

Article

Anomaly Detection in Autonomous Deep-Space Navigation via Filter Bank Gating Networks

Daniel P. Lubey *  and Todd A. Ely 

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

* Correspondence: daniel.p.lubey@jpl.nasa.gov

Abstract: This study investigates methods for autonomous navigation of a deep-space spacecraft where one-way radiometric and on-board optical information are fused to create a fully informed state estimate. The specific focus is on using filter bank methods (i.e., Multiple Model Estimation [MME] and Mixture of Experts [MoE]) to detect when measurement and/or dynamical mis-modeling occurs. We develop a new χ^2 -based gating network for a filter bank that may be used to identify poorly performing filters (i.e., those with low weights), which may be used as a signal for mis-modeling in the system. In addition to defining and deriving this new weighting scheme, numerical simulations based on NASA's InSight mission demonstrate this new algorithm's performance with and without measurement and dynamical mis-modeling present.

Keywords: spacecraft navigation; anomaly detection; Multiple Model Estimation; filter bank; autonomous navigation



Citation: Lubey, D.P.; Ely, T.A. Anomaly Detection in Autonomous Deep-Space Navigation via Filter Bank Gating Networks. *Appl. Sci.* **2022**, *12*, 11161. <https://doi.org/10.3390/app12211161>

Academic Editors: Dah-Jing Jwo and Anna Guerra

Received: 20 September 2022

Accepted: 24 October 2022

Published: 3 November 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 California Institute of Technology. Government sponsorship acknowledged. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Deep-space spacecraft navigation, as it currently stands, is a fairly manual process that involves a team of navigators that process measurements from a dispersed set of sensors with optimal state estimation algorithms, while this method has been shown to be effective for many missions, there is a desire to explore autonomous solutions to the problem of deep-space navigation. Autonomous algorithms would move the measurement processing to be chiefly on-board the satellite. This would reduce the needed navigation resources on the ground while also allowing for mission scenarios that require near real-time estimation—an impossibility for ground-based navigation for missions with large light travel times between the Earth and the satellite. Autonomy also facilitates flying smaller payloads, since it would reduce the investment needed to operate these types of missions in the long-term.

While there are many perceived benefits from autonomous navigation algorithms, there are also many challenges in developing one for operational usage. A prominent challenge is the limited computational resources that are available. On the ground, access to computational resources for processing measurements is effectively limitless, but on-board the spacecraft resources are finite and possibly quite limited. As a result, autonomous algorithms need to be adapted to the flight hardware they are flying on while also dealing with power and data constraints. An autonomous algorithm would also need to be robust to issues that are typically identified and addressed by navigators on the ground such as dynamic and/or measurement mis-modeling and measurement outliers. Furthermore, the autonomous algorithms would also need to be able to fuse together different types of observables that can occasionally conflict (e.g., optical and radiometric data) while also identifying when conflicts occur and what the source of the conflict is. This is especially important as access to one-way radiometric data to support autonomous deep space navigation is becoming feasible via the Deep Space Atomic Clock (DSAC) [1].

This study will focus on the development of a neural network-based anomaly detection algorithm for use in radiometric and optical fused autonomous navigation of a spacecraft

operated in deep-space with a specific focus on cruise and approach mission phases. This work can be viewed as an extension of Ely et al. [2], who quantified the performance of radiometric and optical fused autonomous navigation. Our approach to this problem utilizes filter bank methods (e.g., Multiple Model Estimate [MME] and Mixture of Experts [MoE])—collections of estimators (i.e., filters) that are each modeled slightly differently, and each estimator processes the available information to produce an independent state estimate. On top of the filter bank is a gating network (i.e., weighting scheme) that assigns a time-variable weight to each filter, which indicates how important that filter is to the overall solution. This weighting may also be viewed as an indicator of how well the filters are performing, so it may be used as an indicator of mis-modeling in the system. In this paper, we develop a new optimal gating network, which focuses on supporting radiometric and optical fused estimation for autonomously navigated spacecraft. These weights are used to both identify the presence of mis-modeling in the system as well as characterize the root cause of the mis-modeling—a graphical example of this process is shown in Figure 1.

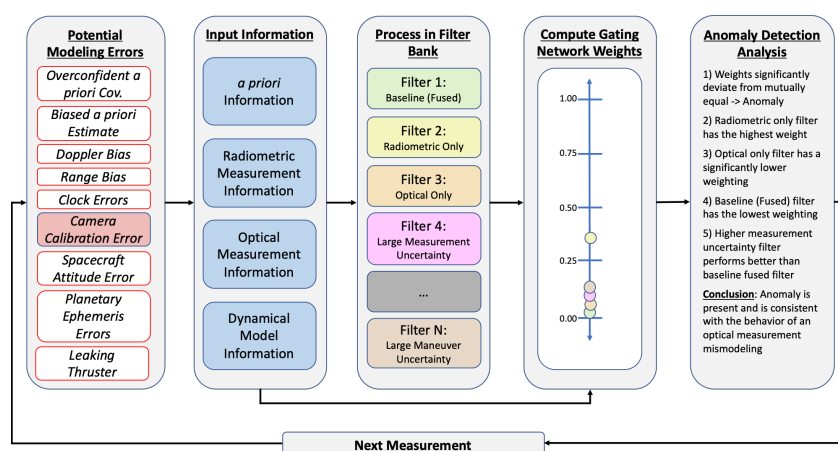


Figure 1. An overview of the anomaly detection and characterization process via filter bank gating network weights with example errors and filter setups.

The general problem we are addressing in this study (anomaly detection and characterization) has many different approaches beyond the neural-network based method we are pursuing in this study. Many approaches deal specifically with the problem of maneuver detection and characterization—i.e., the identification and reconstruction of a dynamic anomaly while tracking a target. A number of classical methods (e.g., dynamic model compensation, polynomial acceleration model estimation, etc.) [3–5] and input estimation techniques [6–8] rely on detecting dynamic anomalies via statistical tests on measurement residuals, and then characterize the anomaly by estimating dynamic parameters associated with the anomaly. System identification methods [9–11] are commonly applied to the characterization of mis-modeled systems. These approaches tend to focus on reconstructing a system by measuring how a system responds to commanded inputs, though some methods [12,13] do not require direct knowledge of the inputs. However, these methods are not always generally applicable especially for deep-space spacecraft navigation problems where there are significant constraints on tracking data and input resources (e.g., power and fuel). A number of spacecraft focused approaches exist [14–17], though they tend to focus on on specific applications which are not related to deep-space navigation. Control distance metrics and related optimal-control-based filtering techniques [18–25] can be used to both detect anomalies and reconstruct them as optimal control policies; however, they tend to focus on dynamic anomalies alone, while there are many different approaches to the more general anomaly detection and characterization problem, we are looking for an approach that accomplishes the following: (1) it is applicable to deep-space navigation problems, (2) it could feasibly be used in an autonomous navigation system, (3) it can both detect an anomaly and be used to characterize and diagnose the cause of the anomaly, (4) it

can deal with many different types of anomalies that impact navigation solutions (e.g., mis-modeled dynamics, measurement biasing, issues with sensor fusion, over-confident uncertainty metrics, etc.), (5) and it can be adapted as necessary to best fit the spacecraft's specific mission criteria. Given these requirements, we pursued a filter bank approach for this research given these methods are naturally adaptable and flexible, so they provide a good foundation from which to address all of our different requirements for this algorithm.

Application of the filter bank methods (e.g., MME and MoE) to spacecraft navigation has been covered to some extent in the existing literature. Various studies that use MME-based methods tend to focus on making algorithms that are robust to model errors [26–30]. Similarly, MoE-based studies tend to focus on developing methods that are robust to anomalies in the system. This includes the work of Chaer, Bishop, and Ghosh [31], who provided a comparison of MoE and MME-based methods and how rapidly each could identify the correct mode of the observed system via gating network weights. Furthermore, the MME and MoE approaches have been explicitly applied to the problems of anomaly detection and characterization [32,33], while the existing literature covers a wide range of applications, they ignore an important component of our presented problem—dealing with the fused-sensor problem. Specifically, we seek a filter bank-based method that supports the inclusion of filters in the bank that only process a subset of the available data types (e.g., radiometric-only, optical-only, etc.). Beyond this we seek a method that can identify and characterize when anomalies occur, including when different data types conflict with each other. This paper will focus on developing an algorithm that meets this description.

Section 2 provides an overview of filter bank methods and their relation to autonomous navigation of deep-space spacecraft. In Section 3, we define and derive a new weighting scheme to use with filter banks, which is designed for anomaly detection and characterization. Section 4 presents the results and analysis from numerical simulations based on the InSight spacecraft Mars approach where different weighting schemes for filter banks are used to identify mis-modelings in the system. Finally, Section 5 provides a summary of this research and some discussion of possible avenues for future work.

2. Related Work

Filter banks provide a platform to test many different filter models on the same input data. The filter models can differ from each other in a number of different ways, including: measurement model parameters, fidelity of the measurement models, measurement weightings (i.e., uncertainties), the set of measurements that are processed by the filter, dynamical model parameters, fidelity of the dynamical model, maneuver models, a priori state uncertainties, etc. By passing this data through different models, we can better understand the nature of the underlying true model by determining which filters in the bank perform best in a given scenario.

The metric for measuring which filters in the bank are the most important is its weighting, which is obtained through a gating network or weighting schemes. Based on each filter's performance (or the properties of the input), it is assigned a weighting between 0 and 1, such that the sum of all weightings is equal to 1. Typically, these weightings are used to weight the state estimate from each filter in a weighted summation that produces a final optimal state estimate that is a linear combination from each filter's output. Alternatively, the weightings can be used as signals for anomalies and mis-modeling in the system. For instance, if a sudden bias appeared in the measurements obtained from a certain camera, then we would expect the weighting for filters that process those measurements to decrease—an indication of anomalous behavior.

How effectively weights respond to anomalies is a function of several things, including: the type of gating network that is used, the filter's sensitivity to the anomaly, and the size of the anomaly. The type of gating network that is used can often differ based on the type of filter bank that is being employed. The two types that have been focused on in the existing literature are the MoE and MME methods.

The MoE method [34] is a method commonly used within the field of neural networks, which are often leveraged in applications like machine learning. The MoE method leverages the idea of divide and conquer, which attempts to break up complicated problems into many smaller and simpler pieces whose individual results can be combined to solve the larger problem. Essentially, the bank is trained such that each filter within the bank becomes an expert that focuses on a specific piece of input (i.e., measurement) state space. When an input comes into that filter's area of expertise, it would be ranked highly by the gating method, thus dominating the final bank's state estimate. It's important to note two traits of MoEs at this point: (1) the gating network values are entirely based on the input values and parameters that are trained ahead of time (via methods like classification and regression tree [CART] with the expectation-maximization [EM] approach [34]); and (2) as a result, the MoE gating network weights do not reflect how the filters are performing, just how much of an expert the filters are for the given input. We are looking for an algorithm that tracks how well filters are performing in real-time so we want to make the gating network focus on filter performance in addition to the input values. This adjustment pushes the algorithm more into the realm of Multiple Model Estimation, but we still utilize concepts from the area of MoEs in this research.

