Feature-aided tracking of ground targets using a class-independent approach

Kevin Sullivan, Craig Agate, and David Beckman
Toyon Research Corporation, 75 Aero Camino, Suite A, Goleta, CA 93117

ABSTRACT

We have developed and implemented an approach to performing feature-aided tracking (FAT) of ground vehicles using ground moving target indicator (GMTI) radar measurements. The feature information comes in the form of high-range resolution (HRR) profiles when the GMTI radar is operating in the HRR mode. We use a Bayesian approach where we compute a feature association likelihood that is combined with a kinematic association likelihood. The kinematic association likelihood is found using an IMM filter that has onroad, offroad, and stopped motion models. The feature association likelihood is computed by comparing new measurements to a database of measurements that are collected and stored on each object in track. The database consists of features that have been collected prior to the initiation of the track as well as new measurements that were used to update the track. We have implemented and tested our algorithm using the SLAMEM™ simulation.

1. INTRODUCTION

Current airborne surveillance radars primarily employ low-range-resolution ground moving target indicator (LRR GMTI) measurements to detect moving ground vehicles and synthetic aperture radar (SAR) measurements to detect stationary vehicles. Human operators observe the LRR GMTI measurements over a period of time to track the location of groups of vehicles, or in some cases, individual vehicles. Humans typically have difficulty tracking more than a few vehicles in a region. Automated tracking algorithms that associate a sequence of LRR GMTI measurements collected over time to form tracks, have been developed and deployed in some situations. Unfortunately, these algorithms have not been able to maintain the identity of a vehicle in track over an extended period of time in situations where there are several maneuvering vehicles that are closely spaced.

Modern radars are being equipped with a high-range-resolution GMTI mode that can provide a kinematic measurement as is done with LRR GMTI, but in addition, a range profile can be obtained that provides a signature of a target vehicle. This can provide the ability to ID targets using the HRR GMTI measurements and it allows for feature-aided association of GMTI measurements. Targets can be identified by comparing measurements collected on them to a library of stored signatures that were previously collected on a set of target types. As long as the targets encountered in operation are similar to the ones used to collect the a priori database, this technique works fairly well [7]. Performance is increased by using multiple measurements that are collected over time and fused to accumulate evidence about the identity of the vehicle. Of course, in order to effectively fuse these profiles, they must have originated from the same vehicle and this requires that they be correctly associated.

GMTI measurements can be better associated when HRR profiles are available by using the signatures and the kinematics to compute an association likelihood using feature-aided tracking (FAT). The signatures can be used to compute association likelihoods in two different ways. First, they can be used to classify a vehicle as belonging to one of a set of target classes and then comparing this classification with the classification of each vehicle in track (class-dependent FAT) [8], [4], [9]. Second, the signature of a single detection can be directly compared to stored signatures of vehicles in track (class-independent FAT).

We have built feature-aided tracking algorithms that use kinematics as well as signatures to associate HRR GMTI detections. This required the development of a ground target tracker that used road information as part of its knowledge base. We have built an interacting multiple model (IMM) tracker that contains three motion models, namely, off-road, on-road, and stopped. The on-road model assumes that the target is traveling on the road at a constant speed and maneuvers to follow the path of the road. The off-road model assumes that the target does not maneuver. The

Signal Processing, Sensor Fusion, and Target Recognition XIII, edited by Ivan Kadar,
Proceedings of SPIE Vol. 5429 (SPIE, Bellingham, WA, 2004)
0277-786X/04/$15 · doi: 10.1117/12.541080

likelihood that each model is valid at any given time is computed by comparing the predicted location of a track using a given motion model to an associated measurement. Our approach to developing this algorithm is described in Section 2.

We have analyzed the performance of our algorithm by modifying the SLAMEM^{TM} simulation so that it can serve as a testbed for the FAT algorithms. Additionally, we have evaluated the performance of our approach using this testbed. We considered a scenario where several vehicles travel together in convoys undergoing regular maneuvers such as turning around on themselves and stopping and starting. We examined our ability to maintain track on these vehicles using kinematics only and when performing FAT. Our approach and the results of our analyses are presented in Section 3.

2. FEATURE-AIDED TRACKING APPROACH

In the traditional approach to the data association problem, an estimate of the target's state (kinematic values such as position, velocity, etc.), is used to predict the value of the next sensor measurement of that target. Given a set of measurements from the sensor, the measurements are assigned (associated) to the various tracks taking into account the *likelihood* that each measurement could have originated from each track. The likelihood is typically a kinematic measurement likelihood, meaning that the elements of the measurement vector are related to the target motion or position parameters and are independent of the *signature or features* of the target itself.

In many tracking scenarios, the signature or feature vector, is used to classify the target (i.e., identify (ID) the target). In this effort we propose a data association algorithm which utilizes the target signature in addition to the kinematic measurements to aid the association of measurements to tracks. The approach taken here parallels that of multi-hypothesis approaches (such as the N-scan algorithm [1], the probabilistic data association (PDA) algorithm [2], or Joint probabilistic data association (JPDA) algorithm [3]) in the sense that hypotheses are formed assigning measurements to tracks. Then, the probability that each hypothesis is correct is calculated. The inclusion of a signature or feature likelihood in the hypothesis probability calculation is what marks the difference between the current approach and kinematic-only approaches. Many of the details of our approach can be found in [4]. Depending on the various target types present, utilization of target signatures should significantly improve data association performance over traditional data association methods.

