AstroSplat: Physics-Based Gaussian Splatting for Rendering and Reconstruction of Small Celestial Bodies

Anonymous CVPR submission

Paper ID *****

Abstract

Image-based surface reconstruction and characterization are crucial for missions to small celestial bodies (e.g., asteroids), as it informs mission planning, navigation, and scientific analysis. Recent advances in Gaussian splatting enable high-fidelity neural scene representations but typically rely on a spherical harmonic intensity parameterization that is strictly appearance-based and does not explicitly model material properties or light-surface interactions. We introduce AstroSplat, a physics-based Gaussian splatting framework that integrates planetary reflectance models to improve the autonomous reconstruction and photometric characterization of small-body surfaces from in-situ imagery. The proposed framework is validated on real imagery taken by NASA's Dawn mission, where we demonstrate superior rendering performance and surface reconstruction accuracy compared to the typical spherical harmonic parameterization.

1. Introduction

Image-based surface reconstruction and photometric characterization are fundamental in the exploration of small airless bodies such as asteroids, comets, and moons. High-fidelity surface maps enable autonomous navigation, landing-site selection, hazard detection and avoidance, and scientific discovery. Though critical to these operations, the extreme illumination variability and similar surface features encountered in the space environment make accurate geometric and appearance modeling challenging.

The traditional approach to reconstructing small body topography is stereophotoclinometry (SPC) [11, 24]. The SPC process combines stereophotogrammetry (SPG) with photoclinometry to align captured images with an existing shape model and estimate *digital terrain maps* (DTMs), local topography and albedo maps. Specifically, SPC combines a reflectance model with accurate *a priori* estimates of the camera pose, illumination direction, surface topography, and albedo to generate renders of a body's surface and employs correlation-based matching to register captured images to the DTM. This registration process is repeated for multiple images to generate dense correspondences, which is then followed by a photoclinometry step that refines surface topography and albedo estimates. Refinements are made to those *a priori* values until convergence criteria is met between the ground truth and rendered images. While SPC has been applied extensively to mapping various small bodies in the solar system, including Mars's moon Phobos [10] and the asteroids 101955 Bennu [2], 4 Vesta [20, 27], and 25143 Itokawa [11], it is limited by its reliance on high-quality *a priori* information and human operators to correct alignment errors. These limitations restrict scalability and constrain the level of autonomy achievable for future applications.

Recent advances in neural scene representations, particularly Gaussian splatting [17], have demonstrated impressive real-time rendering and reconstruction capabilities from multi-view imagery. Namely, the 2D Gaussian splatting (2DGS) [15] variant achieves high-fidelity view synthesis and topographic estimation by representing a scene with a collection of Gaussian disks. The reflectance properties of each Gaussian are typically represented by spherical harmonic (SH) coefficients. However, this parameterization is fundamentally appearance-based: it encodes view-dependent intensity without explicitly modeling the underlying physics of light–surface interactions. As a result, geometry and albedo effects are entangled, resulting in parameter estimates that may attribute illumination variance to albedo instead of topography. To combat this effect, we propose *AstroSplat*, a physics-based Gaussian splatting framework that applies photometric modeling in the intensity computation step. Instead of the purely appearance-based SH parameterization, our method uses planetary reflectance models, which introduce a dependence on illumination direction and surface normal to determine per-Gaussian intensity.

We validate AstroSplat on real imagery acquired by NASA's Dawn mission [28] to the asteroid 4 Vesta and

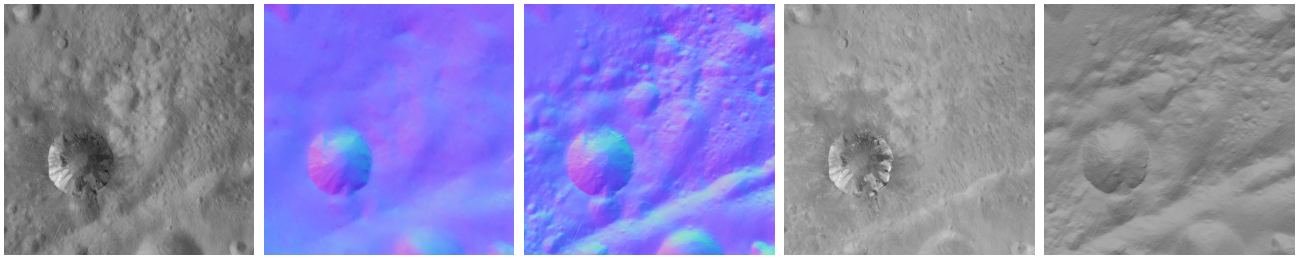

(a) Input Image (b) SH Normals (c) Proposed Normals (d) Proposed Albedos (e) Proposed Mesh

Figure 1. **Our proposed AstroSplat framework compared to the traditional spherical harmonic (SH) parameterization.** The SH parametrization results in (b) smoothed normals maps, while the physics-based reflectance modeling of AstroSplat yields more detailed (c) surface normals, (d) albedos, and (e) meshes.

minor planet 1 Ceres. Experiments demonstrate that incorporating physics-based reflectance models yields superior rendering performance and improved surface topography accuracy compared to the standard SH parameterization (see Figure 1). The result is an autonomous framework that eliminates the need for the human-intensive process of SPC while yielding superior reconstruction quality to traditional 2DGS. In summary, the main contributions of this paper are:

- We integrate *physics-based* planetary reflectance models into the 2DGS framework for topography estimation and rendering of small bodies.
- We analyze the effects of different reflectance models on reconstruction and rendering performance.
- We demonstrate that these reflectance models yield better performance compared to the typical spherical harmonic parameterization.

2. Related Work

To reduce reliance on the human-in-the-loop processes of SPC, pipelines that leverage deep learning have been developed to reconstruct small bodies from imagery. Driver et al. [8] proposed Photoclinometry-from-Motion (PhoMo), an autonomous framework which incorporates photoclinometry techniques into a keypoint-based structure-from-motion (SfM) system for normal and albedo estimation of small celestial bodies. PhoMo uses a deep learning feature detection and matching method [7, 9] as well as the formalism of factor graphs [6] to simultaneously estimate feature positions, camera poses, surface normals, and albedos. This approach shows superior rendering and reconstruction quality over SPC while being fully autonomous and without requiring prior information.

