IEEE Robotics & Automation Magazine (RAM) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RAM paper.

Super-Resolution of Lunar Satellite Images for Enhanced Robotic Traverse Planning

J.I. Delgado-Centeno[1], P. Harder, V. Bickel, B. Moseley, F. Kalaitzis, S. Ganju, M.A. Olivares-Mendez

***Abstract*—Lunar exploration missions require detailed and accurate planning to ensure their safety. Remote sensing data, such as optical satellite imagery acquired by lunar orbiters, is key for the identification of future landing and mission sites. Here, robot- and astronaut-scale obstacles are the most relevant to resolve, however, the spatial resolution of the available image data is often insufficient - particularly in the poorly illuminated polar regions of the Moon -, leading to uncertainty. This work shows how a novel single-image Super-Resolution (SR) application - ANUBIS, Adversarial Network for Uncertainty Based Image Super-resolution - can enhance lunar surface imagery by improving their resolution by a factor of 2, outperforming other approaches and benchmarks. The enhanced images improve the reliability and detail of lunar traverse planning and topographic reconstruction, while providing an estimate of the uncertainty associated with the enhancement process, vital to ensure mission planning integrity. This work demonstrates how machine learning-driven processing can enhance existing data products to maximize their value for science and exploration of the Moon and other celestial bodies.**

***Index Terms*—Image Super-resolution, Space robotics, Remote sensing, Traverse planning.**

I. Lunar missions and remote sensing

Nowadays, the Moon is in the spotlight of numerous companies and space agencies thanks to the benefits to humanity once these missions begin in earnest. In fact, there are several new planned missions for the next decade with many different objectives, including the building of a permanent lunar human settlement [1], or lunar polar exploration for water prospection [2] [3]. The historic success of robots in extraterrestrial bodies reveals that they play a crucial role as an enabling technology to build the lunar economy. Robots can deal with the harsh and vast lunar environment better than humans, allowing us to explore the Moon's surface without facing any of the inherent risks of being present at an extraterrestrial location. Therefore, most of the mentioned missions will involve robotic operations.

Autonomous robotic missions rely on the knowledge collected prior to the mission about the environment where it will be performed. Thus, satellites are sent to other celestial bodies to remotely study and analyze many aspects of the different locations where the robot will perform its tasks. A good understanding of these locations is vital to ensure the success of the mission and the achievement of its end goal. For the Moon, the biggest publicly available remote sensing dataset is provided by NASA's Lunar Reconnaissance Orbiter(LRO) satellite mission [**?**], which for the past 13 years has been collecting data with a variety of sensors, including cameras, an altimeter and radiometers. Such remote sensing data allows to crucially reduce the potential risks that can cause harm to the mission at the same time that is used to set realistic goals and targets. However the available data, for some lunar locations particularly, can be insufficient due to its inherent environmental conditions. In the case of the lunar poles for example, which are one of the main upcoming mission targets, the permanent low lighting conditions won't allow to obtain optimal image resolution of these regions. LRO images taken in these locations can only achieve a 1 meter per pixel resolution, even though the best (nominal) achievable resolution with the Narrow Angle Camera on board of the satellite is 0.5 meters per pixel. Furthermore, the scarcer resolution on satellite images may lead to inaccurate planning and therefore, an increment of the potential hazards of the lunar mission.

For the image resolution problem, there are mainly two alternatives. The first option is to design, build and launch a new satellite with better instrumentation on board. The newly acquired data would present better resolution, which would produce as a result an increment in the mission planning's precision. However, this solution would require enormous time and financial efforts, which makes it challenging. Also, a new mission would need years until a remote sensing dataset as rich as NASA's LRO for example can be built. It is worth mentioning that in-orbit satellites could also acquire more precise data by getting closer to the lunar surface at the cost of spending fuel and reducing drastically the mission duration, which makes it non-viable.

Alternatively, direct image enhancement using image Super Resolution (SR) presents a second, much more viable option that can be directly applied to the huge archives of data already collected by different remote sensing missions. These computer vision methods take an image and increase its resolution several times, producing a higher quality and more detailed version. As an image processing operation, SR can be directly applied to existing image datasets, potentially improving the resolution and detail quality of entire existing collections, with just computational resources as a requirement. Remote sensing has been paired in research with image SR for many years. These either fall into Single Image SR (SISR) [4], [5], or Multi Frame SR (MFSR) [6] techniques depending on the available information and the quality of the curated satellite image dataset. In this collection of recent works, deep learning architectures are state-of-the-art when it comes to image enhancement.

Regarding planetary missions, SR is a valuable asset in image enhancement due to the limitations in data transmission

[1]Correspondence author

IEEE Robotics & Automation Magazine (RAM) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RAM paper.

