

An Adaptive Approach to Cislunar Initial Orbit Determination using Machine Learning

**Juan Ojeda Romero, Wayne Schlei, Gene Whipps,
Nick LaFarge, Gunner Fritsch, and Sean Phillips**
Johns Hopkins University Applied Physics Laboratory

September 20, 2024

ABSTRACT

Initial orbit determination (IOD) is an increasingly relevant problem in the cislunar regime, where chaotic dynamics degrade the performance of classical IOD approaches. In this investigation, a framework for an end-to-end cislunar IOD process is presented that incorporates angles-only observations that simulate a chance detection scenario. Enlisting machine learning techniques to assist the complicated cislunar IOD process, this work presents the Machine Classifier for Cislunar Orbit Determination (MCCLOD) model that employs a neural network for infusing information about known multi-body dynamics structure about the circular restricted three-body problem (CRTBP) into the initial state estimation process. The novel MCCLOD IOD process is compared directly with a classical two-body IOD approach (Gooding) for Earth-Moon L_1 and L_2 halo orbit examples in a high-fidelity dynamics environment. Two neural network (NN) models are implemented with regression and orbit classification to identify a 6D state estimate given angles-only measurements. The resulting simulations indicate drastic improvement in both accuracy (MCCLOD demonstrates at best two orders of magnitude improvement in positional error performance) and batched least-squares convergence consistency. Although a “classification with regression” NN degrades overall MCCLOD IOD performance, simulations indicate that a hybrid classification NN followed by a regression NN framework yields low position error in tested cislunar IOD problems.

1. INTRODUCTION

The proliferation of artificial objects is underway in cislunar space. With increasing global interest in the cislunar regime and accessing lunar resources, both cooperative and non-cooperative vehicles (i.e., spacecraft or objects that share and withhold their ephemerides or measurements, respectively) are quickly populating the space as access to space becomes cheaper. As such, the need for cislunar Space Situational Awareness (SSA) beyond GEO is surging, particularly for detection, Orbit Determination (OD) and tracking capabilities. Furthermore, the large distances (>300,000 km) and faint lighting conditions, relative to a sensor, inherent to the cislunar regime present significant challenges in incorporating optical measurements in the OD process. This investigation addresses this limitation by leveraging a neural network to bootstrap an Initial Orbit Determination (IOD) operation; thus enabling OD despite distant observations and highly nonlinear dynamics.

Given a set of measurements that represent a detection event, a typical OD process proceeds by performing an IOD process, followed by a filter to achieve a suitable OD solution. Classical IOD strategies based on two-body motion (Laplace, Gauss, Gibb’s, Gooding, etc.) are well defined in the vicinity of the Earth and provide an adequate guess to begin the sequential-filtering process, e.g., a Kalman filter. Cislunar applications of these classical IOD methods, however, provide limited success after significant tuning due to the additional gravitational sensitivities induced by the multi-body gravity dynamics (usually modeled with the Circular Restricted Three Body problem or CRTBP) [17, 8]. In fact, six-dimensional state convergence testing with these methods presents non-trivial challenges due to the complex multi-body dynamics where nonlinear sensitivities greatly amplify uncertainty and state errors downstream. One approach to improve IOD in the cislunar regime is to directly employ the knowledge of known multi-body dynamical structure [4], but algorithm strategies employing such structure need improvement in order to merge in numerically sensitive geometry in an effective manner.

Initial orbit determination problems in cislunar space are currently being addressed by advanced algorithms under active investigation. Current cislunar IOD methods include advanced non-linear filtering techniques (such as particle

or Gaussian-mixture filters) [5]. These methods typically incorporate probabilistic admissible regions (or PARs) to form an initial orbital guess and are often constrained by assumptions such as maximum range or orbital energy [2, 7]. While useful, these methods still mostly suffer from the same two-body dynamical assumption, and require a special nonlinear filter to be initialized. Other IOD methods rely on an over-saturation of observation data with multiple phenomenologies [5, 3] to address association and isolation in addition to a state estimate. The common occurrence case of a chance, dim-object detection during a nightly scan, however, yields only a small amount of single-type observation data (typically angles-only for a less than eight-hour pass). Some research approach this type of scenario with PARs [2, 3] or with optimal control theory [9]; other research has also focused on reformulating the cislunar IOD process into an optimization problem albeit yielding the multiple locally-optimal solution basins which requires more investigation [23]. Fortunately, precedence exists in various fields to approach complex state estimation problems with artificial intelligence (AI) and machine learning (ML) that handle previously intractable setups with improvement [14, 17, 1]. Current research applications on the cislunar IOD process have implemented existing ML techniques, e.g., supervised learning algorithms with Neural Networks (NN); however, their efficacy is dependent on the quality of the training datasets, assumptions on problem inputs (i.e., observation type), approach (as in learning dynamics versus classification) as well as the particular procedure. As one example, a Physics-Informed NN (or PINNs which learns dynamics) approach to IOD provided by Scorsoglio et al. [17] showed that the assistance of AI provided an improved outlook towards halo orbit IOD over classical IOD, but the problem is still the difficult sensitivities that generate a strong divergence rate in the filtering process. In the absence of an analytic IOD approach, this paper proposes a state-of-the-art cislunar IOD scheme that benefits from both a knowledge of the non-Keplerian cislunar dynamics and existing machine learning techniques, gathering the benefits of multi-body orbital mechanics understanding with a NN modeling tactic.

In this investigation, a framework for an end-to-end cislunar IOD process is demonstrated that incorporates known orbital structure of the CRTBP with angles-only observations to emulate IOD solutions for a chance detection scenario. Current traditional methods inherit insights from the Two-Body Problem, however, when considering cislunar objects, these approaches become expensive guess-and-check iterations with limited success. In light of this, one of the objectives of this investigation is to provide a six-dimensional state and covariance matrix by training a NN model—named the Machine Classifier for Cislunar Orbit Determination (MCCLOD) model—with periodic orbit information using the dynamics associated with the Earth-Moon CRTBP. Furthermore, some cislunar IOD methods make anticipatory assumptions, such as the availability of space-based observers (which can reduce the IOD problem to triangulation in some cases) or the practical use of niche measurement models [10, 22, 19]. In an effort to produce an IOD tool with a relatively wide applicability, the MCCLOD model operates on angles-only measurements from any Earth-based observer. Datasets based on periodic orbit families (specifically the Earth-Moon CRTBP L_1 and L_2 halo families) train the MCCLOD regression and classification solution NN to provide the appropriate mapping function for cislunar dynamics. For a single measurement representing a small portion of a detection, MCCLOD delivers a two-parameter vector (η and τ) that maps an optical observation data (angles and angular rates) to a unique orbit on a periodic orbit family and a point along the specified orbit, respectively. Thus, the MCCLOD model is a product of existing machine learning techniques and knowledge of complex multi-body dynamics to address the non-trivial cislunar IOD problem.

2. BACKGROUND

2.1 Circular Restricted Three-Body Problem

The Earth-Moon CRTBP dynamics model provides a higher fidelity representation of the cislunar environment from the traditional two-body approximation. Three assumptions are implemented in the derivation: the motion of a massless spacecraft is affected by the gravitational force of the Earth and the Moon, the motion of the isolated Earth-Moon system is assumed to be Keplerian, and its orbit is restricted to a circular orbit. The motion of the spacecraft is defined in a rotating reference frame where: \hat{x} is measured from the Earth to the Moon, \hat{z} is in the direction of the angular momentum vector, and \hat{y} completes the right-handed triad. The Equations of Motion (EOMs) that correspond to the CRTBP are written as,

$$\ddot{x} - 2\dot{y} = \frac{\partial U^*}{\partial x}, \quad \ddot{y} + 2\dot{x} = \frac{\partial U^*}{\partial y}, \quad \ddot{z} = \frac{\partial U^*}{\partial z}, \quad (1)$$

where U^* is the corresponding pseudo-potential defined as,

$$U^* = \frac{1}{2}(x^2 + y^2) + \frac{(1-\mu)}{r_{13}} + \frac{\mu}{r_{23}}, \quad (2)$$

with the dimensionless mass parameter of the system defined as $\mu = m_{Moon}/(m_{Moon} + m_{Earth})$. The vectors \bar{r}_{13} and \bar{r}_{23} correspond to the spacecraft position measured relative to the Earth and Moon, respectively. The dynamical system is non-dimensionalized with a characteristic length, l^* , equal to 384,999.77578 km, and a characteristic time, t^* , defined as: $t^* = \sqrt{l^{*3}/Gm^*}$. Note that G is the gravitational constant and m^* is the sum of the mass of the Earth and the Moon. Two additional advantages of the CRTBP is the time-invariance and the existence of a single energy-like constant, labeled as the Jacobi Constant. Finally, five equilibrium points, i.e., Lagrange points, are computed within the $\hat{x} - \hat{y}$ plane of the Earth-Moon rotating frame. Mission designers frequently leverage insights from the CRTBP model to select mission orbits or construct transfers within the cislunar regime. In this investigation, the objective is to leverage the dynamical insights from a lower fidelity model to improve the cislunar IOD process.

