

# LROC-PANGU-GAN: CLOSING THE SIMULATION GAP IN LEARNING CRATER SEGMENTATION WITH PLANETARY SIMULATORS

Jaewon La<sup>1</sup>, Jaime Phadke<sup>1</sup>, Matt Hutton<sup>4</sup>, Marius Schwinning<sup>4</sup>, Gabriele De Canio<sup>4</sup>, Florian Renk<sup>4</sup>, Lars Kunze<sup>2</sup>, and Matthew Gadd<sup>3</sup>

<sup>1</sup>*Balliol College*, <sup>2</sup>*Cognitive Robotics Group (CRG)*, <sup>3</sup>*Mobile Robotics Group (MRG)*, *University of Oxford*  
{lars,mattgadd}@robots.ox.ac.uk

<sup>4</sup>*European Space Operations Centre, European Space Agency (ESA)*  
Florian.Renk@esa.int\*

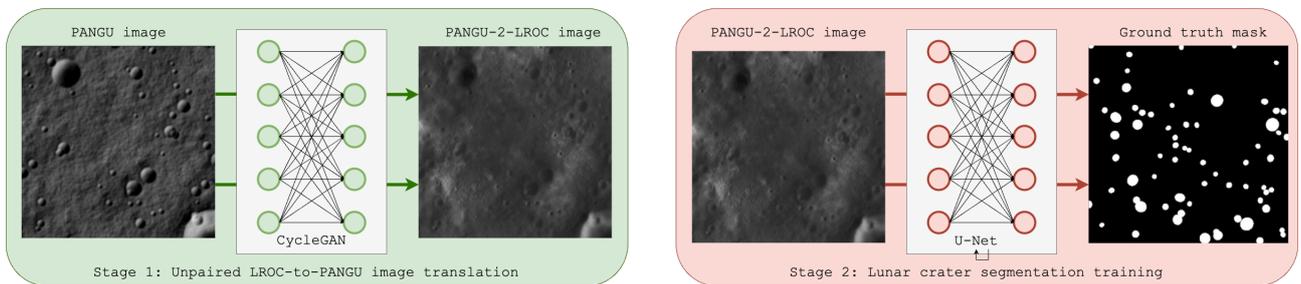


Figure 1: Simplified overview of our two-stage system. **Stage 1** is the core contribution of this work, while **Stage 2** is used to demonstrate the usefulness of the synthesised outputs. By transforming synthetic images of the lunar landscape from the Planet and Asteroid Natural Scene Generation Utility simulator into images which, although synthesised, more faithfully capture the nuances and attributes of the moon’s surface, we produce more realistic-looking lunar data. This is important for training subsequent models e.g. for segmenting and detecting lunar craters, since our synthesised data is derived from precise Planet and Asteroid Natural Scene Generation Utility crater locations, which real crater databases cannot guarantee.

## ABSTRACT

It is critical for probes landing on foreign planetary bodies to be able to robustly identify and avoid hazards – as, for example, steep cliffs or deep craters can pose significant risks to a probe’s landing and operational success. Recent applications of deep learning to this problem show promising results. These models are, however, often learned with explicit supervision over annotated datasets. These human-labelled crater databases, such as from the Lunar Reconnaissance Orbiter Camera (LROC), may lack in consistency and quality, undermining model performance – as incomplete and/or inaccurate labels introduce noise into the supervisory signal, which encourages the model to learn incorrect associations and results in the model making unreliable predictions. Physics-

based simulators, such as the Planet and Asteroid Natural Scene Generation Utility, have, in contrast, perfect ground truth, as the internal state that they use to render scenes is known with exactness. However, they introduce a serious simulation-to-real domain gap – because of fundamental differences between the simulated environment and the real-world arising from modelling assumptions, unaccounted-for physical interactions, environmental variability, etc. Therefore, models trained on their outputs suffer when deployed in the face of realism they have not encountered in their training data distributions. In this paper, we therefore introduce a system to close this “realism” gap while retaining label fidelity. We train a CycleGAN model to synthesise LROC from Planet and Asteroid Natural Scene Generation Utility (PANGU) images. We show that these improve the training of a downstream crater segmentation network, with segmentation performance on a test set of real LROC images improved as compared to using only simulated PANGU images. This will in the future allow researchers to more robustly test in-the-loop autonomous systems when tested in simulators such as PANGU.

\*This work was supported by the Human-Machine Collaboration Programme of the University of Oxford and Amazon Web Services (AWS). Lars Kunze is supported by EPSRC Project RAILS (EP/W011344/1). Matthew Gadd is supported by EPSRC Programme Grant “From Sensing to Collaboration” (EP/V000748/1).

