# Rapid Classification of near-Earth and Cislunar Objects using Electro-Optical Observations

**Joseph Diamond, Michael Goodman, Peter Kent, and Stephen Sawyer**
*Peraton, Herndon, VA*

**Sasidhar Devabhaktuni, Roshan Eapen, Jonathon Hope, and Puneet Singla**
*Pennsylvania State University, University Park, PA*

## ABSTRACT

The proliferation of satellites being launched into cislunar orbits is accelerating, following renewed global interest in lunar missions, a fast-growing commercial space industry, and evolving military ambitions. These objects may occupy orbits that traverse near-Earth, potentially posing hazards or threats to high-value Geosynchronous Earth Orbit (GEO) satellites; however, the ability of current Space Situational Awareness (SSA) systems to observe and track these objects and to assess these potential threats/hazards remains insufficient. Despite a mature and significant sensing apparatus for the observation of geostationary objects within public and private sectors, the processing of these observations is based on 2-body physics that is ineffective for cislunar orbits. Therefore, the extent of observational coverage of near-Earth cislunar objects by existing sensors is difficult to assess; because such observations are not associated with a 2-body catalog, they are destined to become hopeless uncorrelated tracks (UCT) or discarded as 'bad data.' However, if cislunar observations could be classified as such, the potential for existing GEO-focused sensor systems to detect and track cislunar objects could be better understood. Further, once classified, these cislunar observations could apply appropriate 3-body physics for association (to a cislunar catalog), orbit determination, orbit propagation, and conjunction assessment (CA) processing, thereby augmenting the ability of SSA systems to protect high-value GEO assets.

## 1.0 INTRODUCTION

A sophisticated understanding of spacecraft motion in multi-body environments such as the cislunar space is essential for the forecasting and tracking of objects for persistent surveillance, mission safety, collision avoidance, and adversarial behavior. The primary challenges that limit the transferability of tools and techniques from the near-Earth regions to cislunar regions is non-Keplerian dynamics, data sparsity from limited coverage, and sensor availability. Further, the few astrometric observations of objects in the cislunar region result in large measurement cadences, and in many cases, limit the amount of information required to compute a full orbit according to the least squares principle. This paper investigates whether the currently available limited measurements contain significant orbit information to enable rapid classification of trajectory as sub-GEO, GEO, or cislunar.

In this paper, we present a new approach for the rapid classification of near-Earth and cislunar objects using electro-optical (EO) observations. As space operations extend beyond GEO into the cislunar region, traditional SSA techniques are challenged by non-Keplerian dynamics, limited sensor coverage, and sparse observational data. We propose a classification framework based on admissible regions (AR) of orbital parameters that use angular and angular-rate measurements (from a single track or track sequence comprised of multiple optical measurements) to determine whether an object follows 2-body near-Earth dynamics or 3-body cislunar dynamics. Our methodology builds upon principles from asteroid identification and applies them to satellite surveillance to enable immediate classification from observations or improved accuracy with additional measurements [1]. More recently, Dinh et al. [2] have used the ARs for orbit determination of libration point orbits within the Earth-Moon (EM) system.

SSA for deep-space objects (GEO) has been traditionally based on sensor-tasking orchestrated via a 2-body catalog of resident space objects (RSOs). While collections from such actively tasked EO telescopes may produce incidental collections of cislunar objects, the emergence of passive, wide field-of-view sensors provides the potential for the creation of longer, concurrent observation tracks from a single sensor. The emerging passive SSA sensor platforms

are an evolutionary spinoff from networks such as the Global Fireball Observatory [3] network of astronomical cameras. These imagery cameras were installed and operated to detect meteoroid fireballs as they enter Earth's atmosphere [4]. Australia's continental Desert Fireball Network [5] served as a pathfinder imagery network. One example of an emerging passive SSA imagery platform is the FireOPAL (Fireball OPtical ALert) network of observatories operated by Lockheed Martin Australia [6]. FireOPAL fixed camera/sensing arrays collect a mix of EO RSO signatures from distances near and far. Passive EO observations typically require long image exposures for analysis or stacking of multiple sequential images. From these observations, we can readily estimate the angular rates of an object as it traverses space along its trajectory. Passive SSA observatories typically co-locate multiple cameras, each skewed slightly to overlap and produce an extended composite field of view, thereby creating a camera/sensing array that functions as one element within a coherent network of sensing arrays to observe the large number of artificial satellites passing simultaneously overhead during timeframes conducive to EO sensing.

The wide area surveillance and precision tracking capabilities of passive EO systems generate near real-time satellite measurements suitable for the extraction of RSO angle rates, either through measurement post processing or multi-measurement consolidation. Further, if we can constrain the space of ranges and range-rates associated with angle-only measurements, we can distinguish between near-Earth and cislunar RSOs by constructing an AR set of orbits and trajectory types associated with an RSO. Minimization of the AR through iterative analytical processing reveals the orbit type and provides a level of confidence in the classification result. Our work builds on the AR concept applied earlier for asteroid identification [5]. The work presented here used mocked up sensor outputs to simulate UCTs produced by passive EO systems such as FireOPAL. These simulated track segments generated inputs into an orbit classifier produced by Penn State University (PSU) under a Peraton-sponsored research agreement. Researchers then augmented the simulated tracks with as-flown ephemeris from current and past space missions operating in the regions of interest.

The goal of this research is to develop a classifier to distinguish between 2-body and 3-body dynamic objects based on EO measurements. This work is supported by a pipeline we developed that generates physically accurate synthetic observation measurements for testing and validation purposes. For our discussion, 2-body dynamics equates to the terms 'near-Earth,' 'GEO,' or 'sub-GEO,' while 3-body dynamics equates to 'beyond GEO' (xGEO) or 'cislunar'. Overlap between category labels can occur, and RSOs can be subject to both sets of dynamics over portions of their orbits or trajectories.

The main result of this study is a classification framework that improves SDA capabilities. This framework supports a more robust space catalog, improves initial orbit determination (IOD) by supporting sparse measurement conditions, and enables more efficient sensor tasking. We highlight improvements to SSA and persistent tracking, which meet space traffic management and national security goals. We also illustrate the viability and significance of this strategy and highlight the ways in which classification methods can fill in observational gaps in cislunar surveillance and enhance the capacity to monitor and forecast RSO activity across the cislunar domain.