In the absence of explicit feature-matching algorithms, neural scene representations have gained popularity for use in space-based applications. Small body reconstructions have been achieved through the use of neural radiance fields (NeRFs) [4, 12] and their signed-distance function (SDF) [3] extensions. 3D Gaussian Splatting (3DGS) has recently emerged as an efficient alternative to NeRFs [17]. Although applications to asteroid imagery are nascent, early work such as Prosvetov et al. [26] applied 3DGS for detailed reconstruction of the lunar surface, and Aira et al. [1] adapts Gaussian splatting to pushbroom satellite sensors for Earth–observation photogrammetry. These studies demonstrate the promise of neural implicit methods for reconstructing irregular small body geometries. However, existing methods remain predominantly *appearance-based*, lacking explicit modeling of bidirectional reflectance and material properties. Our work builds on these trends by introducing a physics-aware Gaussian splatting framework that integrates planetary reflectance models, enabling autonomous, photometrically consistent reconstruction from *in situ* imagery.

3. Gaussian Splatting

Conventional 3DGS [17] represents a scene with a set of ellipsoids containing learned parameters that are iteratively optimized to minimize the discrepancy between ground truth images and corresponding renders of the scene. 2DGS [15] adjusts the 3DGS method by removing a dimension of the Gaussian representation, such that each ellipsoid flattens into a planar ellipse. In doing so, this representation allows for better geometric reconstructions by encouraging the Gaussians to align with surfaces in the scene. To accomplish this, the 2DGS method utilizes 5 types of learned parameters: each Gaussian's mean $\mathbf{p}_k$, scaling $\mathbf{s}_k$, orientation $\mathbf{q}_{\mathcal{W}\mathcal{S}_k}$, opacity α_k, and spherical harmonic coefficients that encode the view-dependent appearance $\mathbf{c}_k$. The index subscript k is hereon dropped for conciseness.

To begin, a set of n Gaussians are initialized in some world frame $\mathcal{W}$, which are either derived from an initial sparse point cloud (e.g., from SfM) or simply randomized. Each Gaussian has a locally defined frame, referred to as the local splat frame $\mathcal{S}$. The origin is at the Gaussian's center, and the x- and y-axes of this local splat frame

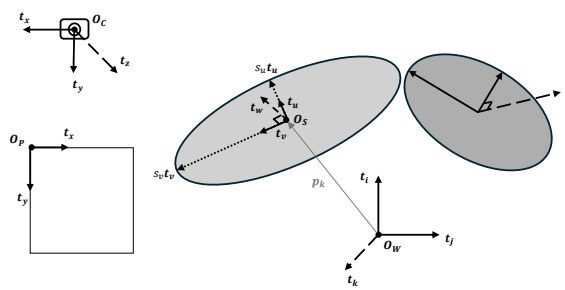

Figure 2. **2DGS frame definitions.** Relative orientation and positions of the local splat $\mathcal{S}$, world $\mathcal{W}$, camera $\mathcal{C}$, and pixel $\mathcal{P}$ frames. The origin of each frame is indicated by the point labeled O and the basis directions are defined by $\mathbf{t}$ vectors.

are defined by the principal directions of the ellipse, $\mathbf{t}_u$ and $\mathbf{t}_v$. The third direction completes the right-hand rule $\mathbf{t}_w = \mathbf{t}_u \times \mathbf{t}_v$, and this direction is, by construction, normal to the Gaussian plane itself. The local splat frame is also scaled along its $\mathbf{t}_u$ and $\mathbf{t}_v$ directions by the scalars s_u and s_v. The basis vectors of the local splat frame are expressed in the world frame such that they define the rotation from local splat to world coordinates by applying the rotation $\mathbf{R}_{\mathcal{WS}} = \begin{bmatrix}\mathbf{t}_u^{\mathcal{W}} & \mathbf{t}_v^{\mathcal{W}} & \mathbf{t}_w^{\mathcal{W}}\end{bmatrix}$. This work uses the passive representation of attitude [5, 36]. A quaternion encoding this rotation $\mathbf{q}_{\mathcal{WS}}$ represents the learned relative orientation of each Gaussian with respect to the world frame. The scaling $\mathbf{s}$ is encoded in the matrix $\mathbf{S} = \text{diag}([s_u,\ s_v,\ 0])$. Figure 2 provides a visual representation of these frames.

For all transformations, points in a given frame will be represented by homogeneous coordinates, in which a one is appended as weight to each 3D coordinate representation. Let $\mathbf{x}^{\mathcal{S}} = \begin{bmatrix}x^{\mathcal{S}} & y^{\mathcal{S}} & 0 & 1\end{bmatrix}^\top$ be the homogeneous coordinate for a point on a Gaussian expressed in its local splat frame. The third coordinate is zero due to the 2D nature of the Gaussian representation. In order to transform a point in this local splat frame to the 3D world frame, the following transformation is applied:

$$\mathbf{T}_{\mathcal{WS}} = \begin{bmatrix}\mathbf{R}_{\mathcal{WS}}\mathbf{S} & \mathbf{p}^{\mathcal{W}} \\ 0 & 1\end{bmatrix} = \begin{bmatrix}s_u\mathbf{t}_u^{\mathcal{W}} & s_v\mathbf{t}_v^{\mathcal{W}} & \mathbf{0}_{3\times1} & \mathbf{p}^{\mathcal{W}} \\ 0 & 0 & 0 & 1\end{bmatrix}. \tag{1}$$

To transform a point in $\mathcal{W}$ to the camera frame $\mathcal{C}$, whose x- and y-axes align with the image height and width directions and z-axis points boresight, the following transformation is applied:

$$\mathbf{T}_{\mathcal{CW}} = \begin{bmatrix}\mathbf{R}_{\mathcal{CW}} & \mathbf{r}_{WC}^{\mathcal{C}} \\ 0 & 1\end{bmatrix}, \tag{2}$$

where $\mathbf{R}_{\mathcal{CW}}$ is the rotation matrix from $\mathcal{W}$ to $\mathcal{C}$, and $\mathbf{r}_{WC}^{\mathcal{C}}$ is the position of the origin of $\mathcal{C}$ with respect to the origin of $\mathcal{W}$ expressed in $\mathcal{C}$. 2DGS assumes these camera poses are known *a priori*. After applying the transformation to the camera frame, homogeneous coordinates can be converted to pixel coordinates $\mathcal{P}$ by applying the conventional camera calibration matrix $\mathbf{K}$ with an appended column of zeros to account for the projection from the $\mathbb{P}^3$ to the $\mathbb{P}^2$ projective spaces. Let (d_x, d_y) denote the focal length in units of pixel length along both axes, and let (u_p, v_p) denote the pixel coordinates of the image center, such that this transformation is defined by

$$\mathbf{T}_{\mathcal{PC}} = \begin{bmatrix}d_x & 0 & u_p & 0 \\ 0 & d_y & v_p & 0 \\ 0 & 0 & 1 & 0\end{bmatrix} = \mathbf{K}\begin{bmatrix}\mathbf{I}_{3\times3} & \mathbf{0}_{3\times1}\end{bmatrix}. \tag{3}$$

The training process for 2DGS is broken into two stages: pre-processing and rendering. In the pre-processing stage, the transformation from each local splat frame to pixel frame is computed by applying the transformation defined above. This transformation is used to compute a bounding box around each Gaussian, which defines which pixels need to include contributions from this Gaussian in the rendering stage. Next, the view-dependent appearance of each Gaussian is determined using the reflectance model of choice. The original 2DGS implementation parameterizes the reflectance model using the coefficients of an mth degree SH expansion.

