Resident Space Object Feature Identification using Hierarchical Mixtures of Experts

Jessica T. Anderson, David E. Gaylor, Dunning Idle 5th
Emergent Space Technologies, Inc.
Ryan P. Russell, Ashley Biria, Maximilian Schadegg
The University of Texas at Austin

As the number of objects in orbit and spacecraft capability grows, the need for space situational awareness (SSA) also increases. The identification and cataloging of resident space objects (RSOs) is a large and complex problem, and new methods of performing these tasks are necessary to keep up with the rapid growth of objects in space.

Emergent Space Technologies, Inc. (Emergent) and the University of Texas at Austin (UT) worked together to develop an algorithm to aid in detecting RSO characteristics using a Hierarchical Mixture of Experts (HME). This algorithm has been tested in simulation, and has been shown to be capable of detecting the size, shape, reflectivity, and maneuvers performed by an RSO.

1. BACKGROUND

Space situational awareness (SSA) involves gaining and maintaining knowledge of objects in Earth orbit, particularly objects in the vicinity of valuable space assets. Hazards to space assets are increasing, largely due to the growing number of resident space objects (RSOs) as well as the growing capabilities of potential adversaries. How to effectively automate, schedule, and manage the decision-making associated with maintaining the U.S. Space Object Catalog is a challenging problem of increasing importance to our national security [1]. The U.S. Space Object Catalog currently lists approximately 15,000 trackable objects accounting for approximately 5,800 tons of on-orbit mass. The total population is thought to exceed 20,000 objects larger than 10 cm [2]. Before objects can be catalogued, they must first be characterized. It is necessary to know the location and, preferably, the identity of these objects.

Distant RSOs of interest may be beyond the resolving capability of current electro-optical sensors. Therefore, RSO feature identification must rely on radiometric, spectral, and polarimetric measurements over time. In general, ground-based telescopes with electro-optical sensors are able to provide three measurements of celestial objects: apparent magnitude, right ascension, and declination. The apparent magnitude of a celestial body is a measure of its brightness as seen by an observer on Earth, corrected to the value it would have in the absence of the atmosphere. Right ascension and declination measurements can be used with traditional methods in order to determine the Cartesian position and velocity or orbital elements of an RSO [3]. Apparent magnitude acts as a weak range measurement and also contains some attitude information, so it is able to provide more RSO state information than only angle measurements can provide.

While accurate knowledge of RSO orbital elements is important, to achieve true SSA we must know much more. We must discern the identity, behavior, health, and ultimately, the intent of the RSO. Only then can we differentiate potentially hostile satellites from their benign neighbors. Electro-optical measurements can be used to estimate important characteristics, such as the size, shape, configuration, rotational dynamics (attitude and angular velocity), and surface properties such as specular and diffuse albedo (reflectivity) of an RSO. In addition, estimated features can be used to match or discriminate with known objects or classes of objects and to identify their behavior.

The U.S. Joint Chiefs of Staff set a goal to enhance U.S. Department of Defense (DoD) effectiveness by constructing a system capable of real-time, persistent tracking of RSOs, augmented by advanced sensor fusion techniques capable of inferring probabilistic intent and detecting dynamic change [4]. As part of meeting this goal, the U.S. Air Force is seeking to develop innovative feature estimation or matching approaches to process measurements produced by new and/or existing electro-optical systems.

Under a Phase I SBIR sponsored by Air Force Research Laboratory (AFRL), Emergent Space Technologies, Inc. (Emergent) and the University of Texas at Austin (UT) investigated the use of Hierarchical Mixtures of Experts (HMEs) to process simulated electro-optical measurements of RSOs to estimate their size, shape, configuration, attitude, angular velocity, and reflectivity. The electro-optical measurements included photometric measurements (apparent magnitude) and angles (right ascension and declination).

We demonstrated that an HME can determine size change, shape change, reflectivity change, and can detect the time, direction, and magnitude of a maneuver. Output from our HME can provide insight into RSO intent, capabilities, and changing behavior. The Joint Directors of Laboratories introduced a model of data fusion that divided the various processes into five levels [5]. Most SSA research focuses only on level 1 processing, which combines data to obtain estimates of an entity's location, motion, attributes, characteristics, and identity. Level 2 processing requires using the data gleaned from level 1 processing to determine how the observed entity and its relationship with its environment impact a changing situation. Level 3 processing involves the prediction of the future situation. The HME algorithm not only provides level 1 processing, as it can detect an RSO's attributes and characteristics such as size, shape, and albedo, but it offers capabilities to perform level 2 and level 3 processing as well because it can detect changes in RSO parameters, which could provide insight into an RSO's behavior and its impact on the RSO's environment. Identifying RSO behavior could also lead into level 3 processing, as current changes to an RSO could be used to predict its future behavior.

