Resident Space Object Feature Identification and Attitude Detection using Hierarchical Mixtures of Experts

David E. Gaylor, Angelica Ceniceros
University of Arizona, Tucson, AZ 85721

Jessica T. Anderson, Elfego Pinon III
Emergent Space Technologies, Inc., Greenbelt, MD 20770

Space Domain Awareness (SDA) involves detecting, tracking, identifying and characterizing resident space objects (RSOs). This paper proposes an algorithm based on a Hierarchical Mixture of Experts (HME) for identifying RSO features and determining an RSO's attitude profile from astrometric and non-resolved photometric observations. This paper discusses the mathematical background of the HME; the assumptions, test scenarios, and results of processing simulated apparent magnitude and angles data. The results show that the HME is capable of detecting the size, shape, reflectivity, and maneuvers performed by an RSO. The HME is also shown to be capable of identifying and distinguishing between nadir-pointing, sun-pointing, and spinning objects even though none of the experts in the HME is directly estimating attitude.

• Introduction

The hazards to operational satellites are increasing due in part to the growing number of resident space objects (RSOs) including orbital debris. 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 [1]. In order to protect valuable space assets from the hazard of collision, it is necessary to observe, understand and predict the behavior of natural and manmade objects in orbit around the Earth. This involves detecting, tracking, identifying and characterizing RSOs.

Filtering algorithms for state estimation, such as the Extended Kalman Filter (EKF)[2] and the Unscented Kalman Filter (UKF)[3] have been used for tracking RSOs. However, Kalman filters require an accurate knowledge of the dynamics and measurement model parameters and uncertainties. Incorrect modeling can lead to large errors or filter divergence. For the case of RSO tracking, the size, shape, attitude and other characteristics that affect the dynamics are often not known. Furthermore, the RSO itself may perform unknown maneuvers. Thus we do not always have accurate knowledge of the dynamics necessary for successfully using the Kalman filter for RSO tracking.

Adaptive filtering attempts to address unknown changes in the environment. A framework for an adaptive filtering algorithm was introduced by Magill [4]. Magill presented an optimal estimation of a sampled, Gauss-Riarkov stochastic process. The process consists of certain parameters being initially unknown, and it assumes that the parameters come from a set that contains a finite number of possibilities which are known a priori. [4] In Magill's adaptive filtering algorithm, which later became known as the Multiple Model Adaptive Estimator (MMAE), measurements are fed into a bank of parallel Kalman filters. Each filter computes its own estimated state based on the provided measurements and models. Subsequently, with the assumption that the state and measurements are Gaussian, the probabilities for each filter are calculated using Bayes rule, thus providing a solution that is the weighted sum of the individual filter outputs. [4]

The use of light curves, the temporal brightness of an observed RSO, can be applied to small

and/or dim objects in high-altitude orbits (e.g. geosynchronous orbit). Light curve data is the time-varying sensor wavelength-dependent apparent magnitude of photons reflected from an object along the line-of-sight to an observer. The apparent magnitude of an RSO is a function of its size, orientation, and surface material properties. Consequently, one should be able to estimate these characteristics using appropriate algorithms. Attitude estimation using light curve data has been already demonstrated by Wetterer and Jah [5]. In Linares et al. [6], light curve data is used to estimate the shape of an RSO along with its rotational and translational states using an MMAE algorithm. The MMAE contained a bank of filters with different hypothetical candidate RSO shape models. By computing the likelihood associated with each hypothesis, the MMAE could determine which of the candidate shape models is most probable given the observations [6].

Although the MMAE has the advantage of having the bank filters work independently of each other, it also presents numerical issues when dealing with steady state operations when applying multiple parameter vectors with near optimal accuracy [7]. Problems or limitations of the MMAE approach include a) numerical underflow behavior over time, b) the fact that the unknown parameter(s) is assumed to be constant, and c) early measurements are weighted equally as later ones. These issues cause the MMAE to be limited in its ability to adapt to a changing environment [8].

Magill's algorithm was modified by Chaer, et al. to use a neural network approach, which is referred to as a mixture of experts (ME) [9]. Similar to the MMAE, the ME is composed of a finite number of filters running in parallel. Each models different realizations of unknown parameters. The filters can include different process models, process noise, and measurement noise. The ME selects the best filter using a learning network, i.e. a gating network capable of assigning a specific weight to each expert in order to provide the optimal mixture of models. It is able to respond rapidly to changing environments based on the most recent data and avoids numerical instabilities of MMAE. Also, it does not assume the optimal filter is in the bank thus good results can be obtained with smaller filter banks. The ME adapts individual filters to better match incoming data via search algorithm in a feedback loop and uses gating network to weight the filter outputs [9].

Within this framework, we define a hierarchical mixture of experts (HME) [7] as an ME that is designed to include multiple layers and groupings. Crain applied the HME as part of a real-time interplanetary orbit determination monitoring algorithm for detecting and identifying changes in the spacecraft dynamic and measurement environments. He conducted experiments of the Mars Pathfinder cruise trajectory environment using simulated range and Doppler measurements. These experiments demonstrated the ability of the HME to identify three environment macromodes, i.e. 1) unmodeled impulsive maneuvers, 2) changes in the solar radiation pressure dynamics, and 3) changes in the measurement noise strength [7].

In this paper we develop a new HME-based algorithm to process astrometric and photometric data to determine the shape, size, reflectivity, and attitude profile on an RSO as well as determine the time, magnitude and direction of maneuvers. The proposed algorithm uses a HME composed of banks of EKFs that model different hypotheses about the RSO. The results of our simulation show that the HME is able to successfully determine the shape, size, reflectivity, and attitude profile of an RSO as well as determine the time, magnitude and direction of maneuvers.

The organization of this paper is as follows. In Section II, the models for RSO shape, orbital dynamics, and attitude dynamics employed by the banks of filters are discussed and the HME structure is outlined and discussed. In Section III, simulation results for nine possible RSO detection scenarios are provided, including shape change detection, maneuver change detection and attitude state detection. Finally, in Section IV, summary and conclusions are reported.

• Methods

• Simulation Overview

The simulation used in this work consists of a truth model and estimation model. The truth model propagates 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) as shown in Fig. 1.

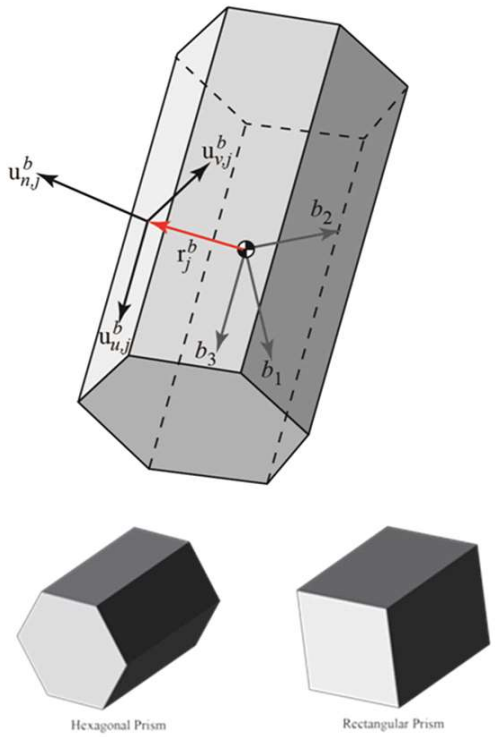

Fig. 1 RSO Modeled as a Collection of Flat Plates

The truth RSO state is used to generate simulated measurements of an RSO, including apparent magnitude, right ascension, and declination. The right ascension and declination are angles that define the location of the object of interest on the celestial sphere, as defined at the J2000 epoch. 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. The details of the apparent magnitude model as provided in section II.B.

The simulation includes the capability to model three different possible attitude profiles for the RSOs modeled. These attitude profiles are major-axis spinning, nadir-pointing, and sun-pointing. These profiles are modeled by calculating a small moment on the spacecraft at each time step, which is then propagated by a 6 degree of freedom propagator to cause the RSO model to point in the direction indicated by the selected attitude profile.

The estimation model consists of the HME with banks of extended Kalman filters (EKFs) that process the simulated measurements and produce orbit estimates and HME gating weights. A description of the HME algorithm is provided in Section II.C.

- **Apparent Magnitude Model**

The "brightness" of a celestial object is actually measured in terms of the radiant flux, denoted by *F*, received at the sensor from the object. The radiant flux is the total amount of light energy of all wavelengths that crosses a unit area oriented perpendicular to the direction of the light's travel in unit time. It can be shown that an object's apparent magnitude *m* is related to the radiant flux *F* received from the object by

$$m = m_{sun} - 2.5\log_{10}\left(\frac{F}{F_{sun}}\right)$$

where the apparent magnitude and radiant flux of the Sun is $m_{sun} = -26.74$ and $F_{sun} = 1368\ W/m^2$ respectively.

