

# Robust Cislunar Initial Orbit Determination

**Sam Wishnek**

*The University of Colorado at Boulder, Smead Aerospace Engineering Sciences Department*

**Marcus J. Holzinger**

*The University of Colorado at Boulder, Smead Aerospace Engineering Sciences Department*

**Patrick Handley**

*Ball Aerospace*

## ABSTRACT

An approach for initial orbit determination outside the constraints of two-body dynamics is developed and assessed. The method formulates a trial solution state as a cost function in order to optimize over the unobserved state space dimensions. The proposed approach can be adapted to work with any dynamics that provide a smooth and continuous cost function to optimize over. The approach is tested for several cislunar orbits under various conditions. These include the Nelder-Mead optimizer parameters and the measurement error. An analytical approach for determining if a derived solution meets necessary conditions for convergence is developed as well.

## 1. INTRODUCTION

As the population of near-Earth objects has grown, more scientific, commercial, and military operators have expanded their utilization of the domain. One opportunity of growing interest to these operators is exploration of the opportunities provided by the cislunar domain and beyond. Methods for maintaining space domain awareness in the cislunar domain are essential to support these current and future missions. One of the essential tools of space domain awareness is initial orbit determination for characterizing the orbits using a minimal set of observations. However, existing initial orbit determination algorithms rely on the two-body assumption for estimating the state [1]. For cislunar missions that operate where moon gravitational effects are significant, new methods for initial orbit determination are required in order to estimate orbit state under multi-body dynamics.

In order to perform initial orbit determination beyond the reasonable limits of the two point-mass approximation, the proposed method applies an optimization-based approach to find the state or set of states that can feasibly match the observations. The approach leverages admissible regions to constrain the search space and improve convergence characteristics. An angles-only method is implemented, but the theory behind an angles and angle rates implementation is discussed as well. The optimization for the angles-only case is run over the four unobserved dimensions while the angles and angle rates case is run over the two unobserved dimensions. These are range and three velocity components for angles-only and range and range rate for angles with angle rates. The optimized cost function is the sum of Mahalanobis distances between predicted and observed measurements based on a trial solution.

The approach implements JPL's SPICE Toolkit to accurately predict the relative positions of the Solar System bodies [2]. The gravitational influence of the Earth, Moon and Sun dominates the dynamics of orbits in cislunar space and beyond. The influence of the Sun is essential for identifying if a target object is in a cislunar, geocentric, or heliocentric orbit as targets sufficiently far from the observer can be difficult to differentiate in terms of orbit domains. The geopotential models of the Earth and Moon are also included up to fourth order using data from the EGM2008 and GRGM1200A respectively [3][4]. These include both the zonal and tesseral components.

The angles-only case uses three observations while the angles with angle rates case uses two observations. The angles-only case is more applicable for larger range observations as the angular velocity of a target in cislunar orbit typically has a much slower inertial velocity than an object in a near-Earth orbit and the significant increase in separation distance between target and observer makes the observed angle rates much smaller. The relative drop in angle rate observability compared to the near-Earth orbits significantly restricts the types of measurements that can accurately

capture angle rates. Accordingly, angles-only measurements are far more applicable for more distant domains. The cost function itself is the sum of six Mahalanobis distances for angles-only observations and two Mahalanobis distances for angles with angle rates. From a single trial value, these are the distances between the predicted and observed measurements for the orbit propagated to each measurement time. The measurement treated as truth is then changed and the process repeated until each measurement is given equal weighting in the estimate. For perfect measurements and an accurate state guess, the cost function returns zero. We find that small measurement errors can still yield accurate estimated states.

The Mahalanobis-based cost function alone tends to form long curves through the solution space that some optimization methods may struggle with [5]. In order to improve the geometry of the underlying contours for optimization, modifications to the solution space and cost function are investigated. A method used to improve convergence for the cislunar domain is the modification of the admissible regions from the two-body implementation. The admissible regions for cislunar space differ substantially from the near-Earth orbits, and these limits are shifted accordingly to penalize the cost function in a way that improves convergence and limits the search domain. Investigated penalty function parameters include geocentric periaapse and apoapse limits, heliocentric periaapse and apoapse, obstruction of line of sight due to the Moon and Earth, occultation of the Sun on the target due to the Moon and Earth, and absolute limits on maximum and minimum velocities and ranges. Well targeted application of these admissible regions can significantly restrict the search space and improve the convergence properties of the optimization [6].

The geometry of the cost function is explored by generating contour plots over the solution space. The results of this initial orbit determination method are investigated in terms of convergence accuracy with respect to measurement error and the relative orbit geometry of the observer and target using a series of Monte Carlo simulations.

The proposed method is an approach for performing initial orbit determination beyond near-Earth space and the algorithm itself is neutral to the dynamics included in the propagator. The proposed method is a new approach that optimizes over a cost function to find feasible states that match observations. Implementations for the angles-only are investigated for relative accuracy along with the implementation theory for the angles and angle rates case. The necessary conditions to determine convergence is derived and implemented to exclude non-global minima from returning false solutions. The efficacy of these methods are demonstrated through a series of Monte Carlo simulations investigating the response of the algorithm across cislunar space and in response to measurement error.

## 2. THEORY

### 2.1 Cost Function

At the core of the optimization-based approach is a cost function that models how well a trial solution fits the collected measurements. The cost function is a metric that should rigorously return a value that represents how well a state estimate fits the data considering the known measurement uncertainty. For this reason, the cost function is chosen as a sum of Mahalanobis distances. The Mahalanobis distance is expressed in its most general form in equation 1 where  $\vec{u}$  is the measurement vector,  $\hat{u}$  is a predicted measurement vector based on the state estimate, and  $S$  is the measurement covariance matrix and is assumed known based on the particulars of the measurement system. The Mahalanobis distance is computed in the measurement space which depends on the observer's state. The observer's state is assumed known and any uncertainty folded into the measurement covariance matrix.

$$D = \sqrt{(\vec{u} - \hat{u})^T S^{-1} (\vec{u} - \hat{u})} \quad (1)$$

Several measurements are required for the state to be observable. For angles-only measurements, three measurements are required for the state to be observable and for measurements with both angles and angle-rates, two measurements are required. Accordingly, the cost function is taken as the sum of Mahalanobis distances relating each pair of measurements to each other both forward and backward in time. This process starts with a trial solution state at one of the measurement times. The choice of measurement time is arbitrary and it will be assumed that the first measurement time,  $t_1$ , is selected.

