

Sensor Tasking for Space Situational Awareness in Cislunar Space



Master Thesis

Keshav Agarwal

Sensor Tasking for Space Situational Awareness in Cislunar Space

Master Thesis

Student Name

Keshav Agarwal

Supervisor: Dr. S. Gehly
Supervisor: L. Ngo
Project Duration: March, 2025 - December, 2025
Faculty: Faculty of Aerospace Engineering, Delft

Preface

This thesis marks the culmination of more than 6 years of my college journey. It has been an insane ride, with more ups and downs than I can count. At this moment, all I feel is relief that I am finally done, and gratitude for all the people in my life.

Most of all, I am grateful to my family. Even though they were thousands of miles away, their constant support was invaluable to me. I have always felt very loved, I know I can always count on them. All that I am today is thanks to my family, they deserve the credit for every good part of me, and I get the blame for every bad part. I love you all, and I consider myself so lucky to have you all. Special credit goes to my grandparents, and I hope they are looking down on me with a smile on their face, and I will continue to strive to make them proud.

Without my friends, these last six years would have been meaningless. I am grateful for the friends I've made along the way, many of whom I know will stay with me all my life, and the countless memories I've made. SGC, TISB, Purdue and Rotterdam - four places on three different continents that I have been fortunate to call home, and it is all because of my friends. Min, my dear friend, I wish you could know how much you've helped me grow, and how much I love and miss you.

Last but not the least, I want to thank my supervisors Steve Gehly and Linh Ngo, for their constant support, guidance and invaluable insight. Thank you for putting up with me and getting me over the line.

Keshav Agarwal
Rotterdam, December 2025

Summary

The rapid expansion of activity in cislunar space creates an urgent need for dedicated Space Situational Awareness (SSA) capabilities beyond Earth orbit. Ground-based sensors and sensors in Earth's orbit suffer from limited coverage and unfavorable viewing geometry for objects in the Earth–Moon system, motivating the deployment of space-based observers in cislunar orbits. However, the performance of such observers depends critically on both their orbital placement and the sensor tasking strategy used to schedule observations. This thesis investigates how cislunar SSA by a space-based observer can be supported through the joint evaluation of observer orbits and sensor tasking reward functions.

A simulation framework is developed using the Earth–Moon Circular Restricted Three-Body Problem (CR3BP) and implemented using the TUDAT astrodynamics toolbox. A catalog of 22 representative target orbits, including Near Rectilinear Halo Orbits (NRHOs), Distant Retrograde Orbits (DROs), Lyapunov, and Halo orbits, is constructed from the JPL periodic orbit database, together with eight candidate observer orbits spanning L1/L2 Lyapunov, L1/L2 Halo, and Distant Prograde orbit families. The observer is simulated as an optical angles-only sensor subject to limiting magnitude and solar, lunar, and Earth exclusion constraints. Target states are estimated using an Extended Kalman Filter, and a greedy tasking algorithm selects at each decision epoch the visible target that maximizes a chosen reward function. The reward functions considered are information gain based metrics (Kullback–Leibler divergence, mutual information, Cauchy–Schwarz divergence), an age-of-information reward that schedules the target that has been unobserved the longest, and a finite-time Lyapunov exponent (FTLE) based reward that considers the expected growth in state uncertainty over a prediction horizon.

Performance is evaluated using an *observed RMSE*, which measures estimation error at measurement update times, and a *complete RMSE*, which captures end-to-end performance over the full trajectory including propagation intervals without measurements, along with a discussion of coverage and fairness metrics across targets. The results show that a suitably chosen cislunar observer yields markedly better tracking performance than Earth-based observers. For the same target catalog, the cislunar observer reduces the mean observation position RMSE from 2851 km (HAPO/LEO) to 5.83 km and significantly lowers the complete RMSE as well. Within cislunar space, an L2 North Halo orbit achieves the lowest complete RMSE, while a Distant Prograde orbit offers the lowest instantaneous (observation) RMSE, highlighting trade-offs between range, visibility, and occultation.

Across reward functions, information-gain-based metrics and age-of-information deliver the most competitive and balanced performance. When an extreme, rarely observable target is removed from the catalog, all three information-gain rewards produce very similar tasking patterns and RMSE, while an age-of-information provides a simple, robust baseline within 12 km of the best complete RMSE and offers highly uniform coverage. In contrast, FTLE-based rewards perform poorly, biasing observations toward distant, geometrically challenging targets and yielding higher average and worst-case errors.

Overall, the thesis demonstrates that effective cislunar SSA requires carefully chosen observer orbits and well-designed scheduling algorithms. It provides a unified framework to evaluate these choices under consistent dynamical and observational assumptions. Observer placement is shown to be key to the overall SSA goal, and observational geometry between the observer and the targets is recognized as a key factor for error reduction. It shows that certain cislunar trajectories may suffer from low visibility even from a cislunar observer, and a myopic scheduling algorithm may have difficulties tracking such a trajectory. Information gain-based rewards are effective in minimizing position errors, and out of the three rewards considered in this work, Shannon mutual information (equivalent to Fisher Information Gain under linear Gaussian assumption) has the best estimation accuracy. The age-of-information reward serves as a robust and stable objective and can be used as a baseline for comparison. The formulation of FTLE-based reward used in this work performs poorly. The work lays the groundwork for future studies with a higher number of potential observer orbits, multi-sensor architectures, advanced tasking algorithms with more realistic constraints and penalties, and higher-fidelity dynamical models.

Contents

Preface	i
Summary	ii
Nomenclature	vi
1 Introduction	1
2 Literature Review	4
2.1 Cislunar Orbits	4
2.2 Sensor Tasking	8
3 Research Questions	11
3.1 Research Gaps and Relevance to This Work	11
3.2 Research Questions	11
4 Methodology	13
4.1 Cislunar Orbital Mechanics	13
4.2 State Estimation	17
4.3 Measurement Model	19
4.4 Scheduling Observations	19
4.5 Reward Functions	20
4.6 Visibility Constraints for a Space-Based Optical Sensor	22
4.7 Generation of Orbital Data	24
4.8 Simulation Framework	29
5 Simulation verification	33
5.1 CR3BP state propagation	33
5.2 Extended Kalman Filter	34
5.3 Sensor Tasking Logic	37
6 Results	41
6.1 Comparison of Observers in Earth's orbit and Cislunar Orbit	41
6.2 Comparison of Different Observer Orbits	45
6.3 Discussion of orbits of low visibility	50
6.4 Comparison of reward functions	51
7 Conclusion and Recommendations	58
7.1 Conclusion	58
7.2 Recommendations and Future Work	60
References	63
A RMSE Computation	67
B Project Plan	68

List of Figures

2.1	Lagrange points in the Earth-Moon system	5
2.2	Flow chart of IG-based tasking (taken from [54])	10
4.1	A representation of the complex non-planar nature of cislunar orbits (reprinted from [16])	14
4.2	Inertial (left) and Synodic (right) reference frame for the CR3BP (taken from [50])	16
4.3	Representation of solar phase angle	23
4.4	Representation of the solar exclusion angle	24
4.5	Evolution of a very stable DRO for 10 synodic periods	26
4.6	Evolution of an unstable Halo orbit for 2 synodic periods	26
4.7	Evolution of an unstable Halo orbit for 3 synodic periods	27
4.8	Evolution of an highly unstable Halo orbit for 2 synodic periods	28
5.1	The norm of the error in TUDAT's propagated position in reference to the position ephemeris from JPL database of Target 1	33
5.2	The norm of the error in TUDAT's propagated position in reference to the position ephemeris from JPL database of Target 10	34
5.3	Measurement Residuals of the EKF	35
5.4	Position Estimate Error and the 3σ covariance boundary	35
5.5	Observability angle of the target from the observer	37
5.6	Range of the target from the observer	37
5.7	Normalized Innovation Squared (NIS)	38
5.8	Figure showing which target is tasking for observation by the sensor tasking algorithm .	39
5.9	Figure showing the epochs at which each of the targets is visible to the observer. It is overlayed with the times at which the target is scheduled for observation	39
5.10	Figure showing target visibility and target tasked over a small window of time	39
5.11	Figure showing target visibility and target tasked over a small window of time	40
6.1	Visibility and observation charts for Observer 3, HAPO observer, and LEO observer. . .	42
6.2	Visibility and observation charts for 6 targets for Observer 3 and HAPO observer	43
6.3	Visibility and observation chart of Observers 2 and 3	46
6.4	Figure showing the targets that are scheduled for observation at each time step, from observers 2 and 3	46
6.5	Visibility and observation chart from observer 4 and 6	48
6.6	Visibility and observation chart of Observer 8	48
6.7	Figure showing the orbit of Target 17 along with Observer 8	50
6.8	Scheduling of targets from IG-based reward functions from observer 8	52
6.9	Scheduling of targets from IG-based reward functions from observer 8	55
B.1	Gantt chart of initial planning	70

List of Tables

4.1	Radius of celestial bodies used in calculation of exclusion angles	24
4.2	Non-dimensionalization constants for the Earth–Moon CR3BP.	25
4.3	Catalog of orbits chosen for observer placement	28
4.4	Catalog of target orbits	29
4.5	Parameters used to create CR3BP system in TUDAT. μ refers to CR3BP mass parameter	30
4.6	Simulation time control	30
4.7	EKF noise covariance matrices	31
5.1	Residual Statistics	34
5.2	NIS Statistics	36
6.1	Keplerian orbital elements of Earth observers	41
6.2	Comparison of the average RMSE of targets from cislunar and Earth observers	43
6.3	Catalog of observer orbits	45
6.4	Complete and Observation RMSE of position and velocity, averaged across all the targets, for each observer	45
6.5	Average range across all targets for each observer	47
6.6	RMSE position and observed-only RMSE position for different reward functions and observers	51
6.7	Complete and observed RMSE in position for individual targets for IG-based reward functions from observer 8. Columns with (C) refer to complete RMSE and columns with (O) refer to observed RMSE	52
6.8	Average range from observer 8 to observed targets for each reward function.	54
6.9	RMSE position and observed-only RMSE position for different reward functions and observers	55
6.10	Observation statistics for observer 8 over 21 targets (Target 17 excluded). \bar{N}_{obs} and $\sigma_{N_{\text{obs}}}$ denote the mean and standard deviation of the number of observations per target; the remaining columns show the median, minimum, maximum and 95th percentile.	57
6.11	Per-target all-times position RMSE statistics for observer 8 over 21 targets (Target 17 excluded). All RMSE values are in km. The last column shows the correlation coefficient between per-target observation counts and RMSE.	57

Nomenclature

Abbreviations

Abbreviation	Definition
3BP	Three-Body Problem
AOI	Age-of-Information
ASO	Anthropogenic Space Objects
CR3BP	Circular Restricted 3-Body Problem
CS	Cauchy-Schwarz
DCS	Cauchy-Schwarz divergence
Dec	Declination
DPO	Distant Prograde Orbit
DRO	Distant Retrograde Orbit
ECI	Earth-Centered Inertial (frame)
EKF	Extended Kalman Filter
E-M	Earth-Moon
FIG	Fisher Information Gain
FTLE	Finite-Time Lyapunov Exponent
GEO	Geostationary Orbit
IG	Information Gain
JPL	Jet Propulsion Laboratory
KL	Kullback-Leibler
KLD	Kullback-Leibler divergence
LEO	Low Earth Orbit
LOS	Line-Of-Sight
LPO	Lagrange-Point Orbit (Libration-Point Orbit)
MI	Mutual Information
NRHO	Near Rectilinear Halo Orbit
n.d.	non-dimensional
RA	Right Ascension
S/C	Spacecraft
SSA	Space Situational Awareness
STM	State-Transition Matrix
TUDAT	TU Delft Astrodynamics Toolbox

Symbols

Symbol	Definition	Unit
a	Semi-major axis (n.d. CR3BP description)	[-]
C	Jacobi constant of the CR3BP	[-]
C_d	Diffuse reflection (albedo) coefficient of the target	[-]
d	Distance from observer to target	[km]
$F(\mathbf{X}, t)$	Dynamical system function for state propagation	[state/s]
$F_{\text{diff}}(\psi)$	Reflected solar flux factor as function of phase angle	[-]
G	Gravitational constant	[km ³ kg ⁻¹ s ⁻²]
$G(\mathbf{X}, t)$	Function for state-measurement space mapping	[state/s]
H	Prediction horizon length for FTLE reward	[-]

Symbol	Definition	Unit
$\tilde{\mathbf{H}}_k$	Linearized measurement matrix at epoch t_k	[-]
i_k^*	Index of selected target at epoch t_k	[-]
\mathbf{I}_n	$n \times n$ identity matrix	[-]
\mathbf{J}^-	<i>A priori</i> and information matrices	[state ⁻²]
\mathbf{J}^+	<i>A posteriori</i> and information matrices	[state ⁻²]
$J_{i,k}$	One-step reward for target i at decision epoch t_k	[-]
\mathbf{K}_k	Kalman gain matrix at epoch t_k	[-]
l^*	Length unit	[km]
m	Generic mass	[kg]
m_E	Mass of the Earth	[kg]
m_M	Mass of the Moon	[kg]
m^*	Mass unit	[kg]
m_{object}	Apparent visual magnitude of the target	[mag]
m_{sun}	Apparent visual magnitude of the Sun	[mag]
n	Mean motion of the E–M system	[rad/s]
$p(x), q(x)$	Probability density functions used in divergence measures	[-]
$\bar{\mathbf{P}}_k$	A priori (predicted) covariance at epoch t_k	[state ²]
$\hat{\mathbf{P}}_k$	A posteriori (estimated) covariance at epoch t_k	[state ²]
\mathbf{Q}_k	Process noise covariance at epoch t_k	[state ²]
R_B	Physical radius of body B	[km]
\mathbf{R}	Radius of the (spherical) target object	[m]
\mathbf{R}_k	Measurement noise covariance at epoch t_k	[meas ²]
\mathbf{r}	Spacecraft position vector in synodic frame	[n.d.]
$\dot{\mathbf{r}}$	Spacecraft velocity vector in synodic frame	[n.d.]
\mathbf{r}^I	Position vector in inertial (ECI) frame	[km]
\mathbf{r}^S	Position vector in synodic (rotating) frame	[n.d.]
\mathbf{r}^M	Position vector in Moon-fixed frame	[km]
$\mathbf{r}_o, \mathbf{r}_t$	Observer and target position vectors in ECI frame	[km]
$\mathbf{r}_S, \mathbf{r}_E, \mathbf{r}_M$	Sun, Earth, Moon position vectors	[km]
\mathbf{r}_{o-t}	Line-of-sight vector from observer to target	[km]
$\hat{\mathbf{r}}_{o-t}$	Unit line-of-sight vector from observer to target	[-]
$\mathbf{s}^I(t)$	Spacecraft state vector in inertial frame	[km, km/s]
$\mathbf{s}^S(t)$	Spacecraft state vector in synodic frame	[km, km/s]
$\mathbf{s}^M(t)$	Spacecraft state vector in Moon-fixed frame	[km, km/s]
\mathbf{S}_k	Innovation covariance at epoch t_k	[meas ²]
t	Continuous time variable	[s]
t^*	Characteristic time unit	[s]
t_k	Discrete observation / decision epoch	[s]
t_{last}	Time of most recent observation of a target	[s]
T	Orbital period	[days]
\mathbf{u}_B	Line-of-sight unit vector from observer to body B	[-]
U	Pseudo-potential function of the CR3BP	[-]
\mathbf{X}	Full state vector of the system	[state]
$X^*(t)$	Reference (nominal) trajectory used in EKF	[state]
$X_{\text{es},k}$	Estimated state at epoch t_k	[state]
x, y, z	Cartesian position components in synodic frame	[-]
$\dot{x}, \dot{y}, \dot{z}$	Velocity components in synodic frame	[-]
\mathbf{Y}_k	Measurement vector at epoch t_k	[rad]
y_k	Measurement residual (innovation in measurement space)	[rad]
\mathcal{F}_k	Feasible target set at decision epoch t_k	[-]
\mathcal{T}	Catalog of targets	[-]
α	Right ascension or phase angle	[rad]

Symbol	Definition	Unit
β_B	Angular separation between LOS and body B (Sun/Earth/Moon)	[rad]
Γ_k	Process noise mapping matrix for covariance propagation	[state/(noise)]
Δt	Time step between epochs t_{k-1} and t_k	[s]
δ	Declination angle in measurement model	[rad]
ϵ_k	Measurement noise at epoch t_k	[rad]
Θ_B	Exclusion half-angle (apparent angular radius) of body B	[rad]
λ_{\max}	Dominant eigenvalue of the monodromy matrix	[-]
μ	Non-dimensional mass parameter	[-]
ν	Stability index of a periodic CR3BP orbit	[-]
ρ_{o-t}	Range (distance) between observer and target	[km]
$\rho_{o-t,xy}$	Projected range in the xy -plane	[km]
Ψ_f	Observability angle	[rad]
Solar angle (Sun-object-observer)	phase (Sun-observer)	[rad]
$\sigma_{\alpha,k}$	Standard deviation of right ascension measurement noise	[rad] or [arcsec]
$\sigma_{\delta,k}$	Standard deviation of declination measurement noise	[rad] or [arcsec]
$\Phi_{k,k-1}$	STM from epoch t_{k-1} to t_k	[-]
$\Phi_k^{(H)}$	Horizon STM over H prediction steps for FTLE reward	[-]
Ω	Angular velocity vector of rotating (synodic) frame	[rad/s]
D_{KL}	Kullback–Leibler divergence between two Gaussian states	[-]
D_{CS}	Cauchy–Schwarz divergence between two Gaussian states	[-]
r_{FIG}	Fisher Information Gain reward	[-]
r_{KL}	Kullback–Leibler divergence reward	[-]
r_{CS}	Cauchy–Schwarz divergence reward	[-]
r_{AOI}	Age-of-Information reward	[s]
r_{FTLE}	FTLE-based reward (max eigenvalue of predicted covariance)	[state ²]

1

Introduction

The last six decades have witnessed an exponential growth of human activity in outer space, and has transformed outer space from a vast, untapped frontier into a heavily utilized and increasingly congested domain [17]. This transformation has led to the emergence of **Space Situational Awareness (SSA)** as a critical field, concerned with the detection, tracking, identification, and characterization of objects in space [17]. While SSA was initially developed in the context of national defense and early warning systems, it has become foundational to ensuring the safety, security, and sustainability of modern space operations [20].

In Earth's orbit, particularly Low Earth Orbit (LEO) and Geostationary Orbit (GEO), the growth of **Anthropogenic Space Objects (ASOs)**, which include active and defunct satellites, spent rocket bodies, and fragments from collisions and explosions, has created a dense and hazardous environment [54]. As of 2025, it is estimated that more than 130 million pieces of debris larger than 1 mm, and over 54,000 tracked objects larger than 10 cm, occupy Earth's orbital space [13]. These objects travel at velocities of up to 28,000 km/h, where even a millimeter-sized fragment can inflict significant damage on functioning spacecraft [13].

The threat posed by this space debris is genuine, and collisions could lead to cascading debris generating events such as the Kessler Syndrome [28]. The Big Sky Theory, which states that the probability of collision of human-made objects is nearly zero due to the immense volume of space and the relatively small size of human-made objects, was proved false after the 2009 Iridium-Cosmos collision [34]. Other notable incidents such as the 2007 Chinese anti-satellite test and the recent fragmentation events in LEO have demonstrated how individual actions can have cascading effects on orbital safety [12]. This has spurred the development of **Space Traffic Management (STM)** systems that integrate cataloging, tracking, maneuver planning, and collision avoidance, each of which is dependent on accurate and timely SSA.

A critical sub-component of SSA is **sensor tasking**, which is the process of dynamically assigning sensing actions to a network of space surveillance instruments, including radars, telescopes, and space-based sensors [54]. In Earth-based SSA, sensor tasking has evolved from deterministic scheduling and heuristic rules to optimization-based algorithms, probabilistic decision-making, and more recently, machine learning approaches such as reinforcement learning and Bayesian optimization [54]. These techniques aim to maximize coverage, reduce uncertainty, ensure frequent revisits to high-priority objects, and adapt to real-time operational constraints. Central to these methods are **reward functions** that translate mission goals and constraints into a scalar objective for each decision: typical terms include expected information gain or entropy reduction on the state estimate, probability of detection or reacquisition, maintenance of custody, coverage of priority objects/regions, and revisit timeliness, while penalizing slews, settle/dwell costs, exclusion-angle violations, and resource usage.

Humanity is now at the cusp of a new phase in space exploration: expansion into the cislunar region, which is the volume of space between Earth and the Moon (and beyond), dominated by gravitational forces from both the Earth and the Moon [16]. This region is becoming increasingly important due to

a surge in planned and ongoing missions aimed at establishing a sustained lunar presence for exploration and commercial utilization. The United States' Artemis program aims to establish a long-term presence on the Moon via the Lunar Gateway. Meanwhile, China's Chang'e missions and Russia's collaboration proposals, along with India's Chandrayaan series and contributions from Europe, Canada, Japan, and South Korea, highlight a shared interest in lunar exploration, leading to a dramatic increase in the number of spacecraft operating in cislunar space [24]. The strategic significance of this domain extends beyond science and exploration and includes infrastructure development, resource extraction, communications relay, and even defense posturing [1]. The Moon's south pole, for example, is a focal point of interest due to the possible presence of water ice in permanently shadowed regions. Likewise, Earth–Moon Lagrange points, particularly L1 and L2, are ideal locations for space stations, observation platforms, and logistics hubs [51][24].

However, the cislunar region is inherently more complex and risk-prone than Earth orbit. It is very large; while GEO lies at close to 36,000 km from Earth's surface, the Moon is nearly 384,000 km away. Also, the cislunar region of interest extends beyond the moon, encompassing a search volume more than 1,700 times larger than the volume around GEO [16]. Additionally, the dynamics in this region are non-intuitive because spacecraft experience the concurrent gravitational influence of both Earth and the Moon. While the full motion is governed by the n-body problem with additional perturbations (e.g., solar gravity, Earth's non-spherical gravity field, and solar radiation pressure etc.), it is often simplified and modeled using the Circular Restricted Three-Body Problem (CR3BP). But even within the CR3BP framework, objects can exhibit quasi-periodic, unstable, and non-planar trajectories, and these solutions are approximations of the true dynamics and typically require regular station-keeping to remain on track in the vast cislunar space [16][1].

Although current cislunar space has relatively few ASOs compared to Earth orbit, this will not remain the case for long. Projections suggest that dozens to hundreds of spacecraft, including crewed vehicles, logistics modules, lunar orbiters, and relay satellites, will be operating in overlapping spatial corridors within the next decade [9][24]. The possibility of collisions and debris-generating events is real and growing. Importantly, fragmentation events in cislunar space could create long-lived debris fields due to weak gravitational and atmospheric decay, posing enduring risks to future missions [48]. In such an environment, the risk of debris must be treated seriously.

Despite these challenges, there is currently no operational SSA infrastructure to monitor cislunar space with sufficient precision or coverage. Earth-based telescopes and radars are largely ineffective due to line-of-sight occlusions, Earth's rotation, weather constraints, and signal attenuation over large distances [1][9]. Even in well-equipped simulations, up to 40 percent of cislunar trajectories, such as that of Artemis I, can be unobservable for days at a time [15].

As a result, the use of sensors in periodic cislunar orbits (e.g., Earth–Moon L1/L2 halo orbits, NRHOs, or DROs) is seen as a necessary next step [15]. These sensors, positioned in carefully chosen cislunar orbits, offer persistent observation capabilities with more favorable viewing geometries. These sensors would have shorter range to the targets, line-of-sight unconstrained by Earth's horizon or weather as well as better solar phase angles for better illumination. However, the cost and logistical complexity of deploying such systems means that each sensor must be tasked as efficiently as possible to deliver maximum SSA coverage. This makes sensor tasking a mission critical component of cislunar missions [11].

The core objective of this thesis is to address the question: How can cislunar Space Situational Awareness by a space-based observer in a cislunar orbit be supported through the evaluation of sensor tasking strategies and observer orbits? This includes investigating how Earth-based sensor tasking strategies can be adapted and extended to operate effectively under cislunar conditions. This thesis assesses the relative effectiveness of different sensor tasking strategies and observer orbit configurations by analyzing their impact on the cumulative uncertainty reduction of tracked targets. The research specifically focuses on: 1) Evaluating different sensor tasking reward functions in the context of greedy tracking algorithms, and how these influence tracking effectiveness. 2) Analyzing how the orbital placement of the observing sensor (e.g., NRHO, DRO, L1/L2 Halo orbits) affects tracking effectiveness.

To accomplish this, the thesis employs dynamical simulations under the Circular Restricted 3-Body Problem using the TUDAT astrodynamics toolbox [42], creating a synthetic environment with multiple

target objects and a single observer. Different observer placement and sensor tasking configurations are tested under comparable conditions, with outputs such as position and velocity errors and coverage metrics analyzed to evaluate tracking effectiveness and identify trade-offs.

This research contributes to bridging the existing gap between Earth-based SSA frameworks and the unique demands of the cislunar regime. Through a comparative analysis of the performance of different tasking strategies in different orbits, it lays foundational groundwork for future single and/or multi-sensor architectures and autonomous SSA constellations designed to safeguard humanity's next steps beyond Earth.

2

Literature Review

2.1. Cislunar Orbits

Cislunar space refers to the region of space between Earth and the Moon, including the vicinity of the Moon and all five Earth–Moon Lagrange points [1]. In this region, the gravitational influences of both Earth and Moon are significant, so the Circular Restricted Three-Body Problem (CR3BP) is a useful dynamical model [16]. In the CR3BP, the Earth and Moon (the two primaries) orbit their common barycenter in circular orbits, and a third body of negligible mass (e.g. a spacecraft) moves under their combined gravity in a rotating frame.

Recent works in cislunar SSA demonstrate the inadequacy of a ground-based SSA infrastructure for comprehensive monitoring of the cislunar region [15][11][51]. Frueh et al. [1] showed that space-based sensors located in favorable orbits, such as NRHO and DRO, offer various observational advantages such as longer tracking windows and better geometry for phase angle minimization. Other works such as [43], [56], [50] also stress the need for a cislunar observational architecture and analyze and compare the advantages offered by different cislunar orbits to a cislunar observer. This section discusses various periodic cislunar orbital families with regard to their advantages and/or disadvantages for cislunar surveillance. Before a discussion of these orbits, a brief discussion of Lagrange points and stability index is given.

Lagrange (Liberation) Points

The five Lagrange points (L1 through L5) are equilibrium points in the rotating synodic frame where the gravitational and centrifugal forces balance. This means that in CR3BP, a spacecraft placed at a Lagrange point remains there indefinitely [16], since the forces from the two primaries cancel out at these points. They appear to be stationary in the synodic frame, and appear to move in a perfectly circular orbit around the system's barycenter in the inertial frame. They can be found numerically by setting the velocity and accelerations in the equations 4.3 to zero. Doing this also results in the z position being zero, meaning that these Lagrange points are coplanar with rotating E-M system. The five Lagrange points for the E-M system are shown in Figure 2.1. Earth–Moon L1, L2, and L3 lie along the line connecting Earth and Moon (L1 between Earth and Moon, L2 beyond the Moon's far side, and L3 beyond Earth opposite the Moon), and are unstable equilibrium points, while L4 and L5 lie at 60° ahead and behind the Moon's orbit and are stable points where small perturbations cause a body to oscillate around the equilibrium [29]. A Lagrange point is said to be:

- **Stable** if small perturbations result in bounded oscillations around the point (i.e., the spacecraft does not drift away).
- **Unstable** if perturbations grow exponentially over time, requiring active station-keeping to remain nearby.

These Lagrange points offer various advantages for cislunar missions. For example, orbits around L2 can maintain continuous line-of-sight to both Earth and the lunar far side (useful for communications),

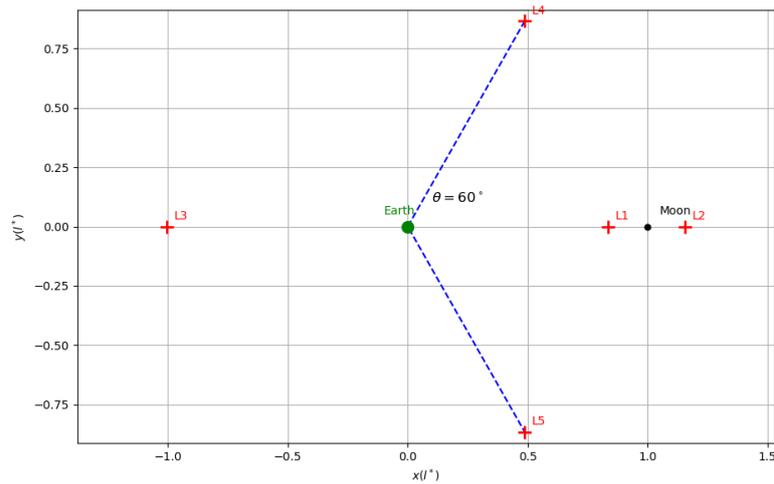


Figure 2.1: Lagrange points in the Earth-Moon system

while orbits around L_1 offer good coverage of the space between Earth and Moon and could host navigation beacons [30].

Most cislunar orbits proposed in literature [56] are generated around the Lagrange points (especially L_1 and L_2). There are databases available, such as JPL's Three Body Periodic Orbit website [25] which provide the parameters (initial conditions, period, Jacobi constant etc.) required to create millions of periodic orbits in the Earth-Moon CR3BP framework. Next, a brief overview description of stability index is given, followed by an overview of the orbits deemed most favorable for cislunar SSA.

Stability index

Another important concept is stability of an orbit. Most periodic orbits in the Earth-Moon CR3BP are unstable, meaning a spacecraft will drift off the orbit without active corrections. Stability can be quantified by a stability index derived from the monodromy matrix [49]. If the stability index is equal to one, the orbit is neutrally stable (meaning small perturbations do not grow or die out), while larger indices indicate exponential divergence and the need for station-keeping. For example, Halo orbits around Lagrange points are typically highly unstable with stability indices ranging from nearly 1 up to 1200, whereas some orbits like distant retrograde orbits can have an index near one (extremely stable).