2. HME ARCHITECTURE

An HME is comprised of a number of experts (state estimators), all of which are estimating the state of the same object. Emergent's HME algorithm is related to an algorithm developed by Crain, Bishop, and Ely to process Mars Pathfinder data [6]. The experts are arranged in banks, each of which can contain any number of experts. In turn, the banks of filters can be organized into multiple layers as necessary. An example of a two layer HME can be seen in Fig. 1.

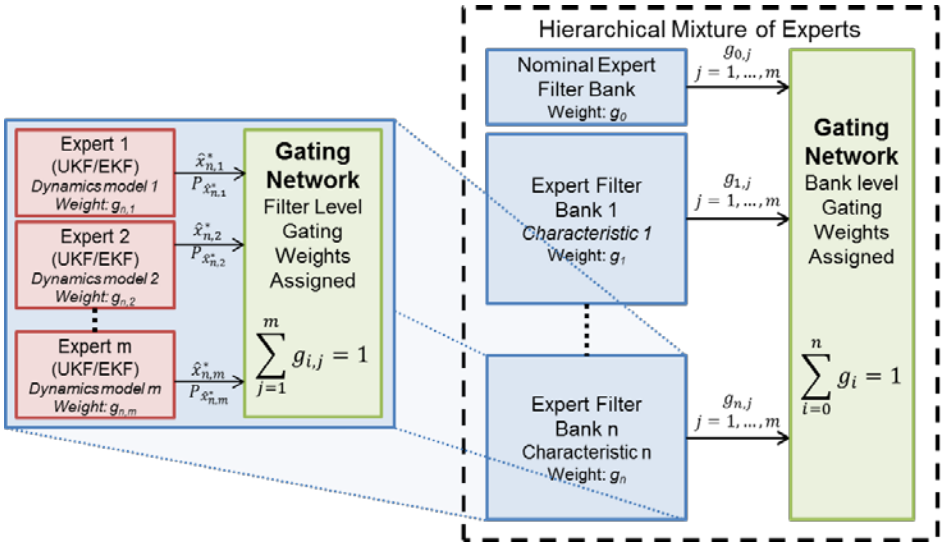

Fig. 1. Two Layer Hierarchical Mixture of Experts

On each layer, a gating network (that is, a simple neural network) acts on the outputs of the experts or banks of experts. The gating network assigns weights to individual filter estimates (or to banks of filters) based on how well the filters predict incoming measurements. A confidence measure, used in adjusting the gating weights for each expert and bank of experts, is provided by the state error covariance. In addition, gating network weights provide a scalar confidence measure that can be used to identify which experts (and, by extension, which dynamic models) best match reality.

Each expert has a slightly different model of the object's state dynamics. All of the experts process the same measurements, but each will produce a different estimate of the object's state. For the purposes of this research, all of the experts used were extended Kalman filters (EKFs). Each EKF uses a different realization of some unknown system parameters, such as size or shape. In the general case, the unknown parameter vector can include different process models as well as different process and measurement noise. Use of higher order filters or other modern filters such as unscented Kalman filters (UKFs) would be possible, and could allow for improved performance over an HME using EKFs. In theory, any type of filter or state estimator could be used as an expert in an HME, as long as it could provide measurement residuals and covariances for those residuals.

The experts and banks of experts in an HME can be arranged to detect micro- and macro-modes in the true state dynamics. In this research, each filter contained a different instance of a micro-mode and each bank of filters modeled a different instance of a macro-mode. An HME can also contain a nominal filter or bank. Nominal filters can help with the detection of RSO features or hinder it, depending on the characteristics of interest. Examples of macro-modes that could be modeled include a change in RSO reflectivity, a change in size, or the existence of measurement noise. Micro-modes are different realizations of the macro-modes, such as individual RSO size models, or specific values of surface reflectivity or measurement noise. An example HME arrangement modeling various micro- and macro-modes can be seen in Fig. 2.

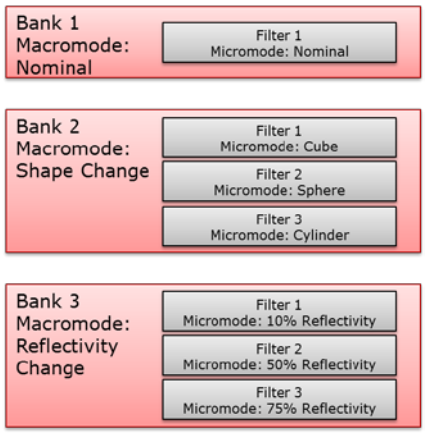

Fig. 2. Examples of micro- and macro-modes in the HME

A recursive weighting function (i.e., the gating function) is then used to assign weights (the *g*'s in Fig.1) to the outputs of the filters. Higher weights correspond to the experts in the HME that most accurately represent the true environment. The weight factors are computed as measurements and are processed by the filter bank. The bank of filters learns which filter is performing best by examining a given performance measure, such as measurement residuals and covariances corresponding to those residuals. When all input processing is completed, the highest weight factor will correspond to the best performing filter.

