

# Relative Estimation in the Cislunar Regime using Optical Sensors

**Jesse A. Greaves**

*University of Colorado Boulder*

**Daniel J. Scheeres**

*University of Colorado Boulder*

## ABSTRACT

Traditional concerns regarding space domain awareness are exacerbated in cislunar space due to highly nonlinear dynamics, limited viewing geometries, and increasingly overloaded Earth-based resources. These concerns motivate the investigation of space-based navigation and observation capabilities which lend themselves to autonomy. This paper expands upon previous results to demonstrate that relative optical measurements provide simultaneous navigation of a primary spacecraft and observation of a secondary spacecraft. First, linear observability analysis proves that relative optical measurements fully estimate the dual spacecraft system over various orbits. Subsequently, the Optimal Control Based Estimator is modified with the Unscented Transformation to account for the strongly nonlinear dynamics and then leveraged to generate a state solution along with an estimated control policy. Finally, the estimated control policy is employed to identify and classify maneuvers. Collectively, this methodology directly enables a chaser craft to navigate to a final docking phase autonomously and provides a basis for further research into autonomous operations in cislunar space.

## 1. INTRODUCTION

Cislunar space is a primary target for development due to the region's suitability as a proving ground and staging point for deep space exploration. NASA's Artemis program is an exemplary model for cislunar development, but NASA is not alone in its vision for inhabiting this expanse [29, 23]. A significant majority of national space agencies, as well as space-focused commercial companies, have expressed interest in sending spacecraft on cooperative and independent ventures to this territory [27, 28, 12, 21]. It is anticipated that the combined efforts of the international and commercial space community will lead to a myriad of spacecraft traversing and occupying cislunar space. Such an environment necessitates the design of cislunar Space Domain Awareness (SDA) capabilities to ensure all vehicles' collective safety and long-term multi-player cooperation in the region.

Currently, cislunar estimation relies on Earth-based resources, which leads to various observational challenges. Earth-based radar, such as the Deep Space Network (DSN), is already oversubscribed, and equitably scheduling DSN time for all future vehicles will be prohibitive. Ground-based optics must avoid exclusion zones around the Moon because excessive light can damage the powerful sensors required to detect objects at cislunar distances. Additionally, all Earth-based observers have a limited viewing geometry of cislunar objects, which results in sparse information and potentially long periods of occultation. Detailed issues concerning cislunar tracking are covered in [18]. To overcome these factors, as well as many other issues, the exploration of new observation platforms is essential. In particular, cislunar observation platforms demonstrate promise to reduce Earth-based resource load, enhance observability, and enable novel missions.

Previous works have studied the benefits of utilizing various cislunar observation platforms for cislunar orbit determination [31, 9, 11, 2, 15, 8]. Generally, sensors in these studies can be divided into passive vs. active and electro-optical vs. radio-frequency categories. Observer location also plays a critical role given the vast array of periodic/quasi-periodic orbit families and the potential to utilize the lunar surface itself. Space-based optical platforms are an appealing solution since they leverage existing commercial technology and provide equivalent accuracy for cooperative and non-cooperative systems. Additionally, studies have shown that optical measurements are a viable option for autonomous cislunar navigation [3, 5, 17, 14]. The combined advantages of space-based optical sensors lead to the possibility that optical measurements can provide simultaneous autonomous navigation and observation. Thus, this

paper seeks to examine relative optical measurements to estimate a primary and secondary spacecraft over various orbits.

It is essential that estimation algorithms consider maneuver detection and higher-order filtering to ensure proper tracking in the cislunar environment. Maneuver detection is necessary because many orbits in the cislunar region require frequent station keeping maneuvers due to the strongly nonlinear dynamics [7, 24]. These dynamics also cause filtering issues because the higher-order moments are non-negligible and can cause linear systems to diverge [19, 7, 25]. Numerous methods accomplish maneuver detection and higher-order filtering, but they often require large computational requirements. The Optimal Control Based Estimator (OCBE) and Unscented Transformation (UT) address each concern and require minimal computation cost compared to alternative methods [22, 20].

This paper begins by utilizing linear observability to prove that relative optical measurements lead to a fully observable system. The applied method is a modification of LiAISON results which previously demonstrated that relative radiometric measurements are sufficient to perform dual navigation [15]. Then, a modified version of the OCBE with Unscented uncertainty propagation is adapted to provide filtering solutions and an estimated control policy. The UT is applied to account for the strongly nonlinear dynamics and leads to more consistent results. Finally, the control policy from the estimator is input into a classification algorithm to identify and label maneuvers. The classification methods separates the control policy into four categories of maneuvers: none, stable manifold, unstable manifold, and generic. These maneuver categories correspond to station keeping, exponential orbit departures, and generalized maneuvers, respectively. In combination, this paper extends previous work to account for uncertainty in the observer state, improves filter performance, and adds maneuver classification [13]. By successfully demonstrating dual spacecraft estimation utilizing only relative optical measurements, this work seeks to further autonomous operations in cislunar space.

## 2. MODELING DEFINITIONS

### 2.1 Dynamics

Two spacecraft are instantiated in the Circular Restricted 3 Body Problem (CR3BP) to simulate relative estimation in this region. Thus the state is composed of the position ( $\mathbf{r}$ ) and velocity ( $\dot{\mathbf{r}}$ ) of both spacecraft as described by Equation 1. The CR3BP is commonly derived in a frame that rotates with the primary and secondary body such that the bodies are constant along the first axis. Figure 1 contains a diagram of the CR3BP in the rotating frame. Equation 3 gives the dimensional CR3BP equation of motion in this frame. The dimensional form is used for filtering because all simulation parameters are in dimensional units.

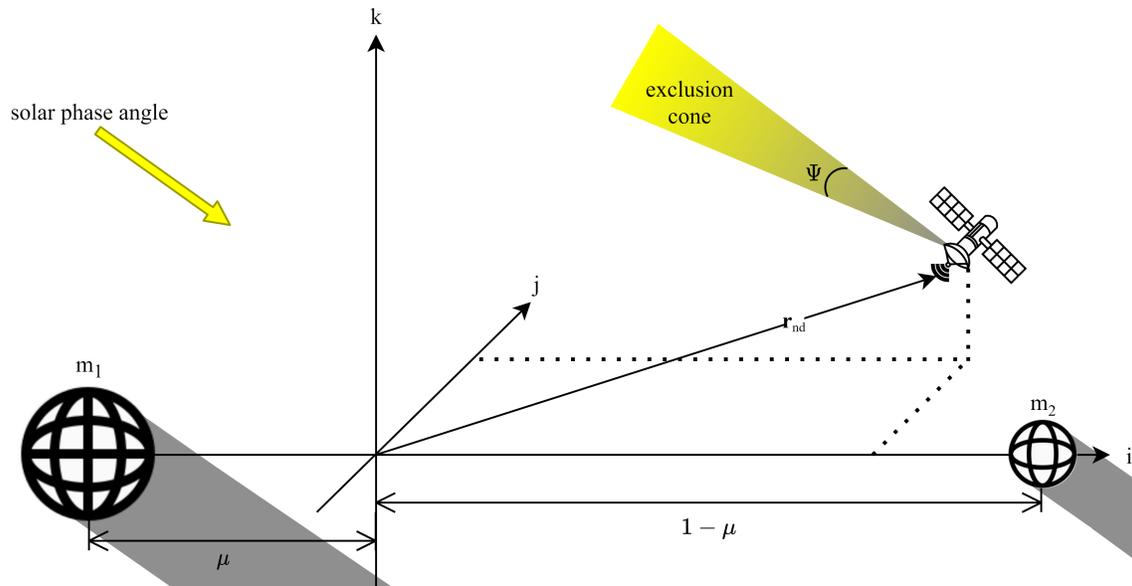


Fig. 1: Diagram of CR3BP with sun shadows, and viewing exclusion zone.

$$\mathbf{x} = [\mathbf{r}_{cha}, \dot{\mathbf{r}}_{cha}, \mathbf{r}_{tar}, \dot{\mathbf{r}}_{tar}]^T \quad (1)$$

$$\mathbf{n} = \left[ 0 \quad 0 \quad \sqrt{\frac{G(m_1+m_2)}{\|\mathbf{r}_2-\mathbf{r}_1\|^3}} \right]^T \quad (2)$$

$$\ddot{\mathbf{r}} = -2\mathbf{n} \times \dot{\mathbf{r}} - \mathbf{n} \times (\mathbf{n} \times \mathbf{r}) - \frac{Gm_1}{\|\mathbf{r}-\mathbf{r}_1\|^3}(\mathbf{r}-\mathbf{r}_1) - \frac{Gm_2}{\|\mathbf{r}-\mathbf{r}_2\|^3}(\mathbf{r}-\mathbf{r}_2) \quad (3)$$

The non-dimensional equation of motion normalizes position by the distance between massive bodies, time by the period of the system, and mass by the total mass of the system. Equation 5 gives the non-dimensional equation of motion. Non-dimensionalizing the system alleviates numerical scaling issues, reduces the parameter space, and allows for easy comparisons between systems with similar parameters. Reference [30] contains a brief but informative overview of the CR3BP, including dynamics, orbits, and manifolds.

$$\mu = \frac{m_2}{m_1 + m_2} \quad (4)$$

$$\ddot{\mathbf{r}}_{nd} = -2\hat{\mathbf{k}} \times \dot{\mathbf{r}}_{nd} - \hat{\mathbf{k}} \times (\hat{\mathbf{k}} \times \mathbf{r}_{nd}) - \frac{1-\mu}{\|\mathbf{r}_{nd}-\mathbf{r}_{nd,1}\|^3}(\mathbf{r}_{nd}-\mathbf{r}_{nd,1}) - \frac{\mu}{\|\mathbf{r}_{nd}-\mathbf{r}_{nd,2}\|^3}(\mathbf{r}_{nd}-\mathbf{r}_{nd,2}) \quad (5)$$

Several Periodic Orbit (PO) families were generated to explore the effects of viewing geometry on tracking solutions. Included families are northern L1 halos (NL1), planar L2 halos (PL2), southern L2 halos (SL2), southern L2 halo Quasi-Periodic Orbits (QPO), and Distant Retrograde Orbits (DRO). This paper focuses on one representative orbit from each family to limit chaser and target combinations. Figure 2 plots the trajectories of the selected orbits. The orbital period and radius of periapsis for each orbit is shown in Table 1. Note that the SL2 orbit has a 9:2 resonance with the lunar synodic cycle, which is the currently selected orbit for the lunar gateway.

The QPO is of particular interest because it has naturally bounded relative motion about a PO with the same base frequency. The resulting relative trajectory produces diverse viewing geometries between the two orbits, which provides observational benefits. Additionally, this yields an ideal staging point for rendezvous and proximity operations due to the zero-cost bounded motion. Figure 3 presents the natural relative trajectory of the QPO about the PO over 15 orbit periods.

