



Kent Academic Repository

Le Corre, Daniel, Mason, Nigel, Bernard-Salas, Jeronimo, Mary, David and Cox, Nick L. J. (2025) *Automated impact melt fracture mapping on the Moon with weakly supervised deep learning*. *Journal of Geophysical Research: Planets*, 130 (12). ISSN 2169-9097.

Downloaded from

<https://kar.kent.ac.uk/112161/> The University of Kent's Academic Repository KAR

The version of record is available from

<https://doi.org/10.1029/2025je009145>

This document version

Publisher pdf

DOI for this version

Licence for this version

CC BY (Attribution)

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in **Title of Journal**, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).



Key Points:

- IMFMapper is a DeepLabV3 deep learning model trained to detect impact melt fractures (IMFs) within LROC NAC imagery of the Moon
- We propose new candidate melt ponds within the western and southern walls of Crookes crater, thanks to the detection of IMFs
- In Schomberger A crater, we present the first mapping via any means of IMFs so close to the Lunar South Pole

Correspondence to:

D. Le Corre,
dl387@kent.ac.uk

Citation:

Le Corre, D., Mason, N., Bernard-Salas, J., Mary, D., & Cox, N. L. J. (2025). Automated Impact melt fracture mapping on the Moon with weakly-supervised Deep Learning. *Journal of Geophysical Research: Planets*, 130, e2025JE009145. <https://doi.org/10.1029/2025JE009145>

Received 13 APR 2025

Accepted 13 NOV 2025

Author Contributions:

Conceptualization: Daniel Le Corre
Data curation: Daniel Le Corre
Formal analysis: Daniel Le Corre
Funding acquisition: Nigel Mason, Jeronimo Bernard-Salas
Investigation: Daniel Le Corre
Methodology: Daniel Le Corre
Project administration: Daniel Le Corre, Nigel Mason, Jeronimo Bernard-Salas
Resources: Daniel Le Corre
Software: Daniel Le Corre
Supervision: Nigel Mason, Jeronimo Bernard-Salas, David Mary
Validation: Daniel Le Corre
Visualization: Daniel Le Corre
Writing – original draft: Daniel Le Corre
Writing – review & editing: Daniel Le Corre, Nigel Mason, Jeronimo Bernard-Salas, David Mary

© 2025. The Author(s).

This is an open access article under the terms of the [Creative Commons Attribution License](#), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Automated Impact Melt Fracture Mapping on the Moon With Weakly Supervised Deep Learning

Daniel Le Corre^{1,2} , Nigel Mason¹ , Jeronimo Bernard-Salas² , David Mary³, and Nick L. J. Cox² 

¹Centre for Astrophysics and Planetary Science (CAPS), University of Kent, Canterbury, UK, ²Centre d'Etudes et de Recherche de Grasse (CERGA), ACRI-ST, Grasse, France, ³Lagrange UMR 7293, Observatoire de la Côte d'Azur (OCA), Université Côte d'Azur, Nice, France

Abstract Cooling fractures found within impact melt deposits have been manually mapped within several craters on the Moon and Mercury, as their distribution can indicate which heat-loss processes were most significant in the periods after impact. However, due to the discovery of melt deposits in Lunar impact craters with sub-km diameters, it is unlikely that the complete mapping of these impact melt fractures (IMFs) on the Moon will be achievable without automation. As such, we have trained a DeepLabV3 semantic segmentation deep convolutional neural network, called IMFMapper, to detect IMFs within Lunar Reconnaissance Orbiter Narrow Angle Camera (LROC NAC) satellite imagery. As a means of maximizing the size of the training data set, “weak” pixel-level labels were generated by buffering line annotations. In testing upon the IMFs found within Ohm crater, IMFMapper achieved an average F₁-score of 69.3%. IMFMapper has also been deployed to map IMFs within the previously surveyed Crookes crater, where we have found new candidate melt deposits within the crater's western and southern walls. In addition, IMFMapper has produced the first map of IMFs within Schomberger A crater, in which IMFs may act as permanently shadowed regions due to the crater's proximity to the Lunar South Pole. The successful mapping of IMFs in Schomberger A also signifies IMFMapper's robustness to extreme solar incidence angles. We also demonstrate that IMFMapper could be implemented for automated mapping of IMFs on Mercury upon the commencement of BepiColombo's science operations.

Plain Language Summary After a meteorite hits a planetary body at hyper-velocity speeds, the surface material is melted into molten rock under the immense energies of an impact event. As the crater forms underneath, this pool of molten rock—often called an impact melt deposit - contracts as it cools over time. This contraction causes linear fractures to pop open on the surface of the impact melt deposit. Previously, these impact melt fractures (IMFs) have been manually mapped by researchers on the Moon and Mercury in order to investigate which cooling mechanisms may be most influential on airless bodies such as these. In this work, we use a deep learning model to automate this process. Our model, called IMFMapper, builds upon the previous manual mapping of Crookes crater on the Moon and finds new impact melt deposits within its western and southern walls. We also use IMFMapper to catalog the IMFs found within the Moon's Schomberger A crater, which could potentially trap reserves of water ice within them due to being deprived of sunlight at such high latitudes.

1. Introduction

Impact melt fractures (herein referred to as IMFs) are cracks found in the melt deposits of impact craters, which can often occur in curvilinear or polygonal complexes. Figure 1 presents several patterns, scales and morphologies of IMFs as found on the Moon. The cracks are thought to form due to extensional stresses as the deposited shock-heated material (either found in the crater's floor or ejected from the site) cools and contracts after the event of an impact (Xiao et al., 2014). The presence of solid debris in a melt pond, such as material falling from the crater's walls, will affect the uniformity of the stresses experienced by the melt (Denevi et al., 2012) - and, by proxy, the patterns of IMFs observed. Due to the conditions under which these processes occur being similar on the two bodies, fractures relating to the cooling of impact melt deposits have been observed on both the Moon and Mercury (Xiao et al., 2014).

Impact melt deposits (also known as melt ponds) and the IMFs found within them have been the subject of several mapping exercises. Xiao et al. (2014) have carried out surveys of IMFs within the young impact melt deposits




Figure 1. Examples of impact melt fractures found in (left) Virtanen F, (middle) Copernicus, (top right) Tycho and (bottom right) Lalande craters on the Moon.

belonging to six complex impact craters on the Moon and Mercury. On the Moon, the Copernican-aged craters Tycho, Aristarchus and Copernicus were surveyed, whilst the Kuiperian-aged Degas, Abedin and Hokusai craters were studied on Mercury. Xiao et al. (2014) categorize their mapped IMFs according to whether they occur within the proximity to the center of the crater floor, or at the margins between the floor and crater walls. From the resulting maps, Xiao et al. (2014) deduced that IMFs were most spatially concentrated near the crater floor-wall boundary and around solid debris within the melt deposit.

As part of their morphological analysis, Thaker et al. (2020) mapped IMFs with Crookes and Ohm craters on the Moon. Crookes and Ohm are both Copernican-aged, complex impact craters found on the far side of the Moon with diameters of 48 and 62 km, respectively (IAU, 2025). Despite also creating a morphological map for Stevinus crater on the Moon's near side, a complete lack of IMFs was observed in this instance. Thaker et al. (2020) also found that both Crookes and Ohm displayed IMFs within the melt ponds situated on the terraces of their crater walls. Figure 8 of Thaker et al. (2020) gives evidence of these IMFs found within the northern melt pond of Crookes. Via similar geologic mapping exercises, IMFs have also been identified within Jackson (Dhingra et al., 2017) and Lowell (Srivastava & Varatharajan, 2016) craters on the Moon, but neither included IMFs within their final maps. The presence of IMFs within the visible southern half of the melt deposit with Schomberger A crater has also been discovered (Lee, 2022). Schomberger A is of particular interest for its proximity to the south pole (78.61°S), since the IMFs within it have the potential to act as permanently shadowed regions (PSRs). PSRs are topographic minima with sufficiently extreme latitudes such that they receive no sunlight, meaning that they can cold-trap volatiles such as water ice (Brown et al., 2022).

Mapping of IMFs has also been completed as part of the Lunar Pit Atlas (LPA). The LPA gives the positions, morphometric parameters, and descriptions of 281 Lunar pits (Wagner & Robinson, 2021). LPA features are subcategorized according to their location, with the majority being found within impact melt deposits. Of the 256 impact melt pits, 8 have been specifically designated as being “fractures,” with more than 90 also being associated with fracture systems. These 8 LPA fractures can be found within the craters: Kepler, King, Messier A, Ohm and Tycho.

The studies discussed up to this point have all focused on multi-kilometer-scale impact craters. However, these are not the only places where impact melt deposits have been found. Melt deposits have been found within small

impact craters situated in the Lunar highlands. Whilst these instances are thought to be rare, melt deposits were found in craters as small as 170 m in diameter, with IMFs also being observed “in a few cases” (Plescia & Cintala, 2012). Melt ponds have also been mapped in the vicinity of five impact craters near the Moon's south pole (Gagnon et al., 2023), which is the proposed landing site for the Artemis III mission.

The identification and morphological analyses of IMFs to date have been enabled by the breadth of high-resolution imagery available for the surfaces of the Moon and Mercury. On the Moon, the Lunar Reconnaissance Orbiter's Narrow Angle Camera (LROC NAC) reveals IMFs with a typical resolution of ≈ 0.5 m/px. Whereas for Mercury, the surface detail is more limited due to MESSENGER's Mercury Dual Imaging System's Narrow Angle Camera achieving resolutions of 12 m/px at best during its lifetime (Hawkins et al., 2007; Solomon et al., 2001).

Automated techniques, such as Machine and Deep Learning (ML/DL), are being increasingly used throughout planetary science, particularly for the detection of planetary surface features within remote-sensing data. Impact craters are the most common subject of studies using ML/DL for surface feature detection on planetary bodies (Chaini & Jha, 2024). Despite this dominance of crater detection algorithms, a DL model has been trained to detect pits/skylights on the Moon and deployed to discover two previously uncatalogued candidate cave entrances (Le Corre et al., 2025). ML/DL algorithms have also been developed to map linear surface features such as sinuous rilles (Zhang et al., 2024) and rockfalls (Bickel et al., 2020) on the Moon, as well as dark slope streaks (Wang et al., 2017) and transverse aeolian ridges (Barrett et al., 2023) on Mars.

