Evaluation of Maneuver Detection within an Autonomous, Heterogeneous Sensor Network

Jonathan Kadan
Virginia Polytechnic Institute and State University
Amit Bala, Kevin Schroeder, PhD. Jonathan Black, PhD.
Virginia Polytechnic Institute and State University

ABSTRACT

The number of Resident Space Objects (RSO) is in the tens of thousands and this problem is being exacerbated by the rapidly accelerating rate of satellite launches. Coinciding with the launch of new RSO is the increasing maneuver capability of satellites. These more agile RSO can perform many types of maneuvers including station keeping and Rendezvous and Proximity Operations (RPO) at small timescales and at close distances to their targets. The increasing capabilities of RSO further emphasizes the need for efficient, autonomous sensor tasking and maneuver detection.

In order to maintain custody of this growing and maneuverable catalog, Space Domain Awareness (SDA) is trending away from human-in-the-loop satellite observation tasking and towards human-on-the-loop autonomous sensor tasking. These human-on-the-loop systems will need to identify, observe, and re-schedule observations of maneuvered satellites as quickly as possible to prevent loss of custody. For an autonomous system to react to changes in the SDA environment, accurate, real-time maneuver detection algorithms must be implemented.

This work focuses on the implementation of various maneuver detection techniques in an autonomous SDA simulation and evaluation of these algorithms for satellite maneuver detection. The standard practice in satellite maneuver detection is to evaluate the standard Normalized Innovations Squared (NIS) metric of an incoming observation. However, this is not the only option. Other potential choices for maneuver detection consist of variations of the standard NIS metric: sliding NIS and fading memory NIS.

From here, different sensor tasking solutions were implemented to determine a computationally efficient method of reward function augmentation that emphasized revisit of maneuvered RSO. To generate the data necessary for this analysis, scenarios with RSO in near Geosynchronous Earth Orbits (GEO) and maneuver events will be generated. The maneuvers modelled were representative of East/West and North/South station keeping maneuvers. Ground Based Electro-Optical (EO) and radar sensors, with their necessary constraints, were modeled. For example, the optical sensors modeled will implement necessary lighting constraints to determine RSO visibility.

An Unscented Kalman Filter (UKF) was used to consolidate observation data and produce RSO estimate ephemerides. A special perturbations dynamics model was used for RSO propagation. Sensor specifications were approximated with open-source documentation of the current Space Surveillance Network (SSN).

Comparison between control group scenarios without maneuvers and those with maneuvers were used to show how the augmented reward functions direct sensor tasking toward RSO w/ detected maneuvers while still maintaining catalog coverage. Maneuver detection failed at high process noise values and false positive maneuver detections were identified in low process noise simulations; the UKF was tuned accordingly.

This work was used to evaluate which maneuver detections algorithms perform best for GEO orbits and determined suitable process noise levels for the UKF to predict measurement residuals for maneuver detections.

For the radar sensing scenarios, both Lambert Initial Orbit Determination (IOD) and Generalized Pseudo-Bayesian Estimator of Order 1 (GPB1) Multiple Model Adaptive Estimation (MMAE) were implemented to reduced RSO residuals to within nominal levels without having to severely modify the tasking solution.

1. INTRODUCTION

Due to the exponential uptick in the commercialization of space, accurate and reliable Space Traffic Management (STM) has become increasingly critical to ensuring safe space operations. While the majority of attention has been dedicated to tracking the ever-growing proliferated Low Earth Orbit (pLEO) constellations [9, 7], the RSO population in Geosynchronous Earth Orbit (GEO) has been steadily increasing as well [27, 23]. The GEO belt consists vital satellite networks for global communication needs [24, 10]. Additionally, the amount debris in GEO has been growing at a consistent rate [1].

Coinciding with the increased congestion in GEO, satellite owners are forced to maintain orbits in increasing smaller longitude slots. The distribution of RSO in various GEO longitude slots can be seen below in Fig. 1.