MME [35,36] is a form of optimal state estimation where instead of operating a single filter to process measurements, a bank of filters is used in which each filter has a different model. The estimators can differ in the dynamical model, the measurement models, the process noise parameters, the measurements' uncertainty parameters, the estimation algorithm, etc. Each estimator is given the same measurement vector, they each process it with their unique setup, and then a weight is computed for each estimator to reflect how likely each estimator is to reflect the true model of the system. Typically, the outputs from each estimator are scaled by their weight, and then they are all linearly combined to form a final state estimate (Figure 2). Alternatively, the weights can be used to select the estimator that is indicated to be the most likely model (i.e., highest weight), and then its solution will be used as the optimal MME solution. Care must be used in both situations since the size of the estimate vector may vary depending on the estimator's setup. However, if the user is only interested in certain parameters for prediction purposes (e.g., target state, dynamics parameters, etc.), then these may be extracted to form the final estimate vector.

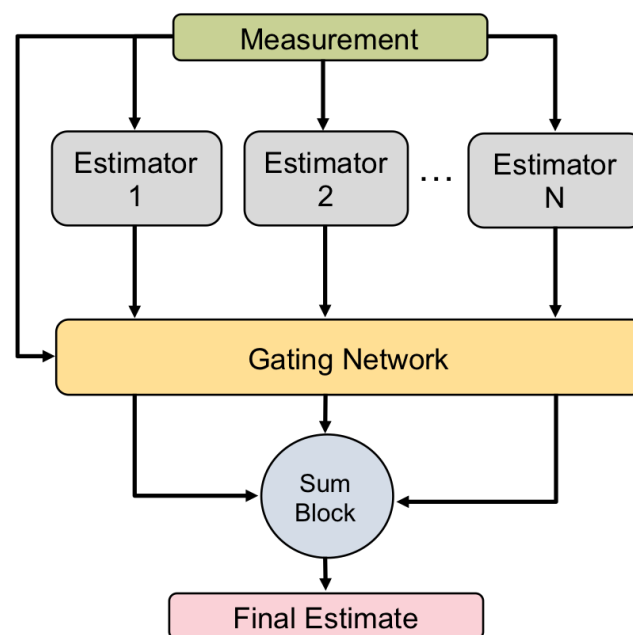


Figure 2. Multiple Model Estimation (MME) filter bank with gating network.

As defined, MME is a method that is very similar to the idea of MoE; however, MME is more focused on estimation problems (whereas MoE is more generally applicable) and MME gating networks can (and usually are) functions of the filters' outputs. This difference is significant because it means an MME does not need to be trained in the same way that an MoE does—though it is still possible to design a gating network that would be trained. A standard MME gating network relies more on statistical theory and the modeling put into the estimators as opposed to the MoE, which can be trained to act as desired. This may tend to mean that an MME is less flexible than an MoE (depending on how many experts are used), but it may be naturally more robust. To achieve robustness with an MoE, the operator needs to understand all of the possible failure scenarios, in order to make the proper adjustments and train the MoE accordingly.

As mentioned, the weights from the standard gating network in the MME algorithm tend to measure the likelihood that their corresponding filter is the correct model. The original metric used by Magill [35] is based on Bayes' rule. Given a set of N estimators in an MME ($\{\alpha_1, \alpha_2, \dots, \alpha_N\}$), Magill's weight quantifies the likelihood that any given model is the correct one. We compute this as shown in Equation (1).

$$p_{k,i} = f(\alpha = \alpha_i | \mathcal{Y}_1^k) = \frac{f(\mathcal{Y} = \mathcal{Y}_1^k | \alpha_i) f(\alpha = \alpha_i)}{\sum_{j=1}^N f(\mathcal{Y} = \mathcal{Y}_1^k | \alpha_j) f(\alpha = \alpha_j)} \tag{1}$$

$$\mathcal{Y}_i^j = \{\vec{y}_i, \vec{y}_{i+1}, \dots, \vec{y}_{j-1}, \vec{y}_j\} \tag{2}$$

In this equation, \mathcal{Y}_i^j is the set of all measurements between indices i and j (Equation (2)), $f(\mathcal{Y} | \alpha_i)$ is the conditional probability density function expressing the likelihood of receiving the measurement set \mathcal{Y} given estimator α_i , and $f(\alpha)$ is the prior distribution describing the likelihood of a given estimator being the true model. It can be shown that $\sum_{i=1}^N p_{k,i} = +1$, and it is required that $\sum_{i=1}^N f(\alpha = \alpha_i) = +1$. Generally, we assume that the measurements contain Gaussian error that is zero mean with known covariance (R_i being the covariance for measurement \vec{y}_i). We also assume that the MME estimators have a uniform prior distribution ($f(\alpha = \alpha_i) = 1/N, i = 1, 2, \dots, N$).

If we further assume that measurements are statistically independent in time (or at least over smaller mini-batches), then we can rewrite this weighting scheme in a more sequential notation as shown in Equation (3).

$$p_{k,i} = f(\alpha = \alpha_i | \mathcal{Y}) = \frac{f(\vec{y} = \vec{y}_k | \alpha_i) f(\alpha = \alpha_i | \mathcal{Y}_1^{k-1})}{\sum_{j=1}^N f(\vec{y} = \vec{y}_k | \alpha_j) f(\alpha = \alpha_j | \mathcal{Y}_1^{k-1})} \tag{3}$$

In this expression, $\mathcal{Y}_j^k = \{\vec{y}_j, \vec{y}_{j+1}, \dots, \vec{y}_{k-1}, \vec{y}_k\}$ with $j \leq k$. If $j > k$, then we use our prior model instead ($f(\alpha = \alpha_i | \mathcal{Y}_j^k) = f(\alpha = \alpha_i)$, if $j > k$). This mode of weighting allows us to sequentially compute our weights at each evaluation point while properly accounting for previous and current measurement information.

Additionally, if we assume that the inputs are Gaussian, then we can write explicit equations for computing the MME weights. The multivariate conditional Gaussian probability density function for measurement $\vec{y}_k \in \mathbb{R}^p$ as a function of estimator α_i with estimator-dependent modeled covariance matrix $R_{i,k} \in \mathbb{R}^{p \times p}$ is given in Equation (4).

$$f(\vec{y} = \vec{y}_k | \alpha_i) = \frac{1}{\sqrt{(2\pi)^p |S_k^{(i)}|}} \exp \left[-\frac{1}{2} (\vec{y}_k - h_{i,k}(t_k, \bar{x}_{i,k}))^T S_{i,k}^{-1} (\vec{y}_k - h_{i,k}(t_k, \bar{x}_{i,k})) \right] \tag{4}$$

$$S_{i,k} = R_{i,k} + \tilde{H}_{i,k} \bar{P}_{i,k} \tilde{H}_{i,k}^T$$

In this definition, $\bar{x}_{i,k}$ and $\bar{P}_{i,k}$ are the a priori state estimate vector and covariance matrix (respectively) for estimator α_i at the k th measurement time (t_k); $h_{i,k}(t, \vec{x})$ is the

function that computes the expected measurement for the k th measurement given the state of the i th estimator; and $\tilde{H}_{i,k}$ is the state vector-derivative of the k th expected measurement function for the i th estimator.

The Magill gating network, as described, outputs weights that favor filters that produce measurement residuals that are small relative to the innovations covariance matrix ($S_{i,k}$) and that have a small innovations covariance matrix. Because it balances two costs, the estimator that it selects as best, does not necessarily have to be behaving nominally—it may just have a very small innovations uncertainty relative to other estimators in the filter bank. Furthermore, the way these weights are defined does not easily lend itself to supporting filter banks that have filters that only process a subset of the available measurement data types. These types of filters will likely have high innovations covariances for the data types that they do not process, which will lead to the estimator’s weight dropping even if the filter is behaving nominally. For these reasons, we are motivated to design a new gating network that will focus on whether the filters in the filter bank are demonstrating nominal performance independent of the data that they are processing. The development of this new gating network is detailed in Section 3.

3. Materials and Methods

In this section, we develop a method to compute optimal weightings for the filters within our filter bank. We start by defining the cost function that defines this optimal weighting scheme and then derive the analytical solution that minimizes this cost function. Next, we discuss the practical application of this weighting scheme including its strengths and weaknesses, and how the method may be adjusted to improve its performance.

3.1. Cost Function Definition and Solution Derivation

The cost function that defines this optimal filter weighting scheme is defined in Equation (5). This equation uses the following notation: N is the number of filters in the filter bank, \vec{p}_k is the vector of N filter weightings at the k th weightings update epoch, M is the number of measurements included in each weightings update batch, \bar{W} is a symmetric positive definite weighting matrix that controls how strictly the previous filter weights are enforced, λ is a Lagrange multiplier that enforces the constraint that the filter weightings sum to unity, and $\vec{1}_N$ is an N -dimensional vector where all entries are one.

$$\begin{aligned} \mathcal{J}(\vec{p}_k, \lambda; M) &= \left(\sum_{i=1}^N \sum_{j=(k-1)M+1}^{kM} \frac{1}{2} p_{k,i}^2 \bar{z}_{i,j}^T \left[R_{i,j} + \tilde{H}_{i,j} \bar{P}_{i,j} \tilde{H}_{i,j}^T \right]^{-1} \bar{z}_{i,j} \right) \quad (5) \\ &+ \frac{1}{2} (\vec{p}_k - \vec{p}_{k-1})^T \bar{W}^{-1} (\vec{p}_k - \vec{p}_{k-1}) + \lambda \left(\vec{1}_N^T \vec{p}_k - 1 \right) \\ \bar{z}_{i,j} &= \vec{y}_j - h_{i,j}(t_j, \vec{x}_{i,j}) \\ \vec{p}_k &= [p_{k,1} \quad p_{k,2} \quad \cdots \quad p_{k,N}]^T \\ \tilde{H}_{i,j} &= \left. \frac{dh_{i,j}(t, \vec{x})}{d\vec{x}} \right|_{(t_j, \vec{x}_{i,j})} \\ \vec{p}_0 &= \left(\frac{1}{N} \right) \vec{1}_N \end{aligned}$$

This first term in this cost function is designed to penalize filters that generate anomalous predicted measurement residuals using a χ^2 statistical metric. The second term ensures memory is retained within this gating network so that new weights do not differ from a priori weights by too much (where “too much” is quantified by \bar{W}). Finally, the third term in this cost function is a constraint that ensures that the weights from each filter in the bank sum to unity. It is worth noting, that this constraint does not preclude weights from becoming negative; however, the cost function’s design should prevent this from occurring. For instance, the first term uses the square of the updated weights on the χ^2 met-

rics, so making the weights negative is not beneficial. In fact, this would increase the cost, because any negative weight implies the other weights must be larger (to ensure that the constraint is met), which would increase the weighted- χ^2 term. Furthermore, the memory term disadvantages negative weights, as long as the a priori weights are positive. This cost function may be written in tensor form as shown in Equation (6).

