LEARNED MULTITARGET TRACKING by FEATURE RECOGNITION

Capt. Erik P. Blasch
AFRL/SNAT
241 Avionics Cir.
WPAFB, OH 45433
blaschep@aa.wpafb.af.mil

Scott Watamaniuk, Ph.D.
Wright State University
Psychology Dept.
Dayton, OH 45435
swatamaniuk@nova.wright.edu

Key Words: Learning, Sensor Integration, Multitarget Tracking

ABSTRACT

A tactical pilot typically experiences difficulty in maintaining an accurate lock on multiple-interacting targets in the presence of clutter. If the pilot is trained to focus on object features and predict the relative motions of targets, the anticipated target tracking methodology allows the pilot to identify multiple targets simultaneously. By assessing features, the paper seeks to demonstrate how feature recognition enhances multisensor integration of multitarget scenarios.

1.0 INTRODUCTION

The problem of multitarget tracking in the presence of clutter has been effectively explored with unknown scenarios [Bar Shalom 95]. If, however, mission scenarios (possible numbers of targets) and target features (identifiable markers) are known, tracking algorithms can be enhanced by learning techniques. Intelligent tracking would make use of learned associations of features to targets. The main advantage of the proposed intelligent algorithm is the reduction in computations needed to perform time-critical, space-limited tracking.

The problem of multitarget tracking is a subset of sensor management which includes selecting sensors, sensor-detection patterns and policies, and tracking algorithms for a given set of mission requirements [Popoli 92]. For example, a typical tactical aircraft contains sensors with different modes to detect different features. These sensors make kinematic and identity measurements to detect, track, and identify objects of interest while reducing pilot workload. In a dynamic and uncertain environment, the onboard sensor manager must select the correct sensor to measure the correct target at a given time. Thus, the automated reasoning of the sensor manager must control the measurement sequencing process for effective tracking. A multi-sensor/multi-target tracking policy is best described as a problem in sequential-decision making and uncertainty. Similar applications include engineering, management science, and biology and a wide variety of mathematical techniques have been developed and applied to some aspect of parameter tracking and classification [Kastella 97]. We propose to leverage human subject experiments for our mathematical model.

Biological studies have discovered methodologies on how humans track multiple objects [Bravo 95]. It was found that a coarse local speed signal is used for object segregation and a precise global speed signal for discrimination. Thus, speed of the targets was found to be a discriminating factor for target classification. As in these experiments, we model the multisensor-multitarget tracking problem by utilizing learned feature recognition as a pilot would do to discriminate between targets. In sensor-target tracking, a sensor is directed to perform a sequence of measurements that will detect and track targets. The challenge is to guide the sensor so it locates the target "efficiently". Learning techniques such as reinforcement and association learning have been applied for searching and detection of targets [Blasch 97]. Utilizing learned feature recognition may offer a means to control some aspects of the computational burdens experienced by analytical multitarget optimization techniques while providing an effective solution for multitarget tracking in the presence of clutter.

2.0 INTELLIGENT TRACKING

Intelligence in tracking is the ability of the agent to discern the salient features in an image to track and classify targets simultaneously. Based on human-subject experiments, where tracking includes finding a coarse local signal for the target and a precise global signal for classification, we model intelligence as shown in Figure 1. Our model for human motion processing includes a motion detector, like those of the primary visual cortex. Other detectors include spatial and temporal frequency, orientation, and direction. The motion detector is used by the association cortex as a feature with which to classify targets using association learning [Blasch 98]. Association learning is a biological technique thought to represent the filtering of multisensor information by the thalamus and the processing of information in the prefrontal association cortex.

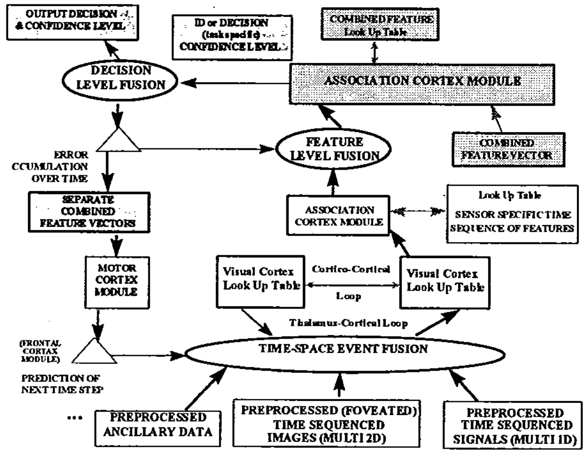

Figure 1. Biological Model of Intelligent Feature-Recognition Tracking.

