SEEING IN THE DARK: FEATURE EXTRACTION AND MATCHING IN LOW-LIGHT REGIONS USING DEEP LEARNING

Dhruv Kapu[1*], Travis Driver[1]; and Panagiotis Tsiotras[1]; [1]Georgia Institute of Technology, Atlanta, GA.
[*][dkapu3@gatech.edu]

Abstract. *Missions to small bodies rely heavily on optical feature tracking for characterization and relative navigation of the target body. While deep learning has led to great advancements in computer vision, feature extraction and matching in imagery of small bodies is still challenging in scenarios involving high illumination angles and shadows. This paper investigates the use of recently proposed "detector-free" matching methods for imaging matching in low-light and completely shadowed regions. The approach is validated on real imagery from the AstroVision dataset, demonstrating the ability of these approaches to precisely match keypoints in shadowed regions.*

Introduction. There has been an increasing interest in missions to small bodies (e.g., asteroids, comets) due to their great scientific value.[1] Image feature tracking is an integral component of small body science missions, serving as the foundation for shape reconstruction and relative navigation methodologies. Challenging illumination conditions are an innate characteristic of small body imagery. Indeed, due to the absence of an atmosphere of most small bodies to diffuse incoming light, images commonly feature rapidly changing lighting and dynamic hard shadowing, especially at high illumination angles.[2] Moreover, mapping perpetually shadowed regions, such as those found at the Lunar South Pole,[3,4] presents additional challenges for optical sensors. State-of-the-practice approaches to small body surface reconstruction, i.e., Stereophotoclinometry (SPC),[5] are sensitive to lighting conditions and cannot operate on perpetually shadowed regions.[6] Despite the success of keypoint detection and feature description methods based on deep learning, these approaches still struggle in scenarios involving significant shadowing.[7]

Approaches to mapping in low-light and perpetually shadowed regions typically rely on range measurements acquired by active sensors such as flash LiDARs. However, methods based on active sensors are constrained by the relatively small effective operating range and the increased size, weight, and power (SWaP) requirements relative to passive sensors such as monocular cameras. For example, the OSIRIS-REx Guidance, Navigation, and Control (GNC) flash LiDAR had a maximum operational range of approximately 1 km.[8,9] Conversely, the OSIRIS-REx Camera Suite (OCAMS)[10] was able to acquire 5 cm GSD images at almost 4 km, providing higher-resolution measurements earlier in the mission than the active sensors onboard and allowing for detailed surface characterization during the early phases of the mission.[11]

In this paper, we investigate the use of recently proposed "detector-free" matching methods, i.e., DKM,[12] for mapping of shadowed regions on the surface of small bodies. We find that training these detector-free methods on imagery of small bodies from the AstroVision dataset,[7] along with the increased receptive field afforded by these methods, enables the detection and matching of keypoints in *completely shadowed regions*. We believe this could enable image-based surface reconstruction in areas featuring perpetually high illumination angles and shadowing.

Detector-Free Feature Matching. Traditional feature matching (e.g., SIFT[13]) follows the *detect-then-describe* paradigm which comprises sparse keypoint detection, feature description, and then matching. This approach degrades in low-texture regions where localization of keypoints is challenging. To that end, we use DKM,[12] a deep-learning model that follows the *detector-free* paradigm. Instead of focusing on matching sparse keypoints, DKM provides a dense warp between views by matching all pixels. DKM leverages a novel *global matching* scheme using a Gaussian Process regression to achieve a large receptive field. Thus, DKM is able to aggregate information from other regions in the image and provide matches in featureless areas (e.g., low-light regions) where traditional methods would fail to detect keypoints.

Implementation Details. We use a batch size of 1 with a learning rate of 4×10^{-4} for the decoder and refiners, and 2×10^{-5} for the backbone. We use the AdamW[14] optimizer with a weight-decay factor of 1×10^{-2}. Using the pre-trained weights, we fine-tune for 75,000 steps on 17,504 images of Asteroid 4 Vesta from the Dawn mission. Training takes approximately 24 hours on an Nvidia GeForce RTX 2080 Ti GPU.

We evaluate the matching performance of the pretrained and finetuned DKM model based on precision, i.e., the ratio of correct to total matches. Matches are verified by projecting keypoints from the first image to the second and from the second image to the first using ground truth poses and depth maps and enforcing that the projected image coordinates are within some distance to the matching keypoint, where we report the precision at distance thresholds of 1, 3, and 5 pixels. Moreover, we report the area under the curve (AUC) of the normalized cumulative pose error at thresholds of 5°, 10°, and 20°, where the pose error is the maximum of the angular errors in rotation and translation.

Results & Discussion. We compare the performance of the pretrained DKM model against the model finetuned on AstroVision[7] data in Table 1, and a qualitative comparison is shown in Figure 1 with image pairs with perpetually shadowed regions. In Figure 1, we filter out matches with a certainty of less than 0.005, since it was empirically

Table 1. AstroVision benchmark results on the 4 Vesta dataset. Best results are bolded.

