Leveraging Fisher Information to Optimize Observation Scheduling for Orbit Determination

Samuel Wishnek
Ball Aerospace
Joshua Wysack
Ball Aerospace

ABSTRACT

We have developed an approach for planning future observations of space objects from a set of potential observers and a range of potential observing times. This approach maximizes the information gain of the next observation so as to minimize the total state uncertainty of the target object. This approach is applied with initial orbit determination algorithms as a method for minimizing error in their resulting estimates. Additional improvements to a dynamics-neutral initial orbit determination algorithm are presented as well. These include adaptations to accept more than the minimal number of observations which is compared against a similar adaptation to the Gooding algorithm. The other is a new approach for generating initial state guesses for the dynamics-neutral algorithm by sampling the admissible space based on the actual density of known space objects.

1. INTRODUCTION

As the number of NGOs and countries interested in launching satellites beyond the near-Earth environment grows, so does the importance of the tools used to detect and track these satellites. Initial orbit determination (IOD) is the first step in the process of establishing an understanding of the state of a space object. It is important for building a first state estimate to inform every subsequent step in orbit determination. Initial orbit determination is an essential step for estimating the state of an unknown space object. It is a preliminary requirement for both batch and sequential filters that can provide precise estimates of the state and a valuable product on its own for building an actionable understanding of the object's orbit. However, the complex dynamics beyond the near-Earth environment make IOD a difficult task. IOD traditionally makes heavy use of the simplifications afforded by pure two-body dynamics, and the complexity of the problem opens-up beyond this environment. We have previously developed a method for IOD beyond the near-Earth, however this method suffers from high sensitivity to measurement error which limits its real-life utility [7]. By taking a new approach to the problem, we have developed a solution to this problem by allowing additional measurements to inform the estimate. Furthermore, we take an information theoretic approach that extends this solution to improving the understanding of space objects with less than the required number of observations to find a full state through IOD.

Several new approaches are taken to improve these results. First, IOD traditionally operates on the minimum number of observations required to estimate a state. For example, three angles-only observations have six constraints to fully determine the six-element dynamical state. These methods include algorithms such as Gooding and Double-R [2]. In practice, this can mean neglecting other available data if more than the minimum number of observations were taken. Observations are often taken in short bursts, and traditionally only a single measurement is used per burst. By folding these additional measurements into the solution, we can improve the resulting estimate's tolerance for measurement error. This places the developed algorithm between traditional IOD and batch least squares. The former accepts only the minimal amount of data while the latter can accept an arbitrary amount but expects a good initial estimate of the state. Henderson, Mortari and Davis have developed modifications to the Gooding IOD algorithm to accept additional observations. This approach is compared against the newly developed modifications to the dynamics-neutral IOD algorithm [4]. In addition, by understanding the observer's capabilities we can also reverse the problem and determine the optimal time to take an additional observation that completes the state space and is most likely to yield a reliable IOD result. For example, if more observations in the near-term would have little impact on the quality of the estimate, it would be more efficient to task the sensor elsewhere and return when a measurement would have more impact. This kind of scoring for the expected impact of an observation is essential for understanding the collection schedule that best uses the available sensors. This approach works by assuming the target is in a circular orbit and propagating forward its information matrix based on that assumption. The feasibility of using the circular assumption to perform IOD and maintain custody of a target was demonstrated by Stoker et al [6]. This work applies these results to plan out additional observations.

Second, artificial space-objects are more likely to inhabit certain favorable orbits. These include special orbit geometries such as geosynchronous, sun-synchronous, Molniya, and the cislunar three-body periodic orbits. By informing the orbit determination algorithm of these orbits we can more quickly and reliably solve for these cases that can otherwise be hard to estimate due to the instability of the orbits. The existing algorithm searches the full-feasible space and begins its optimization by spreading sample points across the full admissible region. By incorporating an understanding of orbits preferred for artificial space objects, those regions can be sampled more densely and potentially converge to a solution quicker than the naïve approach.

Even without sufficient quality or number of measurements to perform IOD, it is still possible to glean useful information about the space object such as orbit class. The information theoretic approach can be leveraged in combination with admissible regions to determine feasible spaces in state space that a space object could exist within. This is applicable to scenarios without enough observations to fully determine the state. Similar to how a human observer can glance at a satellite streaking across the sky and immediately understand whether the object is likely to be in low Earth orbit or something much higher; this kind of understanding can be useful for systems without a human in the loop. While a single solution may not be possible for these cases, by developing an understanding of the space in which the object could exist we can build actionable intelligence that could be useful for scheduling follow-up observations or classifying the target.