Once the parameters of each Gaussian are determined, the rendering stage begins, employing the same rasterization techniques used for 3DGS. This procedure requires ordering all Gaussians contributing to a given pixel by their depth. These depths are computed by finding the point where a ray emitted from the camera through the pixel coordinates, $\mathbf{x}^{\mathcal{P}} = (x, y)$, intersects each Gaussian in its path, $\mathbf{u}^{\mathcal{S}} = \mathbf{u}^{\mathcal{S}}(\mathbf{x}^{\mathcal{P}}) = (u, v)$. Then, the Gaussian value at this intersection is computed according to $\mathcal{G}(\mathbf{u}^{\mathcal{S}}) = \exp\left(-(u^2+v^2)/2\right)$. Sorting by the depths and starting from front to back, an iterative alpha-blending process is used to compute the intensity of each pixel $\bar{\mathbf{c}}(\mathbf{x}^{\mathcal{P}})$ by accumulating the view-dependent intensity c_i of each Gaussian according to its opacity α_i and Gaussian spread $\hat{\mathcal{G}}_i$

$$\bar{\mathbf{c}}(\mathbf{x}^{\mathcal{P}}) = \sum_{i=1} c_i\omega_i, \tag{4}$$

where $\omega_i = \alpha_i\hat{\mathcal{G}}_i(\mathbf{u}^{\mathcal{S}})\prod_{j=1}^{i-1}(1-\alpha_j\hat{\mathcal{G}}(\mathbf{u}^{\mathcal{S}}))$ and $\hat{\mathcal{G}}_i(\mathbf{u}^{\mathcal{S}}) = \max\left\{\mathcal{G}_i(\mathbf{u}^{\mathcal{S}}), \mathcal{G}_i(\sqrt{2}(\mathbf{u}^{\mathcal{P}} - \mathbf{p}^{P})\right\}$. Since small body imagery is monochrome, the computed intensity at each pixel is duplicated across the RGB channels to yield a grayscale image. This alpha-blending procedure is also used to accumulate a normal map and depth map by substituting the value c_i with the third column of the splat-to-world rotation matrix $R_{\mathcal{WS}}$ and the previously computed depth, respectively.

This rasterization step is conducted at each training iteration using a random camera pose from the set of training

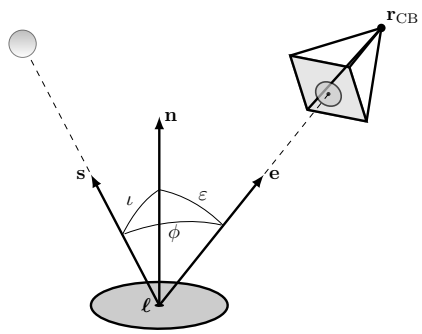

Figure 3. **Photometry conventions.** Photometric angles and their relationships to the Sun vector $\mathbf{s}$, emission vector $\mathbf{e}$, and normal $\mathbf{n}$ with respect to a local patch centered at $\boldsymbol{\ell}$.

data, and the render, normal map, and depth map products are used to define the loss function that the optimization scheme is seeking to minimize. This loss function has two components in our implementation: an intensity loss $\mathcal{L}_I$ and a normal loss $\mathcal{L}_n$. The intensity loss is a weighted sum of the L1 difference and the Structural Similarity Index Measure (SSIM)[35] between the RGB render and the ground truth image. The normal loss encourages the normal $\mathbf{n}_i$ defined by the third axis of $\mathbf{R}_{\mathcal{WS}}$ to align with the normal estimated from the gradient of the depth map $\mathbf{n}_{d_i}$ for every i-th pixel in the normal and depth maps. The total loss $\mathcal{L}$ is defined as a weighted sum of these individual loss terms:

$$\mathcal{L} = \mathcal{L}_I + \beta\mathcal{L}_n,\ \beta \in \mathbb{R}, \quad \text{where} \tag{5}$$

$$\mathcal{L}_I = (1-\lambda)\mathcal{L}_{\text{L1}} + \lambda\mathcal{L}_{\text{SSIM}},\ \lambda \in [0,1], \tag{6}$$

$$\mathcal{L}_n = \sum_i \omega_i (1 - \mathbf{n}_i^\top \mathbf{n}_{d_i}). \tag{7}$$

4. Small Body Photometry

4.1. Background

Photometry seeks to model the observed "brightness" of a scene point in an image as influenced by its surface topography and material properties. While photometry is inherently complex, operation in space presents us with a number of advantages that simplify the photometric modeling process. First, we assume the Sun is a point source (the Sun subtends an angle of $0.5°$ at Earth [32]) delivering approximately collimated light to the surface due to the large distance to the Sun 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) or *a priori* state knowledge. 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 [19].

More formally, an image may be considered as a mapping $I\colon \Omega \to \mathbb{R}_+$ over the pixel domain $\Omega \subset \mathbb{R}^2$ that maps a point in the image, $\mathbf{p} \in \Omega$, to its 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 ($\mathrm{W{\cdot}m^{-2}}$). Specifically, each brightness value $I(\mathbf{p})$ is initially represented by a digital number (DN), which is computed by converting the charge accumulated by the photodetector over some exposure time Δt to an integer DN using an analog-to-digital converter. For example, the Dawn mission to Asteroid 4 Vesta and Minor Planet 1 Ceres digitized the signal from the framing camera to a 14-bit integer DN [34]. An important requirement for deep space imagers is response linearity, i.e., an almost perfectly linear relationship between the incident irradiance on the detector and the quantized charge rate $\mathrm{DN}/\Delta t$ [5, Chapter 7]. This linear relationship is characterized by a rigorous *radiometric calibration* process conducted both on the ground and during flight [29, 31, 34]. It can be shown [14, Chapter 10.3] that the irradiance on the photodetector is proportional to the *radiance*, with units of power per unit solid angle per unit area ($\mathrm{W{\cdot}sr^{-1}{\cdot}m^{-2}}$), 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.