Maneuvers are implemented as impulsive  $\Delta\mathbf{v}$ 's with execution errors given by the Gates model [32]. There are three types of maneuvers considered: station keeping, exponential departure, and generalized. Station keeping is necessary for spacecraft to maintain their desired orbits, and optimal  $\Delta\mathbf{v}$ 's on halo orbits are known to approximately lie along the stable manifold ( $\hat{\mathbf{e}}_s$ ) [10, 1]. The cost of station keeping policies varies depending on navigational accuracy, but typical impulses are on the order of tens of mm/s to several m/s [7, 24]. Exponential departures are maneuvers along the unstable manifold ( $\hat{\mathbf{e}}_u$ ) that lead to low-cost orbit transfers. Finally, generalized maneuvers cover the control space and thus may not indicate long-term behavior.

Maneuvers occur at apoapsis, since it is an optimal control placement [7, 24], and are a function of their class and the applied station keeping policy. The station keeping policy dictates impulse magnitudes for all three maneuver types because it sets the requirements for spacecraft thruster performance. There are two station keeping policies, noisy and quiet, each with their own statistics. The noisy policy has a mean of 1 m/s with a 0.3 m/s standard deviation. The quiet policy has a mean of 40 mm/s with a 10 mm/s standard deviation.

Table 1: Reference orbit details.

Parameter	Northern L1	Planar L2	Southern L2	Southern L2 QPO	DRO
orbit period (days)	11.9	14.6	6.55	6.55	14.9
radius of periapsis (km)	51111	62902	3212	2346	74309

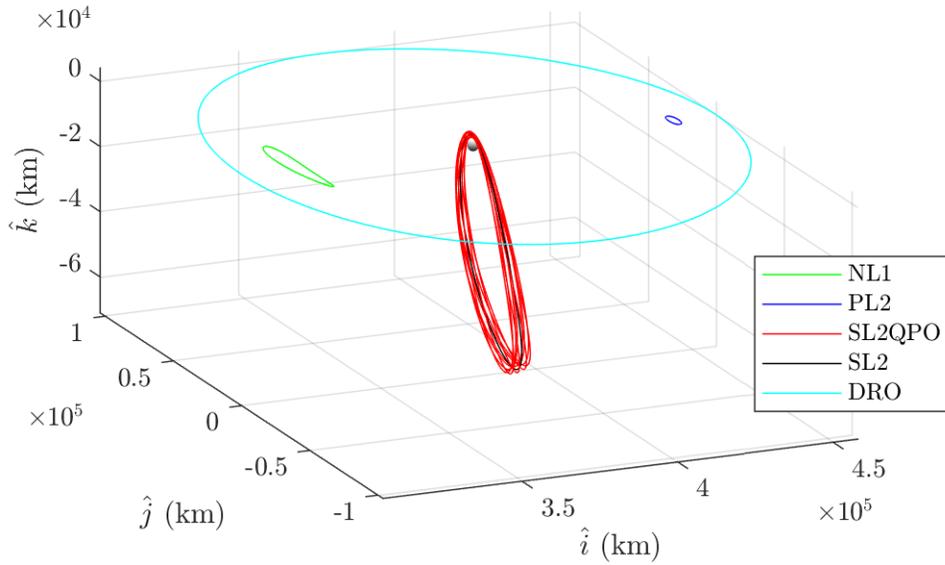


Fig. 2: Reference orbits examined in this paper.

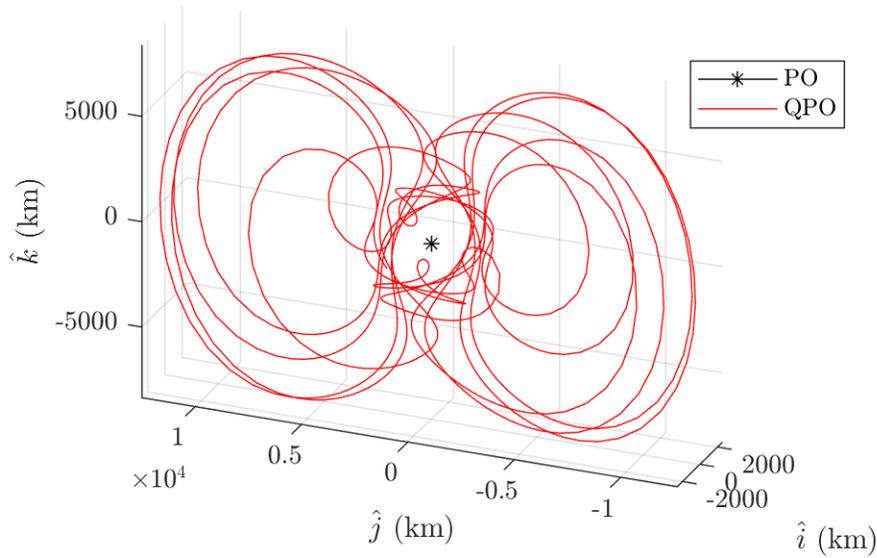


Fig. 3: Relative motion of the southern L2 QPO centered about the PO.

All maneuvers are affected by Gates execution errors. These errors are characterized as zero-mean Gaussians that are oriented in a frame aligned with the maneuver. Equation 6 describes the Gates covariance, which is composed of both fixed and proportional components. The error is rotated from the  $\Delta \mathbf{v}$  frame to the CR3BP via the rotation matrix  $[\mathbf{CV}]$ . Equations 7-9 mathematically define the maneuvers given the previous assumptions. In these equations,  $\mathcal{N}$  is a Gaussian distribution, and  $\mathcal{U}$  is a uniform distribution. Table 2 summarizes the statistics associated with the maneuver model. Note that Cassini’s pre-flight Gates values are adopted as the noisy policy parameters. The quiet policy parameters are a mixture of the 2000 and 2006 models [32]. When applied to the given policies, proportional errors prevail in the noisy case, and fixed errors dominate in the quiet case.

Table 2: Maneuver statistics.

Policy	$\mu_{sk}$	$\sigma_{sk}$	$\sigma_{g1}$	$\sigma_{g2}$	$\sigma_{g3}$	$\sigma_{g4}$
noisy	1 m/s	0.3 m/s	10 mm/s	0.35	17.5 mm/s	10 mrad
quiet	40 mm/s	10 mm/s	6.5 mm/s	0.2	4.5 mm/s	3.5 mrad

$$\mathbf{P}_{gates} = \begin{bmatrix} \sigma_{g1}^2 + \|\Delta\mathbf{v}\|\sigma_{g2}^2 & 0 & 0 \\ 0 & \sigma_{g3}^2 + \|\Delta\mathbf{v}\|\sigma_{g4}^2 & 0 \\ 0 & 0 & \sigma_{g3}^2 + \|\Delta\mathbf{v}\|\sigma_{g4}^2 \end{bmatrix} \quad (6)$$

$$\Delta\mathbf{v}_{SK} = \pm\mathcal{N}(\mu_{sk}, \sigma_{sk})\hat{\mathbf{e}}_s + [\mathbf{CV}]\mathcal{N}(0, \mathbf{P}_{gates}) \quad (7)$$

$$\Delta\mathbf{v}_{ED} = \pm\mathcal{N}(\mu_{sk}, \sigma_{sk})\hat{\mathbf{e}}_u + [\mathbf{CV}]\mathcal{N}(0, \mathbf{P}_{gates}) \quad (8)$$

$$\Delta\mathbf{v}_{Gen} = \mathcal{N}(\mu_{sk}, \sigma_{sk})Sph2Cart(\mathcal{U}(0, 2\pi), \mathcal{U}(-\pi/2, \pi/2)) + [\mathbf{CV}]\mathcal{N}(0, \mathbf{P}_{gates}) \quad (9)$$

## 2.2 Measurements

Measurements are taken between the two spacecraft and processed by a filter to estimate both vehicles' state. The first craft, denoted as the chaser, is treated as the controlled vehicle and is assumed to take the measurements. The second craft, designated the target, is nominally non-cooperative in this setting. Measurements of range, azimuth, and elevation are available to gather information and are given by Equations 11-13 respectively. Note that a single optical measurement comprises azimuth and elevation elements, and that range is separately used to verify and compare LiAISON results.

$$\mathbf{r}_{rel} = \mathbf{r}_{tar} - \mathbf{r}_{cha} \quad (10)$$

$$\rho = \|\mathbf{r}_{rel}\| \quad (11)$$

$$\theta = \tan^{-1}\left(\frac{\mathbf{r}_{rel} \cdot \hat{\mathbf{j}}}{\mathbf{r}_{rel} \cdot \hat{\mathbf{i}}}\right) \quad (12)$$

$$\phi = \sin^{-1}\left(\frac{\mathbf{r}_{rel} \cdot \hat{\mathbf{k}}}{\rho}\right) \quad (13)$$

Range measurements are corrupted by additive Gaussian white noise, while optical measurements incorporate partially resolved object uncertainty effects. This model is derived by setting the angular uncertainty to an analog of the root-mean-square combination of a minimum angular uncertainty ( $\sigma_{\theta/\phi, min}$ ) and a reference resolution uncertainty ( $D$ ). Equation 15 defines the angular uncertainty model. The current reference resolution uncertainty is 4 m, which is approximately the size of lunar gateway modules [6, 4]. Range measurements are taken every 6 minutes with 1 m noise [16], and angular measurements are taken every 2 hours with a minimum noise of 10 micro-rad [13, 3].

$$\rho^2 \sigma_{\theta/\phi}^2 = \rho^2 \sigma_{\theta/\phi, min}^2 + D^2 \quad (14)$$

$$\sigma_{\theta/\phi} = \sqrt{\sigma_{\theta/\phi, min}^2 + \frac{D^2}{\rho^2}} \quad (15)$$

All measurements are limited by line-of-sight, and any occultation will cause a missed observation. In addition, optical measurements are subject to lighting constraints such as solar occultation and a solar exclusion region ( $\psi$ ) of 30 deg [3]. Figure 1 qualitatively depicts these restrictions, and table 3 contains a summary of the measurement parameters.

Table 3: Measurement parameters.

