Figure 23: Banked turn trajectory: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

Table 11: Confusion matrix for banked turn trajectory with noise figure = 65 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	78	10	3	9
T-38A	15	52	8	25
Falcon-20	4	7	70	19
Falcon-100	8	18	17	57

4.2.4 Edwards Trajectory, with True Orientation Angles

The magnitudes of the power profiles resulting when aircraft execute the Edwards trajectory are shown in Figure 24a. The corresponding probability of error curves, obtained when the true orientations are provided to the algorithm, are shown in Figure 24b. Since a wide variety of aspect angles are presented to the receiver and estimation of the orientation angles is unnecessary, the recognition algorithm does extremely well. Even at a noise figure of 65 dB, 20 dB above the maximum level expected in a real system, the algorithm correctly identifies the aircraft in roughly 70% of the Monte Carlo trials. This is clear from the confusion matrix in Table 12.

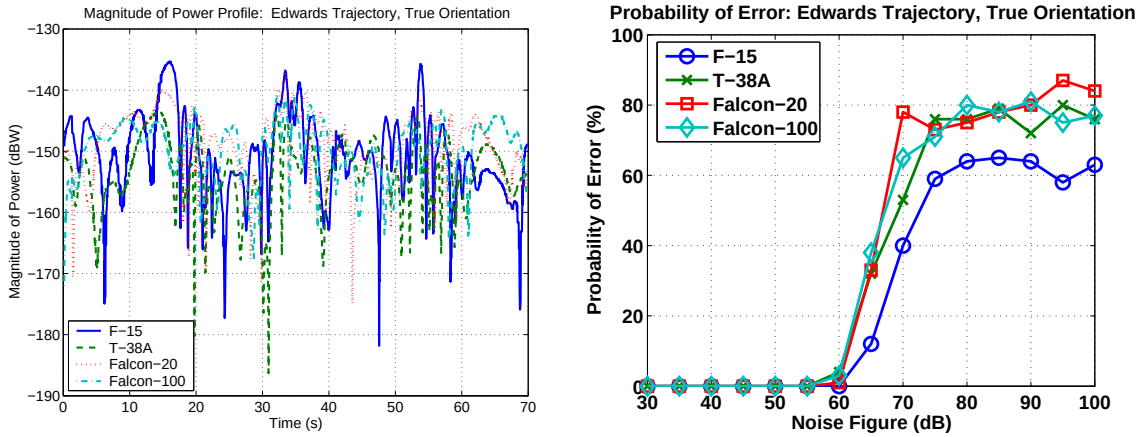

Figure 24: Edwards trajectory: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

Table 12: Confusion matrix for Edwards trajectory (using true orientation angles) with noise figure = 65 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	88	1	10	1
T-38A	1	68	8	23
Falcon-20	10	10	67	13
Falcon-100	3	23	12	62

4.2.5 Edwards Trajectory, with Approximated Orientation Angles

As in the case with the vertically polarized transmitter, the recognition performance degrades rapidly when aircraft execute the Edwards trajectory and the orientation angles are estimated. Because of errors in the estimated orientation angles, the algorithm identifies *all* the aircraft as Falcon-100s, until the noise figure is so large that the noise begins to swamp the signal. The magnitudes of the power profiles for this case are shown in Figure 25a, and the probability of error curves are shown in Figure 25b. A confusion matrix, shown in Table 13, confirms that the algorithm incorrectly identifies the aircraft as Falcon-100s.

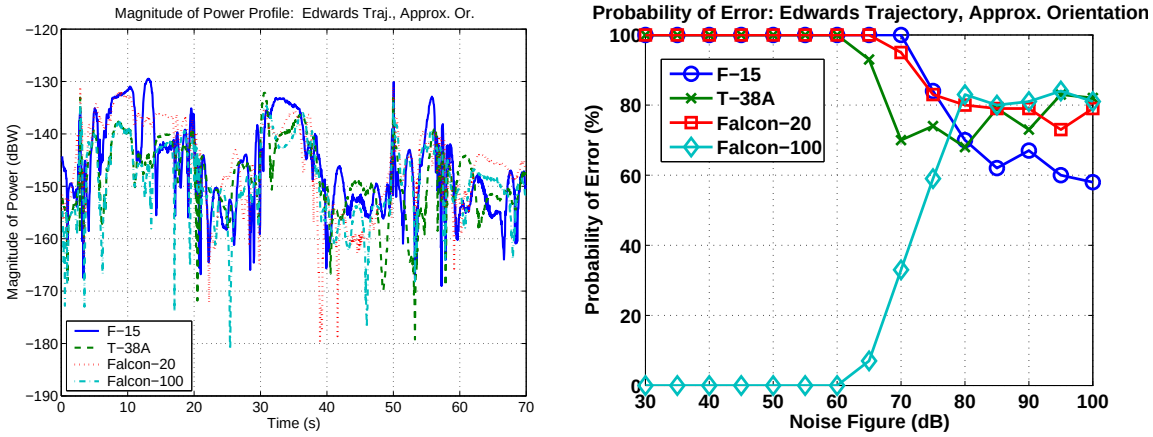

Figure 25: Edwards trajectory with approximated orientation: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

Table 13: Confusion matrix for Edwards trajectory (using approximated orientation angles) with noise figure = 50 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	0	0	0	100
T-38A	0	0	0	100
Falcon-20	0	0	0	100
Falcon-100	0	0	0	100

4.2.6 Edwards Trajectory, with Approximated Orientation Angles and Errors in the Position Estimates

Operating on biased position estimates has little effect on the algorithm that is already having such a difficult time identifying the aircraft. As in Section 4.1, the biased position

estimates are more of a second-order problem and contribute less to the algorithm's problems than the incorrect orientation estimates. The magnitudes of the power profiles are shown in Figure 26a, and the probability of error curves are provided in Figure 26b.

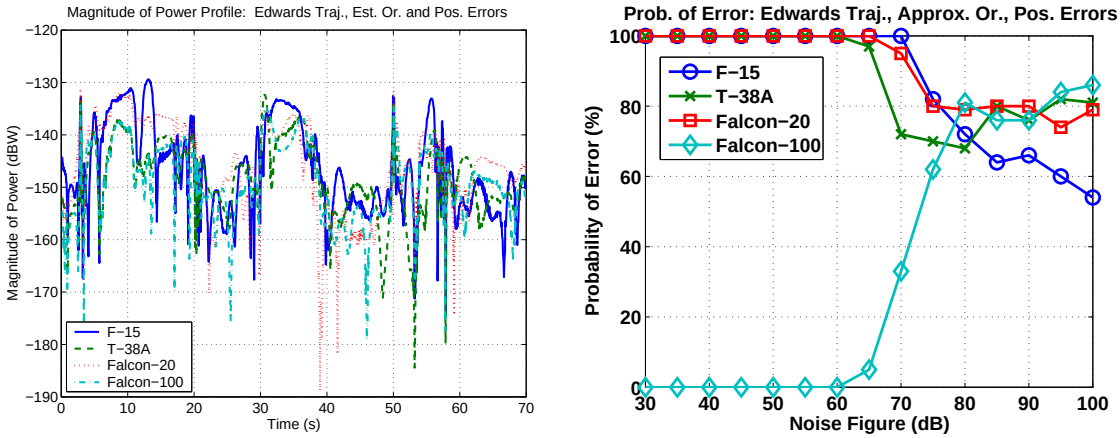

Figure 26: Edwards trajectory with approximated orientation: a) Magnitude of power profile (left), b) Probability of error vs. noise figure (right)

4.3 Summary of Results

This chapter demonstrates that the recognition algorithm has great potential, particularly if the aircraft is not executing a complicated maneuver. Since this sort of technology is frequently envisioned in use with tracking algorithms that covertly follow aircraft in transit, rather than in dog-fighting situations, this algorithm is probably sufficiently reliable for the vast majority of realistic scenarios. Even so, a more robust algorithm will be developed and tested in Chapter 6. The more robust algorithm will jointly estimate the aircraft's type and orientation (from a feasible set determined by the EKF's state and covariance). It will also exploit multiple transmitters, to reduce the uncertainty regarding the aircraft orientation. In doing so, the more robust algorithm will be able to avoid the failings of this algorithm when attempting to identify aircraft executing complicated maneuvers.

CHAPTER V

ESTIMATING ATR PERFORMANCE VIA RELATIVE ENTROPY AND CHERNOFF INFORMATION

The Rician distribution describes a random variable z, where z is related to independent identically distributed Normal random variables, x and y by,

$$z = \sqrt{(x+b)^2 + y^2}. \tag{53}$$

This model, which reduces to a Rayleigh distribution if $b = 0$, arises in a number of engineering problems, from fading multipath channels in communications to modeling the radar cross section of aircraft using low-frequency radar. However, since the Rician pdf is not mathematically convenient, scenarios that ought to be modeled with Rician densities are often modeled with Normal densities [35], leading to increased modeling error.

This is particularly true when computing the relative entropy and Chernoff information for scenarios that are best described by Rician densities. These information measures have great utility in a number of hypothesis-testing problems. For example, Stein's Lemma [14] reveals that in the Neyman-Pearson framework, a bound on the probability of a Type II error (i.e. the probability that the target modeled with q(x) will be misidentified as the target modeled with p(x)), $\beta_{p||q}$, is approximately

$$\beta_{p||q} \approx e^{-D(p(x)||q(x))}, \tag{54}$$

where $D\left(p(x)||q(x)\right)$ is the relative entropy between densities $p(x)$ and $q(x)$. Thus, the existence of a closed-form approximation of the relative entropy between two Ricians would make it possible to approximate the probability of Type II error in the Neyman-Pearson framework. Section 5.1 derives a closed-form approximation and shows its superiority to the usual Normal approximation.

If the problem is set in a Bayesian framework, rather than a Neyman-Pearson framework, the Chernoff information could be used to bound the probability of error, P_E. In particular, the probability of error in a Bayesian hypothesis testing problem is approximated as

$$P_E \approx e^{-C(p(x),q(x))}, \tag{55}$$

where $C\left(p(x), q(x)\right)$ is the Chernoff information [14]. Section 5.2 derives a closed-form approximation and compares its accuracy with that of the usual Normal approximation.

5.1 Estimating Performance Under the Neyman-Pearson Framework Via the Relative Entropy

5.1.1 Derivation of the Relative Entropy Between Two Rician Densities

The relative entropy between two Rician probability density functions, $p(x)$ and $q(x)$, is given by

$$D\left(p(x)||q(x)\right) = \int_0^\infty p(x) \ln\left(\frac{p(x)}{q(x)}\right) dx, \tag{56}$$

where $p(x)$ is defined as

$$p(x) = \frac{x}{\sigma^2} e^{\frac{-(x^2+s_p^2)}{2\sigma^2}} I_0\left[\frac{x s_p}{\sigma^2}\right], \tag{57}$$

and $q(x)$ is defined as

$$q(x) = \frac{x}{\sigma^2} e^{\frac{-(x^2+s_q^2)}{2\sigma^2}} I_0\left[\frac{x s_q}{\sigma^2}\right]. \tag{58}$$

Substituting (57) and (58) into (56) reveals that the relative entropy between two Rician densities with the same σ^2 is

$$D(p(x)||q(x)) = \int_0^\infty \frac{x}{\sigma^2} e^{\frac{-(x^2+s_p^2)}{2\sigma^2}} I_0\left(\frac{x s_p}{\sigma^2}\right) \left\{ \frac{s_q^2 - s_p^2}{2\sigma^2} + \ln\left[I_0\left(\frac{x s_p}{\sigma^2}\right)\right] - \ln\left[I_0\left(\frac{x s_q}{\sigma^2}\right)\right] \right\} dx. \tag{59}$$

This is then broken into three integrals. If $p(x)$ is substituted back in for ease of notation, then this becomes

$$D\left(p(x)||q(x)\right) = \left(\frac{s_q^2 - s_p^2}{2\sigma^2}\right)\int_0^\infty p(x)dx + \int_0^\infty p(x)\ln\left[I_0\left(\frac{xs_p}{\sigma^2}\right)\right]dx - \int_0^\infty p(x)\ln\left[I_0\left(\frac{xs_q}{\sigma^2}\right)\right]dx, \tag{60}$$

which reduces to

$$D(p(x)||q(x)) = \left(\frac{s_q^2 - s_p^2}{2\sigma^2}\right) + \int_0^\infty p(x)\ln\left[I_0\left(\frac{xs_p}{\sigma^2}\right)\right]dx - \int_0^\infty p(x)\ln\left[I_0\left(\frac{xs_q}{\sigma^2}\right)\right]dx. \tag{61}$$

The approximations presented in Sections 5.1.2 and 5.1.3 are suggested as a means of evaluating the two remaining integrals.

5.1.2 A Normal Approximation for the Relative Entropy Between Two Rician Densities

The relative entropy between two Normal distributions, $p(x)$ and $q(x)$, is given by

$$D\left(p(x)||q(x)\right) = \ln\left(\frac{v_p}{v_q}\right) + \frac{1}{2v_q^2}\left[v_p^2 + (\mu_p - \mu_q)^2\right] - \frac{1}{2}, \tag{62}$$

where $p(x) \sim N(\mu_p, v_p)$ and $q(x) \sim N(\mu_q, v_q)$ [32]. If the means and variances of the two Normal densities are set to match the means and variances of the Rician densities, then (62) approximates the relative entropy between two Rician distributions.

The mean of a Rician density is given by

$$E[x] = \sqrt{\frac{\pi\sigma^2}{2}}e^{\frac{-s^2}{4\sigma^2}}\left[\left(1 + \frac{s^2}{2\sigma^2}\right)I_0\left(\frac{s^2}{4\sigma^2}\right) + \left(\frac{s^2}{2\sigma^2}\right)I_1\left(\frac{s^2}{4\sigma^2}\right)\right], \tag{63}$$

and the variance is

$$var(x) = s^2 + 2\sigma^2 - E^2[x]. \tag{64}$$

Thus, the means of the Normal densities should be set such that

$$\mu_p = \sqrt{\frac{\pi\sigma_p^2}{2}}e^{\frac{-s_p^2}{4\sigma_p^2}}\left[\left(1 + \frac{s_p^2}{2\sigma_p^2}\right)I_0\left(\frac{s_p^2}{4\sigma_p^2}\right) + \left(\frac{s_p^2}{2\sigma_p^2}\right)I_1\left(\frac{s_p^2}{4\sigma_p^2}\right)\right], \tag{65}$$

and

$$\mu_q = \sqrt{\frac{\pi\sigma_q^2}{2}} e^{\frac{-s_q^2}{4\sigma_q^2}} \left[\left(1 + \frac{s_q^2}{2\sigma_q^2} \right) I_0 \left(\frac{s_q^2}{4\sigma_q^2} \right) + \left(\frac{s_q^2}{2\sigma_q^2} \right) I_1 \left(\frac{s_q^2}{4\sigma_q^2} \right) \right]. \tag{66}$$

The variances should then be set such that

$$v_p^2 = s_p^2 + 2\sigma_p^2 - \mu_p^2, \tag{67}$$

and

$$v_q^2 = s_q^2 + 2\sigma_q^2 - \mu_q^2. \tag{68}$$

5.1.3 Derivation of a Closed-Form Approximation for the Relative Entropy Between Two Rician Densities

The proposed closed-form approximation uses the Laplace method [15] to evaluate the integrals in (61). Begin with the second integral. Applying the Laplace method transforms

$$\int_0^\infty \frac{x}{\sigma^2} e^{\frac{-(x^2+s_p^2)}{2\sigma^2}} I_0 \left(\frac{x s_p}{\sigma^2} \right) \ln \left[I_0 \left(\frac{x s_q}{\sigma^2} \right) \right] dx, \tag{69}$$

into

$$\int_0^\infty e^{\left\{ \ln\left(\frac{x}{\sigma^2}\right) - \frac{(x^2+s_p^2)}{2\sigma^2} + \ln\left[I_0\left(\frac{x s_p}{\sigma^2}\right)\right] + \ln\left\{ \ln\left[I_0\left(\frac{x s_q}{\sigma^2}\right)\right]\right\}\right\}} dx. \tag{70}$$

Now define $h_{pq}(x)$ such that

$$h_{pq}(x) = \ln \left(\frac{x}{\sigma^2} \right) - \frac{(x^2 + s_p^2)}{2\sigma^2} + \ln \left[I_0 \left(\frac{x s_p}{\sigma^2} \right) \right] + \ln \left\{ \ln \left[I_0 \left(\frac{x s_q}{\sigma^2} \right) \right] \right\}, \tag{71}$$

reducing the integral in (70) to

$$\int_0^\infty e^{h_{pq}(x)} dx. \tag{72}$$

Taking the Taylor Series expansion of $h_{pq}(x)$ around the value of x that maximizes $h_{pq}(x)$,[1] $\hat{x}$, results in

[1] Much like the normal density, the Rician density is unimodal with exactly one global maximum and no global minima (except at $\pm\infty$). Since $h_{pq}(x)$ is merely a scaled version of a Rician density, it is clear that $\hat{x}$ maximizes (rather than minimizes) $h_{pq}(x)$.

$$h_{pq}(x) \approx h_{pq}(\hat{x}) + \frac{1}{2} h''_{pq}(\hat{x})(x - \hat{x})^2. \tag{73}$$

Thus, (72) is approximated by

$$\int_0^\infty e^{h_{pq}(\hat{x}) + \frac{1}{2} h''_{pq}(\hat{x})(x-\hat{x})^2} dx, \tag{74}$$

which reduces to

$$e^{h_{pq}(\hat{x})} \int_0^\infty e^{\frac{h''_{pq}(\hat{x})(x-\hat{x})^2}{2}} dx. \tag{75}$$

This is manipulated into the form of a Normal density as

$$e^{h_{pq}(\hat{x})} \sqrt{\frac{-2\pi}{h''_{pq}(\hat{x})}} \int_{-\infty}^\infty \frac{1}{\sqrt{\frac{-2\pi}{h''_{pq}(\hat{x})}}} e^{\frac{-(x-\hat{x})^2}{2\frac{1}{h''_{pq}(\hat{x})}}} dx, \tag{76}$$

or

$$e^{h_{pq}(\hat{x})} \sqrt{\frac{-2\pi}{h''_{pq}(\hat{x})}}. \tag{77}$$

This is the closed-form approximation of the third integral in (61). Similarly, the second integral in (61) reduces to

$$e^{h_{pp}(\hat{x})} \sqrt{\frac{-2\pi}{h''_{pp}(\hat{x})}}. \tag{78}$$

The closed-form approximation for the relative entropy between two Ricians is then given by,

$$D\left(p(x)||q(x)\right) \approx \frac{s_q^2 - s_p^2}{2\sigma^2} + e^{h_{pp}(\hat{x})} \sqrt{\frac{-2\pi}{h''_{pp}(\hat{x})}} - e^{h_{pq}(\hat{x})} \sqrt{\frac{-2\pi}{h''_{pq}(\hat{x})}}. \tag{79}$$

All that remains to finish this closed-form approximation is to find expressions for $\hat{x}$ and $h''_{pq}(\hat{x})$. The first derivative of $h_{pq}(x)$ is

$$h'_{pq}(x) = \frac{1}{x} - \frac{x}{\sigma^2} + \left(\frac{s_p}{\sigma^2}\right) \left(\frac{I_1\left(\frac{x s_p}{\sigma^2}\right)}{I_0\left(\frac{x s_p}{\sigma^2}\right)} \right) + \left(\frac{s_q}{\sigma^2}\right) \left(\frac{I_1\left(\frac{x s_q}{\sigma^2}\right)}{I_0\left(\frac{x s_q}{\sigma^2}\right) \ln\left[I_0\left(\frac{x s_q}{\sigma^2}\right)\right]} \right). \tag{80}$$

Setting this equal to zero and using the approximation $I_p(z) \approx z$ results in

$$\frac{2}{x} - \frac{x}{\sigma^2} + \frac{s_p}{\sigma^2} = 0. \tag{81}$$

This has a solution for $\hat{x}$ of

$$\hat{x} \approx \frac{1}{2}\left(s_p + \sqrt{s_p^2 + 8\sigma^2}\right). \tag{82}$$

Taking the derivative of (80) produces the expression

$$h''_{pq}(x) = \tfrac{-1}{x} - \tfrac{1}{\sigma^2} + \left(\tfrac{s_p}{\sigma^2}\right)^2 \left(\frac{I_1^2\left(\frac{xs_p}{\sigma^2}\right)}{I_0^2\left(\frac{xs_p}{\sigma^2}\right)}\right) + \left(\tfrac{s_p}{\sigma^2}\right)^2 \left(\frac{I_0\left(\frac{xs_p}{\sigma^2}\right) + I_2\left(\frac{xs_p}{\sigma^2}\right)}{I_0\left(\frac{xs_p}{\sigma^2}\right)}\right) \cdots$$
$$\cdots - \left(\tfrac{s_q}{\sigma^2}\right)^2 \left(\frac{I_1^2\left(\frac{xs_q}{\sigma^2}\right)}{I_0^2\left(\frac{xs_q}{\sigma^2}\right)\ln^2\left[I_0\left(\frac{xs_q}{\sigma^2}\right)\right]}\right) - \left(\tfrac{s_q}{\sigma^2}\right)^2 \left(\frac{I_1^2\left(\frac{xs_q}{\sigma^2}\right)}{I_0^2\left(\frac{xs_q}{\sigma^2}\right)\ln\left[I_0\left(\frac{xs_q}{\sigma^2}\right)\right]}\right) \cdots$$
$$\cdots + \left(\frac{s_q}{\sigma^2}\right)^2 \left(\frac{I_0\left(\frac{xs_q}{\sigma^2}\right) + I_2\left(\frac{xs_q}{\sigma^2}\right)}{I_0^2\left(\frac{xs_q}{\sigma^2}\right)\ln\left[I_0\left(\frac{xs_q}{\sigma^2}\right)\right]}\right). \tag{83}$$

for the second derivative. In summary, the closed-form approximation of the relative entropy is obtained by substituting (82) and (83) into (79).