### 2.1.1 Angles and Angle Rates Measurements

With only two measurement times, the case with measurements containing both the relative angle from the observer to the target and the observed angle rates is covered first. In this case, the measurement vector contains four linearly-independent elements, two for the angles, and two for the angle rates. The solution space is two-dimensional. For this case, the cost function is the sum of two Mahalanobis distances. The trial solution, a vector with range from observer to target and the range rate, is combined with the measurement vector for the  $t_1$  measurement. This yields a full six-element state. This state is propagated to the time of the second measurement under the simulated cislunar dynamics to return an expected full state at the time of the second measurement. This can be combined with the known observer state to extract a predicted measurement at the time of the second measurement. The first Mahalanobis distance in the angles and angle-rates cost function is equation 1 with this propagated predicted measurement as  $\hat{u}$  and the actual measurement for the second time as  $\bar{u}$ .

The second Mahalanobis distance is based on the first measurement. Since the state is built from the measurement at the initial time, directly calculating this distance would return zero and there would be no compensation for the measurement error at  $t_1$ . To ensure this information is included in the cost function, the initial full state is propagated to the time of the second measurement. Then the unobserved parameters of the state at this time are extracted from the predicted state with the known observer state. This range and range-rate are then combined with the second measurement to build a full state estimate. This state can then be propagated back in time to the first measurement. At this point the approach becomes similar to that of the first Mahalanobis distance. The expected measurement is extracted and the Mahalanobis distance is built for the initial measurement using equation 1. These distances are added together to construct the cost function.

### 2.1.2 Angles Only Measurements

For more measurements, the required number of terms in the cost function sum is 2-permutations of  $n$  where  $n$  is the number of measurements. Accordingly, for the three angles-only measurements required to observe a full state, six Mahalanobis distances are included in the cost function. These are constructed similarly to those of the angles and angle-rates case. However, for angles-only data the measurement state has two linearly independent values and the solution space is four dimensional. The first two Mahalanobis distances take the estimated initial state and propagate it to the time of the second and third measurements. The expected measurements at these two times are compared to the observed measurements for the first two Mahalanobis distances.

The other four are generated in a similar manner as the second case in the angles and angle rates case. The initial state estimate is propagated to the second or third measurement time. Then the expected unobserved state is pulled from the propagated state and combined with the measurement at that time. This state is then propagated to one of the other measurement times to find an expected measurement to compare to the observation. Take  $\vec{x}_i$  to represent the full six-element target state at measurement time  $t_i$  and  $\vec{R}_i$  the six-element observer state at measurement time  $t_i$ . Let the function  $\vec{x}_j = \vec{p}(\vec{x}_i, t_i, t_j)$ , represent the propagation of the state at measurement time  $t_i$  to measurement time  $t_j$ . Let the function  $\hat{u} = \vec{f}(\vec{x}, \vec{R})$  extract the expected measurement from state  $\vec{x}$  and let the function  $\vec{v} = \vec{g}(\vec{x}, \vec{R})$  extract the unobserved states from state  $\vec{x}$ . Finally, let the function  $\vec{x} = \vec{h}(\vec{v}, \vec{u})$  build a state  $\vec{x}$  from a measurement,  $\vec{u}$ , and an unobserved state  $\vec{v}$ . With this, first two Mahalanobis distances are  $D_{1,2}$  and  $D_{1,3}$  and expressed by equation 2 while the latter four Mahalanobis distances are  $D_{2,1}$ ,  $D_{2,3}$ ,  $D_{3,1}$ , and  $D_{3,2}$  where  $D_{i,j}$  is given in equation 3.

$$D_{1,j} = \sqrt{\left(\bar{u}_j - \vec{f}\left(\vec{p}(\vec{x}_1, t_1, t_j), \vec{R}_j\right)\right)^T S_j^{-1} \left(\bar{u}_j - \vec{f}\left(\vec{p}(\vec{x}_1, t_1, t_j), \vec{R}_j\right)\right)} \quad (2)$$

$$D_{i,j} = \sqrt{\left(\bar{u}_j - \vec{f}\left(\vec{p}\left(\vec{h}\left(\vec{g}\left(\vec{p}(\vec{x}_1, t_1, t_i), \vec{R}_i\right), \bar{u}_i\right), t_i, t_j\right), \vec{R}_j\right)\right)^T S_j^{-1} \left(\bar{u}_j - \vec{f}\left(\vec{p}\left(\vec{h}\left(\vec{g}\left(\vec{p}(\vec{x}_1, t_1, t_i), \vec{R}_i\right), \bar{u}_i\right), t_i, t_j\right), \vec{R}_j\right)\right)} \quad (3)$$

## 2.2 Extension to Other Measurements

Angles-only and angles with angle rates measurements are possible types of observations with electro-optical systems. The approach can also be extended to radar measurements which return high accuracy values for range and range rate with lower accuracy estimates for the relative angles. Due to the higher uncertainty angle measurements, it would be

reasonable to collect three observations or more to generate a state. However, only two measurements are required to generate a state. The method is similar to the angles-only case. There are 2-permutations of  $n$  projections where  $n$  is the number of measurements. The solution space spans over the angle rates or velocity components orthogonal to the displacement vector between the target and observer. Other than the methods needed to project the state space into the measurement or solution space, the fundamental approach to performing the algorithm with different types of measurements is unchanged.

### 2.2.1 Admissible Region

With only the cost function as described in the previous section, any optimization would be unconstrained and suffer from the associated penalties to computational efficiency as the optimizer attempts to search an unconstrained solution space. By applying known limits to the feasible solution states as a penalty function, the search space can be constrained and the algorithm can receive the corresponding benefits of a smaller search space. The most basic limits that can be implemented are restrictions directly on the feasible values of the solution space. For measurements of angles and angle-rates these would be limits on the feasible range and range-rate between the observer and target. The range can be constrained by a combination of the known capabilities of the observing instrument and the maximum observable cross section of the target as a target sufficiently small and far away eventually becomes undetectable. The range-rate can be constrained by assuming physical limits on the velocity of the object depending on the types of objects intended to be characterized. For example, if the target is assumed to be bound to the Earth-Moon system, range-rate can be limited by the Earth escape velocity. If the measurements are angles-only, then the entire velocity component of the state is a part of the search space, and a similar limit can be applied.

The next type of constraint that can be applied as a penalty function are those that would have prevented detection in the first place. This includes line-of-sight constraints where detection would have been impossible due to the Earth or Moon occluding the target's position from the observer. The line-of-sight constraint can be implemented directly from the geometry of the system bodies and positions of the observer and estimated target position. Similarly, for visible-spectrum optical data, in most cases it can be assumed that the light coming from a target object that enables detection is light reflected from the Sun. Accordingly, a target object in the Earth or Moon's umbra would not be visible. Another geometric constraint can be applied for this case to remove solutions that would place the target object in the Earth or Moon's umbra.

