

Bayesian multiple-look updating applied to the SHARP ATR system

Matthew B. Ressler^{*a}, Rob Williams^{**b}, Dave Gross^c, Adrian Palomino^c

^aToyon Research Corporation, 75 Aero Camino, Suite A, Goleta, CA 93117-3139

^bRob Williams, AFRL/SNAA, 2241 Avionics Circle, Wright-Patterson AFB, OH 45433-7333

^cVeridian Engineering, 5200 Springfield Pike, Suite 200, Dayton, OH 45431-1289

ABSTRACT

This study summarizes recent algorithmic enhancements made to the AFRL/SNAA Systems-Oriented High Range Resolution (HRR) Automatic Recognition Program (SHARP) in the areas of multiple-look updating and sensor fusion. The benefits in improved 1-D Automatic Target Recognition (ATR) performance resulting from these enhancements are quantified. The study incorporates a unique method of estimating Bayesian probabilities by exploiting the fact that 1-D range profiles formed from Moving and Stationary Target Acquisition and Recognition (MSTAR) target chips overlap in azimuth. Thus, multiple samples of range profiles exist for the same target at very similar viewing aspects, but from independent passes of the sensor. ATR performance using the Bayesian technique is characterized first for an updating architecture that fuses probabilities over a fixed number of looks and then makes a “classify or reject” decision. A second proposed architecture that makes a “classify, reject, or take another measurement” decision is also analyzed. For both postulated architectures, ATR performance enhancement over the SHARP baseline updating procedure is quantified.

Keywords: ATR, Multiple-Look Updating, Sensor Fusion, High Resolution Radar, Imagery Exploitation

INTRODUCTION

Analysis of recent military operations in Kosovo has revealed weaknesses in our ability to identify, track, and strike moving ground targets in all weather situations, and at day or night from beyond visual range (BVR). In order to identify targets with a high level of confidence, data fusion from many sources is essential. Data fusion must occur for multiple sensor measurements from the same sensor, as well as for measurements from different sensors. In order for measurement fusion systems to work in a moving target scenario, one must be assured that the measurements are coming from the same target. This implies that the ability to track targets with a high level of confidence is necessary. Thus, tracking of moving targets aids target identification (ID). Recently, many researchers have been intrigued by the notion of using features derived from sensor measurements to help associate tracks over time. Features can help resolve ambiguous tracks in difficult situations such as long sensor revisit times (e.g., platform is turning around), multiple targets passing through an intersection, frequent line of sight blockage, and high traffic densities. Feature Aided Tracking (FAT) is the broad term that has been applied to this type of tracking. One candidate feature is target ID. Thus, it is also important to note that ID can aid tracking. This paper details findings related to one important piece of the “track aiding ID, ID aiding track” problem: the fusion of multiple measurements on a target to improve confidence in target ID.

This study focuses on fusing measurements from an HRR Moving Target Indicator (MTI) radar operating at X-band frequencies. X-band frequencies allow for detection of BVR ground targets in many diverse weather conditions and work during both night and day. Doppler processing of the HRRMTI signal allows for the detection of moving ground targets. Processing yields an estimate of the target’s position and radial velocity, both of which are used to track the target. The HRRMTI radar also yields a high resolution one-dimensional “range profile” signature of the moving ground target. The range profile signatures can be fed to an Automatic Target Recognizer (ATR) to estimate target ID. Due to the unavailability of a statistically significant HRR data set on moving ground targets, the HRR profiles used in this study were formed from measured high resolution SAR images of military ground targets from the MSTAR data set.

* Correspondence: Email: mressler@toyon.com; Phone: 805 968 6787 x152; www.toyon.com

** Correspondence: Email: rwilliam@mbvlab.wpafb.af.mil; Phone: 937 255 1105 x3434

The ATR system used for this study is the AFRL/SN sponsored SHARP baseline. The objective of the SHARP program is to develop and mature advanced air-to-ground HRR ATR capabilities for transition into suitable operational Air Force airborne and space-based platforms. The Bayesian multiple-look fusion algorithm detailed in this study was built on top of the SHARP baseline ATR system. Performance evaluation of the Bayesian multiple-look fusion algorithm was done in accordance with the procedures outlined in reference 1.

TERMS AND DEFINITIONS

This section defines a number of the technical terms and metrics used throughout this document. Note these definitions are consistent with those outlined in reference 1.

- Track** – a ground target being tracked by the tracker system
- Pose** – the viewing aspect of the target relative to the sensor
- Measurement** – a 1-D HRR range profile of a track
- Observation** – an HRR measurement used to test the ATR
- Template** – a stored 1-D HRR range profile of a given target at a given pose
- Known** – a track object in the ATR template (training) set
- Unknown** – a track object not in the ATR template (training) set
- P_{dec}** – Probability that a declaration will be made given a decision opportunity on a known track (said differently, the probability that a known track will be classified either correctly or incorrectly as a member of the template [training] set)
- P_{id}** – Probability that a known track is correctly classified
- P_{fa}** – Probability that an unknown track is classified as an object in the ATR template (training) set