3.0 PROBLEM FORMULATION

Consider Figure 2 as an environment that the pilot is monitoring. The pilot's goal is to effectively detect, track, and classify any target that enters the region.

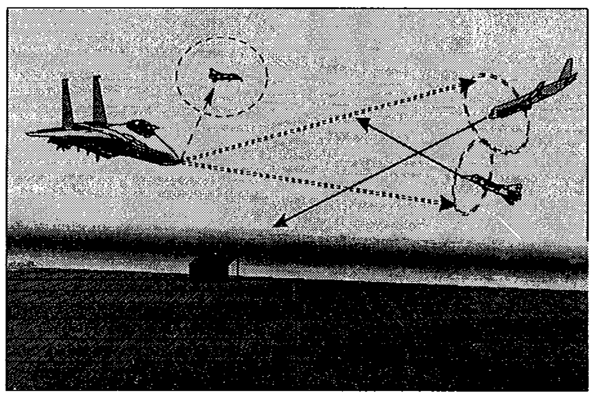

Figure 2. Multitarget Tracking.

Assume that the region in Figure 2, the 2-D frame, is composed of T targets with f features. Dynamic target measurements z, are taken at time steps k, which include target kinematics and features $z(k) = [x_t(k), f_1, \ldots f_n]$. Any sensor can measure independently of the others, and the outcome of each measurement may contain kinematic or feature variables indicating either target. The probability density of each measurement depends on whether the target is actually present or not. Further assume that a fixed number of kinematic and feature measurements will be taken at each time interval, where we model the clutter composing spurious measurements. A final decision is rendered as to which [x, y] measurement is associated with the target, determined from the learned feature recognition.

The *multisensor-multitarget tracking problem* is to determine which measured features should be associated with which kinematics in order to optimize the probability that the targets are tracked correctly after z measurements. The multitarget kinematic tracking problem is formulated and solved by using concepts from probability data association. For the symmetric-target case, the "association rule", - associate the measurement with the highest probability of target - produces the sub-optimal result when making the final decision.

Two methods are chosen. The first, a probability data association technique, which we call *Measurement Tracking*, searches through all the measurements and probabilistically chooses the measurement most likely to be associated with the target. The second method, *Feature Recognition Tracking*, which is described next, is a procedure that uses position measurements as the coarse local signal for believable target measurements and a precise global feature signal for discriminating between targets. In the example, we use speed as the feature since one target has a positive y speed and the other target has a negative y speed when the objects are close together.

4.0 FEATURE TRACKING

4.1 Tracking

The target *state* and *true measurement* are assumed to evolve in time according to:

$$x(k+1) = F(k)\,x(k) + v(k) \qquad (1)$$
$$z(k) = H(k)x(k) + w(k) \qquad (2)$$

where, $v(k)$ and $w(k)$ are zero-mean mutually independent white Gaussian noise sequences with known covariance matrices $Q(k)$ and $R(k)$, respectively. *False measurements* are uniformly distributed in the measurement space. Tracks are assumed initialized at an initial state estimate $x(0)$, contain a known number of targets determined from the scenario, and have associated covariances.

A plausible elliptical validation region V, with a *gate threshold*, γ, is set up at every sampling time around the predicted measurement and is used to select believable correct measurements. Measurements from one target can fall in the validation region of the neighboring target and is *persistent interference.* All feature variables that carry information useful to discern the correct measurement from the incorrect ones are assumed to be included in the measurement vector. The approaches studied differ in how feature measurements are used in the estimation of the kinematic state to the correct target.

4.2 Tracking Belief Filter

The *Tracking Belief Filter* is an intelligent method which devotes equal attention to every believable measurement and cycles through measurement features until an object classification is reached. The filter assumes the *past* is summarized by an *approximate sufficient statistic* - state estimates (approximate conditional mean) and covariances for each target.