We have developed an approach to performing FAT that uses both a class-independent signature matching process and a class-dependent process. The class-dependent process is used in the early stages of the track and it is eventually replaced by the class-independent process as the "on-the-fly" database of signatures develops. This process can be better described by considering Figure 1. The figure shows that a measurement is first fused with terrain and road data to get a better estimate of the aspect of the target. The road information is only used to help with the aspect determination if the onroad motion model has a sufficiently high probability. The better aspect estimate that can usually be obtained by doing this allows for a better calculation of the class-dependent and class-independent feature likelihoods. The class-dependent likelihood is computed using a fixed database of signatures that were collected for each target class prior to the initiation of the track. It is used to update the probability that the object being tracked belongs to one of a fixed number of classes regardless of whether or not it is used to help with data association. The class-independent likelihood is computed using a database of signatures that was constructed only after the start of the track. The database is searched for signatures that are within a specified aspect angle of the sample. If enough signatures are available, a statistical model of the signatures is built to provide a class-independent feature association likelihood. If insufficient signatures are available, only the class-dependent likelihood is used to provide a feature association likelihood.

2.1 Kinematic association likelihood

Toyon has developed an interacting multiple model (IMM) tracker. The multiple models represent different motion models for the targets. We consider an off-road model, an on-road model, and a stopped model. For each of these models, an extended Kalman filter (EKF) is employed to estimate the state of the target assuming that the target is undergoing motion as described by the model. The state estimate for the track is made by weighting each of the state estimates from the individual models. For a complete description of this process refer to [5]. The weighting of each model is made based on how well a measurement agrees with the predicted state of each motion model. Each of the motion models uses a state vector, $\boldsymbol{x}$, that comprises the position and velocity of the target in a local Cartesian coordinate

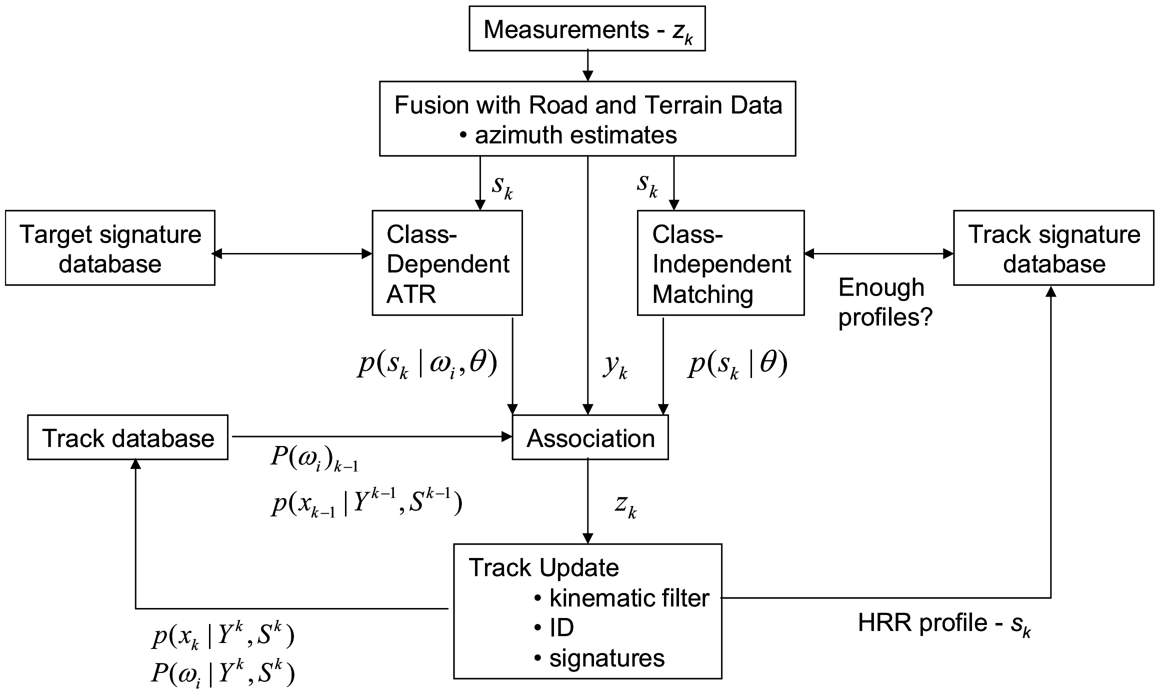

Fig. 1. Block diagram of hybrid feature-aided association process.

system. The differences between the models in our tracker result from differences in the way a track is propagated ahead in time and how it is updated given a new measurement. These differences are briefly described below. Each of the models assumes a dynamic system with a target whose state, $\boldsymbol{x}(t_k)$, is propagated using

$$x(t_k) = F(\delta_k)x(t_{k-1}) + G(\delta_k)v(t_{k-1}) \tag{1}$$

where δ_k is the difference in time between the measurement at time t_k and the measurement at time t_{k-1}, $v(t_{k-1})$ is white Gaussian process noise with covariance Q, F is the transition matrix whose values depend on the particular motion model employed, and G is a matrix that maps the process noise to the state variables (also a function of the motion model employed).

Off-road model

For the off-road model, we assume that the target continues at its current speed in the direction indicated by its current velocity vector. To estimate a future covariance, we assume the process noise is represented by a second-order piecewise-constant model. Thus, the matrices F (transition matrix) and G (process noise model) are given by:

$$F = \begin{bmatrix} 1 & \delta_k & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \delta_k \\ 0 & 0 & 0 & 1 \end{bmatrix}, \qquad G = \begin{bmatrix} \delta_k^2/2 & 0 \\ \delta_k & 0 \\ 0 & \delta_k^2/2 \\ 0 & \delta_k \end{bmatrix} \tag{2}$$

The propagation of the state and covariance are then given by

$$\hat{x}(t_k^-) = F(\delta_k)\hat{x}(t_{k-1})$$
$$P(t_k^-) = F(\delta_k)P(t_{k-1})F(\delta_k)' + G(\delta_k)Q(\delta_k)G(\delta_k)' \qquad \text{(3) and (4)}$$

Given this predicted state and covariance, the standard EKF equations are used to compute an association likelihood and to update the future state given an associated measurement.