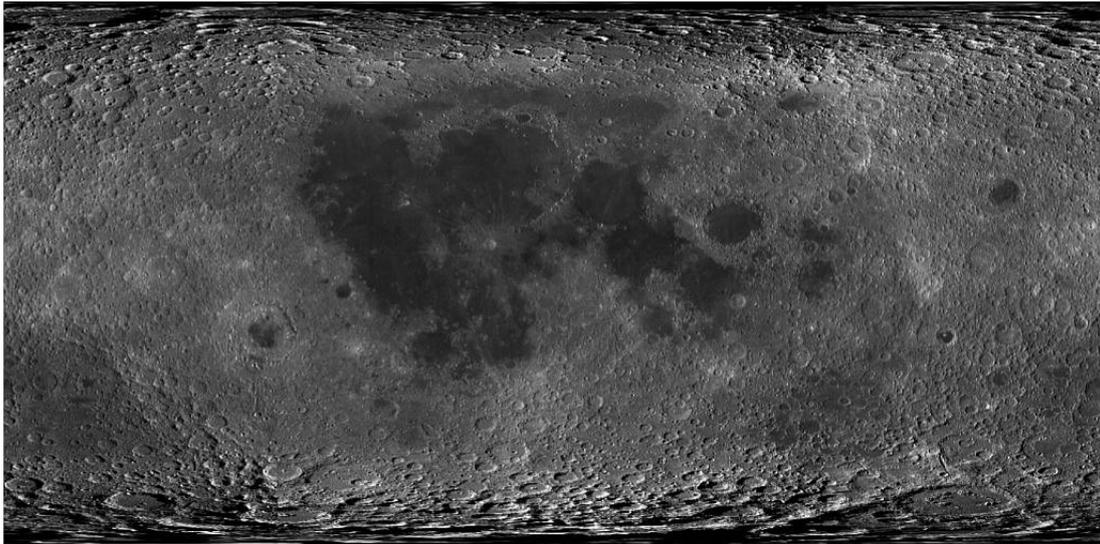


Figure 2: The Lunar Reconnaissance Orbiter Camera (LROC) dataset used in this investigation was obtained through the NASA and European Space Agency (ESA) databases. The dataset consists of a very high resolution image of the entire lunar surface, which was generated by stitching together thousands of images taken by the Lunar Reconnaissance Orbiter (LRO), with corresponding document containing the location, size, and the approximate shape of the craters. The approximate resolution of the LROC image is around 100 m per pixel. In the associated *Lunar Crater Database* [1], labels of craters were generated by human inspection, i.e. human volunteers hand-labelled each crater in the image.

## 1. INTRODUCTION

Accurate representation of extraterrestrial landscapes is important for landing [2] and navigation [3] – for example, terrain understanding can improve the precision of a spacecraft navigation by positioning relative to large-scale features or may inform terrestrial missions [4] after landing – as well as scientific understanding of the geology and evolution of planets and moons [5].

In this area, data from the the Lunar Reconnaissance Orbiter Camera (LROC) [6], a high-resolution imaging system on board the Lunar Reconnaissance Orbiter (LRO) spacecraft [7], is increasingly being used to learn deep models for crater detection [8, 9, 10, 11, 12, 13, 14, 15, 11, 16]. In these, convolutional neural networks (CNNs) are well-suited to learning robustness to variability in cameras arises from illumination conditions, sensor sensitivity, noise, camera settings, and lens characteristics. The lunar surface as imaged by this sensor is shown in Fig. 2.

Alongside this, the Planet and Asteroid Natural Scene Generation Utility [17] and its extensions [18] are popular for in-the-loop testing of spacecraft systems in simulation as well as as a source of data for machine learned systems [19, 20, 21, 22, 23, 24, 25, 26, 27]. Here, simulators are valuable for training deep models as they provide abundant data, are controllable and cost-effective, and allow rapid prototyping – i.e. they are safe and efficient. The lunar surface as simulated by PANGU is shown in Fig. 3,

which – while very detailed – is readily identifiable as simulated.

Comparing Figs. 2 and 3, however, there remains a gap in realism. For this, models trained with simulated data which is synthetically altered to be more realistic perform better than models trained with data generated by simulators as well as with real data that is altered with classical image-processing techniques (blurring, flipping, etc) [28]. Here, learned augmentations can achieve more complex appearance transformations that classical augmentations do not capture – useful when data features intricate patterns or features.

