AAS 21-336

COMPARING OPTICAL TRACKING TECHNIQUES IN DISTRIBUTED ASTEROID ORBITER MISSIONS USING RAY-TRACING

Kaitlin Dennison[*] and Simone D'Amico[†]

Missions to asteroids have traditionally relied on the Deep Space Network and Light Detection and Ranging (LIDAR) for landmark tracking and mapping. Optical tracking and multi-agent structure from motion show great potential for improving autonomy and reducing hardware requirements in such missions. However, there is a lack of testing and simulation in this area, especially because validation is difficult. In this work, keypoint descriptors are compared to a target-specific landmark descriptor (craters) for optical tracking within a single- and multi-satellite space mission orbiting an asteroid. Additionally, a novel method to validate correlations between any landmark descriptors using ray-tracing is developed. Optical tracking methods are evaluated in simulation to identify tracking limitations. Simulated test cases are validated using the novel ray-tracing technique. The optical tracking methods are further validated using images from the Near Earth Asteroid Rendezvous (NEAR) Shoemaker mission to Eros. Results indicate that SIFT is the most accurate optical tracking technique across monolithic and distributed space system simulations. However, SURF has the most robustness to lighting conditions and produces the most accurate 3D points from stereo-vision. Craters are shown to be an inconsistent landmark detection method, and normalized cross-correlation appears to be largely unreliable for crater correlation. Despite issues with some optical tracking methods, matching ray-traced landmarks shows great potential for optical tracking validation.

INTRODUCTION

Scientists have made great strides in increasing autonomy of missions to asteroids and other small bodies. However, the field of space exploration does not currently have an industry standard for autonomous, image-based landmark detection and tracking, otherwise known as autonomous optical tracking, for asteroids. Finding such a method is essential for further research on autonomous on-board navigation, mapping, and small body characterization.

Missions to asteroids are of great interest to the space community for the scientific value, which includes insights into solar system formation,[1,2] potential for resource mining,[3] and Earth-asteroid collision prevention.[4,5] Despite their usefulness, missions to asteroids have proven to be quite technically challenging because little a priori information is available, orbits about asteroids are highly perturbed, and asteroids are typically distant from the Earth. These complexities have forced all orbiter missions to date to depend extensively on the NASA Deep Space Network (DSN).[6–9] Relying on the DSN leads to latency in spacecraft maneuvers and slows down the mission timeline. Since

[*]Doctoral Candidate, Stanford University, Department of Aeronautics and Astronautics, Durand Building, 496 Lomita Mall, Stanford, CA 94305-4035.

[†]Associate Professor, Stanford University, Department of Aeronautics and Astronautics, Durand Building, 496 Lomita Mall, Stanford, CA 94305-4035

these missions commenced, researchers have made great strides in making spacecraft control,[10] path planning,[11] and state estimation[12] methods more autonomous and less reliant on the DSN. The area falling behind, however, is autonomous optical tracking.

There are three main gaps in the domain of autonomous optical tracking for spacecraft missions to asteroids. First, spacecraft missions tend to use target-specific landmarks, such as craters and boulders, for optical tracking because of their unique appearances and ability to be manually selected.[6,7] State-of-the-art autonomous tracking methods from other scientific fields, like keypoint descriptors, have been developed. Although these landmarks are more meaningful to computers than humans and may prove to be more robust, a comparison between keypoints and target-specific landmarks has not yet been performed.[13] Second, groundtruth landmark data is difficult to obtain for asteroids. A groundtruth is necessary to validate landmark correlations, thus compounding the difficulty of testing new optical tracking methods for asteroids. Finally, autonomous optical tracking has not been tested for multi-satellite missions. Spacecraft orbiting small bodies have inherently different dynamics than planetary use cases, which may increase error in state estimates for multi-agent navigation and mapping techniques that use the optical tracking data.

Most missions to small bodies to date have used target-specific landmarks because humans can manually select them as necessary[6] and they are scale- and illumination-invariant.[14] These qualities of invariance are especially important for small body missions because spacecraft often revisit the same section of surface many times with different perspectives and lighting conditions. The same landmarks need to be identified each time to produce accurate navigation and mapping data.[12]

Keypoint descriptors are the most common optical tracking technique in the field of robot autonomy because of their ease of use and scale- and rotation-invariance.[15,16] Takeishi et al. showed that keypoint descriptors can be a viable option for autonomous optical tracking on asteroids.[13] However, techniques that work on a planet surface may not translate perfectly to an asteroid. In addition, Takeishi et al.'s study did not compare keypoint descriptors to target-specific landmarks or evaluate tracking under different lighting conditions, leaving it unclear if keypoint descriptors can be used in place of target-specific landmark tracking. If missions are to switch to fully autonomous optical tracking, an autonomous version of target-specific tracking should be compared to the state of the art of other fields, especially under lighting changes.

Comparing optical tracking techniques is complicated because it is difficult to obtain a groundtruth for 3D locations of landmarks on the surface of an asteroid. Without a groundtruth database of landmarks, the optical tracking statistics cannot be quantified. Typically, optical tracking is verified by tracking the same landmarks in an image and a distorted version of the same image.[16] However, this method would not support rigorous simulation of multiple, sequential images. Takeishi et al. attempted to solve this problem for asteroid optical tracking by calculating overlapping regions of matched landmarks on a model of their target asteroid.[13] A limitation of their region-matching method is that the region must be densely packed with landmarks. In contrast, this paper matches 3D positions obtained from back-propagating the camera perspective through the asteroid 3D model, a process known as ray-tracing. Ray-tracing is simple to implement and can be used for any landmark descriptor, regardless of how densely packed the landmarks are, as long as a 3D asteroid model and high-fidelity spacecraft state information are available.

Ray-tracing landmarks can also be used to validate the final gap addressed by this paper: how does multi-agent 3D mapping perform for spacecraft using autonomous optical tracking? Many researchers are considering distributed space systems (DSS), which is two or more spacecraft co-

operating to complete an objective, for exploration missions. DSS allow more redundancy and flexibility than traditional monolithic spacecraft. Additionally, they can be used to perform multi-agent mapping techniques such as Simultaneous Localization and Mapping (SLAM),[17] structure from motion, and stereo-vision. SLAM is the most popular autonomous navigation and mapping technique because of its speed and ability to adapt to almost any application. Structure from motion and stereo-vision are both global mapping algorithms that cannot perform navigation. They are generally used in place of Light Detection and Ranging (LIDAR) to calculate 3D points of landmarks.[18] All three techniques are common for planetary exploration with ground-based robots[19,20] and unmanned aerial vehicles.[17]

But again, techniques that work on a planet surface may not translate perfectly to an asteroid. The dynamics of a spacecraft orbiting a small body are inherently different than Earth-based dynamics. As such, navigation and mapping algorithms may not work the same and much be tested rigorously in the asteroid environment. Rather than testing all three mapping methods mentioned, this paper focuses on verifying stereo-vision. Stereo-vision is the more basic than SLAM and structure from morion. Thus, stereo-vision can act as an indicator of the performance of the other two: if stereo-vision fails, the other two most likely will too.

This paper addresses each of these gaps by comparing the performance of a target-specific optical tracking method to keypoint descriptors. These comparisons use a novel method to validate correlations between any landmark descriptors by using the known asteroid 3D model. Further evaluation is performed by simulating multi-agent 3D mapping performance from stereo-vision for each optical tracking method.

To determine the best autonomous optical tracking method for asteroid missions, this paper evaluates four different autonomous optical tracking methods in four separate test cases using Asteroid 433 Eros as the target. One target-specific optical tracking method, correlating craters via normalized cross-correlation, is compared to three keypoint descriptors and their associated correlation metrics. The four test cases are: simulated optical tracking robustness to lighting changes; simulated optical tracking robustness to single spacecraft orbit configuration; simulated stereo-vision robustness to 2-agent DSS configuration; and optical tracking accuracy in real images from the Near Earth Asteroid Rendezvous (NEAR) Shoemaker mission to Eros.[21] The accuracy of optical tracking in simulated images is quantified using this paper's novel method of matching 3D point from ray-tracing.

The remainder of this paper is organized as follows. First, the optical tracking methods used are described along with their advantages and disadvantages. Next, the algorithm used for matching ray-traced landmarks is explained in detail. Finally, the four experimental test cases and their results are presented and discussed.

OPTICAL TRACKING METHODS

This work compares four optical tracking methods. Each method is a specific landmark detector and correlation metric pair. Three use keypoint descriptors while the third uses crater detection, which is a target-specific descriptor. The results of all four tracking methods are further refined using modified z-score outlier rejection.