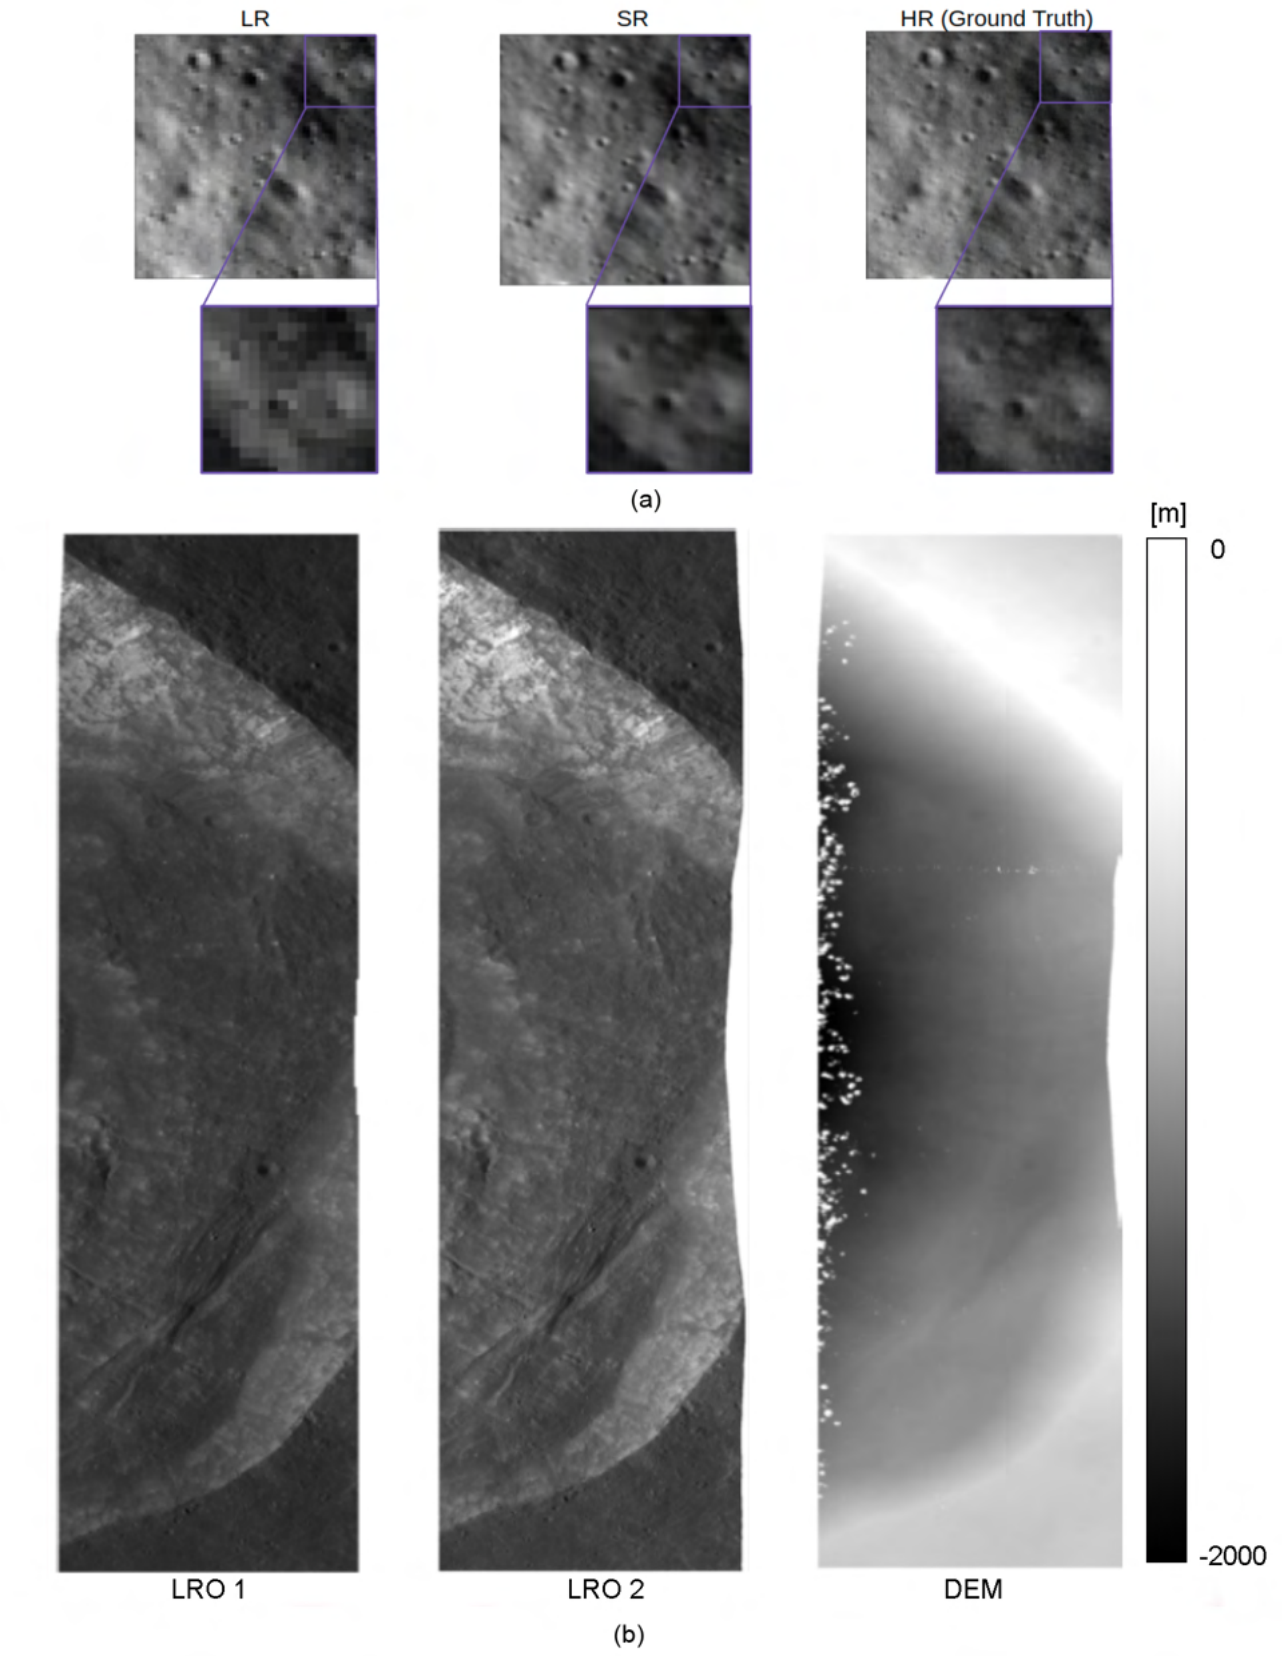

Fig. 1. (a) Direct comparison between lunar images with different resolutions: LR is Low resolution at 1m/px, SR is the product of our image enhancement Super-resolution process at 0.5m/px and HR the ground truth image also at 0.5 m/px. (b) Two map projected LRO images of the Kepler crater (latitude: 8.09, longitude:322) from different satellite orbits (left, center) and their corresponding lunar elevation map (right). The noisy pixels that can be found in the maps are particular locations whose altitude was not possible to calculate.

IEEE Robotics & Automation Magazine (RAM) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RAM paper.

as well. In [7], the authors propose a method that combines SR and inference suppression to improve the resolution of sub-surface data acquired with the Shallow Radar on board NASA's Mars Orbiter Mission (MRO). Furthermore, in [8] is shown how the characterisation of dynamic surface changes on the Martian surface can be enhanced by applying SR techniques to the image data provided by MRO's instrument High-Resolution Imaging Experiment (HiRISE). The closest method to the work presented in this paper is Lunar HighRes-net [9], which provides the first approach of lunar image enhancement with Deep Learning. The HighRes-net MFSR architecture [6] is utilised to increase the resolution of lunar surface images provided by NASA's LRO NAC instrument in it. However, this approach uses synthetically downsampled and shifted imagery. Therefore, it is not possible to use it directly to improve the LRO dataset, as every gathered image by LRO has different lighting conditions, making an MFSR approach not directly applicable. Additionally, in [9] the inherent uncertainty of the output features is not estimated. This fact presents a limiting factor in the applicability of the proposed method, as its usage in critical applications such as lunar robotic missions could introduce unknown errors in the process and therefore compromise the safety of the planned robotic task.

This work introduces a Single Image Super-resolution (SISR) application for enhancing lunar surface images gathered by NASA's Lunar Reconnaissance Orbiter satellite mission. The main goal of this application is to showcase the applicability of these techniques to real lunar data and their viability in improving lunar robotics traverse planning. Our method takes a low resolution lunar surface image from this dataset and super-resolves it, providing a $2\times$ resolution enhancement using a Machine Learning (ML) based technique. Due to the ill-posedness of SISR, we include an uncertainty estimation in our output. Importantly, downstream tasks related to space robotics missions were employed to validate the performance of this application and to showcase its potential reliability. Specifically, we investigate this work's performance on obstacle detection and path planning tasks, emulating lunar mission planning. Lastly, this method was used to produce Digital Elevation Maps (DEM) of the Moon with enhanced resolution, which is the most versatile and exploited data product in robotic traverse planning for extraterrestrial locations. The main contributions of this work can be summarized as:

1) A novel application of generative model based SISR for real lunar surface images, which achieves a $2\times$ resolution enhancement.
2) Uncertainty estimation along with the enhanced images by using a deep ensemble of networks to increase the reliability of the method.
3) Performance validation of the SR method using real downstream tasks related to lunar robotic missions, allowing to reliably estimate its performance and ensure safety by mean of the uncertainty estimation.
4) A proof of concept for the utilisation of image SR to enhance DEMs of the lunar surface.

This application is called ANUBIS - **A**dversarial **N**etwork for **U**ncertainty **B**ased **I**mage **S**uper-resolution.

