

FAST AUTONOMOUS LOST-IN-SPACE CATALOG-BASED OPTICAL NAVIGATION (FALCON)

Antonio Rizza*, Justin Kruger†, Mike Timmerman‡, and Simone D’Amico§

Ensuring the safety and sustainability of the space environment requires comprehensive tracking of the rapidly growing number of Resident Space Objects (RSO) in orbit, fused with robust autonomous navigation capabilities for active satellites. This paper proposes an alternative space-based positioning, navigation, and monitoring system called Fast Autonomous Lost-in-space Catalog-based Optical Navigation (FALCON). Observer satellites use onboard cameras to obtain bearing angle measurements to visible targets and match measurements to existing identities in an onboard RSO catalog. Measurement batches may be used to triangulate the observer’s position, or to refine the observer’s orbit estimate using sequential estimation strategies. This methodology enables lost-in-space optical orbit determination for the observer entirely independent of Global Navigation Satellite Systems (GNSS) or tracking of the observer from the ground. The new methodology is validated using a highly realistic simulation scenario based on the NASA Starling mission. Flight data is combined with the Space-Track catalog to generate realistic measurement inputs. Simulation results display precise matching of measurements to catalog objects in the presence of catalog orbit uncertainties, without observer orbit knowledge, while remaining within feasible computational bounds. Batch and sequential orbit determination demonstrates accurate refinement of the observer’s orbit from a lost-in-space scenario in less than one orbit with sub-kilometer accuracy when the proper set of beacons is selected. The resulting implementation of the FALCON architecture demonstrates strong new strategic capabilities for resilient optical navigation and RSO tracking.

INTRODUCTION

To facilitate the rapid growth of space operations, there is an urgent need to provide resilient Positioning, Navigation, and Timing (PNT) and responsive Space Situational Awareness (SSA) for on-orbit assets. Advances in miniaturized and commercial off-the-shelf satellite hardware, together with reduced costs for satellite launches, have led to a three-fold increase in the number of tracked Resident Space Objects (RSO) since 2010.¹ It is also predicted that up to 3,000 satellites will be launched annually by 2030.¹ This increase is partially fueled by the proliferation of Distributed Space Systems (DSS)² launched by commercial, academic and government operators.

Ensuring the ongoing safety and sustainability of the space environment requires comprehensive tracking of the RSO population, in combination with robust online navigation capabilities for active space assets. RSO may be cooperative or uncooperative, known or unknown, and are typically tracked using passive measurements from ground stations (e.g. radar and optical telescopes).

*Postdoctoral Scholar, Aeronautics and Astronautics, Stanford University, 496 Lomita Mall, Stanford, CA 94305.

†Postdoctoral Scholar, Aeronautics and Astronautics, Stanford University, 496 Lomita Mall, Stanford, CA 94305.

‡M.Sc. Student, Aeronautics and Astronautics, Stanford University, 496 Lomita Mall, Stanford, CA 94305.

§Associate Professor, Aeronautics and Astronautics, Stanford University, 496 Lomita Mall, Stanford, CA 94305.

These ground-based measurements are subsequently used to generate a space object catalog. These procedures are resource-intensive and measurements are sparse, leading to limited responsiveness, accuracy, and coverage of the catalog.^{3,4} Active satellites in Earth orbit instead rely on Global Navigation Satellite Systems (GNSS) for PNT. Although such methods are accurate, they possess limited resilience in scenarios such as hardware failure, signal jamming, or deep space operation.

To address these shortcomings, this work proposes a new space-based PNT/SSA system known as Fast Autonomous Lost-in-space Catalog-based Optical Navigation (FALCON). FALCON is hosted on one or more observer satellites carrying passive optical sensors. Observers obtain bearing angle measurements to visible RSO targets. Detected RSO are matched to existing identities in an on-board RSO catalog and subsequently act as purely optical “beacons” providing known geometric information to the observer. This enables lost-in-space orbit determination for the observer, entirely independent of GNSS. Simultaneously, target measurements obtained on-orbit are fed back into the catalog for improved SSA coverage and timeliness. The only hardware requirement posed by FALCON is the presence of an optical sensor, which are already ubiquitous on modern spacecraft in the form of star trackers. Thus, FALCON may serve as a primary or secondary on-board PNT system with minimal additional hardware requirements.

Realization of FALCON presents fundamental algorithmic challenges related to target identification, batch orbit initialization, sequential orbit refinement, and system design. The contributions of this work address the first three of these aspects. First, it is assumed observers possess an on-board RSO catalog but do not possess a-priori observer orbit information. A Catalog Matching (CAM) algorithm is proposed for robustly matching onboard bearing angle measurements to catalog objects, for a lost-in-space observer, i.e. with no prior on the observer position, within the limited computational envelope of satellite avionics. The identified measurements are subsequently used to produce an observer position estimate. Next, a new Batch Orbit Determination (BOD) method is developed which solves for an initial observer orbit estimate, by exploiting catalog matches over several epochs. Finally, a Sequential Orbit Determination (SOD) framework is proposed which uses new measurements to refine the observer orbit estimate over time. This framework achieves the necessary accuracy to support a range of mission profiles while taking the limitations of angles-only observability into account.

Feasibility of the FALCON concept is verified through simulation results for the CAM, BOD and SOD systems, using a scenario drawn from flight data from the NASA Starling mission.⁵ After this introduction, an overview of the state-of-the-art in optical PNT/SSA is presented, followed by a description of the overall FALCON architecture. Algorithms for lost-in-space target identification, initial observer positioning, and sequential orbit refinement are then described, together with the pipeline used to generate physically-accurate imagery. Finally, each of the algorithms is demonstrated in simulation, followed by conclusions.

STATE OF THE ART

For FALCON, a primary challenge is the matching of target measurements to known identities in the RSO catalog, without observer position knowledge. A similar task is commonly performed with stars during attitude determination via a star tracker.^{6,7} However, stars are effectively stationary in an inertial frame, with angular positions known to sub-arcsecond accuracy; conversely, RSO are dynamic objects and RSO catalogs typically feature large orbit uncertainties of hundreds of arcseconds in Low Earth Orbit (LEO).^{4,8} This prevents the straightforward application of star identification techniques such as Pyramid or TETRA^{9,10} to the FALCON problem. Target identifi-

cation is also a critical component of ground-based RSO tracking, to which a variety of methods are applied, such as multi-hypothesis tracking and filtering,¹¹ statistical correlation in state or measurement space,^{12,13} and pre-sorting of catalog data to avoid costly processing of irrelevant objects.¹⁴ Each assumes good knowledge of the observer state, and limited computational resources are typically not accounted for, which is again incompatible with FALCON.

A second challenge is achieving accurate angles-only orbit estimation. Target relative orbit estimation using bearing angles from a known observer is a well-studied problem in literature, with several prior attempts in flight or on ground. These include the Advanced Rendezvous using GPS and Optical Navigation (ARGON) experiment,^{15,16} Autonomous Vision Approach Navigation and Target Identification (AVANTI) experiment,^{17,18} Starling Formation-Flying Optical Experiment (StarFOX),⁷ and Debris Detection Using Star Trackers (DeDUST)¹⁹ algorithm. Among these, only StarFOX has considered angles-only tracking of multiple targets while performing simultaneous observer absolute orbit determination. However, StarFOX relied upon a cooperative swarm with consistent presence in the camera Field of View (FOV) and its algorithms cannot be directly applied here. The StarFOX+ extended mission has proposed performing on-board optical RSO identification,⁷ similar to FALCON, but it currently assumes observer position knowledge and does not leverage the resulting measurements for observer state updates.

More broadly, recent work in angles-only Initial Relative Orbit Determination (IROD) has applied iterative batch least squares,²⁰ second-order dynamics modeling,^{21,22} and reformulation of the measurement space to maximize observability.²³ Prior research in angles-only sequential orbit determination has typically utilized Extended Kalman Filter (EKF) or Unscented Kalman Filter (UKF) architectures, which feature nonlinear models at a reasonable computational cost.²⁴ However, the inverse problem, i.e. observer absolute orbit estimation using bearing angles to known targets, is less well-studied. It is partially considered by recent angles-only architectures which apply multi-observer measurement sharing to estimate the observer's absolute orbit,²⁵ though in that case, targets are assumed persistent (not transient), and multiple observers are required for long-term convergence. It has also been explored in the context of triangulation where the known targets are celestial objects,²⁶ but these targets are again persistent with no uncertainty in their positions. Therefore, to realize the FALCON concept, the development of new algorithms and architectures is necessary, as described in the following.

FALCON ARCHITECTURE

Figure 1 illustrates core system blocks and data flow for a single instance of the proposed FALCON architecture, which operates on a generic observer spacecraft. There are four core software modules, in green: Image Processing (IMP), Catalog Matching (CAM), Batch Orbit Determination (BOD) and Sequential Orbit Determination (SOD). Figure 2 presents a broader diagrammatic view of the multi-observer target identification, observer localization and catalog refinement task. FALCON is implemented identically on each observer and proceeds via five core steps, which invoke the various algorithms in Figure 1. The steps are summarized as follows.

In Step 1), the onboard Vision Based Sensor (VBS) acquires the observed scene and provides an image to the IMP algorithm. IMP applies state-of-the-art image processing steps to: 1) detect regions of interest, 2) perform centroiding, 3) identify stars, and, 4) estimate the VBS attitude.⁷ The output of IMP is a list of inertial unit vectors to detected objects which may correspond to RSO.

In Step 2), the observer uses the CAM algorithm to identify which objects correspond to catalog

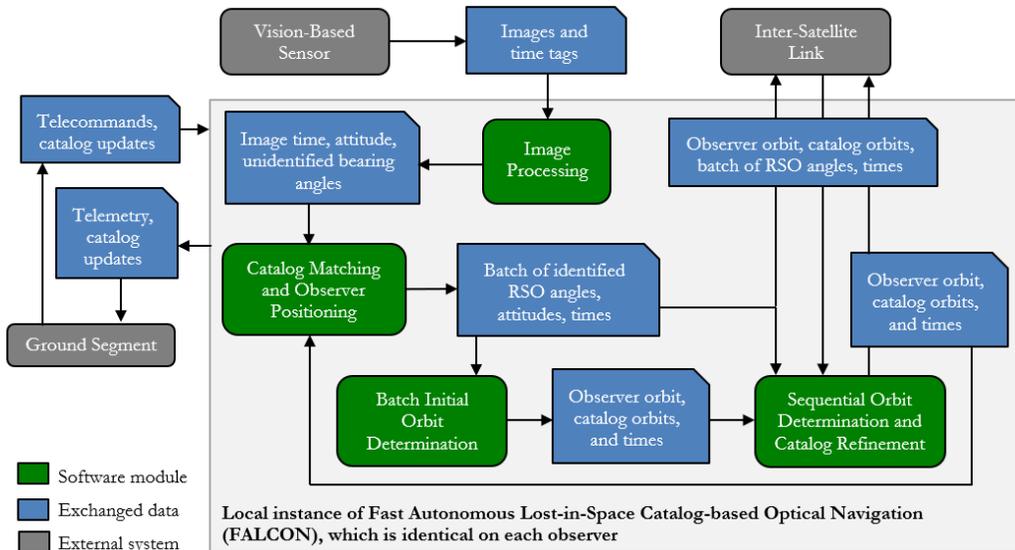


Figure 1: High-level block diagram of the FALCON architecture. Core algorithms are in green; exchanged data is in blue; and data sources are in dark gray.

targets, without a-priori knowledge of its own orbit. It is assumed the observer possesses knowledge of the image epoch, the sensor parameters (e.g., field of view, distortion models), and the sensor attitude, and has access to an onboard RSO catalog. Core outputs are the identities of RSO targets visible in the image and their inertial position vectors according to the catalog.