Keypoint Detection and Matching

The three keypoint descriptors used are Scale-Invariant Feature Transform[22] (SIFT), Speeded-Up Robust Features[23] (SURF), and Oriented FAST* and Rotated BRIEF† (ORB).[24] These three descriptors were chosen for their popularity and their scale- and rotation-invariance. Furthermore, these keypoint descriptors were also studied by Takeishi et al. for mapping Asteroid 25143 Itokowa,[13] which provides a validating comparison.

SIFT keypoints are the most commonly used keypoint descriptor because they work with most applications. These keypoints are the extrema of the difference of Gaussian of the image. SIFT is often called a "blob" detector because of the natural shape of the extrema detected. Each keypoint is described by a 128-dimensional vector that represents directional image intensity gradients around the point.[13,22] Features are matched via the smallest Euclidean distance between each feature's 128-dimensional vector. These matches are refined by the ratio between the distances of the first-closest match and the second-closest match.[22]

SURF keypoints are very similar to SIFT keypoints. They are also blob-shaped but are detected using an approximation of the Hessian matrix, which makes SURF more computationally efficient than SIFT. Each keypoint is described by a 64-dimensional vector representing Haar wavelets in the x and y directions. SURF keypoints are matched using the same method as SIFT keypoints except that sum of squared differences is used to compare vectors instead of Euclidean distance.[13,23]

ORB keypoints use a corner detector as opposed to a blob detector. Since ORB is a corner detector and asteroids do not generally have corners, it is not expected to perform well. However, ORB is still included because it is extremely popular and is more computationally efficient than SIFT and SURF. ORB keypoints look at the intensity difference between pairs of pixels surrounding a point. The signs of the differences are used to define each keypoint by a binary code. Scale- and rotation-invariance are obtained by estimating orientation and scale via intensity gradients. ORB keypoints are matched using the same method as SIFT keypoints except that the Hamming distance is used to compare binary code vectors instead of Euclidean distance.[13,24]

Crater Detection and Matching

A crater detection algorithm was chosen to represent target-specific landmark descriptors because craters are the most prolific natural feature on Eros' surface.[25] If a different type of natural landmark had been more prevalent, such as boulders on Asteroid 101955 Bennu, a descriptor for that type of landmark would have been chosen instead. This paper uses Maass et al.'s edge-free crater detection algorithm, which looks for pairs of shaded and bright patches that remain consistent across many levels of grayscale thresholding.[26] Maass et al. demonstrated that their method performs better than the industry-standard for crater detection: a method designed by Cheng et al. for the Jet Propulsion Lab in 2003.[14]

Once the craters are identified in the images, they must be correlated between the images. The most common crater-matching algorithms take advantage of a pre-defined crater database such as identifying shapes created by multiple craters[27] or matching a crater image synthesized from a computer-generated model.[6] Others use a spacecraft state estimate to determine where matches could be based on changes in the camera perspective.[14] These methods require environment knowledge (a crater database or spacecraft state) that is not always easy or possible to obtain without using

*Features from Accelerated Segments Test

†Binary Robust Independent Elementary Features

the DSN. However, a major benefit of keypoints is that they do not require such environment knowledge to perform correlation. Thus, it is preferable that the target-specific optical tracking method not require environment knowledge either.

There are two popular crater-matching methods that do not require environment knowledge. The first is to identify conic invariants of the ellipse parameters.[28,29] However, these conic invariants are only valid when the relationship between the two image views is a homography, which is not true for sequential images of asteroids. The homography between images is invalidated because the asteroid rotates between them and points detected will not all lie on the same plane in 3D space.[18] The second method is normalized cross-correlation (NCC), which is a way to quantify the pixelwise difference between two images.[6,14] NCC only requires that images be as similar as possible. It is not as robust as any of the other crater-correlation algorithms mentioned here but it is the most similar to the correlation metrics used for keypoint descriptors.

Therefore, this paper uses NCC to correlate craters between two images. Once craters are detected in an image, the cropped images of individual craters are compared. The pair of craters with the highest NCC metric are matched. Spurious matches are eliminated via an empirically determined minimum-NCC threshold and a maximum 2D Euclidean distance between points based on the camera perspective.

Modified Z-Score Outlier Rejection

Optical tracking is typically paired with an outlier rejection method to further refine matches. Random sample consensus[30] (RANSAC) is the most popular but it operates by fitting a model to the data, such as the camera perspective or a homography. Having a state transition model that includes asteroid rotation and spacecraft state would defeat the purpose of this paper since autonomous optical tracking is used to support navigation and asteroid characterization techniques. As stated previously, the homography is invalid for sequential images of an asteroid. Thus, in this work, correlated landmark pair outliers are identified by the modified z-score of the Euclidean distance between the pair's 2D points from the two images. Modified z-score outlier rejection uses the median of a univariate dataset in lieu of the mean to calculate the z-score of each data point. The original author of the z-score recommends using the modified z-score instead of the true z-score when the distribution of the dataset may be skewed.[31] Imaging an oblong asteroid like Eros will likely produce a skewed distribution of 2D distances between matched points because of the camera perspective: the far end of the asteroid will appear to move less between images than the closer end. Given a dataset of 2D points matched between two images, the modified z-score (z_i) of the Euclidean distance (d_i) between two matched points in the image frame is calculated as,

$$z_i = \frac{0.6745\left(d_i - \tilde{d}\right)}{m}. \tag{1}$$

Here, $\tilde{d}$ is the median distance between two matched points and m is the median absolute deviation, which is defined as,

$$m = \text{median}\left|d_i - \tilde{d}\right| \ \forall_i. \tag{2}$$

Datapoints with a modified z-score above 3.5 are considered outliers, as also recommended by the original authors of the modified z-score.[31]

GROUNDTRUTH VALIDATION USING RAY-TRACING

A quantitative metric is essential for validating and comparing each optical tracking method. This paper introduces a novel method for validation that identifies the groundtruth landmark 3D locations on the asteroid surface via ray-tracing and matches them according their Euclidean distance. Ray-tracing converts a 2D image point into 3D space using a known camera perspective and 3D model of the scene. Although Takeishi et al. performed a similar groundtruth process by matching overlapping regions of landmarks with the 3D model surface, their method can only operate with densely packed landmarks.[13] This paper's method of matching ray-traced points relies on the points not being densely packed. This quality is especially desirable for target-specific landmarks because they can be sparsely distributed. The process of ray-tracing is described as follows.

The center of a landmark on a 3D model of an asteroid corresponds to a 2D point in an image of the asteroid. If the camera perspective that relates the image to the 3D scene is fully known, a ray from the camera's optical center through the landmark 3D position can be calculated.[18] For a 2D image point ($\boldsymbol{p}$), an arbitrary homogeneous 3D point ($\boldsymbol{P}^h$) that lies along the ray from the camera through the true 3D point in the asteroid-centered frame is found by

$$\boldsymbol{P}^h = \begin{pmatrix} K \begin{bmatrix} R & \boldsymbol{r} \end{bmatrix} \end{pmatrix}^\dagger \begin{bmatrix} \boldsymbol{p} \\ 1 \end{bmatrix}. \tag{3}$$

Here, K is the camera intrinsic matrix, R is the rotation matrix from the asteroid-centered frame to the camera-centered frame, $\boldsymbol{r}$ is the position vector from the camera to the center of the asteroid in the camera-centered frame, and † indicates the Moore-Penrose pseudoinverse.[18] The ray from the camera through the point in the asteroid-centered frame is then recovered through

$$\hat{\boldsymbol{u}} = \begin{bmatrix} \boldsymbol{P}^h(1) \\ \boldsymbol{P}^h(2) \\ \boldsymbol{P}^h(3) \end{bmatrix} \frac{1}{\boldsymbol{P}^h(4)} + R^\top \boldsymbol{r}. \tag{4}$$

The distance (ν) along the ray from the camera to the true 3D point ($\boldsymbol{P}$) is then recovered by finding the ray-triangle intersection of the ray and the model's triangular mesh.[32] $\boldsymbol{P}$ is fully recovered as

$$\boldsymbol{P} = \nu \hat{\boldsymbol{u}} - R^\top \boldsymbol{r}. \tag{5}$$

In simulated images, K, R, and $\boldsymbol{r}$ are fully known, which produces the camera perspective. This means any point in an image can be ray-traced back through the model to obtain the 3D position of the point in the asteroid-centered frame. Because the landmark represents a point on the asteroid surface, ray-tracing the center of a detected landmark will produce the same 3D position every time, meaning the ray-traced point is the groundtruth of where that landmark is located on the asteroid. One note is that the ray-traced points are not completely accurate. They are subject to how fine the triangular mesh of the model is[32] and to the reliability of centroiding detected landmarks. Subsequently, landmarks can be ray-traced in each image and matched across images via the Euclidean distance between ray-traced points assuming the error produced by the mesh and centroiding is small.