$$\mathcal{J}(\vec{p}_k, \lambda; M) = \frac{1}{2} \vec{p}_k^T G_k \vec{p}_k + \frac{1}{2} (\vec{p}_k - \vec{p}_{k-1})^T \bar{W}^{-1} (\vec{p}_k - \vec{p}_{k-1}) + \lambda (\vec{1}_N^T \vec{p}_k - 1) \quad (6)$$

$$G_k = \begin{bmatrix} G_{k,1} & 0 & \cdots & 0 \\ 0 & G_{k,2} & \cdots & 0 \\ \vdots & \vdots & \cdots & \vdots \\ 0 & 0 & \cdots & G_{k,N} \end{bmatrix}$$

$$G_{k,i} = \sum_{j=(k-1)M+1}^{kM} G_{k,i,j}$$

$$G_{k,i,j} = \bar{z}_{i,j}^T [R_{i,j} + \tilde{H}_{i,j} \bar{P}_{i,j} \tilde{H}_{i,j}^T]^{-1} \bar{z}_{i,j}$$

The solution that minimizes this cost function must satisfy the necessary conditions for optimality as defined in Equation (7).

$$\vec{\mathcal{J}}_p = \left. \frac{d\mathcal{J}(\vec{p}, \lambda; M)}{d\vec{p}} \right|_{(\vec{p}_k, \lambda; M)} = G_k \vec{p}_k + \bar{W}^{-1} (\vec{p}_k - \vec{p}_{k-1}) + \lambda \vec{1}_N = \vec{0} \quad (7)$$

$$\mathcal{J}_\lambda = \left. \frac{d\mathcal{J}(\vec{p}, \lambda; M)}{d\lambda} \right|_{(\vec{p}_k, \lambda; M)} = \vec{1}_N^T \vec{p}_k - 1 = 0$$

Simplifying these conditions results in the solution that is defined in Equation (8).

$$\vec{p}_k = (G_k + \bar{W}^{-1})^{-1} [\bar{W}^{-1} \vec{p}_{k-1} - \lambda \vec{1}_N] \quad (8)$$

$$\lambda = \frac{\vec{1}_N^T (G_k + \bar{W}^{-1})^{-1} \bar{W}^{-1} \vec{p}_{k-1} - 1}{\vec{1}_N^T (G_k + \bar{W}^{-1})^{-1} \vec{1}_N}$$

Having obtained an analytical solution that minimizes our designed cost function, we will next focus on this algorithm’s strengths and weaknesses, how it may be slightly modified to improve its performance in detecting filter mis-modelings, and how it may be applied practically.

3.2. Discussion

As described, the cost function that defines this optimal gating network was designed to prefer filters in the bank that were generating predicted measurement residuals that were not anomalously large relative to the uncertainty associated with the measurements and state estimates. Among its strengths, are its abilities to analytically generate normalized weightings for a filter bank with an arbitrary number of filters, automatically identify and de-weight filters that are producing anomalous results, maintain a memory of how well each filter in the bank is performing, and be tuned to adjust how strong its memory is.

Though this method has its strengths, it has two notable limitations, which include:

- **Dealing with Biased Measurements:** When measurements are biased, a filter that processes them will generally produce solutions that drift away from truth. This may or may not be detectable in measurement-residual space. However, the measurement residuals from a properly functioning filter that does not process the biased measurements will show the bias in its χ^2 residual metrics. Whether it can be detected or not is based on how large the bias is relative to the filter innovations uncertainty. Essentially

this means that properly functioning filters are penalized by the presence of biased measurements even when this filter does not process the biased measurements.

- Preference for Over-Inflated Uncertainty Metrics: When state or measurement uncertainty is inflated (i.e., made to appear more uncertain than the available information implies), the resulting χ^2 metrics will be smaller than they should be in a statistical sense, which will make the filter seem like it is performing better than it really is. This weighting scheme only penalizes residuals when they are too big, not when they are abnormally small.

To ensure that biased measurements do not disadvantage a filter that does not even process them, we can modify the optimal weighting scheme by using a mean χ^2 value for all residual metrics when the associated measurement was not actually processed by the filter in question. This is accomplished by setting $G_{k,i,j} = n_j$ (i.e., the χ^2 distribution mean) where n_j is the dimension of the j th measurement vector and i is in the set Ω_j , which defines all filter indices that do not process the j th measurement vector. All non-processing filters will be represented as performing nominally, while the remaining filters will be penalized if the biased measurements degrade their performance.

The second issue, a preference for over-inflated uncertainty metrics, does not have an obvious solution within the currently defined algorithm. The algorithm's χ^2 form could be adjusted to penalize solutions that produce residuals that deviate from a mean or median estimator, but this would require significant adjustments to the algorithm. We leave this as an avenue for future development of this algorithm. For now, the user would need to manually check that the filters are not producing residuals that are statistically too small too often.

In terms of practical application, the authors provide the following notes:

- Gating network a priori weights: When initializing a filter bank with N filters, we generally assume that the filters are equally weighted (i.e., each has an initial weight of $1/N$).
- Gating network batch length: This optimal gating network is designed to work on batches of data. New weights are computed each time a new set of M measurements are processed. These M measurements can be a collection of different types (e.g., radiometric, optical, etc.). The selected batch length can have a large effect of the gating network's performance. A small batch size makes the network more agile, but it also makes the network less robust to false detections. Generally, it is preferred that each batch size contains a good mixture of the available data types, but this is not a requirement. For the simulations in this paper, we use a batch size of 1000 measurements, while we use equal batch sizes in this analysis, there is nothing that requires this to be the case. Batches can be split up by other means if preferred (e.g., split by pre-defined data arcs). The only requirement is that the batches are the same for each filter in the bank—e.g., even if some batches of data in a larger data arc contain zero optical measurements, optical-only and radiometric-only filters would still produce weightings for each batch of data in the arc.
- Gating network memory: This optimal gating network was designed to retain a memory of previous weightings in order to make it reflect filter performance over time rather than at an instance. The \bar{W} parameter controls the strength of the gating network's memory. For the purposes of this analysis we use the following relation: $\bar{W} = \sigma_W^2 I_{N \times N}$. This is not a requirement, \bar{W} is only required to be a symmetric positive definite matrix. A large value for σ_W results in a forgetful gating network, which tends to favor current performance over previous performance. A small value for σ_W results in a gating network with a strong memory, which tends to favor filters that have performed over long spans of the measurement arc. For the simulations in this analysis we use $\sigma_W = 0.05$.
- Gating network reset: Periodically, the user may implement gating network resets, which reset the weights of all the filters to their a priori values. This is normally performed after an anomaly has been detected, characterized, and addressed or any

other time that adjustments are made to any of the filters in the filter bank. For the simulations in this analysis, we do not implement gating network resets.

It is important to note that this gating network algorithm requires no training being an MME-based algorithm; whereas, training is typically required for MoE-based algorithms. There are some parameters that can be tuned by the user (e.g., \bar{W}) to obtain the desired performance for a given application; however, their selection does not require a formal automated training algorithm. Similarly, the filters in the filter bank are assumed to be fully designed by engineers who setup the estimation framework, while the filters could be trained in MoE style, it is ideal to setup the filters in a way that can isolate common errors that could occur up in actual operations. For instance, having filters that process all or a subset of the available measurement types is valuable in identifying and characterizing sensor-specific mismodeling. It is beyond the scope of this article to identify specific filter bank setups that should be used as this is a very problem dependent issue; however, it is left as future work to establish general practices for setting up a filter bank and automating the process for characterizing the source of anomalies based on the population of filters that exist within the filter bank. At present this characterization process still requires an engineer's intuition.

Given this work is meant to progress the capabilities of autonomous (i.e., on-board) navigation, it is also important to address the computational impacts of this algorithm. Being an MME-based algorithm, our algorithm requires measurement processing by various filters as part of a filter bank. The number of filters that are a part of that filter bank will be limited by the computational resources on-board; however the advantage of a MME-based algorithm is that the number of filters in the bank is completely tunable by the user and can be adjusted across the mission timeline based on variable computational capabilities. For deep-space applications, our specific area of interest, there are often long observation gaps that would enable many different filters to be run as the spacecraft awaits new measurement information, so this would facilitate more filters in the filter bank. Our filter bank gating network algorithm, though, does not require significant computations beyond the filtering algorithms because it is largely based on parameters already computed in the filters (e.g., predicted measurement residuals). For ideal operation, it would require some history of how the weights evolve over time, but this can be tunable based on the allocated computational resources. Our simulations in Section 4 were not produced with a flight-like code architecture, so we do not provide computational load metrics as this would not be representative of an actual on-board implementation. This computational performance is likely to be application-specific too, so it is left as future work to quantify this for specific applications of interest.

Our algorithm, as designed, is meant to identify the best-performing (in a statistical sense) filters within a filter bank. These filters will be given the largest weights, while the poorer performing filters will be assigned lower weightings. If all filters are properly performing, then we expect that the weights for each filter will be essentially equal. When this is not the case, we take this as a signal that some anomaly may be occurring within the system. By investigating which filters are performing best and which are performing worst, we can generate an informed hypothesis about what caused the anomaly. Examples of how this gating network may be applied to the problem anomaly detection and characterization in measurement-fused spacecraft navigation are provided in Section 4.

4. Results

In this section, we demonstrate the performance of the Magill-based gating network and our optimal gating network in the presence of measurement biasing and dynamical model biasing through numerical simulation that is based on the InSight mission. Specifically, the simulation focuses on the last 45 days of the trajectory before Mars entry. The parameters that describe this simulation are defined by Ely et al. [2]. We first summarize the filter bank performance when the system is properly modeled, then address how the filter bank performs with no biasing, when dealing with radiometric and optical

measurement biasing, and finally introduce the gating networks to the problem of detecting dynamical model mis-modeling.

4.1. Simulation Description

Our simulations in this section focus on application of our algorithm to a deep-space navigation problem where different observables are fused together to form a fully-informed navigation solution—the primary use-case for which this algorithm was developed. Spacecraft navigation problems have a unique set of qualities relative to other navigation problems—mostly related to things like the data-sparse nature of the tracking data (i.e., long gaps between observation windows are common), the limitations on controllability of the system due to power and fuel scarcity, the need for highly precise sensors and clocks, and the level of dynamical knowledge that is needed to successfully predict the motion of the spacecraft. Deep-space navigation (i.e., beyond the Earth-moon system) has its own unique challenges including: even longer delays between observation windows, long communication delays with the ground and low-signal communications (due to large Earth-spacecraft distances), highly non-linear dynamics in dynamically unknown systems (e.g., asteroid proximity operations), and the limited ground communication infrastructure there is to communicate with all of the currently operating deep-space missions (one of the major motivations behind autonomous navigation research).

Deep space missions typically rely on Earth-based radiometric data types for tracking (e.g., ranging based on light-travel times, Doppler, and differenced one-way ranging at different receiver locations), but can also use on-board observables from spacecraft-mounted cameras and light detection and ranging (LIDAR) sensors. Typically, the radiometric data is two-way (ground->spacecraft->ground) in order to improve data quality by removing spacecraft clock errors; however, one-way observables are feasible with a highly-stable/accurate on-board clock like DSAC. For this analysis, we will focus on one-way radiometric data and optical data because these better support autonomous spacecraft navigation, which is a prime use-case for this research. Fusing radiometric and optical data types into a single navigation solution can often be difficult due to the data types being so independent with each other. However, this independence is powerful from an information standpoint, so it is desirable to develop a method for robust autonomous fused navigation.