2. METHOD

2.1 Improvements to Dynamics-Neutral IOD

The dynamics-neutral IOD algorithm applies admissible regions in order to constrain the search and more efficiently converge onto a feasible solution. The dynamics-neutral IOD approach from this prior work was fairly sensitive to measurement error [7]. This had a particularly large impact on the measurement error requirements for cislunar orbits where the large distances between observer and target and instability of many orbits could produce particularly high errors in the output state estimates. New additions to this algorithm can reduce the resulting state error by applying basic assumptions or incorporating additional data.

2.1.1 Artificial Bias

The dynamics neutral approach can be extended from a simple limit to remove non-physical and unrealistic parts of the solution space to a more complex series of constraints that can incorporate known information about the target not immediately evident in the raw measurement data or reasonable limits based on the known population of space objects and the limitations of the observing sensor.

The population of artificial space objects is well understood and the states are available in bulk from online sources [1]. Using this census of Earth-orbiting artificial satellites, one can determine which state components are most common and which are uncommon. Since the dynamics-neutral IOD algorithm samples the admissible region, this provides an expected density function for where the solution may lie. Figure 1 shows the relationship between eccentricity and semi-major axis for the public satellite catalog on August 7th, 2023. Clear regions of both lower and higher density of space objects are apparent in the data. From this, the uniform random sampling of the admissible space in the IOD algorithm can be replaced with a density sampling that reflects the true population.

2.1.2 Incorporating Additional Observations

IOD algorithms typically accept the minimal number of observations required to form a complete state estimate. In the angles-only case this would be three distinct observations. However, in practice it isn't unusual to have more

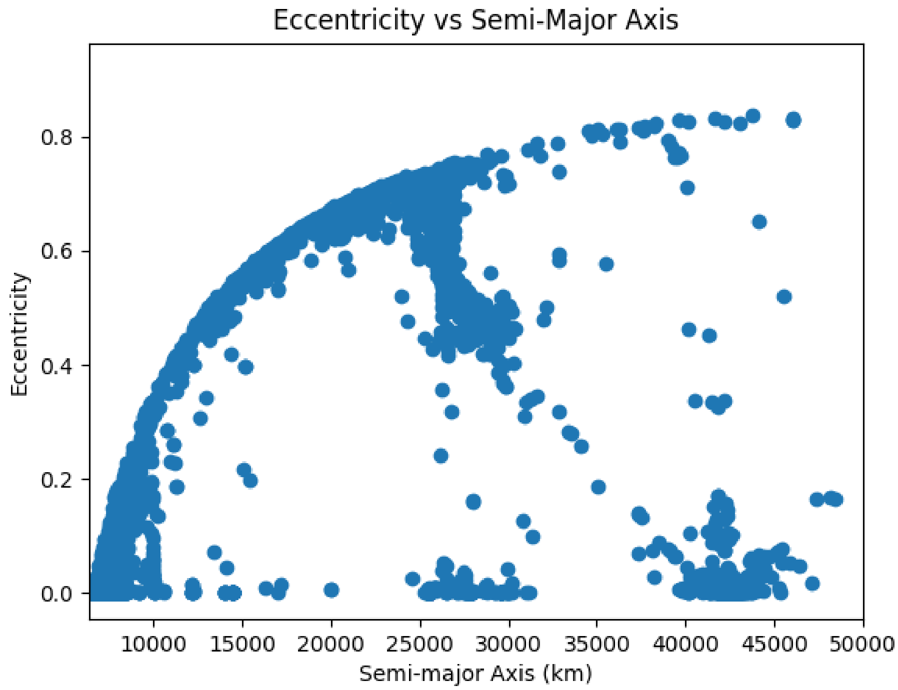

Fig. 1: Eccentricity vs semi-major axis for the full public satellite catalog between semi-major axes of 6378 km and 50,000 km.

than the minimal number of observations to choose from. A sensor may collect more than the bare minimum number of observations and by the time IOD is to be performed there may be more than three observations to choose from. While simply choosing three observations from the data can yield fine results, if additional information is available the quality of the results could potentially be improved by including it.

Henderson, Mortari, and Davis adapted the Gooding IOD algorithm to accept an arbitrary number of observations [4]. The traditional Gooding algorithm iterates over trial values of the first and final ranges from the observer to the target, solves Lambert's problem for these two states and then finds the error between the predicted and actual middle measurement. This is all wrapped in a Newton iteration with a finite-difference gradient to iteratively correct the range values to minimize the middle state error.