Periodic Orbit Families in cislunar space

Cislunar periodic orbits can be defined as trajectories in cislunar space which return to their initial state (position and velocity) after a certain period of time [50]. These orbits occur in the vicinity of Lagrange points due to their stability characteristics, and are typically classified by the orbit's shape and which equilibrium point or primary they are associated with [1]. A comprehensive review and details on numerically creating these orbits can be found in [56]. Vendl et al. [43] provide a detailed description of different orbit types such as NRHOs, DROs and Halo orbits, and compare their visibility to the cislunar region by taking into account the observer's motion, the Sun's motion, apparent brightness of a reference object and other observer constraints such as the location and eclipses of the Earth and Moon. An important finding of this work is that observers in a particular orbit must be "phase optimized" to showcase the ideal illumination characteristics for that orbit. It means that an observer placed in a particular position in an orbit at different times will exhibit different illumination characteristics. Wilmer et al. [51] evaluates 14 periodic cislunar orbits to support different missions, and evaluates them using criteria such as Jacobi constant, manifold structures and accessibility from Earth/GEO. A general overview of different orbit families and their relevance to cislunar SSA is provided below. Note that each family can theoretically contain an infinite number of orbits and the orbits described here demonstrate the general evolution of the family.

Lyapunov Orbits

These are retrograde planar periodic orbits around the L1, L2 and L3 Lagrange points. These orbits stay between Earth and Moon (for L1) or around the far side of the Moon (for L2). They can provide good observation capabilities for cislunar space due to their proximity to the region and solar viewing geometry [43]. Lyapunov orbits are typically unstable, but they serve as the planar foundation from which more complex 3D orbits bifurcate.

Halo Orbits

These are 3-dimensional orbits existing around the L1 and L2 points, named such because they appear as "halos" from an Earth inertial frame. In the Earth–Moon system, there are northern and southern Halo families at both L1 and L2 (the designation refers to which side of the plane the orbit spends more time above) [7]. Halo orbits are unstable, often requiring regular station-keeping, but they provide a variety of operational advantages. For example, an L2 Halo orbit can maintain continuous communication with Earth and the lunar far side, and Halo orbits around L1 can serve as gateways for rapid access to the lunar surface or as hubs for navigation satellites [30]. A notable subset of the Halo family is the Near-Rectilinear Halo Orbits (NRHOs), which are Halo orbits of very high eccentricity that pass close to one lunar pole. NRHOs exist as members of the L1 or L2 Halo families and are relatively stable orbits. In fact, some NRHOs have stability index very close to one, meaning a spacecraft can remain in the orbit with minimal propulsive maintenance. This rare stability (unusual for a Halo orbit) is one reason NASA selected an L2 NRHO for the Lunar Gateway outpost [1]. These have also been used in the past in the Chinese Chang'E missions [55]. NRHOs offer continuous coverage of the lunar south pole (important for expeditions there) and good Earth visibility most of the time, while requiring low ΔV for maintenance. Halo orbits provide good observation capabilities due to proximity to key areas of the cislunar space and consistent solar viewing geometry [43].

Distant Retrograde Family (DRO)

These are orbits around the Moon that are retrograde relative to the Moon's orbital motion around Earth. A DRO encircles the Moon at a large radius and is highly stable. The stability index of DROs can be as low as one, making them essentially neutrally stable in the CR3BP. Their orbital size varies a lot, and certain orbits offer good characteristics for both lunar missions and cislunar observation [1]. They are suitable for long duration missions due to their stability, It was used in the Asteroid Redirect Mission (ARM) and The Artemis I mission (2022) [3]. The DRO family can include orbits that go around just the Moon, or loop out to also encircle the L1/L2 points in a broad retrograde loop. Despite their long periods and distance, DROs are valuable for long-term stays because very little station-keeping is needed [51].

L_4 and L_5 axial orbits

This family possesses a wide range of motion that allows access to a large part of cislunar space. Additionally, they have been studied to make high-quality observations of the solar corona and other astrophysical processes, for deep-space maneuvers and for observing cislunar space [1]. They are unstable.

Vertical Orbits

These unstable orbits have oscillations in primarily the z-direction, and typically exist near the L_2 point. They are used to study the transitions between planar and spatial orbits, and can be useful for out-of-plane surveillance which helps in reducing exclusion zones due to celestial body albedo.

Butterfly orbits

These spatial orbits bifurcate from the NRHO and are often symmetric about the Moon. They are unstable, but provide good lunar south pole visibility, which has been the target for several recent and future missions. They are characterized by z-axis oscillation which offers valuable out-of-plane geometry for non-planar surveillance.

Resonant Orbits

These are large orbits characterized by mean-motion resonances with the Moon. They are stable orbits, which can provide great cislunar observation capabilities for an observer constellation.

Touring class of periodic orbits (TCPOs)

These orbits encompass a wide space in the E-M system, and can traverse L_1, L_2, L_3 Lagrange points. They can enable flybys and some out of plane TCPOs can offer good vantage points for surveillance [51].

Orbits for cislunar SSA

Wilmer et al. [47] suggest that Halo orbits and NRHOs near L1/L2 are likely ideal for an SSA sensor constellation due to their proximity to the “Earth–Moon corridor” through which most transfer trajectories pass. By positioning sensors along this corridor (e.g. one at L1 and one at L2), one can in principle monitor traffic coming from Earth or going behind the Moon.

Another attractive option is the Distant Retrograde Orbit (DRO) around the Moon for an observer. As discussed, a DRO is highly stable and encompasses a large volume around the Moon in a retrograde loop. A space-based telescope in a DRO would essentially orbit the Moon at a distance and could continuously monitor objects in lunar vicinity or the transiting region. Eapen et al. [11] explicitly compared sensor tasking performance using an observer in a DRO versus one in an L1 Halo orbit. In their cislunar SSA simulations, they considered two scenarios: Scenario 1 with a space-based sensor in a Distant Retrograde Orbit, and Scenario 2 with a sensor in an L1 Halo orbit (alongside a lunar ground telescope). Both orbits had their pros and cons. The DRO-based sensor is dynamically stable due to the low stability index of the DRO and experiences a wide panoramic view when above the Moon’s far side. It spends half of each orbit high above the lunar far side, which is excellent for seeing cislunar space beyond the Moon. The L1 Halo sensor stays between Earth and Moon, guaranteeing direct Earth line-of-sight and coverage of the trans-Earth injection paths. However, the L1 Halo is less stable (requiring more orbit maintenance). Their work indicates that a combination of observers can yield the best coverage. For example, a lunar far-side DRO sensor plus a lunar pole-based telescope gave complementary views. In general, one finds that no single orbit covers all of cislunar space all the time, and a hybrid architecture (multiple space sensors and perhaps ground sensors on the Moon) is likely needed. This is supported by Paul et al. [35], who develop an optimization framework to choose multiple satellite orbits from various families for a cislunar SSA constellation. In their approach, candidate orbits from several CR3BP families (Halo orbits, both northern and southern, as well as L1, L2, L3 Lyapunov orbits) are evaluated for observational geometry and coverage advantages. Their results suggest that mixing different families (e.g., one spacecraft on a Halo orbit with significant out-of-plane coverage and another on a planar Lyapunov orbit that sweeps a different region) yields more robust coverage. The inclusion of an L3 orbit in their candidate list is notable. An observer orbiting the far side of Earth (around Earth–Moon L3, which is behind Earth) could in theory detect objects coming from deep space before they reach the Earth–Moon system. While Earth–Moon L3 is a distant and somewhat impractical location for a dedicated sensor, it shows that full cislunar domain awareness might even consider such far-flung vantage points.

Frueh et al. [15] identified the Earth–Moon 2:1 resonant orbit family as especially promising. A spacecraft on a 2:1 resonant orbit (completing two circuits for each lunar month) can weave through all corners of cislunar space over its repeating cycle. They demonstrated that ground-based telescopes could observe a satellite on this orbit throughout its path (with only brief gaps), and adding a Moon-based telescope provided only marginal benefit. This suggests that a constellation of spacecraft on resonant or touring orbits might be able to patrol cislunar space effectively, using their own onboard sensors to detect other objects. Resonant orbits, being more extended, have longer periods and would be slower to revisit a particular area compared to a Halo orbit. However, because they are unstable, they naturally drift through different regions, which aids in scanning the whole space.

Finally, it’s worth mentioning the potential role of Lunar ground-based sensors in augmenting cislunar SSA. A telescope on the Moon’s far side, for instance, can scan the sky without Earth’s light pollution and could track objects in Earth–Moon L2 orbits or deep-space trajectories. Eapen et al. [11] considered a Moon-based optical telescope at the lunar north pole in conjunction with space-based sensors. They found that the lunar surface telescope improved tracking of certain orbits, but it has limitations (fixed location, periods of darkness, etc.). Overall, surface sensors alone are not sufficient for full coverage and must be part of a mixed architecture with space-based assets.

To summarize, Halo orbits, Lyapunov orbits, DRO and DPOs are identified as favorable orbits for single-sensor cislunar SSA architectures.

2.2. Sensor Tasking

Space Situational Awareness (SSA) relies on the use of sensor networks, which include ground based radars and optical telescopes, as well as space based sensors, to detect, track and characterize ASOs [11]. Since the number of ASOs is much larger than the number of sensors, these sensors must be effectively *tasked* to maximize the surveillance outcome of these sensors. *Sensor tasking (or scheduling)* refers to the problem of dynamically allocating and orienting these sensors over time to maximize observation effectiveness [54]. This involves assigning sensing actions (e.g., pointing direction, timing, or tracking priorities) to maximize coverage, minimize uncertainty and/or detect objects to ensure continued safety and sustainability of space operations [54]. This is a complex problem due to various constraints. For example, ground based optical telescopes can only observe at night under clear weather, and can only see well illuminated objects above the horizon. On the other hand, radars do not have visibility constraints but are power-limited and therefore typically have shorter range (up to LEO). Space-based sensors avoid atmospheric limitations and have a much wider field-of-regard (FOR) leading to more continuous coverage.

A comprehensive overview of recent sensor-tasking methodologies is provided by Xue et al. [54], who classify them according to objective function design and algorithmic approach. This paper is used as a baseline for this literature review, with specific topics or papers cited in it studied in more detail if relevant to this thesis. This thesis focuses on the design of a myopic (one-step) sensor tasking algorithm for the purpose of catalog maintenance of ASOs in the cislunar region.

Problem Formulation

Sensor tasking is formulated as an optimization of a control vector representing all sensors' pointing or observation actions in a given time interval. At each decision epoch the goal is to choose a set of sensor actions \mathbf{u} that maximizes a defined *objective function* $E(\mathbf{u})$. Mathematically, the optimal sensor schedule \mathbf{u}^* is obtained as [54]:

$$\mathbf{u}^* = \arg \max_{\mathbf{u}} E(\mathbf{u}) \quad (2.1)$$

The objective function $E(\mathbf{u})$ measures the value of a particular sensor action with respect to the overall SSA goal, and hence is maximized to get the most value out of the sensor network. This general formulation supports both single-sensor and multi-sensor frameworks. In practice, $E(\mathbf{u})$ may represent different mission objectives, such as number of targets detected, reduction in orbit uncertainty, or a combination, and is subject to constraints like sensor field-of-view (FOV), target visibility and slewing time. The sensor tasking process can be applied in a planning horizon of one step (*myopic* or *greedy* scheduling for the next observation) or multi-step (planning a sequence of observations) depending on whether long-term outcome needs to be considered [54].

Objective functions

A critical aspect of sensor tasking is the design of appropriate objective functions. The objective function quantitatively represents the goal of the tasking problem, and guides what observations the sensors should prioritize [54]. In the context of SSA, sensor tasking serves two broad purposes: 1) search for uncataloged or new objects, and 2) maintenance of the catalog of known objects. Poorly defined objectives can mislead the scheduling algorithm. For example, if the goal is precise orbit determination, it is reasonable to have an objective function that takes into account the uncertainty in the state of the objects.

In a probabilistic setting, *uncertainty* refers to the spread of the probability distribution that represents our belief about an object's state. This can be quantified in different ways, for example by the covariance matrix in a Gaussian filter or by information-theoretic quantities such as entropy. For a continuous random variable with probability density function $p(x)$, the (differential) Shannon entropy is defined as [52]:

$$H(X) = - \int p(x) \ln p(x) dx, \quad (2.2)$$

and provides a scalar measure of the “spread” or unpredictability of X . In the Gaussian case, the entropy is proportional to $\ln |P|$, where P is the covariance matrix, so reducing $\ln |P|$ corresponds to reducing uncertainty [21].

Search Oriented Objective Functions

Search refers to the process of discovering new space objects and determining their initial orbits when no prior state is known. A classic approach to search is the use of sky *surveys*, where a sensor sweeps a predefined pattern across the search volume to maximize coverage. For optical telescopes, this often involves scanning in declination or right-ascension “stripe” patterns to cover a belt of the sky where targets are likely. Different survey strategies are discussed in [54]. It must be noted that sky surveys do not have an explicit objective function; they can be considered as predetermined tasking actions. *Coverage* and *target density* are the two categories of search objective functions.

Coverage is the tasking of a sensor to maximize coverage of a surveillance area such that no region goes unobserved for too long. In a step-scan strategy using coverage, the search volume is partitioned into sectors, and at each step the sensor points to the sector that has been unobserved for the longest time. This “minimize revisit time” criterion was used by Moretti et al. [32] in a study comparing static scan patterns to dynamic tasking, where choosing the sector with the oldest timestamp ensured regular visits to all regions and avoided redundant observations.

Target density focuses on pointing sensors where the density of potential targets is highest. Unlike pure coverage (which treats all regions equally), a density-driven approach accepts that some objects may go unobserved longer in order to frequently monitor the busiest regions where new detections are most probable. For example, Siminski [40] studied optimizing declination stripes for GEO search by exploiting the known clustering of GEO object orbits around the Laplace plane. By prioritizing such high-density sky bands, a target-density objective yields more detections overall, at the expense of some low-density areas being observed less often.

In practice, search strategies often combine these considerations: e.g. a sensor might regularly cover a full belt (coverage) while spending extra time on known dense clusters (density) to maximize new discoveries.

Catalog Maintenance Objective Functions

After initial detection, catalog maintenance requires subsequent observations of known objects to preserve sufficient accuracy and precision in their estimated states for SSA purposes. To guide this process, objective functions are formulated to capture the specific goals of catalog maintenance, such as controlling uncertainty growth or ensuring track continuity.

Information gain (IG) based objective functions quantify the value of an observation in terms of how much it increases the information available about the state of the system. More formally, IG measures how much an observation changes a prior probability distribution into a posterior one, thereby reducing uncertainty. These measures consider how much an observation will reduce the uncertainty (entropy or variance, depending on the exact formulation) in the target’s state estimate. Bayesian estimation error covariance is commonly used as a measure of uncertainty. Williams et al. [46] incorporated an Extended Kalman Filter (EKF) into the sensor tasking loop and used the expected reduction in orbital state covariance as the reward for scheduling a sensor. Their algorithm scheduled the observation that results in the greatest reduction in the predicted error in satellite state. This coupling of sensor tasking with state estimation is depicted in Fig. 2.2. After obtaining the predicted states $(\pi_{k|k-1})$ for the entire set of ASOs (through a state estimation model), a set of pseudo-observations $(Z_{u_i}^{\text{PIMS}})$ is generated for each candidate sensor action u_i . Using these pseudo-observations a set of hypothetical posteriors $\pi_{k|i}$ is generated, which are compared to the predicted states $\pi_{k|k-1}$ to compute the information gain for each sensing action. The sensing action u^* that maximizes information gain is then picked as the sensor tasking outcome.

There are different mathematical formulations for the information gain. In general, these are built from either *entropy* (as a measure of uncertainty) or from a *divergence* between two probability distributions. A divergence is a non-negative scalar that quantifies how different two distributions are and it is zero only if the distributions are identical, but it is not necessarily symmetric and may not satisfy all properties of a distance metric [52].

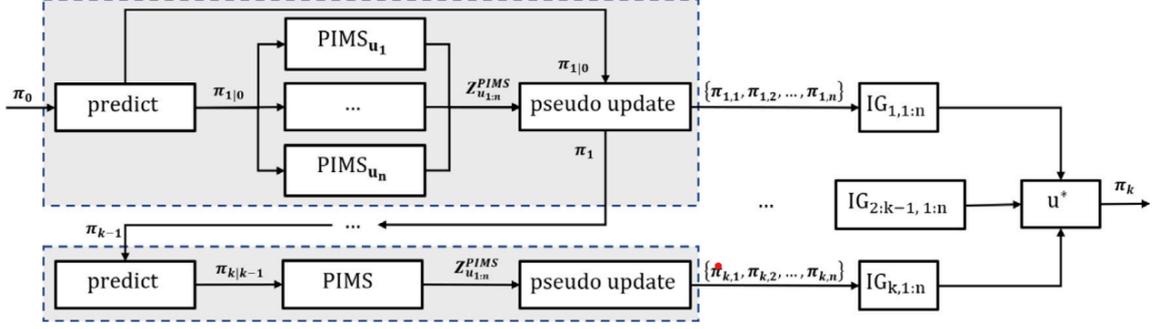


Figure 2.2: Flow chart of IG-based tasking (taken from [54])

Shannon mutual information (MI) is defined as the decrease in entropy of the system (or target state distribution) due to a measurement. Equivalently, it can be written as the expected Kullback–Leibler divergence between the posterior and the prior. In a Gaussian context Shannon entropy is proportional to $\ln |P|$ (the log-determinant of the covariance) [21]. Tasking decisions can be made by predicting which potential observation reduces $\ln |P|$ the most. This mutual information between the predicted measurement and the object’s state has been a popular objective in SSA. For example, Tian and Zhu [27] introduced a Fisher Information Gain (FIG) criterion for sensor scheduling. Fisher information is the inverse of covariance in the linear-Gaussian case. In a linear Gaussian setting, the FIG and MI are mathematically equivalent, and this is also referred to as the D-optimal design [21].

The Cauchy-Schwarz information divergence (CSD) is another commonly used objective function [54]. DeMars et al. [10] proposed a specialized closed form solution of the Cauchy-Schwarz divergence for SSA, and demonstrated its application in a target tracking problem. Since it depicts the difference between two Gaussian distributions, using it as an objective function ensures that we maximize the increase in information between the two states.

Kullback–Leibler Divergence (KLD) is another information-theoretic metric commonly used. It measures the divergence between the prior and posterior state distributions, and is defined as

$$D_{\text{KL}}[p \parallel q] = \int p(x) \ln \left(\frac{p(x)}{q(x)} \right) dx, \quad (2.3)$$

for prior $q(x)$ and posterior $p(x)$. A KLD objective function aims to select observations that make the posterior state estimate as “different” (i.e., more informed) from the prior as possible. It directly measures how much the belief over the target’s state changes due to the observation [19]. Rényi divergence is a generalized form of the KLD which computes the change in the probability distribution with an adjustable order α that controls sensitivity to outliers. One common form is

$$D_{\alpha}[p \parallel q] = \frac{1}{\alpha - 1} \ln \int p(x)^{\alpha} q(x)^{1-\alpha} dx, \quad (2.4)$$

which reduces to the KLD as $\alpha \rightarrow 1$. Principe [36] considered Rényi’s entropy as a more robust information measure where tuning α for a specific mission scenario can improve results.

The consensus in recent literature is that no single information metric is universally best for SSA. FIG (or MI) works well for simpler Gaussian models, CS divergence offers analytic simplicity, and Rényi gain is robust due to parameter tuning. Indeed, a comprehensive quantitative comparison of these measures in an SSA context remains an open research topic. Nevertheless, the trend is to use information-driven objectives because they naturally align sensor tasking with the end-goal of improving the space object catalog’s fidelity. By directly valuing uncertainty reduction, these objectives help avoid schedules that might superficially observe many targets but without significantly improving knowledge (e.g. repeatedly observing an already well-known object is of little information value).

3

Research Questions

3.1. Research Gaps and Relevance to This Work

While SSA research is very advanced for Earth orbits, there is limited understanding of how specific sensor tasking strategies perform under realistic cislunar dynamics and constraints. The overview of orbit families provided in Section 2.1 is compiled from different sources, and the parameters by which favorable orbital geometry or stability characteristics are characterized are often not directly comparable to those from other papers. Different dynamical models (such as CR3BP, ER3BP, full n-body model) used in different works also cause inconsistencies when comparing the results of these works.

Recent work by Eapen et al. [11] evaluates greedy sensor tasking strategies across several cislunar orbits and highlights that identical reward functions can produce varying results depending on orbital placement. However, their study covers only a limited subset of orbit types and does not establish a framework for direct comparison. The interaction between reward functions and orbital geometry—such as whether Near Rectilinear Halo Orbits (NRHOs) support better coverage than Distant Retrograde Orbits (DROs)—remains underexplored.

Moreover, few studies integrate tasking policy design with orbital selection. While algorithmic advances such as Monte Carlo Tree Search (MCTS) [22] and fairness-aware revisit optimization [57] offer sophisticated tasking strategies, they are rarely evaluated alongside the choice of observer orbit or under consistent dynamical conditions.

This thesis addresses these research gaps through a unified approach. By systematically analyzing the coupling between orbital geometry and sensor tasking strategies using a greedy algorithm under CR3BP dynamics, this work provides new insights into the joint optimization of observer placement and tasking logic for future cislunar SSA architectures.

3.2. Research Questions

How can cislunar Space Situational Awareness by a space-based observer in a cislunar orbit be supported through the evaluation of sensor tasking strategies and observer orbits?

The focus of the thesis can be encapsulated by the following sub-questions, which are evaluated in a coupled form:

1. How does the performance achieved by different reward functions for a greedy sensor tasking strategy compare with each other to track objects in cislunar space?
2. How do different observer orbits compare with each other to track objects in cislunar space?

This project simulates the observation of a set of target objects in the cislunar domain by a single space-based observer, evaluating performance across different orbital placements and sensor tasking reward functions. To establish a baseline to build from, sensors will be tasked using a greedy algorithm to

select its action. Different simulation configurations will be compared using the cumulative uncertainty reduction of the tracked targets.

4

Methodology

This chapter presents the methodology used to evaluate space-based cislunar SSA architectures. It first introduces the dynamical framework, formulating the Earth–Moon CR3BP. The state estimation approach is then described, followed by the sensor tasking problem, which is defined through a set of reward functions and scheduling algorithms, together with detailed visibility and brightness constraints for an optical observer. Finally, the generation of orbital data and the end-to-end simulation framework for propagating target/observer trajectories and executing observation schedules are outlined.

4.1. Cislunar Orbital Mechanics

In astrodynamics, a system consisting of two bodies, such as the Earth and a spacecraft, where both can be approximated as point masses, is referred to as a Keplerian or two-body system. In this idealized model, the spacecraft follows a conic section trajectory, which can be circular, elliptical, parabolic, or hyperbolic based on its energy, that is governed by Kepler’s laws. This simplification can be used for a preliminary analysis of orbits around Earth. However, for more accurate orbital propagation, perturbations such as atmospheric drag, solar radiation pressure, Earth’s spherical harmonics, third body forces need to be modeled, which are typically done as accelerations acting on the spacecraft[16][1].

Three-Body Problem

However, for missions operating in the cislunar region, which is the space between Earth and the Moon, the gravitational influence of the Moon becomes significant and cannot be neglected. In such cases, the two-body approximation is no longer sufficient, and a more accurate model is needed. This leads to the **Three-Body Problem** (3BP), where the motion of a spacecraft is influenced by two massive primaries (e.g., Earth and Moon).

The 3BP describes the motion of three point masses. Given masses m_1, m_2, m_3 at positions $\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3 \in \mathbb{R}^3$, their dynamics in an inertial frame can be described as follows: [33]

$$\ddot{\mathbf{r}}_i = G \sum_{\substack{j=1 \\ j \neq i}}^3 m_j \frac{\mathbf{r}_j - \mathbf{r}_i}{\|\mathbf{r}_j - \mathbf{r}_i\|^3}, \quad i \in \{1, 2, 3\}, \quad (4.1)$$

where G is the gravitational constant.

The general 3BP does not have a closed form solution due to its inherent complexity and sensitivity to initial conditions[16][1]. A detailed discussion on multi-body dynamics can be found in [33].

Circular Restricted Three-Body Problem

A widely used simplification of the 3BP is the **Circular Restricted Three-Body Problem (CR3BP)**. This model assumes [16]:

1. Three bodies are present: two primaries (Earth and Moon) and a third body (the spacecraft).

2. The spacecraft's mass is negligible and does not influence the motion of the primaries.
3. All three bodies are treated as point masses.
4. The Earth and Moon revolve in perfect circular orbits around their common barycenter, lying in a fixed plane with constant angular velocity.

In the CR3BP, spacecraft trajectories are no longer simple conic sections but can be highly complex, featuring quasi-periodic, chaotic, or resonant behavior, as shown in Figure 4.1 [16]. The figure shows a non-planar three dimensional view of several orbits in the cislunar space. The CR3BP model allows for the analysis of specific features such as Lagrange points, families of periodic orbits (e.g. Halo, DRO etc), and low-energy transfer pathways. As with the two-body system, the CR3BP is still an approximation and has its limitations. It assumes circular orbits for the primaries and neglects additional perturbations such as solar gravity, non-spherical mass distributions, and solar radiation pressure. These effects are non-negligible in high-precision cislunar navigation and thus require further refinement[16][1].

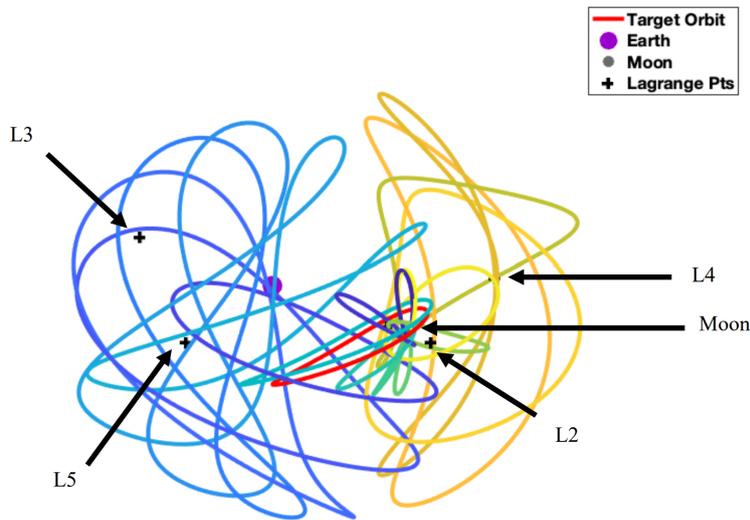


Figure 4.1: A representation of the complex non-planar nature of cislunar orbits (reprinted from [16])

Mathematical Model of CR3BP

The CR3BP has two large primaries (Earth and Moon, or E-M) and an S/C of negligible mass. The primaries orbit around their common barycenter. The equations of motion of the CR3BP are expressed in the Earth-Moon (E-M) rotating frame, also referred to as the **synodic frame** (Figure 4.2), with the origin at the barycenter (discussed in more detail in Section 4.1).

Use of non-dimensional quantities

It is conventional to use non-dimensional quantities when dealing with the CR3BP to reduce numerical complexity [1]. This is done through the use of characteristic length, mass and time.

1. Characteristic length (or length unit) l^* is the average distance between the Earth and Moon,
2. Characteristic mass $m^* = m_E + m_M$ is the sum of Earth and Moon's mass,
3. Characteristic time (or time unit) $t^* \sqrt{\frac{l^*}{m^*}} = \frac{T}{2\pi}$ is defined to ensure the mean motion of the Moon around Earth is equal to unity, where T is the time period of the Moon around the Earth.

where m refers to mass, and the subscripts E and M refer to Earth and Moon respectively. The non-dimensional **mass parameter** (or **mass ratio**) is defined as:

$$\mu = \frac{m_M}{m_E + m_M} \quad (4.2)$$

The non-dimensionalized position of the spacecraft is denoted by $\mathbf{r} = [x, y, z]^T$ and the velocity by $\dot{\mathbf{r}} = [\dot{x}, \dot{y}, \dot{z}]$. Then, the motion of the spacecraft can be described by the following non-dimensional equations:

$$\ddot{x} = 2\dot{y} + \frac{\partial U^*}{\partial x}, \quad \ddot{y} = -2\dot{x} + \frac{\partial U^*}{\partial y}, \quad \ddot{z} = \frac{\partial U^*}{\partial z} \quad (4.3)$$

in which U^* is the pseudo-potential function of the system defined as:

$$U^* = \frac{1 - \mu}{\|\mathbf{r}_{E-s/c}\|} + \frac{\mu}{\|\mathbf{r}_{M-s/c}\|} + \frac{1}{2}(x^2 + y^2) \quad (4.4)$$

where $\mathbf{r}_{E-s/c}$ and $\mathbf{r}_{M-s/c}$ are the position vectors from Earth and Moon to the S/C, respectively.

The **Jacobi constant** (C) is the only constant of motion in the CR3BP. It provides information regarding the energy level associated with a periodic orbit. It is evaluated as:

$$C = 2U^* - (\dot{x}^2 + \dot{y}^2 + \dot{z}^2) \quad (4.5)$$

It is a measure of negative energy, meaning a large Jacobi constant corresponds to less energy. The Jacobi constant can be used to characterize orbits. An orbit with higher C will be a smaller orbit close to the moon, while an orbit with lower C will encompass a larger volume around the Moon.

A detailed derivation of the CR3BP equations of motion can be found in [50].