The filter bank update corresponds to intelligently speculating on different realizations of the unknown parameters. The gating network weights are a natural metric to identify groups or classes of events along with confidence in classification assessment. For example, if one of the filter banks represents an object performing a propulsive maneuver of various magnitudes and/or directions, and if the filter bank has a high weight relative to the other filter banks, then this indicates that the object is probably undergoing a propulsive maneuver. Examining the filters within that particular bank, the filter with the highest weight indicates the magnitude and direction of the propulsive thrust maneuver. Similar filter banks can be realized for RSOs undergoing component articulation, orientation changes, and other characteristics that can be detected and classified. The mixture-of-experts can classify events which are otherwise indistinguishable, but are indirectly detectable in the tracking data through their effects on the trajectory and observable via electro-optical measurements. For example, a noncontinuous property of an RSO (like shape or configuration) cannot be estimated directly in a filter like an EKF, but it can be determined using an HME.

One key feature of the HME framework is that it does not assume that the optimal filter is included in the bank. This lack of complete knowledge is important in this application because we know that practical issues (lack of measurements, non-Gaussian random noise on the measurements, etc.) dominate the estimation error accuracies, and the optimal filter will most likely not be in the bank. The block on the right in Fig. 1 illustrates the overall adaptive filtering structure consisting of a filter bank and a gating network in a forward loop, together with a search algorithm in a feedback loop (not included in Fig.1). The forward loop can be viewed as a multiple hypothesis estimation algorithm, wherein different realizations of the unknown (or uncertain) system parameters are coded into each individual filter in the bank, and the learning network decides which realization provides the best estimate given the available input data. The feedback loop, which contains the search algorithm, is used to periodically update the various filters in the filter bank utilizing the information learned about the system in the forward loop.

The HME uses a single learning parameter for each layer of the HME in order to control its speed and sensitivity. This value can be thought of as a tuning parameter. As it increases, the HME responds more quickly to changes in the incoming measurements, but the HME becomes less able to detect small changes. It is possible for a different value of the learning parameter to be set for each layer of an HME, or to be the same for all layers, depending on the application.

3. SIMULATION CONFIGURATION

In order to test the ability of an HME to characterize RSOs, a 6 degrees-of-freedom (6DOF) simulation was used. The simulator models the translational and rotational motion of RSOs, including effects due to the non-uniformity of Earth's gravity field, third-body perturbations, solar radiation pressure (SRP), gravity-gradient torques and atmospheric drag. In order to compute SRP and drag perturbations, each RSO configuration used in the simulation was modeled as a convex system of flat plates. The forces on each plate were computed individually and summed to compute the total acceleration and torque on the RSO. Several different shapes of RSOs can be simulated in the measurement generator, including flat plates, cuboids, and hexagonal prisms (which are meant to approximate cylinders).

The realistic RSO state modeled in the simulation can then be used to generate high fidelity measurements of an RSO, including apparent magnitude, right ascension, and declination.

The apparent magnitude of an RSO, as measured by an observer on the Earth, is a function of the amount of radiant flux received by the RSO from the Sun and of the fraction of light that is reflected in the direction of the observer. This fraction is computed by summing the amount of light reflected by each of the n flat plates that form the body of the RSO model. In order to compute the amount of light reflected by each plate, the Bidirectional Reflectance Distribution Function (BRDF) is used. The BRDF for an object models the amount of incident light which is reflected by an object, and is defined as the sum of the specular and diffuse reflections, which are both functions of its material properties, as well as the angles of incidence and reflection of incoming light. Thus, RSO apparent magnitude is highly dependent on shape and attitude.

During a previous Air Force Phase 1 SBIR, Emergent made improvements to an apparent magnitude measurement model described by Linares, et al. [7]. These improvements included correcting an observability condition, correcting the amount of reflected light reaching the observer, replacing the specular reflectance term in the BRDF, computing an analytic Jacobian of the measurement for use in an EKF, and modifying the implementation of the BRDF to handle "glints."

In the original Linares model, the apparent magnitude is computed for each surface of the RSO and the value corresponding to the brightest magnitude is accepted as the magnitude measurement. This is valid if and only if one surface of the RSO is illuminated and observable. However, if more than one side of the object is illuminated and observable, this model would be inadequate. In order to make the measurement model truly applicable to a wide variety of objects, both resolved and unresolved, the model was modified to include the contributions of all illuminated and observable reflecting surfaces.

The original model employed the BRDF developed by Ashikhmin and Shirley [8]. However, the highly nonlinear nature of the specular reflectance term contained in this particular BRDF made it less tractable for use in an EKF because the EKF requires the first order derivative or Jacobian of the measurement with respect to the estimated states. A modified version of the original BRDF was found in an unpublished paper by Ashikhmin and Premoze [9], where it was shown that the modified BRDF produced a better overall match to real data. More importantly, it was observed that the spectral reflectance term in this modified BRDF model is of a form that is more amenable to computation of the analytic Jacobian of the measurement. For these reasons, the BRDF was replaced with the modified version described in [7], and the analytic Jacobian of the apparent magnitude measurement model was derived and validated.