2.2 Periodic Orbits

Periodic motion offers insight regarding predictable long-term behavior within a dynamical system. Although no closed-form solution of the generalized CRTBP model exists, various numerical techniques are available to construct periodic orbits along the cislunar space. This investigation implements a multiple-shooting technique with a Newton-Raphson algorithm to create families of periodic orbits that emanate from the L_1 and L_2 Lagrange points as well as the Moon. Fig. 1 contains the set of periodic orbit families for the L_1 and L_2 halo families—both northern and southern groups—that provide a base data set for representing known CRTBP dynamics. The distant retrograde orbit family (DRO) family emanating from the Moon is also included. These periodic orbit families form as 2D surfaces within the multi-body dynamics, but the surface exists entirely as the result of a numerical process. Through this observation, a location on the family surface, corresponding to a six-dimensional state, is parameterized via two non-dimensional η and τ values. From Fig. 1(b), η describes a unique periodic orbit along the family and has a normalized range of $[0, 1]$. After a periodic orbit along the family has been identified, a specific location along the selected orbit is expressed via the τ value; note that the range of τ is also normalized as $[0, 1]$. Numerically defined families, such as the ones found in Fig. 1, easily map to simulated observations through state transformations. Thus, orbit family data provides ample training data for machine learning methods to develop a numerically approximated inversion function developed for this study.

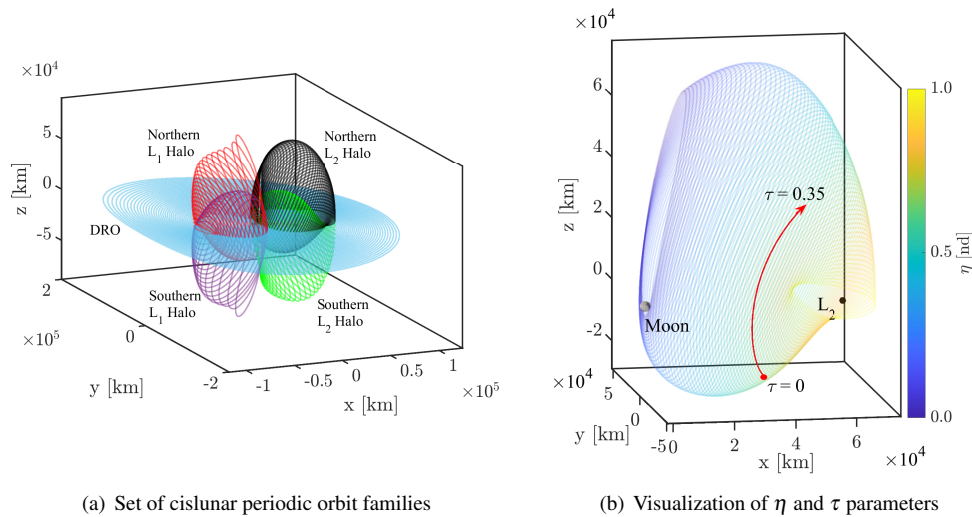


Fig. 1: Cislunar periodic orbit families: (a) Set of Earth-Moon Periodic orbit families used to generate neural network training data. The set includes Northern and Southern L_1 and L_2 halos and a planar Distant Retrograde Orbit family, plotted in the Earth-Moon rotating frame and centered at the Moon. (b) Northern L_2 halo family centered at the Moon. Each periodic orbit in the family is assigned a non-dimensional η value and a unique location along a periodic orbit is described via the non-dimensional τ value.

2.3 Nonlinear Batch Least-Squares

In order to maintain a fair comparison between MCLDOD's IOD performance and that of traditional IOD methods (i.e. Gooding), the results in Section 6 assume all methods are supplied with the same observations when computing

an initial estimate of the object’s position/velocity—three pairs of right-ascension and declination angles. However, many times there may exist more than three observations of an object, and to use all of the available data, it is common to employ a sequential or batch estimator once the IOD has been performed. Recalling that MCCLOD does not inherently produce a representation of uncertainty, selecting a method that does not require an initial covariance is ideal. An information filter is a possible choice, as it allows the user to effectively prescribe an infinite initial uncertainty [6]. However, as with any sequential estimator, the uncertainty of the estimate will be reduced with each consecutive measurement update, such that only at the time of the final observation will the estimate be refined using the entire set of observations. This work is more concerned with refining the estimate at the initial observation time, which requires an additional smoothing operation when processing the observations sequentially[15]. Thus, in order to reduce complexity when comparing IOD methods, the widely-used batch estimator known as nonlinear least-squares is selected, as it is straight-forward to employ, does not require an initial uncertainty, and maps all of the information gained from an observation set to a single time of interest [18].

As Batch Least-Squares (BLS) is fairly common, it is described here in brief. Once the selected IOD method produces an initial state estimate, a high-fidelity dynamics model is used to form the state-transition matrices mapping the remaining observation times to the time of the initial IOD observation. This dynamics model treats Earth as the central body, and includes a point-mass gravity model, and the sun and moon are treated as third bodies to account for their gravitational effects. The measurements are modeled classically as right-ascension and declination with white noise, and are assumed to have $1\text{-}\sigma$ standard deviations of 2 arcseconds. Using these models, nonlinear least squares is performed iteratively until a sufficient convergence has been reached.

3. MACHINE LEARNING

The nonlinear mapping between input observations and output orbit parameters is modeled by a feedforward artificial neural network. The neural network is composed of a base encoder, a regression layer, and, when appropriate, a classification layer. The base encoder is constructed from a Multilayer Perceptron (MLP) and is primarily responsible for learning a representation of the mapping structure between angles-only measurements and cislunar orbit parameters. The regression and classification layers are simply normalized linear layers. The regression layer infers a set of $\eta\text{-}\tau$ orbit parameters that describe a 6D state along a cislunar periodic orbit, while the classification layer determines the discrete orbit family (e.g., a southern L_1 halo family).

Multilayer Perceptrons have been successfully employed to solve regression and classification problems since the late 1950’s [16]. The MLP is at the core of many encoders such as the Variational Autoencoder [11] and even the novel Transformer, the first attention-based sequence transduction model [21]. Here, a relatively simple architecture, i.e., MLP plus linear layers, is found to be sufficient for the low-dimensional orbit state inference task. The overall feedforward artificial neural network is presented in Fig. 2.

Representing periodic functions with neural networks is a known challenge, with prior work suggesting replacing nonlinear activation functions with ones better suited to periodic functions [24]. Instead, this investigation leverages a trigonometric mapping (i.e., $\gamma^c = \cos \tau$ and $\gamma^s = \sin \tau$) to indirectly infer the location along a periodic orbit [12]. Thus, the outputs of the regression layer are $\{\eta, \gamma^c, \gamma^s\}$, where the phasing of the state along the orbit is computed as $\tau = \text{atan2}(\gamma^s, \gamma^c)$. As a result, no bespoke modifications to the neural network are required, and a simple construction remains.

3.1 Model Training

The neural network model is trained via supervised learning techniques and is implemented with the PyTorch machine learning library. To determine the efficacy of neural networks with periodic orbits, models are first trained and evaluated for regression only of periodic orbit states for each individual orbit family. Simulated data, which are generated offline as described in previous sections for each family, are used for model training with random observation noise added at each model update iteration. Simulations generate 25,000 samples for each orbit family and data are randomly split 80%, 10%, and 10% for training, validation, and testing, respectively. Validation and testing samples are excluded from model training. Observation noise is modeled as independent and identically distributed zero-mean Gaussian with $\sigma = 10$ arcsecs standard deviation for the angles and scaled as $\sqrt{2}\sigma$ for the angle rates. Regression-only models are trained to minimize mean squared error (MSE) loss between the simulated orbit state parameters and the inferred parameters. The regression-only model fit results are shown in Table 1 for holdout test data samples. As seen in the table, the neural network architecture performs state parameter inference reasonably well across all families.




Fig. 2: Feedforward artificial neural network for periodic orbit state and family inference. This includes a combined classification with regression combination; classification outputs may be omitted for a pure regression evaluation.

Table 1: Periodic orbit state parameter inference MSE from holdout test data for each individual orbit family.

Family	Observation Noise (arcsec)	MSE AVG (STD)
L ₁ northern halo	10	2.04e-07 (1.02e-07)
L ₁ southern halo	10	3.72e-07 (1.32e-07)
L ₂ northern halo	10	3.72e-07 (3.45e-07)
L ₂ southern halo	10	3.59e-07 (9.03e-07)
DRO	10	3.14e-07 (5.50e-07)

The results in Table 1 assume the periodic orbit family is known. However, the true family is not known *a priori* at runtime, and therefore must be inferred. While the orbit family is not required for inferring orbit state parameters (see Fig. 2), the family is required for downstream processing. Unlike in regression-only model training, where data are separated by family and an individual model is trained for each, one model is trained for both regression and classification across all periodic orbits and their datasets. The same base architecture from the regression-only case is also used for classifying the periodic orbit family, as seen in Fig. 2. In addition to minimizing the orbit state MSE, the combined regression and classification model also aims to minimize classification accuracy. The classification criterion for model training is cross-entropy (CE) loss. To balance the trade-off between optimizing for regression and/or classification, the combined overall loss is a weighted sum of the two respective losses.