The emitted radiance from the surface, $L(\iota, \varepsilon, \phi, a)$ and the incident (collimated) irradiance from the Sun, F, which is inversely related to the square of the distance to the Sun [5, 32], are related by the *bidirectional reflectance function*, r (in units of sr^{-1}):

$$\frac{L(\iota, \varepsilon, \phi, a)}{F} = r(\iota, \varepsilon, \phi, a), \tag{8}$$

where a is the surface albedo, ι is the angle between the incoming light and the surface normal, or the *incidence angle*, ε 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 Fig. 3). A similar measure of reflectance, which is very popular in the context of planetary photometry [19, 32], is the *bidirectional radiance factor*, r_{F}, 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., $\iota = \varepsilon = \phi = 0$):

$$r_{\mathrm{F}}(\iota, \varepsilon, \phi, a) = \pi r(\iota, \varepsilon, \phi, a) = \frac{\pi L(\iota, \varepsilon, \phi, a)}{F}. \tag{9}$$

While the radiance factor is dimensionless, it is common to refer to its value as being in units of reflectance or L/F (or,

more commonly, I/F when I is used to denote the radiance). Henceforth, we will refer to the radiance factor when mentioning the reflectance function.

If we normalize the radiance factor by the radiance factor when observed and illuminated from overhead ($\iota = \varepsilon = \phi = 0°$), we get the *photometric function*, $\rho(\iota, \varepsilon, \phi)$:

$$\rho(\iota, \varepsilon, \phi) = \frac{r_{\mathrm{F}}(\iota, \varepsilon, \phi)}{r_{\mathrm{F}}(0, 0, 0)} \Leftrightarrow r_{\mathrm{F}}(\iota, \varepsilon, \phi) = a_{\mathrm{n}} \rho(\iota, \varepsilon, \phi), \tag{10}$$

where $a_{\mathrm{n}} := r_{\mathrm{F}}(0, 0, 0)$ is the *normal albedo*. Finally, the photometric function may be factorized into the *phase function*, $\Lambda(\phi)$, and the disk function, $d(\iota, \varepsilon, \phi)$ [33]:

$$r_{\mathrm{F}}(\iota, \varepsilon, \phi) = a_{\mathrm{n}} \Lambda(\phi) d(\iota, \varepsilon, \phi). \tag{11}$$

The phase function, normalized to unity at $\phi = 0°$, models the phase-dependent brightness variations that are independent of the incident and emission angles, referred to as the *opposition effect* [13, Chapter 9], where $\Lambda(\phi)$ is often represented as an nth-order polynomial with parameters fit to imagery of the specific target body [30, 31]. The disk function models brightness variations due to the underlying topography (which may also be a function of the phase angle). Henceforth, the term albedo will refer to the normal albedo and be denoted by a.

We consider the case where the image values have not been radiometrically calibrated to units of reflectance. Recalling the response linearity of the photodetector, we may rewrite Equation (9) as

$$\frac{\pi L(\iota, \varepsilon, \phi)}{F} \propto \frac{\mathrm{DN}(\iota, \varepsilon, \phi)}{\Delta t} - \xi = \tilde{a}_{\mathrm{n}} \lambda d(\iota, \varepsilon, \phi). \tag{12}$$

The scale term λ accounts for the phase dependent brightness in the absence of an explicit phase function, the incident solar flux, and the lack of radiometric calibration to radiance. The bias term ξ may be included to account for possible background noise [11, 25]. The *relative* albedo $\tilde{a}_{\mathrm{n}}$ is only proportional to the absolute albedo since compensation for the lack of a radiometric calibration may also be expressed through scaling of the albedos.

4.2. Photometry Implementation in 2DGS

The 2DGS SH reflectance model uses only the viewing direction to fit the coefficients of an SH expansion for intensity computation. Physics-based reflectance models can additionally use light direction (e.g., from a sun sensor measurement) and surface normal in this computation step to estimate pixel intensity. Therefore, replacing the SH approach with a physics-based photometric function adds more geometric information to the estimation of image intensity, intrinsically tying the render to the underlying surface topography.

Photometric functions shown in Equation 12 are defined by their disk functions $d(\iota, \varepsilon, \phi)$. To exemplify the benefit of photometric modeling in this framework, we employ three reflectance models: Lambert, Lommel-Seeliger, and Lunar-Lambert. The *Lambert* model assumes that a surface is perfectly diffuse, such that light scatters equally in all directions. As such, it is only a function of the illumination direction and surface normal,

$$d_{\mathrm{L}}(\iota) = \cos(\iota). \tag{13}$$

The *Lommel-Seeliger* (L-S) model considers that the light reflected by the surface experiences single-scattering along a specified viewing direction, yielding a dependence on illumination direction, viewing direction, and surface normal,

$$d_{\text{L-S}}(\iota, \varepsilon, \phi) = \frac{2\cos(\iota)}{\cos(\iota) + \cos(\varepsilon)}. \tag{14}$$

The *Lunar-Lambert* model[1] is a weighted sum of the Lambert and Lommel-Seeliger models [21, 22], which uses a phase weighting function $g(\phi)$, where ϕ is the phase angle in degrees, to adjust their relative contributions,

$$d_{\text{L-L}}(\iota) = (1 - g(\phi)) \cos(\iota) + g(\phi) \frac{2\cos(\iota)}{\cos(\iota) + \cos(\varepsilon)}, \tag{15}$$

where we use the exponential form of the phase weighting function g proposed by Gaskell et al. [11], i.e., $g(\phi) = \exp(-\phi/60°)$. Derivations of these functions and a discussion of their properties can be found in [5].

In order to apply these reflectance functions to the view-dependent intensity computation step, we add a per-Gaussian relative albedo $\tilde{a}_i$, which captures material property variation across the surface, and a per-image scale and bias correction factor (similar to the learned exposure compensation in hierarchical 3DGS [18]), which accounts for the use of uncalibrated PNG image data.

5. Experimental Setup

To demonstrate the impact of implementing different photometric functions, we utilize three datasets, which were captured by NASA's Dawn mission during its observation of the Cornelia crater on 4 Vesta, as well as the Ikapati crater and Ahuna Mons on 1 Ceres [8, 23, 28, 34]. The train/test split is indicated in Table 1 and Table 2. We trained for 30,000 iterations using the default parameters. Each experiment was initialized with camera poses estimated by SPC and a point cloud constructed with the dense triangulated RoMa [9] keypoints from PhoMo.

