

OPTICAL NAVIGATION FOR AUTONOMOUS APPROACH OF UNEXPLORED SMALL BODIES

Jacopo Villa*, Saptarshi Bandyopadhyay†, Benjamin Morrell‡, Benjamin Hockman‡, Daniel Lubey‡, Alexei Harvard‡, Soon-Jo Chung‡, Shyam Bhaskaran† and Issa A. Nesnas†

State of the practice in navigation toward and around small celestial bodies heavily relies on ground support and human skill, in particular for perception-based operations such as optical navigation and mapping. This leads to longer and more complex mission operations and subsequently higher cost. Furthermore, it imposes limitations for certain missions such as fast fly-bys or multi-agent operations. In this work, we present an autonomous navigation strategy for approaching unexplored small bodies. During the approach, we estimate the body’s physical properties as well as the spacecraft relative trajectory and associated uncertainties. This strategy, which is solely based on optical measurements, begins as soon as the body becomes resolved in the navigation camera and terminates at the start of proximity operations, when the spacecraft makes its first trajectory correction to stay in the vicinity of the body. Our multi-phased approach uses light curve analysis for estimating the body’s rotation rate, Shape-from-Silhouette techniques to reconstruct its 3D shape and estimate its rotation pole, and feature tracking tailored to small-body images for estimating relative navigation parameters. We used the Mission Analysis, Operations, and Navigation Toolkit Environment (MONTE), developed by the Jet Propulsion Laboratory and synthetic images to evaluate the feasibility and performance of the algorithms. As a case study, we reproduce the approach phase of the Rosetta mission. This work is based on the assumption that the spacecraft attitude is known, the body rotates about its principal axis of inertia and *a-priori* estimates of ephemerides and body size are available. Preliminary results show orbit determination performance on par with ground-based navigation from the Rosetta mission; albeit a bias in spacecraft position estimate is observed. This systematic error is likely due to the correlation between surface feature detection and fast-changing lighting and perspective conditions.

INTRODUCTION

There is a growing interest in small bodies, evidenced by recent scientific successes from NASA’s OSIRIS-REx and JAXA’s Hayabusa2 missions, as well as contributions from past missions such as Rosetta and NEAR Shoemaker.¹⁻⁴ In addition to pushing the boundaries of planetary science, these celestial bodies are being targeted in the future to assess planetary defense strategies (e.g. through DART, a mission planned to intercept the binary asteroid system Didymos in 2022)⁵ and are seen as fruitful loci for human exploration and In-Situ Resource Utilization.⁶

Furthermore, the diversity, abundance, and relative accessibility of small bodies make them a compelling target to advance autonomy for space exploration: they offer an adequately challenging environment with a unique balance between an *a-priori* unknown environment and one where the interaction between the spacecraft and the body is dynamic, albeit more forgiving in nature, due to its low gravity.⁶ Learning to autonomously reach and operate in such environments provides a stepping stone toward more complex autonomous missions.

Missions toward small bodies can generally be divided into three phases: cruise, approach and either fly-by or proximity operations. We define the beginning of the approach phase (and hence the end of the cruise)

*KTH Royal Institute of Technology, Brinellvägen 8, 114 28 Stockholm, Sweden

†Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

‡Graduate Aerospace Laboratories of the California Institute of Technology, 1200 E California Blvd., Pasadena, CA 91125, USA

as the moment when the target body becomes resolved by the onboard camera and the end of the approach phase when the spacecraft makes its first trajectory correction to remain in the vicinity of the body, which marks the start of proximity operations. This paper focuses on the approach, particularly on autonomous navigation strategies aimed at fulfilling the most representative requirements related to this mission phase.

When approaching a small body, it is critical to recover an early spacecraft-target relative trajectory estimate and plan potential trajectory correction maneuvers beforehand. Lack of such a preliminary knowledge can yield to severe consequences for the mission, including missing nominal orbit insertion for proximity operations. As such, standard navigation strategies are based on the intensive use of both radiometric tracking measurements, performed from ground stations, and optical measurements of the target body taken by the onboard camera.⁷ Furthermore, shape modeling is typically based on techniques requiring manual selection and tracking of fixed surface landmarks, such as *Stereo Photoclinometry* (SPC).⁸ In the current human-intensive process, different teams of experts must carefully coordinate and continuously iterate as the estimated parameters are strongly coupled. This process drives the complexity of mission operations, imposes limitations on what can be achieved, given the coupling and time-delayed information and increases overall mission cost.

In this paper, we present an optical-based autonomous orbit determination (OD) strategy suitable for the approach phase of an unexplored small body, which could, in principle, run on board without the need for neither human intervention nor radiometric tracking data. We start the estimation process as early as the body becomes resolved in the images of the navigation camera, i.e. at subpixel resolution. In the context of autonomous exploration missions, previous work has focused on cruise phase navigation, small-body proximity operations and landing.⁹⁻¹² However, the approach phase of small body missions is characterized by distinctive challenges, which makes it a unique and separate problem to tackle.

In the following sections, we first define the problem statement. Second, we outline the simulation setup that we used in this study. We then provide an overview of our autonomous approach strategy, followed by a description of all its phases. Finally, we evaluate orbit determination performance to validate the strategy and conclude with a discussion of these results.

PROBLEM STATEMENT

Considering a spacecraft approaching a small body, we aim to estimate navigation parameters with uncertainties comparable to the state of the practice, solely using images acquired by an onboard camera. These parameters are: (1) the spacecraft position and velocity relative to the target, (2) the target body shape, pole, and rotation rate, (3) landmark positions onto the body surface, (4) a scale factor denoting the body size (5) the body's gravitational parameter μ and (6) planned Orbital Control Maneuvers (OCM) for the spacecraft.^{7,13} This process is based on the following assumptions:

1. The spacecraft attitude is known, typically from star trackers, and the camera boresight points toward the body's center of rotation. Centroid estimation is a well-studied problem in optical navigation.¹⁰
2. The target body has a stable spin about its principal axis of inertia. To date, this assumption has held for small bodies visited by spacecrafts.
3. The images are of good quality, have no lens distortion and no pixel noise. In practice, we would have to filter out image artifacts and use image calibration to correct for lens distortion.^{13,14}
4. *A-priori* estimates of spacecraft and target ephemerides, e.g. from cruise-phase knowledge, as well as priors on the body scale, which usually come from ground-based (or space telescopes) albedo measurements performed prior to the mission, are available.¹⁵
5. For shape estimation, the scale is assumed to be known. Knowledge of absolute scale is not necessary for this purpose; the relative scale factor suffices. However, errors in relative scale estimation, e.g. due to image pixelation, are beyond the scope of this paper.