Family classification with regression accuracy are summarized in Table 2 for three evaluation cases. The first row lists results for the 10% test data withheld from all datasets where the model correctly classifies on average the families in 99.28% of the test data. Though, there is an increase in parametric regression error, which is likely due to the additional orbit family uncertainty as well as ambiguities where orbits overlap across families. To demonstrate the impact of this overlap ambiguity, the results in the second row of Table 2 represent metrics for data-points from a family that overlap in physical, and therefore $\eta - \tau$, space of another family and taken from across all datasets. Accuracy results from the table suggest that these orbit regions create significant confusion leading to consequential regression and classification error, rendering the neural network uninformative in these regions. Conversely, the performance is much improved when these samples are excluded from the data, as seen in the last row of Table 2. Lastly, the orbit confusion regions represent a small (1.4%) fraction of the datasets and, therefore, cover a small region of real space. For observations collected over these regions, there is an expected degradation in inference performance.

Table 2: Periodic orbit state parameter inference MSE and family classification accuracy across all orbit families.

Data	Observation Noise (arcsec)	MSE AVG (STD)	Accuracy % AVG (STD)
Test data	10	6.34e-3 (3.06e-2)	99.28 (0.821)
Overlapping	10	1.83e-1 (1.38e-1)	48.59 (13.16)
Non-overlap	10	2.39e-3 (2.13e-2)	99.99 (0.105)

4. NEURAL NETWORK OUTPUT TRANSFORMATION AND VALIDATION

The neural network employed in this investigation estimates a lower-dimensional parameterization of the tracked object's six-dimensional position and velocity state that is required to initialize the IOD process. As visualized in Fig. 1, the $\eta - \tau$ encoding provides a useful way to represent the full state space across a single orbit family using only two variables. However, one obstacle present in this approach is defining and implementing a function capable of inverting the $\eta - \tau$ encoding to translate the NN outputs back into the six-dimensional physical state space. Defining such a function enables the overall decoding process outlined as,

$$\underbrace{\{\eta \quad \gamma^c \quad \gamma^s\}}_{\text{Raw NN output}} \xrightarrow{\tau = \text{atan2}(\gamma^s, \gamma^c)} \underbrace{\{\eta \quad \tau\}}_{\eta - \tau \text{ space}} \xrightarrow{\text{Decoding Function}} \underbrace{\bar{q} = [x \quad y \quad z \quad \dot{x} \quad \dot{y} \quad \dot{z}]^T}_{\text{CRTBP position \& velocity}} \quad (3)$$

where the $\{\gamma^c, \gamma^s\}$ is the cyclic value encoding of τ , as detailed in Section 3.

An additional challenge that results from employing a lower-dimensional state representation is in interpreting the NN loss function MSE metrics, listed in Tables 1 and 2. These values are useful in measuring NN performance during training and in comparing multiple trained models to evaluate potential model architecture and tuning options. However, these figures of merit provide limited insight into real-world behavior. Hence, evaluating results in the context of decoding process, Eq. 3, demonstrates the NN's performance in more intuitive position and velocity units. This investigation evaluates how regression errors translate to physical space by processing the entire training/test datasets through the decoder, and comparing the resulting estimated state vector with the expected one in the training dataset.

4.1 Monotonic Parameter Decoding Process

The overall decoding process leveraged in this investigation is comprised of two basic steps. First, given estimated η and τ values from the NN, the periodic orbit that corresponds to η value is identified by mapping η to a pre-selected orbit parameter ξ that evolves monotonically across the initial state ($\tau = 0$) of each database entry for a particular orbit family, e.g., orbital period or a position component; the parameters leveraged in this study are plotted in Fig. 3. A differential corrections approach is subsequently employed to converge on a periodic orbit at the precise orbit parameter value. The selected orbit parameter forms a one-to-one correspondence with the η value, and is necessary to precisely identify a specific orbits due to the fact that η has no physical significance and, therefore, cannot be implemented in the targeting algorithm. Alternatively, comparing the estimated η value against a sufficiently dense dataset may be sufficient to achieve an sufficiently accurate initial guess for the IOD process; however, a more comprehensive analysis of decoding accuracy requirements remains future work. Finally, the τ parameter is leveraged to select the final CRTBP rotating state vector along the periodic orbit by propagating the converged initial state for $(\tau \cdot \mathbb{P})$ time units, where \mathbb{P} is the orbital period.




Fig. 3: Monotonic parameters ξ for each orbit family, where \mathbb{P} is the orbital period. Y-axis values correspond to the ξ variable listed for each family; all quantities are plotted in nondimensional coordinates.

4.2 Neural Network Decoding Validation

An initial validation step for the training NNs is performed by evaluating each NN across the entire training and test datasets to quantify the NN accuracy and decoding performance, and to identify areas of space that may pose issues

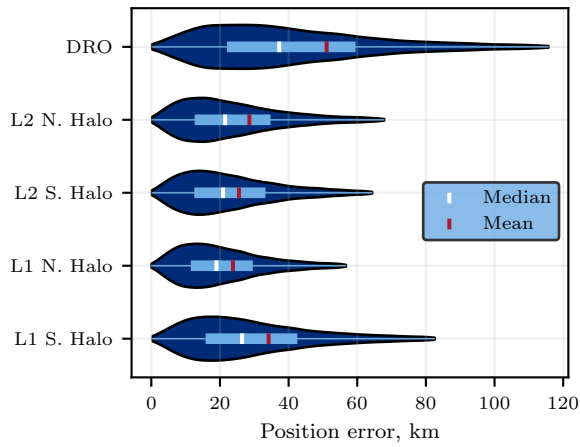
in the proposed IOD approach. The results are divided into three broad categories: 1) “Regression-only” models that perform regression for individual orbit families, assuming a-priori classification knowledge of the correct family, 2) the model that simultaneously performs both classification and regression tasks across all orbit families (Fig. 2), and 3) a “hybrid” approach that leverages the combined classification-regression model for classification, followed by the corresponding regression-only model.

The numerical results of the decoding validation study are listed in Table 3. For each model, statistics are gathered for the position and velocity error magnitudes the result from decoding NN estimation across the orbit dataset, omitting cases where the resulting orbit impacts the Moon. Consistent with the NN loss values in Tables 1 and 2, the regression-only models perform much more accurately than the combined classification and regression model. However, the joint model is able correctly classify observations in the vast majority (99.2%) of cases. This discrepancy between classification and regression performance motivates the third “hybrid” approach; which achieves results similar results to the regression-only models without relying on a-priori knowledge of the correct orbit family.

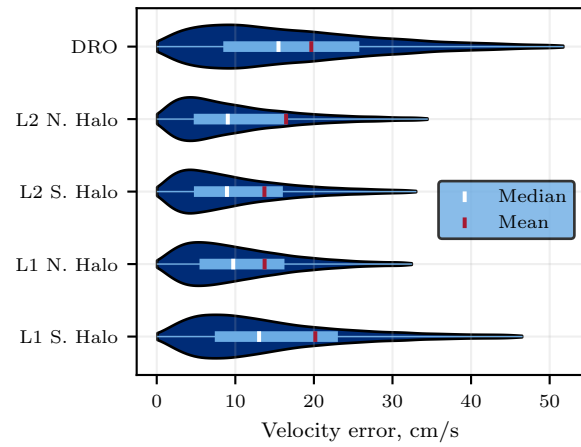
The error statistics summarized in Table 3 are further visualized for each family in Fig. 4. In regression-only cases, the DRO family possesses the highest error values, likely due to observably challenges from the fact that the family exists entirely within the rotating x - y plane of the Earth-Moon system. Furthermore, the relatively large difference between the mean and median values demonstrates the influence of large outlier values that inflate the overall mean, also evident in the maximum error magnitudes listed in Table 3. While the outlier magnitudes in the regression-only models appear large relative to the overall statistics, the “worst-case” scenario of 2718.88 km and 14.4 m/s error magnitudes in the DRO family are still relatively small compared with the vast scale of cislunar space.

Table 3: Decoding error statistics across different model configurations and orbit families. Check mark (✓) indicates statistics that only include correctly classified family values whereas (✗) denotes misclassified results.