	Precision [%]			AUC [%]		
Method	@1 px	@3 px	@5 px	@5°	@10°	@20°
Pretrained	19.80	58.30	88.21	18.10	24.30	30.00
Finetuned	**19.90**	**69.10**	**88.91**	**22.30**	**35.90**	**48.60**

found to work well, and we randomly sample 300 matches from shadowed regions and 100 from other areas for each pair to plot. The finetuned model outperforms the pretrained model on all metrics in Table 1. Additionally, as shown in Figure 1, the model finetuned on Astrovision data demonstrates precise matching in completely shadowed regions where the pretrained model fails to detect any matches in the shadowed regions. We believe that our models could enable image-based reconstruction of perpetually shadowed regions, which is not possible with SPC.

Acknowledgments. This work was partially supported by a NASA Space Technology Graduate Research Opportunity.

References.

[1] M. A. Barucci, E. Dotto, and A. C. Levasseur-Regourd, "Space missions to small bodies: asteroids and cometary nuclei," *Astronomy and Astrophysics Review*, vol. 19, no. 48, pp. 1–29, 2011.

[2] L. K. McCarthy, C. D. Adam, J. M. Leonard, P. G. Antresian, D. Nelson, E. Sahr, J. Pelgrift, E. J. Lessac-Chenen, J. Geeraert, and D. Lauretta, "OSIRIS-REx landmark optical navigation performance during orbital and close proximity operations at Asteroid Bennu," in *AIAA SciTech Forum*, p. 2520, 2022.

[3] D. Bussey, J. McGovern, P. Spudis, C. Neish, H. Noda, Y. Ishihara, and S.-A. Sørensen, "Illumination conditions of the south pole of the Moon derived using Kaguya topography," *Icarus*, vol. 208, no. 2, pp. 558–564, 2010.

[4] P. Gläser, F. Scholten, D. De Rosa, R. M. Figuera, J. Oberst, E. Mazarico, G. Neumann, and M. Robinson, "Illumination conditions at the lunar south pole using high resolution digital terrain models from LOLA," *Icarus*, vol. 243, pp. 78–90, 2014.

[5] R. W. Gaskell et al., "Characterizing and navigating small bodies with imaging data," *Meteoritics & Planetary Science*, vol. 43, no. 6, pp. 1049–1061, 2008.

[6] K. Getzandanner et al., "Small body proximity operations & TAG: Navigation experiences & lessons learned from the OSIRIS-REx mission," in *AIAA SciTech Forum*, pp. 1–23, 2022.

[7] T. Driver, K. Skinner, M. Dor, and P. Tsiotras, "AstroVision: Towards autonomous feature detection and description for missions to small bodies using deep learning," *Acta Astronautica: Special Issue on AI for Space*, vol. 210, pp. 393–410, 2023.

[8] E. Church, T. Bourbeau, J. Curriden, A. Deguzman, F. Jaen, H. Ma, K. Mahoney, C. Miller, B. Short, K. Waldorff, *et al.*, "Flash LIDAR on-orbit performance at Asteroid Bennu," in *AAS Guidance, Navigation and Control (GN&C) Conf.*, 2020.

[9] J. M. Leonard, M. C. Moreau, P. G. Antreasian, K. M. Getzandanner, E. Church, C. Miller, M. G. Daly, O. S. Barnouin, and D. S. Lauretta, "Cross-calibration of GNC and OLA LIDAR systems onboard OSIRIS-REx," in *AAS Guidance, Navigation and Control (GN&C) Conf.*, no. 22-166, 2022.

[10] B. Rizk, C. D. d'Aubigny, D. Golish, C. Fellows, C. Merrill, P. Smith, M. Walker, J. Hendershot, J. Hancock, S. Bailey, D. DellaGiustina, D. Lauretta, R. Tanner, M. Williams, K. Harshman, M. Fitzgibbon, W. Verts, J. Chen, T. Connors, D. Hamara, A. Dowd, A. Lowman, M. Dubin, R. Burt, M. Whiteley, M. Watson, T. McMahon, M. Ward, D. Booher, M. Read, B. Williams, M. Hunten, E. Little, T. Saltzman, D. Alfred, S. O'Dougherty, M. Walthall, K. Kenagy, S. Peterson, B. Crowther, M. Perry, C. See, S. Selznick, C. Sauve, M. Beiser, W. Black, R. Pfisterer1, A. Lancaster, S. Oliver, C. Oquest, D. Crowley, C. Morgan, C. Castle, R. Dominguez, and M. Sullivan, "OCAMS: The OSIRIS-REx Camera Suite," *Space Science Reviews*, vol. 214, no. 26, pp. 1–55, 2018.