II. Methodology and data

A. Lunar surface image dataset

NASA's Lunar Reconnaissance Orbiter [10] is a satellite launched in 2009 with the objective of collecting as much data as possible from the lunar surface. Since its launch, the satellite has gathered more than 2 million images using both its Narrow Angle Camera (NAC) and Wide Angle Camera (WAC) [11]. Mainly, the only surface areas with a lack of representation by the satellite observations are the permanently shadowed regions inside polar craters. Due to the inherent lighting conditions, sunlight never reaches the inside of different polar craters, causing the lack of visual information and data in the LRO dataset.

The best available image resolution of LRO's dataset is provided by the NAC. This instrument consists of two nominally identical cameras that capture 12-bit panchromatic optical images. Each camera consists of a one-dimensional 5064-pixel CCD oriented perpendicularly to the direction of flight, providing a 10 μrad field of view. Two-dimensional images are generated along the satellite's orbit by capturing multiple image lines as the spacecraft moves (push-broom scanning). This means that the spatial resolution of the images depends on both the spacecraft altitude and the exposure time. Both factors can vary; typically, the spatial resolution ranges between $0.5 - 2.0$ meters per pixel in both the in-line and cross-line image directions. Additionally, this instrument has two operational modes: regular and summed mode [12]. When the camera capture images in regular mode, the full 5064-pixel array is used to obtain the best possible image resolution. However, if this regular mode would be used in low light conditions, the photon count for every pixel of the sensor would be low and therefore, the collected images rather dark. As an alternative, the summed mode adds the adjacent pixel intensity values to produce better quality images, where the details of the lunar surface can be seen more clearly. This operational mode comes at the cost of having half of the maximum achievable resolution with the sensor, which is the main reason of why for example,in the darker polar regions the maximum image resolution is 1 meter per pixel for the lunar surface frames.

B. Data augmentation

This work aims to learn to super-resolve the summed mode LRO lunar images back to their original regular mode resolution. To do so, a dataset of Low Resolution (LR, 1m/px) - High Resolution (HR, 0.5 m/px) image pairs was created from a collection of LRO NAC images in their Experimental Data Record (EDR) form, i.e. raw photon count images without any post-processing (calibration) applied. Deep Neural Network (DNN) based remote sensing SR requires both LR and HR data for the training process. In most cases, sufficient quantities of these matching image pairs are not available due to the complexity of collecting matching LR - HR pairs that cover the same location with similar illumination conditions. Therefore, synthetic downsampling modes are commonly used

as an alternative to create the necessary datasets for this matter. The HR images in our dataset are authentic regular mode images, whilst the LR images are their corresponding synthetically down-sampled versions obtained using the real LRO summed mode operation. Many upcoming lunar missions will be focused on the south pole, where the summed mode images are predominant. Therefore, the dataset was curated to represent the real summed mode images taken over the lunar south pole, with the goal of developing a tool to improve the quality of the polar region's data and enhance the information for the missions. Ground truth regular mode HR images were selected for this task from the more equatorial lunar Highland regions, as the South pole is considered a Highland region as well. Images were also selected such that their lighting conditions matched the expected lighting conditions of summed mode images at the South pole. Specifically, only images that had solar incidence angles between 65 and 90 degrees were finally chosen. In total, 7000 EDR regular mode images with 0.5 meters per pixel resolution were selected from the LRO database. Within each selected image, several 32×32 image patches from random locations were extracted. To avoid the presence of completely dark frames (shadowed regions), image patches were discarded if their average intensity was below a certain threshold. This resulted in a total of 220,000 image patch pairs. The summed mode downsampling process was performed on every patch following the guidelines from NASA, replicating every step and operation performed at pixel-level and producing 16x16 pixels (1 meter per pixel resolution) matching LR patches. 10,000 LR-HR patches pairs were randomly sub-selected for the test and validation datasets, and the remaining 200,000 patches formed the training dataset. To perform a qualitative evaluation of the implemented synthetic downsampling operator, the LRO dataset was scanned to find a regular-summed image pair that covered the exact same location and had similar lighting conditions at the capture moment. A synthetically down-sampled regular mode image was compared to a real summed mode image (Fig. 2) at a location where both images happened to overlap, showcasing the accuracy of the downsampling operation used in this work. It is worth noting that these types of real LR-HR pairs were not used in our training data as they are notoriously rare, and accurately aligning them would be a challenging task.

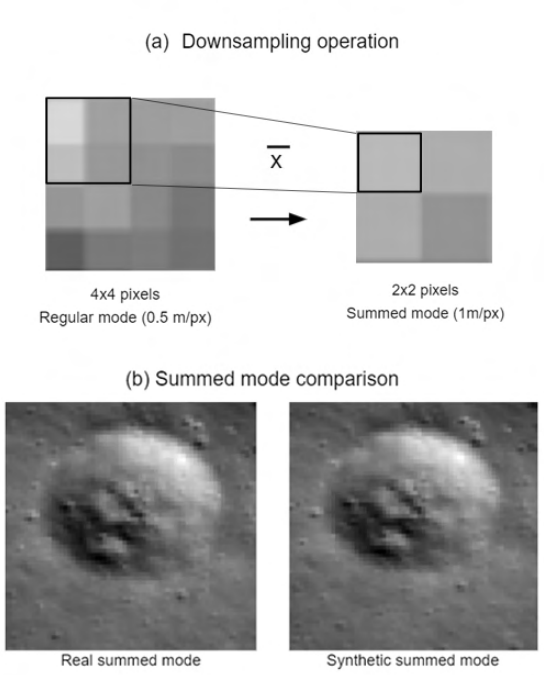

Fig. 2. Summed mode operation. (a) Mean of near neighbour pixels diagram. (b) Perceptual validation of our downsampling approximation. A real Regular-summed mode image pair that cover the same lunar surface area were used to validate our approach.

C. Lunar images super-resolution

This work introduces ANUBIS (Adversarial Network for Uncertainty Based Image Super-resolution), see Fig. 3. ANUBIS is a generative model (with a Generative Adversarial Network, GAN architecture) that maps 1/m pixel lunar surface summed mode images to a $2\times$ higher resolution regular mode equivalent. This provides enhanced quality data of sections of the Moon where the resolution was limited due to environmental conditions. Additionally, the application gives an uncertainty estimation at the same grid as the SR output to provide information on the reliability of the super-resolved image. The uncertainty is calculated by using an ensemble [13] of networks whose outputs can be compared to understand which parts of the super-resolved images are consistently enhanced similarly and therefore, reliable for any application where the images are used.

