

INFORMATION-THEORETIC SENSOR TASKING FOR OPTIMAL SPACE OBJECT CUSTODY

Brighton N. Smith* and Kyle J. DeMars†

Increased congestion across orbital regimes and the growing heterogeneity of sensor networks demand adaptive tasking strategies to ensure effective space domain awareness. This work investigates an information-theoretic approach to sensor scheduling, using the expected Kullback–Leibler divergence between prior and posterior state distributions as the optimization objective. A forecasted formulation is employed, in which measurement utility is evaluated at a common reference time, enabling consistent prioritization across sensing modalities and asynchronous observation timelines. The framework is demonstrated in two representative scenarios: a single-target, single-observer case in cislunar space and a multitarget, multi-observer problem in low Earth orbit, showcasing its flexibility, scalability, and applicability to complex surveillance networks.

INTRODUCTION

As the population of space objects (SOs) grows across orbital regimes – from crowded low Earth orbit (LEO) to the expanding cislunar domain – efficient sensor tasking has become essential for space domain awareness (SDA). Tracking diverse targets such as satellites, debris, and uncooperative objects requires intelligent allocation of limited sensing resources. Traditional methods often rely on fixed schedules or heuristic policies that struggle to adapt to evolving mission priorities like detection, classification, or high-accuracy tracking.¹

Information-theoretic frameworks offer a principled alternative by directly optimizing expected information gain, independent of specific filter designs or data association rules.¹ This work adopts the expected Kullback-Leibler (KL) divergence between prior and posterior target distributions as the tasking objective, enabling selection of measurements that most effectively reduce uncertainty within operational constraints.

Extending the approach in Reference 2, each candidate measurement is forecasted to a common reference time, allowing consistent evaluation across asynchronous, heterogeneous sensor networks. This common-epoch formulation supports batch updates and accommodates various sensing modalities. The reference time can be chosen strategically – for instance, before a visibility gap or expected maneuver – shifting the objective from immediate gain to long-term uncertainty reduction. In doing so, the forecasted framework addresses the limitations of myopic strategies, such as those explored in Reference 3, and enables coordinated, information-driven tasking across complex space surveillance networks.

*Undergraduate Researcher, Department of Aerospace Engineering, Texas A&M University, College Station, TX 77843

†Associate Professor, Department of Aerospace Engineering, Texas A&M University, College Station, TX 77843

The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government

INFORMATION ENTROPY AND DIVERGENCES

To understand the usage of information divergence as a tasking objective, it is first imperative to define the concept of *information entropy*, which is a measure of the uncertainty or spread of a random variable.⁴ Consider a continuous random vector \mathbf{x} defined on the event space \mathcal{X} with probability density function (pdf) $p(\mathbf{x})$. The information entropy, commonly referred to as Shannon entropy, is defined by the quantity $H[\cdot]$ and is given by⁴

$$H[p] = - \int_{\mathcal{X}} p(\mathbf{x}) \log p(\mathbf{x}) d\mathbf{x}. \quad (1)$$

This value quantifies the inherent uncertainty in the random variable's outcome, which bounds the maximum amount of information that could be gained through observation.² A higher entropy value indicates greater uncertainty, while lower entropy signifies a more predictable distribution.

It is often useful to quantify how different two distributions are from one another, leading to the notion of an *information divergence*, which measures the discrepancy between two probability distributions.⁵ Considering two pdfs $p(\mathbf{x})$ and $q(\mathbf{x})$, the information divergence, written as $D[q \parallel p]$, provides an asymmetric measure of how well p approximates q . Unlike symmetric metrics that measure distance in a geometric sense, information divergences reflect directional information loss, often with one distribution interpreted as the true source and the other as an approximation.⁴

Considering an n_x -dimensional random target state vector $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^{n_x}$, the *Kullback-Leibler (KL) divergence* from $q(\mathbf{x})$ to $p(\mathbf{x})$ can be expressed as

$$D_{\text{KL}}[q \parallel p] = H[q, p] - H[q] = \int_{\mathcal{X}} q(\mathbf{x}) \log \frac{q(\mathbf{x})}{p(\mathbf{x})} d\mathbf{x}, \quad (2)$$

where $H[q]$ is the Shannon entropy of $q(\mathbf{x})$, and $H[q, p]$ is the cross-entropy, defined as

$$H[q, p] = - \int_{\mathcal{X}} q(\mathbf{x}) \log p(\mathbf{x}) d\mathbf{x}.$$

This formulation highlights that the KL divergence quantifies the additional information required to represent q using p .⁴

When considering two multivariate Gaussian distributions, defined as

$$\begin{aligned} p(\mathbf{x}) &\triangleq p_g(\mathbf{x}; \mathbf{m}_x^-, \mathbf{P}_{xx}^-) \\ q(\mathbf{x}) &\triangleq p_g(\mathbf{x}; \mathbf{m}_x^+, \mathbf{P}_{xx}^+), \end{aligned}$$

a closed-form solution to Eq. (2) exists as

$$\begin{aligned} D_{\text{KL}}[q \parallel p] &= \frac{1}{2} \left[\log |\mathbf{P}_{xx}^-(\mathbf{P}_{xx}^+)^{-1}| + \text{tr}\{(\mathbf{P}_{xx}^-)^{-1} \mathbf{P}_{xx}^+\} \right. \\ &\quad \left. + (\mathbf{m}_x^+ - \mathbf{m}_x^-)^T (\mathbf{P}_{xx}^-)^{-1} (\mathbf{m}_x^+ - \mathbf{m}_x^-) - n_x \right]. \end{aligned} \quad (3)$$

THE KL DIVERGENCE AS A TASKING OBJECTIVE

To develop the KL divergence as an objective for sensor tasking, consider space objects whose motions follow the continuous-time dynamical system⁶

$$\dot{\mathbf{x}}(t_k) = \mathbf{f}(\mathbf{x}(t_k)),$$

where $\mathbf{x}_k \triangleq \mathbf{x}(t_k)$ is the state of the target at time t_k , defined on the state space $\mathcal{X} \subseteq \mathbb{R}^{n_x}$, and the function $\mathbf{f} : \mathcal{X} \rightarrow \mathcal{X}$ represents the deterministic system dynamics. This formulation omits process noise and assumes no control inputs, making it suitable for non-maneuvering objects whose trajectories can be well-approximated by known physics. Such assumptions are generally appropriate for cataloged space object tracking.⁶

Additionally, measurements $\mathbf{z}_k \triangleq \mathbf{z}(t_k)$ at time t_k defined on the measurement space $\mathcal{Z} \subseteq \mathbb{R}^{n_z}$ are generated according to the discrete-time process²

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{v}_k, \quad (4)$$

where $\mathbf{v}_k \in \mathcal{Z}$ is zero-mean, Gaussian white noise with symmetric and positive definite covariance $\mathbf{P}_{vv,k}$, and $\mathbf{h} : \mathcal{X} \rightarrow \mathcal{Z}$ is the measurement function.

Myopic Objective

The myopic objective seeks to quantify the information gain from a single measurement in terms of how much it reduces the uncertainty in the estimated target state. This is done by evaluating the KL divergence between the *a priori* and *a posteriori* Gaussian state densities, as described in standard Kalman filtering treatments.^{7,8}

Let $\mathbf{m}_{x,k}^-$ and $\mathbf{P}_{xx,k}^-$ denote the mean and covariance of the *a priori* estimate at time t_k , and let $\mathbf{m}_{x,k}^+$ and $\mathbf{P}_{xx,k}^+$ be the corresponding *a posteriori* quantities after a Kalman update. Assuming that the state estimates are described by Gaussian pdfs, the same closed-form solution developed in Eq. (3) can be applied. Substitution of the standard Kalman mean update equation given in Reference 7 yields^{2,9}

$$D_{\text{KL}}[q \parallel p] = \frac{1}{2} \left[\log |\mathbf{P}_{xx,k}^+ (\mathbf{P}_{xx,k}^+)^{-1}| + \text{tr}\{(\mathbf{P}_{xx,k}^-)^{-1} \mathbf{P}_{xx,k}^+\} \right. \\ \left. + (\mathbf{z}_k - \mathbf{m}_{z,k}^-)^T \mathbf{K}_k^T (\mathbf{P}_{xx,k}^-)^{-1} \mathbf{K}_k (\mathbf{z}_k - \mathbf{m}_{z,k}^-) - n_x \right], \quad (5)$$

where \mathbf{z}_k is the measurement taken at time t_k , $\mathbf{m}_{z,k}^-$ is the expected measurement at the same time, \mathbf{K}_k is the Kalman gain, and n_x is the dimension of the state vector. By simultaneously considering the measurement model and the uncertainty in the state estimation itself, this expression gives the relative impact an observation has on the state estimate density.