Let a given RSO be represented as a series of N flat plates, as depicted in Fig. 1, where the position vector of the j^{th} plate $\mathbf{r}_j^b$, expressed in body coordinates, is defined as the distance of the plate's center-of-area relative to the RSO's center-of-mass. The orientation of the j^{th} plate with respect to the body frame is represented by three orthogonal unit vectors: $\mathbf{u}_{n,j}^b$ represents the direction normal to the plate and $\mathbf{u}_{u,j}^b$, $\mathbf{u}_{v,j}^b$ represent the u and v directions, respectively, of the uv-plane, i.e., the plane formed by the plate.

In general, the radiant flux received from a RSO is due to the reflection of sunlight from the illuminated surfaces of the RSO. As illustrated by Fig. 2, the geometrical configuration between the Sun, object, and observer will determine the amount of reflected radiant flux received by the sensor, which in turn determine if the RSO is observable. In Fig. 2, the unit vectors that describe the inertial orientation of the j^{th} plate have been transformed from body to inertial coordinates via the standard transformation

$$\mathbf{u}_{k,j}^i = \left[\mathbf{T}\left(\mathbf{q}_i^b\right)\right]^T \mathbf{u}_{k,j}^b, \quad k = n,u,v$$

where $\mathbf{T}\left(\mathbf{q}_i^b\right)$ is the inertial-to-body transformation matrix (using quaternion parameterization) and $\mathbf{q}_i^b$ is the RSO's inertial-to-body attitude quaternion. The direction to the Sun relative to the RSO is represented by the unit Sun direction vector expressed in inertial coordinates and defined as

$$\mathbf{u}_{sun}^{i} = \frac{\mathbf{r}_{sun}^{i} - \mathbf{r}^{i}}{\left\| \mathbf{r}_{sun}^{i} - \mathbf{r}^{i} \right\|}$$

where $\mathbf{r}_{sun}^{i}$ and $\mathbf{r}^{i}$ are the inertial position vectors of the Sun and RSO, respectively. The direction to the observer relative to the RSO is represented by the unit observer direction vector expressed in inertial coordinates and defined as

$$\mathbf{u}_{o}^{i} = \frac{\mathbf{r}_{o}^{i} - \mathbf{r}^{i}}{\left\| \mathbf{r}_{o}^{i} - \mathbf{r}^{i} \right\|}$$

where $\mathbf{r}_{o}^{i}$ is the inertial position vector of the Earth-based observer. The unit Sun and observer direction vectors form the Sun-observer plane, as illustrated in Fig. 2. The unit half vector $\mathbf{u}_{h}^{i}$ bisects the angle θ between the unit Sun and observer direction vectors and is defined as

$$\mathbf{u}_{h}^{i} = \frac{\mathbf{u}_{sun}^{i} + \mathbf{u}_{o}^{i}}{\left\| \mathbf{u}_{sun}^{i} + \mathbf{u}_{o}^{i} \right\|}$$

From Fig. 2, it follows that the j^{th} surface is illuminated by the Sun and a reflection occurs only when the angle between the unit Sun direction vector $\mathbf{u}_{sun}^{i}$ and the unit vector normal to the surface $\mathbf{u}_{n,j}^{i}$ is less than 90 degrees or (equivalently) when the dot product of these two vectors is positive and greater than zero. Moreover, the reflection is observable by an Earth-based observer only when the dot product between the unit observer direction vector $\mathbf{u}_{o}^{i}$ and the unit vector

normal to the surface $\mathbf{u}_{n,j}^{i}$ is also positive and greater than zero. If these two criteria are met for even one of the N surfaces, then the reflection geometry is such that the RSO is observable and a measurement is obtained. Mathematically, this observability condition ψ is represented by

$$\psi\left(\xi_{1,j},\xi_{2,j}\right)=\begin{cases}1 \quad (observable) & if\ \xi_{1,j}>0\ and\ \xi_{2,j}>0\\ 0 \quad (not\ observable) & else\end{cases}$$

where $\xi_{1,j}=\left[\mathbf{u}_{sun}^{i}\right]^{T}\mathbf{u}_{n,j}^{i}$ and $\xi_{2,j}=\left[\mathbf{u}_{o}^{i}\right]^{T}\mathbf{u}_{n,j}^{i}$.

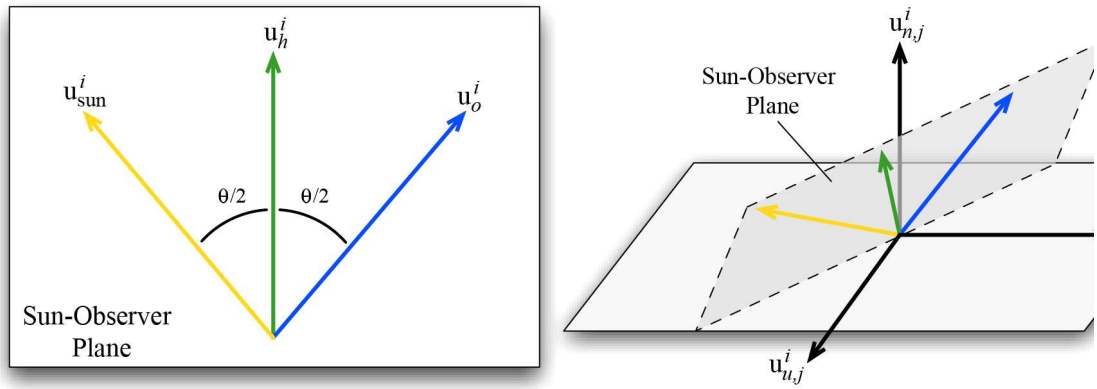

Fig. 2 Reflection Geometry.

The Bidirectional Reflectance Distribution Function (BRDF) models the amount of light reflected by each plate, t, as shown in Fig. 3. More specifically, BRDF is defined as the ratio between reflected (directional) radiance and incoming surface irradiance [10]. The amount of directional radiance reflected by

an object 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 of the incoming irradiance. Thus, RSO apparent magnitude is highly dependent on shape and attitude.

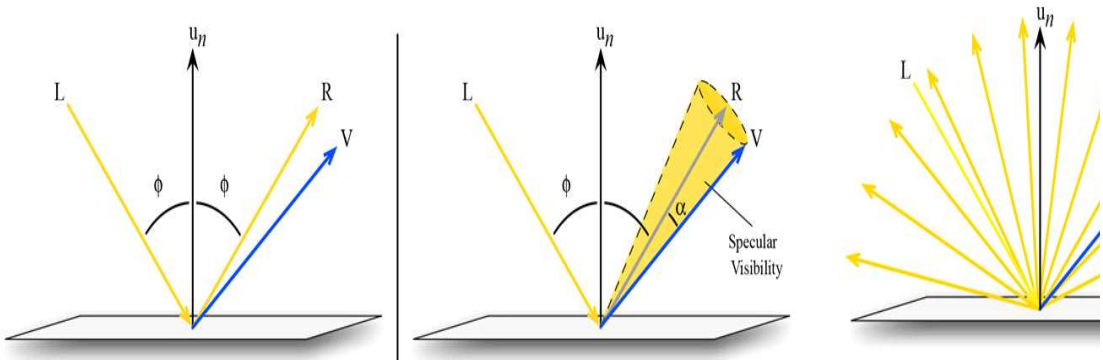

Fig. 3 BRDF Defines How Diffuse and Specular Components are Reflected

The measured radiant flux F_j of the j^{th} Sun-illuminated plate is related to the intensity of light $L_{o,j}$ reflected in the direction of the observer and inversely proportional to the square of the distance d_j from the plate to the observer,

$$F_j = \frac{L_{o,j}}{d_j^2}$$

The distance d_j from the plate to the observer, neglecting the distance from the RSO center-

of-mass to the plate's center-of-area, is simply

$$d_j = \left\| \mathbf{r}_o^i - \mathbf{r}^i \right\|$$

The intensity of reflected light $L_{o,j}$ or radiance is proportional to the intensity of incident sunlight $L_{i,j}$ and the reflectance behavior of the plate,

$$L_{o,j} = \rho_j L_{\mathrm{i},j}$$

where ρ_j is the BRDF. The intensity of sunlight light (irradiance) $L_{i,j}$ incident upon a plate of surface area A_j is related to the solar radiant flux F_{sun} and the cosine of the incident angle,

$$L_{i,j} = F_{sun} A_j \left[\mathbf{u}_{sun}^i \right]^T \mathbf{u}_{n,j}^i$$

Here $\left[\mathbf{u}_{sun}^i \right]^T \mathbf{u}_{n,j}^i$ has been used in place of the cosine of the incident angle. Finally, since multiple plates could conceivably be illuminated and observable at a given time, it then follows that the measured radiant flux *F* of the RSO is given by

$$F = \sum_{j=1}^{N'} F_j = \frac{F_{sun}}{\left\| \mathbf{r}_o^i - \mathbf{r}^i \right\|^2} \sum_{j=1}^{N'} \rho_j A_j \left[\mathbf{u}_{sun}^i \right]^T \mathbf{u}_{n,j}^i$$

where $N' \left(N' < N \right)$ is the number of illuminated plates that meet the observability criteria defined in Eq. 11. It is important to note that this expression does not account for shadowing effects.

The following improvements were made to the apparent magnitude model described by Linares, et al, [6]: 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 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 inaccurate. 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.