To test our algorithm in a realistic deep-space navigation problem, we developed a scenario based on NASA's InSight mission—specifically, the final 45 days of the Martian approach phase. Table 1 defines our nominal filter setup—specifically where it differs from the model described by Ely et al. [2]. For completeness, our nominal filter setup includes the following parameters: spacecraft cartesian position and velocity; solar pressure scale factor; attitude control thruster magnitudes and direction; impulsive burn parameters for trajectory correction maneuvers; Earth pole motion and universal time (UT1); Tropospheric and Ionospheric delay terms; spacecraft clock offset and drift terms; ranging biases; gravitational parameters for the Earth, Moon, and Mars; ephemerides for the Earth, Mars, Deimos, Phobos, and other optical targets (i.e., asteroids); and deep space network station locations. In the measurement biasing analysis, we use the same general filter setup for each filter in the filter bank with only difference being what measurements are being processed by each filter and bias terms that are specific to specific data types (e.g., removing range bias terms from the optical-only filter). For the dynamic mis-modeling analysis, we make specific changes to the attitude thrusting and process noise models in each filter, which are defined described in Section 4.4.

Table 1. Nominal Filter Setup Used in Simulations.

Parameter	Uncertainty	Model Notes
3D Inertial Mars-Relative Position	100.0 km (diagonal)	—
3D Inertial Mars-Relative Velocity	0.1 m/s (diagonal)	—
Solar Pressure Scale Factor	0.11 (Bias, $1-\sigma$) 0.03 (Temporal, $1-\sigma$) 1 day (Batch Length) 7 days (Correlation time)	Modeled with a constant random bias and stochastic varying component that is estimated in discrete batches that are correlated in time
Attitude Thruster Acceleration Magnitude	3% (Bias, $1-\sigma$) 5% (Temporal, $1-\sigma$) 1 day (Batch Length) No Correlation	Modeled with constant bias and temporally varying component that is estimated in discrete batches that are not correlated in time
Attitude Thruster Orientation	3.0 deg	Modeled with a constant random bias
Earth Pole Motion	5.5×10^{-6} deg (Bias, $1-\sigma$) 5.5×10^{-6} deg (Temporal, $1-\sigma$) 1 h (Batch Length) 2 days (Correlation time)	Modeled with a constant random bias and stochastic varying component that is estimated in discrete batches that are correlated in time
Earth UT1	3.0×10^{-3} deg (Bias, $1-\sigma$) 3.0×10^{-3} deg (Temporal, $1-\sigma$) 1 h (Batch Length) 6 h (Correlation time)	Modeled with a constant random bias and stochastic varying component that is estimated in discrete batches that are correlated in time
Tropospheric Dry/Wet Delays	1 cm (Bias, $1-\sigma$) 1 cm (Temporal, $1-\sigma$) 1 h (Batch Length) 6 h (Correlation time)	Modeled with a constant random bias and stochastic varying component that is estimated in discrete batches that are correlated in time
Ionospheric Day/Night Delays	55/15 cm [Day] (Bias, $1-\sigma$) 55/15 cm (Temporal, $1-\sigma$) 1 h (Batch Length) 6 h (Correlation time)	Modeled with a constant random bias and stochastic varying component that is estimated in discrete batches that are correlated in time

4.2. Performance with No Biasing

The filter bank we consider in this subsection, includes three filters. This includes an optical-only filter (only processes optical data), a radiometric-only filter (only processes radiometric data), and a fused filter (processes both optical and radiometric data). The intent of this filter bank is to identify when optical and radiometric are in conflict with one another—a strong indication of mis-modeling in the system. Beyond this we would like to identify which data type is mis-modeled, and ideally be able to determine the cause of the mis-modeling.

We begin our analysis by evaluating the filter performance in the absence of measurement biasing. The position error and uncertainty for each filter in the bank are shown in

Figures 3–5. In these results, we see that the optical-only filter (Figure 3) tracks the spacecraft at a level of 10–50 km ($1\text{-}\sigma$) until right before entry, where uncertainty decreases to the sub-kilometer level over the course of a few days. Conversely, the radiometric-only filter (Figure 4) tracks consistently at the 10 km level ($1\text{-}\sigma$) until immediately before entry where position uncertainty rapidly decreases to the sub-kilometer level, while radiometric-only tracking generally has a smaller square root of the trace of the position covariance (RTPC) level than the optical-only tracking, it should be noted that there is a period right before entry where optical-only tracking begins to outperform the radiometric-only tracking. The fused filter (Figure 5) consistently outperforms the other two filters throughout the duration of the simulation. Position uncertainty ranges between 1 and 10 km ($1\text{-}\sigma$, RTPC), generally decreasing with time, until about 4-days before entry when position tracking improves significantly. It should be noted, that each filter’s position error stays within its position uncertainty bounds, though the radiometric-only filter results show that the error is generally smaller than what the uncertainty metrics dictate that it should be. This may be indicative of over-inflating radiometric data-weighting, which would result in larger uncertainties than the available information would otherwise indicate. However, this is a single realization of error in this system, so it cannot be taken as fully representative of the system’s more general error response behavior.




Figure 3. Position error (blue) and square root of the trace of the position covariance (RTPC, $1\text{-}\sigma$ (orange) and $3\text{-}\sigma$ (red)) for optical-only filter in the no-biasing setup.

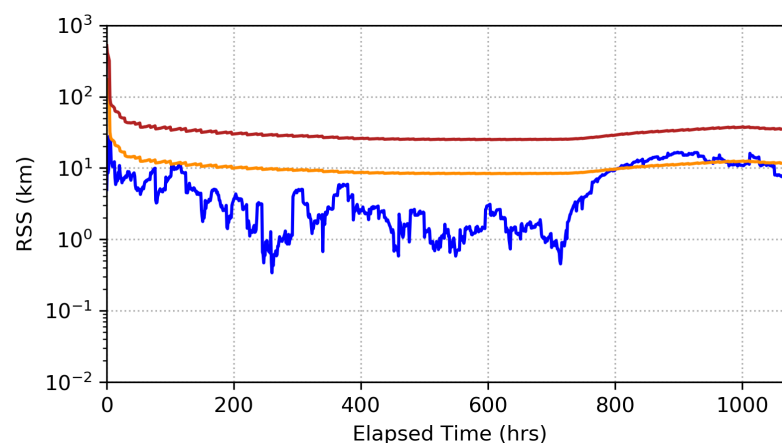


Figure 4. Position error (blue) and RTPC uncertainty ($1\text{-}\sigma$ (orange) and $3\text{-}\sigma$ (red)) for radiometric-only filter in the no-biasing setup.

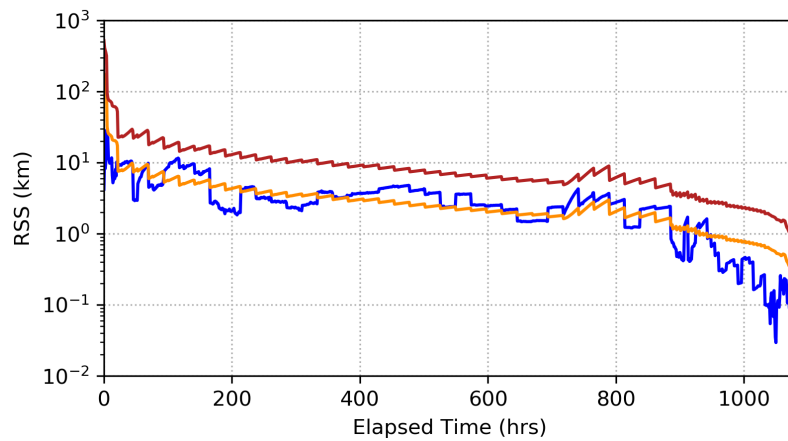


Figure 5. Position error (blue) and RTPC uncertainty ($1\text{-}\sigma$ (orange) and $3\text{-}\sigma$ (red)) for fused filter in the no-biasing setup.

When applying the Magill-based gating network to this filter bank in this nominal scenario, we obtain the results depicted in Figure 6. The results show a clear preference for the fused and radiometric-only filter for the first portion of this approach phase. This makes intuitive sense because each filter is exhibiting nominal behavior (i.e., no anomalous measurement residuals), and the fused and radiometric-only filters produce solutions with significantly smaller state uncertainties. Beyond the first portion of approach, the fused filter begins to dominate the weightings because the aggregate information from both radiometric and optical data types begins to dominate over the less-informed radiometric-only solution. These results essentially tell us that the fused filter is overwhelmingly the best performing filter in the bank under the standard MME criteria (i.e., small measurement residuals and small innovations uncertainty).

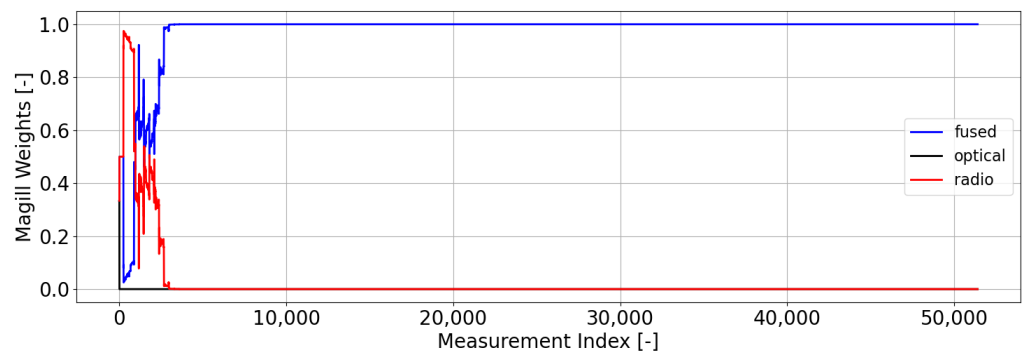


Figure 6. Magill gating network weights with properly modeled measurements.

The weightings that come from our optimal gating network are shown in Figure 7. These weightings use a batch size of 1000 measurements ($M = 1000$). Unlike the weightings from Magill’s gating network, we do not see a clear preference for any of the filters, which is an indication of a healthy (i.e., well-modeled) filter bank. Equal weightings occur when the measurement residuals and uncertainties output from the filters in the bank are statistically consistent with each other. When one or more filters begin to be preferred by the network (i.e., large weighting), this is an indication that at least one of the filters in the bank is not behaving in a statistically consistent manner—i.e., its residuals are either too large relative to its uncertainty (leading to a low weight) or its residuals are too small relative to its uncertainty (leading to a higher weight). Unlike the Magill weighting scheme, our optimal weighting scheme focuses on prioritizing filters that output residuals that are statistically small with no preference for how small the uncertainty of that filter is. Since each filter in this scenario is properly modeled, they will each output measurement residuals that are statistically consistent with their uncertainty metrics, thus resulting in nearly equal weights

across the bank. In these types of scenarios, we would usually select the default filter to use for estimation purposes and conclude that no anomalies are occurring. In this instance, that default filter is the fused filter, since it is the best-informed filter.

