

Modeling and Resolving Tracking Ambiguities using Features for Long-Term Tracking

Craig S. Agate, Charlene S. Ahn and David E. Beckman^a

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

ABSTRACT

Tracking an entity for a long duration allows the gathering of intelligence on a target. While the system comprises a collection of different elements (e.g., tracking, sensor tasking, etc.), the ability to track objects continuously over long periods rests on feature measurements that are collected “on-the-fly” and used to uniquely characterize the target of interest. These features are then used to track the target over extended periods of time and through situations in which the targets can be confused with other moving objects. The collecting of features helps support tracking the target when it becomes kinematically ambiguous with other objects. If the system is unable to avoid ambiguities between the target of interest and other moving objects, features collected post-ambiguity can be used to resolve the ambiguities. A collection of algorithms that model and attempt to resolve any association ambiguity between a target of interest and the tracks in the fusion and tracking database is required to accomplish this task. This module is referred to as the Tracked Object Manager (TOM) and forms the backbone of a system for the continuous tracking of high-value targets. The TOM utilizes the collected features to help correct track switches and, if appropriate, stitch tracks together to maintain continuous track on high-value targets. The algorithms are being incorporated into and evaluated using Toyon’s Intelligence, Surveillance and Reconnaissance (ISR) simulation environment named SLAMEM.

Keywords: long-term tracking, high-range resolution profiles, feature-aided tracking

1. INTRODUCTION

A multi-target tracking system comprises several components that vary depending on the algorithmic approach. One of the most common and matured approaches is that based on the concept of tracks. In a track-based approach, all received measurements are partitioned into sets such that all measurements in a particular set are assumed to originate from the same object. Each set of measurements is processed to yield the probability density function (PDF) of the object state. Typically, the object state, which we desire to characterize from the partitioned measurements, includes position, velocity and possibly other object attributes such as class type. For on-line applications the process of partitioning is not done in one step after all measurements are collected (i.e., a batch process). Rather, the process is recursively implemented so that at any given time, the measurements collected over one or more scans (i.e., multiple frames of data) are partitioned and assigned to existing tracks, where an existing track refers to a collection of past measurements that are assumed to originate from the same object.

In this paper we will briefly discuss the difficulty in correctly assigning measurements to tracks, the consequences of association mistakes in the context of continuous tracking, and an approach to mitigating such effects.

1.1. Measurement-to-track Association

The first problem associated with a track-based approach is that the collected measurements are not *labeled* with their origin. Thus, a track-based approach to tracking multiple targets must employ an association algorithm to first assign measurements to existing tracks, collect measurements for initialization of a new track or ignore measurements as erroneous. Assignment algorithms all require some measure of how well a track’s state (i.e., the state of the object in that track) *matches* a given measurement. Typically, such a measure is based on a weighted distance between a predicted measurement, based on the track state and measurement model, and

Further author information: (Send correspondence to Craig S. Agate: cagate@toyon.com)

the actual measurement. Therefore, the assignment problem is difficult for cases in which two or more tracked objects are “close” in measurement space. Different approaches are employed for assigning measurements to tracks depending on the time over which decisions must be made. For example, in a single-frame approach to associating measurements to tracks, every frame of detections is assigned to tracks independent of other frames. A single frame of detections is defined as a collection of measurements such that no more than one measurement may originate from the same object. In multi-frame approaches, measurements are assigned to tracks by simultaneously considering multiple frames of data^{1,2} and, therefore, should be more robust to measurement outliers. The most complex approach to data association involves considering all possible measurement-to-track associations from the initialization of the system to the current time. Each possible assignment of measurements to tracks over a period of time is a global hypothesis for which the probability of being correct is calculated.^{3,4}

1.1.1. Continuous Tracking

Ideally, the association process correctly partitions the set of measurements into tracks and, therefore, achieves continuous tracking of all objects. Continuous tracking refers to the ability to track vehicles for an extended period without losing them, where losing a vehicle can result from several situations. For example, a vehicle is lost when its track is dropped from the fusion and tracking database, typically because no more detections were assigned to the track. Additionally, a vehicle may be lost due to a track switch, a situation in which after one or more measurement-to-track mis-associations, a vehicle becomes associated with a different track than the one with which it was previously associated. The underlying assumption in most tracking approaches is that the same object is associated with the same track over time since most data association algorithms emphasize correctly associating measurements to tracks. In general, no consideration is given to the coping with mis-associations; we simply accept that mistakes are inevitable and that little can be done.