On-road model

For the on-road model, we assume that the target continues at its current speed, but that it changes direction in accordance with the known road network. Thus, its motion is constrained to be piece-wise linear as described by the road network. Mathematically, this can be thought of as a sequence of transition matrices F_j, that are applied at each road segment that the target state is expected to traverse during δ_k. Furthermore, the velocity vector of the track is rotated at each road segment, producing a propagation model that is given by

$$\hat{x}(t_k^-) = \prod_{j=0}^{Number\ of\ segments} F_j(\delta_j)R_j\hat{x}_j(t_j) \tag{5}$$

where F_j is the transition matrix described in the road-motion model, δ_j is the time that will be spent on road segment j, R_j is a rotation matrix that rotates the velocity vector from the direction aligned with the previous segment to a direction aligned with segment j, and t_j is the time at the start of segment j.

To estimate a future covariance, we assume an along-road acceleration and a cross-road acceleration are possible. These accelerations define the size of the predicted covariance – the orientation of the covariance is modified each time the predicted state switches to a different road segment. Similarly to the state update, this is given by

$$P(\delta_k) = R(F(\delta_k)P_{k-1}^r F(\delta_k)' + G(\delta_k)Q^r G(\delta_k)')R' \tag{6}$$

where P^r is the covariance matrix in the along-road and cross-road directions, Q^r is the acceleration matrix in the along-road and cross-road directions, and R is a rotation matrix that rotates the covariance from the along-road and cross-road directions to the direction of the road segment that the track is on at time t_k.

In addition to using road information for the propagation of the state and covariance, the road information is also used to modify GMTI measurements. The GMTI measurements are fused with the road network information to create a measurement that has a position that is influenced by the road with a reduced covariance. This is done by first computing all of the road segments that lie within a specified likelihood cutoff of a GMTI measurement covariance value as shown in Figure 2. For each of these segments, the point on the road segment with the maximum value of the measurement covariance is calculated. Each of these locations is then treated as a hypothesized location of the target that originated the measurement. Each hypothesized location has a corresponding covariance associated with it that is determined by the original measurement covariance and the expected cross-road errors in the road network. When computing the association likelihood of a measurement to a track, the likelihood is computed using the standard EKF equations for each of the hypothesized target locations and covariances. The hypothesis with the greatest likelihood is chosen as the value for the target to measurement pairing. This hypothesized measurement location and covariance are also used to update the track once a measurement has been associated to it.

Stopped model

For the stopped model, we do not propagate the state and covariance of the target. We assume a very small process noise to allow for some updating of the target's state with additional measurements. The velocity is always assumed to be zero even if the position changes after a measurement update.

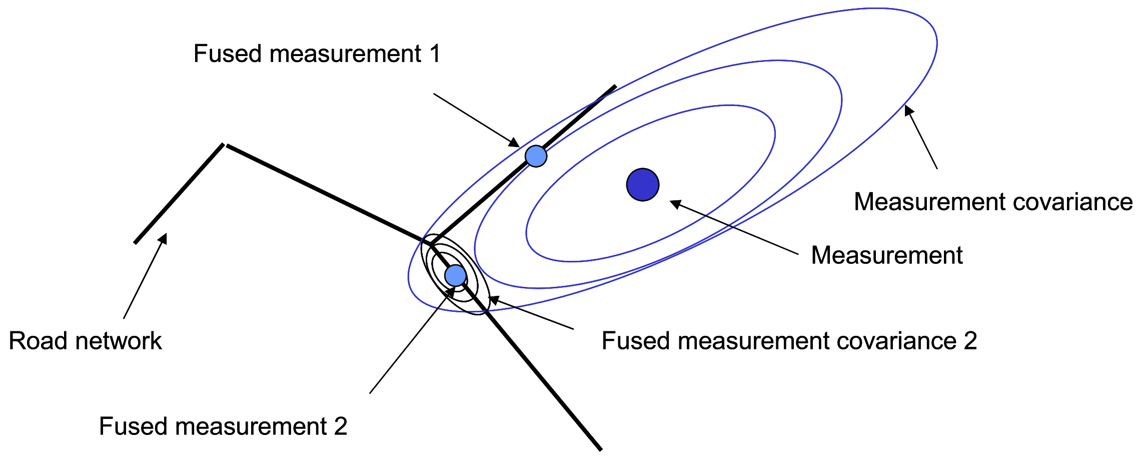

Fig. 2. Road-clamping illustration.

2.2 Feature association likelihood

Once we have feature measurements from a set of vehicles, we need an algorithm for deciding which measurements belong to which tracks based on the features themselves. One approach to calculating the feature likelihood is to compute a correlation coefficient. A second approach, investigated here, is to use one set of profiles, a *template* or *training* set, to build a statistical model for the feature, and compute the likelihoods of the features in the second set, the *test* set, given this model. When performing class-dependent FAT, the templates are created using a set of profiles that are collected prior to the initialtion of a track. For simulation purposes, we separated a set of profiles for this from the set of profiles used to simulate measurements collected while tracking vehicles. When performing class-independent FAT, the templates were created from the same set of profiles that were collected after the initiation of the track. These signatures do not have labels and thus the tracker must decide which profiles belong in the correct database of each track.

