

Cislunar initial orbit determination using CAR-MHF

**Paul Billings¹, Jason Baldwin², Charles J. Wetterer¹, John Gaebler¹, Christopher Craft¹,
Micah Dille¹, Keric Hill¹, and Jill Bruer³**

¹*KBR*, ²*Complex Futures*, and ³*Air Force Research Laboratory*

1. ABSTRACT

High levels of nonlinearity inherent in the cislunar orbital regime present stressing challenges for the estimation and custody of cislunar objects. Despite this, effective orbital determination strategies using Unscented Kalman Filters (UKFs), Gaussian Mixture Models (GMMs)—and, in the most stressing cases, Adaptive Entropy-based Gaussian-mixture Information Synthesis (AEGIS)—have been demonstrated. Before these track maintenance functions can be accomplished, however, initial orbit determination (IOD) must be achieved. The Constrained Admissible Region Multiple Hypothesis Filter (CAR-MHF) was previously developed by AFRL for IOD, employing an iterative filter-smoother for estimate refinement and the use of a multi-hypothesis, joint probabilistic data association (MH-JPDA) framework for multi-target data association. In this paper, simulated Earth-based and space-based angles-only observations of cislunar objects in various orbital families are utilized to evaluate the ability of CAR-MHF to perform IOD successfully in the cislunar regime. Cold-start catalog generation scenarios are considered in the presence of ballistic objects across a diverse set of orbit families, constellations, object separation events, and break-up events. Of critical importance in the cislunar regime is the use of a constrained admissible region to bound the landscape of considered orbits given the presence of a single short-arc tracklet of angles-only measurements. Significant treatment is given to the trade-offs involved in bounding the scope of admissible orbits with a focus on uncertainty evolution, sparseness of the hypothesis sampling, compute load, and other considerations. CAR-MHF shows promising results across the spectrum of evaluated scenarios and opens the door to information-based optimization of sensor tasking.

2. INTRODUCTION

Cislunar space, the orbital regime encompassing the Earth and Moon, is of growing interest and importance in space operations. As deep space capabilities continue to advance, and as commercial, scientific, and national activity increase in this domain, the ability to track and monitor objects in cislunar space becomes vital to addressing both awareness and sustainability concerns. However, the inherent dynamical complexities and non-linearities of the cislunar orbital environment pose significant challenges for conventional orbit determination methodologies. In particular, initial orbit determination (IOD)—the process of first establishing an object’s orbit from a minimal set of observational data—proves notably challenging in this regime.

A number of advanced methods have been proposed to tackle the challenges presented in the tracking and estimation of cislunar objects. A general explanation of these challenges and an in-depth overview of the resulting periodic orbits was provided in [1]. Techniques such as the unscented Kalman filter (UKF) and Gaussian mixture models (GMMs) have shown promising results in the context of cislunar orbit determination, demonstrating their utility in dealing with the non-linearity of the cislunar regime [2]. Furthermore, for extremely challenging cases including extended gaps in observational measurements on the order of weeks, approaches like the Adaptive Entropy-based Gaussian-mixture Information Synthesis (AEGIS) have been demonstrated [3].

It appears that the predominant observational approach for handling cislunar IOD and uncorrelated track (UCT) resolution will be with short-arc electro-optical (EO) measurements, given the significant extent to which this observation type dominates the IOD/OD process for geosynchronous Earth orbit (GEO). Much research has been conducted over the last two decades exploring approaches for improving IOD performance for uncorrelated tracks in GEO. Admissible regions have been popular in providing a foundation for exploring the orbital possibility space in the presence of short-arc EO measurements [4-6]. Multiple hypothesis tracking also has demonstrated success [7]. Other ongoing efforts focus on extensions of these methods to the cislunar regime [8-9]. Use of an optimization scheme bounded by admissible region constraints is also under investigation [10].

The extension of methods developed for the constrained admissible region multiple hypothesis filter (CAR-MHF) is pursued in this work, building upon previous developments at AFRL [5]. Involved in this methodology is the use of an iterative filter-smoother for estimate refinement which, when combined with a multi-hypothesis, joint probabilistic

data association (MH-JPDA) framework for multi-target data association, has proven quite effective for addressing the multiple hypothesis problem encountered when tracking numerous and closely spaced objects. The current investigation addresses 1) how effectively CAR-MHF can be extended to support IOD in the cislunar regime, and 2) the extent to which CAR-MHF (and similar probabilistic methods) can support future operational considerations such as sensor follow-up, high data utilization, etc.

An in-depth study assessing these concerns is presented. Extensive simulations are conducted employing both Earth-based and space-based angles-only observations on a population of objects selected from a set of periodic cislunar orbit families. Trade-offs involved in the application of CAR-MHF are also discussed, considering factors such as uncertainty evolution, data sparseness, hypothesis sampling density, graduation timelines, computational load, and data utilization. A high-level description of the CAR-MHF algorithm, including the extensions allowing for applicability in the cislunar regime, is provided in Section 2. The experiment and the simulations performed to generate the data sets under consideration are outlined in Section 3. Section 4 contains the experimental results and a discussion of the trade-offs. A summary of conclusions and a discussion about necessary future investigations are included in Section 5.

3. EXTENSION OF CAR-MHF INTO CISLUNAR SPACE

In traditional space situational awareness (SSA), an analyst or operator's intuition is based largely on assumptions rooted in perturbed 2-body dynamics (e.g., orbits resemble simple conic sections, like ellipses). Expanding SSA, and in particular IOD, to apply to cislunar space effectively requires that new intuitions be built for the complete dynamical system. Specifically, it is important to understand the behavior or time evolution of objects and phenomena in this more complex domain, where the traditional simplifying assumptions are no longer valid (i.e., where orbits are non-Keplerian).

The question at hand with IOD in cislunar space is the extent to which previous methods (and their assumptions) will hold in their effectiveness in evaluating and resolving UCTs. This section will begin with a look at how the constrained admissible region (CAR) can be extended into the cislunar regime with the loosening of two-body constraint assumptions and with the application of a new constraint to bound the unwieldy vastness of cislunar space. CAR-MHF has a distinct advantage over other approaches in that it immediately forms a probability density function (PDF), using only a single observation track (ideally three or more measurements with sufficient intra-track spacing), to be scrutinized by future observations. This allows for the possibility of tasking sensors within a sensing network to return and collect more data on an object when the information situation, encapsulated by the time evolution of the estimated PDF, dictates such. In addition to these topics, the salient features of CAR-MHF will be discussed in brief in the remainder of this section; readers should consult the references for additional detail.

3.1 Constrained Admissible Region

The premise of utilizing a CAR for IOD is to use two-body orbital element constraints in semi-major axis and eccentricity to bound the search space of possible orbits described by a short-arc track of EO measurements [5, 11]. This short arc of EO measurements, used to obtain not just the angles (right ascension and declination) but also the angular rates (right ascension rate and declination rate), effectively observes four of the six dimensions required to form an initial estimate of orbital position and velocity. The remaining two unknown quantities can be encapsulated in a range value and a range-rate value. Given the angles and their rates, constraints imposed on the semi-major axis and eccentricity (shown in red and black, respectively, in Fig. 1) allow for a bounding of the searchable range and range rate space, and CAR-MHF fills that space with a uniform distribution of range and range rate hypotheses (green in Fig. 1).

Fig. 1 conveys the significant increase in possible ranges and range rates that the cislunar environment enables. Filling this space with a uniform distribution of hypotheses leads to tens of thousands of hypotheses, quickly becoming computationally burdensome when the hypotheses are used in downstream calculations. An additional constraint is proposed in this work, in which only reasonable levels of range rate are considered in the first pass, effectively serving as a first level "sieve" in range rate to capture most periodic, quasi-periodic, and transfer orbits in the cislunar regime. This initial constraint was set at ± 2.5 km/s from the sensor's range rate to target, resulting in almost an order of magnitude fewer hypotheses to consider in some cases.

Also, it is of note that the eccentricity constraint line is outside the Earth 2-body zero energy line (blue) in the cislunar case (Fig. 1, right). The three-body dynamics introduce instances where some members of the periodic orbit families experience osculating eccentricities greater than one. The constrained admissible region was adapted to allow for constraint regions beyond the Earth 2-body zero energy curve. While these high osculating eccentricities can approach double digits, they do so for short periods of time, and selecting a constraint of between two and three for eccentricity does a robust job of capturing these orbits without undue expansion of the number of state hypotheses. Processing complexity is reduced, at the expense of a slight delay in IOD in isolated cases while the eccentricity remains outside the constraint region. The use of this trade-off does not allow 100% successful initiation of orbital estimates, but by restricting the expanse of the constraint region, the forthcoming observations--coupled with the dynamics model--will eventually allow for a successful solution.