Constraints regarding the orbit of the shape of the orbit itself can also be applied, but there is a cost-benefit to consider in their implementation for the cislunar domain. Under the two-body assumption, the orbits of a target object can be assumed to be conic sections and for non-escaping objects, the orbits can be assumed periodic. This is not necessarily the case for cislunar objects and other target objects acted upon by significant forces other than the Earth's gravity. Without the two-body assumption, constraining a target object's orbit based on future and past states requires a full propagation of the orbit to some user-selected time in the past and future to detect violation of orbital constraints. Under the two-body assumption, cases such as Earth-impacting orbits, highly eccentric orbits, or escaping orbits can be detected directly from the orbital elements that can be derived from a single full state estimate. These penalties can still be applied to objects under additional forces, but the implementation must balance the time-cost of an expensive propagation into the past and future to catch violation of these limits and the benefit of a more tightly constrained search space. An optimal solution will depend on assumptions regarding the target object and goals for the implementations speed of convergence.

Turning these various constraints into penalty functions first requires representing violation of these constraints as continuous values rather than Boolean. For search space limits, the violation can be simply implemented as the amount that the estimate violates the condition. For a line-of-sight penalty, the cost can be implemented as equation 4 where  $\vec{r}$  is the estimated position state of the target,  $\vec{r}_B$  is the position state of the occluding body,  $\hat{p}$  is the pointing vector from the observer to the target, and  $R_B$  is the radius of the occluding body. For an umbra penalty, if the angle between the body-relative position state vector of the target and the target-relative position state vector of the Sun is less than ninety degrees, the penalty expressed in equation 5 can be applied. Otherwise, the target object is necessarily not in the umbra and the body. The  $\vec{r}_{\odot}$  is the relative position vector of the Sun to the Earth.

$$f = \max \left[ 0, -(\vec{r} - \vec{r}_B) \cdot \hat{p} - \sqrt{(\vec{r} - \vec{r}_B)^2 - R_B^2} \right] \quad (4)$$

$$f = R_B^{-2} \max [0, R_B - \|(\vec{r} - \vec{r}_\odot) \times (\vec{r} - \vec{r}_M)\|] \quad (5)$$

Each of these penalties may be then be weighted by a multiplicative factor to account for the different scales of the units involved. For example, an angle penalty of one radian is far more severe than a position penalty of one kilometer and the weights can account for these differences. After the penalties are balanced, they can be added together and the penalty can be applied as  $\exp(\text{penalty}) - 1$  and added to the cost function.

### 2.3 Nelder-Mead Optimization

The optimization is performed with a Nelder-Mead algorithm. This is a gradient-free algorithm that moves the vertices of a simplex each iteration to find a minimum of the cost function [7]. For measurements of angles and angle-rates, the search space is two-dimensional so the simplex is a triangle. For angles-only measurements, the search space is four-dimensional, so the simplex is the 4-simplex. The Nelder-Mead optimizer is chosen because, it has been shown that for angles and angle-rates measurements applied to two-body dynamics, the Nelder-Mead algorithm outperforms a particle swarm in both speed and accuracy [5]. The Nelder-Mead algorithm also works well for low-dimension problems which applies for these two and four-dimensional searches [8].

There is no guarantee of uniqueness for the initial orbit determination solution under the additional dynamics of the cislunar domain. Accordingly, rather than searching for a single minimal solution, the goal of the optimization is to find the full set of feasible solutions for the given set of measurements. A set of solutions can then be further narrowed down by collecting additional measurements and eliminating incompatible states. In order to do this, the algorithm is run multiple times for a given set of measurements and the set of distinct feasible minima are reported. Two methods for initializing the Nelder-Mead algorithm are compared in terms of their effectiveness at providing a set of solutions that contain a close estimate to the true state. One is random sampling of the admissible region for generating initial simplex vertices. The other is dividing the admissible region into separate regions and initializing the simplex in each region once.

By constraining the initial simplex vertices to lie within the admissible region, it is known that the algorithm starts with a set of unpenalized hypotheses. This guarantees that the algorithm can converge to a state within the admissible region. The admissible region is defined by the physical limits of the system, so as long as the user-defined limits are valid, the solution will exist within the admissible region.

### 2.4 Orbit Instability

With the addition of the influence of the Moon's gravity, the dynamics in some orbital regimes become chaotic. Small differences in the real and measured initial state grow with time to that point that the propagated predicted state no longer well-approximates the truth. As a consequence, for these chaotic orbits, there exists a time horizon beyond which an initially accurate state estimate cannot predict the actual physical state. This time horizon can be characterized by a Lyapunov exponent [9]. The Lyapunov exponent for a given orbit can be approximated with equation 6. Where  $t_0$  and  $t_1$  are an initial and final time,  $\Phi(t_1, t_0)$  is the state transition matrix from  $t_0$  to  $t_1$ , and the matrix norm is the spectral norm [10]. An in-depth discussion of this approach is available in Ren, Li, and Zheng's 2020 paper.

$$\lambda = \frac{1}{t_1 - t_0} \ln \|\Phi(t_1, t_0)\| \quad (6)$$

### 2.5 Determining Convergence

For a single initialization of the Nelder-Mead algorithm, it will either exit when the maximum number of iterations is exceeded or it converges to a minimum and sequential best vertices on the simplex change by less than some tolerance. For a solution to be feasible, it should be close to zero since a state that perfectly predicts the observed measurements will have a zero cost. Since the cost function is the sum of Mahalanobis distances, it accounts for measurement error so long as the measurement covariance matrix accurately reflects the actual measurement error. If the returned solution is large, it can be immediately rejected as a local minimum but not global minimum. The question then becomes, what is close enough to accept as a feasible solution. The simplest approach would be to consider the cost purely as the sum of Mahalanobis distances and set a threshold based on hypothesis testing and a target false-positive rate. However, this does not account for the impact of the propagation of uncertainty. To include that, the propagator can be modified to find the state transition matrix by computing the Jacobian of the dynamics at each step in the integration. The resulting

state transition matrix can be applied to a vector of the error from the propagated state at the propagation start time to find the error from the dynamics at the propagation end time. This can be used to analytically express the cost function as a function of the initial state estimate and the measurements.