Henderson, Mortari, and Davis proposed altering this to accept additional observations by calculating the sum of the errors between the predicted and actual observations for all observations between the first and last. The Newton iteration was replaced with a gradient-free optimization algorithm and the individual errors were taken as just the angle error between the actual and expected observation vectors [4].

For the developed IOD algorithm, the change to accommodate additional observations is simple and similar to the changes to Gooding. The original cost function was the sum of the Mahalanobis distances between each of the three observation predictions and actual observations. Extending this to accept more observations only requires adapting the algorithm to treat the list of measurements as one of arbitrary length. Since the Mahalanobis distance includes the measurement covariance matrix for each observation, the resulting cost function appropriately weights each error value by the quality of the sensor it was recorded with.

2.2 Planning Future Observations

An established observation architecture is difficult to change. For space-based systems this typically requires launching new observing spacecraft at great expense. Accordingly, getting the most utility from an existing system can be highly beneficial. Given a static system of observers, how should one go about optimizing the choice of observers and when to observe in order to gain the most information about the target object's state?

2.2.1 Circular Orbit Assumption

When collecting angles only observations, three observations are required to form a complete state estimate. There may be cases where an observer collects on a space object but is either unable to collect three measurements, or the observations may be too close together in time to provide the geometric diversity needed to form a reliable IOD estimate. In these cases an additional observation is required before passing the data through an IOD algorithm. One option for planning out the next observation is to use the observed kinematic trajectory and use the same observer to follow up on the target. However, this may not be ideal for several reasons. Another observation from the same observer, particularly after only a brief delay, will provide little geometric diversity and this may result in a poor estimate. Another potential issue is that the original observer may not be available for this additional tasking. Tasking a different observer may be possible, but without a full state solution it isn't clear where this other observer should look. One potential solution is to assume the target object is in a circular orbit and task the second observer based on that assumption. There will be cases where this fails when the eccentricity is large, however a census of the public satellite catalog shows that 72 percent of cataloged space objects have an eccentricity under 0.01 and that rises to 88.5 percent under 0.05. Depending on the field of view of the sensor these satellites could be assumed circular for observation tasking and a more accurate solution can be found based on the results of that final observation.

Applying the circular assumption to two angles-only observations provides enough information to determine a complete six-element translational state. The computation of these state elements is performed using Newton's method. For a circular orbit, the angle between the position state of the target satellite at the first observation and second observation can be computed both directly from the two position states and indirectly based on the orbital radius and the time between the observations. The former can be found using equation 1. The latter uses Kepler's third law to relate the radius and period and then uses the ratio of the actual time between observations over the period to find the angle. This is is expressed in equation 2. The orbital radius is the unknown value to be found. The correct value of the orbital radius will make the difference of these two angles zero. This difference, expressed in equation 3, is the function that Newton's method is applied to. The two position states can be expressed in terms of the known position state of the observer, observation vector toward the target, and the range to the target. This simple equation is shown in equation 4. The ranges at each time can then be determined from computing the intersection point between the ray pointing from the observer toward the target and the sphere with radius equal to the orbital radius. This is computed using equation 5. There is an ambiguity with the plus or minus in this equation. However, in scenarios with the observer looking out the minus solution will produce a negative range value and can be discarded. Knowledge of the observer's capabilities such as visual magnitude capabilities can remove the ambiguity in other scenarios.

At this point, the only unknown is the orbital radius. Newton's method also requires a derivative of the target function with respect to the unknown. This can be taken analytically and is expressed in equation 6.

As described, this approach is sufficient to find the solution. However, in practice the implementation will run into issues as the radius value must be positive for a valid solution and the algorithm may push it negative for an iteration. This problem can be avoided by defining $x^2 = r$ and finding x instead of r. Other than substitution, the only change to the equations is that the derivative must be multiplied by $2x$ as per the chain rule. In this case, negative and positive values of x will still appropriately converge to the same solution for r

$$\theta = \arccos\left(\frac{\vec{r}_1 \cdot \vec{r}_2}{r^2}\right) \tag{1}$$

$$\theta = \Delta t \sqrt{\frac{\mu}{r^3}} \tag{2}$$

$$f = \Delta t \sqrt{\frac{\mu}{r^3}} - \arccos\left(\frac{\vec{r}_1 \cdot \vec{r}_2}{r^2}\right) \tag{3}$$