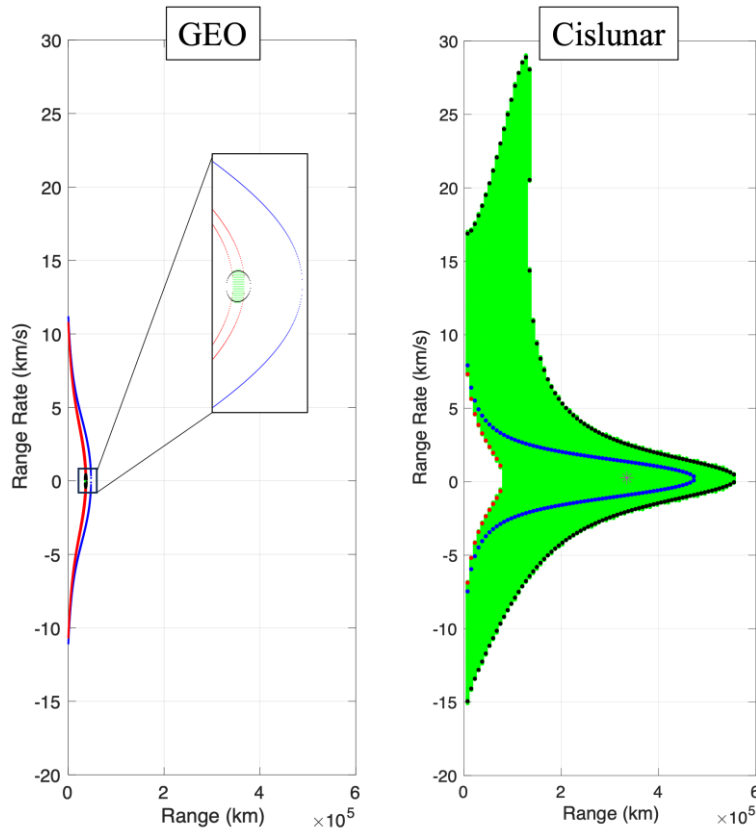


Fig. 1. Comparison of notional CAR range and range rate hypothesis solutions for a GEO orbit (left) and a cislunar orbit (right). Red lines are semi-major axis constraints, blue lines are the Earth 2-body zero energy contours which denote orbits captured by Earth, and black lines are eccentricity constraints. Green is the bounded region for hypothesis generation. Asterisk is location of actual object (only shown on right for cislunar orbit example).

3.2 Propagation of Hypotheses

Once the CAR is generated, the unscented transform is used to map the hypotheses from measurement space to state space, with uncertainties based on the measurement noise, the extracted rate noise, and range and range rate distributed uniformly, to form the associated set of state space hypotheses. These state space hypotheses (represented as Gaussian distributions) are then again converted via the unscented transform into sets of sigma points for processing via a sequential UKF. Fig. 2 shows an example result of this transition into state space for the same orbit example displayed in Fig. 1 (right). This orbit example is a northern Halo (HN) orbit about the $\mathcal{L}1$ Lagrange point with orbit index 100, full designation HN1-100 (see [1-2] for a detailed breakdown of the elemental periodic orbit families and their designations). The large spatial extent spanned by the hypotheses--given that so little is known with just one short-arc EO measurement--is evident.

A key takeaway is the existence of a set of state-space hypotheses which can be propagated through a high-accuracy full-ephemeris dynamics model to provide the time evolution of the PDF for this IOD solution. CAR-MHF currently utilizes Earth (with non-spherical perturbations), Moon, and Sun as gravity bodies, along with solar radiation pressure perturbations. Monitoring of the measurement space PDF as a function of time can then be achieved by mapping the state space hypotheses back into measurement space for a sensor of choice. The presence of a PDF from just a single short-arc EO measurement enables numerous sensing and revisit strategies that are not available when processing UCT measurements in the typical batch manner. We will return to this topic in Section 4.

Given the increased likelihood of direct engagement with the gravity bodies during the IOD process, it was also necessary to implement hypothesis removal logic in the event that excessive proximity to a gravity body is encountered. The implemented solution, while still in the early stages of evaluation, utilizes a user-defined distance from the surface of any gravity body, checks for sigma point collision during propagation, and removes the hypothesis to which that sigma point belongs. Earth-bound IOD problems, and GEO in particular, suffer much less from this possibility, but having this gravity body check has proven essential for cislunar IOD.

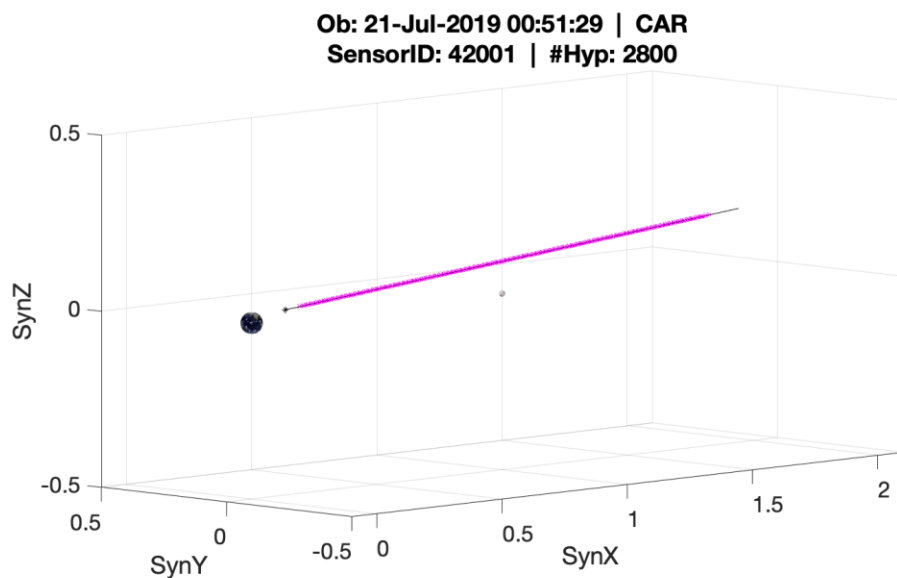


Fig. 2. State space representation of the CAR hypotheses for a GEO sensor collecting on an HN1-100 orbit. Coordinate system is Earth-Moon synodic frame, with Earth-Moon barycenter at the origin and the average Moon distance at approximately [1,0,0]. Earth and Moon size are not to scale (2X to improve visibility).

3.3 Data Association

Details of CAR-MHF and its use of a MH-JPDA filtering scheme are thoroughly discussed in [5]. The present analysis expands upon the power and insight of MH-JPDA to convey how well the sequential UKF is handling confusion and successfully graduating objects. Covariance-based data association, and in particular JPDA, provides the analyst with control of the method by which information from measurements are being marginally associated (soft decisions) amongst different objects (and hypotheses within an object) in the presence of association confusion. This starkly contrasts with hard-decision schemas such as nearest neighbor and “no decision” schemas like multiple hypothesis tracking (which perform all-on-all associations and rely on the correct association becoming increasingly evident with time and further comparisons).

Recall from Section 2.2 that a primary benefit of the CAR-MHF approach is that it allows instant access to a PDF for further investigation and updating. Hard association decisions in the presence of confusion can lead to significant error and estimate divergence, and “no decision” negates the benefits of having a PDF to interrogate. MH-JPDA and its soft association decisions provide a sound balance between 1) updating *a priori* distributions when information is

available and 2) delaying association decisions by not overcommitting to a single interpretation in conditions of high uncertainty.

Below is an introductory example of a multi-class confusion matrix which is helpful in parsing out the associations the filter is making and the knowledge the filter is acquiring as it updates in the presence of high levels of confusion. The X-axis presents the “truth” class, the actual objects that exist in the population. The Y-axis presents the filter assessment, the object estimates that CAR-MHF has generated based on the data available. The example in Fig. 3 is for a closely spaced constellation of objects in a different HN1 periodic orbit, this time HN1-200. Three simulated objects, with designators 90033, 90034, and 90035, are spaced within 100 km, and are observed via measurements that will be described in Section 3.

The color of the grid box designates the strength--or effective number of measurements--with which “truth” observations from objects along the X-axis were associated with the filter estimates in CAR-MHF along the Y-axis. The numbers in each grid box show the actual values. When processing larger populations (see Section 4), the actual values in each cell are removed, and interpretation is best done by a comparison of the relative shading. Perfect performance (no confusion) would be displayed as a case with all dark color along the diagonal, as no cross and/or mis-associations--which would present themselves, as shown in the example, with color in the off-diagonal—are present. In short, for optimal filter performance, it is desired that all mass be concentrated along the diagonal. In this sample problem, CAR-MHF successfully solves for a solution on all three objects, even though there is shared information during the original filter pass through the observations. Despite some mass off-diagonal, most of the mass is still along the diagonal, which informs the analyst that the filter is making sense of the situation correctly. The combination of MH-JPDA for patient, covariance-informed filter updates and an iterative filter/smoothen framework leads to a successful solution in highly challenging examples within the cislunar regime. The synergy between the filter and smoother is the next topic.

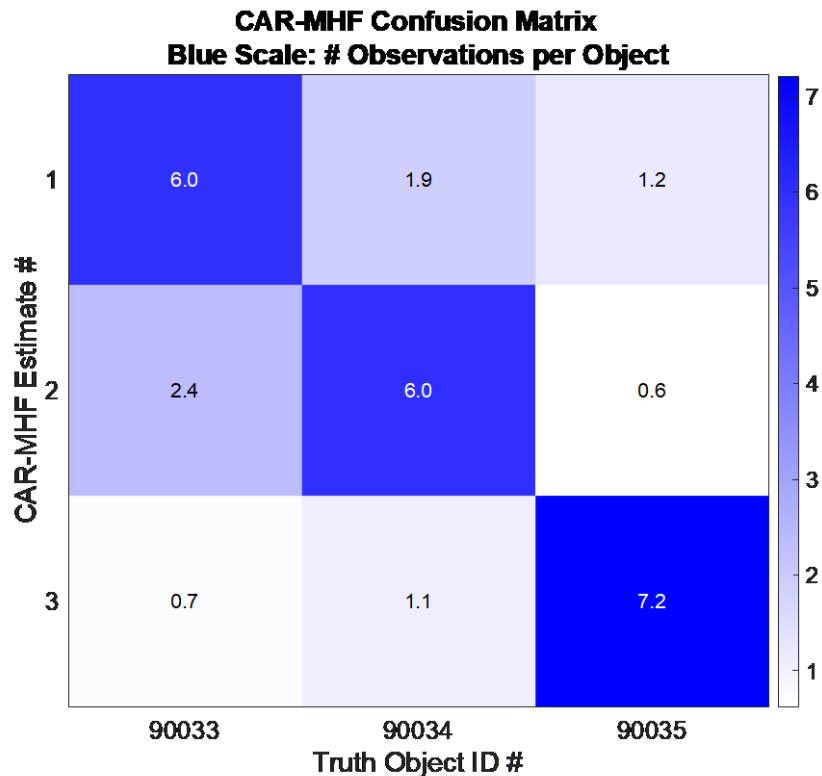


Fig. 3. Confusion matrix for example 3-object constellation in the cislunar regime. X-axis is the "truth" object class, and the Y-axis is the CAR-MHF "estimate" object class. Color bar shows the number of effective measurements associated for each CAR-MHF estimate (object). A total of nine measurements were collected on each object and processed simultaneously at each time step.