Let  $\hat{x}^*$  be the true dynamical state for the set of measurements, and let  $\vec{k}$  be the vector of measurements with some covariance matrix,  $P_{k,k}$ . The calculated cost function can then be written as  $J$ , a scalar function of  $\hat{x}^*$  and the mean of the measurement vector  $\vec{\mu}_k$ . Then one can find the hessian of the cost function with respect to the six measurements for angles-only data or eight measurements for angles and angle-rates. With the hessian of the cost function, it becomes feasible to take a second order Taylor expansion of the cost function. That second order expansion is expressed in equation 7. Start with the Taylor expansion of the cost function  $J$  around the solution. Since the cost function is assumed to be at the solution,  $J$  and  $\frac{\partial J}{\partial \vec{x}} \Big|_{\hat{x}^*} \delta \hat{x}^*$  should be zero. Then just looking at the expansion with respect to the measurement states simplifies to equation 8. Assuming that the  $\vec{k}$  is normally distributed about  $\vec{0}$ , the mean of  $J$  with respect to  $\vec{k}$  will be the expectation value of the second term [11].

The variance of  $J$  can be found similarly. Starting with the expectation value representation in equation 11, the result can be simplified to 13. The simplification is derived in Holzinger, Scheeres, and Alfriend's 2012 paper.

$$J(\hat{x}^* + \delta \hat{x}^*, \vec{\mu}_k + \delta \vec{k}) = J(\hat{x}^*, \vec{\mu}_k) + \frac{\partial J}{\partial \vec{x}} \Big|_{\hat{x}^*} \delta \hat{x}^* + \frac{\partial J}{\partial \vec{k}} \Big|_{\vec{\mu}_k} \delta \vec{k} + \frac{1}{2} \delta \hat{x}^{*\top} \frac{\partial^2 J}{\partial \vec{x}^2} \Big|_{\hat{x}^*} \delta \hat{x}^* + \frac{1}{2} \delta \vec{k}^{*\top} \frac{\partial^2 J}{\partial \vec{k}^2} \Big|_{\vec{\mu}_k} \delta \vec{k} + \delta \hat{x}^* \frac{\partial^2 J}{\partial \vec{x} \partial \vec{k}} \Big|_{\hat{x}^* \vec{\mu}_k} \delta \vec{k} + \dots \quad (7)$$

$$J(\hat{x}^*, \vec{\mu}_k + \delta \vec{k}) = \frac{\partial J}{\partial \vec{k}} \Big|_{\vec{\mu}_k} \delta \vec{k} + \frac{1}{2} \delta \vec{k}^{*\top} \frac{\partial^2 J}{\partial \vec{k}^2} \Big|_{\vec{\mu}_k} \delta \vec{k} = \nabla J \delta \vec{k} + \frac{1}{2} \delta \vec{k}^{*\top} J_{k,k} \delta \vec{k} \quad (8)$$

$$\mu_J = \mathbb{E} \left[ \frac{1}{2} \delta \vec{k}^{*\top} J_{k,k} \delta \vec{k} \right] \quad (9)$$

$$\mu_J = \frac{1}{2} \text{Tr} (J_{k,k} P_{k,k}) \quad (10)$$

$$P_{j,j} = \mathbb{E} \left[ \left( J(\delta \vec{k}) - \mu_j \right)^2 \right] \quad (11)$$

$$P_{j,j} = \mu_j^2 - \mu_j \text{Tr} (J_{k,k} P_{k,k}) + \frac{1}{4} \text{Tr} (J_{k,k} P_{k,k})^2 + \frac{1}{2} \text{Tr} (J_{k,k} P_{k,k} J_{k,k} P_{k,k}) \quad (12)$$

$$P_{j,j} = \frac{1}{2} \text{Tr} (J_{k,k} P_{k,k} J_{k,k} P_{k,k}) \quad (13)$$

With the mean and variance of the cost about the solution found, the necessary condition for convergence is that the proposed solution has a cost that belongs to this distribution. A low-cost, but imprecise test for this is to decide in advance on a maximum number of standard deviations from expected to accept and only accept a result if the cost falls within that limit as defined by the derived mean and standard deviation. This would only require a single cost estimate, and the algorithm could exit immediately when an acceptable solution is found. A more precise test would involve running the optimizer multiple times and collecting those returned solutions that cluster about a value. A statistical hypothesis test could then determine if these solutions fit the expected distribution.

### 3. RESULTS

#### 3.1 Implementation

In order to assess the capabilities and limitations of the proposed algorithm, the approach is verified against a set of simulated scenarios. These scenarios model both an observer and target with a high fidelity orbital model that includes the gravitational influence of the non-dwarf planets, the Sun, and the Moon. As well as the geopotential model of the Earth up to order 360 based on the EGM2008 model [3] and the geopotential model of the Moon up to order 165 based on the GRGM1200A model [4]. Next, the simulated orbits included the influence of the solid tides of both the Earth and Moon. There are also non-gravitational forces included, these are the atmospheric drag and solar radiation pressure on the satellites.

The dynamics of the system used in the IOD algorithm itself were simplified beyond this model with various forces toggle-able to measure the costs and benefits of simplifying the model in various ways. The results of this investigation are discussed in section 3.3. The forces included in the IOD algorithm propagator are the point-mass influence of the non-dwarf planets, Sun and Moon and the geopotential model of the Earth and Moon up to order four.

The test cases were randomly generated along specific orbital regimes. Each observer regime and target regime pair was generated with two hundred cases. The observer regimes are equatorial geosynchronous orbit (GEO), equatorial medium Earth orbit (MEO), equatorial low Earth orbit (LEO), and L1 halo orbits both in and out of plane. The target regimes are a lunar near rectilinear halo orbit (NRHO) based on the planned lunar gateway, L1 halo orbits both in and out of plane, L2 halo orbits both in and out of plane, Molniya orbits, and highly eccentric Earth orbits (HEO). Not every pair of observer and target is considered. Rather interesting cases are chosen from the set of generated test cases.

For each scenario, simulated observations are generated for different time steps between observations and with different levels of measurement noise added in. The measurement noise is Gaussian and symmetric on the unit sphere. Trial cases can be selected by choosing a regime scenario, time interval between measurements, and measurement error.

### 3.2 Cost Function Structure

First, the structure of the solution space passed through the proposed cost function is explored. Previous work has shown that a similarly derived cost function applied solely to the near-Earth domain under two-body dynamics can yield a smooth and continuous contour within the admissible region when the measurements include both angles and angle rates [5]. In cislunar space, the more complex dynamics of the additional forces create additional opportunities for issues to arise with the search space. This could potentially result in multiple global minima for a given scenario or an unfavorable geometry for optimization. In order to determine if these issues occur, contours are generated in the search space for several orbit scenarios. These contours show the behavior of the cost function as a function of the trial state.

Cases for the angles-only scenarios are generated. This has a four-dimensional solution space, and contours can only clearly show the response to changes in two dimensions. As a result, angles-only contours are generated as a series of slices through the solution space in order to show the behavior of the cost function along all four search dimensions.