Even with the advancing field of planetary surface feature detection, there has yet to be any attempt at automatically mapping IMFs with ML/DL. However, there is significant scope for doing so, given how the presence of melt deposits in both complex and sub-km simple craters means that it is unlikely that the complete mapping of Lunar IMFs will be a manually completable task. With the ability to perform global spatial and morphometric analysis of IMFs, the dominant cooling mechanisms of smaller impact melt deposits, as well as those in the floors of complex craters, can be more confidently understood. It is also perceivable that an automated tool for detecting IMFs could be used as a proxy for the identification of uncatalogued melt deposits. The discovery of readily accessible melt deposits would be relevant to Science Objective 3 (“Interpreting the Impact History of the Earth-Moon System”) of the Artemis III mission, since one of its investigations is to collect a sample of impact melt (Weber et al., 2021).

In this work, we have developed a DL model which can map IMFs within LROC NAC imagery of the Moon's surface. This model, named IMFMapper, has been trained on LROC NAC observations of over 4,000 manually labeled IMFs. As explained in further sections, a novelty of IMFMapper is that it was trained upon “weak” labels (pixel-level labels produced from line annotations) in order to increase the speed and ease with which annotations are made to ensure that a large, high-quality data set could be built in acceptable time frames. IMFMapper has been validated upon two data sets, the first drawn from unseen imagery of the same craters as in training, and the second being a subset of IMFs found within the Copernicus and Virtanen F impact craters. To deduce its ability to generalize to new data, IMFMapper has been independently tested against the IMFs seen in the southern half of Ohm crater's floor. IMFMapper has also been applied to self-processed mosaics of LROC NAC images of Crookes and Schomberger A craters. Thus, enabling a qualitative comparison with the mapping of Crookes crater by Thaker et al. (2020), and representing the first attempt at mapping IMFs so close to the Lunar South Pole. A secondary objective of IMFMapper's development was to set the groundwork for automatically mapping IMFs on Mercury upon the arrival of the joint ESA and JAXA BepiColombo mission. The SIMBIO-SYS (Spectrometers and Imagers for Mercury Planetary Orbiter (MPO) Integrated Observatory System) aboard the ESA-provided MPO aims to provide 6 m/px panchromatic imagery for $\approx 20\%$ of the surface within the first year of operation (Benkhoff et al., 2021). Therefore, we demonstrate the possibility of using IMFMapper to take full advantage of this larger and higher-resolution data set of IMFs on Mercury.

2. Data Set

A significant amount of labeled data is required to train a supervised DL model. Therefore, the footprints of LROC NAC images were cross-referenced with the positions of LPA impact melt pits as a proxy for finding high-resolution satellite imagery containing Lunar IMFs. As a result, 25 LROC NAC images were earmarked for labeling as training and validation data after manual inspection revealed that they contained IMFs. These 25 images exhibited sufficient numbers of IMFs such that a high-quality training data set could be built. The IMFs in

these images also exhibited a range of sizes and morphologies under varying sensing conditions. The majority of the solar incidence angles seen in these images range between 25 and 55°, whilst the absolute minimum and maximum were 11.0 and 74.2°, respectively.

For reasons explained in Section 4.1, a second validation data set was necessary to deduce the highest performing combination of model parameters. As such, 6 LROC NAC images containing a selection of the IMFs within Copernicus and Virtanen F craters (neither of which were used for sources of training data) were retrieved. The unambiguous IMFs seen in Copernicus crater, and the complexes of thin (widths less than 10 m) IMFs found in Virtanen F, provided a range of difficulty that was helpful for finding the model that is most capable of generalizing to new data. To avoid confusion, this additional validation data set and the validation data set drawn from 25 training images are herein referred to as Va_2 and Va_1 , respectively.

For an independent assessment of the performance that can be expected upon new data, 8 LROC NAC images covering the IMFs found in the southern half (below 18.35°N) of Ohm crater were also acquired. Whilst one of the 25 images reserved for training also overlapped Ohm, there are relatively few IMFs present and those that are, are all clustered north-east of its central peak. This, along with the fact that the IMFs were not endemic to Ohm, meant that this was one of the largest sources of testing data that could be considered independent.

As mentioned in Section 1, Crookes and Schomberger A are chosen as the sites for applying the trained model as examples of automatically mapping IMFs. The LROC NAC images overlapping the two sites were found using JMARS. JMARS (Java Mission-planning and Analysis for Remote Sensing) is a GIS tool for visualizing and analyzing the satellite data retrieved by various NASA and other agencies' missions (Christensen et al., 2009). JMARS was purposely chosen for its "Mosaic Select Tool", which removes any images from the selection of overlapping LROC NAC products that do not improve the coverage of the crater. Through the use of JMARS, a total of 56 and 55 LROC NAC images were consequently identified for Crookes and Schomberger A craters, respectively. The 56 images of Crookes crater had solar incidence angles between 30 and 60°, whereas the 55 for Schomberger A were between 76 and 80° due to its extreme latitude.

Between September 2009 and December 2011, while LRO was still in its original 50 km circular polar orbit, the majority of LROC NAC images had a resolution of ≈ 0.5 m/px (Wagner et al., 2024). Since then, LRO has been in a fuel-conserving elliptical polar orbit with its periapsis near the Lunar South Pole, meaning that images taken in the northern and southern hemispheres range in resolutions from 1.0 to 2.0 and 0.4 to 1.0 m/px, respectively (Wagner et al., 2024). Therefore, the LROC NAC images described in this section had a range of resolutions which needed to be homogenised. As a result, after the raw LROC NAC image products were acquired from PDS, they were processed and map-projected to a common resolution of 1.5 m/px using the ISIS (Integrated Software for Imagers and Spectrometers) and GDAL (Geospatial Data Abstraction Library) software suites (Laura et al., 2023; GDAL/OGR contributors, 2025). This value was chosen since the lowest resolution of the used images equaled 1.5 m/px when rounded to the nearest 0.1 m. However, in order to avoid repeat detections of the same IMF when mapping Crookes and Schomberger A craters, mosaics were created by stitching together the resulting processed and map-projected ISIS cubes.

Table 1 provides a list of the impact craters containing IMFs, which were used for either training and validation, testing, or inferring the model. Alongside the names of each crater is its location, the number of IMFs which were labeled in this work, and the stage of model development that it was responsible for. The numbers of labeled IMFs are only given to highlight which impact craters would have the most influence on the model's training/validation or testing results—rather than stating the genuine total number of IMFs within each crater. This is because there was occasionally some overlap between two or more images of the same crater, and each crater was not wholly mapped for IMFs.

3. Methodology

3.1. Creating a Weakly Labeled IMF Training Data Set

Creating a high-quality training data set of IMFs as they appear in remote-sensing imagery presents several challenges. First is the linear morphology of these features, which can be kilometers-long and as small as a few meters in width, as shown in Figure 1. This means that a training image which encapsulates an entire IMF, from start to end, could become too large to pass through a DL model due to memory constraints or could have such a low resolution that thinner IMFs would be unresolvable. Therefore, training images will frequently contain IMFs

Table 1

Lunar Impact Craters Used for Model Training (Tr), Validation (Va_1/Va_2), Testing (Te), or Inference (In)

Used for	Crater name	D [km]	Lat, Lon [°]	No. IMFs ^a
Tr+ Va_1	Aristarchus	40	23.73, −47.49	229
	Das	36	−26.49, −137.05	14
	Jackson	71	22.05, −163.32	747
	Kepler	29	8.12, −38.01	45
	King	76	4.96, 120.49	439
	Klute W	31	37.98, −143.31	77
	Lalande	24	−4.46, −8.65	196
	Messier A	11	−2.03, 46.94	42
	Ohm (North)	62	18.32, −113.78	50
	Philolaus	71	72.22, −32.88	13
	Rutherford	50	−61.15, −12.25	43
	Tharp	13	−30.60, 145.63	77
	Tycho	85	−43.30, −11.22	2,118
Va_2	Wiener F	45	41.19, 149.97	151
	Copernicus	96	9.62, −20.08	164
	Virtanen F	11	15.79, 177.32	206
Te	Ohm (South)	62	18.32, −113.78	571
In	Crookes	48	−10.40, −165.10	n/a
	Schomberger A	29	−78.61, 23.52	n/a

Note. Crater diameters and coordinates are from IAU (2025). ^aNumber of IMFs are the amounts labeled, not necessarily the genuine total.

which have been partially cut off by the image's extent. This was one of the motives for utilizing a semantic segmentation model, whereby individual pixels are assigned a class. On the other hand, object detection models, which learn the features that constitute an object, may confuse cut-off IMFs with other non-linear depressions.

In using a semantic segmentation model, another issue arises in that they first require pixel-level labels. As Table 1 foreshadows, the number of IMFs that were seen in the imagery acquired in Section 2 was so high that manually labeling them as polygons would be unachievable within available time frames. Pixel-level labels of some of the smallest IMFs, which could often be just a few pixels wide, would also be an unnecessary degree of detail. Therefore, we explored the possibility of producing weak pixel-level labels of IMFs, which would enable a sufficiently large data set to be built without a significant decrease in their accuracy.

Weakly supervised learning is the principle of training a supervised ML or DL model, but with the labeling of the data being incomplete, artificial or inaccurate. In this instance, instead of labeling IMFs as polygons at the first attempt, we annotate them along their entire path as lines. This labeling was completed in QGIS (QGIS Development Team, 2025) since it has a facility for buffering lines by a certain map unit (i.e., meters or degrees) to produce a linear polygon. Therefore, from these loose labels of the paths of IMFs, weak pixel-level labels can be generated. However, IMFs are observed to have a range of widths, meaning that a buffer value that works for one IMF will not work for all. Therefore, line labels were given additional classifications according to their approximate width in pixels, which were: “tiny” (0–5 pixels), “small” (5–10), “medium” (10–20), “large” (20–30) and “huge” (30 or more). The most frequently labeled classes, along with their average lengths (in degrees), were as follows: tiny (51%, 0.002°), small (33%, 0.004°), medium (13%, 0.006°), large (3%, 0.011°) and huge (<0.01%, 0.021°).