3.4 Iterative Filter/Smoothing

The concepts of convergence and consistency have been little changed from GEO UCT processing with CAR-MHF (discussed in [5]), except for the selection of the user-defined convergence threshold. Over the last decade of processing GEO and cislunar UCTs, the iterative filter/smoothing approach has proven itself to be exceedingly reliable, giving the analyst high confidence that graduated objects are indeed solid tracks. These graduated objects can then be passed on to catalog maintenance functions, either within CAR-MHF or other external workflows. In summary, graduation within CAR-MHF is fully described by three features:

- 1) The estimate must pass a convergence test, defined by its RSS position covariance dropping below a user-defined threshold.
- 2) The estimate must pass an iterative filter/smoothing consistency test, with the current test being the McReynolds' consistency [12].
- 3) The observations within an estimate must be 1-to-1, meaning no observations for the estimate under consideration for graduation can be shared with another estimate.

The 1-to-1 association check is confirmation of a lack of confusion. Only through iterative forward and backwards processing of the observations can a filter overcome any association confusion, which, when solving the IOD problem, is mostly driven by unavoidably high levels of initial uncertainty. These steps, executed sequentially within an iterative filter/smoothing framework, are the forcing function that takes an initial high-uncertainty track solution composed of many hypotheses and transforms it into a solid object track solution for follow-on maintenance.

Additionally, this iterative forward and backwards process is not limited to data within the time frame of the observations currently under consideration within the filter/smoothing. The filter/smoothing can be directed back to past still-unresolved UCTs to check their association levels with the current solution. If associations are found, these observations can be added to the estimate and the iterative filter/smoothing can again be accomplished using all currently associated observations. This iterative approach, able to reach back in time to check for observations that fell through the cracks, leads to high levels of data utilization, minimizing the number of observations in the collection of UCTs carried forward in time and thus decreasing computational load.

3.5 CAR-MHF Example

To close this section, we will show an example of CAR-MHF in which it utilizes a CAR to form an initial set of hypotheses, propagates those hypotheses to the time of the next measurement (time update), and performs a combination of filter update and hypothesis pruning (measurement update) to hone the filter solution down to a solid track on another member of the HN1 orbit family, this time HN1-284. The initial measurement is from a ground-based EO sensor (see Section 3.1 for more information on the simulated sensor network), and Fig. 4 renders the resulting 1,472 hypotheses generated by the CAR in both state space (Earth-Moon synodic frame) and measurement space (angular inertial reference frame from the perspective of the sensor).

The measurement space representation (Fig. 4, right) shows the measurement as a black diamond marker and attempts to convey two different fields of view (FOV) from the sensor. The first is the figure as a whole, which represents a 5-degree FOV. The second is the black rectangle drawn in the figure, which represents a much narrower 0.5-degree FOV. The state space representation (Fig. 4, left) shows the sensor line of sight as the line and the sigma points for all existing hypotheses as magenta asterisks. For this and following figures, a green line-of-sight line indicates a CAR or measurement update has been made, whereas a red line-of-sight line indicates propagation or a time update. Note the radial extent (along the line of sight) of the hypotheses at the moment of the CAR and following all time and measurement updates. Uncertainty in performing cislunar IOD, given the limited observability of objects in the radial direction afforded by EO sensors, is dominated by radial uncertainty.

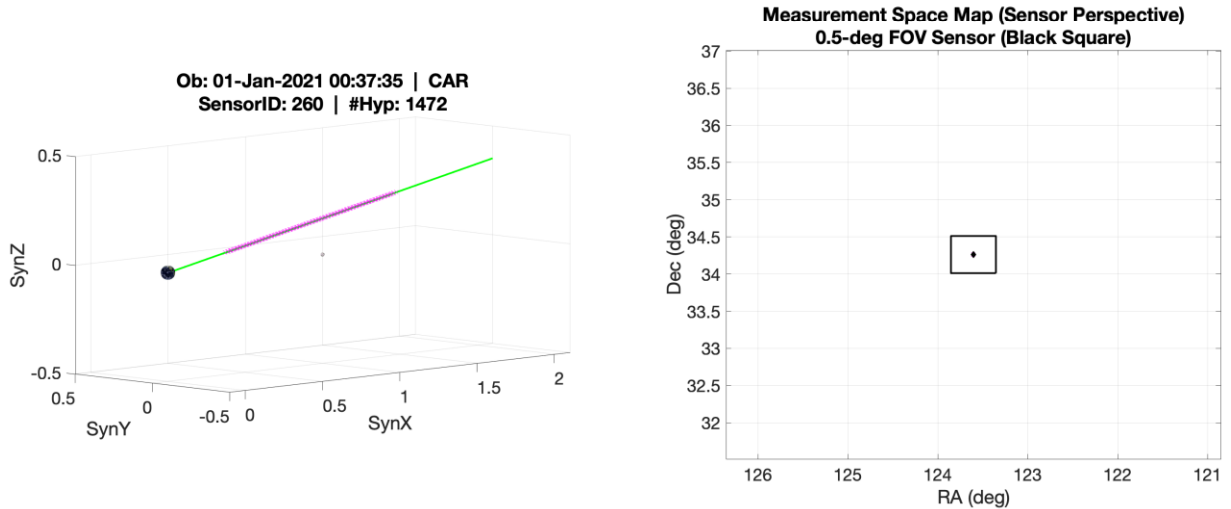


Fig. 4. Sample CAR-MHF solution at the point a CAR is generated. State space (left) and measurement space (right) sigma points for all existing hypotheses are shown. Measurement space sigma points and uncertainty contours aren't visible due to the zoom level of the figure. Actual measurement is displayed as black diamond marker.

Fig. 5 shows the CAR solution propagated for approximately six hours to the time of the first post-CAR measurement. Note the growth in the radial direction evident in the state-space representation. This next measurement is from a different ground-based EO sensor, which partially explains the extensive measurement space uncertainty. Had a sensor with a narrow 0.5-deg FOV been tasked on the highest-probability region of the measurement space PDF (as depicted by the black rectangle in Fig. 5, right), it would have missed the object. The tails of the measurement-space PDF exceed even the 5-deg FOV for this sensor. See Section 4.3 for additional treatment of the measurement space problem.

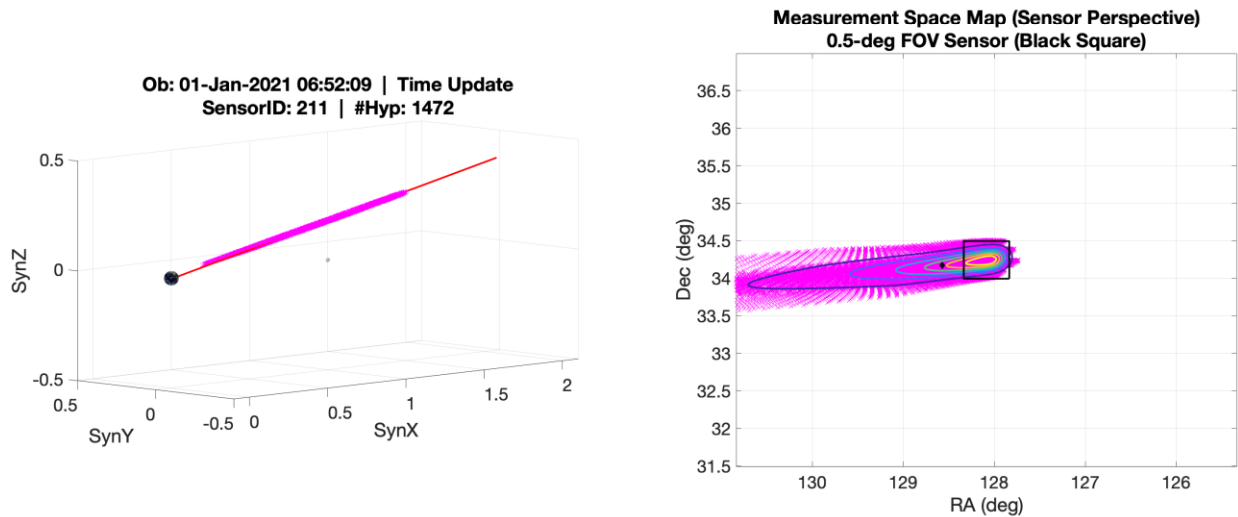


Fig. 5. Sample CAR-MHF solution just prior to the first post-CAR measurement (~6.25 hrs after initiation). State space (left) and measurement space (right) sigma points for all existing hypotheses, along with measurement space uncertainty contours, are displayed. Actual measurement is displayed as black diamond marker.