To compare the Normal approximation with the closed-form approximation, the s_q parameter is swept over a range of values while the s_p and σ^2 parameters are held constant. Figure 27 shows the resulting relative entropy and probability of a Type II error when $s_p = 10$, $\sigma^2 = 4$, and s_q is swept from 0 to 20. The closed-form approximation derived using the Laplace method, however, produces results that are nearly identical to the numerically approximated results. The Normal approximation is clearly not as accurate as the closed-form solution for the smaller values of s_q. The difference becomes even more apparent if we reduce s_p to 5 and sweep s_q from 0 to 10. This case is shown in Figure 28.

5.1.4 Applying the Closed-Form Approximation for the Relative Entropy to the ATR Algorithm

The relative entropy is used to estimate the probability of a Type II error, β, for each aircraft pairing. These are compared to the Monte Carlo runs from Chapter 4. Several trends are compared, including the lowest noise level at which the algorithm makes classification errors, the noise level at which the algorithm's odds of success are reduced to one in four, and the most likely aircraft to be swapped.

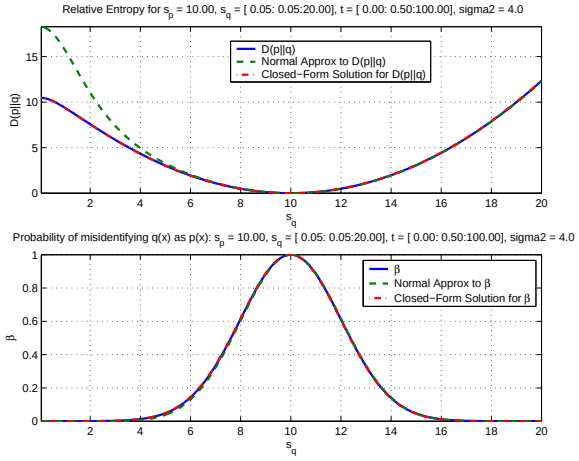

Figure 27: $s_p = 10$, $\sigma^2 = 4$, and s_q sweeps from 0 to 20: top: $D(p(x)||q(x))$, bottom: $\beta_{p||q}$

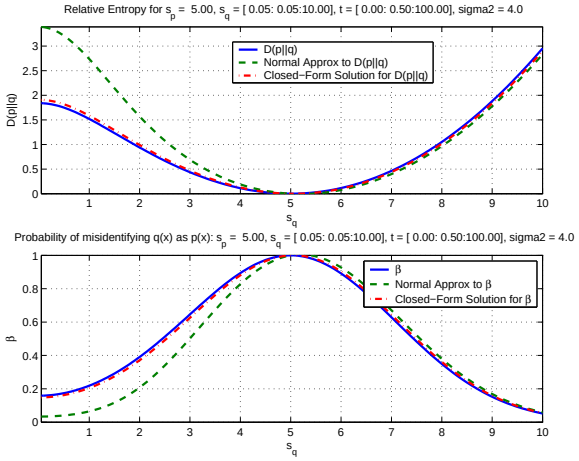

Figure 28: $s_p = 5$, $\sigma^2 = 4$, and s_q sweeps from 0 to 10: top: $D(p(x)||q(x))$, bottom: $\beta_{p||q}$

Results for the straight-and-level trajectories using the vertically polarized transmitter from Section 3.1 appear in Figures 29 and 30. The agreement between the predictions made using the relative entropy and the results observed in the Monte Carlo trials from Sections 4.1.1 and 4.1.2 is quite good. Table 14 summarizes the comparison for the first straight-and-level trajectory. The relative entropy predictions are slightly optimistic regarding the noise figure at which the algorithm makes its first classification error. However, they provide reliable indicators of the noise level at which the algorithm breaks down, as well as the aircraft that are most likely to be swapped. For example, the relative entropy predicts that

if an aircraft is misidentified as being an F-15, it is likeliest to be a T-38A. The Monte Carlo trials corroborate with this result. Similar data is shown for the second straight-and-level trajectory in Table 15. Once again, the relative entropy predictions provide an excellent indicator of the Monte Carlo results, without the need for conducting computationally expensive Monte Carlo trials.

Table 14: Comparison of relative entropy predictions and Monte Carlo results: Straight-and-level trajectory #1, using the vertically polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	55	60
Noise Figure at which Algorithm Breaks	80	80
Most Likely Aircraft to be Misidentified as the F-15	T-38A	T-38A
Most Likely Aircraft to be Misidentified as the T-38A	F-15	F-15
Most Likely Aircraft to be Misidentified as the Falcon-20	Falcon-100	Falcon-100
Most Likely Aircraft to be Misidentified as the Falcon-100	all	all

Table 15: Comparison of relative entropy predictions and Monte Carlo results: Straight-and-level trajectory #2, using the vertically polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	70	75
Noise Figure at which Algorithm Breaks	90	90
Most Likely Aircraft to be Misidentified as the F-15	all	all
Most Likely Aircraft to be Misidentified as the T-38A	Falcon-20 and Falcon-100	Falcon-20 and Falcon-100
Most Likely Aircraft to be Misidentified as the Falcon-20	Falcon-100 and T-38A	Falcon-100 and T-38A
Most Likely Aircraft to be Misidentified as the Falcon-100	Falcon-20 and T-38A	Falcon-20 and T-38A

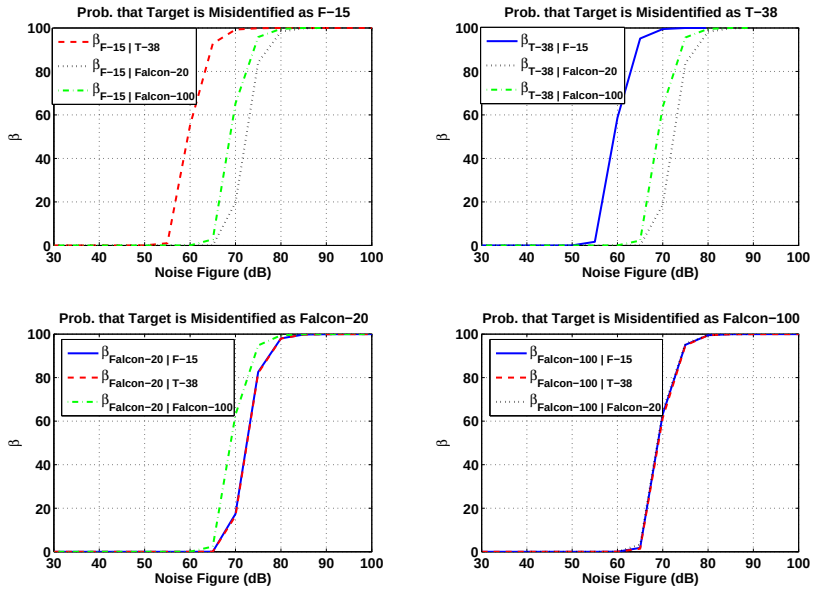

Figure 29: Probability of a type II error for straight-and-level trajectory #1 (vert. pol. transmitter)

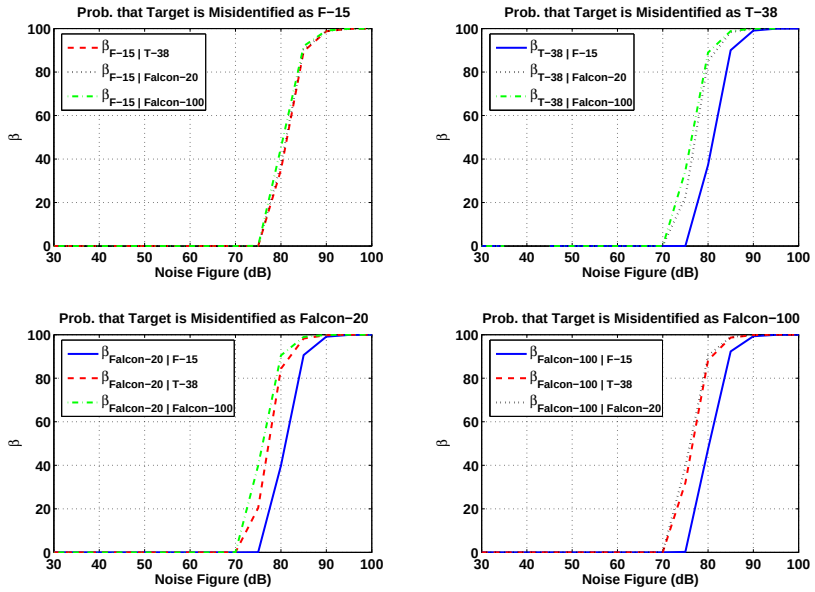

Figure 30: Probability of a type II error for straight-and-level trajectory #2 (vert. pol. transmitter)

Figure 31 shows the results when the banked turn trajectory is used. A comparison between the relative entropy predictions and Monte Carlo trials is shown in Table 16. As with the last two trajectories, the relative entropy predictions are excellent indicators of the Monte Carlo results.

Figures 32 through 34 show the results for the Edwards trajectory with real orientation angles, approximated orientation angles, and approximated orientation angles with incorrect position estimates, respectively. The corresponding comparisons between the relative entropy predictions and Monte Carlo results appear in Tables 17 through 19. Although the relative entropy predictions agree well with the Monte Carlo results obtained when the true orientation angles are provided to the algorithm, they are significantly different once approximated orientation angles are used. The reason is simple. The relative entropy predictions are based upon the assumption that the correct orientation angles are used in the ATR algorithm. Thus, the relative entropy predictions derived in this chapter do not account for the possibility that incorrect orientation angles are fed from the coordinated flight model to the ATR algorithm, causing a breakdown in performance. Such differences are only noticeable when the approximated orientation angles differ significantly from the truth, as in the case of the Edwards trajectory.

Table 16: Comparison of relative entropy predictions and Monte Carlo results: Banked turn trajectory, using the vertically polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	65	65
Noise Figure at which Algorithm Breaks	80	80
Most Likely Aircraft to be Misidentified as the F-15	T-38A	T-38A
Most Likely Aircraft to be Misidentified as the T-38A	F-15	F-15
Most Likely Aircraft to be Misidentified as the Falcon-20	all	all
Most Likely Aircraft to be Misidentified as the Falcon-100	all	all

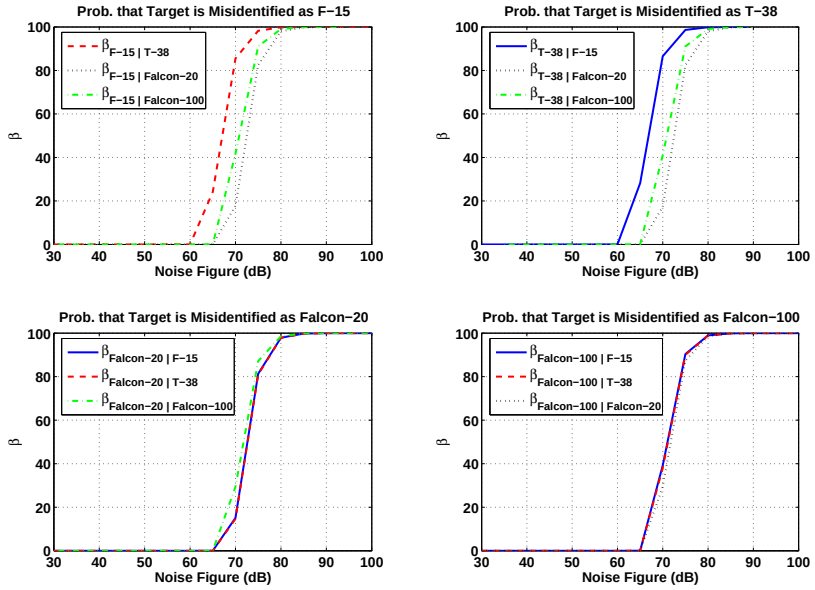

Figure 31: Probability of a type II error for banked turn trajectory (vert. pol. transmitter)

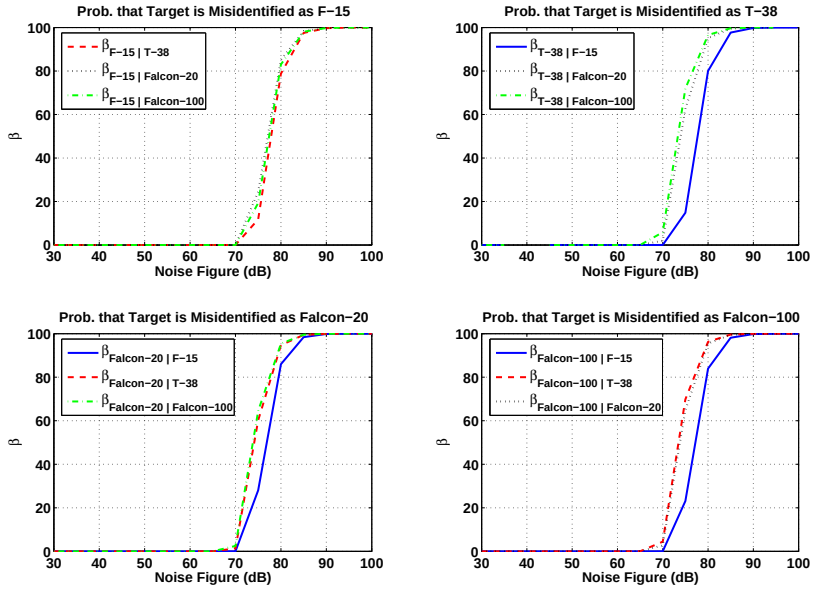

Figure 32: Probability of a type II error for Edwards trajectory with true orientation (vert. pol. transmitter)

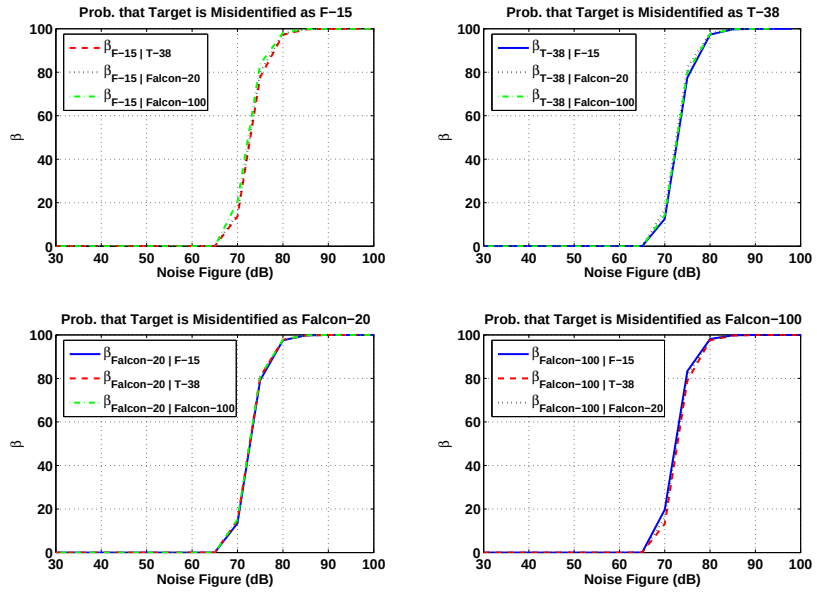

Figure 33: Probability of a type II error for Edwards trajectory with approximated orientation (vert. pol. transmitter)

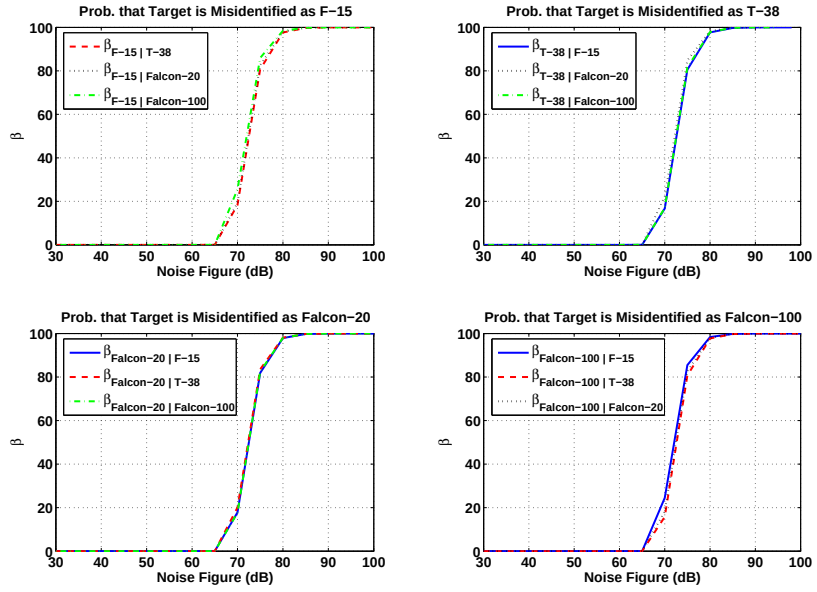

Figure 34: Probability of a type II error for Edwards trajectory with approximated orientation and biased position estimates (vert. pol. transmitter)

Table 17: Comparison of relative entropy predictions and Monte Carlo results: Edwards trajectory with true orientations, using the vertically polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	70	70
Noise Figure at which Algorithm Breaks	90	90
Most Likely Aircraft to be Misidentified as the F-15	all	all
Most Likely Aircraft to be Misidentified as the T-38A	Falcon-20 and Falcon-100	Falcon-20 and Falcon-100
Most Likely Aircraft to be Misidentified as the Falcon-20	Falcon-100 and T-38A	Falcon-100 and T-38A
Most Likely Aircraft to be Misidentified as the Falcon-100	Falcon-20 and T-38A	Falcon-20 and T-38A

Table 18: Comparison of relative entropy predictions and Monte Carlo results: Edwards trajectory with approximated orientations, using th vertically polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	5	70
Noise Figure at which Algorithm Breaks	5	80
Most Likely Aircraft to be Misidentified as the F-15	T-38A	all
Most Likely Aircraft to be Misidentified as the T-38A	Falcon-20 and Falcon-100	all
Most Likely Aircraft to be Misidentified as the Falcon-20	F-15	all
Most Likely Aircraft to be Misidentified as the Falcon-100	Falcon-20 and F-15	all

Table 19: Comparison of relative entropy predictions and Monte Carlo results: Edwards trajectory with approximated orientations and biased position measurements, using the vertically polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	5	70
Noise Figure at which Algorithm Breaks	5	80
Most Likely Aircraft to be Misidentified as the F-15	T-38A	all
Most Likely Aircraft to be Misidentified as the T-38A	Falcon-20 and Falcon-100	all
Most Likely Aircraft to be Misidentified as the Falcon-20	F-15	all
Most Likely Aircraft to be Misidentified as the Falcon-100	Falcon-20 and F-15	all

The comparison is then repeated with the horizontally polarized transmitter from Section 3.1. Figures 35 through 40 show the results predicted by the relative entropy method. Comparisons between these results and the Monte Carlo results from Section 4.2 are provided in Tables 20 through 25. With the exception of the last two cases, in which the aircraft orientation is estimated for the Edwards trajectory, the trends observed from the relative entropy match those from the Monte Carlo trials.