EXPERIMENTAL TEST CASES AND RESULTS

All four optical tracking methods were rigorously tested using simulated and real imagery of Eros. Simulated images were produced from a high resolution triangular mesh model of Eron.[8,33] The camera model used for simulation is based on the GOMSpace NanoCam C1U with an 8 mm focal length.[34] Real images were obtained from the NEAR Shoemaker mission to Eros.[21] Optical tracking was tested for both single spacecraft and a DSS via simulation while only single spacecraft optical tracking was tested using the real images from NEAR Shoemaker. All experiments were generated in MATLAB using VLFeat* for SIFT and the MATLAB Computer Vision Toolbox for SURF† and ORB‡.

Test Case 1: Landmark Detection Lighting-Invariance

Many types of spacecraft navigation require the same landmarks to be detected when the spacecraft returns to the same section of asteroid surface.[12] When a section is revisited, the lighting conditions can be very different from the previous visit because of asteroid rotation. To test how well a landmark detection method identifies the same landmarks under different lighting, a series of 50 images are generated with the same perspective of Eros but with different lighting conditions. The lighting was changed by rotating the vector from the asteroid center to the Sun in an archimedean spiral about a hemisphere centered between the camera and the asteroid model.

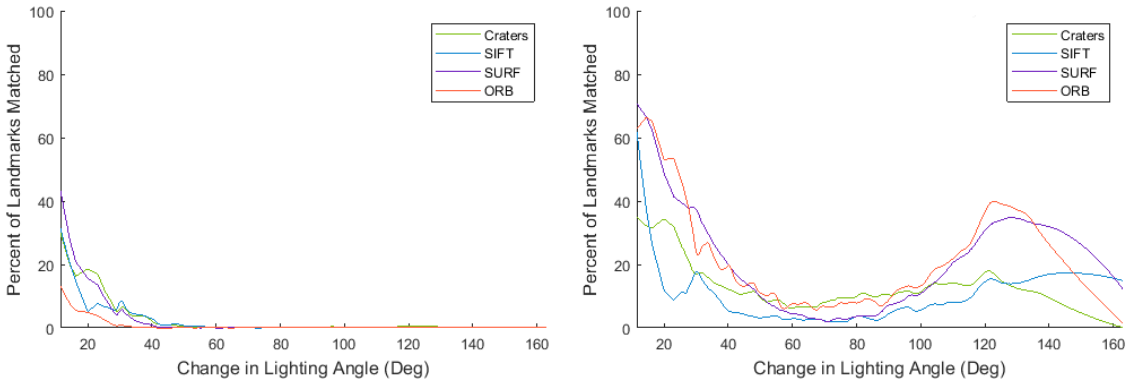

Figure 1. Percentage of landmarks matched between two images of the same scene with lighting from different angles. (Left) Landmarks are matched by their respective, formal optical tracking correlation metric. (Right) Landmarks are matched by the Euclidean distance of their ray-traced 3D positions representing groundtruth correlations. All optical tracking correlation metrics fail under significant lighting changes. However, the same landmarks are still able to be detected under different lighting conditions for some cases of SURF and ORB.

All four optical tracking methods are used to detect and match landmarks between all 50 images (1225 total unique image pairs). Figure 1 shows the percentage of landmarks that were matched between each image and the change in lighting angle between the two images that were matched. The change in lighting angle is the angle between the vector from the asteroid center to the Sun used for each image being compared. The data was smoothed with robust local regression using weighted

*https://www.vlfeat.org/

†https://www.mathworks.com/help/vision/ref/detectsurffeatures.html

‡https://www.mathworks.com/help/vision/ref/detectorbfeatures.html

linear least squares to denoise the data for visual clarity.[35] An ideal optical tracking method is illumination-invariant, meaning it will detect and correlate the same landmarks regardless of the lighting condition (assuming the scene is sufficiently illuminated). Thus, an illumination-invariant tracking method will have a decently high percentage of matched landmarks consistent across all changes in lighting angle.

Figure 1 shows that all four of the optical tracking method's formal correlation metrics fail to correlate the same landmarks when the lighting angle changes more than 20°. This is expected since all four correlation methods are purely based on similarity of image intensity, which depends entirely on the lighting conditions. However, this experiment is still useful because it shows that some detection methods used may have lighting-invariance even if the correlation methods do not. Figure 1 also plots percent of landmarks matched via the Euclidean distance of their ray-traced points, which indicates how many of the same landmarks were actually detected in the two images being compared.

Exclusively looking at landmarks matched via ray-tracing (the right-side plot in Figure 1), Craters (green line) and SIFT (blue line) have consistent but poor detection performance with a match percentage typically below 20%. One difference between the two is that SIFT has a much higher match percentage than craters for a lighting change of less than 20°. This indicates that craters may be generally inconsistent in detection between images irrespective of lighting while SIFT is consistent between images as long at the lighting does not change much.

SURF (purple line) and ORB (orange line) appear to be slightly less prone to detection change due to lighting conditions. They both have two peaks with greater than 20% matching: 60-70% at 0° and 35-45% around 135°. In the lighting angle change range of 40-110°, the shadows produced by landmarks change significantly in size and direction. For instance, a partially-shaded crater may become fully illuminated or a previously-illuminated section of the asteroid might drop into shadow. Almost any image-intensity based landmark detection method can not be expected to perform well with such drastic changes. Furthermore, the 135° peak occurs because the same size shadows appear to be coming from the opposite direction of the image compared, like a mirror image of the shadows. One might expect this to occur at a 180° change in lighting angle but the surface of the asteroid is highly curved and the asteroid body itself begins to block the light after 135°.

From these results, craters are inconsistent in detection regardless of lighting while SIFT is extremely influenced by the lighting condition. SURF and ORB are as robust to lighting changes as any intensity-based landmark detection method could be. In order to obtain this level of lighting-invariance, landmarks can be detected using SURF or ORB and the landmarks should be matched using a metric known to be lighting-invariant, such as Mahalanobis distance,[15] which would generally require a full dynamics model or an asteroid 3D model.

Test Case 2: Sequential Optical Tracking

The second test case is designed to evaluate the sequential tracking performance of the four optical tracking methods and the viability of using ray-traced points as the groundtruth. All images in this case are entirely simulated. Given a chosen asteroid target and a camera model, two tuning parameters are used to define the simulated images. First, the semi-major axis of the spacecraft orbit determines the speed and distance of the spacecraft from the asteroid surface. Second, the time interval between images directly determines what landmarks will enter and leave the frame and how much the perspective will change between images.

The same time intervals cannot be used for each semi-major axis because the spacecraft travels at a different speed and observes a different portion of the surface based on the orbit size. The maximum imaging time interval for a particular orbit is bound by how quickly a landmark leaves the image frame. Thus, the time intervals are based on a factor of the maximum imaging time interval for each semi-major axis.

An rough order of magnitude approximation is used for the maximum imaging time interval for a particular semi-major axis. This approximation is derived as follows, assuming the spacecraft is in a prograde, equatorial, circular orbit about the asteroid. The same landmark needs to be detected multiple images in a row in order to build an accurate estimate of its 3D position for a landmark map.[12] Thus, this approximation assumes a landmark travels no more than one fourth of the frame width between images, meaning it will be detected in at least 3 sequential images. This constrains the maximum imaging time interval by the time is takes a landmark to leave the image frame.