Stability Index

The stability index of a periodic CR3BP orbit is a metric that quantifies a satellite's ability to stay on its periodic trajectory. As described by [49], it is derived from the eigenvalues of the monodromy matrix, which is the state transition matrix evaluated after one orbital period. These eigenvalues describe how small perturbations to the initial state evolve over time. The stability index, denoted ν , is computed as

$$\nu = \frac{1}{2} \left(\lambda_{\max} + \frac{1}{\lambda_{\max}} \right), \quad (4.6)$$

where λ_{\max} is the dominant eigenvalue of the monodromy matrix. A value of $\nu = 1$ indicates a linearly stable orbit (i.e. small perturbations remain bounded), while $\nu > 1$ implies instability, with a higher stability index meaning higher instability. Specifically, a higher stability index implies that a satellite placed in that orbit will be more sensitive to small perturbations, leading to greater deviations from its trajectory over time. It can be used as a preliminary measure of the amount of station keeping and/or propellant required to keep a satellite on course.

Reference Frames

We define two important reference frames for the E-M system: the synodic reference frame and the inertial reference frame.

The **synodic frame**, represented as \hat{S} , has its origin at the E-M barycenter (the center of mass of the E-M system). The positive x-axis points towards the Moon, the positive z-axis perpendicular to the orbital plane and the positive y-axis completes the right hand system. An important feature of this frame is that the Earth and the Moon remain fixed in this frame, and thus the evolution of the orbit of the spacecraft can be seen relative to the primary bodies. This also helps visualize the orbit of the spacecraft more intuitively.

The **inertial frame**, depicted as \hat{I} , has its origin at the center of the Earth, with its axes represented by $\hat{X}, \hat{Y}, \hat{Z}$. In this frame, the moon revolves around the Earth in a circular Keplerian orbit with an angular velocity of 1 (due to non-dimensionalized quantities), and the spacecraft follows a more complicated trajectory that is more difficult to understand visually. These frames can be visualized in Figure 4.2. (Note that in the first image the origin is actually at the center of the Earth, but is displayed like such for better visualization).

Frame Transformation

Schaub et al. [39] discuss kinematics and transformations which are used to derive the following frame transformations. Note that all the quantities in this section are in the CR3BP non-dimensional form unless explicitly stated otherwise.

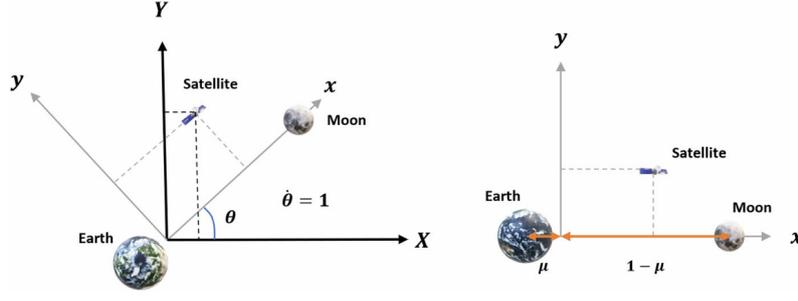


Figure 4.2: Inertial (left) and Synodic (right) reference frame for the CR3BP (taken from [50])

Inertial \rightarrow synodic state transformation

Let $s^I(t) = \begin{bmatrix} \mathbf{r}^I(t) \\ \mathbf{v}^I(t) \end{bmatrix}$ be the spacecraft state in the inertial frame, and let $s_b^I(t) = \begin{bmatrix} \mathbf{r}_b^I(t) \\ \mathbf{v}_b^I(t) \end{bmatrix}$ be the state of a reference body b (e.g. Earth, Moon, or spacecraft) in the inertial frame. The Earth is fixed at the center of the system, with $s_E^I(t) = [0, 0, 0, 0, 0, 0]^T$

Let $\theta(t) = nt$ be the angle of the synodic frame w.r.t. the inertial frame, with mean motion $n = 1$. The rotation matrix about the axis \hat{Z} is:

$$\mathbf{R}_z(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (4.7)$$

and the skew-symmetric cross product matrix of angular velocity Ω (required to apply transport theorem) is:

$$[\Omega \times] = \begin{bmatrix} 0 & -n & 0 \\ n & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}. \quad (4.8)$$

The full 6×6 inertial-to-synodic rotation matrix is

$$\mathcal{T}_{I \rightarrow S}(\theta) = \begin{bmatrix} \mathbf{R}_z(-\theta) & \mathbf{0} \\ -\mathbf{R}_z(-\theta)[\Omega \times] & \mathbf{R}_z(-\theta) \end{bmatrix}, \quad (4.9)$$

so that

$$\mathbf{r}^S = \mathbf{R}_z(-\theta) \mathbf{r}^I, \quad \mathbf{v}^S = \mathbf{R}_z(-\theta) (\mathbf{v}^I - \Omega \times \mathbf{r}^I). \quad (4.10)$$

After adding the translation to the rotation, the synodic state is then

$$\boxed{s^S(t) = \mathcal{T}_{I \rightarrow S}(\theta(t)) s^I(t) + \tau_E} \quad (4.11)$$

where τ_E represents the translation of the origin of the inertial frame from the synodic frame, and is mathematically equal to:

$$\tau_E = [-\mu \ 0 \ 0 \ 0 \ 0 \ 0]^T \quad (4.12)$$

Synodic \rightarrow Inertial state transformation

Similarly, the synodic-to-inertial rotation matrix is

$$\mathcal{T}_{S \rightarrow I}(\theta) = \begin{bmatrix} \mathbf{R}_z(\theta) & \mathbf{0} \\ [\Omega \times] \mathbf{R}_z(\theta) & \mathbf{R}_z(\theta) \end{bmatrix}, \quad (4.13)$$

and the inertial state is obtained by the equation:

$$\boxed{s^I(t) = \mathcal{T}_{S \rightarrow I}(\theta(t)) (s^S(t) - \tau_E)}. \quad (4.14)$$

Moon-fixed frame

A moon-fixed frame \hat{M} can also be defined which is body-fixed at the center of the Moon. Since the Moon is tidally locked with Earth, the moon-fixed frame \hat{M} is a simple translation from the synodic frame given by the following equations:

$$\boxed{\mathbf{s}^M(t) = \mathbf{s}^S(t) - \boldsymbol{\tau}_M.} \quad (4.15)$$

where $\boldsymbol{\tau}_M$ is:

$$\boldsymbol{\tau}_M = [1 - \mu \quad 0 \quad 0 \quad 0 \quad 0 \quad 0]^T \quad (4.16)$$

4.2. State Estimation

In a real-world scenario, the observations of any object provided by a telescope or radar have errors. Additionally, the dynamical model of the system is also an approximation and prone to unmodeled forces and errors. The transformation between the state and measurement space is linearized and may also contain unmodeled elements or noise. An estimator helps fuse the predictions that can be obtained from the known dynamical model of the system with the observations obtained from the sensor to obtain the optimal or near-optimal estimate of the state of an object $\hat{\mathbf{x}}_k$ and its uncertainty \mathbf{P}_k [41].

Sequential Estimation (Extended Kalman Filter)

The specific estimation structure and algorithm implemented is obtained from [41, Sec. 4.7]. The sequential estimation algorithm or the Extended Kalman Filter (EKF) is summarized below.

In this formulation, a “deviation” from the state $\hat{\mathbf{x}}_k$ (hereby referred to as deviation state) and its covariance $\hat{\mathbf{P}}_k$ is updated at each observation epoch t_k based on a propagated reference trajectory $\mathbf{X}^*(t)$ and a linearized measurement model about that reference. The deviation state is mathematically equal to:

$$\hat{\mathbf{x}}_k = \mathbf{X}_{\text{es},k} - \mathbf{X}_k^* \quad (4.17)$$

where $\mathbf{X}_{\text{es},k}$ is the actual estimated state

Propagation of the reference and sensitivities.

Between observation epochs $t_{k-1} \rightarrow t_k$, we integrate the nominal (reference) dynamics as follows:

$$\dot{\mathbf{X}}^*(t) = F(\mathbf{X}^*(t), t) \quad (4.18)$$

and the state-transition matrix (STM) Φ and its derivative $\dot{\Phi}$ as follows:

$$\dot{\Phi}(t, t_{k-1}) = \mathbf{A}(t) \Phi(t, t_{k-1}), \quad \Phi(t_{k-1}, t_{k-1}) = \mathbf{I}, \quad \mathbf{A}(t) \equiv \left. \frac{\partial \mathbf{F}}{\partial \mathbf{X}} \right|_{\mathbf{X}^*(t)}. \quad (4.19)$$

Here, $\mathbf{A}(t)$ is the partial derivative of the function $\mathbf{F}(\mathbf{X}, t)$ at the reference trajectory $\mathbf{X}^*(t)$.

At $t = t_k$ this gives $\mathbf{X}_k^* \equiv \mathbf{X}^*(t_k)$ and $\Phi_{k,k-1} \equiv \Phi(t_k, t_{k-1})$.

A priori time update.

Let $\hat{\mathbf{x}}_{k-1}$ and $\hat{\mathbf{P}}_{k-1}$ be the *a posteriori* estimate and covariance at t_{k-1} . The deviation state and covariance are propagated as:

$$\bar{\mathbf{x}}_k = \Phi_{k,k-1} \hat{\mathbf{x}}_{k-1}, \quad \bar{\mathbf{P}}_k = \Phi_{k,k-1} \hat{\mathbf{P}}_{k-1} \Phi_{k,k-1}^\top + \Gamma_k \mathbf{Q} \Gamma_k^\top. \quad (4.20)$$

(Here $\bar{\mathbf{x}}_k$ denotes the *a priori* predicted deviation and $\bar{\mathbf{P}}_k$ the *a priori* predicted covariance at t_k . Γ_k translates the 3×3 process noise matrix \mathbf{Q} to the 6×6 covariance space.)

$$\Gamma_k = \begin{bmatrix} \frac{1}{2} \Delta t^2 \mathbf{I}_3 \\ \Delta t \mathbf{I}_3 \end{bmatrix}, \quad \Delta t = t_k - t_{k-1}, \quad (4.21)$$

Linearized measurement model at t_k .

Given a nonlinear observation model $\mathbf{Y}_k = \mathbf{G}(\mathbf{X}, t_k) + \varepsilon_k$ with $E[\varepsilon_k] = 0$, $E[\varepsilon_k \varepsilon_j^\top] = \mathbf{R}_k \delta_{kj}$, we define the residual y_k and linearized observation matrix $\tilde{\mathbf{H}}_k$ about the reference:

$$\mathbf{y}_k = \mathbf{Y}_k - \mathbf{G}(\mathbf{X}_k^*, t_k), \quad \tilde{\mathbf{H}}_k \equiv \left. \frac{\partial \mathbf{G}}{\partial \mathbf{X}} \right|_{\mathbf{X}_k^*}. \quad (4.22)$$

(Here \mathbf{Y}_k denotes the actual measurement from the telescope and $\mathbf{G}(\mathbf{X}_k^*, t_k)$ denotes the predicted measurement computed from the reference trajectory at time t_k .)

Innovation, gain, and measurement update.

Define the innovation covariance

$$\mathbf{S}_k = \tilde{\mathbf{H}}_k \bar{\mathbf{P}}_k \tilde{\mathbf{H}}_k^\top + \mathbf{R}_k, \quad (4.23)$$

the Kalman gain

$$\mathbf{K}_k = \bar{\mathbf{P}}_k \tilde{\mathbf{H}}_k^\top \mathbf{S}_k^{-1}, \quad (4.24)$$

and the innovation (prediction residual)

$$\boldsymbol{\beta}_k = \mathbf{y}_k - \tilde{\mathbf{H}}_k \mathbf{x}_k. \quad (4.25)$$

Then the *a posteriori* estimate $\hat{\mathbf{x}}_k$ and covariance $\hat{\mathbf{P}}_k$ are

$$\hat{\mathbf{x}}_k = \bar{\mathbf{x}}_k + \mathbf{K}_k \boldsymbol{\beta}_k, \quad \hat{\mathbf{P}}_k = (\mathbf{I}_6 - \mathbf{K}_k \tilde{\mathbf{H}}_k) \bar{\mathbf{P}}_k (\mathbf{I}_6 - \mathbf{K}_k \tilde{\mathbf{H}}_k)^\top + \mathbf{K}_k \mathbf{R}_k \mathbf{K}_k^\top. \quad (4.26)$$

Equations (4.20) - (4.26) define the recursive sequential algorithm.

Since we are using a deviation state formulation, we need to retrieve the actual estimated state at time t_k , denoted by $\mathbf{x}_{es,k}$. The estimated state is simply the deviation state $\hat{\mathbf{x}}$ added to the reference trajectory \mathbf{X}^* :

$$\mathbf{x}_{es,k} = \hat{\mathbf{x}}_k + \mathbf{X}_k^* \quad (4.27)$$

Extended sequential algorithm (EKF) with adaptive relinearization

To mitigate linearization error, we reset (recenter) the linearization point whenever the estimated error state has converged. Essentially, we update our reference state \mathbf{X}_k^* when the EKF estimate has converged. Mathematically, if $\|\hat{\mathbf{x}}_k\| < T$ (where T is some threshold), we update the reference to the current state estimate and reinitialize the subsequent propagation:

$$\text{if } \|\hat{\mathbf{x}}_k\| < T : \quad \hat{\mathbf{X}}_k = \mathbf{X}_k^* + \hat{\mathbf{x}}_k, \quad \mathbf{X}^*(t_k) \leftarrow \hat{\mathbf{X}}_k, \quad \hat{\mathbf{x}}_k \leftarrow 0. \quad (4.28)$$

Otherwise, the reference is left unchanged. The covariance propagates as in (4.20), and the measurement update uses (4.24). This adaptive relinearization improves convergence when the reference and true states are initially far apart, at the cost of reinitializing the dynamics at epochs where the threshold is triggered.

Algorithm summary (per epoch t_k).

Given $\hat{\mathbf{x}}_{k-1}$, $\hat{\mathbf{P}}_{k-1}$ and the new observation $(\mathbf{Y}_k, \mathbf{R}_k)$:

1. Integrate $\mathbf{X}^*(t)$ and $\Phi(t, t_{k-1})$ from t_{k-1} to t_k to obtain \mathbf{X}_k^* and Φ_k .
2. Time update: $\bar{\mathbf{x}}_k = \Phi_k \hat{\mathbf{x}}_{k-1}$, $\bar{\mathbf{P}}_k = \Phi_k \hat{\mathbf{P}}_{k-1} \Phi_k^\top + \Gamma_k \mathbf{Q} \Gamma_k^\top$.
3. Linearize measurement: $\mathbf{y}_k = \mathbf{Y}_k - \mathbf{G}(\mathbf{X}_k^*, t_k)$, $\tilde{\mathbf{H}}_k = \partial \mathbf{G} / \partial \mathbf{X} |_{\mathbf{X}_k^*}$.
4. Form $\mathbf{S}_k = \tilde{\mathbf{H}}_k \bar{\mathbf{P}}_k \tilde{\mathbf{H}}_k^\top + \mathbf{R}_k$ and $\mathbf{K}_k = \bar{\mathbf{P}}_k \tilde{\mathbf{H}}_k^\top \mathbf{S}_k^{-1}$.
5. Measurement update: $\hat{\mathbf{x}}_k = \bar{\mathbf{x}}_k + \mathbf{K}_k (\mathbf{y}_k - \tilde{\mathbf{H}}_k \bar{\mathbf{x}}_k)$, $\hat{\mathbf{P}}_k = (\mathbf{I} - \mathbf{K}_k \tilde{\mathbf{H}}_k) \bar{\mathbf{P}}_k (\mathbf{I} - \mathbf{K}_k \tilde{\mathbf{H}}_k)^\top + \mathbf{K}_k \mathbf{R}_k \mathbf{K}_k^\top$.
6. Retrieve estimated state: $\mathbf{x}_{es,k} = \hat{\mathbf{x}}_k + \mathbf{X}_k^*$

7. Adaptive relinearization: if $\|\hat{\mathbf{x}}_k\| > 1$

$$\hat{\mathbf{X}}_k = \mathbf{X}_k^* + \hat{\mathbf{x}}_k, \quad \mathbf{X}^*(t_k) \leftarrow \hat{\mathbf{X}}_k, \quad \hat{\mathbf{x}}_k \leftarrow \mathbf{0},$$

8. Proceed to t_{k+1} .

Note that formulations of the EKF can differ slightly depending on the source. In this thesis, EKF refers to the formulation provided in this section, taken from [41, Sec. 4.7].

4.3. Measurement Model

We model each observation from an observer (either Earth-based or space-based) as a pair of topocentric angles (right ascension and declination) formed from the line-of-sight (LOS) vector expressed in the inertial frame. The following mathematical model is obtained from [53] and [31].

Let $\mathbf{r}_t^I(t)$ be the target (object to be observed) position and $\mathbf{r}_o^I(t)$ the observer position, both in the ECI frame. The LOS vector from the observer to the spacecraft is denoted as r_{o-t} . We define the topocentric vector and auxiliary scalars as follows:

$$\mathbf{r}_{o-t} \equiv \mathbf{r}_t^I - \mathbf{r}_o^I = \begin{bmatrix} X_{o-t} \\ Y_{o-t} \\ Z_{o-t} \end{bmatrix}, \quad \rho_{o-t} = \|\mathbf{r}_{o-t}\| = \sqrt{X_{o-t}^2 + Y_{o-t}^2 + Z_{o-t}^2}, \quad \rho_{o-t,xy} = \sqrt{X_{o-t}^2 + Y_{o-t}^2}. \quad (4.29)$$

The measurement function (angles-only) is

$$\mathbf{h}(\mathbf{x}) = \begin{bmatrix} \alpha \\ \delta \end{bmatrix} = \begin{bmatrix} \text{atan2}(Y_{o-t}, X_{o-t}) \\ \arcsin\left(\frac{Z_{o-t}}{\rho_{o-t}}\right) \end{bmatrix}, \quad (4.30)$$

where α is the topocentric right ascension (RA) and δ the topocentric declination (Dec), both in radians. The state \mathbf{x} is $[\mathbf{r}^I \quad \mathbf{v}^I]^\top$.

Jacobian (measurement sensitivity)

The 2×6 measurement matrix is

$$\mathbf{H} = \frac{\partial \mathbf{h}}{\partial \mathbf{x}} = \begin{bmatrix} \frac{\partial \alpha}{\partial \mathbf{r}_{o-t}} & \mathbf{0}_{1 \times 3} \\ \frac{\partial \delta}{\partial \mathbf{r}_{o-t}} & \mathbf{0}_{1 \times 3} \end{bmatrix}, \quad (4.31)$$

with the position partials (using ρ_{o-t} and $\rho_{o-t,xy}$ defined above):

$$\frac{\partial \alpha}{\partial \mathbf{r}_{o-t}} = \begin{bmatrix} -\frac{Y_{o-t}}{X_{o-t}^2 + Y_{o-t}^2} & \frac{X_{o-t}}{X_{o-t}^2 + Y_{o-t}^2} & 0 \end{bmatrix}, \quad (4.32)$$

$$\frac{\partial \delta}{\partial \mathbf{r}_{o-t}} = \begin{bmatrix} -\frac{X_{o-t}Z_{o-t}}{\rho_{o-t}^2 \rho_{o-t,xy}} & -\frac{Y_{o-t}Z_{o-t}}{\rho_{o-t}^2 \rho_{o-t,xy}} & \frac{\rho_{o-t,xy}}{\rho_{o-t}^2} \end{bmatrix}. \quad (4.33)$$

Measurement noise and covariance.

Each angle is perturbed by zero-mean noise:

$$\mathbf{y}_k = \mathbf{h}(\mathbf{x}_k) + \boldsymbol{\varepsilon}_k, \quad E[\boldsymbol{\varepsilon}_k] = \mathbf{0}, \quad \text{cov}(\boldsymbol{\varepsilon}_k) = \mathbf{R}_k = \begin{bmatrix} \sigma_{\alpha,k}^2 & 0 \\ 0 & \sigma_{\delta,k}^2 \end{bmatrix}, \quad (4.34)$$

with $\sigma_{\alpha,k}$ and $\sigma_{\delta,k}$ specified in radians.

4.4. Scheduling Observations

Problem setup.

Let $\mathcal{T} = \{1, \dots, N\}$ be the catalog of N objects and t_k the discrete decision epochs. At each epoch t_k , the sensor may observe at most one target from the feasible set $\mathcal{F}_k \subseteq \mathcal{T}$ defined by hard constraints (visibility, exclusion, occlusion etc.). For each target i , we denote the EKF *a priori* state and covariance by $(\bar{\mathbf{x}}_{i,k}, \bar{\mathbf{P}}_{i,k})$. $\tilde{\mathbf{H}}_{i,k}$ is the measurement matrix $\mathbf{R}_{i,k}$ is the observer noise matrix at time t_k [41].

Greedy (myopic) selection rule.

We choose at each t_k the single target that maximizes a one-step reward (hereby referred to as the “best” target) [44]:

$$i_k^* \in \arg \max_{i \in \mathcal{F}_k} J_{i,k} \quad (4.35)$$

Note that the reward is only evaluated for targets that are visible, i.e. targets that satisfy the visibility criteria (discussed in Sec. 4.6). The formulation of different reward functions are discussed in the next section (Sec. 4.5).

Filter updates at the selected target.

Once i_k^* is chosen and a measurement is taken, the EKF performs the standard measurement update on $(\bar{\mathbf{x}}_{i_k^*,k}, \bar{\mathbf{P}}_{i_k^*,k})$: [41]

$$\mathbf{S}_{i,k} = \tilde{\mathbf{H}}_{i,k} \bar{\mathbf{P}}_{i,k} \tilde{\mathbf{H}}_{i,k}^\top + \mathbf{R}_{i,k}, \quad (4.36)$$

$$\mathbf{K}_{i,k} = \bar{\mathbf{P}}_{i,k} \tilde{\mathbf{H}}_{i,k}^\top \mathbf{S}_{i,k}^{-1}, \quad (4.37)$$

$$\hat{\mathbf{x}}_{i,k} = \bar{\mathbf{x}}_{i,k} + \mathbf{K}_{i,k} (\mathbf{y}_{i,k} - \mathbf{h}(\bar{\mathbf{x}}_{i,k})), \quad (4.38)$$

$$\hat{\mathbf{P}}_{i,k} = (\mathbf{I} - \mathbf{K}_{i,k} \tilde{\mathbf{H}}_{i,k}) \bar{\mathbf{P}}_{i,k} (\mathbf{I} - \mathbf{K}_{i,k} \tilde{\mathbf{H}}_{i,k})^\top + \mathbf{K}_{i,k} \mathbf{R}_{i,k} \mathbf{K}_{i,k}^\top, \quad (4.39)$$

Prediction for unobserved targets (a priori propagation).

For every $j \neq i_k^*$, as well as for targets that are not visible to the observer, no measurement is incorporated at t_k ; only the EKF *time update* (Eq. 4.20) is applied:

$$\hat{\mathbf{x}}_{j,k} = \bar{\mathbf{x}}_{j,k}, \quad \hat{\mathbf{P}}_{j,k} = \bar{\mathbf{P}}_{j,k}, \quad (4.40)$$

and then all targets are propagated to the next epoch with the state-transition matrix $\Phi(\cdot)$ and process noise \mathbf{Q}_k :

$$\bar{\mathbf{x}}_{i,k+1} = \Phi_k \hat{\mathbf{x}}_{i,k}, \quad \bar{\mathbf{P}}_{i,k+1} = \Phi_k \hat{\mathbf{P}}_{i,k} \Phi_k^\top + \Gamma_k \mathbf{Q}_k \Gamma_k^\top, \quad \forall i \in \mathcal{T}. \quad (4.41)$$

In regard to the EKF algorithm in Sec 4.2, this means that for all the unobserved targets ($j \neq i_k^*$), steps 3 to 5 and step 7 are skipped.

Note that in this thesis only one target is observed at once. The observation of other targets in the sensor’s FOV is not within the scope of this thesis.

4.5. Reward Functions

Selecting which observation is “best” depends on the objective of the tasking problem. We quantify this objective through a *reward function* that maps each candidate measurement to a scalar quantity or *reward* [54]. The action with the highest reward is selected. The aim of this thesis is to evaluate and compare different reward functions that are described here.

At each simulation time step we obtain the *a priori* (predicted) covariance $\bar{\mathbf{P}}$ from the time-update equation (4.20) and the *a posteriori* (updated) covariance $\hat{\mathbf{P}}$ from the measurement-update equation (4.26). It can be seen from the equations in Section 4.2, evaluating $\hat{\mathbf{P}}$ for a *candidate* observation does not require the actual or *realized* measurement value. In the linearized Gaussian setting it depends only on the known or assumed measurement model and noise statistics (e.g., $\tilde{\mathbf{H}}$ and \mathbf{R}) together with the *a priori* covariance $\bar{\mathbf{P}}$. Since covariance quantifies the uncertainty in the estimate of the state, the change from $\bar{\mathbf{P}}$ to $\hat{\mathbf{P}}$ can be considered a measure of how much a prospective observation would reduce the uncertainty in the state.

The first class of rewards is *information-theoretic*: the goal is to choose the candidate whose measurement results in the maximum information gain about a system, or equivalently one which would lead to a maximum reduction in uncertainty of a system. Information theory quantifies uncertainty via *entropy*; higher entropy corresponds to greater uncertainty. The *information gain* (IG) is the reduction in entropy produced by an observation, and in our Bayesian setting it corresponds to the expected decrease from prior to posterior uncertainty. Several measures are proposed to quantify the information gain, such as Shannon Mutual Information, Kullback-Leibler Divergence, Fisher Information Gain, Cauchy-Schwarz Divergence etc. [21][54]

In the present estimation structure (Sec. 4.2), we use a deviation state EKF, so both the *a priori* and *a posteriori* error states have zero mean (they represent deviations from a reference trajectory). This simplifies several reward expressions that would otherwise include mean state terms.

The mathematical expressions for different candidates for IG, in the context of our specific problem, is presented below.

Kullback-Leibler (KL) Divergence

The Kullback-Leibler (KL) divergence measures the separation between two distributions. It is a commonly used reward in sensor management algorithms [54][21][19]. The KL divergence between two Gaussian distributions p and q is defined as:

$$D_{\text{KL}}(p||q) = \int p(x) \log \frac{p(x)}{q(x)} dx. \quad (4.42)$$

The scalar KL divergence reward for our zero mean state can be expressed in terms of the *a priori* and *a posteriori* covariances as follows:

$$r_{\text{KL}} = D_{\text{KL}}(p||q) = \frac{1}{2} \left[\text{tr}((\bar{\mathbf{P}})^{-1} \hat{\mathbf{P}}) - n + \log \frac{\det \bar{\mathbf{P}}}{\det \hat{\mathbf{P}}} \right]. \quad (4.43)$$

Fisher Information Gain and Shannon Mutual Information

Fisher Information Gain (FIG) quantifies the increase in estimation precision due to an action as measured by the Fisher information matrix (the inverse covariance in Gaussian models). In the linear Gaussian case, the information matrix update is [54]:

$$\mathbf{J}^+ = (\hat{\mathbf{P}})^{-1}, \mathbf{J}^- = (\bar{\mathbf{P}})^{-1}. \quad (4.44)$$

A scalar reward based on the Fisher Information Gain can be formulated as follows[21]:

$$r_{\text{FIG}} = \frac{1}{2} \log \frac{\det \mathbf{J}^+}{\det \mathbf{J}^-} \quad (4.45)$$

Similarly, Shannon Mutual Information (MI) quantifies the expected reduction in differential entropy due to a measurement. In the linear Gaussian case, Shannon Mutual Information is equivalent to the aforementioned FIG reward [8][37][21]. For the rest of this work, MI is used to refer to this reward formulation.

Cauchy-Schwarz (CS) Divergence

The Cauchy-Schwarz (CS) divergence measures the angular separation between two densities. It is defined as follows[23][26]:

$$D_{\text{CS}}(p||q) = -\log \left(\frac{\int p(x) q(x) dx}{\sqrt{\int p(x)^2 dx \int q(x)^2 dx}} \right). \quad (4.46)$$

The scalar CS divergence reward for our zero mean state can be expressed in terms of the *a priori* and *a posteriori* covariances as follows:

$$r_{\text{CS}} = D_{\text{CS}}(p||q) = \frac{1}{2} \log \det(\bar{\mathbf{P}} + \hat{\mathbf{P}}) - \frac{1}{4} (\log \det \bar{\mathbf{P}} + \log \det \hat{\mathbf{P}}) - \frac{n}{2} \log 2. \quad (4.47)$$

where n is the dimension of the state ($n = 6$ in our case).

Age-of-Information (time since last observed)

A rather simple way to schedule observations is to select the target that has been unobserved the longest. Age of information (AOI) is a simple heuristic reward function that is equal to the time since the last observation of a target.

Let t_{last} be the last time a target was observed.

$$r_{\text{AOI}} = t_k - t_{\text{last}}, \quad (4.48)$$

Finite-Time Lyapunov Exponent (FTLE)

A finite-time Lyapunov exponent (FTLE) based reward function incorporates dynamical instability into the sensor tasking strategy. The objective is to prioritize observations of targets whose state uncertainty is expected to grow the most. This is done by propagating each target's reference trajectory \mathbf{X}^* and STM Φ over a short prediction horizon H and evaluating the projected covariance growth.

For each target i at time t_k , the CR3BP variational equations are propagated forward over H future steps, producing a horizon STM $\Phi_k^{(H)}$

$$\Phi_k^{(H)} = \Phi(t_k + H\Delta t, t_k). \quad (4.49)$$

Neglecting process-noise accumulation over the horizon period, the predicted covariance at the horizon $\bar{\mathbf{P}}_k^{(H)}$ is calculated as follows:

$$\bar{\mathbf{P}}_k^{(H)} = \Phi_k^{(H)} \bar{\mathbf{P}}_k \Phi_k^{(H)\top} \quad (4.50)$$

The scalar reward can be set to the maximum eigenvalue λ_{\max} of the predicted covariance at the horizon $\bar{\mathbf{P}}_k^{(H)}$

$$r_{\text{FTLE}} = \lambda_{\max}(\bar{\mathbf{P}}_k^{(H)}) \quad (4.51)$$

which quantifies how large the target's state uncertainty is expected to become in case of lack of measurement updates.