$\sigma_p$	$\sigma_{\theta,min}$	$\sigma_{\phi,min}$	$D$	$\psi$
1 m	10 $\mu$ rad	10 $\mu$ rad	4 m	30 deg

### 3. OBSERVABILITY

Observability is often assessed linearly via the condition number of the information matrix, and this matrix is defined by Equation 18. Note that  $\mathbf{y}_i$  is the measurement vector at time  $t_i$ , and  $\Phi(t_i, t_0)$  is the State Transition Matrix (STM) from time  $t_0$  to  $t_i$ . If condition number of  $\mathbf{\Lambda}$  is over a given limit, often  $1e16$ , then the matrix is said to be singular, and the system is considered unobservable [15]. This metric stems from the fact that when the information matrix is singular, it cannot be inverted to uniquely determined a linearized state. Thus, the condition number is used to measure how close a matrix is to singular.

$$\mathbf{H}_i = \frac{\partial \mathbf{y}_i}{\partial \mathbf{x}} \Phi(t_i, t_0) \quad (16)$$

$$\mathbf{W}_i = \begin{bmatrix} \sigma_{y,1}^2 & & & \\ & \ddots & & \\ & & \ddots & \\ & & & \sigma_{y,n}^2 \end{bmatrix}^{-1} \quad (17)$$

$$\mathbf{\Lambda} = \sum_{i=1}^{N_m} \mathbf{H}_i^T \mathbf{W}_i \mathbf{H}_i \quad (18)$$

The condition number is defined as the maximum to the minimum singular value of a matrix. Thus, while the condition number of a singular matrix is infinity, the condition number inherently measures the relative scaling of the matrix and not necessarily singularity. Relative scaling is vital to linear estimation because numerical computation can not produce infinite accuracy, and the condition number bounds the resulting accuracy from calculations with a matrix. This loss of accuracy due to scaling was one of the motivations behind the Square-Root Information (SRI) filter. However, while relative scaling is essential to acknowledge, using the condition number leads to various ambiguities when assessing observability.

A simulation with the chaser on the PL2 orbit taking range measurements of the target on the SL2 orbit is generated to demonstrate ambiguities with the condition number. This situation mimics a LiAISION case which was proven to be observable [16]. The information matrix is computed in a dimensional, non-dimensional, and dimensional square-root setting. The corresponding condition numbers were:  $2.3e20$ ,  $3.9e10$ ,  $1.5e10$ .

These results indicate that even though the dynamics and measurements are equivalent, the non-dimensional and square-root information systems are observable, but the standard dimensional is not. This ambiguity arises from the condition number inherently measuring the relative scaling of the information matrix and not the absolute information of the system. Thus, while the condition number metric has tangible impacts on linearized filters, it provides conflicting results due to the nature of the metric and may not truly identify when a system becomes observable.

The standard definition of linear observability is adopted to eliminate these ambiguities, reduce scaling errors, and allow null space analysis. By the standard definition, a system is observable if the observability Gramian defined in Equation 19 has full column rank. This metric is inherently less sensitive to numerical scaling issues since this matrix is similar to a square-root of the information matrix. Additionally, finding the unobservable subspace is trivial because it is simply the null space of  $\mathbf{O}$ . Finally, suppose the chosen filter is nonlinear. In that case, this metric is more appropriate since the filter does not need to invert the information matrix and will not suffer numerical scaling issues associated with the condition number. When this standard definition test is applied to the same simulation used to assess the condition number test, the observability Gramian returned full rank for every case.

$$\mathbf{O} = \begin{bmatrix} \mathbf{H}_0 \\ \vdots \\ \mathbf{H}_{N_m} \end{bmatrix} \quad (19)$$

The standard linear observability test is applied to chasers and targets on every combination of reference orbits, excluding the condition where both crafts inhabit the same orbit. First, this approach is compared to LiAISON results by assessing range measurements. Table 4 displays the time to observability results. Note that swapping target and chaser orbits provide equivalent dynamics and measurements, making the table symmetric. The observability Gramian results agree with LiAISON that relative range measurements lead to observable systems for non-coplanar orbits [16]. It is worth noting that the coplanar PL2 and DRO unobservable subspace was the  $\hat{\mathbf{k}}$  position and velocity for both crafts. This null space agrees with intuition because the lack of out-of-plane motion leads to no information in the  $\hat{\mathbf{k}}$  direction.

Table 4: Hours to reach observability given range measurements.

Target Orbit	NL1	-				
	PL2	61.7	-			
	SL2	32.9	38.1	-		
	QPO	34.5	37.6	36.7	-	
	DRO	48.5	N/O	30.1	29.1	-
		NL1	PL2	SL2	QPO	DRO
	Chaser Orbit					

The observability Gramian test was applied to the same orbit combinations again, but with optical measurements. Table 5 details the hours required to reach observability. The optical measurement results are promising for several reasons. First, optical measurements lead to fully observable systems. Second, observability was achieved faster in all cases, even with a measurement cadence that is 20 times longer. Lastly, the coplanar PL2 to DRO case is observable. Collectively, these results indicate that optical measurements have a significant potential for simultaneous autonomous navigation and observation and are more robust than range measurements.

Table 5: Hours to reach observability given optical measurements.

Target Orbit	NL1	-				
	PL2	16	-			
	SL2	10	10	-		
	QPO	10	10	10	-	
	DRO	12	14	10	10	-
		NL1	PL2	SL2	QPO	DRO
	Chaser Orbit					

#### 4. STATE ESTIMATION

Strongly nonlinear dynamics hinder estimation in the CR3BP and cause filter failure under various circumstances [7, 25, 13]. Recent studies have shown this is likely due to the higher-order moments that significantly impact uncertainty propagation [19]. Thus there is a need for higher-order, or nonlinear, filters in cislunar estimation. There are many higher-order estimation processes, but the UT is often desirable because it requires little additional computation cost for third-order accuracy [20, 33].

The frequent station keeping events required to ensure stability in the cislunar region further inhibit observation in this domain. Therefore, estimation algorithms must account for station keeping through some form of maneuver detection. Again, there are many methods of maneuver detection, but the ballistic OCBE has proven to be an effective filter and maneuver identification algorithm that minimally impacts computation cost [22, 13]. The OCBE achieves this

by estimating costates along with the state in a linearized filter. However, as mentioned previously, linearized filters struggle in the cislunar regime. Thus, the OCBE is modified to improve its success.

The intrinsic single-step smoothing property of the OCBE facilitates easy modification of its algorithm. This property establishes that the costate estimate is a function of the current state and smoothed prior state. Thus, the time and measurement state updates can directly apply the UT transformation to produce a current estimate. Then a single-step smoother is applied to the current estimate after the measurement update. Finally, the costate and control estimates are calculated from the current and smoothed states. The pseudo-code for the UT modified ballistic OCBE is given in Algorithm 1.

**Result:** State, costate, and control estimates.

initialization;

```

while new measurement do
     $\mathbf{y}_{i+1}$  = Read Measurement;
    if detection then
         $q$  = Adaptive ADU Algorithm( $\mathbf{x}_{i|i}$ ,  $\mathbf{y}_{i+1}$ ,  $q$ );
        outlier count = 0;
        detection = false;
    else
         $q$  = Desired Noise Function;
    end
     $\mathbf{x}_{i+1|i}$ ,  $\Phi(t_{i+1}, t_i)$  = UT Time Update( $t_i$ ,  $t_{i+1}$ ,  $\mathbf{x}_{i|i}$ ,  $q$ );
    outlier = OCBE Distance Metric( $\mathbf{x}_{i+1|i}$ ,  $\mathbf{y}_{i+1}$ );
    if outlier then
        outlier count += 1;
        detection = Automated Detection(outlier count);
        if detection then
            return to measurement at first outlier;
        end
    else
        outlier count = 0;
    end
     $\mathbf{x}_{i+1|i+1}$  = UT Measurement Update( $\mathbf{x}_{i+1|i}$ ,  $\mathbf{y}_{i+1}$ );
     $\mathbf{x}_{i|i+1}$  = One Step Smoother( $\mathbf{x}_{i|i}$ ,  $\mathbf{x}_{i+1|i+1}$ );
     $\mathbf{p}_{i|i+1}$  = OCBE Costate Function( $\mathbf{x}_{i|i}$ ,  $\mathbf{x}_{i|i+1}$ );
     $\hat{\mathbf{u}}(t)$  = OCBE Estimated Control( $\mathbf{p}_{i|i+1}$ ,  $\Phi(t_{i+1}, t_i)$ );
end

```

**Algorithm 1:** Ballistic OCBE with UT propagation and measurement update.

Filter solutions are generated to assess the performance of the UT modified OCBE. Only three orbit combinations of interest are selected for evaluation as a means to reduce run time. Each case has the target starting at the periapsis of the SL2 orbit, which is the planned location of the lunar gateway. The chaser is initialized on the NL1, PL2, and QPO orbits to evaluate these promising candidate chaser locations. Each simulation lasts the duration of an SL2 orbit period. Initial uncertainty values assume that the chaser is instantiated with ground-based navigation which it is dropping to start autonomous operations, and the target is known to a minor degree of accuracy. Thus, the chaser 1-sigma is 10 km and 10 cm/s, and the target 1-sigma is 100 km and 100 cm/s. Reference [7] serves as the basis for these assumptions. The simulations also include a maneuver at apoapsis with an 8 hour window of increased process noise around the event to account for the expected maneuver without modeling it. This process noise is always included, regardless if a maneuver occurs or not. The UT hyperparameters that produced the most consistent filters were  $\alpha = 0.5$ ,  $\beta = 3$ ,  $\kappa = 0$ . These hyperparameters imply that the system is strongly nonlinear because of the spread of the particles necessary to capture uncertainty propagation adequately.

First, the chaser is placed on the NL1 orbit and simulated. Figure 4 illustrates the smoothed state estimates for chaser and target. The errors in blue are bounded by the filter covariance in red, confirming a converged filter solution.

Figure 5 plots the OCBE control estimate and metric. The control reflects that uncertainty in the state estimate is most significant on the target during apoapsis as desired, and the metric indicates proper filter performance.

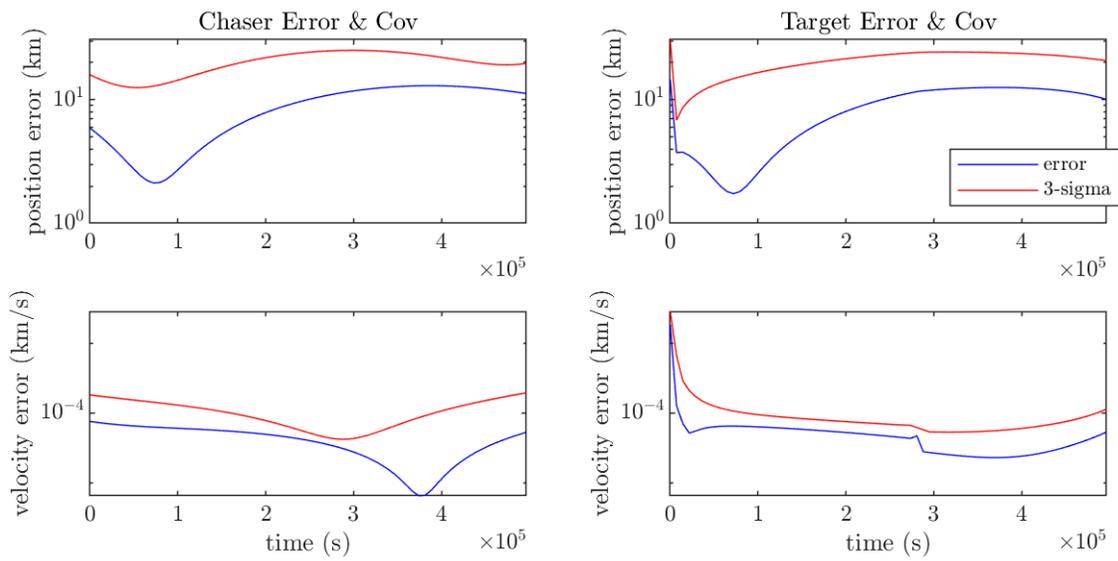


Fig. 4: Chaser on NL1 with target on SL2. Left plot is chaser error and covariance, and the right plot is the target error and covariance.