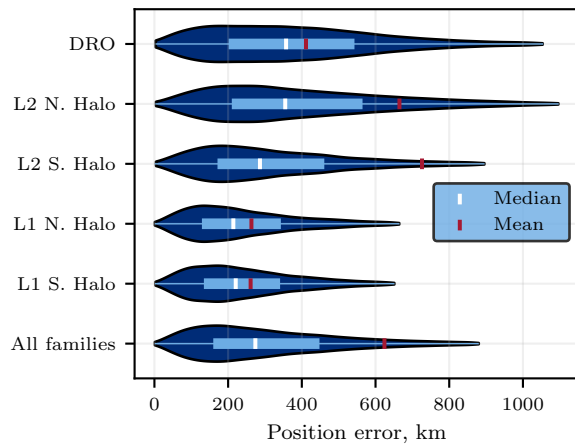
		Regression-only		Classification w/ regression		Hybrid	
		Figs. 4(a-b)		Figs. 4(c-d)		Figs. 4(e-f)	
		Pos. [km]	Vel. [m/s]	Pos. [km]	Vel. [m/s]	Pos. [km]	Vel. [m/s]
L ₁ S. Halo ✓	μ	34.19	0.20	261.05	1.46	33.82	0.20
	σ	29.60	0.25	199.74	1.69	28.53	0.25
	med	26.46	0.13	220.64	1.08	26.42	0.13
	max	369.53	5.57	7458.69	65.04	344.10	5.57
L ₁ N. Halo ✓	μ	23.80	0.14	263.23	1.33	22.77	0.13
	σ	20.65	0.16	227.10	1.60	17.48	0.13
	med	18.99	0.10	214.16	0.98	18.78	0.10
	max	666.14	3.08	12044.55	83.97	250.25	3.08
L ₂ S. Halo ✓	μ	25.54	0.14	726.15	3.65	25.27	0.14
	σ	18.95	0.20	4313.01	20.03	18.55	0.20
	med	20.91	0.09	286.38	1.16	20.80	0.09
	max	274.40	10.39	84091.33	462.81	274.41	10.40
L ₂ N. Halo ✓	μ	28.58	0.16	664.93	3.27	27.22	0.16
	σ	30.61	0.36	3695.69	13.48	23.75	0.33
	med	21.51	0.09	354.99	1.36	21.33	0.09
	max	575.52	17.65	77469.10	384.12	309.42	17.65
DRO ✓	μ	51.08	0.20	411.30	1.30	51.08	0.20
	σ	76.81	0.21	301.42	0.96	76.81	0.21
	med	37.26	0.15	357.29	1.07	37.29	0.15
	max	2718.88	14.40	3384.79	9.09	2718.90	14.40
All Families ✓	μ	-	-	457.17	2.18	31.56	0.16
	σ	-	-	2520.01	10.76	39.54	0.23
	med	-	-	273.00	1.11	23.65	0.11
	max	-	-	84091.33	462.81	2718.90	17.65
All Families ✗	μ	-	-	20959.83	86.09	9189.25	14.93
	σ	-	-	29409.72	137.47	34612.63	80.04
	med	-	-	673.67	3.37	86.62	0.56
	max	-	-	126183.18	1031.30	200245.38	557.07




(a) Regression models: position errors




(b) Regression models: velocity errors




(c) Classification-regression model: position errors



(d) Classification-regression model: velocity errors



(e) "Hybrid" Classification-then-regression approach: position errors



(f) "Hybrid" Classification-then-regression approach: velocity errors

Fig. 4: Error magnitude statistics measuring the performance of the proposed models and run configurations for each orbit family. Outlier values greater than 1.5 IQR points beyond the third quartile are omitted in each visualization.

5. CISLUNAR IOD SIMULATION PROCESS

The cislunar IOD process, i.e., the MCCLOD process, proposed in this investigation implements a trained neural network model within the analysis framework to deliver a representative six-dimensional state vector of a detected object. Results from the novel MCCLOD process are compared to a traditional Gooding process. Measurements from a single optical ground observer, that is, right ascension and declination angles, are assumed throughout this investigation. Recall that the Gooding algorithm assumes two-body mechanics, requires three distinct optical measurements, and a guess for the radial, or range, distance of the detected object. For cislunar objects, the radial guess necessary for Gooding must be carefully selected to be within an approximate vicinity of the object to improve convergence success rate. Although, this is challenging due to large cislunar distances and eventually becomes a computationally expensive guess-and-check process. The framework of an application of both methods are summarized in Fig. 5.

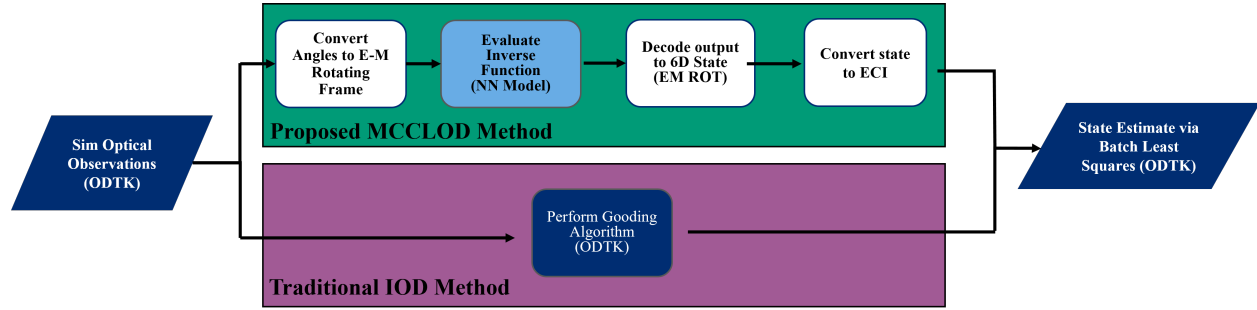


Fig. 5: Framework for both a traditional IOD method using Gooding algorithm and the proposed MCCLOD process using a trained neural network.

The framework with the Gooding algorithm is best summarized in three steps: 1) Simulate/Collect optical measurements, 2) perform Gooding algorithm and estimate initial state vector, and 3) perform Batch Least Squares to generate estimated 6D IOD state of observed object, Alternatively, the MCCLOD process is summarized in the following steps:

- *Simulate/Collect Optical Measurements:* Collecting optical measurements assumes that real data from a presumed cislunar object is available. In this investigation, optical data is simulated via the ODTK software for a numerically converged cislunar trajectory, labeled as the truth, from a single ground site. The list of ground sites considered in this analysis are described in Table 4.

Table 4: List of selected ground sites

Ground Site	Geodetic Latitude [deg]	Longitude [deg]	Altitude [m]
New Mexico	32.82	-106.66	1250
Hawaii	20.71	-156.26	3000
Diego Garcia	-7.41	72.45	0

- *Convert Angles to E-M Rotating Frame:* Optical measurements collected or simulated from a ground site are received in the J2000 topocentric frame with respect to the location of the site. The MCCLOD neural network is trained with angles calculated with respect to the Earth in the Earth-Moon rotating frame. Therefore, a conversion step is necessary to transform the angles to the appropriate frame. The angle conversion process is summarized in Fig. 6. Given a set of right ascension and declination angles in the topocentric frame, α^T and ψ^T , a conversion to the geocentric J2000 frame is performed using the measurement epoch and a geocentric radius guess. Optical ground sensors do not provide geocentric range or radius information, therefore, the conversion strategy from topocentric to geocentric requires a radius or range guess, see Vallado [20] for more details regarding the transition steps. In cislunar space, an appropriate radius distance is either the Moon's distance or the distance of one of the Lagrange points, either L_1 or L_2 . The output of this process is right

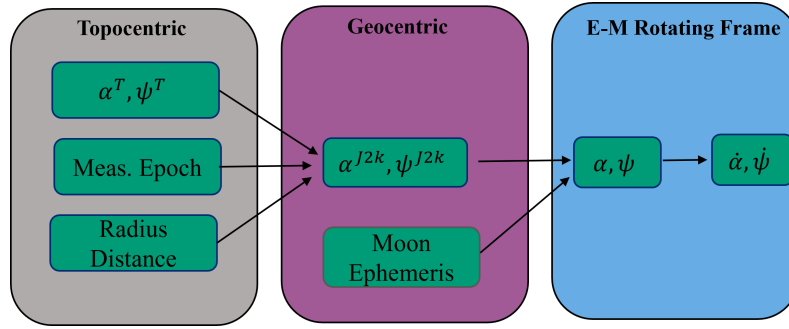


Fig. 6: Diagram describing angle conversion process. Right ascension and declination measurements are received in topocentric frame, NN models require Earth-Moon Earth-Centered defined angles

ascension and declination angles in the geocentric frame, α^{J2K} and ψ^{J2K} , respectively. Next, a conversion into the Earth-Moon rotating frame delivers angles α and ψ , that is, the right ascension and declination with respect to the Earth. Finally, for a set of distinct observation measurements, the angular rates are numerically approximated in two steps: first the sets of observed angles are fit with a cubic polynomial, then the rates for each observed time are calculated by computing the numerical derivative of the angular fit function.

- *Evaluate Inverse Function:* Each transformed measurement vector in the Earth-Moon rotating frame, that is, $Y = [\alpha, \psi, \dot{\alpha}, \dot{\psi}]$, is mapped to a two-vector output, i.e., $[\eta, \tau]$, via a trained neural network model. In this investigation, a regression model and a classification-regression model are constructed with periodic orbit family information. An additional hybrid approach that combines both NN models is considered as an alternative to any single model.
- *Decode Output to 6D State:* Given a periodic orbit family, a single $\eta - \tau$ vector is mapped to a six-dimensional state vector in the Earth-Moon system. The decoding steps are summarized in Section 4.
- *Convert State to ECI:* After a 6D state is decoded from the NN output parameters, a state vector in the Earth-centered J2000 reference frame is computed by performing a rotation from the Earth-Moon rotating frame. The state transformation steps are adapted from Park et al. [13].
- *State Estimate via BLS Filter:* Finally, the 6D state in the ECI frame is ingested into a BLS process via the ODTK software package. Parameters for the BLS method are discussed in later sections of the document.