In Step 3), the observer leverages the geometric relationship between the inertial position of each identified target at that epoch, and target measurements at that epoch, to fix its own inertial position. Figure 3 presents a notional schematic of the problem: the red observer position vector is solved for by examining correspondence between blue unit vector measurements and red catalog position vectors. The BOD algorithm uses a batch of observer position estimates over sequential epochs to initialize the complete observer orbit estimate. Subsequently the SOD algorithm is used to refine

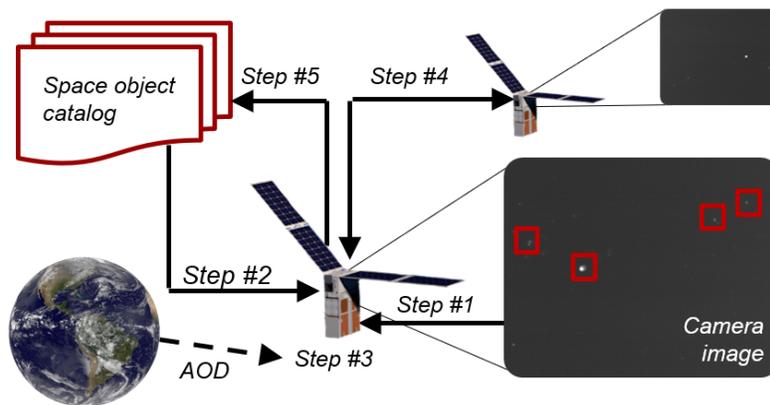


Figure 2: Notional illustration of data flow for target identification, observer localization and space catalog refinement.

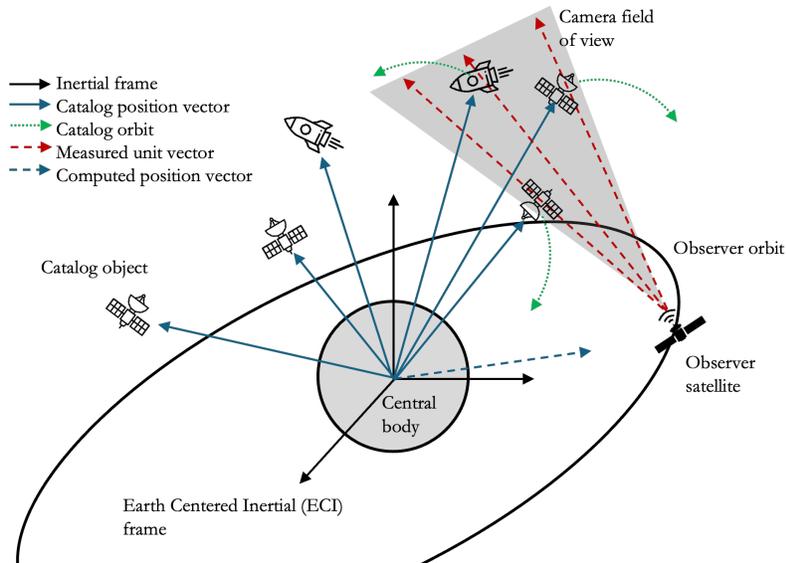


Figure 3: Notional schematic of the PNT problem for six catalog objects and three visible targets orbiting a central body.

the observer orbit estimate over time, using new bearing angle measurements from CAM. Observers may also leverage the onboard orbit estimate to simplify the target identification task in upcoming images.

In Step 4), if multiple observers are present, they may cooperatively share measurement information with one another over an Inter-Satellite Link (ISL). Distributed stereo-vision has been shown to significantly improve angles-only observability in practice.^{7,27}

In Step 5), if the observer has consistently imaged a target not present in the catalog, it may apply the BOD module to compute an initial orbit estimate for that target. Similarly, the observer may apply the SOD module to refine the orbit estimates of observed catalog targets, using on-orbit measurements. Catalog updates are then communicated to other observers in both a centralized and decentralized fashion (i.e., to a common database on ground and/or between observers).

In particular, this paper presents new methods to solve the fundamental algorithmic challenges of Steps 1) to 3). This corresponds to the CAM, BOD and SOD blocks in Figure 1. The new methods can be combined with an image processing system (such as that developed for the StarFOX flight experiment⁷) to enable a complete pipeline for lost-in-space optical positioning for a single observer spacecraft.

METHODOLOGY

Preliminaries

IMP produces angles-only measurements by computing the time-tagged bearing angles to objects detected in VBS images. First, define the observer VBS coordinate frame \mathcal{V} , which consists of orthogonal basis vectors $\hat{\mathbf{x}}^{\mathcal{V}}, \hat{\mathbf{y}}^{\mathcal{V}}, \hat{\mathbf{z}}^{\mathcal{V}}$ where $\hat{\mathbf{z}}^{\mathcal{V}} = \hat{\mathbf{x}}^{\mathcal{V}} \times \hat{\mathbf{y}}^{\mathcal{V}}$ is aligned with the camera boresight. Bearing angles consist of azimuth and elevation $[\alpha, \epsilon]^T$ and subtend the line-of-sight vector $\delta \mathbf{r}^{\mathcal{V}}$ from the observer to the target. The measurement model \mathbf{y} for the bearing angles from observer to

target is²⁴

$$\delta \mathbf{r}^{\mathcal{V}} = \mathbf{r}_{\text{tar}}^{\mathcal{V}} - \mathbf{r}_{\text{obs}}^{\mathcal{V}} = [\delta r_x^{\mathcal{V}}, \delta r_y^{\mathcal{V}}, \delta r_z^{\mathcal{V}}]^{\top} \quad (1)$$

$$\mathbf{y}^{\mathcal{V}}(\delta \mathbf{r}^{\mathcal{V}}) = [\alpha^{\mathcal{V}} \quad \epsilon^{\mathcal{V}}]^{\top} = \left[\arcsin \left(\frac{\delta r_y^{\mathcal{V}}}{\|\delta \mathbf{r}^{\mathcal{V}}\|_2} \right) \quad \arctan \left(\frac{\delta r_x^{\mathcal{V}}}{\delta r_z^{\mathcal{V}}} \right) \right]^{\top} \quad (2)$$

where \mathbf{r} is a Cartesian position vector. The unit vector $\hat{\mathbf{v}}$ corresponding to $[\alpha, \epsilon]^{\top}$ is $\hat{\mathbf{v}}^{\mathcal{V}} = \frac{\delta \mathbf{r}^{\mathcal{V}}}{\|\delta \mathbf{r}^{\mathcal{V}}\|_2}$.

It is also useful to define the Radial/Tangential/Normal (RTN) frame of the observer, denoted \mathcal{R} . It is centered on the observer²⁸ and defined by orthogonal basis vectors with $\hat{\mathbf{x}}^{\mathcal{R}}$ axis directed along the observer's absolute position vector; $\hat{\mathbf{z}}^{\mathcal{R}}$ axis directed along the observer's orbital angular momentum vector; and $\hat{\mathbf{y}}^{\mathcal{R}} = \hat{\mathbf{z}}^{\mathcal{R}} \times \hat{\mathbf{x}}^{\mathcal{R}}$.

Measurements and states are referenced with respect to an inertial reference frame centered on an arbitrary central body, denoted \mathcal{I} . Here, the Earth-Centered Inertial (ECI) J2000 frame is used.²⁸ Bearing angles are related to \mathcal{I} by rotating $\delta \mathbf{r}^{\mathcal{V}}$ into \mathcal{I} , as per $\delta \mathbf{r}^{\mathcal{I}} = {}^{\mathcal{V}}\mathbf{R}^{\mathcal{I}} \delta \mathbf{r}^{\mathcal{V}}$ where ${}^{\mathcal{V}}\mathbf{R}^{\mathcal{I}}$ denotes a rotation from frame \mathcal{V} into frame \mathcal{I} . This rotation is typically computed by performing attitude determination using stars identified by the VBS.

For propagation of catalog orbits within FALCON, analytical models of perturbed orbital motion in an integration constant state space are leveraged. Specifically, the Simplified General Perturbation (SGP4) model is applied to rapidly propagate Two Line Elements (TLE) in the catalog. Though propagation accuracy is somewhat coarse, and higher fidelity models could be used, this enables highly efficient propagation of many catalog orbits on board. Note that TLE are generally defined with respect to the inertial True Equator Mean Equinox (TEME) frame,²⁹ denoted \mathcal{T} , and must be rotated into \mathcal{I} when appropriate.

Synthetic Rendering of Stars and RSO

A radiometrically consistent rendering pipeline is put in place to accurately emulate the scene observed by the VBS on board. Four steps are followed to generate the scene: 1) the apparent magnitude of the star or the RSO is computed, 2) radiometric properties are derived based on magnitude values and a set of empirically-derived hyperparameters, 3) objects are rendered using a Moffat distribution centered on the subpixel centroid of the light source and 4) a Gaussian background noise is added to the whole image. In particular, the local pixel brightness $N_{\text{DN}}(u, v)$ provided by the Moffat distribution is given by

$$N_{\text{DN}}(u, v) = N_{\text{DN}, \text{peak}} \left(1 + \frac{(u - \mu_u)^2 + (v - \mu_v)^2}{\alpha^2} \right)^{-\beta} \quad (3)$$

where (u, v) are discrete pixel coordinates on the image plane, $N_{\text{DN}, \text{peak}}$ is the peak pixel brightness, (μ_u, μ_v) is the subpixel centroid of the source, α is the point-spread width, and β is a sharpness parameter. Specifically, β is expressed as a function of the total pixel brightness $N_{\text{DN}}^{\text{tot}}$ via

$$\beta = 1 + \pi \alpha^2 \frac{N_{\text{DN}}^{\text{tot}}}{N_{\text{DN}, \text{peak}}} \quad (4)$$

The photon flux F arriving at the sensor aperture per unit area and per unit time is determined using radiometric models. The flux is computed relative to a reference star using standard magnitude

scaling laws^{30,31}

$$F_{\text{star}} = \Gamma P_{\text{ref}} \cdot 10^{-0.4(m_{\text{star}} - m_{\text{ref}})} \quad [\text{photons}/\text{m}^2/\text{s}] \quad (5)$$

$$F_{\text{obj}} = K \cdot 10^{-0.4m_{\text{obj}}} \quad [\text{photons}/\text{m}^2/\text{s}] \quad (6)$$

where $K = 5.6 \times 10^{10}$ and $m_{\text{star}}, m_{\text{ref}}, m_{\text{obj}}$ are the apparent magnitudes of the observed star, a reference star and the observed RSO respectively. Γ is the ratio between the camera bandwidth and the photon energy of the signal, and P_{ref} is a reference energy flux. The apparent magnitude of stars is tabulated and retrieved from a star catalog.³² RSO magnitudes are instead calculated as^{30,33}

$$m_{\text{obj}} = m_{\text{sun}} - 2.5 \log_{10} \left(\frac{2\rho}{3\pi^2} \frac{A}{R^2} [(\pi - \Phi_s) \cos \Phi_s + \sin \Phi_s] \right) \quad (7)$$

where $m_{\text{sun}} = -26.74$ is the sun's apparent magnitude, $\rho \approx 0.17$ is the object's bond albedo, R is distance from the observer, A is cross-sectional area and Φ_s is the solar phase angle. Fluxes F_j represent the incoming photon rate per unit aperture area, and the total number of photons collected by the star tracker over the exposure duration Δt is then given by

$$N_\gamma = S \cdot \frac{\pi d_{\text{lens}}^2}{4} \cdot \Delta t \cdot F_j \quad [\text{photons}] \quad (8)$$

where S is a sensitivity factor capturing the quantum efficiency and optical transmittance of the system, and d_{lens} is the lens diameter. Using as reference star α -Lyrae and typical values found in literature³¹ the values reported in Table 1 are applied. For the sensor, a lens diameter of 50 mm and an exposure time of 120 ms are considered.³⁴

Table 1: Selected properties for the reference star.