In practice, Eq. (5) requires a measurement, \mathbf{z}_k , in order to determine the information gain that is realized through the update; as the objective is to determine when to take a measurement, Eq. (5) is not immediately applicable. Instead, it is common practice to consider the average (or expected) information gain that would be realized by a measurement taken at t_k . This is done by computing the expected KL divergence, otherwise known as the first moment of the KL divergence.¹⁰ This first moment can be computed by taking the

expected value of Eq. (5) with respect to the measurement pdf, $p(\mathbf{z}_k)$, giving

$$\begin{aligned} M_{\text{KL}}(t_k) &= \mathbb{E}\{D_{\text{KL}}[q \parallel p]\} \\ &= \int_{\mathcal{Z}} D_{\text{KL}}[q \parallel p]p(\mathbf{z}_k) \, d\mathbf{z}_k \\ &= \frac{1}{2} \log |\mathbf{P}_{xx,k}^-(\mathbf{P}_{xx,k}^+)^{-1}|, \end{aligned} \quad (6)$$

where the argument t_k is the time at which measurement \mathbf{z}_k is “taken.”

This expression is a compact and computationally inexpensive method for assessing the impact of measurements without requiring a measurement to be acquired, as opposed to Eq. (5). This approach is hereafter referred to as “myopic” to signify the emphasis on the information gain at each of the t_k observation times as opposed to at a downstream or upstream time.

Given a set of M candidate measurement time indices $\zeta = \{\delta_1, \dots, \delta_M\}$, the cumulative myopic objective is defined as

$$\begin{aligned} \arg \max_{\zeta} J(\zeta) &\triangleq \arg \max_{\zeta} \sum_{\delta \in \zeta} M_{\text{KL}}(t_{\delta}) \\ &= \arg \max_{\zeta} \frac{1}{2} \sum_{\delta \in \zeta} \log |\mathbf{P}_{xx,\delta}^-(\mathbf{P}_{xx,\delta}^+)^{-1}|. \end{aligned} \quad (7)$$

For a multitarget scenario with N independent targets, the multitarget myopic objective simply accumulates the single-target contributions as

$$\mathcal{J}(\zeta) \approx \sum_{i=1}^N J(\zeta^{(i)}), \quad (8)$$

where $\zeta^{(i)}$ denotes the measurement time indices for the i^{th} target.

Forecasted Objective

While the myopic objective offers a tractable and interpretable tasking objective, it suffers from two key limitations: a bias toward later measurements due to growing uncertainty, and difficulty in comparing heterogeneous observations across time, sensors, and modalities.^{2,10} As uncertainty accumulates over time, the myopic KL divergence tends to prioritize delayed measurements, potentially overlooking earlier data that could have significantly improved estimation. Additionally, because each measurement is evaluated in isolation at its acquisition time, there is no consistent mechanism to compare observations from different sensors or epochs.

To address these limitations, this work adopts a forecasting approach that maps the impact of each observation to a common reference time. By evaluating expected information gain at a strategically chosen epoch, the method enables consistent comparison across measurement types and prioritizes observations based on long-term utility rather than immediate uncertainty reduction.²

Forecasting information to a common reference time naturally facilitates a batch update formulation, allowing multiple measurements to be processed simultaneously within a unified estimation framework. Given a measurement acquired at time t_k , a mapped Jacobian is defined as

$$\tilde{\mathbf{H}}_k = \mathbf{H}_{x,k} \Phi(t_k, t_{\text{ref}}), \quad (9)$$

where t_{ref} is the reference epoch at which all measurements are analyzed, $\mathbf{H}_{x,k} = \mathbf{H}_x(\mathbf{m}_{x,k}^-)$ represents the Jacobian of the measurement model evaluated at the *a priori* mean, and $\Phi(t_k, t_{\text{ref}})$ is the state transition matrix that propagates the state from t_{ref} to t_k .

For a set of M measurements acquired at times t_1, t_2, \dots, t_M , the batch update structure is constructed by concatenating the mapped measurement Jacobians as

$$\bar{\mathbf{H}} = \begin{bmatrix} \tilde{\mathbf{H}}_1 \\ \tilde{\mathbf{H}}_2 \\ \vdots \\ \tilde{\mathbf{H}}_M \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{x,1} \Phi(t_1, t_{\text{ref}}) \\ \mathbf{H}_{x,2} \Phi(t_2, t_{\text{ref}}) \\ \vdots \\ \mathbf{H}_{x,M} \Phi(t_M, t_{\text{ref}}) \end{bmatrix}.$$

Additionally, assuming the measurement noise is uncorrelated in time, the corresponding covariance matrices for all observations can be consolidated into a block diagonal matrix,

$$\bar{\mathbf{P}}_{vv} = \text{diag}(\mathbf{P}_{vv,1}, \mathbf{P}_{vv,2}, \dots, \mathbf{P}_{vv,M}). \quad (10)$$

This formulation allows for the inclusion of heterogeneous measurement sources, as the individual Jacobians $\mathbf{H}_{x,k}$ need not be derived from the same measurement model, and the noise covariance matrices $\mathbf{P}_{vv,k}$ may vary across sensors and observation conditions. By treating all measurements in a batch framework, the covariance update at t_{ref} is performed as¹¹

$$(\mathbf{P}_{xx,\text{ref}}^+)^{-1} = (\mathbf{P}_{xx,\text{ref}}^-)^{-1} + \bar{\mathbf{H}}^T \bar{\mathbf{P}}_{vv}^{-1} \bar{\mathbf{H}}, \quad (11)$$

where $\mathbf{P}_{xx,\text{ref}}^-$ is the *a priori* covariance propagated to the reference time, given by

$$\mathbf{P}_{xx,\text{ref}}^- = \Phi(t_{\text{ref}}, t_0) \mathbf{P}_{xx,0} \Phi^T(t_{\text{ref}}, t_0),$$

where $\mathbf{P}_{xx,0}$ is the initial covariance. This batch processing formulation enables the assimilation of multiple temporally distributed observations into a single, consistent state update at t_{ref} .

The myopic form of the mean KL divergence in Eq. (6) can be reformulated in its forecasted form as²

$$\begin{aligned} \bar{M}_{\text{KL}}(\mathbf{t}) &= \frac{1}{2} \log |\mathbf{P}_{xx,\text{ref}}^- (\mathbf{P}_{xx,\text{ref}}^+)^{-1}| \\ &= \frac{1}{2} \log |\mathbf{I}_{n_x} + \mathbf{P}_{xx,\text{ref}}^- \bar{\mathbf{H}}^T \bar{\mathbf{P}}_{vv}^{-1} \bar{\mathbf{H}}|, \end{aligned} \quad (12)$$

where \mathbf{t} is the vector of measurement times, \mathbf{I}_{n_x} is an n_x -dimensional identity matrix, and the overbar notation denotes that the KL divergence is evaluated at the reference epoch t_{ref} .

Again, given a set of M candidate measurement time indices $\zeta = \{\delta_1, \dots, \delta_M\}$, the forecasted tasking objective for a single target is written as²

$$\begin{aligned} \arg \max_{\zeta} \bar{J}(\zeta) &\triangleq \arg \max_{\zeta} \bar{M}_{\text{KL}}(\mathbf{t}_{\zeta}) \\ &= \arg \max_{\zeta} \frac{1}{2} \log |\mathbf{I}_{n_x} + \mathbf{P}_{xx,\text{ref}}^{-} \tilde{\mathbf{H}}^T \tilde{\mathbf{P}}_{vv}^{-1} \tilde{\mathbf{H}}|, \end{aligned} \quad (13)$$

where $\mathbf{t}_{\zeta} \triangleq \mathbf{t}(\zeta)$ denotes the measurement times selected by the schedule. For multitarget scenarios, the forecasted objective extends naturally as²

$$\bar{\mathcal{J}}(\zeta) \approx \sum_{i=1}^N \bar{J}(\zeta^{(i)}), \quad (14)$$

where individual target objectives are aggregated while preserving the reference time framework.