$$\vec{r}_n = \vec{R}_n + \hat{p}_n \rho_n \tag{4}$$

$$\rho_n = -\hat{p}_n \cdot \vec{R}_n \pm \sqrt{(\hat{p}_n \cdot \vec{R}_n)^2 - \left(\vec{R}_n \cdot \vec{R}_n - r^2\right)} \tag{5}$$

$$\frac{df}{dr} = \frac{\frac{d\vec{r}_1}{dr}\cdot\vec{r}_2 + \vec{r}_1\cdot\frac{d\vec{r}_2}{dr} - \frac{2\vec{r}_1\cdot\vec{r}_2}{r}}{r^2\sqrt{1-(\vec{r}_1\cdot\vec{r}_2)\,r^{-4}}} \tag{6}$$

$$\frac{d\vec{r}_n}{dr} = \hat{p}_n \frac{r}{\sqrt{\left(\hat{p}_n\cdot\vec{R}_n\right)^2 - \left(\vec{R}_n\cdot\vec{R}_n - r^2\right)}} \tag{7}$$

where:

θ = angle between observed target states
$\vec{r}_n$ = target position at observation time n
r = radius of the orbit
Δt = time between first and second observations
μ = Earth gravitational parameter
$\vec{R}_n$ = position of the observer of observation n
$\hat{p}_n$ = unit vector angles measurement from the observer to the target for observation n
ρ_n = range between observer and target for observation n
f = the function being searched for zeros
$\frac{df}{dr}$ = the derivative of the function to be zeroed

2.3 Predicting Information Gain

Here the cases of planning for an additional observation on an already sufficient set of observations and planning for the final observation on non-sufficient set of observations converge. The former case can simply apply one of the many IOD algorithms to generate a rough initial estimate that can be propagated out into the future while the latter can apply the circular assumption and propagate that assumed state.

With this low-accuracy initial guess, the information filter algorithm can be used to estimate the evolution of the information regarding the state. The primary advantage of the information filter over other non-linear filters like the extended or unscented Kalman filters is that it processes the Fisher information matrix rather than the covariance matrix. This is simply the inverse of the covariance matrix, but it has one key advantage for this use case. Where the covariance matrix has infinite eigenvalues for the eigenvectors corresponding to the unknown axes, the Fisher information matrix has zeros. This is much more computation friendly and the unknown states simply correspond to the null space of the information matrix.

Accordingly, the information filter can be run with the low-accuracy initial guess and an initial information matrix of all zeros. This is initialized at the time of the first observation and the measurement update is applied to this low-accuracy state and the zero information matrix. The information filter predict step is then run to the second observation time and the measurement update is run again, but this time with the second observation. At this point there is the low-accuracy state at the time of the second observation and an information matrix with a two-dimensional null space corresponding to the unknown axes of the full six-dimensional state space.

The predict and measurement steps are shown in equations 9 and 10 respectively. The information matrix is Λ, the measurement covariance matrix is R, the measurement

$$C = \Phi^T_{k,k-1}\Lambda_{k-1,k-1}\Phi^{-1}_{k,k-1}\left(\Phi^T_{k,k-1}\Lambda_{k-1,k-1}\Phi^{-1}_{k,k-1} + Q^{-1}\right)^{-1} \tag{8}$$

$$\Lambda_{k,k-1} = (I - C) + CQ^{-1}C^T \tag{9}$$

$$\Lambda_{k,k} = \Lambda_{k,k-1} + H^T R^{-1} H \tag{10}$$

where:

$\Lambda_{k-1,k-1}$ = information matrix before the predict step
$\Lambda_{k,k-1}$ = information matrix before after the predict step but before the update step
$\Lambda_{k,k}$ = information matrix after the update step
$\Phi_{k,k-1}$ = state transition matrix from the time of observation $k-1$ to the time of observation k
Q = process noise covariance matrix
I = identity matrix
H = measurement matrix
R = measurement covariance matrix

From here, the goal is to find the third observation that will not only remove this null space, but maximize the final information. Now, since the information is represented as a matrix, the maximum is ill-defined. Some scalar value needs to be found that represents the total information and preferably one that severely penalizes eigenvalues of zero. For this research, the determinant of the matrix was chosen. The information matrix is positive semi-definite so the determinant should always be zero or positive, and it will only be zero in the case that the matrix retains a null space. The determinant is the product of all of the eigenvalues, so it acts as an approximate measure of the squared volume of the associated information ellipsoid. Just as a very small covariance ellipsoid can show high certainty in a state estimate, a very large information ellipsoid shows high certainty. These ellipsoids have principal axes aligned with the eigenvectors of their associated matrices and principal axis lengths equal to the square root of the corresponding eigenvalues.

