

Supplementary Materials for CVPR Submission “AstroSplat: Physics-Based Gaussian Splatting for Rendering and Reconstruction of Small Celestial Bodies”

Jennifer Nolan
jnolan9@gatech.edu

Travis Driver
travisdriver@gatech.edu

John Christian
john.a.christian@gatech.edu

Georgia Institute of Technology
Atlanta, GA

1. Sensitivity Analysis to Initialization

The experiments in the main paper demonstrate how the AstroSplat framework achieves improved reconstruction quality over the conventional 2D Gaussian Splatting [4] approach by using physics-based reflectance models in place of the traditional spherical harmonics (SH) method. All of the presented results relied on camera poses that were derived using stereophotoclinometry (SPC) and an initial point cloud estimated by Georgia Tech’s Structure-from-Motion(GTSfM) library [1], which used these poses to triangulate keypoints derived by RoMa [3]. Although RoMa is an autonomous method for detecting and matching keypoints, SPC is a human-intensive process, making these results not fully autonomous. Furthermore, SPC pose estimates benefit from the intervention of human operators and are considered a high-fidelity initialization.

To expand upon the results and conclusions drawn from the main paper, the aim of this additional analysis is twofold: 1) to demonstrate a true end-to-end autonomous implementation of the AstroSplat framework by using a completely autonomous initialization strategy, and 2) to show that the improvements in reconstruction achieved by using physics-based reflectance functions remain even with a lower-fidelity initialization. To address both of these objectives, the following experiments will use Georgia Tech’s Structure-from-Motion(GTSfM) library to estimate the camera poses autonomously. Then, the same autonomously detected RoMa [3] keypoints are used to triangulate the initial point cloud for the AstroSplat pipeline. Therefore, this alternate initialization strategy provides both a method to perform surface reconstruction autonomously end-to-end and to compare how the lower-quality initial estimates alter the quality of the reconstruction itself.

1.1. Experimental Setup

Since the goal of this analysis is to isolate the impact of using different initialization products, only the camera poses and initial point clouds deviate from the experiments outlined in the main paper. Therefore, the same three datasets—namely, Cornelia crater on 4 Vesta as well as Ikapati crater and Ahuna Mons on 1 Ceres [2, 5–7]—are used to demonstrate the impact of applying different photometric functions. The same train/test split (shown in Table 1 and Table 2) as well as the same number of training iterations (i.e., 30,000 iterations) are used. The same evaluation metrics and ground truth as in the main paper are used here as well.

1.2. Results

Figure 1 shows the same set of ground truth test images as well as the new associated test renders for each reflectance model. Table 1 provides the overall PSNR, SSIM, and LPIPS values for both train and test image sets. The renders themselves do not appear to show much variance compared to the results shown in the main paper. However, looking at the Ahuna Mons renders, some artifacts can be seen in the top left corner across all reflectance models that were not present when using the SPC-derived initialization. These artifacts indicate regions where the Gaussians do not well model the scene and could result from using the lower-fidelity poses and its limited coverage, leading to lower overlap in those regions. Quantitatively, for nearly all test images, the Lambert reflectance model again yielded the best results, except for the Ahuna Mons PSNR, in which SH and Lunar-Lambert show a slightly improved value. This change could stem from the aforementioned artifacts in the test renders. Overall, the rendering metrics themselves do not seem to significantly change when using this alternate initialization strategy.

Now considering reconstruction quality, Figure 2 shows




Figure 1. 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 \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
SH	36.48/35.07	0.9608/0.9572	0.1420/0.1351	39.66 /32.63	<u>0.9731</u> /0.8756	<u>0.0514</u> /0.1292	37.91/35.26	0.9804/0.9544	<u>0.0399</u> /0.0709
Lambert	39.45 / 38.74	0.9686 / 0.9660	0.1129 / 0.1068	39.11/33.29	0.9740 / 0.8810	0.0469 / 0.1258	38.87 / 36.17	0.9811 / 0.9561	0.0376 / 0.0679
L-S	37.58/36.47	0.9625/0.9577	0.1358/0.1339	38.96/ 33.63	0.9679/0.8780	0.0592/ <u>0.1283</u>	38.00/35.45	0.9762/0.9531	0.0472/0.0722
L-L	<u>37.68</u> / <u>37.29</u>	<u>0.9635</u> / <u>0.9596</u>	<u>0.1308</u> / <u>0.1266</u>	<u>39.21</u> / <u>33.58</u>	0.9694/0.8802	0.0575/0.1296	<u>38.65</u> / <u>35.76</u>	0.9781/ <u>0.9546</u>	<u>0.0429</u> / <u>0.0707</u>