Therefore, in this work, we have developed an approach to generating realistic lunar surface images using a CycleGAN [29]. By converting “fake” images from the PANGU simulator into highly realistic images that accurately represent the features and characteristics of the lunar surface, we are capable of generating more realistic data. This is essential for training downstream segmentation and detection models, as our generated data is accompanied by perfect PANGU ground truth, which are laborious and expensive to perform by hand and are imperfect or incomplete for the *Lunar Crater Database* [1]<sup>1</sup> annotation of LROC dataset, as motivated in Fig. 4. Here, imperfection in human labels is often due to *subjectivity* (different ways of interpreting visual information), *ambiguity* (images with complex or am-

<sup>1</sup>[astrogeology.usgs.gov/search/map/Moon/Research/Craters/lunar\\_crater\\_database\\_robbins\\_2018](https://astrogeology.usgs.gov/search/map/Moon/Research/Craters/lunar_crater_database_robbins_2018)

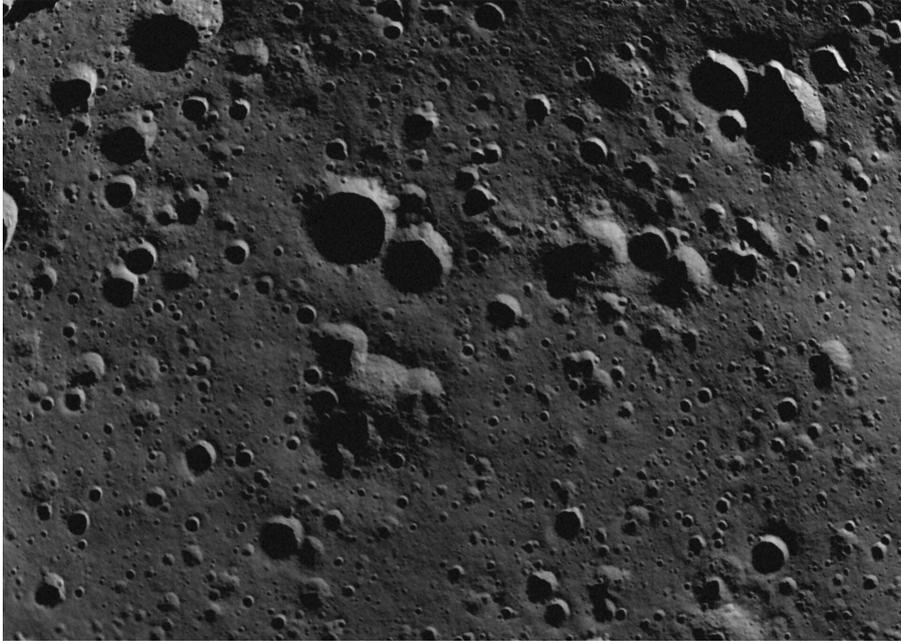


Figure 3: Sample PANGU image, taken from [www.esa.int/Enabling\\_Support/Space\\_Engineering\\_Technology/Shaping\\_the\\_Future/Planet\\_and\\_Asteroid\\_Natural\\_Scene\\_Generation\\_Utility\\_PANGU\\_Tool\\_Enhancement](http://www.esa.int/Enabling_Support/Space_Engineering_Technology/Shaping_the_Future/Planet_and_Asteroid_Natural_Scene_Generation_Utility_PANGU_Tool_Enhancement). We chose PANGU as our baseline simulator due to its ease of use and the numerous options it has to generate various terrains and its widespread adoption in terrain simulation. As the default parameters are set to resemble the lunar surface, parameters for the distribution in size and age of the craters were left untouched. One issue with the PANGU outputs is the artificially well-defined crater edges. This can be somewhat remedied by adjusting the crater edge decay option in the PANGU interface. However, manual fine-tuning of this kind can be avoided in a learned approach, which can additionally capture other non-intuitive factors which are not modelled by the procedural simulator.

biguous content), etc.

## 2. RELATED WORK

Here we discuss existing crater identification systems to identify works for which our method for synthesising more realistic lunar surface imagery will be beneficial – in training and testing.

**Manual identification** involves classification of lunar craters based on their size and morphology [1, 30], also performed in [31] for Mars impact craters by a larger team. However, this method is *time-consuming* can be *subjective* (prone to human error), and does not scale well to large projects.

**Digital image processing** with classical machine vision techniques reduces this effort, which may involve edge detection and ellipse fitting [32]. These methods must often be *tuned by hand* and are especially *prone to image noise*.