SHARP ATR SYSTEM WITH MULTIPLE-LOOK UPDATING

This section describes the SHARP ATR system and the incorporation of a multiple-look updating process into the system architecture. Figure 1 shows a system architecture diagram for the SHARP ATR system with multiple-look updating capability. First, an HRR observation and pose estimate of a track enter the Index and Match portion of the ATR. Candidate templates from the template database are selected based on the pose estimate (nominally templates within +/- 5 degrees of the pose estimate). The observation is matched against all candidate templates. To determine the discriminant or match score, the

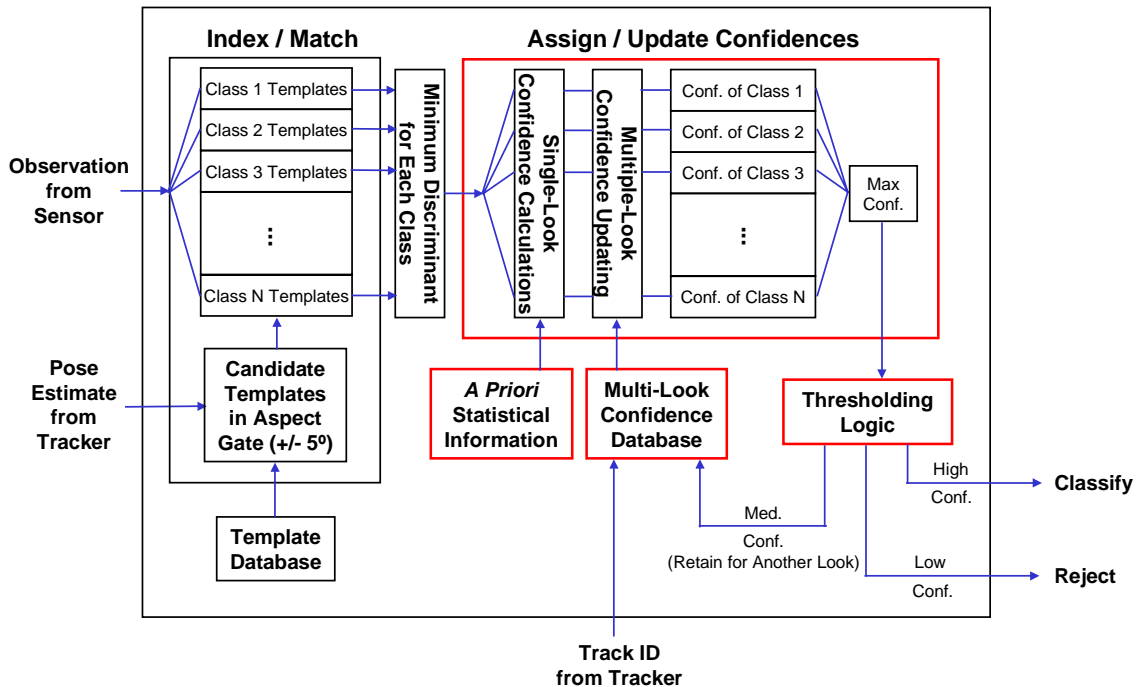


Figure 1. SHARP ATR system with multiple-look updating.

SHARP ATR system first performs a linear regression fit of the template to the observation in order to minimize the effects of scale and offset errors. The Mean Square Error (MSE) is then computed between the fit template and the observation. This MSE becomes the discriminant value (or match score) for this template–observation pair. The template from each class with the minimum discriminant becomes the winning template for that class.

The discriminant from the winning template for each class is then passed to the Assign / Update Confidences portion of the ATR system. Discriminants are converted to “confidences”. The confidences essentially put discriminants from the observation of the current look on a level playing field with discriminants from observations of other looks and from other sensors. For most systems, confidences are generally probabilities or likelihoods. Next, confidences from the current look are fused with the stored confidences for this track that have been updated over previous looks at the target. Note that in order for this fusion process to take place, a track ID must be used to access the stored confidence levels. Thus, correct track association from look to look is crucial to the updating process.

Once confidences have been updated, the maximum confidence over all of the classes is passed to a thresholding block. If the confidence is sufficiently high, a track can be classified (and potentially nominated for attack, etc.). If the confidence drops below a specified level, the track is rejected (or in other words, classified as an unknown). Otherwise, the updated confidence levels for this track are stored in the confidence database and another look at the target is requested.

SHARP BASELINE VS. BAYESIAN MULTIPLE-LOOK UPDATING

This section briefly summarizes the SHARP baseline procedure for fusing multiple HRR measurements, and then details the Bayesian procedure. Both procedures are described in a “fixed number of looks” architecture. In this case, a specified number of measurements is made on a track, and a “classify or reject” decision is made on the track once the specified number of looks have been collected and fused. The Bayesian procedure is then expanded to a “variable number of looks” architecture. In this case, a “classify, reject, or retain for another look” decision is made after each measurement on a track.