The statistical model we have used is based on the beta distribution

$$P_{\mathrm{B}}(x \mid \alpha, \beta) = \begin{cases} \dfrac{x^{\alpha-1}(1-x)^{\beta-1}}{\mathrm{B}(\alpha,\beta)} \equiv \dfrac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1} & : \text{for } x \in (0,1) \\ 0 & : \text{otherwise} \end{cases} \tag{7}$$

with two parameters $\alpha, \beta > 1$ and with mean and variance

$$\mu(\alpha,\beta) = \frac{\alpha}{\alpha+\beta}, \qquad \sigma^2(\alpha,\beta) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}. \tag{8}$$

The beta distribution is nonzero only in the bounded region (0,1) and hence is a reasonable model for normalized profiles such as the ones we are considering here. (For unnormalized profiles a different distribution would be required, to account for potentially unbounded returns.) We model the profile as a product of independent beta distributions, one for each range bin. From the template data set we fit the parameters α_r, β_r for each range bin, either by finding the maximum-likelihood estimator or (for a suboptimal but quicker solution) by inverting the equations for $\mu(\alpha,\beta)$ and $\sigma^2(\alpha,\beta)$. (These two methods produced comparable results in tests.) Then we compute the likelihood of each test profile given this statistical model; that is, from each set of template data (derived from a single target over a small range of aspect angles) we derive a statistical model for the HRR profiles in that set,

$$Y(\mathbf{y} \mid (\alpha),(\beta)) = \prod_r P_{\mathrm{B}}(y_r \mid \alpha_r, \beta_r) \tag{9}$$

For each profile **y** to be tested we compute, for selected templates *t* and aspects Φ, the likelihoods $\Lambda(\mathbf{y} \mid t, \Phi) \equiv Y(\mathbf{y} \mid (\alpha)_{t,\Phi}, (\beta)_{t,\Phi})$.

In these investigations, a separate model is developed for each set of profiles spanning a 3º azimuth range. This value was chosen to be large enough to have good sample population, while still being small enough that profiles do not decorrelate too much over the set. With more data, allowing for smaller azimuth bins, these model regions could be further reduced in size, probably improving model fidelity and resulting performance. In a realistic target-identification scenario, the target aspect may not be known to within three degrees. In such a situation the ATR algorithm must compare a profile to models in a range of azimuths corresponding to the uncertainty in the azimuth measurement uncertainty; a method for doing this is described in [6].

Note that the distribution used to represent the signature pdf is sensitive to range alignment between the profiles in the two sets. Because of the assumption that all range bins have independent distributions, this model cannot account for relative shifts between the template and test data. Poor alignment throughout the data set will effectively act to smear the profile model in range, effectively lowering the range resolution. But with poor alignment between the template and test profile sets, the correct model will not in general form a good match to the profile.

The initial range-alignment techniques used here are fundamentally centroid-alignment techniques. The centroid $C(y_r)$ of a profile y_r is defined as

$$C(y_r) \equiv \frac{\sum_r r y_r}{\sum_r y_r} \tag{10}$$

A set of profiles is *(linearly) centroid aligned* by computing the centroids of each profile, then shifting each profile in range so that the centroids occur at the same range value *R*. Our centroid alignment is only performed to within one range bin: that is, only range shifts by an integer number of bins, round $[R - C(y_r)]$, are performed. Following centroid alignment, a test profile is shifted past a template to compute the maximum likelihood over all shifts. This maximum value is then used as the feature association likelihood.

3. SIMULATION RESULTS

3.1 Feature-aided tracking testbed

We have developed a feature-aided tracker that uses HRRGMTI measurements from airborne or space-based radars to aid data association. In order to develop, test, and demonstrate this module, we have built a testbed. The testbed was built using the existing SLAMEM™ simulation. This testbed operates on a PC. It contains models of all the elements required to study this problem including sensors, aircraft, ground vehicles, terrain, and exploitation systems. We use the Ground Vehicle Simulator (GVS) to simulate the motion of ground vehicles. The state of each vehicle along with terrain masking checks are used to determine if it is detected. The type of each vehicle is used to determine the possible classification outcomes. The detections and signatures are sent to the fusion and tracking module to estimate the type and location of each target. The fusion and tracking module uses a database of terrain, road data, and class signatures to better perform its job. For the analyses in this paper, the terrain was identical to what was used for ground truth, but the road network was not. The ground vehicles typically wavered away from the assumed road network by about 10-20 meters. The motion of the ground vehicles was determined by recording the location of actual vehicles using GPS signals. The recorded positions were then used to determine the state of the ground vehicles as a function of time in the simulation. The road network came from standard geographical databases.

3.2 Scenario

The scenario that we will consider for this analysis involves convoys of vehicles that maneuver in a variety of ways. Two of the convoys only have one vehicle in them. One of the convoys has four vehicles and the remaining thirty vehicles travel in three different convoys. The vehicles traveling in the convoys make maneuvers such as move-stop motion and circling back on themselves, making tracking difficult.

We simulate the observation of these vehicles by two aircraft with advanced radars. The aircraft travel at an altitude of 60 kft and fly about 120 km away from the center of the region of interest. The orbits of the aircraft are arranged so that one of the aircraft can observe the region while the other is turning. The revisit time of each sensor is about ten seconds, but due to the turns of each aircraft, only one aircraft is illuminating the region much of the time. We assumed the radar had an azimuth accuracy of 0.03 degrees, a Doppler accuracy of 0.1 m/s, and a range resolution of 1 ft. We assumed a minimum detectable velocity of the radar to be 1.0 m/s.