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 1.1 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)		
	$\delta\theta$ ($^\circ$) \downarrow	δa \downarrow	d_H \downarrow	$\delta\theta$ ($^\circ$) \downarrow	δa \downarrow	d_H \downarrow	$\delta\theta$ ($^\circ$) \downarrow	δa \downarrow	d_H \downarrow
SH	7.6154/7.6016	--	0.0033	10.8259/10.8262	--	0.0036	8.6959/8.4889	--	0.0040
Lambert	6.4569/6.3051	0.0448 / 0.0456	0.0015	9.1833 / 9.3232	0.0309 / 0.0307	0.0033	<u>7.7929</u> / <u>7.6997</u>	<u>0.0284</u> /0.0297	<u>0.0022</u>
L-S	<u>6.1523</u> / <u>6.1487</u>	0.0730/0.0752	0.0016	<u>9.3400</u> / <u>9.3342</u>	0.0318/0.0313	0.0024	6.9910 / 6.8594	0.0280 / 0.0289	0.0021
L-L	6.0356 / 6.0189	<u>0.0656</u> / <u>0.0690</u>	<u>0.0015</u>	9.3905/9.4373	<u>0.0316</u> / <u>0.0311</u>	<u>0.0030</u>	8.1565/7.9755	0.0285/ <u>0.0296</u>	0.0036

the ground truth images, PhoMo ground truth normal maps, and associated AstroSplat normal maps for a test set of each model, and Figure 3 shows the same for albedo. Table 2 provides normal error, albedo error, and Hausdorff distance values for both train and test image sets. The most stark

difference between the results achieved using SPC-derived initialization versus the GTSfM and RoMa-derived initialization is shown in the normal maps. Especially beyond the green boxed regions, which indicates where the ground truth values are, several craters, ridges, and other surface