Property	Value
m_{ref}	0.03
P_{ref}	$3.44 \times 10^{-2} \text{ W}/\text{m}^2$
Γ	$9.7894 \times 10^{11} \text{ m}/\text{J}$
S	0.4

The analog photon signal is subsequently converted to a digital signal using the full-well capacity $N_{\gamma, \text{max}}$ and the bit depth n_{bits} . The resulting total digital pixel value is thus computed as

$$N_{\text{DN}}^{\text{tot}} = \frac{2^{n_{\text{bits}}} - 1}{N_{\gamma, \text{max}}} \cdot N_\gamma \quad [\text{DN}] \quad (9)$$

assuming $n_{\text{bits}} = 10$. The full-well capacity $N_{\gamma, \text{max}}$, the ratio $r = N_{\text{DN}}^{\text{tot}}/N_{\text{DN}, \text{peak}}$ in Equation 4, and the point-spread width α in Equation 3 are each estimated for the proposed application by means of a calibration process. This approach consists in performing a least square fitting of stars detected in real flight images from the Starling mission⁷ with the Moffat distribution defined in Equation 3, leading to the estimated values reported in Table 2.

The NASA Starling mission obtained flight imagery from Blue Canyon Nano Star Trackers on each satellite, possessing a resolution of 1280×1024 , a FOV of $12^\circ \times 10^\circ$, and a maximum detectable magnitude of ~ 7.5 . Figure 4 presents a side-by-side comparison of a synthetically rendered image and the corresponding flight image with star magnitudes labeled. Similarly, Figure 5 compares a synthetically rendered RSO (a Starling spacecraft) against the same object imaged in flight. In both cases the Moffat model results pretty accurate in reproducing the point spread function behaviour of the flight image.

Table 2: Estimated values of the radiometric hyperparameters used for rendering.

Parameter	Value
$N_{\gamma,max}$	1.6×10^4
r	0.0445
α	1.0222

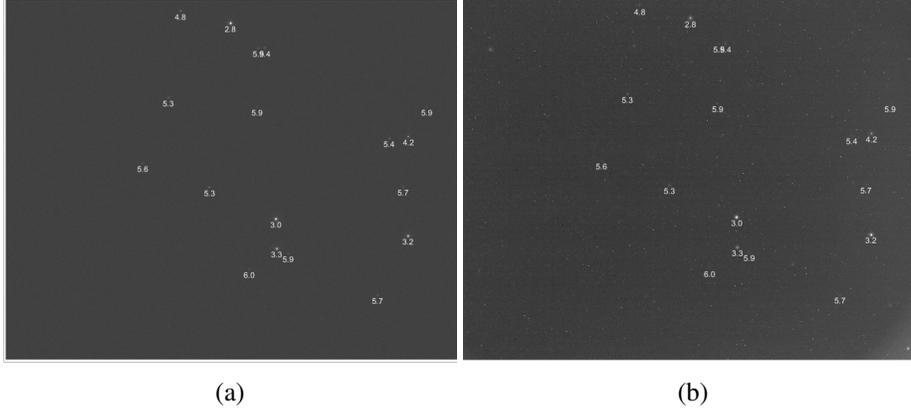


Figure 4: Star rendering comparison between (a) synthetic scene, and (b) flight image.

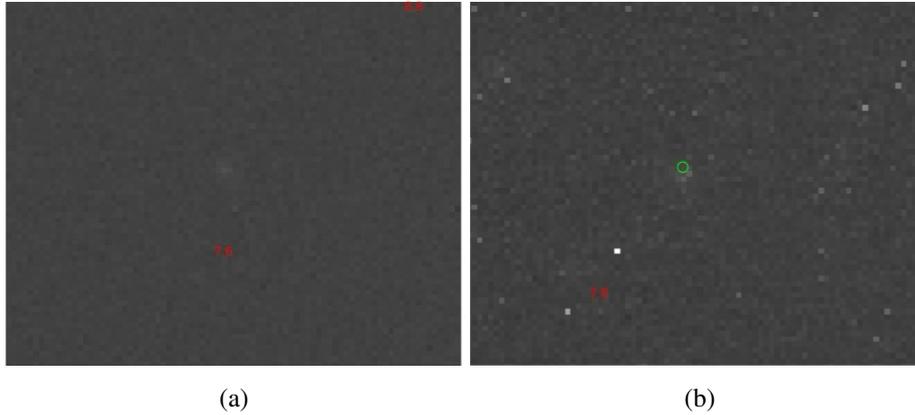


Figure 5: RSO rendering comparison between (a) synthetic scene, and (b) flight image. The green circle indicates the location of the RSO in the flight image.

Image Processing

The objective of the IMP module is to produce batches of time-tagged bearing angle measurements to any potential RSO in the image, so they can be identified by the CAM algorithm. To minimize possible sources of ambiguity, it is desirable to remove other types of objects (e.g. stars and pixel noise) from the list of measurements provided to CAM. This work leverages the IMP system developed for the StarFOX experiment⁷ flown on the NASA Starling mission.

First, an iterative weighted center of gravity³⁵ algorithm is used to simplify the raw VBS image into a list of pixel cluster centroids. Centroids are converted to Line of Sight (LOS) unit vectors in

the VBS frame using a calibrated sensor model. Next, the Pyramid star identification algorithm⁹ is applied to remove stars from the list of vectors. Uncataloged stars are detected by considering objects with unchanging inertial unit vectors between images. Similarly, camera hotspots are removed by considering objects with unchanging pixel coordinates. The VBS attitude is estimated by applying the q-method to identified star LOS vectors in the \mathcal{I} and \mathcal{V} frames.³⁶ The remaining minimalistic set of inertial unit vectors likely corresponds to known or unknown RSO in the FOV.

Lost-in-Space Catalog Matching

Consider the problem of a spacecraft located at $\mathbf{r}_{\text{obs}}^{\mathcal{I}}$ observing two RSO. The IMP extracts from the FOV, at a known epoch t_k two LOS measurements, i.e. $\mathbf{v}_1^{\mathcal{V}}$ and $\mathbf{v}_2^{\mathcal{V}}$, in the \mathcal{V} frame. Assuming an attitude determination solution is available on board, the two LOS can be expressed in \mathcal{I} .

The expected position $\mathbf{c}_i^{\mathcal{I}}$ of catalog objects at time t_k can be grouped in a matrix $\mathbf{C}(t)$

$$\mathbf{C}(t) = [\mathbf{c}_1^{\mathcal{I}} \quad \dots \quad \mathbf{c}_i^{\mathcal{I}} \quad \dots \quad \mathbf{c}_N^{\mathcal{I}}] \quad (10)$$

with N begin the number of objects in the catalogs. The matching thus aims to find the column indices K_i associated with the objects observed in the FOV. Once K_i are found, the observer position \mathbf{r}_{obs} can be estimated as the geometric intersection of the two LOS or through a least squares minimization if more than two LOS are considered.²⁶

In this work, the indices K_i are found by efficiently looping over the catalog objects and exploiting the fact that two LOS, $\mathbf{v}_i^{\mathcal{I}}$ and $\mathbf{v}_j^{\mathcal{I}}$, define a family of observation planes $\Pi_{ij}(\mathbf{v}_i^{\mathcal{I}}, \mathbf{v}_j^{\mathcal{I}})$ with normal $\mathbf{n}_{ij}^{\mathcal{I}} = \mathbf{v}_i^{\mathcal{I}} \times \mathbf{v}_j^{\mathcal{I}}$, where the offset of this plane depends on the observer position. The matching is then performed by finding the two catalog objects that belong to the same plane of this family, and comply with a certain set of visibility constraints, as shown in Figure 6.

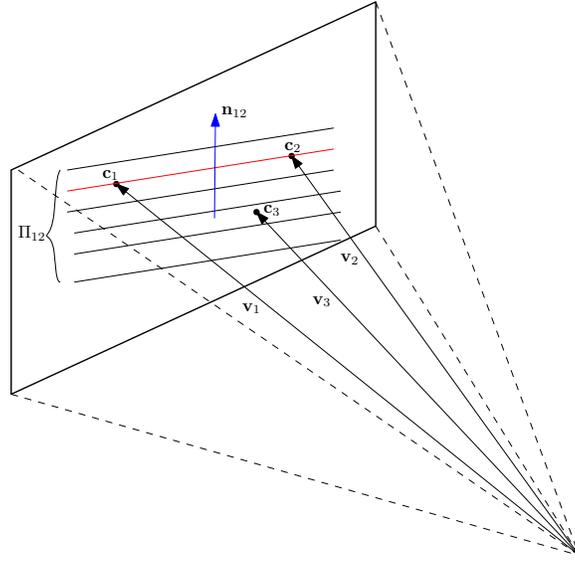


Figure 6: Representative example of the matching algorithm. The RSO c_1 and c_2 are matched with the LOS \mathbf{v}_1 and \mathbf{v}_2 because they both belong to the plane Π_{12} (in red).