The instantaneous speed of a landmark (v_l) is found using the asteroid rotation rate[36] ($\dot{W}$) and average radius (ρ) of the asteroid,

$$\dot{W} = \frac{v_l}{\rho}. \tag{6}$$

The instantaneous speed of the spacecraft (v_s) is found assuming a circular obit using the spacecraft orbit's semi-major axis (a) and the asteroid's gravitational parameter (μ),

$$v_s = \sqrt{\frac{\mu}{a}}. \tag{7}$$

The spacecraft's camera is assumed to always be pointed at the asteroid center of mass and to rotate with the orbit. Thus, there is no contribution from angular velocity. Now, the relative velocity (v_r) of the landmark to the spacecraft is broken into the distance (h_l) the landmark travels relative to the spacecraft in an increment of time (δt),

$$v_r = v_l + v_s = \frac{h_l}{\delta t}. \tag{8}$$

Using the single lens approximation for image magnification,[37] the relationship between h_l and the distance (h_s) that the landmark travels in the image frame is given by,

$$\frac{h_l}{h_s} = \frac{d_l}{d_s} = \frac{a-\rho}{f}. \tag{9}$$

Where d_l is the distance from the camera to the landmark, d_s is the distance from the lens to the image, and f is the focal length. Plugging in equations (6)-(8) into equation (9) and using one quarter of the camera sensor width (w) as h_s, the maximum imaging time interval ($\delta t_{\max}(a)$) for a particular a is,

$$\delta t_{\max}(a) = \frac{(w/4)(a-\rho)}{\left(\dot{W}\rho + \sqrt{\mu/a}\right) f}. \tag{10}$$

Now that $\delta t_{\max}(a)$ has been defined, the experimental parameters can be established. This paper uses an a of $[23, 41, 59, 77, 95]$ km and $\delta t_{\max}(a)$ factors of $\delta t_{\max}(a)/[0.5, 1, 2, 2.5, 3]$ for a total of 25 simulations. A factor of 0.5^{-1} is included to verify the $\delta t_{\max}(a)$ approximation. The spacecraft is propagated in an asteroid-centered-inertial (ACI) frame whose z-axis aligns with the asteroid axis of rotation and whose x-axis aligns with the Vernal Equinox. The satellite orbit elements are: 0.01 eccentricity, 15° inclination, 35° right ascension of the ascending node, 10° argument of periapsis, and 9.98° initial mean anomaly. Eros' position was determined from JPL Ephemerides data* and Eros' rotation was calculated using rotational elements specifications from the IAU.[36] All simulations start at 12:00:00 UTC on Jan 1, 2000 and run for two satellite orbit periods. This date and orbit specifications were chosen to have a decent view and lighting on Eros' Psyche Crater and overlap with the NEAR Shoemaker mission. The satellite orbit is slightly inclined because an equatorial orbit about Eros is unstable.[38]

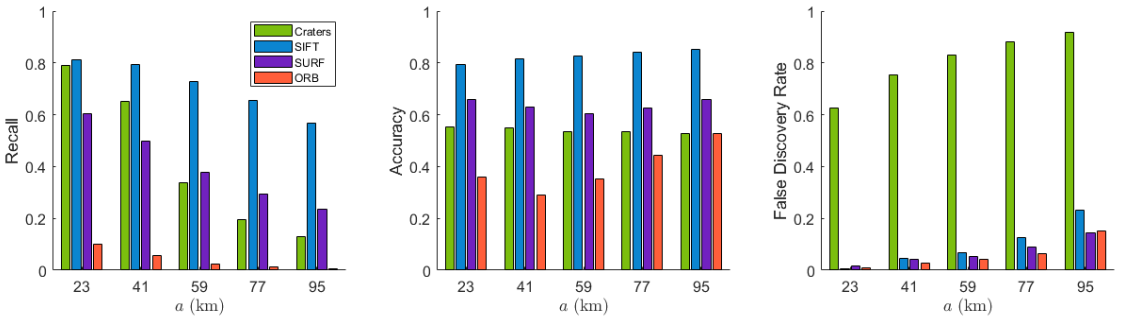

Figure 2. Binary classifier rates for all simulations with a $\delta t_{\max}(a)$ factor of 2^{-1} using matches of ray-traced points as the truth data. All tracking methods are quite sensitive to the semi-major axis with the exception of SIFT.

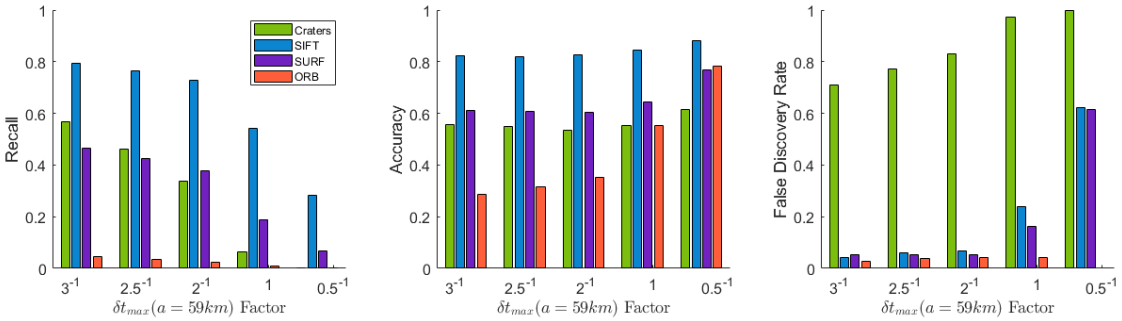

Figure 3. Binary classifier rates for all simulations with $a = 59km$ using matches of ray-traced points as the truth data. SIFT is the most consistent across changes in time interval between images, followed by SURF. There is a definitive upper bound to the time interval between images where all methods fail: $2^{-1}\delta t(a)$.

*https://ssd.jpl.nasa.gov/?ephemerides

Figures 2 and 3 show the binary classifier rates for each optical tracking method with the semi-major axis and time interval factor varied, respectively. The groundtruth data for binary classification is produced by ray-tracing each landmark through the 3D Eros model and matching landmarks via their Euclidean distance. Recall is calculated as,

$$\text{Recall} = \frac{TP}{P} \tag{11}$$

where P is the number of correlated landmarks in the groundtruth and TP is the number of correlations the associated optical tracking correlation metric made that matched the groundtruth correlated pairs. Ideally, $TP = P$ so "good" recall approaches 1. Accuracy is calculated using,

$$\text{Accuracy} = \frac{TP + TN}{P + N} \tag{12}$$

where N is the number of landmarks that were not matched with another landmark according to the groundtruth and TN is the number of the same landmarks that went unmatched as those not matched by the groundtruth. Ideally, $TN = N$ so "good" accuracy approaches 1. Finally, the false discovery rate is defined as,

$$\text{False Discovery Rate} = \frac{FP}{FP + TP} \tag{13}$$

where FP is the number of landmarks that were matched to another via the respective optical tracking metric but that match was not the same as the match (or lack thereof) made by the groundtruth. Ideally, no false discoveries are made.

Looking at the overall trends, recall decreases and false discovery rate increases as the semi-major axis increases. This is because the landmarks are much smaller and less detailed in images the farther away the spacecraft is. Furthermore, optical tracking in general performs better the less time there is between images. This is to be expected because the less time between images, the less perspective change there is. All four tracking methods performed very poorly for $0.5^{-1}\delta t_{\text{max}}$, which shows that the order of magnitude approximation for δt_{max}, Equation (10), is a good estimator for the time interval bound.

SIFT performs the best out of all the optical tracking methods, followed by SURF. SIFT is the only tracking method with consistently "good" classifier rates across all tests, the others are only good for specific cases. SURF performs well for sub-59 km orbits and $2^{-1}\delta t_{\text{max}}$ or faster imaging.

Craters are the most influenced by a and δt. They actually have better recall than SURF for the 23 and 41 km orbits. However, those rates worsen rapidly at 59 km and above. These results are likely because crater detection and NCC is highly dependent on the images being very detailed and similar to each other. Crater's extremely high false discovery rate indicates that crater correlation is not to be trusted and would need a better outlier rejection method than the modified z-score.

ORB is an outlier amongst the four optical tracking methods. As expected, ORB does not perform well for optical tracking on an asteroid because most landmarks are natural, curved shapes rather than straight lines and corners. It mostly picks up many overlapping interest points on the very edge of the asteroid where the illuminated surface meets the black background at odd angles.

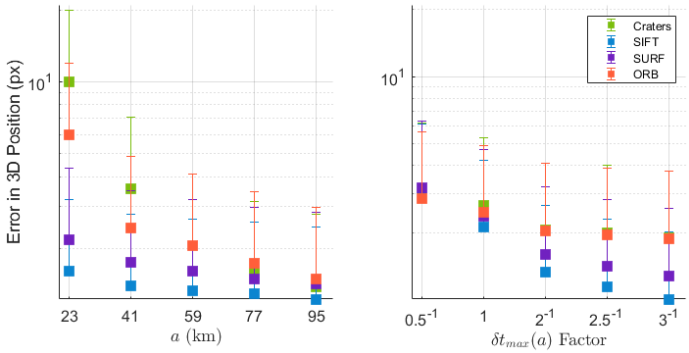

Figure 4. The mean accepted 3D Euclidean distance error when ray-traced landmarks are matched for each optical tracking method during the single spacecraft simulations. All accepted error metrics are significantly smaller than the expected landmark size, which indicates that ray-tracing points produces a reliable groundtruth.