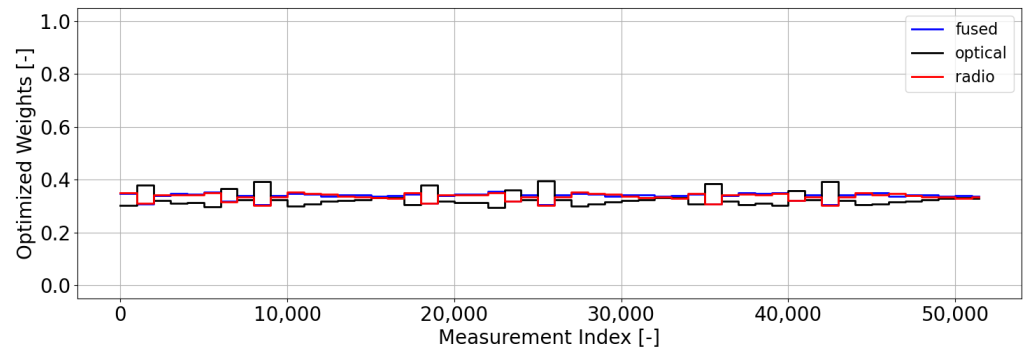


Figure 7. Optimal gating network weights with properly modeled measurements.

4.3. Performance with Measurement Biasing

Having established how these filters and gating networks perform when measurements are properly modeled, we now move toward assessing the weighting schemes' performance in the presence of measurement biasing. In order to assess the sensitivity of the filter banks to bias types, we will run different biasing scenarios where the size of the bias is varied. These scenarios are defined in Table 2. We will evaluate the effects of optical and radiometric biasing independently in the ensuing analysis. These biases are simply added to the actual measurement; however, some are representative of realistic issues in the system. The optical bias can be seen as a slight mis-modeling in the fixed orientation of the camera relative to the spacecraft attitude. The range bias can be seen as either a clock bias (spacecraft or ground-based) or transponder delay mis-model associated with the receiver electronics.

Table 2. Measurement Biasing Used in Simulations.

	Optical Biasing		Radiometric Biasing	
	Sample Bias (pix)	Line Bias (pix)	Range Bias (m)	Doppler Bias (mHz)
No Bias	0.0	0.0	0.0	0.0
Small Bias	0.2	0.3	5.0	0.03
Medium Bias	2.0	3.0	50.0	0.3
Large Bias	20.0	30.0	500.0	3.0

The effects of a small optical bias on the optical-only and fused filters are shown in Figures 8 and 9, respectively. It is clear that even this small bias (on the order of uncertainty in the measurements) has the effect of causing the filters to have significantly higher position errors than predicted by the filter's uncertainty metrics. The solutions do not fully diverge, but larger biases could cause this to occur. The radiometric-only solution is unaffected in this scenario, since it does not process the biased optical measurements.

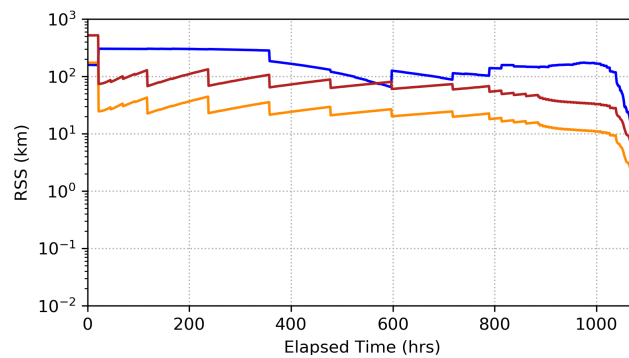


Figure 8. Position error (blue) and RTPC uncertainty (1- σ (orange) and 3- σ (red)) for optical-only filter in Small Optical Bias setup.

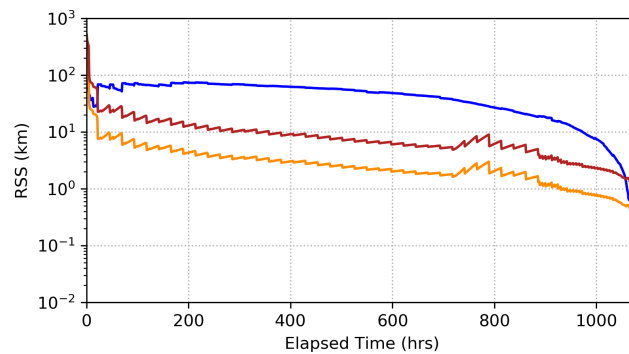


Figure 9. Position error (blue) and RTPC uncertainty (1- σ (orange) and 3- σ (red)) for fused filter in Small Optical Bias setup.

The resulting weights from the Magill gating network for these optical bias scenarios are shown in Figure 10. From these results we can make two clear conclusions: (1) the fused filter is still preferred overwhelmingly and (2) the size of the optical bias does not have a significant effect on how the Magill weights behave. Essentially, the Magill weights are unable to identify that the radiometric-only solution is the best performing filter. This is because the method for computing the weights necessitates that all measurements be used to compute filter weights, even measurements that are not processed by specific filters. This rigidity results in all three filters behaving poorly with the biased optical measurements, but because the fused filter has the smallest uncertainty it ends up rising to the top.

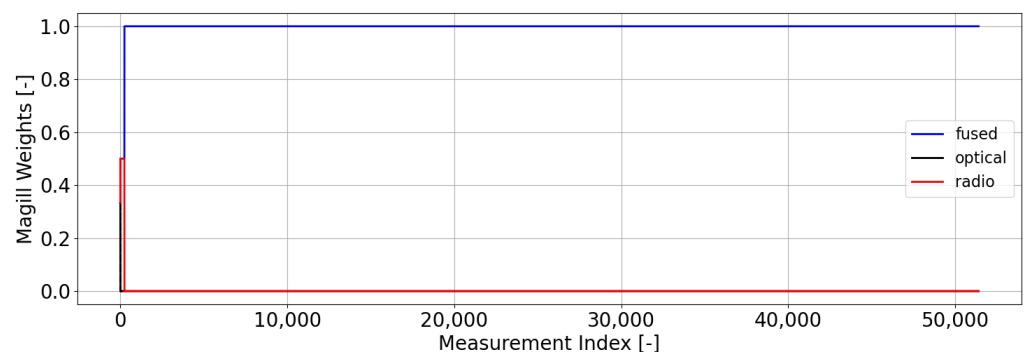


Figure 10. Magill gating network weights with biased optical measurements.

Our optimal gating network’s weights for this scenario with biased optical measurements are shown in Figure 11. The rigidity of the Magill weighting scheme is contrasted by our optimal gating network, which can be retooled to ignore unprocessed measurements when computing filter weights. The weights show a clear preference for the radio solution, and this preference becomes more significant as the size of the bias in the optical

measurement grows. As previously discussed, the fact that the filter weights deviate significantly from their nominal values (i.e., equally weighted at $1/N$) can be viewed as a sign of mis-modeling in the system. Because the radiometric solution is clearly preferred, we can assume that the source of the mis-modeling is something to which the optical-only and fused filters have greater sensitivity (e.g., biasing in the optical measurement model).

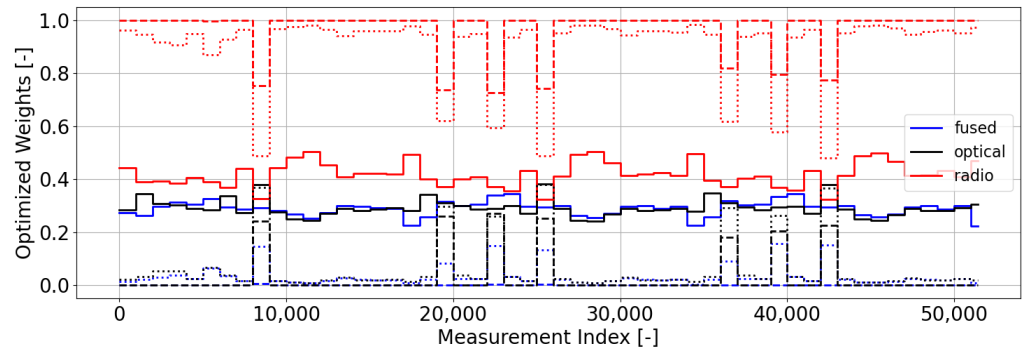


Figure 11. Optimal gating network weights with biased optical measurements. Small (solid lines), medium (dotted lines), and large (dashed lines) biasing cases are included.

These results show that the Magill weights are largely insensitive to optical biasing, while our new gating network is able to clearly identify the presence of a problem while also indicating the potential source of the issue. Moving forward, we will investigate how these gating networks and filters perform in the presence of radiometric measurement biasing. The position uncertainty and error due to a Large Radiometric Measurement Bias (Table 2) for the radiometric-only and fused filters are shown in Figures 12 and 13. Neither filter diverges because a large portion of the radiometric biases are absorbed by different parameters in the filter (e.g., ranging errors can be absorbed by clock error parameters); however, the effect of the biases can still be seen. The radiometric-only filter position error drifts noticeably outside of the $3\text{-}\sigma$ bounds. The biases' effects on the fused filter are more subtle. Essentially, the additional un-biased optical information helps constrain the solution to stay near the true trajectory, which gives an advantage over the radiometric-only solution. However, when comparing the fused filter's performance in the no bias scenario setup, we can see a slight increase in position error that is attributable to the presence of these radiometric measurement biases.

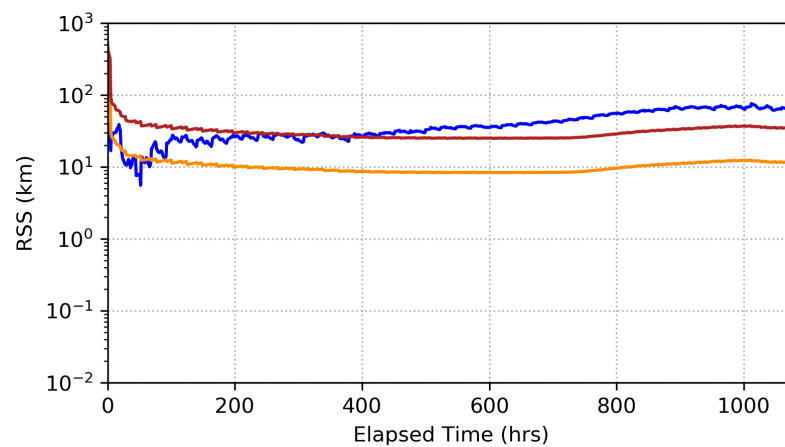


Figure 12. Position error (blue) and RTPC uncertainty ($1\text{-}\sigma$ (orange) and $3\text{-}\sigma$ (red)) for radiometric-only filter in Large Radiometric Bias setup.




Figure 13. Position error (blue) and RTPC uncertainty ($1\text{-}\sigma$ (orange) and $3\text{-}\sigma$ (red)) for fused filter in Large Radiometric Bias setup.

Further investigating the effect of these radiometric biases, the clock phase uncertainty and error statistics from the fused filter are shown in Figure 14. When processing un-biased radiometric data, the fused filter outputs clock errors that are near $0\text{-}\sigma$ in size; however, when the large radiometric bias is introduced the filter now outputs clock error in the range of $1\text{--}2\sigma$. More specifically, the clock error is approximately $2\text{ }\mu\text{s}$, which is what should be expected from the 500 m range bias that we introduced. Essentially, this result shows that the filter absorbed the range bias into the clock error modeling in a way that is not statistically anomalous. This is why the fused filter is still able to output accurate state estimates despite being given biased radiometric data. By turning down the clock phase uncertainty, we would be able to better detect smaller radiometric errors; however, it is important not set the clock uncertainty too low where the filter becomes susceptible to divergence from realistic/anticipated clock errors.