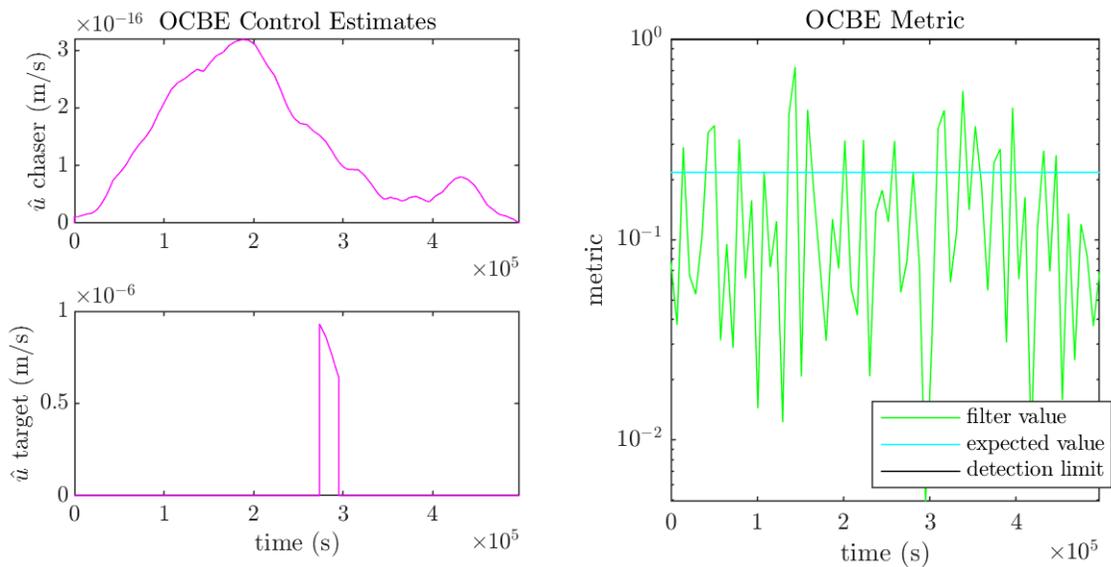


Fig. 5: Chaser on NL1 with target on SL2. Left plot is OCBE estimated control, and the right plot is the OCBE metric.

Next, the simulation with the chaser on the QPO is executed. Figure 6 shows that the covariance properly bounds the errors again. The OCBE control estimate and metric are plotted in Figure 7 and also indicate nominal behavior. It is worth noting that placing the chaser on the QPO produces an order of magnitude better position estimates and marginally better velocity estimates. These results agree with intuition because angular measurements are more informative at closer ranges, and the QPO has a much smaller average range.

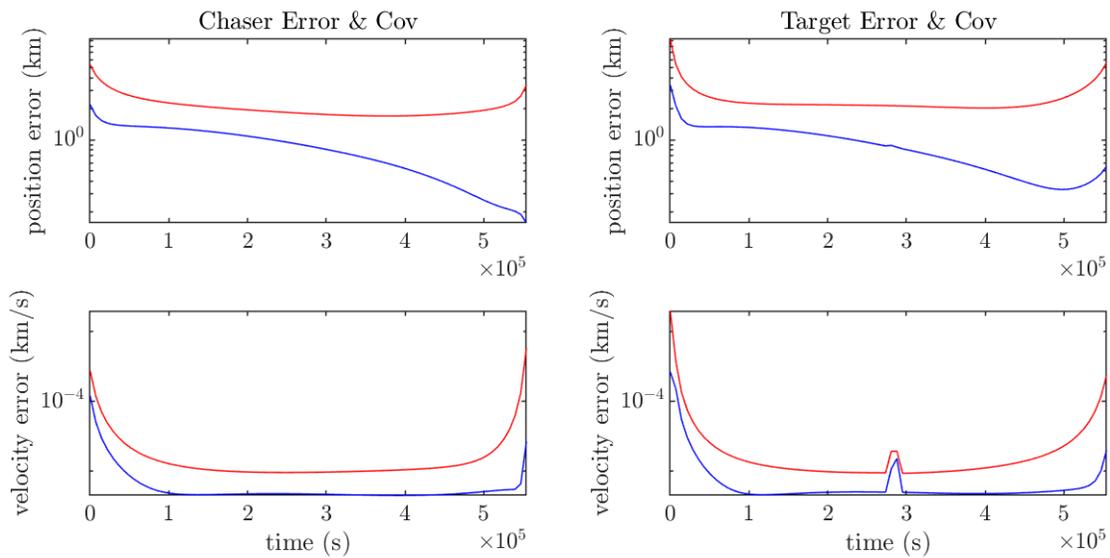


Fig. 6: Chaser on QPO with target on SL2. Left plot is chaser error and covariance, and the right plot is the target error and covariance.

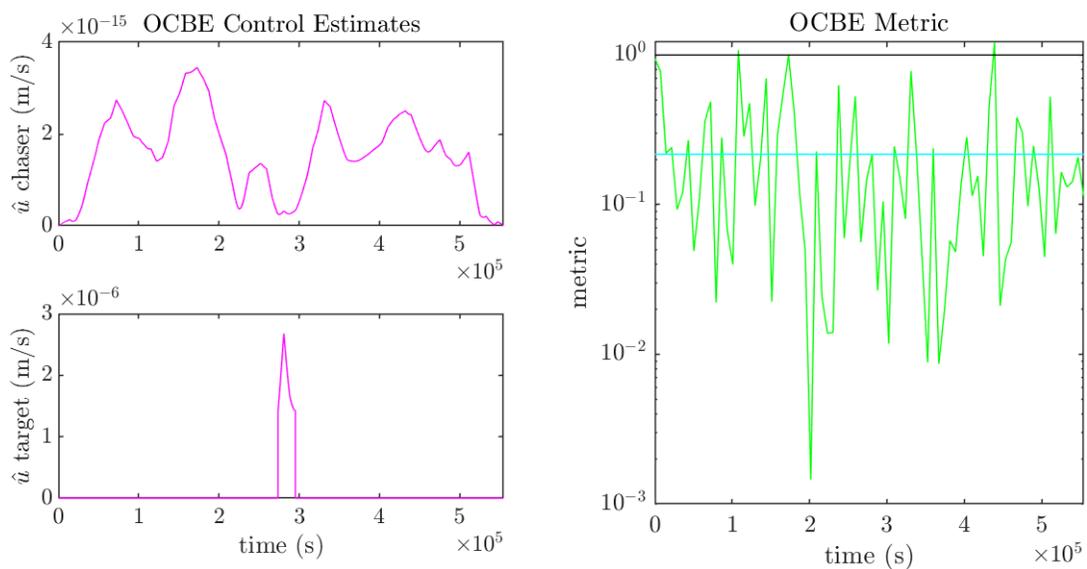


Fig. 7: Chaser on QPO with target on SL2. Left plot is OCBE estimated control, and the right plot is the OCBE metric.

Finally, the chaser is placed on the PL2 orbit and simulated. Figure 8 reveals that the errors are bounded by the covariance again, affirming proper filter behavior. In addition, figure 9 confirms that the OCBE control estimate and metric exhibit no anomalies. These results also highlight that the PL2 estimates are closer to the NL1 results than the QPO results. This similarity is expected since the average range is more comparable between these cases.

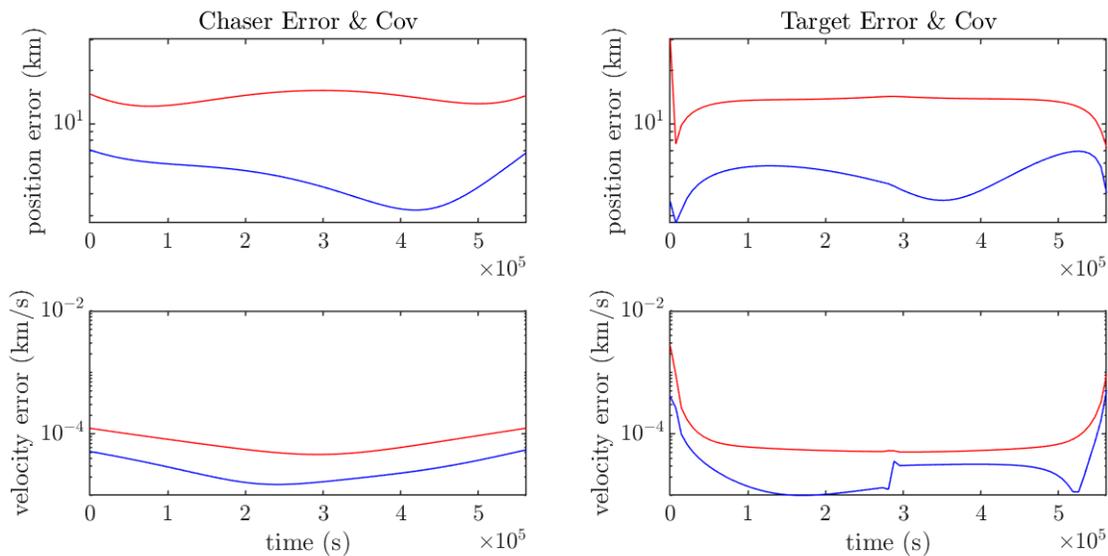


Fig. 8: Chaser on PL2 with target on SL2. Left plot is chaser error and covariance, and the right plot is the target error and covariance.