Figure 4 shows the mean accepted error when ray-traced landmarks are matched for each optical tracking method. In all cases, the mean accepted error for SIFT and SURF is below three pixels, which is at the resolvable limit of a visual landmark. At this limit, it is highly unlikely that correlating ray-traced points is incorrectly matching two landmarks because SIFT and SURF generally do not produce co-located landmarks. However, this is not the case for ORB. While the accepted error is just as low for ORB, ORB tends to produce multiple landmarks in the same location, which will produce spurious matches from Euclidean distance. Craters are much larger than keypoints so their accepted error is expected to be larger. The average crater radius of Eros is 0.5 km,[25] which corresponds to 86 pixels when the spacecraft is at 23 km. Thus, an error on the order of 10 pixels is entirely reasonable.

In this simulated sequential tracking test, SIFT was the most consistent performer and is the most accurate. However, SURF follows closely behind SURF but is more sensitive to orbit parameters. Craters have much too high of a false discovery rate to be useful. Equation (10) did prove to be a good estimator for the maximum time interval between images. Finally, using the ray-traced points as the groundtruth for landmarks proved to work very well as long as the landmarks were not too densely packed, as was expected.

Test Case 3: Sequential Optical Tracking in NEAR Shoemaker Images

Comparing results from simulated images and real images provides insight into how simulation performance translates to a real mission. Two sets of 30 sequential images from NEAR Shoemaker's approach to Eros are selected to test all four optical tracking methods.[21] Ray-tracing cannot be used to validate tracking for these images because of the lack of high-fidelity spacecraft attitude data for each image. Therefore, the number of correct (true positive, TP) and incorrect (false positive, FP) matches in each image is manually determined. The other binary classifiers cannot be determined by manual selection. As a result, only the false discovery rate is used.

Figure 5 shows the false discovery rate for the two NEAR Shoemaker sets of images compared to the false discovery rate for the $a = 23km$, $3^{-1}\delta t$ simulated sequential tracking test from Test Case 2. This simulated test was chosen because the images are most similar to the NEAR Shoemaker images. Optical tracking in real images shows the same trend as in simulated images: crater detection and tracking has an extremely high false discovery rate when compared to the keypoint descriptors.

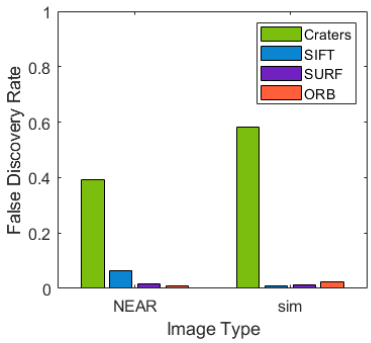

Figure 5. The false discovery rate for the four optical tracking methods in the two NEAR Shoemaker sequential tracking test cases compared to the 23 km, $3^{-1}\delta t$ sequential tracking simulation. The overall trend is the same between real and simulated images: keypoints have relatively low false discovery while craters have exceptionally high false discovery.

Furthermore, the similarity to Test Case 2 in terms of false discovery rate reconfirms that this paper's novel ray-tracing method for groundtruth is accurate. Both manual selection of the groundtruth and matching via 3D Euclidean distance of ray-traced points produced the same data trends.

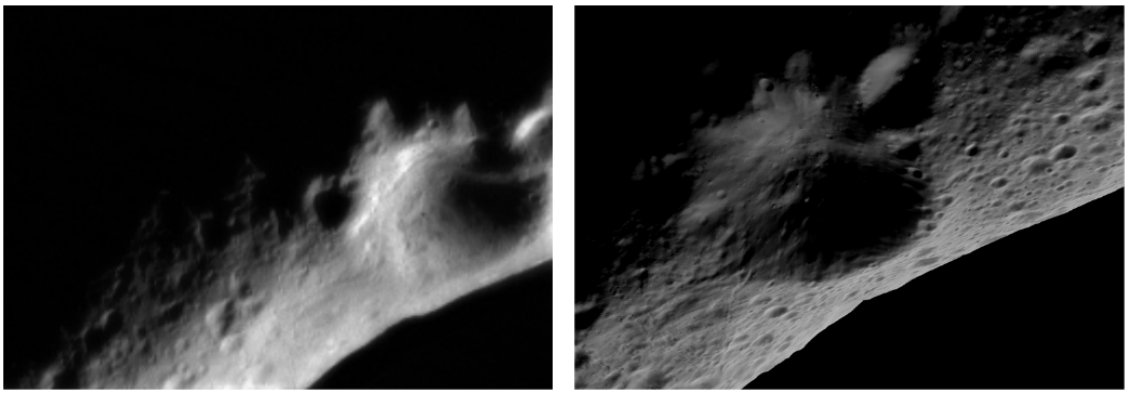

Figure 6. A real image (left) from the NEAR Shoemaker mission to Eros and a simulated image (right) of Eros using the GOMSpace NanoCam C1U as the camera model in a 23 km orbit. The simulated image is significantly more detailed than the real image because the NanoCam is a high-resolution camera and there are no distortions or light scattering modeled. Despite these differences, optical tracking performance is comparable between real and simulated images for all four methods.

One note, however, is that the false discovery rate is slightly lower for craters in the real images. This can be attributed to the quality of the images themselves. A real and a simulated image are compared in Figure 6. The real images are lower resolution, have light scattering, and are slightly blurry, which makes the smaller craters and other features disappear. Thus, only the more obvious, high-contrast craters can be detected in the real images. NCC and crater detection are more accurate if the differences in pixel intensities across the crater are stronger. Thus, NCC and crater detection will perform better on real images where only the most intense craters can be detected.

Test Case 4: Multi-Agent Optical Tracking and Stereo Vision

The fourth test case is designed to evaluate each optical tracking method's robustness to multi-spacecraft orbit configurations and viability for use in 3D mapping. Multi-agent stereo-vision is used as an indicator of the potential performance of other mapping algorithms like SLAM and structure from motion. Test Case 4 uses all the same orbit and camera specifications as Test Case 2: Sequential Optical Tracking. However, instead of varying image timing, the spacecraft along-track separation ($a\delta\lambda$) is varied because the two spacecraft are assumed to take images at the same time. If the spacecraft are too far apart, the images will not overlap and they will not see the same landmarks. Thus, the $a\delta\lambda$ is constrained to ensure a sufficient image overlap.

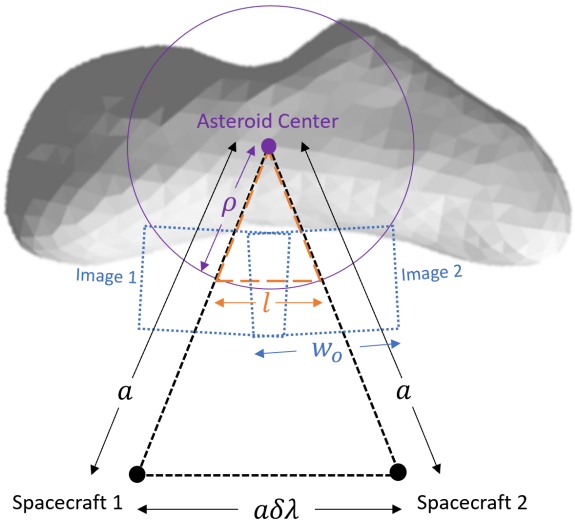

Figure 7. The geometry of two spacecraft and their image frames projected onto the average radius (ρ) of Asteroid 433 Eros. The spacecraft are in a circular, equatorial orbit with semi-major axis, a, and spacecraft separation, $a\delta\lambda$. Each image has a projected width of w_o and the centers of the images are a distance, l, away from each other.