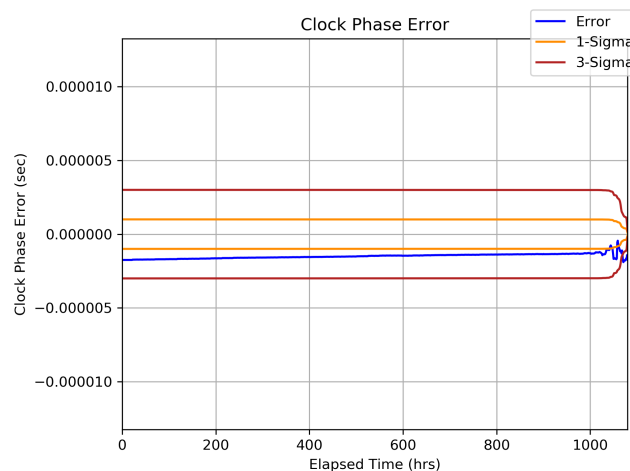


Figure 14. Clock phase error and uncertainty for fused filter in Large Radiometric Bias setup.

The weights from the Magill gating network for these biased radiometric data scenarios are shown in Figure 15. Once again, we find that that the Magill weights overwhelmingly prefer the fused filter despite the introduction of biased measurements. This makes sense, since the fused filter is still able to output statistically healthy estimates and it still produces the smallest uncertainty in its estimate. It is interesting to note that the size of the radiometric bias does have a noticeable effect on how quickly the filter weights converge to their steady-state values. As the size of the bias increases, we find that the filter weights more quickly favor the fused filter over the radiometric-only filter—the optical-only is always low-weighted due to its high uncertainty metrics. This makes intuitive sense because the radiometric-only solution is more significantly compromised by the radiometric data

biasing as compared to the fused filter. Even though this result is intuitive, we would prefer the gating network be able to identify the presence of the mis-modeling and help isolate its root cause. Our optimal gating network better exhibits this type of behavior.




Figure 15. Magill gating network weights with biased radiometric measurements. Small (solid lines), medium (dotted lines), and large (dashed lines) biasing cases are included.

The weights from our optimal gating network for these radiometric data bias scenarios are shown in Figure 16. We clearly see through these results that radiometric measurement biasing is more difficult to detect as compared to optical measurement biasing. The type of radiometric biasing that we introduced is consistent with a clock bias, and because our filter is setup to estimate clock states, the filters themselves are able to deal with the biasing as long as that biasing is not too large of a clock error (relative to our a priori clock uncertainty modeling). When the biasing is detected, the weights are still able to identify that the optical-only filter is the best behaving filter since it does not process the biased data. This result is significant since it can be used in analysis to point to the root cause of the anomaly being radiometric in nature or another similar problem that would primarily disadvantage radiometric-based tracking over optical-based tracking. The effects of smaller radiometric biases are not easily distinguishable, but this goes back to the issue of the clock and other parameters absorbing their effects. By adjusting the filter models, we could make the gating network more sensitive to radiometric biases, but the filter may then be less robust to the biases themselves. As an alternative, we could possibly add new filters to the bank that vary modeling for the parameters that absorb radiometric biases (e.g., clock modeling). This proposed extension to this investigation is left as future work.

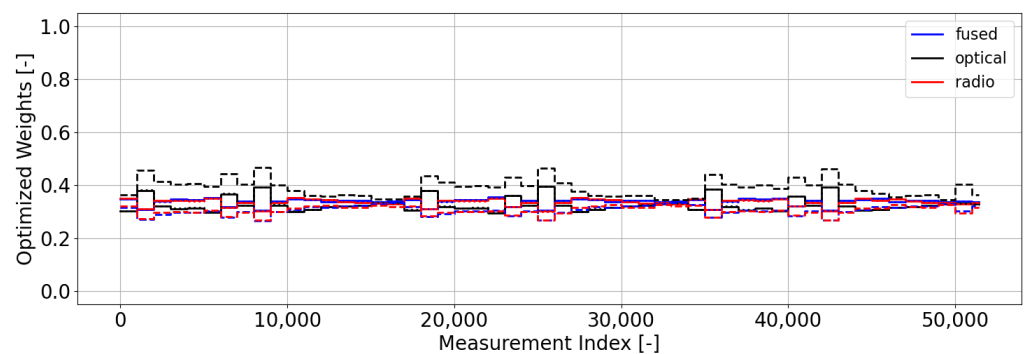


Figure 16. Optimal gating network weights with radiometric measurement biasing. Small (solid lines), medium (dotted lines), and large (dashed lines) biasing cases are included.

For each of the scenarios run to this point, we computed their average optimal gating network weight over the full data arc and summarized those results in Table 3. These average values give a general sense of the priority of the filters across the full data arc. These results really underline the conclusions that we have made up to this point:

- When all filters in the bank are properly modeled, each of the weights stays close to the ideal $1/N$ value, thus indicating nominal filter bank behavior (i.e., no anomaly is detected).
- Optical biasing is detectable via our optimal gating network at all three magnitudes (with detectability increasing with the size of the bias), and the gating network clearly prefers the radiometric-only solution over the the other two (with the fused filter performing the worst of all). This result indicates that the anomaly is likely related to optical data quality and that specifically the optical and radiometric data are inconsistent with one another.
- Given our filter setup, the radiometric biasing is not as easily detectable via our gating network. This makes sense because the filters were setup to estimate parameters that mimic the biasing that we injected, so essentially the biasing is not anomalous under our model unless it becomes too large. As it becomes larger, the optical-only solution does become preferred, which indicates the biasing is becoming too large for what our filters are modeling.

Table 3. Average Optimal Gating Network Weights for Biased Measurement Scenarios.

Case	Radio Filter	Optical Filter	Fused Filter
Ideal	0.3333	0.3333	0.3333
No Bias	0.3354	0.3272	0.3374
Small Bias—Optical	0.4186	0.2923	0.2891
Medium Bias—Optical	0.9008	0.0638	0.0354
Large Bias—Optical	0.9679	0.0316	0.0004
Small Bias—Radio	0.3354	0.3272	0.3374
Medium Bias—Radio	0.3350	0.3280	0.3370
Large Bias—Radio	0.3047	0.3899	0.3054

Through these results, we have found that the Magill gating network is largely insensitive to measurement biasing in radiometric-optical fused deep space navigation. When biasing occurs, the Magill weights consistently prefer the fused filter even when it is corrupted by the biased information. Our new optimal gating network, however, is able to detect statistically significant mis-modeling. Beyond this, its outputs may also be used to isolate the possible root causes behind detected anomalies, which can greatly aid in operational analysis. Moving forward, we will investigate how this new optimal gating network performs in the presence of dynamic mis-modeling.

4.4. Dynamical Model Mis-Modeling

To investigate the efficacy of our new gating network in detecting the presence of dynamic mis-modeling, we will introduce more dynamic variations in the truth data and then slightly adjust the filter bank to include filters with varied parameters for acceleration estimation. In terms of truth data changes, we significantly increase the error in the attitude control thrusting such that each thruster has a random constant bias drawn from a zero-mean normal distribution with a standard deviation of $3 \times$ the nominal thrust level (5×10^{-12} km/s²) and daily temporal variations in the magnitude that are drawn from a zero-mean normal distribution with a standard deviation of $10 \times$ the nominal thrust level. We leave the attitude thruster direction uncertainty untouched relative to the nominal scenario setup.

In order to enable detection of dynamic mis-modeling in our filter bank, we introduce variations of our existing three filters where the new versions have different uncertainties on attitude control thruster magnitude parameters. Our standard filters (radio, optical, and fused) remain unchanged relative to the scenario setup, so they have large errors in

their attitude control thrusting models—they only model a 3% bias and a 5% daily variation in the attitude thruster magnitudes. The new filters (radio-Q, optical-Q, and fused-Q) are all modeled properly—they model a 300% bias and a 1000% daily variation in the attitude thruster magnitudes, while thrusters in reality would never be this poorly behaved, this type of scenario is reminiscent of the Mars Climate Orbiter crash, where a units modeling error led to significant modeling issues in the attitude control system. It should be noted that these modeling errors only occur during attitude thrusting, so the errors are not constant, and thus will not be as detectable as a comparably-sized constant bias in acceleration. The intent of this exercise is to see whether these dynamically mis-modeled filters are de-weighted by our gating network in a way that would allow us to detect the anomaly as well as determine the cause of the anomaly.

Before investigating how the gating networks react to the introduction of this attitude thruster mis-modeling, we first want to quantify the effect of this dynamic mis-modeling on the individual filters. Figures 17–19 show the position error and uncertainty metrics that results from the optical-only, radiometric-only, and fused filters that are dynamically mis-modeled. The optical-only filter (Figure 17) seems relatively unaffected by the presence dynamic mis-modeling—this is a different truth realization from the no-biasing case, but the position error is similarly well-bounded by its uncertainty envelope. Conversely, the radiometric-only filter (Figure 18) is noticeably affected by the dynamic mis-modeling. Position tracking error begins to steadily increase over time (relative to the uncertainty bounds) until it exceeds the $3\text{-}\sigma$ bound for an extended period. The fused filter (Figure 19) exhibits similar behavior—Its position error grows beyond the $3\text{-}\sigma$ bound for an extended period, but it begins to recover near entry as the optical data becomes more information rich. These results tell us that radiometric data is far more sensitive to this type of dynamic mis-modeling than optical data. In initial analysis, this increased sensitivity does not seem to be associated with simulation geometry (i.e., net attitude control thrust orientation with respect to Earth and optical targets), but it may be due to the uncertainties present in the filter. The optical filter has larger state uncertainties, in general, which can make it less sensitive to dynamic mis-modeling and it has fewer parameters than radiometric data (e.g., clock error, transponder bias, etc.) that could improperly absorb the dynamic mis-modeling.




Figure 17. Position error (blue) and RTPC uncertainty ($1\text{-}\sigma$ (orange) and $3\text{-}\sigma$ (red)) for optical-only filter with attitude thruster mis-modeling.

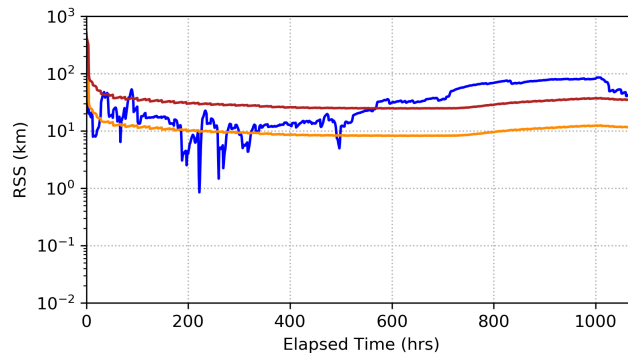


Figure 18. Position error (blue) and RTPC uncertainty (1- σ (orange) and 3- σ (red)) for radiometric-only filter with attitude thruster mis-modeling.




