

Keypoint-based Stereophotoclinometry for Characterizing and Navigating Small Bodies: A Factor Graph Approach

Travis Driver*, Andrew Vaughan†, Yang Cheng‡, Adnan Ansar§, John Christian¶, and Panagiotis Tsiotras||

This paper proposes the incorporation of techniques from stereophotoclinometry (SPC) into a keypoint-based structure-from-motion (SfM) system to estimate the surface normal and albedo at detected landmarks to improve *autonomous* surface and shape characterization of small celestial bodies from *in-situ* imagery. In contrast to the current state-of-the-practice method for small body shape reconstruction, i.e., SPC, which relies on human-in-the-loop verification and high-fidelity *a priori* information to achieve accurate results, we forego the expensive maplet estimation step and instead leverage dense keypoint measurements and correspondences from an *autonomous* keypoint detection and matching method based on deep learning to provide the necessary photogrammetric constraints. Moreover, we develop a factor graph-based approach allowing for *simultaneous* optimization of the spacecraft’s pose, landmark positions, Sun-relative direction, and surface normals and albedos via fusion of Sun sensor measurements and image keypoint measurements. The proposed framework is validated on *real* imagery of the Cornelia crater on Asteroid 4 Vesta, along with pose estimation and mapping comparison against an SPC reconstruction, where we demonstrate precise alignment to the SPC solution without relying on any *a priori* camera pose and topography information or humans-in-the-loop.

I. Introduction

There has been an increasing interest in missions to small bodies (e.g., asteroids, comets) due to their great scientific value, with four currently in operation (OSIRIS-REx, Hayabusa2, Lucy, Psyche) and three scheduled to launch over the next five years (Hera, EMA, Janus). In addition to planetary protection [1] and resource utilization [2, 3], small bodies are believed to be remnants of the solar system’s formation, and studying their composition could provide insight into the evolution of the solar system and the origins of organic life on Earth [4]. These missions currently rely on an extended characterization phase, where a shape model is reconstructed from images acquired during a ground-controlled trajectory around the body, as shape models are essential for characterizing the body and estimating the spacecraft’s relative pose in subsequent phases [5]. However, current state-of-the-practice shape reconstruction methods rely on humans-in-the-loop and accurate *a priori* information to ensure accurate results.

Stereophotoclinometry [6] (SPC) is the current method of choice for 3D reconstruction of small bodies and has been used to model a broad suite of celestial bodies, including the Moon, 433 Eros, 25143 Itokawa [6], 4 Vesta [7, 8], and 101955 Bennu [9, 10] (see Figure 1). While SPC has proven effective, the process requires extensive human-in-the-loop verification and high-fidelity *a priori* information to achieve accurate results. Specifically, SPC attempts to estimate a collection of digital terrain maps (DTMs), high-resolution local topography and albedo maps, through direct alignment of ortho-rectified projections, or *orthoimages*, of a given surface patch from multiple images predicated on an initial shape model and accurate *a priori* estimates of the spacecraft’s pose (position and orientation). Photometric stereo techniques are applied to derive surface gradients and albedos of the imaged surface patches at each pixel of the orthoimages. The local topography solution is fixed upon convergence, typically requiring human input to achieve precise alignment to the images, and used to refine pose and landmark position estimates through a multistep iterative process by rendering the DTM and aligning it across multiple views [11]. Finally, the local DTMs can be collated into a global shape model by exploiting overlap and limb constraints within a separate iterative processing loop. While this

*PhD Candidate, Institute for Robotics and Intelligent Machines, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA.

†Senior Engineer, Mission Design and Navigation, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA

‡Robotics Technologist, Mobility and Robotic Systems, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA.

§Group Supervisor, Mobility and Robotic Systems, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA.

¶Associate Professor, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA.

||David & Andrew Lewis Chair, Professor, Institute for Robotics and Intelligent Machines, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA.




Fig. 1 Examples of shape models produced by SPC.

approach has achieved much success, its reliance on extensive human involvement for extended durations and a complex multistep optimization process limits mission capabilities and increases operational costs [12–14].

In contrast to the traditional maplet-based approach of SPC, this work proposes a keypoint-based framework that leverages dense image keypoint measurements and correspondences within a structure-from-motion (SfM) system. Indeed, SfM and simultaneous localization and mapping (SLAM), which leverage autonomous keypoint detection and matching methods to estimate correspondences between images, have been shown to be promising technologies for *autonomous* optical navigation for missions to small bodies [15, 16]. Consequently, we propose the incorporation of photometric stereo constraints and Sun vector measurements into a graph-based SfM system to estimate surface normals and albedos at estimated landmarks, providing detailed information for surface characterization and shape reconstruction. The proposed framework, which leverages *factor graphs* [17] to model and solve the complex SPC estimation problem, forgoes the expensive local maplet estimation step, streamlines the optimization process, and renders SPC more amenable to recent advances in computer vision, including feature detection, description, and matching methods based on *deep learning* [18] and graph-based incremental smoothing [19]. Specifically, the contributions of this work are as follows:

- 1) we propose the fusion of measurements from onboard monocular cameras and Sun sensors in a SfM system to reconstruct a landmark map of the body and the relative pose of the spacecraft while *simultaneously* estimating the surface normal and albedo at landmarks;
- 2) we model the photometric stereo process of SPC using the formalism of *factor graphs* and analyze its performance with respect to the chosen reflectance model;
- 3) we leverage keypoint measurements and correspondences from a state-of-the-art, *autonomous* keypoint detection and matching method based on deep learning [20] and apply the proposed framework to *real* imagery of the Cornelia crater on Asteroid 4 Vesta, demonstrating precise alignment to an SPC derived map.

II. Preliminaries

In this section, we include some relevant mathematical concepts and notation. Specifically, we summarize the representation of 3D poses as elements of the Special Euclidean group $SE(3)$ (Section II.A), and introduce a minimal representation of unit three-vectors (Section II.B).

A. 3D Rigid Body Transformations

We represent the relative position and orientation of the spacecraft, that is, its relative *pose*, as an element of the Special Euclidean group $SE(3)$ by a matrix

$$T_{\mathcal{BS}} \triangleq \begin{bmatrix} R_{\mathcal{BS}} & \mathbf{r}_{\text{SB}}^{\mathcal{B}} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix}, \quad (1)$$

where $R_{\mathcal{BS}} \in SO(3)$ is the orientation of some body-fixed frame of the small body $\mathcal{B}$ with respect to a spacecraft body-fixed frame $\mathcal{S}$, and $\mathbf{r}_{\text{SB}}^{\mathcal{B}} \in \mathbb{R}^3$ is the position of the spacecraft's origin with respect to the origin of $\mathcal{B}$, expressed in $\mathcal{B}$. Moreover, the fixed pose $T_{\mathcal{SC}}$ of the onboard camera relative to the body-fixed frame of the spacecraft is precisely known *a priori*, which can be used to derive the sequence of *the camera's relative pose* $T_{\mathcal{BC}_k} = T_{\mathcal{BS}_k} T_{\mathcal{S}_k \mathcal{C}_k} = T_{\mathcal{BS}_k} T_{\mathcal{SC}}$ at each time index k .

The uncertainty of the camera's pose, using the exponential map (at the identity) of the $SO(3)$ group of rotations and the *hat* operator (see [21]), can be defined as

$$T_{\mathcal{BC}} \triangleq \bar{T}_{\mathcal{BC}} \begin{bmatrix} \exp([\omega]^\wedge) & \mathbf{v} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix}, \quad (2)$$

where $\omega \sim \mathcal{N}(\mathbf{0}_{3 \times 1}, \Sigma_R)$, $\mathbf{v} \sim \mathcal{N}(\mathbf{0}_{3 \times 1}, \Sigma_r)$, and $\bar{T}_{\mathcal{BS}}$ is the *actual* pose [17, 22]. Therefore, the estimated pose $T_{\mathcal{BC}_k}$ is represented by the uncertain orientation $R_{\mathcal{BC}_k} \triangleq \bar{R}_{\mathcal{BC}_k} \exp([\omega]^\wedge)$ and the uncertain position $\mathbf{r}_{\text{CB}}^{\mathcal{B}} \triangleq \bar{\mathbf{r}}_{\text{CB}}^{\mathcal{B}} + \bar{R}_{\mathcal{BC}_k} \mathbf{v}$. This relates to a *retraction* $\mathcal{R}_T(\gamma, \tau)$ (at a pose $T \in SE(3)$) that maps the local coordinates $\zeta = [\gamma^\top \tau^\top]^\top \in \mathbb{R}^6$ to an element of $SE(3)$ via

$$\mathcal{R}_T(\gamma, \tau) \triangleq T \begin{bmatrix} \exp([\gamma]^\wedge) & \tau \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} = \begin{bmatrix} R \exp([\gamma]^\wedge) & \mathbf{r} + R\tau \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix}. \quad (3)$$

Thus, the optimization of elements of $SE(3)$ may be framed in terms of the coordinates $\zeta \in \mathbb{R}^6$ [22, 23].

B. The Unit 2-Sphere $\mathbb{S}^2$

An important two-dimensional manifold is the unit 2-sphere $\mathbb{S}^2 \triangleq \{\mathbf{x} \in \mathbb{R}^3 \mid \|\mathbf{x}\| = 1\}$, i.e., the topological space composed of all unit vectors in $\mathbb{R}^3$. The tangent space $\mathcal{T}_{\mathbf{x}}(\mathbb{S}^2)$ at a point $\mathbf{x} \in \mathbb{S}^2$ is defined as the set of all three-vectors tangent to $\mathbb{S}^2$ at $\mathbf{x}$, i.e.,

$$\mathcal{T}_{\mathbf{x}}(\mathbb{S}^2) \triangleq \{\mathbf{y} \in \mathbb{R}^3 \mid \mathbf{x}^\top \mathbf{y} = 0, \mathbf{x} \in \mathbb{S}^2\}. \quad (4)$$

For any $\mathbf{y} \in \mathcal{T}_{\mathbf{x}}(\mathbb{S}^2)$, we can write $\mathbf{y} = B_{\mathbf{x}} \xi$ with $\xi \in \mathbb{R}^2$ in the plane tangent to $\mathbb{S}^2$ at $\mathbf{x}$ defined by the basis $B_{\mathbf{x}} \in \mathbb{R}^{3 \times 2}$. With these definitions, we can define another useful retraction $\mathcal{R}_{\mathbf{x}}(\xi)$ as follows [17]:

$$\mathcal{R}_{\mathbf{x}}(\xi) \triangleq \cos(\|B_{\mathbf{x}} \xi\|) \mathbf{x} + \sin(\|B_{\mathbf{x}} \xi\|) \frac{B_{\mathbf{x}} \xi}{\|B_{\mathbf{x}} \xi\|} \in \mathbb{S}^2. \quad (5)$$

This minimal representation allows for the optimization of the unit vector $\mathbf{x} \in \mathbb{S}^2$ with respect to the local coordinates $\xi \in \mathbb{R}^2$ according to the basis $B_{\mathbf{x}}$. Uncertainty in the unit vector can also be defined in the local coordinate system defined by $B_{\bar{\mathbf{x}}}$ at the true value $\bar{\mathbf{x}}$, i.e., $\mathbf{x} = \mathcal{R}_{\bar{\mathbf{x}}}(\varepsilon)$, where $\varepsilon \sim \mathcal{N}(\mathbf{0}_{2 \times 1}, \Sigma_\varepsilon)$. We will use this formulation, implemented in the Georgia Tech Smoothing and Mapping (GTSAM) library [24], to represent and optimize surface normal estimates (detailed in Section IV.D).