In order to quantify the performance of each reflectance model, metrics for both the quality of the render and of

[1] The Lunar-Lambert model was referred to as the McEwen model in the PhoMo study [8] used as the ground truth in this work.

the surface reconstruction were used. To evaluate render quality, we compare the peak signal-to-noise ratio (PSNR), the structural similarity index measure (SSIM) [35], and the learned perceptual image patch similarity (LPIPS) [37]. To evaluate reconstruction quality, we compare the normal and albedo maps produced by the rasterization procedure to values derived from PhoMo results, which will be treated as the ground truth. The normal error is defined as the total angular difference between the ground truth $\mathbf{n}_{ij}$ and rendered normals $\tilde{\mathbf{n}}_{ij}$ at each corresponding pixel in the normal maps: $\delta\theta_{ij} = \cos^{-1}\left(\tilde{\mathbf{n}}_{ij}^{\top}\mathbf{n}_{ij}\right)$. The albedo error is defined as the relative difference between the ground truth albedo a_{ij} and rendered albedo $\tilde{a}_{ij}$, after being linearly fit to the ground truth, at each corresponding pixel in the albedo maps: $\delta a_{ij} = |a_{ij} - \tilde{a}_{ij}|/a_{ij}$. Both the average normal error $\delta\theta$ and average albedo error δa across the overlapping regions of the PhoMo and AstroSplat results for data set are used to quantify the quality of the reconstruction. Finally, the learned products from 2DGS can be exported as triangular meshes using a truncated signed-distance fusion formulation. These meshes are aligned to the PhoMo ground truth point cloud using an iterative closest point algorithm, and the Hausdorff distance [16] normalized by the size of the mesh , d_H, is calculated.

6. Results

A set of ground truth test images as well as the associated test renders for each reflectance model are shown in Figure 4. The overall PSNR, SSIM, and LPIPS values for both train and test image sets are provided in Table 1. Across all datasets and all render metrics, the physics-based reflectance functions consistently outperformed the SH approach. For nearly all test images, the Lambert reflectance model yielded the best results, with Lunar-Lambert yielding the next best performance.

Beyond rendering quality, the normal maps and meshes created from the 2DGS products are used to determine the surface reconstruction quality. The ground truth images, PhoMo ground truth normal maps, and associated AstroSplat normal maps for a test set of each model are shown in Figure 5. The average normal error $\delta\theta$ and symmetric Hausdorff distance d_H metrics over the ground truth region are shown in Table 2 for all reflectance models. All physics-based reflectance models yielded lower mean normal errors compared to the original SH method, with the Lommel-Seeliger model minimizing the error for all datasets. Looking beyond the cropped region in each normal map, it is clear that several minor topographical features, such as smaller craters and ridges, appear in the physics-based approaches whereas those same areas are smoothed out by the SH approach. When comparing the ground truth PhoMo point cloud to the cropped mesh product for each approach, the symmetric Hausdorff distance was also consistently minimized by a physics-based approach.

The use of physics-based reflectance models also adds the benefit estimating albedo, which can be used to assess the quality of the surface reconstruction. The ground truth images, PhoMo ground truth albedo maps, and associated AstroSplat albedo maps for a test set of each model are shown in Figure 6. The mean relative albedo error δa for each physics-based approach for all datasets is shown in Table 2. The Lambert model provides the closest albedo to the ground truth for Cornelia, yielding about a 20% and 30% reduction in error compared to Lunar-Lambert and Lommel-Seeliger, respectively. The difference in relative errors between the models for Ikapati and Ahuna Mons are on the order of thousandths. The underestimation of higher albedo values on the Ceres features yields less variation in albedo value compared to the PhoMo results, resulting in a consistent error across models, especially compared to Cornelia.

With superior rendering quality, more accurate surface normal estimation, and the capability to estimate albedo, the physics-based reflectance models outperform the traditional 2DGS approach on all metrics. Furthermore, these models require fewer parameters than the SH method, which learns 16 parameters for the standard third-degree expansion. The Lambert, Lommel-Seeliger, and Lunar-Lambert models require only one intensity-specific optimized parameter, namely albedo, and leverage the available illumination information and the normal estimate from the Gaussian orientation to accomplish the same task with better results.

7. Conclusion

Deep learning strategies, like Gaussian splatting, have demonstrated promising results for surface reconstructions. However, in the space environment, physics-based methods have mission-proven success and yield reliable solutions, which make them desirable when mapping asteroids and minor planets. This paper introduces AstroSplat, a 2D Gaussian splatting framework that incorporates physics-based small body photometry into the intensity computation step to improve normal estimation of craters and other surface features. Instead of using the traditional Gaussian splatting technique of SH-based appearance-encoding, AstroSplat leverages available illumination information, viewing geometry, and surface normal estimates as well as a new per-Gaussian learned albedo to enable photometric modeling. This approach allows reflectance functions—like the Lambert, Lommel-Seeliger, and Lunar-Lambert models—to prescribe each Gaussian's intensity based on the behavior of light interacting with the topography of the surface, inherently linking the observed brightness in a render to the underlying normal estimates. This adjusted color computation step is validated on real imagery from NASA's Dawn mission and compared to a state-of-the-art surface recon-