The *measurement-to-target association probabilities* are *computed across* the targets and these probabilities are computed only for the *latest set of measurements*. The conditional probabilities of the joint-target association events pertaining to the current time k is defined as $\theta(k)$, where θ_{jt} is the event that measurement j originated from target t, $j = 1\ldots, m(k)$; $t = 0, 1, \ldots, N_t$, where $m(k)$ is the total number of measurements for each time step, k is the time of measurements, and N_t is the known number of targets.

4.3 The Believable Joint Events

A validation gate is used for the selection of the believable joint events, but not in the evaluation of their probabilities.

The *Plausible validation matrix*: $\Omega = | \omega_{jt} |$ is composed of binary elements that indicate if measurement j lies in the validation gate of target t. The index $t = 0$ stands for "none of the targets" and the corresponding column of Ω includes all measurements, since each measurement could have originated from clutter, false alarm, or the true target.

A *joint association event* consists of the values in Ω corresponding to the associations in θ,

$$\hat{\Omega}(\theta) = | \hat{\omega}_{jt}(\theta)| = \begin{cases} 1 & \text{if } \theta_{jt} \in \theta \\ 0 & \text{otherwise} \end{cases} \qquad (3)$$

A *believable association event* has
i) a single measurement source:

$$\sum_{t=0}^{N_T} \hat{\omega}_{jt}(\theta) = 1 \quad \forall j \qquad (4)$$

ii) and at most one measurement originating from a target:

$$\delta_t(\theta) \triangleq \sum_{j=1}^{m} \hat{\omega}_{jt}(\theta) \le 1 \qquad (5)$$

The generation of event matrices, $\hat{\Omega}$, corresponding to believable events can be done by scanning Ω and picking one unit/ row and one unit/column except for $t = 0$.

The binary variable $\delta_t(\theta)$ is called the *target detection indicator* since it indicates whether a measurement is associated with the target t in event θ, i.e. whether it has been detected.

The *measurement association indicator*

$$\tau_j(\theta) \triangleq \sum_{t=1}^{m} \hat{\omega}_{jt}(\theta) \qquad (6)$$

indicates measurement j is associated with the target t in event θ.

The number of *false measurements* in event θ, is

$$\phi(\theta) = \sum_{tj=1}^{m} [\, 1 - \tau_j(\theta) \,] \qquad (7)$$

The **joint association event probabilities** are, using Bayes' Formula:

$$\begin{aligned} P\{\theta(k)|Z^k\} &= P\{\theta(k)|Z(k),m(k),Z^{k-1}\} \\ &= \frac{1}{c}\, p[Z(k) \mid \theta(k),m(k),Z^{k-1}]\, P\{\theta(k) \mid m(k)\} \\ &= \frac{1}{c} \prod_{j=1}^{m(k)-\phi(k)} V\, \{f_{t_t}(k)\, [z_j(k)]\}^{\tau_j} \qquad (8) \end{aligned}$$

where c is the normalization constant.

The number of *measurement-to-target assignment events* $\theta(k)$, is the number of targets to which a measurement is assigned under the same detection event, $[m(k) - \phi]$. The *target indicators* $\delta_t(\theta)$ are used to select the probabilities of detecting and not detecting events under consideration.