Fig. 6 shows the resulting state space and measurement space representations of the CAR-MHF solution after the first measurement update. The number of hypotheses has been pruned by the algorithm to 75, yielding much improvement to the radial uncertainty situation as shown in the state space representation. The measurement space representation effectively returns to the measurement noise Gaussian.




Fig. 6. Sample CAR-MHF solution after the first measurement update. State space (left) and measurement space (right) sigma points for all existing hypotheses are shown. Measurement space sigma points and uncertainty contours aren't visible due to the zoom level of the figure. Actual measurement is displayed as black diamond marker.

Propagation of the CAR-MHF IOD solution to the time of the second post-CAR measurement is depicted in Fig. 7. This measurement is from a GEO-based EO sensor, collected after another time gap of approximately 6 hours. Note the dispersal of the hypotheses evident in both the state space and measurement space representations. However, for this measurement, the extent of the measurement space uncertainty is not as excessive as at the time of the first measurement due to the improved radial uncertainty from that filter update.

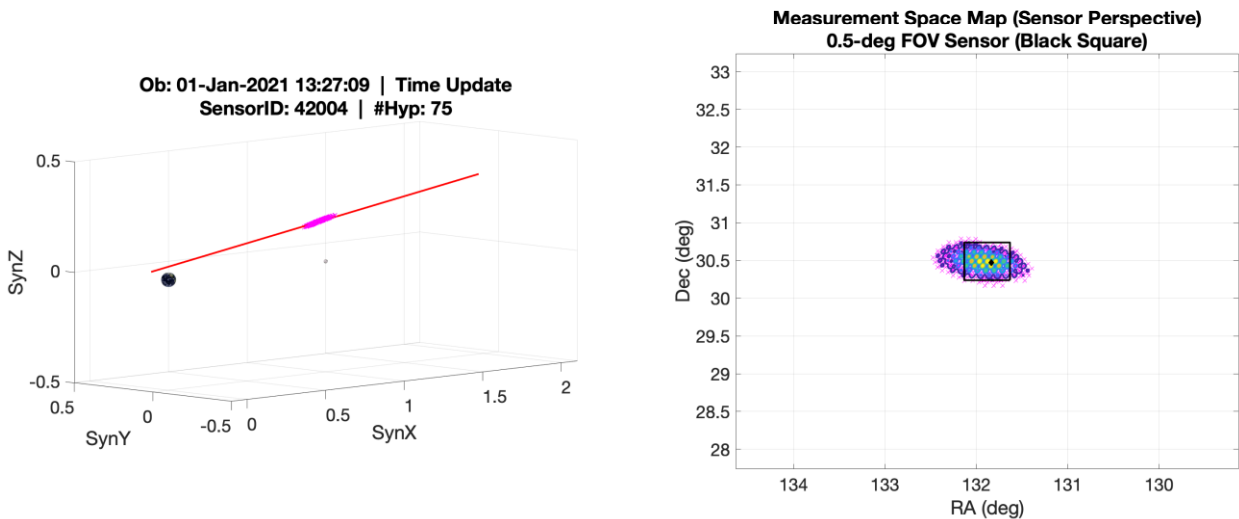


Fig. 7. Sample CAR-MHF solution at the time of the second measurement post-CAR, occurring ~13 hrs after the CAR. State space (left) and measurement space (right) sigma points for all existing hypotheses, along with measurement space uncertainty contours, are displayed. Actual measurement is displayed as black diamond marker.

Finally, Fig. 8 shows the CAR-MHF solution after the second measurement update. The number of hypotheses has been pruned down to three, progressing closer to a graduated solution. The measurement space uncertainty once again effectively returns to the measurement noise Gaussian. CAR-MHF, following the next observation ~6 hours later, takes the resulting 2-hypothesis solution, puts the hypotheses through the convergence and consistency tests (which they pass), and merges the remaining hypotheses into a single graduated object track.




Fig. 8. Sample CAR-MHF solution after the second measurement update. State space (left) and measurement space (right) sigma points for all existing hypotheses are shown. Measurement space sigma points and uncertainty contours aren't visible due to the zoom level of the figure. Actual measurement displayed as black diamond marker.

4. EXPERIMENT

Two simulated populations were generated using differing full-ephemeris dynamics propagation models. The first, Population #1, uses NASA's General Mission Analysis Tool (GMAT) [13] and concentrates on periodic orbit families near the Moon with the addition of some closely spaced clusters. Population #2 uses a KBR-developed dynamics model and focuses on a broader diversity of periodic orbits in cislunar space. Observational data sets were simulated for these populations with notional ground and space-based sensors for filter evaluation. This section will describe the data sets and the pertinent CAR-MHF assumptions and settings.

4.1 Observational Data Sets

Population #1 consists of 48 different orbits, all of which remain in the vicinity of the Moon and were ballistically propagated for about one month with a starting date of 2019-07-20. Table 1 details each orbit with the orbit number corresponding to the nomenclature of [1], in which each periodic orbit family is represented by a continuous sampling of 512 orbits. Fig. 9 displays the positions at the initial time (in right ascension and declination as viewed from center of Earth) of orbits 1-48¹.

Table 1. List of Orbits in Population #1

Object ID	Cislunar family	Orbit number(s)	Comment
1-10	HN1	100	Constellation of 10 halo objects in same orbit separated in phase
11-13	HN1	50/100/150	Constellation of 3 halo objects starting at same phase in same family (Note: orbits 1, 3 and 12 form a cluster of 3 halo objects separated by ~1000 km)
14-16	HS1	50/100/150	Constellation of 3 halo objects starting at same phase in same family
17-19	HN2	50/100/150	Constellation of 3 halo objects starting at same phase in same family
20-22	HS2	50/100/150	Constellation of 3 halo objects starting at same phase in same family
23-32	HS1	100	Constellation of 10 halo objects in same orbit separated in phase

¹ The population developed originally contained 59 objects, as shown in Fig. 9. For this IOD analysis, the maneuver and break-up objects (object IDs 49-59) were not considered.

33-35	HN1	200	Cluster of 3 halo objects separated by ~100 km
36	HN1	200	Object that starts in vicinity of objects 33-35 cluster
37-39	L1	50/100/150	Constellation of 3 Lyapunov objects starting at same phase in same family
40-42	L2	50/100/150	Constellation of 3 Lyapunov objects starting at same phase in same family
43	A1	50	Axial family object that remains in same lunar vicinity as previous objects
44	A2	50	Axial family object that remains in same lunar vicinity as previous objects
45-46	V1	50/100	Vertical family objects that remain in same lunar vicinity as previous objects
47-48	V2	50/100	Vertical family objects that remain in same lunar vicinity as previous objects

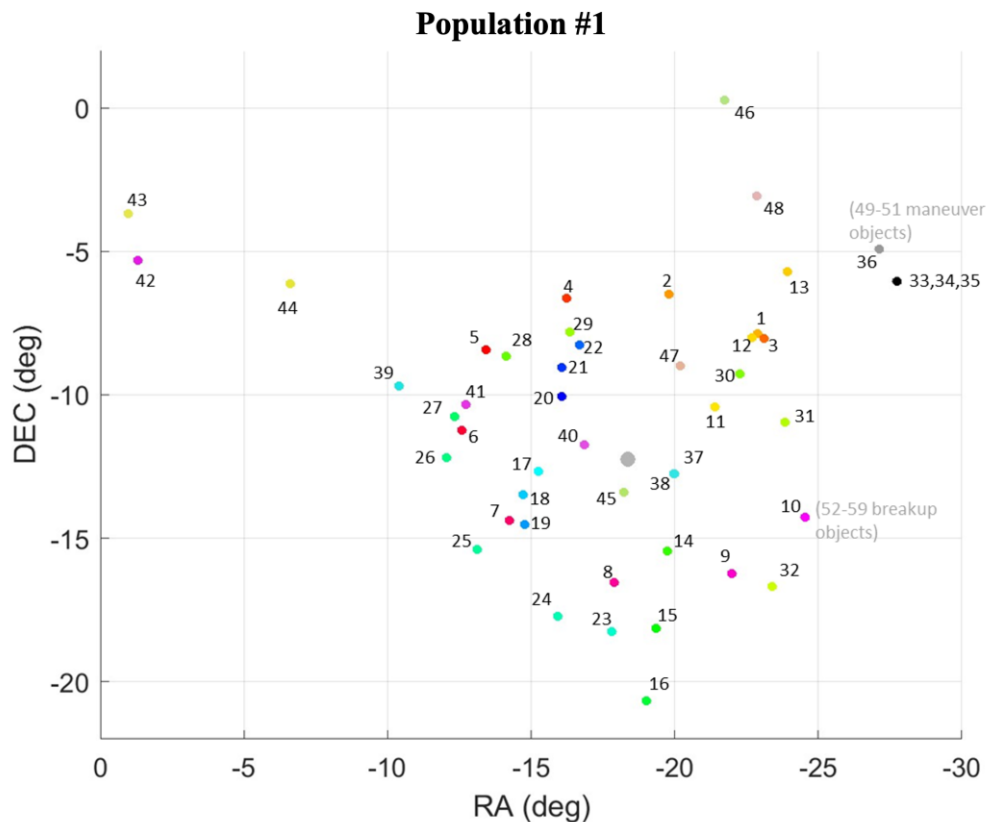


Fig. 9. Orbits 1-48 in Population #1 at initial simulation time in right ascension and declination, color shaded. Objects not under consideration are shaded light gray (see previous footnote). The moon is in gray at -18-deg right ascension, -12-deg declination.