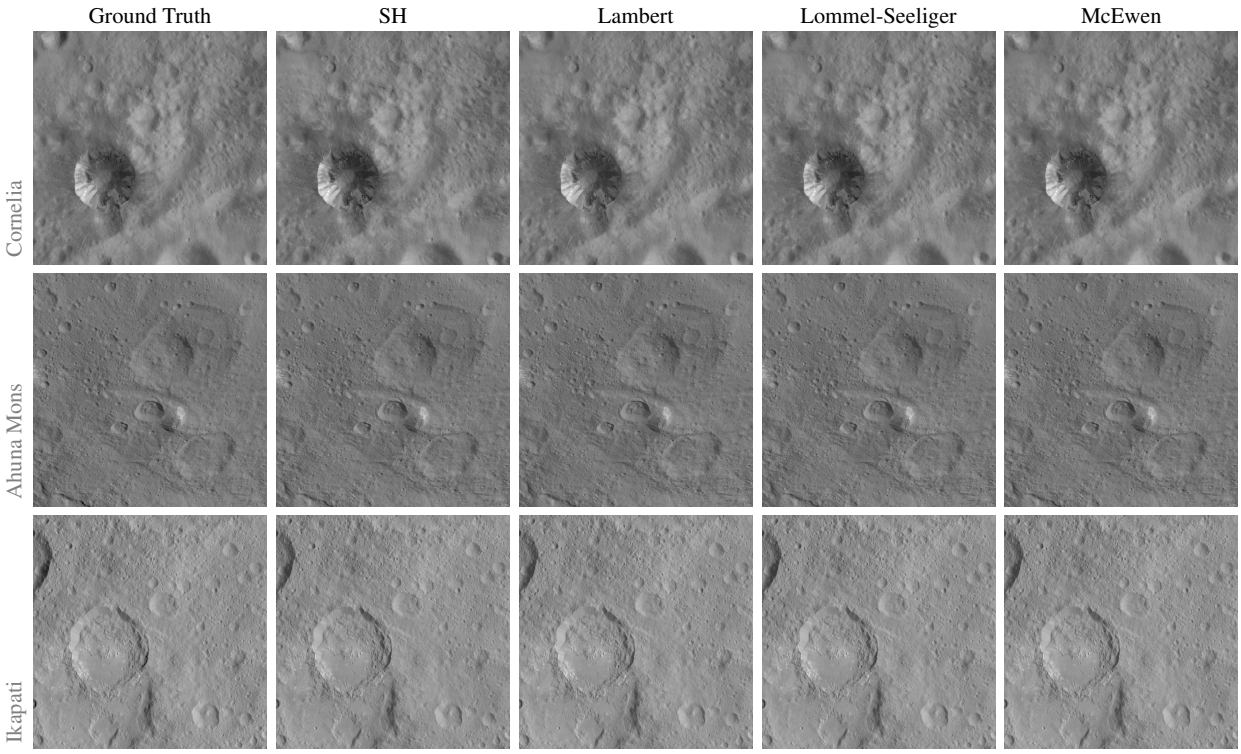

Figure 4. Qualitative comparison of test image renderings for each reflectance model.

Table 1. **Rendering metrics for each reflectance model.** Each entry is of the form [train_value]/[test_value]. The **first** and second best values are bold and underlined, respectively. The number of train and test images used are shown next to the body name.

	Cornelia (27/2)			Ahuna Mons (30/2)			Ikapati (27/2)		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
SH	36.59/35.17	0.9630/0.9614	0.1371/0.1280	37.83/36.73	0.9629/0.9561	0.0685/0.0738	38.46/36.00	0.9805/0.9656	0.0416/0.0651
Lambert	39.23/**39.32**	**0.9702/0.9689**	**0.1091/0.1017**	36.58/**37.93**	**0.9653/0.9587**	**0.0630/0.0687**	39.08/36.93	**0.9837/0.9683**	**0.0348/0.0597**
L-S	38.86/35.68	0.9687/0.9608	0.1174/0.1189	**38.10**/37.10	0.9642/0.9551	0.0660/0.0761	38.68/36.62	0.9795/0.9649	0.0427/0.0656
Lunar-Lambert	**40.15**/38.35	0.9695/0.9651	0.1149/0.1101	37.21/37.88	0.9634/0.9569	0.0654/0.0718	**39.36/37.06**	0.9823/0.9673	0.0389/0.0636

Table 2. **Reconstruction metrics for each reflectance model.** Each entry of the form [train_value]/[test_value]. The **first** and second best values are bold and underlined, respectively. See Section 5 for metric definitions. The number of train and test images used are shown next to the body name.

	Cornelia (27/2)			Ahuna Mons (30/2)			Ikapati (27/2)		
	$\delta\theta$ (°) ↓	δa ↓	d_H ↓	$\delta\theta$ (°) ↓	δa ↓	d_H ↓	$\delta\theta$ (°) ↓	δa ↓	d_H ↓
SH	7.4510/7.4873	–/–	0.0032	10.3486/10.4027	–/–	0.0061	8.4515/8.2710	–/–	0.0033
Lambert	6.4468/6.2974	**0.0278/0.0474**	**0.0015**	9.3419/9.5286	0.0334/**0.0310**	0.0053	7.4277/7.3510	0.0331/0.0295	**0.0021**
L-S	**6.3027/6.0265**	0.0409/0.0672	0.0021	**8.3245/8.2817**	0.0341/0.0312	**0.0021**	**7.0273/6.8529**	0.0332/**0.0292**	0.0026
Lunar-Lambert	6.5053/6.0975	0.0594/0.0614	0.0016	9.3279/9.4170	**0.0316**/0.0313	0.0057	7.9741/7.7830	**0.0282**/0.0293	0.0040

struction method. The results demonstrated both improved rendering quality and normal estimation accuracy when using physics-based approaches compared to the SH-based technique. Furthermore, the use of reflectance functions allowed for the introduction of albedo estimation, which is an important quantity for small body photometry. Future work may include exploring other reflectance models, such as empirical, body-specific methods, and introducing larger data sets to work toward global terrain mapping of a body.

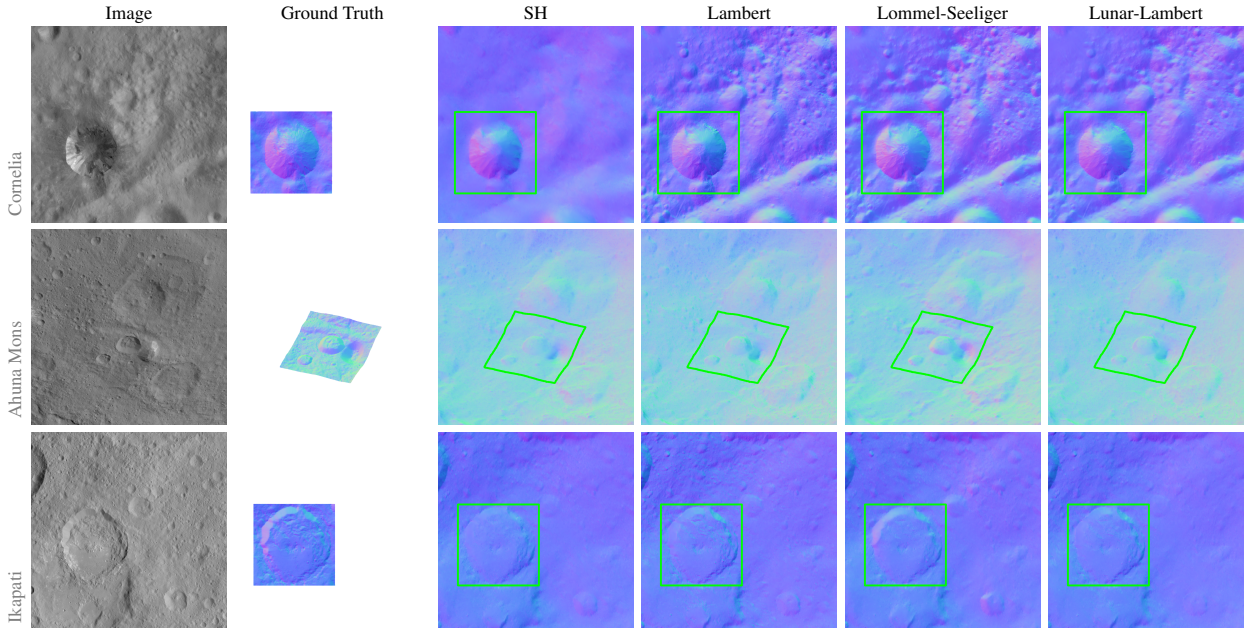

Figure 5. Qualitative comparison of test image normals for each reflectance model.