4. EXPERIMENT CONFIGURATION

For the experimental results given below, a 1000 kg RSO, in a near-geosynchronous orbit, was tracked from a ground station in Chile. The latitude, longitude, and altitude of the ground station are [-30.690800° 289.193730° 2220.00 m].

A second ground station was modeled in Tenerife, Canary Islands; however, the vehicle never came into view from this station. The initial orbital elements for the RSO are given in Table 1. The vehicle was tracked over a 90 minute observation window, with measurements of apparent magnitude, declination, and right ascension available every 15 seconds. Both ground stations had a field of view of 50°. The noise on the measurements received by both ground stations had a strength of 1 arcsec for right ascension and declination and 0.00667 for apparent magnitude.

Table 1. Orbital Elements for tracked RSO

a	41000.0 km
e	0.3
i	7.0°
Ω	212.8°
ω	0.0°
θ	282.0°

The HME was configured with a Learning Parameter, for both layers of the HME, of 1.75. This value was selected through experimentation.

The procedure to run Emergent's HME algorithm is as follows. First, initial conditions and parameters for a truth RSO model and for a number of different experts are defined. These initial conditions include a model of the size, shape, surface reflectivity and initial position and velocity of the RSO for each expert and the truth model. In addition, the parameters for the ground stations (including sensor noise), parameters for the experts (e.g. filter process noise, gravity model specifications, etc.) and any other pertinent and necessary inputs are defined. Then, simulated measurements of the truth RSO are generated using realistic sensor models for the entire length of the observation window. The measurements generated are then sent to the desired number of experts for the HME. For all of the experiments in this paper, all of the experts used were EKFs. The experts use the simulated measurements to generate the inputs for the HME algorithm, which, in this case, are the EKF measurement residuals and their associated measurement covariances.

Finally, the HME processes the measurement residuals and covariances in order to assign gating weights to each expert and bank of experts at each time that measurements are available. These gating weights indicate how well the estimates from each expert agree with the truth model. That is, a higher gating weight means that an expert is more likely to give the "correct" estimate.

Our version of the HME has the following features:

- The code is designed to post-process residuals and innovation covariances from *m* experts in a serial fashion and assumes that all experts processed the same measurements.
- The configuration of the experts into groupings or banks at the top level of the HME is arbitrary and may be configured after data recording. This capability allows for evaluation of macro-mode model change options to be processed after expert data has been generated.
- The current implementation assumes that measurements are uncorrelated for simplicity, but the processing may be modified to consider an innovations covariance with correlation between measurement residuals.
- A single learning rate parameter is used for both layers of the HME, and the learning within each bank is additionally adjusted by the a posteriori probability that that bank is the best fit for the input data stream.

The results of five experiments are included in this paper. These experiments include detecting the change in the shape of an RSO, detecting a change in the surface reflectivity of an RSO, detecting a change in the size of an RSO, detecting the time when an impulsive maneuver was performed, and detecting the direction and magnitude of an impulsive maneuver.

5. RESULTS

Shape Change Experiment

Our 6DOF simulation is capable of modeling three different shapes of RSOs – a flat plate, a cuboid, and a hexagonal prism. In this example, the RSO being tracked by the experts in the HME was initially a 2m x 2m x 2m cube, which changed to a hexagonal prism with rectangular, 2m x 4m sides. Realistically, a change in the shape of an RSO could

occur if solar panels or other auxiliary structures were deployed, or if an RSO discarded a portion of itself, such as a fairing used during liftoff. The arrangement of the experts and expert banks in the HME is given in Table 2.

Table 2. Arrangement of Filters for Shape Change Experiment

Bank	RSO Model	Filter	Dimensions
1	Cuboid	1	1m x 1m x 2m
		2	2m x 2m x 2m
		3	4m x 4m x 2m
2	Hexagonal Prism	1	1m x 2m
		2	2m x 2m
		3	2m x 4m
Truth Model: Cuboid RSO with dimensions 2m x 2m x 2m changing to a hexagonal prism with 2m x 4m rectangular sides at time step 92.			

The gating weights for both the filter banks and the individual experts for the shape change experiment are given in Fig. 4.

At the beginning of the experiment, the highest gating weights were assigned to Bank 1, the bank containing cuboid RSOs, and to Filter 1 in Bank 1, which corresponds to a 2m x 2m x 2m cube RSO model. This matches the truth model. After the change in RSO shape takes place, the HME does recognize that the size of the RSO has changed – Filter 1 no longer is assigned the highest weight. However, it fails to select Bank 2, meaning that it did not determine that the RSO was now a hexagonal prism.

This may be due to the top level learning parameter being set too low, which would keep the top level gating weight from being affected too much by a small change in the measurement residuals. It would be worthwhile to experiment with different learning parameters for the top level and filter level neural networks. It would also be instructive to observe the RSO over a longer observation window.