OPTIMIZATION

The objectives developed in the preceding section offer a principled tasking criterion that maximizes the expected information gain at each measurement instant (Eqs. (7) and (8)) or considering a common reference time (Eqs. (13) and (14)). However, these objectives are in general non-convex, posing challenges for global optimization.

To demonstrate this, consider the forecasted mean KL divergence from Eq. (12) evaluated as a continuous function of time for a single measurement,

$$\bar{M}_{\text{KL}}(t) = \frac{1}{2} \log |\mathbf{I}_{n_x} + \mathbf{P}_{xx,\text{ref}}^{-} \tilde{\mathbf{H}}^T \tilde{\mathbf{P}}_{vv}^{-1} \tilde{\mathbf{H}}|, \quad (15)$$

where $\tilde{\mathbf{H}}$ is a single mapped measurement Jacobian as in Eq. (9), $\mathbf{P}_{vv} \succ 0$ is a single measurement noise covariance matrix, and $\mathbf{P}_{xx,\text{ref}}^{-} \succ 0$ is the initial prior covariance propagated to the reference time. Define the matrix

$$\mathbf{A}(t) \triangleq \mathbf{I}_{n_x} + \mathbf{P}_{xx,\text{ref}}^{-} \tilde{\mathbf{H}}^T \tilde{\mathbf{P}}_{vv}^{-1} \tilde{\mathbf{H}},$$

which is guaranteed to be positive definite due to the similarity form of the information matrix, $\tilde{\mathbf{H}}^T \tilde{\mathbf{P}}_{vv}^{-1} \tilde{\mathbf{H}} \succeq 0$. Since the log-determinant function is twice differentiable over the cone of positive definite matrices,¹² the second derivative of Eq. (15) can be expressed as

$$\begin{aligned} \frac{d^2}{dt^2} \bar{M}_{\text{KL}}(t) &= \frac{d}{dt} \left(\frac{1}{2} \text{tr} \left\{ \mathbf{A}^{-1} \frac{d\mathbf{A}}{dt} \right\} \right) \\ &= \frac{1}{2} \text{tr} \left\{ \mathbf{A}^{-1} \frac{d^2 \mathbf{A}}{dt^2} - \mathbf{A}^{-1} \frac{d\mathbf{A}}{dt} \mathbf{A}^{-1} \frac{d\mathbf{A}}{dt} \right\}. \end{aligned}$$

The first term may be positive or negative depending on the curvature of the Jacobian and state transition matrix, while the second term can be rewritten as the squared Frobenius norm of a whitened matrix,¹³

$$\text{tr} \left\{ \mathbf{A}^{-1} \frac{d\mathbf{A}}{dt} \mathbf{A}^{-1} \frac{d\mathbf{A}}{dt} \right\} = \left\| \mathbf{A}^{-1/2} \frac{d\mathbf{A}}{dt} \mathbf{A}^{-1/2} \right\|_F^2.$$

Since the second term is always non-negative and subtractive and the first term is not guaranteed to dominate it in sign or magnitude, the second derivative may be negative in some regions and positive in others. This confirms that the function is neither convex nor concave.

To illustrate this visually, consider a toy example in which a ground-based radar sensor selects a single measurement time for a geostationary target over a 19-hour period. The geometry of the forecasted KL divergence objective, shown in Figure 1, reveals the presence of multiple local maxima and varying curvature, corroborating the theoretical result. Although this result is derived for the forecasted objective, similar conclusions apply to the myopic objective, which has also been explained to be non-convex in general.²

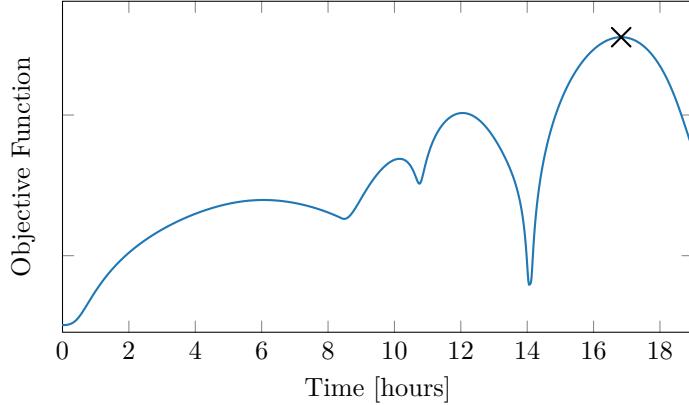


Figure 1: The forecasted KL divergence objective when scheduling a single measurement for a ground-based radar, with the optimal measurement time denoted by \times .

Furthermore, as the number of sensors, targets, and allowable measurements increases, the tasking problem quickly becomes computationally intractable. For instance, evaluating the forecasted objective requires batch matrix multiplications and block inversions whose complexity grows rapidly with the number of candidate measurements. The myopic objective is similarly affected, highlighting the need for efficient, structure-aware optimization strategies. To address this, the following sections introduce a convex relaxation of the forecasted objective based on submodular function analysis. Analogous formulations and efficiency gains apply to the myopic case as well.²

Submodular Function Analysis

Submodularity provides a principled way to handle non-convex optimization problems by leveraging the property of diminishing returns. As shown in numerous applications, such as sensor placement and experimental design,¹⁴ submodular objectives admit efficient greedy algorithms with provable approximation guarantees. Establishing that the forecasted objective is submodular opens the door to such scalable and theoretically grounded optimization methods.

To establish the submodularity of the forecasted objective function in Eq. (13), consider a set function $f : 2^{\mathbb{Z}} \rightarrow \mathbb{R}$ defined over subsets $\zeta \subseteq \mathbb{Z}$. A set function is said to be *submodular*

if it satisfies the diminishing returns property,¹⁴

$$f(\zeta_1 \cup \{\tau\}) - f(\zeta_1) \geq f(\zeta_2 \cup \{\tau\}) - f(\zeta_2) \quad \forall \zeta_1 \subseteq \zeta_2 \subseteq \mathbb{Z}, \tau \in \mathbb{Z} \setminus \zeta_2. \quad (16)$$

Let $\zeta \subseteq \mathbb{Z}$ denote a set of selected measurement indices and τ be a candidate measurement index not yet in ζ . Define a fixed reference time t_{ref} that occurs after all possible measurements, such that the information from each observation is mapped in a consistent direction to this common time. Utilizing the block-diagonal property of $\bar{\mathbf{P}}_{vv}$ in Eq. (10), the forecasted information gain from the set ζ is represented by the matrix

$$\bar{\mathbf{F}}(\zeta) = \sum_{\delta \in \zeta} \mathbf{F}_\delta,$$

where $\mathbf{F}_\delta \triangleq \mathbf{F}(t_\delta, t_{\text{ref}}) = \tilde{\mathbf{H}}_\delta^T \mathbf{P}_{vv, \delta}^{-1} \tilde{\mathbf{H}}_\delta$ is the mapped Fisher information contribution from measurement δ . Each matrix \mathbf{F}_δ is symmetric positive semi-definite, as it is constructed by a similarity transform of positive definite noise covariance $\mathbf{P}_{vv, \delta}$. As a result, the sum $\bar{\mathbf{F}}$ is also symmetric positive semi-definite. Using this notation and adding a new index τ , Eq. (13) becomes

$$\bar{J}(\zeta \cup \{\tau\}) = \frac{1}{2} \log \left| \mathbf{I}_{n_x} + \mathbf{P}_{xx, \text{ref}}^- (\bar{\mathbf{F}}(\zeta) + \mathbf{F}_\tau) \right|. \quad (17)$$

Subtracting $\bar{J}(\zeta)$ from Eq. (17) to get into the form of Eq. (16), the marginal gain is found as

$$\bar{J}(\zeta \cup \{\tau\}) - \bar{J}(\zeta) = \frac{1}{2} \log \left| \mathbf{I}_{n_x} + \left(\mathbf{I}_{n_x} + \mathbf{P}_{xx, \text{ref}}^- \bar{\mathbf{F}}(\zeta) \right)^{-1} \mathbf{P}_{xx, \text{ref}}^- \mathbf{F}_\tau \right|. \quad (18)$$