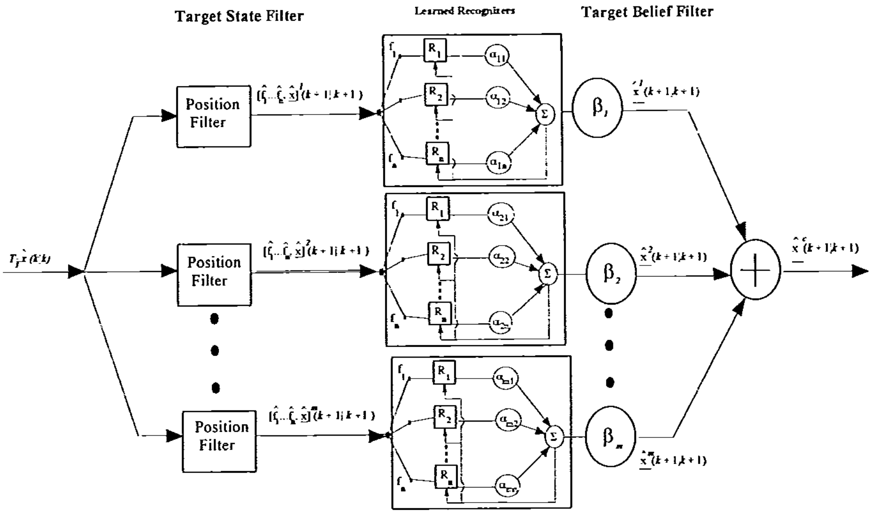

Figure 3. Feature Recognition Tracking Model.

4.4 Feature Recognition Learning

The learning of features is done through a training network. Although algorithms exist for solving image recognition problems using NNs, these algorithms are seldom implemented with tracking algorithms due to the computational burdens arising from the large number of possible hypotheses. By employing *belief states* which incorporate all previous hypotheses, the dynamic-detection trained network is converted to a Markov Decision Problem. The MDP problem is solvable by an association learning technique that uses a function approximator for weights and minimizes computational permutations.

In *feature-recognition tracking,* every measurement in the validation region is evaluated, based on the features, to determine the expected reward $R_k(t)$ for assessing the measurement. The next measurement (with the learned association of features-to-targets) is chosen with the largest expected reward. For a particular measurement z(k*), $R_{k*}(t)$ is evaluated using an learned association-trained network and the number of feature measurements accumulated in k*. The procedure for initializing detection and proceeding through the validated measurement was shown above.

Association learning assigns a *reinforcement* value, $R = \alpha * f(V(\gamma))$, which is proportional to how close the measurement is to the center of the gated validation region. The agent selects validated measurements and cycles through until the feature-to-target belief indicates a true match of measurement to target. A threshold is used to pick an acceptable measurement from which to quit cycling through validated measurements.

4.5 State Estimation

Assuming the targets conditioned on the past observations are *mutually independent*, the decoupled state estimation (filtering) of the *marginal association probabilities*, which are obtained from the joint probabilities, is obtained by summing over all joint events in which the marginal event of interest occurs. The conditional probability of the event (the association probability) is:

$$\beta_{jt} \triangleq P\{\theta_{jt}(k)|Z^k\}$$
$$= \sum_{\theta} P\{\theta|Z^k\}\hat{\omega}_{jt}(\theta) = \sum_{\theta:\theta_{jt}\in\theta} P\{\theta|Z^k\} \quad (9)$$

The algorithm decomposes the estimation with respect to the origin of each element of the *latest set* of validated measurements.

Using the total probability theorem, with respect to the above events, the *conditional mean* of the state at time k can be written as:

$$\hat{x}(k|k) = \sum_{i=0}^{m(k)} \hat{x}_i(k|k)\,\beta_i(k) \qquad (10)$$

where x(*k/k*) is the update state *conditioned on the event that* i^{th} *validated measurement is correct.*

The state estimate, conditioned on measurement i being correct, is:

$$\hat{x}_i(k|k) = \hat{x}_i(k|k-1) + W(k)\nu_i(k) \qquad (11)$$

$$\nu_i(k) = z_i(k) - \hat{z}(k|k-1) \qquad (12)$$

$$W(k) = P(k|k-1)H(k)^T S(k)^{-1} \qquad (13)$$

The *combined* state update equation, combined innovation, and covariance associated with the state are:

$$\hat{x}(k|k) = \hat{x}(k|k-1) + W(k)\nu(k) \qquad (14)$$

$$\nu(k) = \sum_{i=1}^{m(k)} \beta_i(k)\,\nu_i(k) \qquad (15)$$

$$P(k|k) = \beta_0(k)P(k|k-1)+[1-\beta_0(k)]P^c(k|k) \qquad (16)$$

where the covariance of the state is updated with the *correct measurement* is:

$$P^c(k|k) = P(k|k-1) - W(k)S(k)W(k)^T \qquad (17)$$

The *prediction* of the state and measurement to time k+1 is done as in standard filtering techniques.