**Machine learning** techniques, instead, are optimised over examples to distinguish data based characteristic features. Early work in this area applied support vector machines (SVMs) [33] and continu-

ously scalable template models (CSTMs) [34]. These methods, however, still require *feature engineering*.

**Neural networks**, in contrast, learn hierarchical representations directly from large amounts of data. This has been applied to crater detection in various ways, including for bounding box prediction [16, 11, 13], crater rim segmentation [35, 14], and segmentation of the entire crater [10, 15].

In applying these techniques to lunar crater detection, the issue of **incomplete labelled sets** over a range of lunar crater databases is discussed in [36]. Therefore, in this work, we turn to simulators for perfect label sets, and propose learning to apply realism to the simulated samples from unpaired domain adaptation.

**Closest to our work** is [37, 38]. In [37], sim-to-real transfer involves changing image exposure and contrast, additive white Gaussian noise, and blurring – whereas our sim-to-real transfer is learned onto the real data distribution. In [38], domain adaptation is performed at the feature level, meaning that there is no sim-to-real output at the sensor level as in our work. Both [37, 38] focus on spacecraft pose estimation – with views of orbiting craft – rather than bird’s eye views of planetary landscapes.

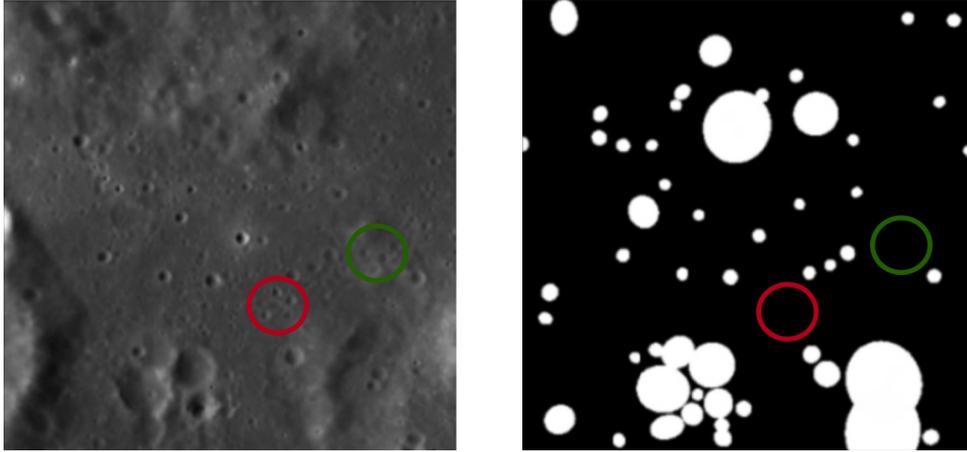


Figure 4: *Lunar Crater Database* sample (left) and corresponding binary segmentation mask (right). Some small craters are missing (e.g. in the top right of the frame) – we indicate two such clusters (green and red), but on inspection find many other missing craters of a similar size. This is detrimental when optimising neural networks (NNs) or other supervised machine learning models.

Finally, while this initial proof-of-concept investigation was carried out with generative adversarial networks (GANs), future work will leverage improvements from recent latent diffusion models (LDMs) [39] which offer better control over the generation process and exhibit improved convergence properties.

### 3. METHOD

Our system is as illustrated in Fig. 1. It is a two-stage process, consisting of first training an image translation model and then using the outputs of the image translation model to train a crater segmentation model. Therefore we are demonstrating the utility of our synthesised images in a *downstream task*. Other downstream tasks which could benefit from this more realistic source of training data may include learned pose estimation for autonomous space landers [25].

#### 3.1. PANGU

The Planet and Asteroid Natural Scene Generation Utility (PANGU) [17] is a simulator used to generate synthetic images of planetary surfaces and asteroids. It is often used in space mission planning and analysis, and it is developed and maintained by ESA. PANGU uses a combination of high-resolution topographic data, multispectral imaging data, and other parameters to create 3D models of planetary surfaces and asteroids, and then generates images of the landscapes from various viewpoints. The resulting images can be used to simulate the appearance of a spacecraft’s view during a flyby, landing, or other mission phases.

One issue with the PANGU dataset is the artificially well-defined crater edges. While this can be somewhat remedied by tuning the crater edge decay option on PANGU, it still does not solve the issues of the differing terrain appearances, and also suggests that there are other non-intuitive factors at play. Therefore, it is more beneficial to take a learned approach to match the general appearance of the terrain to a realistic corpus such as LROC.