When only kinematic measurements are available (i.e., measurements that depend solely on the location or velocity of the object), resigning oneself to the inevitable mis-association of measurements to tracks (and all the related consequences) is quite understandable since it is impossible to do anything about it. When multiple objects are physically close in the state space, the corresponding kinematic measurements of the objects are close in measurement space and, therefore, correctly associating measurements is a matter of chance. However, the availability of feature measurements from more modern sensors offers a potential solution since features are generally a function of non-kinematic object attributes such as the spectral properties of the object’s exterior, the object’s shape, etc. Thus, regardless of the proximity of objects in the state space, the feature measurements may be sufficiently separated in measurement space such that correct association of measurements to tracks is possible.

1.1.2. Feature-Aided Data Association

As alluded to previously, feature measurements offer the potential to facilitate the correct association of measurements to tracks since they depend on attributes of the tracked objects that are (only) weakly dependent on the proximity of the tracked objects. Many researchers have developed feature-aided data association algorithms. To utilize features in a Bayesian framework, some probabilistic characterization of the features is required since Bayesian methods rely on the availability of a measurement likelihood function. Typically, statistical analysis of collected feature data is used to develop PDFs or probability mass functions (PMFs) of the feature data conditioned on variables such as the object type or label and the aspect angle at which the feature data was collected (most features are aspect-dependent).

Developing statistical feature models (required for Bayesian-based feature-aided data association algorithms) depends on the conditions under which features are “collected.” We shall distinguish two approaches and consider only one in this work. The first approach to mention is what we refer to as *class-dependent* approach. In a class-dependent approach, an *a priori* database of features is available on *classes* of objects. For example, we may have a collection of features that characterize Toyota Tundras. The collection is not unique to any particular vehicle but to the class of vehicles known as Toyota Tundras. The second approach, and the one we shall focus on in this work, is what we refer to as object-dependent. In an object-dependent approach to feature-aided tracking (FAT), the features are collected “on-the-fly,” which means that we dynamically build a database of features on a *particular* tracked object while the object is being tracked. Ideally, the collected features uniquely characterize the object and, therefore, can be exploited to continuously track the vehicle.

In the following sections we discuss an approach to tracking vehicles of interest (VOIs) over extended periods of time. When a vehicle is designated as a VOI, it is essential that we continuously track the vehicle for an extended period of time. Generally, this is a challenging problem that requires both special fusion and tracking capabilities as well as intelligent sensor tasking. Our approach has been to develop a module that enhances a fusion and tracking system. We refer to this module as the Tracked Object Manager (TOM).⁵ The fusion and tracking center uses a simple single-frame association technique; that is, with every scan of data, the association algorithm computes a matrix of measurement-to-track likelihoods, which are processed by an Auction algorithm in order to make a “best” assignment of measurements to tracks. The TOM’s objective is to aid the fusion and tracking system in long-term tracking of particular vehicles. In order to support this process, the TOM has several responsibilities:

Feature Database Management Features are dynamically collected in order to uniquely characterize a vehicle to aid continuous tracking. Clearly, keeping the feature database *pure* (i.e., avoid adding features to the database that were actually collected from a different vehicle) is essential for reliable feature-aided tracking.

Model and Register Ambiguities As mentioned, the tracker uses a single-frame association technique. Track switches are an inherent part of single-frame association approaches. The TOM models such track switches as an ambiguity in associating vehicles to tracks and registers them so that we are aware that a track switch may have occurred.