#### 3.2.1 L4

The contours for the a target object in a stable orbit around the L4 Lagrange point have a relatively simple structure. The contours for angles-only measurements with the correct velocities orthogonal to the line-of-sight direction are shown in figure (1). The structure of the contour shows a long valley of low cost for increasing range. This appears due to the ambiguity in the range state as three angles-only measurements taken in relatively quick succession for a distant object have a limited ability to discern the range. Near to the observer, several local minima appear. These represent nearby orbits that could approximate the observed measurements, though they do not perfectly match as the costs do not fall to zero.

Deviations to the line-of-sight orthogonal velocities result in the contours shown in figures (2) through (3). As the orthogonal velocities move away from truth, the cost minima do not fall as deep and the main valley containing the true solution shifts. This indicates that the contour should be optimizable, though there remains the potential for an optimizer to return local minima if initialized to close to another minimum.

#### 3.2.2 Low Lunar Orbit

Figure (4) shows the range and range rate dimensions of the angles-only contour for a target object in low Lunar orbit as seen by an observer in GEO. The other two states are set to the true values. The red dot indicates the location of the true solution. The contours shown are the logarithm of the cost.

Figures (5) through (6) show the contour angles-only contours with the non-plotted dimensions shifted slightly. This exposes the behavior of the solution space along the velocity axes perpendicular to the line of sight. The region of high cost near the solution reflects Moon-impacting orbits that trigger additional penalties on the cost function. There are also additional local minima in the nearby region such as the region at a higher range on the opposite side of the




Fig. 1: Angles-only GEO observer to L4 target log(cost) contours with correct line-of-sight perpendicular velocities.




Fig. 2: Angles-only GEO observer to L4 target log(cost) contours with out of plane velocity increased by 0.05 km/s.




Fig. 3: Angles-only GEO observer to L4 target log(cost) contours with in plane velocity increased by 0.05 km/s.

Moon-impacting lobe. This indicates that there is a significant potential for an optimizer to return minima that are not the solution. These minima also have a relatively low cost, so a relatively low amount of measurement error could make the minima indistinguishable in terms of detecting which is the global minimum and true solution. Comparing

figure (4) to figures (5) through (6) shows how these contours shift in the non-plotted dimensions. The nearby states are similar with slight shifts to the shapes of the contours and an increased cost around the solution minimum. This indicates that an optimization algorithm should be expected to fall into the solution if in the correct region rather than a set of solutions curving through the solution space.




Fig. 4: Angles-only GEO observer to low lunar orbit target log(cost) contours with correct line-of-sight perpendicular velocities.




Fig. 5: Angles-only GEO observer to low lunar orbit target log(cost) contours with out of plane velocity increased by 0.05 km/s.

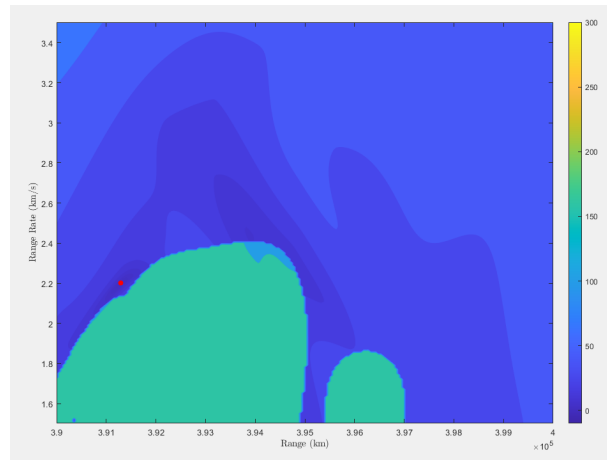


Fig. 6: Angles-only GEO observer to low lunar orbit target log(cost) contours with in plane velocity increased by 0.05 km/s.

### 3.3 Impact of Propagator Fidelity

The orbital propagator used to compare measurements to predictions at different times significantly impacts both the algorithm's computational efficiency and its accuracy. The optimal balance between these parameters will vary based on the goals of the implementation, but understanding the trade-offs with the propagator capabilities is valuable for appropriately choosing the implementation. The propagator is tested by initializing it with the first state in each orbit scenario and measuring the position error and required calculation time after one solar day of motion. Each regime is tested independently to isolate the influence of different propagator elements on different orbital regimes. The final position errors for the full fourth order geopotential model with the inclusion of all major Solar System bodies is shown in figure (7). The distribution of calculation times is shown in figure (8). Most orbits are modeled with less than a single kilometer error after a day of propagation. The halo orbits in particular perform well with the Lagrange point halo orbit position error ending up at around twenty-five meters. The highly elliptical and Molniya orbits have the cases with the largest error. This is due to the unmodeled, higher orders of the geopotential model, solid tides, and atmospheric drag.




Fig. 7: Final position error logarithmic swarm chart for a series of orbital regimes after orbit propagation over a solar day using ODE 113.




Fig. 8: Computation time logarithmic swarm chart for a series of orbital regimes for orbit propagation over a solar day using ODE 113

The previous case used MATLAB's ODE113, a nonstiff differential equation solver. This is compared against ODE45. The position error and time cost for this differential equation solver is provided in figures (9) and (10) respectively. While the accuracy of the two methods are similar, the ODE113 case performed significantly better in terms of the time cost of the propagation. This performance improvement is visible in all tested cases.

Figures (11) and (12) show the results after removing the influence of all bodies save for the Earth, Moon, and Sun. The time cost of the algorithm improved slightly, but there was a significant penalty to the final position state accuracy. Near-Earth MEO orbits were largely unaffected, but near rectilinear, L1, and L2 halo orbits saw over an order of magnitude loss of final state position accuracy to just under single kilometer level error for the Lagrange points and slightly more for the NRHO.

Reducing the order of both the Earth's and Moon's geopotential models to two yields the results shown in figures (13) and (14). The change increased the position error in the orbits closest to Earth by a small amount and had little observable impact on the more distant orbits. The impact was less severe than the loss in accuracy when the influence of other planets was removed. In terms of the computation time, the time cost fell slightly across the board, though not by as much as removing the influence of the planets. Cutting the third and fourth order components of the geopotential model could be a reasonable method for saving computation time in the propagation step without significantly impacting the algorithm's accuracy.