4.6. Visibility Constraints for a Space-Based Optical Sensor

We impose four visibility constraints before scheduling an observation: (i) *minimum brightness* or *limiting magnitude* of the target as seen by the observer, and *exclusion angles* from (ii) the Sun, (iii) the Moon, and (iv) the Earth.

Minimum Brightness (Limiting Magnitude)

In order for an optical sensor to detect a target, it must receive a sufficient amount of light from the target. We assume that the targets do not emit any light of their own, so two important factors in determining if it is visible is the amount of light from the Sun reflected by it (it is also assumed that the Sun is the only body illuminating the target), and the distance to the observer. Cognion [6] discuss a mathematical model to quantify the magnitude of light received by the observer. This model simplifies the problem by focusing on the geometry of the observer, the target, and the Sun, and neglects or makes assumptions to simplify other factors such as the sensor's optical parameters and background sources. This model is stated in Equation 4.52.

Assuming that the target is a spherical Lambertian object with reflectivity C_d and radius R , its apparent visual magnitude m_{object} can be calculated as follows:

$$m_{\text{object}} = m_{\text{sun}} - 2.5 \log_{10}(F_{\text{diff}}(\psi)), \quad (4.52)$$

$$F_{\text{diff}}(\psi) = \frac{C_d R^2}{\pi d^2} \frac{2}{3} (\sin(\psi) + (\pi - \psi) \cos(\psi)), \quad (4.53)$$

where m_{sun} is the apparent visual magnitude of the sun (constant value), $F_{\text{diff}}(\psi)$ is incident solar flux reflected by the spherical satellite as a function of the phase angle. d is the distance of the object from the observer, and the angle ψ is the *phase angle* between the sun, the object and the observer [4][43].

The threshold below which the object is visible is termed as the *limiting magnitude* and is determined by the electro-optical sensor. For this thesis, a limiting magnitude of 20 will be used, which corresponds to high-end space telescope [43]. The object radius $R = 1$ m, the reflection coefficient $C_d = 0.5$ and the Sun's reference magnitude $m_{\text{sun}} = -26.74$ [6].

For an object to be visible, its observed magnitude m_{object} should be *less than* the limiting magnitude.

Phase Angle

The phase angle is the geometric angle between the Sun, the object and the observer (Fig.4.3). It can be calculated as follows:

Let $\mathbf{r}_o, \mathbf{r}_t, \mathbf{r}_S, \mathbf{r}_E, \mathbf{r}_M$ denote the inertial positions of the observer, target, Sun, Earth, and Moon. Define the line-of-sight (LOS) vector \mathbf{r}_{o-t} and range ρ as:

$$\mathbf{r}_{o-t} = \mathbf{r}_t - \mathbf{r}_o, \quad \rho = \|\mathbf{r}_{o-t}\|, \quad \hat{\mathbf{r}}_{o-t} = \frac{\mathbf{r}_{o-t}}{\rho}.$$

The *phase angle* (Sun–target–observer) is

$$\alpha = \arccos\left(\frac{\mathbf{r}_S - \mathbf{r}_t}{\|\mathbf{r}_S - \mathbf{r}_t\|} \cdot -\hat{\mathbf{r}}_{o-t}\right), \quad (4.54)$$

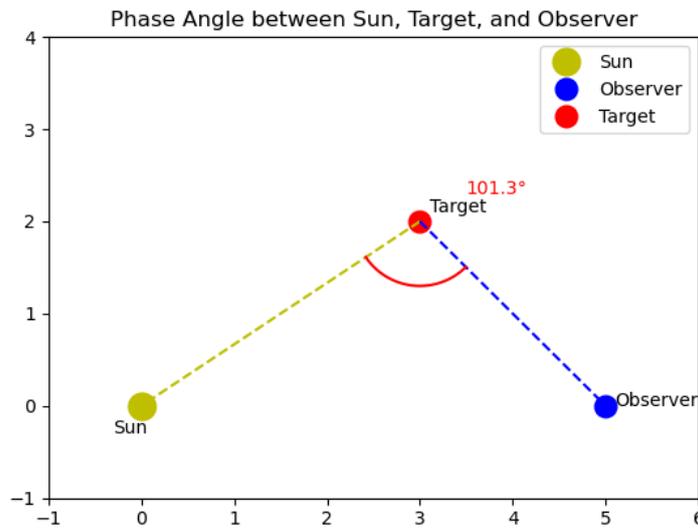


Figure 4.3: Representation of solar phase angle

For a more accurate minimum brightness evaluation, a probability of detection formula is used which takes into account several other parameters such as the sensor's optical parameters, celestial background sources and detector noises [38]. Using the limiting magnitude as a hard limiting constraint neglects the optics aspect of the problem while also keeping it computationally simple. The only variables are the target-observer distance d and the phase angle ψ , since the radius and reflectivity of each object is set to be constant.

Solar, Lunar, and Earth body exclusion

The body exclusion criterion ensures that the telescope's line of sight does not intersect any major celestial body (Sun, Earth, or Moon). This covers the case when the target is directly in front of or behind a much larger celestial body. Typically, the exclusion angle for each celestial body is taken because pointing the sensor in the direction of the bright object can cause damage. However, that is not taken into account in this thesis.

Let the sightline unit vectors for each body $B \in \{S, E, M\}$ (Sun, Earth, and Moon, respectively) be defined as:

$$\hat{\mathbf{u}}_B = \frac{\mathbf{r}_B - \mathbf{r}_o}{\|\mathbf{r}_B - \mathbf{r}_o\|},$$

where \mathbf{r}_o and \mathbf{r}_B denote the position vectors of the observer and body B , respectively. The instantaneous angular separation between the line of sight to the target and the direction to each body is given by

$$\beta_B = \arccos(\hat{\mathbf{r}}_{o-t} \cdot \hat{\mathbf{u}}_B), \quad (4.55)$$

where $\hat{\mathbf{r}}_{o-t}$ is the unit vector from the observer to the target.

The exclusion condition is:

$$\beta_B \geq \Theta_B, \quad B \in \{S, E, M\}, \quad (4.56)$$

where the half-angle Θ_B represents the apparent angular radius of the body as viewed from the observer, which can be determined mathematically as follows::

$$\Theta_B = \arcsin\left(\frac{R_B}{\|\mathbf{r}_B - \mathbf{r}_o\|}\right), \quad (4.57)$$

with R_B being the physical radius of the body B , the values of which are listed in Table 4.1.

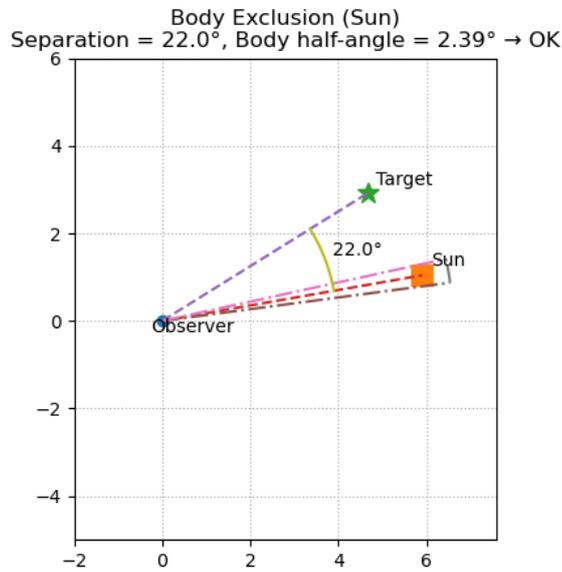


Figure 4.4: Representation of the solar exclusion angle

Parameter	Value	Units
Earth's Radius	6378.137	km
Moon's Radius	1737.1	km
Sun's Radius	695 700	km

Table 4.1: Radius of celestial bodies used in calculation of exclusion angles

Summary of visibility constraints

At each decision epoch, a candidate observation is feasible iff all of the following hold:

Limiting Magnitude: $m_{object} < 20$ (Eq. (4.52));
 Exclusions: $\beta_B \geq \Theta_B, B \in \{S, E, M\}$ (Eq. (4.55),(4.57)).

4.7. Generation of Orbital Data

For the target sensing simulation, it is required to place targets and observers in orbits in the cislunar space. This section discusses the approach taken to generate these orbits.

The open-source Python library TUDAT [42] is used for the simulations in this thesis. To simulate the cislunar environment and then propagate orbit trajectories, Earth and Moon are defined in TUDAT using the non-dimensional mass ratio μ . The Earth is defined as a fixed body at the center of the system, and the Moon revolves around the Earth in a circular Keplerian orbit with semi-major axis $a = 1$, with all other Keplerian elements set to zero. Consistent with the CR3BP dynamics, the forces on the spacecraft from Earth and Moon are defined as point mass forces. An orbital trajectory can be propagated in this TUDAT environment by inputting its non-dimensional initial state. The propagated states are then converted to dimensional units using the length unit and time unit (Table 4.2).

To generate orbital data, the initial states are obtained from the JPL Three-Body Periodic Orbits database [25]. The JPL database provides a catalog of orbits that are periodic in the E-M CR3BP framework. It has several hundred thousands of orbits that are categorized by the periodic orbital family and the Lagrange point they exist around. Each orbit is characterized by its initial state in the synodic frame, its time period, its Jacobi constant and its stability index. Given that the initial orbital states are derived from the JPL periodic-orbit database, the CR3BP constants adopted in the TUDAT simulations in this work are chosen to be identical to the values used in the database's orbit generation, as documented on the database website. These are listed in Table 4.2.

Mass Ratio (μ)	Length Unit [km]	Time Unit [s]
$1.215\,058\,560\,962\,404 \times 10^{-2}$	389 703	382 981

Table 4.2: Non-dimensionalization constants for the Earth–Moon CR3BP.

Once the catalog of target and observer orbits was selected from the JPL database, the corresponding initial states were supplied to the TUDAT propagator. Each orbit was then propagated for a duration of two synodic periods using a timestep of 600 seconds. The first synodic period of the propagated trajectory was discarded, and only the second period was retained for subsequent analysis. This approach is necessary because the periodic orbits in the JPL database are generated by setting the initial y-position component and the x-velocity component in the synodic frame to zero. If the sensor tasking simulation began immediately, the orbits would still be adjusting from this artificial alignment, and their relative positions would not reflect a realistic spatial alignment of satellites in the cislunar space. Propagating the orbits for one synodic period lets the trajectories settle into their normal periodic behavior, resulting in a more realistic spatial distribution of the targets and observers, and adds randomness in the states of the observers and targets. Again, the orbital state data for the first half of the propagation is discarded, leaving the orbital state data for the second half of the propagation which is one synodic period long. This 'true' orbital state data is used for the simulation with the sensor tasking algorithm employing the EKF and target scheduling. The duration of the sensor tasking algorithm simulation is one synodic period with a timestep of 600 seconds.

The stability index listed in the JPL database has a direct effect on the propagation of the orbital trajectories. As discussed in Section 4.1, a larger stability index implies that a satellite on that orbit will be more sensitive to perturbations, and could diverge from its periodic trajectory. All the orbits listed in the database are exactly periodic, but since the orbits are generated numerically and not analytically, certain mathematical techniques and constraints are utilized to ensure that these orbits are exactly periodic [14]. Since only the initial states of the orbits are propagated numerically in TUDAT, small numerical errors arise that cause the orbital trajectory to deviate from its periodic trajectory. The higher the stability index of the orbit, the more it deviates from its periodic trajectory. For example, Figure 4.5 shows the evolution of a DRO with stability index $\nu = 1.0$ and time period $T \approx 13.7$ days for a duration of 10 synodic periods (1 synodic period ≈ 29.5 days). It remains perfectly periodic during this long simulation period. Similarly, Figure 4.6 shows the evolution of a Halo orbit with $\nu \approx 86.6$ and time period $T \approx 14$ days for a duration of two synodic periods. The orbit remains close to its intended trajectory initially; however, when it is propagated for three synodic periods (Figure 4.7), it deviates significantly, eventually departing entirely from the expected path.

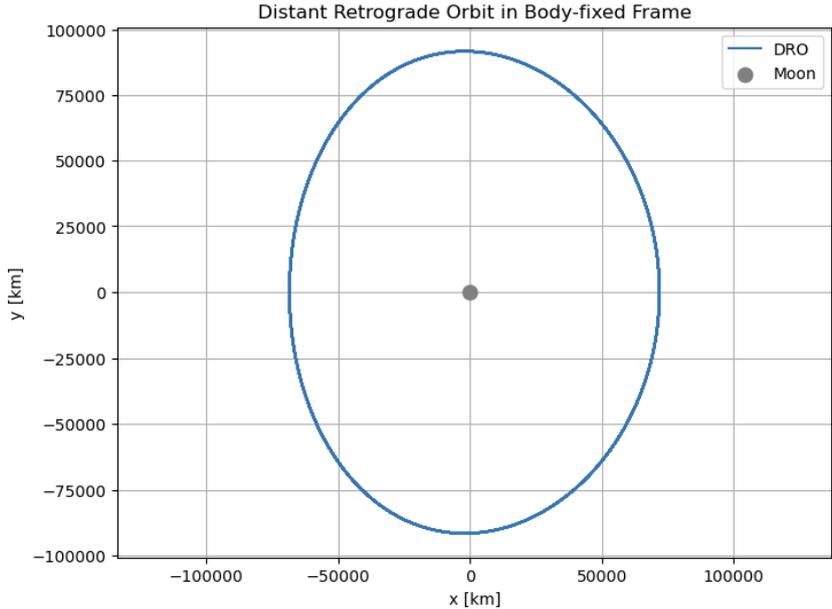


Figure 4.5: Evolution of a very stable DRO for 10 synodic periods

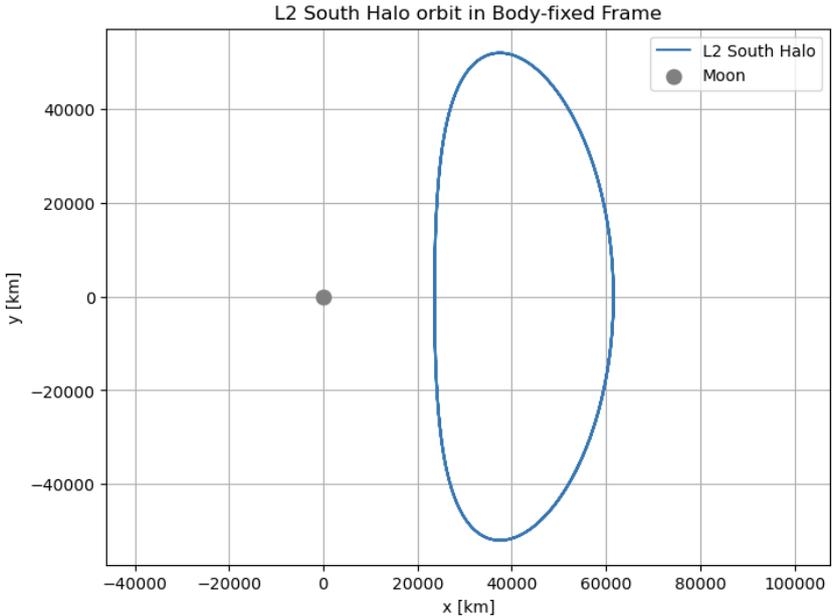


Figure 4.6: Evolution of an unstable Halo orbit for 2 synodic periods

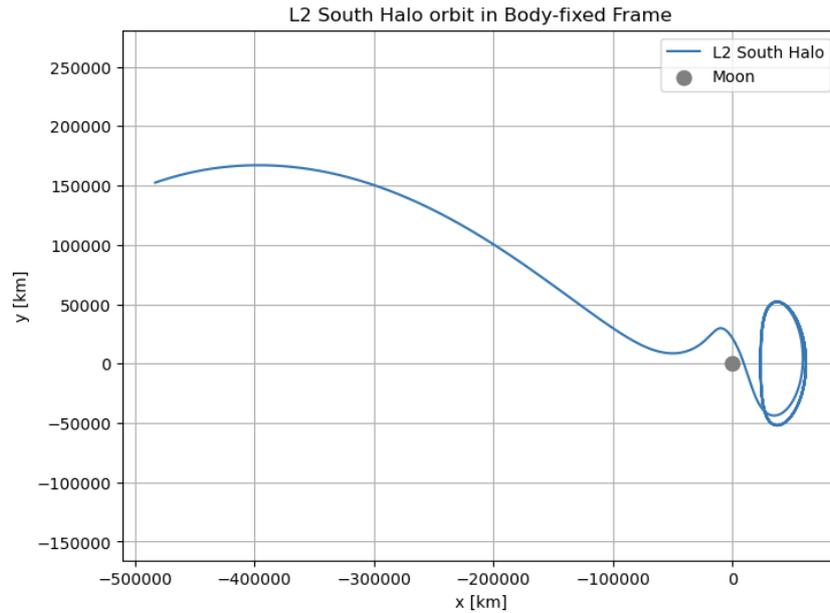


Figure 4.7: Evolution of an unstable Halo orbit for 3 synodic periods

However, since this work propagates the initial state of each orbit for a duration of two synodic periods only, most orbits remain on their intended periodic trajectory for this duration. For example, Figure 4.8 shows the evolution of a Halo orbit with $\nu \approx 599$ for a duration of two synodic periods. It can be seen that it remains on course for the majority of the duration of the simulation. This is the orbit with the highest stability index in the orbital catalog (both observers and targets), and it remains periodic for the majority of the simulation duration. Hence, most orbits remain stable for the duration of the simulation.

Note that the field of cislunar SSA aims to track and catalog all objects in the region, including debris and objects that do not remain on periodic cislunar trajectories. A strongly unstable orbit such as that shown in Figure 4.7 was initially included in trial simulations. When this orbit was added to the catalog, it drifted far from the vicinity of the Moon and the observer, and became effectively unobservable to the sensor for the remainder of the simulation. As a result, the associated state uncertainty grew without bound and began to dominate the overall performance metrics, thereby skewing the results. As a result, any orbits that would deviate highly from its periodic trajectory was removed from the catalog for this study.

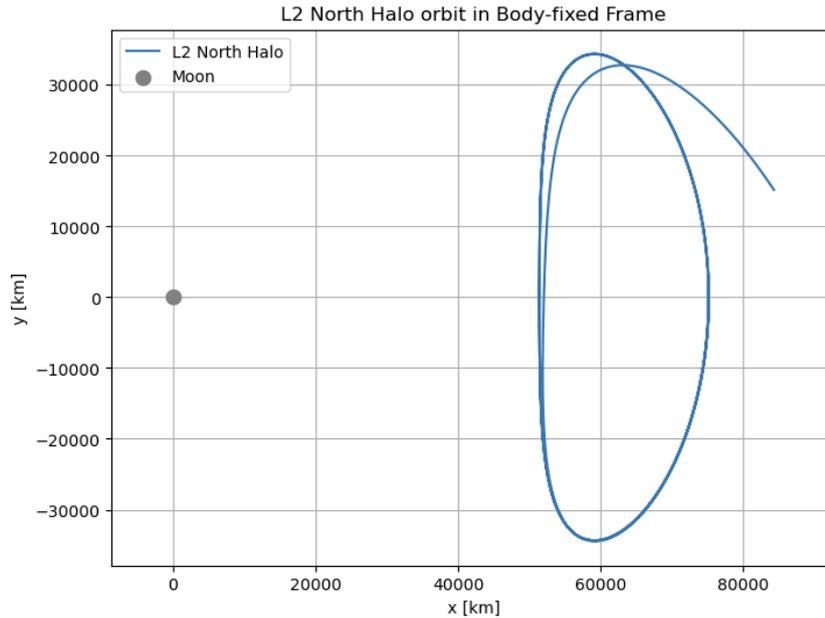


Figure 4.8: Evolution of an highly unstable Halo orbit for 2 synodic periods

Selection of Target and Observer Orbits

Observer Orbits

A study of different observer orbits for cislunar surveillance is an important part of this thesis. The selection of this orbit is a very important factor in enhancing cislunar coverage. Vendl and Holzinger [43] conducted a study comparing several observers in different periodic cislunar orbits and concluded that the most favorable observer orbits (in terms of providing persistent visibility of targets in the cislunar region) are in the L1 and L2 Lyapunov families, L1 and L2 Halo families and in the Distant Prograde Orbit (DPO) families. Based on this, **eight observer orbits** belonging to these orbital families were selected for this work. The time period, stability index, Jacobi constant and orbit family for these orbits are given in Table 4.3.

Index	Orbit family	Time period [days]	Stability index ν	Jacobi constant C
1	L1 Lyapunov	28.94	53.67	2.91
2	L2 Lyapunov	25.44	59.09	2.95
3	DPO	5.33	1.00	3.19
4	DPO	0.21	1.00	4.34
5	L1 North Halo	9.62	1.00	3.00
6	L1 South Halo	10.01	1.29	3.00
7	L2 South Halo	7.17	1.52	3.04
8	L2 North Halo	7.17	1.52	3.04

Table 4.3: Catalog of orbits chosen for observer placement

Target orbits

For comparing the performance of different sensor tasking reward functions as well as different observer orbits, it is necessary to have a consistent catalog of target satellites to track and catalog. For this thesis, a catalog of **22 target orbits** is generated. Since there is little to no publicly available data of precise orbit placement for cislunar missions, the orbits in this catalog were selected from research that proposes favorable orbits for different cislunar missions. For example, [45] discusses the NRHOs as favorable orbits for cislunar missions due to their proximity to and visibility of the lunar surface and stability. In [11] 21 targets in L2 Northern Halo orbits, which are categorized by their Jacobi constant and time period, are part of a cislunar SSA study. Furthermore, to get realistic targets based on past

missions, the proposed orbits of Artemis 1 and the Chang'e relay satellite were obtained from [3] and [55], respectively. The orbits used in this work were initialized from a combination of the proposed orbits from the aforementioned papers, and a catalog of 22 target orbits was finalized, as listed in Table 4.4. Orbits 1 and 2 correspond to those of the Artemis 1 and Chang'e relay mission, respectively.

Index	Orbit family	Time period [days]	Stability index ν	Jacobi constant C
1	DRO	13.65	1.00	2.93
2	L2 South NRHO	6.53	1.24	3.05
3	L2 South Halo	14.00	86.60	3.07
4	L2 South NRHO	7.00	1.45	3.04
5	L2 North NRHO	7.00	1.45	3.04
6	L1 North NRHO	9.85	1.00	3.00
7	L1 South NRHO	9.85	1.00	3.00
8	L1 South NRHO	8.44	2.08	3.00
9	L1 North NRHO	8.44	2.08	3.00
10	L2 North Halo	15.13	598.66	3.15
11	L2 North Halo	14.58	211.14	3.10
12	L2 North Halo	14.22	117.82	3.08
13	L2 North Halo	8.32	1.68	3.03
14	L2 North Halo	10.93	2.17	3.02
15	L1 Lyapunov	27.84	54.71	2.92
16	L2 Lyapunov	27.86	51.81	2.94
17	DRO	27.85	1.00	2.09
18	DRO	0.28	1.00	4.06
19	DRO	14.88	1.00	2.92
20	L2 South Halo	12.00	8.66	3.02
21	LPO	6.54	1.00	3.18
22	DPO	3.48	1.00	3.23

Table 4.4: Catalog of target orbits

4.8. Simulation Framework

This section describes the general simulation framework used in obtaining the results in the next chapter with regard to how the simulation is set up in Python.

Setup of environment

For the propagation of the true states of the observer and target as well as the reference trajectory in the EKF, a CR3BP environment is created with the Earth at the center and the Moon and Sun revolving around it in a Keplerian orbit. The forces from the Sun are not used in the propagation of the states of the orbits; its relative position is used to evaluate the visibility criteria in Section 4.6. The three bodies are in different planes with the Sun's orbit having an inclination of 5.135 degrees for more realistic spatial orientation of the three bodies. However, both orbits are perfectly circular for simplicity and due to the CR3BP conditions. The Moon's rotation is tidally locked with the Earth, implying that its revolution period is equal to its rotation period. Table 4.5 lists the parameters used in the simulation for the setup of the system of bodies in TUDAT.

Propagation of satellite states

The initial states of the selected targets and observers are propagated under the CR3BP conditions listed in 4.1 with the system characterized by the values listed in 4.2 for a duration of two lunar synodic periods, and then the first half of the data is discarded.

Scheduling and observation

The duration of simulation of the sensor tasking scenario is equal to one synodic period. The timestep of both the EKF and the observer schedule is 10 minutes. This means an observation occurs every 10

Parameter	Value	Units
Gravitational parameter of Earth	$1 - \mu$	n.d.
Gravitational parameter of Moon	μ	n.d.
Gravitational parameter of Sun	328899.46	n.d.
Moon's semi-major axis	1	n.d.
Sun's semi-major axis	383.877	n.d.
Moon inclination	0	deg
Sun inclination	5.145	deg

Table 4.5: Parameters used to create CR3BP system in TUDAT. μ refers to CR3BP mass parameter

Algorithm 1 Generation of orbital data

- 1: **Setup of CR3BP environment**
 - 2: Create TUDAT environment with Earth at origin.
 - 3: Model Moon and Sun on circular Keplerian orbits around Earth
 - 4: Use system parameters as in Table 4.5
 - 5: Define Moon rotation tidally locked with Earth

 - 6: **Propagation of true orbital data**
 - 7: **for** each observer $o \in \mathcal{O}$ **do**
 - 8: Propagate true state of observer o under Earth and Moon point mass forces for two lunar synodic periods
 - 9: **end for**
 - 10: **for** each target $i \in \mathcal{T}$ **do**
 - 11: Propagate true state of target i under Earth and Moon point mass forces for two lunar synodic periods
 - 12: **end for**
 - 13: Discard first synodic period of all propagated trajectories
 - 14: Keep second synodic period as “true” state history for simulation
-

minutes and that the EKF propagates the reference trajectory for a period of ten minutes.

Simulation duration (days)	Simulation time-step (minutes)
29.530 589	10

Table 4.6: Simulation time control

At each time step, the reference state $\mathbf{X}_{i,k}^*$ is obtained for each target i in the CR3BP environment created earlier. If the position in the state $\mathbf{X}_{i,k}^*$ passes the visibility checks described in Section 4.6, its *a posteriori* covariance is calculated using the equations in Section 4.2. Then, depending on the reward being evaluated in the simulation (Section 4.5), the “best target” is scheduled for observation. Depending on whether a particular target is scheduled for observation, its states are updated according to the equations 4.36 - 4.39.

The process noise and measurement noise for the EKF are kept constant across all the targets throughout the simulation. Their values are listed in Table 4.7.

The procedure described above outlines the overall simulation framework used throughout this thesis. In subsequent analyses, the observer setup and reward function are adjusted to explore different scenarios and to generate the results discussed in the next chapter.

Note on units Generation of the orbital state data of the observer and target orbits is done using the non-dimensionalized CR3BP model. The CR3BP environment that is setup in TUDAT is done with non-dimensionalized parameters, with Earth, Moon and Sun being defined as point masses. Once the

Matrix	Value
$Q \in \mathbb{R}^{3 \times 3}$	$Q = q_{\text{pos}} I_3 = \begin{bmatrix} 1.0 \times 10^{-15} & 0 & 0 \\ 0 & 1.0 \times 10^{-15} & 0 \\ 0 & 0 & 1.0 \times 10^{-15} \end{bmatrix} \text{ km}^2$
$R \in \mathbb{R}^{2 \times 2}$	$R = \text{diag}(\sigma_\alpha^2, \sigma_\delta^2) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \text{ arcsec}^2$

Table 4.7: EKF noise covariance matrices

orbital time history is generated, the time components are converted to seconds, the position components are converted to km and the velocity components are converted to km/s. This serves as the ‘true’ states in the sensor tasking and estimation algorithm. In the EKF state propagation step (eq. 4.18), the reference trajectory \mathbf{X}_k^* is non-dimensionalized through the use of the time unit and length unit in Table 4.2 and propagated to the next time step, before being converted back to dimensional units for the measurement update step.

Algorithm 2 Sensor Tasking Simulation Framework

Require: Selected reward function $r(\cdot)$ (MI, DCS,KLD, AOI, FTLE)**Require:** Selected observer

```

1: Main sensor tasking loop
2: Set  $N \leftarrow T_{\text{sim}}/\Delta t$ 
3: for  $k = 1$  to  $N$  do
4:    $t_k \leftarrow t_0 + k\Delta t$ 

5:   for each target  $i \in \mathcal{T}$  do
6:     Propagate reference dynamics and STM from  $t_{k-1}$  to  $t_k$ 
7:     Compute predicted deviation state  $\bar{\mathbf{x}}_k$  and covariance  $\bar{\mathbf{P}}_k$ 
8:     Check visibility of target  $i$  from observer
9:     if any visibility constraint is violated then
10:      Set reward  $r_{i,k} \leftarrow -\infty$  (target not observable)
11:      continue to next target
12:     else
13:      Compute estimated deviation state  $\hat{\mathbf{x}}_k$  and covariance  $\hat{\mathbf{P}}_k$ 
14:      Compute reward  $r_{i,k}$ 
15:     end if
16:   end for

17:   Select index of best target (maximum reward)  $i_k^*$ 
18:   for each target  $i \in \mathcal{T}$  do
19:     if  $i = i_k^*$  then ▷ Target is observed at  $t_k$ 
20:       Estimated state  $\mathbf{X}_{\text{es},i,k} = \mathbf{X}_{i,k}^* + \hat{\mathbf{x}}_{i,k}$ 
21:       Update last observation time  $t_{i,\text{last}}$ 
22:     else ▷ Target is not observed at  $t_k$ 
23:       Predicted state  $\mathbf{X}_{\text{pred},i,k} = \mathbf{X}_{i,k}^* + \bar{\mathbf{x}}_{i,k}$ 
24:       Set estimated state and covariance to predicted state and covariance (for next epoch)
        $\hat{\mathbf{x}}_{i,k} \leftarrow \bar{\mathbf{x}}_{i,k}, \hat{\mathbf{P}}_{i,k} \leftarrow \bar{\mathbf{P}}_{i,k}$ 
25:     end if
26:   end for
27: end for

```

5

Simulation verification

The complete simulation framework integrates several components such as the CR3BP state propagation, EKF, observation scheduling, and implementation of the visibility constraints, as discussed in Chapter 4. To ensure that the results produced by this complete system are credible and scientifically reliable, each element was subjected to a series of verification and validation tests. The purpose of this chapter is to present these tests and demonstrate that the underlying algorithms are functioning as intended.