Resolve Ambiguities After registering potential track switches, the TOM processes received feature measurements on the tracks, which may have switched vehicles, in order to *resolve* the potential switch.

Note that the approach makes a distinction between vehicles and tracks. Thus, the *track ID* is merely a temporary “label” associated with a vehicle. Vehicles have different track labels when, for example, they switch track with other vehicles or were associated with a track that was dropped and, subsequently, become associated with a newly formed track.

A block diagram illustrating the relationships between the various elements of an intelligence, surveillance and reconnaissance (ISR) system is shown in Figure 1. From the diagram, note that the Sensor Resource Manager (SRM) directs the sensors to collect information to be sent to an Exploitation Center for processing. The resulting information is then sent to the Fusion and Tracking System and used to track objects and identify objects. The TOM receives information from the Fusion and Tracking System and uses that information to monitor potential track switches, manage the object feature databases and, ultimately, to resolve vehicle-to-track association ambiguity. While not yet implemented, the TOM may also send information to the track to help with data association since the TOM manages the vehicle feature databases. Information on the various entities in track is sent to the SRM, which uses the information to formulate sensor tasks that will “improve” the fusion and tracking database. In the next section we discuss each of the TOM’s responsibilities in more detail.

2. THE TRACKED OBJECT MANAGER

Before discussing the algorithms composing the TOM, we first define two ambiguity states in which a track may be. The first ambiguity is the standard *measurement-to-track* association ambiguity with which all association algorithms must contend. The second ambiguity that we shall define is unique to our approach (which distinguishes vehicles from tracks). A situation in which there may have been a track switch results in what we call *vehicle-to-track* association ambiguity. Suppose we have two tracks each of which is associated with one vehicle. Furthermore, suppose the tracks were initiated and remained isolated so that we are certain that the same vehicle on which we initiated the track is the same vehicle associated with the track. When the tracked objects are widely separated in measurement space, then correctly associating measurements to tracks is easily accomplished. Assume, that the vehicles in the two tracks interact in such a way that associating the sensor measurements to the two tracks becomes difficult. Under these conditions, we say measurement-to-track association ambiguity exists. Now, suppose that the vehicles separate, leading to a period during which correct measurement-to-track




Figure 1. A general diagram of the information flow in an ISR system, and the Tracked Object Manager's location within the system.

association is, once again, a straightforward process. While measurement-to-track association is no longer ambiguous in such a situation, due to the prior measurement-to-track association ambiguity, we can no longer be sure whether the vehicles are still associated with the same tracks as they were prior to the period of ambiguous measurement-to-track association. When uncertainty exists as to which track a particular vehicle is associated, we shall refer to this as vehicle-to-track association uncertainty (or ambiguity). For an MHT approach, this uncertainty manifests as a collection of competing hypotheses of similar probabilities that can be resolved using only features. For a multi-frame assignment approach, the correct association of measurements to tracks may not be possible unless the collection of useful features falls within the window size on the number of scans processed by the assignment algorithm.

2.1. Modeling Vehicle-to-Track Association Ambiguity

As discussed previously, we wish to model the uncertainty associated with track switches. After two or more vehicles interact in such a way that associating collected measurements with the tracks is ambiguous, uncertainty exists as to whether each vehicle is still associated with the same track prior to the ambiguous period. One of our goals is to model this vehicle-to-track association uncertainty and, hopefully, resolve it by exploiting subsequently collected features. Therefore, we must first develop a mathematical structure for modeling the uncertainty as well as algorithms for modifying the uncertainty as new information is acquired. We have looked at representing the association uncertainty by considering the mapping of tracks at one point in time to the set of tracks at a later point in time. The tracks to which the vehicles are associated at any point in time are a permutation of the tracks to which the vehicles were associated at some prior time. If this permutation is the identity permutation, then no track switching has occurred. In reality this permutation is not known with certainty (and we could, in theory, determine the probability mass function (PMF) over the set of permutations).