The following descriptions define the key terms used in this paper and clarify our use of these terms:

- *Epoch:* A specific point in time used as a reference for defining orbital or trajectory parameters or timestamping observed measurements.
- *Measurement:* A single value of a spatial or physical property of an RSO as observed by a sensor at a recorded epoch.
- *Raw measurement:* A single basic measurement used as input into calculations, formulas, or algorithms. The raw measurements used in our work are azimuth ($\alpha$) and elevation ($\delta$), i.e., angles-only.
- *Derived measurement:* A single measurement derived from one or more raw measurements obtained at epoch(s) appropriate to the derivation. The measurements derived in our work are azimuth rate ($\dot{\alpha}$) and elevation rate ($\dot{\delta}$).
- *Measurement set:* A set of raw and derived measurements representing a partial or complete state of an RSO at a single specified epoch. The measurement sets used in our work are [$\alpha$, $\delta$, $\dot{\alpha}$, $\dot{\delta}$]

- *Observation:* A set of measurements of the spatial or physical properties of an RSO, taken by a single sensor at a single epoch.
- *Observation Window:* Period during which a sensor can observe and track an RSO; the time between the start and end of a window, designated as $\Delta t_{ObsWindow.}$
- *Track:* A set of multiple sequential observations of a single RSO collected by a single sensor over the course of a specific period of continuous observation. The time between observations within a track is shown as $\Delta t_{ObsGap}$. The time between the start and end of a track is shown as $\Delta t_{Track}$.
- *Track sequence:* A set of multiple sequential tracks of a single RSO linked together with track-to-track time gaps. The time between the start of the first track and the end of the last track is shown as $\Delta t_{TrackSeq}$. The time between the end of one track and the start of the next track is shown as $\Delta t_{TrackGap}$.
- *Hypothesis:* A candidate orbital state created by combining observed angle measurements, measurement-derived angle rates, and conjecture regarding associated range and range-rates ($\rho$, $d\rho/dt$). Hypothesized states populate the AR grid that supports the classification of UCTs as near-Earth or cislunar RSOs.
- *Residual*: Difference between what is taken as the actual value of an orbital state at some point and its approximate or calculated value (based on measurements, derivations, and hypothetical elements) at that point. If a residual is "too large" in comparison to the applied criteria, the hypothetical range and range-rate used for that point is considered to be a poor match to the actual value and is reflected in further steps of the overall classification algorithm.

We have organized this paper as follows. Section 2 presents the problem formulation which guided our research and development work. Section 3 presents the methods we developed to ingest sensor system output and to classify each RSO as near-Earth (sub-GEO and GEO) or as cislunar, along with an associated classification confidence level. Section 4 presents our results and our evaluation of the results, and finally, Section 5 discusses the results and presents conclusions from the technical, operational, and mission focus perspectives.

## 2.0 PROBLEM FORMULATION

Given RSO observation O, the following constraints apply to physically acceptable solutions of cislunar orbits or trajectories: D1: O is not of a planetary body, asteroid, or comet; D2: the trajectory of O is influenced by Earth and Lunar gravity; and D3: the trajectory of O goes beyond GEO. The union of D1, D2, and D3 then form a set, i.e. a region, that contains admissible solutions. We apply insight from 2-body trajectory representations and a restricted 3-body dynamical model to construct the AR in the range and range-rate space. Linking this data with RSO observation track information on angles and angle-rates provides sufficient information to enable the reconstruction of potential initial conditions to explain the observed track. One can impose further constraints that distinguish between near-Earth and cislunar trajectories with the goal of reducing the AR to achieve higher levels of confidence in the solution.

High-level filtering based on the apparent angular rate ($\omega$) is used to flag likely LEO and MEO tracks as 'near-Earth' on short arcs. Apparent angular speed alone cannot reliably distinguish GEO and xGEO tracks, so the admissible region (AR) approach is applied to those for classification. Track length is an important factor in classification outcomes. SSA observatories must either configure image cadences to produce the necessary track length, and/or merge additional observations into an adequate track. We also explore how the AR can be reduced further by incorporating information available across multiple observation tracks. By incorporating statistical filtering techniques and optimization methods, the process can be further refined to reduce uncertainty and improve prediction capabilities. We use a combination of synthetic data and real data collected from past and current operational missions to validate the developed methodology. Ultimately, these developments can contribute to a more robust framework for rapid-classification and subsequent IOD methods for objects in the near-Earth and cislunar environment.

# 3.0 METHOD

## 3.1 Methods and Algorithms

The methodology for classifying an RSO as near-Earth or cislunar is a multi-stage process that depends on the iterative refinement of an AR. The AR method narrows down possible solutions of equations or systems of equations that use numerical techniques as solvers, and where the available information is incomplete or indefinite. Application of the AR method starts with the definition of physics-based constraints applicable to the current problem. These constraints filter out potential solutions that are not physically acceptable, thereby leaving a solution space that contains those that are. Further constraints typically incorporate additional knowledge or hypotheses about the subject and its environment. For this classification algorithm, ssa-classify, these constraints are:

1. Observation is not of a planetary body, asteroid, or comet - implemented via SSA sensor system raw image processing to remove these objects based on star tracker catalogs or similar astrometric data.
2. The RSO trajectory is influenced by Earth and Lunar gravity - implemented by high-level filtering to remove sub-GEO objects based on radial magnitude thresholds applied to angular rates derived from raw measurements.
3. RSO trajectory extends beyond GEO - implemented by propagation of measurement sets [α, δ, $\dot{\alpha}$, $\dot{\delta}$] augmented by conjecture regarding associated range and range-rates (ρ, dρ/dt) to yield a hypothetical orbit/trajectory state.

The process begins by deriving a complete state measurement from raw optical data, using a pair of tracks to filter dynamically inconsistent trajectories, and concludes by classifying the remaining valid states to produce a final determination with an associated confidence score. The following sections detail the methods and algorithms to achieve this result.

### 3.1.1 Measurement Set Derivation from Observation Tracks

The process input is a track sequence from an EO sensor, where each track contains multiple observations at known epochs. To prepare this data for the AR analysis, we use a single track to generate a complete angular state vector. We fit the time-series of raw measurements (azimuth, α, and elevation, δ) with a cubic spline, which we then analytically differentiate to compute the derived measurements: azimuth rate ($\dot{\alpha}$) and elevation rate ($\dot{\delta}$) at a specific reference epoch within the track. The result is a single, comprehensive measurement set [α,δ,$\dot{\alpha}$,$\dot{\delta}$] that encapsulates the observed angular dynamics of the RSO at that instant.