To maximize the reliability of the performances made upon the Va_2 data set, all features within the 5 LROC NAC images of Copernicus crater were labeled as polygons, whilst the thin IMFs in Virtanen F were still accurately approximated with the weak line labels. Due to the much larger numbers of IMFs in just the southern half of Ohm crater alone, the IMFs within the 8 LROC NAC images reserved for testing were also labeled as weak line labels. Figure 2 gives an example of the weak line labels and how they have been differentially buffered to generate pixel-level polygons ready for training a DL model. Figure 2 also shows two instances (which occurred for approximately 5% of features) where IMFs were labeled as polygons due to being LPA impact melt pits which resembled IMFs or an element of an IMF which was too difficult to label as a line.

Entire LROC NAC image products are far too large to pass through a DL model, as they can be thousands and tens of thousands of pixels in width and height, respectively, even after downscaling to 1.5 m/px. As a result, all processed LROC NAC image products reserved for training and validation were tiled into a coordinate grid of 512×512 px crops with a 10% overlap between consecutive tiles. An overlap was used to minimize the number of cases where an IMF was only ever being seen as cut-off by the edge of a tile, with a value of 10% balancing the risk of over-fitting. Implementing an overlap meant that such instances occurred for less than 3% of labeled IMFs. Any resulting tiles that did not contain any IMFs were also used for training and validation, with the aim of reducing the number of false positives. Nevertheless, where IMFs were present in a given tile, the weak polygon labels were rasterized using facilities within the GDAL in order to produce binary masks of their locations. Additionally, binary masks of the valid data regions within an image tile were also exported, which allowed for the removal of any adverse detections within the no-data regions after the loss calculation but before the performance metrics were calculated (see Section 3.2 for further details).

In order to bolster the size of the training and Va_1 data sets, which were split 4:1, some data augmentation is necessary. As such, rotated versions (clockwise by 90, 180 and 270°) of training and Va_1 tiles, except those which did not contain IMFs, were included in the data sets - thus quadrupling the number of IMFs seen by the model. To




Figure 2. An example of the weak labeling of impact melt fractures (IMFs) in the LROC NAC image M126622485R (at 1.5 m/px). (a) The impact crater “Messier A” within the LROC WAC global mosaic (Speyerer et al., 2011), showing M126622485R and the location of the IMFs shown in (b). (c) The manually produced weak line labels and polygon annotations of the Lunar Pit Atlas (LPA) features. (d) The pixel-level polygon labels created by buffering the lines in (c) and merging with the LPA polygons.

mitigate the risk of over-fitting, training samples also had the possibility of undergoing three further steps of data augmentation, each with a probability of 0.5 and no mutual exclusivity. First, images could undergo Contrast-limited Adaptive Histogram Equalization, which is a variant of Adaptive Histogram Equalization that clips the histogram to a chosen value to minimize the amplification of noise in uniform regions (Pizer et al., 1987). Second, Gaussian noise could also be applied to the images with a σ randomly selected from a uniform distribution between 0.0075 and 0.0125. Lastly, the perspective of the image could also be transformed to simulate different viewing angles. This was done using the `perspective` function within PyTorch and a distortion scale randomly selected from a uniform distribution between 0.5 and 1.5. The values for the severity of each augmentation were chosen by manually inspecting examples of augmented images to see if they were still realistic.

The same procedure as that described above was used to generate tiles for the 6 images used for Va_2 , the 8 testing images, as well as the Crookes and Schomberger A mosaics, except that there was no overlap between consecutive tiles, nor any data rotation or augmentation. Table 2 details the number of image tiles within each of the training, validation, testing and inference data sets. Where applicable, these quantities are given separately as to whether they contained labeled IMFs or not.

3.2. Model Architecture and Evaluation

In this work, we use a DeepLabV3 Deep Convolutional Neural Network (DCNN) for performing semantic segmentation on IMFs in tiled LROC NAC imagery (Chen et al., 2016). Some semantic segmentation models use an Encoder-Decoder structure, whereby images are downsampled to “encode” a condensed representation of the data and then upsampled by the decoder to the original resolution. However, the foundations of the DeepLabV3 model are built on the principle of atrous convolution. Deriving from the French for holes (“trous”), atrous

Table 2
Number of Image Tiles in the Training, Validation, Testing, and Inference Data Sets

Data set	Craters	No. Image Tiles ^a
Training (Tr)	See Table 1	4,500/2,696
Validation (Va ₁)		1,132/667
Validation (Va ₂)	Copernicus	127/2,212
	Virtanen F	24/139
Testing (Te)	Ohm (South)	320/3,038
Inference (In)	Crookes	3,332
	Schomberger A	1,341

^aTr, Va and Te data sets are split by those that do (left) or do not (right) contain IMFs.

convolution is where zeroes are embedded within a kernel to increase the receptive field of a convolutional layer without increasing the number of parameters that require learning (Chen et al., 2017).

DeepLabV3 implements atrous convolution to form an Atrous Spatial Pyramid Pooling (ASPP) block. Traditional Spatial Pyramid Pooling (SSP) structures attempt to solve the issue of detecting several scales of features by performing a range of resampling operations and pooling them together (He et al., 2014). Meanwhile, ASPP replaces the image resampling with atrous convolution operations. The degree to which atrous convolution widens the receptive view is defined by the rate, which is equal to $N_{trous} + 1$ (where N_{trous} is the number of zeroes inserted between two consecutive kernel values in each dimension). This is shown (not to scale) in Figure 3.

The focus on learning feature maps at several scales is particularly advantageous in this work since IMFs can be observed with a wide range of different widths and lengths. Other semantic segmentation structures, such as

U-Net (Ronneberger et al., 2015), are also tailored to multi-scale detection, but DeepLabV3 is more suited to this use case due to its reduction in the number of parameters which require learning. This reduces the risk of overfitting to what is, in this work, a relatively small training data set by DL standards. Another benefit to the use of DeepLabV3 is that it is available within PyTorch (see <https://pytorch.org/vision/main/models/deeplabv3.html>) with pre-trained weights (for both the backbone and remaining DeepLab head layers) upon a subset of the COCO (Common Objects in Context) data set (Lin et al., 2014). The pre-trained DeepLabV3, with ResNet50 and ResNet101 (He et al., 2016) used as the backbones, achieved a mean Intersection over Union (IoU) upon the validation COCO data set of 66.4% and 67.4%, respectively.

In this work, Binary Cross-Entropy (BCE) loss, also known as logarithmic loss, is used as the function whose objective it is to be minimized when calculated upon the training samples. This is achieved by the Adam optimizer (Kingma & Ba, 2017), which changes the model weights accordingly with a standard learning rate of 0.0001. Ideally, a large batch size should be used to limit the effect that a single training sample has on the changing of weights during one pass through the model. Whilst a batch size of 32 could be used in conjunction with the ResNet50 backbone, memory errors occurred due to the addition of more convolutional layers when swapping in the ResNet101 backbone. Therefore, DeepLabV3 models with a batch size of 16 were also trained. All three of the following models were trained for 100 epochs: ResNet50 with a batch size of 32 (ResNet50-32), ResNet50 with a batch size of 16 (ResNet50-16) and ResNet101 with a batch size of 16 (ResNet101-16). A duration of 100 epochs was selected since it was the point at which the training BCE loss began to plateau. However, model checkpoints were saved whenever a new highest validation F_1 was recorded during its training. This means that the final model state occurred at the epoch where the validation F_1 was maximized, not necessarily the state at the end of the 100 epochs.

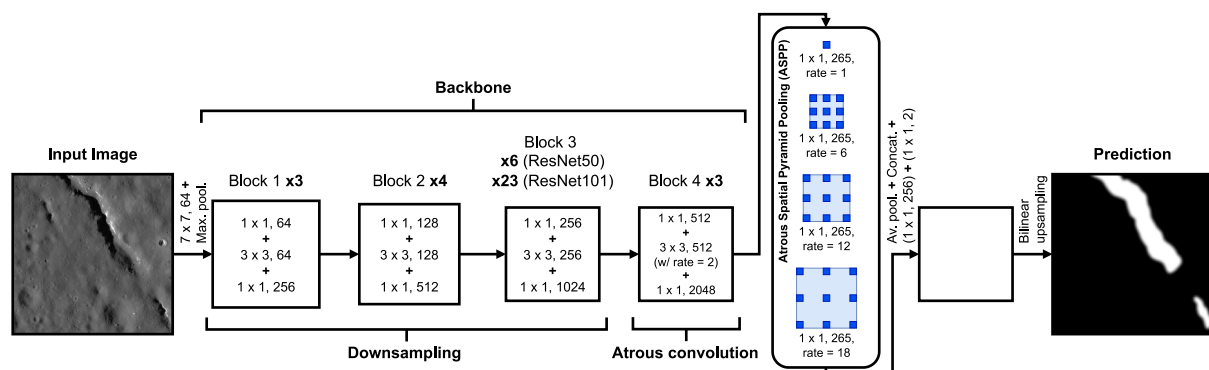


Figure 3. Diagram of the DeepLabV3 Deep Convolutional Neural Network (Chen et al., 2016) with ResNet50 or ResNet101 backbones (He et al., 2016)—within the context of impact melt fracture semantic segmentation. Convolutional layers are represented by the kernel size and the number of output channels (e.g., $3 \times 3, 126$ for a kernel size of 3 and 126 output channels).

Table 3
Average Performance Metrics Upon Impact Melt Fractures in the Va_1 Data Set

Backbone	Batch size	Epoch ^a	P [%]	R [%]	F ₁ [%]
Resnet50	32	47	77.4	62.8	69.4
Resnet50	16	93	80.0	61.9	69.8
Resnet101	16	55	78.1	63.5	70.0

Note. The highest values for each metric are highlighted in bold. ^aThis is the epoch when the F₁-score was at its maximum over the course of training.