The ANUBIS architecture workflow is shown in Fig. 3. The first element of the architecture, the generator, is used for upscaling the 16×16 pixels LR image patch input. This network is formed by 8 residual blocks, 4 convolutional (3 followed by ReLU activation layers), and 1 transposed-convolution layer. Each one of the convolutional layers including the residual blocks contain 64 channels with a stride and padding of 1. The kernel size for the residual layers is 3 while for the rest of the convolutional layers is 1. The ablation study performed on the network led us to remove the batch normalization layers that were present in some GAN of the literature for better performance. The second element of ANUBIS, the discriminator, is used to improve the quality of the SR output of the generator. It uses as input the SR image and tries to discern if it is an actual real lunar surface image or a super-resolved one. The discriminator is composed by 5 convolutional layers followed by ReLU activations and a final pooling followed by sigmoid activations. The architecture generates a single (deterministic) super-resolved image for every LR image input. The network assumes that the missing information needed to infill and enhance the image resolution is implicit in the input pixels, even though the SR is an ill-posed problem. Instead of generating a single SR output from one trained generator, ANUBIS generate one SR from each member of the ensemble, following other recent works [15]. To do so, 24 instances of the GAN model were trained with different random seeds and initial weights, creating a deep

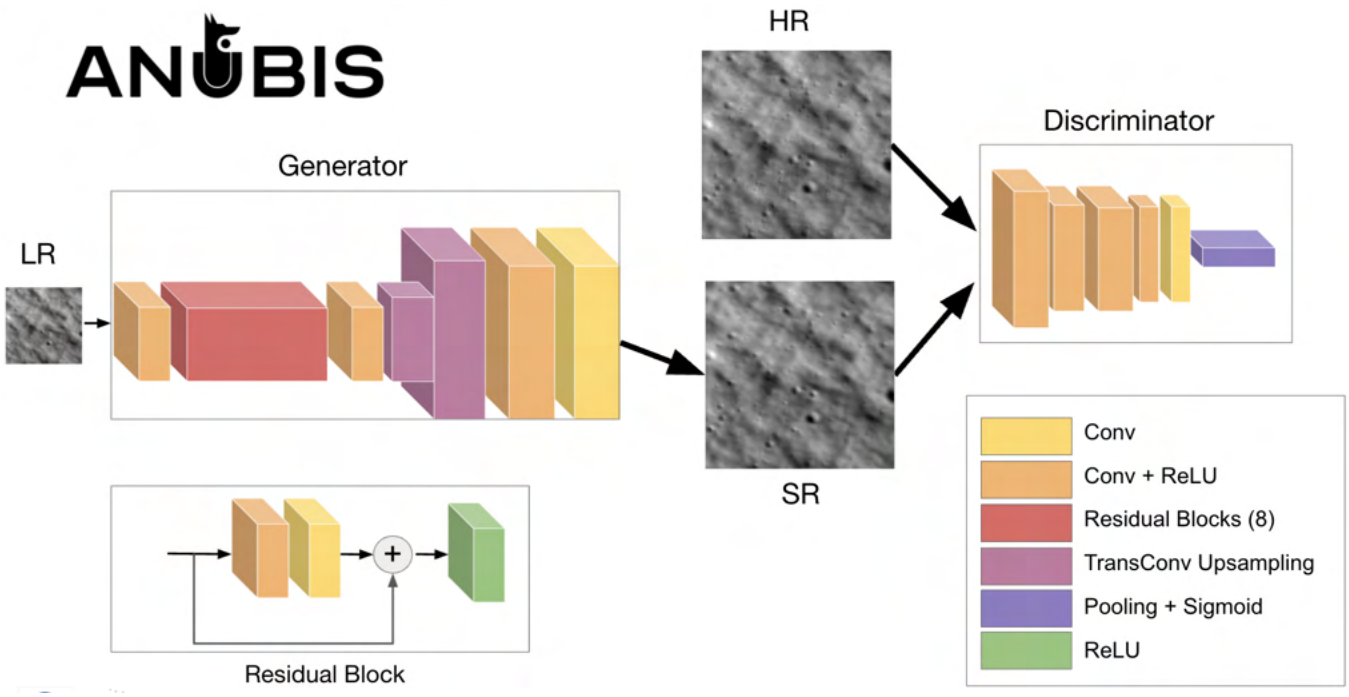

Fig. 3. The ANUBIS architecture consists of a ResNet-based generator and a CNN-based discriminator. Colors indicate the block types.

ensemble of networks that produces various super-resolution images from a single LR image. The empirical distribution of SR samples approximates the true posterior distribution of HR images, conditioned on the entire dataset of LR/HR pairs. Finally, the standard deviation is computed across the SR samples, generating a pixel-wise uncertainty map, see Fig. 4. For every LR image, there are several corresponding HR versions of it due to the missing information in the original one. It is important to explore this distribution to ensure the accuracy of the proposed super-resolved image and its validity for real world applications where safety is critical.

ANUBIS is trained on the LR-HR image pairs previously described in this section. Both architectures were trained using the Adam stochastic gradient descent algorithm for 200 epochs. An MSE loss is used for the generator and a Binary Cross-Entropy (BCE) loss function is used for the discriminator. We use a batch size of 256 image patches and a learning rate of $1 \cdot 10^{-4}$ for the GAN. In the adversarial loss, the importance ratio of the generator vs the discriminator loss is 1000:1. All experiments were run on an NVIDIA A100 GPU with an early stopping mechanism with a patience of 20 epochs.

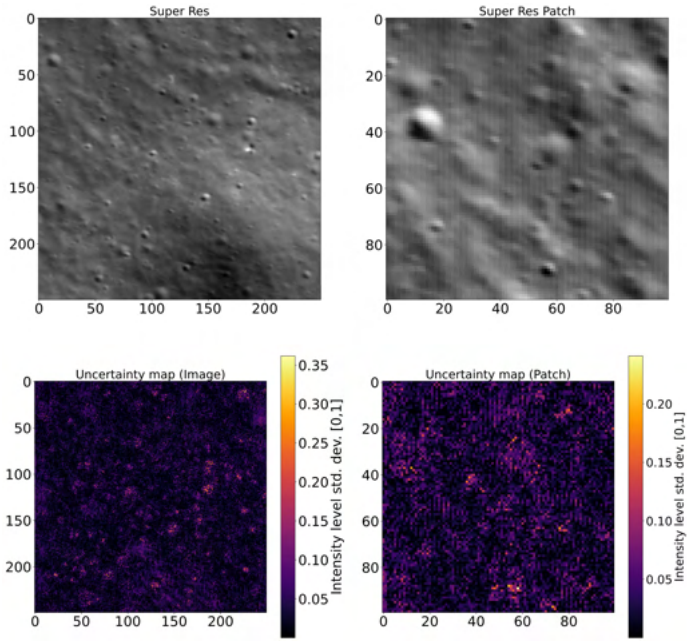

Fig. 4. Estimated uncertainty over the super-resolved image and its upper-left patch. For every pixel (of intensity range [0,1]), we compute the variance over the ensemble set of super-resolved images. The result is a heat map that is associated with the uncertainty of the posterior distribution about the unknown HR.

III. Method validation

A. Metrics

For the quantitative evaluation of ANUBIS, two different metrics have been used: Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM). These two metrics are the most commonly used evaluation methods in SR state-of-the-art works for image enhancement comparison. The results of the application introduced in this paper are compared with different baselines and evaluated directly using the mentioned metrics in this section. By obtaining good image reconstruction, the pixel variance along ANUBIS' output image distribution can be studied to see which pixels are consistent throughout the collection as it was stated in the previous section.

