
Low-light image enhancement of permanently shadowed lunar regions with physics-based machine learning

Ben Moseley
University of Oxford

Valentin Bickel
ETH Zurich/ MPS Goettingen

Ignacio G. López-Francos
NASA Ames Research Center

Loveneesh Rana
University of Luxembourg

Miguel Olivares-Mendez
University of Luxembourg

Dennis Wingo
Skycorp Inc.

Allison Zuniga
NASA Ames Research Center

Nuno Subtil
NVIDIA

Abstract

Finding water(-ice) on the lunar surface is key to enabling a sustainable human presence on the Moon and beyond. There is evidence that water-ice is abundant in and around the Moon's Permanently Shadowed Regions (PSRs), however, direct visual detection has not yet been possible. Surface ice or related physical features could potentially be directly detected from high-resolution optical imagery, but, due to the extremely low-light conditions in these areas, high levels of sensor and photon noise make this very challenging. In this work we generate high-resolution, low-noise optical images over lunar PSRs by using two physics-based deep neural networks to model and remove CCD-related and photon noise in existing low-light optical imagery, potentially paving the way for a direct water-ice detection method.

1 Introduction

In recent years the international space community has gained significant momentum for continuing the exploration of the Moon and establishing a sustainable presence, with multiple national space agencies and private organizations working towards this goal (Neal et al., 2009). A critical objective of these efforts is to identify potential lunar water resources, because water is a key element for life support and a source of fuel (Hydrogen and Oxygen) (ESA, 2019). Lunar water is also a primary target for scientific investigations, such as studies of the evolution of our solar system.

Recent missions have found evidence that suggests water(-ice) exists on the Moon and could be abundant. Amongst others, observations from Chandrayaan-1 (Li et al. (2018); Pieters et al. (2009)), the Lunar Reconnaissance Orbiter (LRO, Hayne et al. (2015); Fisher et al. (2017); Hayne et al. (2020)), the Lunar Crater Observation and Sensing Satellite (LCROSS, Colaprete et al. (2010)), and the airborne SOFIA observatory (Honniball et al. (2020)) indicate that water likely exists in and around topographic depressions at the lunar poles, mainly in the so-called Permanently Shadowed Regions (PSRs).

Whilst these observations are a significant step forward, the direct detection of (surface) water-ice or related physical features in PSRs has not yet been possible. The presence of bright ice could potentially be directly detected using high-resolution optical imagery. Currently the best imagery of the polar regions is produced by LRO's Narrow Angle Camera (NAC) with a nominal (polar) spatial resolution of 1 to 2 m/pixel (Robinson et al., 2010; Humm et al., 2016). However, due to the extremely low lighting conditions over PSRs, Charge-Coupled Device (CCD) sensor-related and photon noise dominate these images, strongly limiting our ability to make meaningful observations.

In this work we generate high-resolution, low-noise images of PSRs by training two deep neural networks to model and remove CCD and photon noise in these images, significantly improving their quality. We name our method Hyper-effective nOise Removal U-net Software (HORUS).




Figure 1: Performance of HORUS over 3 example PSRs, compared to the current calibration routine (ISIS) (Humm et al., 2016; Gaddis et al., 1997). Colour bar shows the estimated surface signal S for the DestripeNet+PhotonNet plot. Raw/ISIS image credits to LROC/ASU/GSFC.

2 Methods

The LRO NAC takes 12-bit panchromatic optical images and consists of two nominally identical, 5064-pixel line-scan CCD cameras oriented perpendicular to the direction of flight of the satellite (Robinson et al., 2010). Each camera has two operating modes, “regular” and “summed”, where the summed mode sums adjacent pixels and is typically used to maximize the signal received over poorly lit regions of the Moon. Over 10 years of operation the LRO NAC has captured over 2 million images, covering the entire surface multiple times over.