### 3.1.2 Hypothesis Testing and AR Refinement

This stage is the core of the classification algorithm, designed to systematically eliminate physically impossible trajectories. The process uses a pair of tracks from the track sequence and operates as follows:

1. **Candidate State Generation:** An initial AR is defined as a discrete grid in the range and range-rate space. Each point on this grid represents a hypothesized range and range-rate. We combine these hypotheses with the measurement set from the first track to generate a large set of candidate six-degree-of-freedom state vectors.
2. **State Propagation:** We transmit each candidate state vector forward in time across the $\Delta t_{TrackGap}$ to the epoch of the second track. This transmission uses dynamic models appropriate for the potential trajectory type, such as 2-Body Problem (2BP) dynamics for near-Earth hypotheses or the Circular Restricted 3-Body Problem (CR3BP) for cislunar hypotheses.
3. **Residual Computation and AR Shrinking:** At the second track's epoch, we use the propagated state to predict the expected angular measurements. We then calculate residuals by comparing these predicted values to the actual measurement set derived from the second track. We consider hypothesized states that yield residuals above a predefined threshold to be dynamically inconsistent and discard these. This filtering process

shrinks the AR to leave only a subset of range and range-rate pairs consistent with the observations from both tracks. To verify the statistical significance of the final confidence score, we apply a termination threshold during the iterative refinement process. If the population count of valid hypotheses within the AR falls below 150, the algorithm terminates.

### 3.1.3 Classification and Confidence Scoring

With the refined AR containing only dynamically valid states, we perform the final classification. In this work, we propagate each remaining state within the shrunken AR for 10 days to determine the nature of its full trajectory. We classify a state as cislunar if its trajectory crosses a predefined altitude threshold; otherwise, we classify it as near-Earth. In this work, we selected a 100,000 km altitude threshold as a conservative demarcation far beyond geosynchronous orbit to establish a practical boundary where lunar gravitational perturbations become too significant for 2-body models to remain effective. This boundary verifies that any trajectory crossing this line is distinctly non-Keplerian, thereby providing a clear and computationally efficient criterion for the algorithm to differentiate between near-Earth and cislunar objects.

Further, the choice of a 10-day propagation window to classify each hypothesis is a strategic trade-off between classification accuracy and computational efficiency. This duration is sufficient to unambiguously determine whether a trajectory will cross the 100,000 km altitude threshold for the great majority of relevant cislunar orbits. For example, highly eccentric trajectories on a direct path from near-Earth to the lunar vicinity cross this boundary rapidly, often within a few days. Further, the orbital periods of key cislunar families, such as resonant orbits, typically range from approximately 7 to 20 days (e.g., 4:1 or 3:1 resonance with the Moon); the outbound leg of these trajectories from perigee breaches the 100,000 km threshold in less than half the orbital period, which places it well within the 10-day window. This duration captures the characteristic behavior of these critical trajectory classes while keeping the one-time computational load of the initial classification step manageable.

The final output is not only a binary classification but also a level of confidence. We compute the classification confidence as the ratio of states of the majority class to the total number of states remaining in the refined AR. For example, if 95 out of 100 valid AR hypotheses are classified as cislunar, the confidence for a cislunar classification is 95%.

This confidence metric is critical because the boundary between near-Earth and cislunar space is not a rigid line; an object in a highly elliptical orbit, for example, can be subject to both dynamical regimes at different points in its trajectory. To accommodate this physical reality, the specific criteria used for classification in this work - reaching a 100,000 km altitude threshold within a 10-day propagation - should be considered tunable in an operational use scenario. Tuning would be driven by the specific mission focus and the operational needs of the user. In this manner, the output of ssa-classify supports RSO custody management driven by object dynamics, which in turn promotes efficiency and effectiveness in RSO catalog construction and maintenance.

## 3.2 Admissible Regions Procedure

Figure 1 presents the AR-based procedure implemented by Peraton and PSU for the classification of RSO observation data as cislunar or near-Earth in nature. This diagram covers the ssa-classify algorithm input, output, and high-level processing flow.

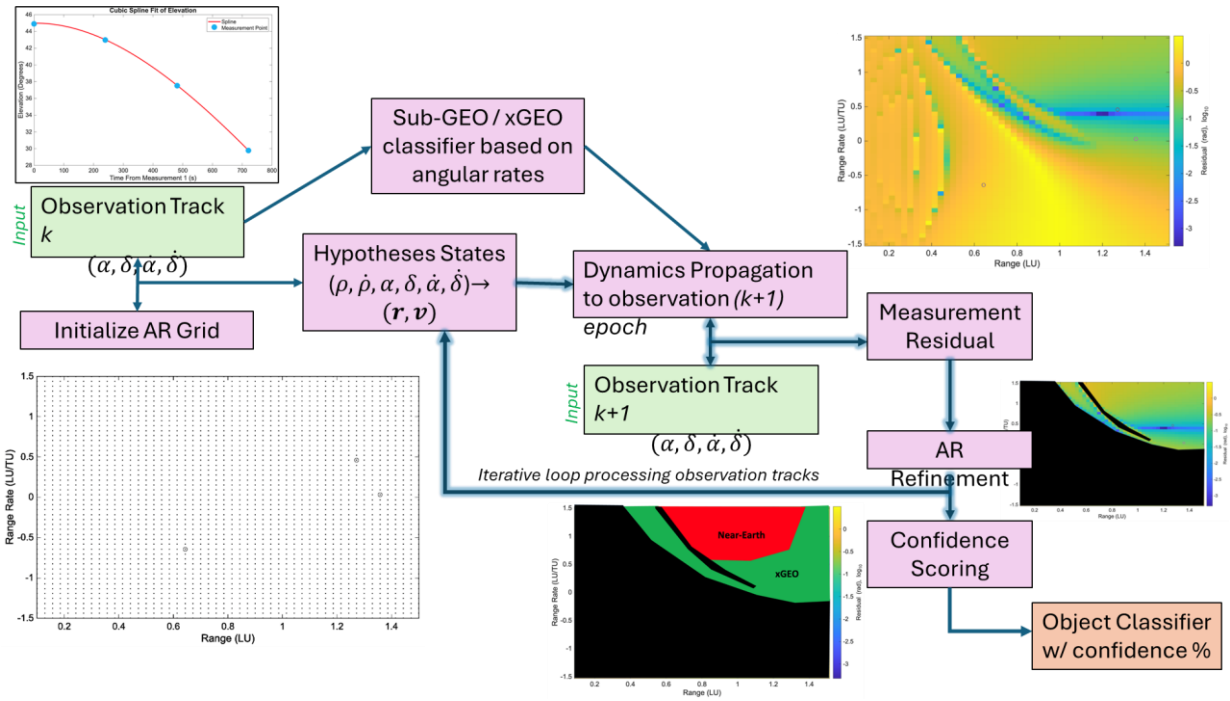


*Fig. 1. AR-based SSA Classification Procedure*

### 3.3 Measurement Realism

To maintain the robustness of the classification algorithm and to simulate realistic sensor performance, we corrupt the synthetic raw measurements with additive white Gaussian noise. We apply a standard deviation ($\sigma_{noise}$) of 1 millidegree (3.6 arc seconds).