3.3 Association performance

We have evaluated the performance of our tracker using the SLAMEM™ testbed. First, consider the performance of the tracker when operating in a mode where only kinemaitc information is used to make association decisions. Figure 3 shows a plot of the performance of our IMM tracker when operating in this scenario. The plot uses a color for each vehicle to represent a detection that originates from that vehicle Each track is represented by a sequence of detections that are shown as a series of colors that go from the left side of the figure to the right. A track that is pure will only have a single color in it. A track that is impure will have many segments of different colors. Note that tracks 7 and 20 have solid colors that extend during the lifetime of the tracks. The colors of these tracks indicate that they were created by measurements from vehicles 21 and 22. These are the two vehicles in the scenario that travel alone. All other vehicles

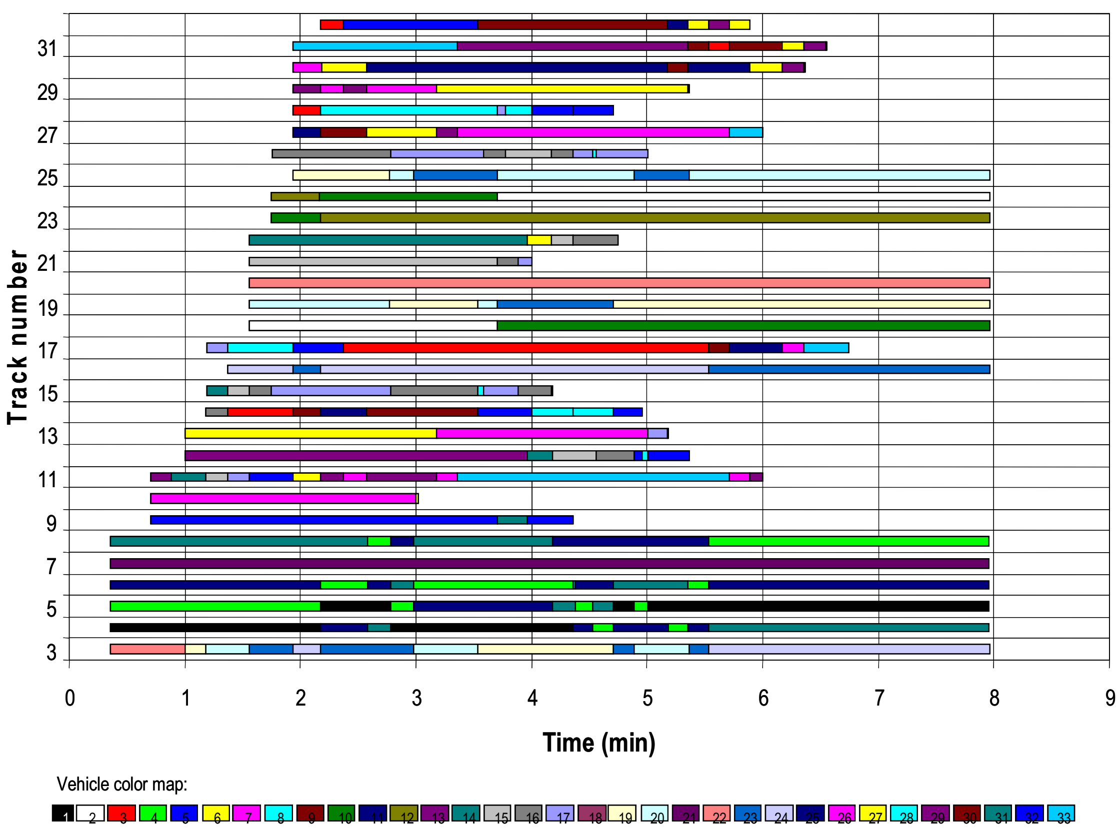

Fig. 3. Track to vehicle measurement mapping – kinematic-only association.

travel in a convoy. It is possible to track the solo vehicles perfectly using only kinematic measurements. Other tracks are not so pristine. For example, consider track 4 where the track starts out with measurements from vehicle 1, then is updated with detections from vehicle 5, then is updated with detections from vehicle 14, then continues to switch from vehicle to vehicle. This occurs because vehicle 1 is the lead vehicle in a convoy of four vehicles that are traveling closely together. Additionally, maneuvers are made that make association difficult and sometimes a further complication is that some of the vehicles are not detected in a scan of the radar. With only kinematics to make association decisions, errors are made. Similar behavior is observed in track 18 where it is originated with measurements from vehicle 2. The tracker is able to maintain a pristine track on this vehicle for about two minutes and then switching occurs. This is because vehicle 2 is the lead vehicle in a convoy of about twelve vehicles and after about two minutes, the convoy loops back around on itself, making data association difficult.

To simulate the use of signature information, we assumed that vehicles 1 and 2 were tanks and the rest of the vehicles were trucks. We used measured signatures from a specific type of tank to represent the signature of all simulated vehicles that belonged to the tank class. Multiple serial numbers of tanks were included in this dataset. Likewise, we used a specific type of truck to represent all trucks in the truck class. If we add the use of the signature data to the data association process, we get a new mapping of measurements to tracks as shown in Figure 4. Note that track 4 begins with measurements from vehicle 1 and is mostly updated with measurements from vehicle 1 during the lifetime of the track. There is some misassociation that occurs at around 4.5 minutes, but this is corrected fairly quickly. The misassociation occurs at a time when the vehicles are circling around themselves, making the kinematic association likelihood difficult to estimate properly. Also note that for the other tank in the scenario (vehicle 2), association performance is much improved. Vehicle 2 is initially tracked as track 18, but then switches to track 3 at about 3.7 minutes. After this switch is made, vehicle 2 is tracked as track 3 for the remainder of the scenario. Although the association performance is much improved for this vehicle, the improvement in terms of track duration are poor because of the track switch. In other Monte Carlo runs this switching did not occur for this vehicle and thus the average duration was significantly improved.