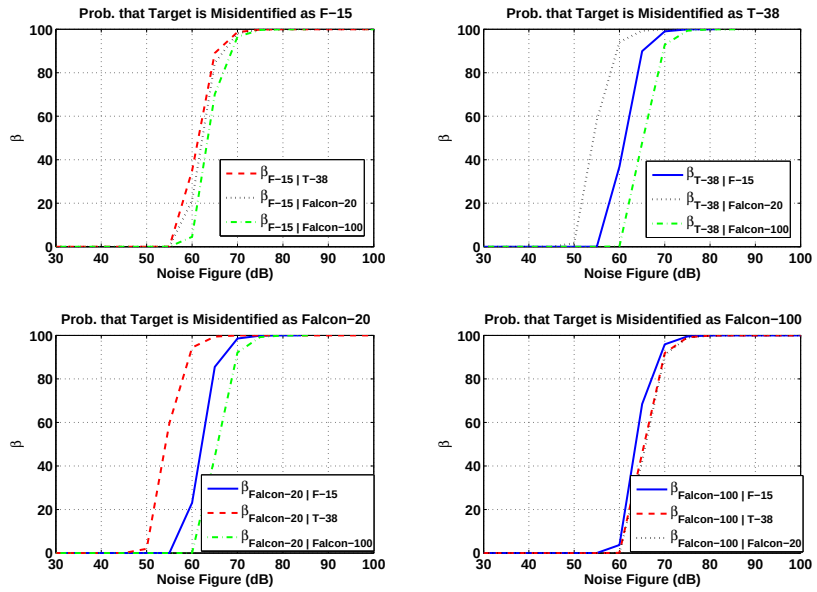

Figure 35: Probability of a type II error for straight-and-level trajectory #1 (hor. pol. transmitter)

Table 20: Comparison of relative entropy predictions and Monte Carlo results: straight-and-level trajectory #1, using the horizontally polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	50	55
Noise Figure at which Algorithm Breaks	70	70
Most Likely Aircraft to be Misidentified as the F-15	all	all
Most Likely Aircraft to be Misidentified as the T-38A	Falcon-20	Falcon-20
Most Likely Aircraft to be Misidentified as the Falcon-20	T-38A	T-38A
Most Likely Aircraft to be Misidentified as the Falcon-100	F-15	F-15

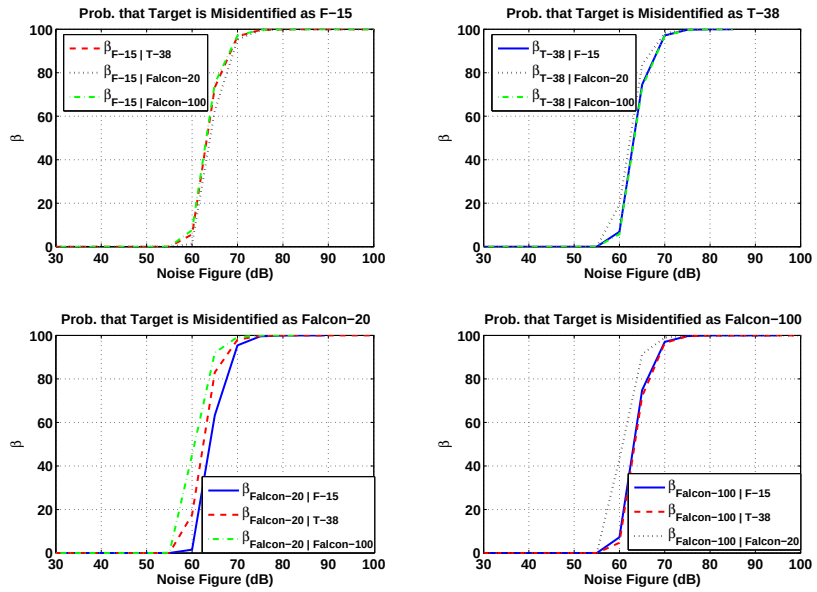

Figure 36: Probability of a type II error for straight-and-level trajectory #2 (hor. pol. transmitter)

Table 21: Comparison of relative entropy predictions and Monte Carlo results: straight-and-level trajectory #2, using the horizontally polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	55	60
Noise Figure at which Algorithm Breaks	70	70
Most Likely Aircraft to be Misidentified as the F-15	all	all
Most Likely Aircraft to be Misidentified as the T-38A	all	all
Most Likely Aircraft to be Misidentified as the Falcon-20	Falcon-100	Falcon-100
Most Likely Aircraft to be Misidentified as the Falcon-100	Falcon-20	Falcon-20

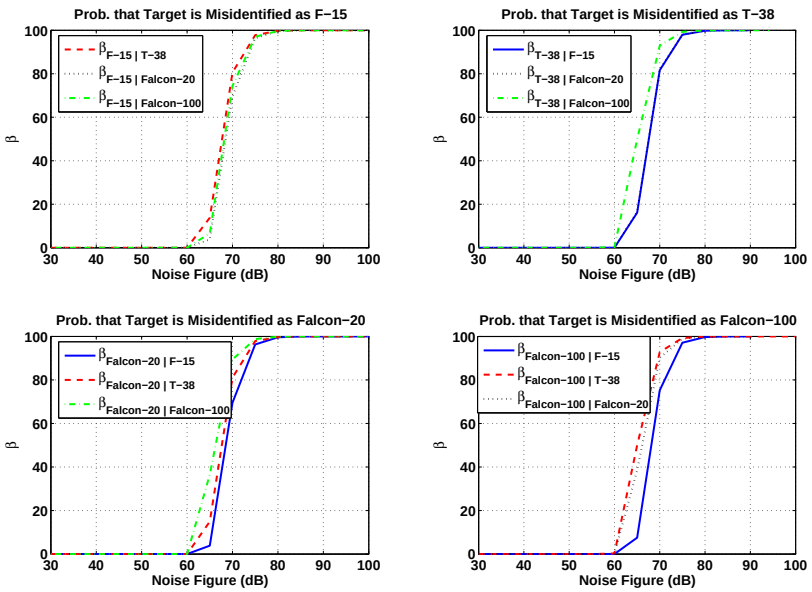

Figure 37: Probability of a type II error for banked turn trajectory (hor. pol. transmitter)

Table 22: Comparison of relative entropy predictions and Monte Carlo results: Banked turn trajectory, using the horizontally polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	60	65
Noise Figure at which Algorithm Breaks	75	75
Most Likely Aircraft to be Misidentified as the F-15	all	all
Most Likely Aircraft to be Misidentified as the T-38A	Falcon-100	Falcon-100
Most Likely Aircraft to be Misidentified as the Falcon-20	Falcon-100	Falcon-100
Most Likely Aircraft to be Misidentified as the Falcon-100	T-38A and Falcon-20	T-38A and Falcon-20

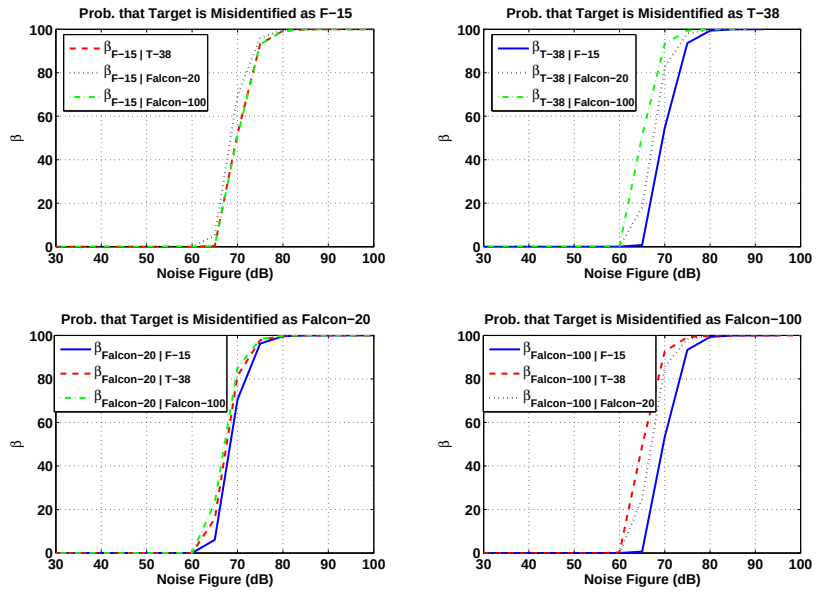

Figure 38: Probability of a type II error for Edwards trajectory with true orientation (hor. pol. transmitter)

Table 23: Comparison of relative entropy predictions and Monte Carlo results: Edwards trajectory with true orientation, using the horizontally polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	60	65
Noise Figure at which Algorithm Breaks	75	75
Most Likely Aircraft to be Misidentified as the F-15	Falcon-20	Falcon-20
Most Likely Aircraft to be Misidentified as the T-38A	Falcon-100	Falcon-100
Most Likely Aircraft to be Misidentified as the Falcon-20	all	all
Most Likely Aircraft to be Misidentified as the Falcon-100	T-38A and Falcon-20	T-38A and Falcon-20

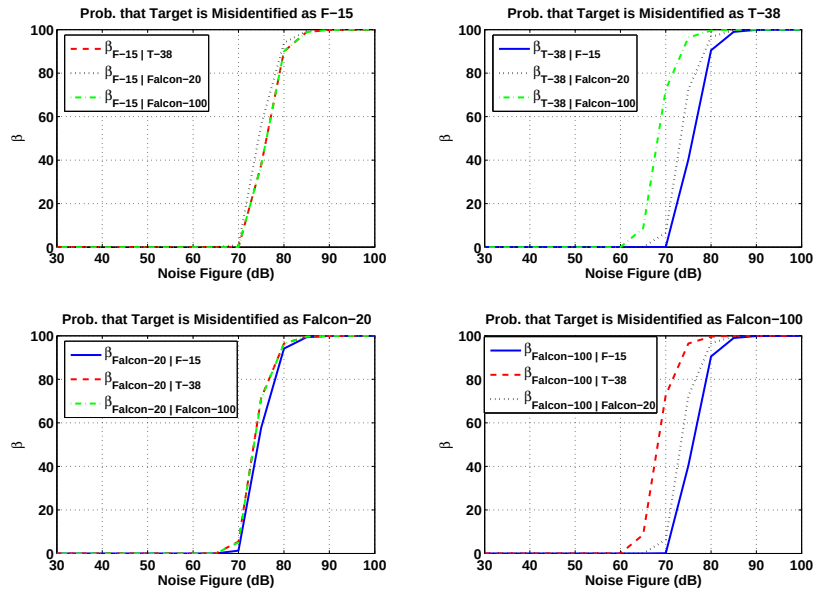

Figure 39: Probability of a type II error for Edwards trajectory with approximated orientation (hor. pol. transmitter)

Table 24: Comparison of relative entropy predictions and Monte Carlo results: Edwards trajectory with approximated orientation, using the horizontally polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	5	65
Noise Figure at which Algorithm Breaks	5	75
Most Likely Aircraft to be Misidentified as the F-15	all	all
Most Likely Aircraft to be Misidentified as the T-38A	F-15	Falcon-100
Most Likely Aircraft to be Misidentified as the Falcon-20	all	all
Most Likely Aircraft to be Misidentified as the Falcon-100	all	T-38A

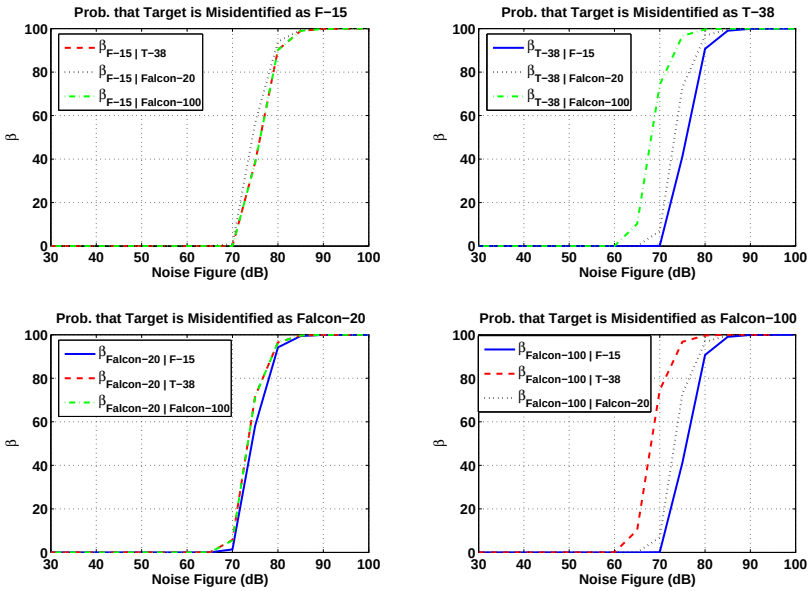

Figure 40: Probability of a type II error for Edwards trajectory with approximated orientation and biased position estimates (hor. pol. transmitter)

Table 25: Comparison of relative entropy predictions and Monte Carlo results: Edwards trajectory with approximated orientation and biased position measurements, using the horizontally polarized transmitter.

	Monte Carlo Results	Relative Entropy Predictions
Noise Figure at which 1st Mistake is Made	5	65
Noise Figure at which Algorithm Breaks	5	75
Most Likely Aircraft to be Misidentified as the F-15	all	all
Most Likely Aircraft to be Misidentified as the T-38A	F-15	Falcon-100
Most Likely Aircraft to be Misidentified as the Falcon-20	all	all
Most Likely Aircraft to be Misidentified as the Falcon-100	all	T-38A

5.2 Estimating Performance Under the Bayesian Framework Via the Chernoff Information

5.2.1 Derivation of the Chernoff Information Between Two Rician Densities

The Chernoff information between two densities, $p(x)$ and $q(x)$ is given by

$$C\left(p(x), q(x)\right) = -min_{0\leq\lambda\leq 1}\left\{\mu(\lambda)\right\}, \tag{84}$$

where $\mu(\lambda)$ is

$$\mu(\lambda) = \ln\left[\int_x q(x)^{\lambda} p(x)^{1-\lambda} dx\right]. \tag{85}$$

Substituting (57) and (58) into (85) results in

$$\mu(\lambda) = \ln\left\{\int_0^{\infty}\left[\frac{x}{\sigma^2}e^{\frac{-(x^2+s_q^2)}{2\sigma^2}} I_0\left(\frac{xs_q}{\sigma^2}\right)\right]^{\lambda}\left[\frac{x}{\sigma^2}e^{\frac{-(x^2+s_p^2)}{2\sigma^2}} I_0\left(\frac{xs_p}{\sigma^2}\right)\right]^{1-\lambda} dx\right\}. \tag{86}$$

Sections 5.2.2 and 5.2.3 seek to evaluate this expression.

5.2.2 A Normal Approximation for the Chernoff Information Between Two Rician Densities

The Chernoff information between Rician densities is often approximated by the Chernoff information between two Normal distributions with different means and variances. This is accomplished by substituting $p(x) \sim N(\mu_p, v_p)$ and $q(x) \sim N(\mu_q, v_q)$ into (85), resulting in

$$\mu(\lambda) = \ln\left\{\int_{-\infty}^{\infty}\left[\frac{1}{\sqrt{2\pi v_q^2}}e^{\frac{-(x-\mu_q)^2}{2v_q^2}}\right]^{\lambda}\left[\frac{1}{\sqrt{2\pi v_p^2}}e^{\frac{-(x-\mu_p)^2}{2v_p^2}}\right]^{1-\lambda} dx\right\}. \tag{87}$$

Expanding this results in,

$$\mu(\lambda) = \ln\left\{\int_{-\infty}^{\infty}\left[\frac{2\pi v_p^2}{2\pi v_q^2}\right]^{\frac{\lambda}{2}}\frac{1}{\sqrt{2\pi v_q^2}}e^{\frac{-(x-\mu_q)^2\lambda}{2v_q^2}}e^{\frac{(x-\mu_p)^2(\lambda-1)}{2v_p^2}} dx\right\}, \tag{88}$$

which reduces to

$$\mu(\lambda) = \frac{\lambda}{2}\ln\left\{\frac{v_p^2}{v_q^2}\right\} + \ln\left\{\int_{-\infty}^{\infty}\frac{1}{\sqrt{2\pi v_p^2}}e^{\frac{(x-\mu_p)^2(\lambda-1)v_q^2-(x-\mu_q)^2v_p^2\lambda}{2v_p^2v_q^2}} dx\right\}. \tag{89}$$

Now focus on the power of the exponential term. Completing the square transforms

$$\frac{(x-\mu_p)^2 v_q^2(\lambda-1)-(x-\mu_q)^2 v_p^2\lambda}{2v_p^2 v_q^2} \tag{90}$$

into

$$\frac{(v_q^2(\lambda-1)-v_p^2\lambda)\left(x+\frac{\mu_q v_p^2\lambda-\mu_p v_q^2(\lambda-1)}{v_q^2(\lambda-1)-v_p^2\lambda}\right)^2}{2v_p^2 v_q^2}-\frac{(\mu_q v_p^2\lambda-\mu_p v_q^2(\lambda-1))^2}{2v_p^2 v_q^2(v_q^2(\lambda-1)-v_p^2\lambda)}. \tag{91}$$

Substituting this expression back into (89) results in

$$\mu(\lambda)=\frac{\lambda}{2}\ln\left\{\frac{v_p^2}{v_q^2}\right\}+\frac{\mu_p^2 v_q^2(\lambda-1)-\mu_q^2 v_p^2\lambda}{2v_p^2 v_q^2}+\ln\left\{\int_{-\infty}^{\infty}\frac{1}{\sqrt{2\pi v_p^2}}e^{\frac{(v_q^2(\lambda-1)-v_p^2\lambda)\left(x+\frac{\mu_q v_p^2\lambda-\mu_p v_q^2(\lambda-1)}{v_q^2(\lambda-1)-v_p^2\lambda}\right)^2}{2v_p^2 v_q^2}}dx\right\}. \tag{92}$$

Now, the integral in (92) is rewritten as

$$\frac{1}{2}\ln\left\{\int_{-\infty}^{\infty}\frac{1}{\sqrt{2\pi v_p^2}}e^{\frac{-\left(x-\frac{\mu_p v_q^2(\lambda-1)-\mu_q v_p^2\lambda}{v_q^2(\lambda-1)-v_p^2\lambda}\right)^2}{\frac{2v_p^2 v_q^2}{v_p^2\lambda-v_q^2(\lambda-1)}}}\right\}. \tag{93}$$

Scaling appropriately reduces this integral to a Normal pdf, resulting in

$$\frac{1}{2}\ln\left[\frac{v_q^2}{v_p^2\lambda-v_q^2(\lambda-1)}\right]. \tag{94}$$

Thus, $\mu(\lambda)$ reduces to

$$\mu(\lambda)=\frac{\lambda}{2}\ln\left\{\frac{v_p^2}{v_q^2}\right\}+\frac{\mu_p^2 v_q^2(\lambda-1)-\mu_q^2 v_p^2\lambda}{2v_p^2 v_q^2}-\frac{(\mu_q^2 v_p^2\lambda-\mu_p^2 v_q^2(\lambda-1))^2}{2v_p^2 v_q^2(v_q^2(\lambda-1)-v_p^2\lambda)}+\frac{1}{2}\ln\left[\frac{v_q^2}{v_p^2\lambda-v_q^2(\lambda-1)}\right]. \tag{95}$$

The Normal approximation is completed by setting μ_p, μ_q, v_p^2, and v_q^2 as suggested in (65) through (68), and substituting the resulting $\mu(\lambda)$ into (84).