4.1 SHARP Baseline Updating Procedure: Fixed Number of Looks

The SHARP baseline procedure averages discriminants from all of the previous looks at a track and then uses the minimum averaged discriminant from each class to determine whether to classify or reject a track. Figure 2 shows a flow diagram of the process. An HRR measurement is made on a track and passed to the Match / Index block where minimum discriminants are computed for each class according to the procedures outlined in Section 3. The minimum discriminant for each class is then added to the track’s discriminant list for each class. Once the specified number of looks at the track has been achieved, the discriminant lists for each class are averaged. The minimum averaged discriminant over all of the classes is then compared to a threshold level to determine whether to classify or reject the track. This procedure is easily implemented, and discriminants are efficiently fused through averaging.

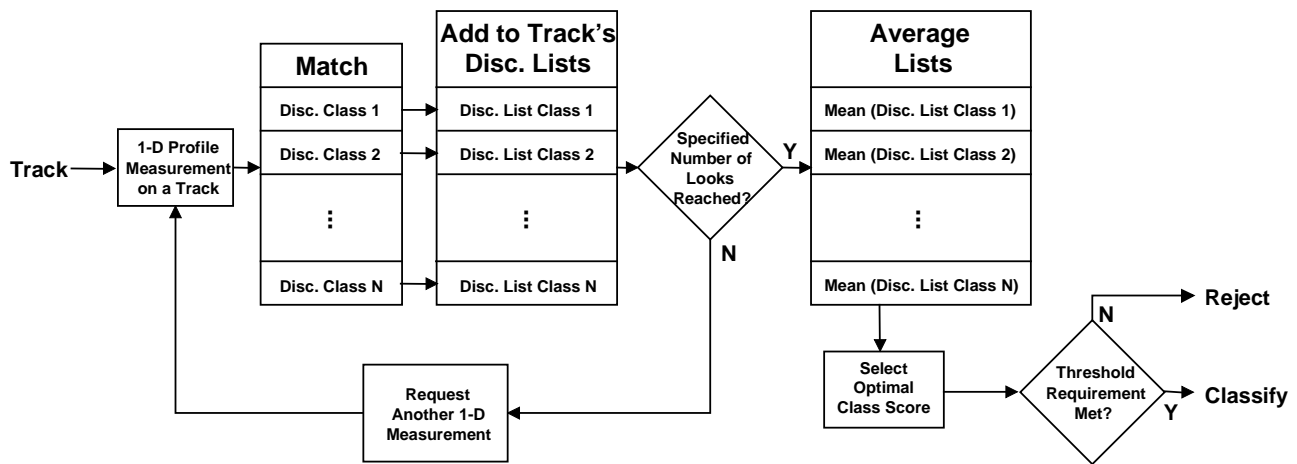


Figure 2. SHARP baseline multiple-look updating.

4.2 Bayesian Updating Procedure: Fixed Number of Looks

Another approach to fusing multiple measurements is to update them according to Bayes rule:

$$P_n(\omega_i) = \frac{P_{n-1}(\omega_i)p(s | \omega_i)}{\sum_{k=1}^N P_{n-1}(\omega_k)p(s | \omega_k)}$$

where ω_i is class i , N is the number of classes in the template database, s is the discriminant score, $P_n(\omega_i)$ is the probability that the track belongs to class ω_i after n updates, and $p(s|\omega_i)$ is the likelihood of getting discriminant score s given the track belongs to class ω_i .

For the ATR problem, the possibility for measurements on unknown targets exists. Thus, an additional unknown class can be added and updated according to Bayes rule in a similar fashion:

$$P_n(\omega_i) = \frac{P_{n-1}(\omega_i)p(s | \omega_i)}{\sum_{k=1}^N [P_{n-1}(\omega_k)p(s | \omega_k)] + P_{n-1}(U)p(s | U)}$$

$$P_n(U) = \frac{P_{n-1}(U)p(s | U)}{\sum_{k=1}^N [P_{n-1}(\omega_k)p(s | \omega_k)] + P_{n-1}(U)p(s | U)}$$

where U represents the unknown class.

In using Bayes rule, initial *a priori* probabilities ($P_o(\omega)$) must be set. If the densities of ground vehicles are known for a given scenario, these densities can be used to calculate the initial probabilities. Typically, however, this is not the case, and some other technique must be used.

For Bayes rule to behave appropriately, measurements must be independent. For fusion of HRR measurements, independence is typically achieved by updating with new measurements only if they are separated from previous measurements by at least the corresponding angular separation at which scatterers decorrelate (e.g., ~ 0.3 degrees for ground target sized objects at X-band).

Figure 3 shows a flow diagram of the process. The flow is similar to that of the baseline approach, except that after passing a measurement through the match block, conditional likelihoods are calculated for each class. The conditional likelihoods are then used with the prior probabilities to update the probability that the track belongs to each class. After the specified number of looks has been achieved, the track is classified as the class that produced the highest updated probability. Note that because an unknown class probability has been updated along with the other N classes, a track will be classified as unknown (i.e., rejected) if the unknown class produces the highest updated probability.