III. Related Work

Vision-based methods, i.e., SPC, have traditionally been applied to 3D reconstruction for small celestial bodies. However, approaches based on active sensors such as Flash-LiDARs have also been proposed. Bercovici et al. [25] proposed a pose estimation and shape reconstruction approach based on Flash-LiDAR measurements by solving a maximum likelihood estimation problem via particle-swarm optimization to refine an initial Bezier surface mesh, followed by a least-squares filter providing measurements for the position and orientation of the spacecraft. Other works

in the field have established proofs-of-concept for batch optimization and graph-based approaches for near-asteroid navigation and shape reconstruction. Notably, Nakath et al. [26] present an active SLAM framework which also employs Flash-LiDAR as the base measurement, with sensor fusion of data from an inertial measurement unit and star tracker, tested with simulated data. However, the limited range of Flash-LiDAR instruments bounds the spacecraft’s orbit to unrealistically small radii, reducing the feasible scenarios to either navigation near very small asteroids or the touchdown phase. For example, the OSIRIS-REx Guidance, Navigation, and Control (GNC) Flash-LiDAR, which is mentioned in both [26] and [25], has only a maximum operational range of approximately 1 km and a relatively small 128×128 detector array [27, 28]

Moreover, the recent OSIRIS-REx mission to Asteroid 101955 Bennu was also equipped with the OSIRIS-REx Laser Altimeter (OLA) [29] to provide an alternative means of shape reconstruction to SPC. The OLA reconstruction process begins by generating local DTMs, similar to SPC, from the range measurements, which are then collated through an iterative closest-point algorithm. However, the OSIRIS-REx camera suite (OCAMS) [30] features a long-range camera that provides higher spatial resolution than OLA at the same distance. Therefore, although the SPC process is more time consuming than the OLA-based approach, SPC products can be available before OLA models, and with a higher resolution. Moreover, testing of OLA-generated DTMs showed that uncertainty in OLA measurements created unacceptable errors in the elevations of smaller features and albedo is not automatically included as part of the solution [9]. Ultimately, the SPC data products were used in the final touch-and-go phase of the mission.

More recently, feature-based methods have been shown to be promising technologies for *autonomous* optical navigation and mapping for missions to small bodies. Most notably, Dor et al. [15] demonstrated precise visual localization and mapping on *real* images of Asteroid 4 Vesta through a feature-based SLAM system based on Oriented FAST and Rotated BRIEF (ORB) features [31]. ORB is a handcrafted method based on Features from Accelerated Segment Test (FAST) keypoints [32] and Binary Robust Independent Elementary Features (BRIEF) descriptors [33] and outputs binary descriptor vectors, enabling more efficient matching. This work was extended in [16] to include known dynamical motion constraints between the small body and the spacecraft to further improve mapping and localization performance. Furthermore, Driver et al. [18] proposed the use of deep learning-based feature detection and description methods, which were shown to significantly outperform traditional handcrafted methods (e.g., ORB), especially in scenarios involving considerable changes in illumination and perspective. Our work capitalizes on this recent success in feature-based SLAM and SfM for autonomous optical navigation by imbuing the traditional photogrammetric framework with added characterization power by incorporating photometric stereo constraints for concurrent estimation of surface normals and albedos.

As stated previously, the proposed approach takes inspiration from the success of SPC. However, our implementation has several key differences. First, SPC represents surface normals implicitly through a local height map parametrization via the integration of computed surface slopes whereby each DTM is optimized *independently* in a local frame of the associated landmark located at the center of the maplet [6, 11]. Conversely, we leverage a minimal representation of the surface normals based on retractions of the tangent space of the $\mathbb{S}^2$ manifold, detailed in Section II.B, which allows for global and simultaneous optimization of all observed landmarks, along with their associated surface normals. Second, and most notably, we forego the dense alignment step used to generate the local DTMs and instead rely on keypoint measurements, which can be generated by autonomous feature detection and description methods, and have been shown to be robust to the significant illumination and perspective changes inherent to small body imagery [18]. Thus, our framework treats keypoints (referred to as *landmark image locations* in the SPC text [11]) as measurements instead of another variable of interest that must be estimated independently. Lastly, we exploit the formalism of factor graphs (see Section IV.B) to model the complex estimation problem, as opposed to the iterative multistep process employed by the traditional SPC implementation [11], allowing for the simultaneous optimization of camera poses, landmark positions, and surface normals and albedos.

SPC borrows many techniques from the process of *photometric stereo* [34], which has been used extensively in terrestrial applications. This is not to be confused with *shape-from-shading* (SfS) whereby the shape of a 3D object may be recovered from shading in a *single image*. However, terrestrial photometric stereo formulations have relied on a number of simplifying assumptions, including Lambertian reflectance [35, 36] or specialized lighting or image capture setups [37–39]. Methods based on deep learning have also been proposed but, as before, require specially designed image acquisition setups [40–43], and thus could not be leveraged in a general multi-view reconstruction scenario. We refer the reader to [44] and [45] for more information about physics-based and data-driven approaches, respectively, to photometric stereo for terrestrial applications. Also worth mentioning are Neural Radiance Fields (NeRFs) [46], which learn a volumetric scene representation via the use of multilayer perceptrons (MLPs) and are primarily applied to view-synthesis. However, NeRFs implicitly model the underlying reflectance and material properties of the target object




Fig. 2 System overview.

or scene through the learned weights of the MLPs. Our work, instead, exploits semi-empirical photometric models of the surface of airless bodies with few free parameters to explicitly estimate the topography and material properties of the imaged surface.

IV. Proposed Approach

First, we introduce the SfM problem (Section IV.A). Next, we formulate SfM as a maximum *a posteriori* (MAP) inference problem using the formalism of *factor graphs* (Section IV.B). Finally, we present the photometric stereo framework and detail its integration into the proposed SPC-SfM system (Sections IV.D, IV.E, and IV.F). An overview of the proposed approach is shown in Figure 2.

A. Structure-from-Motion

The proposed formulation will leverage SfM to estimate the spacecraft’s relative pose and a landmark map of the small body’s surface. Feature-based SfM and SLAM [47] leverage monocular images taken from multiple viewpoints along a trajectory to jointly estimate the robot’s pose with respect to its environment and construct a 3D model of the scene. The SfM architecture is typically comprised of two main components: the *front-end* and the *back-end*. The front-end extracts 2D interest points (*keypoints*) from images, represents each keypoint with a local feature *descriptor*, and matches keypoints between images by comparing their associated descriptors [18]. The front end also performs *data association* by associating the 2D keypoint measurements with specific points in 3D space (the *landmarks*). Finally, the associations from the front-end are used to *simultaneously* reconstruct a map of the environment and resolve the pose of the camera through inference in the *back-end* via maximum *a posteriori* (MAP) estimation.

Formally, let $\mathcal{B}$ denote some body-fixed frame of the small body with origin B, and let C_k denote the camera frame at time index k with origin C_k . Moreover, let $\ell_j^{\mathcal{B}} = [\ell_{x,j}^{\mathcal{B}} \ \ell_{y,j}^{\mathcal{B}} \ \ell_{z,j}^{\mathcal{B}}]^\top \in \mathbb{R}^3$ denote the vector from B to the j th surface landmark expressed in $\mathcal{B}$, let $\mathbf{q}_j^{C_k} = [q_{x,j}^{C_k} \ q_{y,j}^{C_k} \ q_{z,j}^{C_k}]^\top \in \mathbb{R}^3$ denote the vector from C_k to the j th landmark expressed in C_k , and let $\mathbf{p}_{j,k} = [u_{j,k} \ v_{j,k}]^\top \in \mathbb{R}^2$ denote the 2D image coordinates of the j th landmark observed by camera C_k , i.e., the keypoint. SfM seeks the MAP estimate of the camera poses $\mathcal{T} := \{T_{\mathcal{B}C_k} \in \text{SE}(3) \mid k = 0, \dots, n\}$ and a network of landmarks (the *map*) $\mathcal{L} := \{\ell_j^{\mathcal{B}} \in \mathbb{R}^3 \mid j = 1, \dots, m\}$ given the (independent) keypoint *measurements*

$$\mathcal{P} := \{\hat{\mathbf{p}}_{j,k} \in \mathbb{R}^2 \mid k = 1, \dots, n, j = 1, \dots, m\}:$$

$$\mathcal{T}^*, \mathcal{L}^* = \arg \max_{\mathcal{T}, \mathcal{L}} p(\mathcal{T}, \mathcal{L} \mid \mathcal{P}) \quad (6)$$

$$\propto \arg \max_{\mathcal{T}, \mathcal{L}} p(\mathcal{T}, \mathcal{L}) p(\mathcal{P} \mid \mathcal{T}, \mathcal{L}) \quad (7)$$

$$= \arg \max_{\mathcal{T}, \mathcal{L}} p(\mathcal{T}, \mathcal{L}) \prod_k \prod_j p(\hat{\mathbf{p}}_{j,k} \mid T_{\mathcal{B}C_k}, \ell_j^{\mathcal{B}}). \quad (8)$$

Note that the SfM solution is innately expressed in some arbitrary body-fixed frame since most SfM techniques assume operation in a static scene, typically referred to as the “world” frame [47]. By assuming that the measurements $\hat{\mathbf{p}}_{j,k}$ are corrupted by zero-mean Gaussian noise, i.e., $\hat{\mathbf{p}}_{j,k} = \bar{\mathbf{p}}_{j,k} + \boldsymbol{\eta}_{j,k}$ where $\boldsymbol{\eta}_{j,k} \sim \mathcal{N}(\mathbf{0}, \Sigma_{j,k})$, we get

$$p(\hat{\mathbf{p}}_{j,k} \mid T_{\mathcal{B}C_k}, \ell_j^{\mathcal{B}}) \propto \exp \left\{ -\frac{1}{2} \|\Pi(\ell_j^{\mathcal{B}}, T_{\mathcal{B}C_k}; K) - \underline{\mathbf{p}}_{j,k}\|_{\Sigma_{j,k}}^2 \right\}, \quad (9)$$

where $\|\mathbf{e}\|_{\Sigma}^2 := \mathbf{e}^\top \Sigma^{-1} \mathbf{e}$, and the *forward-projection* function Π relates landmarks $\ell_j^{\mathcal{B}}$ to their (homogenous) coordinates $\underline{\mathbf{p}}_{j,k}$ in the k th image, i.e.,

$$\underline{\mathbf{p}}_{j,k} = \Pi(\ell_j^{\mathcal{B}}, T_{\mathcal{B}C_k}; K) = \frac{1}{d_j^{C_k}} [K \mid \mathbf{0}^{3 \times 1}] T_{\mathcal{B}C_k}^{-1} \ell_j^{\mathcal{B}} = \frac{1}{d_j^{C_k}} K \mathbf{q}_j^{C_k}, \quad (10)$$

where $d_j^{C_k} = q_{z,j}^{C_k}$ is the landmark depth in C_k , $\underline{\ell}_j^{\mathcal{B}} = \left[\left(\ell_j^{\mathcal{B}} \right)^\top \ 1 \right]^\top \in \mathbb{P}^3$ and $\underline{\mathbf{p}}_{j,k} = \left[\left(\mathbf{p}_{j,k} \right)^\top \ 1 \right]^\top \in \mathbb{P}^2$ denote the homogeneous coordinates of $\ell_j^{\mathcal{B}}$ and $\mathbf{p}_{j,k}$, respectively, and K is the camera calibration matrix:

$$K = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}, \quad (11)$$

where f_x and f_y are the *focal lengths* in the x - and y -directions of the camera frame, and (c_x, c_y) is the *principal point* of the camera.

Finally, the MAP estimate can be formulated as the solution to a nonlinear least-squares problem by taking the negative logarithm of (8):

$$\mathcal{T}^*, \mathcal{L}^* = \arg \min_{\mathcal{T}, \mathcal{L}} \sum_{k,j} \|\Pi(\ell_j^{\mathcal{B}}, T_{\mathcal{B}C_k}; K) - \underline{\mathbf{p}}_{j,k}\|_{\Sigma_{j,k}}^2, \quad (12)$$

where we have omitted the priors $p(\mathcal{T}, \mathcal{L})$ for conciseness and generality, which can be ignored if no prior information is assumed (i.e., $p(\mathcal{T}, \mathcal{L}) = \text{const.}$) or can encode relative pose constraints via known dynamical models [16]. This process is commonly referred to as *Bundle Adjustment* (BA). Note that the optimization process of SPC decouples estimation of the poses and the landmarks, i.e., *a priori* landmark position and camera pose estimates are passed back-and-forth between the pose determination and DTM construction steps, respectively, until convergence [6].