In our work, a dataset was generated using the PANGU software. Large tiles of size 8192 by 8192 pixels were made and sliced into smaller tiles of size 416 by 416 pixels. With an overlap of 208 pixels between subsequent tiles, each large tile yielded a total of 1369 smaller tiles. The maximum crater side-length was limited to 200 pixels to ensure that no craters are bigger than the tile itself.

#### 3.2. LROC

The Lunar Reconnaissance Orbiter Camera (LROC) is a high-resolution imaging system on board the Lunar Reconnaissance Orbiter spacecraft. Launched by NASA in 2009, LRO is a robotic mission designed to map and study the surface of the Moon.

The LROC system consists of three separate cameras: two narrow-angle cameras (NACs) and one wide-angle camera (WAC). The NACs capture images at a resolution of 0.5 m per pixel, while the WAC captures images at a resolution of 100 m per pixel.

In our work, as broad views suffice to capture the general appearance of the lunar surface and considering it is this visual theme we want to synthesise, we use the WAC. To conform with the resolution in Sec. 3.1, images from this camera were sliced into



Figure 5: The left side of the image shows a fake lunar surface image that is clearly computer-generated and lacks the realism and detail of a genuine image. On the right side of the image, the synthesized lunar image created using the CycleGAN model is highly realistic and visually striking. The model has successfully captured the textures, shadows, and other characteristics of the Moon’s surface

tiles of size  $416 \times 416$  pixels. These cover the full range of longitudes on the moon and the range of latitudes between  $\pm 60^\circ$ . Each tile covers an area of approximately  $1730 \text{ km}^2$ .

### 3.3. LROC2PANGU (CycleGAN)

CycleGAN [29] enables unsupervised image-to-image translation between two domains – it learns to capture style and content information from the input images. It does this without the need for paired training data. Specifically, CycleGAN enforces cycle consistency, which ensures that when an image is translated from one domain to another and then back, it should remain similar to the original image.

Our CycleGAN<sup>2</sup> uses generator networks (which perform the image translation – from one style to another) formed by 9 contracting and 9 expanding ResNet [40] blocks, each block of which in brief consists of either a convolutional (contracting) or deconvolutional (expanding) layer, and which for ResNet are connected via shortcuts to other layers in order to avoid vanishingly small gradients deep in the network.

The discriminators (which the generators must fool into not being able to tell synthesised styles from original images) are formed by 3 convolutional layers with non-linear ReLU activations.

### 3.4. Crater segmentation (U-Net)

Given the success shown in Jackson *et al* [35] for crater detection, a U-Net model was chosen, al-

<sup>2</sup>With implementation taken from [github.com/junyanz/pytorch-CycleGAN-and-pix2pix](https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix)

though our core contribution is in improving the realism of lunar simulators and our method is applicable for other segmentation models.

U-Net features a contracting path to capture context and a symmetric expanding path to precisely localise object boundaries, making it highly effective for pixel-level image segmentation. Specifically, the output of the expanding path is supervised by (optimised to produce) binary segmentation masks as in Figs. 1, 4 and 6, with aerial views input to the contracting path.

Our U-Net implementation<sup>3</sup> consists of 4 down-sampling blocks, each using an instance-normalised convolutional layer. These blocks progressively reduce the input spatial resolution and produce deeper feature representations. This is followed by 4 up-sampling blocks that progressively up-sample the spatial resolution and return to the original image channel depth.

Our implementation of the U-Net model was trained over 30 epochs with a batch size of 15 and learning rate of  $1e^{-4}$ . The model is mainly comprised of a series of  $3 \times 3$  convolutions, with a stride of 1 and 1 pixel of padding. We use 5 convolutional layers for both the contraction and expansion path.

## 4. EXPERIMENTAL SETUP

**Datasets** Our training and testing datasets contain craters smaller than 16km in radius so as to prevent cases where the entire tile is occupied by a single crater. We consider it important to focus on these smaller craters as large craters would other-

<sup>3</sup>Similar to [github.com/milesial/Pytorch-UNet](https://github.com/milesial/Pytorch-UNet)

wise “mask out” many small craters. All of the models were trained with approximately 10000 images in the training dataset and 1000 images in the validation dataset.