Within the proposed MCCLOD IOD process, the trained neural network is either a regression-only model based on an individual orbit family or a classifier and regression model trained with five distinct periodic orbit families. Additionally, a hybrid approach is considered due to the observed accuracy improvement during the validation analysis, see Section 4 for details. All types of NN options are considered in the following simulated cislunar scenarios discussed in the next section. Through the implementation of a trained NN model, the objective of the MCCLOD process is to deliver an improved IOD 6D state vector, compared to the traditional Gooding method, to maintain custody of a newly detected cislunar object.

6. ANALYSIS RESULTS

The proposed MCCLOD methodology is applied to simulated cislunar trajectories and compared to a traditional Gooding IOD method. For IOD test cases, both L_1 and L_2 halo orbit trajectories are numerically converted to the Earth-Moon-Sun point-mass ephemeris dynamical model employing a time-of-flight greater than 5 months. Ephemeris states for the Moon and Sun are collected from the DE440 BSP file from the NAIF SPICE software package. The simulated trajectory represents the "truth" and provides 6D states to compute the position errors for both IOD method types. To investigate realistic tracking scenarios, this investigation assumes an individual optical sensor utilizing a ground site in Table 4 that observes a single cislunar test object. A set of observations for a single pass are generated with a fixed collection cadence. Each observation pass is simulated via the ODTK software package with the ground

Table 5: ODTK simulation ground site settings. Simulation assumes an optical sensor the delivers angles-only measurement, i.e., right ascension and declination for an observed object.

Parameters	Value
Tracking Cadence	1 hr
Angle White Noise (1σ)	2 arcsecs

station settings appearing in Table 5. Recall, that angles-only measurements, i.e., right ascension and declination, are estimated via the ODTK simulator. After a set of measurements are generated, a series of analysis trials are defined given the number of measurements for a single tracking opportunity. Each trial delivers a 6D state solution for each IOD method and the correlated BLS result at a solution epoch, that is, the epoch associated with a simulated measurement. Additionally, trials possess a unique time span for a BLS execution that determines the measurements to consider during the filtering process. A schematic describing the setup of the trials is presented in Fig. 7. The layout of this analysis format is selected for two reasons: (i) direct the examination towards whether the number of measurements in a single pass lowers the efficacy of either the traditional or the MCCLOD IOD methods, and (ii) perform the IOD process at different epochs of a single tracking opportunity representing varying chance detection scenarios. While this investigation focuses on the performance after one tracking opportunity, it is easily extendable to incorporate multiple tracking passes in the process pipeline. The following examples highlight regions along a cislunar trajectory with improved IOD performance from the MCCLOD IOD process compared to the traditional methods.





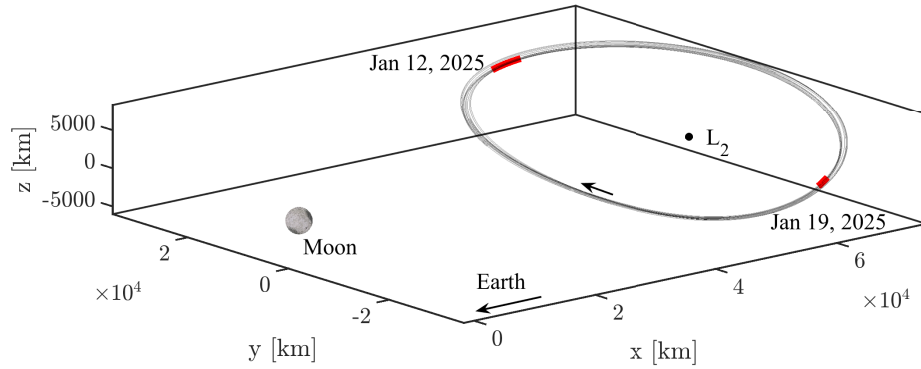
Fig. 7: Schematic for the trial setup of the IOD performance analysis

6.1 Observations for a Northern L_2 Halo Trajectory

The first example case localizes optical observations for a northern L_2 halo orbit in the Earth-Moon system. The Hawaii ground site, see station parameters in Table 4, is selected to view the object given the simulation parameters in Table 5. Two separate locations along the selected L_2 halo trajectory, plotted in red in Fig. 8(b), are observed with an optical sensor. The first opportunity on Jan 12, 2025 produces 13 measurements and the second pass on Jan 19, 2025 produces 9 total measurements, see Fig. 8(a). The analysis trials, see the schematic in Fig. 7, are constructed with the collected measurement data and epochs from each tracking pass. In this example, the orbit type, i.e., a northern L_2 halo, is known and an approximate radius distance of $70 R_e$ is selected, that is, the distance of the L_2 Lagrange point with respect to the Earth given in units of Earth radii. Recall that the geocentric radius guess is an input to both the traditional Gooding and the proposed MCCLOD IOD methods. A regression NN model trained with the northern L_2 halos, see Section 4 for additional details, is selected as the NN in the MCCLOD *Evaluate Inverse Function* step. The position errors with respect to the truth in Fig. 8(b) for both tracking dates are plotted in Fig. 9. In Fig. 9(a), the points in magenta represent the position error of the estimated states from the MCCLOD IOD process with respect to the truth states. Recall that each measurement, at an associated solution epoch, is mapped to a single unique MCCLOD IOD 6D state output. The magenta points in Fig. 9(b) highlight the volatile performance of a traditional IOD method, where large position errors greater than 10,000 km are generated from the Gooding process. The color of each point in Figs. 9(a) and 9(b) represents the total time span of each BLS filter run, recalling that greater time span equals more

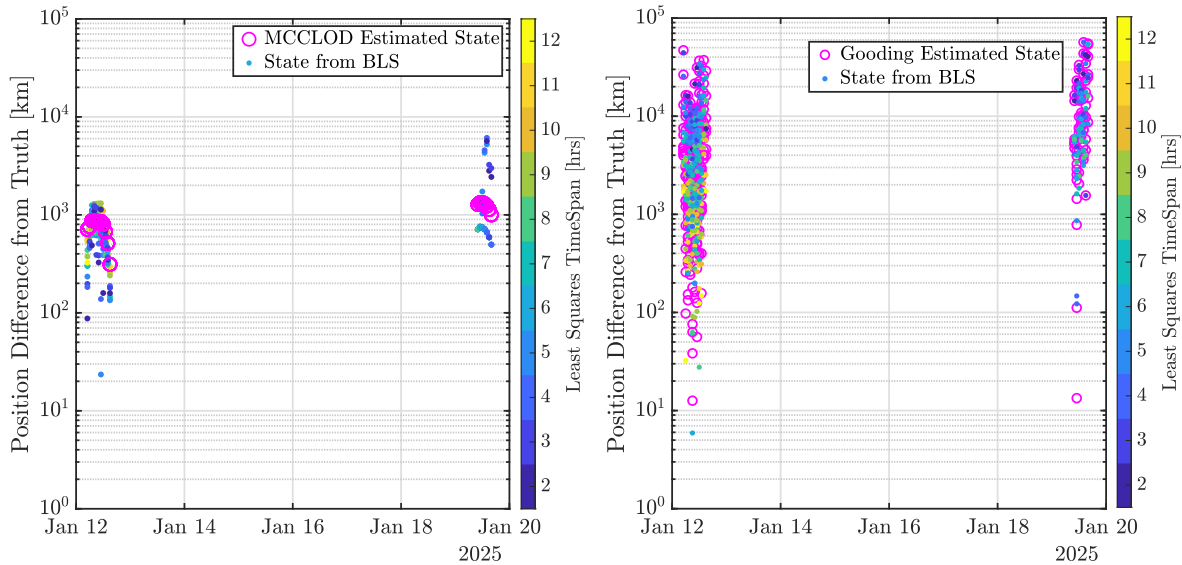


(a) Observation windows for two tracking opportunities



(b) Selected northern L_2 halo trajectory (in grey) and the positions along the orbit observed by the optical sensor (in red)

Fig. 8: Simulated tracking opportunities for an object in a northern L_2 halo trajectory from a ground site in Hawaii



(a) Results from MCCLOD with northern L_2 halo regression model

(b) Results from Gooding algorithm

Fig. 9: Results from IOD analysis: The results from both IOD methods assume an approximated geocentric radius of 70 Re. (a) The position error from the MCCLOD IOD states, in magenta, and the resulting BLS 6D state vector. (b) Position error from Gooding 6D state. Errors of the BLS trials from both methods are color-coded by the time span of the filter run.

measurements during filter step. In Fig. 9(a), the position error of the BLS solutions are not guaranteed to be below the error of the MCCLOD estimated states for either passes; note that this behavior is also visualized in the results from Gooding in Fig. 9(b). An important distinction between the performance of the MCCLOD and the Gooding methods is observed in the range of position errors for each pass. The range of position errors from MCCLOD are an order of magnitude lower compared to the Gooding method, see Table 6 for a summary. Lower position error ranges indicate lower final position dispersions and suggest that the MCCLOD delivered states for the selected tracking opportunities are unaffected by the BLS time span, or rather, the total number of measurements. Note that results for both IOD methods suggests that longer time spans are not correlated to lower position errors (displayed in Fig. 9).