To demonstrate the model, we shall consider an example test case. Consider the diagram in Figure 2 in which three vehicles will undergo two ambiguous periods during which measurement-to-track association is uncertain and, therefore, induces an uncertainty in the mapping of vehicles to tracks. The situation in Figure 2 will be




Figure 2. A notional example showing three vehicles that interact. After the two interactions, the mapping of vehicles to tracks is ambiguous.




Figure 3. An abstract representation of the situation in Figure 2.

abstractly represented using the diagram in Figure 3. The circles represent situations in which the tracks “going into” the circle may switch vehicles.

Utilizing notation from abstract algebra, we can represent the possible switching (or not) of tracks by the following permutation (denoted $\sigma : \{1, 2, 3\} \rightarrow \{1, 2, 3\}$)

$$\begin{bmatrix} 1 & 2 & 3 \\ \sigma(1) & \sigma(2) & \sigma(3) \end{bmatrix} \quad (1)$$

in which $\sigma(i)$ is the track to which the i^{th} vehicle is associated after an ambiguous situation(s). A common representation of a permutation is the *cycle notation*. An example of cycle notation is (1 3 2) which represents the following permutation

$$\begin{bmatrix} 1 & 2 & 3 \\ 3 & 1 & 2 \end{bmatrix} \quad (2)$$

As can be seen from (2), the cycle (1 3 2) represents that “1 is now mapped to 3, 3 is now mapped to 2, and 2 is now mapped to 1.”

For a less abstract description of the permutation in (2), suppose that objects \mathcal{O}_1 , \mathcal{O}_2 , and \mathcal{O}_3 are initially associated with tracks \mathcal{T}_1 , \mathcal{T}_2 , and \mathcal{T}_3 , respectively. After some ambiguous situations, the permutation in (2) indicates that object \mathcal{O}_1 is associated with track \mathcal{T}_3 , object \mathcal{O}_3 is associated with track \mathcal{T}_2 , and object \mathcal{O}_2 is association with track \mathcal{T}_1 . For another example, consider the permutation

$$\begin{bmatrix} 1 & 2 & 3 \\ 3 & 2 & 1 \end{bmatrix} \quad (3)$$

which, using cycle notation, is represented by (1 3)(2), two disjoint permutations, or simply (1 3), since explicitly representing the fixed point, 2, is unnecessary. As a matter of notation, we denote the identity permutation (i.e., no change in the mapping) as “e.”

In the context of multi-target tracking, *many* situations in which measurement-to-track association ambiguity exists for a set of tracks induces a vehicle-to-track association ambiguity (i.e., potential track switches) *. Every situation in which a track switch may have occurred can be represented by a polynomial whose coefficients are real numbers and whose variables are members of the permutation group. If we restrict the coefficients to be real and positive and sum to unity, then the coefficients represent a PMF over the set of possible permutations. As an example of this representation, consider the ambiguity representation for the vehicles in Figure 2 with abstract representation in Figure 3. Consider the first vehicle-to-track ambiguity (i.e., the ambiguity between vehicles \mathcal{O}_1 and \mathcal{O}_2 , which are initially associated with tracks \mathcal{T}_1 and \mathcal{T}_2 , respectively). We can represent the uncertainty using the following polynomial

$$[(1 - q)e + q(1\ 2)] \quad (4)$$

in which $q \in [0, 1]$. Simply stated, there is a probability of $1 - q$ that Tracks \mathcal{T}_1 and \mathcal{T}_2 did not switch and a probability of q that they did switch. We can view tracking multiple vehicles over time as a sequence of ambiguous situations (between ambiguous situations, association of measurements to tracks is relatively straightforward, and we assume that no track switching occurs) and represent this as a product of permutations. For the scenario shown in Figure 3, the association of the vehicles \mathcal{O}_1 , \mathcal{O}_2 , and \mathcal{O}_3 to the tracks \mathcal{T}_1 , \mathcal{T}_2 and \mathcal{T}_3 after the two ambiguous regions is mathematically modeled as the following product of two polynomials