**Performance metrics** We use a range of metrics to determine the performance of the model.

1. *Accuracy* is calculated as the proportion of the pixels that were predicted correctly in the entire image.
2. Intersection over Union (*IoU*) is calculated by dividing the area of intersection between the predicted and ground truth regions by the area of their union.
3. *Precision* indicates how many of the pixels predicted to be a crater were actually crater.
4. *Recall* indicates what proportion of actual craters were predicted as crater.
5.  $F_1$  is the harmonic mean of precision and recall.
6. *Specificity*, or the true negative rate, quantifies how well a classifier can avoid false crater detections by correctly identifying the regions which are not craters.

## 5. RESULTS

In the following we compare segmentation models trained with only simulated examples against models trained with our sim-to-real transfer applied, as a proof that the resulting images are more useful for understanding the lunar surface.

Fig. 5 shows (left) a PANGU sample which is then processed by our trained image translation network to appear more realistic (right). Upon visual inspection we observe that craters themselves are preserved, but appear more realistically “textured” in the synthesised result. Note that these are  $416 \times 416$  samples as prepared in Sec. 3.3, not full-size tiles as in Fig. 2.

Fig. 6 then shows (middle) the crater segmentation mask predicted from an input LROC image (left) when the U-Net model (Sec. 3.4) is only trained with PANGU images, as opposed to the predicted segmentation when U-Net is trained with our sim-to-real transfer applied to all training examples (right). As can be seen, the model is better at identifying craters, with large craters more closely following the true outline (e.g. the top right round crater in particular – with PANGU-only models predicting falsely a pair of craters) and, crucially, a much more comprehensive capture of the small craters.

Table 1: Crater detection performance metrics when training on simulated (PANGU) or synthesised (PANGU2LROC) data.

Method →	PANGU	PANGU2LROC ( <i>Ours</i> )
Accuracy [%] ↑	89.59	<b>92.53</b>
$F_1$ [%] ↑	15.29	<b>23.31</b>
IoU [%] ↑	8.46	<b>13.81</b>
Precision [%] ↑	12.71	<b>22.09</b>
Recall [%] ↑	28.41	<b>33.13</b>
Specificity [%] ↑	92.47	<b>95.37</b>

Tab. 1 lists quantitative segmentation performance metrics to support the qualitative proof of the advantage of our method in Fig. 6. These are averaged over the 1000-image validation set, see Sec. 4. We observe improvements in all performance metrics, indicating that our sim-to-real transfer produces images more useful for training crater detection models than simulation alone.

## 6. CONCLUSION

We have presented a system for closing the realism gap for rendered images from planetary simulators. We have proved the fidelity of its outputs by improving the training of a downstream lunar crater segmentation model. This system will in the future be useful for more robust in-simulator testing of lunar operations, presenting autonomous and otherwise tasks with more realistic lunar data and scenarios.

## REFERENCES

1. S. Robbins, “Developing a global lunar crater database, complete for craters  $\geq 1$  km,” in *47th Annual Lunar and Planetary Science Conference*, no. 1903, p. 1525, 2016.
2. B. Maass, S. Woicke, W. M. Oliveira, B. Razgus, and H. Krüger, “Crater navigation system for autonomous precision landing on the moon,” *Journal of Guidance, Control, and Dynamics*, vol. 43, no. 8, pp. 1414–1431, 2020.
3. L. M. Downes, T. J. Steiner, and J. P. How, “Lunar terrain relative navigation using a convolutional neural network for visual crater detection,” in *2020 American Control Conference (ACC)*, pp. 4448–4453, IEEE, 2020.
4. M. Balme, M. Curtis-Rouse, S. Banham, D. Barnes, R. Barnes, A. Bauer, C. Bedford, J. Bridges, F. Butcher, P. Caballo-Perucha, *et al.*, “The 2016 uk space agency mars utah rover field investigation (murfi),” *Planetary and Space Science*, vol. 165, pp. 31–56, 2019.

