

**Target Detection and Orbit Determination for Mars Sample Return**  
Zubin P. Olikara<sup>(1)</sup>, Tomas J. Martin-Mur<sup>(1)</sup>, Declan M. Mages<sup>(1)</sup>,  
Joseph E. Riedel<sup>(1)</sup>, Margaret M. Rybak<sup>(1)</sup>, Matthew J. Abrahamson<sup>(1)</sup>

<sup>(1)</sup>*Jet Propulsion Laboratory, California Institute of Technology  
Pasadena, California, United States of America  
Email: {zubin.p.olikara, tomas.j.martinmur, declan.m.mages,  
joseph.e.riedel, margaret.m.rybak, matthew.abrahamson}@jpl.nasa.gov*

**Abstract** – The Mars Sample Return campaign is a joint effort being planned by the National Aeronautics and Space Administration (NASA) and the European Space Agency (ESA) to bring back a curated set of samples from the Jezero crater region of Mars to Earth. A fundamental part of this campaign is the Launch, Rendezvous, and Capture phase at Mars, which would begin with the launch of an Orbiting Sample (OS) container aboard NASA’s Mars Ascent Vehicle (MAV). Once in orbit, the MAV upper stage would release the OS. ESA’s Earth Return Orbiter must then determine the state of the OS so that it could maneuver to its vicinity for capture. This paper provides an overview of ongoing analyses performed at NASA’s Jet Propulsion Laboratory to support the development of the baseline MAV and OS detection and orbit determination strategies, which leverage both optical and radio frequency techniques.

## I. INTRODUCTION

Returning samples from Mars has been one of the highest scientific priorities within the planetary science and astrobiology communities for decades, and the focus on this goal continues to this day [1]. The return of samples to the Earth would involve the successful coordination between various elements; an overview of a joint NASA/ESA baseline architecture for the Mars Sample Return (MSR) campaign is available in a companion paper [2]. While this architecture continues to evolve as its design matures, the MSR campaign is already underway with NASA’s Perseverance rover collecting samples on the Martian surface.

Upon the arrival of NASA’s Sample Retrieval Lander at Mars, up to 30 sample tubes would be loaded into the basketball-sized Orbiting Sample (OS) container, which would be mounted aboard the Mars Ascent Vehicle (MAV). Once the MAV launches from the surface of Mars and the OS is delivered into orbit, the Earth Return Orbiter (ERO) would begin the process of finding them with ground in the loop. Development of the strategy for detection and orbit determination has involved close collaboration between NASA, ESA, and ERO’s prime contractor Airbus. This paper provides a look at work undertaken at NASA’s Jet Propulsion Laboratory to support this effort. Successful orbit determination would enable ERO to rendezvous with and capture the OS

before beginning its return journey back to Earth.

The launch of the MAV would provide the initial interface for the detection activities. It is envisioned that the MAV would consist of a guided solid rocket first stage that raises apoapsis above 300-km altitude after which an unguided solid rocket second stage would bring itself and the mounted OS to orbital velocity with periapsis above 300-km altitude. Once in orbit, the MAV would release the OS in such a way so as to minimize the chance of recontact as well as to ensure that the OS would be sufficiently far away from the MAV for safe approach and capture by ERO. Per the current architecture [2], the OS orbit would be subject to the following limits:

1. Periapsis above 300-km altitude to allow for safe ERO capture.
2. Semi-major axis above 330-km altitude to ensure an orbital lifetime of at least 10 years.
3. Semi-major axis below 500-km altitude to reduce size of orbit dispersion and ensure reliable imaging opportunities by ERO.
4. Inclination close to the launch site latitude as well as ERO’s inclination to maximize launch capability and reduce ERO’s inclination change costs for rendezvous.

To facilitate detection and orbit determination analyses, example dispersion scenarios have been constructed that are consistent with these limits. Nominal orbital elements and 3-sigma uncertainties for the MAV are shown in Table 1 in an IAU Mars Pole inertial frame. For reference, the semi-major axis (SMA) altitudes are defined relative to Mars’s equatorial radius of 3396.2 km. Note that since the actual distributions used are not perfectly normal, the reported 3-sigma SMA altitudes in the table are slightly above the 500-km bound. Injection dispersions 14A and 14D have a tighter range of semi-major axis compared to 15A and 15D, which are more stressing. Since the inclination of approximately 22.5° is above the latitude of Jezero crater (18.4° N), there are both ascending (denoted by “A”) and descending (denoted by “D”) injection options.

A combination of radio frequency (RF) and optical approaches are considered to ensure a robust search. It is envisioned that ERO along with any other available assets equipped with ultra-high frequency (UHF) radios would receive telemetry during the MAV’s first stage

Table 1. Example MAV orbit nominal  $\pm$  3-sigma uncertainty injection scenarios.

Injection dispersion	SMA alt. [km]	$e \cos(\omega)$	$e \sin(\omega)$	Inclination [deg.]	RAAN [deg.]	Arg. latitude [deg.]
14A	455.8 $\pm$ 47.7	0.0055 $\pm$ 0.0203	0.0015 $\pm$ 0.0120	22.50 $\pm$ 0.46	24.61 $\pm$ 2.80	73.17 $\pm$ 3.67
14D	456.7 $\pm$ 47.2	-0.0125 $\pm$ 0.0146	0.0109 $\pm$ 0.0183	22.52 $\pm$ 0.83	-48.40 $\pm$ 2.14	142.01 $\pm$ 3.22
15A	420.6 $\pm$ 90.7	0.0076 $\pm$ 0.0217	0.0179 $\pm$ 0.0215	22.49 $\pm$ 0.47	24.44 $\pm$ 2.70	76.56 $\pm$ 2.78
15D	421.3 $\pm$ 90.7	-0.0142 $\pm$ 0.0196	0.0138 $\pm$ 0.0233	22.52 $\pm$ 0.79	-48.43 $\pm$ 2.13	145.21 $\pm$ 2.31

ascent followed by a one-way beacon signal beginning at the second stage separation. This MAV beacon would be run in a continuous mode through the release of the OS after which it would switch to an intermittent mode providing Doppler beeps over the days following launch which could be used for determining the upper stage's orbit. Since the OS is completely passive, direct determination of its orbit would be done using a narrow angle camera (NAC) on ERO. However, knowledge of the MAV's orbit would also provide information on the OS orbit up to the uncertainty in the MAV-OS separation and any unmodeled tail-off impulse from the MAV. This separation uncertainty is expected to be much smaller than the uncertainty associated with the MAV orbital injection. This MAV orbit estimate could be obtained using optical detections of the MAV as well, which would likely be easier to detect given its larger size compared to the OS. Once the MAV orbit is sufficiently well known, ERO could then perform a dedicated optical search for the smaller OS potentially from a more favorable relative orbit to which it would maneuver. Given the importance of optical imaging, the reflective properties of the MAV and OS surfaces are fundamental to ensuring detectability by ERO.

Detectability would also be a function of the relative orbits between ERO and the targets as well as their attitudes. The current analysis assumes that ERO would maneuver itself into a near circular 305-km altitude orbit before the launch of the MAV. This orbit is selected to have a repeat ground track providing a consistent offset relative to the launch site for launch opportunities every three sols. Returning to Table 1, ERO at 305-km altitude would correspond to the dispersed cloud of injected MAV and OS states being above ERO. Having ERO in an orbit below the target could be advantageous for mitigating straylight since targets would pass higher above the Mars limb. At injection, it is also assumed that the ERO would be offset by roughly  $15^\circ$  in its right ascension of ascending node (RAAN) from the injection state. This offset allows the RAAN difference between ERO and the OS to naturally drift closer to alignment over the course of the search and rendezvous due to their relative precession rates.