$$[(1 - p)e + p(2\ 3)][(1 - q)e + q(1\ 2)] \quad (5)$$

Note that when representing a sequence of ambiguous situations as a product of polynomials, the number of polynomials grows linearly with the number of ambiguous situations; however, the number of possible permutations (found by multiplying out the polynomials) grows exponentially. For example, multiplying (5) out yields the following

$$[(1 - p)e + p(2\ 3)][(1 - q)e + q(1\ 2)] = (1 - p)(1 - q)e + (1 - p)q(1\ 2) + p(1 - q)(2\ 3) + qp(1\ 3\ 2) \quad (6)$$

in which we have simplified the last term from (2 3)(1 2). The polynomial product representation is compact since it grows linearly. Furthermore, from the product of polynomials it is fairly easy to determine the set of tracks a vehicle, which was formerly associated with a particular track, is now associated *without* multiplying out the polynomials.

2.1.1. Probabilities as Polynomial Coefficients

In the representation above, the polynomials have permutations as the variables and real numbers as the coefficients. In particular, the coefficients are probabilities that sum to unity. First, we must have a way to determine the probabilities for each possible permutation. Unfortunately, it is not possible to calculate the probabilities in the single-frame association framework. To accurately calculate the probabilities of the various permutations requires using a multi-target tracking algorithm that captures the correlation between tracks that results from measurement-to-track association uncertainty. For example, we could use a multi-hypothesis tracking algorithm to track the two or more vehicles, which become kinematically ambiguous, through the ambiguous situation or a particle-filter based algorithm that tracks the group through the ambiguity, followed by a clustering algorithm to extract the individual tracks and corresponding track switching probabilities.

*there are situations in which measurement-to-track association is temporarily ambiguous, but the tracks cannot switch

2.1.2. Non-Probabilistic Approach to Polynomial Coefficients

An alternative to considering the probabilities of the various permutations is to merely set all coefficients to unity[†]. In other words, the probability of a particular permutation occurring is irrelevant since it is sufficient to know which permutations are possible. The advantage of this approach is in avoiding the determination of the actual track switching probabilities. We also note that the probabilities are only of value if they are utilized in algorithms that process the polynomials or vehicle-to-track association matrices. Additionally, using the probabilities would significantly increase the complexity of the algorithms used to resolve the vehicle-to-track association ambiguities (to be discussed subsequently).

2.2. Resolving Vehicle-to-Track Association Uncertainty

Modeling the vehicle-to-track association ambiguities using the permutation group is useful for being aware of the possible track switches. Modeling the ambiguities is essentially a bookkeeping exercise, albeit a very important element in a tracking system that must track objects through extended periods of time. However, *being aware* of vehicle-to-track association ambiguities is merely the first step in a tracking approach designed to confidently track vehicles over extended periods of time and/or through ambiguous periods in which the vehicle interacts with other vehicles. In addition to the ambiguity modeling, we desire a set of algorithms that allow for the resolution of vehicle-to-track ambiguity. These algorithms would attempt to make vehicle-to-track associations given measured features, as well as infer additional vehicle-to-track associations. In the next several sections, we discuss the algorithms that allow us to make associations of vehicles-to-tracks and the algorithms that allow us to infer the associations of other vehicles to tracks.