In particular, the set of N LOS $\{\mathbf{v}_1, \dots, \mathbf{v}_N\}$ detected by the image processing at time t_k , is grouped in $N - 1$ tuples as $\{\mathbf{v}_1, \mathbf{v}_2\}, \{\mathbf{v}_1, \mathbf{v}_3\}, \dots, \{\mathbf{v}_1, \mathbf{v}_N\}$. For each tuple $\{\mathbf{v}_1, \mathbf{v}_j\}$ the algorithm

loops over the columns of $\mathbf{C}(t_k)$. For each column \mathbf{c}_i , the equation of the observation plane crossing \mathbf{c}_i and defined by \mathbf{n}_{1j} is then given by

$$\Pi_{\mathbf{c}_i}(\mathbf{r}|\mathbf{n}_{1j}) = \{\mathbf{r} \in \mathbb{R}^3 \text{ s.t. } \mathbf{n}_{1j} \cdot \mathbf{r} = -D(\mathbf{c}_i) = \mathbf{n}_{1j} \cdot \mathbf{c}_i\} \quad (11)$$

If no uncertainty is considered, \mathbf{c}_i is then a match for \mathbf{v}_1 if there exist a $\mathbf{c}_j \neq \mathbf{c}_i$ in $\mathbf{C}(t_k)$ that belongs to the plane $\Pi_{\mathbf{c}_i}$. If this is true it means that \mathbf{c}_j is also a match for \mathbf{v}_j . An efficient approach to assess this, is developed by looping, for each tuple $\{\mathbf{v}_1, \mathbf{v}_j\}$, over the columns \mathbf{c}_i of $\mathbf{C}(t_k)$, computing the offset $D(\mathbf{c}_i) = -\mathbf{n}_{1j} \cdot \mathbf{c}_i$ and then sort these values. This leads to a matrix $\mathbf{C}^S(t)$

$$\mathbf{C}^S(t) = \begin{bmatrix} \mathbf{c}_1^{S,\mathcal{I}} & \dots & \mathbf{c}_i^{S,\mathcal{I}} & \dots & \mathbf{c}_N^{S,\mathcal{I}} \end{bmatrix} \quad (12)$$

such that $D(\mathbf{c}_{i+1}^S) > D(\mathbf{c}_i^S)$. The matching is then implemented by looping over the columns of \mathbf{C}^S , computing the distance $d_i = D(\mathbf{c}_{i+1}^S) - D(\mathbf{c}_i^S)$, and then picking the minimum value. To reduce possible ambiguities associated with catalog objects that lie on the plane but are very far away from each other, or otherwise violate visibility constraints, the search space is reduced by adding a range constraint

$$\|\mathbf{c}_{i+1} - \mathbf{c}_i\| < 2\Delta r_{\max} \quad (13)$$

and two orientation-preserving constraints

$$-(\mathbf{v}_1 - (\mathbf{v}_1 \cdot \mathbf{v}_2) \mathbf{v}_2) \cdot (\mathbf{c}_{i+1}^S - \mathbf{c}_i^S) > 0 \quad (14)$$

$$(\mathbf{v}_2 - (\mathbf{v}_1 \cdot \mathbf{v}_2) \mathbf{v}_1) \cdot (\mathbf{c}_{i+1}^S - \mathbf{c}_i^S) > 0 \quad (15)$$

that are associated with the geometric relative orientation of the RSO in the FOV. Two additional constraints are added to reduce the feasible set by requiring the observed target to be below a maximum visibility range Δr_{\max} .

$$\|\mathbf{c}_i - \hat{\mathbf{r}}_{obs}(\mathbf{c}_i^S, \mathbf{c}_{i+1}^S)\| \leq \Delta r_{\max} \quad (16)$$

$$\|\mathbf{c}_j - \hat{\mathbf{r}}_{obs}(\mathbf{c}_i^S, \mathbf{c}_{i+1}^S)\| \leq \Delta r_{\max} \quad (17)$$

The application of these constraints reduces the size of the matrix \mathbf{C}^S and makes the two observed targets consecutive elements.

When uncertainties in the RSO states are considered, the correct match is not necessarily given by the minimum value of $D(\mathbf{c}_{i+1}^S) - D(\mathbf{c}_i^S)$. A more robust approach is thus adopted, returning all sets of feasible matches that are compliant with the constraints in Equations (13) to (17) and whose distance from the reference plane is below a certain threshold ΔD_{tol}

$$D(\mathbf{c}_j^S) - D(\mathbf{c}_i^S) < \Delta D_{\text{tol}} \quad (18)$$

Following this pipeline, for each tuple $\{\mathbf{v}_1, \mathbf{v}_j\}$ a matrix $\mathbf{K}_{1j} \in \mathbb{R}^{N_{1j} \times 2}$ of feasible matches is identified, with N_{1j} being the number of feasible matches corresponding to observations \mathbf{v}_1 and \mathbf{v}_j . Each row of \mathbf{K}_{1j} contains the indexes of the catalog objects associated with the \mathbf{c}_1 and \mathbf{c}_j RSO. The $N - 1$ feasibility matrices are then compared generating the feasible matches for the second column components. The satisfaction of feasibility constraints, Equations (13) to (17), on this new set is checked and the combination that violates them are removed from the list generating $N - 1$ reduced

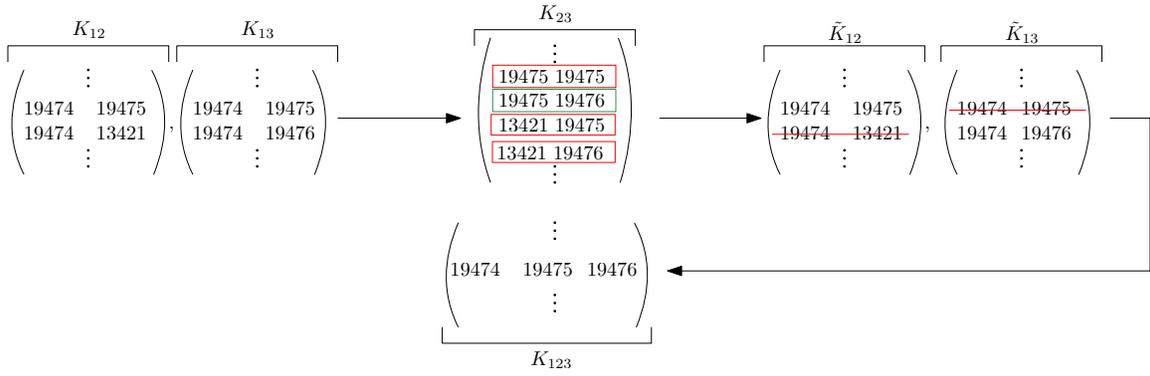


Figure 7: Explanatory example of the consistency check reduction.

feasible tuples $\tilde{\mathbf{K}}_{1j}$. These are then combined to obtain N_f final feasible matches $\mathbf{K}(t_k) \in \mathbb{R}^{N_f, N}$. This process is illustrated for the simplified case of only three LOS in Figure 7.

A prior on the observer position, for example provided by either ground-based or on-board orbit determination can be used to further reduce the search space, but the proposed methodology works well even in a lost-in-space scenario in which no information on the spacecraft position is provided. While reducing by several order of magnitudes the search space, the procedure illustrated in Figure 7 can still lead to a $\mathbf{K}(t_k)$ with more than one row ($N_f > 1$), typically less than five. In this case, under the assumption that the star tracker acquisition frequency is such that the observer can detect the same RSO for a few consecutive frames, sequential observations can be exploited to solve the remaining ambiguity and find a unique solution for the match. An example of this procedure is shown in Figure 8 in which three consecutive observation of the NASA starling spacecraft are sufficient to uniquely identify the correct match starting from a catalog size of more than 20,000 objects with no a priori information on the observer position.

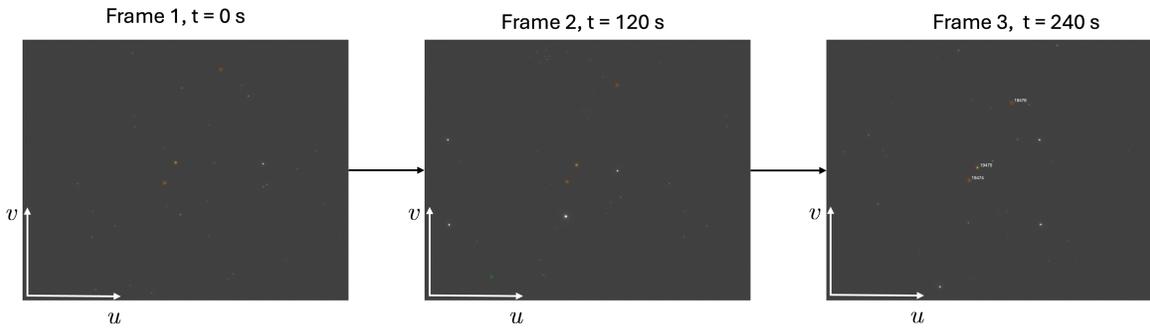


Figure 8: Lost-in-space initial acquisition. After three acquisition the matching algorithm is capable to assign a unique ID to each RSO detected in the FOV.

The computation time associated with this new pipeline clearly depends on the number of observations but, in a prototype implementation in MATLAB R2022a running on a MacBook Air M1 processor clocked at 3.5 GHz, computation time remained below 2 seconds which shows promise for good performance in flight-representative hardware.

Batch Orbit Determination

The Direct Linear Transform (DLT) algorithm^{26,37} is used to triangulate observer position given the outputs from CAM. Consider $n \geq 2$ LOS measurements $\{\hat{\mathbf{v}}_i\}_{i=1}^n$ corresponding to points with known 3D locations $\{\mathbf{r}_i\}_{i=1}^n$. Each LOS measurement must be parallel with the relative position vector of the RSO with respect to the observer at that epoch. Taking the cross product leads to

$$\begin{aligned} \hat{\mathbf{v}}_i \times \delta \mathbf{r}_i &= \hat{\mathbf{v}}_i \times (\mathbf{r}_i - \mathbf{r}_{\text{obs}}) = \mathbf{0} \\ \rightarrow \hat{\mathbf{v}}_i \times \mathbf{r}_{\text{obs}} &= \hat{\mathbf{v}}_i \times \mathbf{r}_i \end{aligned} \quad (19)$$

If many measurements are available, they may be combined into a linear system

$$\begin{pmatrix} [\hat{\mathbf{v}}_1 \times] \\ [\hat{\mathbf{v}}_2 \times] \\ \vdots \\ [\hat{\mathbf{v}}_n \times] \end{pmatrix} \mathbf{r}_{\text{obs}} = \begin{pmatrix} [\hat{\mathbf{v}}_1 \times] \mathbf{r}_1 \\ [\hat{\mathbf{v}}_2 \times] \mathbf{r}_2 \\ \vdots \\ [\hat{\mathbf{v}}_n \times] \mathbf{r}_n \end{pmatrix} \quad (20)$$

where $[\cdot \times]$ is the skew-symmetric cross-product matrix $[\hat{\mathbf{v}} \times] \mathbf{r} = \hat{\mathbf{v}} \times \mathbf{r}$. This linear system can be solved via least squares to estimate observer position \mathbf{r}_{obs} .