In searching the literature for a suitable BRDF model, we set out to find a BRDF that is physically plausible, i.e., based upon physical laws, with reciprocity and conservation of energy being the most important. Reciprocity means that if the incoming and outgoing directions of light are reversed, then the value of the BRDF remains unchanged. Conservation of energy states that the sum of the reflected, absorbed, and transmitted light must always equal the amount of incident light. In other words, the amount of reflected light can never exceed the amount of incident light. The BRDF selected for this research meets the above criteria for being physically plausible. It is also a position-invariant BRDF, which means that it is assumed that the reflectance properties of the material do not vary with spatial position. In other words, it assumes that the material is homogeneous, which is not overly restrictive for this research. The selected BRDF model was originally developed by Ashikhmin and Shirley [11]. It is hereafter denoted by ρ, which has units of inverse steradians (sr^{-1}), and is decomposed into a specular component ρ_s and a diffuse component ρ_d

$$\rho = \rho_s + \rho_d$$

Specular reflection occurs when light incident upon a surface appears to be focused in one direction. In other words, there is a bright spot, called a specular highlight, which is more readily

apparent on shiny surfaces. For the ideal reflector, such as a mirror, the angle of incidence $\phi\phi$ equals the angle of specular reflection, as shown in Fig. 4.

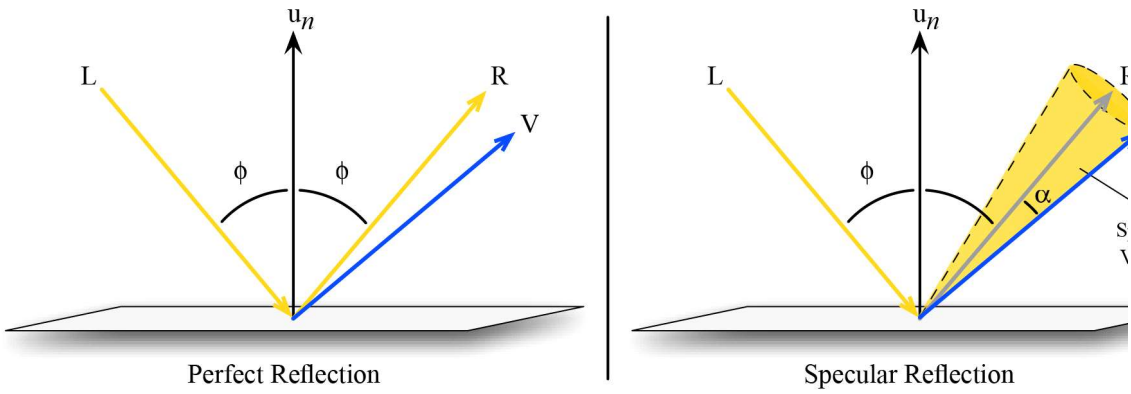

Fig. 4 Perfect or ideal reflection versus specular reflection

If R is the direction of the specular reflection and V is the direction of the viewer, then for an ideal reflector (i.e., perfect reflection) the specular reflection is visible only when R and V coincide. For real objects, however, the specular reflectance can be seen even if R and V do not coincide, i.e., it is visible over a range of values that form a cone about the R direction, as depicted in Fig. 4. In general, the shinier the surface is the smaller the range is for specular visibility. Consequently, a specular reflectance model must have a maximum intensity at R, with an intensity that decreases as a function of α the angle between R and V. The anisotropic specular reflection component in Eq. 12 is defined as

$$\rho_s = \frac{\sqrt{(n_u+1)(n_v+1)}}{8\pi} \frac{\left(\left[\mathbf{u}_h^i\right]^T \mathbf{u}_n^i\right)^z}{\left(\left[\mathbf{u}_{sun}^i\right]^T \mathbf{u}_n^i\right) + \left(\left[\mathbf{u}_o^i\right]^T \mathbf{u}_n^i\right) - \left(\left[\mathbf{u}_{sun}^i\right]^T \mathbf{u}_n^i\right)\left(\left[\mathbf{u}_o^i\right]^T \mathbf{u}_n^i\right)} F_{ref}$$

where z is the Phong exponent given by

$$z = \frac{n_u \left(\left[\mathbf{u}_h^i \right]^T \mathbf{u}_u^i \right)^2 + n_v \left(\left[\mathbf{u}_h^i \right]^T \mathbf{u}_v^i \right)^2}{1 - \left(\left[\mathbf{u}_h^i \right]^T \mathbf{u}_n^i \right)^2}$$

and F_{ref} is the Fresnel fraction given by

$$F_{ref} = R_s + \left(1 - R_s\right) \left(1 - \left[\mathbf{u}_{sun}^i \right]^T \mathbf{u}_h^i \right)^5$$

Moreover, n_u and n_v are parameters that represent the roughness of the reflecting material along the *u* and *v* directions, respectively, of the *uv*-plane. For small values $\left(n_u = n_v = 10\right)$, the material is like rough metal; whereas for large values $\left(n_u = n_v = 10^4\right)$, the material is like a perfect mirror. In Eq. 15, R_s represents the material's specular reflectance for normal incidence, where $0 < R_s < 1$.

Diffuse reflection occurs when light incident upon a surface scatters isotropically (i.e., the same in all directions) such that the apparent brightness of the surface to an observer is the same regardless of the viewing angle. In other words, diffuse reflection is the reflection of light from an uneven or granular surface such that the incident ray appears to be reflected in a number of directions simultaneously, as illustrated in Fig. 5. The non-Lambertian diffuse reflection component is defined as

$$\rho_d = \frac{28R_d}{23\pi}(1-R_s)\left[1-\left(1-\frac{\left[\mathbf{u}_{sun}^i\right]^T\mathbf{u}_n^i}{2}\right)^5\right]\left[1-\left(1-\frac{\left[\mathbf{u}_o^i\right]^T\mathbf{u}_n^i}{2}\right)^5\right]$$

where R_d is the diffuse albedo of the surface, where $0 < R_d < 1$ and a value of 1 represents total reflectance (i.e., no absorption).

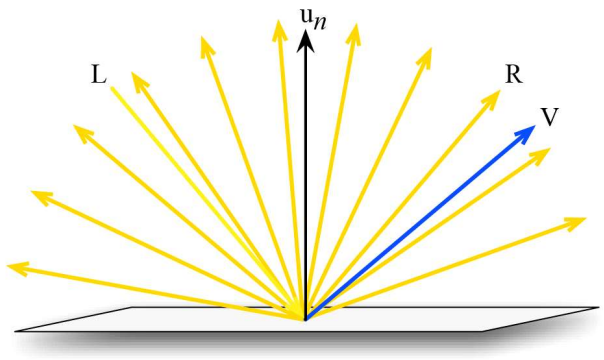

Fig. 5 Diffuse reflection

The Linares model employed the BRDF developed by Ashikhmin and Shirley [11]. 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 Ashikhmin and Shirley BRDF was found in an unpublished paper by Ashikhmin and Premoze [12], 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, we used the BRDF described in [12], and were able to derive, implement and validate the analytic Jacobian of the apparent magnitude measurement model in our algorithms.

- **Hierarchical Mixture of Experts (HME): Theoretical Approach and Algorithm**

Description

- *Mixture of Experts (ME) Algorithm*

The foundation for the HME is the Mixture of Experts (ME) algorithm. A flowchart of the ME with a bank of L experts is presented in Fig. 6. It can be seen that the total output, y, is obtained through Eq. 17, which consists of the output pattern and the gating weight, y_i and g_i respectively, of the ith expert.

$$y = \sum_{i=1}^{L} g_i y_i$$

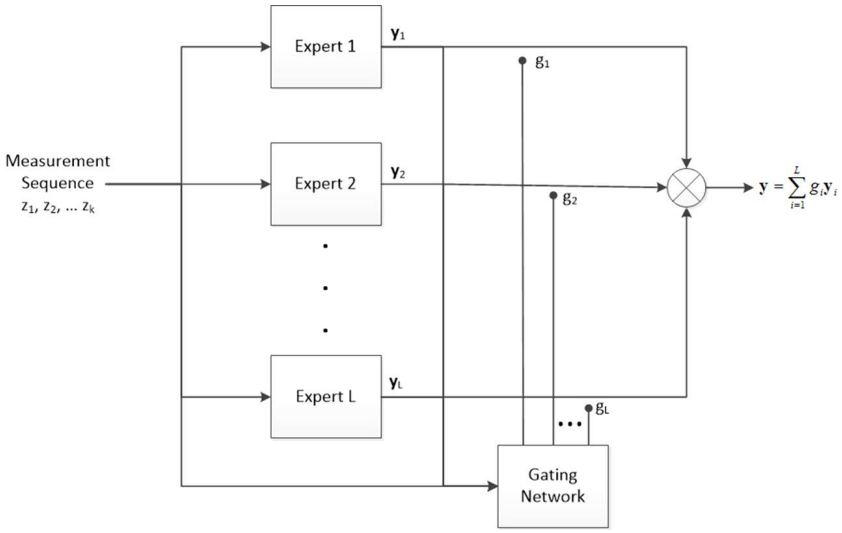

Fig. 6 Schematic of ME with a Bank of *L* Experts

The weighted sum of the conditional probabilities, given in Eq. 18, provides the probability distribution of the input pattern for the bank.

$$f(z) = \sum_{i=1}^{L} f(z \mid i) g_i$$

The optimal gating weight vector, $g = [g_1 \quad g_2 \quad \text{L} \quad g_L]$ maximizes the distribution in Eq. 18. [9]

The schematic of a gating network as shown in Fig. 7a shows that the gating weights activate a fully connected single layer neural network. In addition, Fig. 7b shows the action process of the *i*th neuron. The *i*th neuron consists of the gating network component associated with the *i*th expert.