While the above represents a complete list of the assumptions made in support of the whole approach scenario, not all of them are necessary for the whole approach, as denoted in the following sections.

APPROACH STRATEGY OVERVIEW

During the approach, the size of the target in the image grows from occupying less than a single pixel at the beginning of the phase to occupying the entire image when in the proximity of the body. As such, we use

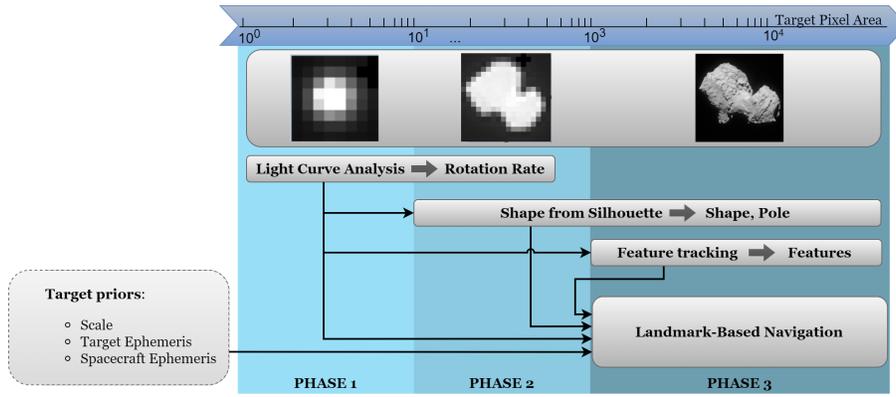


Figure 1: A phased strategy with overlapping phases; priors from each phase feed into the next one to estimate navigation parameters during the small body approach.

an approach strategy divided into phases, so to match the resolution of the imaged body, as shown in Figure 1. In Phase 1, the body’s visible area spans from subpixel resolution to tens of pixels. Here, we estimate the body’s rotation rate using a standard light curve analysis method. In Phase 2, where the visible area goes from tens to thousands of pixels, we jointly estimate its pole orientation and shape. In this phase, the body’s limb has distinctive characteristics. We use a *Shape-from-Silhouette* (SfS) algorithm to estimate pole and shape based on the rotation rate knowledge from Phase 1. In Phase 3, the visible area is typically over thousands of pixels. In this phase, we identify and track surface landmarks for precise relative navigation. These optical measurements are then processed by a navigation filter to determine the relative orbit. The estimates from Phases 1 and 2 initialize the *a-priori* dynamics model, which is used in Phase 3 for navigation.

It is worth noting that we do not assess full-state estimation throughout Phases 1 and 2, whereby landmarks are not visible yet. Here, we assume that we can use legacy autonomous-optical-navigation techniques, based on center finding, such as *AutoNav*.¹⁰

SIMULATION SETUP

In order to test and validate our estimation pipeline, we constructed a simulation environment to reproduce realistic approach scenarios, evaluate performance and compare results obtained with reference data.

Reference Mission

The reference data selected for this study are from Rosetta, the European Space Agency mission that visited comet 67P Churyumov-Gerasimenko between August 2014 and September 2016.¹⁶ A rich dataset is available from this mission, which we used to reconstruct the Rosetta’s approach scenario. In particular, we employed data of the NAIF SPICE kernel repository to reconstruct spacecraft and comet trajectories and attitudes, coordinate systems and clock data.¹⁷ These were used to compose the ground truth for our simulations. Furthermore, the bilobed shape of comet 67P makes it a particularly challenging and interesting target for autonomous navigation, in terms of shape modeling, landmarks estimation and lighting conditions.

Image Rendering

To test an optical-based navigation pipeline, it is key to generate realistic images to reproduce those acquired by the onboard camera. While the use of images from the Rosetta mission would ensure realism, the existing dataset would make it difficult to test different operational scenarios, such as varying sampling frequency, pointing, and image specifications. Furthermore, operations’ sampling frequency is limited by ground-in-the-loop requirements. When more autonomy is enabled, images can be acquired and processed more frequently so to improve the performance of several image processing algorithms. Given these considerations, we use high-fidelity synthetic images to test our navigation pipeline. For rendering, we reproduce camera and target body poses as well as sunlight vector from SPICE kernels data. We then use a state-of-the-art shape model of comet 67P¹⁸ for surface triangulation and Rosetta’s navigation camera (NAVCAM) specifications⁷ to reproduce the camera calibration matrix. We finally employ *Blender Cycles* as a rendering

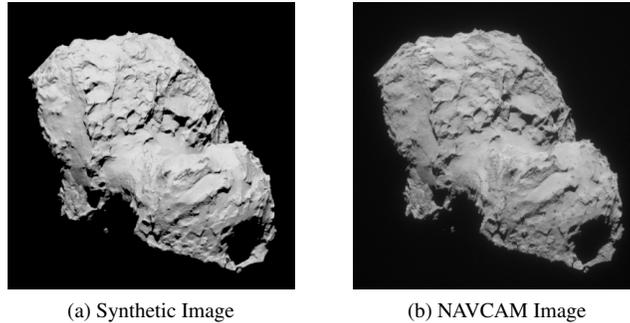


Figure 2: Comparison between a synthetic and a real image, taken at a distance of 145 km. One key difference is given by surface albedo: a constant albedo map is chosen for synthetic images.

engine, which is based on *ray tracing* of light, to obtain realistic lighting conditions considering multiple reflections from the surface.¹⁹ A comparison between a sample synthetic image and the corresponding real image* is shown in Figure 2. One key difference is that no albedo map of the comet is used in rendering the images and a constant albedo over the whole surface is assumed. Therefore, whilst synthetic images resemble true ones in terms of surface patterns from shading, they differ in photometric properties. Apart from an arguably minor effect on feature tracking, this difference does not present relevant implications on the other image processing algorithms utilized.

Approach Simulation

We use the *Mission Analysis, Operations, and Navigation Toolkit Environment* (MONTE), developed by the Jet Propulsion Laboratory, to simulate the approach scenario and perform orbit determination throughout it.²⁰ MONTE offers state-of-the-art astrodynamics and navigation libraries, which we use to model orbital dynamics and the optical navigation process; it is then used to set up and run the navigation filter, obtain orbit determination solutions, process measurements and analyze results.