After the second update to the information filter, the singular information matrix can be propagated forward to potential observation times and the measurement update applied to find a predicted final state information matrix. It is then a simple optimization problem to maximize this final information matrix determinant. This can be solved with either a brute force approach in the case of small number of discrete times and observers or with an optimization algorithm.

While the example explained above used the premise of the IOD case, this same method can be implemented with few changes for sequential and simultaneous filters. One potential use for this method is planning an observation schedule while minimizing the number or frequency of follow up times and maintaining custody to some required level. In this case it may make sense to target a different metric rather than the matrix determinant. For example, if the requirement is to maintain one sigma position uncertainty to within a certain value, it may be better to target maximizing the minimum eigenvalue of the three by three position block of the covariance matrix. This approach would allow the implementer to determine which observer and when to schedule follow up observations of a space object while keeping the position uncertainty under some defined limit.

3. RESULTS

3.1 Incorporating Additional Observations

It has been previously explained how both Gooding IOD and the dynamics-neutral IOD can be adapted to accept an arbitrary number of angles-only observations. However, the question of how this impacts the quality of the results remains. To investigate this, the IOD algorithms are tested against a series of randomly generated target object orbits with measurement error on the observations. Each case is tested with a different number of observations to evaluate the impact of adding these additional measurements.

Test cases used three, four, or five measurements and each trial case was tested under all three scenarios. The first and last observations remained the same in all trial cases. The other observations were evenly distributed in time between the first and last. This ensured that the higher number of observation cases did not benefit from a longer arc length subtended by the full series of observations.

The target object orbits were generated as near geosynchronous objects with randomly generated deviations to the initial state. The accuracy was assessed based on the position error of the central state estimate output by the IOD function. For cases with an even number of observations the middle state was chosen by rounding down. The magnitude of the position error of this middle state was recorded for the full Monte-Carlo simulation and a cumulative distribution function (CDF) of the position error was generated for each case. The results for the modified Gooding method are shown in figure 2 and the results for the dynamics-neutral IOD algorithm are shown in figure 3.

Adding in these additional observations reduced the error magnitude for the majority of cases. Both cases show a

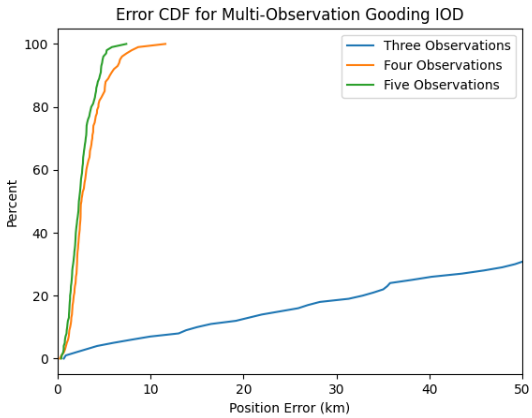

Fig. 2: CDF of position errors for the modified Gooding IOD method for 3, 4, and 5 measurements.

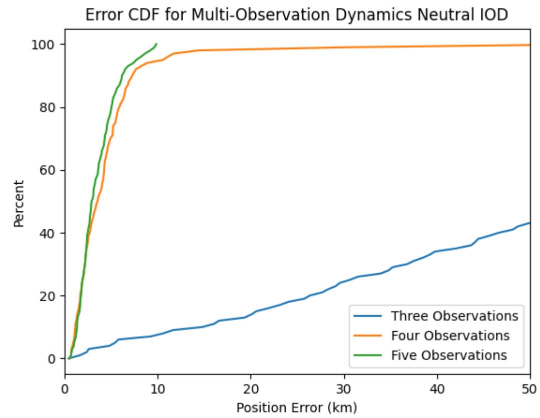

Fig. 3: CDF of position errors for the dynamics-neutral IOD method for 3, 4, and 5 measurements.

significant improvement with the addition of more than three measurements. However, the change from four observations to five is much weaker. In both cases, five observations slightly outperformed four for the trial states. There was a small number of cases for the dyanmics-neutral algorithm with four observations that performed significantly worse than others. Without the two-body assumption, the dynamics neutral algorithm is more sensitive to error in its input observations. It will generally not perform as well as Gooding and other algorithms that assume two-body dynamics in cases where two-body dynamics accurately dominate the physics. However, it is applicable in a more broad set of cases than Gooding.

