Automated Target Recognition Using Passive Radar and Coordinated Flight Models

Lisa M. Ehrman and Aaron D. Lanterman

Center for Signal and Image Processing
School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, GA 30332, USA

ABSTRACT

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. Many existing passive radar systems estimate the locations and velocities of targets. This paper focuses on adding an automatic target recognition (ATR) component to such systems.

Our approach to ATR compares the Radar Cross Section (RCS) of targets detected by a passive radar system to the simulated RCS of known targets. To make the comparison as accurate as possible, the received signal model accounts for aircraft position and orientation, propagation losses, and antenna gain patterns. The estimated positions become inputs for an algorithm that uses a coordinated flight model to compute probable aircraft orientation angles. The Fast Illinois Solver Code (FISC) simulates the RCS of several potential target classes as they execute the estimated maneuvers. The RCS is then scaled by the Advanced Refractive Effects Prediction System (AREPS) code to account for propagation losses that occur as functions of altitude and range. The Numerical Electromagnetic Code (NEC2) computes the antenna gain pattern, so that the RCS can be further scaled. The Rician model compares the RCS of the illuminated aircraft with those of the potential targets. This comparison results in target identifcation.

Keywords: Automatic Target Recognition, Passive Radar, Coordinated Flight Model, Radar Cross Section

1. INTRODUCTION

Passive radar systems which exploit "illuminators of opportunity" such as commercial television and FM radio signals have been under development for several decades. However, it is only within the last few years that the available computer power has grown to the point where such passive radar systems may be used in real-time applications. The cost dynamics of such systems are attractive. A large portion of the cost of a traditional active radar system is in the transmitter, whereas designers of "hitchhiking" or "parasitic" radars have high-power, wide-coverage transmitters already available.

This paper focuses on adding automatic target recognition (ATR) components to passive radar systems. Two parallel approaches to the recognition of airborne targets have been pursued in the literature. One approach is to form images, such as two-dimensional inverse synthetic aperture radar (ISAR) images or a sequence of one-dimensional range profiles[1] from the raw data, and then classify targets based on these images. Another approach has been to bypass the image formation process and attempt recognition directly from the received data. Herman[2,3] takes this second approach to ATR with passive radar data.

Since phase information is difficult to exploit in a passive radar system, our study focuses on classification using the received power only. To ensure robust classification in the presence of noise and errors in estimates of position and orientation, it is helpful if the Radar Cross Section (RCS) of the targets vary "slowly" with small changes in these components of the state vector. As noted by Herman,[2,3] the variation in RCS, as characterized by the number of nulls encountered as a target's aspect changes, is proportional to the electrical length of the target. At FM-band frequencies (100 MHz), a fighter-sized aircraft is approximately five wavelengths long. In

(Tel: 404-385-2548. Fax: 404-894-8363. Email: ehrman@ece.gatech.edu, lanterma@ece.gatech.edu)

contrast, at the X-band frequencies used by many active radars (10 GHz), the same aircraft would be 500 wavelengths long, making the ATR system very sensitive to small changes in target orientation.

In the late 70's and early 80's, a series of papers[4–6] illustrated that low frequencies are quite natural for target classification. Those papers had active radars in mind, but low frequency radar did not catch on in the West since most of the desired spectrum had been allocated to communications. Passive radar systems, which directly exploit existing long wavelength emissions that are good for target recognition, circumvent the frequency allocation problem faced by active radars.

1.1. History and Examples of Passive Radar Systems

In the United Kingdom in the mid-80s, Griffiths and Long[7] 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, Howland[8] 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.

The most widely tested manifestation of this kind of technology is probably the Silent Sentry series developed by Lockheed Martin Mission Systems of Gaithersburg, MD, which can exploit both analog television and FM radio signals. Another well-known system is the Manastash Ridge Radar* developed by John Sahr at the University of Washington,[9, 10] which uses FM radio-based passive radar for atmospheric studies. Other passive radar testbed systems exploiting television or radio broadcasts have been developed by Dynetics,[11] ONERA,[12, 13] and Avtec†.

1.2. Our Approach

Our approach to ATR continues along the same route as Herman's work. Rather than creating images or targets, we attempt ATR directly from the received data. Using RCS in place of images to conduct ATR poses one major challenge. Since RCS is highly aspect-dependent, accurate estimation of target orientation is crucial. Herman[2, 3] met this challenge with a complex, computationally intensive particle filter[14] which jointly estimated target position, orientation, and target type. We attempt a simpler, less computationally intensive approach in which we consider the positions and velocities output by the existing passive radar system to be "true," and employ a coordinated flight model[15] to estimate aircraft orientations from aircraft positions.