5.2.3 Derivation of a Closed-Form Approximation for the Chernoff Information Between Two Rician Densities

The Bessel functions in (86) render an analytic evaluation of the integral quite difficult. However, an approximation can be made by applying the Laplace method to the integral. Doing so transforms (86) into

$$\mu(\lambda) = \ln\left\{ e^{\frac{\lambda(s_p^2 - s_q^2)}{2\sigma^2}} \int_0^\infty e^{\ln\left(\frac{x}{\sigma^2}\right) - \frac{(x^2+s_p^2)}{2\sigma^2} - (\lambda-1)\ln\left[I_0\left(\frac{xs_p}{\sigma^2}\right)\right] + \lambda\ln\left[I_0\left(\frac{xs_q}{\sigma^2}\right)\right]} dx \right\}, \tag{96}$$

which is equivalent to

$$\mu(\lambda) = \ln\left[e^{\frac{\lambda(s_p^2 - s_q^2)}{2\sigma^2}} \int_0^\infty e^{h(x,\lambda)} dx \right], \tag{97}$$

where

$$h(x,\lambda) = \ln\left(\frac{x}{\sigma^2}\right) - \frac{(x^2+s_p^2)}{2\sigma^2} - (\lambda-1)\ln\left[I_0\left(\frac{xs_p}{\sigma^2}\right)\right] + \lambda\ln\left[I_0\left(\frac{xs_q}{\sigma^2}\right)\right]. \tag{98}$$

Identifying a Gaussian form lets us approximate $\mu(\lambda)$ as

$$\mu(\lambda) \approx \frac{\lambda(s_p^2 - s_q^2)}{2\sigma^2} + h(\hat{x},\lambda) + \frac{1}{2}\ln\left(\frac{-2\pi}{h''(\hat{x},\lambda)}\right), \tag{99}$$

where $\hat{x}$ is the value of x found by setting the derivative of $h(x,\lambda)$ equal to zero.[2] This results in two solutions. However, it is trivial to show that, given the limits of integration of the Rician density, the only valid solution is

$$\hat{x} = \frac{1}{2}\lambda(s_q - s_p) + \frac{1}{2}s_p + \frac{1}{2}\sqrt{\lambda^2(s_q^2 + s_p^2 - 2s_p s_q) + 2\lambda s_p s_q - 2\lambda s_p^2 + s_p^2 + 4\sigma^2}. \tag{100}$$

The second derivative of $h(x,\lambda)$ is given by

$$h''(x,\lambda) = -\frac{1}{\sigma^2} - \frac{1}{x^2} - \left(\frac{\lambda s_q^2}{\sigma^4}\right)\left[\frac{I_1^2\left(\frac{xs_q}{\sigma^2}\right)}{I_0^2\left(\frac{xs_q}{\sigma^2}\right)}\right] + \left(\frac{\lambda s_q^2}{2\sigma^4}\right)\left[\frac{I_0\left(\frac{xs_q}{\sigma^2}\right) + I_2\left(\frac{xs_q}{\sigma^2}\right)}{I_0\left(\frac{xs_q}{\sigma^2}\right)}\right] \ldots. \tag{101}$$

[2] Much like the normal density, the Rician density is unimodal with exactly one global maximum and no global minima (except at $\pm\infty$). Since $h(x,\lambda)$ is merely a scaled version of a Rician density, it is clear that $\hat{x}$ maximizes (rather than minimizes) $h(x,\lambda)$.

$$\ldots - \left(\frac{(\lambda-1)s_p^2}{\sigma^4}\right)\left[\frac{I_1^2\left(\frac{xs_p}{\sigma^2}\right)}{I_0^2\left(\frac{xs_p}{\sigma^2}\right)}\right] + \left(\frac{(\lambda-1)s_p^2}{2\sigma^4}\right)\left[\frac{I_0\left(\frac{xs_p}{\sigma^2}\right) + I_2\left(\frac{xs_p}{\sigma^2}\right)}{I_0\left(\frac{xs_p}{\sigma^2}\right)}\right]. \tag{102}$$

Thus, the probability of error in a binary Bayesian hypothesis test is approximated by

$$P_E \approx e^{-min_{0\leq\lambda\leq 1}\left[\frac{\lambda(s_p^2-s_q^2)}{2\sigma^2}+h(\hat{x},\lambda)+\frac{1}{2}\ln\left(\frac{-2\pi}{h''(\hat{x},\lambda)}\right)\right]}. \tag{103}$$

To compare the Normal approximation from Section 5.2.2 to the closed-form approximation developed in this section, the s_q parameter is swept over a range of values while the s_p and σ^2 parameters are held constant. Figure 41 shows the values of μ and the probability of error that result if s_p is set to 10, σ^2 is set to 4, and s_q is swept from 0 to 20. The closed-form approximation derived in this section matches the numerically approximated result with higher accuracy than does the Normal approximation, which deviates significantly for small values of s_q. This result is typical. The Normal approximation breaks down for small values of s_q, relative to s_p and σ^2.

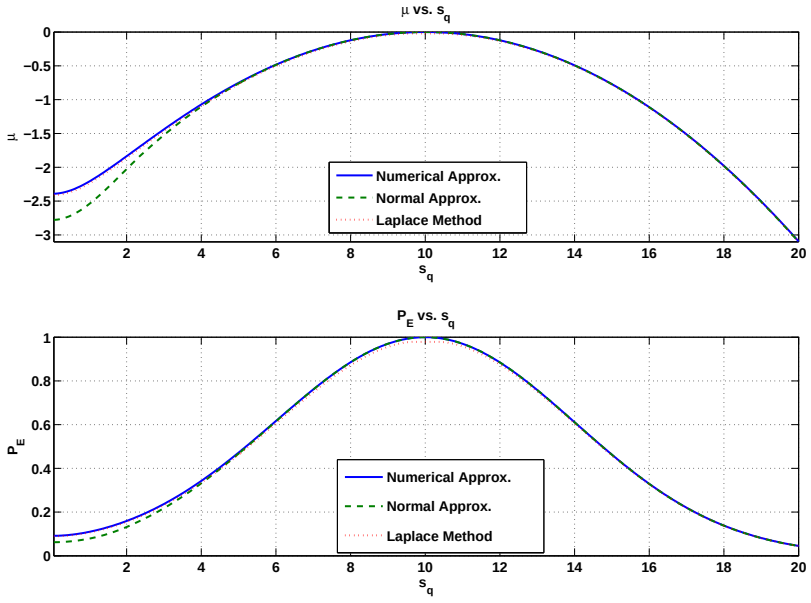

Figure 41: Comparison of methods for computing μ and P_E, if $s_p = 10$, $\sigma^2 = 4$, and s_q is swept from 0 to 20

5.2.4 Applying the the Closed-Form Approximation for the Chernoff Information to the ATR Algorithm

The Chernoff information is then applied to the ATR runs to approximate the probability of error for each possible aircraft pairing. The predictions obtained from the Chernoff information are compared to the Monte Carlo results from Chapter 4. As is the case with the relative entropy, the Chernoff information assumes that the true orientation angles are used by the ATR algorithm. Thus, it is not surprising that the Chernoff information predictions are a good match with the Monte Carlo trials as long as the orientation angles estimated by the coordinated flight model are close to the truth.

The probability of error curves predicted by the Chernoff information for the straight-and-level maneuvers (using the vertically polarized transmitter from Section 3.1) are shown in Figures 42 and 43. These results are then compared to those from Chapter 4 in Tables 26 and 27. Since the coordinated flight model has no trouble estimating the orientations when the aircraft execute straight-and-level maneuvers, the Monte Carlo trials corroborate the Chernoff information predictions quite well.

Figure 44 shows the probability of error curves predicted by the Chernoff information when the aircraft execute the banked turn trajectory. A comparison is made in Table 28 to the Monte Carlo results from Section 4.1.3. Even though the coordinated flight model must estimate a non-zero roll angle, its predictions are still very close to the truth. Thus, the predictions from the Chernoff information are corroborated by Monte Carlo results. Note that both the Monte Carlo trials and Chernoff information predictions indicate a near tie between the pairs of aircraft least likely to be confused by the algorithm.

The probability of error curves predicted by the Chernoff information for the Edwards trajectory, with true orientation angles, approximated orientation angles, and approximated orientation angles with biased position measurements, are shown in Figures 45 through 47, respectively. Tables 29 through 31 compare the predictions to the Monte Carlo results. As in Section 5.1.4, the predictions are a good match for the Monte Carlo data, as long as the estimated orientation angles are similar to the truth. When the estimated orientation angles are significantly different from the truth, the performance predictions are overly optimistic.

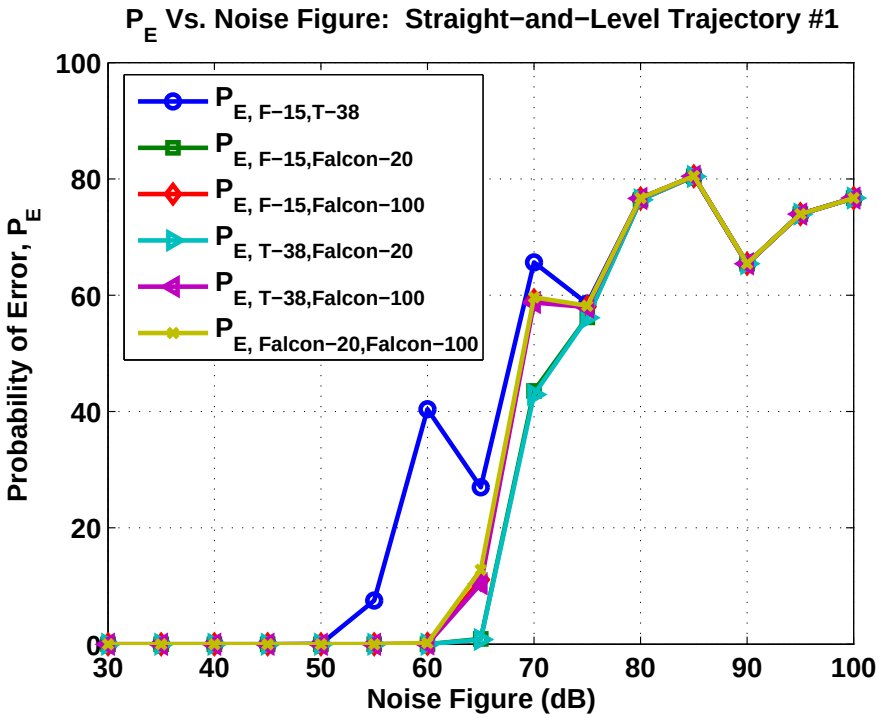

Figure 42: Predicted probability of error in straight-and-level trajectory #1 (vert. pol. transmitter)

Table 26: Comparison of Chernoff information predictions and Monte Carlo results: straight-and-level trajectory #1, using the vertically polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	55	55
Noise Figure at which Algorithm Breaks	80	75
Most Likely Aircraft to be Swapped	F-15 and T-38A	F-15 and T-38A
Leas Likely Aircraft to be Swapped	Falcon-20 and T-38A	Falcon-20 and T-38A

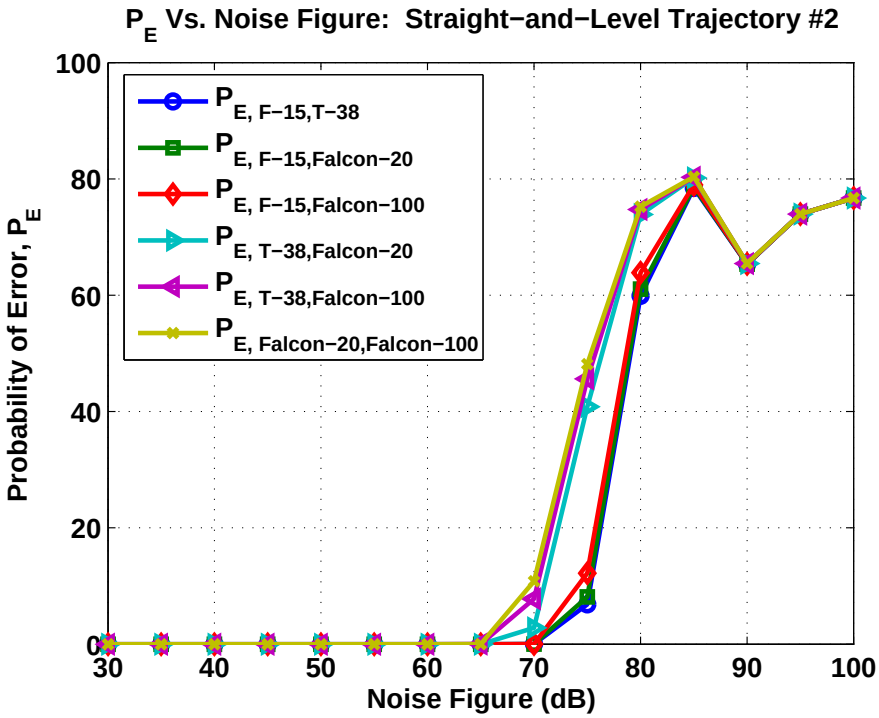

Figure 43: Predicted probability of error in straight-and-level trajectory #2 (vert. pol. transmitter)

Table 27: Comparison of Chernoff information predictions and Monte Carlo results: straight-and-level trajectory #2, using the vertically polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	70	70
Noise Figure at which Algorithm Breaks	90	85
Most Likely Aircraft to be Swapped	Falcon-20 and Falcon-100	Falcon-20 and Falcon-100
Leas Likely Aircraft to be Swapped	F-15 and Falcon-20, F-15 and T-38A	F-15 and Falcon-20, F-15 and T-38A

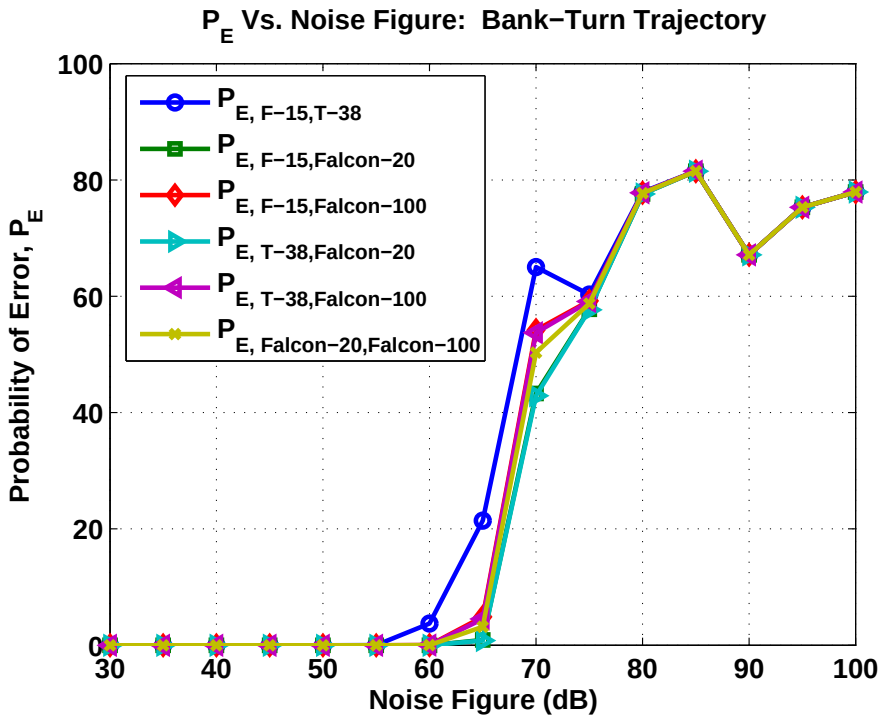

Figure 44: Predicted probability of error in banked turn trajectory (vert. pol. transmitter)

Table 28: Comparison of Chernoff information predictions and Monte Carlo results: Banked turn trajectory, using the vertically polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	65	60
Noise Figure at which Algorithm Breaks	80	75
Most Likely Aircraft to be Swapped	F-15 and T-38A	F-15 and T-38A
Leas Likely Aircraft to be Swapped	F-15 and Falcon-20, Falcon-20 and T-38A	F-15 and Falcon-20, Falcon-20 and T-38A

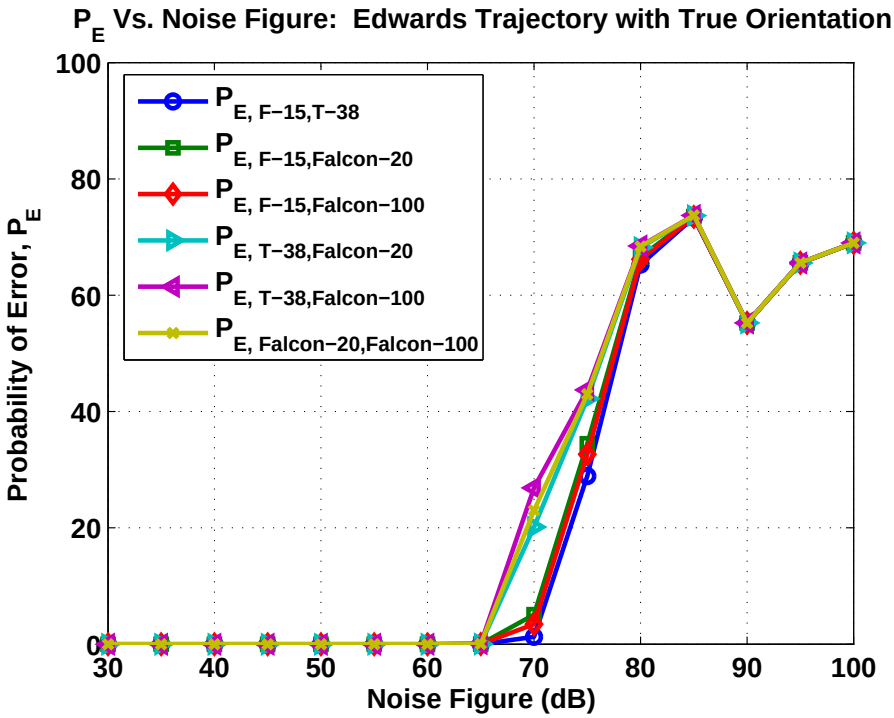

Figure 45: Predicted probability of error in Edwards trajectory with true orientation (vert. pol. transmitter)

Table 29: Comparison of Chernoff information predictions and Monte Carlo results: Edwards trajectory with true orientation, using the vertically polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	70	70
Noise Figure at which Algorithm Breaks	90	80
Most Likely Aircraft to be Swapped	T-38A and Falcon-100	T-38A and Falcon-100
Leas Likely Aircraft to be Swapped	F-15 and T-38A	F-15 and T-38A