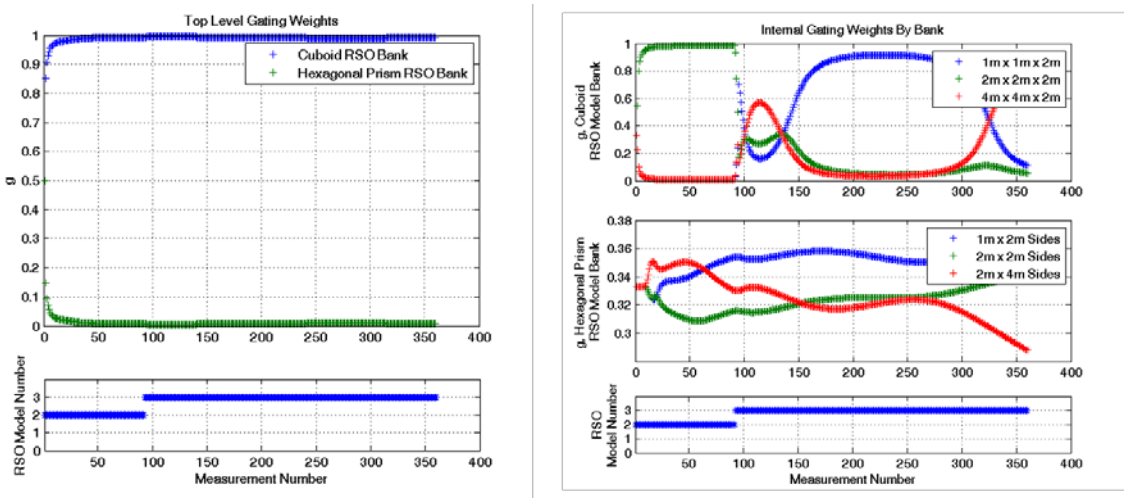

Fig. 4. Gating weights for the shape change detection experiment

Size Change Experiment

The following HME was carefully selected to be used for the size change experiment and the reflectivity change experiment in order to capture typical high-level capabilities of HMEs. Not only does this HME test the individual goals of each experiment, but it also demonstrates that multiple hypotheses can be tested with one HME.

In this HME arrangement, there were three filter banks, including a nominal bank, one with various values of reflectivity, and one with models of several different sizes of RSOs. The purpose of this experiment was to determine the ability of the HME to detect the correct RSO size model and to subsequently detect the correct size of

the RSO after a change in the truth model of RSO size. A change in the size of an RSO could be due to the deployment of solar panels or a solar sail, or the deployment of a threat device by a hostile adversary spacecraft.

For the size change experiment, the RSO was initially modeled as a 3m x 3m x 1m cuboid, which changed to a 5m x 5m x 1m at data point 92. The arrangement of the HME for this experiment is given in Table 3.

Table 3. Arrangement of Filters for Size and Reflectivity Change Experiments

Bank	Filter	Size	Reflectivity
1 (Nominal)	1	Nominal (3m x 5m x 1m)	Nominal (mR = 1)
2 (Reflectivity Change)	1	Nominal	mR = 0.1
	2	Nominal	mR = 0.5
	3	Nominal	mR = 1.5
3(Size Change)	1	1m x 1m x 1m	Nominal
	2	3m x 3m x 1m	Nominal
	3	5m x 5m x 1m	Nominal
Truth Models: • Reflectivity Detection: Nominal Cuboid (3m x 5m x 1m) with 50% reflectivity (mR = 0.5) changes to 150% reflectivity (mR = 1.5) at step 92 • Size Detection: 3m x 3m x 1m Cuboid changes to 5m x 5m x 1m cuboid at step 92			

The response of the HME in this case is given in Fig. 5. This experiment demonstrates how the HME can successfully detect both the size of an RSO, and a change in the size of the RSO. The HME assigns the highest gating weight to the bank containing filters with different-sized RSO models (the correct bank), indicating that it determined that the RSO being tracked by the size experts matched one of the models in that bank. This bank retains the highest gating weight throughout the experiment. Before the change in RSO size, the HME also selects the correct filter (Filter 2 in Bank 3), with a RSO size model matching the truth model (3m x 3m x 1m). After the change, the HME also selects the correct filter – Filter 3 of Bank 3, with a 5m x 5m x 1m cuboid model. In this experiment, the HME, in addition to selecting the correct bank and filter, was able to determine which filter was the closest to the truth after only three measurements.