To prove submodularity, define two sets $\zeta_1 \subseteq \zeta_2$ and consider the marginal gain of adding a new index τ to each. Since ζ_2 contains all elements of ζ_1 and possibly more, the corresponding information matrices satisfy $\bar{\mathbf{F}}(\zeta_1) \preceq \bar{\mathbf{F}}(\zeta_2)$, meaning that their determinants obey $|\bar{\mathbf{F}}(\zeta_1)| \leq |\bar{\mathbf{F}}(\zeta_2)|$.¹³ Using the same “information never hurts” argument,⁴ it follows that

$$\left| \left(\mathbf{I}_{n_x} + \mathbf{P}_{xx, \text{ref}}^- \bar{\mathbf{F}}(\zeta_1) \right)^{-1} \mathbf{P}_{xx, \text{ref}}^- \mathbf{F}_\tau \right| \geq \left| \left(\mathbf{I}_{n_x} + \mathbf{P}_{xx, \text{ref}}^- \bar{\mathbf{F}}(\zeta_2) \right)^{-1} \mathbf{P}_{xx, \text{ref}}^- \mathbf{F}_\tau \right|. \quad (19)$$

It remains to show that adding the identity matrix inside the determinant in Eq. (18) preserves the inequality. Although the matrices in Eq. (19) are not necessarily symmetric, they are similar to symmetric positive semi-definite matrices. Specifically, under the assumption that $\mathbf{P}_{xx, \text{ref}}^- \succ 0$ and both $\bar{\mathbf{F}}(\zeta)$ and \mathbf{F}_τ are symmetric positive semi-definite, the matrices on the left and right sides of Eq. (19) are similar to

$$\mathbf{P}^{-1/2} \left(\mathbf{I} + \mathbf{P}^{1/2} \bar{\mathbf{F}}(\zeta) \mathbf{P}^{1/2} \right)^{-1} \left(\mathbf{P}^{1/2} \mathbf{F}_\tau \mathbf{P}^{1/2} \right) \mathbf{P}^{1/2},$$

where $\mathbf{P} \triangleq \mathbf{P}_{xx, \text{ref}}^-$. This matrix is symmetric and positive semi-definite, so its eigenvalues are real and non-negative. Therefore, the original matrix has real, non-negative eigenvalues. Let $\lambda_i^{(1)}$ and $\lambda_i^{(2)}$, for $i = 1, \dots, n_x$, denote the eigenvalues of the left and right sides

of Eq. (19), respectively. Since $\lambda_i^{(1)} \geq \lambda_i^{(2)}$, adding one to each eigenvalue preserves the ordering,

$$\prod_{i=1}^{n_x} (1 + \lambda_i^{(1)}) \geq \prod_{i=1}^{n_x} (1 + \lambda_i^{(2)}).$$

This confirms that the forecasted objective in Eq. (13), along with its multitarget extension in Eq. (14), satisfies the diminishing returns property in Eq. (16), and therefore constitutes a submodular set function.

Convex Closure Relaxation

Recall that to formulate the forecasted objective in Eq. (13) as a submodular set function, the optimization variables are a set of discrete measurement indices ζ , and $\mathbf{t}(\zeta)$ is used to select the times from these indices. While each precise value in \mathbf{t} is technically continuous, the scheduling problem naturally admits discretization as a finite set of candidate measurement times.

To enable convex optimization, the discrete selection process is reformulated using binary indicator variables. For each sensor $s \in \{1, \dots, S\}$ and target $i \in \{1, \dots, N\}$, let $\mathbf{t}_s^{(i)} = \{t_1, \dots, t_K\}$ denote the candidate measurement times. Each candidate time t_k is assigned an indicator variable

$$\gamma_{s,k}^{(i)} = \begin{cases} 1 & \text{if sensor } s \text{ observes target } i \text{ at time } t_k, \\ 0 & \text{otherwise.} \end{cases}$$

Let $\boldsymbol{\gamma}_s^{(i)} \in \{0, 1\}^K$ denote the full indicator vector for the s - i sensor-target pair, and define the full scheduling vector as $\boldsymbol{\gamma} \in \{0, 1\}^d$, where $d = S \cdot N \cdot K$ is the total number of scheduling opportunities.

The forecasted information objective $\bar{J}(\zeta)$ from Eq. (13) can now be interpreted as a set function $\bar{J}(\boldsymbol{\gamma})$ defined over $\{0, 1\}^d$. To obtain a tractable relaxation of the combinatorial problem, the domain is extended to $\boldsymbol{\gamma} \in [0, 1]^d$, and the convex closure of $\bar{J}(\boldsymbol{\gamma})$ is considered. The convex closure, or *Lovász extension*, is a continuous extension of a submodular function from $\{0, 1\}^d$ to $[0, 1]^d$, and it preserves convexity if and only if the original set function is submodular.¹⁴

Formally, the Lovász extension $f : [0, 1]^d \rightarrow \mathbb{R}$ of a submodular function $F : \{0, 1\}^d \rightarrow \mathbb{R}$ is defined as¹⁴

$$f(\boldsymbol{\gamma}) = \int_0^1 F(\{k \in \{1, \dots, d\} : \gamma_k \geq \tau\}) \, d\tau,$$

where $\boldsymbol{\gamma} \in [0, 1]^d$ is a vector of relaxed indicator variables, and $\tau \in [0, 1]$ is a threshold parameter. For each τ , the set $\{k : \gamma_k \geq \tau\}$ selects a subset of elements for which the relaxed weights exceed the threshold. This subset is evaluated using the original set function F , and the result is averaged over all τ from 0 to 1. Intuitively, this process interpolates the values of F across the continuous hypercube, yielding the tightest convex function that matches F

at all binary vectors. This extension allows the reinterpretation of the optimization problem as a convex program

$$\boldsymbol{\gamma}^* = \arg \max_{\boldsymbol{\gamma} \in [0,1]^d} \bar{J}(\boldsymbol{\gamma}),$$

where $\bar{J}(\boldsymbol{\gamma})$ is now the Lovász extension of the original submodular objective.

This reparameterization yields a convex objective over the continuous hypercube, enabling efficient optimization while preserving the structure of the original submodular problem. By relaxing the binary scheduling variables to lie in $[0, 1]$, the resulting formulation supports global optimization and admits meaningful linear constraints:¹⁵

$$\begin{aligned} \gamma_{s,k}^{(i)} &\in [0, 1] & \forall s \in \{1, \dots, S\}, i \in \{1, \dots, N\}, k \in \{1, \dots, K\}, \\ \sum_{k=1}^K \gamma_{s,k}^{(i)} &= M_s^{(i)} & \forall s \in \{1, \dots, S\}, i \in \{1, \dots, N\}, \\ \sum_{i=1}^N \gamma_{s,k}^{(i)} &\leq 1 & \forall s \in \{1, \dots, S\}, k \in \{1, \dots, K\}. \end{aligned}$$

The first constraint defines the feasible domain of the relaxed problem as the continuous hypercube. The second constraint enforces that each sensor-target pair (s, i) is assigned exactly $M_s^{(i)}$ measurements across the scheduling horizon. The final constraint ensures that, at any given candidate time t_k , each sensor s is assigned to at most one target, preventing simultaneous tasking of multiple targets by the same sensor. These conditions define a feasible region over which the relaxed problem remains convex.

In this formulation, the measurement model in Eq. (4) used during optimization is augmented with the indicator variable,

$$z'_s(\mathbf{x}_k^{(i)}) = \gamma_{s,k}^{(i)} \cdot \mathbf{h}(\mathbf{x}_k^{(i)}) + \mathbf{v}_k.$$

The indicator applies directly to the batch update in Eq. (11) as

$$(\mathbf{P}_{xx,\text{ref}}^+)' = \left((\mathbf{P}_{xx,\text{ref}}^-)^{-1} + \boldsymbol{\Gamma} \bar{\mathbf{H}}^T \bar{\mathbf{P}}_{vv}^{-1} \bar{\mathbf{H}} \boldsymbol{\Gamma} \right)^{-1},$$

where the diagonal matrix of indicator variables is given by

$$\boldsymbol{\Gamma} = \text{diag}(\gamma_{s,1}^{(i)}, \dots, \gamma_{s,K}^{(i)}).$$

The single-target forecasted convex optimization problem can then be written as

$$\begin{aligned} \arg \max_{\boldsymbol{\gamma}} \bar{J}(\boldsymbol{\gamma}) &\triangleq \bar{M}_{\text{KL}}(\mathbf{t}_{\boldsymbol{\gamma}}) \\ &= \frac{1}{2} \log \left| \mathbf{I}_{n_x} + \mathbf{P}_{xx,\text{ref}}^- \boldsymbol{\Gamma} \bar{\mathbf{H}}^T \bar{\mathbf{P}}_{vv}^{-1} \bar{\mathbf{H}} \boldsymbol{\Gamma} \right|, \end{aligned} \quad (21)$$

where $\mathbf{t}_{\boldsymbol{\gamma}} \triangleq \mathbf{t}(\boldsymbol{\gamma})$ denotes the set of measurement times corresponding to nonzero entries in the relaxed indicator vector. When $\gamma_k = 0$, the k -th row of $\bar{\mathbf{H}}$ is effectively ignored, ensuring that unselected measurements do not influence the update. The multitarget extension

follows the same structure as Eq. (14), summing contributions across all targets to yield a globally consistent scheduling objective.