Once an aircraft flight profile has been supplemented with estimated orientations, a set of received power profiles are constructed. This set contains the received power of each aircraft model in the target class, as though each member of the target class was participating in the encounter. Target RCS is extracted from a database of FISC results to compile the power profile which would arrive at the receiving antenna. This profile is then scaled for propagation losses and antenna gain of the receiver using AREPS and NEC2, respectively. Finally, the profiles corresponding to members of the target class are compared to the profile of the actual target, resulting in target identification.

Thus far, our work is simulation-based. We simulate the target profiles arriving at the receiving antenna by adding white Gaussian noise to the profile. One hundred such noisy profiles are created for each member of the target class for each scenario to create a statistically significant set of simulations. It is worth noting that the aircraft comprising the target classes are chosen based upon the availability of acceptable CAD models, rather than likelihood of occurrence near a real system. One goal of this research is to model a system currently being developed by NATO/NC3A with the intent of an eventual comparison with real data. At that point, more realistic target classes will be implemented.

*www-rcs.ee.washington.edu/~radar/Projects/Manastash

†www.avtec.com/asd/solutions/movingtarget.htm

2. MODELING AIRCRAFT RCS

2.1. System Description

The system parameters used in this work are selected to model a passive radar demonstration currently being developed by NATO. Table 1 shows the relevant parameters for the transmitter and receiver modeled in this research.

Table 1. Transmitter and Receiver Parameters

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

2.2. Using a Coordinated Flight Model

Given that target RCS is the sole factor used to classify the aircraft, its accurate representation is paramount. This is not a trivial task. RCS is defined as the ratio of the power reflected from the target to the power incident upon the target. Thus, it is inherently dependent upon the incident and observed angles, which in turn are dependent upon the yaw, heading, pitch, and roll of the aircraft. Our response to this challenge implements a coordinated flight model to estimate these angles, given a set of time-correlated aircraft positions. This process is described at length in other sources[15]; thus, only the results will be presented here.

Since this research is primarily concerned with fast-moving fixed-wing aircraft, yaw is always assumed to be zero. This does not imply that aircraft direction goes unmeasured; it simply means that the aircraft nose is assumed to be oriented in the direction of motion. Aircraft heading, a far more revealing parameter, is expressed as,

$$\xi = \begin{cases} 90 - \arctan(\frac{v_Y}{v_X}), x > 0 \\ 270 - \arctan(\frac{v_Y}{v_X}), x < 0 \end{cases}, \tag{1}$$

where v_X and v_Y are the x and y components of aircraft velocity. The equation of pitch also follows from the assumption that the aircraft nose is pointed in the direction of motion, and is expressed as,

$$\theta = \arctan\left(\frac{dz}{\sqrt{dx^2 + dy^2}}\right), \tag{2}$$

where (dx, dy, dz) is the difference in aircraft position in the x, y, and z directions over an incremental period of time, dt. Finally, the aircraft roll is estimated as,

$$\phi = \arctan\left(\frac{|v|^2 \cos(\theta)}{Rg}\right), \tag{3}$$

where $|v|$ is the magnitude of the aircraft velocity, R is the instantaneous radius of curvature of the aircraft turn, and g is the acceleration of gravity.

The estimated aircraft orientation angles are appended to the aircraft positions, creating a supplemented flight profile. This is then used to determine the incident and observed angles, upon which RCS depends.

2.3. Modeling the Power at the Receiver

Modeling the power profile that arrives at the receiving antenna is a multi-step process. First, the RCS corresponding to the particular set of incident and observed angles is modeled. This is accomplished by accessing a database of FISC results, which are available for each aircraft model in the target class. Additional scaling is required to convert this RCS from FISC into a power profile representing what arrives at the receiver due to the illuminated target. Some significant factors which 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. In our case, as is likely to be the case for most passive radar applications, the transmitting antenna is omni-directional. Thus, the only antenna gain pattern that must be modeled corresponds to the receiver; this modeling is done using 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 criteria. In particular, the angular sampling of the RCS should satisfy,

$$\Delta\theta \leq \frac{c}{(2f_0)(size)}, \tag{4}$$

where c is the speed of light, f_0 is the radar frequency, and $size$ is the longest dimension of the aircraft.[16] The minimum angular sampling required for each aircraft when the frequency is 100 MHz is shown in Table 2. Thus,

Table 2. Minimum Angular Sampling Required for Each Aircraft in the Study

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

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 virtually any desired power profile.

Although the RCS simulated by FISC are sampled according to the Nyquist frequency, the amount of time required to create a sufficiently large database remains daunting. At this point, the database has been completed for all combinations of observed azimuths and elevations, and incident azimuths. However, the incident elevations available in the database are currently limited to $\pm 20^o$[‡]. More sophisticated techniques will eventually be implemented to circumvent this limitation.