The synodic period between ERO and a dispersed MAV (or OS) sample drives the visibility periods during which imaging would be possible. Since ERO has a shorter orbital period than the target sample, a visibility

period starts with the target ahead of ERO after which ERO passes below the target and the target falls behind ERO before being obstructed by Mars's surface. For ERO to provide support over the full MAV launch sequence, this means that the MAV would be behind ERO (generally by approximately  $20^\circ$  in true longitude) by the time of injection. The highest altitude samples have the shortest synodic periods, which corresponds to shorter but more frequent visibility periods. The lowest altitude samples (i.e., those with a semimajor axis most similar to ERO) have the longest synodic periods, which corresponds to longer visibility periods but potentially days or weeks during which the target is occulted behind Mars.

Superimposed upon these visibility periods is the target illumination conditions that vary over each orbit revolution, which takes about two hours, as a function of solar phase angle as viewed from ERO. Most notably, for up to about 40% of each revolution, the target is in eclipse and cannot be observed optically. A successful search strategy must take all these conditions into account to ensure a robust orbit determination not just for the nominal target states but across the full MAV and OS dispersions. Since it is assumed that the NAC and the UHF antenna share a boresight direction, the optical and RF searches are necessarily coupled. However, since the NAC would have a significantly narrower field of detection, its pointing needs would generally be driving.

In the following sections, the optical detection modeling environment will first be discussed. Then potential search strategies for optically locating the MAV will be presented. This search design informs the subsequent MAV orbit determination results using optical-only and RF-only detections as well as a combination of the two. A scheme for optically searching for the OS will then be provided and used for analyzing OS orbit determination performance.

## II. OPTICAL DETECTIONS

The target viewing geometry, camera characteristics, and image processing are all important components for obtaining an optical measurement that can be fed into an orbit determination algorithm. The current analysis assumes a circular  $4.5^\circ$  field of view (FOV) for ERO's NAC with a 1020 x 1020 detector matrix [3]. Whether a

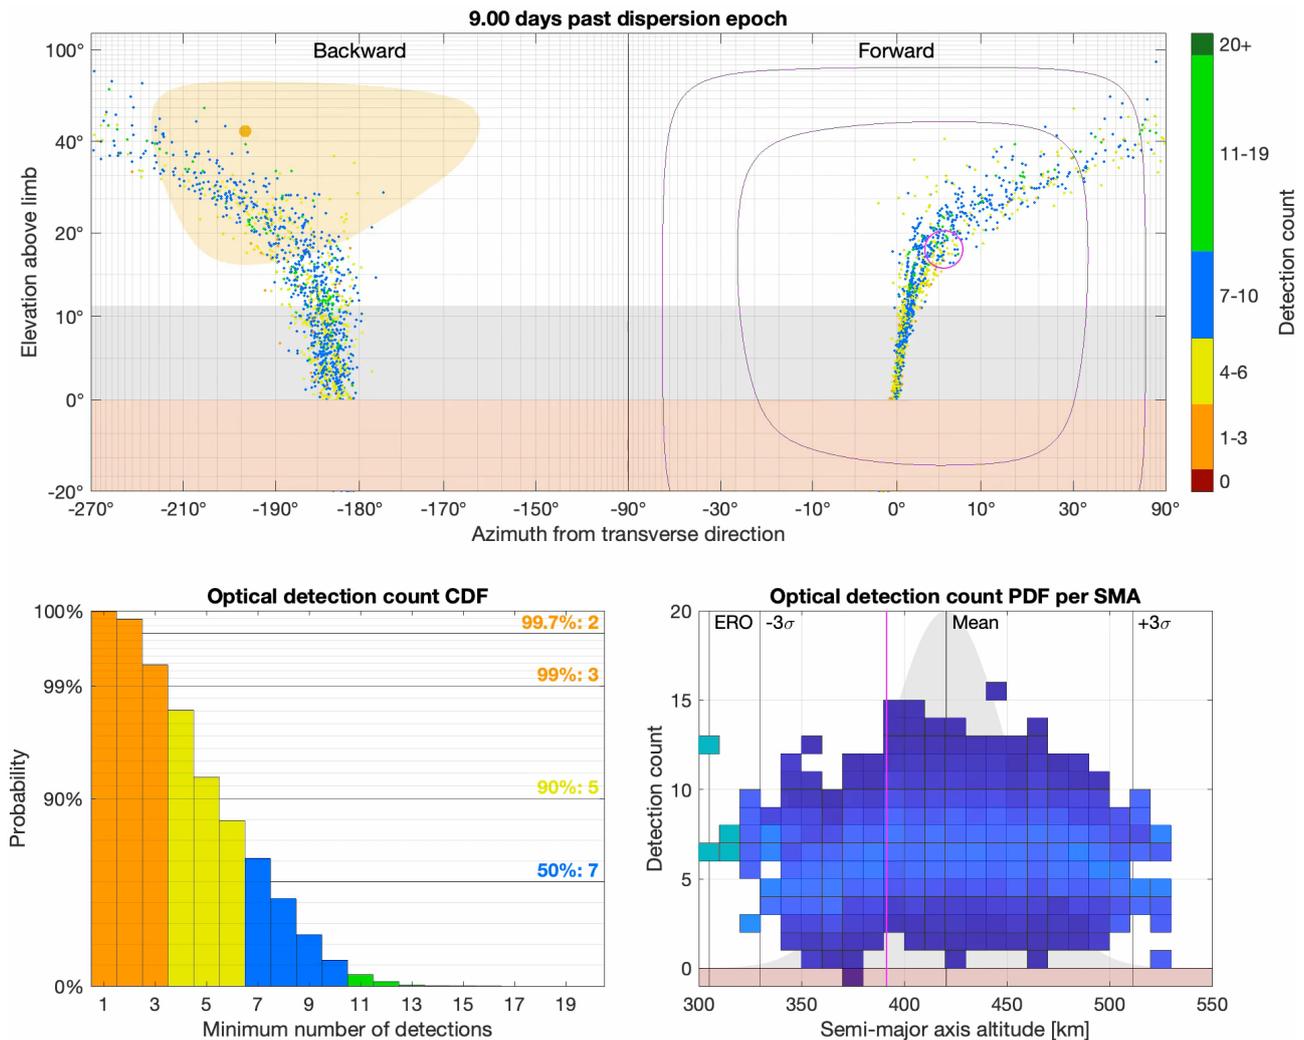


Fig. 1. Dashboard for tracking optical search progress during Monte Carlo simulation.

target in the field of view could be detected is a function of its signal-to-noise ratio (SNR) and the image processing algorithm, which effectively sets a minimum SNR for detectability in the brightest pixel as a function of the number of pixels the target streaks across. To concentrate as much energy as possible from the target into a given pixel (i.e., to increase the signal), it is assumed that the NAC would track the motion of a fictitious candidate target over the course of a 2-second-long exposure. Since the actual motion of the target would come from an unknown dispersed state associated with the MAV injection or OS separation, the target could streak across multiple pixels but much fewer than would be streaked across if the camera was pointed in a fixed inertial direction. To account for the various sources of uncertainty, a Monte Carlo approach is used.