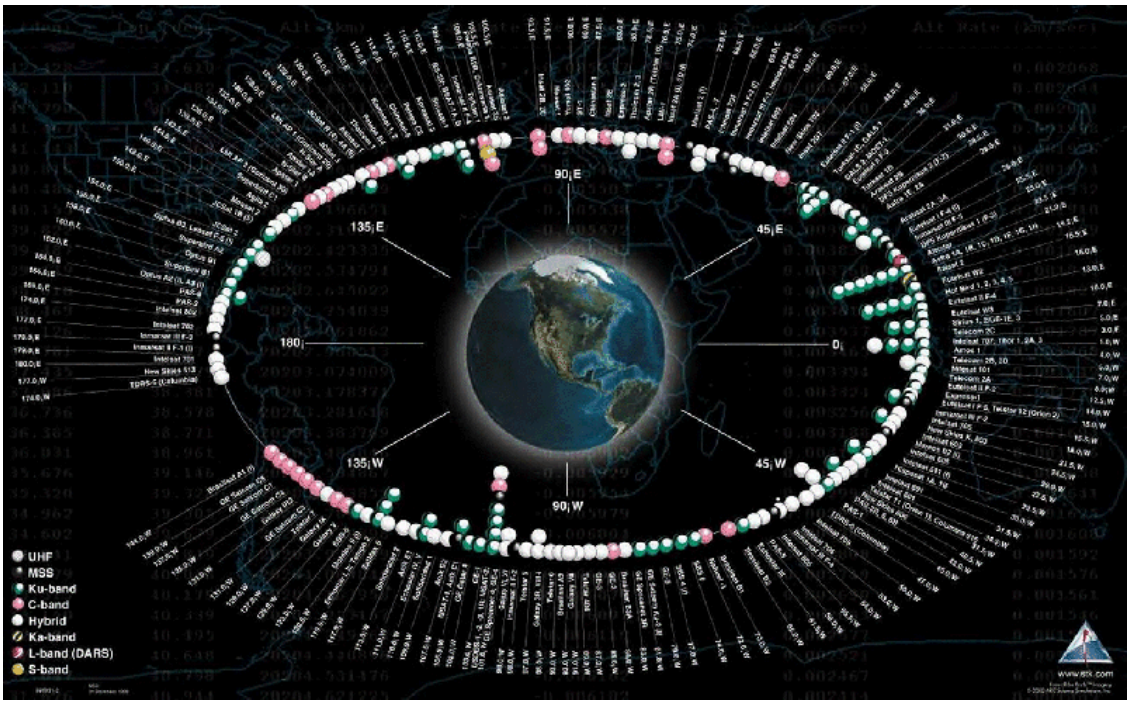

Fig. 1: GEO RSO Population [10]

This congestion forces the operators of these satellite networks to perform station keeping maneuvers at a more frequent rate.

1.1 GEO Satellite Maneuvers

RSO in GEO drift out of their prescribed slots at different rates due to the zonal gravitational harmonics of the Earth gravitational field [5]. GEO RSO maintain their orbital slot (longitude) via East-West Station Keeping (EWSK) maneuvers and required inclination (latitude) via North-South Station Keeping (NSSK). While it depends on the longitude at which an RSO's GEO orbital slot is located, EWSK must be performed, on average, every few weeks and NSSK every few months [5]. For NSSK, the average annual ΔV consumed is 47 m/s and for EWSK it is about a factor of 20 smaller, or about 2.3 m/s annually [5].

As can be seen in equations 1 and 2, NSSK is driven by the need to keep the RSO inclination within prescribed bounds (usually $\pm 0.1\ \Delta i$) where θ is a combination of inclination and change in Right Ascension of the Ascending Node (RAAN) and V is the magnitude of the velocity of the RSO in question.

$$\Delta V = 2V sin(\frac{\Delta\theta}{2}) \tag{1}$$

$$\Delta\theta = cos^{-1}[sin^2 icos(\Omega_2\Omega_1) + cos^2 i] \tag{2}$$

While NSSK ΔV is driven by the inclination needed to maintain, EWSK ΔV requirements are dominated by the semi-major axis (SMA) that a GEO RSO is needed to maintain.

$$\Delta V = \frac{1}{2} n \Delta a \tag{3}$$

The further an RSO's SMA differs from the exact semi-major axis, 42,164 km, needed for Geostationary orbit, the more ΔV is required to maintain its orbital slot. This relationship can be seen via the dependence on Δa in equation 3.

1.2 Maneuver Detection

To account for the increasing number of GEO RSO maneuvers, accurate and reliable maneuver detection algorithms are needed to identify changes in RSO states and assess potential conjunction events. Three maneuver detection algorithms were investigated in this work: the standard Normalized Innovations Squared (NIS), Sliding NIS, Fading Memory NIS.

The standard NIS is a one-sided chi-squared (χ^2) test with n_z degrees of freedom (DoF) involving the measurement residual ν and innovation covariance S_ν against some predefined threshold [2]. The key feature of the standard NIS is that it is completely Markovian. It is evaluated against a single measurement so that it is the most reactive to change, but at the cost of not considering past measurements.

$$\varepsilon_\nu(k) = \nu(k)^T S_\nu(k)^{-1} \nu(k) \approx \chi^2_{n_z} \tag{4}$$

The sliding NIS is a modification of the standard NIS that averages the NIS over the last s timesteps. This gives the sliding NIS sn_z additional degrees of freedom, where z is the DoF of the measurement space [2]. Unlike the standard NIS, it is traditionally more biased towards non-detection in the event of abrupt change, such as an impulsive maneuver.

$$\varepsilon^s_\nu(k) = \sum_{j=k-s+1}^{k} \varepsilon^s_\nu(j) \approx \chi^2_{sn_z} \tag{5}$$

Fading memory NIS averages NIS over a selected number of timesteps, but with more priority given to more recent timesteps. The modifying term δ where $0 < \delta < 1$ can be tuned to give more preference to past NIS values or the most recent term [2].

$$\varepsilon^\delta_\nu(k) = \delta\varepsilon^\delta_\nu(k-1) + \varepsilon_\nu(k) \approx \frac{1}{1+\delta}\chi^2_{n_\delta} \tag{6}$$

As can be seen in equation 6, a δ value of 0 would reduce to the standard NIS equation. This allows the fading memory NIS to have preference to the recent baseline while also being reactive to changes in the environment.