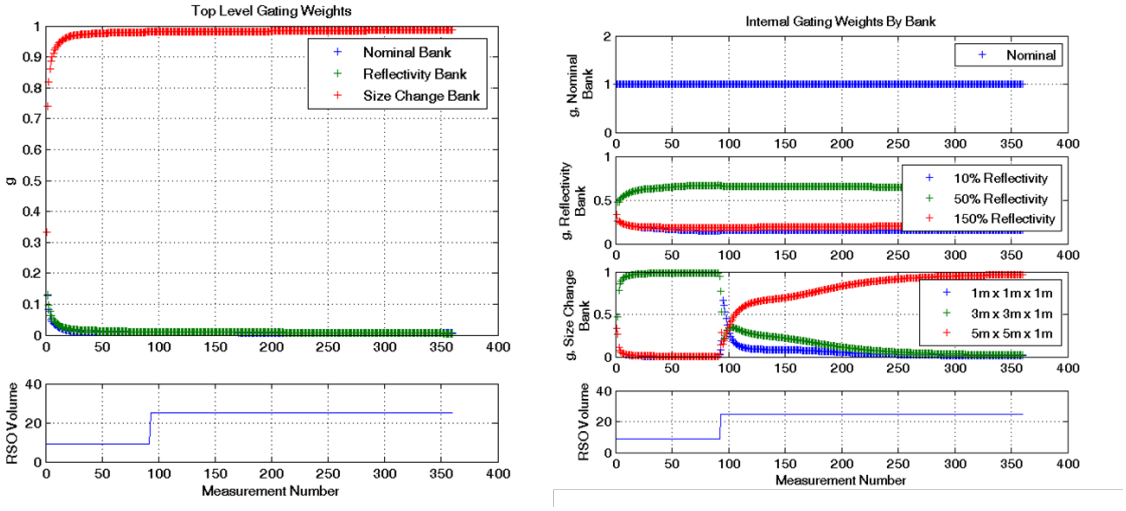

Fig. 5. Gating Weights for the size change detection experiment

Reflectivity Change Experiment

The filter bank arrangement for the experiment was the same as the one for the RSO size change experiments, and is described in Table 3. This experiment was meant to characterize the ability of the HME to detect the reflectivity of an RSO and changes in the RSO's reflectivity.

In this experiment, the truth model RSO initially has reflectivity equal to 50% of the nominal. After 25% of the observation window has elapsed, the reflectivity of the truth RSO model changes to 150% of its nominal value.

The gating weights for the reflectivity change experiment are given in Fig. 6. The HME performs ideally in this experiment. It selects Bank 2 as the correct bank, indicating that it determined the RSO reflectivity had deviated from the nominal. It also selects the correct filters both before and after the change in RSO reflectivity (Filter 2 and Filter 3 in Bank 2, respectively). It was also able to respond to the change in the RSO very quickly, after only three measurements.

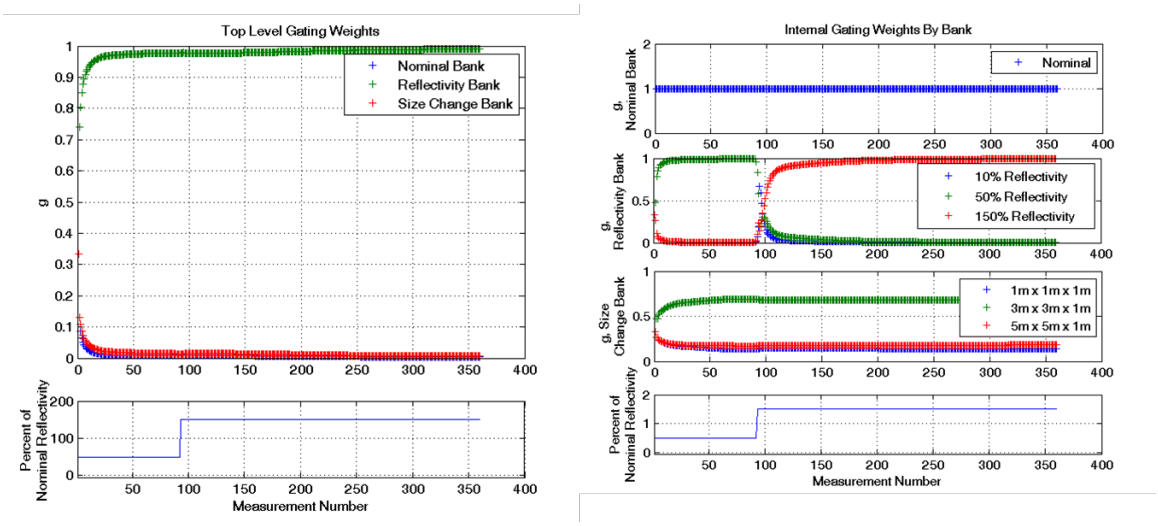

Fig. 6. Gating Weights for the reflectivity change detection experiment

Maneuver Time Detection Experiment

The final two experiments detailed in this paper were intended to test the ability of the HME to detect maneuvers performed by an RSO. In both of these experiments, the maneuvers performed were modeled as instantaneous changes in the translational velocity of the RSO. A maneuver by an RSO could indicate hostile intent, so being able to detect such a maneuver is important to SSA.

In the first maneuver detection experiment, the truth RSO performed a 0.1 km/s maneuver in the [0 1 0] direction (downrange) at step 92. The filter banks were arranged as explained in Table 4. This arrangement was designed to test the ability of the HME to determine at what time a maneuver was performed.