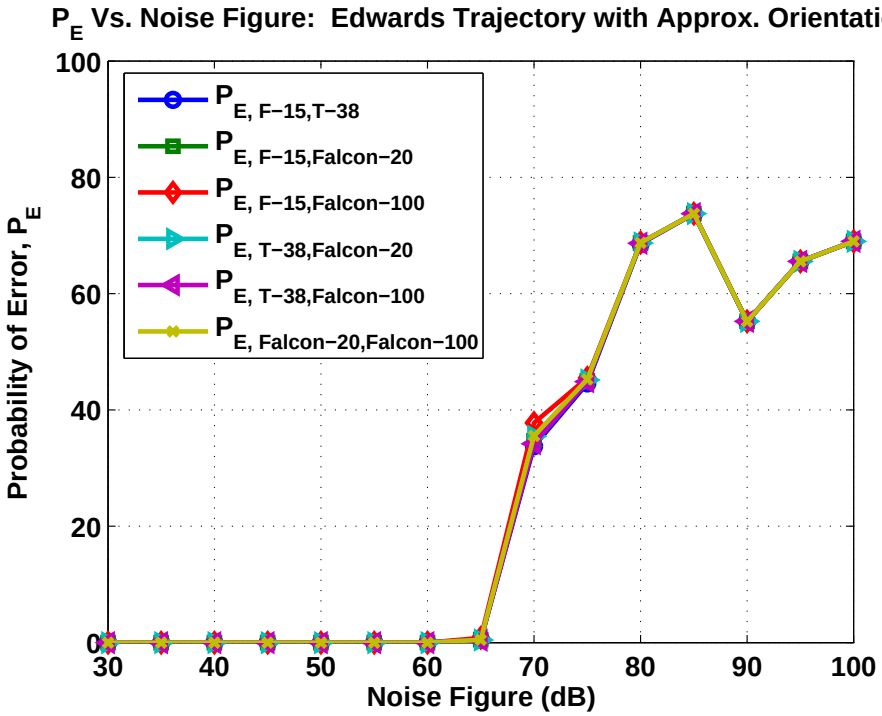

Figure 46: Predicted probability of error in Edwards trajectory with approximated orientation (vert. pol. transmitter)

Table 30: Comparison of Chernoff information predictions and Monte Carlo results: Edwards trajectory with approximated orientation, using the vertically polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	5	70
Noise Figure at which Algorithm Breaks	5	80
Most Likely Aircraft to be Swapped	all	all
Leas Likely Aircraft to be Swapped	all	all

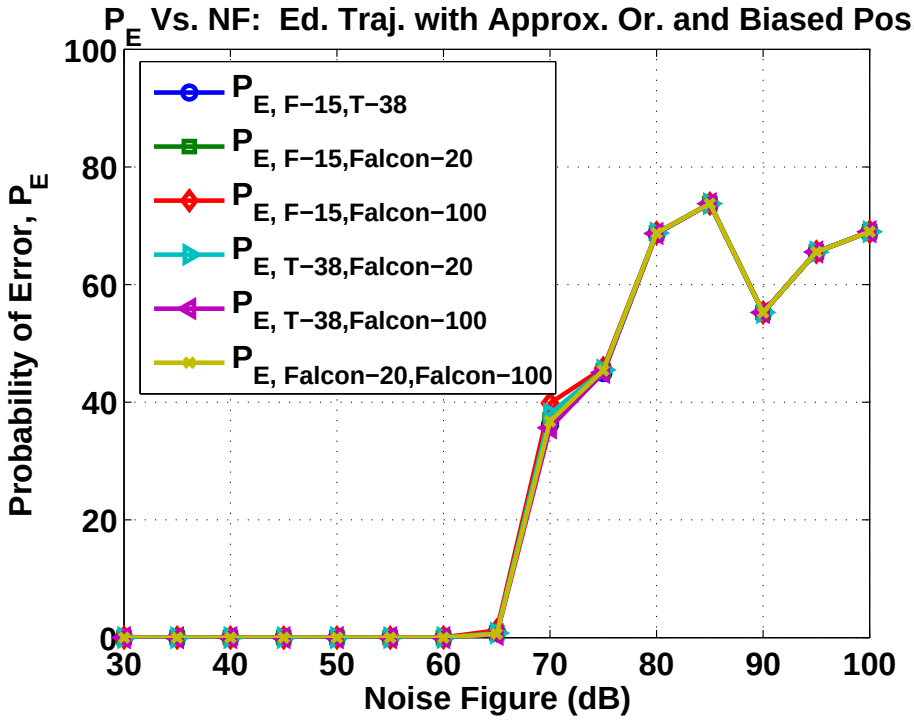

Figure 47: Predicted probability of error in Edwards trajectory with approximated orientation and biased position estimates (vert. pol. transmitter)

Table 31: Comparison of Chernoff information predictions and Monte Carlo results: Edwards trajectory with approximated orientation and biased position measurements, using the vertically polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	5	70
Noise Figure at which Algorithm Breaks	5	80
Most Likely Aircraft to be Swapped	all	all
Leas Likely Aircraft to be Swapped	all	all

The Chernoff information is then used to predict the ATR algorithm's performance when the horizontally polarized transmitter from Section 3.1 is used. Although the trends are slightly different than they were when the vertically polarized transmitter was exploited, the Chernoff information predictions still match the Monte Carlo results with a high level of accuracy. Results for the first straight-and-level trajectory are shown in Figure 48 and Table 32. Similar results appear in Figure 49 and Table 33 for the second straight-and-level trajectory. Figure 50 and Table 34 correspond to the case in which the aircraft execute the banked turn trajectory. Finally, results using the Edwards trajectory with true orientation angles, approximated orientation angles, and approximated orientation angles with biased position measurements appear in Figures 51 through 53 and Tables 35 through 37.

As before, the only cases with discrepancies are those in which the aircraft execute the Edwards maneuver and the orientation angles are estimated. In these two cases, the orientation estimates from the coordinated flight model are not sufficiently close to the truth to result in predictable performance in the ATR algorithm.

Once again, it should be noted that the only reason for using both horizontally and vertically polarized antennas is to provide a baseline for the work in Chapter 6. It is not intended to suggest that one polarization provides better results than the other, and the reader should also bear in mind that the transmitters are at different locations.

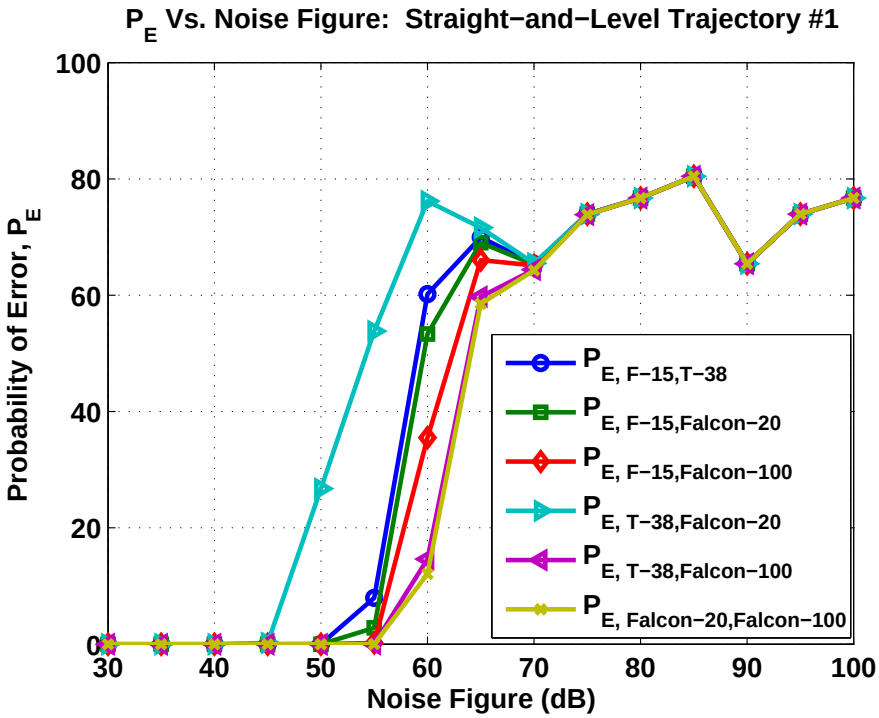

Figure 48: Predicted probability of error in straight-and-level trajectory #1 (hor. pol. transmitter)

Table 32: Comparison of Chernoff information predictions and Monte Carlo results: straight-and-level trajectory #1, using the horizontally polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	50	50
Noise Figure at which Algorithm Breaks	70	70
Most Likely Aircraft to be Swapped	T-38A and Falcon-20	T-38A and Falcon-20
Leas Likely Aircraft to be Swapped	Falcon-20 and Falcon-100	Falcon-20 and Falcon-100

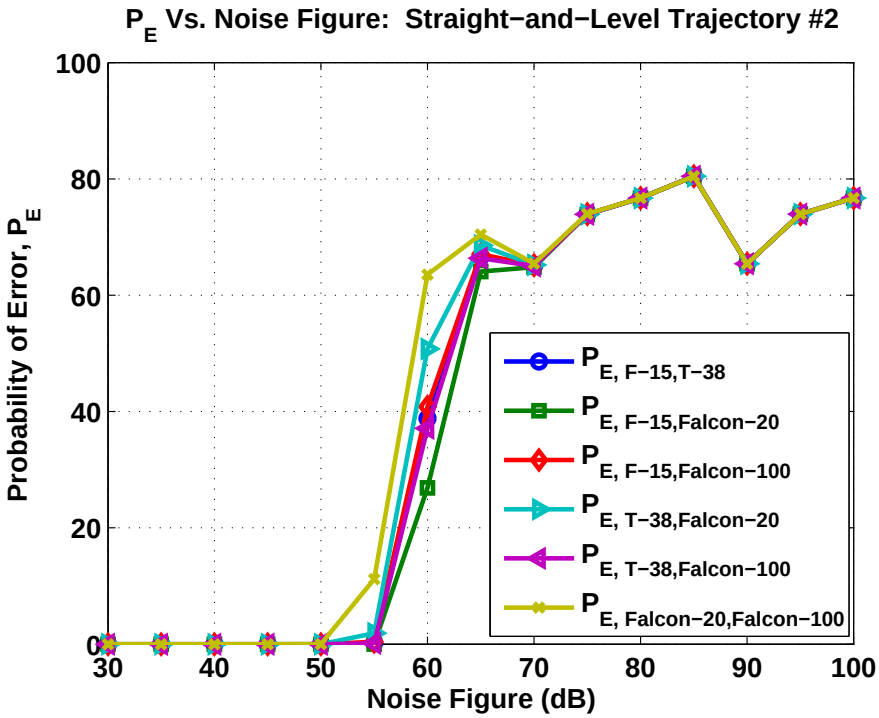

Figure 49: Predicted probability of error in straight-and-level trajectory #2 (hor. pol. transmitter)

Table 33: Comparison of Chernoff information predictions and Monte Carlo results: straight-and-level trajectory #2, using the horizontally polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	55	55
Noise Figure at which Algorithm Breaks	70	70
Most Likely Aircraft to be Swapped	Falcon-20 and Falcon-100	Falcon-20 and Falcon-100
Leas Likely Aircraft to be Swapped	F-15 and Falcon-20	F-15 and Falcon-20

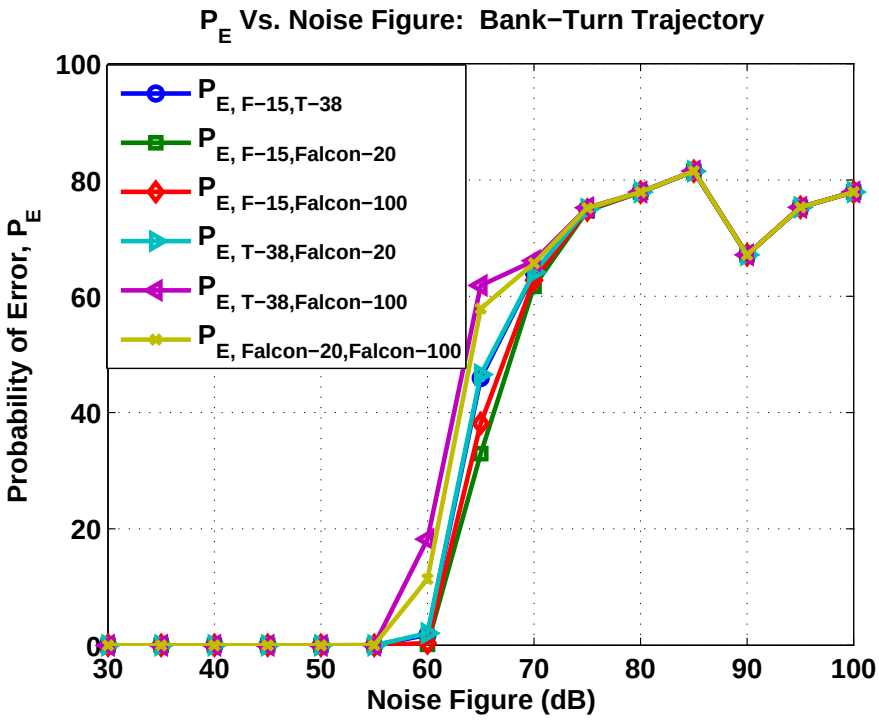

Figure 50: Predicted probability of error in banked turn trajectory (hor. pol. transmitter)

Table 34: Comparison of Chernoff information predictions and Monte Carlo results: Banked turn trajectory, using the horizontally polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	60	60
Noise Figure at which Algorithm Breaks	75	75
Most Likely Aircraft to be Swapped	T-38A and Falcon-100	T-38A and Falcon-100
Leas Likely Aircraft to be Swapped	F-15 and Falcon-20	F-15 and Falcon-20

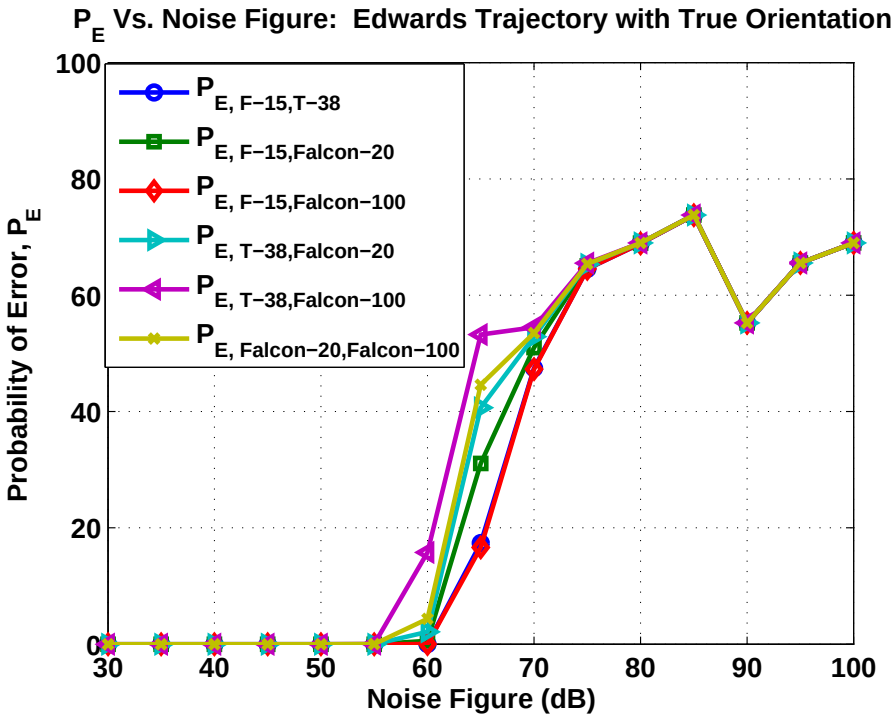

Figure 51: Predicted probability of error in Edwards trajectory with true orientation (hor. pol. transmitter)

Table 35: Comparison of Chernoff information predictions and Monte Carlo results: Edwards trajectory with true orientation, using the horizontally polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	60	60
Noise Figure at which Algorithm Breaks	75	75
Most Likely Aircraft to be Swapped	T-38A and Falcon-100	T-38A and Falcon-100
Leas Likely Aircraft to be Swapped	F-15 and Falcon-100	F-15 and Falcon-100

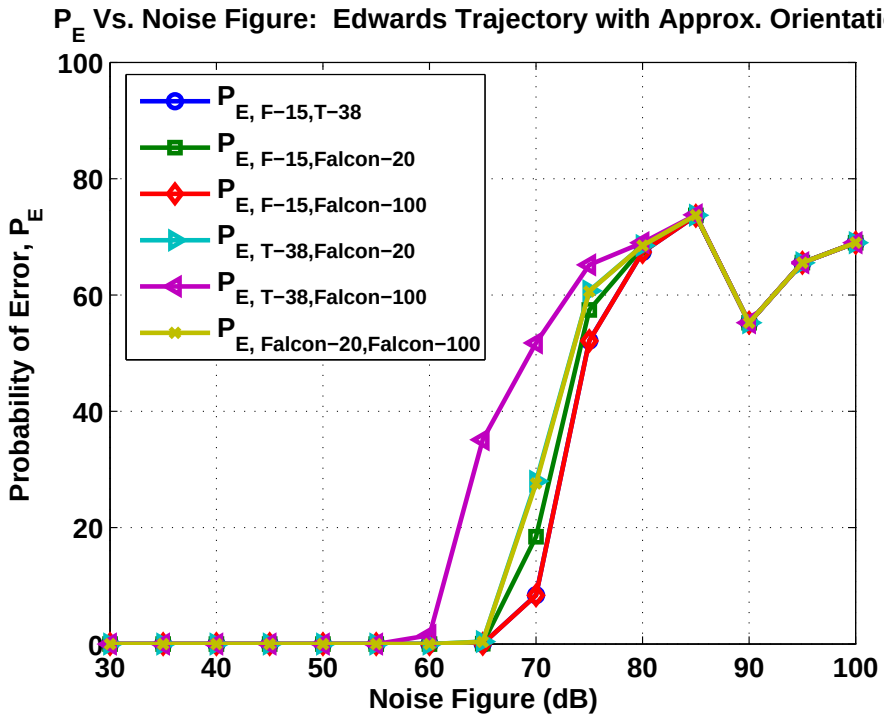

Figure 52: Predicted probability of error in Edwards trajectory with approximated orientation (hor. pol. transmitter)

Table 36: Comparison of Chernoff information predictions and Monte Carlo results: Edwards trajectory with approximated orientation, using the horizontally polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	5	60
Noise Figure at which Algorithm Breaks	5	80
Most Likely Aircraft to be Swapped	all	T-38A and Falcon-100
Leas Likely Aircraft to be Swapped	all	F-15 and Falcon-100

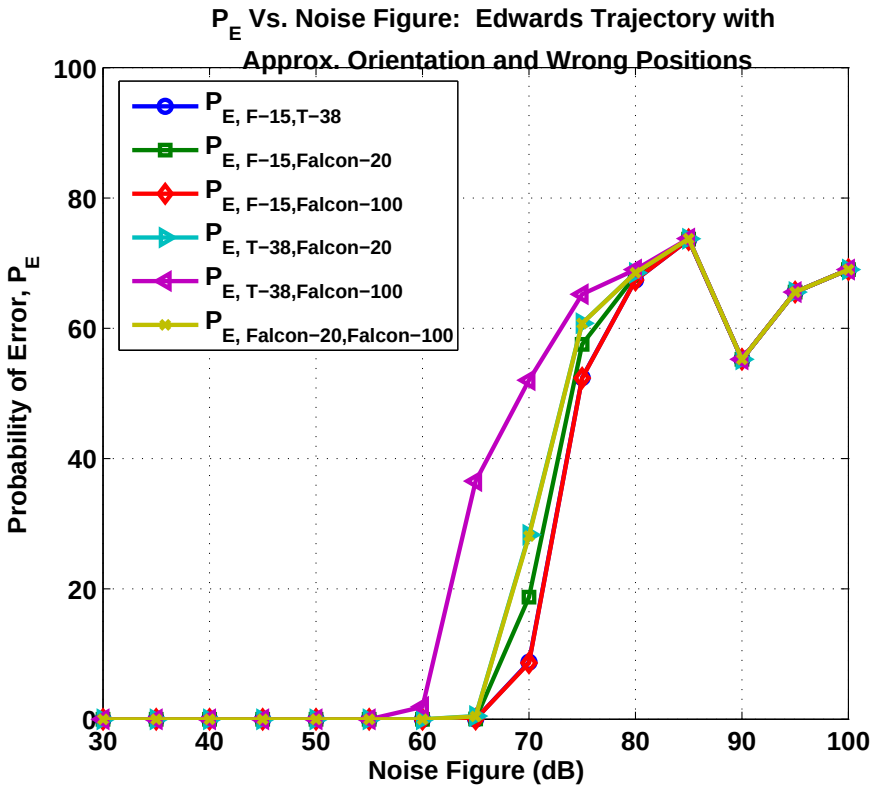

Figure 53: Predicted probability of error in Edwards trajectory with approximated orientation and biased position estimates (hor. pol. transmitter)

Table 37: Comparison of Chernoff information predictions and Monte Carlo results: Edwards trajectory with approximated orientation and biased position measurements, using the horizontally polarized transmitter.