Figure 19. Position error (blue) and RTPC uncertainty (1- σ (orange) and 3- σ (red)) for fused filter with attitude thruster mis-modeling.

The Magill gating network weights for our filter bank with 6 filters are shown in Figure 20. These results show that the Magill gating network overwhelmingly favors the properly modeled fused filter (fused-Q). This makes sense because this filter is properly modeled and fully-informed, so it would have small residuals and small uncertainties. By strongly favoring the fused-Q filter over the fused filter, we can postulate that there is mis-modeling in the system because the lower uncertainty filter (the fused) should be preferred if no mis-modeling were present—it is hard to validate this fully, though, because the Magill weights have been shown to be unable to reliably detect mis-modeling through our previous examples. There is a tradeoff between increasing dynamic uncertainty parameters in the filter. By increasing these parameters, we essentially lessen statistical constraints on the data, which allow the filter to absorb more error, which can yield smaller measurement residuals; however, this comes with higher estimate uncertainties, generally. The Magill gating network accounts for the size of the measurement residual and the innovations' uncertainty, which means adding dynamic uncertainty to the system can yield results that are difficult to predict.




Figure 20. Magill gating network weights with attitude thruster mis-modeling.

The weights from our optimal gating network are shown in Figure 21. The results clearly indicate an anomaly is occurring due to the fact the weights are consistently different from one another. We see that this gating network indicates a preference primarily for the optical filters, and secondarily to the properly modeled fused and radiometric-only filters. The dynamically mis-modeled fused and radiometric-only filters have the least priority. These results are fairly intuitive based on our previous analysis. We previously saw that the optical filters were generally unaffected by the dynamic mis-modeling, so we would expect them to rank highly in the gating network. Similarly, the properly modeled radiometric-only and fused filters, should perform well because they are not subject to any unaccounted biasing. However, two results are odd: (1) the optical filters outperform the properly modeled radiometric-only and fused filters and (2) the effects of the dynamic mis-modeling seem to diminish over time. In the former case, this seems to imply that the optical filters produce lower residuals than the other two properly modeled filters in a χ^2 -sense. This is somewhat confirmed by the position tracking results (Figure 17), which shows position error for the optical filter is generally below the $1\text{-}\sigma$ bound. This might imply that the uncertainty metrics from the optical filter are too large in this scenario. In the latter case, this result is somewhat un-intuitive because we saw that position error in the fused and radiometric-only filters actually grows over time, so we would expect their residuals to grow commensurately. This result would lead to drops in the weights for these two filters, not increases as we have observed. This type of result warrants more study and should be addressed in future work.

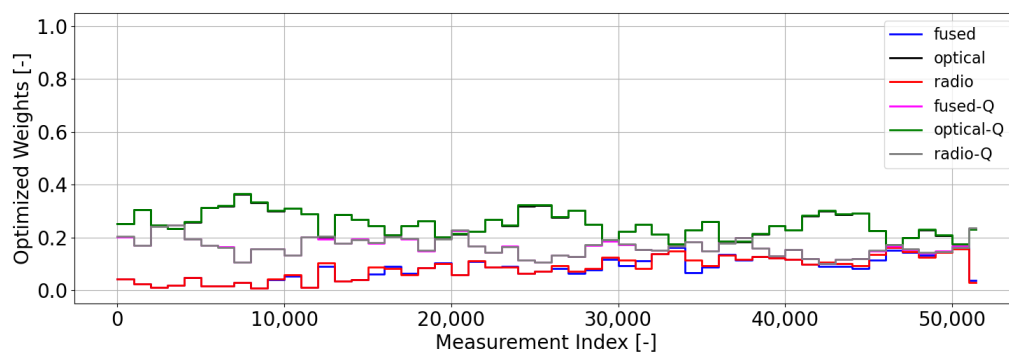


Figure 21. Optimal gating network weights with attitude thruster mis-modeling.

Overall, the weights from this gating network do tell us that an anomaly is occurring in the system. The gating network's preference for the optical filters seems to indicate that the problem is something to which the optical data is relatively insensitive. Furthermore, the gating network's secondary preference for the higher dynamic uncertainty filters seems to indicate the issue may be dynamic in nature. By combining these inferences, we can hypothesize that there is a dynamic mis-modeling in the system that acts in a manner to which radiometric data is especially sensitive. This once again shows the utility of this gating network in its abilities to both identify the presence of mis-modeling and provide a means for characterizing what that mis-modeling may be.

As a final test of this gating network's capabilities, we consider an extension to the dynamic mis-modeling problem. Our previous example focused on whether the gating network could identify mis-modeling when half of the filter's in the bank were properly modeled and the other half under-represented dynamic variations in the attitude thruster model. In this next example, we mis-model all six filters in the same manner—we increase the true nominal attitude thruster magnitude by a factor of ten, while keeping the original values in the filters. This results in a modeling error on the order of $1 \times 10^{-11} \text{ km/s}^2$. We then remove the attitude thruster magnitude-specific parameters from the filters as well as the temporal estimation of the SRP parameter (just estimated as a bias now), and instead we add parameters for a batched (1-day batch length) polynomial acceleration model to the filters to account for mis-modeled accelerations (i.e., process noise). The standard

filters (i.e., fused, radio, and optical) all use an a priori sigma of 5×10^{-12} km/s² for this polynomial acceleration model (termed low-uncertainty filters going forward) while the high-uncertainty filters (i.e., fused-Q, radio-Q, and optical-Q) all use an a priori sigma of 5×10^{-11} km/s². This filter setup is designed to detect dynamic mis-modeling when the filter weights begin to favor the high uncertainty models.

The error and uncertainty behaviors for this setup are similar to what we have seen previously. The low-uncertainty radiometric and fused filters have position errors that quickly diverge outside the $3\text{-}\sigma$ tracking envelopes due to the mis-modeling. The low-uncertainty optical filter does experience some degraded tracking (especially right before entry), but it is clearly less sensitive to the mis-modeling than the filters that use radiometric data. The high-uncertainty filters perform much better than their low-uncertainty counterparts. The error remains bounded by the uncertainty envelopes, which indicates that the higher uncertainty polynomial acceleration model is appropriately compensating for the mis-modeled attitude control thruster models. The Magill gating network metrics for this scenario are also similar to the previous example. This gating network overwhelmingly favors the high-uncertainty fused filter, since it is the only fully-informed filter that outputs measurement residuals that are not too large relative to its expected level of uncertainty. Once again, however, this result is hard to draw a conclusion from, since we have already shown the Magill weights do not reliably detect the presence of anomalies.

The filter weights from our optimal gating network are shown in Figure 22. Other than in the middle of the arc, several things are apparent in these results: (1) anomalous behavior is detectable due to the filters deviating from equal weighting, (2) the low-uncertainty fused and radiometric-only solutions are the worst performing filters, (3) the filter that is the most consistently highly-weighted is the optical high-uncertainty filter, and (4) all three low-uncertainty filters have low weights right before entry. These results seem to indicate that an anomaly is occurring in the system to which radiometric data is more sensitive than optical data except right near entry. Because the high uncertainty filters generally perform better, we might assume that this means the source of the anomaly is dynamic in nature, but as already mentioned this gating network currently is vulnerable to filters with over-inflated uncertainty metrics. This motivates us to address this vulnerability through future research. The performance of the gating network in the middle of the arc (near measurement index 30,000) is also an area of future research, as this is an unexpected result that does not have an obvious answer.

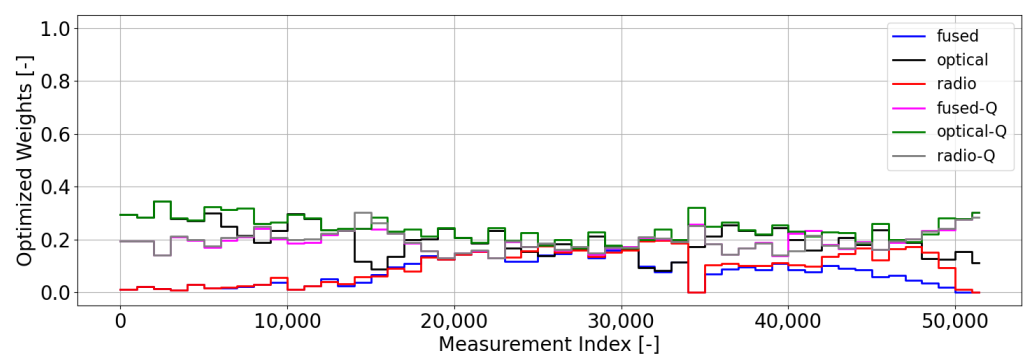


Figure 22. Optimal gating network weights with mis-modeled dynamics and varied process noise models.

As with the measurement biasing analysis, we computed the average optimal network weightings for each filter in the bank for both dynamic mis-modeling scenarios (Table 4). These results reinforce the following conclusions we have previously made:

- An anomaly is detectable in both dynamic mis-modeling scenarios due to the gating network's deviation from the ideal weights
- The source of the anomaly in both instances more significantly affects radiometric-based solutions than optical-based solutions. Radiometric and fused filters perform

- similarly, so this seems to indicate that the radiometric and optical data are not fighting each other
- The higher uncertainty filters are preferred in both scenarios, which seems to indicate that the source of the mis-modeling may be dynamic in nature—specifically, mis-modeling to which optical data is less sensitive than radiometric data.

Table 4. Average Optimal Gating Network Weights for Mis-Modeled Dynamics Scenarios.

Filters	Ideal	Attitude Thruster Scenario	Process Noise Scenario
Radio Filter	0.1666	0.0844	0.0961
Optical Filter	0.1666	0.2510	0.1992
Fused Filter	0.1666	0.0815	0.0742
Radio-Q Filter	0.1666	0.1658	0.1945
Optical-Q Filter	0.1666	0.2520	0.2431
Fused-Q Filter	0.1666	0.1653	0.1929

The simulations and results discussed in this section have shown that our optimal gating network has the ability to identify and characterize filter anomalies related to measurement and dynamic mis-modeling. These results show promise and warrant future research into this method with the ultimate goal of being a method that can autonomously navigate a spacecraft while detecting and diagnosing problems as they occur.

5. Conclusions

In this study, we investigated the usage of filter bank gating networks for the purposes of anomaly detection in deep space autonomous navigation with fused observable types. We specifically investigated methods from the areas of Multiple Model Estimation (MME) and Mixture of Experts (MoE). Existing methods did not tend to deal with the problem of measurement-fused estimation, so our work focused on developing an MME-based gating network that could be used to detect anomalies in a system, especially those relating to conflicting information from different measurement data types. This new gating network was designed to reward filters that produce measurement residuals that are consistent with the filter’s uncertainty metrics while balancing this against a term that incorporates memory of previous filter weightings and ensures that the weightings are properly normalized. We then expanded this algorithm to deal with filters that only process a subset of the available measurement types, which is vital to dealing with the fused-sensor problem.