Table 4. Arrangement of Filters for Maneuver Time Detection Experiment

Bank	Time of Maneuver	Filter	Direction and Magnitude of Maneuver
1	Step 36	1	[0.1, 0, 0] km/s
		2	[0, 0.1, 0] km/s
		3	[0, 0, 0.01] km/s
2	Step 92	1	[0.1, 0, 0] km/s
		2	[0, 0.1, 0] km/s
		3	[0, 0, 0.01] km/s
3	Step 180	1	[0.1, 0, 0] km/s
		2	[0, 0.1, 0] km/s
		3	[0, 0, 0.01] km/s
Truth Model: [0,0.1,0] km/s maneuver performed at step 92 by a nominal RSO			

The results of this experiment can be seen in Fig. 7. Initially, before any of the experts model a maneuver by the RSO, all of the filters and banks are assigned equal weights. Once Bank 1, which is incorrect, begins to model various maneuvers after measurement 36, it is quickly de-weighted. After the correct bank (Bank 2), models various maneuvers at time step 92, it is correctly assigned the highest gating weight. The weight assigned to Bank 2 increases even further after Bank 3 models various maneuvers. The filter modeling the maneuver in the correct direction (Filter 2 of Bank 2) is also assigned the highest weight after the truth RSO performs a matching maneuver. The HME responded successfully within three measurements of the maneuver, indicating that it can detect the time at which a maneuver was performed.

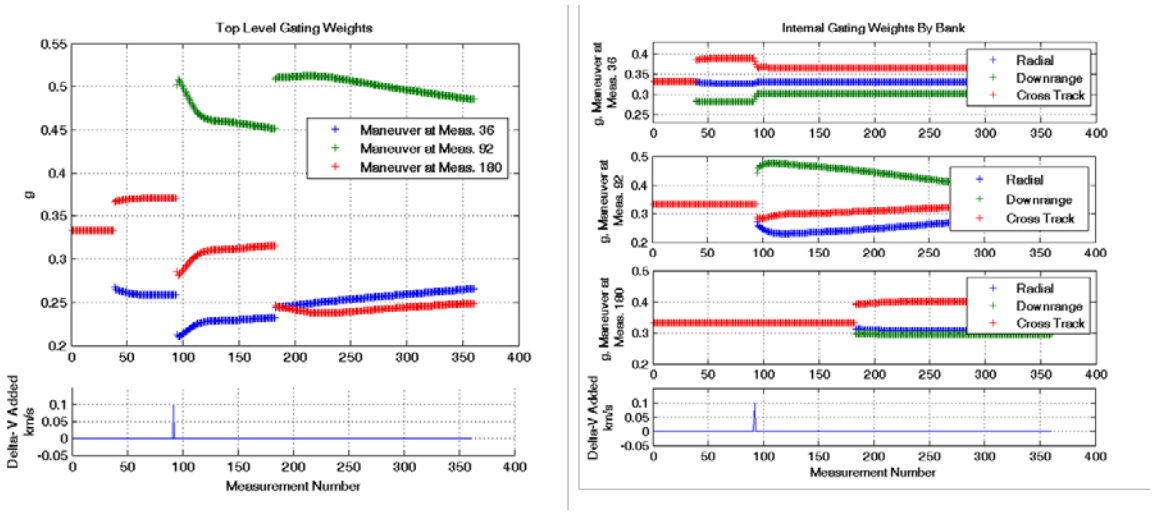

Fig. 7. Gating Weights for the Maneuver Timing Detection Experiment

Maneuver Direction and Magnitude Detection Experiment

The final experiment detailed in this report was designed to show that the HME could detect the direction and magnitude of a maneuver. The filter bank arrangement for this experiment is in Table 5. As in the previous experiment, the truth RSO performed a 0.1 km/s maneuver in the downrange direction after approximately 25% of the observation window, or 92 measurements, had elapsed.

Table 5. Arrangement of Filters for Maneuver Direction Detection Experiment

Bank	Direction	Filter	Magnitude of Maneuver
1	Radial	1	0.1 km/s
		2	0.01 km/s
		3	0.001 km/s
2	Downrange	1	0.1 km/s
		2	0.01 km/s
		3	0.001 km/s
Truth Model: [0,0.1,0] km/s maneuver performed at step 92 by a nominal RSO			