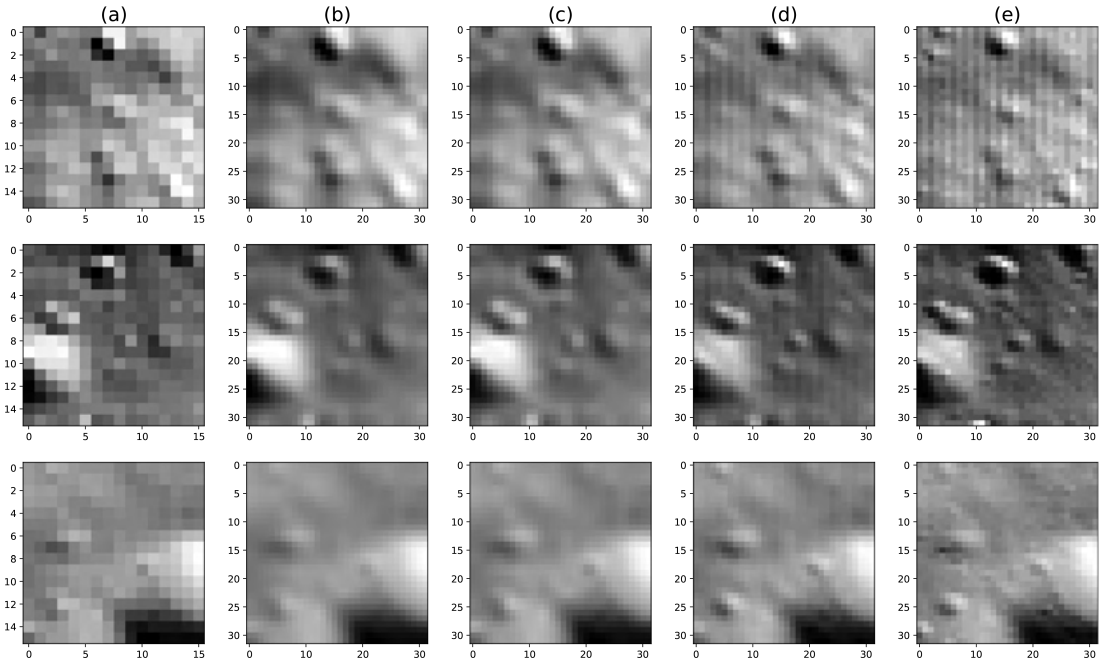

Fig. 5. Three examples of the lunar surface images super resolution comparison: (a) Low Resolution, (b) Bilinear, (c) Bicubic, (d) **ANUBIS (our work)** (e) Ground truth. The selected images were randomly taken out of the test set of images.

1) Peak Signal-to-Noise Ratio: Peak Signal-to-Noise Ratio, one of the most used metrics in image SR, estimates the quality of the reconstruction performed by a DNN when enhancing an LR image. It is directly derived by the Mean Squared Error (MSE). Given a monochromatic high-resolution image I_H of size U, V and its super-resolution counterpart I_S, the PSNR is defined as:

$$MSE = \frac{1}{UV} \sum_{i=0}^{U-1} \sum_{j=0}^{V-1} (I_H(i,j) - I_S(i,j))^2 \quad (1)$$

$$PSNR = 20 \cdot log_{10}(MAX_{I_H}) - 10 \cdot log_{10}(MSE) \quad (2)$$

with MAX being the maximum possible pixel value of the image.

2) Structural Similarity Index Measure: The Structural Similarity Index Measure metric evaluates the degradation as a variation in the structural information between I_S and I_H. It complements the use of PSNR as a metric, instead of being based on the absolute error, it focuses on the structures of the images represented by the correlation between spatially close pixels. For images I_S and I_H of size U, V, the SSIM can be defined as:

$$SSIM(I_H, I_S) = \frac{(2\mu_{I_S}\mu_{I_H})(2\sigma_{I_S I_H} + c_2)}{(\mu_{I_S^2} + \mu_{I_H^2} + c_1)(\sigma_{I_S^2} + \sigma_{I_H^2} + c_2)} \quad (3)$$

where: μ is the mean, σ^2 is the variance, $c_1 = (K_1 L)^2$, $c_2 = (K_2 L)^2$ are variables for stabilization of the denominator, L being the dynamic range of the pixel values and lastly $K_1 = 0.01$, $K_2 = 0.03$

B. Results

For the evaluation of the enhancement process, we have chosen four different baselines. The first two are traditional computer vision techniques for image upscaling, the bilinear and bicubic interpolation. The other two selected upscaling methods are the Fast Super-resolution Convolutional Neural Network (FSRCNN) [17] and the Enhanced Deep Residual Networks for single-image super-resolution (EDSR) [18]. These DL-based techniques are two of the best performing networks metric wise according to one of the latest survey paper [19] on the topic. The result images obtained after performing the upscaling operation for all 10,000 LR image patches from the test dataset with all five methods were compared against their corresponding ground truth, evaluating the quality of the super-resolved images with the metrics chosen. Table I displays the obtained average metric values across the total amount of image LR-HR pairs from the test set. As ANUBIS generates a super-resolved image distribution, the PSNR and SSIM values are calculated as an average across this distribution. ANUBIS outperforms the baseline methods in both of the selected metrics, presenting the best results in enabling the summed mode LRO image upscaling needed for planetary science applications. In terms of PSNR, the reconstruction and upscaling provided by ANUBIS display marginally better results than the baselines. Simultaneously, it exhibits superior performance with respect to the SSIM metric, which is more fitting for assessing the realism and fidelity of the output. Along with these results, a measurement in the reliability of the upscaling process of our application could be obtained by analyzing the similarity of the image collection produced by the deep ensemble of networks that form ANUBIS, which is also of vital importance when using generative models in real world applications to avoid the so-called ”hallucinations” and misinformation that can be added in the process.. Fig. 5 showcases three example comparisons between some baseline, our method, and both LR and HR

image patches. The LR and HR pairs were taken from the test set and the LR image patch was upscaled with the different methods for the visual display of the results, following the procedure that was used for the metric evaluation of ANUBIS. It can be seen that ANUBIS is able to retrieve more details from the LR version of the image than the baselines, even in the case where the ground truth images present some distortion due to specific lighting conditions at the moment of capture. As the images used in the training of the proposed application are raw images captured by the satellite with no preprocessing, the calibration process used by NASA in LRO images could also be applied to the super-resolved imagery to remove the mentioned distortions. The reference LR image presented here and in the section, are obtained by applying the summed mode operation to original LRO images.

TABLE I
COMPARISON OF ANUBIS (OURS) VS. BASELINE METHODS USING BOTH PSNR AND SSIM METRICS AVERAGED OVER THE 10,000 IMAGE PATCHES THAT FORM THE TEST DATA.

	Bilinear	Bicubic	FSRCNN	EDSR	**ANUBIS***
PSNR (dB)	23.01	24.01	25.32	25.54	**26.15 $\pm$ 0.97**
SSIM	0.785	0.83	0.86	0.86	**0.91 $\pm$ 0.03**

C. Lunar robotics downstream tasks

In order to further validate ANUBIS, space robotics related tasks were used to showcase the performance of our application. Even though this validation acts as a check of the potential of this application, it can be seen the improvements that image SR can provide to enhance lunar robotics traverse planning based on remote sensing data.