Figure 1 shows example raw NAC images captured over 3 example PSRs at the lunar south pole. Whilst they receive no direct sunlight, PSRs can be illuminated by indirect light scattered from their surroundings. In this extremely low-light setting NAC images are dominated by CCD noise and photon noise, which is due to the inherent randomness of photons arriving at the CCD. Informed by the work of Humm et al. (2016) and standard CCD theory (Chromey, 2016), in this work we assume that this noise can be described by the following physical model,

$$I = S + N_p + N_d + N_b + N_r, \quad (1)$$

where I is the raw image detector count, S is the mean photon signal, N_p is the stochastic photon noise, which obeys a Poisson distribution with an expected value equal to S , N_d is stochastic dark current noise, which is generated by thermally-activated electrons in the CCD and obeys a Poisson distribution, N_b is the dark bias, which is deterministic and due to a voltage offset which varies pixel-to-pixel, and N_r is the readout noise. Of these noise sources, the photon noise and dark bias are considered the most dominant (Humm et al., 2016) and are the focus of this study.

The current calibration routine for the NAC images (ISIS, or Integrated Software for Imagers and Spectrometers (Humm et al., 2016; Gaddis et al., 1997)) only attempts to remove the dark bias from each image, by subtracting a cyclic trend from its set of 60 masked pixels and a single library dark calibration frame. Examples images after this calibration are shown in Figure 1: we observe that photon noise and a significant amount of dark bias (residual stripes) are still present.

In this work we remove both photon and dark bias noise from the images by using two deep neural networks, collectively named HORUS. The first network is called DestripeNet and it models the dark bias. The input to the network is the camera metadata available at the time of image capture (listed in Figure 2) and its output is a prediction of the dark bias frame. We train this network using a set of 100,000+ calibration images as labels, taken on the non-illuminated side of the Moon, where negligible photon signal is present. We use a convolutional decoder network design, shown in Figure 2 (top). Thus, this network learns a time-varying physics-based model of the CCD sensor.

The second network is called PhotonNet and it models photon noise in the images. The input to the network is the image after DestripeNet has been applied and its output is a prediction of the mean surface signal S . This network is trained using synthetically generated noisy images by taking 1 million randomly selected




Figure 2: Physics-based HORUS denoising approach and noise model. Top: DestripeNet and PhotonNet networks. DestripeNet uses a convolutional decoder design with 13 layers and 512 hidden channels to model the dark bias frame N_b . PhotonNet uses a U-Net design with 4 down-sampling steps to remove photon noise N_p . Both networks are trained using an L2 loss function. PhotonNet operates on images patches, whilst DestripeNet predicts by image line. Bottom: noise model used in this work, showing an example dark calibration frame (N_b). Image credits to LROC/ASU/GSFC.




Figure 3: 3D ray tracing and training distribution matching. Left: histogram of the secondary incidence angle of scattered rays for 4 PSRs simulated using 3D ray tracing, compared to the distribution of solar incidence angles in the training set for PhotonNet. Middle: A render of the secondary illumination radiance for the Shackleton crater PSR; here the crater rim is a strong source of scattered light. Right: distribution of photon counts over the re-scaled bright images used to train PhotonNet. Typically, PSRs have raw detector counts I in the range 0-50 and we re-scale to this level.

bright, sunlit patches of the Moon, re-scaling them to low-lighting photon counts, adding Poisson noise to them, and then learning to remove this noise. We use a U-Net design (Ronneberger et al., 2015), similar to other low-light denoising approaches (Chen et al., 2018).

An important consideration is that the predictions of PhotonNet will be biased towards the distribution of images it has been trained on. Whilst it is trained on images with direct solar illumination, it is intended for use on images of PSRs, which are illuminated by secondary (backscattered) light only. In order to match the training data to the test-time data as best we can, we carry out 3D ray tracing of 4 PSRs at the south pole and plot the expected distribution of secondary illumination angle under typical illumination conditions. We then match the distribution of solar incidence angles in our training images for PhotonNet to this distribution. Ray tracing is performed using 30 m/pixel spatial resolution LRO Lunar Orbiter Laser Altimeter elevation data (Riris et al., 2017) and a Lambertian bidirectional reflection function. Finally, to account for possible differences in the two cameras and two operating modes, we train separate networks for each possible configuration.