The acceptance threshold for residuals is not a fixed value but dynamically linked to the statistical properties of this measurement noise. We compute a residual vector for each candidate state that incorporates the differences between the predicted and observed angular measurements at the second track's epoch. We then normalize this vector by the noise standard deviation ($\sigma_{noise}$). We consider a candidate state to be dynamically consistent, and we add this to the refined AR only if the magnitude of its normalized residual vector is below a specified threshold, typically 3σ. This statistical gating verifies that only hypotheses that fit the observations to within the expected measurement uncertainty are accepted for the final classification stage. Figure 3 presents the gating procedure implemented for noisy measurements.

## 4.0 RESULTS AND EVALUATION

### 4.1 Evaluation Test and Validation Cases

We conducted testing and validation of our described classification methods through a set of use cases supported by data sets that represent the output of an earth-based passive EO sensor system. We ran the developed algorithms and procedures against each case and evaluated the results with respect to classification performance (correctness), confidence level produced, solution rate (AR convergence), and algorithm run time.

A combination of purpose-built synthetic data and as-flown spacecraft-state data (from earth-based SSA observation systems or from spacecraft telemetry messages collected during past and/or current operational missions) formed a

set of UCT Track Sequences (TS) used to evaluate the ssa-classify algorithm. We constructed TSs to represent RSO orbits or trajectories within each region of the E-M system. The construction process proceeded as follows:

- From each RSO ephemeris, we extracted a set of TSs by randomly choosing a starting epoch along the RSO orbit or trajectory for each TS.
- For each TS, we extracted 5 to 10 neighboring tracks from the ephemeris which reflect what a sensor viewing the object at the TS start epoch observes and measures as right ascension and declination, for all the observations that comprise each track in the TS.
- We run TSs constructed in this way through the ssa-classify algorithm.

Figure 2 illustrates the structure of test data sets used to assess the performance and accuracy of the ssa-classify algorithm.

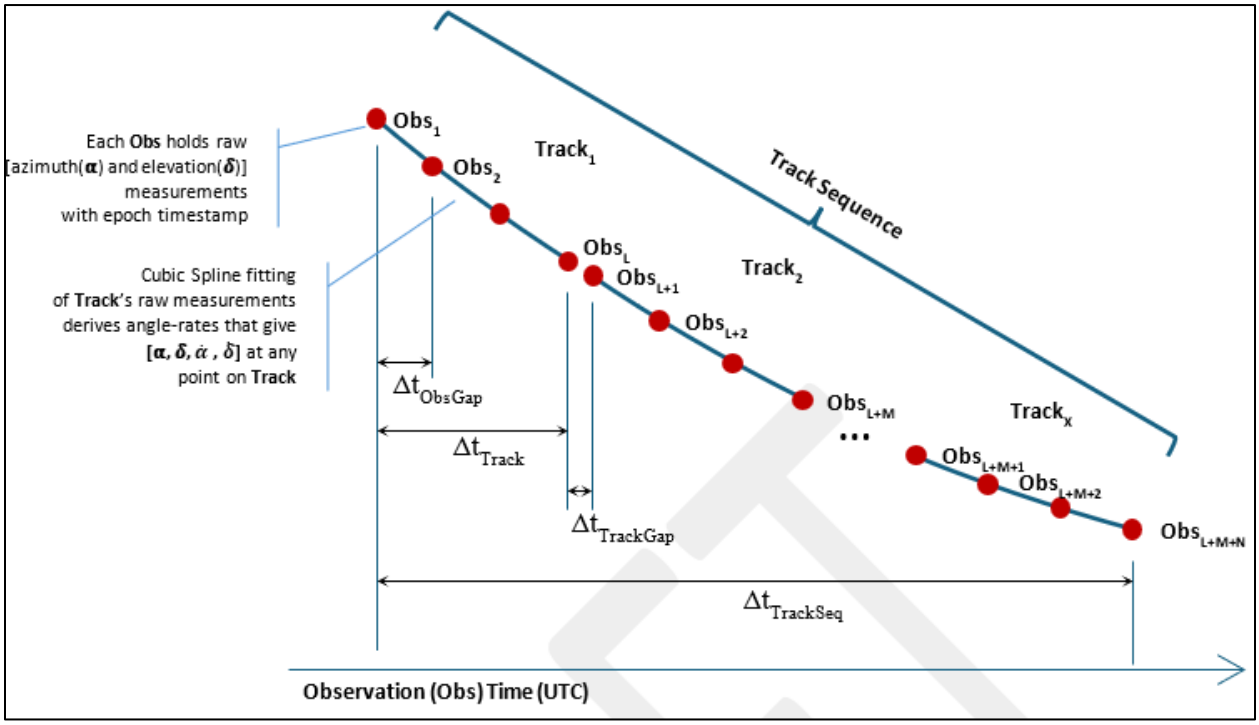


*Fig. 2:Test Data Set Construction Approach and Structure*

Table 1 presents the test and validation cases constructed to evaluate ssa-classify performance.

*Table 1. SSA-Classify Algorithm Test and Validation Cases*

| # | Family/Object | Region | Track Data | Semi-Major Axis (km) | Eccentricity |
|---|---|---|---|---|---|
| 1 | GTO | LEO/GEO | Actual | 43,878 | 0.5128 |
| 2 | GEO/A | GEO | Actual | 42,000 | 0.001 |
| 3 | xGEO | Cislunar | Actual | 70,000 | 0.001 |
| 4 | 2:1 Resonant | Cislunar | Actual | note 1 | note 2 |
| 5 | 3:1 Resonant/A | Cislunar | Actual | note 1 | note 2 |
| 6 | DRO | Cislunar | Actual | note 1 | note 2 |
| 7 | Molniya | MEO | Actual | 26,600 | 0.732 |
| 8 | MEO/A | MEO | Actual | 20,181 | 0.0116 |
| 9 | LEO | LEO | Synthetic | 7,500 | 0.0667 |
| 10 | MEO/B | MEO | Synthetic | 23,500 | 0.5319 |

| # | Family/Object | Region | Track Data | Semi-Major Axis (km) | Eccentricity |
|---|---|---|---|---|---|
| 11 | xGEO/HEO | Cislunar | Actual | 70,686 | 0.8064 |
| 12 | GEO/B | GEO | Actual | 42,248 | 0.0001 |
| 13 | 3:1 Resonant/B | Cislunar | Actual | note 1 | note 2 |

Notes:

1. Varies across orbit; a 3-body propagator used for classification
2. Varies due to 3-body perturbations
3. LEO Low Earth Orbit
   MEO Medium Earth Orbit
   GEO Geosynchronous or Geostationary Earth Orbit
   GTO Geostationary Transfer Orbit
   HEO Highly Elliptical Orbit
   DRO Distant Retrograde Orbit
   xGEO refers to orbits that extend beyond GEO
   (including the region of space out to cislunar, cislunar itself and, beyond to the E-M Libration points)

We used the following three test cases to evaluate the performance of the ssa-classify algorithm.