In this work, we will leverage tools from the Georgia Tech Structure-from-Motion (GTSfM) library [48] to generate a sparse SfM solution (described in the following subsection), followed by a densification step to construct a dense map of the imaged surface. The keypoint measurements and correspondences, $\hat{\mathbf{p}}_{j,k}$, will be computed by a state-of-the-art, *autonomous* keypoint detection and matching method based on deep learning, i.e., DKM [20]. Image brightness measurements will be extracted at observed keypoints and combined with a prior reflectance model (Section IV.D) to determine the surface normal and albedo at the estimated landmarks. The back-end will leverage factor graphs to represent the MAP estimation problem, which will be discussed in Section IV.B.

1. Georgia Tech Structure-from-Motion (GTSfM)

GTSfM [48] is an end-to-end global SfM pipeline based on the Georgia Tech Smoothing and Mapping (GTSAM) library [24], an estimation toolbox based on factor graphs (see Section IV.B). GTSfM was designed from the ground-up to natively support parallel computation to allow for fast pose estimation and mapping. We will leverage GTSfM to




Fig. 3 DKM dense matching examples.

generate initial pose estimates and a sparse map, which will then be subsequently densified and used to initialize the graph-based SPC smoothing step (Section IV.D). Given a collection of images, GTSfM executes the following stages to generate a map (the set of landmarks $\mathcal{L}$) of the imaged target and the pose of the camera when each image was taken (the poses $\mathcal{T}$):

Correspondence Generation Keypoint measurements and correspondences are generated using a state-of-the-art, deep keypoint detection and matching method, Dense Kernelized Feature Matching (DKM) [20]. DKM leverages a deep convolutional feature encoder based on ResNet [49] to generate coarse features, which are subsequently matched using a Gaussian Process regression scheme. The established coarse matches are then refined at the native resolution of the image using stacked feature maps and large depth-wise kernels to provide dense, *pixel-wise* matches. Since these “detector-free” methods do not provide a consistent set of keypoints per image, and instead compute a dense mapping between the pixel coordinates of each image pair, we choose a reference image and match all images with respect to this image only. We provide examples of the dense matching results on *real* images of the Cornelia crater on Asteroid 4 Vesta in Figure 3. Only the 2,048 most confident matches are used in the two-view estimation step described next.

Two-View Estimation The two-view estimation module attempts to estimate the relative rotations, $R_{C_l C_k}$, and translations, $\mathbf{r}_{C_k C_l}^{C_l}$, according to the keypoint correspondences from the front-end. Specifically, given corresponding keypoints $\mathbf{p}_{j,k}$ and $\mathbf{p}_{j,l}$ observed by cameras C_k and C_l , respectively, the *essential matrix* $E := [\mathbf{r}_{C_k C_l}^{C_l}]_{\times} R_{C_l C_k}$ satisfies $\mathbf{p}_{j,l}^{\top} K^{-\top} E K^{-1} \mathbf{p}_{j,k} = 0$, where $[\cdot]_{\times}$ denotes the skew-symmetric cross-product matrix. The well-known five-point algorithm [50] can be used to solve for E given five or more correspondences. The essential matrix is robustly estimated using Graph-cut RANSAC [51], implemented in OpenCV, with a 0.5 pixel outlier threshold. Finally, $R_{C_l C_k}$ and $\mathbf{r}_{C_k C_l}^{C_l}$ (up to some unknown scale) can be estimated via singular value decomposition (SVD) of the essential matrix and imposing the *cheirality constraint*, i.e., triangulating the keypoints and enforcing that the associated landmark lies in front of the cameras [52].

Pose Averaging Pose averaging seeks to estimate the “global” poses $T_{\mathcal{B} C_k}$ from the previously computed relative poses. This module is composed of rotation averaging followed by translation averaging. Nominally, GTSfM leverages the Shonan rotation averaging method [53]. However, since we match all images to a common reference image, this step is trivial as the global orientation of the reference image is taken to be the identity, and the orientations of the remaining images are computed relative to the reference using the previously computed relative orientations. Next, given the estimated camera orientations, $R_{\mathcal{B} C_k}$, from the rotation averaging step, along with the relative translations, $\mathbf{r}_{C_k C_l}^{C_l}$, from the two-view estimation module, we recover the position, $\mathbf{r}_{C_k \mathcal{B}}^{\mathcal{B}}$, of each camera (up to an unknown scale) using the method proposed in [54]. We transform the relative translations to a body-fixed frame using the estimated rotations, project them (and landmark rays) along several random directions, and find estimates consistent with the minimum-feedback arc set (MFAS) problem [54], and remove the inconsistent estimates. Finally, we solve for the

translations $\mathcal{T} \triangleq \{\mathbf{r}_{C_k B}^{\mathcal{B}} \in \mathbb{R}^3 \mid k = 0, \dots, n\}$ via minimization of a chordal distance $d_{\text{ch}}(\mathbf{u}, \mathbf{v}) \triangleq \|\mathbf{u}/\|\mathbf{u}\| - \mathbf{v}/\|\mathbf{v}\|\|_2$:

$$\mathcal{T}^* = \arg \min_{\mathcal{T}} \sum_{k,l} d_{\text{ch}} \left(R_{\mathcal{B} C_l} \mathbf{r}_{C_k C_l}^{C_l}, \mathbf{r}_{C_k B}^{\mathcal{B}} - \mathbf{r}_{C_l B}^{\mathcal{B}} \right)^2, \quad (13)$$

where the tuple (k, l) denotes the image indices which were matched in the correspondence generation step.

Data Association The keypoint measurements $\hat{\mathbf{p}}_{j,k}$ are associated with landmarks from the pairwise matches using a disjoint-set forest (DSF) algorithm. Then, the initial 3D position of the landmarks, ℓ_j , are triangulated from the 2D keypoint measurements using the Direct Linear Transform (DLT) [55, Chapter 4.1], along with RANSAC (with a 0.5 pixel error threshold) to filter possible outlier measurements that made it through the two-view verification process.

Bundle Adjustment Finally, the initial camera poses and triangulated point cloud are refined using a global bundle adjustment step, as described by Equation 12, to obtain the final sparse SfM solution.

The reasoning behind the initial sparse solution is two-fold: (1) leveraging *sparse* correspondences, as opposed to the per-pixel matches, significantly reduces the number of computations in the redundant two-view estimation step; (2) by leveraging only the most confident matches, we reduce the risk of incorporating any outlier matches into the estimation scheme and dense matches can be more reliably verified during the densification step before being added to the map. We refer the reader to [48] for more details about the GTSfM pipeline. Finally, densification of the sparse GTSfM solution is conducted by computing the squared Sampson error [56] of each of the putative correspondences and adding the match if its error is below 1. The dense map is then triangulated using the same scheme described in the data association module. This dense map, and the associated camera poses from the sparse solution, are used to initialize the graph-based SPC step described in Section IV.D.

B. Factor Graphs

We leverage the formalism of factor graphs to facilitate the fusion of measurements from onboard sensors (i.e., monocular cameras and Sun sensors) through prior photometric models to *simultaneously* estimate the camera poses, landmark positions, and surface normals and albedos at each landmark. Formally, a factor graph is a bipartite graph $G = (\mathcal{F}, \Theta, \mathcal{E})$ with *factor nodes* $f_i \in \mathcal{F}$ that abstract the measurements and prior knowledge $z_i \in \mathcal{Z}$ as generalized probabilistic constraints between *variable nodes* $\theta_j \in \Theta$, the unknown random variables, where *edges* $e_{i,j} \in \mathcal{E}$ define the interdependence relationships between a factor f_i and a variable θ_j . With these definitions, a factor graph G defines a factorized function

$$f(\Theta) = \prod_i f_i(\Theta_i), \quad (14)$$

where each measurement factor $f_i(\Theta_i) = l(\Theta_i; z_i)$ of the variables $\Theta_i = \{\theta_j \in \Theta \mid e_{i,j} \in \mathcal{E}\}$, with likelihood $l(\Theta_i; z_i) \propto p(z_i \mid \Theta_i)$, and each prior factor $f(\Theta_i) = p(\Theta_i)$ represents a term in the joint probability density function (PDF), i.e., $f(\Theta) \propto p(\Theta \mid \mathcal{Z})$. We seek the variable assignment Θ^* through maximum *a-posteriori* (MAP) inference over the joint probability distribution encoded by the factors in the factor graph:


$$\Theta^* = \arg \max_{\Theta} \prod_i f_i(\Theta_i). \quad (15)$$

Assuming a zero-mean Gaussian noise model with measurement covariance Σ_i gives factors of the form


$$f_i(\Theta_i) \propto \exp \left\{ -\frac{1}{2} \|h_i(\Theta_i) - z_i\|_{\Sigma_i}^2 \right\}, \quad (16)$$

where $h_i(\cdot)$ is a measurement prediction function. Moreover, assume that priors $p(\Theta_i)$ can be written as $p(\Theta_i) \propto \exp \left\{ -\frac{1}{2} \|h_i(\Theta_i) - z_i\|_{\Sigma_i}^2 \right\}$ with prior mean and covariance z_i and Σ_i , respectively. Therefore, solving (15) is equivalent to minimizing a sum of nonlinear least-squares via

$$\arg \min_{\Theta} (-\log f(\Theta)) = \arg \min_{\Theta} \frac{1}{2} \sum_i \|h_i(\Theta_i) - z_i\|_{\Sigma_i}^2. \quad (17)$$



(a) Typical factor graph formulation of the SfM problem.



(b) With proposed factors f_{SPC} relating to photometric stereo constraints, f_{SS} relating to Sun sensor measurements, and f_{Smooth} relating to local smoothness constraints.

Fig. 4 Variable nodes are camera poses $T_k \in \text{SE}(3)$, landmarks $\ell_j \in \mathbb{R}^3$, sun vectors $s_l \in \mathbb{S}^2$, surface normals $\mathbf{n}_i \in \mathbb{S}^2$, and surface albedos $a_i \in [0, 1]$. Factor nodes f_L and f_P relate to keypoint-based landmark measurements and possibly a prior factor, respectively.

This formulation allows factor graphs to support PDFs or cost functions of any number of variables [19], allowing for the inclusion of multiple sensor modalities, as well as prior knowledge and constraints to uniquely determine the MAP solution for the unknown variables Θ^* . The typical factor graph formulation of SfM is shown in Fig. 4a, where the factors $f_L(\ell_k^B, T_{\mathcal{BC}_k}; K)$ relate to the forward-projection error function defined in Equation (12).

Solving the nonlinear least-squares problem in (17) typically involves repeated linearization. For nonlinear measurement functions $h_i(\cdot)$, nonlinear optimization approaches such as the Levenberg-Marquardt algorithm (LMA) leverage repeated first-order linear approximations to (17) to approach the minimum. In addition, the interdependence relationships encoded by the edges of the factor graph capture the factored nature of the PDF and sparsity of the underlying information matrix, allowing for *exact* nonlinear optimization in an *incremental* setting by exploiting the sparse edge connections to identify the variables to be optimized when a new measurement becomes available [19]. The factor graph formulation of the proposed keypoint-based SPC problem is implemented using the Georgia Tech Smoothing and Mapping (GTSAM) library [24], an estimation toolbox based on factor graphs pioneered at Georgia Tech. The toolbox provides a fully customizable framework for factor graph construction and a suite of nonlinear optimization methods.

C. Small Body Photometry

Photogrammetry, serving as the theoretical foundation for contemporary techniques such as SfM, primarily concentrates on establishing geometric relationships between points in an image (the *keypoints*) and the corresponding points in the scene (the *landmarks*). In contrast, *photometry* attempts to determine the “brightness” of a point within the scene as it appears in the observed image. Photogrammetry has historically garnered more attention in the computer vision community, due, in part, to the relatively low complexity as compared to photometry. Indeed, photometric modeling is inherently complex. This is especially true for terrestrial applications, which typically require the use of deep neural networks to achieve accurate photometric reconstructions [58]. However, operation in space presents us




Fig. 5 Photometry conventions.

Table 1 Investigated reflectance functions. The coefficients for the Schröder model are listed in Table 2.