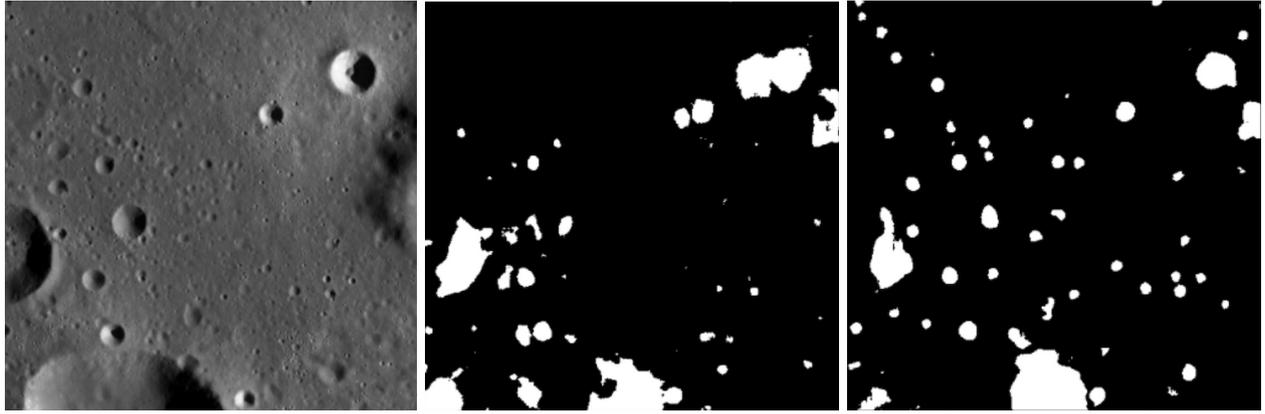


Figure 6: *Top*: Example image of PANGU trained model on LROC small crater dataset. *Bottom*: Example image of transferred PANGU trained model on LROC small crater dataset.

5. A. Lagain, K. Servis, G. Benedix, C. Norman, S. Anderson, and P. Bland, "Model age derivation of large martian impact craters, using automatic crater counting methods," *Earth and Space Science*, vol. 8, no. 2, p. e2020EA001598, 2021.
6. S. Brylow, M. Tschimmel, D. Humm, S. Lawrence, P. Thomas, E. Bowman-Cisneros, J. Zerr, M. Ravine, *et al.*, "Lunar reconnaissance orbiter camera (lroc) instrument overview," *Space science reviews*, vol. 150, pp. 81–124, 2010.
7. R. Vondrak, J. Keller, G. Chin, and J. Garvin, "Lunar reconnaissance orbiter (lro): Observations for lunar exploration and science," *Space science reviews*, vol. 150, pp. 7–22, 2010.
8. H. R. Kerner, K. L. Wagstaff, B. D. Bue, P. C. Gray, J. F. Bell, and H. B. Amor, "Toward generalized change detection on planetary surfaces with convolutional autoencoders and transfer learning," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 10, pp. 3900–3918, 2019.
9. "Deep learning applied to lunar crater counting," in *AGU Fall Meeting Abstracts*, vol. 2020, pp. P004–0004, 2020.
10. R. Del Prete, A. Saveriano, and A. Renga, "A deep learning-based crater detector for autonomous vision-based spacecraft navigation," in *2022 IEEE 9th International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, pp. 231–236, IEEE, 2022.
11. H. Seo, D. Kim, S.-M. Park, M. Choi, *et al.*, "Lunar crater detection using deep-learning," *Journal of Space Technology and Applications*, vol. 1, no. 1, pp. 49–63, 2021.
12. H. Chen, X. Hu, P. Gläser, H. Xiao, Z. Ye, H. Zhang, X. Tong, and J. Oberst, "Cnn-based large area pixel-resolution topography retrieval from single-view lroc nac images constrained with sladem," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 9398–9416, 2022.
13. R. La Grassa, G. Cremonese, I. Gallo, C. Re, and E. Martellato, "Yololens: A deep learning model based on super-resolution to enhance the crater detection of the planetary surfaces," *Remote Sensing*, vol. 15, no. 5, p. 1171, 2023.
14. L. Downes, T. J. Steiner, and J. P. How, "Deep learning crater detection for lunar terrain relative navigation," in *AIAA SciTech 2020 Forum*, p. 1838, 2020.
15. I. Giannakis, A. Bhardwaj, L. Sam, and G. Leontidis, "Deep learning universal crater detection using segment anything model (sam)," *arXiv preprint arXiv:2304.07764*, 2023.
16. Y. Jin, F. He, S. Liu, and X. Tong, "Small scale crater detection based on deep learning with multi-temporal samples of high-resolution images," in *2019 10th International Workshop on the Analysis of Multitemporal Remote Sensing Images (MultiTemp)*, pp. 1–4, IEEE, 2019.
17. S. Parkes, I. Martin, M. Dunstan, and D. Matthews, "Planet surface simulation with pangu," in *Space ops 2004 conference*, p. 389, 2004.
18. M. McCrum, S. Parkes, I. Martin, and M. Dunstan, "Mars visual simulation for exomars navigation algorithm validation," in *i-SAIRAS 2010-International Symposium on Artificial Intelligence, Robotics and Automation in Space*, pp. 283–290, 2010.
19. O. Dubois-Matra and S. Parkes, "Testing and validation of planetary vision-based navigation systems with pangu," in *International Symposium on Space Flight Dynamics*, 2009.