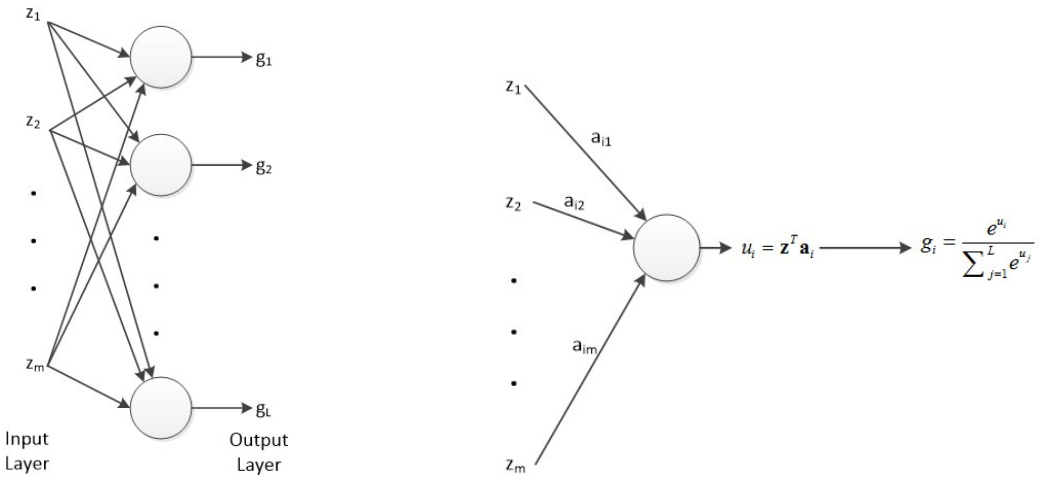

(a) (b)

Fig. 7 General Gating Network

Through the use of the softmax function, each neuron, u_i is activated by Eq. 19, which calculates the gating weights:

$$g_i = \frac{e^{u_i}}{\sum_{j=1}^{L} e^{u_j}}$$

The *i*th filter's conditional probability is given by

$$f(z_k \mid \alpha_i) = \frac{1}{\sqrt{2\pi W_i}} e^{-\frac{r_i^2}{2W_i}}$$

Eq. 20 consists of r_i and W_i which are the scalar prefit residual and the innovations covariance for the *i*th filter, respectively. The a posteriori probability of the *i*th filter is calculated by

$$h_i = \frac{f(z_k \mid \alpha_i) g_i}{\sum_{j=1}^{L} f(z_k \mid \alpha_j)\, g_j}$$

The posterior probability, h_i is then placed back into the gating network calculations. The gradient learning rule is given by

$$a_i^+ = a_i + \eta\left(h_i - g_i\right)$$

where η is the learning rate parameter which is adjustable [7].

The ME learns by updating the synaptic weights after each set of measurements is processed. The new synaptic weight is computed using Eq. 22. The filter pre-fit residual and innovation covariance go into h as well as the gating weight. The learning rate parameter is a user-selected parameter which controls how quickly the ME reacts to change.

- *HME algorithm*

An HME is comprised of a number of experts (estimation filters, in our algorithm), arranged in banks regulated by gating networks. The proposed HME algorithm is based on an algorithm developed by Crain to

process Mars Pathfinder data [7]. 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.

The HME can be used to estimate the state of an object, even when some of the characteristics or forces acting on that object are unknown. Each of the experts in the HME processes the same measurements, but uses different dynamics and/or measurement models in order to predict future measurements. The residuals between the processed measurements and the predicted measurements are used, along with the associated innovation covariance, to assign weights to each of the experts and banks using gating networks. Higher weights are given to the banks and experts that better predict the incoming measurements. A diagram of the HME arrangement is shown in Fig. 8.

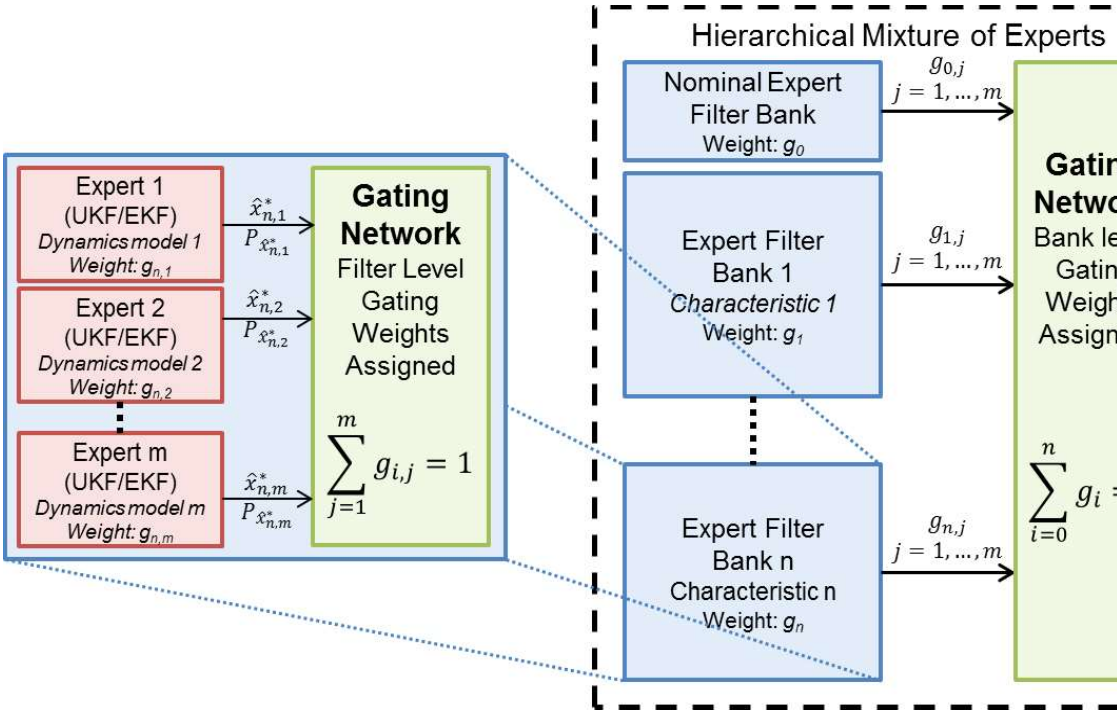

Fig. 8 Two Layer Hierarchical Mixture of Experts

For the HME, a gating network acts on the outputs of the experts or banks of experts on each layer. 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 can have a different model of the object's state dynamics. For the proposed algorithm, Extended Kalman filters (EKFs) were used as the experts. Each EKF uses a different realization of some unknown RSO system parameters, such as object 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 filters such as unscented Kalman filters (UKFs) is possible, and could potentially improve 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 can 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. Examples of macro-modes include a change in RSO reflectivity, a change in size, attitude profile, 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. Here, each filter modeled 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. An example HME arrangement modeling various micro- and macro-modes is shown in Fig. 9.

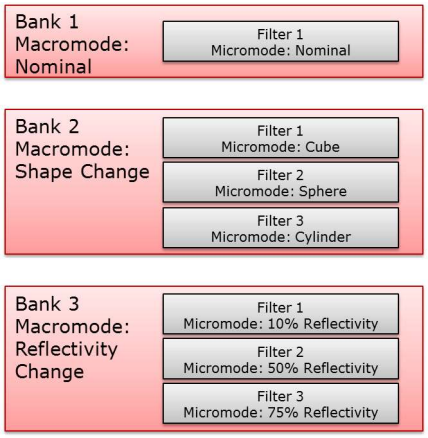

Fig. 9 Examples of micro- and macro-modes in the HME

A recursive weighting function (i.e., the gating function) is then used to assign weights 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 contains filters modeling various sizes of cube shaped objects, and if that filter bank has a high weight relative to the other filter banks modeling other shapes, then this indicates that

the object is most-likely cube-shaped. Similarly, 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 probable size of the object in the first example, and the magnitude and direction of the propulsive thrust maneuver in the second example.

Similar filter banks can be realized for RSOs undergoing component articulation, orientation changes, and other characteristics that can be detected and classified. The HME 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, configuration, or attitude profile) 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. 8 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.8). 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, can be 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. The selection of this parameter in our HME algorithm is currently a matter of trial and error, although it has been observed that using a single value between 1 and 3 for both layers of the HME offers good results for most of the cases studied.