Model training, testing and inference were performed using the specialist and high-performance computing systems provided by Information Services at the University of Kent, of which a total of 13 nodes (sharing an Nvidia A100 40 GB PCIE GPU and 256 GB of RAM) were available. With these resources, the training of a single DeepLabV3 model took approximately 10–12 hr.

For the purposes of assessing the model's training, validation and testing performance, we calculate the precision (P), recall (R) and F₁-score (F₁). P and R are defined in Equations 1 and 2, with the true/false positives/negatives (TP, FP, TN, FN) corresponding to the number of pixels which have been correctly/falsey detected/missed. Since this is a binary classification scenario, a confidence threshold of 50% was used for any pixel to be classified

as an IMF or not. To avoid division by zero errors, we avoid calculating R for tiles where there were no IMFs to recall. The same is true for P for tiles where no detections were made. The F₁, given in Equation 3, is the harmonic mean between the average P and R across all batches. Additionally, the IoU between the detections made upon the Va_2 and testing data sets and the associated labels was also calculated as a means of assessing whether the performance differed between the three craters. Whilst IoU is typically described as the ratio between the areas of intersection and union, in the context of semantic segmentation, this can be defined by Equation 4. Definitions for popular ML metrics, such as those used in this work, can be found in Lindholm et al. (2022).

$$P = \frac{TP}{TP + FP} \quad (1)$$

$$R = \frac{TP}{TP + FN} \quad (2)$$

$$F_1 = 2 \times \frac{P \times R}{P + R} \quad (3)$$

$$IoU = \frac{TP}{TP + FP + FN} \quad (4)$$

4. Results and Discussion

4.1. Training and Validation Performance

Table 3 lists the average performance metrics that were reached by each combination of backbone and batch size on the Va_1 data set at the epoch when the validation F₁ was at its highest over the whole 100 epochs of training. This indicates that there was very little to separate the performance of the three models, with only a 0.6% difference in their F₁. The ResNet50-32 model reached its maximum F₁ within the fewest number of epochs (47), suggesting that it has had less time to over-fit to the training data set than the other two models.

As the three models all exhibited very similar profiles, Figure 4 visualizes the various performance metrics achieved by solely the ResNet50-32 model on the training and validation data sets during the 100 epochs of training. Figure 4 highlights a significant difference between the training and validation R by the end of the training process, which may be a sign that the combination of data augmentation steps taken in Section 3.2 was too strong. With each having a probability of occurring of 50%, a bias toward detecting IMFs within augmented data as opposed to non-augmented data may have arisen, since the models have been fed that the majority of the time. The slowly increasing validation BCE loss from around 10 epochs (as experienced by all models) may also be an indication that they would benefit from the addition of more training data, or perhaps a smaller learning rate.

Due to the comparable performances (in terms of F₁-score) of the three models, combined with the suspicion that the ResNet50-32 will have been less over-fitted to the training data, the three models required comparison against an additional validation data set (Va_2) that included IMFs from craters that were not seen during training. From this, it can be deduced which model is best suited to generalizing to new data, and thus should be the model applied to the mosaics of Crookes and Schomberger A craters. Section 4.2 presents the results of this analysis.

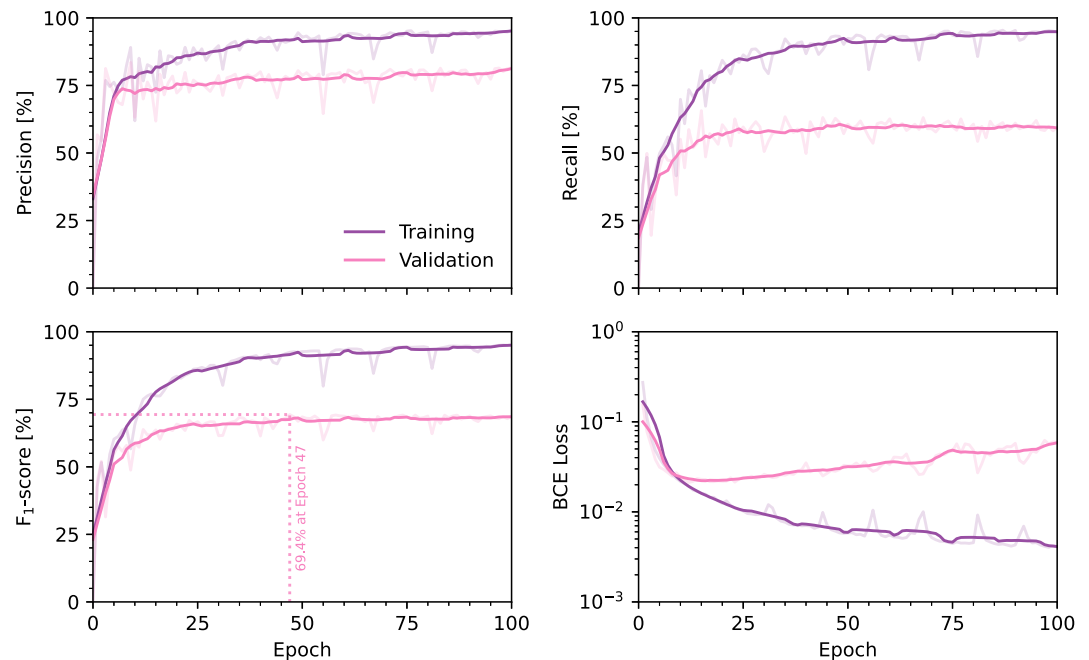


Figure 4. Performance metrics calculated upon the training and validation (Va_1) data sets when during the training of the ResNet50-32 model - called IMFMapper. The solid lines represent the rolling average across 10 epochs, with the semi-transparent ones being the genuine values. The highest validation F_1 -score achieved has been annotated with a dotted line.

4.2. Validation on IMFs in Copernicus and Virtanen F Craters

Table 4 details the average performance metrics calculated across the entire validation (Va_2) data set of IMFs in Copernicus and Virtanen F craters for the same models in Table 3. This shows that, despite the ResNet101-16 model performing marginally better on the Va_1 data set, the ResNet50-32 model performed notably better in terms of R and F_1 in this instance compared to the other two models. Moreover, ResNet101-16 performed significantly worse across all metrics on the Va_2 data set.

The fact that ResNet50-32, which was trained with a larger batch size and reached its maximum F_1 in the fewest epochs, had the highest performance on the Va_2 data set suggests that the others (ResNet50-16 and ResNet101-16) were likely over-fitted to the training data. A possible explanation for why a different model has performed optimally on the two validation data sets is that some LROC NAC images labeled for training were overlapping, meaning that the same IMF could be in both the training and Va_1 data set, just under varying sensing conditions. Therefore, the models which were over-fitting were perhaps becoming progressively more likely to detect those IMFs over the course of the training process. As a result, we select the ResNet50-32 model for use upon the mosaics of Crookes and Schomberger A craters—and herein nickname the model “IMFMapper.”

Figure 5 presents examples of the visual performance of IMFMapper upon the IMFs within Copernicus and Virtanen F craters. The detections made by IMFMapper within the images of Copernicus (all those in Figure 5 except M156435317L) show that IMFs with prominent shadows and enclosed debris were most frequently detected, with more ambiguous sections being missed. From M156435317L,

it can also be seen that IMFMapper can pick out some of the thinnest IMFs, with a sufficient recall for finding IMF complexes. However, where such features were detected, it was occasionally with sparse extents, leading to incorrectly predicted widths and/or connected sections.

As the values in Tables 3 and 4 suggest, there is some discrepancy between the average F_1 of IMFMapper upon the validation and testing data sets (69.4% and 55.9%, respectively). This is a consequence of the Va_2 data set deliberately representing a more difficult distribution of IMFs. For example, features similar to the large bowl pits cataloged in the LPA within Copernicus, and

Table 4
Average Performance Metrics Upon Impact Melt Fractures in the Va_2 Data Set

Backbone	Batch size	IoU [%]	P [%]	R [%]	F_1 [%]
Resnet50	32	38.7	55.7	56.1	55.9
Resnet50	16	36.2	56.9	50.1	53.3
Resnet101	16	34.0	50.0	51.6	50.8

Note. The highest values for each metric are highlighted in bold.

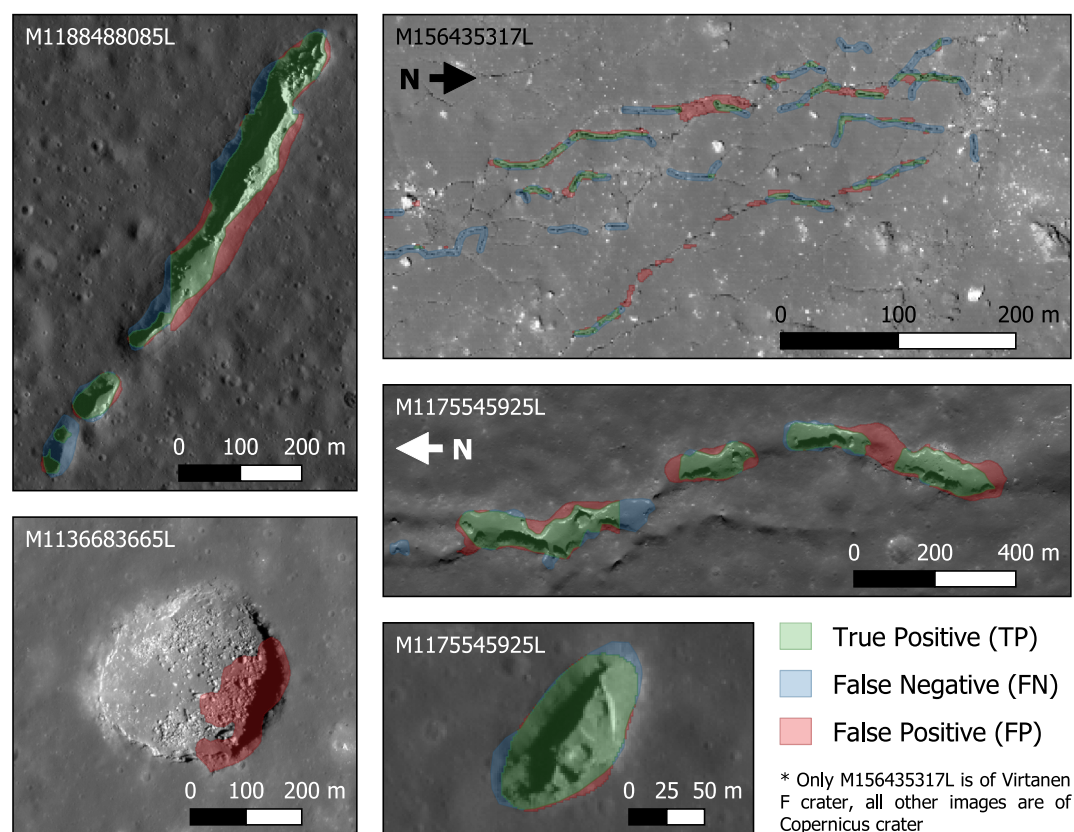


Figure 5. Visual performance of IMFMapper upon the Va_2 data set, containing impact melt fractures (IMFs) from Copernicus and Virtanen F craters. All of the LROC NAC images shown contain IMFs from Copernicus crater, except M156435317L, which is of Virtanen F. The bottom of the two maps from M1175545925L shows the Lunar Pit Atlas feature “Copernicus 4” being detected.