After solving the convex relaxation, the optimal schedule γ^* typically contains fractional entries. A final rounding procedure is used to extract a feasible binary schedule. Following Reference 15, a simple and effective heuristic is to select the $M_s^{(i)}$ entries in $\gamma_s^{(i)}$ with the largest weights for each sensor-target pair. This ensures feasibility while preserving the structure of the relaxed solution. While this convex closure formulation requires a higher-dimensional decision space, it enables globally consistent optimization with guarantees inherited from submodularity and convex analysis.¹² Additionally, since each $\mathbf{t}_s^{(i)}$ is static, expensive operations such as computing the state transition matrix $\Phi(t_k, t_{\text{ref}})$ can be precomputed and cached at each candidate time t_k , allowing the optimization to proceed efficiently without redundant evaluations.

CISLUNAR SIMULATION

Scenario Description

To evaluate the effectiveness of the forecasted tasking approach, a representative single-target, single-observer scenario is considered in the cislunar environment. The target is placed in a near-rectilinear halo orbit (NRHO) about the Earth–Moon L_2 point, while the observer is assumed to reside in a distant retrograde orbit (DRO).¹⁶

The target and observer are both propagated according to the dynamics of the circular restricted three-body problem (CR3BP), with the governing equations of motion given as in Reference 6. The initial uncertainty in the target’s state is represented by a diagonal covariance matrix,

$$\mathbf{P}_{xx,0} = \begin{bmatrix} \sigma_r^2 \mathbf{I}_3 & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \sigma_v^2 \mathbf{I}_3 \end{bmatrix}, \quad (22)$$

with position and velocity uncertainties given by $\sigma_r = 10$ km and $\sigma_v = 0.1$ m/s, respectively. Figure 2 illustrates the orbital geometries of the target NRHO and observer DRO.

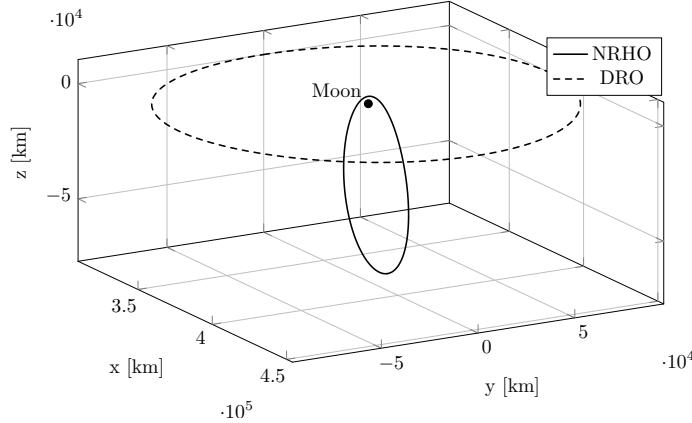


Figure 2: Geometry of the target NRHO and observer DRO orbits in the CR3BP synodic frame.

The observer is capable of obtaining angles-only measurements of the target's position, with azimuth α and elevation ϵ defined as

$$\alpha = \arccos \left(\frac{\gamma_x \rho_x + \gamma_y \rho_y}{\sqrt{\gamma_x^2 + \gamma_y^2} \sqrt{\rho_x^2 + \rho_y^2}} \right)$$

$$\epsilon = \arccos \left(\frac{\gamma_x \rho_x + \gamma_z \rho_z}{\sqrt{\gamma_x^2 + \gamma_z^2} \sqrt{\rho_x^2 + \rho_z^2}} \right),$$

where $\boldsymbol{\gamma}$ is the vector from the observer to a reference body (e.g., the Moon), and $\boldsymbol{\rho}$ is the vector from the observer to the target:

$$\boldsymbol{\gamma} = [\gamma_x \ \gamma_y \ \gamma_z]^T = \mathbf{x}_{\text{ref}} - \mathbf{x}_{\text{obs}}$$

$$\boldsymbol{\rho} = [\rho_x \ \rho_y \ \rho_z]^T = \mathbf{x} - \mathbf{x}_{\text{obs}}.$$

The output of the measurement function is given as $\mathbf{h} = [\alpha \ \epsilon]^T$. These angles are illustrated in Figure 3. The measurement noise is considered to be constant and is represented by the covariance matrix

$$\mathbf{P}_{vv} = \begin{bmatrix} \sigma_\alpha^2 & 0 \\ 0 & \sigma_\epsilon^2 \end{bmatrix}, \quad (23)$$

with $\sigma_\alpha = \sigma_\epsilon = 3''$.

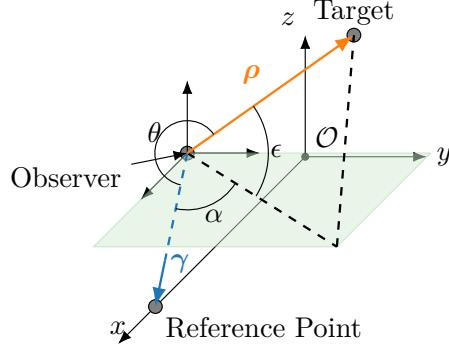


Figure 3: Schematic of the angles and vectors in the measurement model for the cislunar scenario.

To ensure measurement validity, a lighting exclusion constraint is enforced. Specifically, the observer refrains from taking measurements when the target appears too close to either the Earth or the Sun, due to their brightness. This is handled by requiring that the angle θ , depicted in Figure 3, between the vectors $\boldsymbol{\gamma}$ and $\boldsymbol{\rho}$ satisfies

$$\theta = \arccos \left(\frac{\boldsymbol{\gamma} \cdot \boldsymbol{\rho}}{\|\boldsymbol{\gamma}\| \|\boldsymbol{\rho}\|} \right) > \phi,$$

where the threshold angle is defined as $\phi_\odot = 20^\circ$ for the Sun and $\phi_\oplus = 10^\circ$ for the Earth.¹⁷ The target is propagated for one orbital period, approximately 6.5 days in duration. Within the visibility region, the observer is heuristically assigned 50 measurement opportunities.

To initialize the optimizer, a representative candidate set of measurement times is selected using a Sundman-like transformation that reparameterizes time to promote uniform spatial sampling along the trajectory.¹⁸ Rather than sampling uniformly in physical time, which can lead to measurement clustering near apoapsis and sparsity near periapsis,¹⁹ the trajectory is reparametrized according to arc length. This transformation is given by

$$s(t) = \int_{t_0}^t \|\dot{\mathbf{r}}(\tau)\| \, d\tau,$$

where s denotes the cumulative arc length along the trajectory and $\dot{\mathbf{r}} = \mathbf{v}$ is the velocity vector in the CR3BP rotating frame.

The optimization proceeds by iteratively adjusting the measurement times to maximize the expected information gain under either the myopic or forecasted tasking procedure, while enforcing the aforementioned visibility constraints. For the forecasted objective, a reference time of interest is set to the point of closest lunar approach, approximately 3.25 days into the simulation.

Results

The optimized measurement schedules for both the myopic and forecasted approaches are projected onto the true target trajectory in Figure 4. A notable feature of this comparison is the wait-and-see behavior exhibited by the myopic strategy. Although visibility is achieved earlier, the myopic schedule defers its first measurement by approximately 7 hours, in contrast to the forecasted schedule, which initiates measurements as soon as visibility begins. Interestingly, both schedules avoid taking measurements within roughly 25 minutes of the reference time t_{ref} , which corresponds to the point of closest lunar approach.