The HME is similar to the MMAE algorithm in that it uses banks of filters. However the HME gating network does not learn in a recursive multiplicative manner like the MMAE filter bank weighting function. In the MMAE, this is reflected in the computation of the conditional densities for each filter. This is due to the assumption that all possible realizations of the unknown parameter vector (including the optimal) are in the bank. This assumption is not made in the HME.

Typically, the HME gating network output for the best performing filter quickly reaches a value close to unity. Moreover, the gating network readily switches its selection to another filter in the bank when performance of the current best filter degrades due to a change in the dynamics or measurements. The switch occurs quickly due to the gating network property of geometrically decreasing the impact of past measurements. In other words, the gating network has a fading memory. In contrast, the MMAE filter bank never forgets. It gives as much weight for early measurements as later ones. This is a direct result of the way the conditional densities are computed which leads to a longer time for it to switch filters.

Another advantage when using the gating network is numerical stability. In the MMAE, a numerical underflow problem can be encountered in the computation of conditional densities. This results from multiplying exponential terms repeatedly. This problem does not exist when

using a gating network because learning is achieved recursively in an additive manner.

• Results and Discussion

• Experiment Configuration

For the experimental results given below, a 100 kg RSO in a near-geosynchronous orbit is simulated with orbital elements specified in Table 1.

Table 1 Orbital Elements for tracked RSO

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

There were six ground stations modeled in the simulation, although measurements were primarily available from the station in Chile. The latitude, longitude and altitude of the used ground stations are given in Table 2.

Table 2 Available ground stations in the HME simulation

Ground Station	Latitude	Longitude	Altitude
Chile	-30.6908°	289.1937°	2220.00 m
Tenerife	28.1887°	-17.5864°	2400.00 m
Maui	20.7081°	-156.2575°	3060.74 m
Texas	30.6714°	-104.0225°	2070.00 m
South Africa	-32.3783°	20.8105°	1798.00 m
Australia	-31.2270°	149.0673°	1164.00 m

All ground stations had a field of view of 50° about the zenith. The simulated measurements included right ascension, declination and apparent magnitude. The noise on the measurements received by all ground stations had a 1-sigma strength of 1 arcsec for right ascension and

declination and 0.0067 for apparent magnitude. For each experiment run, the measurements were available at intervals of 15 seconds to 2 minutes. The RSO was observed over periods ranging from 1 to 8 hours in duration. All experiments were performed with the HME learning parameter set to 1.75 for both the top level and filter level neural networks, unless otherwise noted. This parameter was determined experimentally.

The procedure to run the 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. The EKFs process the simulated measurements to generate the inputs for the HME algorithm, which 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.

The HME code 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 ten 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, detecting the direction and magnitude of an impulsive maneuver, discriminating whether an RSO is spinning or nadir pointing and detecting the correct spin rate for a spinning RSO, discriminating between a nadir or sun pointing RSO, detecting an attitude profile change from nadir-pointing to spinning, and detecting a change in spin rate.

- **Size Change Experiment**

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. The HME configuration was carefully selected to be used for the size change experiment and the reflectivity change experiment. 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. The HME used three filter banks, including a nominal bank, one with various values of reflectivity, and one with models of several different sizes of RSOs. The arrangement of the HME for this experiment is given in Table 3. 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.

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 shown in Fig. 10. 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.

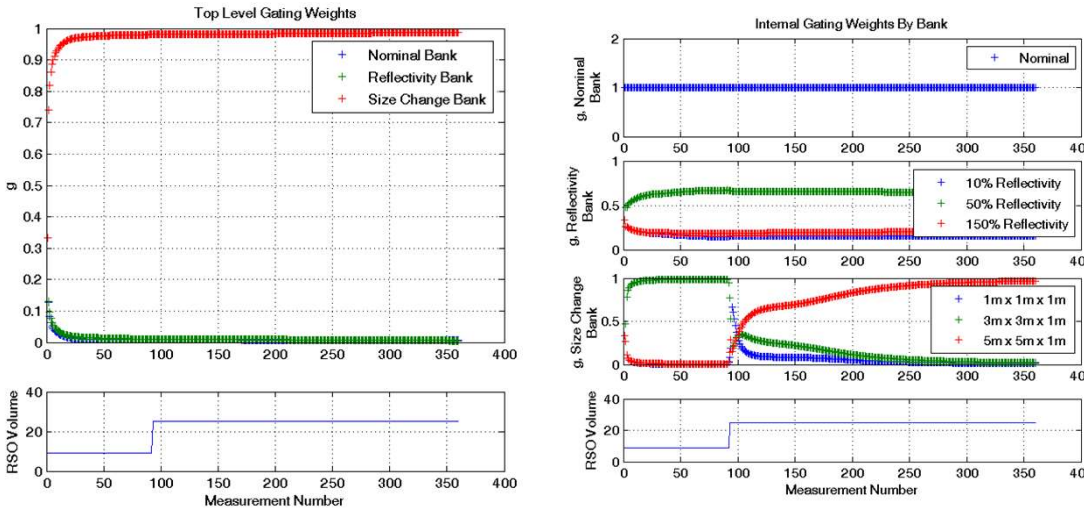

Fig. 10 Gating Weights for the size change detection experiment

- **Reflectivity Change Experiment**

The purpose of this experiment was to characterize the ability of the HME to detect the reflectivity of an RSO and changes in the RSO's reflectivity. 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. 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. 11. 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.

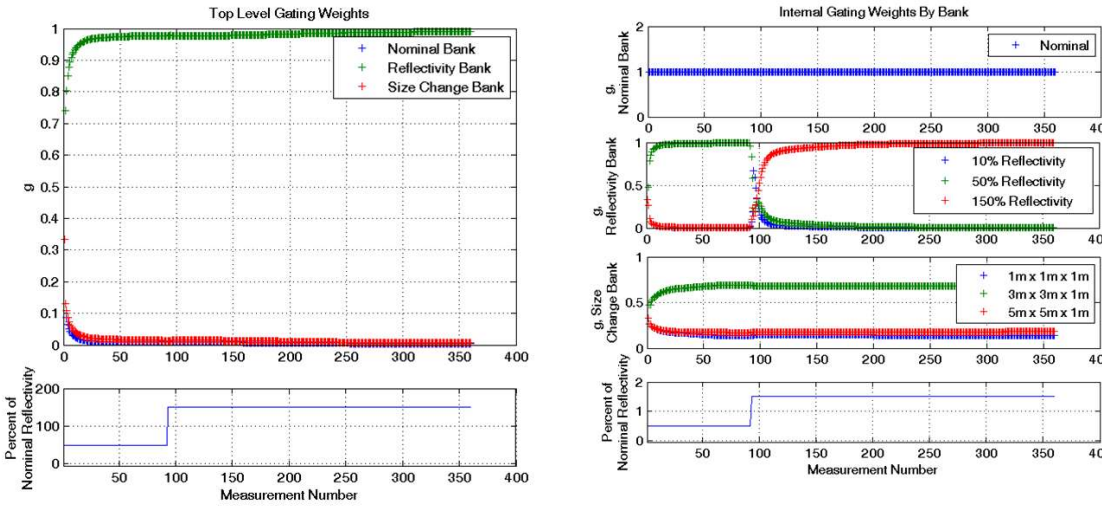

Fig. 11 Gating Weights for the reflectivity change detection experiment

- **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 experiment, 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. 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 4. The gating weights for both the filter banks and the individual experts for the shape change experiment are given in Fig. 12.

Table 4 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 with 2m x 4m rectangular sides at time step 92.			

At the beginning of the experiment, the highest gating weights were assigned to Bank 1, the bank containing cuboid RSOs, and to Filter 2 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.

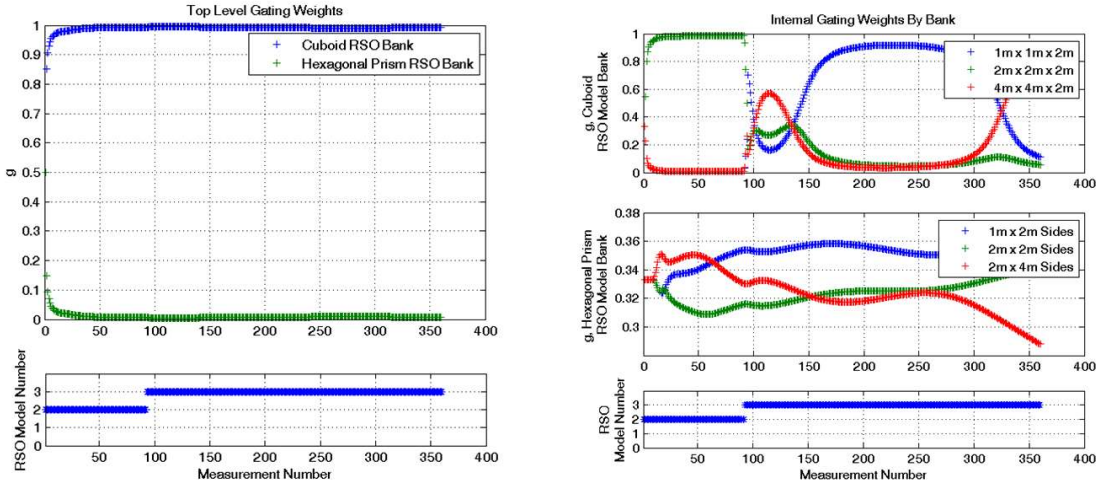

Fig. 12 Gating weights for the shape change detection experiment

- **Maneuver Time Detection Experiment**

The next two experiments 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 along track direction at step 92. The filter banks were arranged as shown in Table 5. This arrangement was designed to test the ability of the HME to determine at what time a maneuver

was performed.