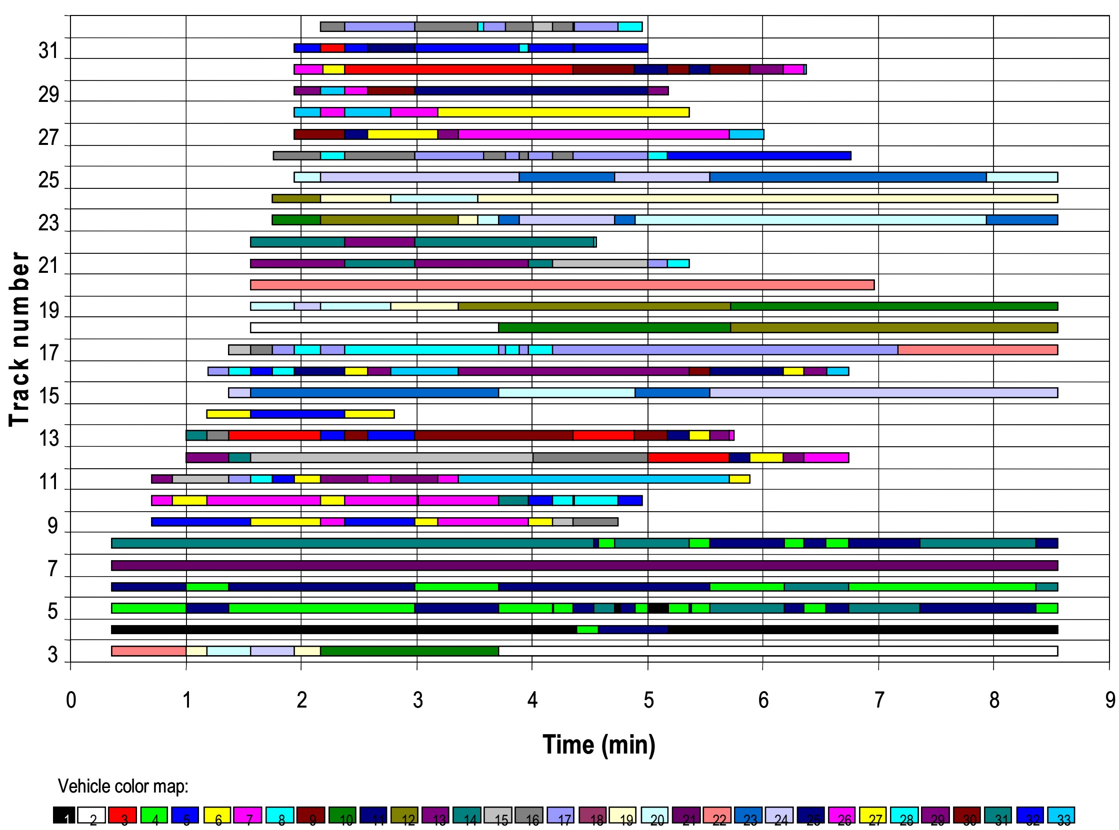

Fig. 4. Track to vehicle measurement mapping – feature-aided association.

For tracks that do not contain one of the tanks, the association performance is similar to the kinematic-only performance. This is because all of these vehicles have signatures that are from trucks and thus there is no benefit when trying to perform feature-aided association. In fact, there may be a small detrimental effect because we are altering the association likelihood assuming that we have some difference in features, when in fact, we do not.

We will now consider the duration of each track using two different metrics. A perfect track time is defined as the time measured from the track initialization of a target vehicle to the moment that a detection from some other vehicle is used to update the track. In other words, the perfect time metric is a measure of the track purity. The other measure of track duration that was used allows for track impurity in the following way. Here, the duration allows for some misassociations so long as they are corrected within a specified time. We assumed a four-minute time window for this purpose in the results presented here, thus, the duration represents the time from the track initialization of a target vehicle to the point at which the track has not been updated by a target detection in four minutes minus the four-minute window. Thus, for the four-minute window track duration, detections from other vehicles are allowed to update the track of the target vehicle as long as detections from the target vehicle are subsequently used to update the target vehicle's track within four minutes of being updated by the detection of another vehicle.

A plot of the track duration for ten Monte Carlo runs is shown in Figure 5 for tracks that are initiated with detections from vehicles 1 and 2. This plot is for a case in which only kinematic information is used to make association decisions. The left side of the plot assumes the perfect association metric is used and the right side of the plot assumes a four-minute window is allowed for association errors. Each bar indicates the track duration in a particular Monte Carlo run for either vehicle 1 or 2. Note that most of the time, the track duration is fairly short and the tracker is confused the first time the vehicles loop around on themselves. This occurs around 200 seconds into the simulation.