*Table 2. SSA-Classify Algorithm Test and Validation Parameters*

| | Test Case 1 | Test Case 2 | Test Case 3 |
|---|---|---|---|
| Number of track sequences run | 100 | 100 | 100 |
| Track Sequence Length $\Delta t_{TrackSeq}$ | Varies | 75 minutes | 570 minutes |
| Number of Tracks per Sequence | Varies | 10 | 10 |
| Track Length $\Delta t_{Track}$ | 3 minutes | 3 minutes | 3 minutes |
| Track to Track Gap $\Delta t_{TrackGap}$ | Varies | 5 minutes | 60 minutes |
| Number of Obs per Track | 4 | 4 | 4 |
| Obs to Obs Gap $\Delta t_{ObsGap}$ | 1 minute | 1 minute | 1 minute |

### 4.2 SSA-Classify Performance

We evaluated the ssa-classify algorithm against the results for accuracy and confidence level, how the algorithm responded to assignments of ranges and range-rates to create hypotheses, how the algorithm responded to measurement noise, and how it performed speed-wise.

#### 4.2.1 Test Case 1: SSA-Classify Performance Based on Track-to-Track Gap

We show the results of running ssa-classify against the cases in Table 1 in Figure 6. To construct the track sequencies for Test Case 1, we selected ten starting epochs for each object from the ephemeris data. For each starting epoch, we repeated the process outlined in 4.1 and varied the track gap across the set intervals: 5 Minutes, 30 Minutes, 1 Hour, 2 Hours, 3 Hours, 4 Hours, 5 Hours, 12 Hours, 18 Hours, and 24 Hours, generating 100 test cases.

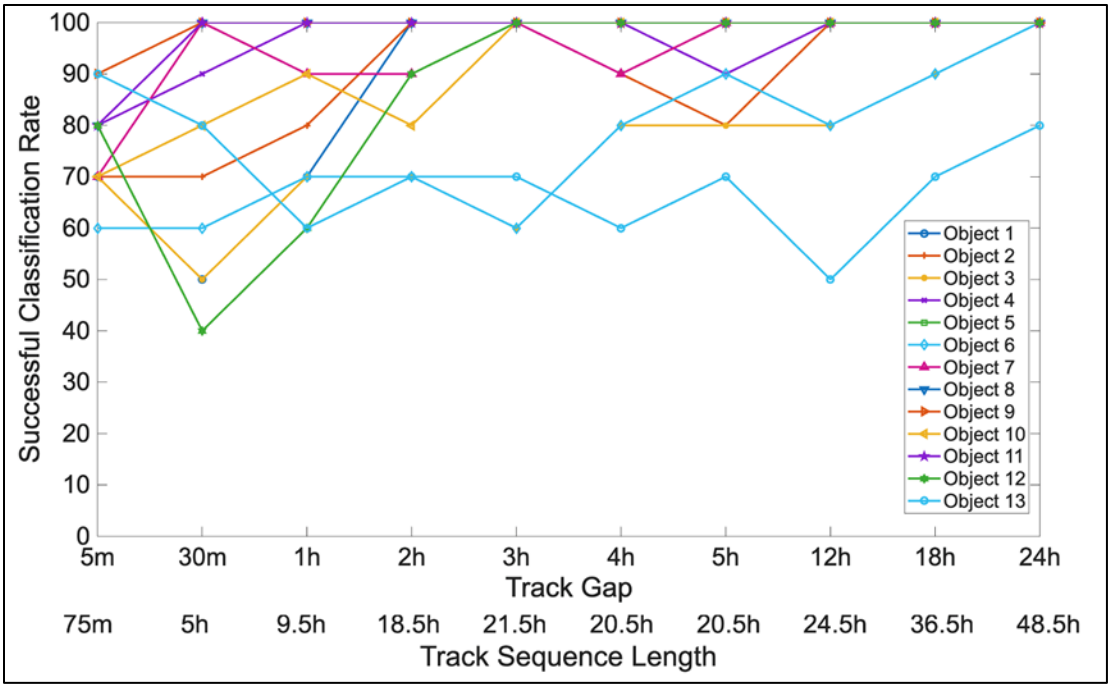


*Fig. 3. Test Case 1 Results*

A clear trend emerges in Figure 3, where the successful classification rate increases with the track-to-track gap ($\Delta t_{TrackGap}$). This behavior is rooted in the basic physics of orbit determination from angles-only data. Short time gaps (5-60 minutes) provide a robust estimate of the object's immediate angular position and rates but offer poor observability of the trajectory's curvature. It is this curvature that contains critical dynamic information that distinguishes a predictable Keplerian ellipse from a complex trajectory perturbed by a third body.

As $\Delta t_{TrackGap}$ increases, the object traverses a longer arc, which makes its true dynamic signature observable in the data and enables the algorithm to effectively reject large portions of the hypothesized AR space based on an incorrect dynamic model. The number of tracks processed to reach a final classification provides a direct measure of the observational effort required. Given the total time spanned by a track sequence ($\Delta t_{TrackSeq}$), the duration of each individual track ($\Delta t_{Track}$), and the time gap between them ($\Delta t_{TrackGap}$), the number of tracks ($N_{Tracks}$) can be calculated. For example, we used the values from Test Case 1 above - a 75-minute track sequence length, 3-minute track lengths, and a 5-minute track gap - 10 tracks.

The performance curve also begins to plateau around a 4–5-hour track gap, which is a significant operational finding. This "knee" in the curve represents a point of diminishing returns, where waiting longer for another track does not substantially improve classification accuracy but increases the time-to-decision. This situation suggests an operational sweet spot for sensor tasking. We also note that with a track separation of 4-5 hours, we needed only 2-3 tracks for a confident classification, compared to upwards of 9 tracks for gaps of an hour or less.