shown in M1136683665L of Figure 5, had not been seen during training. The small widths of the IMFs in Virtanen F meant that they would also be significantly trickier to detect than the majority of those seen during training/validation. This is confirmed by the fact that the IoU calculated between the detections and labels made by IMFMapper within Virtanen F was 21.2%, as opposed to 40.1% for Copernicus.

As previously mentioned in Section 2, the IMFs in Copernicus crater were labeled as polygons, as opposed to the weak line labeling of Virtanen F (as was the case for the training and validation data). The purpose of this was to be able to visualize whether the weak labels have influenced the geometries of the detections made by IMFs. Crucially, the shapes of the detections within Copernicus crater indicate that the weakly supervised training of IMFMapper has succeeded in developing a model which can delineate the pixel-level extents of IMFs without the need for equally detailed labeling.

4.3. Testing on IMFs in Ohm Crater

Following the validation of the selection of model parameters in Section 4.2, IMFMapper was applied to the testing data set containing the IMFs found within the southern half of Ohm crater. This resulted in an average P, R and F_1 -score of 67.7%, 71.0% and 69.3%, respectively, across the entire set of over 3,300 image tiles. This gives future users of IMFMapper an independent assessment of the levels of accuracy that can be expected when it is applied to new, unseen regions. As a direct comparison between the performances in Section 4.2, IMFMapper achieved an IoU of 52.9% upon the testing data set.

Figure 6 displays the visual performance of IMFMapper upon several examples taken from the testing data set. Within 3 of the examples (M1137295029L, M1318456507L, and M181917051L), Figure 6 reveals that there were instances where IMFMapper outperformed the weak line labeling in terms of accurately delineating the




Figure 6. Visual performance of IMFMapper upon the testing data set, containing impact melt fractures within the southern half of Ohm crater's floor. The colors here have the same meaning as in Figure 5.

widths of the IMFs. This is due to some IMFs falling between the width categories discussed in Section 3.1, meaning that too much or too little was labeled as an IMF. However, this is further evidence that providing a sufficiently large training data set, even if weakly labeled, has led to IMFMapper being able to correctly predict IMF widths.

Figure 6 also puts forward two examples of common false positives. First, ridges that are casting prominent shadows, as shown in M1132582219R, were often misclassified as IMFs - especially when they are degraded or fragmented. Another source of confusion for IMFMapper was in genuine fractures, albeit in non-impact melt material. M1257328772R shows a case of this where a boulder has fragmented, and the resulting fracture down the middle has been detected as an IMF.

4.4. Mapping of IMFs in Crookes and Schomberger A Craters

Now that IMFMapper has been independently tested in Section 4.3, the next step was to apply it to the self-processed mosaics of Crookes and Schomberger A craters. As stated in Table 2, the tiling of these mosaics produced 3,332 and 1,341 image tiles for Crookes and Schomberger A, respectively. Inferring IMFMapper on these quantities of tiles took just 1.01 and 0.70 min for Crookes and Schomberger A, respectively—far surpassing what would be achievable through manual means.

The detections made by IMFMapper were converted from raster format (binary images where 1 = IMF and 0 = background) to ESRI shapefiles using GDAL, wherein detections made at the edges of adjacent tiles could be merged. Since the resolution of the tiles remained constant, the surface area (in m^2) of each merged IMF was also be recorded. This allowed for the filtering out of any detections ($<150 \text{ m}^2$ in area) which were too small to be resolvable and identifiable as IMFs. Similarly, adverse detections made along the seams of the mosaics were also removed since these were not indicative of IMFMapper's performance but of the quality of the data to which it has been applied. Furthermore, only the IMF detections which fell within the craters' floors and, in the case of Crookes, any melt ponds found within the terraces of its walls were considered for inclusion in the maps presented in Sections 4.4.1 and 4.4.2.

Figure 7 gives examples of some of the detections which were filtered out, either due to insufficient size, mosaic seams, or location, as well as some common false positives which remained in the maps. For example, despite being filtered out due to their locations on the crater wall, the majority of trails left by rockfalls within Crookes crater were misclassified as IMFs. Therefore, there is potential for IMFMapper to be transfer-learned to satellite imagery of rockfalls and compared to previous efforts in literature (Bickel et al., 2019, 2020). This would also be the first use of semantic segmentation for Lunar rockfalls, since Bickel et al. (2019) trained a RetinaNet object detection DCNN model.

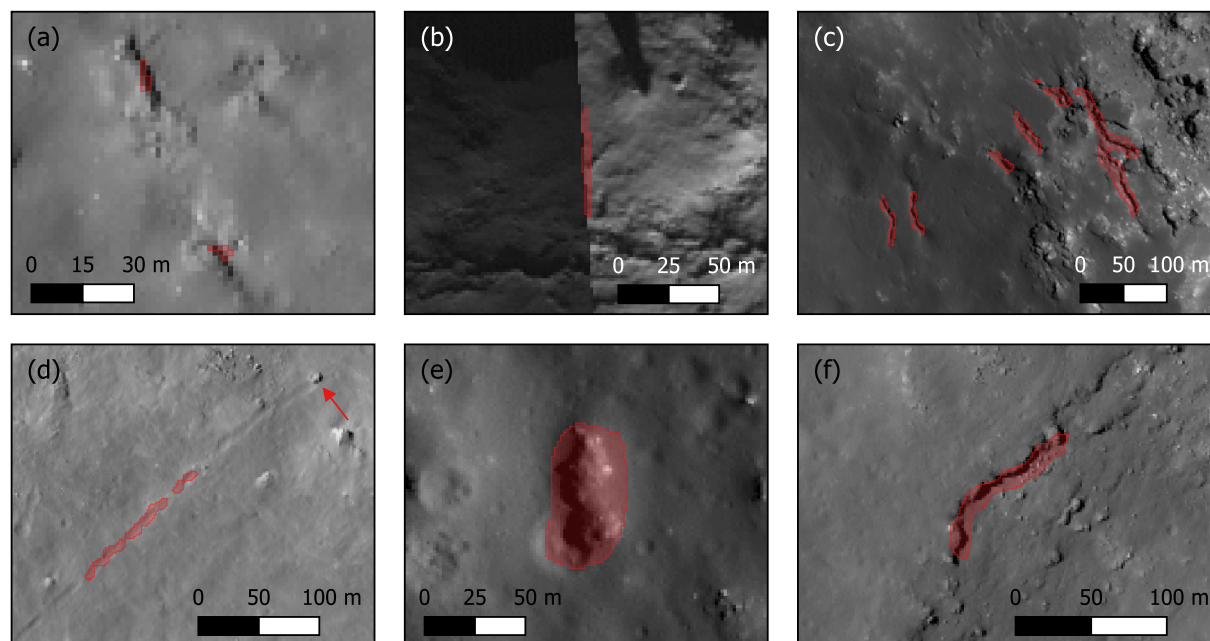


Figure 7. Examples of some of the detections made by IMFMapper within Crookes and Schomberger A craters before some were filtered out due to (a) insufficient area, (b) mosaic seams or (c) and (d) being located on the crater's wall or rim. (d) A rockfall trail left behind by the boulder marked with a red arrow. (e) and (f) Examples of positive elevation features, such as (e) boulders or (f) ridges, being misclassified as impact melt fractures due to their prominent shadows.

4.4.1. Crookes Crater

Figure 8 presents the automated mapping of IMFs within Crookes crater by IMFMapper. The center map of the figure plots the coordinates of the detected IMFs (in green) against the LROC WAC global mosaic with 50% transparency (Speyerer et al., 2011). Within this map, the crater's floor is outlined in red, which is also overlaid upon the mosaic of 56 LROC NAC images produced as part of this work.

The melt ponds identified by Thaker et al. (2020) in the northern and western walls of Crookes are also highlighted in blue. However, the presence of IMF detections within smooth, low albedo terrain suggests the discovery of additional melt ponds: two smaller melt ponds in the western wall, as well as another in the south. Maps (a) and (c) of Figure 8 show IMF detections being made within the melt ponds identified by Thaker et al. (2020) - with the IMFs in (a) being the same ones spotted in Figure 8 of Thaker et al. (2020). Whereas, map (b) exhibits parallel IMFs in the candidate melt pond situated approximately 2 km north of the western melt pond mapped by Thaker et al. (2020). Map (d) displays the largest of the IMFs detected in the proposed southern melt pond in Crookes crater, which is ≈ 600 m-long and between 15 and 40 m in width. In total, 70 IMFs were detected by IMFMapper within the melt ponds of Crookes crater, of which most were found within the western melt pond shown in (c).