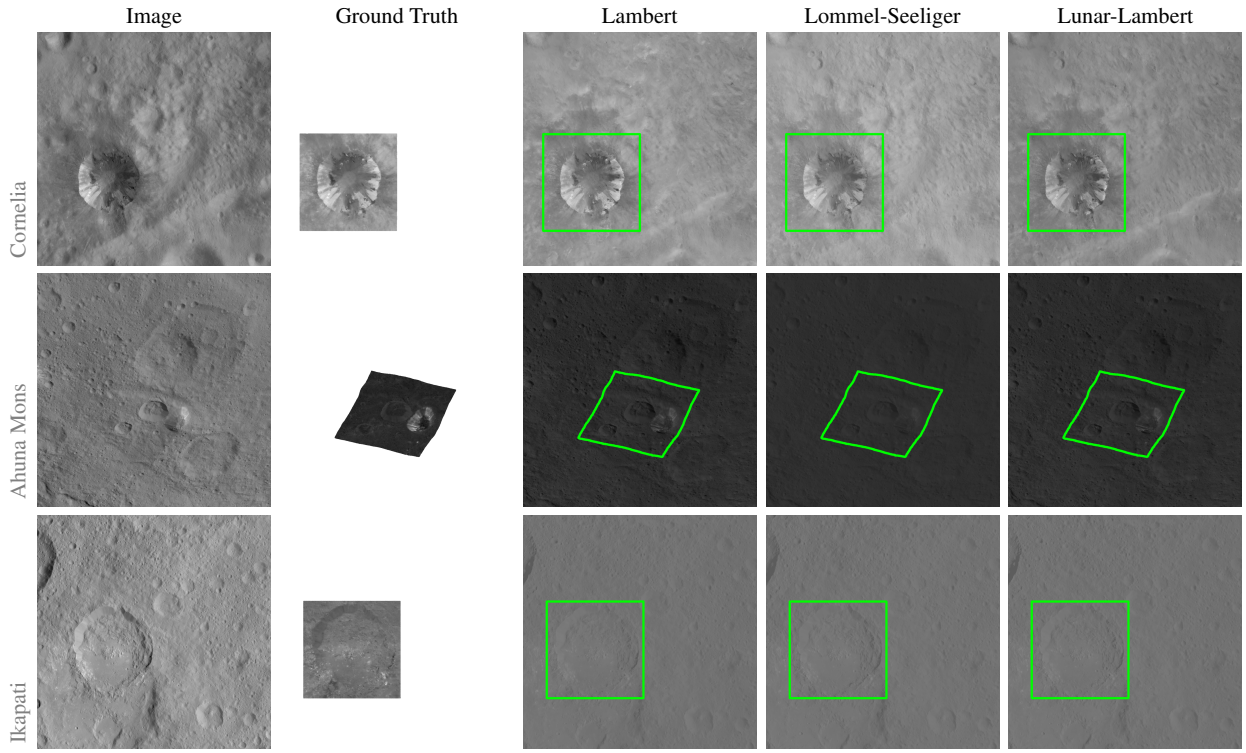

Figure 6. Qualitative comparison of test image albedos for each reflectance model (except spherical harmonics).

References

[1] Luca Savant Aira, Gabriele Facciolo, and Thibaud Ehret. Gaussian splatting for efficient satellite image photogrammetry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Nashville, TN, 2025. 2

[2] O. S. Barnouin, M. G. Daly, E. E. Palmer, C. L. Johnson,

R. W. Gaskell, M. Al Asad, E. B. Bierhaus, K. L. Craft, C. M. Ernst, R. C. Espiritu, H. Nair, G. A. Neumann, L. Nguyen, M. C. Nolan, E. Mazarico, M. E. Perry, L. C. Philpott, J .H. Roberts, R. J. Steele, J. Seabrook, H. C. M. Susorney, J. R. Weirich, and D. S. Lauretta. Digital terrain mapping by the OSIRIS-REx mission. *Planetary and Space Science*, 180 (104764):1–16, 2020. 1

[3] Hao Chen, Xuanyu Hu, Konrad Willner, Zhen Ye, Friedrich Damme, Philipp Gläser, Yongjie Zheng, Xiaohua Tong, Hauke Hußmann, and Jürgen Oberst. Neural implicit shape modeling for small planetary bodies from multi-view images using a mask-based classification sampling strategy. *ISPRS Journal of Photogrammetry and Remote Sensing*, 212:122–145, 2024. 2

[4] Shihan Chen, Bo Wu, Hongliang Li, Zhaojin Li, and Yi Liu. Asteroid-nerf: A deep-learning method for 3d surface reconstruction of asteroids. *Astronomy & Astrophysics*, 2024. in press. 2

[5] John A. Christian. *Fundamentals of Spacecraft Optical Navigation*. John Wiley & Sons, 2026. 3, 4, 5

[6] Frank Dellaert and Michael Kaess. Factor graphs for robot perception. *Found. Trends Robotics*, 6(1-2):1–139, 2017. 2

[7] Travis Driver, Katherine Skinner, Mehregan Dor, and Panagiotis Tsiotras. AstroVision: Towards autonomous feature detection and description for missions to small bodies using deep learning. *Special Issue on AI for Space, Acta Astronautica*, 210:393–410, 2023. 2

[8] Travis Driver, Andrew Vaughan, Yang Cheng, Adnan Ansar, John Christian, and Panagiotis Tsiotras. Stereophotoclinometry revisited. J. of Guidance, Control, and Dynamics (JGCD), 2025. 2, 5