5.1. CR3BP state propagation

The orbit propagation by the TUDAT propagator is verified by comparing the orbital data obtained through the TUDAT propagation to the orbital data downloaded from the JPL website.

The norm of the position error, defined as the difference between the JPL and TUDAT position, is shown in Figure 5.1. The error is only calculated for the length of the orbital period of Target 1, which is about 14 days, since that is the duration of the orbital data available on the JPL website. The root mean square error in position during this time is 3.45×10^{-7} km. This value effectively verifies the propagation of the orbital states done by the TUDAT propagator used in the simulations in this work.

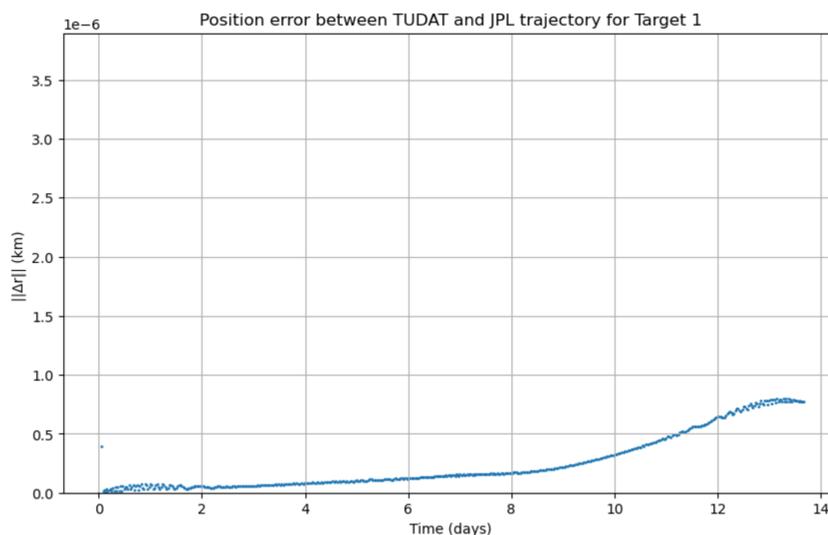


Figure 5.1: The norm of the error in TUDAT's propagated position in reference to the position ephemeris from JPL database of Target 1

Additionally, this figure also shows that TUDAT trajectory diverges over time due to numerical propagation errors. Compare this with Figure 5.2, which shows the position error for Target 10, which has

a stability index of about 599. Since its stability index is larger, it diverges more quickly from the JPL trajectory. This supports the discussion on the stability index of the periodic orbits. However, as shown in Section 4.7 the divergence of the trajectory is small for the duration of the simulations conducted in this work.

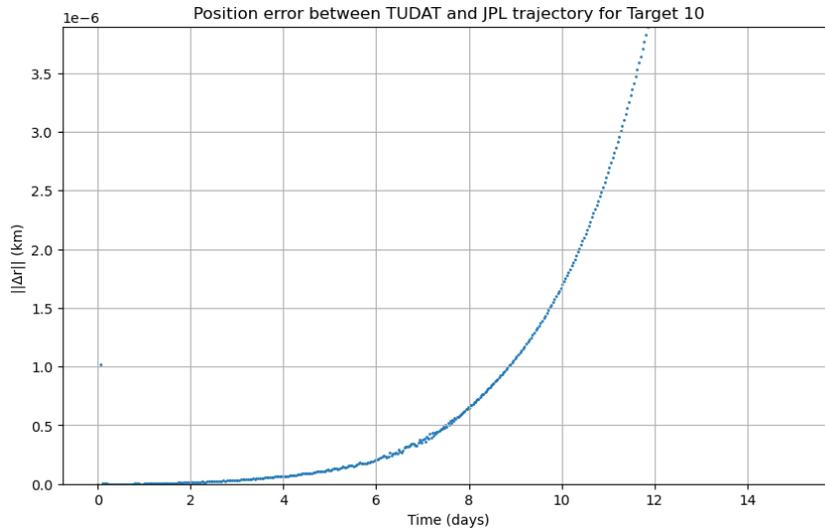


Figure 5.2: The norm of the error in TUDAT's propagated position in reference to the position ephemeris from JPL database of Target 10

5.2. Extended Kalman Filter

The working of the EKF is integral to the tasking algorithm. To verify the functioning of the EKF, the simulation is conducted with only one target. This section shows the results obtained from the configuration of Observer 3 from Table 4.3, and Target 9 from Table 4.4. The observer-target combination was chosen such that the target is visible to the observer at all times to isolate the focus of analysis on the correct implementation of the EKF. Since there is only one target which is visible to the observer at all times, the scheduler selects this target for observation at all of the epochs. Several plots and outputs from the simulation, presented here, are analyzed and discussed to verify the implementation of the EKF.

Figure 5.3 shows the post-fit measurement residuals from the EKF ε_k , calculated as follows:

$$\varepsilon_k = \mathbf{y}_k - \tilde{\mathbf{H}}_k \hat{\mathbf{x}}_k \quad (5.1)$$

where \mathbf{y}_k is the difference between actual measurement and predicted measurement, $\tilde{\mathbf{H}}_k$ is the linearized measurement matrix (Equation 4.22) and $\hat{\mathbf{x}}_k$ is the *a posteriori* state estimate. The horizontal red lines correspond to the standard deviation in the data. The figure shows that the residual data points are symmetric about zero, with most of the data points within the one standard deviation value. Additionally, the precise values of some statistical values corresponding to the post-fit residual data are shown in Table 5.1. The magnitude of the majority of the measurement residuals are within the expected range of ± 1 arcsec, which is the standard deviation of the measurement noise (tab. 4.7). The values combined with the figure shows that the post-fit measurement residuals are zero-mean Gaussian. This shows that the EKF is consistent and works effectively to reduce the differences between the true and computed measurements from the estimated orbit [5].

Metric	RA	DEC
Mean (arcsec)	0.00	0.00
Standard deviation (arcsec)	0.77	0.82
Within 3σ (%)	99.58	99.72

Table 5.1: Residual Statistics

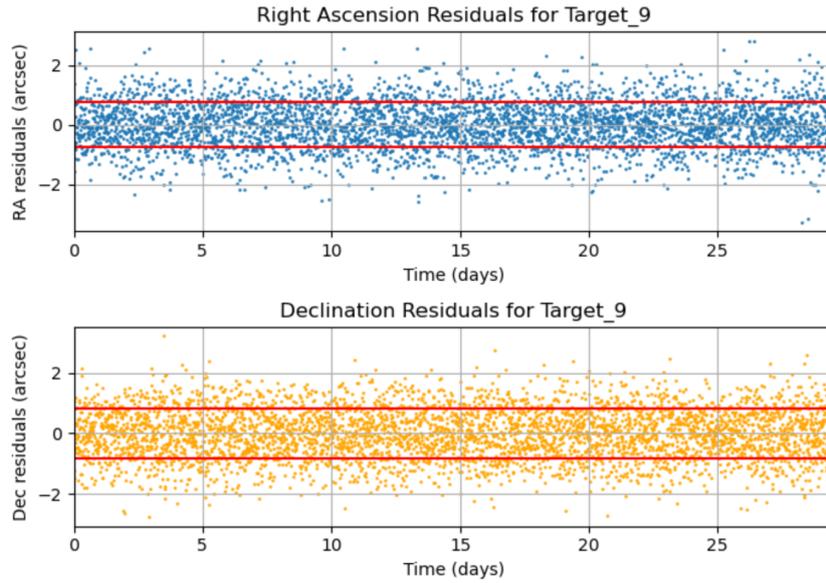
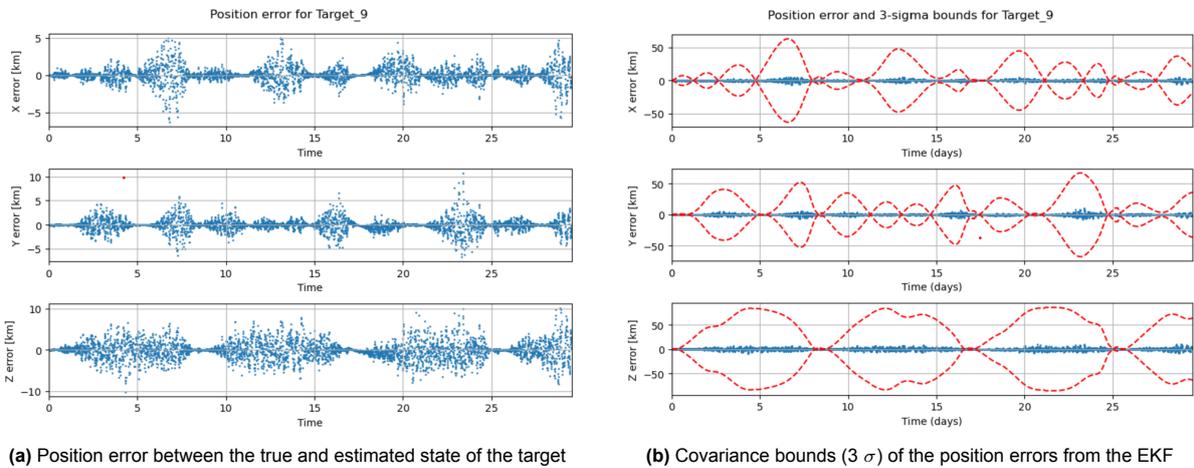


Figure 5.3: Measurement Residuals of the EKF

Figure 5.4a shows the position errors between the estimated position and true position of the target, and Figure 5.4b shows the $3\text{-}\sigma$ bounds of the errors. It can be seen that the errors for all three axes remain well within the $3\text{-}\sigma$ bounds, and that there is no unbounded growth in the error or its covariance. This shows that the filter does not diverge. The covariance boundary expands and contracts as the position error expands and contracts, which shows that the EKF covariance reflects the actual errors. However, the $3\text{-}\sigma$ appear much larger than the errors, which may indicate possible underconfidence in the filter [5][18].



(a) Position error between the true and estimated state of the target (b) Covariance bounds (3σ) of the position errors from the EKF

Figure 5.4: Position Estimate Error and the 3σ covariance boundary

The position errors of the EKF show an oscillatory trend, where they expand and contract over time. This trend can be explained by the relative geometry of the target and the observer. Butcher et al. [5] discuss the observability angle Ψ_f which provides a measure of how rapidly the LOS sweeps the target. Ψ_f characterizes the instantaneous observability of the angles-only relative navigation problem. A large Ψ_f corresponds to strong geometric information about the relative state, while $\Psi_f \approx 0$ indicates near-radial motion and poor observability [5][18].

Figure 5.5 shows the time history of the observability angle Ψ_f computed as follows:

$$\Psi_f = \left| \sin^{-1} \left(\frac{a_f^\top (\boldsymbol{\rho} \times \mathbf{v}_{\text{rel}})}{\|a_f\| \|\boldsymbol{\rho} \times \mathbf{v}_{\text{rel}}\|} \right) \right|, \quad (5.2)$$

where $\boldsymbol{\rho}$ is the observer-to-target relative position, \mathbf{v}_{rel} is the relative velocity, and a_f is the full-nonlinear vector defined in [5] as follows:

$$\mathbf{a}_f = \boldsymbol{\omega} \times \left(\boldsymbol{\omega} \times \mathbf{r}_o + \frac{\mu}{\|\mathbf{r}_t\|^3} \mathbf{r}_o \right) + \frac{3\mu}{\|\mathbf{r}_t\|^5} (\mathbf{r}_t \mathbf{r}_t^\top) \boldsymbol{\rho} \quad (5.3)$$

where \mathbf{r}_o and \mathbf{r}_t denote the inertial positions of the observer and target, respectively, and the observer's orbital angular velocity is $\boldsymbol{\omega} = [0 \ 0 \ n]^\top$, with skew-symmetric matrix $\tilde{\boldsymbol{\omega}}$ satisfying $\tilde{\boldsymbol{\omega}}\mathbf{x} = \boldsymbol{\omega} \times \mathbf{x}$.

Comparing Figure 5.5 with Figure 5.4, a correlation can be seen. The peaks of observability angle generally coincide with very low position errors and covariance, while the troughs in observability angle coincide with high position errors and covariance. This correlation is strong in the X and Y axes but very weak in the Z axis. This is a characteristic of angles-only observation where the covariance grows larger during periods of unfavorable geometry[5][18].

However, the observability angle alone is not enough to explain the oscillatory nature of the errors and covariance. Figure 5.6 shows the range of the target from the observer. Due to the complex CR3BP dynamics, the distance between the observer and the target vary during the duration of the simulation. In an angles-only observation scenario, the position error is directly influenced by the observer-target range because small angular uncertainties translate into larger position uncertainties as distance increases. A measurement noise of magnitude $\delta\theta$ in right ascension or declination corresponds to an approximate transverse position uncertainty of $\rho\delta\theta$, where ρ is the line-of-sight range [5][18]. Thus, for the same angular measurement accuracy, a target located farther away produces proportionally larger errors in the reconstructed position.

The trends in the error and σ bounds in the Z axis in Figure 5.4 very closely follows the range of the target from the observer in Figure 5.6. On the other hand, the trends in the X and Y axis follow the observability angle. For example, peaks in the observability angle at ≈ 5 and ≈ 21 days correspond to low error covariance.

The oscillatory behavior of the EKF errors is thus driven by two factors: (i) instantaneous geometric observability, quantified by Ψ_f which primarily governs the in-plane (X, Y) errors; and (ii) range-driven scaling of angular noise, which is particularly evident in the out-of-plane (Z) component. When the target is far and the observability angle is small, even modest angular noise translates into large position uncertainty, and the EKF covariance inflates accordingly. When the target is closer and Ψ_f is large, the same angular accuracy produces much tighter position estimates, and the covariance rapidly contracts. The combined evolution of $\Psi_f(t)$ and $\rho(t)$ therefore explains not only the instantaneous correlation between geometry and error, but also the quasi-periodic expansion and contraction of the EKF error during the simulation.

Figure 5.7 shows the normalized innovation squared (NIS) across the simulation. The NIS provides a chi-square distributed metric of whether the filter's predicted innovation covariance, S_k , correctly represents the uncertainty in the actual measurement residuals [2]. For an angles-only measurement, the NIS should follow a χ_2^2 distribution, with individual samples lying predominantly within the corresponding confidence bounds and an expected mean value of approximately 2 [2].

NIS Metric	Value
Mean	1.72
Outside 99% confidence bounds (%)	2.52

Table 5.2: NIS Statistics

For this simulation, 2.52 % of the NIS values lie outside the 99 % bounds, which is greater than the expected 1 %. Additionally, the mean is 1.72, which is slightly lower than the expected mean of 2. This

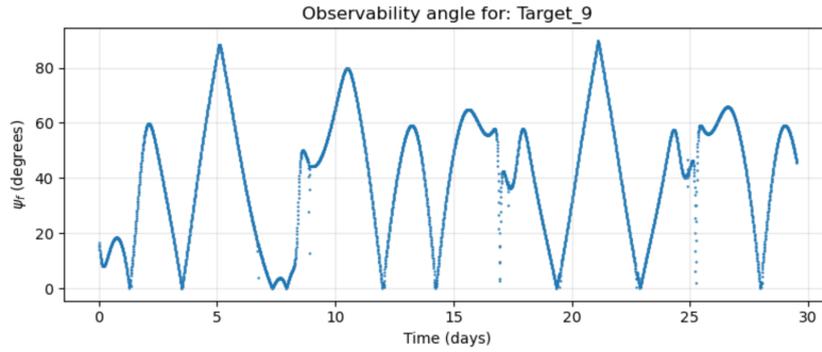


Figure 5.5: Observability angle of the target from the observer

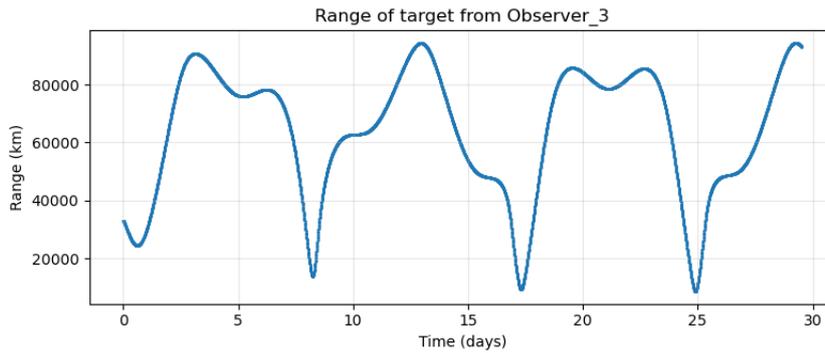


Figure 5.6: Range of the target from the observer

suggests that the filter may be underconfident, consistent with the large $3 - \sigma$ bounds. However, Bar-Shalom et al. [2] note that small deviations from ideal chi-square behavior are expected in nonlinear sequential estimation and do not necessarily indicate filter malfunction. Overall, the NIS analysis shows that the filter may be slightly underconfident. However, since filter tuning is not the main focus of this thesis this underconfidence is deemed to be small enough to proceed with sensor tasking.

5.3. Sensor Tasking Logic

For the sensor tasking algorithm to operate correctly, several behaviors are expected during the scheduling process. Target selection should occur only when the target lies within the sensor's visibility region, and no scheduling should take place when all targets are outside this region. At each observation epoch, the algorithm is required to assign at most a single target, and scheduling observations between simultaneously visible targets should follow the maximum value of the selected reward function. These expectations form the basis for assessing whether the implementation behaves as intended.

To verify the implementation of the sensor tasking logic, a small subset of target orbits were analyzed. A configuration of Observer 3 from Table 4.3, and Targets 15, 16, 17 from Table 4.4 were selected. The targets were selected such that each of them had periods of non-visibility from the observer. For the sensor tasking reward, KL Divergence is used in this simulation.

Figure 5.8 shows the targets that are scheduled for observation by the sensor tasking algorithm. It can be seen that there are periods of time where none of the targets are scheduled for observation. This should correspond to times where none of the targets are visible to the observer and is verified by Figure 5.9 which shows the epochs at which the targets are visible. It can be seen that the times at which none of the targets are tasked correspond with the times at which none of the targets are visible. This verifies the sensor tasking logic. A target is eligible to be scheduled only when it is visible. Additionally, the algorithm also appropriately handles the case where none of the targets are visible.

Figure 5.9 also overlays the times at which the targets are scheduled on its visibility epochs. It may appear that multiple targets are being scheduled for observation simultaneously. However, this is be-

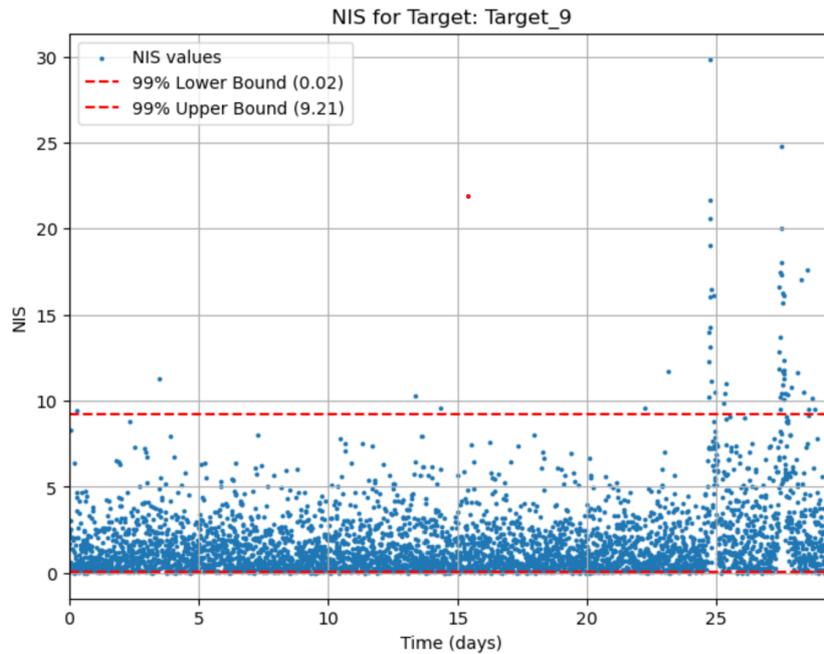


Figure 5.7: Normalized Innovation Squared (NIS)

cause the figure compresses approximately 4,200 observation epochs into a relatively narrow width, causing individual events to overlap visually. This can be verified by only plotting a small subset of the duration of the simulation during which all three targets are visible, as seen in Figure 5.10.

It is worth noting that the algorithm rapidly switches between the available targets, which implies that the reward (KL Divergence) for these targets are close to each other. This causes the maximum reward to constantly switch between the targets, leading to the algorithm toggling between these targets. The absence of any penalties imposed on switching or slewing also facilitates this behavior. To better understand this, consider Figure 5.11, which is a subset of Figure 5.9 but at a time window when Target 15 becomes visible again after a sustained period of non-visibility from the observer. It can be seen that for seven consecutive epochs the algorithm schedules Target 15. After a long period of not being observed, the covariance of the target has grown considerably and hence needs multiple consecutive observations to shrink to a magnitude similar to Target 16, which has been regularly observed while Target 15 was not visible. Due to this large covariance, each observation predicts a high increase in information about the system, of which KL divergence is a scalar measure. Once the covariance has reduced, the observer resumes to rapidly toggle between the two available targets. In the absence of penalties related to switching targets and because of the use of an IG-based reward, the rapid toggling between available targets is an expected behavior from the sensor tasking algorithm, and indicates that the reward function is implemented correctly.

To summarize, Figures 5.8 - 5.11 show the following:

1. a target is only scheduled when it is visible,
2. if none of the targets are visible, no target is scheduled,
3. only one target is scheduled at a single epoch,
4. rapid toggling between available targets aligns with expected behavior of selected reward function function.

Overall, these plots and analysis serve as verification of the sensor tasking algorithm.

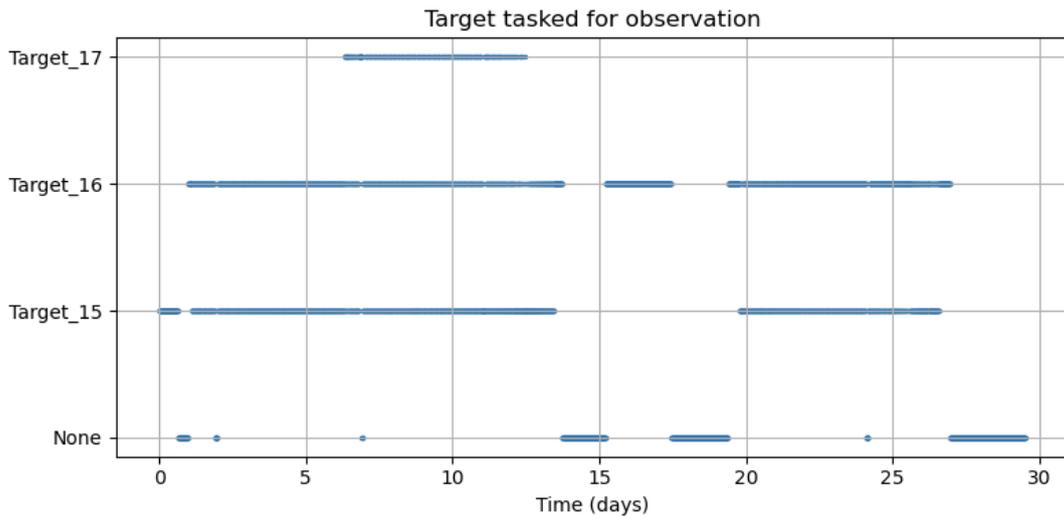


Figure 5.8: Figure showing which target is tasking for observation by the sensor tasking algorithm

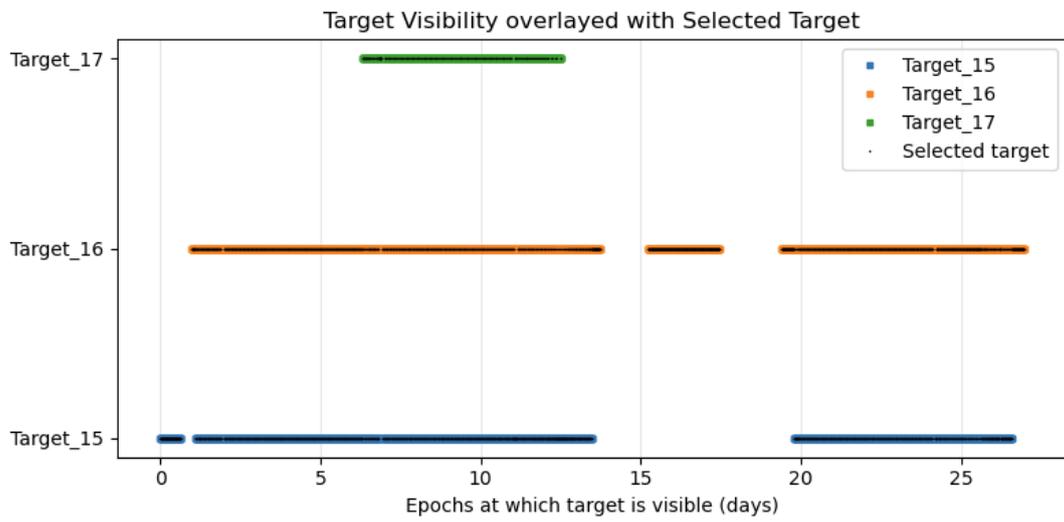


Figure 5.9: Figure showing the epochs at which each of the targets is visible to the observer. It is overlaid with the times at which the target is scheduled for observation

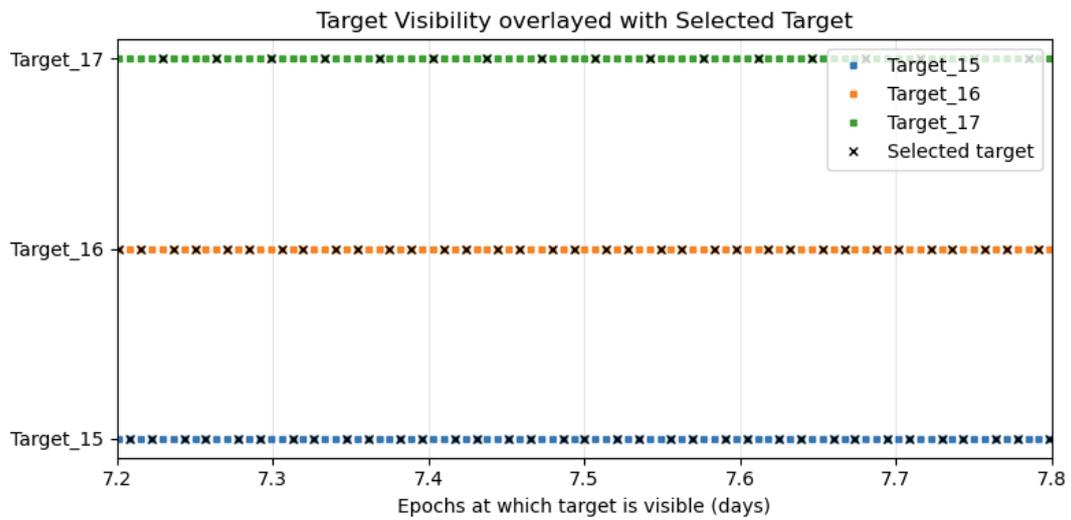


Figure 5.10: Figure showing target visibility and target tasked over a small window of time

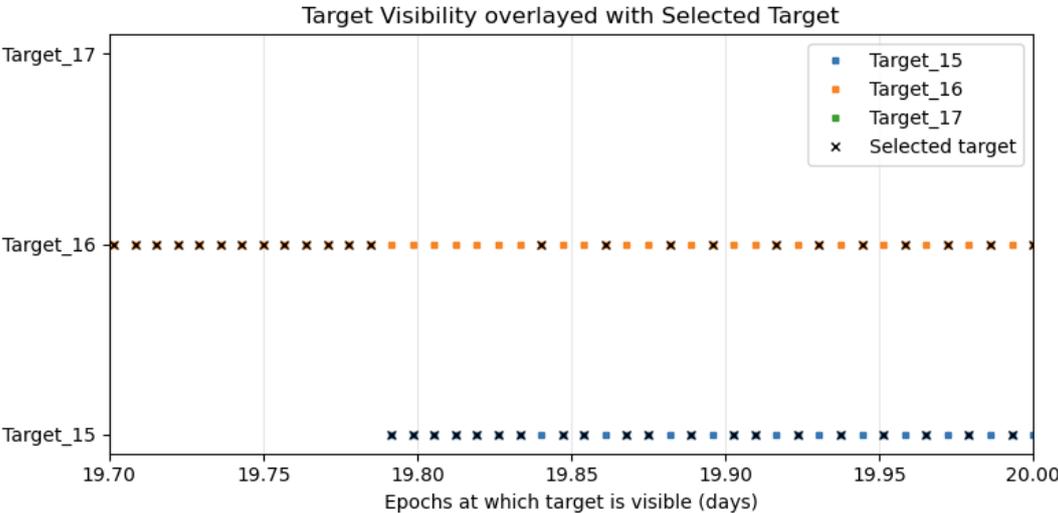


Figure 5.11: Figure showing target visibility and target tasked over a small window of time

6

Results

6.1. Comparison of Observers in Earth's orbit and Cislunar Orbit

The need for a sensor in a cislunar orbit for effective cislunar surveillance is well established in the field of cislunar SSA [1][11]. This is due to the superior visibility geometry and significantly smaller range between the targets and the observer that an observer in an appropriately chosen cislunar orbit offers. This section shows and discusses some results that are obtained through the simulations that support this statement.