Population #2 consists of 60 different orbits propagated ballistically with a starting date of 2021-01-01. These orbits are in a variety of cislunar periodic families to include those about the $\mathcal{L}4$ and $\mathcal{L}5$ Lagrange points, and as such, are observed all throughout the volume of cislunar space and not just near the Moon. Table 2 details each orbit with the orbit number (again corresponding to what was done in previous papers [1-2], with the addition of the Distant Retrograde Orbit [DRO] family, sampled in a similar fashion). Fig. 10 displays the initial positions of these objects in right ascension and declination.

Table 2. List of Orbits in Population #2

Object ID	Cislunar family	Orbit number(s)
101-106	DRO	002, 104, 206, 308, 410, 512
107-112	HN1	002, 096, 190, 284, 378, 472
113-118	HN2	002, 104, 206, 308, 410, 512
119-122	L1	002, 094, 185, 277
123-128	L2	002, 052, 102, 152, 202, 252
129-134	L4	002, 104, 206, 308, 410, 512
135-140	L5	002, 104, 206, 308, 410, 512
141-150	V4V5	002, 059, 115, 172, 228, 285, 341, 398, 454, 511
151-160	W4W5	006, 061, 116, 171, 226, 280, 335, 390, 445, 500

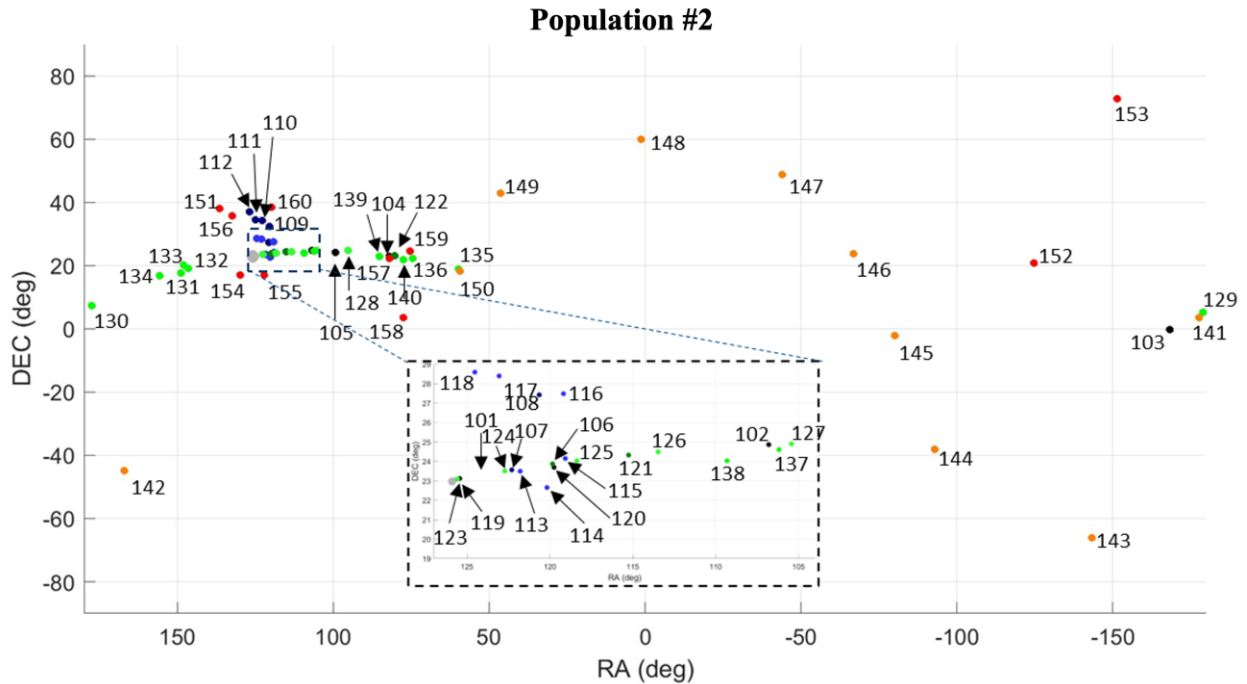


Fig. 10. Orbits 101-160 in Population #2 at initial time in right ascension and declination.

For each orbit set, observations were generated using two different sensor arrays. The first was an Earth ground-based electro-optical (EO-ground) set of five longitudinally distributed sensors to ensure the moon vicinity is continuously visible from at least one of the sites: New Mexico (sensor# 211; 33.8° N, 106.7° W), Maui (sensor# 231; 20.7° N, 156.3° W), Diego Garcia (sensor# 241; -7.4° N, 72.5° E), Spain (sensor# 260; 37.1° N, 5.6° W), and lastly Australia (sensor# 272; -21.9° S, 114.1° E). The second array (EO-space) consisted of three notional geostationary Earth orbit: longitudes of 120° E, 0°, and 120° W. The EO-space sensor IDs for the two populations are different, for Population #1 (sensor# 42001-42003) and for Population #2 (sensor# 42004-42006).

Tasking of observations was done on a sensor-by-sensor basis so that the object with the longest gap in time since its last observation from that sensor and currently observable by that sensor was selected. For EO-ground observations, the exclusion zones were defined such that an object was observable when the object was at least 10 degrees above

the horizon, the Sun was at least 10 degrees below the horizon, the object was at least one degree from the Moon, and the object was brighter than 24th visual magnitude (where the object was modeled as a diffuse sphere with a 1-m² cross-sectional area and reflectance of 20%). For EO-space observations, the exclusion zones were defined such that the object was required to be at least 30 degrees from the Sun, at least 5 degrees from the Earth, 2 degrees from the Moon, and again brighter than 24th magnitude. Once selected by a sensor, the object was observed five times with a time gap of 60 seconds between each measurement to form a track. Other objects within 0.1 degrees of the target object were also observed simultaneously. If any of the objects were within 0.2 arcsec of one another, the measurements were combined by averaging and a single measurement reported. In addition to the exclusions and tasking priority listed above, 60 seconds were allotted to move from one object to the next and no observations of another object were made during gaps between measurements for a particular track. All EO measurements were simulated with additive Gaussian noise with a 2 arcsec standard deviation.

In Fig. 3, it was noted that the object ID numbers are 90000 numbers, while the population objects as described thus far in this subsection have used truncated numbers for brevity (i.e., IDs 1-48 and 101-160). Moving forward, as in Fig. 3, 90000 numbers will be used. Object #1 from Population #1 is actually object #90001, object #101 from Population #2 is actually object #90101, so on and so forth. The use of 90000 numbers in the results section will help the reader distinguish between CAR-MHF estimate numbers and “truth” object ID numbers.

4.2 Experiment Settings and Assumptions

CAR-MHF is being evaluated from a cold start, meaning that an empty database of estimates is provided and CAR-MHF must initialize, update tracks, and graduate objects based on processing the incoming measurements and no *a priori* information other than the admissible region assumptions. The tests are evaluating the tool’s ability to perform the IOD function, and only the IOD function, from scratch; all other steady-state track maintenance functions are present but will not be the focus of the discussion. This approach aims to establish a potential IOD baseline for comparison with other techniques.

The admissible region constraints and other CAR parameters used for all runs presented in the results are displayed in Table 3, mirroring the examples shown in Section 2.1 for all objects:

Table 3. Admissible Region Constraints and Parameters

CAR Parameter	Value	Description
Energy	true	Use the Earth 2-body energy=0 curve as a possible constraint
Semi-Major Axis	45,000 km	Select the lower allowable semi-major axis to effectively set GEO+ as the lower bound
Eccentricity	2	Select the maximum eccentricity constraint to handle high osculating eccentricity scenarios
Range	30,000 km minimum; 800,000 km maximum	Select range bounds to limit the values of range from the sensor; capture much of cislunar space
Range Rate	2.5 km/s	Select +/- range rate bounds to limit possible range rate hypotheses
Hypotheses Range Spacing	7500 km	Separation of hypotheses in range
Hypotheses Range Rate Spacing	0.15 km/s	Separation of hypotheses in range rate
Smoother Interval	4 hours	The cadence over which CAR-MHF will attempt an iterative filter/smoothing refinement and graduation

The variable conditions for each test run are described in Table 4. Each population was processed twice, once with the entire set of measurements collected over the designated time span, and again with the measurements decimated. The decimation strategy kept only measurement tracks which were separated by some number of hours. For this analysis, six hours was chosen as the decimation level for comparison. To visualize the impact of the decimation strategy on the available data, Fig. 11 shows the measurement track separation and total number of measurement

tracks in histogram form for the Population #1 cases, with the Population #2 characteristics (not shown) being very similar. This experiment setup tests the functionality of CAR-MHF in dense and sparser data environments and illuminates any opportunities for adding robustness to CAR-MHF and its ability to perform the IOD function within the cislunar regime.

For Population #2, it was desired to maximize the number of objects available for the cold-start evaluation. The time that maximized the observable number of the original 60 objects was at the beginning of the time span. However, two objects (90119 and 90123) are unobserved over the entire time span due to exclusion constraints, and object 90106 is only sparsely observed on the first day of the time span, due again to exclusion constraints. The resulting test set evaluated below has these three objects removed, resulting in 57 objects for Population #2.