3.2 Predicting Information Gain

A Monte-Carlo test is used to investigate the impact of information-based observation scheduling on IOD. For each trial case, the two initial observations are generated with white Gaussian noise added onto the expected measurements to simulate measurement error with a one sigma magnitude in each direction of five arcseconds. White Gaussian process noise is also added to the propagation step to simulate unmodeled forces in the dynamics model. The observer locations and time between the first observations are fixed. From there, the third observation is planned by either randomly choosing an observer and observation time from a fixed set of potential observers or that same set of observers and the maximum allowable time is fed into the algorithm that determines the time and observer that would maximize the information matrix determinant and that time and observer is chosen for the third observation plan.

When estimating the propagated information matrix, the process noise is modeled using static noise compensation and the process noise state covariance matrix, Q, is the same covariance matrix as the one used to scale the actual process noise between the observations. In this way the observation planning algorithm is blind to the actual values chosen for the process noise, but it understands the distribution the noise was chosen from. This simulates an accurate understanding of the magnitudes of the unmodeled forces in the state propagator. The process noise covariance matrix used is shown in equation 11 and it depends on the chosen scaling factor, q, and the time of the propagation Δt.

$$Q = q^2 \begin{bmatrix} \frac{\Delta t^3}{3} & 0 & 0 & \frac{\Delta t^2}{2} & 0 & 0 \\ 0 & \frac{\Delta t^3}{3} & 0 & 0 & \frac{\Delta t^2}{2} & 0 \\ 0 & 0 & \frac{\Delta t^3}{3} & 0 & 0 & \frac{\Delta t^2}{2} \\ \frac{\Delta t^2}{2} & 0 & 0 & \Delta t & 0 & 0 \\ 0 & \frac{\Delta t^2}{2} & 0 & 0 & \Delta t & 0 \\ 0 & 0 & \frac{\Delta t^2}{2} & 0 & 0 & \Delta t \end{bmatrix} \tag{11}$$

In either case, this third observer and time is used to model the third observation. Again, the process noise is applied to the propagated state and then measurement noise to the measurement value. This yields a set of three observations which are then fed into the chosen IOD algorithm. The IOD algorithm is setup to find the state associated with the

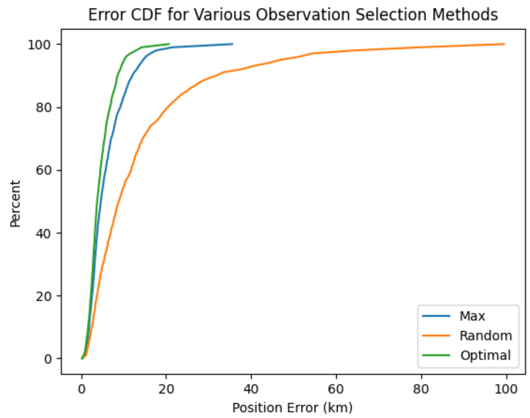

Fig. 4: Empirical CDF for position errors under different observation planning approaches for Gooding IOD.

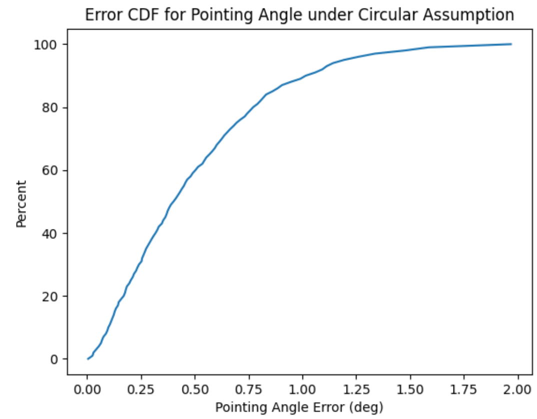

Fig. 5: Empirical CDF for angle errors for fully simulated GEO target orbits.

middle observation time. The error for this result is the magnitude of the position difference between the actual state of the target object at the second observation time and the IOD estimated state at the second observation time. For each trial case in the Monte-Carlo simulation, this scalar error is recorded and the distribution of these errors over the full run is used to assess the efficacy of the method for planning the final observation.