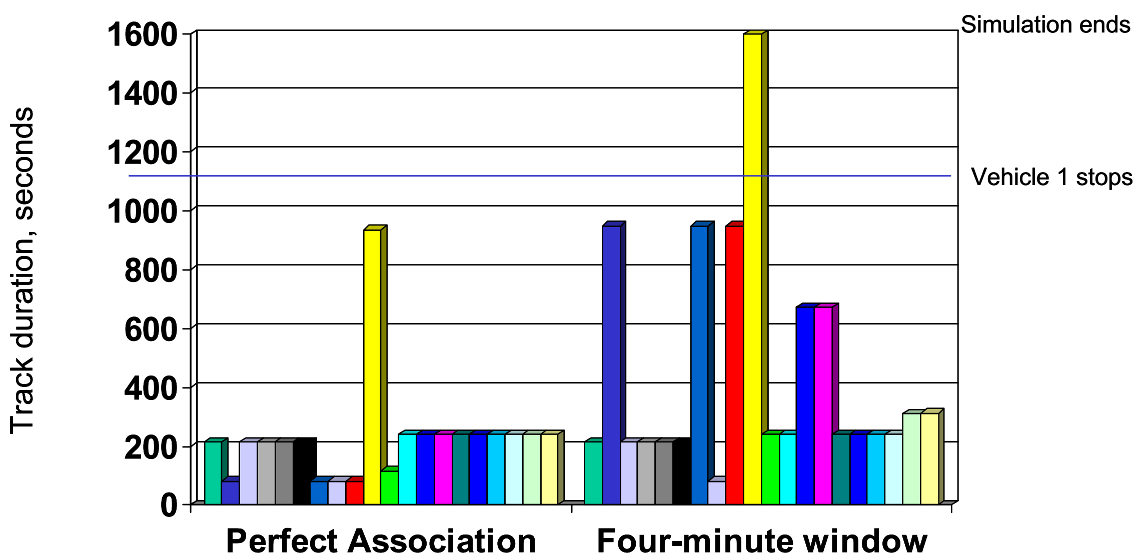

Fig. 5. Track duration for kinematic-only association.

If we add the use of features to aid the association process, the track duration increases as shown in Figure 6. This plot indicates the duration of the tracks that are initiated with measurements that originate from vehicles 1 and 2. Note that the duration of tracks with the perfect association criteria are not much longer than with the kinematic-only case, but that with the four-minute window, the duration of tracks is much longer. This effect was also noticed in the class-dependent analyses described in [2]. It indicates that with features, the tracker may still make mistakes, but they can be corrected and the identity of the track can be preserved. In fact, with the four-minute window, the track duration is capped by the end of the simulation and by the fact that vehicle 1 stops at around 17 minutes into the simulation. Since vehicle 1 receives no more measurements after 17 minutes, the track duration can not extend beyond this time.

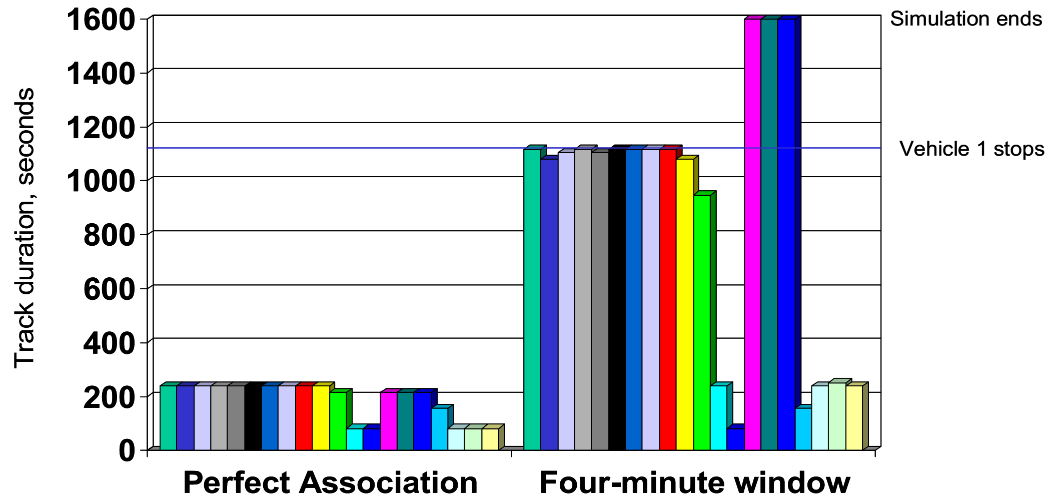

Fig. 6. Track duration – feature-aided association.

We also considered a case where the signatures from vehicles 1 and 2 were simulated using measurements from two different types of vehicles. Additionally these measurements were collected at a different location and time. The type of vehicle used to represent tanks was a tracked armored vehicle, while the type of vehicle used to represent the trucks was a wheeled armored vehicle. In this case, we used the same Beta ATR that was described in earlier sections of this report. For this case, we obtained the track durations shown in Figure 7. Note that performance is similar to what we obtained using the previous signatures. However, for the perfect association metric, the performance is actually a bit worse. This is due to the fact the Beta ATR (described earlier in this paper) has a large dynamic range of values that are created for signature association likelihoods. When the ATR makes an error, the effect can be catastrophic and the kinematic association likelihood can be wiped out.

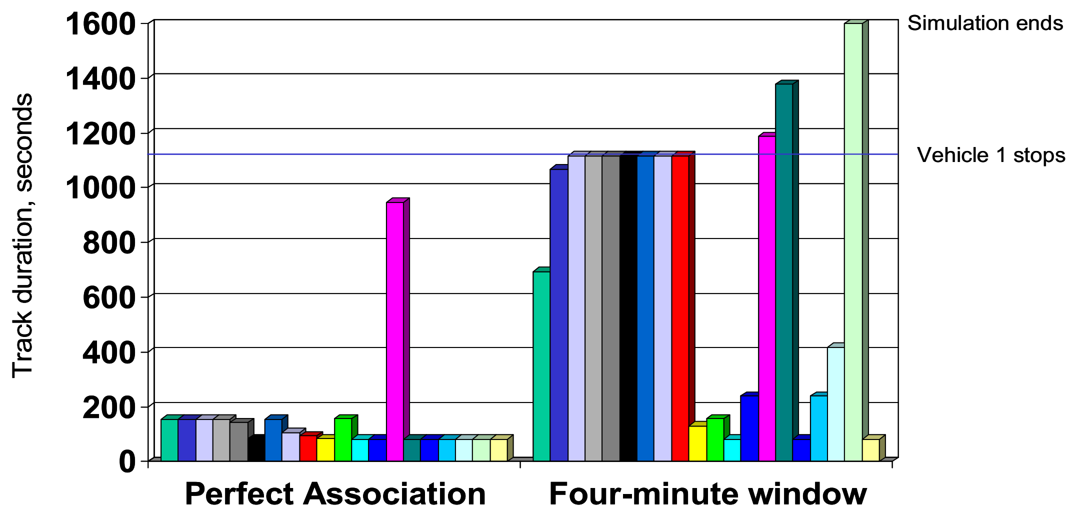

Fig. 7. Track duration – feature-aided association.