2.2.1. Exploiting Features to Resolve Vehicle-to-Track Association Ambiguity

After a period of measurement-to-track association ambiguity, uncertainty exists as to which vehicles are associated to which tracks. With varying degrees of complexity, traditional approaches to tracking and data association attempt to determine the correct measurement-to-track association but do not attempt to note potential mistakes for possible subsequent verification or correction. As we have discussed, the approach taken by the Tracked Object Manager (TOM) is to note any potential mistakes in association (that lead to vehicle-to-track association ambiguity) and attempt to verify that no track switch occurred or correct any track switch when sufficient evidence has been collected. When the proximity of two tracks is such that measurement-to-track association is ambiguous, we can no longer be confident that the tracks are associated with the same vehicles prior to the ambiguous period. Suppose that a particular track may be associated with multiple vehicles after a period of ambiguity. To address this uncertainty, the TOM notes the ambiguity and looks for an opportunity to utilize newly collected features on the track in order to either verify that no track switch occurred or correct a track switch if, in fact, one did occur.

To understand the problem addressed in this section. Consider the three-vehicle example shown in Figure 2. Suppose that after the vehicles interact (in the two ambiguous periods), the three tracks are isolated so that no more ambiguities occur (for some time). Due to the separation of the tracks, features measurements collected on each track are unambiguously assigned to the correct track. Due to the feature database management, there exist three feature databases that correspond to the three vehicles. By learning a statistical model for each of the three vehicles, we should be able to compare newly collected features to the vehicle databases and, perhaps, decide which vehicle is associated with a particular track. The question is what type of statistical test to apply. Ideally, we would employ a Bayesian classifier, which would require a vehicle-conditional feature likelihood for each vehicle that could be associated with the particular track in question. Unfortunately, since we are addressing problems in which the feature databases are being built on-the-fly, it is often the case that vehicle-conditional likelihoods are *not* available for all vehicles due to incompleteness of the feature databases (the aspect dependence of features exacerbates this problem). The result is that our decisions (i.e., deciding which vehicle is associated with a particular track) are based on a variety of heuristic conditions. The key is that the decisions are timely and confident. Obviously, we desire confident decisions since the objective is continuous tracking, and we do not want to swap ambiguity in data association for ambiguity in identifying a vehicle. On the other hand, the

[†]technically, the coefficients are all $\frac{1}{n}$ in which n is the number of possible permutations in a particular interaction. However, as the coefficients are not used, setting them to 1 is convenient.




Figure 4. A graph depicting a sequence of ambiguities for four vehicles that interact over a period of time (with polynomial product that models the ambiguities).

decisions must be timely since ambiguous situations occur regularly, and without resolving some ambiguities, the number of ambiguities will grow and be more difficult to resolve.

Currently, we collect features on a track that may be associated with the VOI and calculate a feature likelihood that the collected feature originated from the VOI (based on a learned statistical model). The feature likelihood is compared to an upper threshold and a lower threshold. Any likelihood exceeding the upper threshold is considered a “vote” in favor of the track containing the VOI, any likelihood falling below the lower threshold is a vote in favor of the track *not* containing the VOI, and any likelihood between the thresholds is considered inconclusive and is not considered. We require four votes either for or against the track containing the VOI before a decision is made. More sophisticated tests are currently under consideration.

2.2.2. Incorporating Hard Associations

An interesting question to ask is whether, after associating a vehicle to a particular track through feature collection, we can deduce other vehicle-to-track associations using our polynomial model for registering ambiguities. In this section we shall discuss how to deduce other associations using algorithms that examine the history of ambiguities, taking into account any resolutions. Suppose we have a sequence of interactions between four vehicles, which is abstractly depicted in Figure 4. Note the product of polynomials that mathematically represents the ambiguities (repeated here)

$$[e + (1\ 2)][e + (2\ 3)][e + (3\ 4)][e + (2\ 4)] \quad (7)$$

We have dropped the coefficients since we do not calculate the probability of each permutation. Suppose through the collection of features we are (confidently) able to determine that vehicle \mathcal{O}_1 is associated with \mathcal{T}_2 . Can any vehicle-to-track associations be deduced from this association? We have algorithms that identify the possible “paths” through the graph in Figure 4 that are consistent with a given association. The two possible paths consistent with the given association are highlighted in Figure 5. Corresponding to each path is a polynomial that represents both known and unknown ambiguities. For example, the two polynomials corresponding to the paths in Figure 5 are

$$\begin{aligned} &(1\ 2)(e)(e + (3\ 4))(e) \\ &(1\ 2)(2\ 3)(3\ 4)(2\ 4) \end{aligned} \quad (8)$$