The maximum spacecraft separation for the 2-agent DSS shown in Figure 7 is calculated using an order of magnitude approximation, which is derived as follows. Equation (9) is used to calculate the width (w_o) of the image projected to a point on the asteroid surface by replacing h_l with w_o and h_s with w, the camera sensor width. This results in,

$$w_o = \frac{w(a-\rho)}{f}. \tag{14}$$

In order to have a significant number of landmarks overlap between images but not constrain spacecraft separation too much, a minimum overlap of $\frac{3}{4}w_o$ is used. The distance (l) between the centers of the two spacecrafts' image frames projected onto the asteroid surface is $\frac{1}{4}w_o$. This is also analogous to the landmark not traveling more than $\frac{1}{4}$ of the image frame between images, which constrained the time interval between images in Test Case 2. As seen in Figure 7, a triangle is formed between the asteroid center and the two spacecraft (black, dashed triangle). The projected image frame centers lie along the triangle sides between the asteroid center and the spacecraft (orange, dashed triangle). The distance from the asteroid center to either spacecraft is a and from the

asteroid center to either projected image center is ρ (the average radius of the asteroid shown as a purple circle). Thus, the relationship between a, ρ, and $a\delta\lambda$ is found using similar triangles:

$$\frac{a}{\rho} = \frac{a\delta\lambda}{l}. \tag{15}$$

Combining equations (14) and (15) to find the maximum $a\delta\lambda$ for a particular a, given as

$$a\delta\lambda_{\text{max}}(a) = \frac{aw(a-\rho)}{4f\rho} \tag{16}$$

Note that equation (16) results in a $a\delta\lambda$ of 32 km for a semi-major axis of 41 km. This is much larger than a typical DSS would use for navigation or gravity recovery.[39,40] Therefore, instead of varying $a\delta\lambda$ with a, the $a\delta\lambda$ appropriate for a 23 km orbit, which has a $a\delta\lambda_{\text{max}}$ of 8.15 km, is used for all DSS simulations. The $a\delta\lambda$'s used are $[0.1, 3, 6, 9, 12]$ km, which is on a similar spacecraft separation scale to that of the GRACE[39] and GRAIL[40] missions. The largest separations, 9 and 12 km, are included to verify the $a\delta\lambda_{\text{max}}$ approximation, at least for the 23 km orbit.

25 multi-spacecraft simulations were completed, one for each combination of $a\delta\lambda$ and a. The simulations were run for two orbits but the landmarks were only matched between the images taken simultaneously by each spacecraft (sequential tracking). The optical tracking matches were quantified using the matches from ray-tracing, just as with the single spacecraft simulation. A linear triangulation method[18] was used with the known spacecraft state, known camera parameters, and matched 2D landmark points to produce the stereo-vision 3D point estimate for each landmark. The error in the stereo-vision 3D estimates was quantified using ray-tracing.

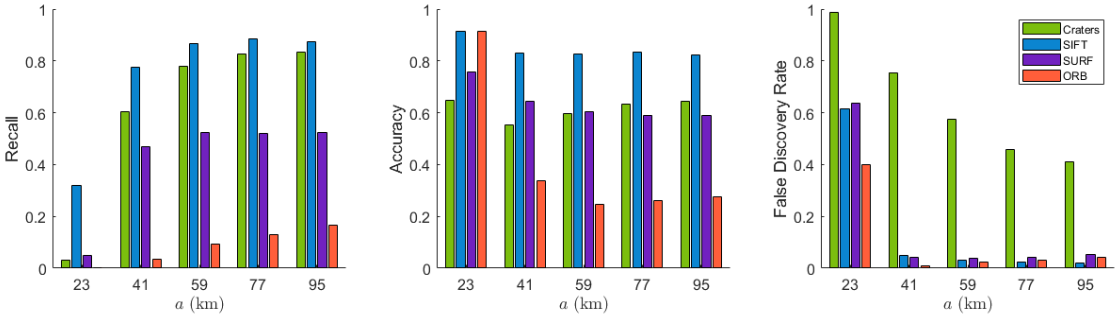

Figure 8. Binary classifier rates for all multi-agent simulations with a $a\delta\lambda = 6$ km using matches of ray-traced points as the truth data. Recall is low and false discovery rate is high for the 23 km orbit because the spacecraft separation is too large for the orbit. Otherwise, all optical tracking methods are quite consistent across changes in semi-major axis and show the similar trends to sequential tracking with the fastest imaging rate (Figure 3).

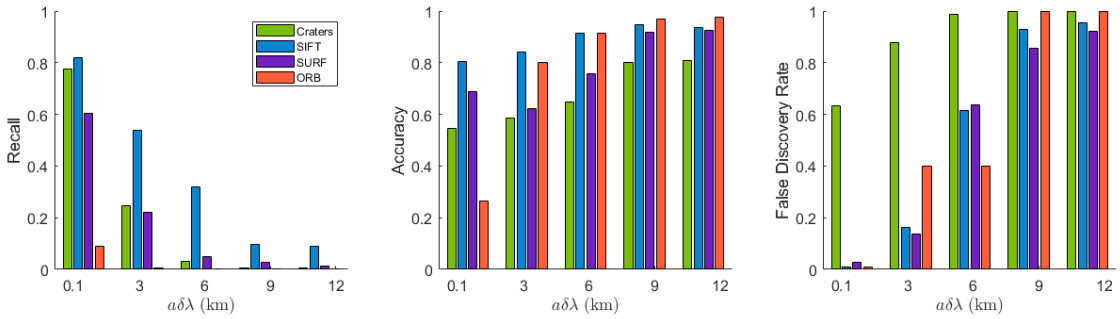

Figure 9. Binary classifier rates for all multi-agent simulations with $a = 23km$ using matches of ray-traced points as the truth data. Equation (16) proved to not be conservative enough to obtain a reliable upper bound on spacecraft separation.

Looking at Figure 8, which shows the performance across all semi-major axes for a separation of 6 km, the 23 km orbit stands out as far less reliable than the other orbits. This is because Equation (16) was not actually conservative enough to maintain the correct perspective, which is seen in Figure 9 by the sharp decrease in accuracy above a 0.1 km separation. The maximum radius of the asteroid should have been used instead of the average radius. Eros has an oblong shape with an average radius of 8.42 km and a maximum radius of 17.2 km.[41] This creates significant perspective shifts when the spacecraft observes Eros down its longest axis. Using the maximum radius in Equation (16) results in a $a\delta\lambda_{\text{max}}$ of 0.6 km for a 23 km orbit, which does correspond well to the data trend seen in Figure 9.

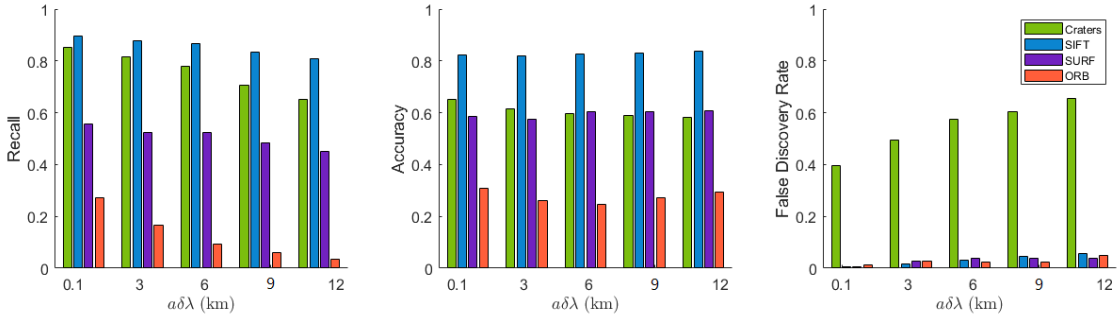

Figure 10. Binary classifier rates for all multi-agent simulations with $a = 59km$ using matches of ray-traced points as the truth data. All four optical tracking methods were consistent across all spacecraft separation tested and were equivalent to sequential tracking with the fastest imaging rate (Figure 3).

The other semi-major axes are all well under their respective $a\delta\lambda_{\text{max}}$ so the rest of this discussion focuses on Figures 8 and 10 and ignores the 23 km case. Figure 10 shows the performance of the 59 km orbit across all spacecraft separations. SIFT (blue) is consistently the best performing across all metrics. SIFT is followed by craters (green) and then SURF (purple), which correlates to the Test Case 2 simulation with the fastest imaging rate (Figure 3). This means that as long as the spacecraft abide by Equation (16) (subject to its assumptions) and use the maximum asteroid radius, SIFT and SURF are reliable optical tracking methods. Craters and ORB, however, retain the same troubles as they did in Test Cases 2 and 3.