PHASE 1: PERIODICITY ESTIMATION

As previously mentioned, the pipeline starts with the estimation of the target body’s rotation rate, which is observable at subpixel resolution. Periodicity estimation is a well-studied astrometric problem, involving the analysis of the so-called *light curve*.¹³ The light curve is, for a given set of images sorted by time, the evolution of image brightness (i.e. the sum of all pixel values) over time. For rotating bodies, this signal usually contains key information to recover the rotation rate. In this work, we rely on results obtained by Bandyopadhyay *et al.*, who studied autonomous periodicity estimation of comet 67P based on the Rosetta approach data.²¹ They use a *Fast Fourier Transform* (FFT) algorithm to estimate the rotation rate from a given set of images. This method can be fully automated and relies neither on the priors of spacecraft and target ephemerides nor on priors of the target scale; it is entirely based on optical data. Furthermore, the assumptions on perfectly known attitude and center-pointing boresight axis are not necessary, as long as the entirety of the target body is visible in the image: the light curve intensity mostly comes from the pixel cluster of the target body and can be reliably computed throughout most of the field of view.

Bandyopadhyay *et al.* successfully applied FFT to recover comet 67P’s rotation rate by using both real images from Rosetta and synthetic images in two different case studies. In the first case, which is the worst in terms of performance, the solution converges within 10 days from the beginning of the approach. This estimation time range is in line with periodicity estimation obtained during Rosetta operations (around 9 days for the longest observation case).¹³ The estimation error is less than 1% with the proposed sampling frequency and number of samples. While the said paper does not exhaustively evaluate the impact of different parameters (e.g. sampling frequency) on periodicity estimation uncertainty, the rotation rate provided by these results is adequately accurate to initialize subsequent approach phases.

*NAVCAM’s image downloaded from the Rosetta Image Archive: <https://imagearchives.esac.esa.int/picture.php?/6887/category/63>

PHASE 2: SHAPE AND POLE ESTIMATION

Once the body’s limb is sufficiently resolved to detect any traits (usually starting from tens-of-pixel in area), shape and pole orientation can be autonomously estimated as well. Although no distinctive surface landmarks for relative navigation are usually visible at this stage, the body’s silhouette and its evolution over time can be observed. We base our work on the *Shape-from-Silhouette* (SfS) algorithm presented by Forbes and on a modified version of the pole estimation method proposed by Bandyopadhyay *et al.* in the aforementioned paper.^{21,22} SfS is a well-known technique in shape reconstruction and has also been successfully employed during ground-based operations (e.g. during the Rosetta mission).^{23,24}

Shape Estimation

The process of shape reconstruction involves observing the target body from multiple points of view and extracting the body silhouette from each of them. In our case, this can be done by leveraging the natural rotation of the body during the approach. The algorithm uses a voxel-carving method to reconstruct the visual hull²⁵ starting from a set of observed silhouettes and known target-camera relative attitude*. In our case, the relative attitude is unknown *a priori*: while the body rotation rate is well-characterized at this stage, the rotation pole needs to be determined. In this section, we assume that the pole, and hence the relative attitude, is known. In the next section, we show how to relax this assumption and include the shape estimation method presented here within an attitude estimation process.

Here, we summarize the visual hull estimation routine. First, the silhouette is extracted from the image foreground. Our silhouette definition is conservative: it is the group of pixels with a non-zero value. Second, a set of volumetric elements (i.e. *voxels*) representing the discretized initial guess for the target shape is defined and oriented according to the known relative attitude. For each image in the considered set, voxels are then projected onto the camera plane and all voxels lying outside the silhouette for all images, i.e. always lying in the background, are erased from the shape volume. To resolve the status of voxels belonging to the image foreground only in some of the images, an octree-based search is performed. Voxel size is adaptive and increases with image resolution. Finally, the image sampling frequency is controlled to maximize the observability of the body silhouettes. It can be shown that a frequency in accord with Equation (1) allows an arbitrarily low visual hull estimation error, up to image pixelation limits. D_{px} is the body diameter in pixels, ω_B its rotation rate and δ_{px} the maximum visual hull estimation error, in pixels. K_S is a correction factor to account for body’s non-sphericity: $K_S = 1$ for a sphere, otherwise $K_S < 1$; we chose $K_S = 0.25$. Changes in approach latitude and body apparent size are considered negligible over a rotation-period time frame.

$$f_S = \left(2\cos^{-1}\left(\frac{D_{px}}{D_{px} + 2K_S\delta_{px}}\right) \right)^{-1} \cdot \omega_B \quad (1)$$

We simulated visual hull estimation for comet 67P by reproducing Rosetta’s relative trajectory, lighting conditions, and camera specifications. In this case study, the sunlight phase angle is non-zero for the entire approach.¹³ The simulation considers a trajectory arc spanning between approximately 10^4 km and 10^3 km away from the body. Over this time frame, the visual hull is estimated at 20 subsequent steps. For each visual hull estimation, an image set covering only one body’s full rotation, sampled with frequency f_S , is used; hence, these 20 shape estimation steps are independent of each other. As a shape modeling error metric, we compute the Hausdorff Distance between the 20 estimated visual hulls and the ground truth shape.

The evolution of shape estimation performance is captured in Figure 3 and the last estimated shape (at 826 km) is visualized in Figure 4. The Root Mean Square (RMS) of the Hausdorff Distance monotonically decreases during the approach. The initially high RMS value is due to both the low resolution of images and the strategy adopted for silhouette extraction, which leads to body oversizing especially at very low resolutions. Future work will assess more accurate silhouette extraction methods. Furthermore, the visual hull’s volume is, by definition, an overestimate of the actual shape, as local concavities cannot be captured by the SfS process. This is also evidenced by local errors (in dark red) in Figure 4.²⁵

*In fact, Forbes’ algorithm requires the relative pose as input rather than mere relative attitude. However, since we are not interested in estimating the target scale at this approach stage (as the latter comes from full-state orbit determination performed in Phase 3), we accept to use relative pose up to a scale factor, which is equal to relative attitude.

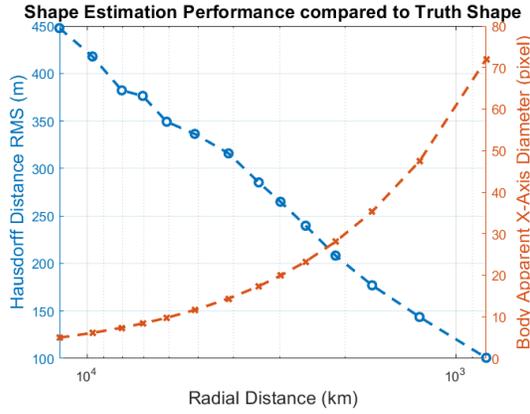


Figure 3: Statistics of Shape-from-Silhouette performance, compared to the X-axis body apparent diameter, in pixels.

