UNIVERSITÉ DU LUXEMBOURG

PhD-FSTM-2023-062
The Faculty of Science, Technology and Medicine

DISSERTATION

Defence held on 07/07/2023 in Luxembourg

to obtain the degree of

DOCTEUR DE L'UNIVERSITÉ DU LUXEMBOURG EN INFORMATIQUE

by

Jose Ignacio DELGADO CENTENO

Born on 29 November 1991 in Spain

LUNAR REMOTE SENSING DATA ENHANCEMENT FOR PRECISE ROBOTIC MISSION PLANNING

Dissertation defence committee

Prof. Dr Miguel Angel OLIVARES MENDEZ, dissertation supervisor
Professor, Université du Luxembourg

Prof. Dr Carlos Perez PEREZ DEL PULGAR MANCEBO
Professor, University of Málaga

Prof. Dr Holger VOOS, Chairman
Professor, Université du Luxembourg

Dr Abigail CALZADA
ESRIC-LIST

Dr. Carol MARTINEZ LUNA, Vice Chairman
Université du Luxembourg

"Once the storm is over, you won't remember how you made it through, how you managed to survive. You won't even be sure, whether the storm is really over. But one thing is certain. When you come out of the storm, you won't be the same person who walked in. That's what this storm's all about."

Haruki Murakami

"In the depth of winter, I finally learned that within me there lay an invincible summer."

Albert Camus

Acknowledgements

Firstly, I would like to thank Miguel Olivares for giving me the amazing opportunity to pursue a PhD in his then-new Space Robotics Research group. He displayed total trust in me when it came to selecting the path for the research that I have undertaken, all the while providing support and the necessary resources to facilitate it. I would also like to seize this opportunity to thank my other two co-supervisors, Carol Martinez and Holger Voos. Carol has provided me with immense support when it came to venturing into this new research field for me. She also demonstrated how to properly manage expectations and handle the difficult aspects of conducting a PhD. Holger gave me valuable feedback that contributed significantly to the current state and quality of the research I have conducted.

Secondly, I would like to express immense appreciation for my colleagues at the Space Robotics Research Group. It has been an absolute pleasure to work alongside all of them. Not only in our day-to-day work at the office, but also outside of it, they have made my stay here in Luxembourg an extraordinary adventure. I am genuinely happy to refer to them as friends, a true gift that I will carry with me wherever I go. During this time, I have been also given invaluable opportunities to work with fantastic people not only from Luxembourg but also from many different universities all over the world. I would like to extend special thanks to all those involved in the Frontier Development Lab (FDL), both to the organizers for giving me the chance to participate in their program, and to the truly incredible researchers I have met and worked with. Their efforts have pushed me and my research to the next level.

Lastly, I would sincerely like to thank my family, my girlfriend and my friends back home. They have always wanted the best for me and have supported every decision I have made throughout my career, consistently providing me with the motivation and confidence I needed to achieve my goals. They have all made this journey possible.

Abstract

In recent times, there has been a resurgence of interest in not only revisiting the Moon but also establishing a lasting and sustainable human presence there. Numerous agencies and private corporations, gearing up for future lunar missions, have established key goals involve both scientific exploration and the emerging lunar economy. These objectives include carrying out scientific experiments, harvesting and utilizing lunar resources on-site, and examining potential methods of using the Moon as an efficient launchpad to go further into the Solar System. In facilitating these lunar missions, robotics emerges as the main disruptive technology. It can allow the execution of various mission-critical activities in harsh environments, ensuring mission success without jeopardizing human safety.

This thesis addresses the challenge presented by the limited resolution of lunar data and the overwhelming amount of non-processed information provided by the different remote sensing missions. Robotic operations on the Moon require tremendous precision in the mission planning phase. For this purpose, the remote sensing data collected by various satellites must be processed and relayed to the mission planning teams in the highest possible resolution and in an easily understandable and digested format. To address this issue, the research presented on this thesis adopts a Machine Learning (ML) approach to enhance and process lunar data gathered by different satellite sensors, providing more detailed and comprehensive insights into the lunar surface, which is crucial for future robotic mission planning. ML provides the necessary tools to efficiently handle and analyze vast volumes of data, a critical aspect in deriving meaningful results, reaching valid conclusions, and deepening our understanding of the subject under study. In this thesis, the author suggests two distinct methods for enhancing and increasing the resolution of lunar images to facilitate improved robot navigation on the lunar surface. The first method involves creating a training dataset using a digital analog environment and utilizing multiple frames of the same location for image enhancement. The second approach proposes a unique architecture that utilizes a single capture from the lunar surface for resolution upscaling, accompanied by an uncertainty estimation of the process. Lastly, a thermophysical analysis of the lunar surface is conducted, which involves processing lunar thermal data in the search for recent asteroid impacts on the lunar surface.