The main advantage of using a Bayesian updating technique is the potential for better performance because more separable measurements are implicitly “rewarded” through the normalization term. For example, suppose that the initial priors are set to a uniform value over all classes ($P_o = 1/N$ for all N classes). If an HRR measurement on a track results in a high likelihood of belonging to one particular class, and low likelihoods of belonging to all other classes, the denominator term will be only slightly larger than the numerator term, resulting in a high probability that the track belongs to that particular class. Another advantage is that since outcomes are probabilities, measurements from other sensor types can be easily fused (so long as those measurements have been converted to probabilities as well). The main disadvantage of a Bayesian technique is that it requires the computation and storage of statistical parameters from which to compute conditional likelihoods ($p(s|\omega)$).

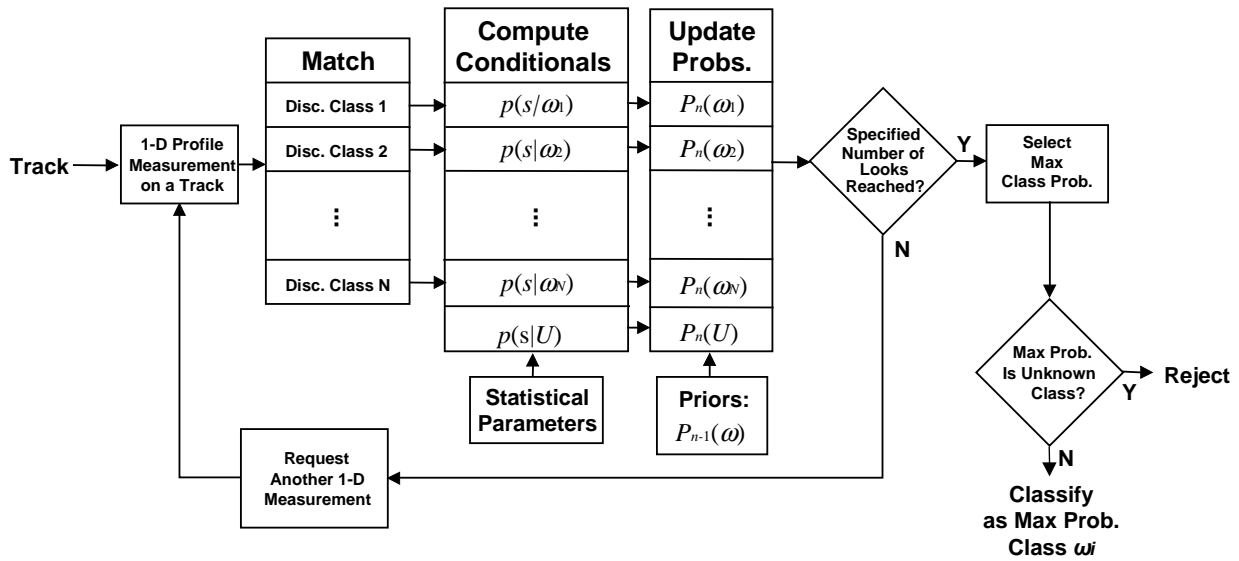


Figure 3. Flow diagram for Bayesian updating.

4.3 Bayesian Updating Procedure: Variable Number of Looks

The two previously discussed updating procedures were described under a “fixed number of looks” architecture in which a track is updated with a specified number of measurements, and then a “classify or reject” decision is made on the track. In general, tracks on more separable targets will require fewer measurement updates to reach a probability level necessary for classification. Thus, the thresholding logic can be modified to make a decision to either classify, reject, or make another measurement on a track after each update.

Figure 4 shows the flow diagram for a “variable number of looks” Bayesian updating system. Note that after each update is made on a track, the maximum class probability is compared to a classification threshold. If the classification threshold is surpassed for one of the target classes, the track is classified. If classification threshold is surpassed for the unknown class, the track is rejected (i.e., classified as unknown). If neither of these events occur, another measurement on the track is requested.

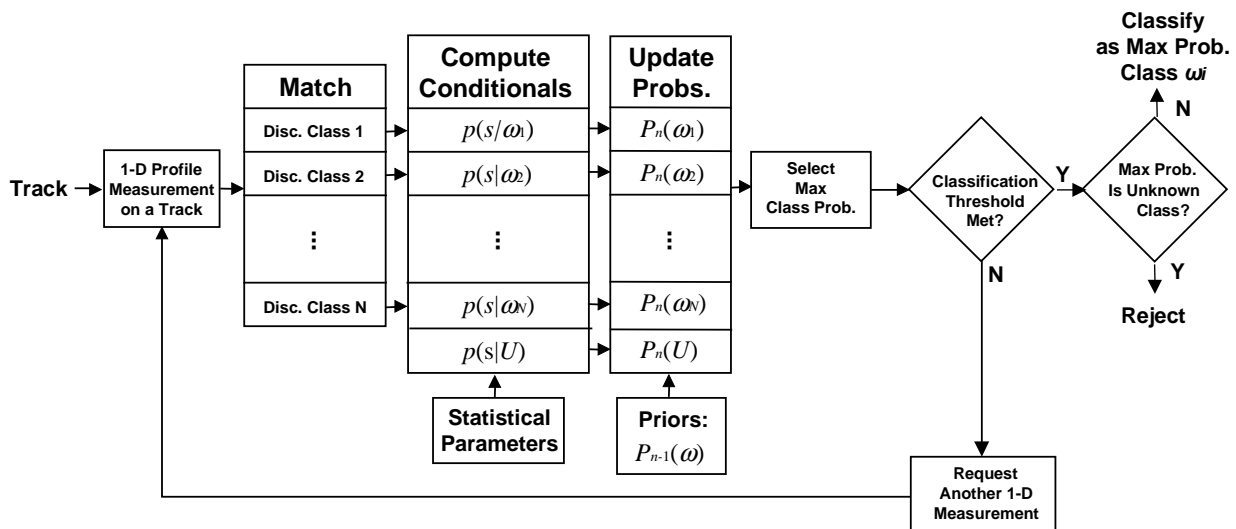


Figure 4. Bayesian updating with variable number of looks.