Table 4. Test Run Descriptions

Run #	Population	Decimation	Scenario Time Span
1	#1 (48 objects)	None	2019-07-20 – 2019-07-27 (1 week)
2	#1 (48 objects)	6 hours	Same as above
3	#2 (57 objects)	None	2021-01-01 – 2021-01-08 (1 week)
4	#2 (57 objects)	6 hours	Same as above

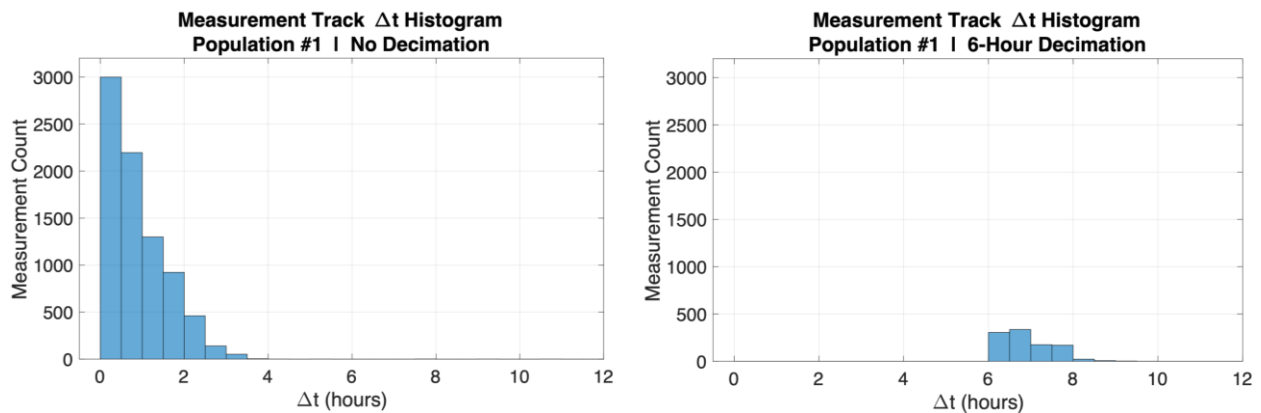


Fig. 11. Time gaps in measurement tracks for Population #1. The 6-hour data decimation case (right) is ~8X sparser in overall measurements than the case with no data decimation (left).

5. RESULTS

The following subsections present the results of processing the test sets discussed in Section 3. Section 4.1 discusses the results on Population #1, and Section 4.2 discusses the results for Population #2. Section 4.3 then provides a high-level overview of the measurement mapping results and the corresponding growth in measurement space uncertainty after CAR has generated an IOD solution.

5.1 Population #1 Results

The results for Population #1 are summarized in Table 5. The decimation strategy of keeping measurements separated by at least six hours led to an effective ~8× reduction in the number of measurement tracks (shown in Fig. 11 above). CAR-MHF achieved a 100% data utilization rate in both cases, with all observations correctly associated to distinct CAR-MHF estimates and 48 unique tracks being graduated, corresponding to the 48 truth objects. The run times were similar, with both runs taking 36 minutes on a MacBook Pro 6-core workstation.

Table 5. Population #1 Results

Metric	No Decimation	Decimation of 6 Hours
Total Measurement Tracks	8132	1072
Graduated Objects	48	48
Correct Associations (True Positives)	8132 / 8132 (100%)	1072 / 1072 (100%)
Run Time	35.6 min	35.9 min

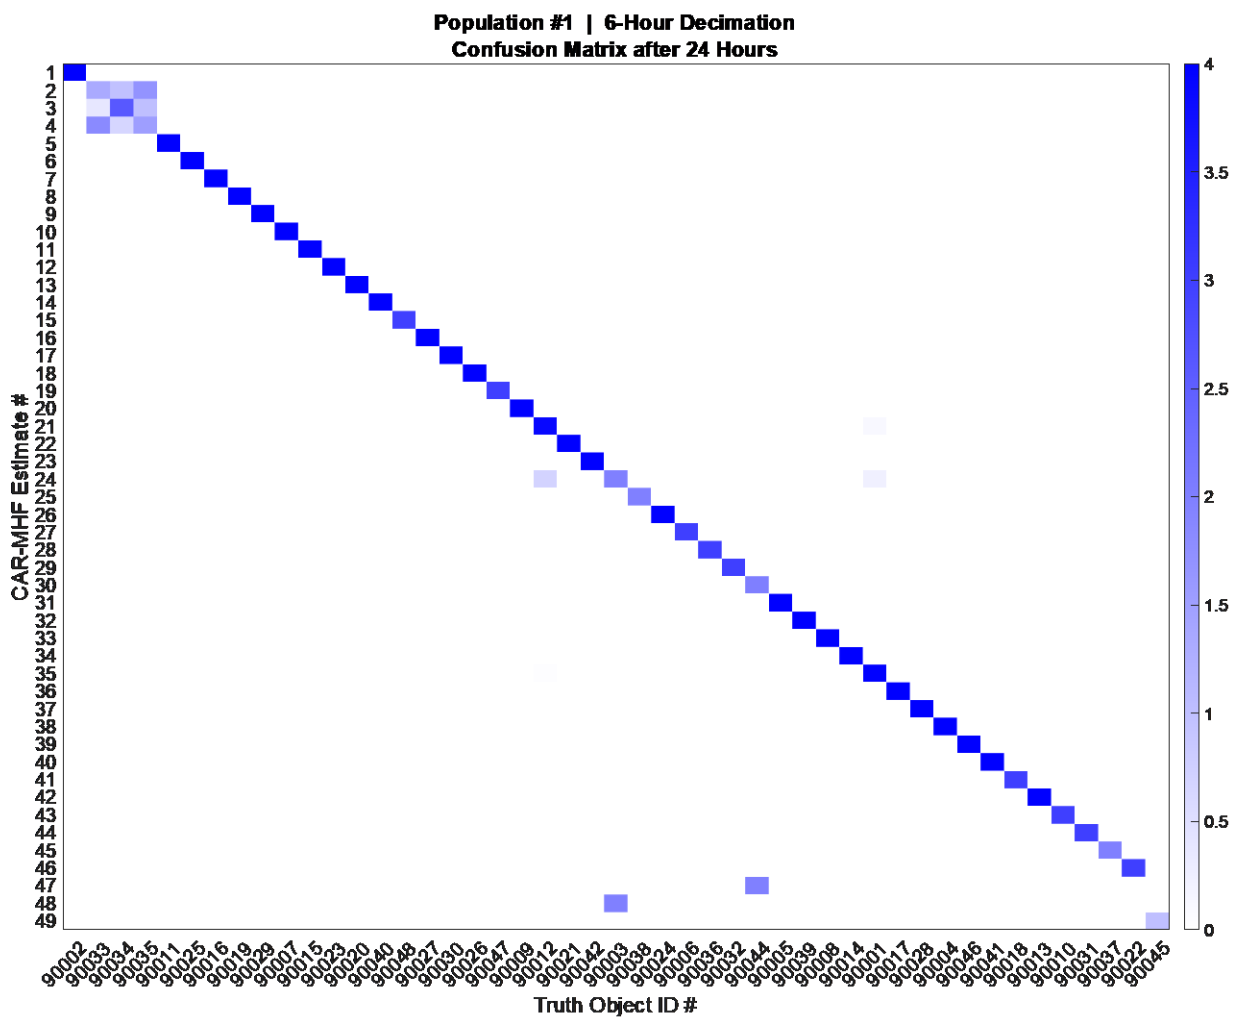


Fig. 12. Confusion matrix for the CAR-MHF run on Population #1 with a 6-hour decimation of the data. This confusion matrix was generated after processing data for 24 hours. The blue color scale designates the number of effective measurements associated across truth objects and CAR-MHF estimates.

The confusion matrix displayed in Fig. 12 highlights the levels of association confusion (after 24 hours) encountered by CAR-MHF while processing the more challenging 6-hour decimation case. Note that the objects in the three-by-three block in the upper left are the closely spaced 90033-90035 cluster (~100 km spacing) as discussed previously in

Section 2.3 and described in Table 1. CAR-MHF has created three estimates (estimates 2, 3, and 4 on the y-axis) for these objects, but there is significant sharing of measurement information across the estimates. There is so much sharing, in fact, that the majority of the mass is not along the diagonal for this cluster. However, with time and more measurements, CAR-MHF eventually resolves this confusion and successfully graduates these objects (graduation overview shown in Fig. 13). There is also confusion among the 90001, 90003, and 90012 objects, which Table 1 highlights as being separated by ~1000 km.

The last intriguing observation from the confusion matrix is that CAR-MHF generated two estimates for object 90044 (estimates 30 and 47). This can happen for multiple reasons, but it is typically caused by noise in the measurements leading to a CAR that doesn't quite capture the object's true orbit. Despite this, CAR-MHF was eventually able to graduate a single estimate for this object, utilizing the iterative filter/smoothing to combine the measurements from these estimates and recognize them as being one object.

The graduation rates given the different data decimations are compared in Fig. 13. The blue solid line shows the results from processing the entire data set, and the red dotted line shows results from processing the decimated data set. The impact of the sparser data is evident, as it takes 3.7 days (88 hours) to graduate the last few objects, whereas with the denser data it takes a little over a day (28 hours).