3 Results and Discussion

Histograms of the incidence angle of scattered rays from ray tracing are shown in Figure 3. We find that these rays typically have incidence angles in the range 60 to 100 degrees (normal to the Moon’s sphere) and our matched distribution of solar incidence angles in the training set is shown.

Example real-data test images of PSRs after training DestripeNet and PhotonNet are shown in Figure 1. We find that DestripeNet removes more of the dark bias “stripes” in the images compared to the current calibration routine, and PhotonNet significantly improves the image quality, removing the “graininess” of the photon noise. We generate a test set of 100,000 unseen synthetic noisy images by adding samples of

Mode	Camera	Baseline	ISIS	DestripeNet	DestripeNet+PhotonNet
		L1 / Peak signal-to-noise ratio (PSNR) / Structural similarity index measure (SSIM)			
Normal	Left	8.12 / 29.89 / 0.56	3.79 / 35.23 / 0.79	3.67 / 35.63 / 0.80	1.08 / 47.49 / 0.98
	Right	9.23 / 29.18 / 0.53	3.79 / 35.31 / 0.79	3.70 / 35.62 / 0.80	1.09 / 47.62 / 0.98
Summed	Left	7.75 / 30.33 / 0.59	3.73 / 35.37 / 0.80	3.63 / 35.75 / 0.81	1.20 / 46.76 / 0.97
	Right	9.19 / 29.29 / 0.54	3.79 / 35.39 / 0.80	3.72 / 35.65 / 0.81	1.25 / 46.89 / 0.97

Table 1: Test set performance of HORUS. We compare to a baseline approach, where simply the mean masked pixel value is subtracted from the raw image, and the current calibration routine (ISIS) (Humm et al., 2016; Gaddis et al., 1997). The average performance over all images is reported.




Figure 4: Ground truth verification. The comparison of a non-HORUS sunlit and a HORUS shadowed image (contrast and brightness-adapted) in a temporarily shadowed region indicates that our denoising does not remove or add any physical features. Non-HORUS image credit to LROC/ASU/GSFC.

dark calibration images and photon noise to re-scaled sunlit images and report the performance of HORUS on this dataset in Table 1. This table shows that HORUS significantly outperforms the current calibration routine under all tested metrics for these images.

A key concern is whether our networks “hallucinate” or remove physical features from the images. To help investigate this, we plot a non-HORUS sunlit image of a Temporarily Shadowed Region (TSR) and compare this to a shadowed image with HORUS applied in Figure 4. Encouragingly, we find that HORUS does not remove or add any incorrect physical features compared to the sunlit image.

4 Conclusion and Further Work

We have shown that it is possible to significantly enhance the quality of extremely low-light images of PSRs on the Moon by using two physics-based deep neural networks to remove CCD and photon noise. Our approach provides high-resolution, low-noise images which could help to directly detect water-ice. We have run HORUS over 20+ PSRs and future work will analyse these images for surface-ice related signals, as well as quantitatively assess the performance of HORUS using ground truth TSR images and investigate joint training of the networks to further improve performance.

5 Broader impact

This work could significantly impact humanity’s plans for exploration of the lunar poles, helping to enable a sustainable, long term presence on the Moon and beyond. First, our work could directly inform robotic scouting missions to the poles, enabling identification of water(-ice)-related science targets and reducing uncertainty when planning traverses into PSRs. Secondly, our work could support the planning and execution of the upcoming NASA Artemis missions, which will see the first woman and the next man operating around and potentially within lunar polar PSRs.

Acknowledgements. This work was completed as part of the 2020 NASA Frontier Development Lab (FDL). The authors would like to thank Julie Stopar and Nick Estes for insightful comments on the NAC instrument, Eugene D’Eon for advice on the ray tracing, and all of the FDL reviewers and partners (Luxembourg Space Agency, Google Cloud, Intel AI and NVIDIA).