Such a modification could potentially save radar and ATR system resources by eliminating HRR measurements on tracks that can be classified or rejected more rapidly. The downside is that classification threshold levels must be specified.

RESULTS

ATR performance results were generated for each of the three updating architectures. Note that ATR performance was assessed according to the baseline SHARP procedures outlined in reference 1.

5.1 Data Set

To evaluate the updating procedures, HRR profiles were formed from MSTAR collection 1 target chips for the nine-class SHARP baseline target set at a depression angle of 29 degrees. The targets included: BTR70, M109, M110, M113, M1, M2, M35, M548, and T72. Note that these are all similarly sized military targets.

The MSTAR data set provides a large volume of high quality ground-truthed data for a wide selection of military vehicles. MSTAR data is at X-band frequencies with a range resolution of 0.3 meters. It should be noted that MSTAR data is in the form of 2-D SAR images of stationary targets. Hence, a process is needed to form representative 1-D HRR range profiles from the 2-D images. It should be noted that the MSTAR data does not capture HRR signature effects related to the movement of the target (rotating parts, structural vibrations, range walk issues, time-varying ground-plane interaction, etc.).

5.2 HRR Profile Formation

A process is needed for forming 1-D HRR profiles from 2-D MSTAR target chips. First, target pixels are segmented from clutter pixels to represent the suppression of ground clutter through Doppler processing that would normally occur for moving targets. Clutter segmented (complex) chips in range, cross-range space are (inverse) Fourier transformed in the cross-range dimension. This results in chips in range, angle space, with every angle sample representing a 1-D HRR profile measurement. Each MSTAR target chip produces 101 1-D HRR profile samples spanning an azimuth range of about 3.5 degrees (each sample is separated by approximately 0.035 degrees).

HRR templates are formed by non-coherently averaging profile samples that lie within 0.5-degree azimuth bins. HRR observations (measurements) are simply taken as single HRR profile samples from the range, angle chips (101 observations can be formed from each chip). It should be noted that MSTAR target chips from different looks of the sensor (different spotlight images) overlap in azimuth. This makes it possible to create templates and observations that have similar azimuth and elevation angles but are from independent looks of the sensor. Figure 5 illustrates this situation.

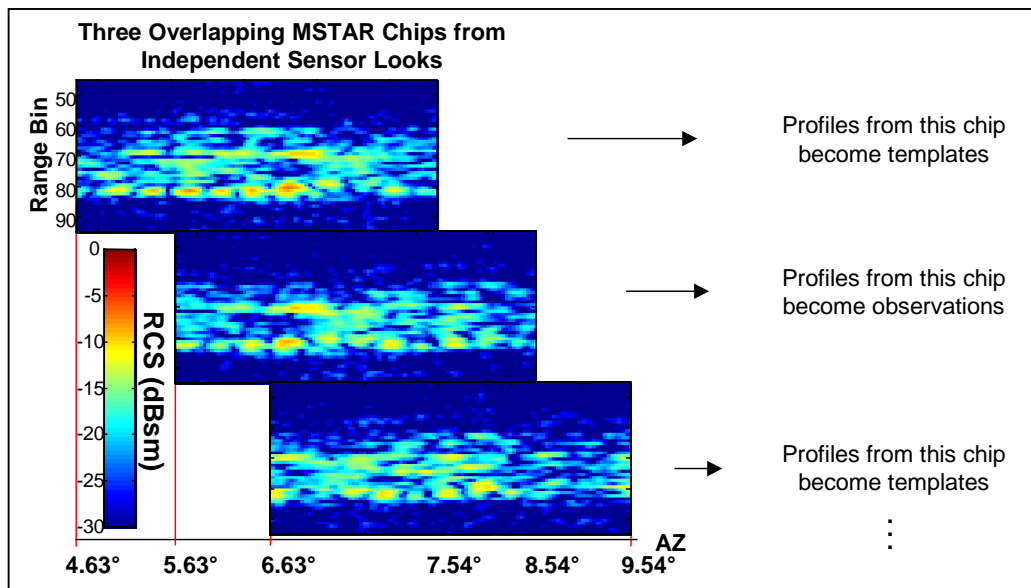


Figure 5. Making independent template and observation sets.

5.3 Leave One Out Method

To evaluate ATR performance in the presence of unknown targets (targets that are not in the template database), the Leave One Out Method (LOOM) is used. For an N -class ATR, a single class is left out of the template database. The resulting $N-1$ class ATR is then tested against observations from all N classes. When an observation is encountered that is from the same class as the template class being excluded, that observation is considered unknown. This process is repeated N times, each time leaving a different class out of the template database. The performance results are then computed by averaging over the N different $N-1$ class ATRs.