Constructing a high-level simulation environment requires integrating component models that range all the way from orbital scales (positions of the ERO, MAV, OS, and Sun) down to proton and electron counts at the

detector. This integration necessarily requires simplification but ideally in a manner that still captures the expected performance at Mars. At the trajectory-level, Monte Carlo dispersed states for the MAV (e.g., consistent with values in Table 1) and the OS (based on MAV OD and release knowledge errors) are ingested and propagated forward in time. The trajectory for the ERO is also required for computing the viewing geometry, and the solar ephemeris is needed for determining illumination conditions including eclipses. Optical properties, specifically the effective reflective area, of the MAV and OS are stored in tables as a function of solar phase angle as well as in terms of favorability of the target's attitude (e.g., looking at the long side of the MAV as opposed to down its nozzle would reflect more light towards ERO).

Modeling the camera performance depends on a diverse set of parameters. Rather than directly generating simulated images during a run, an SNR for the target can instead be computed that is then compared against

detectability thresholds. Background on the SNR computation can be found in [4]. Incorporating noise from straylight is an important component of the SNR assessment. It can enter the camera directly from the Sun or be reflected by the surface or atmosphere of Mars [3]. To help mitigate its impact, a Sun exclusion angle of  $40^\circ$  and a Mars limb exclusion angle of  $13.5^\circ$  from the camera boresight is baselined in the current analysis.

Analyzing detectability also depends on the target imaging plan, and search profiles for the MAV and OS are presented in the subsequent sections. Detectability and false-positive rejection can greatly benefit by taking and analyzing a set of images as a burst rather than as standalone pictures. The current work nominally assumes that a burst of 20 pictures would be taken in succession with bursts occurring every 10 minutes. Each image in the burst can be pre-processed before being checked for potential detections. Cleaning the images during the pre-processing step involves flattening the image by estimating and removing the background as well as normalizing the detector sensitivity variations across the field of view. It can also involve removing stars and cosmic rays. A variety of techniques are then available for identifying whether a target exists in an image burst; two possibilities considered are human-in-the-loop and synthetic tracking.

A human-in-the-loop approach could leverage a “blink comparator” strategy where the sequence of pictures in the burst is cycled through in succession and leveraging the human eye’s ability to detect subtle differences and patterns of motion. This strategy has been a mainstay for astronomers for over 100 years and was the means by which Clyde Tombaugh at Lowell Observatory discovered Pluto in 1930. Experience suggests that detections can reliably be made down to an SNR of around 3.

Synthetic tracking automates the identification by testing various velocity hypotheses of the target’s motion across the burst of images. By shifting the images by a given motion hypothesis and then adding them together, an otherwise faint set of target observations can be made bright. This type of approach could potentially allow detections even when the per image SNR is around 1. Synthetic tracking has been successfully used for detecting asteroids and Earth-orbiting objects [5].

The complete Monte Carlo simulation environment can be run either for a single scenario or in a batch mode sweeping various parameters such as the ERO orbit offsets, camera properties, and image processing settings. When a single scenario is being investigated, a dashboard is often used to illustrate the search progress as shown in Fig. 1. The top view in the figure represent a “fish eye”-like perspective from ERO looking both

ahead and behind. The points shown represent Monte Carlo dispersed states for the MAV in this case. At the example time shown, only part of the dispersion is within line of sight of ERO with the remainder occulted behind Mars. The magenta circle represents the NAC FOV in which detections can be made if the SNR is sufficiently high. The NAC must remain above the Mars limb exclusion shown in gray and outside the Sun exclusion shown in yellow. When a successful detection is made, the color of the corresponding Monte Carlo sample is updated to reflect how many times it has been detected. This coloring allows regions of the dispersion that have been detected fewer times to be more readily apparent. The lower left plot provides counts on number of detections across the sampled distribution, and the lower right plot shows number of detections (more samples with a given number of detections corresponds to brighter colors) in different SMA bins. This latter plot helps verify that the search strategy, which will be discussed in the next section, is covering the full dispersion relatively uniformly.

### III. MAV SEARCH STRATEGY

Given its superior optical properties, it would generally be easier to detect the MAV as opposed to the OS. Therefore, an initial broad search for the MAV is conducted, after which a satisfactory MAV orbit determination would have removed most of the uncertainty associated with the OS injection. While this lower uncertainty after MAV detection enables a much more focused OS search, the MAV search must be robust enough to detect the MAV across the full range of possible injected states (such as the cases listed in Table 1). Furthermore, multiple detections are required for a successful orbit determination. Since the set of dispersed states would expand beyond a single FOV as viewed from ERO soon after injection, a search pattern must be designed. Given the highly nonlinear evolution of the full dispersion over the days and weeks after MAV launch, a Monte Carlo approach is used to evaluate the search performance.

While all six orbital elements are dispersed using a full  $6 \times 6$  covariance as part of the Monte Carlo simulation, the search strategy is designed around a reduced 1-dimensional parameterization. Specifically, the dispersion spreads out most significantly along track due to variation in semi-major axis. The dispersion is parameterized from  $-3.5$  to  $+3.5$  sigma in semi-major axis, and at each value, the most likely values of the other orbital elements are computed by reducing the MAV delivery covariance matrix. This creates a segment of candidate target locations that initially fits within a single FOV as viewed from ERO but then evolves over time to wrap multiple times around Mars. When designing an open-loop search strategy, the goal is not to maximize number of detections of the nominal

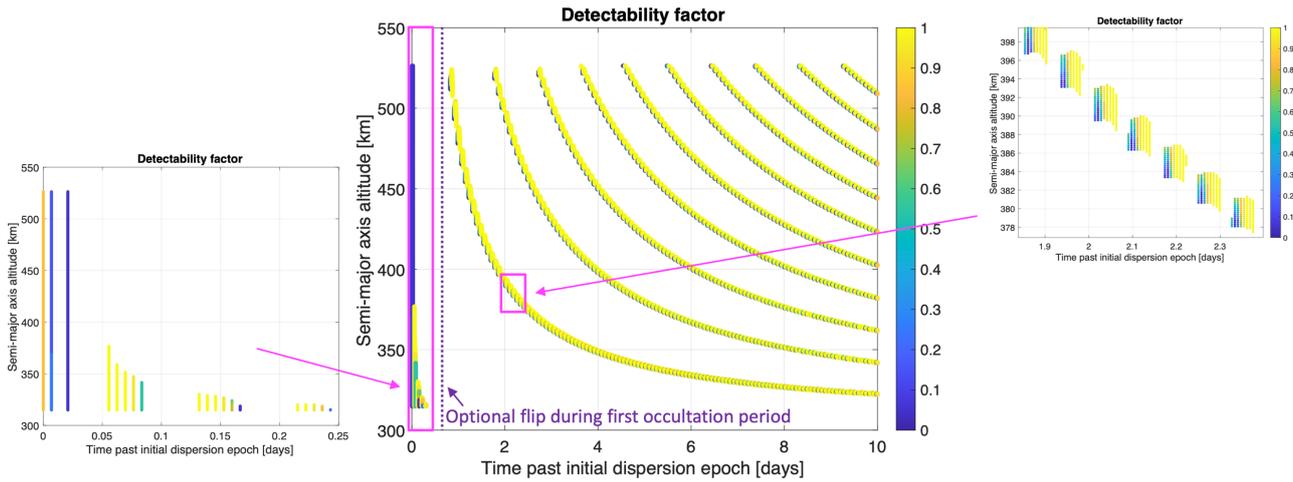


Fig. 2. Example MAV detectability likelihood as a function of semi-major axis and time.

target location but rather to maximize the minimum number of detections expected across the parameterized target set. This objective ensures that the full range of possible semi-major axes are covered by the search. The actual detectability is computed for the fully dispersed Monte Carlo states, but variations in the other five elements appear to remain within a single FOV when the camera itself is pointed at a particular semi-major axis parameterized candidate target.