Table 6: Range of position errors from both IOD methods for the selected northern L_2 halo in km

Date of Pass	MCCLOD	Gooding
Jan 12, 2025	5600	56,600
Jan 19, 2025	1287	44,400




Fig. 10: Position of BLS output state vector from MCCLOD IOD process for a revolution of the northern L_2 halo orbit. The color of each solution corresponds to the BLS filter time span.




Fig. 11: Position of BLS filter output state vector from a traditional Gooding IOD process for a revolution of the northern L_2 halo orbit. The color of each solution corresponds to the BLS time span.

Finally, the positions of the converged BLS states from both IOD methods are visualized in Figs. 10 and 11 along with the tracked cislunar northern L_2 trajectory. Figure 10 displays the final IOD positions from the MCCLOD process and reveals how each BLS solution remains within the vicinity of the truth states, plotted in red. However, in Fig. 11, the estimated BLS state vectors after Gooding are dispersed along a radial line between the Earth and the observed trajectory arc. The outcome of this example suggests that, for these two tracking opportunities, MCCLOD delivered states that are within 2000 km of the true trajectory with information from just a single tracking pass. Of note, two major assumptions during the processing of these observed passes govern such a conclusion: (i) the tracked cislunar

object is correctly assumed to be operating within the vicinity of the Earth-Moon L_2 , and (ii) the mean radii of the observed positions closely approximates the location of the L_2 point.

6.1.1 Performance of MCCLOD with different radial guesses

The traditional Gooding IOD method incorporates optical measurements and requires an initial guess for the geocentric radius. The proposed cislunar IOD method with MCCLOD also assumes an initial geocentric radius, see more details in Section 5. Recall that the results presented in Fig. 9(a) assumes a radius distance of $70 R_e$, i.e., 70 times Earth radius, and equal to the location of the Earth-Moon L_2 Lagrange point. However, for any assumed cislunar object, the region of interest is not limited to just the Lagrange points. Therefore, the selection of the radial guess also has a significant effect on the results of the IOD process. The position error plots in Fig. 12 reveals the performance of the MCCLOD process for the northern L_2 halo in Fig. 8(b) with three different radius approximations.

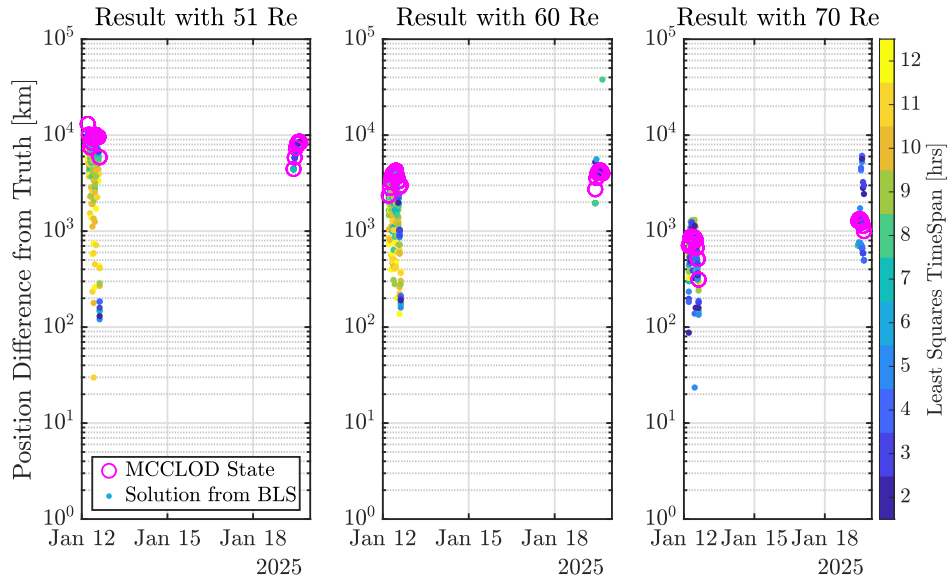


Fig. 12: Position error difference of the MCCLOD IOD and BLS output 6D states for varying radii approximations.

Note that the radii of $51 R_e$, $60 R_e$, and $70 R_e$ represent the distances of the L_1 Lagrange point, the Moon, and the L_2 Lagrange point, respectively. Recall that the truth halo trajectory is contained near the vicinity of the L_2 Lagrange point and, therefore, the position errors of the MCCLOD output for $70 R_e$ from Fig. 12 are lower for during both tracking opportunities. Further examination of the BLS results from the tracking opportunity of Jan 12 suggests that selecting a radius closer to the truth decreases the estimated position error. Although, this behavior is not consistent for the tracking pass starting on Jan 19 as the selection of $70 R_e$ radius produces position errors as high as the other radial guesses. This behavior could be an outcome correlated to the observed location along the trajectory during the tracking pass and the position errors associated with this region, estimated via MCCLOD, are inherently higher regardless of the radial guess.

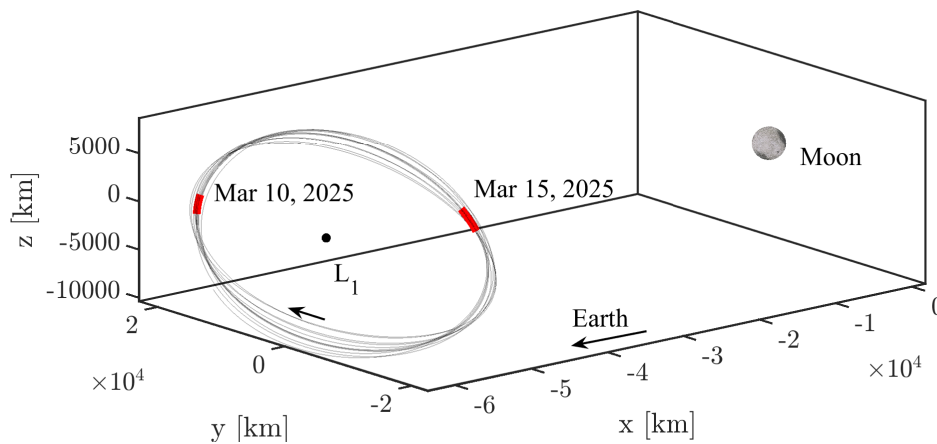
6.2 Observations for a Southern L_1 Quasi-Halo Trajectory

The second example in this investigation demonstrates the performance of both MCCLOD and a traditional Gooding method in tracking an Earth-Moon L_1 trajectory. A New Mexico ground site, see Table 4, is selected to observe a reference cislunar trajectory, defined as the truth, operating in a southern L_1 quasi-halo orbit. A quasi-periodic trajectory was selected for two reasons: 1) all realistic, i.e., higher fidelity, cislunar trajectories have quasi-periodicity characteristics rather than purely periodic, and 2) to test the efficacy of the MCCLOD process as the implemented NN models are only infused with simulated periodic motion. Right ascension and declination angles are generated via the ODTK simulator with white noise parameters from Table 5. Two tracking opportunities in Mar 10 and Mar 15, 2025 are selected to observe the object along the L_1 trajectory visualized in Fig. 13(b). The regions in red in Fig. 13(b) highlight the position of the object during the scheduled tracking times. Ten total measurements are collected from

the first tracking pass on Mar. 10, 2025, and 11 measurements are collected during the pass on Mar 15, 2025. Section 4 introduces two NN models, i.e., the regression and classification-regression models, along with a hybrid approach which are implemented to construct actionable IOD states for this example. The regression NN is trained with an individual periodic orbit family, e.g., the southern L_1 halo family. The classification-regression model is a NN model trained with five different periodic orbit families, more details provided in Sections 3-4, and delivers a predicted periodic orbit family for each input measurement. Finally, the hybrid approach uses the classification-regression NN output to predict a periodic orbit family for each measurement, then transfers each measurement through the regression model trained with identified orbit family. Within the MCCLOD IOD process, either a NN model or the hybrid approach is implemented in the *Evaluate Inverse Function* step, see Section 5 for details. The IOD results



(a) Observation windows for two tracking opportunities



(b) Southern L_1 quasi-halo trajectory (in grey) and positions along the orbit observed by the optical sensor (in red)

Fig. 13: Observation times for two separate tracking opportunities via an optical sensor. (b) Selected southern L_1 quasi-halo trajectory. The trajectory positions at the selected tracking times are highlighted in red.

for this example incorporate all three previously defined NN approaches during the tracking opportunities defined in Fig. 13(a). First, the classification results from the classification-regression model is plotted in Fig. 14, note that results from this model are also equal to the results of the classification step of the hybrid approach. The classification-regression NN successfully identifies all measurements from the tracking pass on Mar 10, 2025 as members of the southern L_1 halo family. Although the truth is more strictly defined as a southern L_1 quasi-halo orbit, the trajectory is




Fig. 14: Classification results of the tracking opportunities for the southern L_1 quasi-halo trajectory in Fig. 13(b).

near the vicinity of the southern L_1 periodic orbit family. Alternatively, 30% of the measurements during the second pass on Mar. 15 are misclassified as northern L_2 halos. Throughout the MCCLOD process, the computation of the geocentric angles and the transition to Earth-Moon right ascension and declination angles and rates introduces numerical errors. Recall that, these steps are guided by the measurement epoch and a geocentric radius guess, selected to be the distance of the L_1 Lagrange point, i.e., 51 Re. The numerical projection of the right ascension and declination rates are especially sensitive to the selection of the measurement radius guess, which, in turn, could lead to faulty evaluation of angular rate values. In an operations scenario, an analyst will not have any orbit parameter information of the discovered object, therefore, successful orbit classification of optical measurements is a necessary step to aid the cislunar IOD process. The results for the tracking opportunities presented in this example demonstrate regions of the orbit with successful classification essential to provide an improved IOD state via the MCCLOD process.