While we have algorithms to determine what deductions are possible given knowledge of the possible paths, the reader can see from Figure 5 that \mathcal{O}_2 *must* be associated with \mathcal{T}_1 . For the example given, no more associations are possible through deduction, and we must wait for more features to be collected and used to identify any vehicles (\mathcal{O}_3 and \mathcal{O}_4 that are ambiguously associated with the remaining tracks (\mathcal{T}_3 and \mathcal{T}_4)).




Figure 5. Given that \mathcal{T}_2 is associated with \mathcal{O}_1 , we determine the possible paths through the graph that are consistent with the association.



Figure 6. Two vehicles travel down different road segments but are converging on the same intersection.

3. RESULTS

In this section we examine the results of applying the approach to two problems. The first problem is a simple problem in which we track two vehicles through an interaction that induces an uncertainty as to whether a track switch occurs. In the second example, we track a vehicle of interest for an extended period of time among a collection of background (confuser) vehicles.

3.1. Two Vehicle Example

To demonstrate the algorithms under development, we examined the output of a case in which we track two vehicles through an ambiguous period. The vehicles are first isolated as they approach an intersection (shown in Figure 6). During this period, the collected feature measurements are stored in each vehicle's feature database. Due to uncertainty as to what happens at the intersection, the TOM declares that there was a potential track




Figure 7. Plot of the number of feature profiles in the track and vehicle databases versus time.

switch, and both vehicle feature databases are frozen since measurement-to-track association is ambiguous, and we want to maintain database purity. The vehicles travel along the same road segment until they reach another intersection at which point they separate and travel down different road segments. While the vehicles are separated sufficiently such that assigning measurements to tracks is straightforward, the collected features are used to determine which vehicle is associated with which track. Finally, features are collected which indicate that Vehicle 1 is associated with Track 2 (in reality, the tracker switched tracks at the first intersection and, fortuitously, switched back at the second intersection). A plot of the vehicle and track databases over time is shown in Figure 7. From the plot we see that initially, since the vehicles are separated and no ambiguity exists in associating measurements to tracks, any collected features are added to the vehicle databases. However, at approximately 1600 seconds, the TOM indicates that measurement-to-track association ambiguity exists and no more features are added to the vehicle databases. During the period in which the vehicles remain in proximity, the collected features (which are at an aspect angle for which the original vehicle feature databases are too sparse to utilize the features) are discarded. After the vehicles separate at the next intersection, the two tracks are not ambiguous anymore, though we do not know whether a track switch has occurred. Any collected features are easily assigned to a particular track since the vehicles in those tracks are physically separated. If we knew which vehicle was associated with each track (i.e., whether a track switch occurred or not), we could put the corresponding features into each vehicle’s feature database. Anticipating that we will finally collect features at an aspect angle for which the original vehicle databases are sufficiently rich, we temporarily store the collected features into the appropriate *track* feature database. Note from the plot that at approximately 3100 seconds collected features indicate that Vehicle 1 is associated with Track 2. At that point, all the features temporarily stored in the Track 2 database are added to the feature database for Vehicle 1. At the time these results were generated, the algorithms that deduce other associations were not implemented, and we see that the TOM never deduces that if Vehicle 1 is associated with Track 2, then Vehicle 2 must be associated with Track 3.