The structure of the camera pointing design space for an example scenario with ERO at 305-km altitude is illustrated in Fig. 2. The dispersion parameterization (semi-major axis) is shown on the y-axis, and the x-axis shows the time past the delivery epoch. Potentially detectable candidate targets are shown as points on the plot with a time step along the x-axis corresponding to cadence of image bursts (assumed to be every 10 minutes). At high altitudes the short but frequent visibility periods are apparent as are the longer but infrequent visibility periods at low altitudes. In addition, regular gaps in time due to eclipsed targets can be seen. Each data point shown meets filters based on the direction the camera is pointed (e.g., only consider targets visible when pointing behind ERO to avoid needing to frequently switch between pointing ahead and behind) as well as pitch limits and yaw limits ( $\pm 30^\circ$  azimuth from ahead/behind direction). Since the attitude of the target is unknown, a probability of detection is computed by sampling the various Effective Reflective Area (ERA) percentiles at the associated solar phase angle. This “detectability factor” is shown by the color of the point in the figure. A detectability factor of 1 (yellow) means that the target is expected to be detectable at any attitude for the image processing algorithm, such as synthetic tracking, that was assumed. A detectability factor near 0 (blue) means that the target is only detectable if the target is in the best possible attitudes, i.e., those with the highest ERA.

A simple search strategy such as fixing elevation of the camera boresight above the Mars limb leads to a “banding” behavior where targets at regular steps in semi-major axis are missed because of the varying illumination conditions over the orbit. These gaps can be filled in by moving the camera boresight target between image bursts. This strategy, which is used in the subsequent MAV orbit determination results, works by stepping through the detectability factor dataset shown in Fig. 2 in time. Using data accumulated over previous time steps, statistics on the current number of detections of each discretized semi-major axis target is available. It is also possible to predict how easy it would be to detect each semi-major axis target in the future. This prediction is generated by looking ahead in the detectability dataset over a time horizon. Combining this information, candidate targets that have the greatest need of imaging at the current time step can be determined while considering potential overlap between targets in the FOV, i.e., a target may be opportunistically detectable if it lies within the FOV when imaging another target. At the current time step, the best candidate target is selected, and the pointing attitude rates are set to match the target’s motion. When checked against the Monte Carlo dispersed sample set, the samples generally have a streak length less than a pixel over a 2 second exposure, and this streaking is accounted for in the computation of SNR, which ultimately determines the MAV’s detectability over a given image burst.

The current strategy is run open loop over the entirety of the MAV search, which introduces conservatism. In practice, once detections of the MAV are made on the ground, an updated search could be uploaded to ERO that is more localized to the orbital regions, particularly in terms of semi-major axis, consistent with the previous detections.

#### IV. MAV ORBIT DETERMINATION

Orbit determination performance is often characterized via covariance and sensitivity analyses of a nominal trajectory. This relies on the assumption that the true in-flight trajectory variance from nominal does not meaningfully alter the information contained in the measurements and their effect on the uncertainty evolution. However, the MAV trajectory dispersions around Mars are large enough that optical and RF measurements can contain meaningfully different state information depending on the dispersion cases. To fully model this, the MAV orbit determination is analyzed here via a Monte Carlo of covariance analyses performed on 2000 dispersed MAV trajectories. Three different variations are considered: orbit determination using (a) only optical, (b) only RF, and (c) both optical and RF measurements.

The objective of the MAV search subphase is to obtain a sufficiently good orbit determination of the MAV such that the ERO could then maneuver to an orbit offset about 30-km below the MAV from which to search for the OS. The uncertainty goals used are provided in Table 2 for all of the orbital parameters mapped to the initial epoch, though semi-major axis and inclination are driving for updating the future search. Therefore, these latter two parameters are the focus of the following analysis.

Table 2. Maximum 3-sigma orbital element uncertainties for MAV orbit determination success.

SMA alt. [km]	0.450
$e \cos(\omega)$	$1.5 \times 10^{-4}$
$e \sin(\omega)$	$1.0 \times 10^{-4}$
Inclination [deg.]	0.140
RAAN [deg.]	0.007
Arg. latitude [deg.]	0.150

##### A. MAV Optical Orbit Determination

First the standalone performance of optical detections for estimating the MAV state is investigated. This provides insight into the power of even a few optical detections in narrowing down the MAV's orbit and also represents a contingency scenario in case of a MAV UHF failure. The input to this process is the set of dispersed MAV states and the detection times, which are determined via modeling of the optical chain for each of the samples.

This produces a statistical representation of possible uncertainty evolution and determines the 99<sup>th</sup>-percentile times at which ERO would have obtained a sufficient number of optical detections of the MAV to proceed to setting itself up for the OS search. The results shown here focus on a stressing dispersion 15D, but include

high level sensitivity results for 14A, 14D, and 15A as well as additional parameter variations. The full 2000 sample MAV uncertainties over time for 15D are plotted in Fig. 3 and are slightly worse but nearly identical to results in the 14A, 14D, and 15A cases. For these cases the 99<sup>th</sup>-percentile of cases reach the threshold for semi-major axis roughly 8 days post release, and the threshold for inclination 6 days post release. The times from release at which 75%, 90%, and 99% of samples reach the threshold for proceeding are listed for these cases along with all sensitivity cases in Table 3.

Table 3. Time past OS release for successful optical-only MAV orbit determination at various percentiles.

Inject. disp.	Parameter variation	75 <sup>th</sup> [days]	90 <sup>th</sup> [days]	99 <sup>th</sup> [days]
14A	Baseline	3.3	4.5	6.8
14D	Baseline	3.8	4.8	7.2
15A	Baseline	3.5	4.5	7.9
15D	Baseline	3.8	5.0	8.3
15D	2-day outage	5.9	7.3	12.6
15D	2.5-min. cadence	1.3	1.5	2.8
15D	Tight OpNav wt.	3.7	4.7	7.9
15D	Loose OpNav wt.	3.9	5.2	8.7

The 2-day outage case represents a scenario where ERO would not take any pictures during the first two days. The 2.5-minute cadence corresponds to taking pictures more frequently but with only 5 pictures in a burst. The tight and loose optical navigation weight cases correspond to 0.1 and 5.0 pixel, respectively, target location sample and line measurement weights compared to the baseline 1.0 pixel weight.

Also important to understand is the relationship between number of detections and the improvement in uncertainty. This relationship is plotted for case 15D in Fig. 4 and shows that with image bursts acquired every 10 minutes, the 90<sup>th</sup> percentile of samples is ready for orbit matching (i.e., begin maneuvering for the OS search after achieving thresholds from Table 2) after only 3 measurements; though to reach 99<sup>th</sup> percentile in inclination it can take up to 7 measurements. While useful for generally characterizing the required number of detections, it is important to note that this behavior is not universal and does not capture the important connection between number of images and observation arc length. It relies heavily on the baseline assumption that image bursts are acquired every 10 minutes when, with a faster a burst cadence such as Case 15D with a 2.5-minute burst cadence in Table 3, multiple detections might be acquired in quick succession and thus not provide as much additional information as detections spaced hours apart.