3. TARGET IDENTIFICATION

3.1. Simulating the Power Profiles at the Receiver

Recall that one goal of this research is to model a system currently being developed by NATO. Eventually, real power profiles may be available from the NATO receiver. Until then, it is necessary to simulate these profiles.

[‡]FISC requires a great deal of computational time to model the aircraft for a particular set of incident angles; however, once the model has been created for the set of incident angles, all of the possible pairs of observed angles can be computed relatively quickly. For this reason, the limited parameter in our database is the incident elevation.

First, the simulation is run with the real aircraft orientations used 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. Along these lines, the simulated received profile is expressed as,

$$P_{SIM} = (\sqrt{P_R} + w_R)^2 + w_I^2, \tag{5}$$

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.[3] The noise power is computed using,

$$P_N = \frac{kT_0N_F}{CPI}, \tag{6}$$

where k is Boltzmann's constant, T_0 is temperature in Kelvin, N_F is the noise figure, and CPI is the coherent processing interval of the system.[17] To match the NATO system, the CPI is set equal to 0.5 seconds, and T_0 is set equal to 290 K. In this case, an appromiate conversion from noise figure to noise power is found by subtracting 200 dB from the noise figure.

Selection of the noise figure, N_F, is more difficult. We have thus far ignored the effects of both out-of-band and direct-path interference; however, these factors are not negligible. In the absence of data indicating the extent to which interference affects the system being modeled, we have opted to sweep the noise figure over a range of values. Using this approach, it is possible to quantify system performance as a function of noise power.

To recap, the simulated received profiles are created by adding noise to the profiles generated using real orientation angles. In contrast, the library profiles used for target recognition result from using orientation angles estimated by the coordinated flight model.

3.2. Target Classification

The process described in Section 3.1 is repeated 100 times for each member of the target class, resulting in a set of simulated received profiles. Target identification results from comparing these received profiles to the library of profiles. The Rician model naturally lends itself to this comparison.[3] Its probability density function is given by,

$$p_x(x) = \frac{x}{\sigma_w^2} e^{\left(\frac{x^2+s^2}{2\sigma_w^2}\right)} I_0\left[\frac{xs}{\sigma_w^2}\right]. \tag{7}$$

To apply the Rician to the simulated and truth profiles, associate x with the magnitude of the received profile, and s with the magnitude of the library profile. The noise power, which equals the noise variance, is then equated with σ_w^2. If each point in time is considered an independent sample from a process, then the data loglikelihood is,

$$\ln(p_x(\bar{x})) = \sum_{i=1}^{n} \ln\left(I_0\left[\frac{x_i s_i}{\sigma_w^2}\right]\right) - \left(\frac{x_i^2 + s_i^2}{2\sigma_w^2}\right). \tag{8}$$

Loglikelihoods are computed for each member of the target class; the target is identified as the member of the target class with the largest loglikelihood.

4. RESULTS

Three scenarios of increasing complexity have been presented to the algorithm. The first scenario consists of a simple straight-and-level flight path. This does little to test the coordinated flight model, but gives a good indication that the remainder of the algorithm is working. Next, a constant-altitude, circular banked turn is implemented. Although this simple circular banked turn is not a strenuous test, it indicates that the algorithm is able to perform when the power profiles become more complex. Finally a flight profile recorded on-board a

maneuvering F-15[§] is used to provide a difficult test of the algorithm. The F-15 trajectory came complete with aircraft orientation, providing us with a unique opportunity to quantify the performance degradation induced by estimating 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 orienation angles.

The locations of these three maneuvers relative to the system transmitter and receiver are shown in Figure 1. Note that the area currently covered by our databases is represented as the intersection of the two circles shown in the figure.