1) Obstacle detection: The first downstream task is lunar surface obstacle detection based on satellite images for lunar rover navigation. First, image gradients are calculated to act as edge detectors. Then a dark spots filter is used to obtain pixel-level information about lunar surface features. This information was processed to create a binary obstacle mask. Every extracted pixel with the image gradients and the low-intensity filter were assigned a value of 1, having the remaining pixels in the mask a 0. This mask is later used as an input for the second downstream task, path planning. In this operation, the goal is to find the same obstacles as in the LR while displaying an amount of obstacles closer to the ground truth in comparison. Fig. 6 shows an example of the results of the obstacle detection task performed over LR, SR, and HR images. Additionally, Table II shows the results of applying the segmentation algorithm to the 10,000 image patches of the test dataset. Rows (1), (2), (3) display the total pixel obstacle count across the test dataset for each LR, SR, and HR capture (ANUBIS obstacle count is performed over the average of each produced distribution). Rows (4) and (5) show the amount of obstacle matches between the HR, SR and LR masks respectively. Row (6) contains the total amount of obstacles with high uncertainty (low reliability) in the SR mask. Row (7) displays the amount of hallucinated obstacles by ANUBIS (compared to the ground truth). 89.233 % of the ground truth obstacles are detected in ANUBIS super-resolved image. In contrast, the percentage of obstacle matches between the ground truth and the LR (enhanced with bicubic interpolation for a more direct and straightforward comparison) is 78.648 %. The uncertainty map generated by ANUBIS is also used to analyze the reliability of the detected obstacles. 92,367 pixel obstacles (2.5 %) have high uncertainty, being 5,847 (0.16 %) of them hallucinated by the GAN. As there is always a tiny chance of an obstacle being a hallucinated feature, with the uncertainty estimation, these areas could be avoided in real-world critical operations such as path planning for rover navigation. It can be seen how ANUBIS output provides better model accuracy for this downstream task when compared to the performance based on LR images.

TABLE II
OBSTACLE DETECTION COMPARISON (TEST DATASET).

	# obstacles	% obstacles
(1) Low Res	3,198,783	-
(2) ANUBIS	3,591,247	-
(3) High Res	3,645,228	-
(4) **HR-SR Matches**	**3,252,770**	**89.233** %
(5) HR-LR Matches	2,866,850	78.648 %
(6) High Uncertainty	92,367	2.5 %
(7) Hallucinated Obstacles	5,874	0.16 %

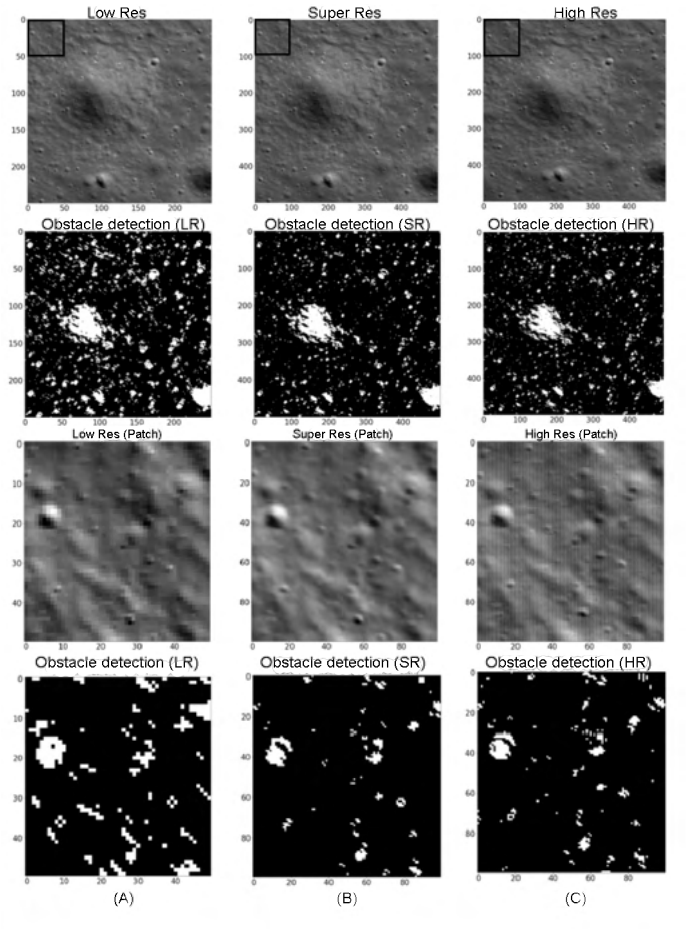

Fig. 6. Obstacle detection image and patch example. (A) Low resolution, (B) ANUBIS, (C) Ground truth. The super-resolved image (B) is the image average of the distribution generated by ANUBIS.

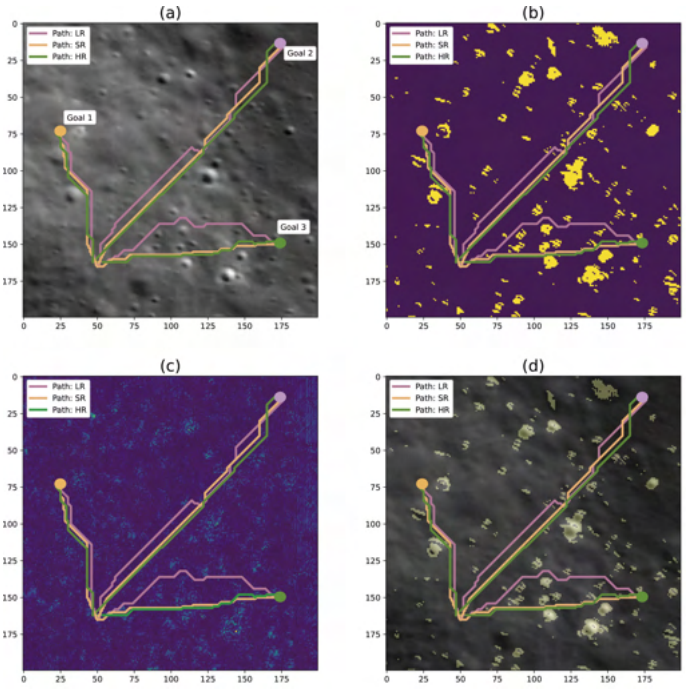

Fig. 7. Path planning downstream task example. 3 different trajectories are overlayed in (a) ANUBIS super-resolved image, (b) Obstacle mask (uncertainty obstacles included), (c) Uncertainty map, (d) ANUBIS super-resolved image + obstacle mask.

2) Traverse planning with uncertainty map: The second downstream task defined for ANUBIS was path planning using super-resolved images. Typically, the image resolution required for this type of analysis is a minimum of 1 meter per pixel or better to include all surface features relevant for rover-/astronaut-scale navigation. The previously defined obstacle detection algorithm was employed to create an obstacle mask. A generated binary mask represents every detected potential obstacle with a pixel value of 1. Then, both starting points and end goals are defined. The goal of the path planning was to define a traversable route for a virtual rover to navigate, avoiding any obstacle by using an A* algorithm with a forward heuristic (shortest path). The ANUBIS uncertainty map can also be a valuable resource for path planning. A certain threshold can be set to exclude high uncertainty regions from the traversable areas. In some of the cases, the distance for the super-resolved image-based trajectory will be greater, but the safety of the task increases with it if the uncertainty estimate is used. In future works, the uncertainty maps estimation will be further explored to assess their potential in super-resolved images based trajectory planning. Quantitative evaluation of this task was done by checking the similarity between distance traveled and trajectories in the test cases. Although this metric might not be optimal for evaluating the improvement obtained with the super-resolved images, it can act as a simple test to see the difference between using an LR or an SR image for the task compared to the HR ground truth. A 100×100 px image (50×50 m area) was used as the operational terrain for the downstream task, with a total of 6 paths generated on it. Table III shows the path's total distance for the different

TABLE III
DISTANCE (M) FOR EVERY GENERATED PATH. FOR EVERY CASE, EQUIVALENT LR/ANUBIS/HR IMAGE INPUTS WERE USED.