Table of Contents

I Introduction

In recent years, the Moon has captured the attention of both space agencies and private companies, as they seek to conduct various activities on its surface. The Moon's proximity to Earth compared to other celestial bodies in our solar system makes it an ideal location for diverse tasks such as planetary scientific research, in-situ space resource utilization, and even the establishment of a remote base for human activities, which can provide highly valuable insights for future human missions in other extraterrestrial locations such as Mars. NASA's ARTEMIS program (Smith et al., 2020) exemplifies international collaboration in achieving complex objectives and further expanding the limits of space missions.

Future lunar missions hold significant importance for several reasons. The Moon's unique environment provides an opportunity for researchers to test equipment and conduct experiments under conditions unachievable on Earth. This testing ground is crucial for advancing our understanding of how technology and materials behave in extreme environments, such as those encountered in deep space exploration. It also allows for the development and refinement of technologies necessary for sustaining human life during long-duration space missions. By establishing infrastructure and refining the techniques needed for lunar landings, resource extraction, and habitation, space agencies can apply these lessons to more ambitious missions. The Moon's proximity to Earth makes it an ideal location for testing these technologies and strategies that will be required for venturing further into our solar system. It could also function as a logistical hub for deep space missions. The lower gravity and lack of atmosphere make it more energy-efficient to launch spacecraft from the lunar surface than from Earth. In addition, resources such as water ice, which can be converted into rocket fuel, have been detected

on the Moon, particularly in crater regions where the Sun never reaches. Lunar missions can also offer valuable information about the formation and development of our solar system. The Moon's surface has remained relatively unchanged for billions of years, acting as a time capsule for the early history of our celestial neighborhood. By studying lunar samples and geological features of its surface, researchers can learn more about the processes that led to the formation of Earth and other planets, potentially unveiling new information about the origins of life. Robotics has emerged as a fundamental enabler of lunar research in the current era of space exploration. Operating in the harsh and unpredictable conditions of the lunar surface, robotic technologies provide a means to conduct detailed and continuous study without risking human life. Uncrewed lunar rovers, for instance, can traverse various terrains, capturing high-resolution imagery and collecting valuable geological samples. These samples can be analyzed onboard using integrated laboratory instruments or returned to Earth for more extensive study. Moreover, robotic technologies are instrumental in identifying and assessing in-situ resources, such as lunar water ice, which is pivotal for future sustainable human presence on the Moon. They can also carry out intricate tasks like setting up infrastructure for future crewed missions, including habitats, power systems, and communication equipment. The evolution of robotics, therefore, plays a critical role in accelerating our understanding of the Moon, paving the way for the next phase of lunar exploration.

To ensure the success of the lunar missions, it is crucial to possess accurate knowledge of the intended landing sites. During the mission planning phase, data from extraterrestrial locations is collected and analyzed to comprehend the conditions under which the mission will be executed. Remote sensing is a key technology that enables the study of distant locations with extreme environments, without the need for direct physical presence. It involves using satellites and other spacecraft equipped with sensors to collect data on the target location. These sensors can include cameras, spectrometers, radar systems, and more, each designed to gather specific types of information. The satellite may be placed in a low lunar orbit or a more distant orbit, depending on the objectives and requirements of the mission. The choice of orbit affects factors such as data resolution, coverage, and the frequency with which the satellite passes over a particular area. Researchers on the ground then process and analyze the data to create detailed maps, identify potential landing sites, and understand the composition and characteristics of

the lunar surface. This information is crucial for planning mission trajectories, designing landers and rovers, and determining the best locations for experiments and resource extraction.