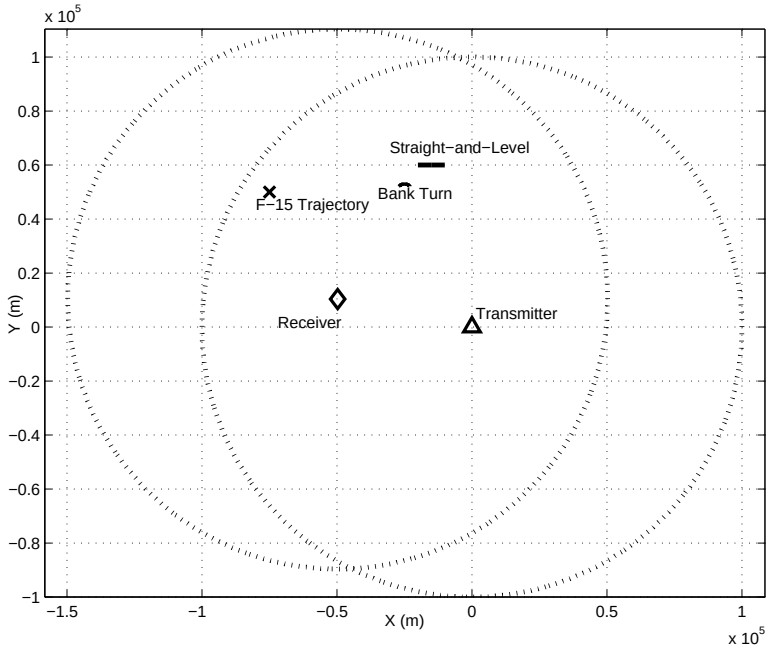

Figure 1. Maneuver Location Relative to System Transmitter and Receiver

4.1. Straight and Level Flight

The first maneuver used to test the algorithm is a straight-and-level flight. The target in this encounter has a velocity of 200 m/s and an altitude of 8000 m. The library profiles resulting from this manuever are shown in Figure 2a.

The noise figure is then swept from 30 dB to 70 dB, in increments of 5 dB, to gauge performance degradation as a function of noise power. The percentage of incorrect identifications obtained for each aircraft model at each noise level is shown in Figure 2b. This data is obtained from a set of Monte Carlo runs; thus the jumps in the curves are artificial, and are expected to smooth out as the number of runs is increased.

As is clear from Figure 2b, the algorithm performs perfectly when the noise figure is 30 dB, and only experiences minor problems when the noise figure increases to 35 dB. Nevertheless, a rapid decline in performance is noted once the noise figure rises above 35 dB.

The decline in algorithm performance corresponds to the point at which the noise power becomes greater than the signal power; in this case, the noise power surpasses the signal power when the noise figure is approximately 35 dB. In summary, the algorithm is able to correctly identify aircraft executing this maneuver, as long as SNR ≥ 1. As the library profile becomes buried in the noise, the algorithm performance declines.

[§]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 Forc Base in May 2000.

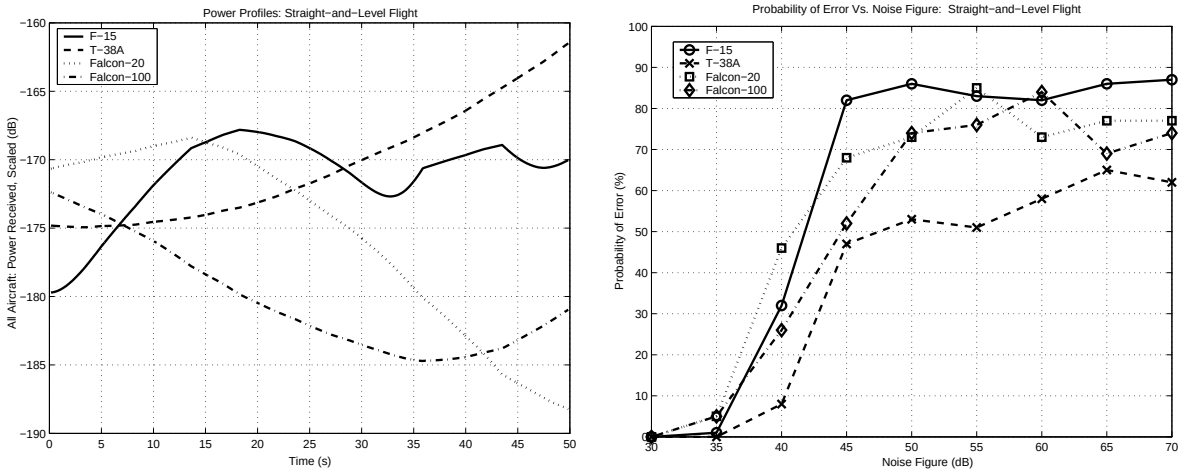

Figure 2. Straight-and-Level Flight Profile: a.) Power Profiles (left), b.) Probability of Error Vs. Noise Figure (right)

4.2. Constant Altitude Banked Turn

The constant-altitude circular banked turn provides a more difficult test of the algorithm. In this case, the target velocity is 100 m/s, and the altitude is 8000m. Like the straight-and-level profile this is an idealized maneuver, rather than one based upon an actual recorded profile. Unlike the straight-and-level maneuver, this trajectory requires the coordinated flight profile to estimate non-zero aircraft roll. The trajectory and its location relative to the antennas also provide more complicated power profiles than the previous test. The library profiles resulting from this maneuver are shown in Figure 3a.