1.3 Sensing

The measurement residuals used to perform the maneuver detection methods mentioned in the previous section can be predicted by various different sensing technologies. Traditionally, Electro-Optical (EO) sensors have been used to track GEO RSO due to their ability to view dim objects in deep space and the difference in sidereal rate of satellites compared to that of the background stars [3].

The pitfall of traditional metric-based EO sensors is that they only measure two DoF, the angular components of the 2D image. These measurements are usually recorded as Right Ascension and Declination or Azimuth and Elevation.

Table 1: Optical Sensor specification modeled

R Matrix [1]	Slew Rate (deg/s)	Azimuth Mask (deg)	Elevation Mask (deg)	FoV (deg)	Limiting magnitude
$\begin{bmatrix} \sigma_1 & 0 \\ 0 & \sigma_1 \end{bmatrix}$	15.0	[0, 360]	[0, 90]	1 x 1	25.0

Because of this, it is not unheard of to use radar sensors to track GEO RSO as well. This is because radar sensors contain the capability to measure RSO range and range rate as well as the angular components, albeit usually less accurately than EO sensors. [21].

Table 2: Radar Sensor specification modeled

Tx Power (W)	Tx Frequency (Hz)	Slew Rate (deg/s)	Azimuth Mask (deg)	Elevation Mask (deg)	FoV (deg)
2.5e6	1.5e9	3.0	[0, 360]	[0, 90]	1 x 1

Since there are benefits to each of these sensing modalities, both have been simulated as a part of this work.

1.4 Estimation

The Kalman Filter (KF) was developed to combine predicted measurements and the effect they will have on reducing the uncertainty surrounding an state vector [14]. Later, the Unscented Kalman Filter (UKF) was introduced with a modification to the linear Kalman Filter to supports systems consisting of non-linear dynamics, such as RSO state vector propagation [12].

The UKF consists of three main phases: prediction, forecast, and update. The prediction phase is used to propagate estimate state vectors and covariances out to the next time step. In the forecast step, potential observations are predicted in order to determine what their impact would be in reducing uncertainty. This step is critical in relation to the maneuver detection methods mentioned in section 1.2, because the innovations covariance S is calculated here. In equation 7 below, ΔY represents the predicted measurement residuals, W_P is the covariance weight, and R is the measurement noise.

$$S_{k+1|k} = [\Delta Y] W_P [\Delta Y]^T + R_{k+1|k} \tag{7}$$

Finally, during the update step, the innovations (the I in NIS) are calculated. The innovations are the measurement residuals from the predicted and performed observations, as can been seen in equation 8.

$$\varepsilon_{k+1} = \tilde{y}_{k+1} - \hat{y}_{k+1|k} \tag{8}$$

The innovations covariance and innovations in equations 7 and 8 are then passed in the the maneuver detection algorithms directly at the end of the UKF update step. The UKF used in this work was generalized to be able to forecast observations from arbitrary DoF sensors, allowing it to be used for both optical and radar scenarios.

2. SENSOR TASKING

The art of allocating sensing resources to track the ever-growing catalog of Resident Space Objects in an efficient and responsible manor is an intractable problem. Efficient tasking requires long horizon planning, centralized control of all sensors in a network, and extensive computational resources. It is becoming increasing evident that human-in-the-loop sensor tasking and satellite maneuvering decision making is not fast enough for the timelines needed to identify maneuvers and implement collision avoidance procedures [4]. Fully autonomous sensor tasking frameworks have already been demonstrated on hardware systems [8], This work seeks to implement a similar approach where a fully autonomous systems will task sensors and modify the tasking solution in the event of maneuver detection.

In order to reduce the computational complexity of autonomous sensor tasking, the problem is modeled as a Markov Decision Process (MDP). The MDP is a process where all the information necessary to make a decision is stored in the state of the previous time step. Specifically, this simulation implemented Partially Observable Markov Decision Process (POMDP). The POMDP is an extension of the standard MDP where not all of the possible states are known to the decision maker, making the optimal solution harder to predict [22]. This was necessary because both optical and radar sensors cannot observe all the 6 degrees of freedom of satellite state vectors. Since the tasking algorithm is Markovian, it will be able to adjust in real time to changes in the space environment (maneuver detection) and schedule rapid follow-up observations if possible [19].