The objective of the BOD system is to produce an initial observer orbit estimate $\boldsymbol{\alpha} \in \mathbb{R}^6$ and accompanying covariance $\boldsymbol{\Sigma}_{\boldsymbol{\alpha}} \in \mathbb{R}^{6 \times 6}$ at estimation epoch t_{est} . The method applied in this work mirrors the BOD approach developed for StarFOX.³⁸ The algorithm operates on bearing angles $\mathbf{y}_i(t_j)$ collected over epochs $j = 1, \dots, N_t$, collectively denoted \mathbf{t} . Define $N_o(t_j)$ as the number objects visible at epoch j , such that the total number of measurements in the batch is $N_y = \sum_{j=1}^{N_t} N_o(t_j)$. The concatenated measurement batch vector $\mathbf{z}(\boldsymbol{\alpha}(t_{\text{est}}), \mathbf{t})$ is

$$\mathbf{z}(\boldsymbol{\alpha}(t_{\text{est}}), \mathbf{t}) = \begin{pmatrix} \mathbf{y}_1(\boldsymbol{\alpha}(t_{\text{est}}), t_1) \\ \mathbf{y}_2(\boldsymbol{\alpha}(t_{\text{est}}), t_1) \\ \vdots \\ \mathbf{y}_{N_o(N_t)-1}(\boldsymbol{\alpha}(t_{\text{est}}), t_{N_t}) \\ \mathbf{y}_{N_o(N_t)}(\boldsymbol{\alpha}(t_{\text{est}}), t_{N_t}) \end{pmatrix} \in \mathbb{R}^{2N_y} \quad (21)$$

The sensitivity matrix for the bearing angle measurements with respect to state components is

$$\mathbf{S}(\boldsymbol{\alpha}(t_{\text{est}}), \mathbf{t}) = \left. \frac{\partial \mathbf{z}(\boldsymbol{\alpha}(t_{\text{est}}), \mathbf{t})}{\partial \boldsymbol{\alpha}_{\text{est}}} \right|_{\boldsymbol{\alpha}(t_{\text{est}})} \in \mathbb{R}^{2N_y \times 6} \quad (22)$$

BOD initializes the state estimate at $t_{\text{est}} = t_j$ by using the $\mathbf{r}_{\text{obs}}(t_j)$ solution for position, and $[\mathbf{r}_{\text{obs}}(t_{j+1}) - \mathbf{r}_{\text{obs}}(t_{j-1})] / (t_{j+1} - t_{j-1})$ for velocity. BOD then performs iterative batch least squares refinement for each sample. Define the batch of measured angles from IMP as \mathbf{z}_{meas} , and a batch of modeled angles $\mathbf{z}_{\text{model}}$ computed by propagating $\boldsymbol{\alpha}$ and the onboard catalog to each epoch in \mathbf{t} . The difference between measured and modeled angles is

$$\Delta \mathbf{z} = \mathbf{z}_{\text{meas}} - \mathbf{z}_{\text{model}}(\boldsymbol{\alpha}(t_{\text{est}}), \mathbf{t}) \quad (23)$$

Subsequently, the state estimate update

$$\Delta \mathbf{z} = \mathbf{S}(\boldsymbol{\alpha}(t_{\text{est}}), \mathbf{t}) \Delta \boldsymbol{\alpha} \quad (24)$$

$$\boldsymbol{\alpha} \leftarrow \boldsymbol{\alpha} + \Delta \boldsymbol{\alpha} \quad (25)$$

is computed and the unknown $\Delta\alpha$ is solved for via least squares. Refinement continues until $\Delta\alpha$ is smaller than a convergence threshold.

It is then necessary to estimate the state uncertainty. The equation applied is $\mathbf{R} = \mathbf{S}\Sigma\mathbf{S}^\top$ which maps measurement batch covariance $\mathbf{R} \in \mathbb{R}^{2N_y \times 2N_y}$ to state covariance Σ_α via sensitivity \mathbf{S} .²⁰ The covariance \mathbf{R}_{meas} associated with individual bearing angles is estimated via

$$\mathbf{R}_{\text{meas}} = \frac{1}{N_y} \sum_{j=1}^{N_y} \left((\mathbf{z}_{\text{meas}}(t_j) - \mathbf{z}_{\text{model}}(t_j))(\mathbf{z}_{\text{meas}}(t_j) - \mathbf{z}_{\text{model}}(t_j))^\top \right) \in \mathbb{R}^{2 \times 2} \quad (26)$$

This estimation of sensor noise using post-fit measurement residuals allows operation even when an accurate a-priori model of sensor noise is unavailable. BOD is then able to estimate the initial state uncertainty from known terms via $\Sigma_\alpha = \mathbf{S}^* \mathbf{R} \mathbf{S}^{*\top}$ where $*$ denotes a pseudoinverse.

Sequential Orbit Determination

The SOD module continually refines the orbit estimate of the observer by utilizing new bearing angle measurements from CAM. SOD is based on an adaptive, efficient UKF with fully nonlinear dynamics and measurement models. In comparison to an EKF, the UKF captures higher order effects in the probability distribution, which enables improved angles-only navigation performance without excessive computational cost.^{7,24}

SOD leverages quasi-nonsingular absolute Orbit Elements (OE) as the observer state parametrization. The absolute orbit α is defined as

$$\alpha = [a \ e_x \ e_y \ i \ \Omega \ u]^\top = [a \ e \cos \omega \ e \sin \omega \ i \ \Omega \ \omega + M]^\top \quad (27)$$

where $a, e, i, \Omega, \omega, M$ are the Keplerian orbit elements. SOD propagates the observer orbit using fourth-order Runge-Kutta integration of Gauss' Variational Equations (GVE). The osculating OE of the observer evolve according to $\dot{\alpha} = G(\alpha)\mathbf{d}^\mathcal{R}$ where $G \in \mathbb{R}^{6 \times 3}$ is the GVE matrix³⁹ and $\mathbf{d}^\mathcal{R}$ is the perturbing acceleration expressed in \mathcal{R} . Depending on the orbit regime, common perturbations include spherical harmonic gravity terms, atmospheric drag, third-body gravity and solar radiation pressure.⁴⁰ Onboard, FALCON applies a 10x10 GGM01S⁴¹ spherical harmonic gravity model, a Harris-Priester atmosphere model with cannonball drag, and a 30 second RK4 step.

At each filter epoch t_{est} , the observer orbit estimate is propagated to the epoch of the new image t_{im} , and the uncertainty associated with each RSO bearing angle from IMP is computed. For a single bearing angle $\mathbf{y} \in \mathbb{R}^2$, its measurement covariance $\mathbf{R}_{\text{meas}} \in \mathbb{R}^{2 \times 2}$ must encompass two contributions: sensor noise \mathbf{R}_{VBS} , and catalog TLE error $\mathbf{R}_{\text{cat}}(t_{\text{prop}})$, such that

$$\mathbf{R}_{\text{meas}} = \mathbf{R}_{\text{VBS}} + \mathbf{R}_{\text{cat}}(t_{\text{prop}}) \quad (28)$$

Here, constant Gaussian noise of 20'' (1σ) is assumed for the diagonal azimuth and elevation terms in \mathbf{R}_{VBS} . This emulates the performance of a modern CubeSat star tracker.^{6,38}

The effects of TLE uncertainties are more difficult to capture given their precise components are not well-known and vary depending on the RSO in question. To address this, FALCON combines model-agnostic and empirical methods. Recall that the IMP module computes an attitude estimate ${}^{\mathcal{I}}\vec{\mathbf{R}}_{\text{qm}}^\mathcal{V}$ for the image, i.e. the rotation such that $\hat{\mathbf{v}}^\mathcal{V} = {}^{\mathcal{I}}\vec{\mathbf{R}}_{\text{qm}}^\mathcal{V} \hat{\mathbf{v}}^\mathcal{I}$. The 'qm' subscript denotes the q-method algorithm³⁶ which provides an optimal estimate for ${}^{\mathcal{I}}\vec{\mathbf{R}}^\mathcal{V}$. Typically, the set of $\{\hat{\mathbf{v}}^\mathcal{V}\}_{i=1}^n$

and $\{\hat{\mathbf{v}}^I\}_{i=1}^n$ vectors is drawn from identified stars in the FOV. The resulting attitude is considered to be accurate within the bounds of star tracker noise, i.e. within tens of arcseconds.

Similarly, this attitude can be computed using LOS vectors to identified RSO. The TRIAD algorithm³⁶ provides a method to estimate ${}^I\vec{\mathbf{R}}_r^\nu$ from two $[\hat{\mathbf{v}}^\nu, \hat{\mathbf{v}}^I]$ sets. Consider $[\hat{\mathbf{v}}_i^\nu, \hat{\mathbf{v}}_i^I]$ provided by an identified star, and $[\hat{\mathbf{v}}_j^\nu, \hat{\mathbf{v}}_j^I]$ provided by an identified RSO. The majority of error in ${}^I\vec{\mathbf{R}}_r^\nu$ is introduced by the catalog error inherent in the $\hat{\mathbf{v}}_j^I$ term. The resulting angle θ_{ij} between RSO-derived attitude ${}^I\vec{\mathbf{R}}_r^\nu$ and reference attitude ${}^I\vec{\mathbf{R}}_{qm}^\nu$ prescribes a cone around $\hat{\mathbf{v}}_j^\nu$, which is the angular uncertainty for that RSO measurement. For a specific RSO j and all identified stars $i = 1, \dots, n$, the uncertainty $\theta_j = \sum_{i=1}^n \theta_{ij}/n$.

This provides a general model for the magnitude of angular uncertainty from catalog errors, but assumes a symmetric covariance shape. To capture shape properties, an empirical model is used. Define TLE uncertainties as a vector $\sigma_{\delta\alpha}$, where $\delta\alpha \in \mathbb{R}^6$ are the quasi-nonsingular Relative Orbital Elements (ROE) introduced by D'Amico.⁴² The ROE vector is $\delta\alpha = [\delta a, \delta\lambda, \delta e_x, \delta e_y, \delta i_x, \delta i_y]^\top$ where δa is the relative semimajor axis, $\delta\lambda$ is the relative mean longitude, $\delta e = [\delta e_x, \delta e_y]^\top$ is the relative eccentricity vector and $\delta i = [\delta i_x, \delta i_y]^\top$ is the relative inclination vector. The uncertainties applied here, derived from NASA Starling flight data,⁷ are

$$\sigma_{\delta\alpha} = [150, 500 + 1500t_{\text{prop}}, 300, 300, 300, 300]^\top \text{m} \quad (29)$$

where t_{prop} is catalog propagation time in days between the TLE epoch and the image epoch. The ROE uncertainty model is then converted to \mathbf{R}_{cat} in bearing angle space via an unscented transform. Finally, $\mathbf{R}_{\text{cat},j}$ is scaled such that the semimajor axis of the bearing angle covariance ellipse has the same magnitude as the attitude-derived uncertainty θ_j . This formulation aims to capture the time- and shape-dependence of TLE uncertainties (which are typically largest in the velocity direction and grow with time) and their projection onto the image plane.