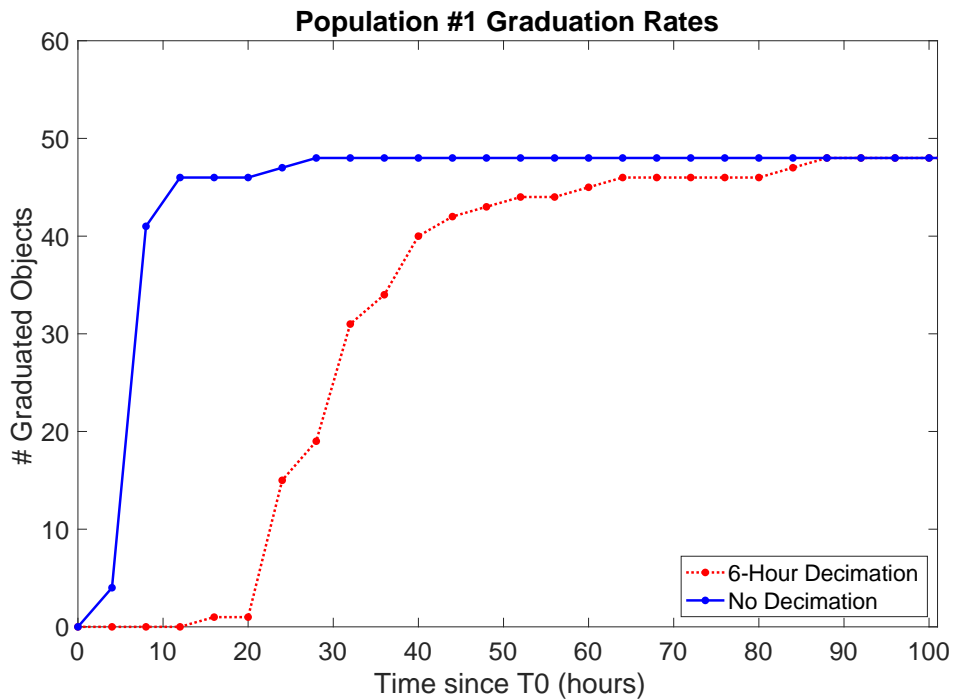


Fig. 13. Comparison of CAR-MHF object graduation rates for Population #1 with different data decimations. Graduations are attempted at the conclusion of each 4 hr smoother interval.

5.2 Population #2 Results

The results for the second population are displayed in Table 6. The decimation strategy of keeping measurements separated by at least six hours led to an effective ~7× reduction in the number of measurements. CAR-MHF achieved a 100% data utilization rate in both cases, with all observations correctly associated to distinct CAR-MHF estimates and 57 unique tracks being graduated. The run times were similar, with both runs taking ~40 minutes on the same MacBook Pro 6-core workstation.

Table 6. Population #2 Results

Metric	No Decimation	Decimation of 6 Hours
Total Measurement Tracks	9779	1339
Graduated Objects	57	57
Correct Associations (True Positives)	9779 / 9779 (100%)	1339 / 1339 (100%)
Run Time	42.3 min	37.9 min

The association confusion after 20 hours is starkly different for the Population #2. As shown in Fig. 14, there is no confusion across the 56 estimates that have been created up to this point in the data processing. This lack of confusion is evidently the result of 1) the lack of closely space objects in this population, and 2) an inference that having observations every six hours or so for this population was sufficiently dense to prevent uncertainty cross-over from occurring. Recall that the PDFs for most of the 56 objects with measurements up to this point (20 hours) have been propagated and updated with the filter multiple times, and that the first ~6-hour data gap after the CAR is where the largest uncertainty growth will occur (because there is the least information at this point in the process). If confusion is going to occur, it will most likely occur following the first post-CAR data gap. This lack of confusion also likely contributes to the improved graduation timelines covered in Fig. 15.




Fig. 14. Confusion matrix for the CAR-MHF run on Population #2 with a 6-hour decimation of the data. This confusion matrix was generated after processing data for 20 hours. The blue color scale designates the number of effective observations associated across truth objects and CAR-MHF estimates.

The graduation rates given the different data decimations are compared in Fig. 15. The blue solid line shows the results from processing the entire data set, and the red dotted line shows results from processing the decimated data set. Once again, the impact of the sparser data can be seen, as it takes about twice as long to graduate the last few objects (36 hrs vs. 68 hrs).

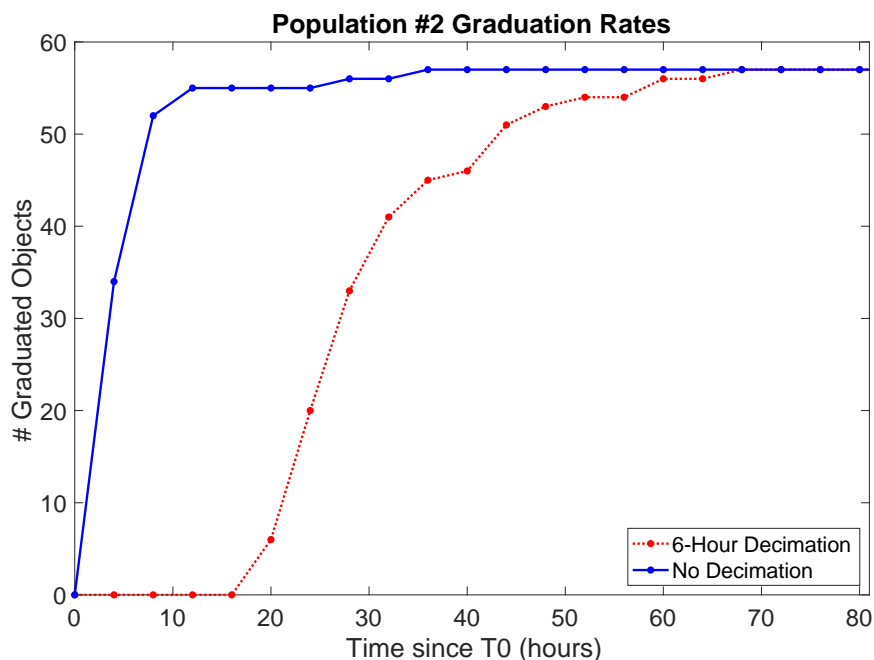


Fig. 15. Comparison of CAR-MHF object graduation rates for Population #2 with different data decimations. Graduations are attempted at the conclusion of each 4 hr smoother interval.

5.3 Time Evolution of Measurement Space Uncertainty

The instant formation of an estimate PDF is perhaps the most significant contribution of techniques such as CAR-MHF. This last part of the analysis looks at a sample propagation of that single-measurement PDF. In this scenario, an HN2-308 object (HN family orbit about the \mathcal{L}_2 Lagrange point) was observed solely by the EO-ground simulated sensor network described in Section 3.1. The following results show that, as time passes and the Earth rotates, a global network of sensors is needed to successfully monitor the near-Moon vicinity for periodic orbits and associated transfers. The example in Fig. 16 shows the time evolution of the measurement space PDF as it progresses through this EO-ground network over 17.5 hours. This example assumes no filter updates are occurring to demonstrate the impact of not obtaining follow-up observations on a notional UCT in cislunar space.

The CAR is generated by sensor ID 211, resulting in a tight Gaussian measurement space PDF. Propagation of the solution for two hours occurs before the object is next observed by sensor ID 231. Portions of the PDF are already exceeding the narrow 0.5-deg FOV boundary (the black rectangle). As time passes, more and more of the mass of the PDF begins to depart from the narrow FOV box, and after eight hours an observation by sensor ID 272 would likely miss the actual object (black diamond marker) if it was targeting its initial collect at the highest probability region of the PDF.

Propagating further, the measurement space PDF begins to exceed the larger 5-deg FOV represented by each figure window, and eventually the PDF contour lines (representing 15% increments in marginal probability of the overall PDF) begin to exceed the 5-deg FOV. In addition, the actual object becomes more difficult to find as it drifts further into the low-probability regions of the PDF. These propagations provide valuable insight into the effects that the complex dynamics are having on a PDF of the nature generated during the IOD process. Critical revisit timelines are measured in hours and are mostly driven by the profound radial uncertainties experienced with observations of objects in the cislunar regime.