Initial Orbit Determination and BLS output states computed via both the MCCLOD and traditional methods demonstrate the efficacy of each NN model and approach. Position errors are generated with respect to the L_1 quasi-halo orbit, i.e., the truth, visualized in Fig. 13(b). The results presented in Fig. 15 are based on the trial methodology described in Fig. 7. The position errors of the output MCCLOD IOD states, in Fig. 15, using the classification-regression NN model are higher compared to the state output from the regression model or the hybrid approach. Note that, this behavior is consistent with the validation results from Section 4 which demonstrated that the regression step of the classification-regression model produces higher errors compared to the single family regression models. Observe that position errors from the classification-regression model for the pass in Mar. 15 are orders of magnitudes higher than the other two approaches. Additionally, results for the hybrid approach for the second pass show three measurements that produce IOD guess states two orders magnitude higher than all other measurements, a consequence of misclassification. The trials for the pass in Mar 15 misclassifies three measurements as L_2 halos, and therefore use

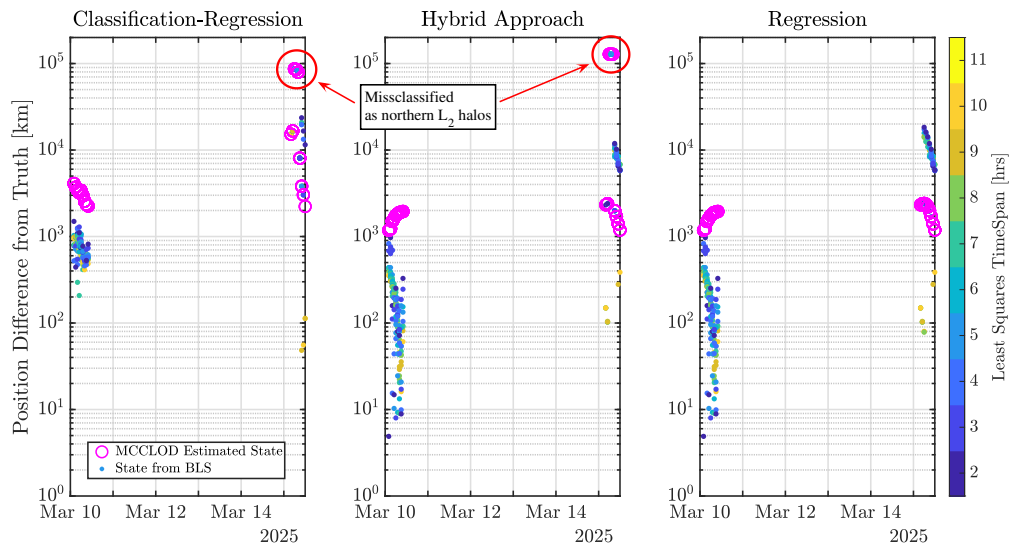


Fig. 15: Position errors from the Neural Networks "regression-only" and "classification with regression", along with a "hybrid" approach. The results reflect a radial guess of 51 Re.

the regression NN model trained with the northern L_2 halos to deliver an initial IOD state to pass to the BLS filter. This likely misclassification pushes MCCLOD into delivering IOD states on the other side of the Moon; however, this type of erroneous state selection is only observed for 30% of all the measurements in the second pass. During the first tracking opportunity, the position error from the BLS filter using the classification-regression NN model, in Fig. 14, are contained in a range of 300 – 1500 km. Yet, the regression model and the hybrid approach deliver position errors in the range of 10 – 1000 km. Successful classification of the entire first pass as a southern L_1 halo leads to identical results from the hybrid approach and the regression NN model. In the second tracking opportunity, all NN models produce position errors with an order of magnitude higher than the results from the Mar. 10 opportunity. This suggests that this region is extremely sensitive and the MCCLOD IOD process struggles to generate errors that are lower than 2000 km. Finally, the results from the regression NN model is compared to the traditional Gooding process in Fig. 16. Results from the L_1 southern halo regression model are selected for comparison because it assumes that successful

classification was achieved for all collected measurements during the tracking pass. Evaluation of the Gooding method for all trials, see diagram 7 for details, for both tracking opportunities is performed by assuming a geocentric distance of 51 Re, i.e., the distance of the L₁ Lagrange point. A summary of the position error ranges for the BLS filter output using MCCLOD and Gooding is encapsulated in Table 7. The error ranges for the Gooding output are an order of magnitude higher than the outputs from MCCLOD. Similar to the northern L₂ example, large error dispersions suggest that there is higher variability in the IOD state vector solution from the Gooding methods compared to MCCLOD. Finally, the positions of the converged BLS states from the MCCLOD regression model and the traditional Gooding method are visualized in Figs. 10 and 11, respectively, along with the observed cislunar southern L₁ quasi-halo trajectory. Figure 10 demonstrates that the MCCLOD southern L₁ halo regression model delivers states within 1000 km of the truth (in grey) for a Mar. 10 tracking pass. However, there is a significant radial offset, with respect to the Earth, in the resulting IOD state vectors for the Mar. 15 opportunity. Comparatively, the 6D states delivered from the Gooding method are dispersed, in the geocentric radial direction, in the region between the L₁ Lagrange point and the Moon. In summary, a southern L₁ quasi-halo trajectory was selected to examine the performance of two different NN models, a hybrid NN approach, and the traditional Gooding method. Successful classification was observed for one tracking opportunity, but misclassified measurements were delivered for a pass four days later. Additionally, the MCCLOD regression model delivered errors that were an order of magnitude lower than Gooding for the first tracking pass, but encountered challenges during the second pass.

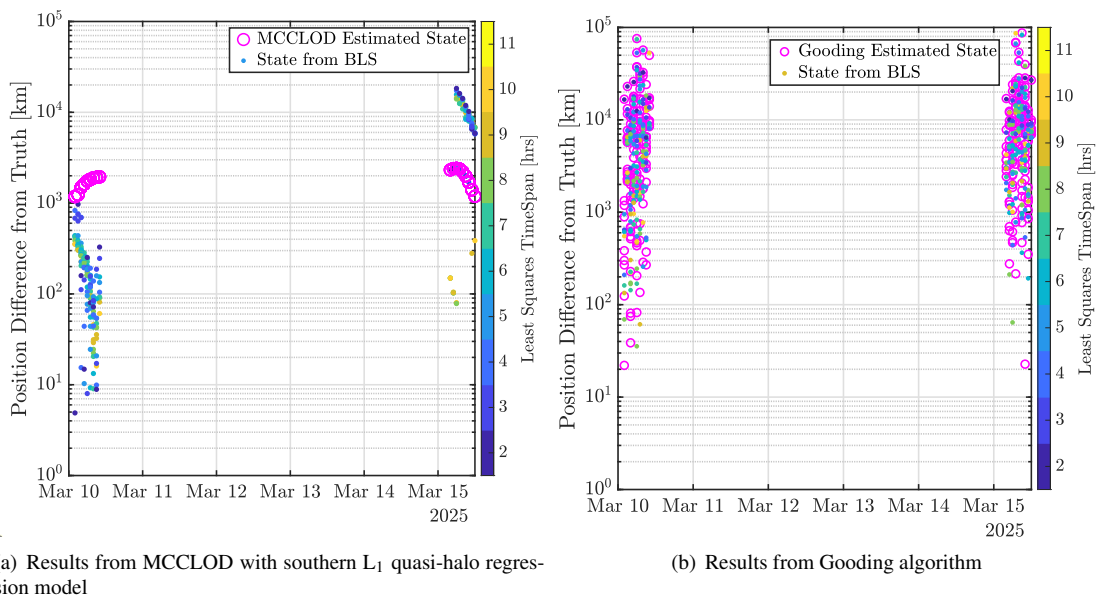


Fig. 16: Results from IOD analysis: The results from both IOD methods assume an approximated geocentric radius of 51 Re. (a) The position error from the MCCLOD IOD states, in magenta, and the resulting BLS 6D state. (b) Position error from Gooding 6D state. Errors of the BLS trials from both methods are color-coded by the time span of the filter run.

Table 7: Range of position errors of the BLS filter for all IOD methodologies for the selected southern L₁ quasi-halo in km.