Figures (15) and (16) show the propagation results when only the first order gravitational influence of the Sun, Earth, and Moon is included. While the computation time falls significantly for all cases, so does the position accuracy. All modeled GEO's and most modeled MEOS have final position errors over ten kilometers with some MEO's approaching just under 100 kilometers of error. The Lagrange point halo orbits are the only orbits with error consistently below

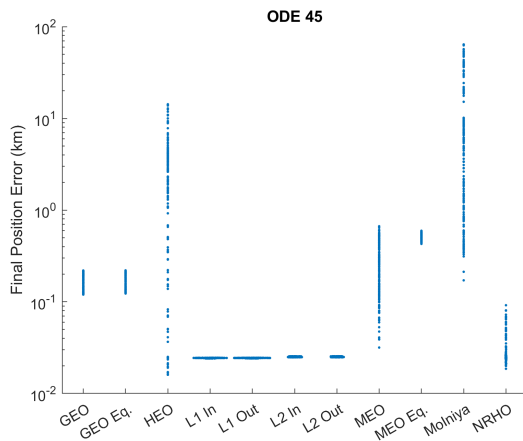


Fig. 9: Final position error logarithmic swarm chart for a series of orbital regimes after orbit propagation over a solar day using ODE 45




Fig. 10: Computation time logarithmic swarm chart for a series of orbital regimes for orbit propagation over a solar day using ODE 45.




Fig. 11: Final position error logarithmic swarm chart for a series of orbital regimes after orbit propagation over a solar day. Propagated without the influence of planets outside the Earth system.

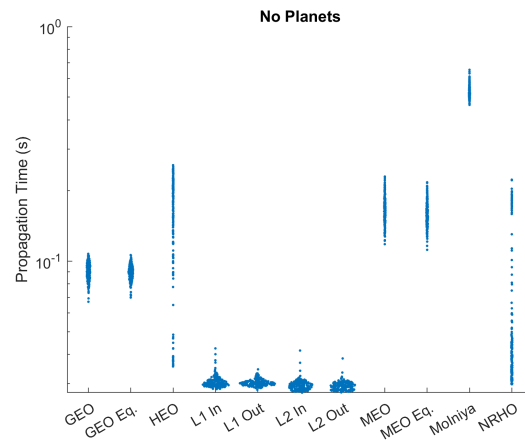


Fig. 12: Computation time logarithmic swarm chart for a series of orbital regimes for orbit propagation over a solar day. Propagated without the influence of planets outside the Earth system.

one kilometer, but still perform far worse than the twenty-five meter error shown in the fourth order model.

### 3.4 Response to Measurement Error

Any measurement of the relative angle between an observer and target will have some amount of measurement error. In order to perform orbit determination on a series of measurements of an object hundreds of thousands of kilometers away, the algorithm applied must be capable of handling some level of measurement error. The response of the proposed algorithm to measurement error is an important metric for assessing its capabilities under realistic conditions. Simulated scenarios are generated with a high-fidelity orbit propagator for several orbital regimes.

Each regime has two hundred cases randomly generated within the defining bounds of the regime. Due to the large distances involved in orbit determination for cislunar space, the scenarios are generated with two different amounts of measurement error standard deviations, 0.1 arcseconds and 1 arcsecond. The errors are Gaussian distributed about the truth, and the error ellipse is a circle projected onto the unit sphere. The time delay between each simulated observation is six hours. The Nelder-Mead algorithm is set to converge to a tolerance between the minimum vertex cost and the maximum vertex cost of 0.01, and the Nelder-Mead algorithm is set to reinitialize up to 10 times or until a result has

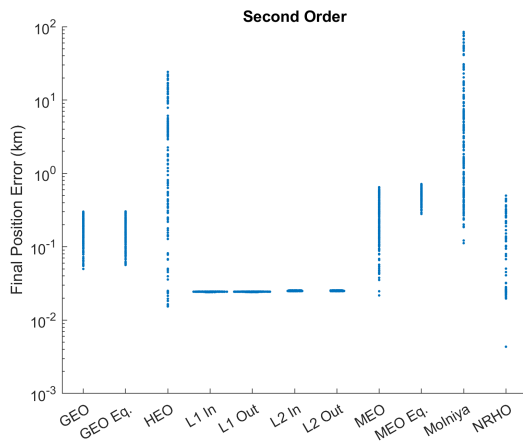


Fig. 13: Final position error logarithmic swarm chart for a series of orbital regimes after orbit propagation over a solar day. Propagated with second order geopotential models of the Earth and Moon.




Fig. 14: Computation time logarithmic swarm chart for a series of orbital regimes for orbit propagation over a solar day. Propagated with second order geopotential models of the Earth and Moon.




Fig. 15: Final position error logarithmic swarm chart for a series of orbital regimes after orbit propagation over a solar day. Propagated with only the point source influence of the Sun, Moon, and Earth.




Fig. 16: Computation time logarithmic swarm chart for a series of orbital regimes for orbit propagation over a solar day. Propagated with only the point source influence of the Sun, Moon, and Earth.

a minimal cost below 1. This is a much tighter constraint than the developed necessary condition for convergence to a potential solution, and considering that the cost is a sum of Mahalanobis distances, a cost below 1 is indistinguishable from the truth. This is because if the sum of Mahalanobis distances is less than one, then at most, a single projection contributes one to the sum. This would imply that a single projection differs from expected by one standard deviation, and an error of only one standard deviation could reasonably be expected for an accurate estimate. This constraint allows the algorithm to exit early when a likely solution is found. If no run of the Nelder-Mead algorithm yields a result with cost below one, the lowest cost solution is returned.

The cumulative distribution functions for the estimate errors are shown in figures (17) through (20). Figures (17) and (18) show the 0.1 arcsecond measurement error case for both position and velocity initial state estimate errors respectively while figures (19) and (20) show the position and velocity errors for the 1 arcsecond case. The five plotted regimes are for a geostationary observer to a L1 halo in and out of plane orbit and a Lunar Gateway inspired near rectilinear halo orbit. As well as L1 halo in plane to L2 halo out of plane and L1 halo out of plane to L2 halo in plane.

With a measurement error of 0.1 arcseconds, 52% of GEO to L1 halo in plane, 63% of GEO to L1 halo out of plane,

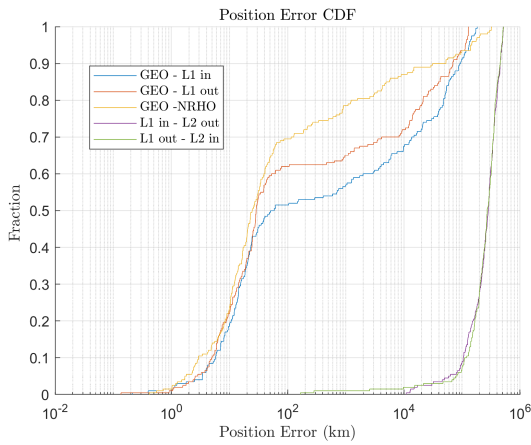


Fig. 17: Cumulative distribution function of the position errors under 0.1 arcsecond measurement error, 10 max reinitializations, and a convergence tolerance of 0.01.