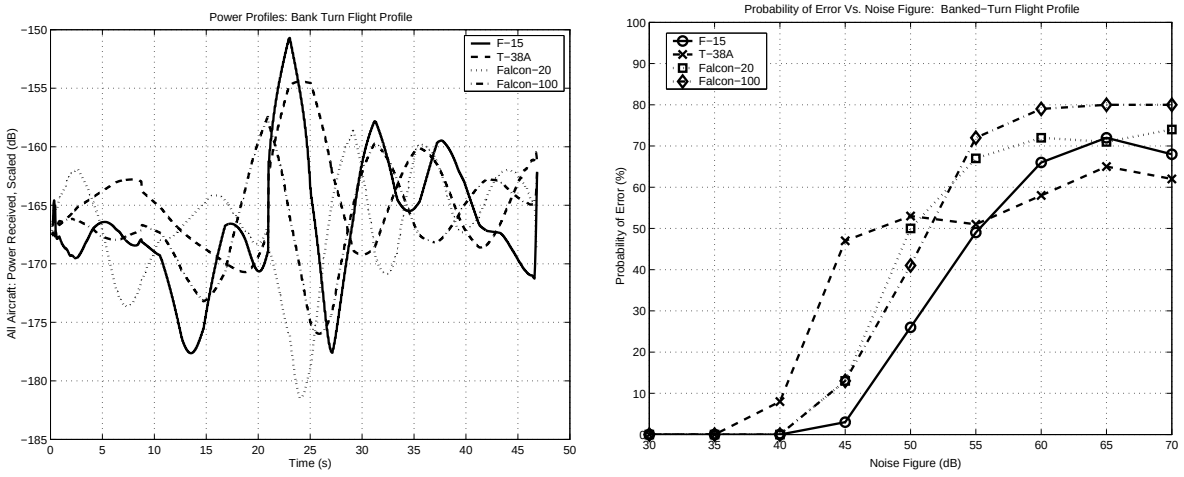

Figure 3. Constant-Altitude Circular Banked Turn: a.) Power Profiles (left), b.) Probability of Error Vs. Noise Figure (right)

Once again, the noise figure is swept from 30 dB to 70 dB. The probability of error is shown as a function of the noise figure in Figure 3b, and confusion matrices corresponding to noise figures of 45, 50, and 55 dB are given in Tables 3 through 5.

As was the case using the straight-and-level profile, the algorithm performs quite well until the noise power surpasses the signal power. At that point, the probability of error approaches 75%, which could be attained

Table 3. Confusion Matrix for Banked-Turn Test with Noise Figure = 45 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	97	3	0	0
T-38A	3	96	0	1
Falcon-20	0	1	87	12
Falcon-100	0	1	12	87

Table 4. Confusion Matrix for Banked-Turn Test with Noise Figure = 50 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	74	16	4	6
T-38A	21	59	9	11
Falcon-20	11	10	50	29
Falcon-100	15	9	17	59

through chance. The only significant difference in performance from the previous test is due to the larger amplitude of the library profiles obtained from the banked-turn maneuver. Since the maximum power of these library profiles is approximately -155 dB, the noise power is expected to overtake the signal power for noise figures greater than or equal to 45 dB. The degradation in performance for noise figures above 45 dB is evident in Figure 3b.

4.3. F-15 Flight Profile

To gauge the performance of the algorithm when the target is maneuvering, the third test uses a F-15 trajectory. During the maneuver, the F-15 changes altitudes and executes turns with varying degrees of curvature. This maneuver is shown in Figure 4 from both 3-D and 2-D perspectives. The maneuver clearly provides an opportunity to exercise the coordinated flight model.

To generate a baseline for comparison, this algorithm is first run using the known aircraft orientation angles. This essentially removes the coordinated flight model, and any error it induces, from the process. The library profiles generated in this fashion are shown in Figure 5a. Since the maximum amplitude of these profiles is approximately -140 dB, it is likely that the algorithm will have difficulty identifying the aircraft for noise figures above 60 dB. The probability of classification error for this set of tests is shown in Figure 5b. The confusion matrices corresponding to noise figures of 55, 60, and 65 dB are given in Tables 6 through 8. As predicted, the probability of error begins to climb as the noise figure climbs above 60 dB.

This test is then redone using the full algorithm, including the coordinated flight model. The new library profiles are displayed in Figure 6a. The corresponding probabilities of error are shown in Figure 6b. The confusion matrices generated for noise figures of 55, 60, and 65 dB are shown in Tables 9 through 11.

With the exception of the F-15, the algorithm performance is very similar to the baseline test. The degradation in performance in identifying the F-15 when using the estimated orientation angles suggests an intriguing idea; it seems as though the F-15 profile obtained with estimated orienation angles more closely resembles the simulated received profile for other aircraft than it does the correct F-15 profile!