Optical Only MAV, 2000 Samples

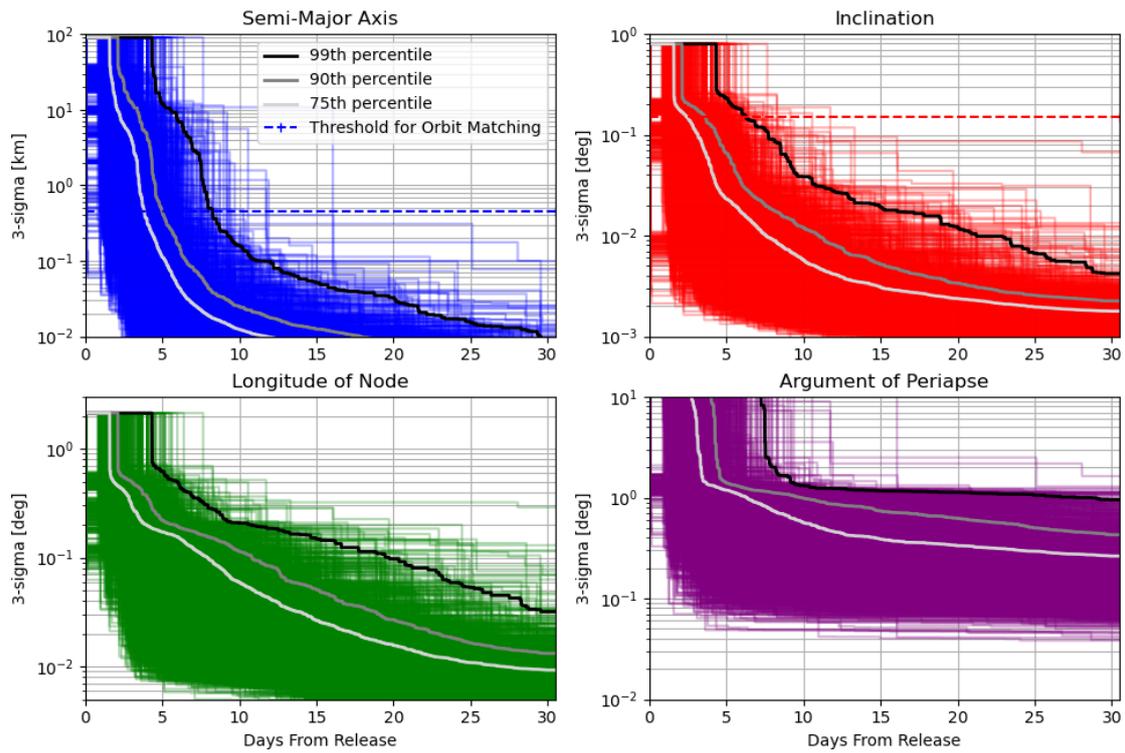


Fig. 3. Optical-only 3-sigma MAV orbital element uncertainties over time for dispersion 15D.

Optical Only MAV, 2000 Samples

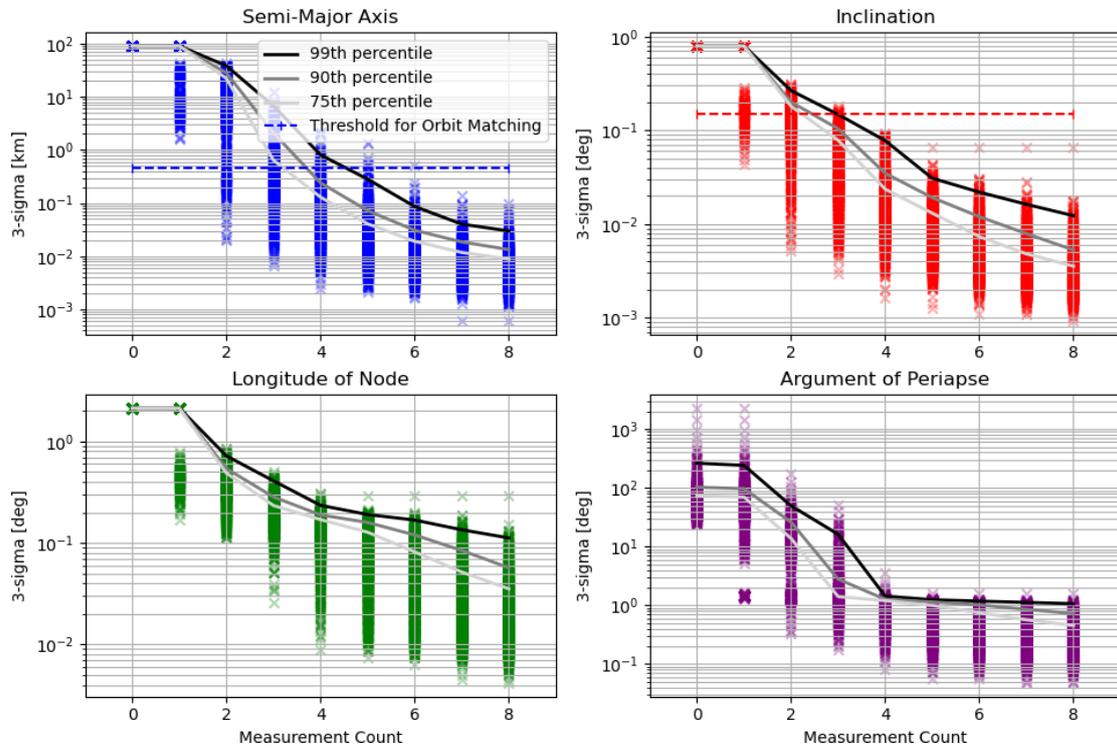


Fig. 4. Optical-only 3-sigma MAV orbital element uncertainties versus number of detections for dispersion 15D.

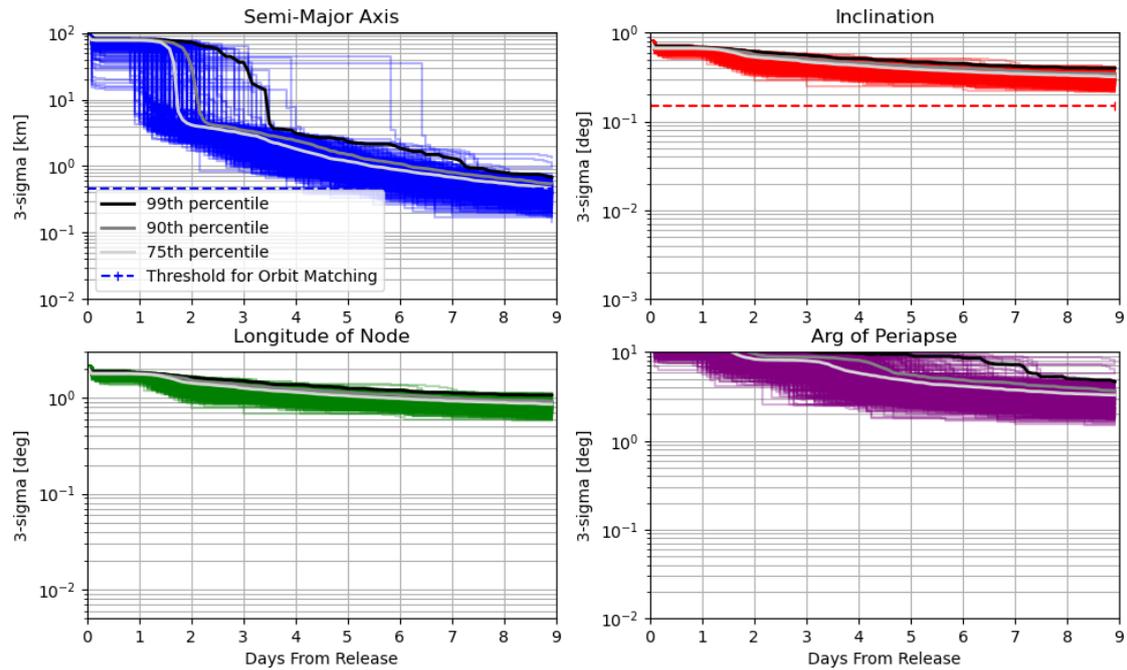


Fig. 5. RF-only 3-sigma MAV orbital element uncertainties over time for dispersion 15D.