This approach was first used with the traditional, three-observation Gooding IOD algorithm [3]. The Monte-Carlo analysis was setup in two different ways. First the target satellite was entirely simulated and placed into a geosynchronous orbit with randomly generated deviations to the state. The position error magnitudes for this case were collected and the empirical cumulative distribution function was plotted and is shown in figure 4. The figure shows three curves for the three different approaches used to determine the third observation. One used random a random observer and a random time between the second observation and a set maximum time. One used a random observer and the maximum time as the third observation time, and finally was the information optimizing approach. The figure shows the optimized approach outperforming both of the other approaches with the maximum time approach performing worse than the random time case. To keep the comparison fair, the random number generator was seeded so that the applied process noise and measurement errors were the same for all three cases [5]. The random time case used a separate random number generator for choosing the time so as not to interfere with this consistency between cases.

One potential issue with the circular assumption for planning the third observation is that if the orbit is more eccentric, the propagated state will differ from the expected case with zero eccentricity. For tasking a real sensor this runs the risk of the target being out of frame for the third observation. To quantify this risk, the angle error between the expected and actual observations were recorded and a CDF of these errors is shown in figure 5. One can interpret the plot as the field of view required to successfully observe a certain portion of the targets at the third time. For building a system based on this circular assumption approach to IOD, generating expected error distributions like this can be a useful tool for driving hardware requirements. In this case, the generated scenarios differed from the expected observation angle by up to two degrees with about sixty percent less than half a degree off.

The second approach to the Monte-Carlo test did not use random generation of orbits but rather sampled the real satellite catalog. The full public catalog for August 7th, 2023 was downloaded and sampled to produce a pseudo-random sampling of real objects. This was done by simply taking every tenth object, propagating its state to a fixed set time and and then simulating two observations and different approaches for planning the third observation identically to the process described previously.

The results for this case are shown in figure 6. The Gooding IOD algorithm could not converge for some cases, and these cases continue off the figure. The range of the plot was limited to highlight the errors for cases where the IOD algorithm did converge to a solution successfully. Once again, we see that the optimized approach outperformed the random selection and fixed time approaches. This shows that the information maximizing approach does consistently provide more accurate estimates for the Gooding IOD algorithm. The angle errors under the circular assumption for the sampled portion of the satellite catalog is shown in figure 7. While most errors are small, highly eccentric orbits

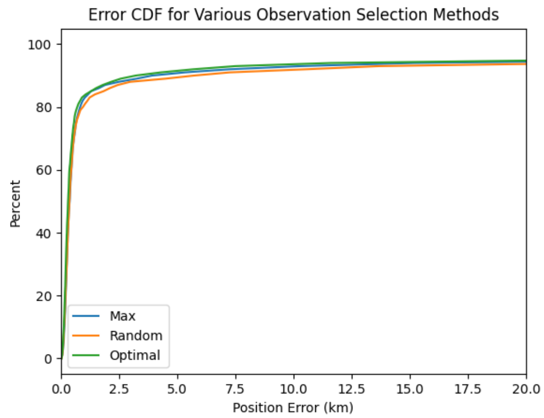

Fig. 6: Empirical CDF for position errors under different observation planning approaches for Gooding IOD and catalog data.

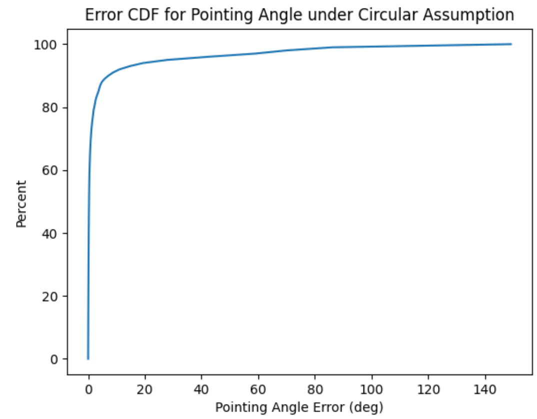

Fig. 7: Empirical CDF for angle errors for a random sampling of the satellite catalog.