Table 5. Confusion Matrix for Banked-Turn Test with Noise Figure = 55 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	34	28	21	17
T-38A	28	30	24	18
Falcon-20	25	28	28	19
Falcon-100	26	28	25	21

Table 6. Confusion Matrix for F-15 Trajectory, Using Real Orientation Angles, with Noise Figure = 55 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	100	0	0	0
T-38A	1	97	0	2
Falcon-20	0	0	100	0
Falcon-100	0	1	0	99

Table 7. Confusion Matrix for F-15 Trajectory, Using Real Orientation Angles, with Noise Figure = 60 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	73	17	0	10
T-38A	15	68	0	17
Falcon-20	0	1	95	4
Falcon-100	17	16	2	65

Table 8. Confusion Matrix for F-15 Trajectory, Using Real Orientation Angles, with Noise Figure = 65 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	50	18	10	22
T-38A	28	40	9	23
Falcon-20	3	6	86	5
Falcon-100	26	18	14	42

Table 9. Confusion Matrix for F-15 Trajectory, Using Estimated Orientation Angles, with Noise Figure = 55 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	31	36	0	33
T-38A	0	94	0	6
Falcon-20	0	0	100	0
Falcon-100	0	1	0	99

Table 10. Confusion Matrix for F-15 Trajectory, Using Estimated Orientation Angles, with Noise Figure = 60 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	29	33	1	37
T-38A	6	61	0	33
Falcon-20	0	3	92	5
Falcon-100	5	19	3	73

Table 11. Confusion Matrix for F-15 Trajectory, Using Estimated Orientation Angles, with Noise Figure = 65 dB

Aircraft	F-15	T-38A	Falcon-20	Falcon-100
F-15	31	26	14	29
T-38A	22	33	14	31
Falcon-20	6	7	81	6
Falcon-100	19	21	19	41

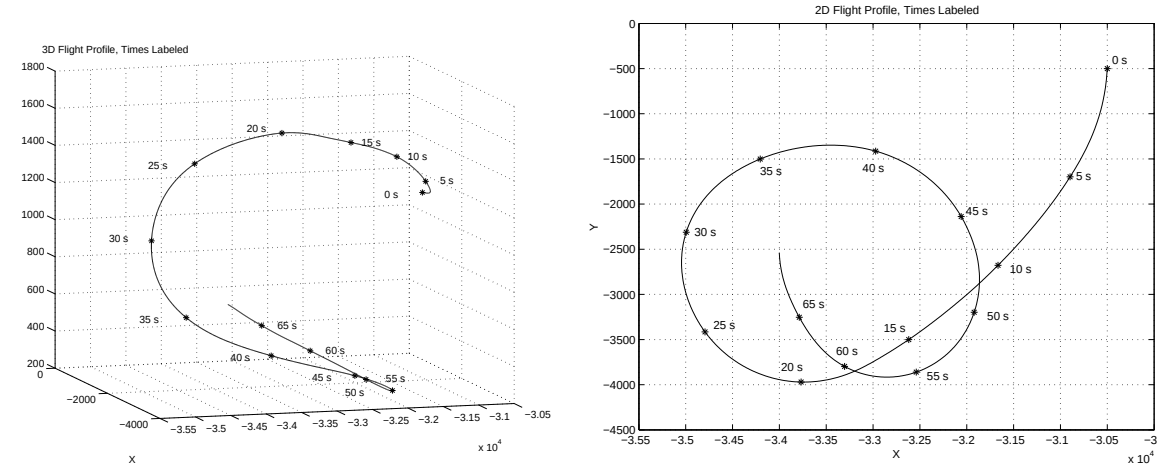

Figure 4. F-15 Maneuver Using Real Aircraft Orientation a.) 3-D View (left), b.) Top View (right)