	$r_F(\alpha, \beta, \phi)$	$g(\phi)$	$\Lambda(\phi)$	
Lunar-Lambert	$a \left((1 - g(\phi)) \cos \alpha + g(\phi) \frac{2 \cos \alpha}{\cos \alpha + \cos \beta} \right)$	$\exp(-\phi/60)$	—	(18)
Schröder	$a \Lambda(\phi) \left((1 - g(\phi)) \cos \alpha + g(\phi) \frac{2 \cos \alpha}{\cos \alpha + \cos \beta} \right)$	$C_0 + C_1 \phi$	$\sum_{i=0}^4 B_i \phi^i$	(19)

with a number of advantages that simplify the photometric modeling process. First, we may treat the Sun as a point source delivering collimated light to the surface owing to the large distance to the Sun (i.e., the Sun subtends an angle of 0.5° at Earth [59]) and the lack of atmosphere of most small bodies to scatter the incoming light. Second, the direction of the incoming light can be precisely measured using typical onboard instrumentation (e.g., Sun sensors, star trackers). Third, previous photometric modeling of small bodies has demonstrated that *global* reflectance functions, as opposed to *spatially-varying* reflectance functions, can precisely estimate the observed brightness across the surface of a target small body. Moreover, these global reflectance properties are similar across asteroids of the same taxonomic class [60].

More formally, an image may be considered as a mapping $I: \Omega \rightarrow \mathbb{R}_+$ over the pixel domain $\Omega \subset \mathbb{R}^2$ that maps points in the image, $\mathbf{p} \in \Omega$, to their corresponding “brightness” value, $I(\mathbf{p}) \in \mathbb{R}_+$. Here, “brightness” refers to the fact that the image values correspond to the amount of light falling on the photodetector inside the camera, referred to as the image *irradiance* with units of power per unit area ($\text{W} \cdot \text{m}^{-2}$). It can be shown [61, Chapter 10.3] that the image irradiance is proportional to the *radiance*, with units of power per unit solid angle per unit area ($\text{W} \cdot \text{sr}^{-1} \cdot \text{m}^{-2}$), reflected towards the camera from the surface. Thus, each point $\mathbf{p} \in \Omega$ corresponds to the radiance emitted from a point (or, more precisely, an infinitesimal patch) on the surface of the body. As such, the raw image brightness values, digitized to a

Table 2 Coefficients for the Schröder reflectance function [57], where we have normalized $B_i, i = 1, 2, 3, 4$, such that $\Lambda(0) = 1$.

C_0	C_1	B_0	B_1	B_2	B_3	B_4
0.830	-7.22×10^{-3}	1	-1.7160×10^{-2}	1.8306×10^{-4}	-1.0399×10^{-6}	2.3223×10^{-9}

14-bit “digital number” (DN) for the Dawn Framing Camera [62], may be converted to radiance by normalizing by the image exposure time and utilizing an appropriate *radiometric calibration* process [63, 64].

The emitted radiance from the surface, $L(\alpha, \beta, \phi, a)$ and the incident (collimated) irradiance from the Sun, F , which is a function of the distance from the target body to the Sun [59], are related by the *bidirectional reflectance function*, r (in units of sr^{-1}):

$$\frac{L(\alpha, \beta, \phi, a)}{F} = r(\alpha, \beta, \phi, a), \quad (20)$$

where a is the surface albedo, α is the angle between the incoming light and the surface normal, or the *incidence angle*, β is the angle between the emitted light and the surface normal, or the *emission angle*, and ϕ is the angle between the emitted light and the incoming light, or the *phase angle* (see Figure 5). A similar measure of reflectance, which is very popular in the context of planetary photometry [59, 60], is the *bidirectional radiance factor*, r_F (in units of I/F , which are dimensionless), which is the ratio of the bidirectional reflectance of a surface to that of a perfect Lambertian surface illuminated and viewed from overhead (i.e., $\alpha = \beta = \phi = 0$):

$$r_F(\alpha, \beta, \phi, a) = \pi r(\alpha, \beta, \phi, a). \quad (21)$$

Henceforth, we are referring to the radiance factor when mentioning the reflectance function.

We consider two reflectance functions in this investigation, defined in Table 1. The first model, which we refer to as the *Lunar-Lambert* function, defined in Equation (18), features a combination of Lambert and Lommel-Seeliger photometric functions weighted according to an exponential function of the phase angle, $g(\phi)$. The model in Equation (18) is an approximation to McEwen’s reflectance function [65, 66], which was fit to lunar imagery captured by the Galileo spacecraft. This model was chosen because it is nominally used by SPC and has been shown to be well-suited for photometry on a wide range of small bodies [6, 7, 10, 11]. The second, which we refer to as the *Schröder* model [57], defined in Equation (19), again features a combination of Lambert and Lommel-Seeliger photometric functions, but is weighted according to an *affine* function of the phase angle and also features a *surface phase function*, $\Lambda(\phi)$. The *phase function* describes the change in brightness with phase that are independent of the incidence and emission angles. The Schröder model in Equation (19) was fit to approach imagery of Asteroid 4 Vesta captured during the Dawn mission. By investigating this reflectance model, we will demonstrate the utility of specialized reflectance models fit to relevant imagery.

We will discuss how these photometric principles can be incorporated into a feature-based SfM system in the following section. We will consider both calibrated and uncalibrated imagery in this investigation. For the uncalibrated case, we use corrected DN values where various error sources such as dark current and readout smear have been removed [63, 64]. For the calibrated case, the images have been converted to units of I/F according to the radiometric calibration process detailed in [64].

D. Photometric Stereo Constraints

Photometric stereo techniques are integrated into the feature-based SfM system to estimate surface normals and albedos at estimated landmarks. Photometric stereo [34] is the process of determining surface gradients of an object by observing it from different viewpoints and lighting conditions and is leveraged by SPC to facilitate dense surface reconstruction. As before, let an image taken at time index k be denoted by $I_k: \Omega \rightarrow \mathbb{R}_+$ over the pixel domain $\Omega \subset \mathbb{R}^2$. The measured image brightness $\hat{I}_k(\mathbf{p}_{j,k})$ (calibrated to units of I/F) at a keypoint $\mathbf{p}_{j,k} \in \Omega$ in image I_k associated with a landmark $\ell_j \in \mathbb{R}^3$ can be modeled by an appropriate reflectance function, as detailed in Section IV.C:

$$I(\alpha_{j,k}, \beta_{j,k}, \phi_{j,k}, a_j) = r_F(\alpha_{j,k}, \beta_{j,k}, \phi_{j,k}, a_j), \quad (22)$$

where a_j is the albedo at landmark ℓ_j and $\alpha_{j,k}, \beta_{j,k}$, and $\phi_{j,k}$ are the incidence, emission, and phase angles, respectively, at landmark ℓ_j in the k^{th} image. When considering *uncalibrated* imagery, a scale, λ_k , and bias, ξ_k , term are typically included in Equation (21) to account for factors such as distance to the Sun, exposure times, and background noise for each image [6, 11, 67], i.e.,

$$I(\alpha_{j,k}, \beta_{j,k}, \phi_{j,k}, a_j) = \lambda_k r_F(\alpha_{j,k}, \beta_{j,k}, \phi_{j,k}, a_j) + \xi_k. \quad (23)$$

In this case, a_j is the *relative* surface albedo, which refers to the fact that, unless considering radiometrically calibrated imagery [63, 64], the albedo a_j is only proportional to the absolute albedo.

The Sun-relative direction $\mathbf{s}_k^{\mathcal{B}} \in \mathbb{S}^2$ in I_k , expressed in the body-fixed frame of the small body $\mathcal{B}$, can be estimated using measurements from typical onboard instrumentation (e.g., Sun sensors, star trackers), detailed in Section IV.E. The emitted light vector $\mathbf{e}_{k,j}^{\mathcal{B}} = \mathbf{r}_{\mathcal{C}\mathcal{B}}^{\mathcal{B}} - \boldsymbol{\ell}_j^{\mathcal{B}}$ can be determined from the estimates of $T_{\mathcal{B}\mathcal{C}_k}$ and $\boldsymbol{\ell}_j^{\mathcal{B}}$ provided by SfM. Finally, dropping the superscripts and letting T_k denote $T_{\mathcal{B}\mathcal{C}_k}$ for conciseness, Equation (22) can be written in terms of $\mathbf{s}_k$, $\mathbf{e}_{k,j}$, and the surface normal $\mathbf{n}_j \in \mathbb{S}^2$ at $\boldsymbol{\ell}_j$ (see Fig. 5) by noticing that $\cos \alpha_{j,k} = \mathbf{s}_k^\top \mathbf{n}_j$, $\cos \beta_{j,k} = \mathbf{e}_{k,j}^\top \mathbf{n}_j / \|\mathbf{e}_{k,j}\|$, and $\phi_{j,k} = \cos^{-1}(\mathbf{s}_k^\top \mathbf{e}_{k,j} / \|\mathbf{e}_{k,j}\|)$. For example, the Lunar-Lambert model (Equation (18)) becomes

$$I(T_k, \mathbf{s}_k, \boldsymbol{\ell}_j, \mathbf{n}_j, a_j) = a_j \left((1 - g(\mathbf{s}_k, \mathbf{e}_{k,j})) \mathbf{s}_k^\top \mathbf{n}_j + g(\mathbf{s}_k, \mathbf{e}_{k,j}) \frac{2\mathbf{s}_k^\top \mathbf{n}_j}{\mathbf{s}_k^\top \mathbf{n}_j + \mathbf{e}_{k,j}^\top \mathbf{n}_j / \|\mathbf{e}_{k,j}\|} \right). \quad (24)$$

We can now define a factor f_{SPC} corresponding to the presented photometric stereo constraints (assuming zero-mean Gaussian noise) as follows:

$$f_{\text{SPC}}(T_k, \mathbf{s}_k, \boldsymbol{\ell}_j, \mathbf{n}_j, a_j; \Sigma_k) \propto \exp \left\{ -\frac{1}{2} |I(T_k, \mathbf{s}_k, \boldsymbol{\ell}_j, \mathbf{n}_j, a_j) - \hat{I}_k(\hat{\mathbf{p}}_{j,k})|_{\Sigma_k}^2 \right\}. \quad (25)$$

This allows for estimation of $\mathbf{n}_j$ and a_j using the measurements $\hat{I}_k(\hat{\mathbf{p}}_{j,k})$, while also further constraining the landmark's position $\boldsymbol{\ell}_j$, Sun-relative direction $\mathbf{s}_k$, and the position of the spacecraft $\mathbf{r}_{\mathcal{C}\mathcal{B}}$. The corresponding factor graph diagram is shown in Fig. 4b.

E. Sun Vector Measurements

Our framework assumes knowledge of the Sun-relative direction $\mathbf{s}_k$. This direction is usually determined from the navigation solution and ephemeris files, and onboard attitude estimates (e.g., from a star tracker) allow this direction to be expressed in the camera frame. The Sun direction may also be measured directly by a Sun sensor. For example, the OSIRIS-REx spacecraft featured multiple coarse Sun sensors, each with an accuracy of $\pm 1^\circ$ (3σ) [68]. Moreover, fine Sun sensors have accuracy on the order of $\pm 0.01^\circ$ (3σ). For simplicity, we will assume that the measurements $\hat{\mathbf{s}}_k^C \in \mathbb{S}^2$ are available at each time index k and are expressed in the camera frame $\mathcal{C}$ via the fixed relative pose $T_{\mathcal{U}\mathcal{C}}$ of the Sun sensor frame $\mathcal{U}$ with respect to $\mathcal{C}$. Recalling that T_k denotes $T_{\mathcal{B}\mathcal{C}_k}$ and $\mathbf{s}_k$ is expressed in the $\mathcal{B}$ frame, a measurement prediction function $\mathbf{s}^C(T_k, \mathbf{s}_k)$ can be defined to predict the measured incident light direction $\hat{\mathbf{s}}^C$ in the $\mathcal{C}$ frame from the current estimates of T_k and $\mathbf{s}_k$:

$$\mathbf{s}^C(T_k, \mathbf{s}_k) \triangleq R_{\mathcal{B}\mathcal{C}_k}^{-1} \mathbf{s}_k. \quad (26)$$

We define a factor f_{SS} to incorporate (simulated) Sun sensor measurements into the estimation problem as follows:

$$f_{\text{SS}}(T_k, \mathbf{s}_k; \Sigma_{j,k}) \propto \exp \left\{ -\frac{1}{2} \|\mathbf{s}^C(T_k, \mathbf{s}_k) - \hat{\mathbf{s}}_k^C\|_{\Sigma_{j,k}}^2 \right\}. \quad (27)$$

This further constrains the orientation of the camera $R_{\mathcal{B}\mathcal{C}_k}$ and the Sun-relative direction $\mathbf{s}_k$.