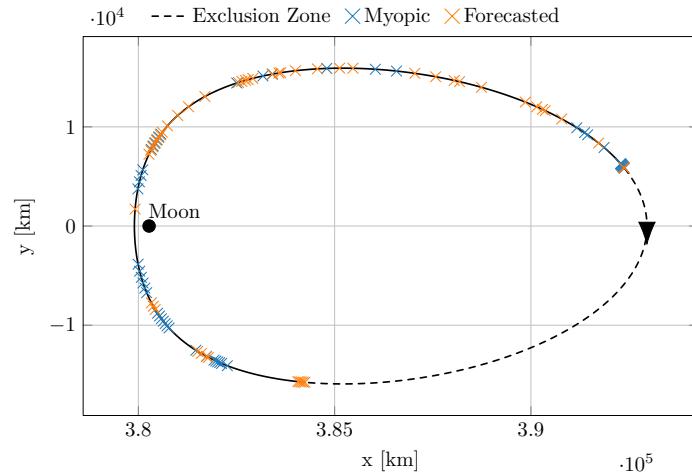


Figure 4: Optimized measurement schedules projected onto the target trajectory in the CR3BP synodic frame, with the direction and beginning of the propagation window denoted by ►.

To evaluate the resulting schedules, a set of 1,000 Monte Carlo trials is performed, with the position and velocity errors plotted in Figure 5. In each trial, the initial state is sampled using the covariance specified in Eq. (22), and independent measurement noise realizations

are applied at the fixed schedule times. This allows the statistical consistency and uncertainty reduction performance of each tasking strategy to be assessed. These results show that the average filter uncertainty $\pm 3\sigma$ for the forecasted schedule consistently lies within that of the myopic case. This is particularly pronounced in the velocity channel, where the forecasted schedule anticipates the sharp nonlinear dynamics around perilune and dramatically reduces uncertainty at the time of this event. In contrast, the myopic schedule – unable to foresee the impending change – suffers a noticeable growth in uncertainty thereafter.

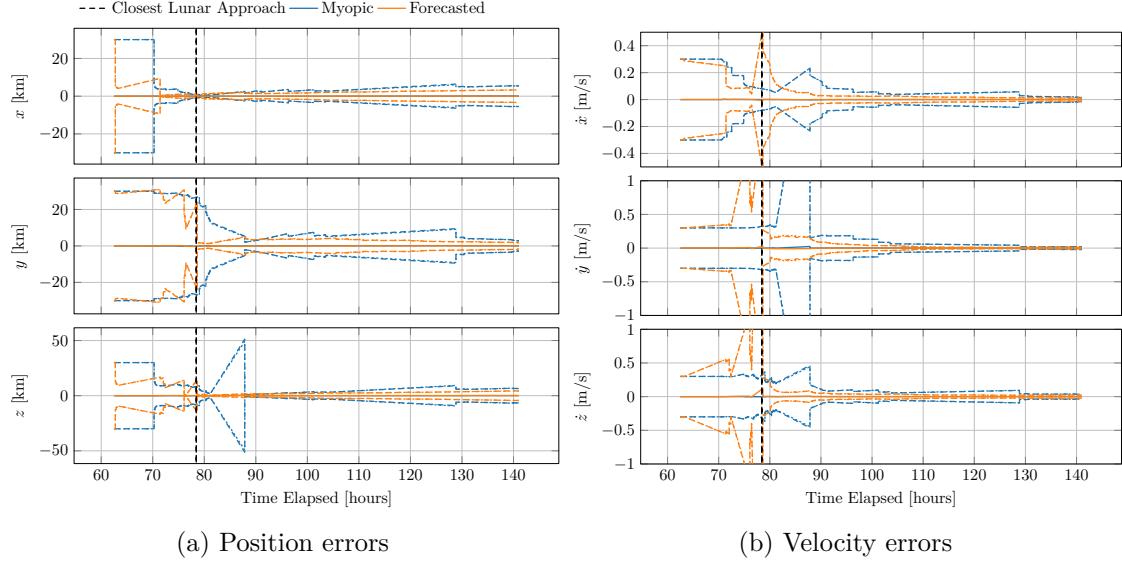


Figure 5: Monte Carlo position and velocity error statistics for the resulting schedules, including average filter error (—), average filter $\pm 3\sigma$ (---), and Monte Carlo $\pm 3\sigma$ (.....).

To assess the informativeness of each schedule, the Shannon entropy over time is plotted in Figure 6a, where the entropy is defined as in Eq. (1). Additionally, the time-averaged Shannon entropy is computed as

$$\tilde{H} = \frac{1}{t_f - t_0} \sum_{\delta=1}^{M-1} H_\delta(t_{\delta+1} - t_\delta), \quad (24)$$

where M is the number of measurements, and t_0 and t_f denote the initial and final times of observability, respectively. This quantity, shown in Figure 6b, provides a scalar measure of overall uncertainty reduction across the observation window. From these plots, it is evident that the forecasted schedule achieves a more consistent and timely reduction in uncertainty, reflected in lower entropy values across the observation window.

To further compare performance, the position and velocity root-sum-square (RSS) errors at both the reference time and the final time of observability are reported in Table 1. Interestingly, while the myopic schedule slightly outperforms the forecasted schedule in velocity at the reference time, the forecasted approach yields significantly lower position error at that time and dominates in both position and velocity at the final time.

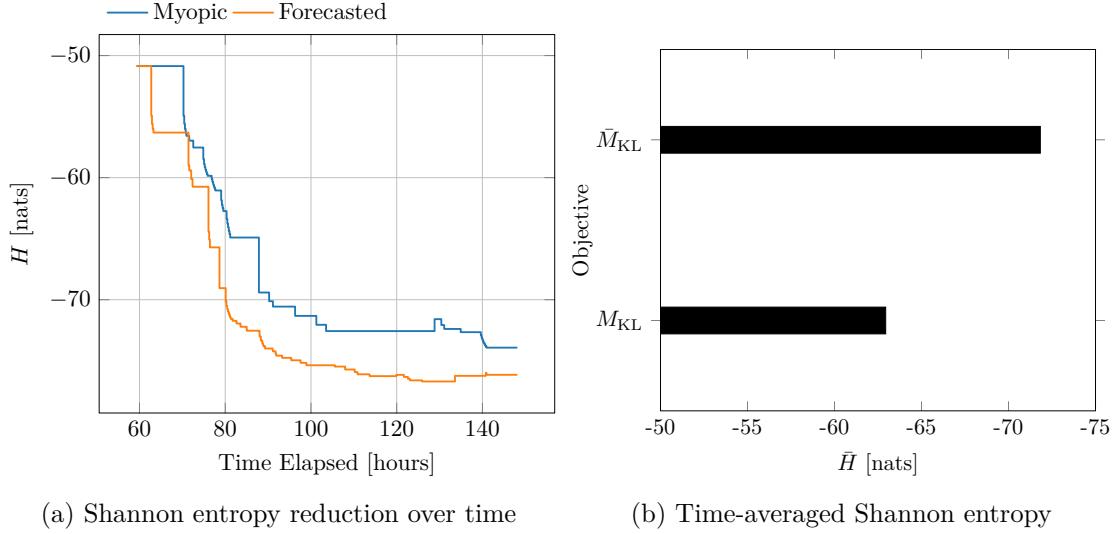


Figure 6: Shannon entropy reduction and average Shannon entropy for the optimized myopic (M_{KL}) and forecasted (\bar{M}_{KL}) schedules.

Table 1: The position and velocity RSS values at the reference time (closest lunar approach) and the final time (end of the observation window) for the myopic (M_{KL}) and forecasted (\bar{M}_{KL}) schedules.

	\mathbf{r} [km]		\mathbf{v} [m/s]	
	M_{KL}	\bar{M}_{KL}	M_{KL}	\bar{M}_{KL}
t_{ref}	9.080	3.054	1.376×10^{-1}	2.230×10^{-1}
t_f	2.981	1.818	1.423×10^{-2}	8.831×10^{-3}

LOW EARTH ORBIT SIMULATION

Scenario Description

Having demonstrated the effectiveness of the forecasted tasking framework relative to the myopic approach, a new simulation scenario is introduced to assess the forecasted approach’s performance in a multitarget, multi-observer context. This scenario highlights two key advantages of the forecasted methodology: its natural extension to heterogeneous sensor networks – where the myopic objective lacks a direct analog – and its compatibility with convex relaxation techniques for scalable optimization. By exploring this more intricate environment, the utility and flexibility of the forecasted objective are further emphasized.

Specifically, 12 space objects are initialized in low Earth orbit (LEO) and propagated using two-body dynamics over a one-hour analysis window, with the governing equations of motion provided in Reference 6. The initial uncertainty in the state of each target is assumed to be identical and is specified by Eq. (22), with position and velocity standard deviations of $\sigma_r = 16.67$ m and $\sigma_v = 6.67$ cm/s, respectively.