Fig. 18: Cumulative distribution function of the velocity errors under 0.1 arcsecond measurement error, 10 max reinitializations, and a convergence tolerance of 0.01.

and 70% of GEO to NRHO cases returned estimates within 100 kilometers of truth and after this point there is a clear point of inflection where the results transition from near estimates of the truth and cases where it never approached the truth. The two regimes observing L2 points performed poorly with only one near estimate to the truth. These results are mirrored nearly exactly before the point of inflection in the velocity distribution. This is because these accurate position estimates tended to also have accurate velocity estimates as the full state is estimated.

For the case with measurement angle standard deviation of 1 arc second, figures 19 and 20 show a reduced capability to reliably converge to the solution. The GEO to L1 in plane case has 29% below 100 kilometers while GEO to L1 out of plane and GEO to NRHO are at 36% and 39% respectively. The L2 cases remain inaccurate, and the general behavior follows the same overall trends. The algorithm shows an expected response to measurement error where measurement error is inversely correlated with estimate accuracy.




Fig. 19: Cumulative distribution function of the position errors under 1 arcsecond measurement error, 10 max reinitializations, and a convergence tolerance of 0.01.

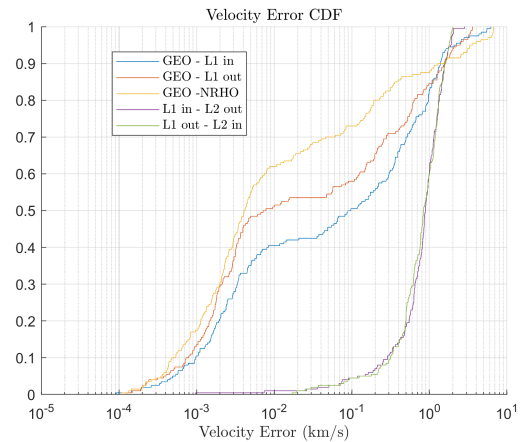


Fig. 20: Cumulative distribution function of the velocity errors under 1 arcsecond measurement error, 10 max reinitializations, and a convergence tolerance of 0.01.

### 3.5 Impact of Optimizer Parameters

The Nelder-Mead optimizer has several customizable parameters that impact the algorithm's performance in both speed and accuracy. Since multiple minima are expected in the cost function, the algorithm is also run multiple times for each case so more of the cost function can be explored in case early runs fall into local minima. Both the number

of reinitialization and the convergence tolerance for determining when to exit a run are explored for their impact on the algorithm's accuracy. From the contour plot study, it is known that in some orbital scenarios the main cost valley is long and narrow, so a convergence tolerance that is too high can lead to exiting a run too early while a tolerance too low can waste time when a solution is already found. The investigated parameters are the convergence tolerance and the number of reinitializations. Since the appropriate balance between time cost and accuracy depends on the implementation goals, these results are intended to show the trade-offs in choosing these values rather than provide a single, universal solution.

Figures (21) and (22) show cumulative distribution functions for the position errors for cases with five and twenty reinitializations respectively. Contrast these with figure (17) which shows the results for ten reinitializations. The case with five reinitializations shows a sharp drop-off in the algorithm's accuracy. The GEO to L1 halo in plane regime has 31% of solutions within one hundred kilometers. The GEO to L1 halo out of plane and GEO to NRHO have 42% and 58% respectively. By contrast, the case with twenty reinitializations shows an improvement in accuracy and a convergence of the three cases observed from GEO. The percent of results under one hundred kilometers from the true initial position for GEO to L1 halo in plane, GEO to L1 halo out of plane, and GEO to NRHO are 73%, 78%, and 75% respectively. The cases observing L2 orbits from L1 orbits remain inaccurate.




Fig. 21: Cumulative distribution function of the position errors under 0.1 arcsecond measurement error, 5 max reinitializations, and a convergence tolerance of 0.01.




Fig. 22: Cumulative distribution function of the position errors under 0.1 arcsecond measurement error, 20 max reinitializations, and a convergence tolerance of 0.01.

The other investigated optimizer parameter is the convergence tolerance. This is implemented as a difference between the maximal and minimal cost of a simplex's vertices. When this falls below the given tolerance, the Nelder-Mead optimizer exits the run and keeps the minimal vertex as the solution. Tolerances of 0.1, 0.01, and 0.001 are investigated. The results for the case with a tolerance of 0.01 are shown in figure (17) while the results for tolerances of 0.1 and 0.001 are shown in figures (23) and (24) respectively. A smaller tolerance will require more iterations to converge, but may converge to a value closer to the truth. This is shown in these results. When the tolerance is 0.1, the point of inflection in the plot moves toward higher position error states. This is due to a subset of Nelder-Mead runs falling into the appropriate cost well, but exiting before they reach the bottom due to the looser constraint. Similarly, with the tighter constraint of 0.001, the point of inflection moves toward smaller position errors. It becomes just above ten kilometers in error rather than a bit above one hundred kilometers as in the case with a tolerance of 0.01. The portion of near-truth results also increases for the tighter cost case and decreases for the looser cost case.

In all assessed cases, the scenarios with observers in L1 orbits and targets in L2 orbits fail to converge to accurate state estimates. This issue persists even if the number of reinitializations is increased to fifty and the full set of solutions is saved. Figure (25) shows the results for such a case with the closest to truth solution from the returned set of fifty kept as the case's solution. However, despite the failure to return a set of solutions containing the truth, when the algorithm is manually initialized near the truth it converges to it. The dynamics of this regime complicate the search. Orbits about Lagrange points one and two are unstable, and small differences in the dynamics for the truth and estimate can lead to divergent results as time increases. In addition, the solution space has several minima deep enough to appear

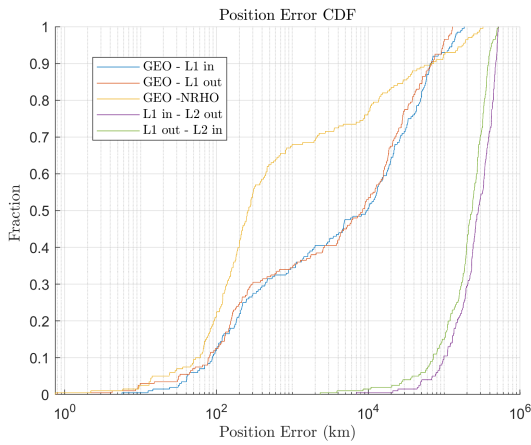


Fig. 23: Cumulative distribution function of the position errors under 0.1 arcsecond measurement error, 10 max reinitializations, and a convergence tolerance of 0.1.