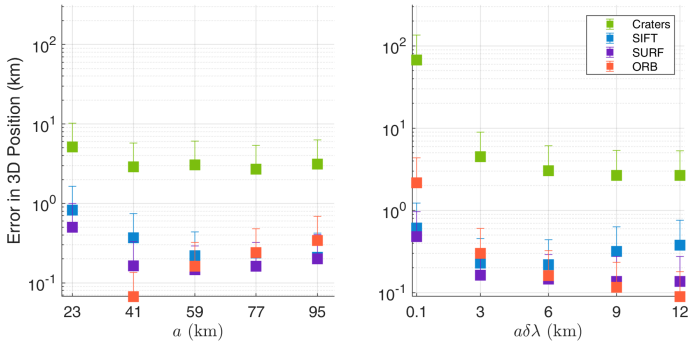

Figure 11. The mean difference between the 3D points calculated from ray-tracing and stereo-vision for each optical tracking method. The error bar for one standard deviation from the mean is included. Error produced by craters is too high to be used but keypoint descriptors can be reliably used for 3D mapping.

Finally, Figure 11 shows the mean error in stereo-vision calculated as the Euclidean difference between the ray-traced 3D points and the 3D points from stereo-vision. The error when using craters is above 1 km in all cases. Recall that the average radius of craters on Eros is 0.5 km.[25] This is quite significant and indicates that crater detection or NCC is too erroneous to use for 3D point recovery from stereo-vision.

Additionally, the error when using keypoint descriptors is around 0.25 km. While SIFT may have had the advantage for optical tracking, SURF outperforms SIFT with an error closer to 0.15 km. If the error is smaller, that means the centers of the landmarks in the image frame were more consistent between images. Thus, this shows that SURF is more consistent than SIFT in picking up the same landmarks in different images, even if SIFT may be better at matching landmarks.

What does this mean for 3D mapping performance? The best case error is 0.15 km. Structure from motion would definitely produce 3D points that have a larger error than this because it is using a lot less information to produce the 3D estimate.[18] SLAM, on the other hand, could press this lower bound depending on what the 3D point is initialized to because SLAM is designed to refine 3D points for mapping and navigation.[15]

In summary of multi-spacecraft optical tracking and 3D mapping, SIFT is the most consistent optical tracking method, followed by SURF. Crater detection and NCC is unreliable for both optical tracking and for 3D mapping. SURF performs better than SIFT for 3D mapping. Equation (16) is a reliable way to constrain multi-spacecraft configuration but needs to be modified from the original derivation.

CONCLUSIONS

This paper identifies and addresses three gaps in the field of autonomous optical tracking for use in asteroid missions. Addressing these gaps will allow missions to small celestial bodies to rely less on the Deep Space Network and reduce the amount of hardware (and, thus, resource constraints) on-board the spacecraft. First, target-specific optical tracking has not been compared to keypoint descriptors for use in navigation and mapping algorithms in asteroid missions. To address this, one target-specific optical tracking method (matching craters[26] via normalized cross-correlation (NCC)) is compared to three keypoint descriptors: Scale-Invariant Feature Transform[22] (SIFT), Speeded-Up Robust Features[23] (SURF), and Oriented FAST and Rotated BRIEF (ORB).[24] Second, current

optical tracking validation methods either do not support rigorous simulation[16] or are designed for densely-packed landmarks.[13] This paper introduces a novel correlation-validation method that uses the 3D points from ray-tracing the landmarks through the known asteroid model as the groundtruth, which is best for sparse landmarks in simulated images. Third, multi-agent stereo-vision has not been studied in orbits of small celestial bodies. Each optical tracking method in this paper is evaluated for its performance in multi-agent stereo-vision for 3D mapping.

Four test cases are performed using all four optical tracking methods to evaluate lighting-invariance, robustness to single- and multi-spacecraft orbit parameters, performance in multi-agent stereo-vision for 3D mapping, and continuity between sequential-tracking in simulated and real images. Detecting craters and matching them via NCC was inconsistent across all tests. This method is highly sensitive to orbit parameters, could not reliably identify correct matches, and produces errors in stereo-vision-recovered 3D points that are an order of magnitude larger than the average crater size. Crater correlation is more commonly performed using conic invariants[42] or pattern recognition[27] and it would be useful to compare one of those methods to keypoints in the future.

The three keypoint descriptors perform as expected according to Takeishi et al.[13] SIFT is the most accurate and consistent across changes in orbit parameters. However, SIFT is closely followed by SURF, and SURF out-performs SIFT in terms of lighting-invariance and performance in multi-agent stereo-vision. ORB did not perform well, most likely because it is a corner detector being used on a natural scene. Therefore, if the user is willing to constrain orbit parameters to minimize perspective changes between images, SURF is the recommended optical tracking method for 3D mapping. Otherwise, SIFT is the recommended method out of all four optical tracking methods tested in this paper.

The error in accepted Euclidean distance between ray-traced points is consistently below one tenth the size of the landmark with very small standard deviation. Manual correlation-validation in real images follows the same trends as ray-trace validation in simulated imagery, further validating the ray-trace method. Therefore, ray-tracing is a robust method to validate optical tracking in simulated images. Future developments could match the 3D model to real images[13] to extend ray-tracing validation to non-simulated imagery.

In conclusion, this paper presents valuable advances for the use and validation of autonomous optical tracking methods, enabling their use in future asteroid missions. This work portrays the potential of easy-to-use keypoint descriptors for navigation and 3D mapping in contrast to an industry standard of target-specific landmark tracking. Finally, the novel validation method presented in this paper paves the way for future studies to evaluate new optical tracking techniques.

ACKNOWLEDGEMENTS

This research is part of the Autonomous Nanosatellite Swarming (ANS) Using Radio-Frequency and Optical Navigation project supported by the NASA Small Spacecraft Technology Program cooperative agreement number 80NSSC18M0058.

REFERENCES

[1] A. Coradini, D. Turrini, C. Federico, and G. Magni, "Vesta and ceres: Crossing the history of the solar system," *Space Science Reviews*, Vol. 163, Dec. 2011, pp. 25–40. Publisher: Springer Netherlands ISBN: 9781461449034, 10.1007/978-1-4614-4903-4.

[2] Y. Takahashi, *Gravity Field Characterization around Small Bodies*. PhD thesis, University of Colorado at Boulder, 2013.

[3] A. M. Hein, M. Saidani, and H. Tollu, "Exploring Potential Environmental Benefits of Asteroid Mining," Oct. 2018.

[4] D. D. Mazanek, R. G. Merrill, J. R. Brophy, and R. P. Mueller, "Asteroid Redirect Mission concept: A bold approach for utilizing space resources," *Acta Astronautica*, Vol. 117, Dec. 2015, pp. 163–171. Publisher: Pergamon, 10.1016/J.ACTAASTRO.2015.06.018.

[5] A. F. Cheng, A. S. Rivkin, P. Michel, J. Atchison, O. Barnouin, L. Benner, N. L. Chabot, C. Ernst, E. G. Fahnestock, M. Kueppers, P. Pravec, E. Rainey, D. C. Richardson, A. M. Stickle, and C. Thomas, "AIDA DART asteroid deflection test: Planetary defense and science objectives," *Planetary and Space Science*, Vol. 157, Aug. 2018, pp. 104–115, 10.1016/j.pss.2018.02.015.

[6] D. A. Lorenz, R. Olds, A. May, C. Mario, M. E. Perry, E. E. Palmer, and M. Daly, "Lessons learned from OSIRIS-REx autonomous navigation using natural feature tracking," *2017 IEEE Aerospace Conference*, Big Sky, MT, USA, IEEE, Mar. 2017, pp. 1–12, 10.1109/AERO.2017.7943684.

[7] S. i. Watanabe, Y. Tsuda, M. Yoshikawa, S. Tanaka, T. Saiki, and S. Nakazawa, "Hayabusa2 Mission Overview," *Space Science Reviews*, Vol. 208, July 2017, pp. 3–16. Publisher: Springer Netherlands, 10.1007/s11214-017-0377-1.

[8] W. M. Owen, T. C. Wang, A. Harch, M. Bell, and C. Peterson, "Near optical navigation at Eros," 2001. ISSN: 00653438.

[9] N. Mastrodemos, B. Rush, A. Vaughan, and O. William Jr, "Optical Navigation for the Dawn mission at Vesta," Oct. 2012.

[10] C. Lippe and S. D'Amico, "Spacecraft swarm dynamics and control about asteroids," *Advances in Space Research*, July 2020, 10.1016/j.asr.2020.06.037.