Table 5 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 108	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 are shown in Fig. 13. 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.

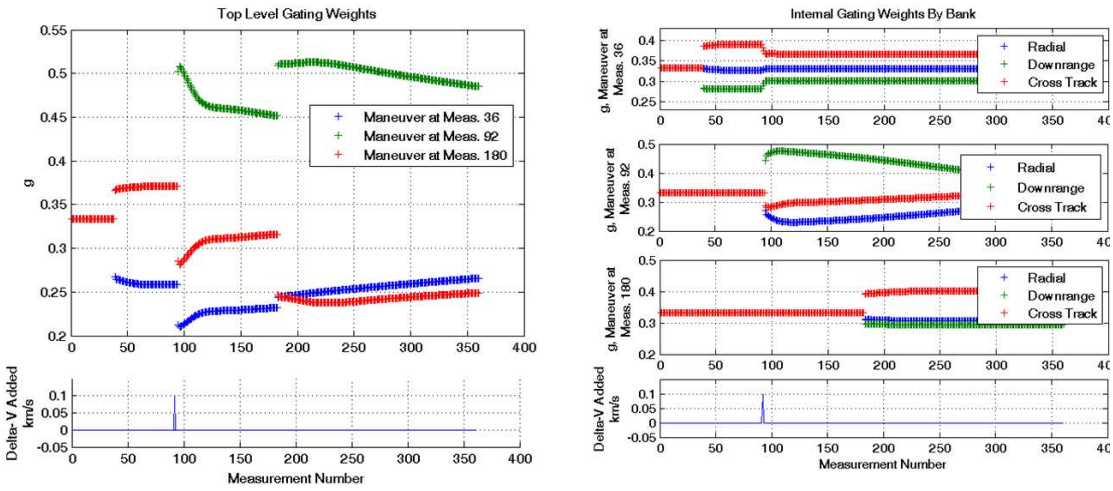

Fig. 13 Gating Weights for the Maneuver Timing Detection Experiment

- **Maneuver Direction and Magnitude Detection Experiment**

This experiment 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 6. 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 6 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

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			

The results of this experiment are shown in Fig. 14. 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.

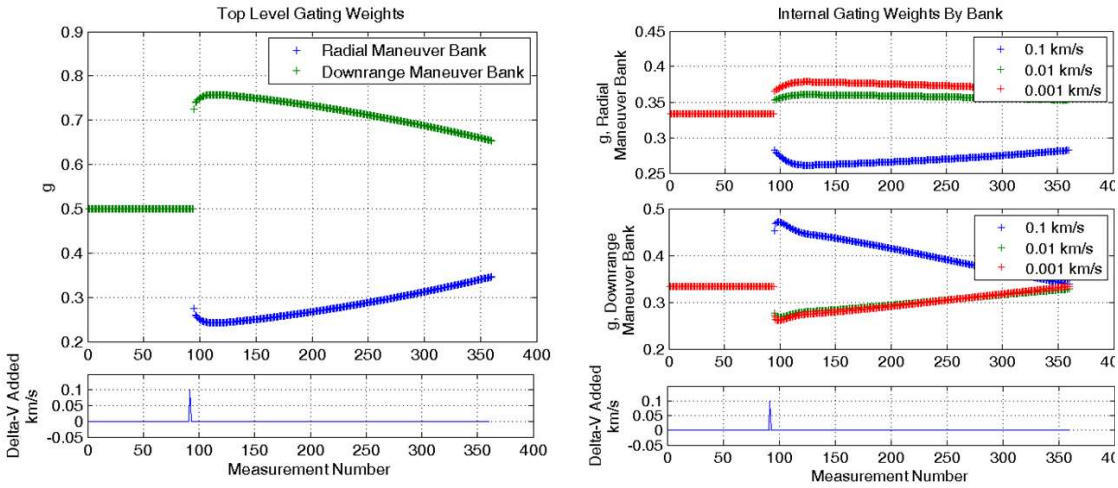

Fig. 14 Gating Weights for the Maneuver Direction and Magnitude Detection Experiment

- **Nadir-Pointing vs Spinning RSO Detection Experiment**

The purpose of this experiment was to determine if the HME could discriminate whether an RSO was spinning or nadir pointing and detect the correct spin rate for a spinning RSO. The HME configuration for this experiment, shown in Table 7, included one bank modeling nadir-pointing RSOs, containing a single expert and one bank modeling major-axis spinning RSOs, which contained experts using different spin rates for the RSO. The truth model for this experiment was a hexagonal prism with 2m x 4m sides which was spinning at a constant rate (0.3 rad/s) about the major axis. The learning parameter for both layers of the HME was 1.75. The top

level gating weights (top) and the filter level gating weights (bottom) for a spinning hexagonal prism RSO truth model indicate that the HME determined the RSO was spinning and selected the correct spin rate.

Table 7 Configuration of HME to Detect Nadir-Pointing or Spinning RSOs

Bank	Attitude Profile	Expert	Spin Rate
1	Nadir-Pointing	1	N/A
2	Major Axis Spin	1	(0, 0, 0) rad/s
		2	(0, 0, 0.3) rad/s

Fig. 15 shows the top and filter level gating weights for the attitude detection HME. The top level gating weights (top) and the filter level gating weights (bottom) for a spinning hexagonal prism RSO truth model indicate that the HME successfully determined the RSO was spinning and selected the correct spin rate. In both cases, an observation window of 1 hour was used, with measurements every 15 seconds. The top level gating weights for the case with the spinning truth RSO show that the highest weight is immediately given to correct bank, modeling major-axis spinning RSOs. The highest weight remains with this bank throughout the simulation. The highest expert-level weight is also immediately given to the expert modeling an RSO with the correct spin rate of 0.3 rad/s. This expert is given the highest weight throughout the simulation.

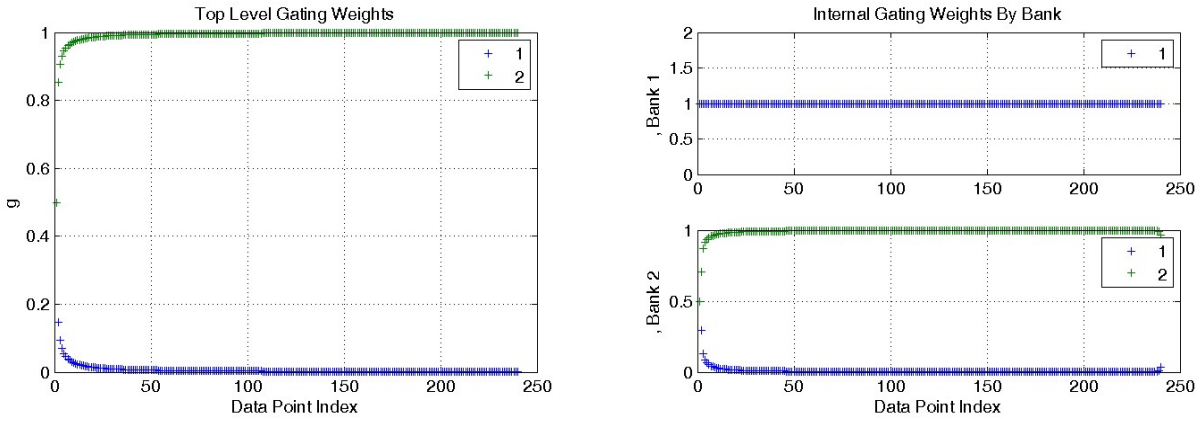

Fig. 15 Gating Weights for the Nadir-Pointing vs Spinning RSO Detection Experiment

- **Unmodeled True Spin Rate Experiment**

We performed an experiment where the correct spin rate was not modeled by any of the filters in the HME to determine how the HME would perform in this case. The spin rate of the modeled RSO was between two of the values modeled in the HME. Ideally, the HME would split the highest gating weights between those filters modeling values on either side of the "truth" spin rate. The configuration of the HME is given in Table 8. For this experiment, the truth RSO is a 2m x 1m x 3m cuboid spinning about its major axis with a rate of 0.5 rad/s.

Table 8 HME Configuration for Unmodeled True Spin Rate Experiment

Bank	Attitude Profile	Expert	Spin Rate
1	Major Axis Spin	1	(0, 0, 0.3) rad/s
		2	(0, 0, 0.7) rad/s
2	Inertially Pointed	1	(0, 0, 0) rad/s

The gating weights for this case are shown in Fig. 16. The HME was able to identify that the RSO was spinning and assigned almost identical weights to the two banks that modeled RSOs with spin rates of 0.2 rad/s less and more than the spin rate of the truth model. It is not expected that the HME would assign exactly equal weights due to the exponential softmax function used within the HME algorithm to calculate the gating weights.