Fig. 24: Cumulative distribution function of the position errors under 0.1 arcsecond measurement error, 10 max reinitializations, and a convergence tolerance of 0.001.

indistinguishable from the truth, and the region around the true state in the solution space is small enough that even with fifty runs, the algorithm is more likely to miss the true state than to capture it. Since the other solutions that are discovered have sufficiently low cost, these states correspond to other orbits that could feasibly produce similar measurements. Accordingly, eliminating these solutions would require ruling out the corresponding states, despite being otherwise feasible. Without a priori knowledge that can restrict the search space further, these other solutions keep the results of a search in this regime ambiguous.

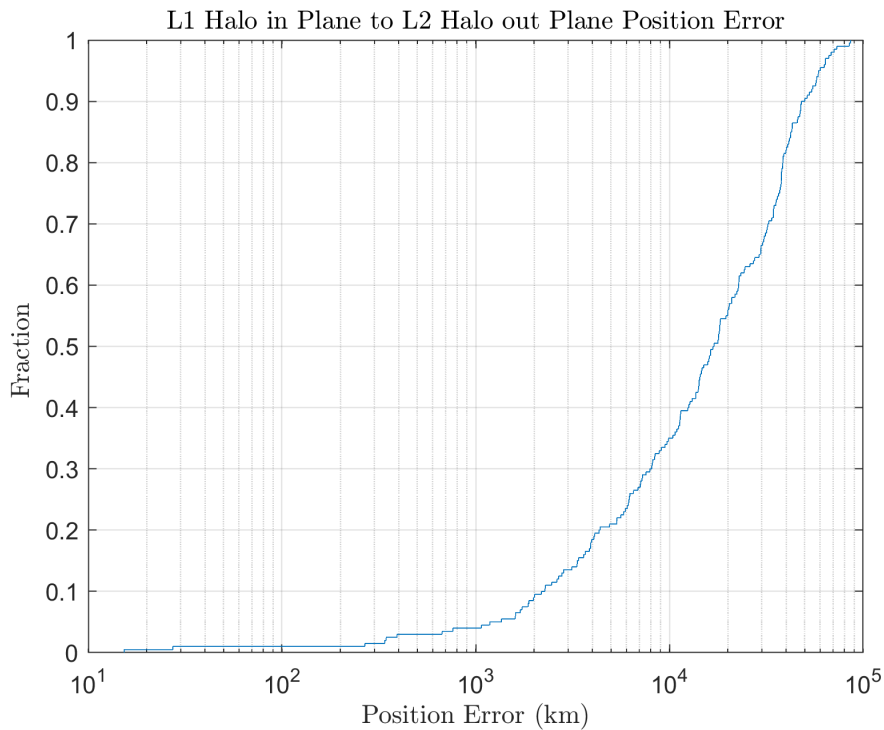


Fig. 25: CDF of the best solution over 50 reinitializations for 200 cases of L1 halo in plane observers to L2 halo out of plane targets.

## 4. CONCLUSION

An optimization-based algorithm for performing initial orbit determination under arbitrary dynamics has been developed and its accuracy assessed under simulated scenarios. In addition, an analytical method for determining if a proposed solution satisfies necessary but not sufficient conditions to be a solution for a given set of measurements has been developed. To explore the cost-benefit ratio of several parameters in the dynamics, the time-cost and propagated state accuracy were assessed. This was done by propagating the orbits of randomly generated objects in various regimes under several simplifications to the dynamics. The results were compared against a high-fidelity propagation and it was discovered that a second order geopotential model for the Earth and Moon with the dynamics of all major Solar System bodies included represented a good balance between propagation accuracy and computational cost.

Contours of the cost function of the algorithm were then developed for a few interesting cases, and their geometry was explored for the impact that they may have on the optimization. It was discovered that the contour had several local minima which needed to be accounted for the algorithm implementation. The algorithm was then run over randomly generated scenarios within a set list of orbital regimes. The behavior of the algorithm in response to adjustments to optimizer parameters as well as measurement error was explored yielding a series of plots showing the relative success of the algorithm subject to these adjusted scenarios. With low measurement error and sufficiently generous optimizer parameters, the proposed IOD algorithm can return a close estimate to the truth as the best fit over 70% of the time. Future work includes exploring the capabilities and limits of the proposed algorithm when returning a full set of the discovered feasible solutions and determining reasonable limits on the number of runs needed to reliably return a set containing the true state.

## REFERENCES

- [1] David A Vallado. *Fundamentals of astrodynamics and applications*, volume 12. Springer Science & Business Media, 2001.
- [2] Charles H Acton Jr. Ancillary data services of nasa's navigation and ancillary information facility. *Planetary and Space Science*, 44(1):65–70, 1996.
- [3] Nikolaos K Pavlis, Simon A Holmes, Steve C Kenyon, and John K Factor. The development and evaluation of the earth gravitational model 2008 (egm2008). *Journal of geophysical research: solid earth*, 117(B4), 2012.
- [4] S Goossens, FG Lemoine, TJ Sabaka, JB Nicholas, E Mazarico, DD Rowlands, BD Loomis, DS Chinn, GA Neumann, DE Smith, et al. A global degree and order 1200 model of the lunar gravity field using grail mission data. In *Lunar and Planetary Science Conference*, number 1903, page 1484, 2016.
- [5] Sam Wishnek, Marcus J Holzinger, Patrick Handley, and Sue Hagerty. Robust initial orbit determination using streaks and admissible regions. *The Journal of the Astronautical Sciences*, pages 1–42, 2021.
- [6] Johnny L Worthy III, Marcus J Holzinger, and Kohei Fujimoto. Optical sensor constraints on space object detection and admissible regions. *Paper AAS*, pages 13–707, 2013.
- [7] John A Nelder and Roger Mead. A simplex method for function minimization. *The computer journal*, 7(4):308–313, 1965.
- [8] Lixing Han and Michael Neumann. Effect of dimensionality on the nelder–mead simplex method. *Optimization Methods and Software*, 21(1):1–16, 2006.
- [9] John Michael Tutill Thompson and H Bruce Stewart. *Nonlinear dynamics and chaos*. John Wiley & Sons, 2002.
- [10] Jing Ren, Mingtao Li, and Jianhua Zheng. Families of transfers from the moon to distant retrograde orbits in cislunar space. *Astrophysics and Space Science*, 365(12):1–21, 2020.
- [11] Marcus J Holzinger, Daniel J Scheeres, and Kyle T Alfriend. Object correlation, maneuver detection, and characterization using control distance metrics. *Journal of Guidance, Control, and Dynamics*, 35(4):1312–1325, 2012.
- [12] John E Prussing and Bruce A Conway. *Orbital mechanics*. Oxford University Press, USA, 1993.
- [13] Bob Schutz, Byron Tapley, and George H Born. *Statistical orbit determination*. Elsevier, 2004.