Path #	LR	**ANUBIS**	HR
1	74.183	**65.606**	66.435
2	104.681	**100.888**	101.181
3	49.384	**50.177**	50.177
4	54.071	**54.571**	54.571
5	55.384	**55.177**	55.177
6	64.899	**64.571**	64.571

cases.It highlights how the increased resolution allows more precise and closer to the ground truth planning compared to the one estimated with the LR image. Additionally, as the high uncertainty super-resolved pixels are treated as obstacles, the safety is potentially increased if the generated path would be used for real-world navigation. For this downstream task, the utilisation of the test set was not viable as the image patches contained in it were too small (32×32 px) to generate and plan a relevant trajectory. Further research in this field to explore more realistic applications of ANUBIS could combine the information of both distance and uncertainty estimation provided by this work to find the optimal path to the set goals, as it is done for example in [20].

3) Super-resolution based digital elevation map enhancement: Lastly, a proof of concept of how image SR techniques can also enhance elevation maps used for lunar missions is presented as a final downstream task. Typically, digital elevation maps can be obtained through remote sensing images of the lunar surface taken from different orbits (stereophotogrammetry). NASA's Ames Stereo Pipeline (ASP) [16] for example, is a set of software tools that allows, among other things, to perform elevation map generation of planetary locations out of satellite images of missions such as LRO (Moon) and MRO (Mars). Using diverse information about the satellite position, the software framework estimates the stereo features and therefore, the altitude of the different surface elements. The elevation maps display the required data needed for robotics trajectory planning among other tasks, as they present slopes and terrain roughness along with surface obstacles such as craters and boulders.

To showcase the utility of image SR with Digital Elevation Maps (DEM), an additional step was included in NASA's ASP DEM generation process. This new preprocessing step consisted of introducing ANUBIS into the pipeline to enhance the resolution of the images before they were used to generate lunar surface DEMs. This technique indirectly produces better maps that can increase the accuracy of the lunar mission planning and thus, lower its risks.

To validate this method, two different lunar regions were selected that each feature a image stereo pair selected from NASA's LRO NAC dataset. These images were used as the high resolution ground truth (HR) for the experiments. As low resolution images with similar lighting conditions of these areas were unavailable due to the overall scarcity of the NAC dataset, we downsampled the images by using the summed mode operation as introduced above, reducing the resolution by half. Then, ANUBIS was utilised to super-resolve (SR) these synthetically downsampled LR images. Finally, we used

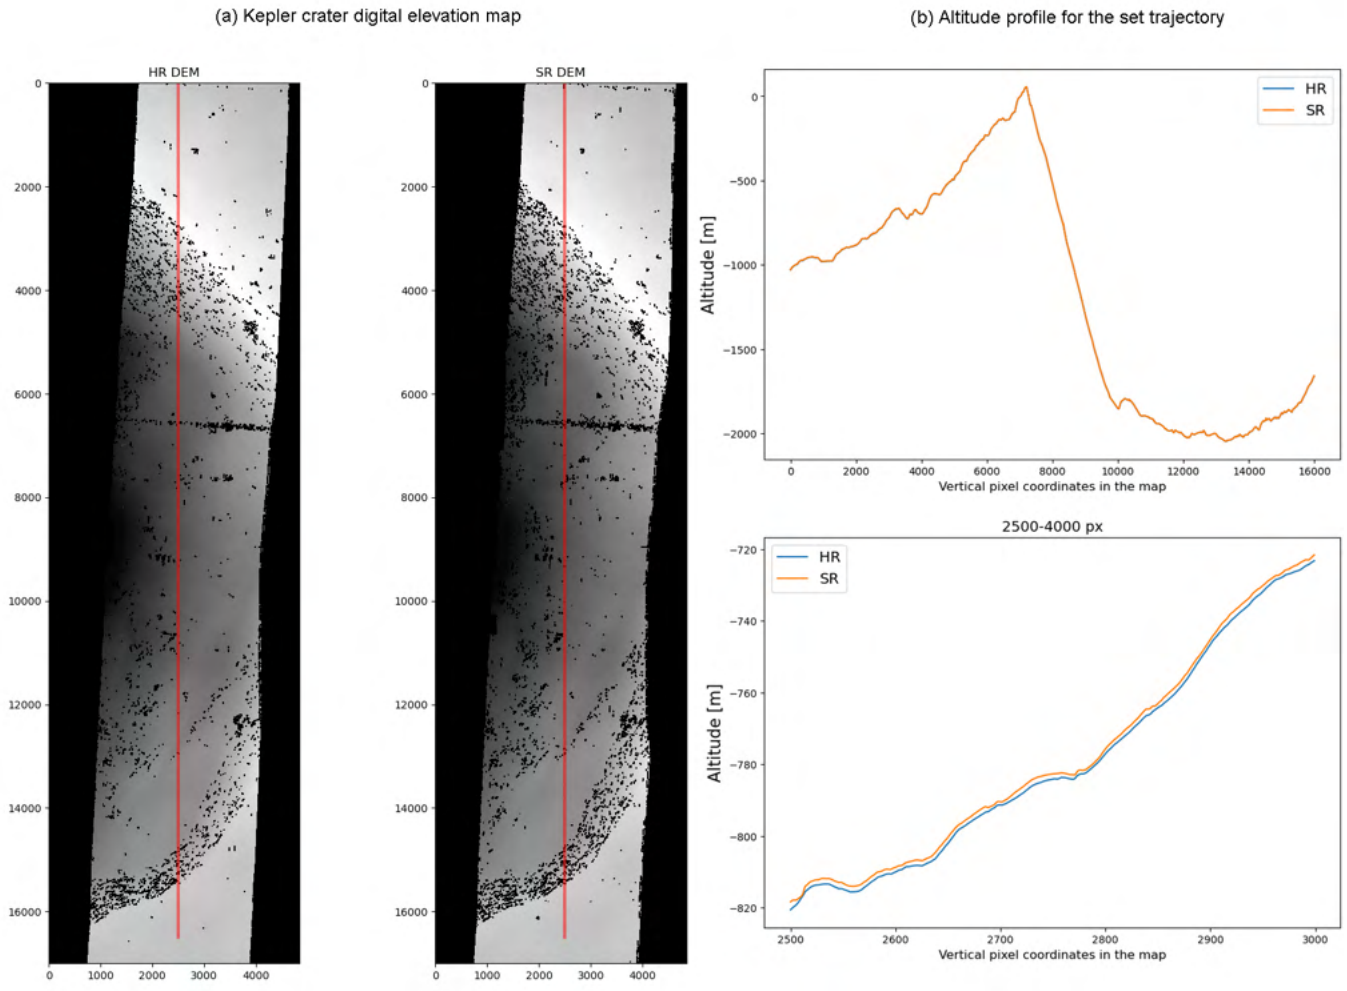

Fig. 8. (a) Digital elevation maps comparison of the Kepler crater between the high resolution ground truth (HR, left) and the super-resolved (SR, right) maps. The red line indicates the selected section to estimate the altitude profile. (b) Altitude profile of the selected region of the Kepler crater. The top figure shows overlapped profiles due to their similarity.