5.4 SHARP Baseline Updating: Fixed Number of Looks

Figure 6 shows the predicted ATR performance for the baseline SHARP updating procedure. In this case, looks were separated by 5.0 degrees. Note that performance increases with increasing look number, but that P_{fa} begins to bottom out at higher look numbers.

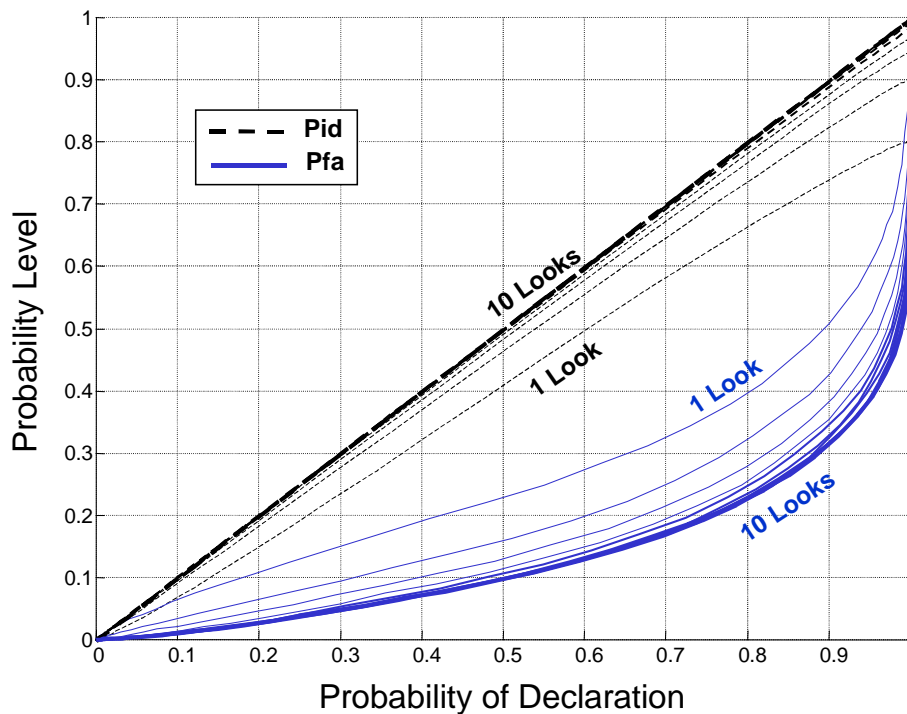


Figure 6. SHARP baseline performance, fixed number of looks.

5.5 Bayesian Updating: Fixed Number of Looks

Figure 7 shows the predicted ATR performance for the Bayesian updating procedure for the “fixed number of looks” architecture (details of the Bayesian implementation follow). The Bayesian procedure outperforms the SHARP baseline updating procedure as look number increases and the probability of declaration decreases. For example, at $P_{dec} = 0.9$, P_{fa} for the baseline method is 0.32 compared to only 0.08 for the Bayesian method after ten looks. The implicit separability weighting in Bayes rule accounts for the reduction in false alarms, as well as for the better performance at a lower P_{dec} . At a lower P_{dec} , the ATR can be more selective about which tracks it chooses to classify, and hence tracks that score well against only one template type and poorly against the others are the first to be classified. The baseline did outperform the Bayesian method for small look numbers (≤ 3) when P_{dec} was high (> 0.6). The explanation for this is the noise introduced in transforming discriminants into conditional probabilities for each look. This noise, however, is integrated out in a few looks and allows the Bayes procedure to produce much better results at higher look numbers.