[11] D. A. Lorenz, R. Olds, A. May, C. Mario, M. E. Perry, E. E. Palmer, and M. Daly, "Lessons learned from OSIRIS-REx autonomous navigation using natural feature tracking," in *IEEE Aerospace Conf.*, pp. 1–12, 2017.

[12] J. Edstedt, I. Athanasiadis, M. Wadenbäck, and M. Felsberg, "DKM: Dense kernelized feature matching for geometry estimation," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 17765–17775, 2023.

[13] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proceedings of the seventh IEEE international conference on computer vision*, vol. 2, pp. 1150–1157, Ieee, 1999.

[14] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2019.

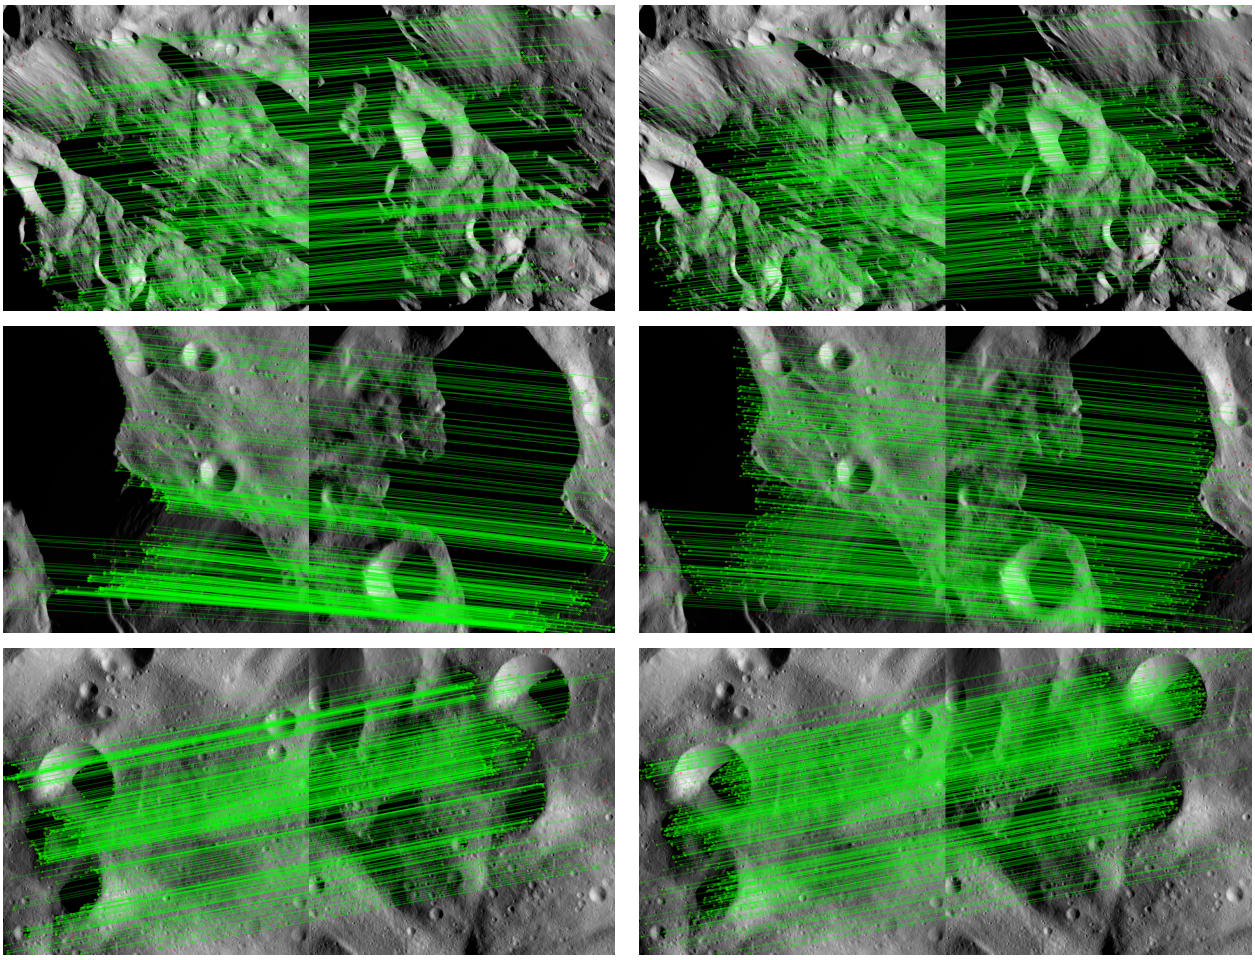

Figure 1. Qualitative comparison of matching performance between the pretrained DKM model (left) and the model finetuned on AstroVision data (right). Correct matches are drawn in green and the keypoints of incorrect matches are drawn in red.