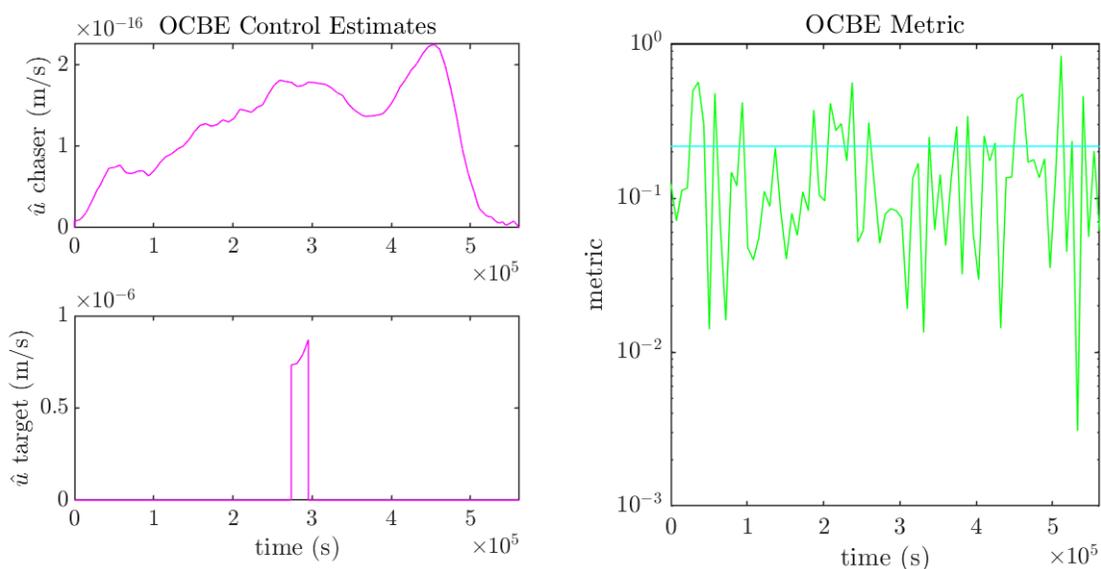


Fig. 9: Chaser on PL2 with target on SL2. Left plot is OCBE estimated control, and the right plot is the OCBE metric.

The UT modified OCBE generally produces proper states estimates given a dual spacecraft system and accounting for a maneuver at apoapsis. While filter results are generally consistent, there are occasions that filters will still fail. Therefore, a Monte Carlo analysis is performed to examine filter convergence rates. In the Monte Carlo analysis, maneuvers are specified as noisy and quiet generalized maneuvers. The simulations were run for 1.5 SL2 orbit periods to determine if the filter converged before the next expected maneuver. The Extended Kalman Filter (EKF) is also assessed as a baseline. Table 6 summarizes convergence rates given 100 runs for each case.

The UT modification enables solutions that linearized filters are not capable of producing. This issue is well known and one of several reasons that other papers adopted additional measurements such as angular-rates. It is also interesting to note that the QPO has more failures than the NL1 or PL2 orbits. This failure is likely due to the highly dynamic

relative motion of the QPO. Through examination, the majority of failures for the UT modified OCBE were caused by large initial errors at periapsis or divergence when coming back around to periapsis. These failures emphasize that a precise periapsis estimate is critical to filtering, even with the UT modification. It is possible that higher-order filters such as higher-order EKFs/UKFs, Gaussian Mixture Model (GMM) filters, or particle filters could address this issue.

Table 6: Filter convergence rates.

Policy	Filter	NL1	QPO	PL2
Noisy	EKF	0/100	1/100	0/100
	UT-OCBE	100/100	94/100	100/100
Quiet	EKF	1/100	0/100	1/100
	UT-OCBE	100/100	93/100	100/100

## 5. MANEUVER CLASSIFICATION

While POs and QPOs are perfectly periodic in the CR3BP, these orbits require regular station keeping in reality. Optimal station keeping maneuvers on halo orbits are known to approximately lie along the stable manifold [10, 1]. Conversely, maneuvers along the unstable manifold will lead to exponential departures from an orbit allowing for low-cost transfers. Thus, maneuvers along the stable and unstable manifold are desirable to identify because they indicate long-term behavior. If these maneuvers are not differentiated, they could easily be conflated and lead to loss of tracking. Therefore, a maneuver classification algorithm based on the OCBE control policy is employed to distinguish maneuvers and maintain custody of a target.

The OCBE control policy is highly sensitive to dynamical mismodeling and makes a powerful maneuver detection indicator. While the OCBE quickly identifies maneuvers, its reconstruction capabilities rely on the optimal control policy, which may not reflect the actual event. This disconnect between optimality and reality could lead to classification difficulties if the control policy is poorly understood or behaved. The OCBE control policy requires some form of uncertainty quantification to ensure classification performance. Attempting to quantify the whole, continuous, policy is not only intractable but not desirable. Instead, the integral of the control profile as defined by Equation 20 is adopted. This value is more appropriate for classification as it represents the total effort applied to the system.

$$\mathbf{z} = \int_{t_0}^{t_f} \hat{\mathbf{u}}(t) dt \quad (20)$$

### 5.1 OCBE Control Quantification

There are many forms of uncertainty quantification, but most suffer from the curse of dimensionality. This issue is particularly concerning for estimation problems because of the sheer quantity of stochastic and variable inputs such as initial conditions, initial uncertainty, measurements, and maneuvers. Monte Carlo analysis breaks the curse of dimensionality because its error convergence rate is  $\sqrt{N_s}$ , where  $N_s$  is the number of samples, which is conveniently independent of the input dimension. This advantage makes Monte Carlo an ideal choice for filter based uncertainty quantification. As such, a Monte Carlo analysis is performed for the three selected orbit combinations from the previous section with both noisy and quiet maneuvers.

First, the chaser is placed on the QPO orbit with a target which is subjected to noisy maneuvers. Figure 10 plots the resulting integral of the estimated control policy and true control. This figure shows that the control integral very closely approximates the actual maneuver, including pointing information. This precise reconstruction of true events means that classification on this system should be highly accurate. Next, the estimated control vectors' norm is plotted in Figure 11 to determine if the OCBE correctly differentiates maneuvering and non-maneuvering cases. The histogram not only confirms that non-maneuvering cases have an order of magnitude smaller estimated control but that the estimated maneuvering distribution closely resembles the designed station keeping mean and covariance. Note that the histogram has been focused on the maneuvering cases and that the non-maneuver plot (grey) goes beyond the axis limit.

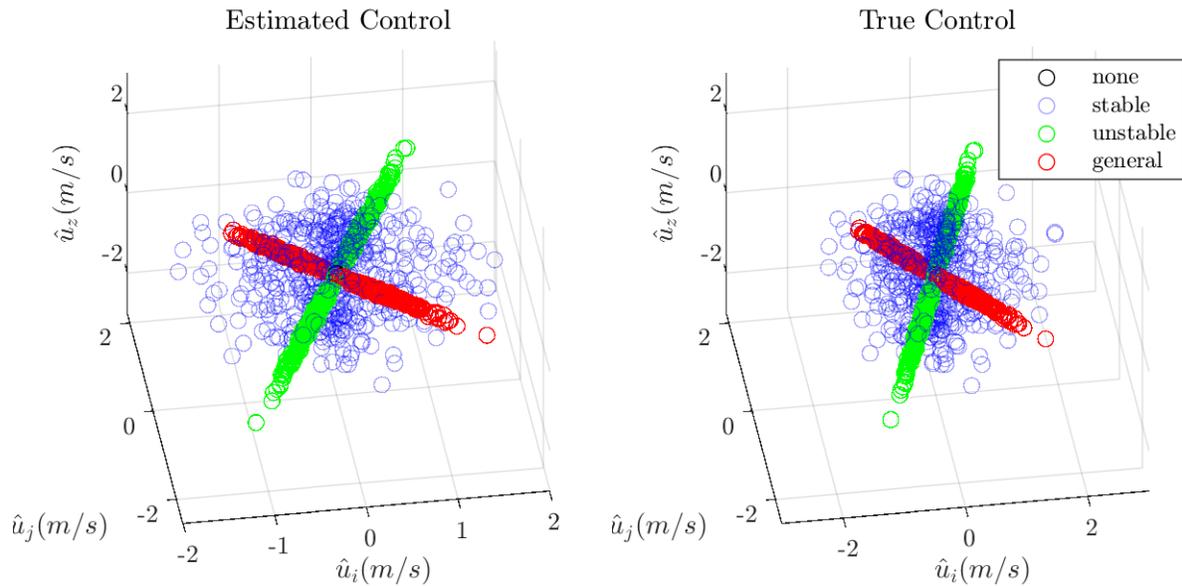


Fig. 10: Chaser on QPO with target on SL2 with noisy maneuvers. Left plot is integral of OCBE estimated control, and the right plot is the true maneuver.

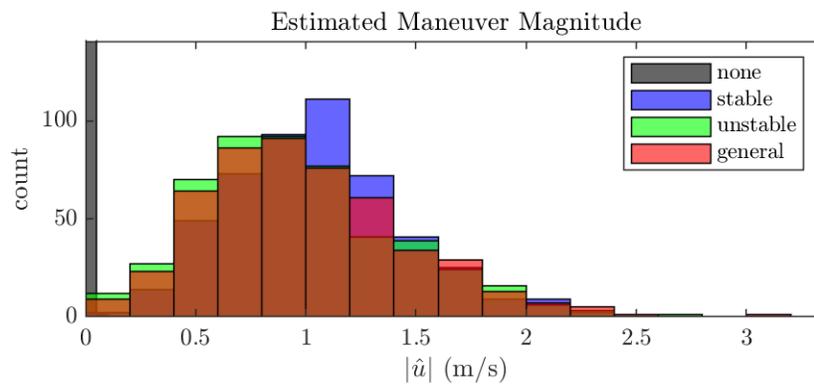


Fig. 11: Chaser on QPO with target on SL2 with noisy maneuvers. Histogram of the norm of the integral of OCBE estimated control.

Next, the same geometry was analyzed with quiet maneuvers. Figure 12 presents the integral of the OCBE control policy and actual control. From the true scatter plot, it is clear that the maneuver execution error has a much more significant impact on the quiet maneuvers. This additional noise implies that classification should be more difficult on the quiet case even though the OCBE does a good job recreating the events. Figure 13 shows the norm of the estimated control. Once again, the non-maneuvering cases are consistently an order of magnitude smaller, and the maneuvering cases closely approximate the true distribution.

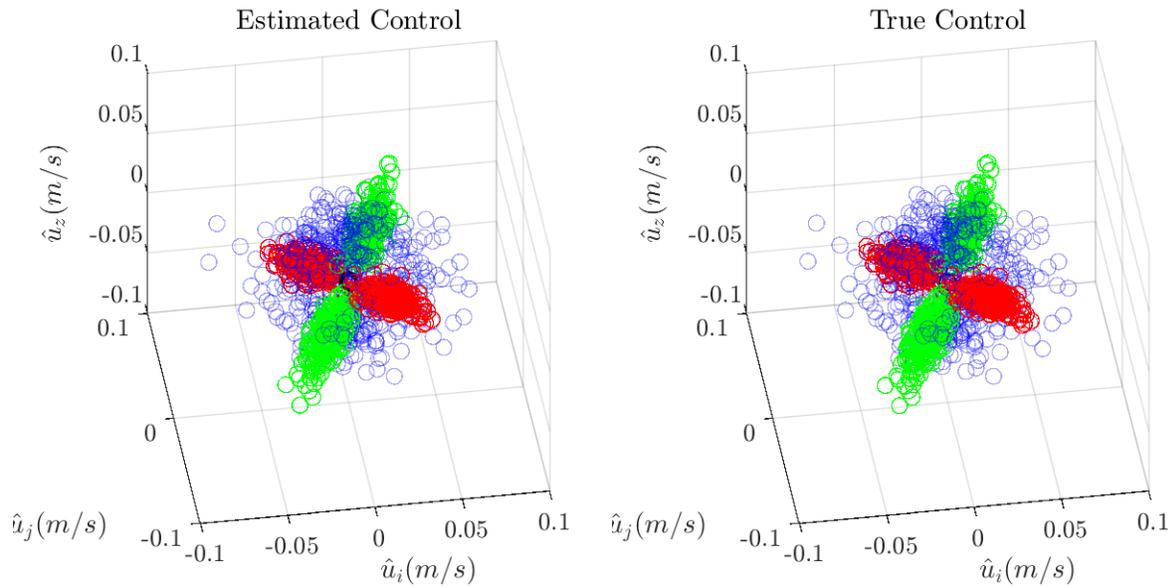


Fig. 12: Chaser on QPO with target on SL2 with quiet maneuvers. Left plot is integral of OCBE estimated control, and the right plot is the true maneuver.