[1] σ_1 = 1.008 arcsec

2.1 Information Gain

The Shannon Information Gain (SIG) metric, $I_{S_{k+1}}$, was chosen to be the measure of uncertainty reduction for the predicted observations in this study. SIG, shown below in equation 9, is a function of the predicted covariance $\hat{P}_{k+1}$ and estimated covariance $\bar{P}_{k+1}$, calculated in the prediction and update step of the UKF, discussed in section 1.4.

$$I_{S_{k+1}} = \frac{1}{2} log(\frac{det(\hat{P}_{k+1})}{det(\bar{P}_{k+1})}) \tag{9}$$

Previous works have demonstrated SIG as a viable metric for tasking sensors [11, 18, 13]. While Fisher Information Gain (FIG) is the more popular method of POMDP sensor tasking, SIG was chosen because it is computationally cheaper [15].

2.2 Decision Algorithms

Before modifying decision making algorithms specifically for the space domain, the most popular decision algorithms for tasking fall into two main categories, greedy and Munkres (Hungarian Algorithm) [13]. The greedy decision operates on each sensor independently, tasking it to the satellite that provides the highest reward. While this is computationally efficient, it may result in redundant sensor tasking, but more accurately models the real world where most sensing networks do not communicate with each other [13]. The Munkres decision algorithm solves the *NxM* (targets x sensors) decision matrix with a bipartite graph restriction, preventing no two sensors to observe the same RSO at the same time [16, 6]. This functionality can be solved easily by passing the *NxM* reward Matrix (R) into scipy's 'linear_sum _assignment' function [26].

While this decision algorithm maximizing the distribution of sensors, it requires knowledge of the state and tasking control of every sensor being utilized. Since that is often not the case, the greedy decision algorithm is the most practical option because each sensor tasking solution is solved independently.

2.3 Reward Function Normalization

While Shannon Information Gain, discussed in section 2.1, is a viable metric for predicting uncertainty reduction for performing observations on its own, it is not reactive to changes in the environment.

For example, if revisit rate reduction is desired but there is a particular RSO with a highly unstable orbit, SIG might skew the tasking towards that RSO at the cost of ignoring more stable orbits. In that case, it may become necessary to augment the SIG metric with a metric focused on minimizing the time between observations.

The combination different metrics into a single value for the decision algorithms described in section 2.1 is called a Reward Function. In order to operate on metrics with vastly different units and magnitudes, some type of normalization must be done. Research has been on done on metric weights for Reward Functions [19].

In this work, each metric (m_i) in a reward function was normalized from $[0,1]$ where 1 represented the maximum predicted value (M) for each metric at each time step.

$$m_i = m_i / max(M) \tag{10}$$

From there, reward functions are calculated by summing all n metrics prescribed.

$$R_{k+1} = \sum_{i=1}^{N} m_N \tag{11}$$

For example, a reward function with two metrics would have a maximum value of 2.

2.4 Event Based Sensing

A focus of this work was to find a computationally cheap, effective metrics to method to modify the reward functions mentioned in section 2.3 above to skew tasking towards satellites that had observed maneuver detections without over-allocating sensing resources towards that RSO. A simple Boolean maneuver detection metric was created that would be value 1 if a maneuver had been detected at the most recent time step for an RSO and 0 if not. This served as a simple, but effective method to modify tasking rewards in the event maneuver detection. Given this modification, it is

key to ensure that false positive maneuver detections and additional maneuver detections on the correct RSO are kept to a minimum to prevent improper tasking.

For the radar sensor cases, if a maneuver was detected for an RSO, hooks were put in the UKF to instantiate and estimate error reduction method, the two chosen for this work being Initial Orbit Determination (IOD) and Multiple Model Adaptive Estimation (MMAE).

3. SIMULATION DESIGN

For this research, a Cowell's method special Perturbations propagation method was used [25]. Based on the perturbations most prominent at GEO a 4th Order, 4th Degree EGM96 gravitational model as well as Solar and Lunar third body gravitational effects was used for both RSO truth and UKF estimate state vector propagation [17].

An Unscented Kalman Filter with an α of 0.05, β of 2.0 and a κ of 3.0 were used. The estimate state vectors were created with initial estimate error standard deviation (std) of 0.01 km in position and 0.1 m/sec. The process noise for the UKF was chosen to be $1e^{-11}$. The research supporting that choice can be found in section 4.1 of the Results.

The orbital Elements of the RSO used in this study are shown below in table 3, with the variation in initial position created by incrementing true anomaly (ν). These values were chosen so that each RSO would not be purely Geostationary, giving some challenge to the Kalman Filter.

Table 3: RSO Classical Orbital Element Sets

RSO	α (km)	e	i (deg)	Ω (deg)	ω (deg)	ν (deg)
1	42117.7	0.0619	0.0	0.0	291.2	123.51
2	42117.7	0.0619	0.0	0.0	291.2	124.51
3	42117.7	0.0619	0.0	0.0	291.2	125.51
4	42117.7	0.0619	0.0	0.0	291.2	126.51
5	42117.7	0.0619	0.0	0.0	291.2	127.51

A single sensor was used to generate the measurements in this work. The sensor was placed in a location and scenario time chosen such that every RSO was visible for the entirety of simulation time.