For this simulation, a small subset of targets is chosen from the complete target set listed in Tab. 4.4, namely Target 15, 16, 17. For the cislunar observer, Observer 3 is selected from Tab. 4.3. The three targets chosen for analysis correspond to those with the poorest visibility from the selected observer. Focusing on these worst-case scenarios highlights how significantly a cislunar based observer can enhance the ability to monitor cislunar targets. For comparison, two observers are placed in two orbit around the Earth. Their orbital parameters are listed in Table 6.1. An arbitrary high altitude polar orbit (HAPO) is initialized as well as an arbitrary low Earth orbit (LEO).

Element	HAPO	LEO
a (semi-major axis)	30 738 km	7038 km
e (eccentricity)	0.0002	0.0002
i (inclination)	90°	90°
ω (arg. periapsis)	0°	0°
Ω (RAAN)	0°	0°
ν (true anomaly)	80°	80°

Table 6.1: Keplerian orbital elements of Earth observers

Figures 6.1a, 6.1b, and 6.1c show the visibility charts of the three targets mentioned above from the cislunar observer, HAPO observer and LEO observer respectively. The top part of each figure shows the percentage of time for which a target is observed, visible and not visible. The percentage for which it is not visible is shown in green. The bottom part shows the percentage of time for which a particular visibility constraint is violated. Both observers in Earth orbit suffer from poor brightness conditions. It can be seen that the LEO observer provides very poor viewing geometry; it cannot observe targets 15, 16 and 17 for 71, 76, and 90 percent of the simulation time, respectively. Similarly, the HAPO sensor cannot observe targets 15, 16 and 17 for 43, 51, and 80 percent of the simulation time, respectively. Compared to the HAPO observer, the LEO observer experiences long periods of earth exclusion due to its low altitude. The cislunar observer cannot observe targets 15, 16 and 17 for 34, 25, and 80 percent of the simulation time, respectively. It may seem that the cislunar observer offers only slightly better visibility conditions than the HAPO observer based on these figures and values. For this particular set

of targets, that is indeed true, since they were the 3 targets from tab. 4.4 that had the worst viewing geometry with respect to Observer 3. This was done to show that even a cislunar observer which has particularly unfavorable viewing geometry has better visibility of cislunar targets than a space-based observer in Earth's orbit.

explain the figures a bit more, each subfigure has top and bottom, the top provides the percentage of overall time that object is visible, not visible, observed, the bottom plot shows the percentage of time that each visibility constraint is violated, etc

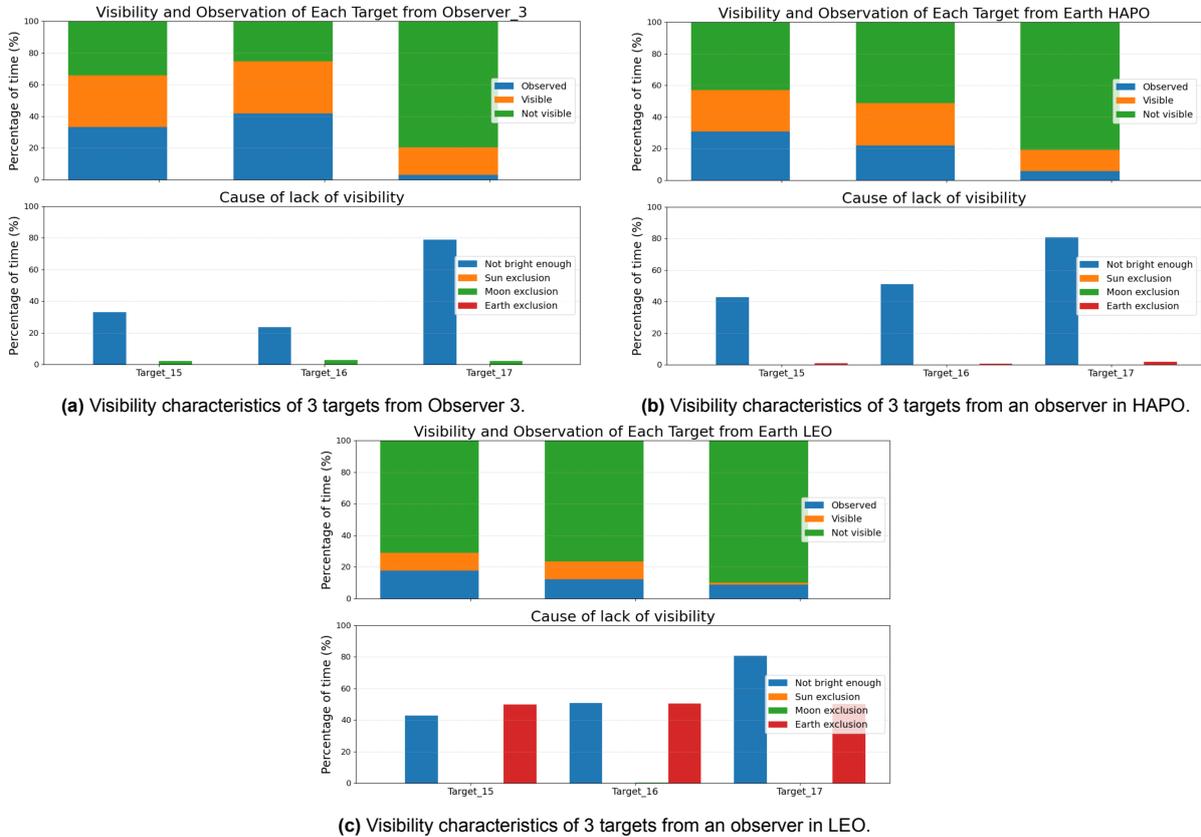


Figure 6.1: Visibility and observation charts for Observer 3, HAPO observer, and LEO observer.

To definitively show the superiority of a cislunar observer with respect to a HAPO observer, three more targets are added to the target set, namely Target 12, 13, 14 from Tab.4.4. Figures 6.2a and 6.2b show the visibility of the appended target set from Observer 3 and the HAPO observer, respectively. Based on these figures it can be seen that the cislunar observer offers much better viewing geometry. The three added targets are visible to the observer for almost the entirety of the simulation, while they are not visible to the HAPO observer for about 35 percent of the simulation duration. Due to the high altitude of the HAPO observer it does not have long periods of Earth exclusion, but it does suffer from poor apparent perceived brightness due to the significantly larger range to the target. Thus, there may exist certain cislunar observer-target combinations from which the sensor has poor visibility geometry for observation, but *in general* a cislunar observer will provide much more favorable viewing geometry than an observer in Earth's orbit. A ground-based Earth observer is not simulated here, but that would certainly have even worse viewing geometry due to Earth horizon constraints, not to mention poor visibility due to light pollution.

Note that these visibility characteristics are determined using the *predicted* target state obtained after the *a priori* time update in the EKF prediction step (eq. 4.20) and not the true states. Thus, it indicates the visibility of the target as determined by the algorithm, not the actual geometric visibility of the target.

Following the discussion and analysis of the visibility of cislunar targets from a cislunar observer as

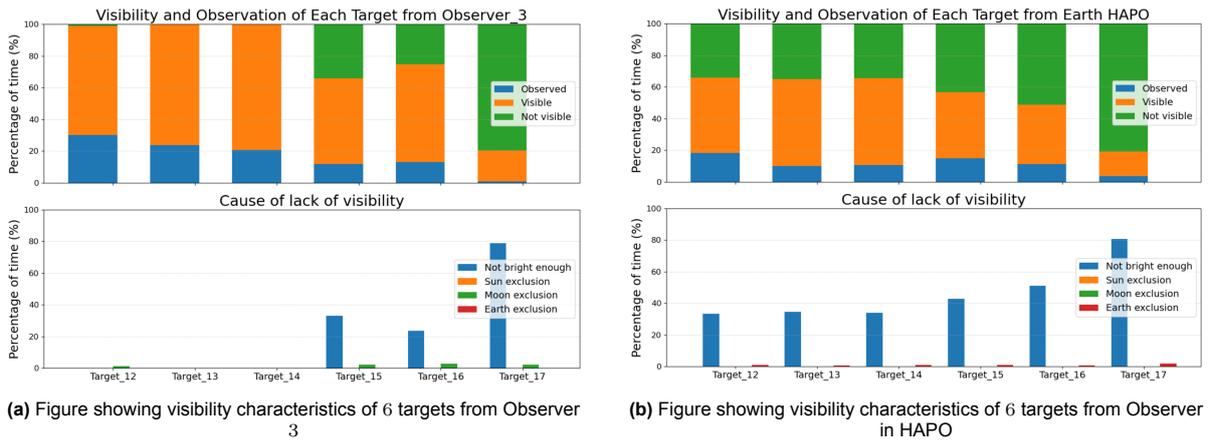


Figure 6.2: Visibility and observation charts for 6 targets for Observer 3 and HAPO observer

compared to observers in Earth orbit, the focus now shifts to the performance of the orbital estimation process carried out by the different observers. This is assessed by evaluating the root mean-squared error (RMSE) of the estimated position and velocity states. Throughout this chapter, two types of RMSE are considered.

The *complete RMSE* (referred to by RMSE) for each target i is computed by calculating the error between the true target states $\mathbf{x}_{i,k}^{true}$ and the estimated state $\hat{\mathbf{X}}_{i,k}$ or the predicted state $\bar{\mathbf{X}}_{i,k}$, depending on if the target i is tasked for observation at a given epoch. As a result, the complete RMSE reflects the overall accuracy of the full estimation process, including both periods when measurements are available and periods when the observer relies solely on prediction through the dynamical model. It therefore provides a global measure of estimation quality and is sensitive to how rapidly estimation errors grow during intervals without observations.

The *observation RMSE* for each target i is computed by calculating the error between the true target states $\mathbf{x}_{i,k}^{true}$ and the estimated state $\hat{\mathbf{X}}_{i,k}$ at epochs when the target i is observed. This metric isolates the performance of the filter when measurement updates are being applied and thus reflects the accuracy of the estimation process under direct observational constraints. In contrast to the complete RMSE, it does not capture error growth during unobserved intervals, and typically provides a more favorable assessment of estimation quality. Comparing the observation RMSE for different observer-reward configurations can help identify which configuration leads to a higher *quality of measurement*. The description of the computation of the two kinds of RMSE can be found in Section A in the Appendix.

The results of the simulation with all the targets in tab.4.4, as observed from the three different observers, are shown in Table 6.2. The RMSE values are averaged across the entire target set.

Observer	RMSE Pos (km)	RMSE Vel (km/s)	Obs. RMSE Pos (km)	Obs. RMSE Vel (km/s)
Observer 3	61.98	8.8×10^{-4}	5.83	6.5×10^{-5}
HAPO Observer	101.63	2.4×10^{-3}	28.22	3.2×10^{-4}
LEO Observer	157.12	3.2×10^{-3}	50.96	4.8×10^{-4}

Table 6.2: Comparison of the average RMSE of targets from cislunar and Earth observers

The values in the table show that a cislunar observer is better for orbit determination for targets in the cislunar region. The observed RMSE for the cislunar observer is lower, since the error in position in an angles-only observation scenario grows proportionally with the range between the target and observer [18]. Additionally, the complete RMSE throughout the simulation period is also lower. This can be attributed partly to the more frequent observations of the target due to better visibility geometry, and partly to the lower error during each observation due to the proximity to the target.

To summarize, it has been shown in this section that an observer in a cislunar orbit exhibits superior

viewing geometry when it comes to observing objects in the cislunar space. Additionally, due to its proximity to the targets, in an angles-only observation scenario, the magnitude of the error in its estimates are also smaller. This supports research in cislunar SSA that emphasizes the need for a space-based observer in a cislunar orbit for effective cislunar surveillance.

6.2. Comparison of Different Observer Orbits

This section compares and analyzes the performance by different observers in estimating the orbits of the complete set of targets as listed in Table 4.4. The description of orbits corresponding to the observers can be seen in Table 6.3. The complete position RMSE is used as the overall metric to compare the performance of the observers and determine the “best” observer. KL Divergence (chosen arbitrarily) is used as the sensor tasking reward in these simulations.

Index	Orbit family	Time period [days]	Stability index ν	Jacobi constant C
1	L1 Lyapunov	28.94	53.67	2.91
2	L2 Lyapunov	25.44	59.09	2.95
3	DPO	5.33	1.00	3.19
4	DPO	0.21	1.00	4.34
5	L1 North Halo	9.62	1.00	3.00
6	L1 South Halo	10.01	1.29	3.00
7	L2 South Halo	7.17	1.52	3.04
8	L2 North Halo	7.17	1.52	3.04

Table 6.3: Catalog of observer orbits

The complete and observed RMSE values are listed in Table 6.4. Observer 8 has the lowest complete position RMSE, while observer 3 has the lowest observation position RMSE. From the results it can be seen that an observer having relatively lower observation RMSE does not necessarily mean it has a relatively lower complete RMSE. The behavior of observer 2 exemplifies this phenomenon, since it has the second highest observation RMSE while having the second lowest complete RMSE. This contrasts with the results of observer 3, which has the lowest observation RMSE and the third lowest complete RMSE. This result can be understood by looking at the visibility chart of the two observers. From Figure 6.3 it is clear that Observer 2 has better visibility geometry with respect to the target set than Observer 3, which allows it to take more regular measurements than observer 3. This implies that observer 3 is better in terms of error reduction when it observes a target (low observation RMSE), but since observer 2 offers better visibility characteristics it provides a better overall RMSE through more observations. That can be supported by Figure 6.4, which shows the tasking action of the two observer’s sensor tasking algorithm. Due to observer 3’s relatively worse visibility to the targets, there are long periods of time when Targets 15, 16, 17 go unobserved. In fact, Target 17 goes unobserved for more than half the simulation. On the other hand, this does not happen with observer 2, which is able to take regular measurements of all the targets due to its superior visibility geometry.

Observer	RMSE Pos (km)	RMSE Vel (km/s)	Obs. RMSE Pos (km)	Obs. RMSE Vel (km/s)
Observer_1	56.99	2.6×10^{-3}	22.79	2.7×10^{-3}
Observer_2	15.00	2.7×10^{-4}	16.05	2.0×10^{-4}
Observer_3	18.14	2.4×10^{-4}	5.91	6.7×10^{-5}
Observer_4	64.05	8.7×10^{-4}	6.47	6.2×10^{-5}
Observer_5	84.69	1.2×10^{-3}	8.52	7.7×10^{-5}
Observer_6	58.32	8.4×10^{-4}	7.43	7.5×10^{-5}
Observer_7	23.33	2.7×10^{-4}	7.50	6.9×10^{-5}
Observer_8	13.94	1.6×10^{-4}	8.50	7.1×10^{-5}

Table 6.4: Complete and Observation RMSE of position and velocity, averaged across all the targets, for each observer .

This, however, does not explain why the observed RMSE of observer 3 is lower than observer 2. Since both observers employ the same myopic reward, they should have similar observation RMSE. This can be explained by Table 6.5 which shows the average distance to all the targets from observer 2 and 3 during the simulation. It can be clearly seen that Observer 3 is relatively closer to the targets than Observer 2. As mentioned earlier, in an angles-only measurement scenario, the magnitude of the

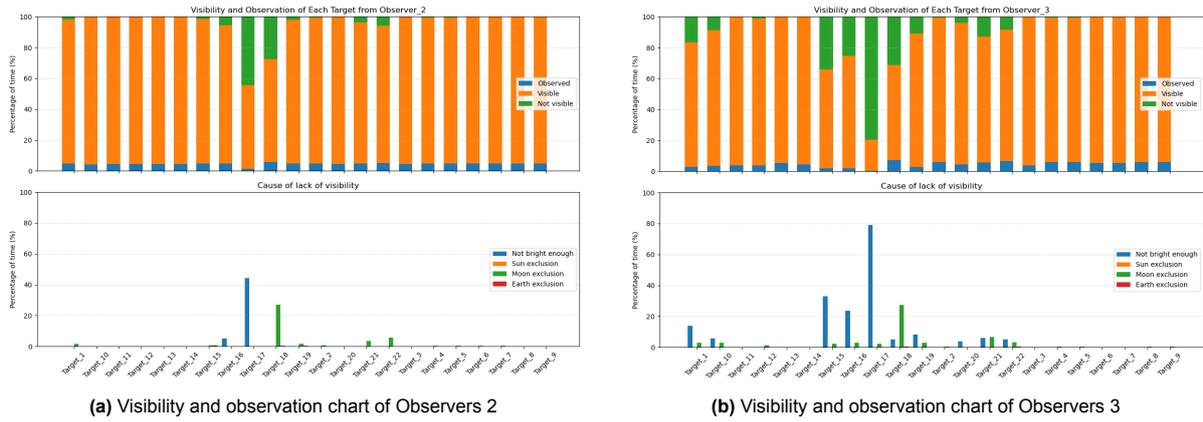


Figure 6.3: Visibility and observation chart of Observers 2 and 3

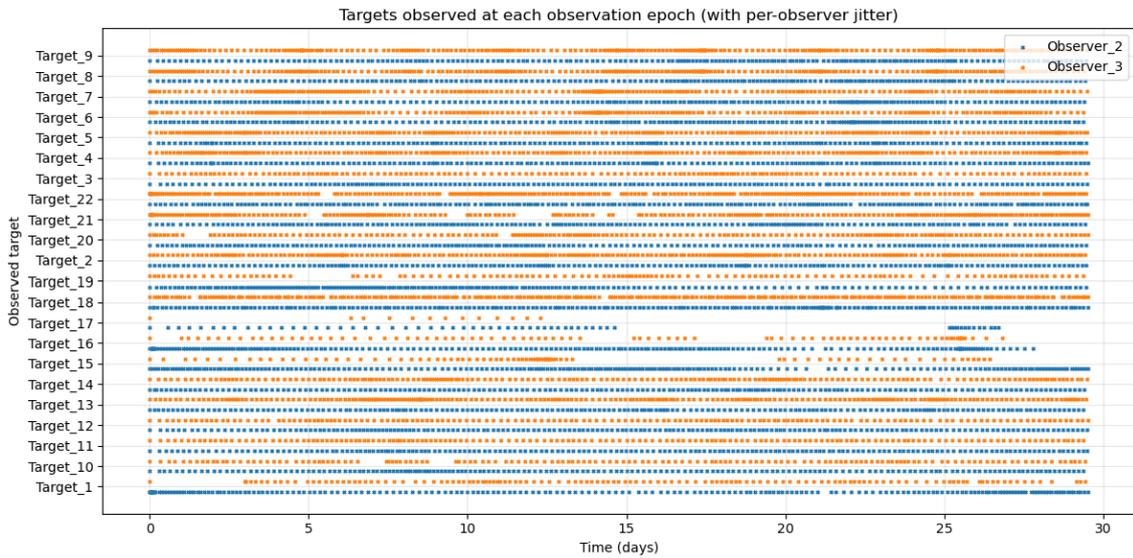


Figure 6.4: Figure showing the targets that are scheduled for observation at each time step, from observers 2 and 3

error in the estimated position grows proportionally with the range between the observer and the target. Thus, due to its proximity to the targets, observer 3 has a smaller observation RMSE. Due to its low observation RMSE, even though it has poorer visibility than observer 2, its complete RMSE is similar to observer 2.

Note that the observation RMSE of observer 2 is higher than complete RMSE. If the observation RMSE for a simulation is higher than the complete RMSE, this indicates that the estimation errors at the epochs where measurements are incorporated are, on average, larger than the errors during periods without observations. Since the complete RMSE includes both observed and unobserved epochs, while the observation RMSE isolates only the observed ones, this situation implies that the filter performs worse during measurement updates than during prediction only intervals. This is a counterintuitive result, as measurement updates are expected to reduce estimation error.

A possible explanation for this behavior is related to the structure of the EKF. In the simulations for this work, the filter propagates its state using the same dynamical model employed to generate the true trajectory, with only a small process noise term applied. As a result, the prediction step is inherently very accurate. However, when the observer is far from the target, the measurement noise in range can translate into relatively large position uncertainties during the update step. In such cases, the incorporation of noisy measurements may degrade the estimate, causing the errors at observed epochs to exceed those during prediction-only intervals. As shown in Section 5.2 there are some inconsistencies

with the

Despite this, it must be mentioned the the observed RMSE in velocity is indeed lower than the complete velocity RMSE, and that the difference between the complete and observation position RMSE is not large (15 and 16.05).

Observer	Average Range (km)
Observer 1	187 263.84
Observer 2	160 866.65
Observer 3	96 242.08
Observer 4	89 674.11
Observer 5	115 517.72
Observer 6	114 368.89
Observer 7	107 999.46
Observer 8	106 658.68

Table 6.5: Average range across all targets for each observer

It must be noted that although the average range to the target shows some correlation to the observation RMSE, a lower average range does not necessarily mean a lower observation RMSE. This can be confirmed by comparing the range and observation RMSE values of observer 3 and 4. The observation RMSE is the result of a variety of factors in a complex physical system involving relative motion, visibility, sensor tasking and state estimation. As discussed earlier, the observability angle Ψ_f has a significant impact on estimation accuracy as well.

Looking at Table 6.4 and 6.3, a loose correlation can be established between the period of the orbit and its observation RMSE. An orbit with a smaller period will be closer to the Moon, and thus will have a smaller average range to the targets in the cislunar space, leading to better estimation accuracy (reflected by observation RMSE). However, having a smaller orbit also means the observer is closer to the moon and more likely to have its LOS to the target obstructed by the Moon. This can be seen in the visibility chart of observer 4 (Figure 6.5a), which is in a DPO with a very small period.

It has been discussed why Observer 2 has a significantly higher observation RMSE but a very low complete RMSE. Inversely, some of the observers (4, 5, 6) have low observation RMSE but very high complete RMSE. To analyze this behavior, observer 6 is compared with observer 3. The visibility chart of observer 6 is shown in Figure 6.5. Comparing this figure to the visibility chart of Observer 3 (in fig. 6.3b), it is not very clear why observer 6 has an observation RMSE close to that of Observer 3, but a much larger complete RMSE. The higher observation RMSE can be explained by the higher average range of observer 6 to the targets as shown in tab. 6.5. However, since they both have a similar visibility to the targets, the significantly higher complete RMSE of observer 6 cannot be attributed to the slightly larger range only. It seems that it is difficult to explain this data with a physical interpretation of the orbits due to the complex nature of the simulation involving the sensor tasking algorithm as well as the interaction between the observer and the targets. This highlights the importance of getting empirical data through the development of increased-fidelity simulation frameworks.

Observer 8 has the lowest complete RMSE in the observer set. Its visibility chart is shown in Figure 6.6. Although it has a relatively high observation RMSE compared to observers 3 – 7, due to its visibility geometry it performs well overall. This orbit is part of the L2 Northern Halo orbit family. The Halo orbit family has also been proposed as a suitable cislunar surveillance orbit by various works [43][51][1]. However, the fact that an orbit family contains an infinite number of individual orbits, and different orbits in the same family can have very different observational abilities, should be acknowledged. This is also highlighted by the results in tab. 6.4, since some Halo orbits (observers 5 and 6) show poor performance in orbit estimation.

Vendl et al. [43], through their analysis of the visibility of the cislunar space from different observer orbits, concluded that the L1 and L2 Lyapunov families exhibit the most favorable conditions for cislunar surveillance. This statement is supported by these simulations (observer 2 is in the L2 Lyapunov orbit).

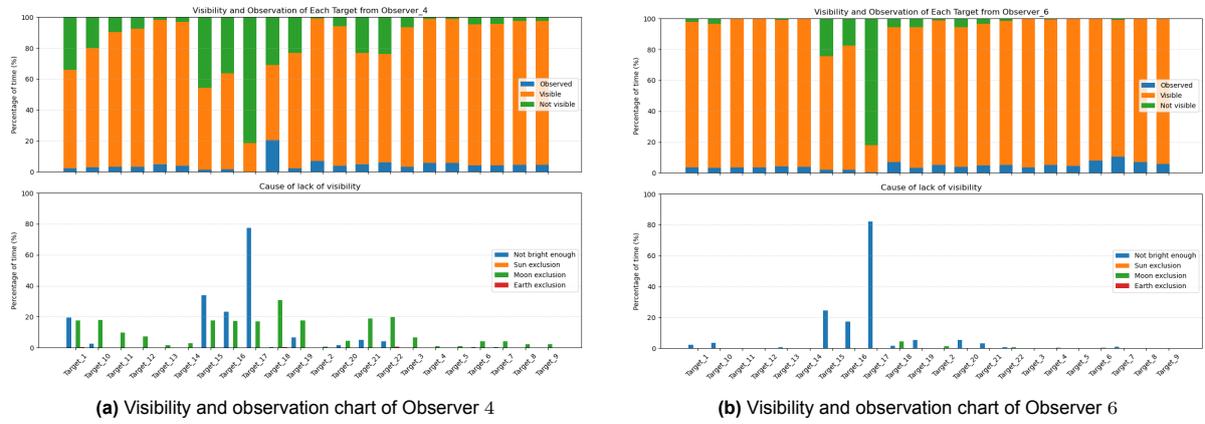


Figure 6.5: Visibility and observation chart from observer 4 and 6

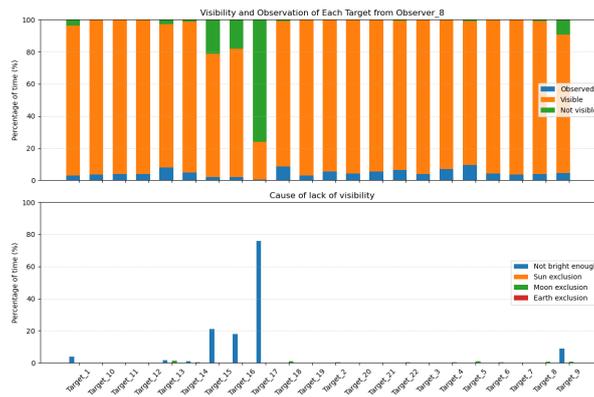


Figure 6.6: Visibility and observation chart of Observer 8

However, due to the relatively large range of an observer from targets in the cislunar space, an observer capable of angles-only measurements placed in this orbit family suffers from estimation disadvantages due to its range. Despite this disadvantage the observer has a good overall performance, and may still offer a good observer placement opportunity. It must also be mentioned the Lyapunov orbits have a high stability index compared to the other observer orbits in my catalog, which would mean higher station keeping costs. Their analysis also showed that Halo orbits exhibit good visibility to the cislunar space, which is also supported by these simulations.

Summary

This section evaluated the orbit determination performance obtained from eight different cislunar observer orbits. Two complementary metrics were used: the *observed* RMSE, which reflects estimation quality during measurement updates, and the *complete* RMSE, which measures overall performance including propagation between observations.

The results show that an observer may achieve excellent instantaneous estimation accuracy while performing poorly overall, and vice versa. For example, Observer 3 has the lowest observed RMSE due to its close proximity to the targets, but its complete RMSE is worse than that of Observer 2, which has a much larger range but superior visibility geometry. Observer 2's placement in an L2 Lyapunov orbit provides regular, unobstructed view to the targets enabling better long-term estimation performance. Similar behavior is seen across the catalog. Short-period orbits near the Moon yield low observed RMSE but suffer from frequent lunar occultation, while distant Lyapunov and Halo orbits offer high-quality overall performance despite less accurate individual measurements.

The results also highlight the complex interaction between geometry, and sensor tasking in angles-only orbit determination, and emphasize the importance of empirical evaluation using simulations that reflect realistic dynamical and measurement models. Study of the evaluation of the observability angle Ψ_f alongside the relative range in sensor tasking simulations can help with understanding the role of relative geometry better and aid in the design of sensor tasking strategies. Overall, the analysis supports previous findings in the literature that Lyapunov, DRO and Halo orbits provide strong visibility for cislunar surveillance, while also demonstrating that individual members of the same orbit family may yield substantially different estimation performance. A more thorough and complete analysis of observer orbits would include a larger set of observer orbits.

6.3. Discussion of orbits of low visibility

Through the visibility charts and tasking figures of different observers in the previous section, it can be seen that the observers have a good visibility to most of the targets. Most of the targets are visible to the observers for more than 90% of the duration of the simulation. However, target 17 is not visible to the observers for roughly 90% of the simulation duration. The reason for the relatively high period of invisibility for this target is that the target is not bright enough.

Figure 6.7, shows the orbit of target 17 alongside the orbit of observer 8. This is the largest orbit in the catalog of target orbits, and due to its size its distance from the observer is very large, leading to the target being too dim for observation. The reason for its inclusion in the target catalog is that it covers a large portion of the cislunar space and is very stable, which makes it a favorable orbit for cislunar communication and navigation spacecraft [1].

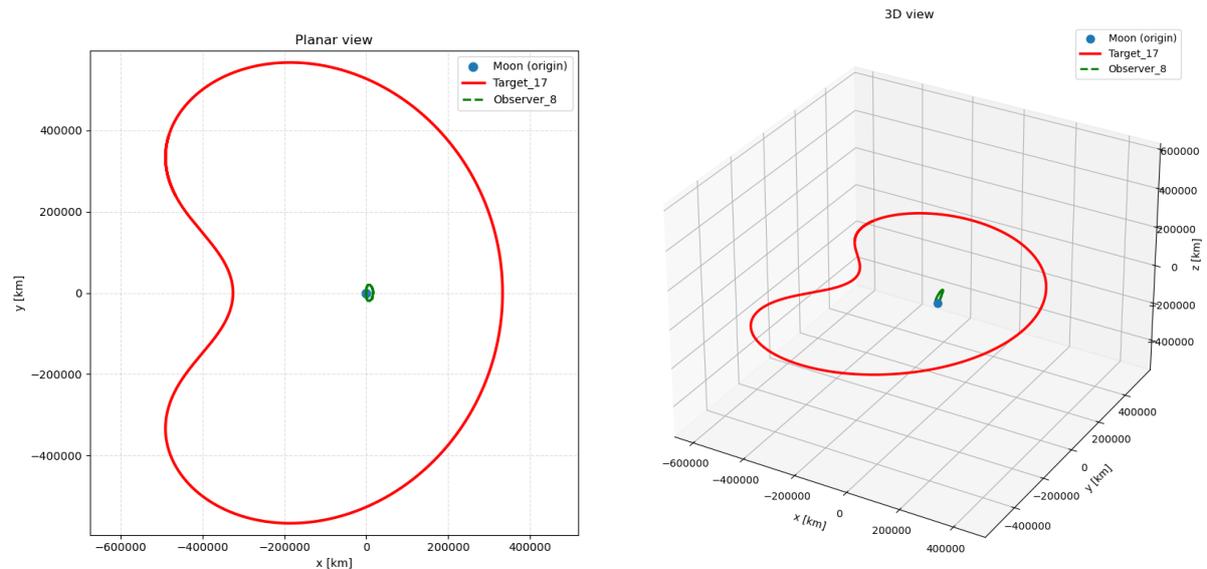


Figure 6.7: Figure showing the orbit of Target 17 along with Observer 8

6.4. Comparison of reward functions

This section discusses how different reward functions perform in the simulation framework. The simulation is run with the reward functions discussed in sec 4.5, namely KL divergence (KLD), Cauchy–Schwarz divergence (DCS), mutual information (MI), finite time Lyapunov exponent (FTLE) and age-of-information (AOI). Additionally, the horizon of FTLE is also varied, and is represented by “FTLE x ” where x is the number of horizons. The simulations are conducted with observer 8 (since it had the best overall performance) as well as observer 3. Two observers are used to study if the trends of reward function performance remain the same from one observer to another.