1 Remote sensing for the Moon

The Moon has a rich history of orbiters sent for study and analysis purposes since 1959, with Luna, Lunar Orbiter, Surveyor, and Apollo being examples of early lunar observation missions. In recent years, with the resurgence of interest in space exploration, national agencies from various countries have started investing in research and preparation for lunar missions. Japan Aerospace Exploration Agency (JAXA) launched Kaguya (Kato et al., 2010), its second lunar orbiter, in 2007. This mission placed a primary satellite 100 km above the Moon's surface, accompanied by two smaller orbiters in polar orbit, each equipped with different instruments. Kaguya's objective was to gather data to enhance understanding of the Moon's history and develop the necessary technology for future lunar exploration. India Space Research Organization (ISRO) developed the Chandrayaan (Goswami and Annadurai, 2009 and Sundararajan, 2018) program over the past decade, which includes two lunar orbiters launched in 2008 and 2019. The first mission aimed to create a 3D topography of the lunar surface, while the second one sought to study variations in surface composition and determine the location and abundance of lunar water. China National Space Administration (CNSA) launched Chang'e 1 (Ouyang et al., 2010) and 2 (Zhao et al., 2011) in 2007 and 2010 respectively, with the goal of mapping the Moon and analyzing its regolith chemical composition. These missions represent some of the most well-known examples of remote sensing for the Moon, allowing various space agencies to gain a better understanding of the lunar surface, its composition, and the requirements for continued robotic and human exploration. NASA's Lunar Reconnaissance Orbiter (LRO), the mission responsible for providing the largest publicly available dataset of the Moon's surface was launched in 2009. The LRO satellite carried seven different instruments and sensors into lunar orbit for a one-year exploration mission. Its primary objective was to scout the lunar surface for high-value resources and potential landing sites to support future human presence on the Moon. Following this initial phase, a two-year scientific mission began, studying aspects ranging from the history of the solar system to the mineralogy and geology of

the lunar surface. Since then, the satellite's orbit has been adjusted to minimize fuel consumption, and the scientific mission has been extended five times. As a summary, the following table contains the different lunar orbiters throughout the history of space missions:

Table I.1: Remote sensing lunar missions

Name	Year	Country	Description
Luna 3	1959	URSS	First satellite to return images of the far side of the Moon
Luna 10	1966	URSS	First spacecraft to orbit the Moon
Luna 12	1966	URSS	Photography mission
Lunar Orbiter 2	1966	USA	Photographic mapping mission
Lunar Orbiter 3	1967	USA	Photographic mapping mission
Lunar Orbiter 4	1967	USA	Photographic mapping mission
Explorer 35	1967	USA	Study of interplanetary plasma and magnetic fields
Lunar Orbiter 5	1967	USA	Photographic survey
Luna 14	1968	URSS	Tested communication from lunar orbit
Luna 22	1974	URSS	Study of magnetic fields
Clementine	1994	USA	Lunar observation and imaging at various wavelengths
Lunar Prospector	1998	USA	Mapping of lunar surface composition
SMART-1	2003	EU	Photographic mapping mission
THEMIS-ARTEMIS	2007	USA	Magnetospheric research
Kaguya	2007	Japan	Study lunar surface environment
Chang'e 1	2007	China	Obtaining 3D images of lunar geological structures
Chandrayaan-1	2008	India	Surveying the lunar surface
LRO	2009	USA	Detailed mapping and resource localization
Chang'e 2	2010	China	Photographic mapping mission
GRAIL	2011	USA	Gravitational field mapping of the Moon
LADEE	2013	USA	Lunar exosphere analysis
Chang'e 3	2013	China	Photographic mapping mission
Chang'e 4	2014	China	Photographic mapping mission
Queqiao	2018	China	Communication experiments from lunar orbit
Chandrayaan-2	2019	India	Lunar topography and mineralogy study
Danuri	2022	South Korea	Lunar resources survey