	Monte Carlo Results	Chernoff Information Predictions
Noise Figure at which 1st Mistake is Made	5	60
Noise Figure at which Algorithm Breaks	5	80
Most Likely Aircraft to be Swapped	all	T-38A and Falcon-100
Leas Likely Aircraft to be Swapped	all	F-15 and Falcon-100

5.2.5 An Application of the Chernoff Information

Section 5.2.4 demonstrates that the approximation of the algorithm's performance using the Chernoff information is very similar to the performance obtained using Monte Carlo trials, as long as the estimated orientation angles are close to the truth. As such, the Chernoff information can confidently be used to address questions that would be cumbersome to address via Monte Carlo trials. For example, a useful piece of information is the length of time that the aircraft must be tracked in order to identify it with a desired probability of error. The number of Monte Carlo trials required to address this question is staggering, as a complete set of trials would be required for each period of time tested. However, this problem is easily addressed using the Chernoff information. Note that a noise figure of 45 dB is used in this section, as that is the maximum noise figure anticipated in a real system.

Figure 54 shows the probability of error as a function of the length of time that the target is tracked while executing the first straight-and-level trajectory, using the vertically polarized transmitter. Similar results are given in Figures 55 through 56 for the second straight-and-level trajectory and the banked turn trajectory, respectively.

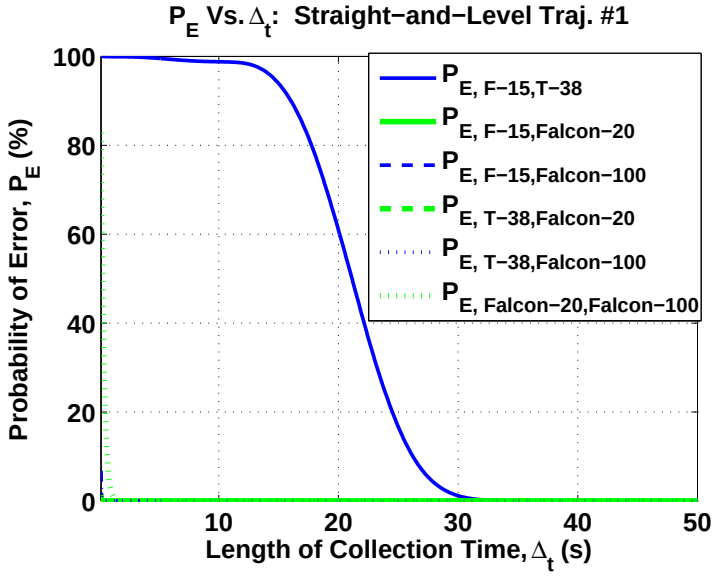

Figure 54: Probability of error vs. time: Straight-and-level trajectory #1, vertically polarized transmitter

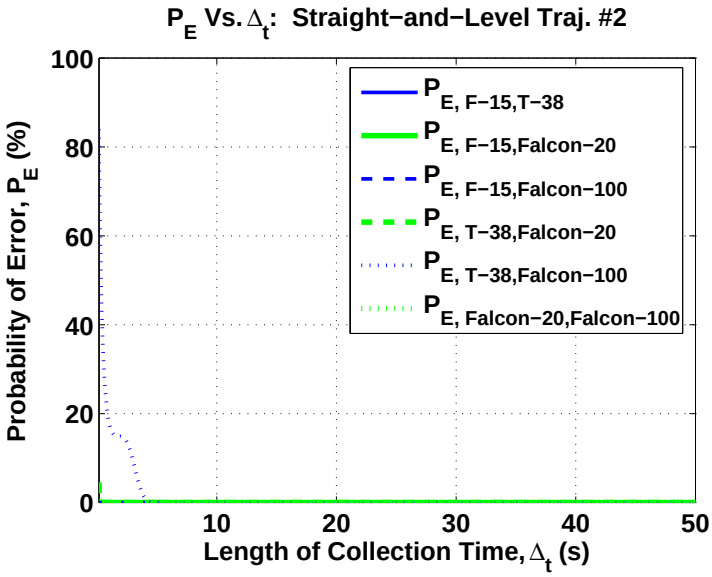

Figure 55: Probability of error vs. time: Straight-and-level trajectory #2, vertically polarized transmitter

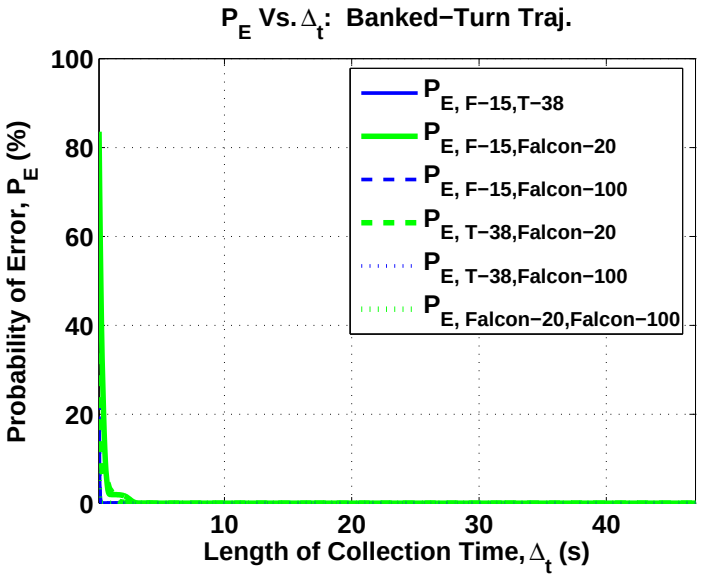

Figure 56: Probability of error vs. time: Banked turn trajectory, vertically polarized transmitter

Several points are worth stating. Consider the first straight-and-level trajectory. The most difficult comparison for the ATR algorithm is between the F-15 and T-38. This matches the result from Section 5.2.4. Another noteworthy point is that the probability of switching any pair of aircraft drops below 5% 22.5 seconds after the start of the trajectory. Using the second straight-and-level trajectory, the ATR algorithm can correctly distinguish between all pairs of aircraft with a probability of error below 5% within 4 seconds, instead of 22.5. This improvement is attributed to the broader range of aspect angles that are presented to the receiver in the second straight-and-level maneuver. The ATR algorithm performance improves further still against aircraft executing the banked turn maneuver. In this case, the algorithm is able to correctly distinguish between all possible pairs of aircraft with a probability of error below 5% within 1 second of tracking.

These tests are then repeated using the horizontally polarized transmitter from Section 3.1. As in Section 4.2, the results are somewhat worse than when the vertically polarized transmitter is exploited. The probability of error curves are shown for the three trajectories in Figures 57 through 59. Now, the radar must maintain a track for 27.5 seconds in order to correctly discern between each pair of targets executing the first straight-and-level maneuver. This drops to 12.5 seconds when the aircraft executes the second straight-and-level trajectory, and 7 seconds when it executes the banked turn maneuver.

Although the ATR performance is expected to degrade somewhat when all aircraft are considered, rather than when they are taken two at a time, two major conclusions can be taken from this work. First, regardless of the transmitter's polarization, if the radars in this study have tracked the aircraft for at least 28 seconds and the estimated orientation angles are good, then the probability of error in the ATR algorithm should be fairly low at the maximum anticipated noise level. Furthermore, depending on the trajectory and its orientation relative to the sensors, the length of time required to achieve this level of performance may be much smaller. Thus, tracking for 28 seconds constitutes a worst-case scenario, in which the noise level is extremely high and the aircraft is presenting a very small range of aspect angles to the receiver.

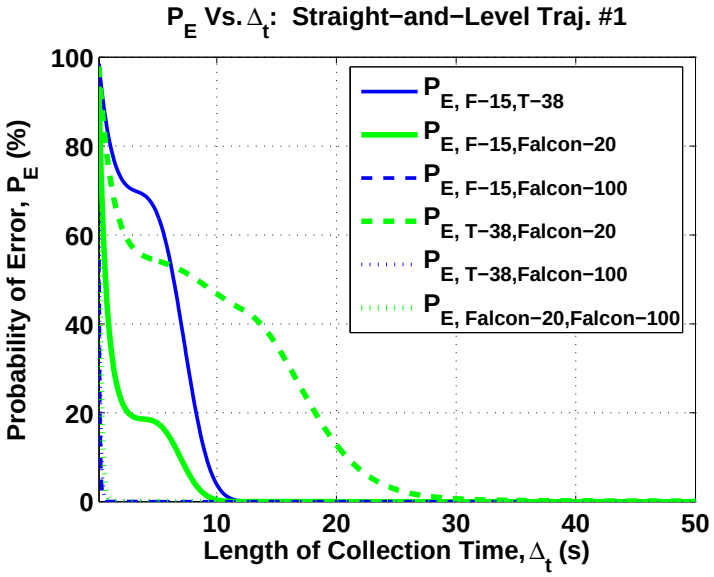

Figure 57: Probability of error vs. time: Straight-and-level trajectory #1, horizontally polarized transmitter

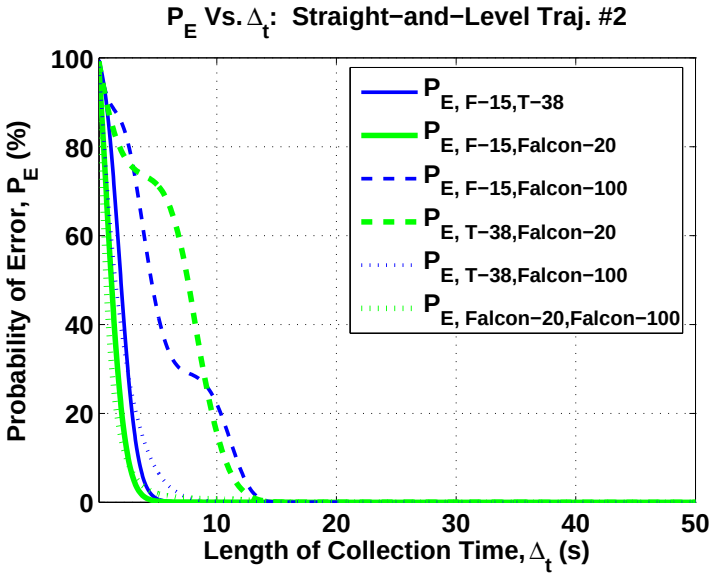

Figure 58: Probability of error vs. time: Straight-and-level trajectory #2, horizontally polarized transmitter

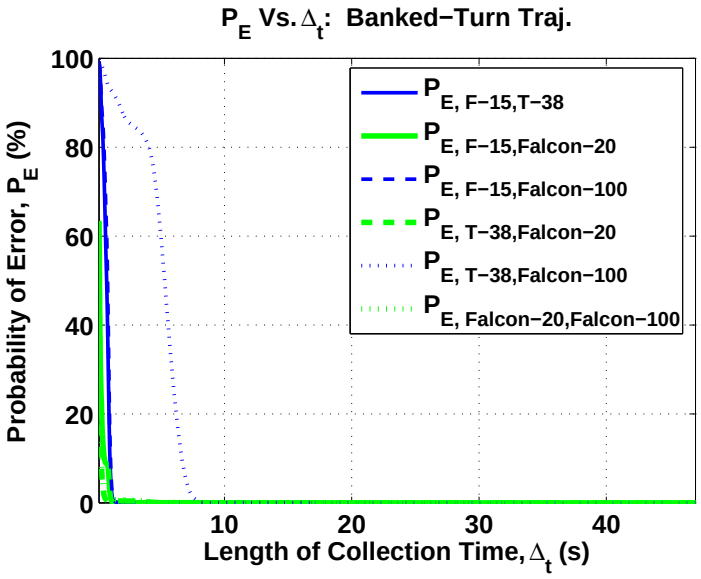

Figure 59: Probability of error vs. time: Banked turn trajectory, horizontally polarized transmitter

CHAPTER VI

A PASSIVE RADAR ATR ALGORITHM THAT EXPLORES THE SPACE OF ORIENTATIONS

6.1 Motivation

The ATR algorithm described in Chapter 3 performs quite well when aircraft travel in straight-and-level trajectories, or in banked turns with large radii of curvature. Since these two types of trajectories comprise the vast majority of trajectories likely to be witnessed by a passive radar algorithm when attempting to identify targets, the ATR algorithm's occasional failings when aircraft execute dog-fighting style maneuvers may not be of major concern in most likely scenarios.

Even so, modifications can make the simple ATR algorithm from Chapter 3 more robust. Since the algorithm's errors (at the anticipated noise levels of a real system) occur when the aircraft execute dog-fighting style maneuvers and are attributed to the error in the estimated roll angle of the aircraft, a better method for estimating the aircraft roll is needed.

6.2 Revisions to the ATR Algorithm

In response to these problems, the EKF that estimates aircraft orientation from velocity measurements is implemented in place of the simple coordinated flight model. The major advantage of this approach is that it provides not only an estimated state, but also a covariance. The revised ATR algorithm then uses the one-sigma error bars on the states (yaw, pitch, and roll) obtained from the covariance matrices to determine what set of orientations are feasible for a given maneuver. The feasible yaw, pitch, and roll sets are then sampled every two degrees,[1] beginning at the estimated state and working out until the one-sigma error bars are reached. All possible combinations of these sampled angles

[1] In this case, two degrees appears to provide a good tradeoff between computational complexity and accuracy.

are taken at each of the radar's dwell times. The total number of combinations depends on the size of the covariance, but it is not unusual to have one feasible yaw, three feasible pitch values, and five feasible rolls. Each of the combinations is compared to the profile collected by the receiver, using the scoring function in (51). The scores are computed for each aircraft in the target library, so that the orientation and aircraft type can be jointly estimated.

Once each combination of the sampled set of feasible orientations and aircraft types has been assigned a score, the set can be collectively thought of as comprising a trellis structure. The number of nodes in this trellis at time t_k is determined by the size of the covariance matrix of the EKF. The connections from each node at time t_k to time t_{k+1} are limited such that each of the orientation angles cannot change by more than two degrees;[2] this constraint is intended to prevent unrealistic orientation profiles, which may score well, from dominating the process. Having structured the problem of jointly estimating the target type and orientation as a trellis where each node has an assigned score, dynamic programming provides a logical solution for finding the optimal path through the trellis [9]. In doing so, it finds the optimal set of orientation angles for a particular aircraft type. The scores of each node along the optimal path are summed to find the total score for the aircraft type. This is repeated for all aircraft in the target library. The one with the largest overall score is deemed the winner.

This problem formulation extends quite naturally to multiple transmitters. If only one transmitter is available, the score of each node in the trellis is based only on that transmitter. When the passive radar system can exploit multiple transmitters, the scores for each transmitter are summed at each node in the trellis. Thus, the dynamic programming algorithm finds the optimal set of orientations for each aircraft type, given data from multiple transmitters.

[2]This constraint seems reasonable given the real flight data that is available. Of course, it should be proportional to the time between measurements.

6.3 Results

Although a variety of maneuvers are tested in Chapter 4 using the simple ATR algorithm, it is unnecessary to repeat all of the results when using the revised ATR algorithm. The EKF results for the straight-and-level trajectories in Section 2.3 reveal that the estimated states are almost identical to the results of the simple coordinated flight model. Since the covariance matrices in these cases are sufficiently small that only the estimated state need be considered feasible, the results of the revised algorithm are identical to those of the simple algorithm. Thus, repeating those results here would be an exercise in tedium. Furthermore, although Section 2.3 reveals that the covariance matrices corresponding to the banked turn trajectory are large enough to encompass a few sets of feasible angles, it also reveals that the simple coordinated flight model is able to correctly estimate the orientation. Since the revised algorithm assigns the highest scores to the cases with the correct orientation (in this case), the identification results for this test remain unchanged. As such, there is no need to repeat them in this section. In short, the only maneuver with new and interesting results using the revised algorithm is also the only maneuver that gives the simple ATR algorithm trouble. Thus, in the interest of brevity, only the results from the Edwards maneuver are shown in this section.

The revised ATR algorithm from Section 6.2 is compared to the simple one from Chapter 3 under three different cases. First, the algorithms are compared under the assumption that the passive radar system only exploits the vertically polarized transmitter from Section 3.1. The second case is identical to the first, but assumes that the passive radar system only exploits the horizontally polarized transmitter from Section 3.1. The final case compares the algorithms assuming that both transmitters from Section 3.1 are simultaneously exploited by the passive radar system. Expanding the simple ATR algorithm to exploit multiple transmitters is straightforward. Rather than summing likelihoods from all transmitters at each node in the trellis (like the revised ATR algorithm), the simple one sums the likelihoods from all transmitters at the end of the scenario. These multiple transmitter results are neglected in Chapter 4 merely because the performance is no better than that obtained using a single transmitter, (in this case).

Figure 60 shows the probability of error curves obtained when the passive radar system exploits only the horizontally polarized transmitter. Since the simple ATR algorithm's odds of a correct identification are reduced to chance by the time the noise figure is 80 dB (using this maneuver and averaging over all four aircraft types), the noise figure is now only swept between 30 and 80 dB. Note that the curves corresponding to the simple ATR algorithm are derived from subsets of the data used to create Figure 25b. The catastrophic misidentifications of the simple ATR algorithm at realistic noise levels (in which all aircraft are identified as the Falcon-100) no longer occur when the revised ATR algorithm is employed. Using the simple ATR algorithm, the aircraft are all identified as the Falcon-100. Thus, the Falcon-100 is the only aircraft whose probability of error does not improve using the revised algorithm. At extremely high noise levels, the revised algorithm has a tendancy to identify all aircraft as the F-15. However, the difference between the best and worst loglikelihoods at these high noise levels is approximately 0.1% (as opposed to 350% at a noise figure of 30 dB). Thus, in any given trial with an unrealistically high noise figure, the revised ATR algorithm only favors the F-15 by a narrow margin. Confusion matrices at a noise figure of 40 dB are given in Tables 38 and 39, for the simple and revised ATR algorithms, respectively. Although performance is not perfect when the revised ATR algorithm is used, it is significantly better than that which is obtained using the simple ATR algorithm. This is especially true for the realistic noise figures of 30 and 40 dB.

Similar results are shown in Figure 61 when the vertically polarized transmitter is exploited. The curves corresponding to the simple ATR algorithm are derived from subsets of the data used to create Figure 19b. Once again, the revised ATR algorithm significantly outperforms the simple one and avoids making catastrophic misidentifications at realistic noise levels. As before, the loglikelihoods computed by the revised algorithm for unrealistically large noise figures differ by less than 0.1% amongst the four aircraft types. The confusion matrices shown in Tables 40 and 41 verify that the revised algorithm is much more reliable at realistic noise levels. When both transmitters are used, the revised ATR algorithm outperforms the simple one by an even greater margin. Probability of error curves for this case are shown in Figure 62; confusion matrices are given in Tables 42 and 43.