### B. MAV RF Orbit Determination

The MAV UHF beacon would emit periodic beeps that would be tracked by ERO. The Doppler measurement error of these beeps would be dominated by the stability of the oscillator used to drive the transmitter in the MAV and the receiver in ERO since the measurement would be of a one-way link. To characterize the performance of MAV OD using RF measurements from ERO a covariance analysis is performed for 500 MAV samples for the 15D injection dispersion case. This analysis assumes a MAV beacon lifetime of up to 9 days. Table 4 lists the oscillator properties assumed for the MAV and ERO. The 500 sample MAV uncertainties over time, assuming a MAV beep every 300 seconds are plotted in Fig. 5. By the end of beacon lifetime, the threshold for successful MAV orbit determination is not met for any orbital elements. Even with a shorter beep period of 150 seconds, there is only a slight improvement; only the 75th percentile of cases reaches the threshold for semi-major axis at the end of the beacon lifetime.

Table 4. Assumed oscillator properties.

Orbiter	Sigma White [km/s]	Sigma Random Walk [km/s]
MAV	$1.85 \times 10^{-4}$	$1.00 \times 10^{-2}$
ERO	$3.70 \times 10^{-5}$	$2.28 \times 10^{-4}$

While the MAV UHF data cannot achieve the required uncertainties for orbit matching in the absence of optical

data, UHF can be a powerful tool to assist the future collection of optical data if early detection fails. After 3 or 4 days of beep data, there is a significant improvement in the long-term predictability of the MAV location; the 3-sigma angular uncertainty of the MAV as viewed from ERO drops to just 2 to 3 NAC FOVs over the following 10 days. Thus, future imaging could be concentrated in a much tighter region of space.

### C. MAV Combined Optical / RF Orbit Determination

Combining the optical detection and RF data, a Monte Carlo covariance analysis of 2000 samples is repeated for the 15D dispersion case. The resulting uncertainty evolution is plotted in Fig. 6. Including both data types predictably produces the best results, with 75<sup>th</sup> / 90<sup>th</sup> / 99<sup>th</sup>-percentile transition to orbit matching times 3.3 / 3.9 / 5.5 days post separation, respectively. As listed in Table 5, compared to the baseline optical-only case this is an approximately 3-day improvement at the 99<sup>th</sup> percentile. Also with both data types, more MAV samples are ready for orbit matching with less optical measurements. With optical only, 4 measurements were required to match the 75<sup>th</sup> and 90<sup>th</sup> percentile sample and by adding RF the 75<sup>th</sup>-percentile samples are ready for orbit matching after just 3 measurements. With RF-only, only the SMA could be determined at the 75<sup>th</sup>-percentile at the desired threshold and is thus marked with an asterisk.

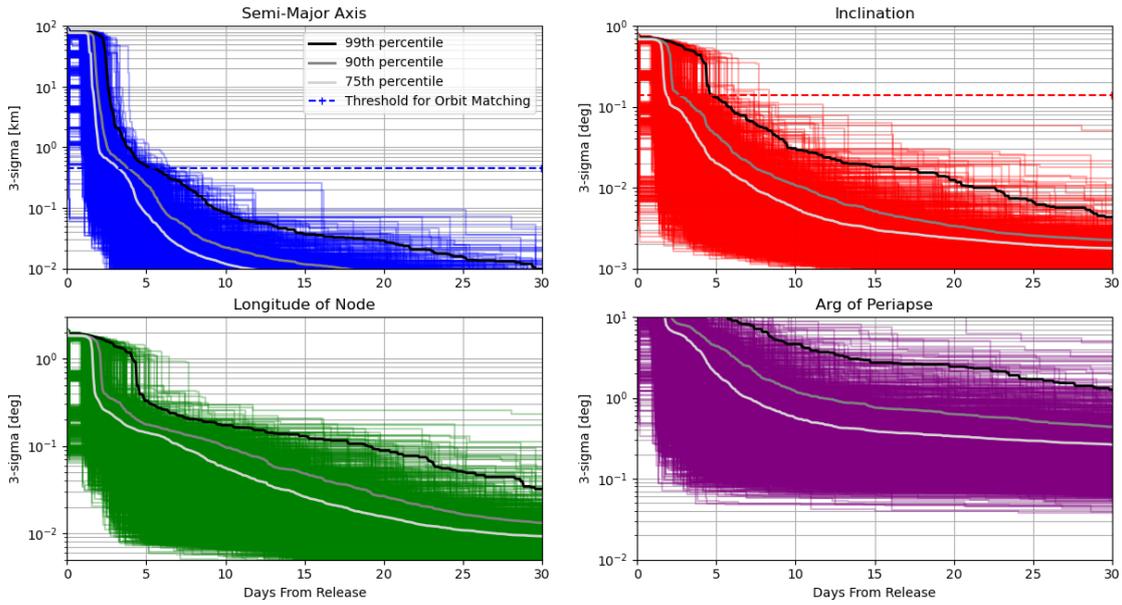


Fig. 6. Combined optical and RF 3-sigma MAV orbital element uncertainties over time for dispersion 15D.

Table 5. Time past OS release for successful MAV orbit determination at various percentiles using optical and/or RF data.

Inject. disp.	Data types	75 <sup>th</sup> [days]	90 <sup>th</sup> [days]	99 <sup>th</sup> [days]
15D	Optical	3.8	5.0	8.3
15D	RF	8.8*	—	—
15D	Optical + RF	3.3	3.9	5.5

## V. OS SEARCH STRATEGY

The baseline rendezvous phase approach considered in this work is to perform a dedicated search for the OS, though it can be analyzed using the same framework as for the MAV. While it is possible that the OS would be imaged as part of the MAV search, the full injection uncertainty combined with the variety of imaging ranges needed to completely cover the dispersion means that the OS may be imaged in conditions that are unfavorable for detection (noting its inferior optical properties as compared to the MAV). However, once the MAV orbit is determined, the OS orbit uncertainty is greatly reduced. It is a function of the MAV orbit uncertainty along with uncertainty associated with the OS separation event and residual tail-off of the MAV. It is expected that the semi-major axis uncertainty before the OS is detected would be on the order of a few kilometers at most.

Given the relatively constrained orbit range for the OS, ERO can maneuver to an orbit favorable for making detections. In the current analysis it is assumed that ERO would maneuver to a circular orbit 30 km below the

nominal OS orbit in semi-major axis. The  $-30$ -km displacement in semi-major axis sets the relative synodic period between ERO and the OS to be approximately one week. It is also assumed that ERO would have matched the nominal OS inclination but that a  $+10^\circ$  offset in RAAN remains at the beginning of the OS search. Since the purpose of the RAAN offset is to have orbital planes ideally naturally drift into approximate alignment at the end of the orbit matching phase, it is expected that there would be a non-zero offset during the OS search.