We have computed the median value of the track duration from each of the last three figures and plotted them in Figure 8. Note once again that with the feature-aided association process, the tracks are not perfect and there is little to no improvement when considering a perfect association metric. However, the tracker is able to maintain track on the correct vehicle if a small window of confusion is allowed when the vehicles are behaving in a way that makes tracking difficult. In this scenario, there were multiple times that the vehicles circled around themselves in a fashion that makes kinematic association unreliable, but with features, the proper identity of the vehicle was maintained for an extended period of time.

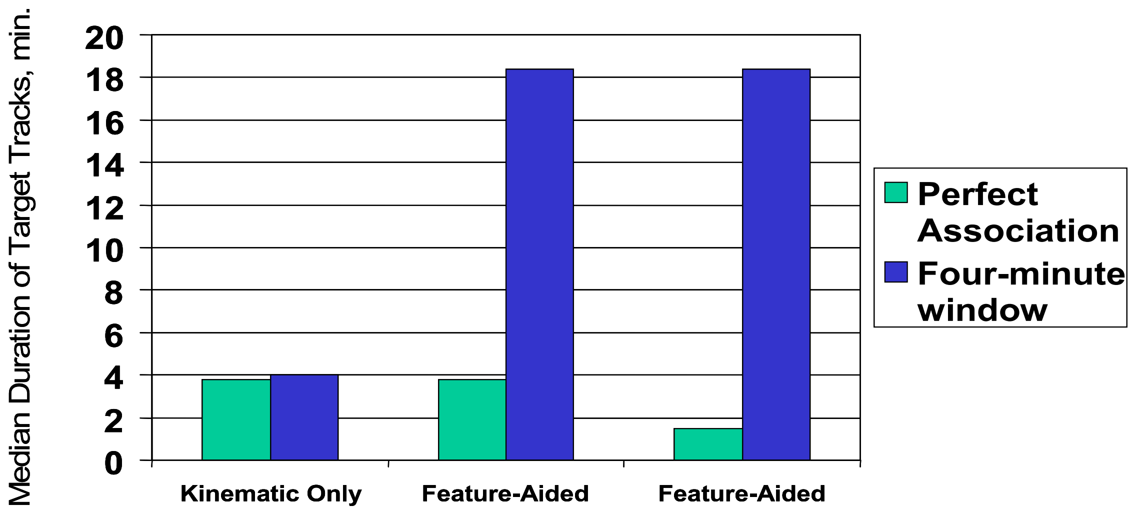

Fig. 8. Median track duration.

4. SUMMARY

We have developed an approach for performing feature-aided tracking (FAT) using HRRGMTI profiles of ground vehicles. Our approach uses a class-independent technique when sufficient profiles have been collected on a vehicle and a class-dependent technique when insufficient profiles have been collected. We tested our algorithm on a difficult scenario involving maneuvering vehicles that traveled in close proximity to each other. We used the SLAMEM™ simulation to model the sensors, ground vehicle motion, and terrain. We used a library of collected HRRGMTI profiles to simulate the signatures of the ground vehicles. An evaluation of the association performance of the tracker indicates that track duration can be significantly improved by using the HRRGMTI profiles.

ACKNOWLEDGEMENTS

The authors would like to thank Alan Wood (AFRL/SNA) for his technical input as well as his sponsorship of this research.

REFERENCES

1. R. Singer, R. Sea, and K. Housewright, "Derivation and evaluation of improved tracking filters for use in dense multitarget environments," *IEEE Trans. on Information Theory*, IT-20, pp. 423-432, July 1974.
2. T. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Sonar tracking of multiple targets using joint probabilistic data association," *IEEE Journal of Oceanic Engineering*, OE-8, pp. 173-184, July 1983.
3. Y. Bar-Shalom and T. Fortmann, *Tracking and Data Association*, Academic Press, Boston, MA, 1988.
4. Agate, C., Sullivan, K., "Signature-Aided Tracking Using Association Hypotheses," *Proceedings of SPIE AeroSense 2002.*

5. Y. Bar-Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, ISBN 0-9648312-0-1, 1995.
6. Beckman, D., Sullivan, K., Agate, C., et al., “Continuous Identification of Ground Vehicles Using HRR GMTI Measurements,” Toyon Research Corp., October 2003.
7. M. Ressler, R. Williams, D. Gross, A. Palomino. “Bayesian Updating Applied to the SHARP ATR System,” *Proceedings of SPIE AeroSense 2000 Conference on Algorithms for Synthetic Aperture Radar Imagery VII,* Volume 4053, 2000.
8. Bar-Shalom, Y., Kirubarajan, T., “Tracking with Classification-Aided Multiframe Data Association,” Proceedings of SPIE Conference on Signal and Data Association of Small Targets 2003, vol. 5204.
9. Blasch, E., “Derivation of a Belief Filter for High Range Resolution Radar Simultaneous Target Tracking and Identification,” Wright State University, 1999.