20. S. Willis, D. Izzo, and D. Hennes, "Reinforcement learning for spacecraft maneuvering near small bodies," in *AAS/AIAA Space Flight Mechanics Meeting*, vol. 158, pp. 1351–1368, 2016.
21. S. Woicke and E. Mooij, "A stereo-vision hazard-detection algorithm to increase planetary lander autonomy," *Acta Astronautica*, vol. 122, pp. 42–62, 2016.
22. I. Martin, M. Dunstan, S. Parkes, and M. S. Gestido, "Testing vision-based guidance and navigation systems for entry descent and landing operations," in *69th International Astronautical Congress*, pp. 1–9, International Astronautical Federation, 2018.
23. O. Sikorski, D. Izzo, and G. Meoni, "Event-based spacecraft landing using time-to-contact," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1941–1950, 2021.
24. M. Paschero, *Assessment of the performance of a visual payload mounted on a small satellite involved in an in-orbit servicing mission during the observation phase*. PhD thesis, Politecnico di Torino, 2022.
25. Z. Chekakta, A. Zenati, N. Aouf, and O. Dubois-Matra, "Robust deep learning lidar-based pose estimation for autonomous space landers," *Acta Astronautica*, vol. 201, pp. 59–74, 2022.
26. Y. Zhou, X. Li, and B. Hua, "Crater identification simulation using lidar on lunar rover," *Measurement*, vol. 210, p. 112550, 2023.
27. D. Izzo, N. Weiss, and T. Seidl, "Constant-optic-flow lunar landing: Optimality and guidance," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 5, pp. 1383–1395, 2011.
28. J. Blumenkamp, A. Baude, and T. Laue, "Closing the reality gap with unsupervised sim-to-real image translation," in *RoboCup 2021: Robot World Cup XXIV*, pp. 127–139, Springer, 2022.
29. J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, pp. 2223–2232, 2017.
30. R. Povilaitis, M. Robinson, C. Van der Bogert, H. Hiesinger, H. Meyer, and L. Ostrach, "Crater density differences: Exploring regional resurfacing, secondary crater populations, and crater saturation equilibrium on the moon," *Planetary and Space Science*, vol. 162, pp. 41–51, 2018.
31. A. Lagain, S. Bouley, D. Baratoux, C. Marmo, F. Costard, O. Delaa, A. P. Rossi, M. Minin, G. Benedix, M. Ciocco, *et al.*, "Mars crater database: A participative project for the classification of the morphological characteristics of large martian craters," 2021.
32. J. R. Kim, J.-P. Muller, S. van Gasselt, J. G. Morley, and G. Neukum, "Automated crater detection, a new tool for mars cartography and chronology," *Photogrammetric Engineering & Remote Sensing*, vol. 71, no. 10, pp. 1205–1217, 2005.
33. M. Ding, Y. Cao, and Q. Wu, "Novel approach of crater detection by crater candidate region selection and matrix-pattern-oriented least squares support vector machine," *Chinese Journal of Aeronautics*, vol. 26, no. 2, pp. 385–393, 2013.
34. P. G. Wetzler, R. Honda, B. Enke, W. J. Merline, C. R. Chapman, and M. C. Burl, "Learning to detect small impact craters," in *2005 Seventh IEEE Workshops on Applications of Computer Vision (WACV/MOTION'05)-Volume 1*, vol. 1, pp. 178–184, IEEE, 2005.
35. A. Silburt, M. Ali-Dib, C. Zhu, A. Jackson, D. Valencia, Y. Kissin, D. Tamayo, and K. Menou, "Lunar crater identification via deep learning," *Icarus*, vol. 317, pp. 27–38, 2019.
36. S. Zang, L. Mu, L. Xian, and W. Zhang, "Semi-supervised deep learning for lunar crater detection using ce-2 dom," *Remote Sensing*, vol. 13, no. 14, p. 2819, 2021.
37. P. F. Proença and Y. Gao, "Deep learning for spacecraft pose estimation from photorealistic rendering," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6007–6013, IEEE, 2020.
38. J. I. B. Pérez-Villar, Á. García-Martín, and J. Bescós, "Spacecraft pose estimation based on unsupervised domain adaptation and on a 3d-guided loss combination," in *Computer Vision—ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part I*, pp. 37–52, Springer, 2023.
39. R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
40. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.