Concerning the IMFs mapped within the crater floor, map (e) of Figure 8 contains the largest IMF found anywhere in Crookes due to it possessing a surface area, length and maximum width of more than 0.05 km^2 , 700 and 90 m, respectively. For reference, the detections in Crookes crater had an average area of $1,854 \text{ m}^2$, with a standard deviation of $4,015 \text{ m}^2$. Map (f) also presents the IMFs detected by IMFMapper within the region toward the top right of that shown in Figure 7 of Thaker et al. (2020), proving that it can detect complexes of sub-parallel IMFs.

4.4.2. Schomberger a Crater

Figure 9 is the first map of the IMFs found within Schomberger A crater, as completed automatically by IMFMapper. The top left map overlays the coordinates of the IMF detections (in green) within the visible portion of Schomberger A's floor (in red) upon the mosaic of 55 LROC NAC images created in this work (with 50% transparency). Maps (a) to (c) display the patterns of IMF detections at the "complex" scale, and maps (d) to (f) zoom in to show individual IMFs in greater detail.

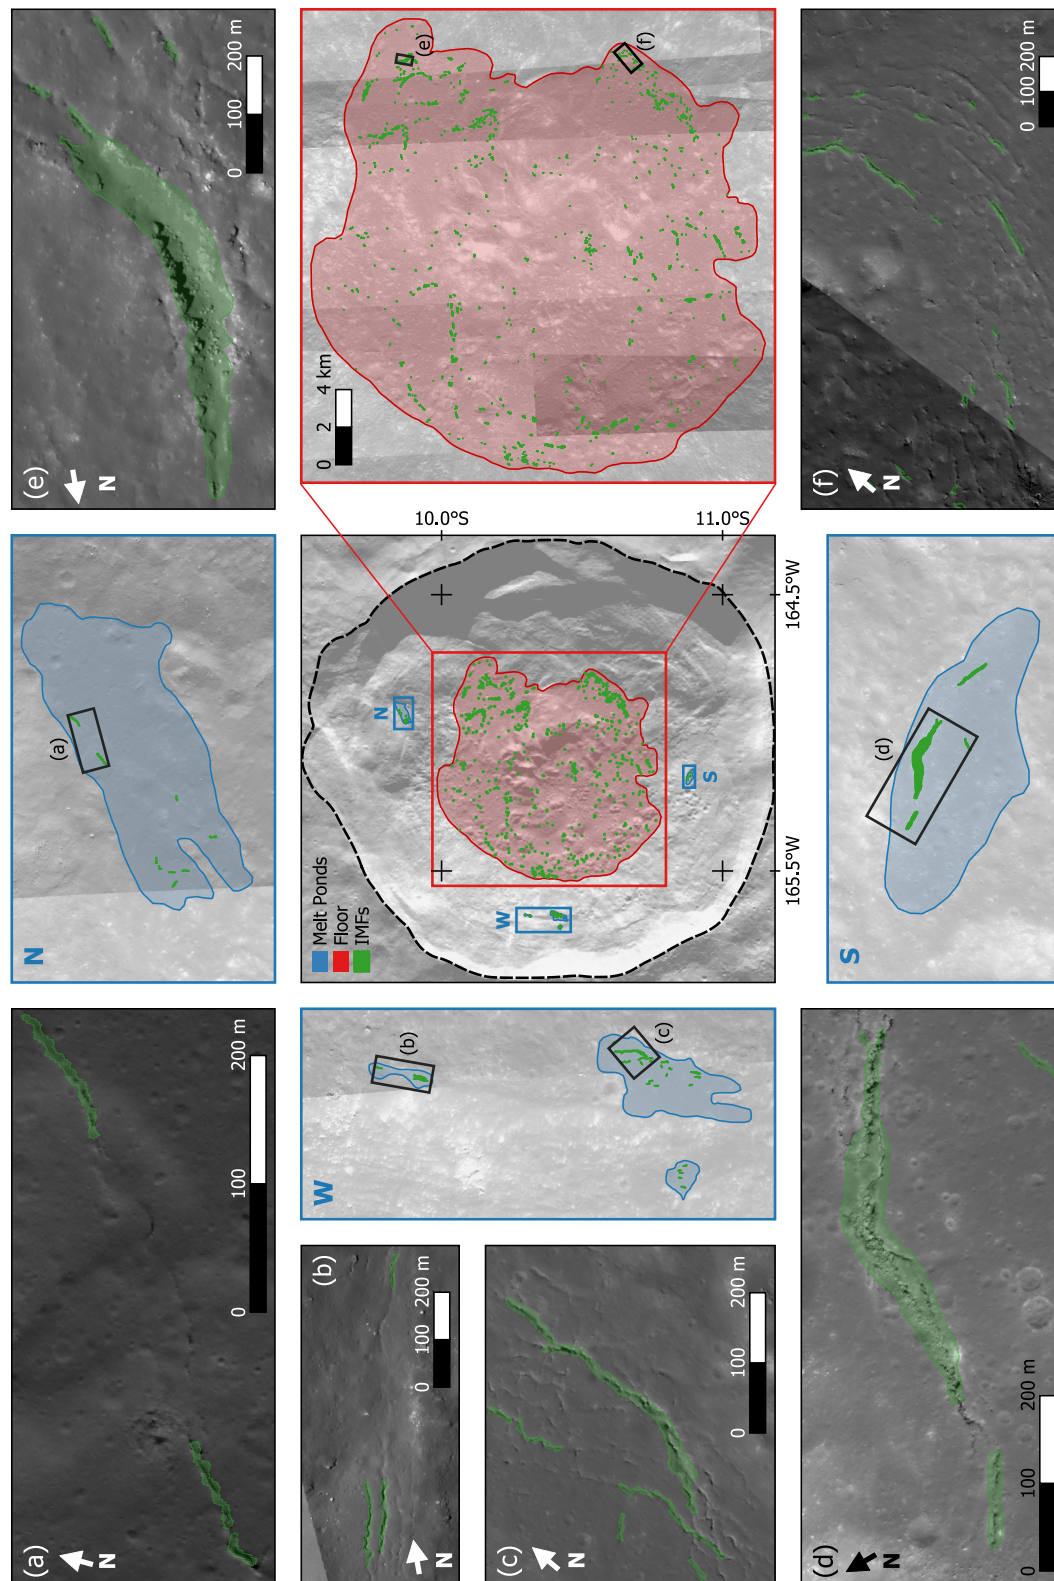


Figure 8. Detections of impact melt fractures (IMFs) made by IMFMapper within Crookes crater. (Center) Detections within the floor (red) and melt ponds (blue) of Crookes, overlaid upon the LROC WAC global mosaic (Speyerer et al., 2011). (a, c and f) Detections made within the same regions shown in Figures 8, 11 and 7 of Thaker et al. (2020), respectively. (b and d) Discovery of candidate melt ponds in the western and southern walls due to detected IMFs and smooth terrain. (e) The largest detection (by area) made by IMFMapper in Crookes.