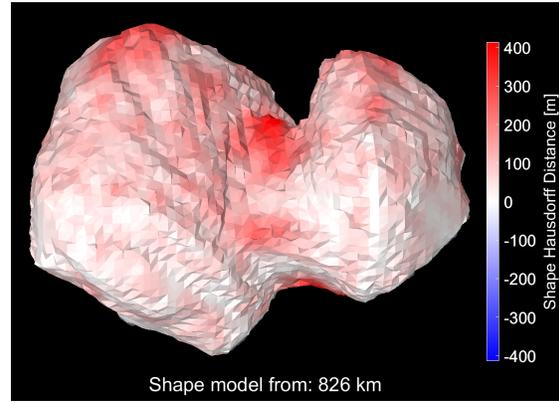


Figure 4: Reconstructed visual hull and error (i.e. Hausdorff Distance) distribution at the last simulation step; the distance from the body is 826 km.

Pole Estimation

Now we remove the above assumption that a spacecraft-target relative attitude is known. In the case of a stable spin about its principal axis of inertia, the target-body pose evolution is entirely given by its pole orientation and its rotational rate, as the body’s Prime Meridian¹⁴ (PM) can be arbitrarily defined. Since the rotation rate is well-characterized during Phase 1*, this problem reduces to estimating the pole orientation.

Here, we propose a strategy to recover the pole using the SfS-based convex hull. This process, which was introduced by Bandyopadhyay *et al.* in the aforementioned paper, is here summarized. A pole orientation hypothesis is defined and the estimated convex hull is rotated according to such a hypothesis. Then, for each frame, the hull is further rotated according to the angular velocity estimate provided by Phase 1. At every time step, the hull is projected into the camera plane, and its resulting silhouette is compared to the observed one. The overlap error ϵ between such visual hull silhouette (\mathcal{S}_h) and the silhouette observed in the image (\mathcal{S}_0) is defined by Equation (2), where \mathcal{A} denotes the area of its argument, in pixels. This is a reformulation of the overlap metric presented by Bandyopadhyay *et al.* The goal is to use this overlap as a measure of pole likelihood, for each hypothesized pole.

$$\epsilon(\mathcal{S}_h, \mathcal{S}_0) = \frac{\mathcal{A}(\mathcal{S}_0 \cup \mathcal{S}_h) - \mathcal{A}(\mathcal{S}_0 \cap \mathcal{S}_h)}{\mathcal{A}(\mathcal{S}_0)} \quad (2)$$

By definition, ϵ goes to zero while the pole hypothesis goes to the true one. Now, we claim that this overlap error is correlated with the likelihood of the corresponding pole hypotheses. For the purpose of providing a pole initial estimate, which will be refined during Phase 3, we can define the heuristic presented in Equation (3), where \mathcal{L} is the pole estimate likelihood, function of pole’s Right Ascension (α) and Declination (δ).

$$\mathcal{L}(\alpha, \delta) \sim \epsilon^{-1}(\mathcal{S}_h, \mathcal{S}_0) \quad (3)$$

This relation is used to merely evaluate relative likelihood, so to compare the different pole hypotheses analyzed. An example of pole likelihood distribution for comet 67P is shown in Figure 5a and 5b. This is performed at 826 km, aligned with results presented in Figure 4. We can empirically verify that the maximum value of \mathcal{L} is close to the truth pole orientation to within a two-degree error, at this resolution. This offset is considered sufficiently small to initialize the pole for subsequent landmark-based navigation. Ongoing work is underway for increasing the accuracy of pole estimated through this method as well as analytically formulating a relation between overlap error and pole likelihood, including related uncertainties. Finally, efficient methods to perform pole search are currently being explored.

*Given the precision of the rotation rate estimate obtained in Phase 1, the maximum phase error of the rotated convex hull is in the order of 10^{-3} degrees, which is considered negligible here.

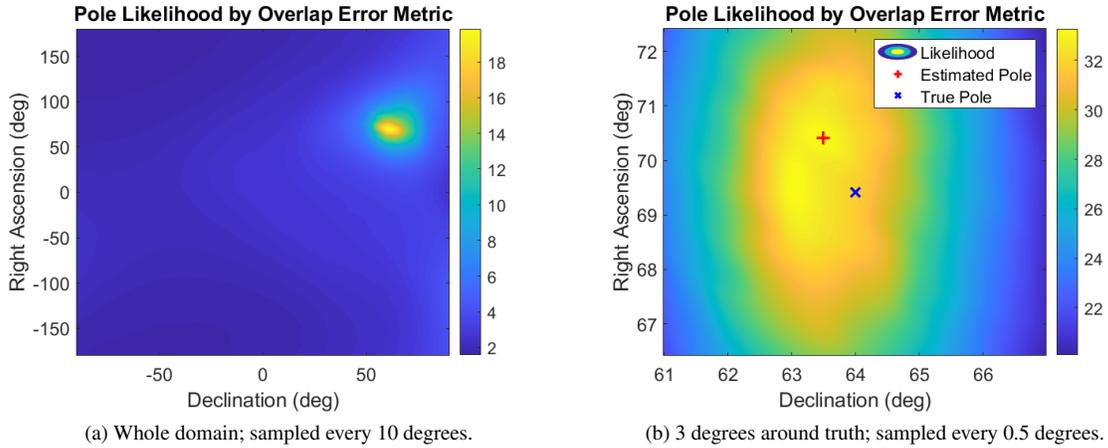


Figure 5: Pole likelihood interpolated data; evaluated at a distance of 826 km from the body.

PHASE 3: LANDMARK-BASED NAVIGATION

Phase 2 terminates when landmarks on the surface of the body become clearly observable, i.e. they are distinguishable from each other and trackable over multiple images. In this paper, we define a *landmark* as the 3D position of a reference point lying on the surface of the target body. We then define a *feature* as the reprojection of a landmark into the 2D camera plane. In other words, a feature is the observable of the corresponding landmark, from the camera perspective.

In Phase 3, landmarks are used for relative navigation. We adopted the process shown in Figure 6: first, surface features are detected and tracked in the image sequence and treated as optical measurements. Second, both a dynamics and a camera model are defined with available knowledge. Third, a preliminary OD solution from the navigation filter is computed. This is used for a-posteriori outliers rejection and hence for updating the measurement set, erasing low-quality feature tracks. Then, the updated measurements are fed into an iterative filtering process and a definitive OD solution is finally computed.