The expected dispersion would be almost entirely along-track at the time of the OS search, potentially wrapping several times around Mars but being less than a FOV in “thickness” as viewed from ERO. As with the MAV search, it is beneficial to parameterize a candidate set of pointing targets via semi-major axis with the remaining five orbital elements set to match their most likely values obtained by reducing associated initial OS covariance in combination with their mean values. While a more sophisticated search strategy like the one used for the MAV could be used for the OS, the relative motion between ERO and the OS dispersion is slow enough to enable a simple search strategy. In particular, for each image burst a parameterized candidate target is selected that matches a specified slant range (e.g., 1200 km) that is kept fixed over the course of the search. For each 20-image burst, the camera boresight tracks the candidate target motion to minimize the streaking when detections are evaluated against the Monte Carlo dispersed OS samples. The process then repeats at 10-minute time steps. Since the pointing strategy does not account for illumination conditions, some “banding”

behavior is present in the number of detections across the span of semi-major axes. However, the dispersion passes through the FOV slowly enough to often still be able to image the OS in a favorable geometry. Furthermore, 15 days are nominally allocated for the OS search, which allows the entire dispersion to cycle through twice as viewed from ERO.

## VI. OS ORBIT DETERMINATION

In the MAV orbit determination analysis, each MAV sample is in a different orbit and would have a different OS search start time and ERO orbit, thus varying the OS search space. Future work would focus on connecting each MAV sample to a dispersion of OS samples for an end-to-end Monte-Carlo analysis. For now, to simplify and separate the MAV-OS orbit determination analysis, a few discrete MAV orbit cases with OS searches starting at 10 and 20 days post separation are considered with full initial assumptions listed in Table 6.

Table 6. Assumed 3-sigma OS uncertainties.

	Initial	Requirement
SMA alt. [km]	1.653	0.419
$e \cos(\omega)$	0.0183	1.5e-4
$e \sin(\omega)$	0.0251	1.4e-4
Inclination [deg.]	9.26e-3	7.91e-3
RAAN [deg.]	0.0284	8.42e-3
Arg. lat. [deg.]	0.029	0.058

Using the OS imaging strategy presented in the previous section another Monte Carlo of covariance analyses is performed on the MAV at 415 km and ERO at 385 km OS dispersion cases with search starting at 10 and 20 days post separation. To provide extra conservatism in the mitigation of straylight, it is assumed that the camera boresight stays at least  $17^\circ$  above the Mars limb. These results are plotted in Fig. 7 and Fig. 8.

With a single detection the OS uncertainty significantly collapses from an SMA uncertainty of several kilometers to on the order of 10 meters. Other orbital parameters take around 3 detections to collapse by an order of magnitude. Fig. 9 shows the most important uncertainty metric, namely the OS uncertainty when mapped into the FOV. From this it is clear that after only one detection the OS 3-sigma uncertainty fits within about a single FOV for the next two synodic periods, which allows for accumulating a large number of detections and enabling a precise OD solution.

## VII. CONCLUSION

Locating a basketball-sized OS in orbit around Mars is an undoubtedly challenging problem that requires an interdisciplinary solution. It involves interfaces with a variety of elements of the Mars Sample Return

campaign including the MAV, OS, and ERO. This paper provides a snapshot of a subset of the work being done to demonstrate that its orbit can in fact be determined robustly even when subject to various sources of uncertainty arising from the many components involved.

The procedure presented can leverage both optical and RF detections of the MAV, which can then be used to perform a more focused search for the OS. The current analysis could continue to evolve and mature as the Mars Sample Return architecture takes shape. Though not presented in this work, other assets in Mars orbit can listen for the MAV's beeps to contribute to the orbit determination solution when combined with ERO's detections as well as independently provide a coarse solution verifying successful delivery into orbit. One particular area of ongoing development is an end-to-end simulation environment that would allow updates to the search based on previous detections; currently the search for each is simulated using an open loop strategy. Other areas of ongoing work are in the discrimination between the MAV and OS in images and the improved modeling of the targets' optical properties and straylight.

The decision to implement Mars Sample Return will not be finalized until NASA's completion of the National Environmental Policy Act (NEPA) process. The information presented in this paper is pre-decisional and provided for planning and discussion purposes only.

## VIII. ACKNOWLEDGEMENTS

Developing the Mars Sample Return campaign has involved close collaboration between NASA, ESA, and industry, and that collaboration has been particularly evident in the design of the Launch, Rendezvous, and Capture phase with significant input from ERO's prime contractor Airbus.

Some of the High Performance Computing resources used in this investigation were provided by funding from the JPL Information and Technology Solutions Directorate.

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. © 2024 California Institute of Technology. Government sponsorship acknowledged.

Optical Only OS, 2000 Samples

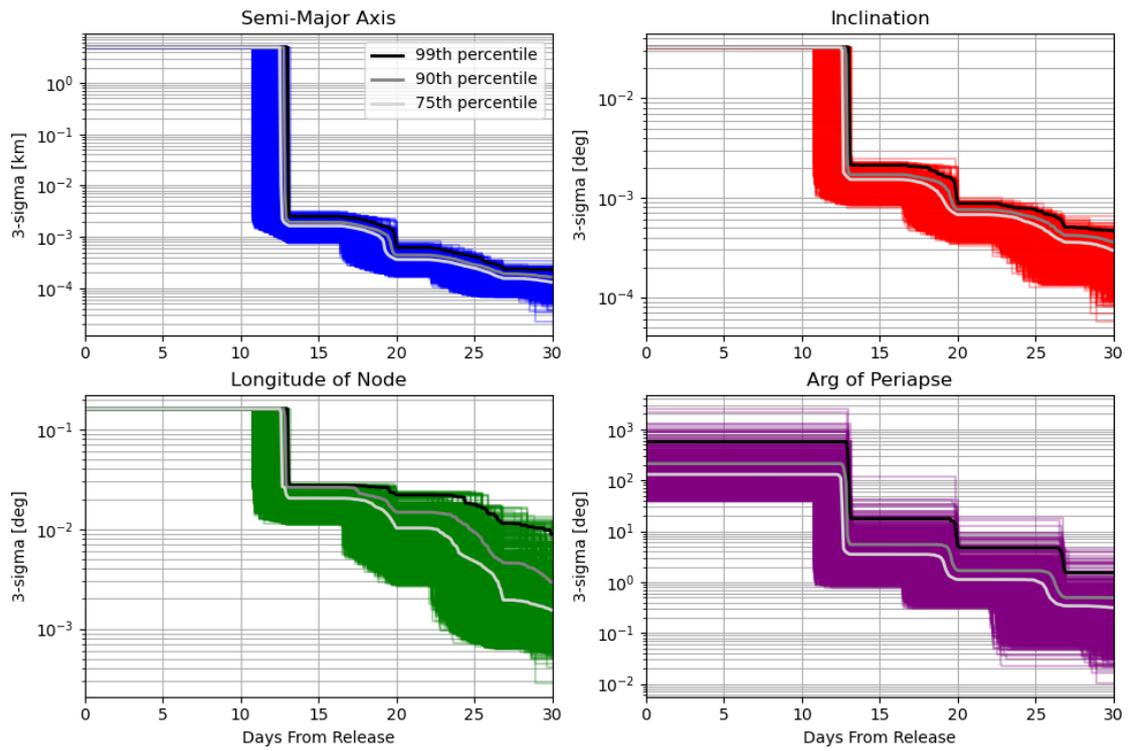


Fig. 7. 3-sigma OS orbital element uncertainties over time for an OS search starting 10 days after release.