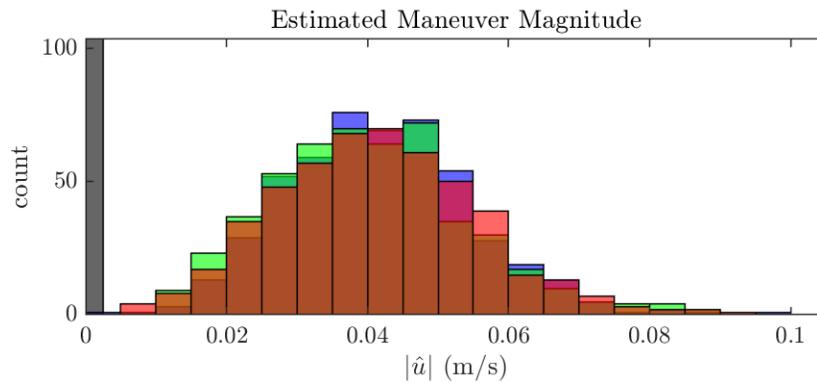


Fig. 13: Chaser on QPO with target on SL2 with quiet maneuvers. Histogram of the norm of the integral of OCBE estimated control.

Now the chaser located on the PL2 orbit is examined. Figure 14 gives the noisy and quiet OCBE control estimates. Note that the true maneuvers are not displayed because they are equivalent to the QPO cases shown in Figures 10 and 12. Once again, the OCBE estimate very closely approximates truth in the noisy case. However, unlike the QPO, the quiet case has significant overlap not seen in the true maneuvers. Fortunately, the structure of the quiet case still has stable and unstable maneuvers loosely aligned with the manifold directions. This association between estimated and true control means that classification is still possible even if it is more difficult. Figure 15 plots the norms of the two cases. The noisy case increases estimated control with a maneuver, while the quiet case shows increased overlap as expected from the scatter plots.

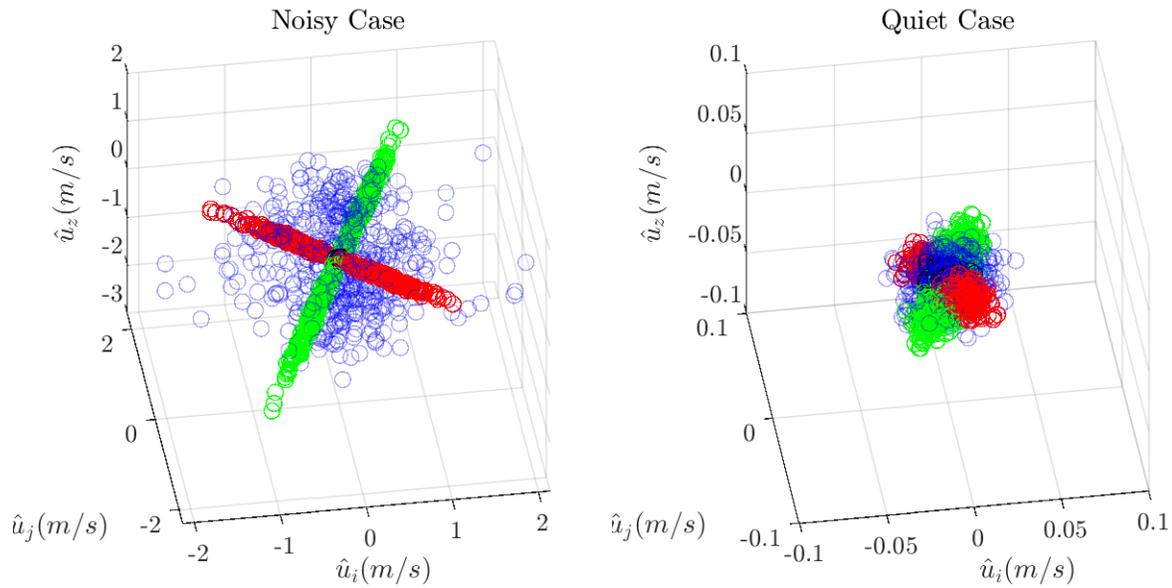


Fig. 14: Integral of estimated OCBE with chaser on PL2 and target on SL2. Left plot is noisy case and the right plot is the quiet case.

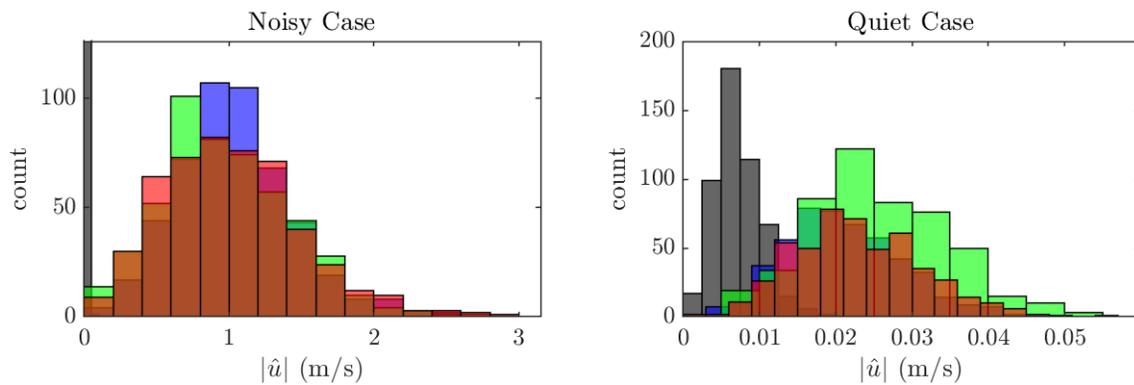


Fig. 15: Norm of the integral of OCBE estimated control with chaser on PL2 with target on SL2 with quiet maneuvers. Left plot is noisy case and the right plot is the quiet case.

Finally, the chaser is placed on NL1 orbit and simulated. Figure 16 reveals that the OCBE control estimates generally represent the true control for the noisy case, with a minor spread along one dimension. The dispersion closely aligns with the average range vector between the crafts, stemming from the relatively static and far viewing geometry between the NL1 and SL2 orbits. Figure 16 also exposes that the quiet estimated control is significantly smaller than the other orbit combinations. This poor reconstruction is because the NL1 orbit has the longest average range, and optical measurements are weaker over longer ranges. The histograms of the norms of the estimated control are plotted in Figure 17. As expected, the noisy case is well separated, and the quiet case is not.

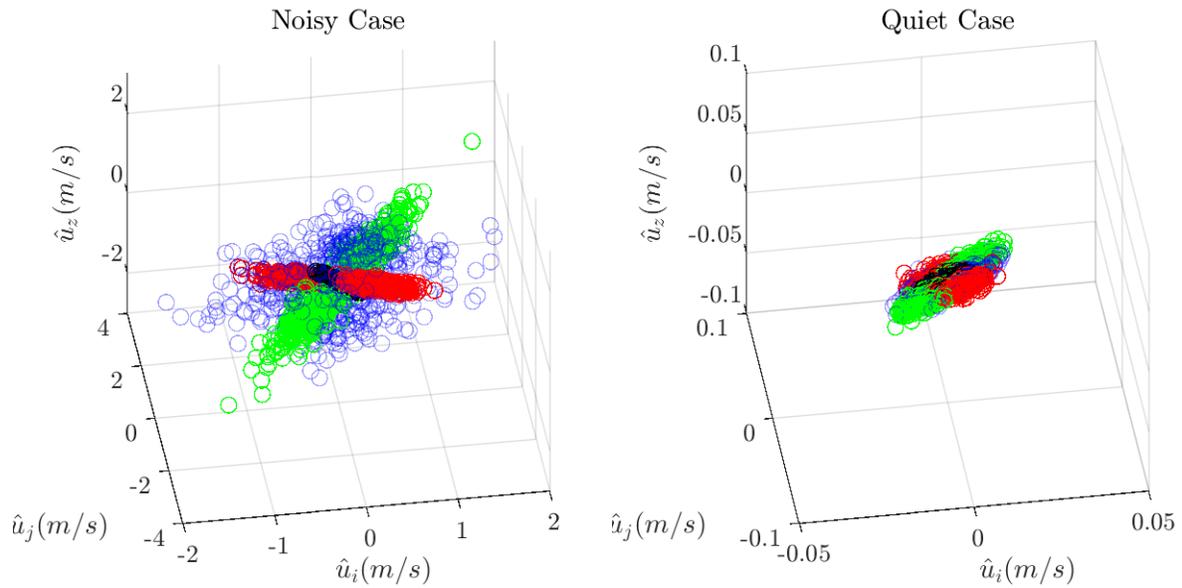


Fig. 16: Integral of estimated OCBE with chaser on NL1 and target on SL2. Left plot is noisy case and the right plot is the quiet case.

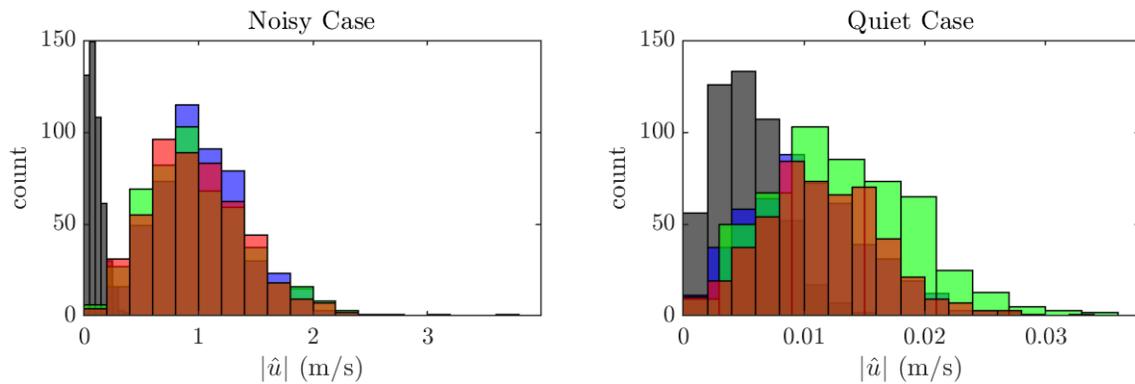


Fig. 17: Norm of the integral of OCBE estimated control with chaser on NL1 with target on SL2. Left plot is noisy case and the right plot is the quiet case.