NASA's ASP tools to create two different versions of the elevation map for each region (HR-SR). The exact algorithms used were the Semi Global Matching (SGM) and More Global Matching (MGM) implementations from the ASP.

An example of the resulting elevation maps (Kepler crater) can be seen in fig. 8, along with an altitude profile of a selected region that allow for a comparison of the altitude variation across the maps. Table IV shows the direct altitude error compared between the reference ground truth and the SR version of the elevation map in average and also the maximum error discrepancy. It can be seen how adding image enhancement to the DEM generation pipeline produces elevation maps with altitude values close to the ground truth and its same resolution in each selected region. This map enhancement process was performed in a Future Tech's Digital Workstation powered by an NVIDIA's RTX Quadro 5000.

TABLE IV
ABSOLUTE ALTITUDE ERROR (M) BETWEEN THE HIGH RESOLUTION GROUND TRUTH ELEVATION MAP (HR) AND ITS SUPER-RESOLVED AND DOWNSAMPLED COUNTERPARTS (SR,LR).

Lunar 6.2727	Avg. Error (HR-SR)	Max Error (HR-SR)	Avg. Error (HR-LR)	Max Error (HR-LR)
Kepler Crater	0.6006	12.4150	3.1290	12.8857
Vitello Scarp 1	3.2819	12.4248	6.2727	14.4853

IV. CONCLUSIONS

We have presented a novel application of state-of-the-art single image super-resolution and uncertainty estimation for lunar satellite images. We used Generative Adversarial Networks to double the resolution of images from NASA's Lunar Reconnaissance Orbiter satellite from 1 m/pixel to 0.5 m/pixel. By using a deep ensemble of these networks, we were additionally able to provide uncertainty estimation of the super-resolved images. Our quantitative and qualitative results show that our method outperforms all baseline approaches. We also evaluated the performance of our method on three lunar exploration related tasks, a) surface obstacle detection, b) lunar path planning and c) Digital Elevation Maps generation enhancement. The performance of the application in these tasks showcase how SR can enable enhanced mission planning with the lunar remote sensing data gathered up until now. It

is worth noting that even if the upscaling process is enabling more accurate lunar data processing, the enhanced data should be always handled carefully as there might be missing information worth considering in the mission planning.

The majority of all upcoming lunar missions will be focused on the lunar south pole, which is generally covered by low resolution imagery due to illumination limitations. Our work is the first step in the direction of improving this data, providing an enhanced data products for future exploration missions.

REFERENCES

[1] M. Smith, D. Craig, N. Herrmann, E. Mahoney, J. Krezel, N. McIntyre, K. Goodliff, "The Artemis Program: An Overview of NASA's Activities to Return Humans to the Moon," 2020 IEEE Aerospace Conference, Big Sky, MT, USA, 2020, pp. 1-10.

[2] T. Hoshino, W. Sachiko, O. Makiko, K. Yuzuru, H. Takahiro, M. Hitoshi, S. Hiroaki et al., "Lunar polar exploration mission for water prospection-JAXA's current status of joint study with ISRO." Acta Astronautica, 2020.

[3] A. Colaprete, D. Andrews, W. Bluethmann, R.C. Elphic, B. Bussey, J. Trimble, K. Zacny and J.E. Captain. "An overview of the volatiles investigating polar exploration rover (viper) mission." In AGU Fall Meeting Abstracts, vol. 2019, pp. P34B-03. 2019.

[4] Z. Pan, J. Yu, H. Huang, S. Hu, A. Zhang, H. Ma and W. Sun, "Super-resolution based on compressive sensing and structural self-similarity for remote sensing images". In IEEE Transactions on Geoscience and Remote Sensing, 2013, pp. 4864-4876.

[5] K. Jiang, Z. Wang, P. Yi, G. Wang, T. Lu and J. Jiang, "Edge-enhanced GAN for remote sensing image super-resolution". In IEEE Transactions on Geoscience and Remote Sensing, 2019, pp. 5799-5812.

[6] M. Deudon, A. Kalaitzis, I. Goytom, M.R. Arefin, Z. Lin, K. Sankaran, V. Michalski, S. E. Kahou, J. Cornebise, and Y. Bengio, "Highres-net: Recursive fusion for multi-frame super-resolution of satellite imagery". In arXiv preprint, 2020.

[7] M.C. Raguso, M. Mastrogiuseppe, R. Seu, and L. Piazzo, "Super resolution and interferences suppression technique applied to SHARAD data". In IEEE International Workshop on Metrology for AeroSpace (MetroAeroSpace), 2018.

[8] Y. Tao and J.P. Muller, "A novel method for surface exploration: Super-resolution restoration of Mars repeat-pass orbital imagery". In Planetary and Space Science, 2016, pp. 103-114.

[9] J.I. Delgado-Centeno, P.J. Sanchez-Cuevas, C. Martinez, and M.A. Olivares-Mendez, "Enhancing Lunar Reconnaissance Orbiter Images via Multi-Frame Super Resolution for Future Robotic Space Missions. In IEEE Robotics and Automation Letters, 2021, pp. 7729-7735.

[10] G. Chin, S. Brylow, M. Foote, J. Garvin, J. Kasper, J. Keller, M. Litvak, I. Mitrofanov, D. Paige, K. Raney, and others, "Lunar reconnaissance orbiter overview: The instrument suite and mission". In Space science reviews, 2007, pp. 391-419.

[11] M.S. Robinson, S.M. Brylow, M. Tschimmel, D. Humm, S.J. Lawrence, P.C. Thomas, B.W. Denevi, E. Bowman-Cisneros, J. Zerr, M.A. Ravine and others, "Lunar reconnaissance orbiter camera (LROC) instrument overview". In Space science reviews, 2010, pp. 81-124.

[12] D.C. Humm, M. Tschimmel, S.M. Brylow, P. Mahanti,T.N. Tran, S.E. Braden, S. Wiseman, J. Danton, E.M. Eliason and M.S. Robinson, "Flight calibration of the LROC narrow angle camera". In Space Science Reviews, 2016, pp. 431-473.

[13] L. Wang, Z. Huang, Y. Gong and C. Pan, "Ensemble based deep networks for image super-resolution". In Pattern recognition, 2017, pp. 191-198.

[14] L. Wang, Z. Huang, Y. Gong and C. Pan, "Simple and scalable predictive uncertainty estimation using deep ensembles". In arXiv preprint, 2016.

[15] B. Lakshminarayanan, A. Pritzel and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles". In arXiv preprint, 2016.

[16] R. A. Beyer, O. Alexandrov, and S. McMichael, "The Ames Stereo Pipeline: NASA's open source software for deriving and processing terrain data". In Earth and Space Science, 5, 2018.

[17] C. Dong, C.C. Loy and X. Tang, "Accelerating the super-resolution convolutional neural network". In European conference on computer vision, 2016, pp. 391-407.

[18] B. Lim, S. Son, H. Kim, S. Nah and K. Mu Lee, "Enhanced deep residual networks for single image super-resolution". In Proceedings of the IEEE conference on computer vision and pattern recognition workshops, 2017, pp. 136-144.

[19] S. M. A. Bashir, Y. Wang, M. Khan and Y. Niu, "A comprehensive review of deep learning-based single image super-resolution". In PeerJ Computer Science 7, 2021.

[20] S. Gautam, B. S. Roy, A. Candela and D. Wettergreen, "Science-aware exploration using entropy-based planning". In IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2017