The UKF then has access to all necessary measurement and uncertainty information with which to perform a measurement update. Sigma points are computed from the observer orbit estimate and covariance, and for each sigma point, a vector of modeled measurements $\mathbf{z}(t_{\text{est}})$ is computed using Equation 2. For N_y measurements in the image, $\mathbf{z} = [\mathbf{y}_1^\top, \dots, \mathbf{y}_{N_y}^\top]^\top \in \mathbb{R}^{2N_y}$ and

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_{\text{VBS}} + \mathbf{R}_{\text{cat},1} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \mathbf{R}_{\text{VBS}} + \mathbf{R}_{\text{cat},N_y} \end{bmatrix} \in \mathbb{R}^{2N_y \times 2N_y} \quad (30)$$

The modeled and true measurement vector are used to compute the state innovation and complete the standard UKF update.⁴³

VALIDATION SCENARIO

The proposed methodology is tested within the framework of the NASA Starling mission.⁷ Starling consists of four propulsive 6U CubeSats launched in July 2023, denoted SV1, SV2, SV3 and SV4. Figure 9 displays swarm evolution in the RTN frame over one orbit, from flight data obtained on March 24, 2024. The SV3 VBS remains pointed in the velocity or tangential direction to ensure visibility of the other three swarm agents, and is thus selected as the observer spacecraft. The half-FOV of the onboard VBS is approximately 5 degrees, as presented in Figure 9.

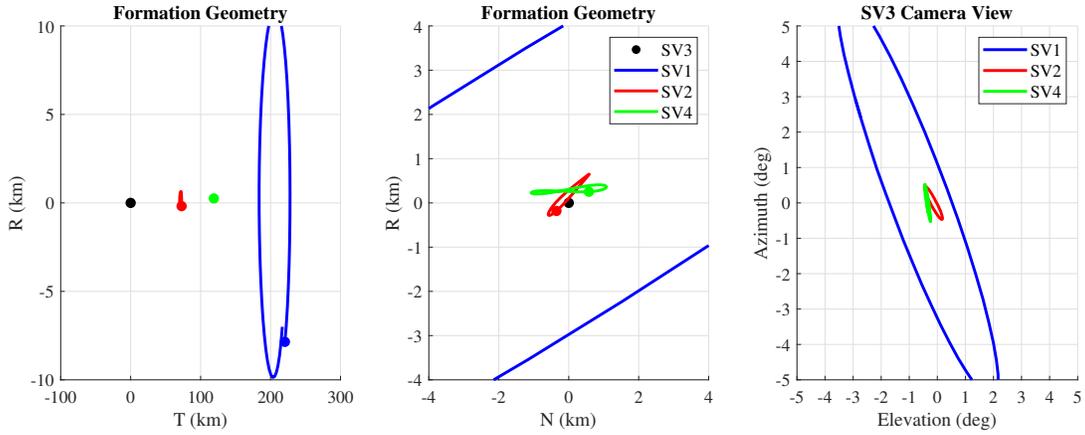


Figure 9: Geometry of the Starling swarm on March 24 2024, at 00:00:00 UTC. RT and RN motion is plotted in curvilinear coordinates with respect to SV3.

Post-processed GNSS flight data is used to define the initial Cartesian state of the observer spacecraft as described in Table 3. A catalog of 21200 space objects is downloaded from SpaceTrack.org⁴⁴ in the form of: 1) TLE sets defined at a given epoch t_{cat} , and 2) object mean cross-sectional area to be used in Equation 7.

Table 3: Initial position and velocity in ECI.

$r_x^{\mathcal{I}}$ (km)	$r_y^{\mathcal{I}}$ (km)	$r_z^{\mathcal{I}}$ (km)	$v_x^{\mathcal{I}}$ (km/s)	$v_y^{\mathcal{I}}$ (km/s)	$v_z^{\mathcal{I}}$ (km/s)
1041.845	4752.220	-4962.418	0.589	-5.511	-5.155

The ground truth is simulated via numerical integration of the GVE by means of an high-fidelity astrodynamics simulator.⁴⁵ Modeled perturbations include 60×60 GGM05S spherical harmonic gravity; an NRLMSISE00 atmosphere model;⁴⁶ solar radiation pressure; and third-body lunisolar gravity, with a variable RK45 step of 1-10 seconds. The observer attitude is defined to always point the camera boresight in the velocity direction. To simulate catalog uncertainties, RSO states at t_{cat} are converted from TLE to RTN coordinates, are perturbed in position and velocity, then reconverted to TLE and propagated to the current simulation epoch using SGP4.²⁹ A Gaussian distribution with uncertainties of 1σ position uncertainties of 1000 m (along-track) and 300 m (radial and cross-track) and 1σ velocity uncertainties of 0.1 m/s (along-track) and 0.01 m/s (radial and cross-track) is used to perturb the TLE. Subsequently, the ground truth trajectory, nominal attitude profile, and perturbed catalog states are fed into the synthetic rendering pipeline to generate synthetic VBS images. The VBS provides an image to the IMP system every 120 seconds. The unperturbed catalog is also provided to FALCON for onboard use.

RESULTS

Catalog Matching Results

Performance of the matching algorithm is shown in Figure 10 by visualizing the RSO detected in the camera FOV over 10 orbits, across a total of 474 images. As anticipated, the majority of the RSO detected belong to the Starling swarm, as shown in Figure 9. However, some other objects

appear in the FOV at sparse intervals along the orbit. The CAM algorithm is capable of converging from a lost-in-space condition with only three consecutive observations (blue markers in the figure) of the Starling swarm. This allows approximate triangulation of the observer’s position in BOD (and initialization of a navigation filter in SOD), and subsequently, this position knowledge can be used to inform the algorithm and achieve more reliable performance in future epochs. Overall the CAM algorithm shows a success rate of 88 % when informed with the observer position, shown by green markers in the figure. There is also a small percentage of cases in which no feasible match is found (red markers). This typically happens when more than two LOS vectors lie on the same plane, within the bounds associated with measurement and orbit uncertainties, generating an ambiguity that the algorithm is incapable of handling. This suggests the need to pre-process measurements so the algorithm only operates on observables with a minimum out-of-plane separation (i.e. one larger than the uncertainty bound) to further improve its robustness.

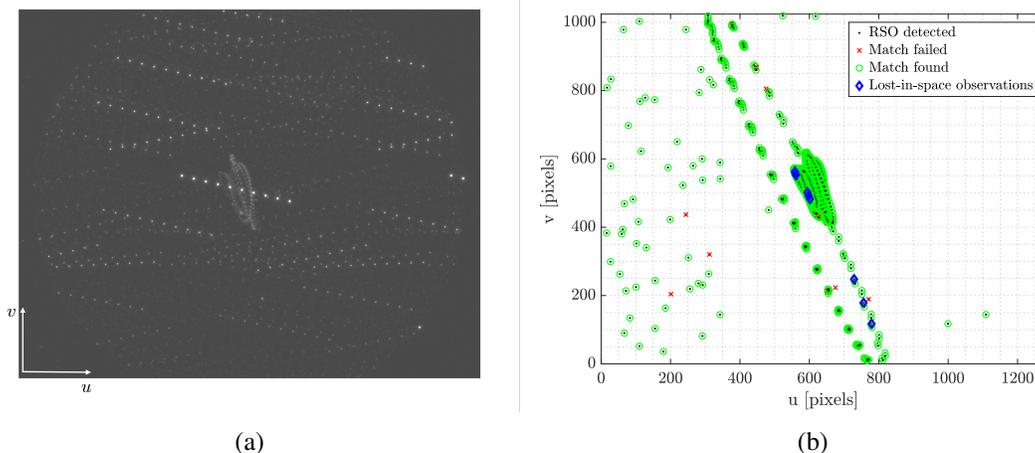


Figure 10: Matching algorithm performances. (a) shows the timelaps of the observed scene. (b) highlights the detected RSO (black), the ones for which the matching found a correct correspondence (green), the ones for which the matching failed (red) and the ones used for the lost-in-space initialization (blue).

Batch Orbit Determination Results

Table 4 presents errors in the DLT position solution for the first six images in which matches were found by CAM. In all cases, the three visible RSO corresponded to the Starling swarm members. Position error varies significantly from less than 1 km in Image 3 to more than 60 km in Image 6, and is heavily dependent on the Geometric Dilution of Precision (GDOP) of the current measurement set. When all Starling targets are in close proximity in the FOV, their bearing angles are similar, which results in larger ambiguities during triangulation and a larger GDOP value. More diverse measurement geometry is therefore beneficial for positioning, and the swarm scenario explored here presents a particularly challenging case.

Table 5 presents orbit initialization errors from BOD using the first 3-6 images in which matches were found. The OE are normalized by the observer’s semi-major axis for geometric interpretation. Errors remain bounded to approximately 10 km, and errors are well-characterized by the state uncertainty. The largest error component is in u or the along-track direction, as expected from swarm

Table 4: Position solution errors in BOD.

Image	RSO	Position Error (km)			GDOP
		r_x	r_y	r_z	
3	3	0.17	0.57	-0.45	6.25
4	3	0.14	4.87	1.08	8.20
5	3	0.34	15.99	2.99	12.44
6	3	2.05	64.74	5.67	26.28
7	3	1.90	38.70	-2.02	34.29
8	3	-1.04	-17.92	2.21	13.11

geometry; this is the primary direction in which GDOP manifests. Increasing the number of epochs in the batch tends to reduce overall error and uncertainty, though errors in u remain relatively unchanged. This is likely a bias caused by catalog errors. Removing catalog TLE errors altogether results in output state errors of less than 200 m in all OE, which indicates that typical catalog orbit errors have a larger effect on solution accuracy than measurement errors or dynamics modeling errors. Despite this, an effective initialization can be computed with as few as three images.

Table 5: Orbit initialization errors and 1σ uncertainties in BOD.

Epochs	a (km)	ae_x (km)	ae_y (km)	ai (km)	$a\Omega$ (km)	au (km)
3	-3.43 ± 18.48	1.35 ± 6.70	3.87 ± 17.40	-0.04 ± 0.20	-0.16 ± 0.50	-9.21 ± 11.18
4	-1.30 ± 8.19	0.48 ± 2.50	1.89 ± 7.93	-0.05 ± 0.12	-0.21 ± 0.34	-10.73 ± 6.95
5	-0.64 ± 4.77	0.25 ± 1.19	1.28 ± 4.73	-0.06 ± 0.08	-0.24 ± 0.25	-11.97 ± 5.95
6	-0.25 ± 3.21	0.14 ± 0.64	0.88 ± 3.25	-0.05 ± 0.05	-0.22 ± 0.19	-10.88 ± 5.80

Sequential Orbit Determination Results

The three-image BOD solution in Table 5 is used to initialize the UKF. At each epoch, new images are processed by IMP; resident space objects are identified by CAM; and identified measurements are used to refine the observer's orbit in SOD. Estimation performance across 10 orbits is shown in Figure 11, together with values for steady-state performance averaged across the final orbit. Angular OE are again scaled by the observer's semi-major axis for interpretation.

Rapid convergence of the observer orbit estimate is observed, despite varying numbers of visible RSO. After two orbits, uncertainties in $[a, ae_x, ae_y, ai, a\Omega]$ are on the order of 100 m and have approached steady-state values. Uncertainties in au are instead on the order of kilometers. This is because the VBS is aligned with in the flight direction, as are any measured LOS vectors. Thus, the system experiences some dilution of precision for observer position/velocity in the flight direction, as primarily captured by u . It is therefore recommended a sensor with a larger FOV or a variable attitude profile is used to improve estimation performance for this state component.