F. Local Smoothness Constraints

While the photometric minimization and Sun vector terms modeled by f_{SPC} and f_{SS} , respectively, are sufficient to estimate the surface normal and albedo, Horn [69] indicates that the solution tends to be unstable and gets stuck in local minima, especially if starting far from the solution. This has also been demonstrated in other works on small body shape reconstruction [67, 70]. Thus, Horn proposed the use of local smoothness constraints which minimize the “departure from smoothness.” Our smoothness constraint *factors* are defined as follows:

$$f_{\text{Smooth}}(\boldsymbol{\ell}_j, \mathbf{n}_j, \boldsymbol{\ell}_{j'}, \eta) \propto \exp \left\{ -\frac{1}{2} \eta \left| \cos^{-1}(\mathbf{d}_{j',j}^\top \mathbf{n}_j) - 90^\circ \right|^2 \right\}, \quad (28)$$

where η weights the local smoothness penalty, and $\mathbf{d}_{j',j} = (\boldsymbol{\ell}_{j'} - \boldsymbol{\ell}_j) / \|\boldsymbol{\ell}_{j'} - \boldsymbol{\ell}_j\|$. In words, f_{Smooth} encourages landmarks to be locally smooth with respect to the reference landmark's surface normal ($\mathbf{n}_j$) by enforcing that $\mathbf{d}_{j',j}$, i.e., the vector pointing from the reference landmark ($\boldsymbol{\ell}_j$) towards a neighboring landmark ($\boldsymbol{\ell}_{j'}$), be perpendicular to $\mathbf{n}_j$. We will demonstrate that adding these smoothness constraints will result in more feasible surface normal estimates and lower photometric errors.




Fig. 6 Image sequence used in this investigation. Camera poses and *global* shape model are drawn according to the SPC solution.




Fig. 7 Baselines for comparison.

V. Experimental Setup

Our experiments leverage imagery of the Cornelia crater on Asteroid 4 Vesta captured by NASA’s Dawn mission [62, 71] for evaluation of the proposed approach. These images are publicly available on NASA’s Planetary Data System (PDS) [72] and maintained by NASA’s Navigation and Ancillary Information Facility (NAIF), each with a resolution of 1024×1024 pixels. Cornelia is an approximately 15 km in diameter crater located at 15.57°E and 9.37°S (in the Claudia double-prime coordinate frame [73]) and has been the subject of numerous studies [57, 74, 75]. Cornelia was also chosen in part due to its interesting albedo features, which represent a challenging scenario for photometric reconstruction methods such as SPC and the proposed approach. The full image sequence used in this investigation, captured during the High Altitude Mapping Orbit (HAMO) phase, is visualized in Figure 6, which features 29 images captured at a mean orbital radius of ~ 950 km and a ground sample distance of approximately 60 m per pixel.

A. Experimental Baselines

We compare our pipeline against four different baselines (see Figure 7), including the SPC solution and two different methods based on stereophotogrammetry (SPG):

SPC An important point of comparison is with respect to the traditional SPC [6, 11] pipeline, which will provide a baseline for our surface normal and albedo estimates, along with pose and landmark position values. The SPC

reconstruction leverages the Schröder reflectance model as opposed to the nominal Lunar-Lambert function. A bigmap of the Cornelia region was created by merging 289 overlapping sub-maps, or maplets. The maplets, and the resulting bigmap, have a spatial resolution of 30 meters. Approximately 50 pictures were used to model the terrain, captured during both the High Altitude Mapping Orbit (HAMO) and Low Altitude Mapping Orbit (LAMO) phases. The HAMO images, which constituted most of the imagery, had a ground sample distance (GSD) of approximately 60 meters per pixel, while the images from the LAMO phase featured a GSD of 20 meters per pixel. Not all pictures were used in every maplet. *A priori* pose information was sampled from a global orbital dynamics solution developed for the purpose of gravity science. *A priori* height information for the maplets was sampled from a previously converged set of maps at a spatial resolution of 50 meters, which were in turn derived from 100 meter spatial resolution maps. More details about the SPC reconstruction process at Vesta can be found in [8].

SPG(1) The first method based on SPG leverages multi-baseline stereo matching and triangulation. The pose estimates from the SPC solution were used, and a reference image was chosen (the first image in Figure 6) to find the pixel correspondents from the rest of images. Because the image poses from SPC are very accurate, it allowed for precise epipolar rectification and 1D disparity search. To improve the matching robustness, the images were also normalized [76]. Then, the 3D position of each point is triangulated [55, Chapter 4.1] to obtain the final map.

SPG(2) A second multi-baseline stereo approach was also used. Again, the pose estimates from SPC were leveraged. Rather than using a reference image, all images were rectified to a series of densely spaced, parallel planes that sweep through the volume of interest. Then, given a point in any image, its 3D location is determined by the rectifying plane which maximized photo-consistency across the entire dataset [77].

SfM This is the dense SfM solution, i.e., the the SPC-SfM process without the f_{SPC} , f_{SS} , and f_{Smooth} factors.

Since our method does not assume any priors on the camera poses or landmark positions, resulting in a scale ambiguity, we must align our solution to the baselines before comparison. To do this, we estimate a Sim(3) transformation between the estimated and ground truth camera poses and a subset of 180 of the landmarks, by Karcher mean and [78] (implemented in GTSAM’s `Similarity3.align` function), and apply this transformation to our reconstructed poses, landmarks, and surface normals. Since the SPG solutions use the poses estimated by SPC, we align to the SPC solution before comparing to the SPG(1) and SPG(2) solutions.

B. Implementation Details

Keypoint measurements and matches are computed using DKM, which provides dense, per-pixel correspondences, where the first image (the top left image in Figure 6) is taken to be the reference image, as described in Section IV.A.1. The matching is constrained to a 500×500 pixel region, drawn in red in Figure 6, as this contains the extent of the baseline SPC and SPG maps. Thus, each of the maps estimated by our approach contains approximately 250,000 points. The keypoint measurements are assigned a covariance of $\Sigma_{j,k} = \mathbf{I}_2$. Image brightness values are (bilinearly) interpolated at the keypoint measurements $\hat{\mathbf{p}}_{j,k}$ to derive the measurements $\hat{I}_k(\hat{\mathbf{p}}_{j,k})$ used in the proposed SPC factors f_{SPC} defined in Equation (25). The image brightness measurements are assigned a standard deviation of $\sigma_I = 0.5$ for the uncalibrated case and $\sigma_I = 0.01$ for the calibrated case, which we found to work well empirically. Only landmarks with ≥ 8 keypoint measurements are inserted into the graph.

The (simulated) Sun sensor measurements $\hat{\mathbf{s}}_k^C \in \mathbb{S}^2$ were derived from the normalized ground truth values from SPC of the Sun’s relative position to the origin of the camera frame $\mathbf{C}_k$ expressed in the camera frame $\mathbf{C}_k$, i.e., $\mathbf{r}_{\text{IC}_k}^{C_k}$, and assigned an uncertainty of $\Sigma_\xi = \sigma_\xi^2 \mathbf{I}_2$ where $\sigma_\xi = 1 \times 10^{-3}$. Next, the surface normals are initialized by finding the 32 closest neighbors to each point in the point cloud and fitting a plane to this local terrain, and the normal to the plane is taken as the initial surface normal. These initial surface normals are then used to initialize the albedo by independently computing the albedo in each image using the initial camera poses and landmark positions, where the initial albedo of each landmark is taken to be the average albedo computed over all views from which it was seen.

The smoothness factors f_{Smooth} (Equation (28)) are inserted into the graph according to the keypoints in the reference image (corresponding to $k = 0$). Specifically, a smoothness factor is inserted between a landmark, ℓ_j (and its associated surface normal, $\mathbf{n}_j$) and a neighboring landmark, $\ell_{j'}$, if $\|\hat{\mathbf{p}}_{j,0} - \hat{\mathbf{p}}_{j',0}\| \leq 1$. Thus, a smoothness factor will be inserted between a landmark and at most four other landmarks in the map since DKM computes *per-pixel* matches between an

image pair. We found a very small value for the local smoothness weight to work well for our experiment, where we used a value of $\eta = 10^{-4}$. As previously discussed, we leverage the GTSAM library [24] to model the proposed keypoint-based SPC problem using factor graphs and optimize the resulting nonlinear least-squares using the Levenberg-Marquardt algorithm and the analytical partial derivatives of the measurement functions for the respective factors (see Appendix A).

C. Performance metrics

We define the following metrics to measure the performance of the proposed approach:

$$\delta \ell_j \triangleq \|\ell_j - \bar{\ell}_j\|_2, \quad (29)$$

$$\delta \epsilon_j \triangleq \cos^{-1} \left(\mathbf{n}_j^\top \bar{\mathbf{n}}_j \right), \quad (30)$$

and

$$\delta a_j \triangleq |a_j - \bar{a}_j| / \bar{a}_j. \quad (31)$$

As before, $\bar{\ell}$, $\bar{\mathbf{n}}$, $\bar{a}$ denote the ground truth values of the landmark position, surface normal, and albedo, respectively. These ground truth values are assigned by finding the closest point in our reconstructed map to that of the baseline (after the Sim(3) alignment step), and taking the position of that landmark, as well as the associated normal and albedo for the SPC baseline, as the ground truth. For quantitative evaluation of the estimated trajectory, for each time index $k = 0, \dots, n$, we compute the error

$$\begin{bmatrix} \delta \boldsymbol{\kappa}_k^\top & \delta \mathbf{r}_k^\top \end{bmatrix} \triangleq \left[\log \left(T_{\mathcal{B}C_k}^{-1} \bar{T}_{\mathcal{B}C_k} \right) \right]^\vee \in \mathbb{R}^6, \quad (32)$$

between the ground truth relative pose $\bar{T}_{\mathcal{B}C_k}$ and the estimated pose $T_{\mathcal{B}C_k}$, where $\log : \text{SE}(3) \rightarrow \mathfrak{se}(3)$ is the SE(3) logarithm map at the identity and $\vee$ is the *vee* operator, as detailed in [21], which gives

$$\|\delta \boldsymbol{\kappa}_k\| = \cos^{-1} \left(\frac{\text{trace}(R_{C_k \mathcal{B}} \bar{R}_{\mathcal{B}C_k}) - 1}{2} \right), \quad (33)$$

$$\|\delta \mathbf{r}_k\| = \|R_{C_k \mathcal{B}} \bar{\mathbf{r}}_{C_k \mathcal{B}} + \mathbf{r}_{\mathcal{B}C_k}^{C_k}\|. \quad (34)$$

Finally, the root mean squared error between the measured, $\hat{I}_k(\hat{\mathbf{p}}_{j,k})$, and estimated, $I(T_k, \mathbf{s}_k, \ell_j, \mathbf{n}_j, a_j)$, image brightness values, normalized by the average measured brightness, is taken to be the photometric error, as in [57]:

$$\delta I_j \triangleq \left(\frac{1}{|\mathcal{K}_j|} \sum_{k \in \mathcal{K}_j} \hat{I}_k(\hat{\mathbf{p}}_{j,k}) \right)^{-1} \sqrt{\frac{1}{|\mathcal{K}_j|} \sum_{k \in \mathcal{K}_j} \left(I(T_k, \mathbf{s}_k, \ell_j, \mathbf{n}_j, a_j) - \hat{I}_k(\hat{\mathbf{p}}_{j,k}) \right)^2}, \quad (35)$$

where $\mathcal{K}_j$ denotes the set of indices of the images from which the j^{th} landmark was viewed.