Table 4: Sensor Location

Sensor	Latitude (deg)	Longitude (deg)	Altitude (km)
1	5.7295	63.0254	1.0

A 5-minute time step was used for the simulations in this work to ensure that slew time was not a limiting factor in RSO observability and to account for image processing and object detection. Within this scenario, one 5m/s in-track impulsive maneuver was injected for RSO 3 without informing the Kalman Filter and sequential tasking engine. The value of 5m/s was chosen based on the standard magnitudes of GEO station keeping maneuvers discussed in section 1.1 and the direction represents an eastward station keeping maneuver. Previous work has investigated optical maneuver detection and found that station keeping maneuver magnitudes were sufficient for testing [20].

For each of the maneuver detection methods, a threshold of 99.9% confidence value was set. For the Sliding NIS, a time step window of 4 was chosen. For the Fading Memory NIS, the δ value was set to 0.8.

4. RESULTS

4.1 Maneuver detection and filter tuning Results

Before focusing on which maneuver detection method performed best for each sensor type, it was import to tune the Process Noise value (Q) on the UKF so that the maneuver could be detected.

Table 5: Optical Sensor Maneuver detection compared to Process Noise

Process Noise:	1e-14	1e-13	1e-12	1e-11	1e-10	1e-9	1e-8	1e-7	1e-6
Standard NIS	True*	True	True	True	True	False	False	False	False
Sliding NIS	True	True	True	True	True	False	False	False	False
Fading Memory NIS	True	True	True	True	True	False	False	False	False

The asterisk in Tables 5 and 6 denote that there were false positive maneuver detections for those process noise values. Notably, for the optical sensor case tuning the process noise to the very sensitive $1e^{-14}$ value resulted in False positive maneuver detections for the Standard NIS only.

Table 6: Radar Sensor Maneuver detection compared to Process Noise

Process Noise:	1e-14	1e-13	1e-12	1e-11	1e-10	1e-9	1e-8	1e-7	1e-6
Standard NIS	True*	True*	True	True	True	False	False	False	False
Sliding NIS	True*	True*	True	True	True	False	False	False	False
Fading Memory NIS	True*	True*	True	True	True	False	False	False	False

For the radar sensing evaluation, all three maneuver detection methods record false positive detections if the process noise was $1e^{-13}$ or smaller.

4.2 Sensor Tasking Reward Function Modification

For the control cases in this study, RSO were propagated without any maneuvers. The POMDP cycles through the RSO for routine observations at a repeated cadence, as shown for the optical and radar cases in Fig. 2. The boxes in Fig. 2 below represent the times at which a particular RSO was observed, the different colors and y axis values denoting a specific RSO.

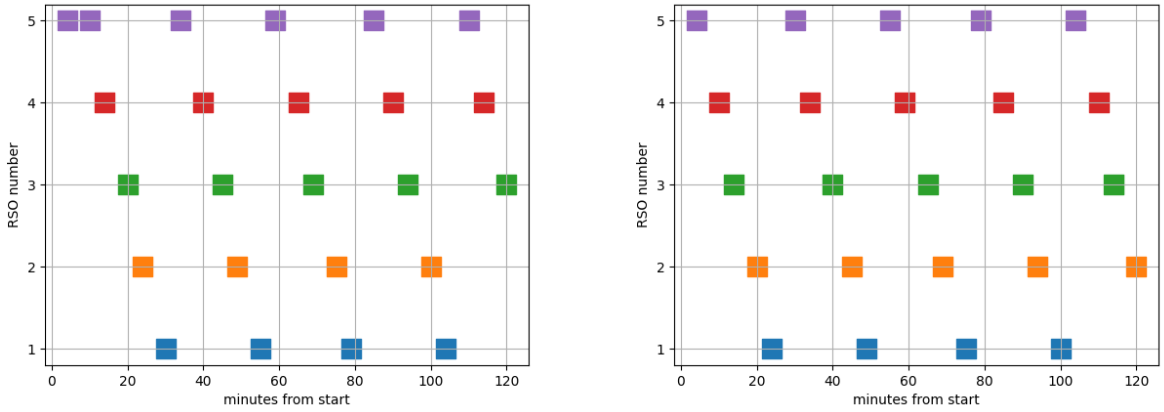

(a) Optical tasking without maneuvers

(b) Radar tasking without maneuvers

Fig. 2: Sensor Schedule without injected maneuvers

In Fig. 4 below, a 5m/s impulsive maneuver was inject 50 minutes into the scenario for RSO 3.