### 4.2.2 Test Case 2&3: SSA-Classify Performance Comparison Between two Different Track Gaps

Figure 4 compares the outcome of 100 TSs between the track-gap conditions of test cases 2 and 3: 5-minute track separation and 1-hour track separation, respectively.

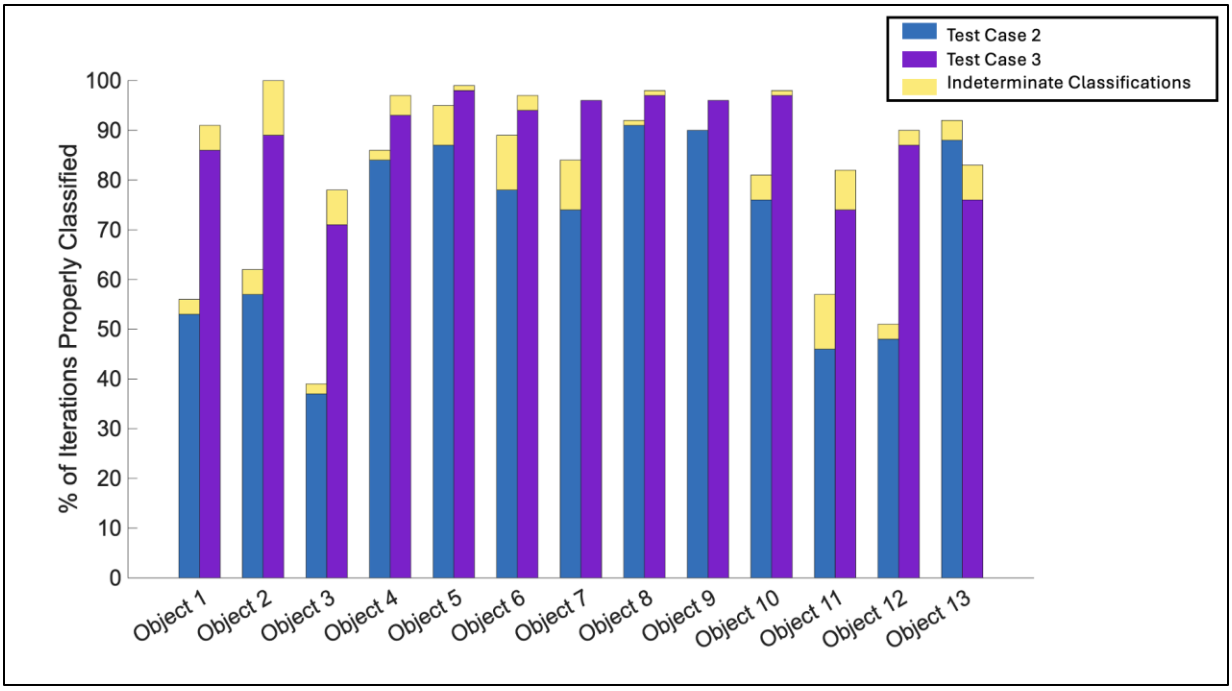


*Fig. 4. Test Case 2 and 3 Results*

The results show a clear and consistent improvement in the classification success rate for the longer 1-hour track gap (Test Case 3) across all object types and reinforces the conclusion that a wider temporal separation between tracks is critical to enable the algorithm to effectively shrink the AR and to converge on the correct classification. Further, the degree of improvement is not uniform across all objects, which suggests that the algorithm's sensitivity to track separation is also influenced by the object's specific orbital regime, with some trajectories revealing their true dynamical nature more quickly than others.

### 4.2.3 AR Refinement for Two Representative Cases

The evolution of the AR structure provides a direct visualization of the information gain during the classification process. The initial AR is characteristically elongated along the range ($\rho$) axis, which reflects the fundamental ambiguity of angles-only observations where an object's distance is almost entirely unknown. This elongated region represents the manifold of all dynamically plausible states in the ($\rho, \dot{\rho}$) space consistent with the first observation track. With each subsequent track, the algorithm propagates every hypothesis forward and trims those that produce residuals inconsistent with the new observation, thereby shrinking the AR.

The figures below illustrate this process for a cislunar and a near-Earth object, respectively. For a 2:1 resonant object (Figure 5), correctly classified as cislunar, the process consistently eliminates near-Earth (blue) hypotheses, thereby retaining a high ratio of cislunar (green) states as the AR shrinks. This process leads to a final, confident classification of 'cislunar' with 90% confidence.