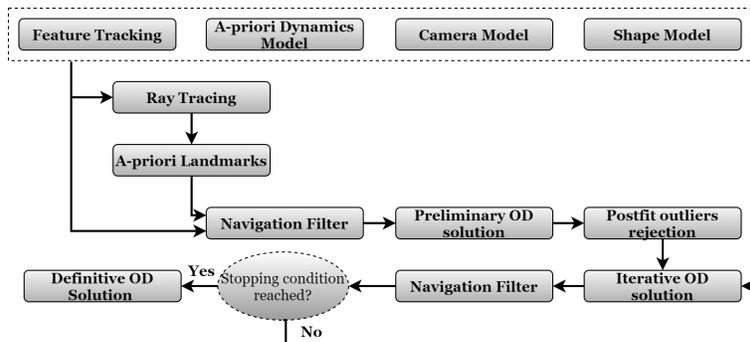


Figure 6: Autonomous Orbit Determination Process Schematic

Feature Tracking

Because our estimation strategy is based on images only, without any radiometric tracking triangulation, it is key to rely on good-quality optical measurements. Ideally, we aim to obtain measurements that are unbiased, have a Gaussian error distribution, and are observed over a relatively large number of images to increase estimation performance. Furthermore, the natural rotation of the body makes features alternate periods of visibility and eclipse (when features lie on the far side of the body), needing to recover them over multiple body rotations, i.e. performing *loop closure*. In seeking this, optical navigation is challenged by the surface appearance of a small body, which is subject to notable high-frequency changes, while observed by the approaching spacecraft camera: the body’s rotational motion with respect to the Sun and the observer

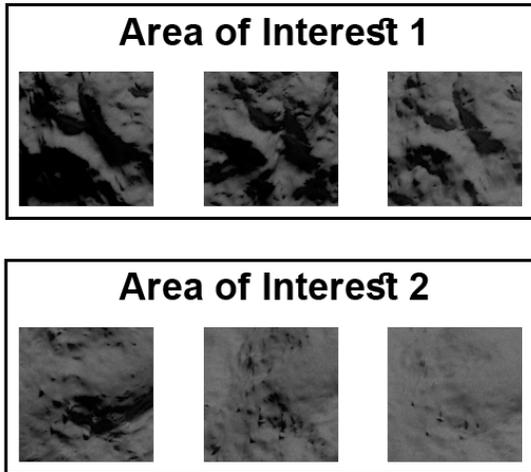


Figure 7: Image patches surrounding two surface features of a procedurally-generated small body, as seen under different lighting and perspective conditions. Such patches are used for CNN training.

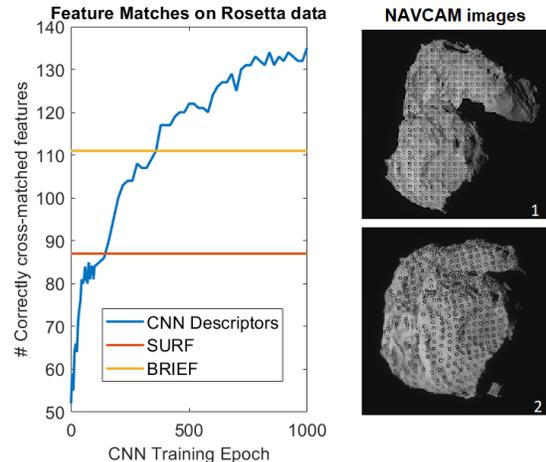


Figure 8: Correct matches on two NAVCAM’s images, from CNN, *SURF* and *BRIEF* descriptors (left). Features (grey dots) selected from image 1 and matched in image 2 from the CNN (right).

produces variations in both lighting and perspective conditions of the surface, as shown in Figure 7. These phenomena decrease the performance of many feature tracking algorithms. Most existing feature trackers are designed to work with terrestrial images and so were designed to detect and describe features that exist due to local albedo variations, i.e. sharp differences in gradients such as from corners or sudden changes in colors. State-of-the-art feature descriptors are also designed to work with these image types.²⁶ In particular, they can handle variations in gradient, as well as overall illumination changes, but can fail when other sharp changes are introduced, such as the presence of a moving shadow region (e.g. cast by a boulder on the surface) within the feature’s image patch. These kinds of events, which are typical under said dynamic lighting of a small body scenario, heavily bias the feature’s histogram, affecting matching performance.

To address these deficiencies, we trained a Convolutional Neural Network (CNN) to process the image patch surrounding a feature and produce a descriptor vector that is robust to changes in shadows, lighting, and surface affine transformations. We used a base *VGG-16* network architecture, reduced to take in 32×32 image patches, with the final tensor layer modified with a maximum activation layer along with L-2 normalization to output 64 floating values (the same size as a *SURF* descriptor vector).²⁶ More details on network architecture and training can be found on the work from Harvard *et al.* where a similar neural network was designed and trained using the same contrastive loss function.²⁷

The dataset for neural network training was produced by randomly generating small-body shapes with procedural methods for terrain generation. We then rendered images of these bodies, defined equally-spaced features (from the camera perspective) on their surface and observed image patches centered in such features, under different lighting and perspective conditions, defining positive and negative matches from ground truth knowledge. An example of such patches is shown in Figure 7. For validation, real Rosetta images were used, together with SPICE-based NAVCAM poses and ground-truth shape model. Regularly spaced features were sampled across images and description algorithms were run on them; resulting descriptors were then used to find matches through standard cross-matching. Since the features were generated with ray tracing, feature correspondences were known *a priori*, so the matching efficacy could be evaluated. It was observed that the CNN descriptors substantially outperformed all conventional descriptor algorithms we tested, even when using real data for cross-matching. An example of performance comparison between CNN-based, *SURF* and *BRIEF* descriptors is shown in Figure 8. The graph reports the number of correctly matched features between the two images from the Rosetta approach. CNN-based matching performance improves over training time and outperforms both *SURF* and *BRIEF* starting from about 400 training epochs.

For the feature detection process, *SIFT* is chosen as the best-performing candidate. One should note that, for reasons similar to the ones discussed before, feature detection (as well as description) presents correlation with dynamical lighting and perspective conditions and hence the whole feature tracking process is still affected by some deterministic error. For example, a feature detector may extract a feature in correspondence of a shadow, and keep detecting such a moving shadow in subsequent frames, albeit the corresponding physical point on the body surface is moving. However, this appears to have a smaller effect on tracking robustness and loop closure capabilities. Future work will address light- and perspective-robust detection as well.