VI. Results

The resulting albedo map and surface normals, along with the photometric errors, from our proposed SPC-SfM solution are provided in Figures 9 and 10 for the Lunar-Lambert and Schröder models, respectively. For the *uncalibrated* cases, we see that our method is able to reconstruct maps with very low photometric errors of 0.0147 and 0.0122 for the Lunar-Lambert and Schröder model, respectively, albeit with numerous artifacts in the surface normal maps. The addition of the smoothness constraints is shown to generate coherent surface normals while also significantly reducing the average photometric error and bringing the photometric error for the Lunar-Lambert model in line with the Schröder model at 0.0108. These extremely low photometric errors result in renderings that are virtually indistinguishable from the captured images, as shown in Figure 8. However, the *calibrated* case paints a slightly different picture in terms of the fidelity of the examined reflectance functions, where the Lunar-Lambert model results in an average photometric error of 0.0227, over twice that of the Schröder model on the same data. This may be unsurprising, as the Lunar-Lambert equation was fit to images of the Moon while the Schröder model leveraged imagery of Vesta. However, this phenomenon has not been demonstrated in previous works on SPC that leverage only uncalibrated imagery, where these

Table 3 Mean values of the performance metrics (outlined in Section V.C) with respect to each baseline. Pose errors were not computed for the SPG solutions, as they leverage the pose estimates from SPC (see Section V.A).

Model	SIM			SPG(1)			SPG(2)			SPC					
	$\delta\ell$ [km]	$\ \delta\mathbf{r}\ $ [km]	$\ \delta\kappa\ $ [°]	$\delta\ell$ [km]	$\ \delta\mathbf{r}\ $ [km]	$\ \delta\kappa\ $ [°]	$\delta\ell$ [km]	$\ \delta\mathbf{r}\ $ [km]	$\ \delta\kappa\ $ [°]	$\delta\ell$ [km]	$\ \delta\mathbf{r}\ $ [km]	$\ \delta\kappa\ $ [°]	$\delta\epsilon_s$ [°]	$\delta\epsilon_n$ [°]	δa [%]
Lunar-Lambert															
<i>Uncalibrated</i>	0.0017	0.0882	0.0082	0.0425	—	—	0.0563	—	—	0.0881	1.0850	0.0966	0.6061	13.72	6.05
<i>Uncal. + f_{smooth}</i>	0.0048	0.0904	0.0080	0.0419	—	—	0.0579	—	—	0.0867	1.0879	0.0966	0.5743	4.80	3.85
<i>Cal. + f_{smooth}</i>	0.0209	1.3522	0.1497	0.0507	—	—	0.0626	—	—	0.0822	1.9542	0.1988	4.8529	7.92	6.96
Schröder															
<i>Uncalibrated</i>	0.0037	0.1589	0.0159	0.0408	—	—	0.0576	—	—	0.0885	1.0985	0.0987	0.6222	13.40	4.31
<i>Uncal. + f_{smooth}</i>	0.0052	0.1285	0.0135	0.0432	—	—	0.0582	—	—	0.0867	1.0850	0.0977	0.5389	4.80	3.85
<i>Cal. + f_{smooth}</i>	0.0057	0.1977	0.0226	0.0433	—	—	0.0585	—	—	0.0867	1.0965	0.1013	0.5170	4.78	3.85




Fig. 8 Per image photometric error and qualitative comparison between measured and estimated brightness values for the Schröder (*uncalibrated*) model with smoothness constraints. The legend in (a) is formatted according to $k : (\lambda_k, \xi_k)$ where the biases ξ_k have been normalized by the average brightness of each image.

works indicate that the choice of reflectance model has little effect on the quality of the final DTM [11]. Moreover, this suggests that the scale and bias variables estimated in the uncalibrated case are able to “explain away” many of the sources of error due to the underlying reflectance model and motivates the estimation and use of specialized reflectance models for different small bodies, especially when dealing with radiometrically calibrated images.

Next, we compare our solution to the SPC reconstruction in Figures 11 and 12 for the Lunar-Lambert and Schröder models, respectively. Again, we see that the Lunar-Lambert model with the radiometrically calibrated imagery results in

higher errors as compared to the other solutions. The uncalibrated cases without the smoothness factors display surface normal errors of $\sim 13^\circ$ with respect to SPC, but these errors are virtually rectified with the addition of the smoothness factors, which bring the average error to $< 5^\circ$. The relative albedo errors are also exceedingly low, almost all of which are below 4%, except for the calibrated case for the Lunar-Lambert model and the uncalibrated cases without smoothness constraints, which are still below 7% on average. However, all solutions exhibit regions with relatively large errors with respect to the landmark positions (on the order of 350 m). Yet, these large landmark error regions are not present when comparing against the SfM and SPG solutions, as shown in Figures 13 and 14. This suggests that the errors in the landmarks arise from errors in the SPC solution, as opposed to errors in our approach, purportedly due to errors in the associated surface normal, which seem to originate at the dark albedo feature near the northwest edge of the crater. Indeed, since the landmark heights in the SPC solution are computed by integrating the slopes [8], errors in the slope translate to errors in the landmark positions that propagate from the points in the map where the slope errors arise towards the direction of the integration. Conversely, since our landmarks are not explicitly derived from the surface normal estimates, and instead are independently estimated and constrained by the keypoint measurements, our pipeline does not seem to be as susceptible to surface normal errors as the traditional SPC pipeline.

These results, along with the associated pose errors with respect to the SPC and SfM solutions, are also summarized in Table 3. Except for the Lunar-Lambert model on calibrated imagery, all position errors and orientation errors are approximately 1 km and 0.1° , respectively, as compared to the SPC solution. Thus, we achieve positions errors of $\sim 0.1\%$ relative to the mean orbital radius of 950 km without requiring *a priori* estimates of the camera’s pose or an initial topography solution. Moreover, our pipeline also refines the Sun-relative direction, where our Sun vector estimates are within 0.6° of the SPC estimates. When comparing against the SfM solution, our position and orientation estimates align even better, with position errors of < 0.2 km (< 0.1 km) and orientation errors of $< 0.025^\circ$ ($< 0.01^\circ$) for the Schröder (Lunar-Lambert) model. It is unclear why the Lunar-Lambert SPC-SfM solution seems to demonstrate marginally lower pose errors than the Schröder model as compared to the SfM solution. In any case, the position errors for the Schröder model as compared to the SfM solution are $\sim 0.02\%$ relative to the mean orbital radius. The Lunar-Lambert model on calibrated imagery exhibits inflated errors with respect to almost all metrics, demonstrating how incorrectly modeled reflectance properties can perturb the solution.


VII. Conclusion

This paper proposes the incorporation of photometric stereo constraints leveraged by stereophotoclinometry (SPC) into a keypoint-based structure-from-motion (SfM) system to estimate the surface normal and albedo at landmarks to improve surface and shape characterization of small celestial bodies from *in-situ* imagery. The proposed framework was validated on real imagery of the Cornelia crater on Asteroid 4 Vesta using dense keypoint measurements and correspondences from a deep learning-based method, i.e., DKM [20]. The results demonstrate that the proposed framework is able to estimate accurate surface normal and albedo estimates as compared to the baseline SPC values. Moreover, our approach does not rely on high-fidelity *a priori* estimates of the camera’s pose or an initial topography solution, and can operate *autonomously* thanks to the accurate keypoint correspondences from the front-end. While the proposed approach assumes a global reflectance function (and phase function for the Schröder model), this assumption may not be sufficient for more heterogeneous bodies such as Asteroid 101955 Bennu [79]. Future work will include continued validation and analysis of the proposed approach on other celestial bodies (e.g., the Moon, Bennu) and incorporating reflectance function parameter estimation. Our reconstructed maps and camera poses will be made available at <https://github.com/travisdriver/spc-factor-results>.


Appendix

A. Measurement Function Partial Derivatives


In this section, we explicitly define the partial derivatives for the measurement function $I(T, \mathbf{s}, \boldsymbol{\ell}, \mathbf{n}, a)$, with the Lunar-Lambert model (Equation (18)), for the factor f_{SPC} (Equation (25)), which are used by GTSAM [24] to solve the nonlinear least-squares problem defined in Equation (17). The nonlinear least-squares objective is reparameterized according to the outlined retractions, i.e., Equations (3) and (5). This reparameterization is referred to as *lifting* [23] and frames the optimization problem in terms of coordinates in the tangent space of the variables. Recall that $\mathbf{e} = \mathbf{r}_{\text{CB}} - \boldsymbol{\ell}$.



(a) L-Lambert (*uncalibrated*)




(b) L-Lambert (*uncalibrated*) with smoothness constraints




(c) L-Lambert (*calibrated*) with smoothness constraints


Fig. 9 Reconstructed surface normal and albedo maps, along with the photometric errors, for the Lunar-Lambert reflectance function. The (relative) albedos for the uncalibrated case are normalized by the maximum albedo in the map.



(a) Schröder (*uncalibrated*)



(b) Schröder (*uncalibrated*) with smoothness constraints



(c) Schröder (*calibrated*) with smoothness constraints

Fig. 10 Reconstructed surface normal and albedo maps, along with the photometric errors, for the Schröder reflectance function. The (relative) albedos for the uncalibrated case are normalized by the maximum albedo in the map.









Fig. 11 Quantitative comparison between our solution and the SPC solution with respect to landmark error, normal error, and relative albedo error for the Lunar-Lambert reflectance function.

Fig. 12 Quantitative comparison between our solution and the SPC solution with respect to landmark error, normal error, and relative albedo error for the Schröder reflectance function.







(a) L-Lambert (*uncalibrated*)(a) Schröder (*uncalibrated*)(b) L-Lambert (*uncalibrated*) with smoothness constraints(b) Schröder (*uncalibrated*) with smoothness constraints(c) L-Lambert (*calibrated*) with smoothness constraints(c) Schröder (*calibrated*) with smoothness constraints

Fig. 13 Quantitative comparison between our solution and the SPC, SfM, and SPG solutions with respect to landmark error for the Lunar-Lambert model.

Fig. 14 Quantitative comparison between our solution and the SPC, SfM, and SPG solutions with respect to landmark error for the Schröder model.

Then,

$$\frac{\partial \mathbf{e}}{\partial \zeta} = \begin{bmatrix} \mathbf{I}_{3 \times 3} & R \end{bmatrix}, \quad \frac{\partial \mathbf{e}}{\partial \ell} = -\mathbf{I}_{3 \times 3}, \quad (36)$$

where we have applied the retraction $r_{\text{CB}} \leftarrow r_{\text{CB}} + R\tau$. Let $\mathbf{d} = \mathbf{e}/\|\mathbf{e}\|$, $f = \mathbf{s}^\top \mathbf{n}$, $w = \mathbf{d}^\top \mathbf{n}$, and $h = \mathbf{s}^\top \mathbf{d}$. Then

$$\begin{aligned} \frac{\partial f}{\partial \xi_s} &= \mathbf{n}^\top B_s, & \frac{\partial f}{\partial \xi_n} &= \mathbf{s}^\top B_n, \\ \frac{\partial w}{\partial \xi_d} &= \mathbf{n}^\top B_d, & \frac{\partial w}{\partial \xi_n} &= \mathbf{d}^\top B_n, \\ \frac{\partial h}{\partial \xi_s} &= \mathbf{d}^\top B_s, & \frac{\partial h}{\partial \xi_d} &= \mathbf{s}^\top B_d, \\ \frac{\partial \xi_d}{\partial \mathbf{e}} &= B_d^\top D_{\text{norm}}(\mathbf{e}), \end{aligned} \quad (37)$$

where we have applied the retraction $\mathbf{x} \leftarrow \mathbf{x} + B_x \xi_x$ with B_x and ξ_x representing the basis for and local coordinates in the tangent space at a unit vector $\mathbf{x} \in \mathbb{S}^2$, respectively, and

$$D_{\text{norm}}(\mathbf{v}) = \begin{bmatrix} v_y^2 + v_z^2 & -v_x v_y & -v_x v_z \\ -v_x v_y & v_x^2 + v_z^2 & -v_y v_z \\ -v_x v_z & -v_y v_z & v_x^2 + v_y^2 \end{bmatrix} (v_x^2 + v_y^2 + v_z^2)^{-3/2}. \quad (38)$$