3.2. Tracking a Vehicle of Interest among 100 Confuser Vehicles

In this section, we show the results of tracking a vehicle of interest among 100 background vehicles. Again, no *a priori* database of features exists, and we must collect features and build a unique characterization of the vehicles




Figure 8. Scenario in which a vehicle of interest is tracked among 100 confuser vehicles.

of interest. A snapshot of the scenario is shown in Figure 8. The background vehicles comprised four vehicle types (i.e., 25 of each vehicle type). The vehicle of interest is shown in red and had a unique feature database. In order to collect enough features prior to being confused with the background vehicles, the vehicle of interest was isolated for 18 minutes while features were collected (at the point of first interaction with a background vehicle, the feature database of the VOI contained approximately 270 aspect-dependent feature measurements[‡]).

The TOM declared ambiguities whenever the vehicle of interest interacted with background vehicles and caused measurement-to-track association ambiguity. Subsequently, the TOM would exploit collected features to locate the vehicle of interest among the two or more tracks that could be associated with the vehicle of interest. The TOM was able to keep the VOI in track for more than 1 hour before mistakenly associating the vehicle of interest with a particular track. Due to the ambiguous periods, the TOM does not always know the location of the TOI. A plot of the TOM's association of the VOI with a track is shown in Figure 9 along with the *actual* track to which the VOI is associated. Whenever the TOM was uncertain as to which track the VOI was associated[§], the plot shows the TOM's track label as "0." In other words, track ID "0" indicates the TOM is attempting to decide with which of several (or, possibly, more) tracks the VOI could be associated.

4. SUMMARY

We have presented an approach to utilizing a simple single-frame assignment approach to multi-target tracking in combination with a module we call the *Tracked Object Manager* (TOM) to achieve continuous tracking of vehicles of interest by exploiting features that *uniquely* characterize a vehicle. We have focused on the case in which no *a priori* feature database is available on any target types and, therefore, the features are dynamically collected (the approach presented in this paper can also be applied when an *a priori* feature database is available). The approach exploits features collected on the vehicle of interest to locate the vehicle within the track database after ambiguities induce potential track switches. As the success of locating the VOI in the track database relies

[‡]Note that without at least some features on the vehicle of interest prior to confusion, no algorithm can hope to continuously track the vehicle for any extended period of time.

[§]Note that "uncertain" means that the TOM only knows the association of the VOI to a subset of the known tracks.




Figure 9. Plot of the *actual* track associated with the vehicle of interest as well as the TOM’s estimate of which track is associated with the vehicle of interest. When the TOM has declared the mapping as ambiguous, the plot shows the track label as “0.”

on feature “matching,” the TOM manages the feature database for the vehicle of interest in order to keep the database as pure as possible. The approach shows potential based on simulation results and more analysis and testing is currently underway.

ACKNOWLEDGMENTS

The work presented here was supported under contract No. F8650-04-C-1683 by the Air Force Research Labs (AFRL). We would like to thank Dr. Djuana Lea (AFRL/SNAT), Pete Trautman (formerly of AFRL/SNAT) and Dr. Devert Wicker (AFRL/SNAT) for their support and encouragement of this work.

REFERENCES

1. A. Poore and A. Robertson III, “A new lagrangian relaxation-based algorithm for a class of multi-dimensional assignment problems,” *Journal of Computational Optimization and Applications* **8**, pp. 129–150, 1997.
2. S. Deb, M. Yeddanapudi, K. Pattipati, and Y. Bar-Shalom, “A generalized s-d assignment for multisensor-multitarget state estimation,” *IEEE Transactions on Aerospace and Electronic Systems* **AES-33**(32), pp. 523–536, 1997.
3. D. Reid, “An algorithm for tracking multiple targets,” *IEEE Transactions on Automatic Control* **AC-24**, pp. 843–854, Dec. 1979.
4. R. Singer, R. Sea, and K. Housewright, “Derivation and evaluation of improved tracking filters for use in dense multitarget environments,” *IEEE Transactions on Information Theory* **IT-20**, pp. 423–432, July 1974.
5. K. Ezal, C. Agate, and J. Gambill, “Continuous feature-aided intelligent and dynamic sensor tasking,” tech. rep., Toyon Research Corporation, 2004.