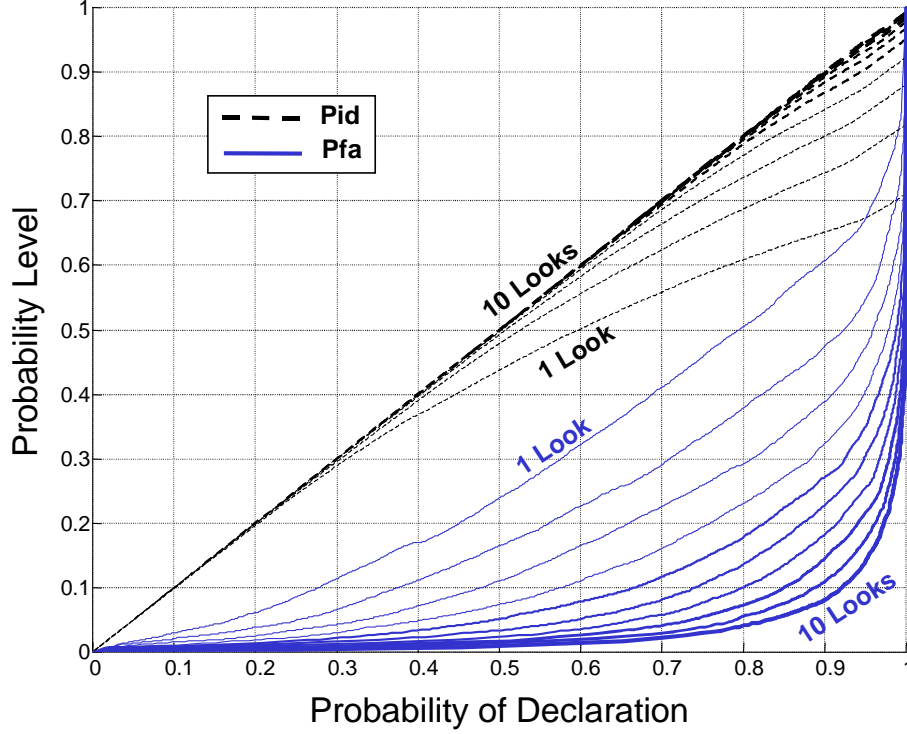


Figure 7. Bayesian updating performance, fixed number of looks.

5.5.1 Estimating conditional density functions

Conditional density functions for each template class ($p(s|\omega)$) are created by taking advantage of the fact that MSTAR target chips formed from independent spotlight images overlap in angle space. Figure 8 illustrates the process. HRR measurements that lie within the same azimuth limits as a given template are extracted from an independent MSTAR chip. Each HRR measurement is matched against the template, and a list of discriminant values is formed. To simplify implementation, distributions of discriminant lists are assumed Gaussian, and mean and variance parameters are computed. This process is repeated for every template in the template database. When an observation enters the ATR, it is matched against all templates from all classes within a specified azimuth window. The discriminant value from each template match is then converted to a density value (likelihood) using the Gaussian parameters for that template. The maximum density value is chosen for each class, and these maximum density values become $p(s|\omega)$ for each class ω .

Estimation of the conditional density function for the unknown class ($p(s|U)$) is typically a difficult process. This function is dependent on the distributions and different types of ground vehicles that move within a given region, and hence will change from scenario to scenario. In light of this, a simple (albeit *ad hoc*) technique for computing $p(s|U)$ is implemented which is not reliant on scenario information. After the $p(s|\omega)$ values have been computed for each class, the highest likelihood value over all classes is chosen. The statistical parameters are then retrieved for the template that produced this highest likelihood value, and the maximum possible likelihood for this template is computed:

$$p_{\max} = \frac{1}{\sigma\sqrt{2\pi}},$$

where σ is the standard deviation stored for this template. The likelihood $p(s|U)$ is then computed according to:

$$p(s | U) = p_{\max} - p(s | \omega_k),$$

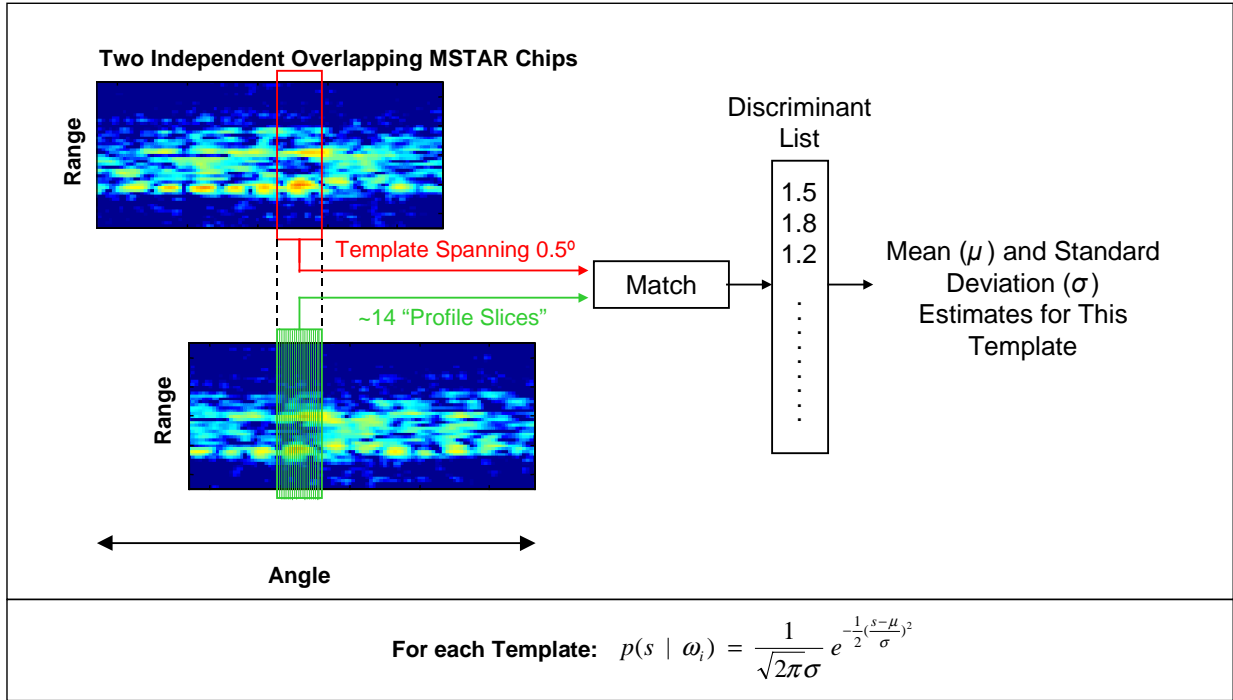


Figure 8. Estimating conditional density functions $p(s|\omega)$.

where k is the class index for the template producing the highest likelihood value. Essentially, this technique merely bases the likelihood that an observation belongs to the unknown class on the degree of likelihood associated with the observation belonging to class k . For example, if the likelihood of the observation belonging to class k is significantly lower than the maximum possible likelihood that an observation belongs to class k , then the likelihood that the observation belongs to the unknown class will be a larger value, and vice versa.