## 5.2 Classification

The integral of the OCBE control profile has a structure related to the true maneuvers in every case examined, making it an appropriate input for maneuver classification. In noisy cases, the estimated control vectors align with the corresponding manifold from which the true maneuver is generated. The quiet cases have structure associated with the stable and unstable manifolds, but there is a significant overlap between maneuver types. Thus, two classification algorithms are examined to determine which performs the best. The first method is a simplistic one-hot optimization approach. The other method utilizes Support Vector Machines (SVMs).

The one-hot method optimizes two parameters ( $a, b$ ) to minimize the cost function given by Equation 21. In this equation  $T_i$  is the truth probability of sample  $i$ . The truth vector takes a value of 1 for the true class, and 0 for all other classes, making it a one-hot vector. The variable  $L_i$  is the predicted probability, a one-hot vector that chooses the most likely class. If the class probability vector is ordered as in Equation 22, then the predicted probability is calculated by Equation 23.

In general, this one-hot method functions on the same fundamental principle as an SVM; it projects the input data into

a domain where it is easily separable. The one-hot method takes advantage of the accurate reconstruction capabilities of the OCBE control policies to define class boundaries in a hierarchical order. An incoming control estimate must be above  $a$  to be classified as a maneuver. Then, if the angle between the estimated control and a manifold is less than  $\cos^{-1}(b)$ , it is assigned to that manifold. Given training data, the parameters  $a, b$  can be optimized by an algorithm of choice and used to classify broader data sets.

$$J = -\sum_{i=1}^{N_L} \mathbf{T}_i^T \mathbf{L}_i \quad (21)$$

$$\mathbf{L}_i = [p(\text{none}), p(\text{stable}), p(\text{unstable}), p(\text{general})]^T \quad (22)$$

$$\mathbf{L}_i = \begin{cases} [1, 0, 0, 0]^T & \text{if } \|z\| \leq a \\ [0, 1, 0, 0]^T & \text{else if } \text{abs}(\hat{\mathbf{z}}^T \hat{\mathbf{e}}_s) \geq b \\ [0, 0, 1, 0]^T & \text{else if } \text{abs}(\hat{\mathbf{z}}^T \hat{\mathbf{e}}_u) \geq b \\ [0, 0, 0, 1]^T & \text{otherwise} \end{cases} \quad (23)$$

The SVM approach leverages well-defined machine learning techniques and libraries. SVMs function by projecting the input data into a higher dimension using a kernel function where the data is easily separable by a hyperplane. This work leverages the package Sci-Kit Learn which has well-developed SVM methods and models [26]. The specific algorithm of choice is an SVM with a Radial Basis Function (RBF) as its kernel. RBF is an appropriate kernel because it is clear from the uncertainty quantification that the classes are well clustered and RBF acts as a similarity measure which will take advantage of this clustering.

The data generated during the OCBE control quantification is used to train and test the classification methods. Half of the dataset is randomly selected and used to optimize the one-hot parameters or train the SVM. Once the training is complete, the entire data set is passed to the classifier to be labeled. The subsequent output is compared to truth values using confusion matrices.

First, the case with the chaser located on the QPO is examined. Figure 18 gives the confusion matrices for both algorithms and station keeping policies to evaluate their effectiveness. This figure confirms that both algorithms perform exceedingly well, but the SVM surpasses the one-hot classifier for quiet generalized maneuvers. It is also worth noting that the minimum accuracy is 86% across all cases and algorithms. Thus, given a chaser on a QPO, both algorithms are sufficient and produce accurate results.

Next, the results with the chaser on the PL2 orbit are processed. The confusion matrices in Figure 19 again show that the two classifiers are comparable for the noisy case, but the quiet case bears a more significant gap for generalized label accuracy. The PL2 dynamics exacerbate the QPO results because PL2 is further away and suffers more significant OCBE control estimate inaccuracies. It is also worth noting that the one-hot method prefers properly separating the maneuver vs. non-maneuver subsets. In general, for the PL2 case, the SVM appears superior.

Finally, the confusion matrices with the chaser located on the NL1 orbit are shown in Figure 20. Given these results, the one-hot method evidently prefers correct maneuver detection but struggles with over-labeling manifold maneuvers. On the other hand, the SVM performs uniformly across the label types and does surprisingly well given how poor the OCBE control estimates are in the quiet case. Thus, the SVM is the preferred method given a chaser on NL1.

From the cumulative results of this section, the SVM classifier is the superior classifier as it equally spreads success rates across each maneuver type. However, the one-hot method is still appealing for maneuver detection or if the relative orbits are close. Unfortunately, the one-hot method degrades as range increases because the OCBE reconstruction deviates from reality, and thus the classification problem is higher-dimension. It is also clear that closer relative orbits produce better classification in general, which is expected since optical measurements are stronger over shorter ranges. Thus, placing the chaser on a QPO about a target on a PO is ideal as it ensures close and bounded relative motion with diverse viewing geometries. Furthermore, QPO to PO relative estimation was able to reconstruct true maneuvers almost one-to-one. These advantages for the QPO lead to a minimum of 94% accuracy with the SVM, even in the quiet case. It is worth noting that across all cases/orbits the one-hot method is never worse 33% inaccurate for maneuver vs. non-maneuver classification, and the SVM is never worse than 27% inaccurate for the same subset.

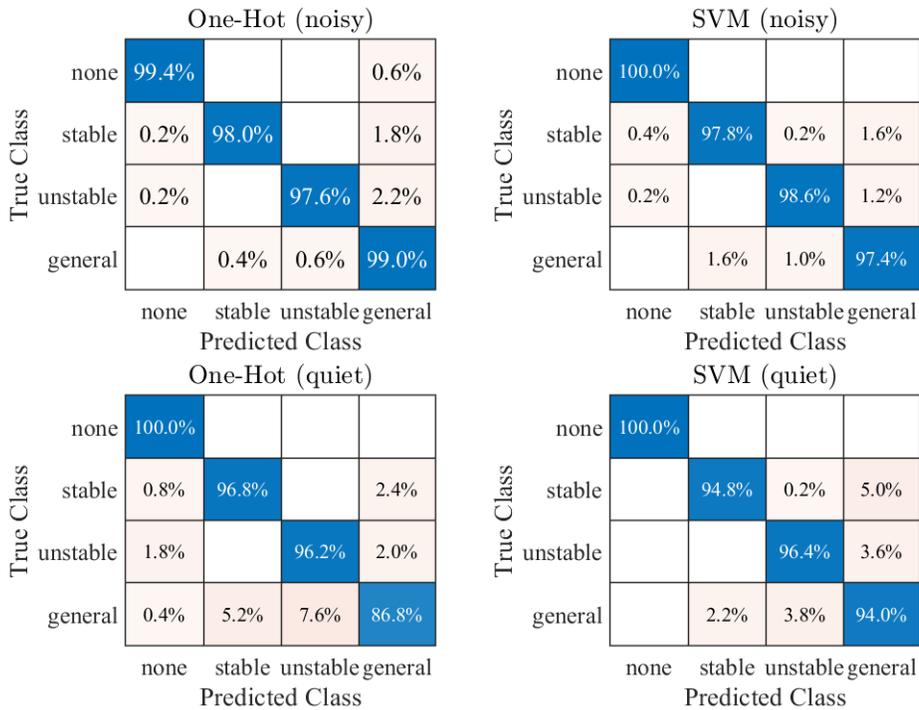


Fig. 18: Confusion matrices for a chaser on the QPO. The left plots are the one-hot algorithm, and the right are the SVM. Top plots are noisy station keeping, bottom are quiet.

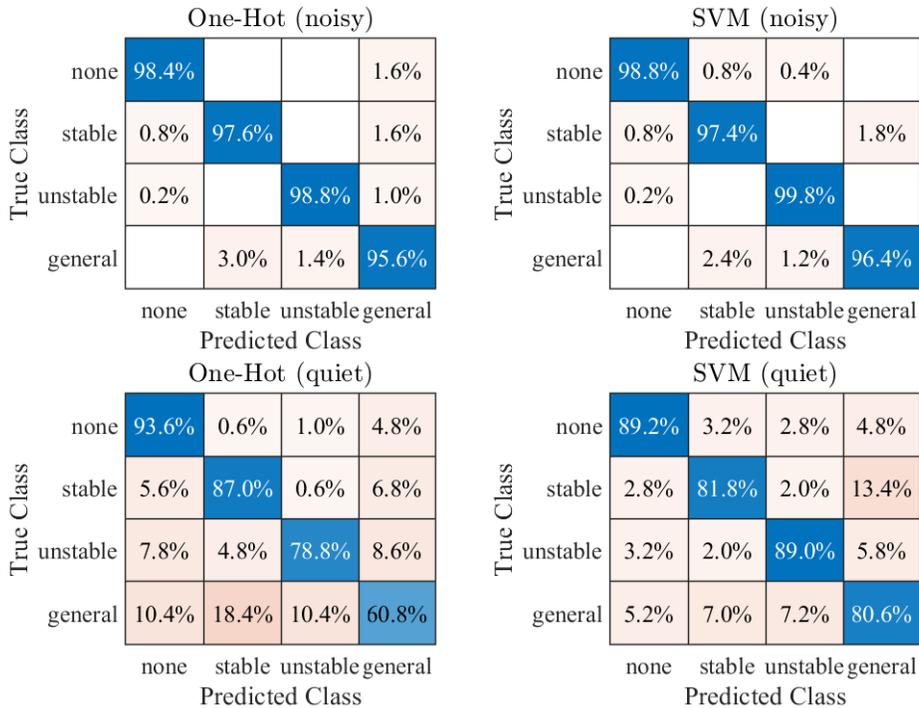


Fig. 19: Confusion matrices for a chaser on the PL2. The left plots are the one-hot algorithm, and the right are the SVM. Top plots are noisy station keeping, bottom are quiet.