Optical Only OS, 2000 Samples

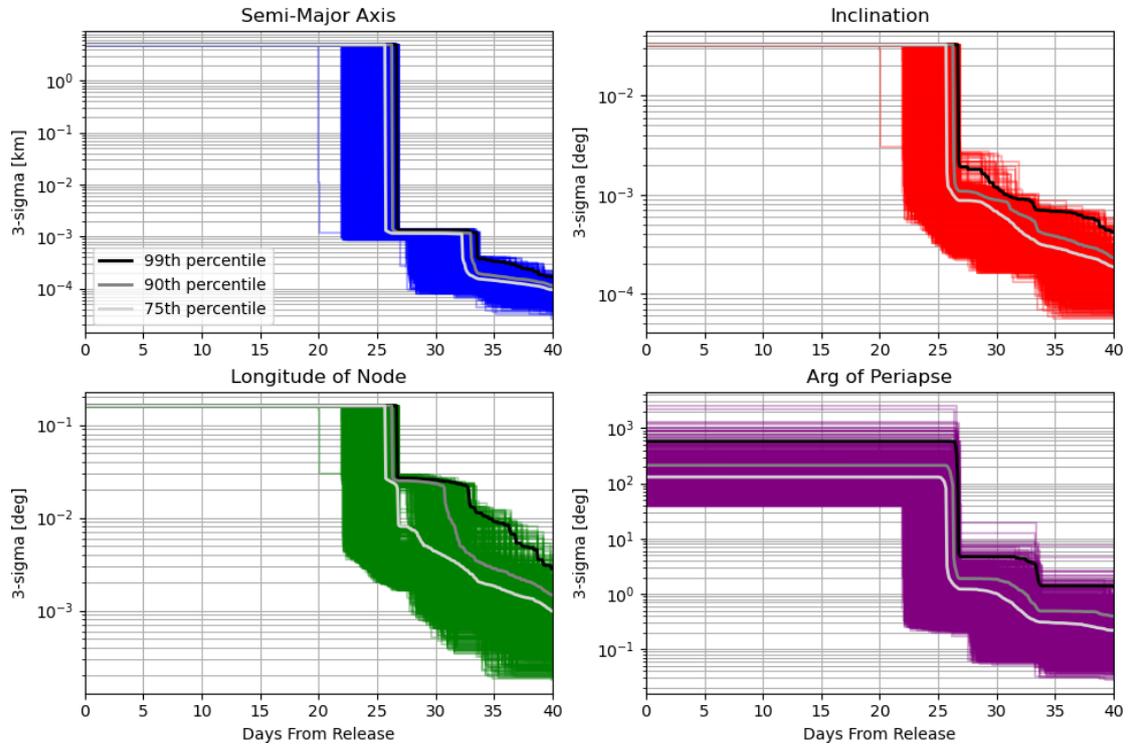


Fig. 8. 3-sigma OS orbital element uncertainties over time for an OS search starting 20 days after release.

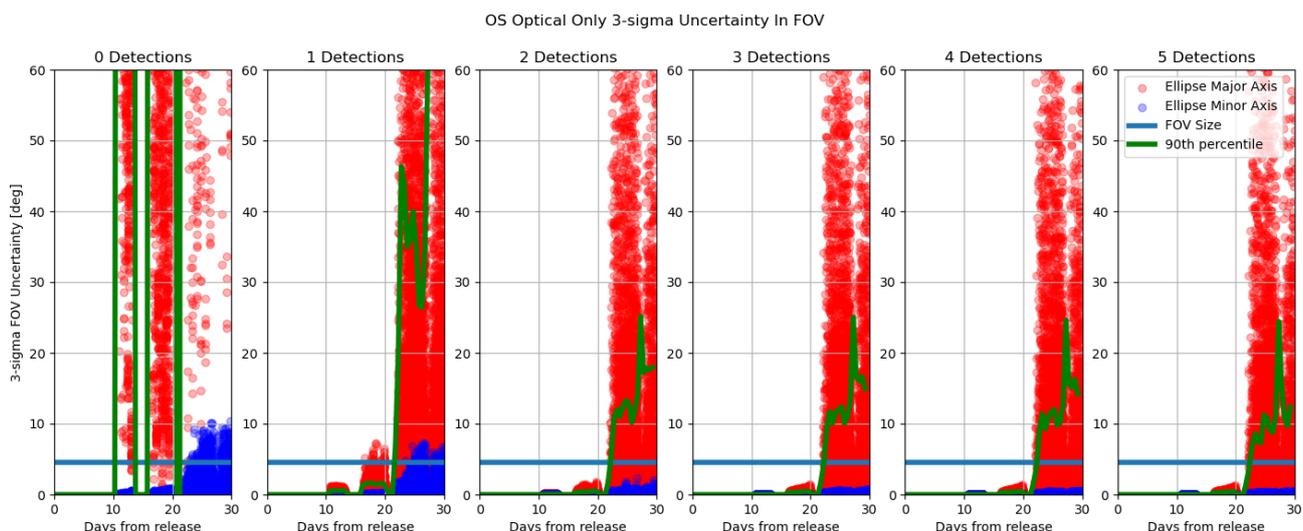


Fig. 9. 3-sigma OS uncertainty in FOV versus detection count for an OS search starting 10 days after release.

## IX. REFERENCES

- [1] National Academies of Science, Engineering, and Medicine, *Origins, Worlds, and Life: A Decadal Strategy for Planetary Science and Astrobiology 2023-2032*. Washington, DC: The National Academies Press, 2022.
- [2] T. Martin-Mur, J. M. Sánchez Pérez, R. Wilson, Z. Olikara, M. Jesick, R. Patel, and D. Spencer, "The NASA/ESA Mars Sample Return Campaign," 29<sup>th</sup> International Symposium on Space Flight Dynamics, Darmstadt, Germany, April 22-26, 2024.
- [3] M. Slipski, A. Kleinböhl, D. Tirsch, G. Kminek, G. Jonniaux, K.-D. Matz, A. Määttänen, A. Nicholas, F. Montmessin, S. N. Madsen, M. Abrahamson, M. Sanchez-Gestido, M. A. Mischna, N. P. Murray, M. J. Wolff, P. Blanc-Paques, F. Cipriani, C. F. Wilson, D. Titov, and R. Zurek, "The radiometric environment for Mars limb observations by the Mars Sample Return Earth Return Orbiter," *Advances in Space Research*, vol. 72, pp. 4048-4063, July 2023.
- [4] A. M. Didion, A. K. Nicholas, J. E. Riedel, R. J. Haw, and R. C. Woolley, "Methods for passive optical detection and relative navigation for rendezvous with a non-cooperative object at Mars," 2018 AAS/AIAA Astrodynamics Specialist Conference, Snowbird, Utah, USA, August 20-23, 2018.
- [5] M. Shao, R. Trahan, C. Zhai, N. Saini, and S. G. Turyshev, "Synthetic tracking on a small telescope," 19<sup>th</sup> Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), Wailea, Hawaii, USA, September 11-14, 2018.