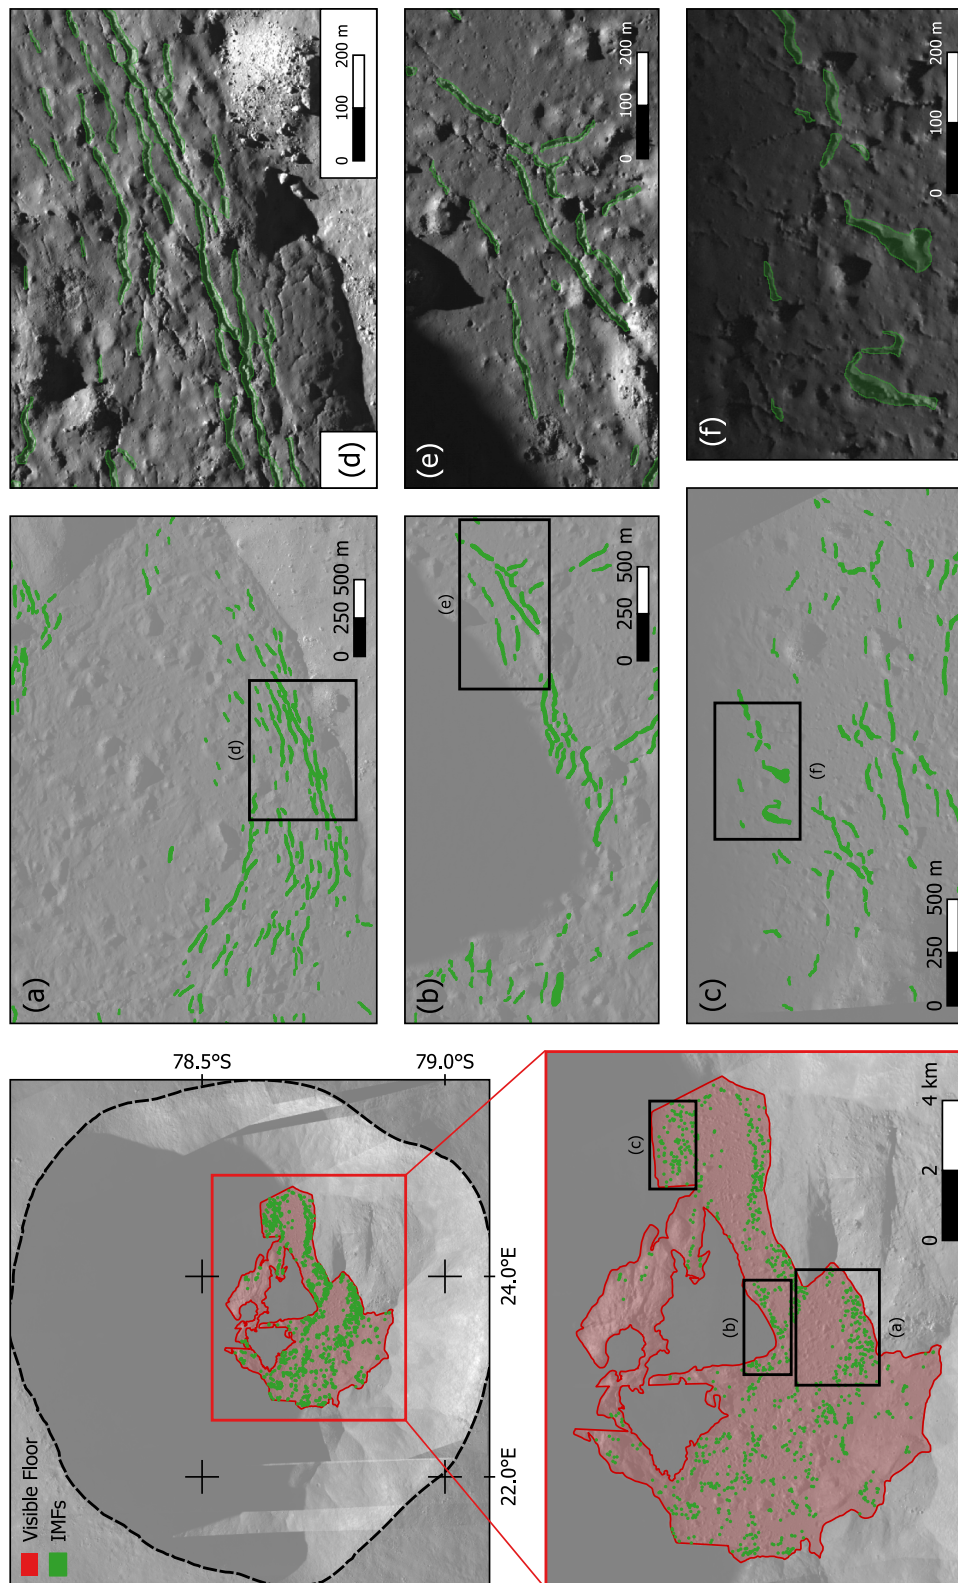


Figure 9. Detections of impact melt fractures made by IMFMapper within Schomberger A crater. (Top left) Detections within the visible floor (shown in red) of Schomberger A, overlaying the self-processed mosaic of 55 LROC NAC images. (a and d) Detections made within the same regions as those shown in Figure 2 of Lee (2022). (b and e) Some of the largest detections made by IMFMapper in terms of area, shown at the “complex” and pixel-level scales. (c and f) Detections made close to the terminus of the visible crater floor.




Figure 10. Results of applying IMFMapper to the LROC NAC mosaics of (a) Crookes and (b) Schomberger A craters after downscaling to 6 m/px—the resolution expected to be achieved by the SIMBIO-SYS instrument aboard BepiColombo (Benkhoff et al., 2021).

Maps (a) and (e), in particular, correspond to the same impact melt flow shown in Figure 2 of Lee (2022), showing how the spacing between the sub-parallel IMFs appears to get closer toward the margin between melt and wall. Maps (b) and (e) focus on some of the largest IMF detections (by area) with lengths between 350 and 400 m and widths of ≈ 15 m. As a comparison with those in Crookes crater, the IMFs detected in Schomberger A had an average area of 795 m², with a standard deviation of 969 m². This agrees with the observations in Xiao et al. (2014) whereby the size of an IMF scales with the depth of the hosting melt deposits, which can be expected to be deeper for a larger impact crater. Maps (c) and (f) show how IMFs are increasingly missed by IMFMapper due to the shadow cast by the crater's rim encroaching upon its floor.

4.5. Demonstration of IMFMapper at SIMBIO-SYS Resolution

In this section, we investigate the feasibility of utilizing IMFMapper in its current state (i.e., having been trained on 1.5 m/px LROC NAC imagery) to map IMFs on Mercury upon the arrival of BepiColombo. In the absence of satellite imagery of Mercury at such scales to experiment with, IMFMapper has been re-applied to the mosaics of Crookes and Schomberger A craters, but this time with them downscaled to 6 m/px in order to match the expected SIMBIO-SYS resolution. This means that a single 512 × 512 px tile at 6 m/px now covers approximately 16 times more area than it would have done at 1.5 m/px.

Figure 10 re-plots the maps of (a) Crookes and (b) Schomberger A craters, with the new detections made at this lower resolution overlaid as green points. This proves that, while experiencing a reduction in the quantities being recalled, IMFs were still detectable in both craters at 6 m/px. However, the number of IMF detections made decreased significantly more for Schomberger A than for Crookes crater. This is likely because, at this resolution and incidence angle, the IMFs are mostly covered in shadow and don't exhibit a clear opposite rim. Figure 10 focuses on two examples of IMFs still being detected at 6 m/px. In Crookes crater, the same IMF as shown in map (e) of Figure 8 has also been clearly delineated by IMFMapper without needing retraining. A segment of the complex of IMFs displayed in map (a) of Figure 9 is also shown as being detected in Figure 10.

5. Conclusion

IMFMapper is a DeepLabV3 DCNN with a ResNet50 backbone trained to perform the semantic segmentation of IMFs within LROC NAC imagery of Lunar impact melt deposits. Due to the quantities of IMFs available to label and the difficulty with creating pixel-level annotations for linear features, IMFMapper was trained in a weakly supervised manner (i.e., upon labels which were produced from a series of buffered line labels). During validation upon the Va₁ data set, it was found that IMFMapper achieved an average F₁ of 69.4% after just 47 out of the total 100 epochs of training. This was marginally less than the highest average F₁ reached by a DeepLabV3 model with a ResNet101 backbone and a smaller batch size (70.0%). However, during further validation of the three combinations of model backbone and batch size upon IMFs within Copernicus and Virtanen F craters, it was found that the configuration with a ResNet50 backbone and a batch size of 32, which we call IMFMapper, returned the highest average F₁ of 55.9%. Lastly, IMFMapper was then tested upon IMFs within Ohm crater, where it achieved an average F₁ of 69.3% and revealed that it can accurately delineate IMFs despite having been trained upon weak labels.

By applying IMFMapper to two self-processed mosaics of LROC NAC imagery, two maps have been created of the IMFs found within the Lunar impact craters: Crookes and Schomberger A. The former has been the subject of a previous manual mapping exercise (Thaker et al., 2020), whilst the latter is confirmed to host IMFs, which themselves have the potential to act as PSRs. Without the ability to quantitatively compare the two, the map of Crookes crater generated by IMFMapper visually agrees with the distributions of IMFs in Thaker et al. (2020). However, additional candidate melt ponds have been found in the western and southern walls of Crookes due to the smooth, low albedo terrain and the presence of IMFs detected by IMFMapper. In Schomberger A, we have produced the first mapping of IMFs in a crater so close to the Lunar South Pole. This indicates that the performance of IMFMapper is robust to the effects of highly oblique solar incidence angles, such as low contrast images and small elevation changes casting shadows. With the planned arrival of the Artemis missions, whose objective it is to sample an impact melt deposit (Weber et al., 2021), the ability to map IMFs at latitudes similar to Schomberger A and within scales (both spatially and temporally) unachievable by humans will be a useful tool - either to assess their potential as PSRs or to use this as a proxy to search for uncatalogued melt deposits.

IMFMapper is the first use of automated methods to detect IMFs on either the Moon or Mercury. As was the focus for Xiao et al. (2014), the mapping of IMFs within impact melt deposits can indicate which cooling mechanism is the most dominant in the periods after the impact event. Therefore, IMFMapper opens the door to scaling up this investigation from the six impact craters in Xiao et al. (2014) to as many as those having corresponding high-resolution imagery. The detection of IMFs in reachable areas of the Lunar surface would also be useful for any future mission which may want to investigate whether they could provide access to underground cavities similar to terrestrial analogs (Lemaire et al., 2024). Even a common false positive detection made by IMFMapper presents an interesting avenue for further work, in that Lunar rockfalls were mostly classified as IMFs. Therefore, the learned model weights within IMFMapper may prove advantageous for the transfer learning of a DeepLabV3 model trained to detect rockfalls.

Having been discussed in Section 1, IMFMapper was partly developed in anticipation of IMFs on Mercury being resolved at greater detail, thanks to the arrival of BepiColombo. Due to MESSENGER's elliptical orbit, resolving IMFs is only currently possible in Mercury's northern hemisphere, with a best resolution of 12 m/px (Hawkins et al., 2007; Xiao et al., 2014). Whereas, the SIMBIO-SYS instrument will provide 6 m/px imagery for roughly 20% of Mercury's surface within its first year of operation (Benkhoff et al., 2021). Despite the difference in the resolution between LROC NAC and what SIMBIO-SYS is expected to achieve, IMFMapper demonstrated that it can still detect some of the larger IMFs that are still clearly resolvable at 6 m/px. IMFMapper would likely still benefit from being trained from scratch upon LROC NAC imagery downsampled to match the SIMBIO-SYS resolution. Albeit, this will result in IMFs with widths of less than ≈ 30 m no longer being resolvable, and thus usable as training data. But with retraining upon resolutions closer to that of MESSENGER, there may be the possibility of including labeled data of IMFs on Mercury within the training data set. Applying IMFMapper, either in its current state or retrained, to SIMBIO-SYS high-resolution images would open up Mercury's southern hemisphere to the automated mapping of IMFs, as well as the detection of smaller, marginal IMFs within the northern hemisphere examples already mapped by Xiao et al. (2014).

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Data Availability Statement

The LROC NAC Experimental Data Record (EDR) image products used throughout this work can be acquired from NASA's Planetary Data System (Robinson, 2009). A publicly available Zenodo repository contains both IMFMapper's PyTorch model checkpoint and the IMFs detected within Crookes and Schomberger A craters as geo-referenced shapefiles (Le Corre, 2025b). Instructions on how to infer IMFMapper on new LROC NAC data, as well as a figure showing the effects of the various data augmentations, are given on GitHub (Le Corre, 2025a). BASH scripts for processing and map-projecting batches of raw LROC NAC image products into GIS-ready GeoTiffs are available on GitHub (Le Corre, 2025c).

Acknowledgments

This project is part of the Europlanet 2024 RI which has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 871149. Model training, testing and inference, as well as batch image processing, was enabled thanks to the Specialist and High Performance Computing systems provided by Information Services at the University of Kent.