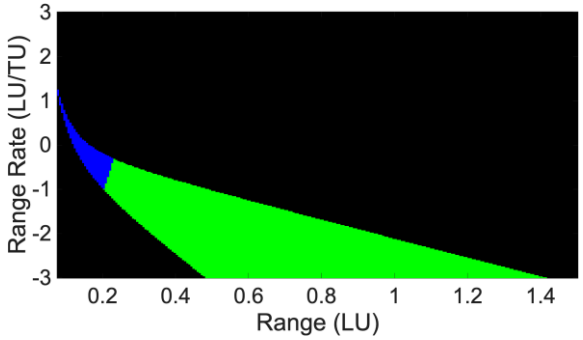


(a) Iteration 1: 94%

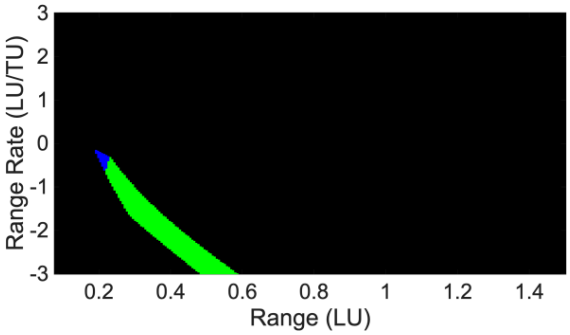


(b) Iteration 2: 95%

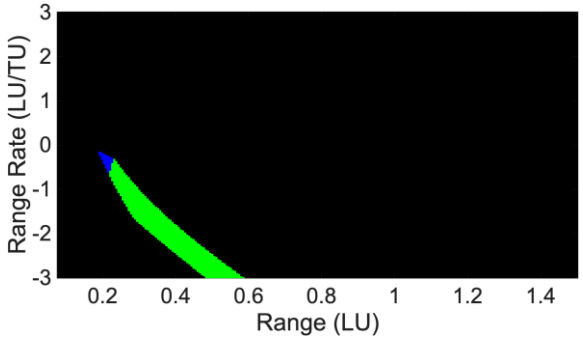


(c) Iteration 3: 91%

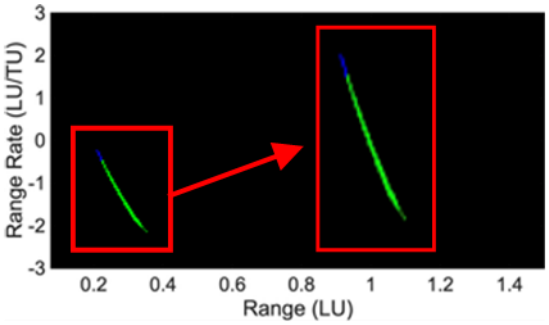


(d) Iteration 4: 90%

*Fig. 5. Successive Admissible Regions Post Propagation, 2:1 Resonant*

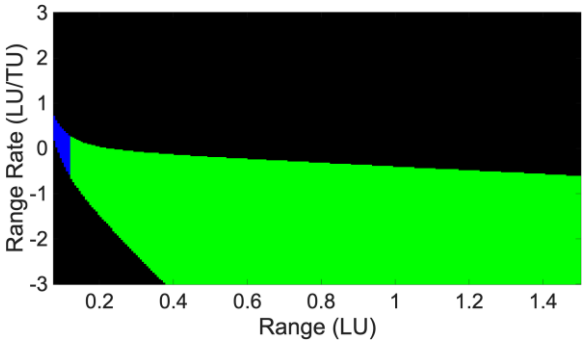


(a) Iteration 1: 99%

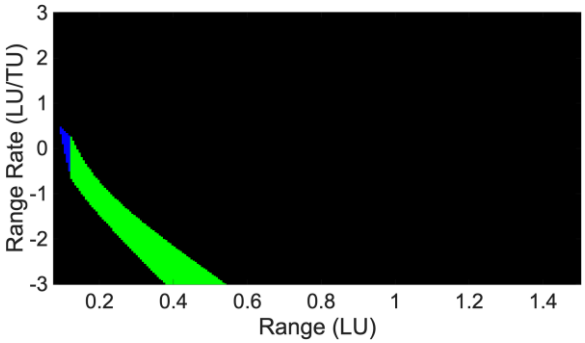


(b) Iteration 2: 95%

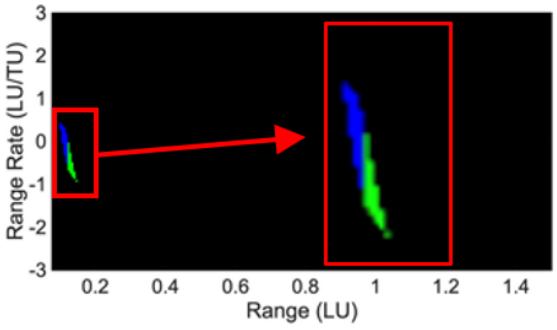


(c) Iteration 3: 46%

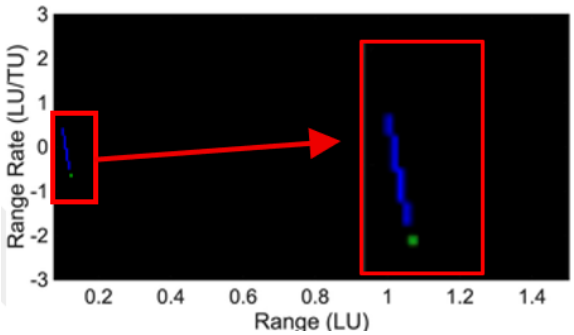


(d) Iteration 4: 7%

*Fig. 6. Successive Admissible Regions Post Propagation, GEO/B*

Conversely, Figure 6 demonstrates the power of the algorithm to resolve the initial ambiguity for a near-Earth object. The GEO/B satellite (Object 12) is correctly classified as "near-Earth" (blue). The percentages shown in the figure reflect the proportion of states classified as cislunar (green). Therefore, the confidence in the correct, near-Earth classification is increasing with each iteration: from 1% (100-99) in Iteration 1, to 5% (100-5), then 54 (100-46), and finally converging to a confident 93% (100-7) in Iteration 4. The dramatic shift from a cislunar-dominated AR to a near-Earth dominated one demonstrates the ability of the algorithm to effectively converge on the correct dynamical regime. We discard a hypothesis only if the residual exceeds a predefined statistical threshold, in this case, 3σ, where σ represents the standard deviation of the measurement noise. This statistical gating verifies the strength of the shrinking process and retains all dynamically plausible states that fall within the expected error bounds of the sensor, rather than making a brittle binary decision based on a perfect match.

### 4.3 SSA-Classify Algorithm Speed

We executed the ssa-classify algorithm on an Apple MacBook Pro with an M4 Pro Chip and 24 GB RAM. We parallelized functions across 10 cores when possible to reduce execution time. On average, the algorithm processes a single-track sequence in 25 seconds, with runtimes ranging from 6 to 60 seconds.

A key aspect of the algorithm's design contributes to this efficiency. The most computationally intensive step is the initial 10-day propagation of every state in the original AR grid to determine its classification (near-Earth or cislunar). We perform this determination only once at the beginning of the process. For all subsequent iterations, as the AR shrinks from new measurements, the algorithm does not need to re-propagate the remaining states, but references the pre-computed classification for each valid hypothesis. The result is that subsequent iterations are significantly faster, and the overall runtime scales largely with the size of the shrunken AR, making the approach suitable for near-real-time analysis.

## 5.0 DISCUSSION AND CONCLUSIONS

We described and analyzed a method for binary classification of RSOs observed by passive E-O sensor systems. We evaluated and described the accuracy and repeatability of the AR-based methods and algorithms under simulated sensor input. Our evaluation shows the AR classification algorithm to be most applicable in use-cases where tracks or track sets can be accessed or assembled to provide observations and measurements covering a temporal span exceeding 30 minutes. In these cases, the number of observations, and their spatial-temporal spread, supported construction of a range and range-rate AR space conducive to iterative shrinking of the space and production of an accurate discrimination between near-Earth (GEO and sub-GEO) and cislunar orbit and trajectory. The AR algorithm performed best in cases with track-to-track gaps of four hours or more and resulted in confidence levels that typically ranged between 90% and 95%. Specific customer needs will drive the use of an AR algorithm such as ssa-classify based on operational use cases.