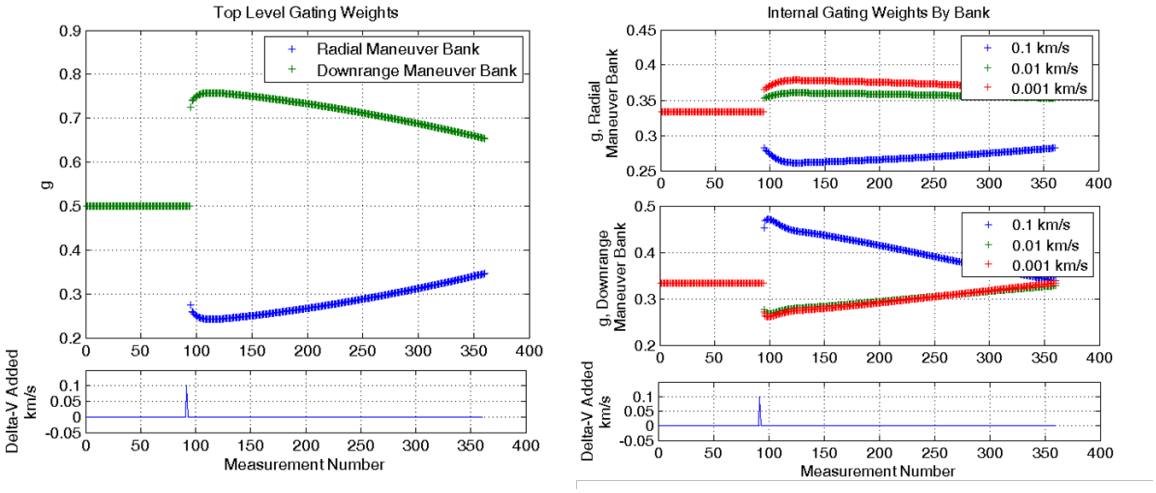

Fig. 8. Gating Weights for the Maneuver Direction and Magnitude Detection Experiment

The results of this experiment are in Fig. 8. Again, the HME was successful in detecting the maneuver. After the various experts modeled different maneuvers, the bank containing filters modeling maneuvers in the correct

direction (Bank 2), is assigned the highest weight. Also, the highest weight is assigned to the filter within Bank 2 (Filter 1) which matched the size of the maneuver performed by the truth RSO.

6. CONCLUSION

Table 6. Summary of Experiments described

RSO Characteristic	Detected Initial Value?	Detected that a change occurred	Determined the value of a characteristic after a change occurred?
Shape	Yes	Yes	No
Size	Yes	Yes	Yes
Reflectivity	Yes	Yes	Yes
Maneuver	The magnitude, direction, and timing of a maneuver were detectable.		

The performance of the HME was explored using a number of different experiments. These experiments included testing the ability of the HME to determine the size, shape and reflectivity of an RSO, detect changes in these characteristics, and to detect maneuvers performed by an RSO. The HME performed remarkably well in these experiments, demonstrating that it is feasible to detect RSO size, shape, reflectivity, and maneuvers using an HME.

The results obtained using the HME algorithm all used EKFs as the experts in the HME. The EKF may not do well in all cases, however, especially if propagation of the state of the RSO is highly nonlinear. In such a case, better results would probably be obtained using our UKF or even another, more advanced filter, such as a Gaussian Mixture Filter or Multiple Model Adaptive Estimator.

Also, as all of these results were for satellites in GEO, the measurements (and therefore the measurement residuals) varied only a very small amount during the 3-hour observation window. It would therefore be interesting and useful to test the performance of the HME in tracking an object in another orbit regime, such as LEO, or processing measurements of the same RSO at GEO taken over multiple nights. Another area suggested for further research is testing the performance of the HME in detecting the attitude and attitude rate of RSOs.

7. REFERENCES

1. Miller, J. G., “A New Sensor Allocation Algorithm for the Space Surveillance Network,” Military Operations Research, v. 12, n. 1, pp. 57 – 70, 2007.
2. Payne, T., Morris, R., “The Space Surveillance Network (SSN) and Orbital Debris”, AAS 10-012, 33rd Annual AAS Guidance and Control Conference, Breckenridge, Colorado, Feb 6-10, 2010.
3. Bate, R. R., Mueller, D.D., and White, J. E., *Fundamentals of Astrodynamics*, Dover Publications, Inc., Mineola, New York, 1971.
4. U.S. Joint Chiefs of Staff, Joint Publication 3-14, Space Operations, dated January 6, 2009, p. 32.
5. White, F.E., Data Fusion Lexicon, Joint Directors of Laboratories, Technical Panel for C3, Data Fusion Sub-Panel, Naval Ocean Systems Center, San Diego, 1991.
6. Crain, T. P., Bishop, R. H., and Ely, T., “Event Detection and Identification during Autonomous Interplanetary Navigation,” AIAA Journal of Guidance, Control, and Dynamics, Vol. 25, No. 2, 2002, pp. 394-403.
7. Linares, R., Crassidis, J. L., Jah, M. K., and Kim, H., “Astrometric and Photometric Data Fusion for Resident Space Object Orbit, Attitude, and Shape Determination via Multiple-Model Adaptive Estimation,” AIAA Guidance, Navigation, and Control Conference, Toronto, Canada, Paper No. AIAA 2010-8341, Aug. 2-5, 2010.
8. Ashikhmin, M. and Shirley, P., “An anisotropic Phong BRDF model,” J. Graph. Tools, Vol. 5, No. 2, 2000, pp. 25-32.
9. Ashikhmin, M. and Premoze, S., “Distribution-based BRDFs,” Unpublished Technical Report, University of Utah, 2007.