5.0 RESULTS

The two dynamic tracking methods discussed are compared. The method for evaluating performance is a Monte Carlo simulation and the performance metric is normalized probability of state error. It is assumed that the feature in question is speed. As detailed in the Figures 4 and 5, by the true trajectory, the targets 1) start with position X = {(2000,10200), (2000,9900)} and speeds of +10 x *m/s*, 2) pass by each other at a distance of 5 meters and speeds of -2 and + 2 y *m/s*, and 3) finish with a speed of +10 x *m/s*.

Table 1: Normalized Square Errors

State Error	X1	X2	Y1	Y2
Measurement	101.84	98.45	6.63	3.50
Feature Recognition	98.59	97.79	3.40	2.61

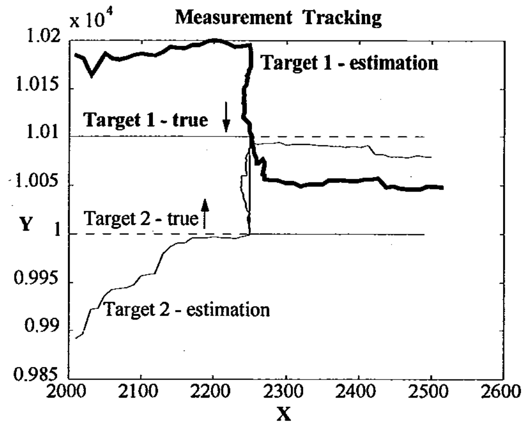

Figure 4. Measurement Tracking.

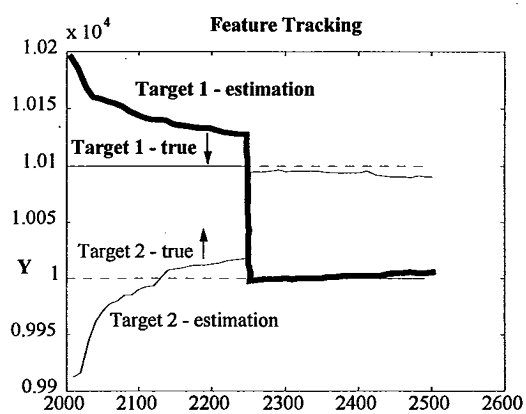

Figure 5. Feature Recognition Tracking.

6.0 DISCUSSION & CONCLUSIONS

The figures above show that measurement tracking incorrectly associates some of the measurement data from the second target with that of the first target. The feature-recognition tracking algorithm, which uses a speed sensor with measurement uncertainty, detects the speed feature of each target and correctly assigns measurements to the targets.

This research included training a network using the association learning for feature recognition to guide an imperfect sensor or a perfect sensor in the presence of clutter to find a set of dim targets in a region. In a series of simulation experiments, the feature recognition network performed well resulting in a desirable solution, and at a faster rate than conventional mutlitarget-multisensor tracking methodologies. The presented feature recognition learning technique demonstrates promise for multitarget tracking problems and warrants further exploration in problems where environmental effects, occlusions, and lost sensor data can be modeled that are not readily handled by current tracking algorithms.

Acknowledgments

This research was supported under task 2312R1 by the Mathematics and Geosciences Directorate of the United States Air Force Office of Scientific Research.

REFERENCES

Bar-Shalom, Y. and X Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, YBS, New York, 1995.

Popoli, R., "The Sensor Management Imperative", Chapter 10 in *Multitarget-Multisensor Tracking: Advanced Applications*, Vol. II, Y. Bar-Shalom, Ed., Artech House, 1992.

Kastella, K., "Joint multitarget probabilities for detection and tracking," *SPIE AeroSense '97*. 21 - 25 April 1997.

Blasch, E. "Physio-Associative Temporal Sensor Integration," *Proceedings of the SPIE AeroSense '98*, Orlando, FL, April 20 - 28, 1998.

Blasch, E., R. Malhotra, and S. Musick, "Using Reinforcement Learning for Target Search in Sensor Management," *IRIS National Symposium on Sensor and Data Fusion*, Boston, MA, April 12 - 18, 1997

Bravo. M, and S. Watamaniuk. "Evidence for Two Speed Signals: A Coarse Local Signal for Segregation and a Precise Global Signal for Discrimination," *Vision Research*, Vol. 35, pp. 1691-1697, 1995.