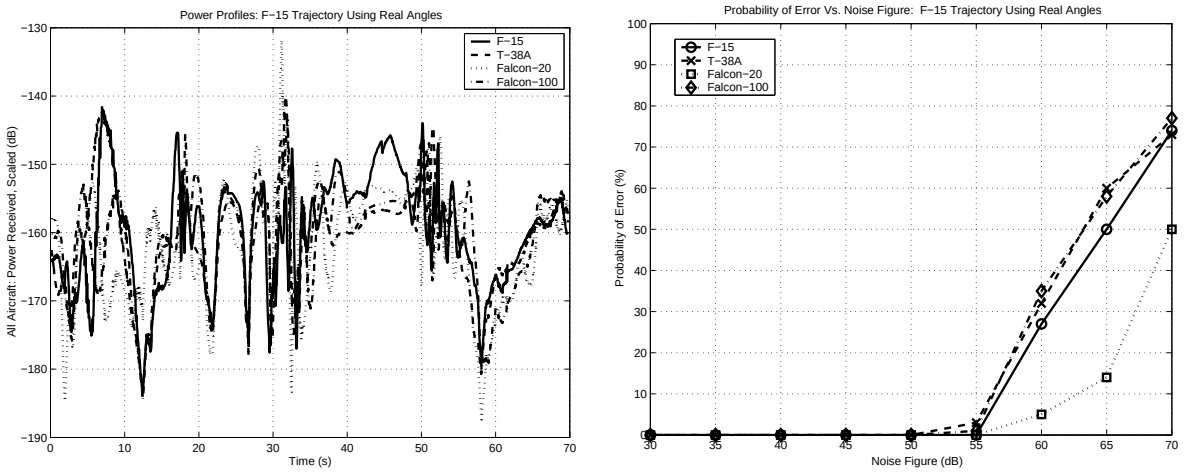

Figure 5. F-15 Trajectory Using Real Orientation Angles: a.) Power Profiles (left), b.) Probability of Error Vs. Noise Figure (right)

4.4. Summary of Results

The results presented in Sections 4.1 through 4.3 provide a good deal of insight into the viability of this algorithm. Whether the maneuver is simple or complex, the ability to identify aircraft using this algorithm is largely dependent upon having SNR ≥ 1.

Additionally, the results suggest that if performance begins to decline when while SNR ≥ 1, the coordinated flight model is to blame. Examination of the results of the coordinated flight model support this notion. Figure 7 compares the real heading, pitch, and roll of the F-15 trajectory to those estimated using the coordinated flight model. Although the estimates for heading and pitch are very good, the estimation of aircraft roll is not as accurate. As RCS is highly dependent upon these angles, particularly roll, this is probably the primary cause of mis-classification in the algorithm when SNR ≥ 1. If sufficiently high SNR cannot be obtained to reach the desired performance level, a more sophisticated means of estimating aircraft orientation may be required.

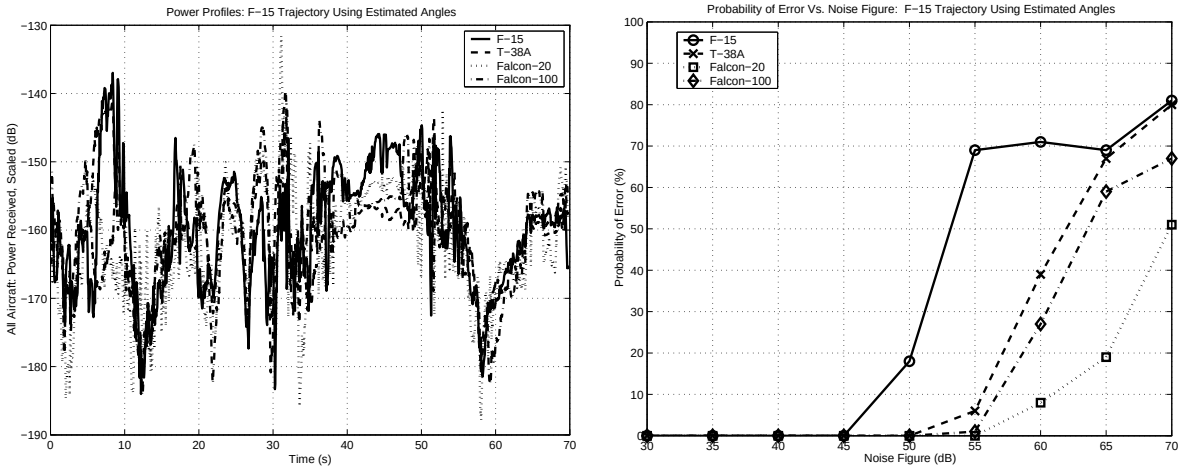

Figure 6. F-15 Trajectory Using Estimated Orientation Angles: a.) Power Profiles (left), b.) Probability of Error Vs. Noise Figure (right)