Furthermore, there is a steady-state bias in e_y and u . This is a consequence of the swarm scenario. Starling spacecraft are visible in the majority of images and each of their measurements is associated with a TLE error. This error was produced via an initial perturbation to their TLE and is propagated over time. Thus, measurement errors for Starling targets become correlated and the assumption of zero-mean Gaussian measurement noise in the UKF is violated. Figure 12 plots results from an idealized scenario in which the filter has access to the reference truth catalog. As expected, the

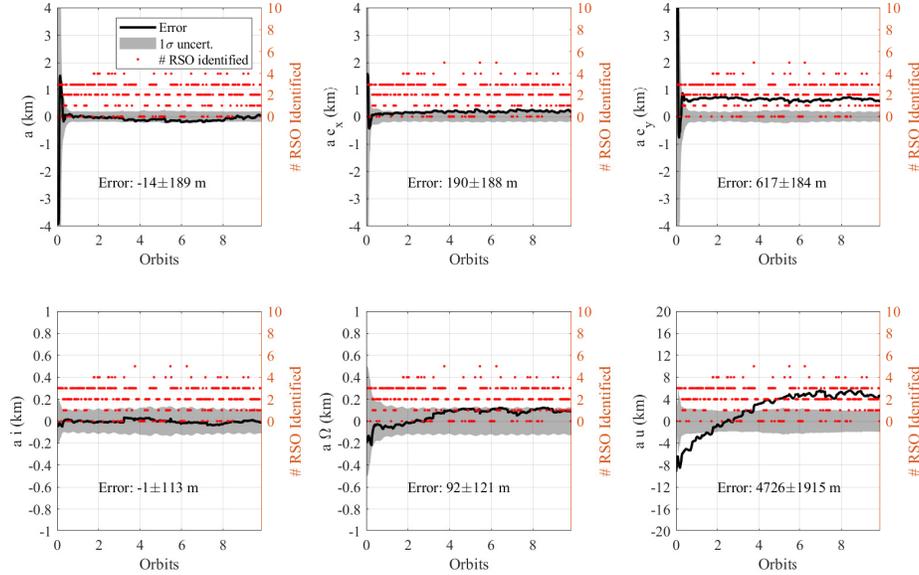


Figure 11: Observer orbit estimation errors and formal 1σ covariances in the validation scenario.

bias is removed and estimation errors are near zero. In particular, in this case, the estimate of the semi-major axis is accurate up to 0.008 %, performance comparable with deep-space angles-only navigation solutions exploiting LOS to well-known inner and outer planets of the Solar System.⁴⁷ Alternately, the bias could be eliminated over time by implementing the additional capability of simultaneous state refinement for RSO on board.

In general, most non-swarm scenarios will feature transient targets visible in single images only, producing catalog errors that are much less correlated. Figure 13 illustrates this outcome by ignoring all measurement updates featuring Starling swarm members. The resulting more varied set of targets eliminates obvious biases in the orbit elements even with TLE errors present. Estimation errors generally remain within 1σ bounds, indicating good health of the filter. However, the much lower frequency of measurement updates results in larger state uncertainties and longer convergence times.

Overall, results indicate resilient lost-in-space PNT can be achieved using exclusively optical measurements from single observer. The UKF framework supports robust estimation, provided measurement errors remain uncorrelated. Furthermore, the validation scenario presents sparse measurement availability and poor measurement geometry, and more accurate positioning could likely be achieved given a larger FOV, more frequent images, and more sensitive camera. For example, similar simulations using a $40^\circ \times 30^\circ$ FOV, 20-second image interval, and consistent detection of RSO up to 1500 km away produce an output observer position error of 113 ± 188 m (1σ).

Finally, it is useful to compare absolute orbit estimation performance to that demonstrated during the StarFOX experiment.²⁵ Instead of an RSO catalog, StarFOX applied distributed stereovision: multiple observers imaged common swarm targets and shared measurements over a crosslink, to achieve absolute orbit observability using only inter-satellite bearing angles. StarFOX used the same flight data and a direct comparison is provided in Table 6. In general, the use of RSO in FALCON provides stronger estimation performance for e_x, e_y, i, Ω , whereas the stereovision applied in

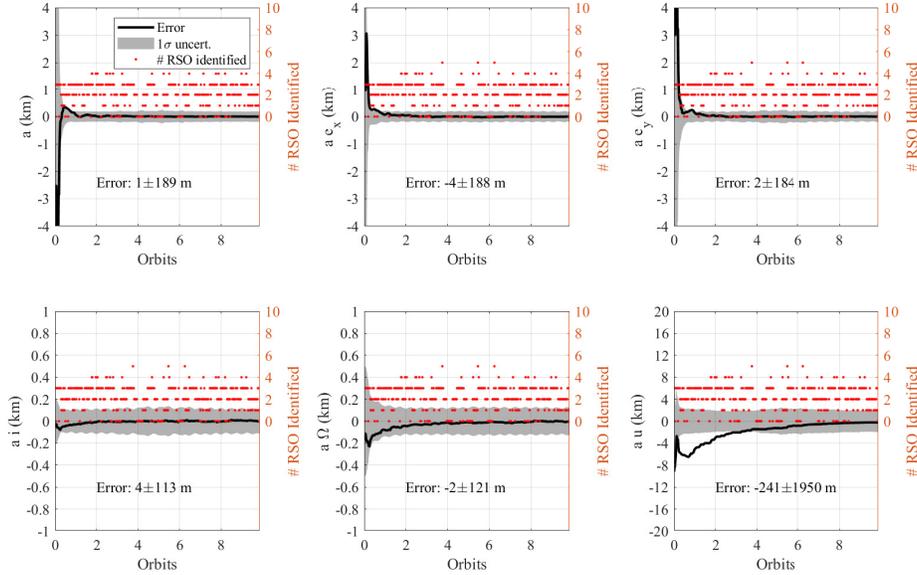


Figure 12: Orbit estimation errors and formal 1σ covariances without catalog TLE errors.

StarFOX provides stronger estimation performance for a, u . The catalog information also facilitates much faster convergence times; StarFOX had not yet achieved convergence after 36 hours (the epoch of the quoted values). However, the lack of steady-state biases for StarFOX demonstrates the benefits of estimating the states of targets as well as the observer.

Table 6: Orbit estimation errors and 1σ uncertainties comparison for StarFOX and FALCON.

Case	a (m)	ae_x (m)	ae_y (m)	ai (m)	$a\Omega$ (m)	au (m)
StarFOX	-3 ± 52	54 ± 502	575 ± 1434	69 ± 1647	1056 ± 1700	-260 ± 547
FALCON	-14 ± 189	190 ± 188	617 ± 184	-1 ± 113	92 ± 121	4726 ± 1915

CONCLUSION

This paper proposes a novel algorithmic pipeline to perform lost-in-space absolute orbit determination for an observer spacecraft, using only bearing angle measurements of resident space objects (RSO) imaged by an onboard vision-based sensor. The new architecture is known as Fast Autonomous Lost-in-space Catalog-based Optical Navigation (FALCON). Tracked RSO are matched to existing identities in a space object catalog and subsequently act as purely optical “beacons” providing known geometric information to the observer. The on-board data fusion of a-priori orbit information of the beacons (from the space catalog) in conjunction with the newly acquired optical bearing angles to the beacons (from the observer camera) provides a means to independently determine the orbit of the satellite observer. RSO measurements may also be fed back into the catalog for improved space situational awareness outcomes.

Key components of FALCON introduced in this paper are the lost-in-space catalog matching (CAM) algorithm, batch orbit determination (BOD) system, and sequential orbit determination

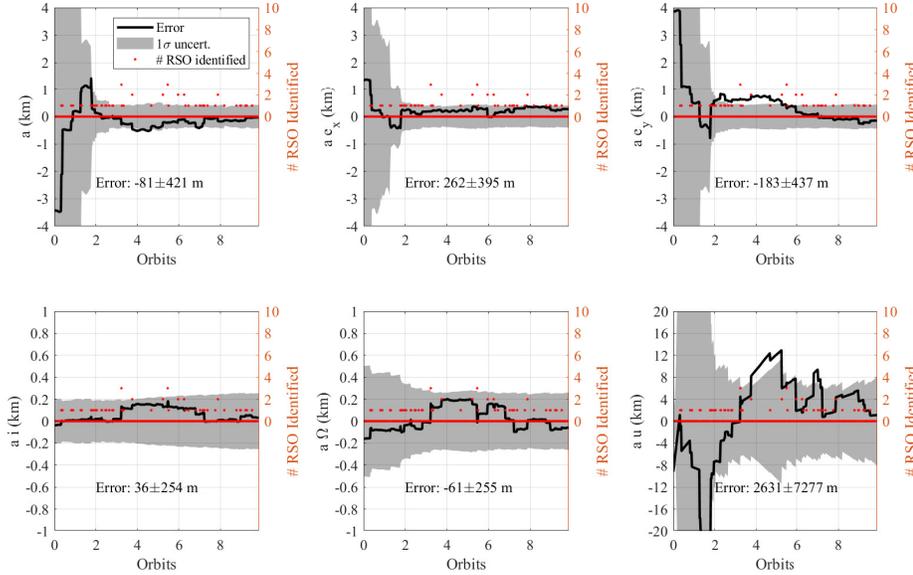


Figure 13: Orbit estimation errors and formal 1σ covariances using non-Starling RSO only.

(SOD) system. CAM determines the optimal match between observed unit vectors and the space object catalog, by computing and sorting the distances between catalog objects and the plane formed by pairs of line-of-sight measurements. The observer’s position may then be estimated with minimal computational cost, and without a-priori observer orbit information. BOD utilizes multiple catalog matches across multiple epochs to produce an initial observer orbit estimate. SOD leverages an unscented Kalman filter, in combination with catalog uncertainty models, to robustly refine the observer’s absolute orbit using only inter-object bearing angles.

A simulation scenario leveraging flight data from the NASA Starling mission is used to validate the robustness of the methodology. The scenario incorporates high-fidelity modeling of uncertainties in dynamics (via different dynamics models in the reference truth and onboard algorithms); catalog orbit knowledge (via perturbation of the catalog between the reference truth and onboard algorithms); and sensing (via physically accurate simulation of synthetic images in the loop).

Simulation results display promising performance. CAM is able to find a unique correspondence between visible objects and catalog identities after only three consecutive detections of the target spacecraft, under lost-in-space conditions. Subsequently, CAM displays a successful identification rate of 88% while remaining within reasonable computational bounds. BOD is able to initialize a navigation filter using only three sequential images. SOD demonstrates effective refinement of the observer orbit from a coarse initial state, achieving convergence in a single orbit. Steady-state position errors of 200 m (radial and normal directions) and 2 km (tangential direction) are achieved with image processing in the loop under challenging measurement conditions. The FALCON prototype therefore demonstrates strong new strategic capabilities for resilient optical navigation.