Contributions. VTB and BM were involved in the conceptualization of this project. VTB, BM, ILF and LR were involved in the methodology, software, validation, formal analysis, investigation, resources, data curation, writing, preparation, visualization. VTB, BM, ILF, LR, AZ, DW and MOM were involved in the review, administration, supervision. ILF, AZ, DW and the FDL management team were involved in funding acquisition. NS and EDE implemented the 3D ray-tracing.

References

- Neal et al. The lunar exploration roadmap. exploring the moon in the 21st century: Themes, goals, objectives, investigations, and priorities, 2009. *LEAG Roadmap*, 2009.
- ESA. Esa space resources strategy. *ESA Report*, 2019.
- Li et al. Direct evidence of surface exposed water ice in the lunar polar regions. *PNAS*, 2018.
- Pieters et al. Character and spatial distribution of oh/h₂o on the surface of the moon seen by m3 on chandrayaan-1. *Science*, 2009.
- Hayne et al. Evidence for exposed water ice in the moon's south polar regions from lunar reconnaissance orbiter ultraviolet albedo and temperature measurements. *Icarus*, 2015.
- Fisher et al. Evidence for surface water ice in the lunar polar regions using reflectance measurements from the lunar orbiter laser altimeter and temperature measurements from the diviner lunar radiometer experiment. *Icarus*, 2017.
- Hayne et al. Micro cold traps on the moon. *Nature Astronomy*, 2020. doi: <https://doi.org/10.1038/s41550-020-1198-9>.
- Colaprete et al. Detection of water in the lcross ejecta plume. *Science*, 2010.
- Honniball et al. Molecular water detected on the sunlit moon by sofia. *Nature Astronomy*, 2020. doi: <https://doi.org/10.1038/s41550-020-01222-x>.
- M. S. Robinson, S. M. Brylow, M. Tschimmel, D. Humm, S. J. Lawrence, P. C. Thomas, B. W. Denevi, E. Bowman-Cisneros, J. Zerr, M. A. Ravine, M. A. Caplinger, F. T. Ghaemi, J. A. Schaffner, M. C. Malin, P. Mahanti, A. Bartels, J. Anderson, T. N. Tran, E. M. Eliason, A. S. McEwen, E. Turtle, B. L. Jolliff, and H. Hiesinger. Lunar reconnaissance orbiter camera (LROC) instrument overview. *Space Science Reviews*, 150(1-4):81–124, jan 2010.
- D. C. Humm, M. Tschimmel, S. M. Brylow, P. Mahanti, T. N. Tran, S. E. Braden, S. Wiseman, J. Danton, E. M. Eliason, and M. S. Robinson. Flight Calibration of the LROC Narrow Angle Camera. *Space Science Reviews*, 200(1-4):431–473, 2016.
- L. Gaddis, J. Anderson, K. Becker, T. Becker, D. Cook, K. Edwards, E. Eliason, T. Hare, H. Kieffer, E. M. Lee, J. Mathews, L. Soderblom, T. Sucharski, J. Torson, A. McEwen, and M. Robinson. An Overview of the Integrated Software for Imaging Spectrometers (ISIS). In *Lunar and Planetary Science Conference*, Lunar and Planetary Science Conference, page 387, March 1997.
- Frederick R. Chromey. *To Measure the Sky*. Cambridge University Press, oct 2016.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *ArXiv e-prints*, may 2015.
- Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to See in the Dark. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3291–3300, 2018. ISBN 9781538664209.
- H. Riris, G. Neuman, J. Cavanaugh, X. Sun, P. Liiva, and M. Rodriguez. The Lunar Orbiter Laser Altimeter (LOLA) on NASA's Lunar Reconnaissance Orbiter (LRO) mission. In Naoto Kadowaki, editor, *International Conference on Space Optics — ICSO 2010*, volume 10565, page 77. SPIE, nov 2017. ISBN 9781510616196.