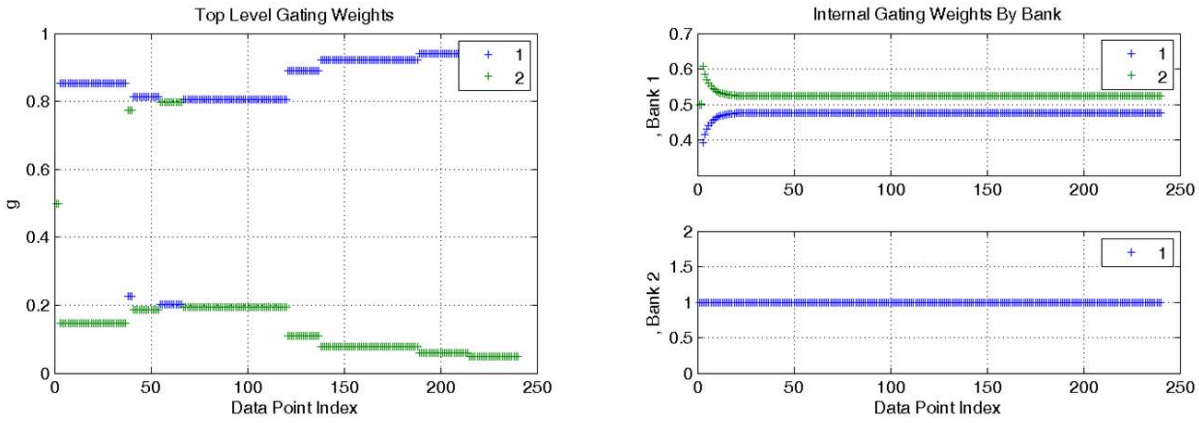

Fig. 16 Gating Weights for the Unmodeled True Spin Rate Experiment

- **Nadir-Pointing vs Sun-Pointing RSO Detection Experiment**

The purpose of this experiment was to determine if the HME could discriminate whether an RSO was nadir or sun pointing. The configuration of the HME is given in Table 9. The truth model for this experiment was a nadir pointing hexagonal prism with 2m x 4m sides. The learning parameter for both layers of the HME is 1.75.

Table 9 Configuration HME to detect nadir-pointing and sun-pointing RSOs

Bank	Attitude Profile	Expert	Spin Rate
1	Nadir-Pointing	1	N/A
2	Sun-Pointing	1	N/A

Fig. 17 shows the top and filter level gating weights for the attitude detection HME. For a nadir-pointing, hexagonal prism RSO, the top level (top) and filter level (bottom) gating weights indicate that the HME selected the correct nadir-pointing bank most of the time. In both cases, an observation window of 1 hour was used, with measurements every 15 seconds. For most of the second experiment, the highest weight is given to the bank modeling nadir pointing RSOs. However, the highest weight is sometimes assigned to the major-axis spin bank in error. While the highest weight is incorrectly given to the major axis spin bank, the weights assigned to filters in that bank are temporarily reassigned.

The incorrect bank selection seen in this experiment is likely due to the radial symmetry about the major axis in the hexagonal prism model. In similar experiments with a nadir-pointing 2m x 1m x 3m cuboid RSO and a nadir pointing 2m x 2m x 2m cube RSO, the HME was more likely to detect that the RSO was nadir-pointing for the cuboid, and less likely to detect that the cube was nadir-pointing. This supports the idea that a high degree of symmetry negatively affects attitude detection. Top level gating weights for a nadir-pointing 2m x 1m x 3m cuboid and a

nadir pointing 2m x 2m x 2m cube indicate that the HME performs better for the less symmetric cuboid than for the cube. In both cases, an observation window of 1 hour was used, with measurements every 15 seconds.

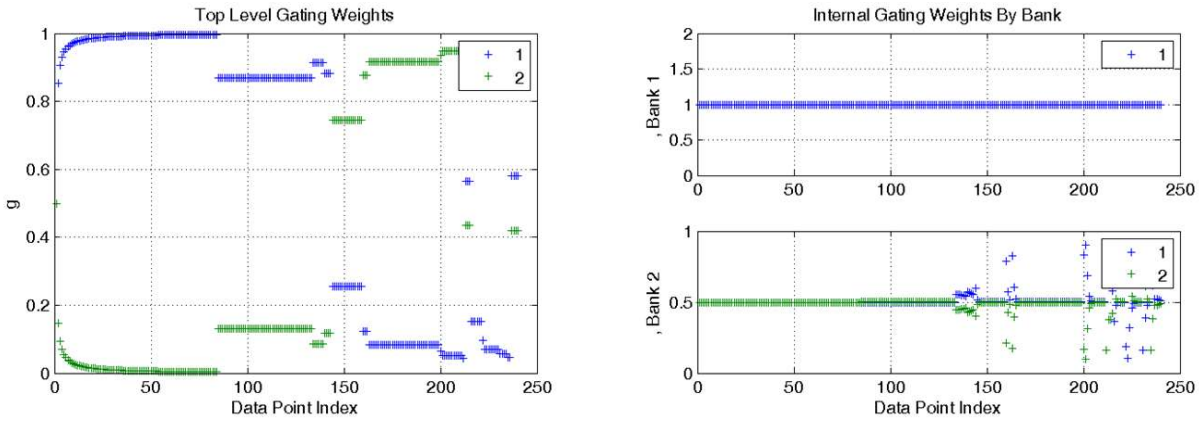

Fig. 17 Gating Weights for the Nadir-Pointing vs Sun-Pointing RSO Detection Experiment

- **Nadir-Pointing to Spinning Attitude Profile Change Detection Experiment**

In addition to investigating whether an HME could identify the initial attitude profile of an RSO, we also investigated whether an HME could detect a change in the attitude profile of an RSO. In one experiment, the RSO being tracked started out with a nadir-pointing attitude profile and then switched to a spinning attitude profile with a spin rate of 0.3 rad/s about the major axis halfway through the simulation. The "truth" RSO model used for this experiment was a 3m x 2m x 1m cuboid. The gating weights for this experiment are shown in Fig. 19 and the configuration of the HME is given in Table 10. For this experiment, the RSO was observed for 1 hour, with measurements available every 3 seconds.

Table 10 HME Configuration to detect attitude profile change

Bank	Attitude Profile	Expert	Spin Rate
1	Sun-Pointing	1	N/A
2	Nadir-Pointing	1	N/A
3	Major Axis Spin	1	(0, 0, 0.3) rad/s

3	Major Axis Spin	1 2	(0, 0, 0.3) rad/s (0,0,0.5) rad/s

As seen in Fig. 18, the HME quickly assigned the highest weight to Bank 2, which was modeling a nadir-pointing attitude profile. After the halfway point of the simulation, when the attitude profile of the RSO is changed to a major-axis spinning profile, with a spin rate of 0.3 rad/s, the highest weight is usually assigned to Bank 3, modeling spinning RSOs. Thus, the HME was able to not only detect the initial attitude profile of an RSO, but was able to detect when a change in the attitude profile occurred and correctly identify the new attitude profile of the RSO.

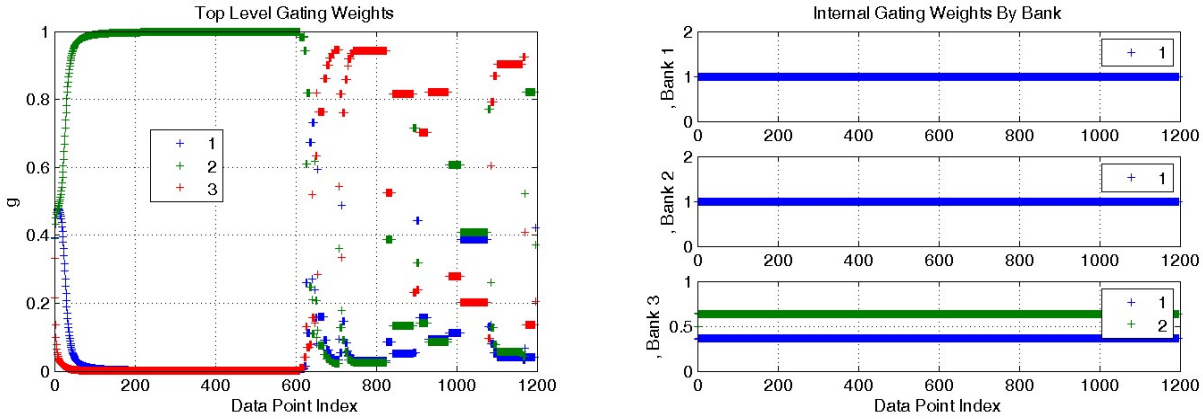

Fig. 18 Gating Weights for the Nadir to Spinning Attitude Profile Experiment

- **Spin Rate Change Detection Experiment**

In another experiment, the RSO being tracked had a spinning attitude profile throughout the entire experiment, but halfway through the observation window, the spin rate about the major axis was changed from 0.1 rad/s to 0.7 rad/s. The "truth" RSO model used for this experiment was a hexagonal prism with 2m x 4m rectangular sides. For this experiment, the RSO was observed for 1 hour, with measurements available every 15 seconds. For the spin rate change experiment, the HME configuration is given in Table 11 and the top level and filter level gating

weights are shown in Fig. 19.