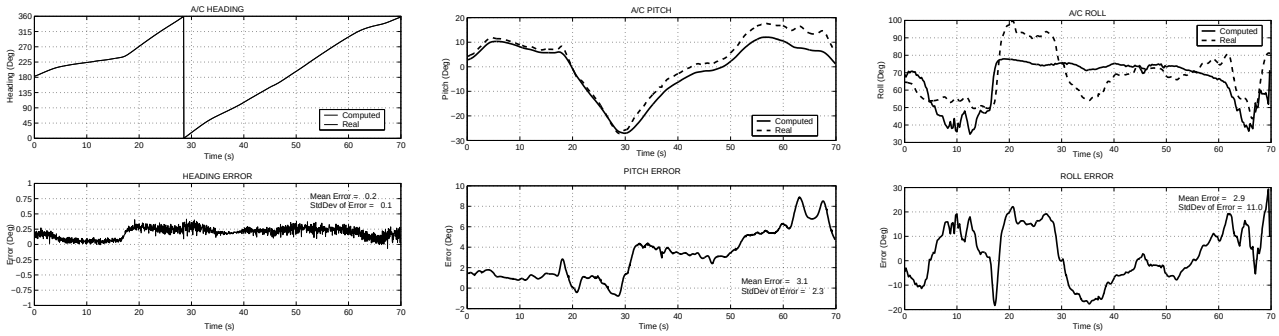

Figure 7. Results of the Coordinated Flight Model on the F-15 Trajectory: a) Heading, b) Pitch, c) Roll

5. SUMMARY

The results presented in this paper suggest that the process is viable, but would benefit from more work. Future plans include a more precise means of accounting for the out-of-band and direct-path interference, a method for handling variations induced by poor approximation of aircraft roll, and extension of the FISC database for incident elevations beyond $\pm 20^o$. Other plans for future work include accounting for errors in aircraft position estimates. Recall that we assumed that the positions output by the host passive radar system were accurate. In reality, the x-y position of the aircraft often has errors on the order of hundreds of meters. The altitude is extremely difficult to estimate at all. Future work will measure the impact of these errors and seek ways of combating them in the ATR algorithm. We also plan to include more aircraft in our study, with the goal of eventually operating on data from the NATO passive radar system.

ACKNOWLEDGMENTS

This work ws supported by NATO NC3A and start-up funds from the School of Electrical and Computer Engineering at the Georgia Institute of Technology. The authors would like to thank Dr. Paul Howland and Dr. Rene van der Heiden for their assistance.

REFERENCES

1. S. Jacobs and J. O'Sullivan, "Automatic target recognition using sequences of high resolution radar range-profiles," *IEEE Trans. on Aerospace and Electronic Systems* **36**(2), pp. 364–382, 2000.
2. S. Herman, *A Particle Filtering Approach to Joint Passive Radar Tracking and Target Classification*, Doctoral Dissertation, Department of Electrical and Computer Engineering, Univ. of Illinois at Urbana-Champaign, Urbana, IL, 2002.
3. S. Herman and P. Moulin, "A particle filtering approach to joint radar tracking and automatic target recognition," in *Proc. IEEE Aerospace Conference*, (Big Sky, Montana), March 10-15 2002.
4. Y. Lin and A. Ksienski, "Identification of complex geomtrical shapes by means of low-frequency radar returns," *The Radio and Electronic Engineer* **46**, pp. 472–486, Oct. 1976.
5. H. Lin and A. Ksienski, "Optimum frequencies for aircraft classification," *IEEE Trans. on Aerospace and Electronic Systems* **17**, pp. 656–665, Sept. 1981.
6. J. Chen and E. Walton, "Comparison of two target classification techniques," *IEEE Trans. on Aerospace and Electronic Systems* **22**, pp. 15–21, Jan. 1986.
7. H. Griffiths and N. Long, "Television-based bistatic radar," *IEE Proceedings, Part F* **133**, pp. 649–657, December 1986.
8. P. Howland, "Target tracking using television-based bistatic radar," *IEE Proc. F: Radar, Sonar, and Navigation* **146**, pp. 166–174, June 1999.
9. J. Sahr and F. Lind, "The Manastash ridge radar: A passive bistatic radar for upper atmospheric radio science," *Radio Science* , pp. 2345–2358, Nov.-Dec. 1997.
10. J. Sahr and F. Lind, "Passive radio remote sensing of the atmosphere using transmitters of opportunity," *Radio Science* , pp. 4–7, March 1998.
11. C. Zoeller and J. M.C. Budge, "Passive coherent location radar demonstration," in *Proc. of the IEEE Radar Conference*, pp. 358–362, 2002.
12. M. Lesturgie and J.-L. Zolésio, "Low frequency radar design trade-offs," in *Third ONERA-DLR Aerospace Symposium*, (Paris, France), June 20-21 2001. avaiable from www.onera.fr.
13. D. Poullin, "On the use of cofdm modulation (dab, dvb) for passive radar applications,"
14. A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*, Springer-Verlag, 2001.
15. L. Ehrman and A. Lanterman, "Estimation of aircraft orientation from flight paths using a coordinated flight model," *submitted to IEEE Transactions on Aerospace and Electronic Systems* , November 2002.
16. D. Mensa, "Radar imaging," *International Journal of Imaging Systems and Technology* **4**, pp. 148–163, 1992.
17. D. Barton, *Modern Radar System Analysis*, Artech House, 1988.