Next, let $b = -\cos^{-1}(h)/60$ and $g = \exp(b)$. Then,

$$\begin{aligned} \frac{\partial b}{\partial h} &= \left(60\sqrt{1-h^2}\right)^{-1}, & \frac{\partial g}{\partial \xi_s} &= \frac{\partial g}{\partial b} \frac{\partial b}{\partial h} \frac{\partial h}{\partial \xi_s}, \\ \frac{\partial g}{\partial b} &= \exp(b), & \frac{\partial g}{\partial \xi_d} &= \frac{\partial g}{\partial b} \frac{\partial b}{\partial h} \frac{\partial h}{\partial \xi_d}. \end{aligned} \quad (39)$$

This gives

$$\frac{\partial I}{\partial \mathbf{e}} = a \left(-f \frac{\partial g}{\partial \xi_d} + 2f \left(\frac{1}{f+w} \frac{\partial g}{\partial \xi_d} - \frac{g}{(f+w)^2} \frac{\partial w}{\partial \xi_d} \right) \right) \frac{\partial \xi_d}{\partial \mathbf{e}}. \quad (40)$$

Finally,

$$\frac{\partial I}{\partial \zeta} = \frac{\partial I}{\partial \mathbf{e}} \frac{\partial \mathbf{e}}{\partial \zeta}, \quad (41)$$

$$\frac{\partial I}{\partial \xi_s} = a \left(\frac{\partial f}{\partial \xi_s} - \left(\frac{\partial g}{\partial \xi_s} f + g \frac{\partial f}{\partial \xi_s} \right) + \frac{2}{f+w} \left(\frac{\partial g}{\partial \xi_s} f + g \frac{\partial f}{\partial \xi_s} \right) - \frac{2gf}{(f+w)^2} \frac{\partial f}{\partial \xi_s} \right), \quad (42)$$

$$\frac{\partial I}{\partial \ell} = \frac{\partial I}{\partial \mathbf{e}} \frac{\partial \mathbf{e}}{\partial \ell}, \quad (43)$$

$$\frac{\partial I}{\partial \xi_n} = a \left((1-g) \frac{\partial f}{\partial \xi_n} + 2g \left(\frac{1}{f+w} \frac{\partial f}{\partial \xi_n} - \frac{f}{(f+w)^2} \left(\frac{\partial f}{\partial \xi_n} + \frac{\partial w}{\partial \xi_n} \right) \right) \right), \quad (44)$$

$$\frac{\partial I}{\partial a} = (1-g)f + g \frac{2f}{f+w}. \quad (45)$$

Acknowledgments

This work was supported by NASA Space Technology Graduate Research Opportunity 80NSSC21K1265. A portion of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). The authors would like to thank Kenneth Getzandanner and Andrew Liounis from NASA Goddard Space Flight Center for several helpful discussions and comments. We also thank Stefanus E. Schröder from Luleå University of Technology for providing support for the radiometric calibration process [63, 64].

References

- [1] Cheng, A. F., Rivkin, A. S., Michel, P., Atchison, J., Barnouin, O., Benner, L., Chabot, N. L., Ernst, C., Fahnestock, E. G., Kueppers, M., Pravec, P., Rainey, E., Richardson, D. C., Stickle, A. M., and Thomas, C., “AIDA DART asteroid deflection test: Planetary defense and science objectives,” *Planetary and Space Science*, Vol. 157, 2018, pp. 104–115.
- [2] Mazanek, D. D., Merrill, R. G., Brophy, J. R., and Mueller, R. P., “Asteroid Redirect Mission concept: A bold approach for utilizing space resources,” *Acta Astronautica*, Vol. 117, 2015, pp. 163–171.
- [3] Rivkin, A. S., and DeMeo, F. E., “How Many Hydrated NEOs Are There?” *J. of Geophysical Research: Planets*, Vol. 124, No. 1, 2019, pp. 128–142.
- [4] Barucci, M. A., Dotto, E., and Levasseur-Regourd, A. C., “Space missions to small bodies: asteroids and cometary nuclei,” *Astronomy and Astrophysics Rev.*, Vol. 19, No. 48, 2011, pp. 1–29.
- [5] Bhaskaran, S., Nandi, S., Broschart, S., Wallace, M., Cangahuala, L. A., and Olson, C., “Small body landings using autonomous onboard optical navigation,” *J. of the Astronautical Sciences*, Vol. 58, No. 3, 2011, pp. 1365–1378.
- [6] Gaskell, R. W., Barnouin-Jha, O. S., Scheeres, D. J., Konopliv, A. S., Mukai, T., Abe, S., Saito, J., Ishiguro, M., Kubota, T., Hashimoto, T., Kawaguchi, J., Yoshikawa, M., Shirakawa, K., Kominato, T., Hirata, N., and Demura, H., “Characterizing and navigating small bodies with imaging data,” *Meteoritics & Planetary Science*, Vol. 43, No. 6, 2008, pp. 1049–1061.
- [7] Raymond, C. A., Jaumann, R., Nathues, A., Sierks, H., Roatsch, T., Preusker, F., Scholten, F., Gaskell, R. W., Jorda, L., Keller, H.-U., Zuber, M. T., Smith, D. E., Mastrodemos, N., and Mottola, S., “The Dawn Topography Investigation,” *Space Science Rev.*, Vol. 163, No. 1, 2011, pp. 487–510.
- [8] Mastrodemos, N., Rush, B. P., Vaughan, A. T., and Owen, W. M., “Optical Navigation for the Dawn Mission at Vesta,” *Int. Symp. on Space Flight Dynamics (ISSFD)*, Pasadena, CA, USA, 2012, pp. 1–26.
- [9] Lorenz, D. A., Olds, R., May, A., Mario, C., Perry, M. E., Palmer, E. E., and Daly, M., “Lessons learned from OSIRIS-REx autonomous navigation using natural feature tracking,” *IEEE Aerospace Conf.*, Big Sky, MT, USA, 2017, pp. 1–12.
- [10] Barnouin, O. S., Daly, M. G., Palmer, E. E., Johnson, C. L., Gaskell, R. W., Al Asad, M., Bierhaus, E. B., Craft, K. L., Ernst, C. M., Espiritu, R. C., Nair, H., Neumann, G. A., Nguyen, L., Nolan, M. C., Mazarico, E., Perry, M. E., Philpott, L. C., Roberts, J. H., Steele, R. J., Seabrook, J., Susorney, H. C. M., Weirich, J. R., and Lauretta, D. S., “Digital terrain mapping by the OSIRIS-REx mission,” *Planetary and Space Science*, Vol. 180, No. 104764, 2020, pp. 1–16.
- [11] Gaskell, R., Barnouin, O., Daly, M., Palmer, E., Weirich, J., Ernst, C., Daly, R., and Lauretta, D., “Stereophotoclinometry on the OSIRIS-REx Mission: Mathematics and Methods,” *The Planetary Science J.*, Vol. 4, No. 4, 2023.
- [12] Quadrelli, M., Wood, L., Riedel, J., McHenry, M., Aung, M., Cangahuala, L., Volpe, R., Beauchamp, P., and Cutts, J., “Guidance, Navigation, and Control Technology Assessment for Future Planetary Science Missions,” *J. of Guidance, Control, and Dynamics*, Vol. 38, No. 7, 2015, pp. 1165–1186.
- [13] Nesnas, I., Hockman, B. J., Bandopadhyay, S., Morrell, B. J., Lubey, D. P., Villa, J., Bayard, D. S., Osmundson, A., Jarvis, B., Bersani, M., and Bhaskaran, S., “Autonomous Exploration of Small Bodies Toward Greater Autonomy for Deep Space Missions,” *Frontiers in Robotics and AI*, Vol. 8, No. 650885, 2021, pp. 1–26.
- [14] Getzandanner, K. M., Antreasian, P. G., Moreau, M. C., Leonard, J. M., Adam, C. D., Wibben, D., Berry, K., Highsmith, D., and Lauretta, D., “Small Body Proximity Operations & TAG: Navigation Experiences & Lessons Learned from the OSIRIS-REx Mission,” *AIAA SciTech Forum*, San Diego, CA, USA, 2022, pp. 1–23.
- [15] Dor, M., Driver, T., Getzandanner, K., and Tsiotras, P., “Visual SLAM for Asteroid Relative Navigation,” *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops*, Virtual, 2021, pp. 2066–2075.
- [16] Dor, M., Skinner, K., Driver, T., and Tsiotras, P., “AstroSLAM: Autonomous Monocular Navigation in the Vicinity of a Celestial Small Body—Theory and Experiments,” *arXiv:2212.00350*, 2022.
- [17] Dellaert, F., and Kaess, M., “Factor Graphs for Robot Perception,” *Found. Trends Robotics*, Vol. 6, No. 1-2, 2017, pp. 1–139.
- [18] Driver, T., Skinner, K., Dor, M., and Tsiotras, P., “AstroVision: Towards Autonomous Feature Detection and Description for Missions to Small Bodies Using Deep Learning,” *Special Issue on AI for Space, Acta Astronautica*, Vol. 210, 2023.
- [19] Kaess, M., Johannsson, H., Roberts, R., Ila, V., Leonard, J., and Dellaert, F., “iSAM2: Incremental Smoothing and Mapping Using the Bayes Tree,” *Intl. J. of Robotics Research (IJRR)*, Vol. 31, No. 2, 2012, pp. 217–236.