5.5.2 Initial *a priori* probabilities

For this study initial *a priori* probabilities were set uniformly for all classes:

$$P_o(\omega_i) = \frac{1}{N + 1},$$

where $N + 1$ is the number of template classes plus the unknown class.

5.6 Bayesian Updating: Variable Number of Looks

Figure 9 shows the predicted ATR performance for the Bayesian updating procedure with the “variable number of looks” architecture. For this case, a classification or rejection decision on a track was forced after 10 looks had been taken. Results show a slight improvement in both Pfa and Pid, especially at lower look numbers.

CONCLUSIONS

A Bayesian updating procedure was applied to the baseline SHARP ATR system, and performance results were compared to the baseline SHARP updating procedure. Results showed that the Bayesian method significantly improved ATR performance (particularly in Pfa), especially at higher look numbers. The Bayesian procedure does require the added complication of estimating and storing additional statistical parameters for each class. Further performance enhancement was achieved by incorporating an architecture that allowed a “classify, reject, or take another measurement” decision after each update.

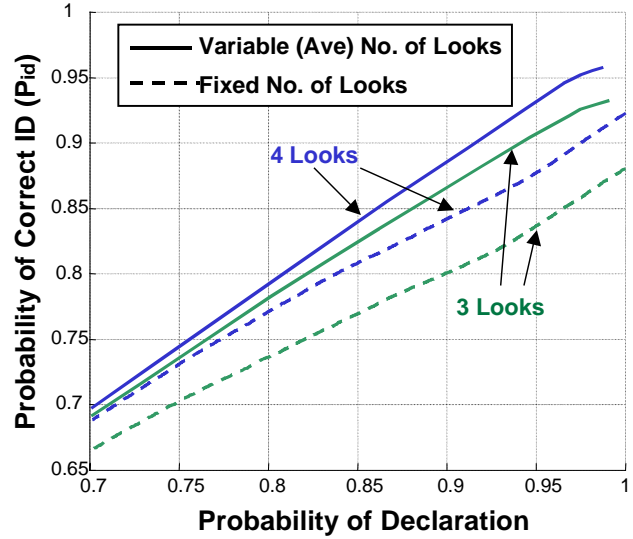
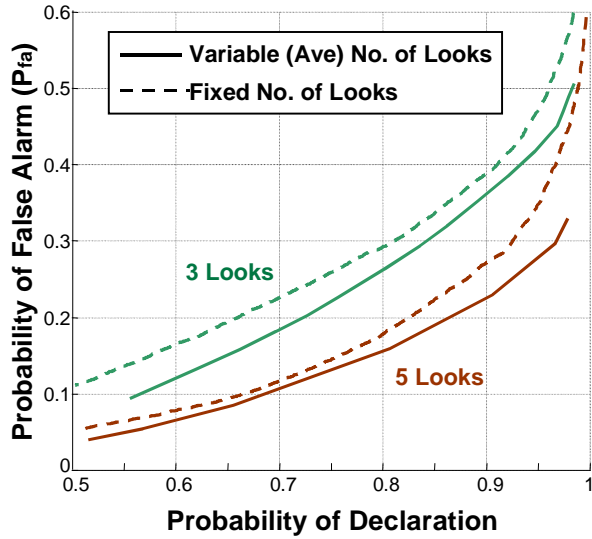


Figure 9. Bayesian updating performance, variable number of looks.

ACKNOWLEDGMENTS

This research was funded through a Phase II Small Business Innovation Research (SBIR) contract from the Sensors Directorate of Air Force Research Labs (AFRL/SN). The authors would like to thank Toyon employee Kevin Sullivan for his technical assistance and guidance.

REFERENCES

1. J. Westerkamp (AFRL/SNAT), "A Framework for 1-D HRR ATR Evaluation", 1998.
2. R. Williams, J. Westerkamp, D. Wardell, S. Worrell, and M. Ressler, "Robustness Issues for 1-D ATR", Proceedings of the Sixth ATRWG Symposium, Redstone Arsenal, AL, October 1997.
3. B. Gragg, "High-Confidence Target Classification Using Multiple Measurements", Strategies for Locating and Targeting Time Critical Targets, Annual Report, 30 September 1995, Toyon Research Corporation (T151D0900).
4. K. Fukunaga, *Statistical Pattern Recognition*, Second Edition, Academic Press, 1990.
5. E. Rich and K. Knight, *Artificial Intelligence*, McGraw-Hill, Inc., 1991.