Observations are collected from two platforms: a space-based observer in a 400 km circular orbit with a 75° inclination, and a ground-based observer located at the Goldstone

Deep Space Communications Complex (GDSCC) in California. The space-based platform provides full observability of all targets via relative right ascension (α) and declination (δ) measurements, while the ground-based observer collects range (ρ) and range-rate ($\dot{\rho}$) measurements. The measurement models for each modality are implemented following the formulations in Reference 2. Measurement noise is assumed to be constant for each observer and is defined by Eq. (23), with $\sigma_\alpha = \sigma_\delta = 3''$ for the space-based observer and $\sigma_\rho = 15$ m, $\sigma_{\dot{\rho}} = 4$ mm/s for the ground-based observer. Each observer-target pair is heuristically assigned 5 measurement opportunities, yielding a total of 120 scheduled measurements over the analysis window.

The resulting measurement geometry from the space-based observer is illustrated in Figure 7a. Nine of the targets remain within the orbital plane of this observer, while the remaining three exhibit out-of-plane motion. Due to the close proximity of the target objects to the space observer, the visibility window for the ground-based observer is driven by its geometric relationship with the space platform. Specifically, the elevation angle θ between the ground and space observers is computed as

$$\theta = \frac{\pi}{2} - \arccos(\hat{\rho} \cdot \hat{k}),$$

where $\hat{\rho}$ is the unit vector from the ground observer to the space observer, and \hat{k} is the unit vector normal to the local horizon at the ground site. Applying a 10° elevation mask constraint yields approximately 5 minutes of ground-based visibility during the analysis window. The resulting elevation angle profile is shown in Figure 7b.

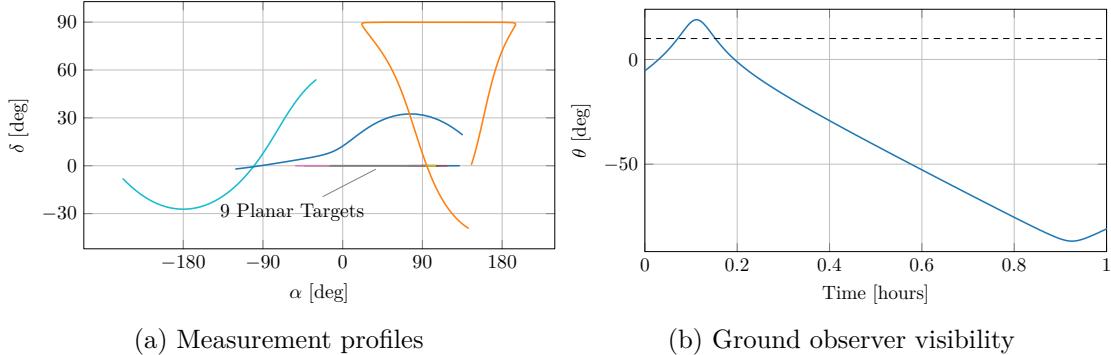


Figure 7: Measurement geometry for the LEO scenario. (a) Angular measurement profiles from the space-based observer. (b) Elevation angle between the ground observer and the space-based observer, indicating visibility windows.

Unlike the cislunar scenario, where a natural reference time such as perilune emerges from the dynamics, this LEO-based simulation lacks a dominant orbital feature. As such, the reference epoch for forecasted tasking is selected to coincide with the end of the observation window. The optimization proceeds by applying both the standard (Eq. (13)) and convexified (Eq. (21)) forms of the forecasted multitarget objective to generate measurement schedules, while enforcing visibility constraints for the ground-based observer. Due to the high dimensionality of the scheduling problem, a greedy-in-target (GIT) and greedy-in-observer (GIO) algorithmic structure is adopted for both formulations. Justification for the use of these greedy strategies is provided in Reference 2.

Results

The optimized measurement schedules produced by both forecasted objectives are shown in Figure 8, with target color associations matching those in Figure 7a. For the convexified case, the optimization is performed over a discrete candidate time grid with 20-second spacing. This yields a total of 15 candidate times for the ground-based observer and 180 candidate times for the space-based observer.

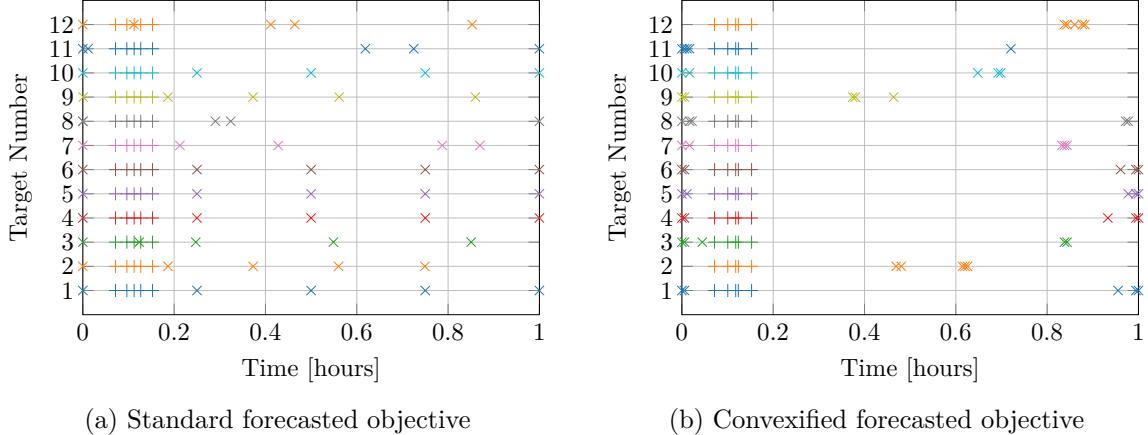


Figure 8: Optimized schedules from both the standard and convexified forms of the forecasted objective, with colors matching Figure 7a. Ground observer measurements are denoted with +, and space observer measurements are denoted with \times .

A qualitative comparison between the two schedules reveals a degree of corroboration in the tasking behavior, particularly for Targets 2 and 9, where both methods prioritize similar measurement epochs. Notably, in the convexified schedule, the space-based observer consistently avoids taking measurements during the ground-based visibility window. Similarly, the standard forecasted objective includes only two space-based measurements that overlap with the ground observer's visibility – those for Targets 3 and 12. This scheduling behavior suggests that the forecasted framework, in both forms, assimilates the relative value of information across sensing platforms. By minimizing redundant observations and spreading measurement effort across complementary epochs, the optimization exploits the strengths of each sensor and leverages their collective contribution to information gain.

As in the cislunar scenario, the reduction in Shannon entropy over time is plotted for both forecasted objectives in Figure 9. Entropy is accumulated collectively across all targets, treating the system as a whole, which is a valid assumption under the condition that there is no ambiguity in target identity.^{2,4} The resulting curves show an impressive degree of agreement, both in the overall magnitude of entropy reduction and in the timing and trajectory of that reduction. Despite originating from fundamentally different optimization processes, the similarity in performance underscores the effectiveness of the convex relaxation in replicating the essential behavior of the full forecasted formulation.

This alignment is particularly striking in light of the computational disparity between the two approaches. The optimization is performed under identical operating conditions: a quad-threaded GIT and GIO scheduling framework, solved using a gradient-based sequential

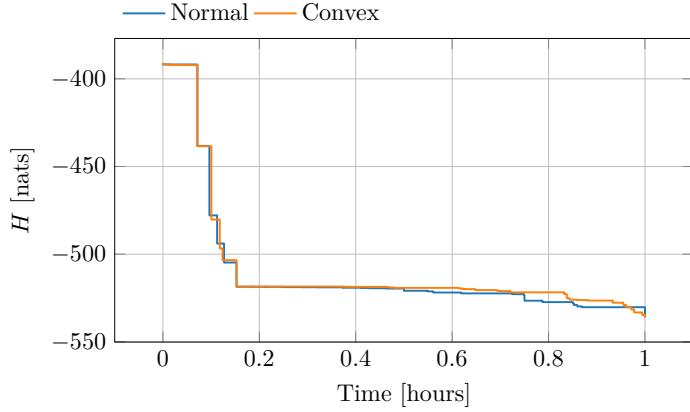


Figure 9: Shannon entropy reduction accumulated across all targets for the optimized schedules obtained using the standard and convexified forms of the forecasted objective.

quadratic programming (SQP) optimizer with strict function and parameter tolerances ($f_{\text{tol}} = 10^{-12}$, $x_{\text{tol}} = 10^{-8}$). Under these conditions, the convexified forecasted objective converged approximately $350\times$ faster than the standard formulation.