- [20] Edstedt, J., Athanasiadis, I., Wadenbäck, M., and Felsberg, M., “DKM: Dense kernelized feature matching for geometry estimation,” *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, Canada, 2023, pp. 17765–17775.
- [21] Chirikjian, G. S., *Stochastic Models, Information Theory, and Lie Groups: Analytic Methods and Modern Applications*, Vol. 2, Springer Science+Business Media, 2011.
- [22] Forster, C., Carlone, L., Dellaert, F., and Scaramuzza, D., “On-manifold preintegration for real-time visual–inertial odometry,” *IEEE Trans. on Robotics (T-RO)*, Vol. 33, No. 1, 2016, pp. 1–21.
- [23] Absil, P.-A., Baker, C. G., and Gallivan, K. A., “Trust-region methods on Riemannian manifolds,” *Foundations of Computational Mathematics*, Vol. 7, 2007, pp. 303–330.
- [24] Dellaert, F., “Factor graphs and GTSAM: A hands-on introduction,” Tech. Rep. GT-RIM-CP&R-2012-002, Georgia Institute of Technology, 2012.
- [25] Bercovici, B., and McMahon, J. W., “Robust Autonomous Small-Body Shape Reconstruction and Relative Navigation Using Range Images,” *J. of Guidance, Control, and Dynamics*, Vol. 42, No. 7, 2019, pp. 1473–1488.
- [26] Nakath, D., Clemens, J., and Rachuy, C., “Active Asteroid-SLAM,” *J. of Intelligent and Robotic Systems (JINT)*, Vol. 99, 2020, pp. 303–333.
- [27] Church, E., Bourbeau, T., Curriden, J., Deguzman, A., Jaen, F., Ma, H., Mahoney, K., Miller, C., Short, B., Waldorff, K., et al., “Flash Lidar On-Orbit Performance at Asteroid Benu,” *AAS Guidance, Navigation and Control (GN&C) Conf.*, Virtual, 2020.
- [28] Leonard, J. M., Moreau, M. C., Antreasian, P. G., Getzandanner, K. M., Church, E., Miller, C., Daly, M. G., Barnouin, O. S., and Lauretta, D. S., “Cross-Calibration of GNC and OLA LIDAR Systems Onboard OSIRIS-REx,” *AAS Guidance, Navigation and Control (GN&C) Conf.*, Breckenridge, CO, USA, 2022.
- [29] Daly, M. G., Barnouin, O. S., Dickinson, C., Seabrook, J., Johnson, C. L., Cunningham, G., Haltigin, T., Gaudreau, D., Brunet, C., Aslam, I., Taylor, A., Bierhaus, E. B., Boynton, W., Nolan, M., and Lauretta, D. S., “The OSIRIS-REx Laser Altimeter (OLA) Investigation and Instrument,” *Space Science Rev.*, Vol. 212, No. 924, 2017, pp. 899–924.
- [30] Rizk, B., d’Aubigny, C. D., Golish, D., Fellows, C., Merrill, C., Smith, P., Walker, M., Hendershot, J., Hancock, J., Bailey, S., DellaGiustina, D., Lauretta, D., Tanner, R., Williams, M., Harshman, K., Fitzgibbon, M., Verts, W., Chen, J., Connors, T., Hamara, D., Dowd, A., Lowman, A., Dubin, M., Burt, R., Whiteley, M., Watson, M., McMahon, T., Ward, M., Booher, D., Read, M., Williams, B., Huntten, M., Little, E., Saltzman, T., Alfred, D., O’Dougherty, S., Walthall, M., Kenagy, K., Peterson, S., Crowther, B., Perry, M., See, C., Selznick, S., Sauve, C., Beiser, M., Black, W., Pfisterer, R., Lancaster, A., Oliver, S., Oquest, C., Crowley, D., Morgan, C., Castle, C., Dominguez, R., and Sullivan, M., “OCAMS: The OSIRIS-REx Camera Suite,” *Space Science Rev.*, Vol. 214, No. 26, 2018, pp. 1–55.
- [31] Rublee, E., Rabaud, V., Konolige, K., and Bradski, G., “ORB: An efficient alternative to SIFT or SURF,” *IEEE Int. Conf. on Computer Vision (ICCV)*, Barcelona, Spain, 2011, pp. 2564–2571.
- [32] Rosten, E., and Drummond, T., “Machine learning for high-speed corner detection,” *European Conf. on Computer Vision (ECCV)*, Graz, Austria, 2006, pp. 430–443.
- [33] Calonder, M., Lepetit, V., Strecha, C., and Fua, P., “BRIEF: Binary Robust Independent Elementary Features,” *European Conf. on Computer Vision (ECCV)*, Heraklion, Crete, 2010, pp. 778–792.
- [34] Woodham, R. J., “Photometric method for determining surface orientation from multiple images,” *Optical Engineering*, Vol. 19, No. 1, 1980, p. 1309–1332.
- [35] Hayakawa, H., “Photometric stereo under a light source with arbitrary motion,” *J. of the Optical Society of America (JOSA)*, Vol. 11, No. 11, 1994, pp. 3079–3089.
- [36] Logothetis, F., Mecca, R., and Cipolla, R., “A differential volumetric approach to multi-view photometric stereo,” *IEEE Int. Conf. on Computer Vision (ICCV)*, Seoul, Korea, 2019, pp. 1052–1061.
- [37] Shi, B., Tan, P., Matsushita, Y., and Ikeuchi, K., “Bi-polynomial modeling of low-frequency reflectances,” *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, Vol. 36, No. 6, 2013, pp. 1078–1091.
- [38] Ikehata, S., and Aizawa, K., “Photometric stereo using constrained bivariate regression for general isotropic surfaces,” *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, Columbus, OH, USA, 2014, pp. 2179–2186.

- [39] Cho, D., Matsushita, Y., Tai, Y.-W., and Kweon, I. S., “Semi-calibrated photometric stereo,” *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, Vol. 42, No. 1, 2018, pp. 232–245.
- [40] Santo, H., Samejima, M., Sugano, Y., Shi, B., and Matsushita, Y., “Deep photometric stereo network,” *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops*, Honolulu, HI, USA, 2017, pp. 501–509.
- [41] Chen, G., Han, K., and Wong, K.-Y. K., “PS-FCN: A flexible learning framework for photometric stereo,” *European Conf. on Computer Vision (ECCV)*, Munich, Germany, 2018, pp. 3–18.
- [42] Bi, S., Xu, Z., Sunkavalli, K., Kriegman, D., and Ramamoorthi, R., “Deep 3D capture: Geometry and reflectance from sparse multi-view images,” *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, Virtual, 2020, pp. 5960–5969.
- [43] Kaya, B., Kumar, S., Oliveira, C., Ferrari, V., and Van Gool, L., “Multi-View Photometric Stereo Revisited,” *IEEE/CVF Winter Conf. on Applications of Computer Vision (WACV)*, Waikoloa, HI, USA, 2023, pp. 3126–3135.
- [44] Ackermann, J., Goesele, M., et al., “A survey of photometric stereo techniques,” *Foundations and Trends in Computer Graphics and Vision*, Vol. 9, No. 3-4, 2015, pp. 149–254.
- [45] Ju, Y., Lam, K.-M., Xie, W., Zhou, H., Dong, J., and Shi, B., “Deep learning methods for calibrated photometric stereo and beyond: A survey,” *arXiv:2212.08414*, 2022.
- [46] Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R., “NeRF: Representing scenes as neural radiance fields for view synthesis,” *Communications of the ACM*, Vol. 65, No. 1, 2021, pp. 99–106.
- [47] Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., and Leonard, J. J., “Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age,” *IEEE Trans. on Robotics (T-RO)*, Vol. 32, No. 6, 2016, pp. 1309–1332.
- [48] Baid*, A., Lambert*, J., Driver*, T., Krishnan*, A., Stepanyan, H., and Dellaert, F., “Distributed Global Structure-from-Motion with a Deep Front-End,” *arXiv:2311.18801*, 2023. *Authors contributed equally to this work.
- [49] He, K., Zhang, X., Ren, S., and Sun, J., “Deep residual learning for image recognition,” *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, 2016, pp. 770–778.
- [50] Nistér, D., “An efficient solution to the five-point relative pose problem,” *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, Vol. 26, No. 6, 2004, pp. 756–770.
- [51] Barath, D., and Matas, J., “Graph-Cut RANSAC,” *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, UT, USA, 2018, pp. 6733–6741.
- [52] Forsyth, D., and Ponce, J., *Computer Vision: A Modern Approach*, Prentice Hall, 2011.
- [53] Dellaert, F., Rosen, D. M., Wu, J., Mahony, R. E., and Carlone, L., “Shonan Rotation Averaging: Global Optimality by Surfing $SO(p)^n$,” *European Conf. on Computer Vision (ECCV)*, Glasgow, UK, 2020, pp. 292–308.
- [54] Wilson, K., and Snavely, N., “Robust Global Translations with 1DSfM,” *European Conf. on Computer Vision (ECCV)*, Zurich, Switzerland, 2014, pp. 61–75.
- [55] Hartley, R., and Zisserman, A., *Multiple View Geometry in Computer Vision*, Cambridge University Press, 2003.
- [56] Sampson, P. D., “Fitting conic sections to “very scattered” data: An iterative refinement of the Bookstein algorithm,” *Computer Graphics and Image Processing*, Vol. 18, No. 1, 1982, pp. 97–108.
- [57] Schröder, S., Mottola, S., Keller, H., Raymond, C., and Russell, C., “Resolved photometry of Vesta reveals physical properties of crater regolith,” *Planetary and Space Science*, Vol. 85, 2013, pp. 198–213.
- [58] Mildenhall*, B., Srinivasan*, P. P., Tancik*, M., Barron, J. T., Ramamoorthi, R., and Ng, R., “NeRF: Representing scenes as neural radiance fields for view synthesis,” *European Conf. on Computer Vision (ECCV)*, Virtual, 2020. *Authors contributed equally to this work.
- [59] Shepard, M. K., *Introduction to Planetary Photometry*, Cambridge University Press, 2017.
- [60] Li, J.-Y., Helfenstein, P., Buratti, B. J., Takir, D., and Clark, B. E., “Asteroid Photometry,” *Asteroids IV*, University of Arizona Press, 2015.

- [61] Horn, B. K., *Robot Vision*, The MIT Press, 1986.
- [62] Sierks, H., Keller, H., Jaumann, R., Michalik, H., Behnke, T., Bubenhausen, F., Büttner, I., Carsenty, U., Christensen, U., Enge, R., et al., “The Dawn framing camera,” *Space Science Rev.*, Vol. 163, 2011, pp. 263–327.
- [63] Schröder, S., Maue, T., Marqués, P. G., Mottola, S., Aye, K., Sierks, H., Keller, H., and Nathues, A., “In-flight calibration of the Dawn Framing Camera,” *Icarus*, Vol. 226, No. 2, 2013, pp. 1304–1317.
- [64] Schröder, S., Mottola, S., Matz, K.-D., and Roatsch, T., “In-flight calibration of the Dawn Framing Camera II: Flat fields and stray light correction,” *Icarus*, Vol. 234, 2014, pp. 99–108.
- [65] McEwen, A. S., “Photometric functions for photoclinometry and other applications,” *Icarus*, Vol. 92, No. 2, 1991, pp. 298 – 311.
- [66] McEwen, A. S., “A precise Lunar photometric function,” *Lunar and Planetary Science*, Vol. 27, 1996, pp. 841–842.
- [67] Alexandrov, O., and Beyer, R. A., “Multiview shape-from-shading for planetary images,” *Earth and Space Science*, Vol. 5, No. 10, 2018, pp. 652–666.
- [68] Bierhaus, E. B., Clark, B. C., Harris, J. W., Payne1, K. S., Dubisher, R. D., Wurts, D. W., Hund, R. A., Kuhns, R. M., Linn, T. M., Wood, J. L., May, A. J., Dworkin, J. P., Beshore, E., and Lauretta, D. S., “The OSIRIS-REx Spacecraft and the Touch-and-Go Sample Acquisition Mechanism (TAGSAM),” *Space Science Rev.*, Vol. 214, No. 107, 2018, pp. 1–46.
- [69] Horn, B. K., “Height and gradient from shading,” *Int. J. of Computer Vision (IJCV)*, Vol. 5, No. 1, 1990, pp. 37–75.
- [70] Capanna, C., Gesquière, G., Jorda, L., Lamy, P., and Vibert, D., “Three-dimensional reconstruction using multiresolution photoclinometry by deformation,” *The Visual Computer*, Vol. 29, 2013, pp. 825–835.
- [71] Russell, C., and Raymond, C., *The Dawn mission to minor planets 4 Vesta and 1 Ceres*, Springer Science & Business Media, 2012.
- [72] NASA, “Planetary Data System (PDS),” <https://pds.nasa.gov/>, 2022.
- [73] Li, J.-Y., “Body-Fixed Coordinate Systems for Asteroid (4) Vesta,” https://sbnarchive.psi.edu/pds3/dawn/vir/DWNVVIR_I1B_v2/DOCUMENT/VESTA_COORDINATES/VESTA_COORDINATES_131018.HTM, 2012.
- [74] Denevi, B. W., Blewett, D., Buczkowski, D., Capaccioni, F., Capria, M., De Sanctis, M., Garry, W., Gaskell, R., Le Corre, L., Li, J.-Y., et al., “Pitted terrain on Vesta and implications for the presence of volatiles,” *Science*, Vol. 338, No. 6104, 2012, pp. 246–249.
- [75] Williams, D. A., O’Brien, D. P., Schenk, P. M., Denevi, B. W., Carsenty, U., Marchi, S., Scully, J. E., Jaumann, R., De Sanctis, M. C., Palomba, E., et al., “Lobate and flow-like features on asteroid Vesta,” *Planetary and Space Science*, Vol. 103, 2014, pp. 24–35.
- [76] Briechle, K., and Hanebeck, U. D., “Template matching using fast normalized cross correlation,” *Optical Pattern Recognition XII*, Vol. 4387, SPIE, 2001, pp. 95–102.
- [77] Xiong, Y., Olson, C. F., and Matthies, L. H., “Computing depth maps from descent images,” *Machine Vision and Applications*, Vol. 16, 2005, pp. 139–147.
- [78] Zinßer, T., Schmidt, J., and Niemann, H., “Point set registration with integrated scale estimation,” *Int. Conf. on Pattern Recognition and Information Processing (PRIP)*, Minsk, Belarus, 2005, pp. 116–119.
- [79] Golish, D., Li, J.-Y., Clark, B., DellaGiustina, D., Zou, X.-D., Rizos, J., Hasselmann, P., Bennett, C., Fornasier, S., Drouet d’Aubigny, C., et al., “Regional Photometric Modeling of Asteroid (101955) Bennu,” *The Planetary Science J.*, Vol. 124, No. 2, 2021, pp. 298 – 311.