References

- Barrett, A. M., Favaro, E. A., & Balme, M. R. (2023). Detecting and measuring transverse Aeolian ridges (TARs) on Mars using deep learning. In *54th lunar and planetary science conference* (Vol. 2806).1534
- Benkhoff, J., Murakami, G., Baumjohann, W., Besse, S., Bunce, E., Casale, M., et al. (2021). BepiColombo - Mission overview and science goals. *Space Science Reviews*, 217(8), 90. <https://doi.org/10.1007/s11214-021-00861-4>
- Bickel, V. T., Conway, S. J., Tesson, P.-A., Manconi, A., Loew, S., & Mall, U. (2020). Deep learning-driven detection and mapping of rockfalls on Mars. *Ieee Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 13, 2831–2841. <https://doi.org/10.1109/JSTARS.2020.2991588>
- Bickel, V. T., Lanaras, C., Manconi, A., Loew, S., & Mall, U. (2019). Automated detection of lunar rockfalls using a convolutional neural network. *IEEE Transactions on Geoscience and Remote Sensing*, 57(6), 3501–3511. <https://doi.org/10.1109/TGRS.2018.2885280>
- Brown, H. M., Boyd, A. K., Denevi, B. W., Henriksen, M. R., Manheim, M. R., Robinson, M. S., et al. (2022). Resource potential of lunar permanently shadowed regions. *Icarus*, 377, 114874. <https://doi.org/10.1016/j.icarus.2021.114874>
- Chaini, C., & Jha, V. K. (2024). A review on deep learning-based automated lunar crater detection. *Earth Science Informatics*, 17(5), 3863–3898. <https://doi.org/10.1007/s12145-024-01396-2>
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., & Yuille, A. L. (2016). DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *arXiv e-prints*. <https://doi.org/10.48550/arXiv.1606.00915>
- Chen, L.-C., Papandreou, G., Schroff, F., & Adam, H. (2017). Rethinking Atrous convolution for semantic image segmentation. *arXiv e-prints*. <https://doi.org/10.48550/arXiv.1706.05587>
- Christensen, P. R., Engle, E., Anwar, S., Dickenshied, S., Noss, D., Gorelick, N., & Weiss-Malik, M. (2009). JMARS - A planetary GIS. In *AGU fall meeting abstracts* (Vol. 2009, p. IN22A-06).
- Denevi, B. W., Koeber, S. D., Robinson, M. S., Garry, W. B., Hawke, B. R., Tran, T. N., et al. (2012). Physical constraints on impact melt properties from Lunar Reconnaissance Orbiter Camera images. *Icarus*, 219(2), 665–675. <https://doi.org/10.1016/j.icarus.2012.03.020>
- Dhingra, D., Head, J. W., & Pieters, C. M. (2017). Geological mapping of impact melt deposits at lunar complex craters Jackson and Tycho: Morphologic and topographic diversity and relation to the cratering process. *Icarus*, 283, 268–281. <https://doi.org/10.1016/j.icarus.2016.05.004>
- Gagnon, S., Lemelin, M., Neish, C., & Diotte, F. (2023). Localisation and Study of impact melts in the south polar Region of the Moon. In *54th lunar and planetary science conference* (Vol. 2806, p. 2655).
- GDAL/OGC contributors. (2025). GDAL/OGC Geospatial Data Abstraction software Library. <https://doi.org/10.5281/zenodo.5884351>
- Hawkins, S. E., Boldt, J. D., Darlington, E. H., Espiritu, R., Gold, R. E., Gotwols, B., et al. (2007). The Mercury dual imaging System on the MESSENGER spacecraft. *Space Science Reviews*, 131(1–4), 247–338. <https://doi.org/10.1007/s11214-007-9266-3>
- He, K., Zhang, X., Ren, S., & Sun, J. (2014). Spatial Pyramid pooling in deep convolutional networks for visual recognition. *arXiv e-prints*. <https://doi.org/10.48550/arXiv.1406.4729>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *2016 IEEE conference on computer vision and pattern recognition (CVPR)* (p. 1). <https://doi.org/10.1109/CVPR.2016.90>
- IAU. (2025). Gazetteer of planetary nomenclature. In *International Astronomical Union (IAU) Working Group for planetary System nomenclature*. Retrieved from <http://planetarynames.wr.usgs.gov/>
- Kingma, D. P., & Ba, J. (2017). Adam: A method for stochastic optimization. *arXiv e-prints*. <https://doi.org/10.48550/arXiv.1412.6980>
- Laura, J., Acosta, A., Addair, T., Adoram-Kershner, L., Alexander, J., Alexandrov, O., et al. (2023). *Integrated software for imagers and spectrometers*. Zenodo. <https://doi.org/10.5281/zenodo.7644616>
- Le Corre, D. (2025a). *IMFMapper GitHub repository*. Zenodo. <https://doi.org/10.5281/zenodo.17460144>
- Le Corre, D. (2025b). IMFMapper PyTorch model checkpoint and shapefiles of IMFs in crookes and schomberger A craters. *Zenodo*. <https://doi.org/10.5281/zenodo.15101175>
- Le Corre, D. (2025c). *Planetary image processing GitHub repository*. Zenodo. <https://doi.org/10.5281/zenodo.17215911>
- Le Corre, D., Mason, N., Bernard-Salas, J., Mary, D., & Cox, N. (2025). New candidate cave entrances on the Moon found using deep learning. *Icarus*, 441, 116675. <https://doi.org/10.1016/j.icarus.2025.116675>
- Lee, P. (2022). Impact melt lava inflation fractures in Schomberger-A crater, Moon: New permanently shadowed regions and possible Caves near the lunar south pole. In *53rd lunar and planetary science conference* (Vol. 2678, p. 2685).
- Lemaire, T., Morgavi, D., Petrosino, P., Calvari, S., Repola, L., Esposito, L., et al. (2024). Lava flow field development and lava tube formation during the 1858–1861 eruption of Vesuvius (Italy), unravelled by historical documentation, lidar data and 3D mapping. *Journal of Volcanology and Geothermal Research*, 455, 108197. <https://doi.org/10.1016/j.jvolgeores.2024.108197>
- Lin, T.-Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., et al. (2014). Microsoft COCO: Common objects in context. *arXiv e-prints*, 740–755. https://doi.org/10.1007/978-3-319-10602-1_48
- Lindholm, A., Wahlström, N., Lindsten, F., & Schön, T. B. (2022). *Machine learning - A first course for engineers and scientists*. Cambridge University Press.
- Pizer, S. M., Amburn, E. P., Austin, D. J., Robert, C., Geselowitz, A., Greer, T., et al. (1987). Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*, 39(3), 355–368. [https://doi.org/10.1016/S0734-189X\(87\)80186-X](https://doi.org/10.1016/S0734-189X(87)80186-X)
- Plescia, J. B., & Cintala, M. J. (2012). Impact melt in small lunar highland craters. *Journal of Geophysical Research (Planets)*, 117(E12), E00H12. <https://doi.org/10.1029/2011JE003941>
- QGIS Development Team. (2025). *QGIS geographic information System*. QGIS Association. Retrieved from <https://www.qgis.org>
- Robinson, M. S. (2009). *LRO-L-LROC-2-EDR-V1.0*. NASA Planetary Data System (PDS). <https://doi.org/10.17189/1520250>
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. *arXiv e-prints*. <https://doi.org/10.48550/arXiv.1505.04597>
- Solomon, S. C., McNutt, R. L., Gold, R. E., Acuña, M. H., Baker, D. N., Boynton, W. V., et al. (2001). The MESSENGER mission to Mercury: Scientific objectives and implementation. *Planetary and Space Science*, 49(14–15), 1445–1465. [https://doi.org/10.1016/S0032-0633\(01\)00085-X](https://doi.org/10.1016/S0032-0633(01)00085-X)
- Speyerer, E. J., Robinson, M. S., & Denevi, B. W., & LROC Science Team. (2011). Lunar reconnaissance Orbiter camera global morphological map of the Moon. In *42nd annual lunar and planetary science conference* (p. 2387).
- Srivastava, N., & Varatharajan, I. (2016). Geomorphology of Lowell crater region on the Moon. *Icarus*, 266, 44–56. <https://doi.org/10.1016/j.icarus.2015.11.013>
- Thaker, A. D., Patel, S. M., & Solanki, P. M. (2020). Morphological analysis and mapping of complex craters of Copernican age: Crookes, Ohm and Stevinus. *Planetary and Space Science*, 184, 104856. <https://doi.org/10.1016/j.pss.2020.104856>

- Wagner, R. V., Henriksen, M. R., Manheim, M. R., Meyer, H. M., Banks, M. E., Malaret, E., et al. (2024). Where is that crater? Best practices for obtaining accurate coordinates from LROC NAC data. *Planetary Science Journal*, 5(7), 157. <https://doi.org/10.3847/PSJ/ad54c6>
- Wagner, R. V., & Robinson, M. S. (2021). Occurrence and origin of lunar pits: Observations from a new catalog. In *52nd lunar and planetary science conference* (p. 2530).
- Wang, Y., Di, K., Xin, X., & Wan, W. (2017). Automatic detection of Martian dark slope streaks by machine learning using HiRISE images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 129, 12–20. <https://doi.org/10.1016/j.isprsjprs.2017.04.014>
- Weber, R. C., Lawrence, S. J., Cohen, B. A., Bleacher, J. E., Boyce, J. W., Collier, M. R., et al. (2021). The Artemis III science definition team report. In *52nd lunar and planetary science conference* (p. 1261).
- Xiao, Z., Zeng, Z., Li, Z., Blair, D. M., & Xiao, L. (2014). Cooling fractures in impact melt deposits on the Moon and Mercury: Implications for cooling solely by thermal radiation. *Journal of Geophysical Research (Planets)*, 119(7), 1496–1515. <https://doi.org/10.1002/2013JE004560>
- Zhang, S., Liu, J., Michael, G., Zhu, K., Lei, D., Zhang, J., et al. (2024). Detecting lunar linear structures based on multimodal semantic segmentation: The case of sinuous rilles. *Remote Sensing*, 16(9), 1602. <https://doi.org/10.3390/rs16091602>