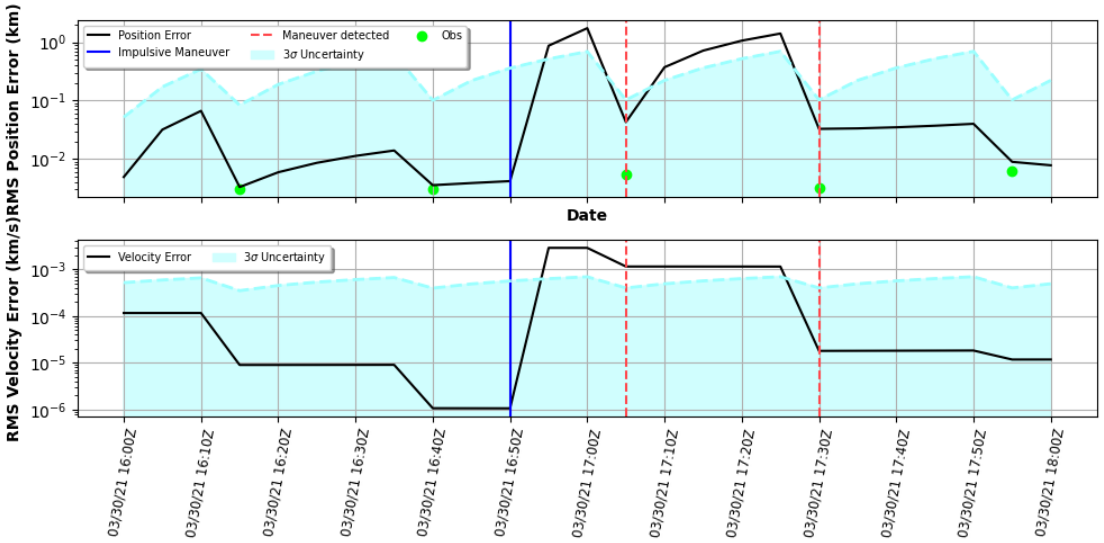

Fig. 3: Estimate Error for Maneuvered RSO 3

As can be seen in Fig. 3 above, the estimate position and velocity error falls outside of the 3σ uncertainty bounds and a maneuver is detected. However, the sensor tasking solution remains the same as the control case above in Fig. 2. This shows clear proof that the sensor tasking solution that works well without maneuvers needs additional event-based sensing modifications in order to prioritize observing an RSO it knows to have maneuvered.

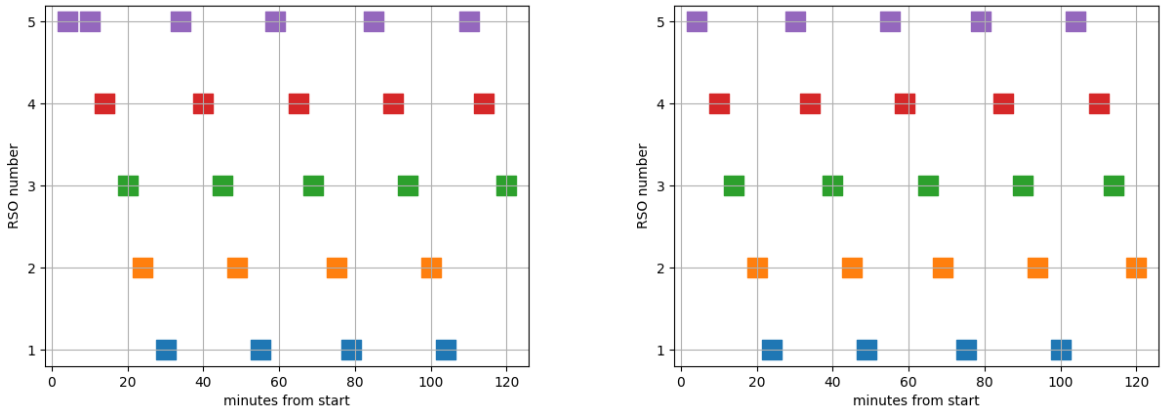

(a) Optical tasking with maneuver

(b) Radar tasking with maneuver

Fig. 4: Sensor Schedule with an injected maneuver

In contrast to the previous tasking solutions, Fig. 5 below shows a clear break in the default tasking modalities from the control for the RSO that has maneuvered.