Complete catalog (22 targets)

Table 6.6 shows the complete and observed RMSE in position from different configurations of observers and reward functions.

Reward	Observer	RMSE position [km]	Observed-only RMSE [km]
KLD	8	13.94	8.50
MI	8	32.42	7.13
DCS	8	42.16	7.76
FTLE1	8	29.74	16.61
FTLE2	8	36.05	17.10
AOI	8	16.82	9.11
KLD	3	18.14	5.91
MI	3	62.38	4.84
DCS	3	8.37	5.22
FTLE1	3	18.54	17.40
FTLE2	3	22.59	18.38
AOI	3	21.06	6.86

Table 6.6: RMSE position and observed-only RMSE position for different reward functions and observers

Since KL divergence, CS divergence and mutual information are different measures of information gain, their tasking outputs should be similar to each other. This implies that for the same observer, their complete and observed RMSE should be similar. Although their observation RMSEs are similar for both observer 3 and 8, the complete RMSE values of the IG rewards are highly irregular. For observer 8, KLD achieves the smallest complete RMSE, while MI and DCS produce substantially larger errors. For observer 3, the trend is essentially reversed. DCS attains the smallest complete RMSE, and MI and KLD perform much worse. Additionally, since the three rewards are IG-based, they should result in broadly similar tasking decisions. This is shown in Figure 6.8. This figure shows that the tasking output for the three configurations are almost identical. Thus, the three rewards generally schedule the same target for observation. Somehow, despite having very similar sensing outputs as well as measurement quality (reflected by observed RMSE), the end-to-end performance (reflected by the total RMSE) is completely different.

Analyzing the RMSE for individual targets, as shown in Table 6.7, helps identify the cause of this unexpected behavior. The complete and position RMSE are similar for all the targets across the three rewards, except for Target 17. Since Target 17 is only visible for a small period of time during the simulation, very few observations are taken (about 12 by each reward). The three rewards lead to the observation of Target 17 on slightly different epochs, as seen in fig. 6.8. However, once Target 17 becomes invisible again and its position error begins to increase, it does so much more in the algorithm with MI and DCS as rewards as compared to the algorithm with KLD reward. This is reflected in the significant difference in the total RMSE of Target 17 from the three rewards (tab. 6.7). Since the RMSE is averaged across all the targets without factoring in the number of observations made, this skews the final results. Since the complete RMSE is computed by averaging the per-target errors uniformly across all targets (without weighting by the number of observations per target), a single target with very few observations but a large estimation error can disproportionately skew the final result. A discussion of the results excluding Target 17 is given in Section 6.4.

This highlights a disadvantage of IG-based reward functions. Once Target 17 becomes visible after

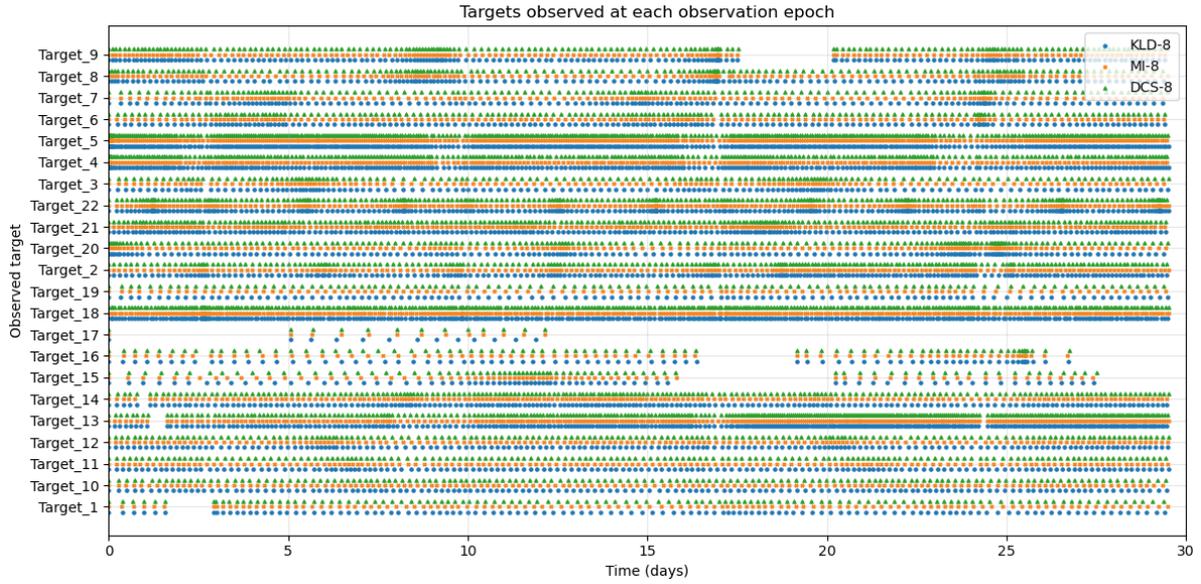


Figure 6.8: Scheduling of targets from IG-based reward functions from observer 8

Target	KLD-8 (C)	KLD-8 (O)	MI-8 (C)	MI-8 (O)	DCS-8 (C)	DCS-8 (O)
1	6.485	6.353	8.921	8.673	9.688	9.578
2	8.100	8.446	8.238	8.708	7.576	7.770
3	16.872	15.561	11.993	11.028	18.449	16.776
4	5.373	5.156	8.516	8.325	5.802	5.541
5	4.234	4.180	2.689	2.547	2.552	2.391
6	5.305	4.877	6.499	6.181	6.443	5.992
7	10.733	9.729	9.384	8.429	8.331	7.784
8	12.118	12.361	12.191	10.827	9.783	8.912
9	5.213	4.896	5.289	4.724	4.227	3.948
10	6.632	6.451	4.565	4.278	5.168	4.832
11	6.936	6.941	5.005	4.920	6.486	6.212
12	6.317	6.139	7.660	7.388	6.910	6.615
13	8.544	9.663	6.623	7.757	6.578	7.500
14	10.238	10.933	4.486	4.340	6.235	5.943
15	23.665	17.867	17.653	15.989	15.753	14.158
16	19.349	19.764	15.976	14.780	17.297	16.766
17	45.154	36.930	146.801	23.382	193.317	46.670
18	0.602	0.374	0.637	0.403	0.644	0.429
19	10.545	10.346	9.860	9.678	7.980	7.776
20	10.961	10.318	7.209	6.763	12.374	11.633
21	1.951	1.728	2.047	1.837	1.992	1.794
22	1.345	1.149	1.314	1.102	1.448	1.184

Table 6.7: Complete and observed RMSE in position for individual targets for IG-based reward functions from observer 8. Columns with (C) refer to complete RMSE and columns with (O) refer to observed RMSE

a period of invisibility, a well designed algorithm should prioritize observing this target. This does not happen, and since the algorithm is myopic, it cannot predict that this target will become invisible after some time. Thus, the algorithm schedules very few observations of Target 17 during its short period of visibility to the observer, leading to bad overall performance in minimizing the error of this particular target.

It must be mentioned that even though in this particular case using KLD reward leads to significantly less complete RMSE of Target 17 than DCS and MI, it does not mean that KLD reward is better at handling cases where visibility is low and very few observations can be made. For the observer 3 configuration results (Table. 6.6), the DCS reward has the lowest complete RMSE while KLD and MI RMSE blows up. This was again due to the low observation period of Target 17, but this time the algorithm with DCS happens to manage this better. In fact, in this configuration, DCS does extremely well since the complete RMSE is very low (8.37 km). The number of observations of Target 17 taken by each IG-based reward algorithm is similar, but since they are very low the future growth in its predicted uncertainty during its period of invisibility is very sensitive to the exact epochs at which its measurements are taken. The number of observations of Target 17 taken by each IG-based reward algorithm is similar, but since they are very few, the growth of its predicted uncertainty during its invisibility period is highly sensitive to the exact measurement times. In this case, one reward only appears better because it happens, by chance, to place the sparse observations at slightly more favorable epochs, leading to a smaller prediction error after the target becomes unobservable. This advantage is therefore not systematic and cannot be expected to generalize to other cases. Thus, the fact that KLD happens to handle this extreme case better for the observer 8 configuration and DCS for the observer 3 configuration is largely coincidental, and these specific configurations should not be expected to yield consistently good performance in general.

A multi-step tasking algorithm could potentially be more effective in the scheduling of orbits with low visibility, as it could predict its future state of invisibility and schedule a higher number of observations when it is visible. Another possible improvement would be to weight the information gain with a scalar that represents the magnitude of the current uncertainty \bar{P}_k of the target. Purely information gain based rewards only consider the expected *change* in the uncertainty through a prospective measurement. Due to this, when Target 17 re-enters the observer's visibility region, its a priori uncertainty is very large due to the long gap without observations, but it is still scheduled only sparsely. Since Target 17 is a very distant target, as shown in Section 6.3, its poor observational geometry does not allow for a measurement that would significantly reduce its uncertainty. Other targets must offer larger expected reductions in covariance and therefore dominate the ranking. By scaling the information gain with a measure of the current uncertainty, the scheduler could explicitly favor highly uncertain targets, reducing the risk that objects like Target 17 get scheduled infrequently even when they are finally observable. Another improvement would be to add realistic slewing cost and penalties related to switching targets, which would lead to observations over consecutive decision epochs, helping reduce the estimation error quickly.

The algorithm with the FTLE reward performs poorly. Not only does it have a high complete RMSE, it has a very high observation RMSE. The reason for this becomes clear looking at Table 6.8, which lists the average of the range to the targets at the time they are scheduled to be observed, for each reward.

It seems that the FTLE algorithm schedules targets that are at larger distances compared to other reward algorithms. The FTLE reward, defined in Sec. 4.5, measures how large the state uncertainty of target i is expected to become after H prediction steps in the absence of additional measurements. The FTLE reward r_{FTLE} depends only on the dynamical propagation through $\Phi_k^{(H)}$ and the current predicted covariance \bar{P}_k . As a result of this, the FTLE algorithm prioritizes targets whose uncertainty grows rapidly under the CR3BP dynamics, irrespective of whether the current observer-target geometry is favorable for obtaining an informative measurement. This is because the measurement model H_k and measurement noise covariance R_k do not appear in this reward. Additionally, the uncertainty \bar{P}_k is likely to be larger for targets that are further (estimation accuracy is proportional to range in angles-only observation scenario, as shown in Section 5.2), which leads to the FTLE algorithm prioritizing targets that are further away.

On the other hand, the IG-based rewards (KLD, MI, and DCS) compute their scalar score from the *pos-*

Reward	Average range to observed targets [km]
KLD-8	70,806.54
MI-8	70,679.30
DCS-8	70,873.02
FTLE1-8	201,102.07
FTLE2-8	200,902.26
AOI-8	87,332.68

Table 6.8: Average range from observer 8 to observed targets for each reward function.

terior covariance after a hypothetical measurement update. Because the update explicitly involves H_k and R_k , these rewards naturally favor configurations where the observational geometry is well conditioned (shorter observer-target range) and the measurement significantly reduces the state uncertainty. This is why the average observer-target range during an observation is smaller for the IG-based rewards compared to FTLE.

An improvement in the formulation of the FTLE-based reward could potentially lead to better scheduling. Since the FTLE reward used in this simulation was biased to targets at larger distances, scaling the FTLE-based reward with a factor of the distance to the target would eliminate the bias.

The age-of-information (AOI) reward selects at each decision epoch the visible target that has remained unobserved for the longest time. From Table 6.6 it can be seen that the AOI reward achieves a complete RMSE of 16.82 km for observer 8, which is slightly worse than KLD (13.94 km) but substantially better than MI and DCS (32.42 km and 42.16 km, respectively), as well as both FTLE configurations (29.74 km and 36.05 km). The observed RMSE for AOI (9.11 km) is comparable to that of KLD (8.50 km) and somewhat worse than MI and DCS (7.13 km and 7.76 km). A similar trend is observed for observer 3.

The AOI reward also exhibits intermediate behavior in terms of observer-target geometry. For observer 8, the average range at the time of observation is about 87,333 km, larger than for the IG-based rewards (70,70070,900 km) but much smaller than for FTLE1 and FTLE2 (201,000 km). This is due to the fact that AOI reward does not take into account the distance to the target or the observer-target geometry in any way.

However, through a “longest-unobserved first” policy at each step, AOI avoids the extreme behavior of FTLE, which tends to schedule distant targets leading to worse observation RMSE. Additionally, it is better at handling the extreme case of target 17, as shown by its tasking decisions in Figure 6.9. For the AOI and observer 8 configuration, the target 17 only observed RMSE is 29.02 km and complete RMSE is 66.89 km, while those of the AOI and observer 3 configuration is 35.43 and 93.38 km. Thus, although it is not as low as the lowest IG-reward configuration, it does remain more stable across both the observers due to the larger number of scheduled measurements of the target.

Overall, a scheduler based on AOI provides a robust, low-complexity option, which could be used as a baseline to compare other rewards. It does not perform as well as the best IG-based reward in terms of complete RMSE, but it consistently outperforms the FTLE-based reward and is much less sensitive to extreme behaviors such as that of Target 17.

Truncated catalog (21 targets)

Since the extreme behavior of Target 17 skews the complete RMSE of the different observer reward configurations, Table 6.9 shows the averaged RMSE values across all targets except target 17.

Excluding Target 17 from the averaging restores the expected behavior of the IG-based rewards. From Table 6.9 it can be seen that, for observer 8, the complete RMSE of the three IG-based rewards (KLD, MI, and DCS) is between 8.6 km and 10.3 km, with MI achieving the smallest complete RMSE (8.64 km), followed closely by DCS (9.06 km) and AOI (9.13 km). The observed RMSE values are even closer, ranging from 7.02 km to 8.62 km, with MI again performing best. The FTLE configurations remain significantly worse, with complete RMSE values above 12.5 km and observed RMSE larger than 15 km.

The results from observer 3 behave similarly. The complete RMSE of the IG-based rewards is tightly

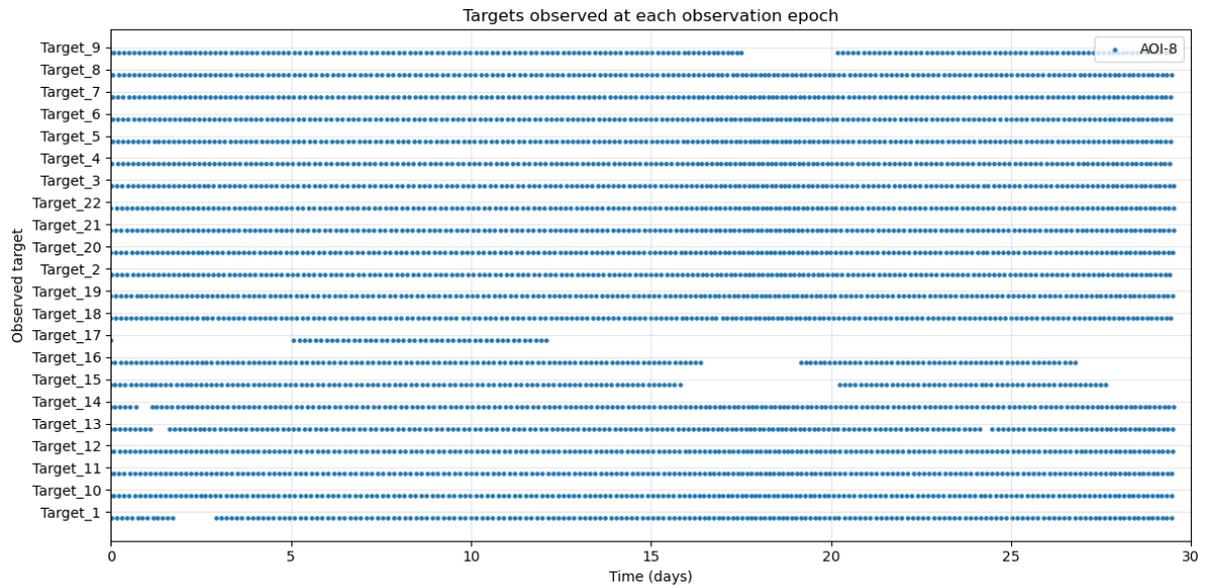


Figure 6.9: Scheduling of targets from IG-based reward functions from observer 8

Reward	Observer	RMSE position [km]	Observed RMSE [km]
KLD	8	10.326	8.27
MI	8	8.64	7.02
DCS	8	9.06	7.33
FTLE1	8	12.66	15.13
FTLE2	8	12.88	16.31
AOI	8	9.13	8.62
KLD	3	6.08	4.70
MI	3	6.15	4.49
DCS	3	6.72	4.90
FTLE1	3	15.04	16.51
FTLE2	3	10.97	17.87
AOI	3	7.025	5.92

Table 6.9: RMSE position and observed-only RMSE position for different reward functions and observers

clustered between 6.08 km and 6.72 km, with KLD and MI very close to each other (6.08 km and 6.15 km, respectively) and DCS only slightly worse (6.72 km). The observed RMSE follows the same pattern. MI has the smallest value (4.49 km), KLD and DCS are only marginally larger (4.70 km and 4.90 km), and AOI is larger at 5.92 km. Once again, the FTLE-based rewards perform poorly for observer 3, with complete RMSE of 10.97 km and 15.04 km and observed RMSE in excess of 16 km for both FTLE2 and FTLE1.

Comparing these truncated results with the full catalog values in Table 6.6 highlights the dominant influence of Target 17 on the complete RMSE. For instance, for observer 8 the complete RMSE of MI and DCS reduces from 32.42 km and 42.16 km to 8.64 km and 9.06 km, respectively, once Target 17 is excluded. However, the observed RMSE values change very little between the full and truncated catalogs. This confirms that Target 17 is rarely observed and that its large estimation error during its prediction period disproportionately affects the average over all targets. The measurement quality, however, reflected by the observed RMSE, remains comparable across the full and truncated catalogs.

The truncated catalog therefore provides a clearer view of the intrinsic behavior of each reward function. The three IG-based rewards behave in a consistent manner and deliver the best overall performance in terms of both complete and observed RMSE for both observers, with only minor differences between KLD, MI, and DCS. The AOI reward remains competitive, typically within about 1–2 km of the best IG-based configuration, and continues to offer robust performance across the two observers. FTLE-based rewards, on the other hand, perform substantially worse than the other rewards even after the removal of Target 17, indicating that their poor performance is not solely due to the extreme behavior of this single target but is also related to their tendency to prioritize distant targets leading to poor measurement quality.

Overall, the comparison based on the truncated catalog supports the initial intuition that when extreme cases such as Target 17 are excluded, IG-based rewards produce very similar tasking decisions and estimation performance, AOI provides a simple and robust baseline that is only slightly worse on average, and FTLE-based rewards are consistently dominated by the other approaches.

Fairness metrics across targets (truncated catalog)

To add to the analysis based on RMSE values, examination of how each reward function distributes sensing effort and estimation accuracy across the catalog is also done. For each reward r and target i , let $N_{\text{obs}}^{(r,i)}$ denote the number of measurement epochs assigned to that target. $\text{RMSE}^{(r,i)}$ denotes the complete position RMSE of target i for reward r . The distribution of these quantities is determined, over all $\mathcal{T} = \{21\}$ targets (Target 17 excluded), using the mean, standard deviation, median, minimum, maximum, and 95th percentile. In addition, we compute the Pearson correlation coefficient between $N_{\text{obs}}^{(r,i)}$ and $\text{RMSE}^{(r,i)}$ across targets, which indicates whether additional observations tend to be spent on easier or harder targets.

Table 6.10 reports coverage statistics for observer 8 in the truncated catalog. The IG-based rewards (KLD, DCS, MI) dedicate on average about 203 observations per target, with standard deviations around 85 and medians slightly below the mean. AOI-8 again provides the most uniform coverage, with a mean of 201.24 observations, a very small standard deviation ($\sigma_{N_{\text{obs}}} \approx 13$), and a narrow range between minimum and maximum counts. In contrast, the FTLE-based rewards allocate substantially fewer observations on average (about 160 per target) and exhibit very uneven coverage: FTLE1-8 and FTLE2-8 have standard deviations above 200, low medians (36 and 43 observations, respectively) and maxima above 750 observations, indicating that a few targets are observed very frequently while many others receive relatively few measurements.

Table 6.11 summarizes the distribution of per-target complete position RMSE values for the truncated catalog. Here $\overline{\text{RMSE}}$ and σ_{RMSE} are the mean and standard deviation across targets, while RMSE^{max} and $\text{RMSE}^{95\%}$ characterize the behavior of the worst targets. The last column reports the correlation coefficient

$$\rho(N_{\text{obs}}, \text{RMSE})$$

between $N_{\text{obs}}^{(r,i)}$ and $\text{RMSE}^{(r,i)}$ for each reward r .

Removing Target 17 leads to a much more homogeneous distribution of per-target RMSE values for the IG-based rewards and AOI. Their mean per-target RMSE lies between roughly 7.5 and 8.6 km, with

Reward	\bar{N}_{obs}	$\sigma_{N_{\text{obs}}}$	median	N_{obs}^{\min}	N_{obs}^{\max}	$N_{\text{obs}}^{95\%}$
KLD-8	202.90	85.66	177.00	79	409	358
DCS-8	202.90	86.18	177.00	79	404	364
MI-8	202.90	85.00	179.00	80	410	357
FTLE1-8	159.95	210.39	36.00	23	767	526
FTLE2-8	160.00	203.61	43.00	23	771	494
AOI-8	201.24	13.21	207.00	160	208	207

Table 6.10: Observation statistics for observer 8 over 21 targets (Target 17 excluded). \bar{N}_{obs} and $\sigma_{N_{\text{obs}}}$ denote the mean and standard deviation of the number of observations per target; the remaining columns show the median, minimum, maximum and 95th percentile.

Table 6.11: Per-target all-times position RMSE statistics for observer 8 over 21 targets (Target 17 excluded). All RMSE values are in km. The last column shows the correlation coefficient between per-target observation counts and RMSE.

Reward	$\overline{\text{RMSE}}$	σ_{RMSE}	median	RMSE^{\max}	$\text{RMSE}^{95\%}$	$\rho(N_{\text{obs}}, \text{RMSE})$
KLD-8	8.644	5.644	6.936	23.665	19.349	-0.645
DCS-8	7.701	4.782	6.578	18.449	17.297	-0.673
MI-8	7.465	4.357	7.209	17.653	15.976	-0.689
FTLE1-8	11.028	6.213	9.801	25.883	23.885	+0.214
FTLE2-8	11.185	6.380	11.483	25.467	23.298	+0.189
AOI-8	7.840	4.683	7.919	18.093	15.709	-0.540

95th-percentile values below 20 km and relatively modest worst-case errors. The FTLE-based rewards remain clearly inferior, with mean per-target RMSE around 11 km and significantly larger 95th-percentile and maximum errors, indicating that several targets still suffer from comparatively poor estimation accuracy.

The correlation coefficients provide additional insight into how each reward allocates effort. For KLD, DCS, MI and AOI, the correlation between N_{obs} and RMSE is negative, meaning that targets which are observed more often tend to achieve lower estimation error, as expected. In contrast, FTLE1 and FTLE2 exhibit a weakly positive correlation, suggesting that additional observations are frequently spent on geometrically challenging targets without fully compensating their higher estimation error. Overall, the fairness metrics confirm that, in the truncated catalog, the IG-based rewards and AOI provide both competitive average performance and relatively balanced accuracy across the target set, whereas FTLE-based tasking remains poor in both coverage and worst-case error.

7

Conclusion and Recommendations

7.1. Conclusion

This thesis investigated how cislunar Space Situational Awareness (SSA) by a space-based observer can be supported through the combined evaluation of sensor tasking strategies and observer orbits in the Earth–Moon Circular Restricted Three-Body Problem (CR3BP). A simulation framework was developed using the TUDAT toolbox, with a catalog of 22 representative cislunar target orbits and 8 candidate observer orbits drawn from the JPL periodic orbit database. The observer was equipped with an optical angles-only sensor subject to realistic visibility constraints, and state estimation was performed using an Extended Kalman Filter (EKF). Sensor tasking decisions were generated by a myopic (greedy) scheduler optimizing different objective functions, and performance was assessed via complete and observation-only root-mean-square errors (RMSE) in position, together with coverage and fairness metrics.

Answer to Research Question 2: Observer orbit selection

The first major conclusion concerns the placement of the space-based cislunar observer. Comparing an observer in a cislunar orbit (Observer 3) with two Earth-centered orbits, a high-altitude polar orbit (HAPO) and a low Earth orbit (LEO), shows that a suitably chosen cislunar orbit decisively outperforms Earth-based observers for tracking objects in the cislunar region. For the same target catalog, the cislunar observer achieves substantially lower complete and observation-only position RMSE than both the HAPO and LEO observers. For example, the average observation position RMSE improves from 50.96 km (LEO) and 28.22 km (HAPO) to 5.83 km for the cislunar observer, while the complete position RMSE decreases from 157.12 km and 101.63 km to 61.98 km.

This improvement stems from the shorter average range to targets and more favorable visibility geometry in cislunar orbits, which reduce angular measurement errors and enable more frequent measurement updates.

Within cislunar space, the comparison of eight candidate observer orbits reveals that geometry, range, and occultation interact in non-trivial ways. Using a Kullback–Leibler divergence (KLD) reward, Observer 8, an L2 North Halo orbit, achieves the lowest complete position RMSE across the catalog, while Observer 3, a Distant Prograde Orbit (DPO), attains the lowest observation RMSE. Short-period orbits closer to the Moon (e.g. DPOs) provide excellent instantaneous estimation accuracy due to shorter ranges, but suffer from frequent lunar occultations that degrade end-to-end performance. More distant Lyapunov and Halo orbits deliver less accurate individual measurements but benefit from superior long-term visibility and more uniform coverage.

These results confirm and refine existing insights from the literature: Lyapunov, Distant Retrograde Orbits (DROs), and Halo orbits provide strong visibility for cislunar surveillance, but the performance of individual orbits within a family can differ significantly. Consequently, observer orbit design for cislunar SSA should not rely solely on orbit family selection but must consider specific orbital parameters such as period, stability index, and phase relative to the target population.

Answer to Research Question 1: Reward function performance

The second major conclusion concerns the relative performance of different reward functions for greedy sensor tasking. The following reward families were evaluated: three information-gain (IG) based rewards (Kullback–Leibler divergence (KLD), mutual information (MI), and Cauchy–Schwarz divergence (DCS)), an age-of-information (AOI) heuristic reward, and finite-time Lyapunov exponent (FTLE) based reward with varying horizons (FTLE1, FTLE2).

When the full catalog of 22 targets is used, the complete position RMSE for the IG-based rewards appears highly irregular, especially for observers 3 and 8. For example, with Observer 8, KLD yields a complete RMSE of 13.94 km, whereas MI and DCS exhibit much larger complete RMSE values (32.42 km and 42.16 km), despite all three rewards producing almost identical tasking decisions and very similar observation RMSE. This discrepancy is traced to Target 17, a very large, distant DRO that is rarely visible and therefore poorly observed. Due to its long range, Target 17 is too dim to be observed most of the time, causing its state uncertainty to grow and its error during prediction epochs to dominate the catalog-averaged complete RMSE.

To mitigate the outlier dominated scenario, a truncated catalog excluding Target 17 is considered. Under this configuration, the behavior of the reward functions is much more consistent with theoretical expectations. For both observers 3 and 8, the three IG-based rewards produce tightly clustered complete position RMSE values (about 610 km) and very similar observation RMSE (about 4.58.5 km), indicating that they indeed generate similar tasking and measurement quality. MI narrowly outperforms KLD and DCS across both observer configurations. AOI emerges as a strong, low-complexity baseline, with its complete RMSE typically within 1 – 2 km of the best IG-based configuration, and it offers very uniform coverage with low variance in the number of observations per target.

The FTLE-based reward performs substantially worse across all metrics. Even after removing Target 17, FTLE configurations exhibit higher mean and worst-case RMSE, allocate fewer observations on average, and distribute observations very unevenly, with a small number of targets receiving hundreds of measurements while many are rarely observed. The positive correlation between number of observations and RMSE for FTLE indicates that additional measurements are often spent on inherently difficult, distant targets without sufficiently improving their estimation accuracy.

Overall, the results lead to three key conclusions about reward functions:

Information-gain metrics (KLD, MI, DCS) provide the best overall performance and are largely interchangeable in a greedy framework, as long as performance is not dominated by extreme, sparsely observed targets.

Age-of-information offers a simple heuristic that yields robust, near-optimal average performance and particularly uniform coverage, making it an attractive choice when computational resources or modeling fidelity are limited.

FTLE-based rewards, at least in the configuration tested here, are not suitable as standalone tasking objectives for angles-only cislunar SSA, as they encourage observations at long ranges with poor geometry, leading to degraded measurement quality and fairness.