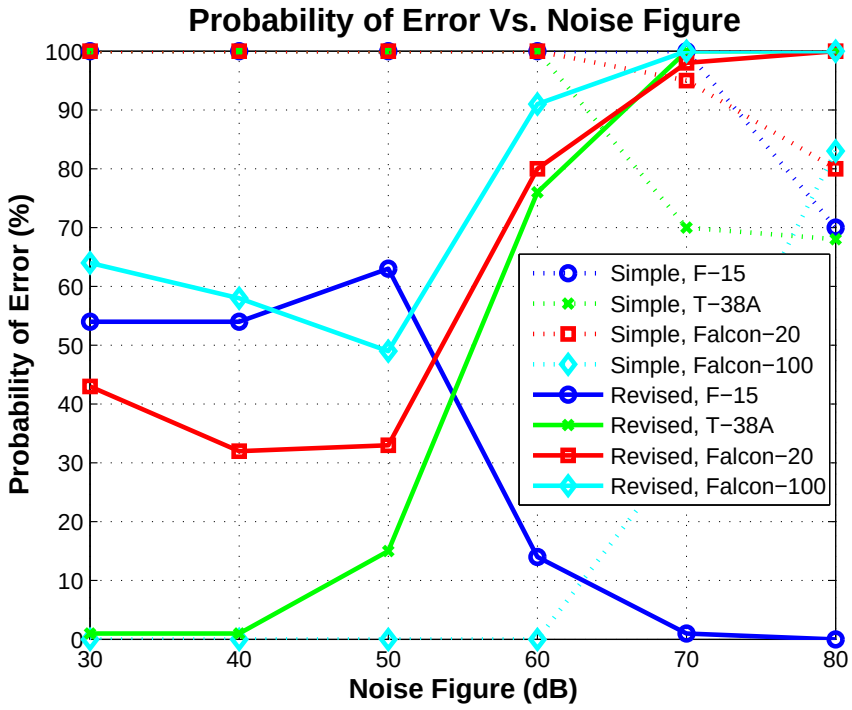

Figure 60: Probability of error vs. noise figure: Edwards trajectory using both revised and simple ATR algorithms, with the horizontally polarized transmitter

Table 38: Confusion matrix: Simple ATR algorithm, horizontally polarized transmitter, noise figure = 40 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	0	0	0	100
T-38A	0	0	0	100
Falcon-20	0	0	0	100
Falcon-100	0	0	0	100

Table 39: Confusion matrix: revised ATR algorithm, horizontally polarized transmitter, noise figure = 40 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	46	1	43	64
T-38A	0	99	0	1
Falcon-20	0	18	68	14
Falcon-100	1	43	14	42

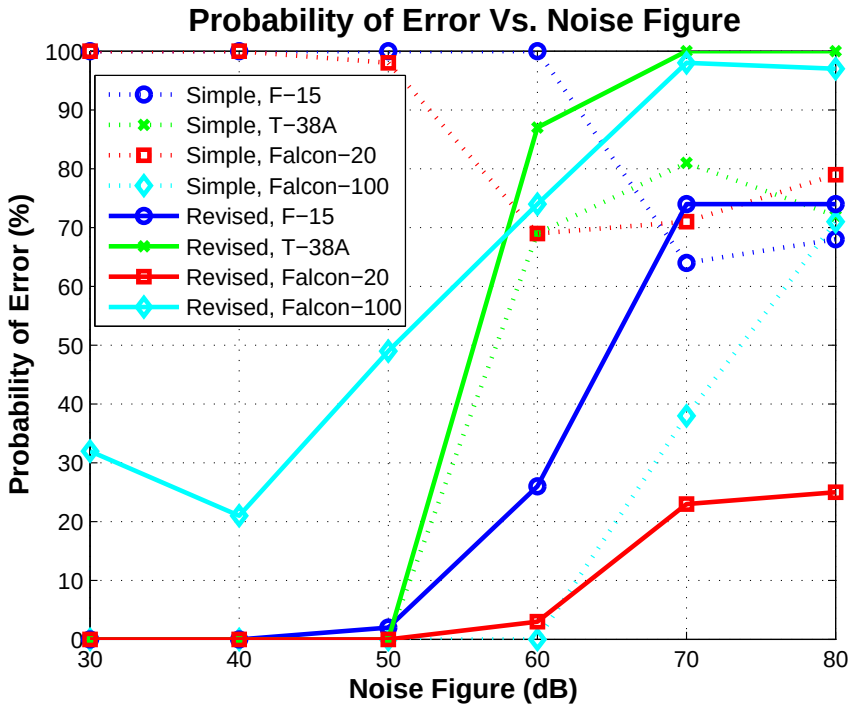

Figure 61: Probability of error vs. noise figure: Edwards trajectory using both revised and simple ATR algorithms, with the vertically polarized transmitter

Table 40: Confusion matrix: Simple ATR algorithm, vertically polarized transmitter, noise figure = 40 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	0	0	0	100
T-38A	0	100	0	0
Falcon-20	0	0	0	100
Falcon-100	0	0	0	100

Table 41: Confusion matrix: Revised ATR algorithm, vertically polarized transmitter, noise figure = 40 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	100	0	0	0
T-38A	0	100	0	0
Falcon-20	0	0	100	0
Falcon-100	0	0	21	79

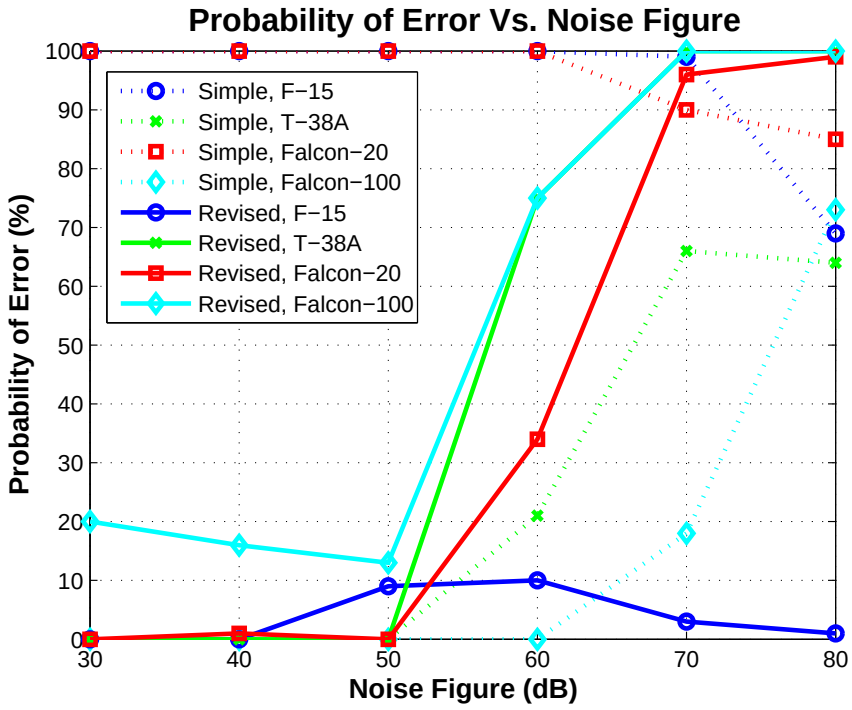

Figure 62: Probability of error vs. noise figure: Edwards trajectory using both revised and simple ATR algorithms, using both transmitters

Table 42: Confusion matrix: Simple ATR algorithm, both transmitter, noise figure = 40 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	0	0	0	100
T-38A	0	100	0	0
Falcon-20	0	0	0	100
Falcon-100	0	0	0	100

Table 43: Confusion matrix: Revised ATR algorithm, both transmitters, noise figure = 40 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	100	0	0	0
T-38A	0	100	0	0
Falcon-20	0	0	99	1
Falcon-100	0	0	16	84

The reason that the revised ATR algorithm outperforms the simple one is quite clear. The simple ATR algorithm incorporates the coordinated flight model from Section 2.1 without accounting for the possibility that the estimated orientations are incorrect. Although this is a reasonable assumption for simple maneuvers, such as straight-and-level flight or banked turns with slowly-varying radii of curvature, it leads to inaccurate orientation estimates when the aircraft executes dog-fighting style maneuvers. Of the three elements of orientation, the roll is the most likely to be incorrectly estimated by the coordinated flight model. Figure 63 shows the true roll, as well as the roll estimated by the coordinated flight model, the roll estimated by the EKF, and the roll (associated with the true aircraft) computed by the revised ATR algorithm when the noise figure is 30 dB. The roll computed by the revised ATR algorithm is closer to the truth than the roll estimated by the coordinated flight model. Since this translates into more accurately estimated power profiles, it is not surprising that the loglikelihoods computed under the revised algorithm more frequently select the correct aircraft.

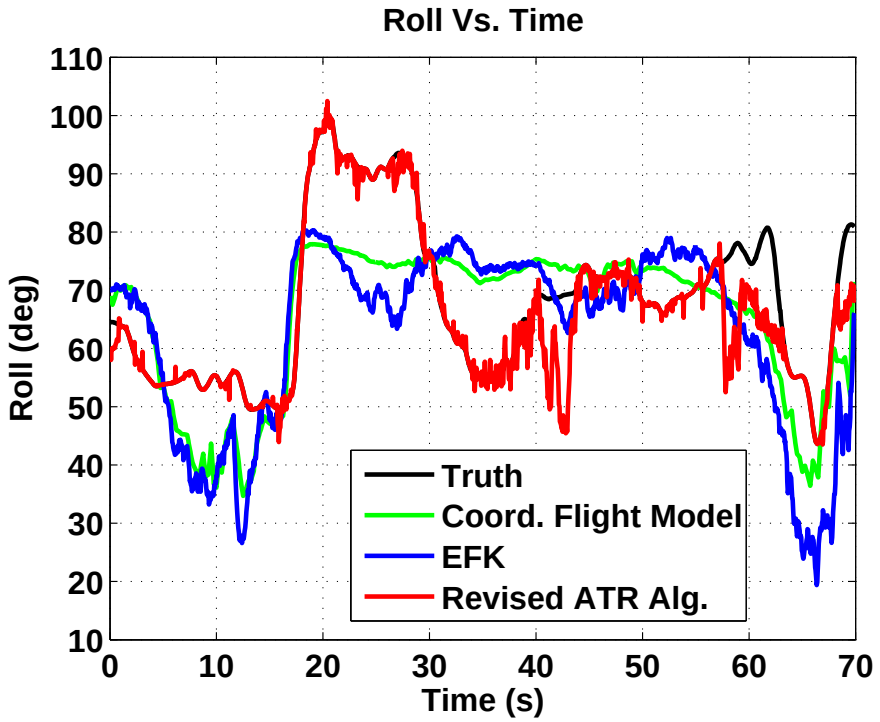

Figure 63: Comparison of aircraft roll profiles computed using coordinated flight model, EKF, and revised ATR algorithm

When the revised algorithm jointly estimates the orientations and aircraft type, the rolls associated with the aircraft not being tracked are typically farther from truth than those associated with the correct aircraft. An example is shown in Figure 64. Since the F-15 is being tracked by the passive radar system, its estimated roll profile matches the truth better than that of the Falcon-100.

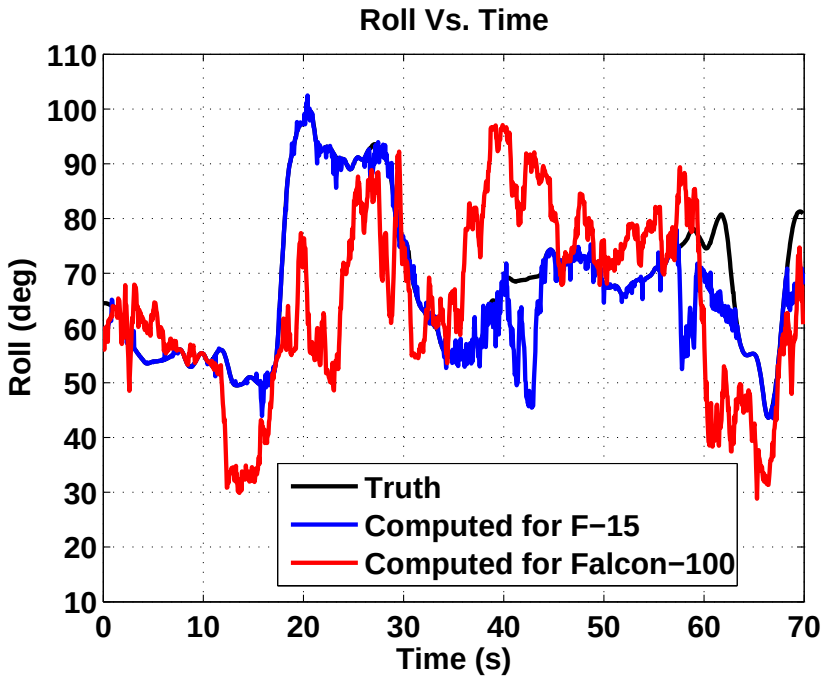

Figure 64: Comparison of aircraft roll profiles of the F-15 and Falcon-100 when the F-15 is the true target

To provide a more concise comparison of the two ATR algorithms, the probability of error curves, averaged over all four aircraft, are shown in Figure 65 for the case in which the horizontally polarized transmitter is exploited. Similar results are presented in Figure 66 for the case in which the passive radar system only exploits the vertically polarized transmitter from Section 3.1. Finally, Figure 67 provides the averaged probability of error curves for the case in which both transmitters are used. For a realistic noise figure of 40 dB, the revised algorithm correctly identifies aircraft 96% of the time, while the simple algorithm only does so in 37.5% of the Monte Carlo trials. Regardless of the transmitter selection, the revised ATR algorithm is clearly a great deal more robust than the simple one.

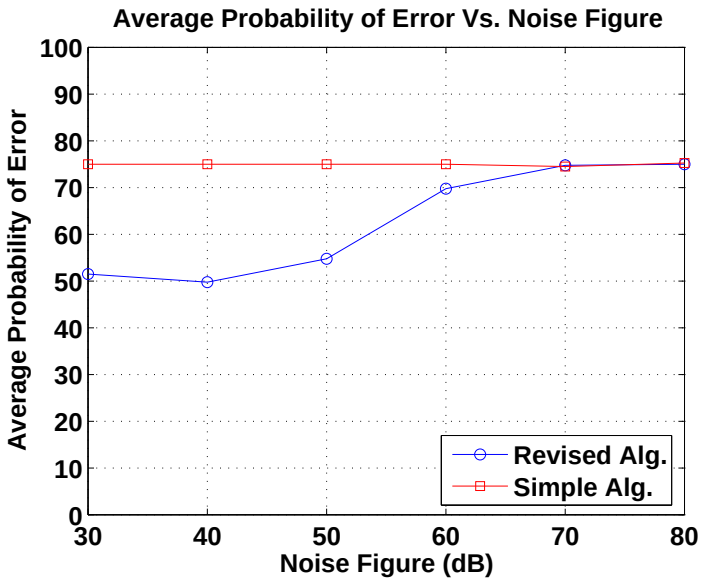

Figure 65: Probability of error vs. noise figure: Edwards trajectory using both revised and simple ATR algorithms, with the horizontally polarized transmitter

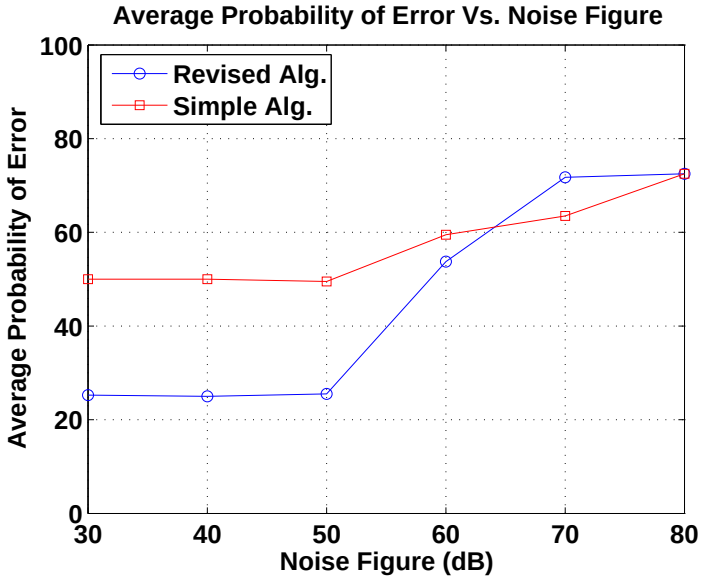

Figure 66: Probability of error vs. noise figure: Edwards trajectory using both revised and simple ATR algorithms, with the vertically polarized transmitter

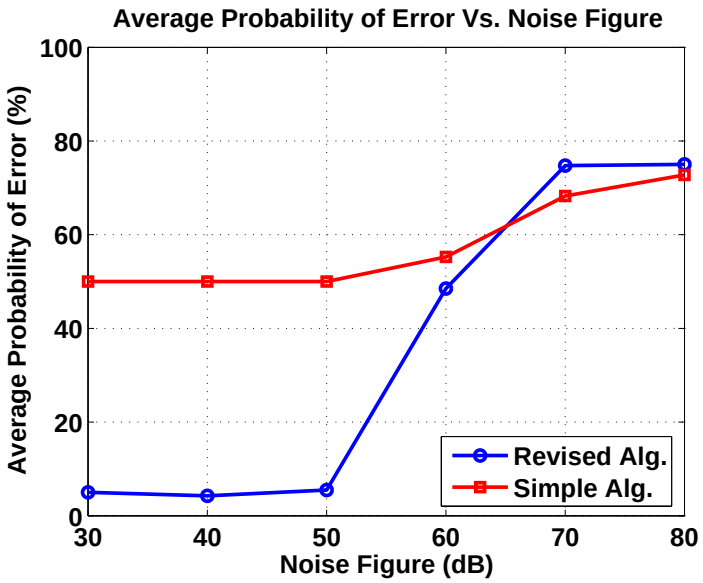

Figure 67: Probability of error vs. noise figure: Edwards trajectory using both revised and simple ATR algorithms, with both transmitters

A comparison of the revised ATR algorithm with the simple one would not be complete without noting the run times of each. The simple ATR algorithm typically executes in roughly one-tenth of the total flight time, while the revised ATR algorithm tends to need roughly twice the total flight time. For example, the 70 second Edwards trajectory has a run time of roughly seven seconds using the simple algorithm, and 140 seconds using the revised one. This run time difference is not insurmountable, as most of the run time needed for the revised ATR algorithm is spent in the dynamic programming portion of the code. Since this is an iterative process, it is coded using "for" loops. MATLAB, which runs most quickly when code is vectorized, is notorious for lengthy run times when "for" loops are used. Rewriting the dynamic programming portion of the code in C and linking it in with the rest of the MATLAB code would most likely reduce the run time by at least an order of magnitude. Furthermore, were the revised ATR algorithm ever imported into a real passive radar system, it is likely that the designer would rewrite the algorithm using optimized code in a lower-level language, further speeding run time. In short, although the revised ATR algorithm's current implementation in MATLAB has a significantly longer run time than the simple one, it is still highly likely that it could be made to run in real-time.

CHAPTER VII

ROBUSTNESS EXPERIMENTS

The process used to model the power (magnitude only) at the passive radar receiver, described in Chapter 3, uses a number of state-of-the-art models including AREPS and FISC. The use of such models begs the question: how much fidelity is really necessary for the ATR algorithm to perform well? Experiments documented in this chapter attempt to address this question.

7.1 Sensitivity of the ATR Algorithm to the AREPS Database

This section investigates the impact on the new ATR algorithm from Chapter 6 if the magnitudes of the power profiles for each aircraft in the target library are computed without incorporating the AREPS model. Instead, the propagation losses are assumed to be constant over the entire scenario. Three sets of tests are conducted. The first set considers the impact of this assumption if the constant propagation loss is reasonably estimated to be the mean propagation loss from AREPS. The next two sets of tests assess the impact if the constant propagation loss is 6 dB too low, or 6 dB too high. In all cases, the profiles used to generate the noisy test data are unaltered; they are still intended to be as realistic as possible, so they still include the propagation losses modeled by AREPS, rather than the more simplistic constant propagation loss model.