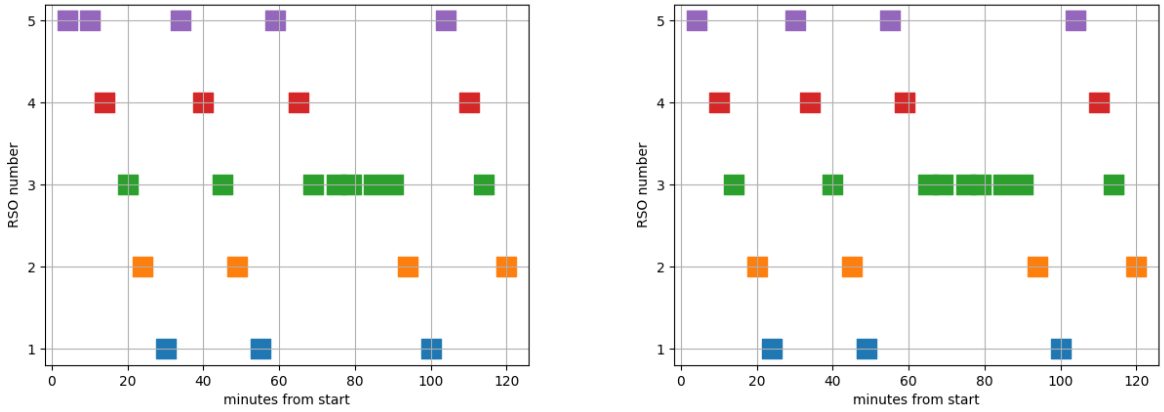

(a) Optical tasking

(b) Radar tasking

Fig. 5: Sensor Schedule with maneuver detection tasking modification

Most promising is the fact that after an adequate number of observations have been made of the maneuvered RSO, the tasking solution returns back to the nominal state of cycling through each target. This clearly shows that standard Shannon Information Gain sensor tasking modified with a maneuver detection metric can prioritize observations for a recently maneuvered RSO without permanently skewing tasking towards that satellite.

4.3 Orbit Determination

For the radar sensing cases, IOD and MMAE orbit determination algorithms were implemented to attempt to reduce the number of observations after maneuver detection. In Fig. 5 above, both the optical and radar tasking algorithms made six consecutive observations of the maneuvered RSO after maneuver detection.

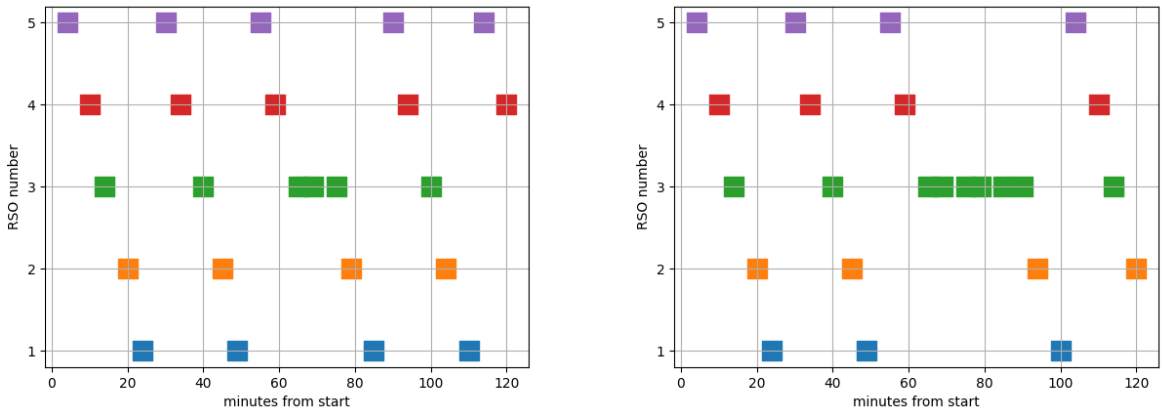

(a) Radar IOD tasking solution

(b) Radar MMAE tasking solution

Fig. 6: Radar Schedule with orbit determination

While the MMAE results in Fig. 6 also produced six consecutive observations of the post-maneuver RSO, the IOD scenario utilizing Lambert's Universal Variables orbit determination was able to return to cycling through targets after only three consecutive observations.

5. CONCLUSIONS

In this work, it was show that Kalman Filter tuning was the dominant factor in deciding whether or not accurate maneuver detection was achieved, regardless of the maneuver detection method used for both optical and radar sensors. Each of the three methods investigated performed accurately if the process noise was set between $1e^{-10}$ and $1e^{-12}$. Improper tuning of the standard normalized innovations squared resulted in false positive maneuver detections for the optical sensor that were not present in the sliding NIS or the fading memory NIS. For the radar sensor, poor tuning resulted in false positive maneuver detections for all three algorithms.

Event-based sensing modification to the tasking reward function proved crucial in allocating resources to recently maneuvered RSO without skewing the tasking solution permanently toward that target. Adding a normalized maneuver detection metric to the Shannon Information Gain metric produced a reward function adaptive to change in the short term but flexible enough to return to its nominal state.

Initial Orbit Determination (IOD) and Multiple Model Adaptive Estimation (MMAE) methods in the event of maneuver detection were implemented to attempt to reduce the amount of observations of the maneuvered RSO post-maneuver. MMAE did not generate any improvement but IOD was able to reduce the number of consecutive observations from six to three.