Hence, the selected feature tracking strategy uses CNN-based description and *SIFT* detection. Furthermore, it includes a prefit outliers rejection step, by solving the standard *Perspective-n-Point* (PnP) problem through RANSAC, with a cutoff value of two pixels. In this way, all features farther than two pixels from estimated epipolar lines are filtered out. The latter process was implemented using *OpenGV*.²⁸ Finally, to increase matching performance, the rotation rate estimate from Phase 1 is utilized to predict the time frame where loop closure for a given feature is expected to happen (as shown in the scheme of Figure 1).

Dynamics and Camera Model

In this implementation, MONTE’s dynamics model accounts for two forces: (1) gravity from Solar System’s planets, Sun and target body, and (2) spacecraft OCMs. Given the limited time frame of the chosen simulation, other non-gravitational effects, such as Solar Radiation Pressure, were not considered; they will be included in future work. The Solar System’s gravity is based on planetary ephemerides *DE430*,¹⁷ the comet’s gravity on *a-priori* estimates (e.g. ground observations and density assumptions), whereas OCMs are based on *a-priori* data from Rosetta operations to reproduce realistic errors given by pre-mission knowledge.

An *a-priori* value for all estimated parameters is defined. Both spacecraft and target orbit priors are given. Spacecraft attitude is known, as stated in our assumptions. The target body is assumed to have a stable rotation about its principal inertia axis, whose rotation rate and pole orientation result from Phases 1 and 2, respectively. Landmark estimates are given by the following process. Given the spacecraft-target relative pose estimate, an observed feature and the camera model (described below), ray tracing is used to project the feature from the image onto the 3D shape model estimated from SfS, by finding the intersection between the ray (passing through the feature on the camera plane and the camera optical center) and said shape. This process can be performed regardless of the relative scale factor used, which is employed to determine the absolute landmark coordinates thereafter. The intersection represents the initialization for the observed landmark. This ray tracing process is repeated for all observed features. Including these priors, the dynamics model is completely defined. The dynamics is then propagated and used by the navigation filter.

A pinhole camera model is also defined to properly transform optical measurement coordinates from pixel to length unit and vice versa.¹⁴ For this purpose, Rosetta’s NAVCAM specifications are employed.⁷ The camera has a 5-degree field of view and a 1024×1024 resolution. The only difference from NAVCAM is that we assume images with zero optical distortion, as specified in the assumptions set.

Filter Setup

A batch least-squares filter is used to compute the orbit determination solution. We define three different types of estimated parameters: *dynamic*, *bias* and *consider*. The full setup is presented in Table 1. The filtering process includes a postfit outliers rejection step (described below) and multiple fit iterations (Figure 6). The stopping condition in the iterative process is given by Equation (4), where \mathcal{F}_{RMS} is the RMS of feature postfit residuals, i is the filter iteration index, and δ_{RMS} is a positive user-set tolerance.

$$|\mathcal{F}_{RMS}(i) - \mathcal{F}_{RMS}(i - 1)| \leq \delta_{RMS} \quad (4)$$

Postfit Outliers Rejection

After the first filter iteration, postfit measurement residuals are processed to detect any evident outliers. Here, we consider outliers as all features whose residuals exceed the $3\text{-}\sigma$ standard deviation at least once over the whole feature track. The idea behind this is that good-quality tracks should always be consistent from each other and with the estimated relative dynamics because they lie on the surface of the same rigid body.

Table 1: Setup and Results of OD Parameters ^a

Parameter Name	Type	Prior 1- σ	Posterior 1- σ	Units
S/C Position	Dynamic	[30,30,30]	[0.8,0.8,0.2]	km
S/C Velocity	Dynamic	[1,1,1]	[0.3,0.8,0.5]	cm/s
Comet Orbit	Consider	Full Set-III Lower Matrix	Not Estimated	N/A
Comet Scale	Bias	0.2	0.02	–
Pole RA,Dec	Bias	[2,2]	[0.2,0.05]	deg
Comet PM	Bias	2	0.8	deg
Comet Spin Rate	Bias	$1 \cdot 10^{-5}$	$3 \cdot 10^{-6}$	deg/sec
Comet μ	Bias	$2 \cdot 10^{-7}$	$2 \cdot 10^{-7}$	km ³ /s ²
OCM Δv	Bias	[1.6,12,11]	[1.1,0.9,0.6]	cm/s
OCM Epoch	Bias	2.0	1.2	s
Landm. Sample	Bias	[355,341,207]	[33,32,34]	m

Table 2: OD Simulation Setup

Parameter Name	Value
Start Epoch (t_0)	02-AUG-2014 23:21:22.812 UTC
End Epoch (t_f)	03-AUG-2014 16:02:22.812 UTC
OCM Name	CAT pre-insertion
Image Sampling Period	2.5 minutes
Filter Mapping Period	1 hour
Comet distance at t_0	397 km
Comet distance at t_f	278 km
Image Resolution at t_0	34 m/pixel
δ_{RMS}	$1 \cdot 10^{-2}$ pixel
Number of Landmarks	63 (before fit)

^aAll values are expressed in the EME2000 frame, except for landmarks that are in comet-fixed frame. Posterior 1- σ values are from one Monte Carlo sample.

Hence, residuals computed through the predicted dynamics should be relatively small. As such, relatively high residuals are likely a sign of low-performance tracking.

Orbit Determination Setup and Results

To evaluate orbit determination performance during Phase 3, given the priors obtained during Phases 1 and 2, we simulated part of the Rosetta approach phase in MONTE. Table 2 provides an overview of the simulation setup. We chose a trajectory that includes the transition between the *Far Approach Trajectory* (FAT) and the *Close Approach Trajectory* (CAT), including the *CAT-preinsertion* firing.⁷ In this way, we can also assess the effect of spacecraft maneuvers on orbit determination. Image resolution within this trajectory suffices to detect and track surface landmarks. The spacecraft state is estimated relative to the comet-centered *EME2000* inertial frame. The *a-priori* covariance values are reported in Table 1 and selected according to the following criteria. For spacecraft position and velocity, these are set in accordance with Rosetta’s navigation analysis ones at the same epoch; ours are set slightly higher to provide some generality.⁷ For comet’s Set-III orbital elements,²⁹ OCM and gravitational parameter μ , covariances originate from Rosetta operations data. For pole Right Ascension, Declination and Prime Meridian, these are defined based on the errors obtained from SFS during Phase 2, at 826 km of distance, whereas for comet’s rotation rate it is given by errors from light curve analysis in Phase 1.²¹ Landmarks covariance is defined as two times the Hausdorff Distance RMS in shape estimation, at 826 km, in boresight (i.e. radial) direction, whereas it is defined based on image pixelation error in the camera plane (i.e. transverse and normal) directions. This covariance is then rotated into the comet-fixed frame, in which landmarks are estimated. Scale’s covariance is 20%, to remain conservative on knowledge of the body size. Finally, a conservative 0.5-pixel value has been selected for the 1- σ uncertainty of optical measurements, based on former feature tracking performance testing.