Role of extreme orbits and metric choice

The analysis of Target 17 and the truncated catalog underscores that performance conclusions can be highly sensitive to a small number of extreme orbits. Distant, dim, and rarely observed targets can dominate catalog-averaged metrics such as complete RMSE, obscuring the intrinsic behavior of tasking policies. Distinguishing between observation-only RMSE and complete RMSE, and complementing them with fairness measures across targets, proves essential to properly interpret results and avoid misleading conclusions about reward function performance.

Methodological contributions and limitations

Methodologically, this work provides a unified simulation framework to jointly study observer orbit placement and sensor tasking strategies under consistent CR3BP dynamics, visibility constraints, and estimation assumptions. It demonstrates how subtle interactions between geometry, measurement noise and filter parameters, and scheduling can produce counter-intuitive behaviors (e.g., higher observation

RMSE than complete RMSE for some observers), and highlights the importance of empirical evaluation rather than relying solely on geometric intuition.

At the same time, several limitations should be acknowledged:

The dynamical model is restricted to the CR3BP with no additional perturbations.

Only a single optical angles-only sensor and a single-sensor greedy scheduling policy are considered.

The simulation horizon is one synodic period, and station-keeping and operational constraints (e.g. propellant, pointing limits) are not explicitly modeled.

The target catalog includes only periodic orbits that remain relatively well observable over the simulation horizon, with highly unstable orbits removed to prevent numerical divergence.

Within these assumptions, the thesis addresses the research questions by showing that (i) observer placement in cislunar orbits is crucial to achieving effective SSA, and (ii) among reasonable reward functions, information-gain-based metrics and AOI provide strong performance, while FTLE-based rewards are ill-suited for greedy angles-only tracking in the cislunar regime.

7.2. Recommendations and Future Work

Building on the conclusions above, several recommendations can be formulated for the design of future cislunar SSA architectures and for extending this research.

Recommendations for cislunar SSA system design

Prioritize cislunar-based observers for cislunar SSA

The substantial improvement in both complete and observation RMSE relative to Earth-based observers supports the deployment of at least one dedicated cislunar sensor for catalog maintenance and custody of cislunar objects.

Treat orbit selection and sensor tasking as a coupled design problem

A lot of work that evaluates potential observer orbits focuses on viewing geometry. It was shown in this work that an observer with a superior viewing geometry may have poorer estimation due to its poor observational geometry. Observer orbits need to be evaluated in combination with potential sensor tasking algorithms with different estimation structures and tasking policies. Study of the evaluation of the observability angle Ψ_f alongside the relative range in sensor tasking simulations can further the understanding of the role of relative geometry better and aid orbit selection and the overall SSA objective.

Use multiple, complementary performance metrics

Evaluation of the observation RMSE in this work led to valuable insights into the coupling of the geometry, tasking and estimation. The study and evaluation of multiple metrics can help in the system design. Metrics such as observation RMSE, per-target statistics, worst-case errors, and fairness metrics should be jointly considered to aid in the design of a robust system.

Future work on orbit selection and environment modeling

Several extensions would strengthen and generalize the findings:

Expanded orbit catalogs and phase optimization

The analysis could be extended to a denser sampling of orbit families (NRHOs, DROs, Lyapunov and Halo families at multiple Jacobi constants and phases) to identify robust regions in the orbital parameter space rather than isolated “good” orbits. Incorporating phase-optimization, as suggested in previous work, would help understand how initial conditions within the same orbit affect long-term illumination and visibility.

More realistic dynamical models

Future studies could move beyond the CR3BP to extended models (e.g. elliptic restricted 3-body problem, n-body models with solar perturbations, non-spherical gravity) to better capture long-term behavior and stability requirements, especially for orbits with high stability index.

Inclusion of station-keeping and operational constraints

The stability index and station-keeping cost could be incorporated into orbit selection, trading tracking performance against propellant usage and operational complexity.

Finding complementary observer locations

Identifying observer orbits that complement each other in terms of viewing geometry and estimation performance is an important step toward effective cislunar SSA. In particular, combining cislunar space-based observers with ground-based assets on Earth and/or the Moon could provide more robust orbit estimation than any single sensor alone.

Future work on sensor tasking strategies

On the tasking side, several promising directions emerge:

Alternative reward formulations

The results highlight several opportunities for improving the reward design. First, purely information-gain-based rewards only consider the expected *change* in uncertainty due to a prospective measurement. As observed for Target 17, a target may re-enter the visibility region with very large a priori covariance after a long gap without observations, yet still be scheduled sparsely because other targets offer slightly larger expected reductions in covariance. A possible improvement is to weight the information gain by a scalar that represents the current uncertainty magnitude \bar{P}_k . This would bias the scheduler toward highly uncertain targets and reduce the risk that they are neglected even when observable.

Second, the reward could be augmented with realistic slewing costs and penalties for switching targets. Such terms would favor sequences of observations on the same target over consecutive decision epochs, allowing the filter to reduce the estimation error quickly while also reflecting physical constraints of the sensor.

Third, this work also showed the critical role of relative geometry between the observer and target in determining its estimation performance. Due to the immensity of cislunar space and the complexity of the trajectories of objects in this space, the geometry between two objects can vary dramatically over time. An investigation into the role of geometry through analysis of range, observability angle and observability matrix could provide an important insight into when a target is well conditioned for observation, and could help guide sensor tasking strategies. Study of the evaluation of the observability angle Ψ_f alongside the relative range in sensor tasking simulations can further the understanding of the role of relative geometry better and aid in the design of sensor tasking strategies.

Lastly, the poor performance of FTLE-based rewards suggests a need to re-examine dynamical sensitivity metrics as tasking objectives. In this work, the FTLE reward is biased toward distant targets. Thus, scaling the FTLE-based reward by an appropriate function of the observer–target range could mitigate this bias.

Multi-step algorithms

The current study focuses on myopic (single-step) policies. A natural extension is to investigate multi-step strategies such as Monte Carlo Tree Search, or reinforcement learning techniques within the same simulation framework. In particular, a multi-step tasking algorithm could explicitly account for upcoming visibility gaps. By predicting that a given target will soon become unobservable, the scheduler could allocate a higher number of observations while it is visible. This would be especially beneficial for challenging orbits with sparse visibility, such as Target 17, where a myopic policy fails to account for its upcoming long period of invisibility. Multi-step tasking could therefore reduce both average and worst-case errors by planning over a longer horizon instead of greedily selecting the best single-step information gain. Multi-step algorithms are common for SSA [54], but they could be particularly effective in cislunar SSA due to the large periods of invisibility of certain objects as well as the rapidly evolving geometry.

Multi-sensor architectures

Finally, extending the framework to multiple cooperating observers (e.g., a mix of cislunar and Earth-based sensors) would enable evaluation of distributed tasking algorithms and network-level performance, bridging the gap between single-sensor studies and operational SSA networks.

Validation with real or planned missions

Finally, future work should seek closer alignment with existing or planned missions. Incorporating realistic orbit and measurement profiles for systems such as Lunar Gateway, Artemis, or relay satellites, and using real or high-fidelity simulated measurement data, would help validate the modeling assumptions and assess how the proposed tasking strategies perform under operational constraints.

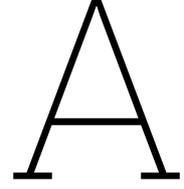
References

- [1] Brian Baker-McEvilly et al. “A comprehensive review on Cislunar expansion and space domain awareness”. In: *Progress in Aerospace Sciences* 147 (2024), p. 101019. DOI: 10.1016/j.paerosci.2024.101019. URL: <https://doi.org/10.1016/j.paerosci.2024.101019>.
- [2] Yaakov Bar-Shalom, X. Rong Li, and Thiagalingam Kirubarajan. *Estimation with Applications to Tracking and Navigation: Theory, Algorithms, and Software*. New York, NY: Wiley, 2001. ISBN: 978-0471416555.
- [3] Amelia L Batcha et al. “Artemis I trajectory design and optimization”. In: *2020 AAS/AIAA Astrodynamics Specialist Conference*. AAS 20-649. 2020.
- [4] Surabhi Bhadauria and Carolin Frueh. “Optical observation regions in cislunar space using the bi-circular restricted four body problem geometry”. In: *Proceedings of the 2022 Advanced Maui Optical and Space Surveillance Technologies Conference*. 2022.
- [5] Eric A Butcher, Jingwei Wang, and T Alan Lovell. “On Kalman filtering and observability in non-linear sequential relative orbit estimation”. In: *Journal of Guidance, Control, and Dynamics* 40.9 (2017), pp. 2167–2182.
- [6] Rita L Cognion. “Observations and modeling of GEO satellites at large phase angles”. In: *AMOS proceedings* (2013).
- [7] Kathleen Connor Howell. “Three-dimensional, periodic, ‘halo’ orbits”. In: *Celestial mechanics* 32.1 (1984), pp. 53–71.
- [8] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. 2nd ed. Hoboken, NJ: Wiley, 2006. ISBN: 978-0471241959.
- [9] Alexander Cropp et al. “Strategies for Cislunar SSA”. In: *Proceedings of the 9th European Conference on Space Debris*. Paper 225. European Space Agency (ESA). 2023. URL: <https://conference.sdo.esoc.esa.int/proceedings/sdc9/paper/225/SDC9-paper225.pdf>.
- [10] Kyle J. DeMars et al. “The Cauchy-Schwarz divergence for assessing situational information gain”. In: *2012 15th International Conference on Information Fusion*. 2012, pp. 1126–1133.
- [11] Arvind Eapen, Patrick Wilmer, and Carolin Frueh. “Sensor Tasking Strategies for Space-Based Observers in the Cislunar Environment”. In: *AIAA Scitech 2024 Forum*. American Institute of Aeronautics and Astronautics. 2024. DOI: 10.2514/6.2024-0964. URL: <https://doi.org/10.2514/6.2024-0964>.
- [12] European Space Agency. *About Space Debris*. https://www.esa.int/Space_Safety/Space_Debris/About_space_debris. Accessed: 2025-07-15. 2024.
- [13] European Space Agency. *ESA Space Environment Report 2025*. Tech. rep. Accessed: 2025-08-12. European Space Agency (ESA), 2025. URL: https://www.esa.int/Space_Safety/Space_Debris/ESA_Space_Environment_Report_2025.
- [14] Carter J Franz and Ryan P Russell. “Database of planar and three-dimensional periodic orbits and families near the moon”. In: *The Journal of the Astronautical Sciences* 69.6 (2022), pp. 1573–1612.
- [15] Carolin Frueh et al. “Cislunar space situational awareness”. In: *31st AIAA/AAS Space Flight Mechanics Meeting*. Vol. 144. 2021.
- [16] Peter Garretson and John M. Olson. *A Primer on Cislunar Space*. Technical Report PA2021-1271. Distribution A: Approved for public release. Air Force Research Laboratory (AFRL), 2021. URL: https://www.afrl.af.mil/Portals/90/Documents/RV/A%20Primer%20on%20Cislunar%20Space_Dist%20A_PA2021-1271.pdf?ver=vs6e0sE4PuJ51QC-15DEfg%3D%3D.

- [17] Giovanni Gasparini and Valérie Miranda. "Space situational awareness: an overview". In: *The Fair and Responsible Use of Space: An International Perspective*. Ed. by Wolfgang Rathgeber, Kai-Uwe Schrogl, and Ray A. Williamson. Vienna: Springer Vienna, 2010, pp. 73–87. ISBN: 978-3-211-99653-9. DOI: 10.1007/978-3-211-99653-9_7. URL: https://doi.org/10.1007/978-3-211-99653-9_7.
- [18] Baichun Gong et al. "Observability criterion of angles-only navigation for spacecraft proximity operations". In: *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering* 233.12 (2019), pp. 4302–4315.
- [19] MJ Gualdoni and Kyle J Demars. "Utilizing Information Statistics in Multi-Observation Sensor Tasking". In: (2018).
- [20] Loretta Hall. "The History of Space Debris". In: *Space Traffic Management Conference* (2014). URL: https://commons.erau.edu/stm/2014/thursday/19?%20utm_source=commons.erau.edu%2Fstm%2F2014%2Fthursday%2F19&utm_medium=PDF&utm_campaign=PDFCoverPages.
- [21] Alfred Olivier Hero et al. *Foundations and applications of sensor management*. Springer Science & Business Media, 2007.
- [22] Marcus Holzinger et al. "Sensor Tasking via Monte Carlo Tree Search in the Cislunar CR3BP". In: *AAS/AIAA Astrodynamics Specialist Conference* (2021).
- [23] Robert Jenssen et al. "The Cauchy–Schwarz divergence and Parzen windowing: Connections to graph theory and Mercer kernels". In: *Journal of the Franklin Institute* 343.6 (2006), pp. 614–629. ISSN: 0016-0032. DOI: 10.1016/j.jfranklin.2006.04.007.
- [24] Kaitlyn Johnson. *Fly Me to the Moon: Worldwide Cislunar and Lunar Missions*. <https://www.csis.org/analysis/fly-me-moon-worldwide-cislunar-and-lunar-missions>. Accessed: 2025-08-12. Feb. 2022.
- [25] *JPL Three-Body Periodic Orbit Catalog*. https://ssd.jpl.nasa.gov/tools/periodic_orbits.html#/periodic. Accessed: 18 August 2025; part of NASA Jet Propulsion Laboratory Solar System Dynamics.
- [26] Kamil Kampa, Enis Hasanbelliu, and Jose C. Principe. "Closed-form Cauchy–Schwarz PDF divergence for mixture of Gaussians". In: *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*. IEEE. 2011, pp. 2572–2579. DOI: 10.1109/IJCNN.2011.6033341.
- [27] Tian Kangsheng and Zhu Guangxi. "Sensor management based on fisher information gain". In: *Journal of Systems Engineering and Electronics* 17.3 (2006), pp. 531–534.
- [28] Donald J. Kessler and Burton G. Cour-Palais. "Collision frequency of artificial satellites: The creation of a debris belt". In: *Journal of Geophysical Research: Space Physics* 83.A6 (1978), pp. 2637–2646. DOI: 10.1029/JA083iA06p02637.
- [29] Chang Liu and Lixian Dong. "Stabilization of Lagrange points in circular restricted three-body problem: A port-Hamiltonian approach". In: *Physics Letters A* 383 (2019). Section A: General, Atomic and Solid State Physics, pp. 1907–1914. DOI: 10.1016/j.physleta.2019.03.033. URL: <https://doi.org/10.1016/j.physleta.2019.03.033>.
- [30] Martin Lo and Shane Ross. *The lunar L1 gateway: portal to the stars and beyond*. Pasadena, CA: Jet Propulsion Laboratory, National Aeronautics and Space ..., 2001.
- [31] M. Losacco and et al. "Differential Algebra-Based Orbit Determination for Optical and Radar Measurements". In: *Proceedings of the 8th ESOC ESA Space Debris Conference*. Eqs. (19)–(20) give topocentric RA/Dec: $\alpha = \text{atan2}(Y, X)$, $\delta = \arcsin(Z/\rho)$. 2021. URL: <https://conference.sdo.esoc.esa.int/proceedings/sdc8/paper/241/SDC8-paper241.pdf>.
- [32] Nicholas Moretti, Mark Rutren, and Travis Bessell. "Space Situational Awareness Sensor Tasking: A Comparison Between Step-Scan Tasking and Dynamic, Real-Time Tasking". In: *2018 21st International Conference on Information Fusion (FUSION)*. 2018, pp. 1339–1344. DOI: 10.23919/ICIF.2018.8455851.
- [33] Carl D. Murray and Stanley F. Dermott. *Solar System Dynamics*. Cambridge: Cambridge University Press, 1999.

- [34] Lauri K. Newman, Ryan Frigm, and David McKinley. "It's Not a Big Sky After All: Justification for a Close Approach Prediction and Risk Assessment Process". In: *2009 AAS/AIAA Astrodynamics Specialist Conference*. AAS 09-369. AAS Paper No. AAS 09-369. 2009.
- [35] Smriti Nandan Paul, Siwei Fan, and Igor Panfil. "Advancing Cislunar Space Domain Awareness Through Robust Optimization Framework for Optical Sensors-Based Autonomous Satellite Systems". In: *AIAA SCITECH 2025 Forum*. 2025, p. 1218.
- [36] Jose C Principe. *Information theoretic learning: Renyi's entropy and kernel perspectives*. Springer Science & Business Media, 2010.
- [37] Friedrich Pukelsheim. *Optimal Design of Experiments*. Philadelphia, PA: SIAM, 2006. ISBN: 978-0898716047. DOI: 10.1137/1.9780898719109.
- [38] Francois Sanson and Carolin Frueh. "Noise estimation and probability of detection in non-resolved images: application to space object observation". In: *Advances in Space Research* 64.7 (2019), pp. 1432–1444.
- [39] Hanspeter Schaub and John L. Junkins. *Analytical Mechanics of Space Systems, 4th Ed*. AIAA Educational Series, 2018. DOI: 10.2514/4.867231. URL: <https://doi.org/10.2514/4.867231>.
- [40] Jan Siminski. "Object correlation and orbit determination for geostationary satellites using optical measurements". PhD thesis. Universitätsbibliothek der Universität der Bundeswehr München, 2016.
- [41] Byron D. Tapley, Bob E. Schutz, and George H. Born. *Statistical Orbit Determination*. Amsterdam: Academic Press / Elsevier, 2004. ISBN: 978-0-12-683630-1.
- [42] Tudat Development Team. *Tudat Documentation*. <https://docs.tudat.space/en/latest/>. Accessed: 2025-09-12. 2025.
- [43] Jacob K. Vendl and Marcus J. Holzinger. "Cislunar Periodic Orbit Analysis for Persistent Space Object Detection Capability". In: *Journal of Spacecraft and Rockets* 58.4 (2021), pp. 1174–1185. DOI: 10.2514/1.A34909. URL: <https://doi.org/10.2514/1.A34909>.
- [44] J. L. Williams. "Information Theoretic Sensor Management". PhD thesis. MIT, 2007. URL: https://willsky.lids.mit.edu/ssg/ssg_theses/ssg_theses_2000_2009/Williams_PhD_2_07.pdf.
- [45] Jacob Williams et al. "Targeting cislunar near rectilinear halo orbits for human space exploration". In: *AAS/AIAA Space Flight Mechanics Meeting*. JSC-CN-38615. 2017.
- [46] Patrick S. Williams, David B. Spencer, and Richard S. Erwin. "Coupling of Estimation and Sensor Tasking Applied to Satellite Tracking". In: *Journal of Guidance, Control, and Dynamics* 36.4 (2013), pp. 993–1007. DOI: 10.2514/1.60496.
- [47] Adam P Wilmer, Robert A Bettinger, and Bryan D Little. "Cislunar periodic orbits for Earth–Moon L 1 and L 2 Lagrange point surveillance". In: *Journal of Spacecraft and Rockets* 59.6 (2022), pp. 1809–1820.
- [48] Adam P Wilmer, Nathan R Boone, and Robert A Bettinger. "Debris propagation and spacecraft survivability assessment for catastrophic mishaps occurring in CisLunar periodic orbits". In: *Journal of Space Safety Engineering* 9.2 (2022), pp. 207–222.
- [49] Adam P Wilmer et al. "Sun–Venus CR3BP, part 1: periodic orbit generation, stability, and mission investigation". In: *Archive of Applied Mechanics* 94.4 (2024), pp. 921–941.
- [50] Adam P. Wilmer and Robert A. Bettinger. "Lagrangian dynamics and the discovery of cislunar periodic orbits". In: *Nonlinear Dynamics* 111.1 (2023). U.S. Government work; not under copyright protection in the US, pp. 155–178. DOI: 10.1007/s11071-022-07829-1. URL: <https://doi.org/10.1007/s11071-022-07829-1>.
- [51] Adam P. Wilmer et al. "Preliminary investigation and proposal of periodic orbits and their utilization for logistics in the cislunar regime". In: *Space Policy* 71 (2025), p. 101635. DOI: 10.1016/j.spacepol.2024.101635. URL: <https://doi.org/10.1016/j.spacepol.2024.101635>.
- [52] Trevor Nathaniel Wolf. "Advancements in space object detection, tracking, and management using information-theoretic approaches". PhD thesis. 2025.

-
- [53] James R. Wright et al. *Orbit Determination Tool Kit: Theory & Algorithms*. Tech. rep. See Sec. 24.3 for RA/Dec (unit range vector) formulas. Analytical Graphics, Inc. (AGI), 2020. URL: <https://help.agi.com/odtk/ODTK/references/MathSpec.pdf>.
- [54] Chenbao Xue et al. “Review of sensor tasking methods in Space Situational Awareness”. In: *Progress in Aerospace Sciences* 147 (2024), p. 101017. DOI: 10.1016/j.paerosci.2024.101017.
- [55] Lihua Zhang. “Development and prospect of Chinese lunar relay communication satellite”. In: *Space: Science & Technology* (2021).
- [56] Renyong Zhang. “A review of periodic orbits in the circular restricted three-body problem”. In: *Journal of Systems Engineering and Electronics* 33.3 (2022), pp. 612–646. DOI: 10.23919/JSEE.2022.000059.
- [57] Yu Zhang et al. “A Review of Sensor Scheduling and Tasking for Large-Scale Satellite Constellations”. In: *Advances in Space Research* (2021).



RMSE Computation

To evaluate the accuracy of the filter, the root-mean-square error (RMSE) is computed for position and velocity. Let $x_k^{\text{true}} = (r_k^{\text{true}}, v_k^{\text{true}})$ denote the true state at time index k , and let $\hat{x}_k = (\hat{r}_k, \hat{v}_k)$ be the corresponding estimated state. The instantaneous position and velocity errors are defined as

$$e_{r,k} = \|r_k^{\text{true}} - \hat{r}_k\|_2, \quad (\text{A.1})$$

$$e_{v,k} = \|v_k^{\text{true}} - \hat{v}_k\|_2. \quad (\text{A.2})$$

For a given observer–target pair, we distinguish between epochs at which the target is observed and epochs at which the filter only propagates the state. Let \mathcal{O} denote the set of observed time indices and \mathcal{K} the set of all indices under consideration.

The *observed RMSE*, which measures estimation accuracy at update times, is computed as

$$\text{RMSE}_r^{(\text{obs})} = \sqrt{\frac{1}{|\mathcal{O}|} \sum_{k \in \mathcal{O}} e_{r,k}^2}, \quad (\text{A.3})$$

$$\text{RMSE}_v^{(\text{obs})} = \sqrt{\frac{1}{|\mathcal{O}|} \sum_{k \in \mathcal{O}} e_{v,k}^2}. \quad (\text{A.4})$$

The *complete RMSE*, which reflects end-to-end performance including propagation gaps, is defined as

$$\text{RMSE}_r^{(\text{com})} = \sqrt{\frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} e_{r,k}^2}, \quad (\text{A.5})$$

$$\text{RMSE}_v^{(\text{com})} = \sqrt{\frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} e_{v,k}^2}. \quad (\text{A.6})$$

These metrics allow comparison of filter performance both during measurement updates and across the full trajectory.

B

Project Plan

This section shows the project plan created at the end of the literature review part of the thesis.

Methods

The data collection for this thesis will be entirely through simulations created in TUDAT. There is limited public data about any current objects in the cislunar space, and future missions have undecided orbits and/or undisclosed orbit locations. However, there is a fair amount of research into which orbits are suitable for different types of missions. These 'predictions' can be used to create an environment with several objects to be observed.

Once an environment has been created, a single observer will be placed on different cislunar orbits to compare their merits for cislunar observation, and each observer will be tested with different reward functions to conduct a greedy tracking objective.

The performance across the different orbits and reward functions will be compared using consistent metrics such as estimation error. Analysis of observers based on their viewing geometry will also be done.

Tools and Data Collection

TUDAT will be used as the platform to create the environment, place the observer as well as implement the algorithm and simulate observations. These simulated observations will be used as the data for this thesis. TUDAT is an astrodynamics toolbox which provides several built in functions useful for creating the environment, propagating bodies as well as conducting observations. JPL's Periodic CR3BP database will be used for obtaining the initial states of orbits.

Expected Results

It is unlikely that a single observer orbit and reward function combination will provide significantly superior observation capabilities than all the others. The expected outcome is to show that certain sensor tasking functions and observer orbits show consistent advantages, which could then be used as an initial starting point to create an observer network with a more advanced sensing algorithm.

Planning

The thesis has been divided into multiple phases, as shown below. Each phase ends with a deliverable that will be reviewed and approved by the thesis supervisor.

Literature Review (Week 1 - Week 7)

- Identify the need for research into cislunar Space Situational Awareness.
- Overview of cislunar dynamics and orbit families

- Overview of sensor tasking functions

Research Phase I (Week 8 - Week 18)

- Conduct an observation of an object in Earth's orbit using a ground based sensor.
- Create an observer model and implement a greedy algorithm to track multiple objects in Earth orbit.
- Begin creating an environment of objects to be tracked in cislunar space.
- Finalize the different orbit families and reward functions to be tested.
- Create observer in cislunar orbit and obtain observations.
- Vary reward function and/or observer orbit to obtain dataset for comparison.

Research Phase II (Week 19 - Week 29)

- Finalize any changes to the base simulation environment/algorithms and observer.
- Run simulation for different combinations of observer/reward function and collect data.
- Begin data analysis
- Work on writing thesis

Research Dissertation (Week 30 - 35)

- Finish writing the thesis
- Green light review
- Prepare for thesis defence

The planned thesis schedule can be seen in the Gantt chart in Figure B.1.

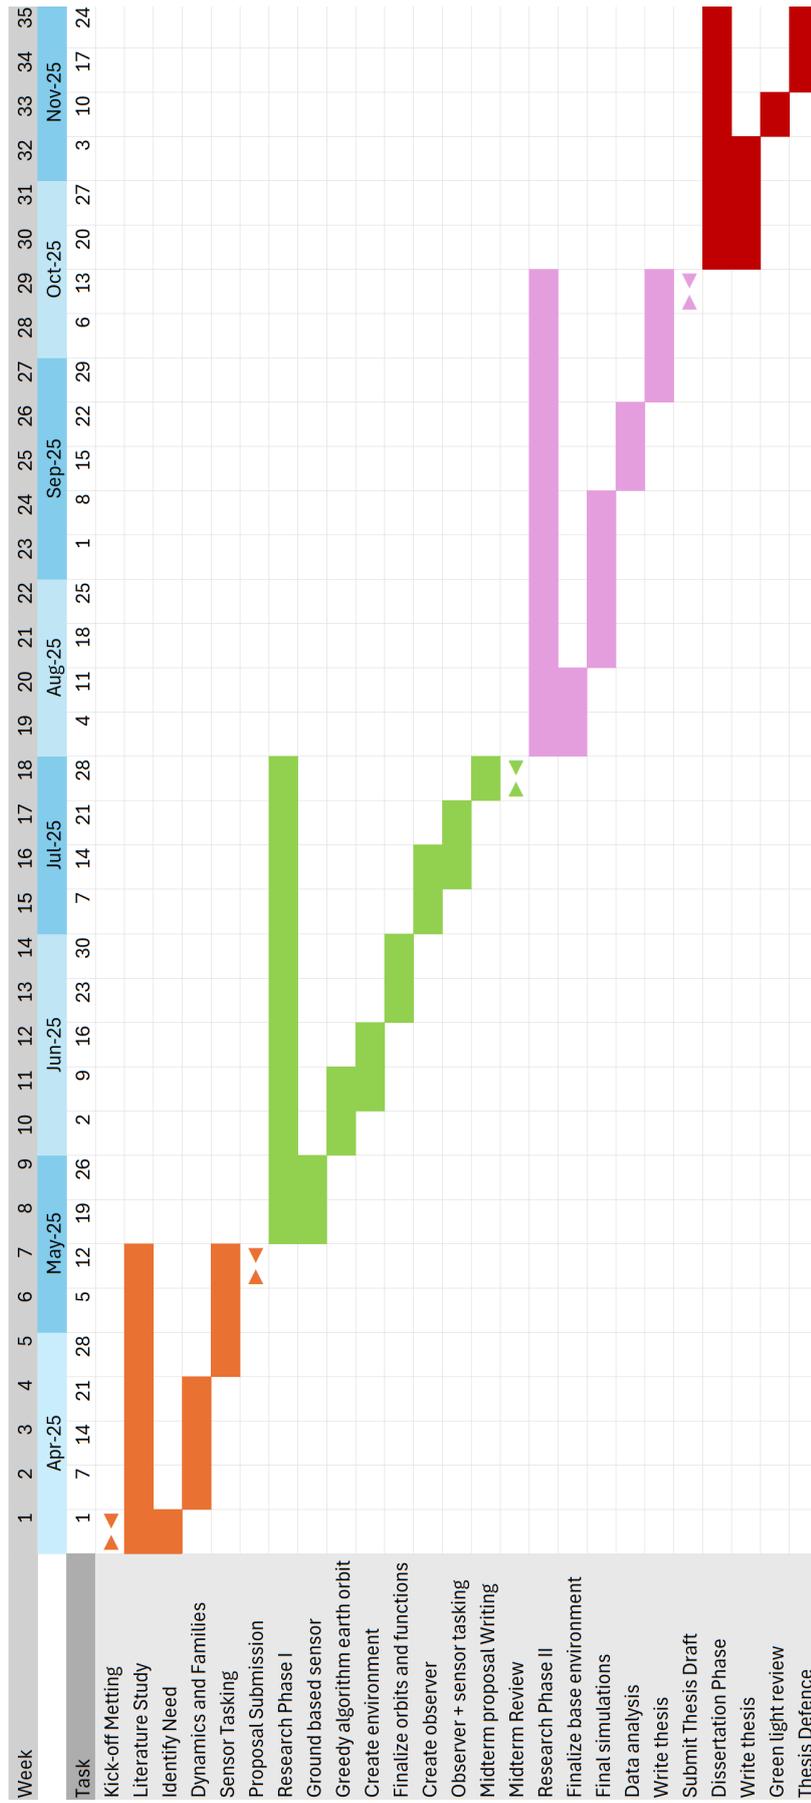


Figure B.1: Gantt chart of initial planning