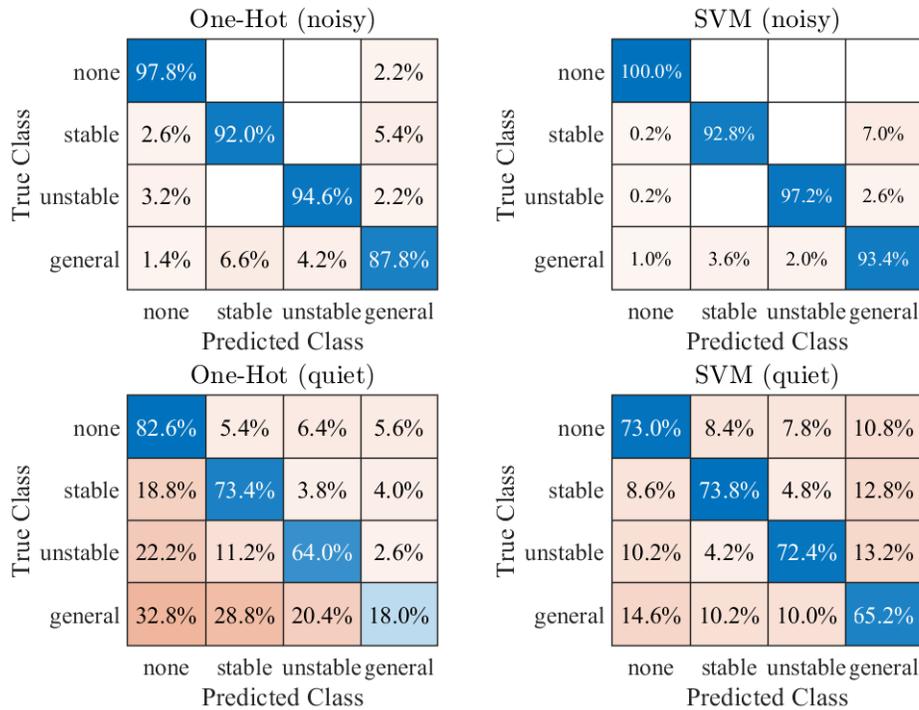


Fig. 20: Confusion matrices for a chaser on the NL1. Left plots are one-hot, right plots are SVM. Top plots are noisy station keeping, bottom are quiet.

## 6. CONCLUSION

This work set out to determine if optical measurements are sufficient for autonomous dual craft estimation in the cislunar regime. The first step in achieving this goal is to prove that optical measurements lead to a fully observable system. Observability is proven using the standard definition of linear observability, which improves on the traditional filtering observability technique that employs the information matrix. The standard definition is beneficial because it is less numerically sensitive and enables trivial null space analysis. The observability results show that optical measurements lead to observable systems and do so faster and more consistently than range measurements.

With observability proven, filtering solutions are generated using a newly developed UT modified OCBE. The UT modification, or other higher-order methods, are necessary to enable estimation in the cislunar environment due to the strongly nonlinear dynamics. With the UT modification, the OCBE can generate proper filter and control estimates. However, even with the UT modification, the filter occasionally struggles with QPO to PO estimation as the relative motion is highly dynamic. It is likely that higher-order, GMM, or particle filters could appropriately address this issue and should be studied for future work.

Finally, maneuver classification methods are generated and tested to label OCBE control estimates as non-maneuvering, stable, unstable, or general families. Uncertainty quantification of the control estimate is necessary to ensure that classification is possible. The quantification results show that OCBE event reconstruction strongly depends on relative orbit geometry, with the QPO to PO estimation leading to highly accurate event reconstruction. Two classification methods are then tested: one-hot and SVMs. The one-hot method is a simplistic model that leverages proper event reconstruction and proves to be efficient for maneuver detection. The SVM method was generally more accurate and provided effective classification even with poor event reconstruction. The SVM with the QPO to PO observation scheme achieved the best accuracy with a minimum of 94% classification accuracy, even given an average maneuver size of only 40 mm/s. Future work should study alternative classification techniques to determine if they can produce better results for poor event reconstructions.

In conclusion, optical measures are power measurements that are capable of autonomous dual craft estimation. Optical measurements produce linearly observable systems that should be estimated with high-order or nonlinear filters. The

newly developed UT modified OCBE provides sufficient consistency when generating state and control estimates with these measurements. The OCBE estimated control policy is then leveraged to classify maneuvers, with the best accuracy resulting from QPO to PO relative motion. Collectively, the developed methods facilitate autonomous operations in cislunar space and engender further research in this area.

## REFERENCES

- [1] Manfred Bester, Daniel Cosgrove, Sabine Frey, Jeffrey Marchese, Aaron Burgart, Mark Lewis, Bryce Roberts, Jeremy Thorsness, John McDonald, Deron Pease, Gregory Picard, Martha Eckert, and Renee Dumlao. Artemis operations - experiences and lessons learned. In *2014 IEEE Aerospace Conference*, pages 1–20, 2014.
- [2] Mark Bolden, Timothy Craychee, and Erin Griggs. An evaluation of observing constellation orbit stability, low signal-to-noise, and the too-short-arc challenges in the cislunar domain. 2020.
- [3] Nicholas Bradley, Zubin Olikara, Shyam Bhaskaran, and Brian Young. Cislunar navigation accuracy using optical observations of natural and artificial targets. *Journal of Spacecraft and Rockets*, 57(4):777–792, 2020.
- [4] Lyndon B. Johnson Space Center. Constellation program: America’s spacecraft for a new generation of explorers the orion crew exploration vehicle. Technical Report FS-2006-08-022-JSC.
- [5] John A Christian and E Glenn Lightsey. Review of options for autonomous cislunar navigation. *Journal of Spacecraft and Rockets*, 46(5):1023–1036, 2009.
- [6] Jason Crusan, Jacob Bleacher, Joe Caram, Douglas Craig, Kandyce Goodliff, Nicole Herrmann, Erin Mahoney, and Marshall Smith. Nasa’s gateway: An update on progress and plans for extending human presence to cislunar space. In *2019 IEEE Aerospace Conference*, pages 1–19. IEEE, 2019.
- [7] Diane Davis, Sagar Bhatt, Kathleen Howell, Jiann-Woei Jang, Ryan Whitley, Fred Clark, Davide Guzzetti, Emily Zimovan, and Gregg Barton. Orbit maintenance and navigation of human spacecraft at cislunar near rectilinear halo orbits. *27th AAS/AIAA Space Flight Mechanics Meeting*, 2017.
- [8] Anaïs Delépaut, Pietro Giordano, Javier Ventura-Traveset, Daniel Blonski, Miriam Schönfeldt, Philippe Schoonejans, Sarmad Aziz, and Roger Walker. Use of gnss for lunar missions and plans for lunar in-orbit development. *Advances in Space Research*, 66(12):2739–2756, 2020.
- [9] Samuel J Fedeler, Marcus J Holzinger, and William W Whitacre. Sensor tasking in the cislunar regime using monte carlo tree search.
- [10] David Folta, Mark Woodard, and Daniel Cosgrove. Stationkeeping of the first earth-moon libration orbiters: The artemis mission. In *AAS/AIAA Astrodynamics Specialist Conference, Girdwood, Alaska*, 2011.
- [11] E Fowler, D Paley, and S Hurtt. Observability metrics for space-based cislunar domain awareness. In *31st AIAA/AAS Space Flight Mechanics Meeting*, 2021.
- [12] Elizabeth A Frank, Casey Dreier, Benton Clark, and Clive R Neal. Why and how to leverage the commercial space sector for the benefit of planetary science and its community. *Bulletin of the American Astronomical Society*, 53(4):403, 2021.
- [13] Jesse A Greaves and Daniel J Scheeres. Observation and maneuver detection for cislunar vehicles using optical measurements and the optimal control based estimator. 2021.
- [14] Joseph R Guinn, Shyam Bhaskaran, Todd A Ely, Brian M Kennedy, Tomas J Martin-Mur, Ryan S Park, Joseph E Riedel, Duane C Roth, and Andrew T Vaughan. The deep space positioning system (dps)-navigator concept for the lunar gateway. 2019.
- [15] Keric Hill and George H Born. Autonomous interplanetary orbit determination using satellite-to-satellite tracking. *Journal of guidance, control, and dynamics*, 30(3):679–686, 2007.
- [16] Keric Hill, George H Born, and Martin W Lo. Linked, autonomous, interplanetary satellite orbit navigation (liaison) in lunar halo orbits. 2006.
- [17] M Hinga. Cis-lunar autonomous navigation via implementation of optical asteroid angle-only measurements. In *Proceedings of the 2003 AMOS Technical Conference*, 2020.
- [18] MJ Holzinger, CC Chow, and P Garretson. A primer on cislunar space. 2021.
- [19] Erica L Jenson and Daniel J Scheeres. Semianalytical measures of nonlinearity based on tensor eigenpairs.
- [20] Simon J Julier and Jeffrey K Uhlmann. Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3):401–422, 2004.
- [21] Stéphanie Lizy-Destrez, Paolo Guardabasso, Laurent Beauregard, and Emmanuel Blazquez. Towards the moon and beyond: preparing for the future of cislunar and solar system exploration. 2021.

- [22] Daniel P Lubey. *Maneuver Detection and Reconstruction in Data Sparse Systems with an Optimal Control Based Estimator*. PhD thesis, University of Colorado at Boulder, 2015.
- [23] NASA. *Artemis Plan: NASA's Lunar Exploration Program Overview*, September 2020.
- [24] Clark P Newman, Diane C Davis, Ryan J Whitley, Joseph R Guinn, and Mark S Ryne. Stationkeeping, orbit determination, and attitude control for spacecraft in near rectilinear halo orbits. 2018.
- [25] Nathan L Parrish, Matthew J Bolliger, Ethan Kayser, Michael R Thompson, Jeffrey S Parker, Bradley W Cheetham, Diane C Davis, and Daniel J Sweeney. Near rectilinear halo orbit determination with simulated dsn observations. In *AIAA Scitech 2020 Forum*, page 1700, 2020.
- [26] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [27] Allyson Reneau. *An International Perspective*, pages 69–81. Springer International Publishing, Cham, 2021.
- [28] Antonino Salmeri and Matej Poliacsek. All for one and one for all: Recommendations for sustainable international lunar base utilization and exploration approaches. In *Proceedings of the 71st International Astronautical Congress-The Cyberspace Edition*. International Astronautical Federation, 2020.
- [29] Marshall Smith, Douglas Craig, Nicole Herrmann, Erin Mahoney, Jonathan Krezel, Nate McIntyre, and Kandyce Goodliff. The artemis program: An overview of nasa's activities to return humans to the moon. In *2020 IEEE Aerospace Conference*, pages 1–10. IEEE, 2020.
- [30] David A Vallado. *Fundamentals of astrodynamics and applications*, volume 12. Springer Science & Business Media, 2001.
- [31] Jacob K Vendl and Marcus J Holzinger. Cislunar periodic orbit analysis for persistent space object detection capability. *Journal of Spacecraft and Rockets*, pages 1–12, 2021.
- [32] Sean V Wagner and Troy D Goodson. Execution-error modeling and analysis of the cassini-huygens spacecraft through 2007. 2008.
- [33] Eric A Wan and Rudolph Van Der Merwe. The unscented kalman filter for nonlinear estimation. In *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373)*, pages 153–158. Ieee, 2000.