Table 11 HME Configuration to detect a change in spin rate

Bank	Attitude Profile	Expert	Spin Rate
1	Major Axis Spin	1	(0,0,0.1) rad/s
		2	(0,0,0.5) rad/s
		3	(0, 0, 0.7) rad/s
2	Inertially Pointed	1	(0, 0, 0) rad/s

Throughout most of the experiment, the highest top level gating weight is assigned to Bank 1, modeling major-axis spinning RSOs as seen in Fig. 19. The other bank in the experiment, Bank 2, was modeling an inertially-pointed RSO. The highest filter level weight is initially assigned to Filter 1, modeling an RSO with a 0.1 rad/s spin rate. After the change in the "truth" RSO's spin rate, the highest filter weight is then assigned to Filter 3, modeling an RSO with a 0.7 rad/s spin rate. Thus, the HME successfully detected a change in spin rate and determined the correct spin rate value.

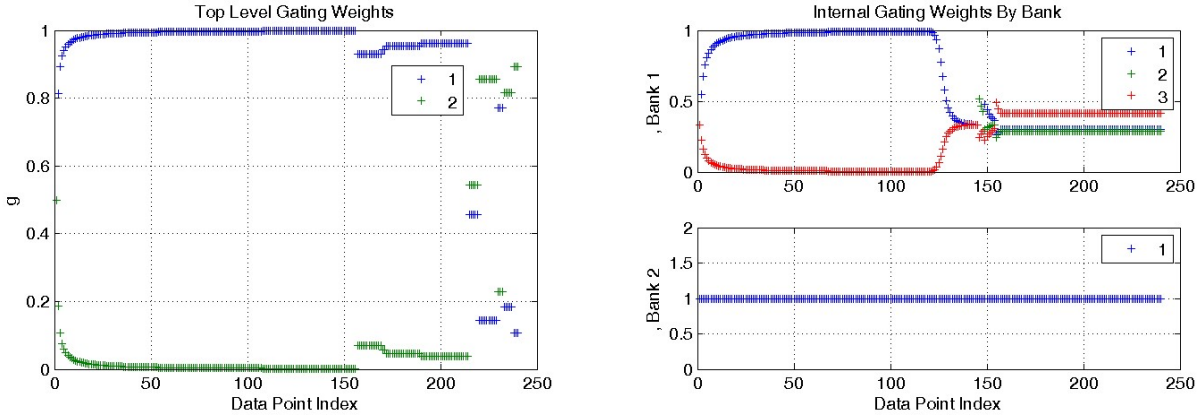

Fig. 19 Gating Weights for Spin Rate Change Experiment

- **Summary and Conclusions**

We developed a new algorithm for RSO feature identification and characterization based on the HME and tested it using simulated astrometric and photometric data. Our simulations have shown the HME is

capable of determining and detecting changes in the size, shape, reflectivity, attitude profile, and maneuvers performed by an RSO. The HME was also shown to be capable of identifying and distinguishing between nadir-pointing, sun pointing, and spinning objects even though none of the experts in the HME is directly estimating attitude. The results of our testing are summarized in Tables 12 and 13. The HME performed well in these simulations, demonstrating the feasibility of the proposed approach.

In addition, over a 1 hour observation window, with measurements every 15 seconds, the HME is able to assign the highest weight to the correct bank for a spinning or a sun-pointing object for 100% of time steps, and is usually able to distinguish a nadir-pointing object. These results show the HME is capable of distinguishing between three different RSO attitude profiles: nadir-pointing, sun-pointing, and major-axis spinning at a constant rate. When the truth RSO model is spinning, the HME was able to recognize that the RSO was not nadir-pointing or sun-pointing virtually all of the time. It was also able to identify that a sun-pointing object was not spinning or nadir-pointing essentially all of the time. For a nadir-pointing RSO, the HME is usually able to recognize that the object is not spinning or sun-pointing, although the results are not as clear cut as for the spinning or sun-pointing RSO. Identification of a nadir-pointing object is even more difficult when the object of interest displays some degree of symmetry (such as a cube or an axially symmetric hexagonal prism); although the HME is still usually able to distinguish that the RSO is nadir-pointing most of the time. The percentages given in Table 13 are the percentage of time steps (out of approximately 250) where the correct bank (that is, the bank with an attitude profile model matching the truth RSO model) was assigned the highest weight by the HME.

Table 12 Summary of Feature Identification Results

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.

Table 13 Summary of Attitude Profile Identification Results

RSO Attitude Profile	RSO Model		
	2m x 1m x 3m Cuboid	**Hexagonal Prism with 2m x 4m sides**	**2m x 2m x 2m Cube**
Spinning at ω = (0, 0, 0.3) rad/s	100% detection	100% detection	100% detection
Nadir-Pointing	85.4% detection	62.5% detection	57.9% detection
Sun-Pointing	100% detection	100% detection	100% detection

In addition, the HME is capable of detecting when changes occur in the attitude profile of an RSO. Our HME correctly identified when a nadir-pointing RSO began spinning about its major axis. Also, the HME was able to detect the correct spin rate of a major-axis spinning RSO both prior to and after a change in the spin rate. We have shown that RSO attitude profile detection is possible despite the fact that the experts, all EKFs, in the HME only estimate position and velocity and do not estimate attitude directly. All information on RSO attitude obtained by the HME comes from the measurements processed in the experts, through the measurement residuals. While the HME does not estimate the attitude state of an RSO, the weights assigned to the various banks and experts in the HME indicate which attitude profile best matches the behavior of the observed RSO. While others have been able to determine attitude from brightness data, as in [5], it is significant that we could do so without running attitude estimation filters, which would require a higher processing rate than estimating translational motion alone.

- **Acknowledgments**

This work was funded by three Phase I SBIRs (FA9451-10-M-0087, FA9451-12-M-0312, FA9453-13-M-0169) from the Air Force Research Laboratory.

• References

[1] Payne, T. P., and Morris, R. F., "The Space Surveillance Network (SSN) and Orbital Debris," *Advances in the Astronautical Sciences*, Vol. 137, No. 4, 2010.

[2] Montenbruck, Oliver, and Eberhard Gill. Satellite orbits: models, methods and applications. Springer Science & Business Media, 2012.

[3] Wetterer, C. J., Crassidis, J. L., Linares, R., Chow, C. C., and Jah, M. K., "Simultaneous Position, Velocity, Attitude, Angular Rates, and Surface Parameter Estimation Using Astrometric and Photometric Observations," Proceedings of the FUSION Conference, Istanbul, Turkey, July 2013.

[4] Magill, D., "Optimal Adaptive Estimation of Sampled Stochastic Processes," *IEEE Transactions on Automatic Control*, Vol. 10, Oct. 1965, pp. 434–439.
doi: 10.1109/TAC.1965.1098191

[5] Wetterer, C. J., and Jah, M. K., "Attitude Determination from Light Curves," *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 5, 2009, pp. 1648-1651.
doi: 10.2514/1.44254

[6] Linares, R., Jah, M. K., Crassidis, J. L., and Nebelecky, C. K., "Space Object Shape Characterization and Tracking Using Light Curve and Angles Data." *Journal of Guidance, Control, and Dynamics, Vol* 37, No. 1, 2013, pp. 13-25.
doi: 10.2514/1.62986

[7] Crain, T., "Adaptive Interplanetary Orbit Determination," Ph.D. Dissertation, Aerospace Engineering and Engineering Mechanics Department, University of Texas at Austin, Aug 2000.

[8] Brown, R. G. and Hwang, P. Y. C., *Introduction to Random Signals and Applied Kalman Filtering*, John Wiley & Sons, New York, NY, 3rd ed., 1997, pp. 353–361.

[9] Chaer, W., Bishop, R., and Ghosh, J., "A Mixture of Experts Framework for Adaptive Kalman Filtering," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 27, No. 3, June 1997.
doi: 10.1109/3477.584952

[10]Bishop, M. P., Bush, A. B., Furfaro, R., Gillespie, A. R., Hall, D. K., Haritashya, U. K., & Shroder Jr, J. F., "Theoretical Foundations of Remote Sensing for Glacier Assessment and Mapping," *Global Land Ice Measurements from Space*, Springer Berlin Heidelberg, pp. 23-52.
doi: 10.1007/978-3-540-79818-7_2

[11] Ashikhmin, M. and Shirley, P., “An Anisotropic Phong BRDF model,” *Journal of Graphics Tools*, Vol. 5, No. 2, 2000, pp. 25-32.
doi: 10.1080/10867651.2000.10487522

[12] Ashikhmin, M. and Premoze, S., “Distribution-based BRDFs,” Unpublished Technical Report, University of Utah, 2007.