A natural question arising from the results presented thus far concerns the scalability of the convexified forecasted objective relative to the standard formulation – both in terms of computational runtime and the potential for performance degradation as the optimization space is refined. To examine this, an experiment is conducted in which the separation time between candidate measurement epochs is varied from 40 seconds down to 5 seconds. For each case, the optimization runtime is normalized against that of the standard forecasted objective, and the average Shannon entropy – defined in Eq. (24) – is computed and normalized by the value obtained using the standard formulation. The resulting trends are shown in Figures 10.

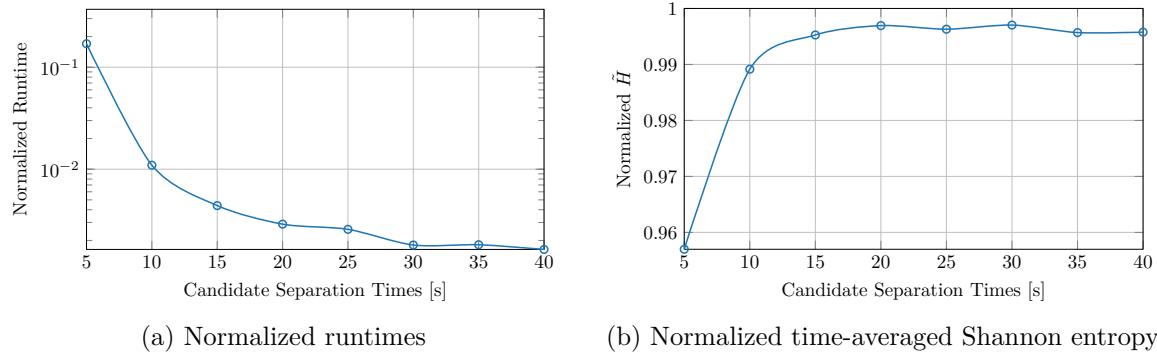


Figure 10: Scalability analysis of the convexified forecasted objective. (a) Optimization runtimes for convexified schedules, normalized by the corresponding runtime from the standard objective. (b) Time-averaged Shannon entropies from convexified schedules, normalized by their standard counterpart. Both are plotted as functions of candidate separation time.

Several important behaviors emerge. First, Figure 10a reveals that coarsening the candidate time grid leads to a dramatic decrease in runtime. This behavior can be attributed to two factors: (1) many of the computationally intensive evaluations involved in Eq. (13) are precomputed, eliminating redundant operations during optimization, and (2) the gradient-based optimizer efficiently navigates the convex landscape, often converging within 15–20 iterations due to rapid identification of the steepest descent direction.

Second, Figure 10b shows that finer candidate spacing does not necessarily yield better performance. In fact, the schedules produced using 5- and 10-second separations exhibit the worst performance in terms of average Shannon entropy, despite requiring significantly longer runtimes. This counterintuitive result arises because, as the optimization space becomes increasingly fine-grained, the marginal utility of additional measurement times diminishes. Many candidate times are nearly indistinguishable in their expected information gain, diluting the effectiveness of greedy selection. An emergent artifact of this over-refinement is that selected measurement times begin to cluster very closely, a phenomenon observable even in Figure 8b. This behavior underscores the need to introduce additional proactive constraints – such as limiting the number of measurements allowed within a given time span – to maintain meaningful diversity in measurement timing and avoid information redundancy.

CONCLUSIONS AND FUTURE WORK

The utility of the proposed tasking framework lies in its flexibility and broad applicability to any operational domain where space domain awareness (SDA) is required. Unlike approaches tied to specific orbital environments or sensor types, the method presented and explored in this work generalizes across regimes and adapts to any sensor network. Its reliance on expected information gain enables principled tasking without requiring fully specified downstream processing models, and its low computational overhead makes it suitable for real-time scheduling and catalog maintenance at scale.

The cislunar scenario highlights how optimizing for future information can yield unintuitive but strategically effective measurement schedules, emphasizing the value of early, informative sensing. Additionally, the low Earth orbit (LEO) scenario demonstrates the framework’s applicability to heterogeneous observer networks and validates the scalability of its convexified formulation – even under high-dimensional scheduling problems – without significant performance loss.

Together, these case studies show that the forecasted approach not only produces robust and efficient schedules, but also facilitates intelligent coordination across sensing assets. This enables a shift toward more agile, globally optimized surveillance strategies – an essential capability in increasingly dynamic space environments. Future work will focus on more rigorous enforcement of practical constraints, such as sensor-specific limitations on measurement timing, as well as exploring hybrid optimization pipelines that combine the standard and convexified formulations to balance optimality with computational efficiency.

REFERENCES

- [1] A. O. Hero, D. A. Casta  n  n, D. Cochran, and K. Kastella, eds., *Foundations and Applications of Sensor Management*. New York: Springer, 2008.

- [2] M. J. Gualdoni, *Applications of Information Theory in Filtering and Sensor Management*. Doctoral dissertation, Missouri University of Science and Technology, 2020.
- [3] C. M. Kreucher, A. O. Hero, K. D. Kastella, and M. R. Morelande, “An Information-Based Approach to Sensor Management in Large Dynamic Networks,” *Proceedings of the IEEE*, Vol. 95, No. 5, 2007, pp. 978–999.
- [4] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, Inc., 2nd ed., 2006.
- [5] A. Rényi, *Foundations of Probability*. Courier Corporation, 2007.
- [6] D. A. Vallado and W. D. McClain, *Fundamentals of Astrodynamics and Applications*. Hawthorne, CA: Microcosm Press, fourth ed., 2013. Space Technology Library, Analytical Graphics Inc.
- [7] A. Gelb, ed., *Applied Optimal Estimation*. Cambridge, MA: MIT Press, 1974.
- [8] D. Simon, *Optimal State Estimation: Kalman, H ∞ , and Nonlinear Approaches*. Hoboken, NJ: John Wiley & Sons, 2006.
- [9] C. M. Kreucher, *An Information-Based Approach to Sensor Resource Allocation*. Doctoral dissertation, The University of Michigan, 2005.
- [10] M. J. Gualdoni and K. J. DeMars, “An Improved Representation of Measurement Information Content Via the Distribution of the Kullback-Leibler Divergence,” *Advances in the Astronautical Sciences*, Vol. 162, Aug 2018, pp. 1889–1908.
- [11] B. D. Tapley, B. E. Schutz, and G. H. Born, *Statistical Orbit Determination*. Amsterdam, Netherlands: Elsevier, 2004.
- [12] J. Dattorro, *Convex Optimization & Euclidean Distance Geometry*. USA: Meboo Publishing, updated ed., 2008. Revised & Enlarged International Paperback Edition III.
- [13] R. A. Horn and C. R. Johnson, *Matrix Analysis*. New York, NY, USA: Cambridge University Press, 2nd ed., 2012.
- [14] F. Bach, “Learning with Submodular Functions: A Convex Optimization Perspective,” *Foundations and Trends in Machine Learning*, Vol. 6, No. 2-3, 2013, pp. 145–373, 10.1561/2200000039.
- [15] S. Joshi and S. Boyd, “Sensor Selection via Convex Optimization,” *IEEE Transactions on Signal Processing*, Vol. 57, No. 2, 2009, pp. 451–462, 10.1109/TSP.2008.2007095.
- [16] Y. Wang, R. Zhang, C. Zhang, and H. Zhang, “Transfers between NRHOs and DROs in the Earth-Moon system,” *Acta Astronautica*, Vol. 186, 2021, pp. 60–73, <https://doi.org/10.1016/j.actaastro.2021.05.019>.
- [17] T. H. Clareson, M. C. Fox, D. K. Amato, and H. W. Lee, “Embedded State Estimation for Optimization of Cislunar Space Domain Awareness Constellation Design,” *AAS/AIAA Astrodynamics Specialist Conference*, Big Sky, MT, American Astronautical Society (AAS), 2023. Substantially revised version of Paper AAS 23-189.
- [18] K. F. Sundman, “Mémoire sur le problème des trois corps,” *Acta Mathematica*, Vol. 36, 1913, pp. 105–179.
- [19] J. Roa, *Regularization in Orbital Mechanics*. Berlin, Boston: De Gruyter, 2017.