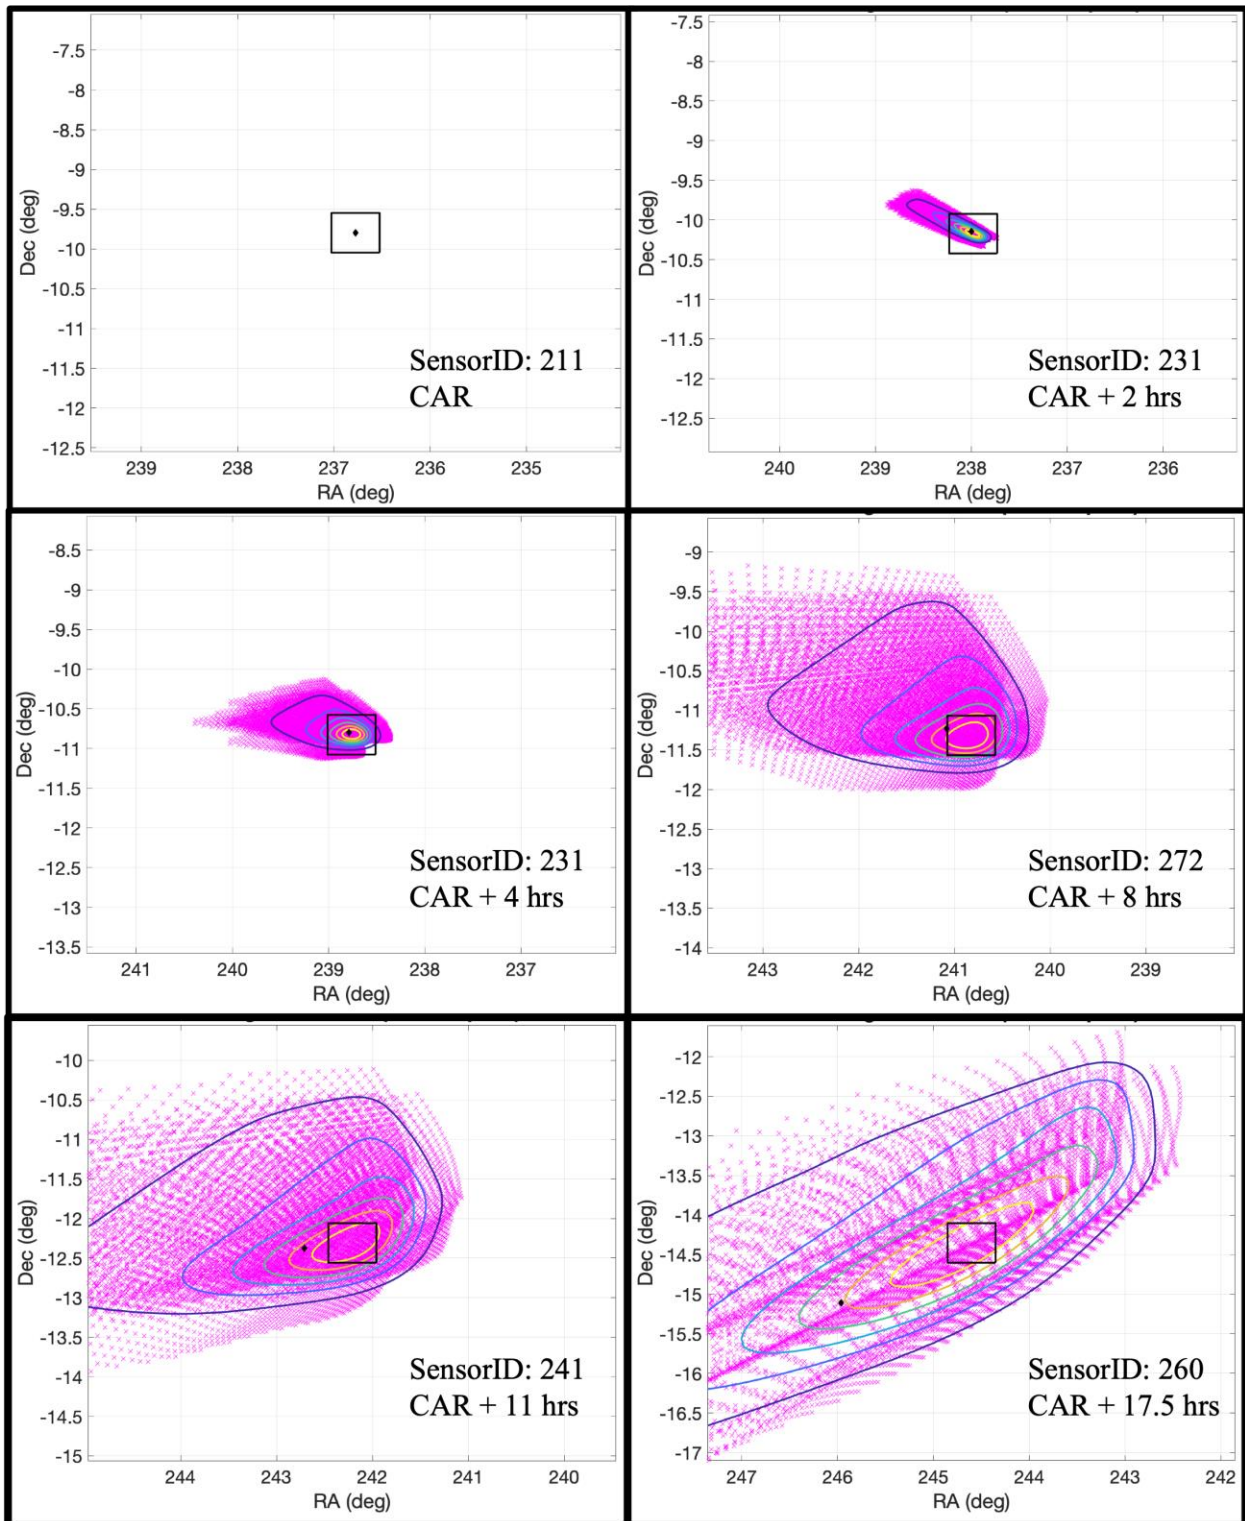


Fig. 16. Sequence of sensor perspectives on a CAR solution for HN2-308. The perspectives are across a simulated global network of EO-ground sensors spanning 17.5 hours. The magenta asterisks are the CAR hypothesis state sigma points mapped into the inertial measurement reference frame for each sensor. The contours represent 15% increments in marginal probability for the overall PDF (sum over all equally weighted hypotheses).

6. CONCLUSIONS AND DISCUSSION

CAR-MHF has been adapted and demonstrated as highly effective in generating IOD solutions for different representative populations in the cislunar regime. The methods employed by CAR-MHF provide a solid balance of performance in dense and sparse data environments, robustness to association confusion, and high data utilization. Continuing work will look to push further into sparse data sets to understand the limits of this approach, while also building an understanding of the sensitivities to measurement noise and other measurement artifacts present with real-world measurements of this orbital regime. Harnessing the available measurement information (angles and rates, visual magnitudes, etc.) to make inferences about the range and range rate information and improve upon uniformly distributed representations in range and range rate is a possibility. However, as a baseline, this work suggests that these potential improvements would likely be accretive but not essential.

There are numerous other features built into CAR-MHF that were not the focus of this assessment, including follow-on orbit maintenance, maneuver handling, processing of lagged measurements, etc. While these capabilities have been previously demonstrated to have utility, the challenges of achieving foundational SSA in the cislunar regime will benefit from higher-fidelity features providing rich, individualized assessments executing within a modular but integrated SSA workflow. CAR-MHF is a tremendous IOD tool, and its focus should remain where its generalized formulations are most performant.

Perhaps the most impactful feature of CAR-MHF and similar methods is the potential for tactical utilization, due to the implicit PDF that is formed. A sensor tasking schema based on information gain to minimize the possibility of failed UCT reacquisition could be immensely valuable within an environment which is, by its nature, difficult to observe. This possibility would improve cislunar space situational awareness by further informing the optimal balance of task-based and search-based sensing requirements across multiple observation locations. To reacquire a UCT with high confidence using narrow-FOV EO sensors, a critical timeline of a few hours for follow-up observations was demonstrated across numerous orbit families in the cislunar regime. In cases where reacquisition proves difficult, the PDF could be used to inform a search strategy for more efficient reacquisition. This use case has limits warranting further investigation due to the rapid growth and distortion of the PDF given the dynamics of the environment and large uncertainties associated with IOD.

7. REFERENCES

- [1] C. Chow, C. Wetterer, K. Hill, C. Gilbert, D. Buehler, and J. Frith, "Cislunar Periodic Orbit Families and Expected Observational Features," in *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, September 15-18, 2020.
- [2] C. Chow, C. Wetterer, J. Baldwin, M. Dilley, K. Hill, P. Billings, and J. Frith, "Cislunar Orbit Determination Behavior: Processing Observations of Periodic Orbits with Gaussian Mixture Model Estimation Filters," in *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, September 14-17, 2021.
- [3] C. Chow, J. Baldwin, C. Wetterer, M. Dilley, K. Hill, P. Billings, C. Craft and J. Frith, "Cislunar Orbit Determination: Improvements in Uncertainty Realism and Data Fusion," in *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, September 27-30, 2022.
- [4] K. Hill. "Maneuver Detection and Estimation with Optical Tracklets," in *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2014.
- [5] J. Stauch, M. Jah, J. Baldwin, T. Kelecy, and K. Hill, "Mutual Application of Joint Probabilistic Data Association, Filtering, and Smoothing Techniques for Robust Multiple Space Object Tracking," in *Proceedings of the 2014 AIAA/AAS Astrodynamics Specialist Conference*, San Diego, CA, August 2014.
- [6] I. Hussein, C. Roscoe, M. Mercurio, M. Wilkins and P. Schumacher, Jr., "Probabilistic Admissible Region for Multi-Hypothesis Filter Initialization," in *Journal of Guidance, Control and Dynamics*, vol. 41, no. 3, 2017.

- [7] J. Aristoff, J. Horwood, N. Singh, A. Poore, C. Sheaff, and M. Jah. "Multiple hypothesis tracking (MHT) for space surveillance: theoretical framework," in *Proceedings of the 2013 AAS/AIAA Astrodynamics Specialist Conference*, Hilton Head, SC, August 2013.
- [8] M. Bolden, I. Hussein, H. Borowski, R. See, and E. Griggs, "Probabilistic Initial Orbit Determination and Object Tracking in Cislunar Space Using Optical Sensors," in *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, September 27-30, 2022.
- [9] J. Gaebler, K. Hill, C. Wetterer and J. Bruer, "Cislunar Initial Orbit Determination with Optical Tracklets," in *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, September 19-22, 2023
- [10] S. Wishnek, M. J. Holzinger and P. Handley, "Robust Cislunar Initial Orbit Determination," in *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, September 14-17, 2021.
- [11] K. DeMars and M. Jah, "Probabilistic Initial Orbit Determination using Gaussian Mixture Models," in *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 5, pp. 1324-1335, 2013.
- [12] J. Wright, "Optimal orbit determination," in *AAS/AIAA Spaceflight Mechanics Meeting*, AAS 02-193, San Antonio, Texas, 27-30 Jan. 2002.
- [13] GMAT Development Team, 2020. *User's Guide*.