We performed a preliminary Monte Carlo analysis with 5000 orbit determination trials. For each trial, all estimated parameters were randomly sampled from the gaussian distribution with mean and standard deviation equal to parameters’ ground truth and *a-priori* 1- σ definition, respectively. Converged postfit residuals for one of such samples are shown in Figure 9: residuals mean is $7 \cdot 10^{-6}$ pixel, whereas standard deviation and RMS are approximately 0.22 pixel. These statistics are in line with maplet tracking performance of Rosetta operations, which present subpixel magnitude as well.²⁴ Overall Monte Carlo results are captured in Figure 10, comparing the evolution of posterior covariances with sample distributions, for the spacecraft position in Radial-Transverse-Normal (RTN) frame. The filter always converges in the vicinity of the ground truth. As expected, estimation precision in the camera plane directions is orders of magnitude higher than in the camera boresight direction. RTN posterior uncertainties computed by the estimator are smaller than the ones from the ESA’s Rosetta Navigation Analysis predictions (which are 13.3 km and 0.97 km in Longitudinal and Transversal direction, respectively, after *CAT-preinsertion*).⁷ However, no landmarks were used during the approach phase from ESA’s Rosetta team, whereby comet center observation was solely employed for navigation.²⁴ Posterior uncertainties for all the other estimated parameters are reported in Table 1.

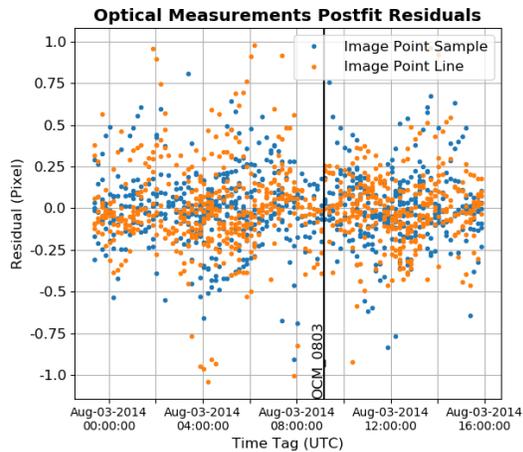


Figure 9: Converged residuals for optical measurements, from one Monte Carlo sample.

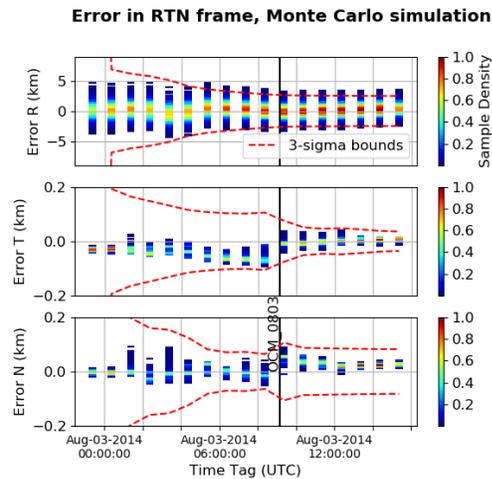


Figure 10: Distribution of RTN position error over 5000 Monte Carlo trials and posterior covariances.

The error distribution in Radial position fits relatively well within covariance bounds; however, the OD solution presents a deterministic bias, visually evident in Transverse and Normal position errors*. Monte Carlo statistics show that the worst-case position bias is around 0.5% of the radial distance from the comet. This phenomenon can be linked to feature detection: as discussed, there arguably is a correlation between feature detection and the dynamical behavior of lighting and perspective conditions characterizing the small body environment. Such bias does not particularly undermine feature matching consistency, which is ensured by trained descriptors and outliers rejection steps, bounding the drift within two pixels. However, further work needs to be done to obtain unbiased results.

CONCLUSIONS

We presented a multi-phase autonomous navigation pipeline for approaching small bodies that adapts to large-scale changes, from subpixel-resolution to images favorable to landmark tracking. We start by estimating the target body’s rotation rate, followed by pole and shape. We use such parameters as inputs for landmark-based navigation thereafter. We assessed the performance of the pipeline which, from preliminary results and under our assumptions, demonstrated performance compatible with Rosetta navigation requirements. However, we note that state-of-the-art visual-feature trackers are challenged by lighting and geometry changes, affecting orbit determination and resulting in a biased solution. Several improvements are planned for future work. First, developing an end-to-end feature tracking strategy robust to lighting and perspective changes during detection, description and outliers rejection. Second, relaxing remaining assumptions, producing orbit determination analyses on longer time frames and generalizing to different *a-priori* knowledge and small-body typologies. Third, include closed-loop guidance and control into the autonomous pipeline.

ACKNOWLEDGMENTS

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. ©2020 California Institute of Technology. Government sponsorship acknowledged.

REFERENCES

- [1] D. DellaGiustina, J. Emery, D. Golish, B. Rozitis, C. Bennett, K. Burke, R.-L. Ballouz, K. Becker, P. Christensen, C. D. d’Aubigny, *et al.*, “Properties of rubble-pile asteroid (101955) Bennu from OSIRIS-REx imaging and thermal analysis,” *Nature Astronomy*, Vol. 3, No. 4, 2019, pp. 341–351.
- [2] K. Kitazato, R. Milliken, T. Iwata, M. Abe, M. Ohtake, S. Matsuura, T. Arai, Y. Nakauchi, T. Nakamura, M. Matsuoka, *et al.*, “The surface composition of asteroid 162173 Ryugu from Hayabusa2 near-infrared spectroscopy,” *Science*, Vol. 364, No. 6437, 2019, pp. 272–275.

*It is noteworthy that the posterior covariance for spacecraft position in Transverse and Normal directions is enlarged (with respect to the sample distribution) by the conservative measurement uncertainty we selected for the features.