Insight is gained by comparing these three sets of tests using the Edwards maneuver and a noise figure of 40 dB, which is quite reasonable given the noise figure analysis in Section 3.1. When propagation losses from AREPS are included in the precomputed power profiles, the algorithm is correct in 95.75% of the Monte Carlo runs. When the propagation loss is treated as a constant and the constant is equal to the mean of the AREPS propagation loss profile (indicating a good estimate), the algorithm's performance only suffers slightly; the algorithm is correct in 95.5% of the trials, which equates to one more error than was previously

observed. The reason for the similar performance is straightforward. The propagation losses observed during the encounter, over the path from either transmitter to the aircraft to the receiver, vary by only ±0.75 dB. Using the constant has almost no effect on the shape of the power profile (magnitude only) that arrives at the receiver. However, if the constant is not accurately estimated, the ATR algorithm's performance declines sharply. For example, if the assumed propagation losses are 6 dB too low, the algorithm's probability of a correct identification drops from 95.5% to 44.5%. If the assumed losses are 6 dB too high, the probability of a correct identification drops to 35.75%.

If the true propagation losses are fairly constant throughout the scenario, then using an accurately selected constant propagation loss may not have a large impact on the ATR result. However, it can lead to dramatic decreases in performance if the assumed propagation loss differs from truth by a few dB. The geographical location of the aircraft, as well as its altitude, should also factor into a decision to replace AREPS by a constant propagation model. The aircraft in this particular maneuver fly at an altitude of 8000 meters over relatively flat terrain. Were the test repeated at a lower altitude over more jagged terrain, multipath effects might cause the propagation losses modeled by AREPS to play a much more significant role in defining the shape of the power profile.

7.2 Sensitivity of the ATR Algorithm to the FISC Database

One drawback of the ATR approach developed in Chapters 3 and 6 is that FISC databases must be created for each aircraft in the target library, at the frequencies of each exploited transmitter. Although this processing is done once up front and does not need to be recomputed every time an aircraft is tracked, a user of the algorithm may still wish to limit the size of the FISC database. With that in mind, this section addresses the question of whether or not separate FISC databases are necessary for each transmitter if the transmitters have reasonably close operating frequencies.

The two transmitters used in Chapters 4 and 6, for example, operate at 96 MHz and 104 MHz. This section briefly addresses how similar the radar cross sections are at these two frequencies, as modeled by FISC. Examples are given in Figures 68 and 69; these show the

absolute value of the difference between the HV and VV radar cross sections of the F-15, taken at 96 and 104 MHz, for incident and observed azimuths between zero and 360 degrees with incident and observed elevations of zero degrees. The similarities are striking. The magnitudes of the HV radar cross sections computed by FISC at the two frequencies have a mean difference of 0.51 dBsm with a standard deviation of 0.41 dBsm. The difference in the magnitudes of the VV radar cross sections is slightly larger, with a mean of 1.20 dBsm and a standard deviation of 1.03 dBsm. Since the power profiles in Sections 4.1 and 4.2 frequently vary by as much as 30 dB, modeling signals from both transmitters as though they came from a 100 MHz transmitter is not likely to impact the ATR algorithm's outcome.

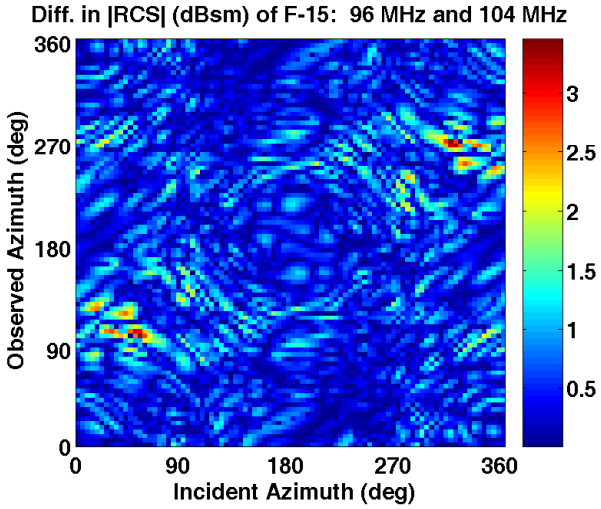

Figure 68: Difference in magnitude of HV RCS for the F-15, collected at 96 MHz and 104 MHz

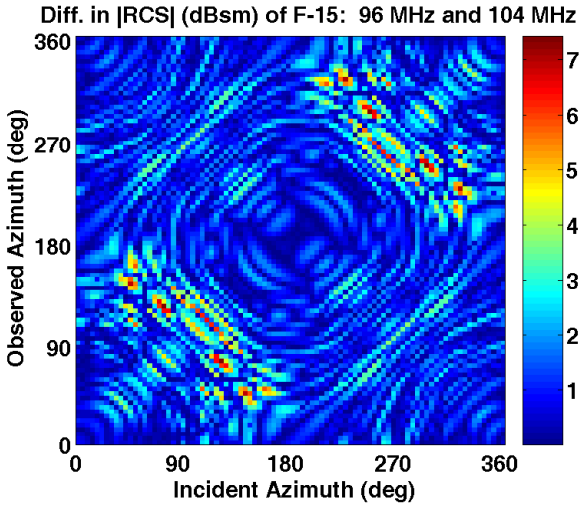

Figure 69: Difference in magnitude of VV RCS for the F-15, collected at 96 MHz and 104 MHz

CHAPTER VIII

CONCLUSIONS

8.1 Contributions

The contributions of this research are three-fold. The most notable contribution is the development of a robust ATR procedure, described in Chapter 6, that can be used to covertly identify aircraft executing maneuvers ranging from straight-and-level trajectories to dog-fighting trajectories. At the anticipated noise figure of 40 dB, the ATR algorithm never makes a mistake if the aircraft being tracked execute predictable straight-and-level or banked turn trajectories. Even when the aircraft under track execute dog-fighting style trajectories, the algorithm's probability of error is less than 5%. This is not meant to suggest that the algorithm would be perfect (or nearly so) in practice, as even the best models fail to account for all effects seen in the real world. However, the simulations suggest that the ATR algorithm has a great deal of potential for reliable performance if imported into a real system. Furthermore, the algorithm could be applied (with superficial changes) to low-frequency *active* radar systems.

The second major contribution of this research is the development of an EKF that estimates aircraft orientation from velocity measurements. This EKF plays a role in the ATR algorithm discussed in the previous paragraph, but has the potential to be used in other applications, as well. Although this dissertation applies the EKF to passive radar velocity measurements generated by a tracker such as the one described by Howland [29], there is no reason it cannot be applied to measurements generated by other trackers. The measurements and their error models would simply need to change to accommodate a different tracker and/or source.

The final contribution of this work is the derivation of closed-form approximations for the relative entropy and Chernoff information between two Rician distributions. Although a

wide variety of problems are well-modeled by the Rician distribution, the Gaussian distribution has been frequently used in its place because of mathematical convenience. Perhaps the development of a closed-form approximation for the relative entropy and Chernoff information between two Rician distributions will cause this practice of substituting the Gaussian distribution for the Rician distribution to eventually wane in popularity.

8.2 Future Work

The work presented in this dissertation encapsulates the academic portion of the problem. The ability to covertly identify aircraft has been transformed from a mere idea into a full-fledged, robust ATR algorithm that shows significant promise against aircraft executing a wide variety of maneuvers. The simulations used to test the ATR algorithm use the best RCS-modeling software currently available and account for propagation losses and antenna gain, rendering them of higher fidelity than many of the passive radar simulations used to generate results in the open literature.

The nature of next step, which falls outside the scope of this research, has more of an engineering flavor than an academic one. One such course of future work involves actually testing the ATR algorithm against a wider variety of real targets using a fully developed passive radar system. Such systems will likely be available in the next few years, making this task feasible. Another avenue for future work could explore the process noise tuning and covariance inflation used by the EKF from Section 2.3 that estimates aircraft orientation. These issues have been addressed in Section 2.3 using the limited amount of aircraft data available. If more position/orientation ground truth data becomes available in the future, (for this application or for another one), these practical implementation issues could be revisited. A third area for future work is hinted at by the results in Section 7.1, in which a constant propagation loss is used instead of results from AREPS. The ATR algorithm only performs well if the constant propagation loss is close to the mean losses estimated by AREPS. If the assumed constant is too high or too low, the algorithm's performance significantly degrades. The work in this dissertation has assumed that the receiver is well-calibrated. If this is not the case, performance degradation like that in Section 7.1 may

occur. A fourth avenue for future work could address this issue by jointly estimating this scaling factor, along with the aircraft orientation and type. A final suggested topic for future work would thoroughly explore the sensitivity of FISC to changes in frequency, as well as the aircraft CAD models. Although this final topic is expansive enough to constitute an entire thesis, the closed-form approximations of the relative entropy and/or Chernoff information from Chapter 5 could be implemented as a good starting point.

REFERENCES

[1] *http://sunspot.spawar.navy.mil/*, October 23, 2005.

[2] *http://www.nec2.org/*, October 23, 2005.

[3] *http://www.digimation.com/models/default.aspx*, October 23, 2005.

[4] *User's Manual for FISC: Fast Illinois Solver Code*. Center for Computational Electromagnetics at the University of Illinois, and DEMACO, INC., 1996.

[5] Alessandri, A., Cuneo, M., Pagnan, S., and Sanguineti, M., "On the convergence of ekf-based parameters optimization for neural networks," in *Proceedings of the 42nd IEEE Conference on Decision and Control*, (Maui, HI), pp. 6181–6186, December 2003.

[6] Bajcsy, P. and Chaudhuri, A., "Benefits of high resolution sar for atr of targets in proximity," in *Proceedings of the IEEE Radar Conference, 2002*, (Long Beach, CA), pp. 29–34, April 2002.

[7] Baltes, R. and van Keuk, G., "Tracking multiple maneuvering targets in a network of passive radars," in *Proceedings of the IEEE International Radar Conference, 1995*, (Alexandria, VA), pp. 304–309, May 1995.

[8] Barton, D., *Modern Radar System Analysis*. Artech House, 1988.

[9] Brogan, W. L., *Modern Control Theory*. Prentice Hall, 1991.

[10] Casasent, D. and Chen, X., "Mine and vehicle detection in hyperspectral image data: Waveband selection," in *Automatic Target Recognition XIII* (Sadjadi, F., ed.), vol. 5094, (Orlando, FL), pp. 228–241, April 2003.

[11] Cetin, M., Karl, W., and Castanon, D., "Feature enhancement and atr performance using nonquadratic optimization-based sar imaging," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 39, pp. 1375–1395, Oct. 2003.

[12] Cetin, M. and Lanterman, A., "Region-enhanced passive radar imaging," *IEE Proceedings on Radar, Sonar, and Navigation*, vol. 153, pp. 185–194, June 2005.

[13] Chen, J. and Walton, E., "Comparison of two target classification techniques," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 22, pp. 15–21, Jan. 1986.

[14] Cover, T. M. and Thomas, J. A., *Elements of Information Theory*. John Wiley & Sons, 1991.

[15] DeBruijn, N., *Asymptotic Methods in Analysis*. Dover Publications, Inc., 1981.

[16] Dole, C., *Flight Theory and Aerodynamics*. John Wiley & Sons, 1981.

[17] Ehrman, L. and Lanterman, A., "Automated target recogntion using passive radar and coordinated flight models," in *Automatic Target Recognition XIII*, vol. SPIE Proc. 5094, (Orlando, FL), April 2003.

[18] Ehrman, L., *Automatic Target Recognition Using Passive Radar and a Coordinated Flight Model*. Atlanta, GA: Master's Thesis, School of Electrical and Computer Engineering, Georgia Institute of Technology, 2003.

[19] Ehrman, L. and Lanterman, A., "Target identification using modeled radar cross sections and a coordinated flight model," in *Proceedings from the Third Multi-National Conference on Passive and Covert Radar*, (Seattle, WA), October 2003.

[20] Ehrman, L. and Lanterman, A., "A robust algorithm for automated target recognition using passive radar," in *Proceedings from the IEEE Southeastern Symposium on System Theory*, (Atlanta, GA), March 2004.

[21] Ehrman, L. and Lanterman, A., "A robust algorithm for automatic target recognition using passive radar," in *Automatic Target Recognition XIV*, vol. SPIE Proc. 5126, (Orlando, FL), April 2004.

[22] Greneker, E. and Geisheimer, J., "The use of passive radar for mapping lightning channels in a thunderstorm," in *Proceedings of the 2003 IEEE Radar Conference*, (Huntsville, AL), pp. 28–33, May 2003.

[23] Griffiths, H. and Baker, C., "Measurement and analysis of ambiguity functions of passive radar transmissions," in *Proceedings of the IEEE International Radar Conference, 2005*, (Arlington, VA), pp. 321–325, May 2005.

[24] Griffiths, H. and Baker, C., "Passive coherent location radar systems: Part 1: Performance prediction," *IEE Proceedings on Radar, Sonar, and Navigation*, vol. 152, pp. 153–159, June 2005.

[25] Griffiths, H. and Long, N., "Television-based bistatic radar," *IEE Proceedings, Part F*, vol. 133, pp. 649–657, December 1986.

[26] Grossman, W., "Connect the pings: Stealth radar from cell-phone radiation," *Scientific American*, pp. 26–28, March 2003.

[27] Herman, S. and Moulin, P., "A particle filtering approach to joint radar tracking and automatic target recognition," in *Proc. IEEE Aerospace Conference*, (Big Sky, Montana), March 10-15 2002.

[28] Herman, S., *A Particle Filtering Approach to Joint Passive Radar Tracking and Target Classification*. Urbana, IL: Doctoral Dissertation, Department of Electrical and Computer Engineering, Univ. of Illinois at Urbana-Champaign, 2002.

[29] Howland, P., "Target tracking using television-based bistatic radar," *IEE Proc. F: Radar, Sonar, and Navigation*, vol. 146, pp. 166–174, June 1999.

[30] Jacobs, S. and O'Sullivan, J., "Automatic target recognition using sequences of high resolution radar range-profiles," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 36, no. 2, pp. 364–382, 2000.

[31] Kuethe, A. and Chow, C.-Y., *Foundations of Aerodynamcis: Bases of Aerodynamic Design, 3rd Edition*. John Wiley & Sons, 1976.

[32] Kullback, S., *Information Theory and Statistics*. John Wiley & Sons, 1959.

[33] Lin, H. and Ksienski, A., "Optimum frequencies for aircraft classification," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 17, pp. 656–665, Sept. 1981.

[34] Lin, Y. and Ksienski, A., "Identification of complex geomtrical shapes by means of low-frequency radar returns," *The Radio and Electronic Engineer*, vol. 46, pp. 472–486, Oct. 1976.

[35] Marzetta, T., "EM algorithm for estimating the parameters of a multivariate complex rician density for polarimetric SAR," in *Proceedings of the 1995 Internatinal Conference on Acoustics, Speech, and Signal Processing*, vol. 5, pp. 3651–3654, May 1995.

[36] Mensa, D., "Radar imaging," *International Journal of Imaging Systems and Technology*, vol. 4, pp. 148–163, 1992.

[37] Meyer, C. and Gibson, D., *Route Surveying and Design, 5th Edition*. Harper & Row, Publishers, 1980.

[38] Munson, D. and Wu, Y., "Multistatic passive radar imaging using the smoothed pseudo wigner-ville distribution," in *Proceedings of the 2001 International Conference on Image Processing*, vol. 3, pp. 604–607, Oct. 2001.

[39] Nandhakumar, N., Michel, J., Arnold, D., Tsihrintzis, G., and Velten, V., "Robust thermophysics-based interpretation of radiometrically uncalibrated ir images for atr and site change detection," *IEEE Trans. on Image Processing*, vol. 6, pp. 65–78, Jan. 1997.

[40] Nguyen, D., Benitz, G., Kay, J., Orchard, B., and Whiting, R., "Superresolution hrr atr with high definition vector imaging," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 37, pp. 1267–1286, Oct. 2001.

[41] Port, O., "Super radar, done dirt cheap," *Business Week Online*, October 2003.

[42] Ringer, M. A. and Frazer, G. J., "Waveform analysis of transmitters of opportunity for passive radar," in *Proceedings of the Fifth International Symposium on Signal Processing and its Applications*, (Brisbane, Australia), pp. 511–514, August 1999.

[43] Ringer, M. A., Frazer, G. J., and Anderson, S. J., *Waveform Analysis of Transmitters of Opportunity for Passive Radar, DSTO-TR-0809*. Surveillance Systems Division, Electronics and Surveillance Research Laboratory, June 1999.

[44] Robila, S. and Varshney, P., "Further results in the use of independent components analysis for target detection in hyperspectral images," in *Automatic Target Recognition XIII* (Sadjadi, F., ed.), vol. 5094, (Orlando, FL), pp. 186–195, April 2003.

[45] Ruel, S., English, C., Melo, L., Berube, A., Aikman, D., Deslauriers, A., Church, P., and Maheux, J., "Field testing of a 3-D automatic target recognition and pose estimation algorithm," in *Automatic Target Recognition XIV* (Sadjadi, F., ed.), vol. 5426, (Orlando, FL), pp. 102–111, April 2004.

[46] Sadjadi, F., "Three-dimensional invariants and classification of ladar signatures," in *Automatic Target Recognition XIII* (Sadjadi, F., ed.), vol. 5426, (Orlando, FL), pp. 124–132, April 2004.

[47] Sahr, J. and Lind, F., "The Manastash ridge radar: A passive bistatic radar for upper atmospheric radio science," *Radio Science*, pp. 2345–2358, Nov.-Dec. 1997.

[48] Sahr, J. and Lind, F., "Passive radio remote sensing of the atmosphere using transmitters of opportunity," *Radio Science*, pp. 4–7, March 1998.

[49] Verly, J. and Delanoy, R., "Model-based automatic target recognition (atr) system for forwardlooking groundbased and airborne image laser radars (ladar)," *Proceedings of the IEEE*, vol. 84, pp. 126–163, Feb. 1996.

[50] Wagner, R., Waagen, D., and Cassabaum, M., "Image super-resolution for improved automatic target recognition," in *Automatic Target Recognition XIV* (Sadjadi, F., ed.), vol. 5426, (Orlando, FL), pp. 188–196, April 2004.

[51] Ye, J., Bresler, Y., and Moulin, P., "Cramer-rao bounds for 2-D target shape estimation in nonlinear inverse scattering problems with application to passive radar," *IEEE Trans. on Antennas and Propagation*, vol. 49, pp. 771–783, May 2001.

[52] Zoeller, C. and M.C. Budge, J., "Passive coherent location radar demonstration," in *Proc. of the IEEE Radar Conference*, pp. 358–362, 2002.

VITA

Lisa M. Ehrman received a B.S. in electrical engineering from the University of Dayton in May of 2000. She worked for the following two years for MacAulay Brown, Inc., in both the Dayton, OH and Huntsville, AL offices, supporting the Suite of Integrated Infrared Countermeasures (SIIRCM) program lead by the United States Army. Her main role was in the Test and Evaluation group, but she also supported the Modeling and Simulation side of the program. In pursuit of a research-oriented career, Lisa left MacAulay Brown in 2002 and began attending the Georgia Institute of Technology. She received her M.S. in electrical and computer engineering in May 2004, and immediately began working full-time for the Georgia Tech Research Institute (GTRI). Her research at GTRI has included development of launch point estimation and impact point prediction algorithms for small ballistic targets, fire-control radar system design, tracking unresolved separating radar targets using monopulse radar, and comparing tracking algorithms for ballistic missile defense. She plans to receive her Ph.D. in electrical and computer engineering in December 2005.