### 5.1 Technical Summary

We demonstrated that the ssa-classify algorithm works with observation inputs that reflect today’s emerging passive staring EO systems. While short tracks and short TSs readily rejected sub-GEO objects and classified GEO objects with confidence levels exceeding 93%, classification of cislunar objects requires significantly longer tracks and TSs, depending on the specific orbit family and orbital path with respect to the observer. Because RSOs in distant cislunar orbits or trajectories move significantly slower than objects in near-Earth orbits, they require track lengths of multiple hours or even days that are long enough to capture discernable movement. Tracks of this nature place an operational requirement on SSA observatories to deliver a single track of this duration or to construct a track sequence of similar duration.

The computational complexity of the ssa-classify algorithm scales linearly with the number of initial hypotheses, which can be expressed as $O(N_{grid})$, where $N_{grid}$ is the number of discrete points in the initial AR grid, because the primary computational driver is the numerical propagation performed for each candidate state. While subsequent

iterations operate on a progressively smaller (shrunken)AR, the overall runtime remains dominated by the size of the initial hypothesis set.

We observed the execution speed to reduce and collapse the AR to produce accurate, confident, classifications to be less than 60 seconds when running on software-development class workstations using multi-core configurations and the Matlab parallel computing toolbox [7]. We used the circular restricted 3-body problem model (CR3BP) of the E-M system throughout our research. Nothing in the performance of ssa-classify or the results it produced indicates that 4-body modeling or inclusion of perturbations due to a $4^{th}$ body are necessary.

### 5.2 Operational Summary

The motivation behind the development of an algorithm such as ssa-classify is to deliver capabilities that operational SSA sensor systems can use to extend and/or augment current capabilities, or to serve as a template for capabilities designed into new SSA sensor systems. This confidence-based classifier serves as a critical decision gate for operators by dictating whether to employ standard 2-body SSA pipelines or more complex 3-body models for crucial downstream tasks such as orbit determination, data association, and conjunction analysis. The operational requirements associated with using ssa-classify apply to several areas described below. An SSA sensor system must satisfy requirements such as those presented in Section 3 to successfully deploy and operate ssa-classify:

- Observations, measurements, track length, track gaps, and TSs
- Time dependent noise modeling and measurement uncertainty within specified bounds
- Observation processing workflows or pipelines supporting AR constraints

The benefits of deploying a capability to provide timely classification of observations/tracks as cislunar depend on the mission of a specific SSA sensor system, but can include the following:

- Enable use of existing GEO sensing to find/track cislunar objects to potentially reduce the need for dedicated cislunar observatory/sensor coverage.
- Enable association of otherwise uncorrelated cislunar observations to convert GEO UCTs into associated cislunar observations.
- Support threat/hazard assessment from near-Earth cislunar objects that pose a potential or real threat to high-value GEO assets.
- Inform the development of new cislunar sensing capabilities by applying understanding of existing sensing capability to already-available GEO sensing capabilities.
- Inform the development of capabilities to build cislunar RSO candidates from a body of uncorrelated observations or super TSs to populate a cislunar catalog.

### 5.3 Misson Focus

Successful application of the ARC techniques to the following mission focus arena are feasible, based on our evaluation: (a) Developing a quick classifier to identify xGEO objects from the observations near the Earth, (b) developing a framework to generate synthetic optical data that captures true dynamics and camera and telescope properties, and to evaluate a probability of detection (which can be used as an additional parameter in the classification), (c) Use of UDL data- a distributed network of standalone, fixed mount optical sensors designed to monitor a large region of the sky- for validation of developed technologies, (d) a robust pipeline for cislunar space domain awareness and cislunar catalog population and maintenance.

### 5.4 Future Investigation

Additional work is required to further understand the relationship between $\Delta t_{TrackSeq}$, the total number of tracks in the track sequence, $\Delta t_{Track}$, the number of observations per track, $\Delta t_{ObsGap}$ and the subset of the orbit observed in the track sequence used in the construction and assessment of the AR, and the probability of successful classification.

- Expanding on our initial investigation, we expect analysis will show the $\Delta t_{TrackSeq}$ term as a function of the observed object's orbital period to have the greatest effect on success. Such investigations would include evaluations and trade studies designed to optimize the relevant parameters and thresholds (number of tracks, track length, track separation, number of data points per track, angle rate thresholds, and cutoff altitude).
- We expect further analysis to evaluate the suitability of the AR configuration thresholds or checks, such as the use of the 100,000 km distance as a threshold for determining an object is cislunar, to provide meaningful positive impacts on operational outcomes.
- Consideration of additional checks to increase the confidence of the classification of an object determined to be xGEO through the algorithmic consideration of potential edge cases.
- AR performance should also be assessed in depth for sensitivity to measurement noise. This assessment can then drive the requirements that SSA sensor systems must satisfy to deliver observations that can be analyzed using the AR techniques described here.
- Further development of algorithms that can search through a large body of UCT Data to find TSs suitable for assessment using AR.
- Subsequent work to assess the feasibility of conducting multi-telescope collections to hand off tracks from one telescope to another in sequence, to obtain a longer $\Delta t_{TrackSeq}$ than can be obtained by a single site. Additionally, the AR approach could be extended to support improved accuracy in the assessment of range and range-rates ($\rho$, $d\rho/dt$) by fusing data collected simultaneously from two geographically separated telescopes. A multi-sensor AR approach should provide classification results more quickly.

## REFERENCES


[1] Milani, A. & Gronchi, Giovanni Federico & De' Michieli Vitturi, Mattia & Knezevic, Zoran. (2004). Orbit determination with very short arcs. I - Admissible regions. Celestial Mechanics and Dynamical Astronomy. 90. 57-85. 10.1007/s10569-004-6593-5.

[2] Dinh, Q., Scheeres, D. and Holzinger, M., Cislunar Initial Orbit Determination Using Sensor and Measurement-Centric Admissible Regions. https://www.amostech.com. Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference, September 2024.

[3] Devillepoix, HAR. Global Fireball Observatory. 2018-2025. https://gfo.rocks. Accessed 8 August 2025.

[4] Towner, Martin C et al. "Fireball streak detection with minimal CPU processing requirements for the Desert Fireball Network data processing pipeline. *Publications of the Astronomical Society of Australia* **37**, e008, 1–12. https://doi.org/10.1017/pasa.2020.48

[5] Devillepoix, HAR. Desert Fireball Observatory. 2025. https://dfn.gfo.rocks. Accessed 8 August 2025.

[6] Lockheed Martin Corporation. Space Domain Awareness. 2018-2025. https://www.lockheedmartin.com/en-au/products/space-domain-awareness.html. Accessed 8 August 2025.

[7] The MathWorks, Inc. Parallel Computing Toolbox. 1994-2025. https://www.mathworks.com/help/parallel-computing/index.html. Accessed 18 August 2025.