[11] T. Guffanti and S. D'Amico, "Integration Constants as State Variables for Optimal Path Planning," *2018 European Control Conference (ECC)*, Limassol, IEEE, June 2018, pp. 1–6, 10.23919/ECC.2018.8550448.

[12] N. Stacey and S. D'Amico, "Autonomous Swarming for Simultaneous Navigation and Asteroid Characterization," Snowbird, UT, Aug. 2018.

[13] N. Takeishi, A. Tanimoto, T. Yairi, Y. Tsuda, F. Terui, N. Ogawa, and Y. Mimasu, "Evaluation of Interest-region Detectors and Descriptors for Automatic Landmark Tracking on Asteroids," *TRANSACTIONS OF THE JAPAN SOCIETY FOR AERONAUTICAL AND SPACE SCIENCES*, Vol. 58, No. 1, 2015, pp. 45–53, 10.2322/tjsass.58.45.

[14] Y. Cheng and J. K. Miller, "Autonomous landmark based spacecraft navigation system," Feb. 2003. Accepted: 2004-09-16T21:13:33Z.

[15] R. Siegwart, I. Nourbakhsh, and D. Scaramuzza, *Introduction to Autonomous Mobile Robots*. Cambridge: Massachusetts Institute of Technology, 2nd ed., 2004.

[16] E. Karami, S. Prasad, and M. Shehata, "Image Matching Using SIFT, SURF, BRIEF and ORB: Performance Comparison for Distorted Images," St. John's, Canada, Nov. 2015.

[17] P. Schmuck, "Multi-UAV collaborative monocular SLAM," *Proceedings - IEEE International Conference on Robotics and Automation*, 2017, pp. 3863–3870. ISBN: 9781509046331, 10.1109/ICRA.2017.7989445.

[18] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge: Cambridge University Press, 2004.

[19] M. J. Schuster, K. Schmid, C. Brand, and M. Beetz, "Distributed stereo vision-based 6D localization and mapping for multi-robot teams," *Journal of Field Robotics*, Vol. 36, Mar. 2019, pp. 305–332. Publisher: John Wiley & Sons, Ltd, 10.1002/rob.21812.

[20] K. T. Song, C. Y. Tsai, and C. H. C. Huang, "Multi-robot cooperative sensing and localization," *Proceedings of the IEEE International Conference on Automation and Logistics, ICAL 2008*, IEEE, Sept. 2008, pp. 431–436, 10.1109/ICAL.2008.4636190.

[21] E. Speyerer and M. Robinson, "NEAR Shoemaker Image Database," 2003.

[22] D. G. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," *International Journal of Computer Vision*, Vol. 60, Nov. 2004, pp. 91–110, 10.1023/B:VISI.0000029664.99615.94.

[23] H. Bay, T. Tuytelaars, and L. Van Gool, "SURF: Speeded Up Robust Features," *Computer Vision – ECCV 2006*, Vol. 3951, Berlin, Heidelberg, Springer Berlin Heidelberg, 2006, pp. 404–417. Series Title: Lecture Notes in Computer Science, 10.1007/1174402332.

[24] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," *2011 International Conference on Computer Vision*, Barcelona, Spain, IEEE, Nov. 2011, pp. 2564–2571, 10.1109/ICCV.2011.6126544.

[25] C. R. Chapman, W. J. Merline, P. C. Thomas, J. Joseph, A. F. Cheng, and N. Izenberg, "Impact History of Eros: Craters and Boulders," *Icarus*, Vol. 155, Jan. 2002, pp. 104–118, 10.1006/icar.2001.6744.

[26] B. Maass, H. Kruger, and S. Theil, "An edge-free, scale-, pose- and illumination-invariant approach to crater detection for spacecraft navigation," IEEE, 2011, pp. 603–608.

[27] W. Park, Y. Jung, H. Bang, and J. Ahn, "Robust Crater Triangle Matching Algorithm for Planetary Landing Navigation," *Journal of Guidance, Control, and Dynamics*, Vol. 42, Feb. 2019, pp. 402–410, 10.2514/1.G003400.

[28] Yang Cheng and A. Ansar, "Landmark Based Position Estimation for Pinpoint Landing on Mars," *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, Barcelona, Spain, IEEE, 2005, pp. 1573–1578, 10.1109/ROBOT.2005.1570338.

[29] K. Kim, Y. Jung, and H. Bang, "Projective invariant based crater matching for visual navigation in planetary landing," *AIAA SPACE 2016*, Long Beach, California, American Institute of Aeronautics and Astronautics, Sept. 2016, 10.2514/6.2016-5531.

[30] A. Cunningham, K. M. Wurm, W. Burgard, and F. Dellaert, "Fully distributed scalable smoothing and mapping with robust multi-robot data association," *Proceedings - IEEE International Conference on Robotics and Automation*, 2012, pp. 1093–1100. Publisher: IEEE ISBN: 9781467314039, 10.1109/ICRA.2012.6225356.

[31] B. Iglewicz and D. Hoaglin, "Volume 16: How to Detect and Handle Outliers," *The ASQC Basic References in Quality Control: Statistical Techniques*, 1993.

[32] T. Möller and B. Trumbore, "Fast, Minimum Storage Ray-Triangle Intersection," *Journal of Graphics Tools*, Vol. 2, Jan. 1997, pp. 21–28, 10.1080/10867651.1997.10487468.

[33] G. Neumann, J. D. Giorgini, J. K. Miller, A. Konopliv, and E. Grayzeck, "NEAR EROS NLR DERIVED PRODUCTS - SHAPE MODEL V1.0, NEAR-A-NLR-5-EROS/SHAPE/GRAVITY-V1.0," 2001.

[34] GOMSpace, "NanoCam C1U Datasheet," 2016.

[35] W. S. Cleveland and S. J. Devlin, "Locally Weighted Regression: An Approach to Regression Analysis by Local Fitting," *Journal of the American Statistical Association*, Vol. 83, Sept. 1988, pp. 596–610. Publisher: Taylor & Francis _eprint: https://www.tandfonline.com/doi/pdf/10.1080/01621459.1988.10478639, 10.1080/01621459.1988.10478639.

[36] B. A. Archinal, M. F. A'Hearn, E. Bowell, A. Conrad, G. J. Consolmagno, R. Courtin, T. Fukushima, D. Hestroffer, J. L. Hilton, G. A. Krasinsky, G. Neumann, J. Oberst, P. K. Seidelmann, P. Stooke, D. J. Tholen, P. C. Thomas, and I. P. Williams, "Report of the IAU Working Group on Cartographic Coordinates and Rotational Elements: 2009," *Celestial Mechanics and Dynamical Astronomy*, Vol. 109, No. 2, 2011, pp. 101–135. ISBN: 0923-2958 1572-9478, 10.1007/s10569-010-9320-4.

[37] M. Levoy, "Optics I: lenses and apertures," tech. rep., 2010.

[38] D. J. Scheeres, B. G. Williams, and J. K. Miller, "Evaluation of the Dynamic Environment of an Asteroid: Applications to 433 Eros," *Journal of Guidance, Control, and Dynamics*, Vol. 23, No. 3, 2000, p. 16, 10.2514/2.4552.

[39] 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. _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2004GL019920, https://doi.org/10.1029/2004GL019920.

[40] M. T. Zuber, D. E. Smith, M. M. Watkins, S. W. Asmar, A. S. Konopliv, F. G. Lemoine, H. J. Melosh, G. A. Neumann, R. J. Phillips, S. C. Solomon, M. A. Wieczorek, J. G. Williams, S. J. Goossens, G. Kruizinga, E. Mazarico, R. S. Park, and D.-N. Yuan, "Gravity Field of the Moon from the Gravity Recovery and Interior Laboratory (GRAIL) Mission," *Science*, Vol. 339, Feb. 2013, pp. 668–671, 10.1126/science.1231507.

[41] D. K. Yeomans, J.-P. Barriot, D. W. Dunham, R. W. Farquhar, J. D. Giorgini, C. E. Helfrich, A. S. Konopliv, J. V. McAdams, J. K. Miller, W. M. Owen, D. J. Scheeres, S. P. Synnott, and B. G. Williams, "Estimating the Mass of Asteroid 253 Mathilde from Tracking Data During the NEAR Flyby," *Science*, Vol. 278, Dec. 1997, pp. 2106–2109, 10.1126/science.278.5346.2106.

[42] J. A. Christian, H. Derksen, and R. Watkins, "Lunar Crater Identification in Digital Images," *arXiv:2009.01228 [cs]*, Sept. 2020. arXiv: 2009.01228.