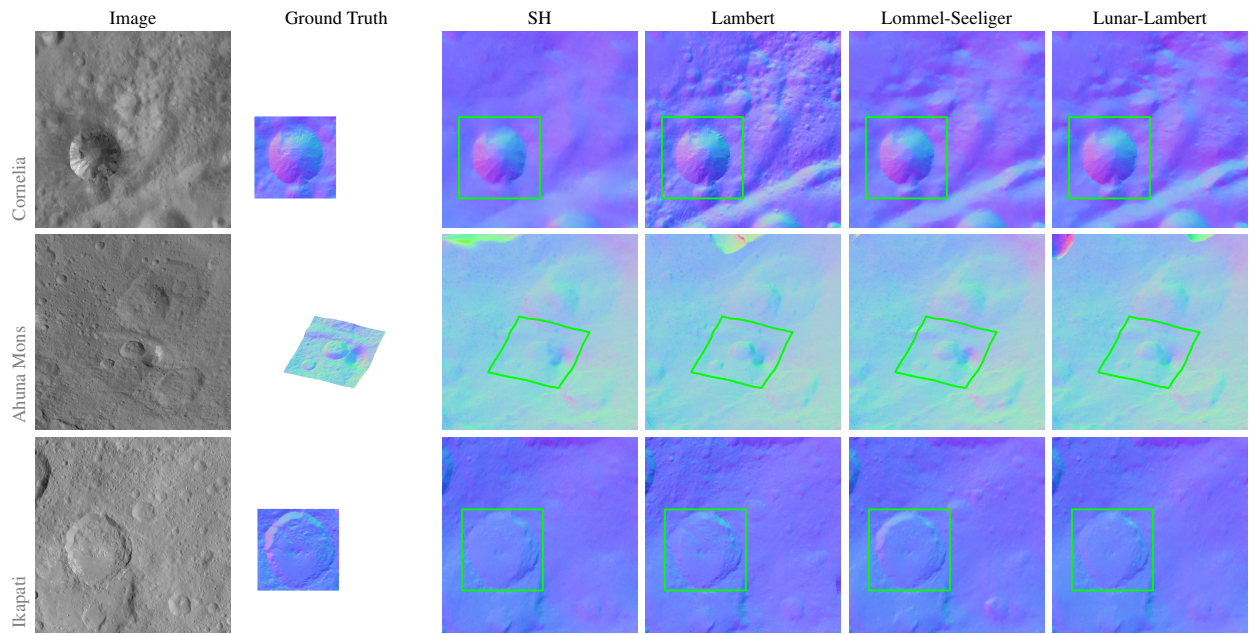


Figure 2. Qualitative comparison of test image normals for each reflectance model.




Figure 3. Qualitative comparison of test image albedos for each reflectance model (except spherical harmonics).

details that were present in the results from SPC-derived experiments are now absent. However, the physics-based reflectance models still show greater detail than the SH-

based approach, corroborating the demonstrated improvement shown in the main paper.

Quantitatively, the mean normal errors and albedo errors

remained approximately unchanged, indicating that while certain sharper details are lost, the averaged values overshadow that effect in the metrics. The loss in detail outside of that region is not captured in the metrics since ground truth data is restricted to the area near the feature of interest. Similar to the experiment in the main paper, the physics-based reflectance models yielded lower mean normal errors than the original SH method across all images, though the specific model that best estimated the surface did change. Specifically, Ahuna Mons saw significant changes, including normal errors being minimized by the Lambert model rather than Lommel-Seeliger. Albedo errors saw little change for Ahuna Mons and Ikapati but were significantly larger in these experiments compared to those from SPC-derived initialization. As seen in the main paper, the Hausdorff distance was again consistently minimized by a physics-based approach. Though the improvements seen when employing physics-based reflectance models over the traditional 2DGS approach are still present in this analysis, further exploration and adjustments to initialization would help reveal where the finer differences between these experiments may arise.

Ultimately, this analysis accomplished its twofold goals. First, it showed that the AstroSplat framework can work as an end-to-end autonomous reconstruction pipeline when autonomously-derived products are used to initialize training. Second, it supported the findings of the main paper, demonstrating that even with lower-fidelity initialization, the physics-based reflectance functions yield a more accurate reconstruction than the SH-based approach.

References

- [1] Ayush Baid*, John Lambert*, Travis Driver*, Akshay Krishnan*, Hayk Stepanyan, and Frank Dellaert. Distributed global structure-from-motion with a deep front-end. *arXiv:2311.18801*, 2023. *Authors contributed equally to this work. 1
- [2] Travis Driver, Andrew T. Vaughan, Yang Cheng, Adnan Ansar, John A. Christian, and Panagiotis Tsiotras. Stereophotoclinometry revisited. *J. of Guidance, Control, and Dynamics*, 2025. 1
- [3] 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. 1
- [4] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D Gaussian splatting for geometrically accurate radiance fields. In *SIGGRAPH*. Association for Computing Machinery, 2024. 1
- [5] NASA. Planetary Data System (PDS). <https://pds.nasa.gov/>, 2022. 1
- [6] Christopher T. Russell and Carol A. Raymond. *The Dawn mission to minor planets 4 Vesta and 1 Ceres*. Springer Science & Business Media, 2012.
- [7] Holger Sierks, Horst U. Keller, Ralf Jaumann, Harald Michalik, Thomas Behnke, Frank Bubenhausen, Irene Büttner, Uri Carsenty, Ulrich Christensen, Rainer Enge, Björn Fiethe, Pablo Gutiérrez Marqués, Hermann Hartwig, Harald Krüger, Wolfgang Kühne, Thorsten Maue, Stefano Mottola, Andreas Nathues, Kay-Uwe Reiche, Michael L. Richards, Thomas Roatsch, Stefan E. Schröder, Istvan Szemerey, and Matthias Tschentscher. The Dawn framing camera. *Space Science Rev.*, 163:263–327, 2011. 1