Date of Pass	Regression	Classification-regression	Hybrid approach	Gooding
Mar 10, 2025	970	1288	970	108,728
Mar 15, 2025	18,252	87,530	128,000	116,590

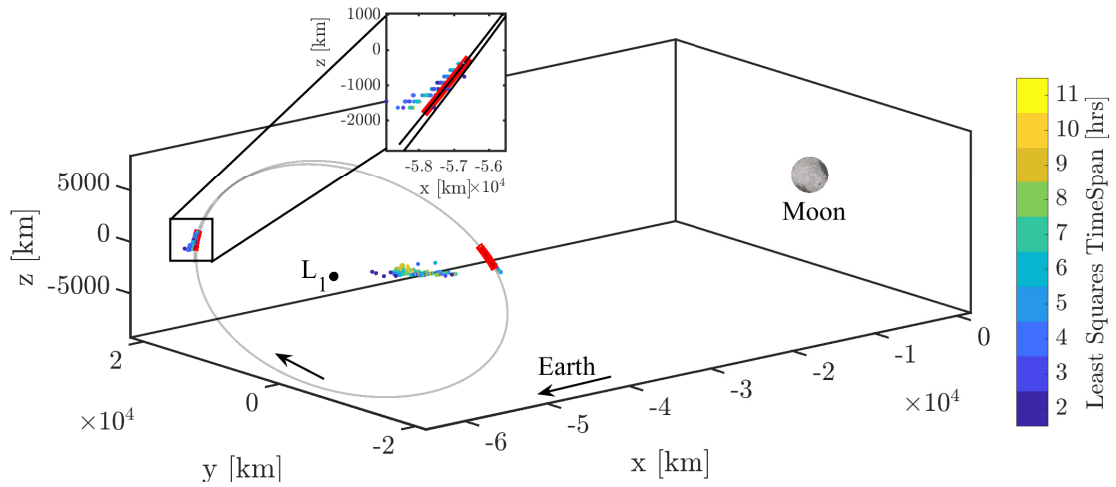


Fig. 17: Position of BLS filter output state vector for a traditional Gooding IOD process for a revolution of the southern L_1 quasi-halo orbit. The color of each solution corresponds to the BLS time span. Large position dispersions are revealed for the Mar 15 pass




Fig. 18: Position of BLS filter output state vector from a traditional Gooding IOD process for a revolution of the southern L_1 quasi-halo orbit. The color of each solution corresponds to the BLS time span. Large position dispersions, assuming a radius guess of $51 R_e$, revealed for both tracking windows.

7. CONCLUDING REMARKS

The chaotic nature of the cislunar regime introduces many challenges to the IOD problem in a chance detection scenario. The complex dynamics often limits the efficacy of traditional methods, i.e., the Gooding algorithm for angles-only observations. This leads to computationally expensive trial-and-error processes in an attempt to deliver an actionable 6D state. To improve the cislunar IOD process, this investigation presents two objectives: to implement machine learning techniques to attempt cislunar classification and regression from optical measurements and to construct a novel IOD process framework incorporating trained neural networks. Two neural network models are trained with Earth-Moon L_1 and L_2 orbit families: a regression-only model and a classification-regression model. An additional hybrid approach is introduced that first performs classification via the classification-regression NN model, followed by regression with a single family regression model. Validation results suggest that the classification-regression model delivers higher state errors compared to the regression-only model and the hybrid approach is a promising alternative for both regression and classification. Future work will focus on the formulation of a hybrid approach with a classification-only model and a set of regression-only models. A novel cislunar IOD process, i.e., MCCLOD, is defined that begins with collection of a set of observed angles-only measurements and outputs a 6D state from a Batch-Least-Squares filter. A northern L_2 halo orbit and a southern L_1 quasi-halo orbit, converged in a high fidelity ephemeris dynamics model, are selected to compare the efficacy of the MCCLOD process and a traditional Gooding angles-only method. Results for two separate tracking opportunities of two cislunar trajectories indicate that MCCLOD consistently provides position errors an order of magnitude lower than Gooding. Although, results from MCCLOD are sensitive to the tracking location along a cislunar trajectory which leads to higher position error values. Additionally, successful classification was demonstrated for one tracking pass of the southern L_1 quasi-halo orbit, however, a 30% misclassification of the measurements, based on a set of nine collected measurements, was observed for a subsequent pass four days later. This investigation has successfully trained a NN model by infusing Earth-Moon periodic orbit family information to successfully classify optical measurements and estimate an actionable 6D state vector guess for a BLS filter.

- [1] Gregory P Badura et al. “Identifying Cislunar Orbital Families via Machine Learning on Light Curves”. In: *The Journal of the Astronautical Sciences* 71.3 (2024), p. 27.
- [2] Paul Billings et al. “Cislunar Initial Orbit Determination using CAR-MHF”. In: *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*. 2023.
- [3] Mark Bolden et al. “Probabilistic initial orbit determination and object tracking in cislunar space using optical sensors”. In: *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*. 2022, pp. 27–30.
- [4] C Channing Chow et al. “Cislunar orbit determination behavior: processing observations of periodic orbits with gaussian mixture model estimation filters”. In: *The Journal of the Astronautical Sciences* 69.5 (2022), pp. 1477–1492.
- [5] Kyle J Craft and Kyle J DeMars. “Initial Cislunar Orbit Determination using Gaussian Mixture Approximations”. In: *John Hopkins Applied Physics Lab Cislunar Security Conference*. Dec. 2023.
- [6] John L Crassidis and John L Junkins. *Optimal Estimation of Dynamic Systems*. Chapman and Hall/CRC, 2011. DOI: 10.1201/b11154.
- [7] Kyle J. DeMars and Moriba K. Jah. “Probabilistic Initial Orbit Determination Using Gaussian Mixture Models”. In: *Journal of Guidance, Control, and Dynamics* 36.5 (2013), pp. 1324–1335. DOI: 10.2514/1.59844.
- [8] J Gaebler et al. “Cislunar Initial Orbit Determination with Optical Tracklets”. In: *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*. 2023, p. 101.
- [9] Casey R Heidrich and Marcus J Holzinger. “Universal Angles-Only Cislunar Orbit Determination Using Sparse Collocation”. In: *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*. 2023, p. 17.
- [10] James E Hippelheuser. “A Novel Multi-Observer Orbit Determination and Estimation Framework for Cislunar Space Domain Awareness”. PhD thesis. University of Central Florida, 2023.
- [11] Diederik Kingma and Max Welling. “Auto-Encoding Variational Bayes”. In: *International Conference on Learning Representations*. Dec. 2014.

- [12] Nicholas B. LaFarge, Kathleen C Howell, and David C. Folta. “Adaptive Closed-Loop Maneuver Planning for Low-Thrust Spacecraft using Reinforcement Learning”. In: *73rd International Astronautical Congress*. IAF. Paris, France, Sept. 2022.
- [13] Beom Park and Kathleen C. Howell. “Leveraging Intermediate Dynamical Models for Transitioning From the Circular Restricted Three-Body Problem to an Ephemeris Model”. In: *AAS/AIAA Astrodynamics Specialist Conference*. American Astronautical Society. Charlotte, North Carolina, Aug. 2022.
- [14] Guy Revach et al. “KalmanNet: Neural network aided Kalman filtering for partially known dynamics”. In: *IEEE Transactions on Signal Processing* 70 (2022), pp. 1532–1547.
- [15] Simo Särkkä and Lennart Svensson. *Bayesian Filtering and Smoothing*. 2nd ed. Cambridge University Press, 2023. DOI: 10.1017/9781108917407.
- [16] Juergen Schmidhuber. *Annotated History of Modern AI and Deep Learning*. 2022. arXiv: 2212.11279 [cs.NE]. URL: <https://arxiv.org/abs/2212.11279>.
- [17] Andrea Scorsoglio, Luca Ghilardi, and Roberto Furfaro. “A physic-informed neural network approach to orbit determination”. In: *The Journal of the Astronautical Sciences* 70.4 (2023), p. 25.
- [18] Byron Tapley, Bob Schutz, and George Born. *Statistical Orbit Determination*. Elsevier, 2004.
- [19] Michael R Thompson et al. “Cislunar Orbit Determination and Tracking via Simulated Space-Based Measurements”. In: *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*. 2021.
- [20] D. A. Vallado. *Fundamentals of Astrodynamics and Applications*. 4th. Microcosm Press, 2013.
- [21] Ashish Vaswani et al. “Attention is All you Need”. In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon et al. Vol. 30. Curran Associates, Inc., 2017. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- [22] Kullen W Waggoner, David Curtis, and Thomas A Lovell. “Cislunar Orbit Determination Using Multi-Receiver Doppler Ratios”. In: *AIAA SciTech Forum*. Jan. 2024. DOI: 10.2514/6.2024-0425.
- [23] Samuel Wishnek, Marcus J Holzinger, and Patrick Handley. “Robust Cislunar Initial Orbit Determination”. In: *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*. 2021.
- [24] Liu Ziyin, Tilman Hartwig, and Masahito Ueda. “Neural Networks Fail to Learn Periodic Functions and How to Fix It”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle et al. Vol. 33. Curran Associates, Inc., 2020, pp. 1583–1594. URL: https://proceedings.neurips.cc/paper_files/paper/2020/file/1160453108d3e537255e9f7b931f4e90-Paper.pdf.