[9] Johan Edstedt, Qiyu Sun, Georg Bökman, Mårten Wadenbäck, and Michael Felsberg. RoMa: Robust dense feature matching. In *CVPR*, pages 19790–19800, Seattle, WA, USA, 2024. 2, 5

[10] Carolyn M. Ernst, R. Terik Daly, Robert W. Gaskell, Olivier S. Barnouin, Hari Nair, Benjamin A. Hyatt, Manar M. Al Asad, and Kielan K. W. Hoch. High-resolution shape models of phobos and deimos from stereophotoclinometry. *Earth, Planets and Space*, 75(1):103, 2023. 1

[11] R. W. Gaskell, O. S. Barnouin-Jha, D. J. Scheeres, A. S. Konopliv, T. Mukai, S. Abe, J. Saito, M. Ishiguro, T. Kubota, T. Hashimoto, J. Kawaguchi, M. Yoshikawa, K. Shirakawa, T. Kominato, N. Hirata, and H. Demura. Characterizing and navigating small bodies with imaging data. *Meteoritics & Planetary Science*, 43(6):1049–1061, 2008. 1, 5

[12] Matthew W. Givens. Applying neural radiance fields to asteroid shape modeling. In *Space Imaging Workshop*, pages 1–4, Atlanta, GA, USA, 2024. 2

[13] Bruce Hapke. *Theory of Reflectance and Emittance Spectroscopy*. Cambridge University Press, 2012. 5

[14] Berhold K.P. Horn. *Robot Vision*. The MIT Press, 1986. 4

[15] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D Gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, pages 1–11, 2024. 1, 2

[16] Daniel P Huttenlocher, Gregory A. Klanderman, and William J Rucklidge. Comparing images using the hausdorff distance. *IEEE Transactions on pattern analysis and machine intelligence*, 15(9):850–863, 2002. 6

[17] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. 2023. 1, 2

[18] Bernhard Kerbl, Andreas Meuleman, Georgios Kopanas, Michael Wimmer, Alexandre Lanvin, and George Drettakis. A hierarchical 3d gaussian representation for real-time rendering of very large datasets. *ACM Transactions on Graphics*, 43(4), 2024. 5

[19] Jian-Yang Li, Paul Helfenstein, Bonnie J. Buratti, Driss Takir, and Beth Ellen Clark. Asteroid photometry. In *Asteroids IV*, pages 129–150. University of Arizona Press, 2015. 4

[20] Nickolaos Mastrodemos, Brian P. Rush, Andrew T. Vaughan, and W. M. Owen. Optical navigation for the Dawn mission at Vesta. In *Int. Symp. on Space Flight Dynamics (ISSFD)*, pages 1–26, Pasadena, CA, USA, 2012. 1

[21] Alfred S McEwen. Exogenic and endogenic albedo and color patterns on europa. *Journal of Geophysical Research*, 91 (B8):8077 – 8097, 1986. 5

[22] A. S. McEwen. Photometric functions for photoclinometry and other applications. *Icarus*, 92(2):298 – 311, 1991. 5

[23] NASA. Planetary Data System (PDS). `https://pds.nasa.gov/`, 2022. 5

[24] Eric E Palmer, Robert Gaskell, Michael G Daly, Olivier S Barnouin, Coralie D Adam, and Dante S Lauretta. Practical stereophotoclinometry for modeling shape and topography on planetary missions. *Planetary Science*, 3(102):1–16, 2022. 1

[25] R.S. Park, A.T. Vaughan, A.S. Konopliv, A.I. Ermakov, N. Mastrodemos, J.C. Castillo-Rogez, S.P. Joy, A. Nathues, C.A. Polanskey, M.D. Rayman, J.E. Riedel, C.A. Raymond, C.T. Russell, and M.T. Zuber. High-resolution shape model of Ceres from stereophotoclinometry using Dawn imaging data. *Icarus*, 319:812–827, 2019. 5

[26] Art Prosvetov, Alexander Govorov, Maxim Pupkov, Alexander Andreev, and Vladimir Nazarov. Illuminating the Moon: Reconstruction of lunar terrain using photogrammetry, neural radiance fields, and Gaussian splatting. *Astronomy and Computing*, 52:100953, 2025. 2

[27] C. A. Raymond, R. Jaumann, A. Nathues, H. Sierks, T. Roatsch, F. Preusker, F. Scholten, R. W. Gaskell, L. Jorda, H.-U. Keller, M. T. Zuber, D. E. Smith, N. Mastrodemos, and S. Mottola. The Dawn topography investigation. *Space Science Rev.*, 163:487–510, 2011. 1

[28] Christopher Russell and Carol Raymond. *The Dawn mission to minor planets 4 Vesta and 1 Ceres*. Springer Science & Business Media, 2012. 1, 5

[29] SE Schröder, T Maue, P Gutiérrez Marqués, S Mottola, KM Aye, H Sierks, HU Keller, and A Nathues. In-flight calibration of the Dawn Framing Camera. *Icarus*, 226(2):1304–1317, 2013. 4

[30] SE Schröder, S Mottola, HU Keller, CA Raymond, and CT Russell. Resolved photometry of vesta reveals physical prop-

erties of crater regolith. *Planetary and Space Science*, 85: 198–213, 2013. 5

[31] SE Schröder, S Mottola, K-D Matz, and Thomas Roatsch. In-flight calibration of the Dawn Framing Camera ii: Flat fields and stray light correction. *Icarus*, 234:99–108, 2014. 4, 5

[32] Michael K Shepard. *Introduction to Planetary Photometry*. Cambridge University Press, 2017. 4

[33] Yu Shkuratov, V Kaydash, V Korokhin, Y Velikodsky, N Opanasenko, and G Videen. Optical measurements of the moon as a tool to study its surface. *Planetary and Space Science*, 59(13):1326–1371, 2011. 5

[34] H. Sierks, H. U. Keller, R. Jaumann, H. Michalik, T. Behnke, F. Bubenhagen, I. Büttner, U. Carsenty, U. Christensen, R. Enge, B. Fiethe, P. Gutiérrez Marqués, H. Hartwig, H. Krüger, W. Kühne, T. Maue, S. Mottola, A. Nathues, K. U. Reiche, M. L. Richards, T. Roatsch, S. E. Schröder, I. Szemerey, and M. Tschentscher. The Dawn framing camera. *Space Science Rev.*, 163:263–327, 2011. 4, 5

[35] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. 4, 6

[36] Renato Zanetti. Rotations, transformations, left quaternions, right quaternions? *The Journal of the Astronautical Sciences*, 66(3):361–381, 2019. 3

[37] Richard Zhang, , Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 6