Results also reveal limitations of the current implementation. For the CAM algorithm, these include its computational scalability to larger catalog sizes and difficulties in handling measurement ambiguities when inter-object geometry is poor. Future work will explore usage of kd-trees, hash-

ing, or data-driven priors to improve catalog search speed, as well as measurement pre-processing to eliminate ambiguous cases. For the SOD algorithm, the state estimate is affected by steady-state biases when catalog orbit errors produce correlated measurement errors, and time is not estimated on board. Future work will investigate simultaneous refinement of RSO orbit estimates to reduce catalog errors over time, and simultaneous positioning and time synchronization within the filter.

ACKNOWLEDGMENTS

This work was supported by the Air Force Office of Scientific Research and Dr. Andrew Sinclair (award number FA9550-24-1-0314) for the project titled ‘Resilient Positioning, Navigation and Timing using Angles-Only Measurements to Cataloged Resident Space Objects’.

REFERENCES

- [1] R. Brukardt, J. Klempner, D. Pachtod, and B. Stokes, “The role of space in driving sustainability, security, and development on Earth,” tech. rep., McKinsey & Company, 2022.
- [2] F. Baccelli, S. Candel, G. Perrin, and J.-L. Puget, “Large Satellite Constellations: Challenges and Impact,” tech. rep., Académie des sciences, 2024.
- [3] M. J. Holzinger and M. K. Jah, “Challenges and Potential in Space Domain Awareness,” *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 1, 2018, pp. 15–18, 10.2514/1.G003483.
- [4] T. Flohrer, H. Krag, H. Klinkrad, B. B. Virgili, and C. Fruh, “Improving ESA’s collision risk estimates by an assessment of the TLE orbit errors of the US SSN catalogue,” *European Conference on Space Debris*, Darmstadt, Germany, 2009.
- [5] S. Miller, C. Adams, N. Alem, H. Cannon, *et al.*, “Starling CubeSat Swarm Technology Demonstration Flight Results,” *Proceedings of the 38th Small Satellite Conference*, Logan, Utah, 2024.
- [6] S. Palo, G. Stafford, and A. Hoskins, “An Agile Multi-Use Nano Star Camera for Constellation Applications,” *Proceedings of the Small Satellite Conference*, Logan, UT, 2013.
- [7] J. Kruger and S. D’Amico, “On-Orbit Performance and Lessons Learned for Autonomous Angles-Only Navigation of a Satellite Swarm,” *IEEE Aerospace Conference*, Big Sky, Montana, 2025.
- [8] H. Bizalio, A. Guillot, A. Petit, and R. Lucken, “Systematic TLE data improvement by neural network for most cataloged resident space objects,” *Advances in Space Research*, Vol. 72, No. 7, 2023, pp. 2649–2659.
- [9] D. Mortari, M. A. Samaan, C. Bruccoleri, and J. L. Junkins, “The Pyramid Star Identification Technique,” *Navigation*, Vol. 51, No. 3, 2004, pp. 171–183, 10.1002/j.2161-4296.2004.tb00349.x.
- [10] J. Brown, K. Stubis, and K. Cahoy, “TETRA: Star Identification with Hash Tables,” *Small Satellites Conference*, 2017.
- [11] B. A. Jones, D. S. Bryant, B.-T. Vo, and B.-N. Vo, “Challenges of multi-target tracking for space situational awareness,” *International Conference on Information Fusion*, Washington, DC, 2015.
- [12] S. Yu, R. Sun, and C. Zhang, “A method for catalogue correlation of space objects,” *NEO and Debris Detection Conference*, Darmstadt, Germany, 2019.
- [13] A. Siminski and T. Flohrer, “Comparison-space selection to achieve efficient tracklet-to-object association,” *Advances in Space Research*, Vol. 64, No. 7, 2019, pp. 1423–1431.
- [14] L. Penafiel, W. Dupree, and T. Gemmer, “Clustering-based uncorrelated track association,” *Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, Hawaii, 2021.
- [15] S. D’Amico, J.-S. Ardaens, G. Gaias, and others., “Noncooperative Rendezvous Using Angles-Only Optical Navigation: System Design and Flight Results,” *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 6, 2013, pp. 1576–1595, 10.2514/1.59236.
- [16] G. Gaias, S. D’Amico, and J.-S. Ardaens, “Angles-Only Navigation to a Noncooperative Satellite Using Relative Orbital Elements,” *Journal of Guidance, Control, and Dynamics*, Vol. 37, No. 2, 2014, pp. 439–451, 10.2514/1.61494.
- [17] G. Gaias and J.-S. Ardaens, “Flight Demonstration of Autonomous Noncooperative Rendezvous in Low Earth Orbit,” *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 6, 2017, pp. 1337–1354.
- [18] J.-S. Ardaens and G. Gaias, “Angles-Only Relative Orbit Determination in Low Earth Orbit,” *Advances in Space Research*, Vol. 61, No. 11, 2018, pp. 2740–2760, 10.1016/j.asr.2018.03.016.
- [19] L. Kazemi, A. Vandenberghe, M. Samson, B. Vandoren, and T. Delabie, “Debris Detection Using Star Tracker Concept Verification,” *Small Satellites Conference*, Logan, Utah, 2024.

- [20] A. W. Koenig and S. D’Amico, “Observability-Aware Numerical Algorithm for Angles-Only Initial Relative Orbit Determination,” *2020 AAS/AIAA Astrodynamics Specialist Conference*, Lake Tahoe, CA, 2020.
- [21] M. Willis and S. D’Amico, “Fast Angles-Only Relative Navigation Using Polynomial Dynamics,” *11th International Workshop on Satellite Constellations and Formation Flight*, Milano, Italy, 2022.
- [22] J. Kulik and D. Savransky, “State Transition Tensors for Passive Angles-Only Relative Orbit Determination,” *33rd AAS/AIAA Space Flight Mechanics Meeting*, Austin, TX, 2023.
- [23] A. J. Sinclair and T. Alan Lovell, “Optimal Linear Orbit Determination,” *Journal of Guidance, Control, and Dynamics*, Vol. 43, No. 3, 2020, pp. 628–632, 10.2514/1.G004182.
- [24] J. Sullivan, A. W. Koenig, J. Kruger, and S. D’Amico, “Generalized Angles-Only Navigation Architecture for Autonomous Distributed Space Systems,” *Journal of Guidance, Control, and Dynamics*, Vol. 44, No. 6, 2021, pp. 1087–1105, 10.2514/1.G005439.
- [25] J. Kruger, S. Hwang, and S. D’Amico, “Starling Formation-Flying Optical Experiment: Initial Operations and Flight Results,” *Proceedings of the 38th Small Satellite Conference*, Logan, Utah, 2024.
- [26] S. Henry and J. A. Christian, “Absolute Triangulation Algorithms for Space Exploration,” *Journal of Guidance, Control, and Dynamics*, Vol. 46, No. 1, 2023, pp. 21–46, 10.2514/1.G006989.
- [27] J. Kruger and S. D’Amico, “Observability Analysis and Optimization for Angles-Only Navigation of Distributed Space Systems,” *11th International Workshop on Satellite Constellations and Formation Flying*, Milano, Italy, 2022.
- [28] D. A. Vallado and W. D. McClain, *Fundamentals of Astrodynamics and Applications*. Hawthorne, California: Microcosm Press, 4 ed., 2013.
- [29] D. Vallado and P. Crawford, “SGP4 Orbit Determination,” *AIAA/AAS Astrodynamics Specialist Conference*, 2008, 10.2514/6.2008-6770.
- [30] J. Shell, “Optimizing orbital debris monitoring with optical telescopes,” *Advanced Maui Optical and Space Surveillance Technologies Conference* (S. Ryan, ed.), 2010.
- [31] E. Andreis, P. Panicucci, and F. Topputo, “Autonomous Vision-Based Algorithm for Interplanetary Navigation,” *Journal of Guidance, Control, and Dynamics*, Vol. 47, No. 9, 2024, p. 1792–1807, 10.2514/1.g007926.
- [32] M. Perryman, L. Lindegren, J. Kovalesky, *et al.*, “The HIPPARCOS catalogue,” *Astronomy and Astrophysics*, Vol. 323, 1997, pp. L49–L52.
- [33] M. D. Hejduk, “Specular and Diffuse Components in Spherical Satellite Photometric Modeling,” *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, 2011.
- [34] S. Palo, G. Stafford, and A. Hoskins, “An agile multi-use nano star camera for constellation applications,” 2013.
- [35] V. Akondi, M. B. Roopashree, and R. P. Budihala, “Improved iteratively weighted centroiding for accurate spot detection in laser guide star based Shack Hartmann sensor,” *Atmospheric and Oceanic Propagation of Electromagnetic Waves IV* (O. Korotkova, ed.), Vol. 7588, International Society for Optics and Photonics, SPIE, 2010, p. 758806, 10.1117/12.841331.
- [36] J. R. Wertz, *Spacecraft Attitude Determination and Control*. Springer Science & Business Media, 2012, 10.1007/978-94-009-9907-7.
- [37] R. Hartley and A. Zisserman, *Multiple View Geometry*. Cambridge, United Kingdom: Cambridge University Press, 2003.
- [38] J. Kruger, A. W. Koenig, and S. D’Amico, “The Starling Formation-Flying Optical Experiment (StarFOX): System Design, Experiment Design and Pre-Flight Verification,” *Journal of Spacecraft and Rockets*, 2023, pp. 1–23, 10.2514/1.A35598.
- [39] K. T. Alfriend, ed., *Spacecraft Formation Flying: Dynamics, Control, and Navigation*. Butterworth-Heinemann/Elsevier, 2010, 10.1016/C2009-0-17485-8.
- [40] O. Montenbruck and E. Gill, *Satellite Orbits: Models, Methods and Applications*. Springer, 2012, 10.1115/1.1451162.
- [41] B. D. Tapley, S. Bettadpur, M. Watkins, and C. Reigber, “The Gravity Recovery and Climate Experiment: Mission Overview and Early Results,” *Geophysical Research Letters*, Vol. 31, No. 9, 2004, 10.1029/2004GL019920.
- [42] S. D’Amico, *Autonomous Formation Flying in Low Earth Orbit*. PhD thesis, Delft University, 2010.
- [43] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
- [44] Space-Track. <https://space-track.org/>.
- [45] T. Bell and S. D’Amico, “Event-Driven Simulation for Rapid Iterative Development of Distributed Space Flight Software,” 2025.

- [46] J. Picone, A. Hedin, D. P. Drob, and A. Aikin, "NRLMSISE-00 Empirical Model of the Atmosphere: Statistical Comparisons and Scientific Issues," *Journal of Geophysical Research: Space Physics*, Vol. 107, No. A12, 2002, pp. SIA-15, 10.1029/2002JA009430.
- [47] E. Andreis, V. Franzese, and F. Topputo, "Onboard Orbit Determination for Deep-Space Cube-Sats," *Journal of Guidance, Control, and Dynamics*, Vol. 45, No. 8, 2022, pp. 1466–1480, 10.2514/1.G006294.