- [3] F. Capaccioni, A. Coradini, G. Filacchione, S. Erard, G. Arnold, P. Drossart, M. De Sanctis, D. Bockelee-Morvan, M. Capria, F. Tosi, *et al.*, “The organic-rich surface of comet 67P/Churyumov-Gerasimenko as seen by VIRTIS/Rosetta,” *Science*, Vol. 347, No. 6220, 2015, p. aaa0628.
- [4] D. Yeomans, P. Antreasian, J.-P. Barriot, S. Chesley, D. Dunham, R. Farquhar, J. Giorgini, C. Helfrich, A. Konopliv, J. McAdams, *et al.*, “Radio science results during the NEAR-Shoemaker spacecraft rendezvous with Eros,” *Science*, Vol. 289, No. 5487, 2000, pp. 2085–2088.
- [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, *et al.*, “AIDA DART asteroid deflection test: Planetary defense and science objectives,” *Planetary and Space Science*, Vol. 157, 2018, pp. 104–115.
- [6] T. D. Swindle, I. A. Nesnas, and J. C. Castillo-Rogez, “Design Reference Missions (DRMs) for Advancing Autonomy in Exploration of Small Bodies,” *AGUFM*, Vol. 2019, 2019, pp. IN34A–08.
- [7] U. Herfort and C. Casas, “Trajectory Preparation for the Approach of Spacecraft Rosetta to Comet 67P/Churyumov-Gerasimenko,” *ISSFD2015 paper, International Symposium of Space Flight Dynamics, Munich, Germany*, 2015.
- [8] R. Gaskell, O. Barnouin-Jha, D. J. Scheeres, A. Konopliv, T. Mukai, S. Abe, J. Saito, M. Ishiguro, T. Kubota, T. Hashimoto, *et al.*, “Characterizing and navigating small bodies with imaging data,” *Meteoritics & Planetary Science*, Vol. 43, No. 6, 2008, pp. 1049–1061.
- [9] S. B. Broschart, N. Bradley, and S. Bhaskaran, “Kinematic Approximation of Position Accuracy Achieved Using Optical Observations of Distant Asteroids,” *Journal of Spacecraft and Rockets*, Vol. 56, No. 5, 2019, pp. 1383–1392.
- [10] S. Bhaskaran, “Autonomous navigation for deep space missions,” *SpaceOps 2012*, p. 1267135, 2012.
- [11] C. Cocaud and T. Kubota, “Autonomous navigation near asteroids based on visual SLAM,” *Proceedings of the 23rd International Symposium on Space Flight Dynamics, Pasadena, California*, 2012.
- [12] S. Bhaskaran, S. Nandi, S. Broschart, M. Wallace, L. A. Cangahuala, and C. Olson, “Small body landings using autonomous onboard optical navigation,” *The Journal of the Astronautical Sciences*, Vol. 58, No. 3, 2011, pp. 409–427.
- [13] F. Castellini, D. Antal-Wokes, R. P. d. Santayana, and K. Vantournhout, “Far approach optical navigation and comet photometry for the rosetta mission,” *Proceedings of the 25th International Symposium on Space Flight Dynamics (ISSFD’15)*, 2015.
- [14] W. M. Owen Jr, “Methods of optical navigation,” 2011.
- [15] P. L. Lamy, I. Toth, H. Weaver, L. Jorda, M. Kaasalainen, and P. Gutiérrez, “Hubble Space Telescope observations of the nucleus and inner coma of comet 67P/Churyumov-Gerasimenko,” *Astronomy & Astrophysics*, Vol. 458, No. 2, 2006, pp. 669–678.
- [16] K.-H. Glassmeier, H. Boehnhardt, D. Koschny, E. Kührt, and I. Richter, “The Rosetta mission: flying towards the origin of the solar system,” *Space Science Reviews*, Vol. 128, No. 1-4, 2007, pp. 1–21.
- [17] C. H. Acton Jr, “Ancillary data services of NASA’s navigation and ancillary information facility,” *Planetary and Space Science*, Vol. 44, No. 1, 1996, pp. 65–70.
- [18] F. Preusker, F. Scholten, K.-D. Matz, T. Roatsch, S. Hviid, S. Mottola, J. Knollenberg, E. Kührt, M. Pajola, N. Oklay, *et al.*, “The global meter-level shape model of comet 67P/Churyumov-Gerasimenko,” *Astronomy & Astrophysics*, Vol. 607, 2017, p. L1.
- [19] B. R. Kent, *3D scientific visualization with blender*. Morgan & Claypool Publishers, 2014.
- [20] S. Evans, W. Taber, T. Drain, J. Smith, H.-C. Wu, M. Guevara, R. Sunseri, and J. Evans, “MONTE: The next generation of mission design and navigation software,” *CEAS Space Journal*, Vol. 10, No. 1, 2018, pp. 79–86.
- [21] S. Bandyonadhyay, I. Nesnas, S. Bhaskaran, B. Hockman, and B. Morrell, “Silhouette-Based 3D Shape Reconstruction of a Small Body from a Spacecraft,” *2019 IEEE Aerospace Conference*, IEEE, 2019, pp. 1–13.
- [22] K. Forbes, *Calibration, recognition, and shape from silhouettes of stones*. PhD thesis, University of Cape Town, 2007.
- [23] M. Lauer, S. Kielbassa, and R. Pardo, “Optical measurements for attitude control and shape reconstruction at the Rosetta flyby of asteroid Lutetia,” *ISSFD2012 paper, International Symposium of Space Flight Dynamics, Pasadena, California, USA*, 2012.
- [24] R. P. d. Santayana and M. Lauer, “Optical measurements for rosetta navigation near the comet,” *Proceedings of the 25th International Symposium on Space Flight Dynamics (ISSFD), Munich*, 2015.
- [25] A. Laurentini, “The visual hull concept for silhouette-based image understanding,” *IEEE Transactions on pattern analysis and machine intelligence*, Vol. 16, No. 2, 1994, pp. 150–162.
- [26] S. A. K. Tareen and Z. Saleem, “A comparative analysis of sift, surf, kaze, akaze, orb, and brisk,” *2018 International conference on computing, mathematics and engineering technologies (iCoMET)*, IEEE, 2018, pp. 1–10.
- [27] A. Harvard, V. Capuano, E. Y. Shao, and S.-J. Chung, “Pose Estimation of Uncooperative Spacecraft from Monocular Images Using Neural Network Based Keypoints,” *AIAA Scitech 2020 Forum*, 2020, p. 1874.
- [28] L. Kneip and P. Furgale, “OpenGV: A unified and generalized approach to real-time calibrated geometric vision,” *2014 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2014, pp. 1–8.
- [29] T. D. Moyer, *Formulation for observed and computed values of Deep Space Network data types for navigation*, Vol. 3. John Wiley & Sons, 2005.