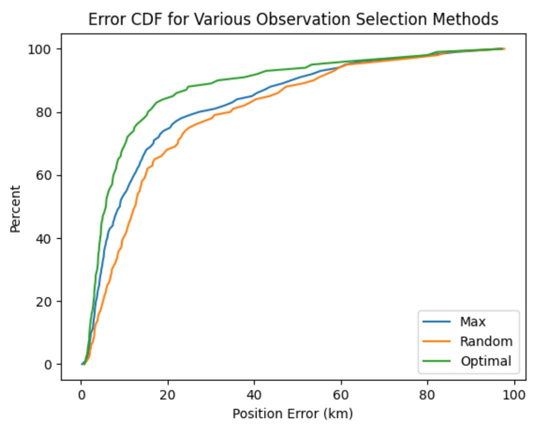

Fig. 8: Empirical CDF for position errors under different observation planning approaches for dynamics-neutral IOD and simulated GEO orbits.

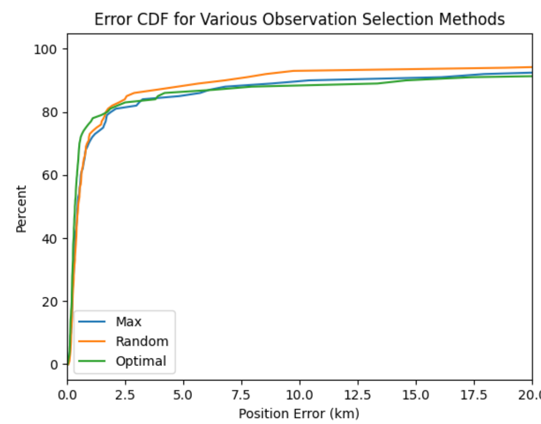

Fig. 9: Empirical CDF for position errors under different observation planning approaches for dynamics-neutral IOD and catalog data.

can have very large errors as the circular assumption inaccurately predicts its future position and the target object ends up far from the expected position. The observers in this case were taken as very low Earth orbiting satellites to ensure there was no ambiguity in the circular assumed orbits as there may be in the case with observer orbits above the targets.

This full analysis was repeated for both Monte-Carlo approaches with the dynamics-neutral IOD algorithm. The only change in the execution was the substitution of the IOD algorithm itself. This change had no impact on the angle errors as those are dependent on the circular assumption and independent of the IOD algorithm. The position error results for the simulated GEO targets are shown in figure 8 while the results for the sampled catalog are shown in figure 9.

The proposed optimal method outperforms the others in the fully simulated case. It performs better in the majority of cases for the catalog-based data as well. However, random takes the lead for the cases with high position error for around twenty percent of the cases. These are the cases where the target object has a higher eccentricity and include cases with high enough error that the IOD algorithm can be considered to have failed. The circular assumption used to perform the prediction breaks down in these cases.

4. CONCLUSION

We have developed a method for scheduling information-maximizing observations for orbit determination. The approach has been shown to out-perform uninformed methods for planning future observations in initial orbit determination when the core assumptions used to propagate the state into the future are approximately accurate. Since the vast majority of artificial space objects have very low eccentricity, the circular assumption on two angles-only observations approximates the truth well enough to optimally select a future observer and observation time in more than eighty percent of tested cases. By verifying the method against a random sampling of the public satellite catalog, we have demonstrated the utility of the approach for a realistic set of targets. The method was demonstrated for both the Gooding IOD algorithm and a dynamics-neutral optimization approach that can be extended to work in dynamics beyond the near-Earth environment. Along with this, we have described improvements to the dynamics neutral IOD algorithm to both improve its ability to find initial estimates for the optimization it performs and to include additional data. This result was compared against a similar adaptation to the Gooding IOD algorithm by Henderson, Mortari, and Davis [4].

REFERENCES

[1] 18th Space Defense Squadron. Space-track.

[2] RH Gooding. *A new procedure for orbit determination based on three lines of sight (angles only)*. Defence Research Agency Farnborough, UK, 1993.

[3] CS GROUP. Orekit source code: Iodgooding.java.

[4] Troy A Henderson, Daniele Mortari, and Jeremy Davis. Modifications to the gooding algorithm for angles-only initial orbit determination. In *AAS/AIAA Space Flight Mechanics Meeting*, 2010.

[5] Manuel Schubert, Christopher Kebschull, and Steven Horstmann. Analysis of different process noise models in typical orbitdetermination scenarios. In *8th European Conference on Space Debris*, volume 8. ESA Space Debris Office Darmstadt, German, 2021.

[6] Kyle Stoker, James Johnson, Kira Abercromby, S Liam, and Y Morgan. Angles-only orbit determination accuracies with limited observational arc. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2020.

[7] Samuel Wishnek, Marcus J Holzinger, and Patrick Handley. Robust cislunar initial orbit determination. In *AMOS Conf. Proc*, 2021.