In order to continue to test the event based response ability of EO sensors, angles-only IOD methods such as Gooding's method will be implemented. Unlike Lambert Orbit Determination, Gooding's Method generated the 6 DoF necessary for OD by using 3 angles-only measurements. It does this by iterating to solve for the range of the RSO at the first and third observation time, allowing Lambert's method to be solved from that information.

REFERENCES

[1] Phillip D Anz-Meador. Orbital debris quarterly news. *Orbital Debris Quarterly News (ODQN)*, 24(JSC-E-DAA-TN77633), 2020.

[2] Yaakov Bar-Shalom, X Rong Li, and Thiagalingam Kirubarajan. *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2001.

[3] Han Cai, Yang Yang, Steve Gehly, Changyong He, and Moriba Jah. Sensor tasking for search and catalog maintenance of geosynchronous space objects. *Acta Astronautica*, 2020.

[4] R Cardin, Dustin Burchett, and H G Reed. Snare (sensor network autonomous resilient extensible): Decentralized sensor tasking improves sda tactical relevance, 2021.

[5] Chia-Chun (George) Chao and Felix Hoots. *Applied orbit perturbation and maintenance*. The Aerospace Press, 2 edition, 2018.

[6] David F Crouse. On implementing 2d rectangular assignment algorithms. *IEEE Transactions on Aerospace and Electronic Systems*, 52(4):1679–1696, 2016.

[7] ESA. Space environment statistics, 2022.

[8] J Fletcher, D Archambeault, J Schmidt, R Peterson, P Sydney, and S Hunt. The dynamic optical telescope system: Collaborative autonomous sensing for space domain awareness. *JDRE*, 4:10–18, 2021.

[9] Jeff Foust. Spacex passes 2,000 starlink satellites launched, 2022.

[10] Norbert Frischauf. *Outer Space in Society, Politics and Law - Satellite Telecommunication*. Springer, 09 2011.

[11] Ashton Harvey, Kathryn Laskey, and Kuo-Chu Chang. Bridging heuristic and deep learning approaches to sensor tasking. In *2021 IEEE 24th International Conference on Information Fusion (FUSION)*, pages 1–8. IEEE, 2021.

[12] Simon J Julier, Jeffrey K Uhlmann, and Hugh F Durrant-Whyte. A new approach for filtering nonlinear systems. *Proceedings of 1995 American Control Conference-ACC'95*, 3:1628–1632, 1995.

[13] Jonathan Kadan, Dylan Thomas, Cameron Harris, Kevin Schroeder, and Jonathan Black. Parametric analysis of an autonomous sensor tasking engine for spacecraft tracking. In *AIAA Scitech 2021 Forum*, page 1397, 2021.

[14] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *The Journal of Fluids Engineering*, 1960.

[15] Tian Kangsheng and Zhu Guangxi. Sensor management based on fisher information gain. *Journal of Systems Engineering and Electronics*, 17(3):531–534, 2006.

[16] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.

[17] Oliver Montenbruck, Eberhard Gill, and Fh Lutze. Satellite orbits: models, methods, and applications. *Appl. Mech. Rev.*, 55:B27–B28, 2002.
[18] Kevin M Nastasi and Jonathan T Black. Adaptively tracking maneuvering spacecraft with a globally distributed, diversely populated surveillance network. *Journal of Guidance, Control, and Dynamics*, 42(5):1033–1048, 2019.
[19] Kevin Michael Nastasi. *Autonomous and Responsive Surveillance Network Management for Adaptive Space Situational Awareness*. PhD thesis, Virginia Tech, 2018.
[20] Alejandro Pastor, Guillermo Escribano, and Diego Escobar. Satellite maneuver detection with optical survey observations. *21st Advanced Maui Optical and Space Surveillance Technologies Conference*, 2020.
[21] Lorenzo Porcelli. Satellite manoeuvre detection and estimation with radar observations. *Polytechnic University of Milan*, 2022.
[22] Martin L Puterman. Markov decision processes. *Handbooks in operations research and management science*, 2:331–434, 1990.
[23] A Serebryanskiy, Ch Akniyazov, B Demchenko, A Komarov, Ch Omarov, I Reva, M Krugov, and V Voropaev. Statistical analysis of object congestion in the geostationary region. *Acta Astronautica*, 182:424–431, 2021.
[24] Mark A Skinner, Thomas M Kelecy, Stephen A Gregory, Joseph P Toth, Dennis Liang, Dean Yamanaka, Stan Kent, Rodney Tjoelker, Dragos Margineantu, Audrey L Allison, et al. Commercial space situational awareness: an investigation of ground-based ssa concepts to support commercial geo satellite operators. *Proc. AMOS*, 2013.
[25] Eduard Stiefel and DG Bettis. Stabilization of cowell's method. *Numerische Mathematik*, 13:154–175, 1969.
[26] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
[27] P Wegener, J Bendisch, H Krag, M Oswald, and S Stabroth. Population evolution in the geo vicinity. *Advances in Space Research*, 34:1171–1176, 2004.