After developing this new gating network, we tested it and a standard gating network against a numerical simulation based on the Insight spacecraft’s Mars approach phase navigation with fused, optical-only, and radiometric-only data. In these results, we found that our new optimal gating network was able to detect statistically significant measurement mis-modeling—indicated by the gating network weights deviating from equality. Optical sensor biasing stood out clearly (especially as the size of the bias grew) due to there being no modeling in the filters that accounted for this type of biasing. Radiometric biasing was not as easily detectable, though, because the filters already accounted for clock errors, which were consistent in size with the introduced biasing. Essentially, the radiometric biasing was not statistically significant because our models already accounted for the possible presence of these level of errors. For dynamic mis-modeling, we introduced mismodeling into the attitude control thrusting and added new filters to the filter bank that better modeled the introduced level of dynamic uncertainty. We found that the gating network weights clearly identified the presence of the dynamic mis-modeling. We were also able to characterize the mis-modeling as dynamic in nature due to the network’s preference for the large dynamic uncertainty versions of the filters and the fused filters did not indicate any problems with radiometric and optical information consistency. In each of our simulations we compared our results against the existing Magill gating network and found that it provided no means

to consistently identify mis-modeling and often selected filters that were mis-modeled as the optimal solution. Not being able to consistently detect the presence of mis-modeling, the Magill method also had no means to characterize anomalies like our optimal gating network is able to do. This characterization process is not automated at this point, but our gating network provides a good foundation that is flexible enough to adapt to specific mission setups and different types of anticipated anomalies.

These results demonstrated that our new algorithm shows significant promise for its intended use-case, but we did identify areas to expand and improve this algorithm in the future. Specifically, work should be commenced to adjust the algorithm such that it does not prefer filter solutions with over-inflated uncertainty metrics. Furthermore, we need to further push this algorithm toward its autonomous use-case by defining an automated method for triggering an anomaly detection by the gating network, developing a fully-defined decision tree to characterize the anomaly using the filter weightings as a guide, and assessing the computational load of this algorithm using a flight-like code framework. Finally, developing simulations that more realistically represent a broader set of expected anomalies in a navigation solutions will better demonstrate the capabilities of this algorithm.

Author Contributions: Conceptualization, D.P.L. and T.A.E.; Methodology, D.P.L. and T.A.E.; Software, D.P.L. and T.A.E.; Validation, D.P.L. and T.A.E.; Formal Analysis, D.P.L.; Investigation, D.P.L.; Resources, T.A.E.; Data Curation, D.P.L. and T.A.E.; Writing—Original Draft Preparation, D.P.L.; Writing—Review & Editing, D.P.L. and T.A.E.; Visualization, D.P.L. and T.A.E.; Supervision, T.A.E.; Project Administration, T.A.E.; Funding Acquisition, T.A.E. All authors have read and agreed to the published version of the manuscript.

Funding: This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). The authors would also like to thank Shyam Bhaskaran for his input on this work.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Ely, T.A.; Burt, E.A.; Prestage, J.D.; Seubert, J.M.; Tjoelker, R.L. Using the Deep Space Atomic Clock for Navigation and Science. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control.* **2018**, *65*, 950–961. [[CrossRef](#)]
2. Ely, T.A.; Seubert, J.; Bradley, N.; Drain, T.; Bhaskaran, S. Radiometric Autonomous Navigation Fused with Optical for Deep Space Exploration. *J. Astronaut. Sci.* **2021**, *68*, 300–325. [[CrossRef](#)]
3. Schutz, B.; Tapley, B.; Born, G. *Statistical Orbit Determination*; Elsevier Science: Amsterdam, The Netherlands, 2004.
4. Jazwinski, A. *Stochastic Processes and Filtering Theory*; Dover Books on Electrical Engineering Series; Dover Publications: New York, NY, USA, 2007.
5. Rong Li, X.; Jilkov, V. Survey of maneuvering target tracking. Part I. Dynamic models. *IEEE Trans. Aerosp. Electron. Syst.* **2003**, *39*, 1333–1364. [[CrossRef](#)]
6. Chan, Y.; Hu, A.; Plant, J. A Kalman Filter Based Tracking Scheme with Input Estimation. *IEEE Trans. Aerosp. Electron. Syst.* **1979**, *AES-15*, 237–244. [[CrossRef](#)]
7. Bar-Shalom, Y.; Birmiwal, K. Variable Dimension Filter for Maneuvering Target Tracking. *IEEE Trans. Aerosp. Electron. Syst.* **1982**, *AES-18*, 621–629. [[CrossRef](#)]
8. Goff, G.M.; Black, J.; Beck, J.A. Orbit Estimation Of A Continuously Thrusting Satellite Using Variable Dimension Filters. In Proceedings of the AIAA Guidance, Navigation, and Control Conference, Kissimmee, FL, USA, 5–9 January 2015. [[CrossRef](#)]
9. Ljung, L. *System Identification: Theory for the User*; Prentice Hall Information and System Sciences Series; Prentice Hall: Hoboken, NJ, USA, 1999.

10. Ahmed, J.; Coppola, V.; Bernstein, D. Asymptotic tracking of spacecraft attitude motion with inertia matrix identification. In Proceedings of the 36th IEEE Conference on Decision and Control, San Diego, CA, USA, 12 December 1997; Volume 3, pp. 2471–2476. [[CrossRef](#)]
11. Chandrasekar, J.; Bernstein, D. Position Control Using Acceleration-Based Identification and Feedback With Unknown Measurement Bias. *J. Dyn. Syst. Meas. Control. Trans. ASME* **2008**, *130*, 014501. [[CrossRef](#)]
12. Wang, D.; Haldar, A. System identification with limited observations and without input. *J. Eng.-Mech. ASCE* **1997**, *123*, 504–510. [[CrossRef](#)]
13. Brincker, R.; Zhang, L.; Andersen, P. Modal identification of output only systems using Frequency Domain Decomposition. *Smart Mater. Struct.* **2001**, *10*, 441. [[CrossRef](#)]
14. Patera, R.P. Space Event Detection Method. *J. Spacecr. Rocket.* **2008**, *45*, 554–559. [[CrossRef](#)]
15. Hill, K. Maneuver Detection and Estimation with Optical Tracklets. In Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference, Maui, HI, USA, 27–30 September 2014; Ryan, S., Ed.; The Maui Economic Development Board: Kihei, HI, USA, 2014; p. 26.
16. Lemmens, S.; Krag, H. Two-Line-Elements-Based Maneuver Detection Methods for Satellites in Low Earth Orbit. *J. Guid. Control. Dyn.* **2014**, *37*, 860–868. [[CrossRef](#)]
17. Folcik, Z.; Cefola, P.; Abbot, R. Geo maneuver detection for space situational awareness. *Adv. Astronaut. Sci.* **2008**, *129*, 523–550.
18. Holzinger, M.J.; Scheeres, D.J.; Alfriend, K.T. Object Correlation, Maneuver Detection, and Characterization Using Control Distance Metrics. *J. Guid. Control. Dyn.* **2012**, *35*, 1312–1325. [[CrossRef](#)]
19. Singh, N.; Horwood, J.; Poore, A. AAS 12-159 Space Object Maneuver Detection via a Joint Optimal Control and Multiple Hypothesis Tracking Approach. In Proceedings of the 22nd AAS/AIAA Space Flight Mechanics Meeting, Charleston, SC, USA, 29 January–2 February 2012; Volume 143.
20. Jaunzemis, A.D.; Mathew, M.V.; Holzinger, M.J. Control Cost and Mahalanobis Distance Binary Hypothesis Testing for Spacecraft Maneuver Detection. *J. Guid. Control. Dyn.* **2016**, *39*, 2058–2072. [[CrossRef](#)]
21. Lubey, D.; Scheeres, D. Identifying and Estimating Mismodeled Dynamics via Optimal Control Policies and Distance Metrics. *J. Guid. Control. Dyn.* **2014**, *37*, 1512–1523. [[CrossRef](#)]
22. Rao, C.; Rawlings, J.; Mayne, D. Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations. *IEEE Trans. Autom. Control.* **2003**, *48*, 246–258. [[CrossRef](#)]
23. Crassidis, J.; Junkins, J. *Optimal Estimation of Dynamic Systems*; Chapman & Hall/CRC Applied Mathematics & Nonlinear Science, CRC Press: Boca Raton, FL, USA, 2004.
24. Mook, D.J.; Junkins, J.L. Minimum model error estimation for poorly modeled dynamic systems. *J. Guid. Control. Dyn.* **1988**, *11*, 256–261. [[CrossRef](#)]
25. Lubey, D.P. Maneuver Detection and Reconstruction in Data Sparse Systems with an Optimal Control Based Estimator. Ph.D. Thesis, University of Colorado, Boulder, CO, USA, 2015.
26. Hanlon, P.; Maybeck, P. Characterization of Kalman filter residuals in the presence of mismodeling. *IEEE Trans. Aerosp. Electron. Syst.* **2000**, *36*, 114–131. [[CrossRef](#)]
27. Nebelecky, C.K.; Crassidis, J.L.; Singla, P. A Model Error Formulation of the Multiple Model Adaptive Estimation Algorithm. In Proceedings of the 17th International Conference on Information Fusion, Salamanca, Spain, 7–10 July 2014.
28. Alsuwaidan, B.N.; Crassidis, J.L.; Chen, Y. Generalized Multiple-Model Adaptive Estimation Using an Autocorrelation Approach. In Proceedings of the 9th International Conference on Information Fusion, Florence, Italy, 10–13 July 2006.
29. Marschke, J.M.; Crassidis, J.L.; Lan, Q.M. Multiple Model Adaptive Estimation for Inertial Navigation During Mars Entry. In Proceedings of the 2008 AIAA/AAS Astrodynamics Specialists Conference, Honolulu, HI, USA, 18–21 August 2008.
30. Xiong, K.; Wei, C.L.; Liu, L.D. Robust multiple model adaptive estimation for spacecraft autonomous navigation. *Aerosp. Sci. Technol.* **2015**, *42*, 249–258. [[CrossRef](#)]
31. Chaer, W.S.; Bishop, R.H.; Ghosh, J. A Mixture-of-Experts Framework for Adaptive Kalman Filtering. *IEEE Trans. Syst. Man, Cybern.* **1997**, *27*, 452–464. [[CrossRef](#)]
32. Lee, S.; Hwang, I. Interacting Multiple Model Estimation for Spacecraft Maneuver Detection and Characterization. In Proceedings of the 2015 AIAA Guidance, Navigation, and Control Conference, Kissimmee, FL, USA, 5–9 January 2005.
33. Crain, T.P.; Bishop, R.H.; Ely, T.A. Event Detection and Characterization During Autonomous Interplanetary Navigation. *J. Guid. Control. Dyn.* **2002**, *25*, 394–403. [[CrossRef](#)]
34. Haykins, S. *Neural Networks: A Comprehensive Foundation*, 2nd ed.; Prentice Hall: Upper Saddle River, NJ, USA, 1999.
35. Magill, D.T. Optimal Adaptive Estimation of Sampled Stochastic Processes. *IEEE Trans. Autom. Control.* **1965**, *10*, 434–439. [[CrossRef](#)]
36. Li, X.R.; Jilkov, V.P. Survey of Maneuvering Target Tracking. Part V: Multiple Model Methods. *IEEE Trans. Aerosp. Electron. Syst.* **2005**, *41*, 1255–1321.