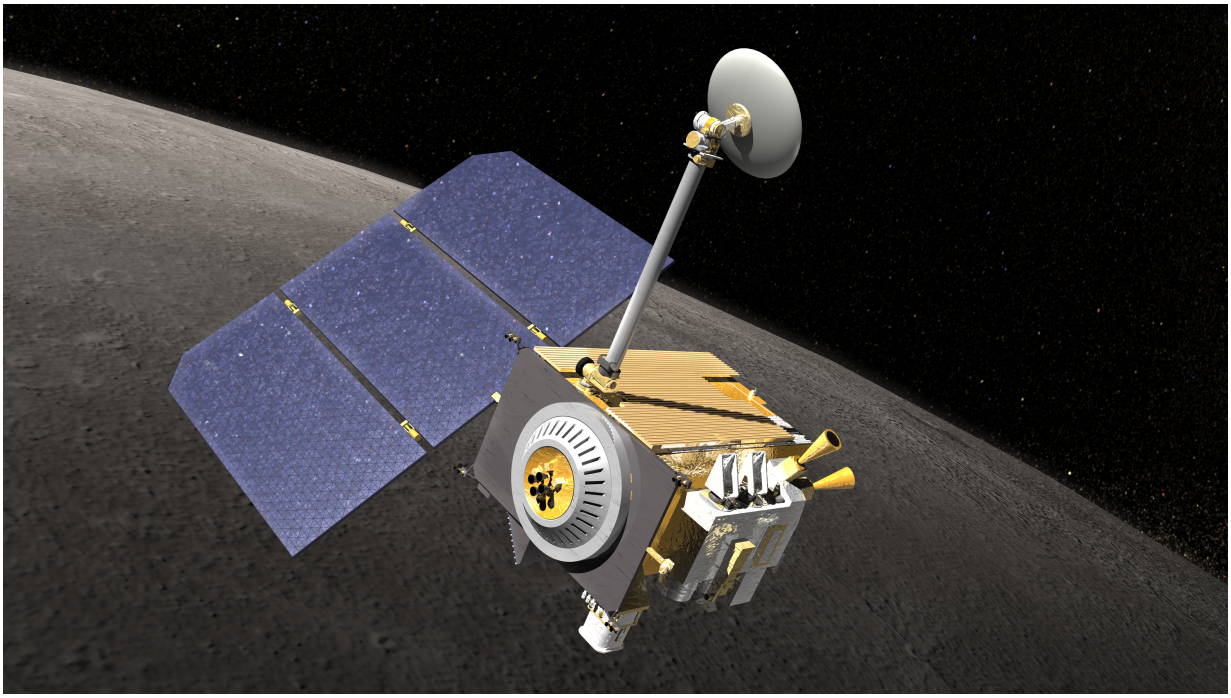

Figure I.1: NASA's Lunar Reconnaissance Orbiter satellite. Source: NASA

1.1. Lunar science

NASA's LRO mission (Fig. I.1) has built since 2009 the most complete collection of lunar data so far thanks to the following seven instruments that have been analyzing the surface since then:

- **Cosmic Ray Telescope for the Effects of Radiation (CRaTER)**: Evaluates the global lunar radiation environment and assesses the biological impact of cosmic radiation by using a "human tissue-equivalent" plastic to simulate human exposure (Schwadron et al., 2012).

- **Diviner Lunar Radiometer (DLRE)**: Collects thermal data of the lunar surface, enabling the creation of detailed day and night surface temperature models. This data also offers insights into ice deposits at the Moon's poles, surface composition, and subsurface temperatures (Bandfield et al., 2011a).

- **Lyman Alpha Mapping Project (LAMP)**: Maps the lunar surface in the far ultraviolet spectrum, aiding in the search for surface ice and frost in polar regions. Additionally, it captures images of permanently shadowed regions (PSRs) that are

illuminated only by starlight, which cannot be imaged by the main camera instrument (Gladstone et al., 2010b).

- **Lunar Exploration Neutron Detector (LEND)**: Detects the presence of hydrogen and hydrogen-bearing compounds on the lunar surface, providing information for the search for water ice on the Moon's surface and supplying space radiation data to better prepare for future human exploration missions (Mitrofanov et al., 2010a).
- **Lunar Orbiter Laser Altimeter (LOLA)**: Measures the altitude of the lunar surface at each measured point, helping to analyze landing site slopes and surface roughness. This data can also be used to generate 3D maps of the Moon, offering valuable information for studying PSRs (Smith et al., 2017).
- **Lunar Reconnaissance Orbiter Camera (LROC)**: Captures high-resolution black and white images of the lunar surface, covering nearly the entire lunar surface except for PSRs, where optical information cannot be retrieved due to the inherent environmental lighting conditions (Robinson et al., 2010).
- **Mini-RF technology**: Supplies additional information on the presence of water ice in the lunar poles and serves as a testing ground for future communication technology demonstrations (Nozette et al., 2010).

The data is accessible through NASA's archive, the Planetary Data System (PDS). This extensive repository contains information not only about the Moon but also about Mars and other celestial bodies studied through remote sensing missions. Other space agencies, such as JAXA and ISRO, have also contributed to this enormous archive with data from missions like Kaguya and Chandrayaan-1. As a result, the PDS holds a vast array of datasets for planetary science research, with the LRO mission being one of the most significant contributors concerning the Moon.

To visually explore these lunar datasets, the Quickmap tool can be utilized. This application enables users to overlay data maps on top of an image mosaic of the lunar surface, which was captured using the LRO Camera instrument. Quickmap allows for the examination of all available data products in PDS for specific coordinates or regions and facilitates downloading the selected content. This makes it an invaluable resource

for researchers and enthusiasts alike, providing easy access to a wealth of lunar data. Researchers utilize this information to deepen their understanding of the history of our solar system, including Earth. They also examine potential space resources present on the lunar surface. Additionally, space agencies are using this data to meticulously plan for future robotic and human lunar missions, such as those outlined in the Artemis program. These endeavors underscore the significance of analyzing lunar data for various purposes, from enhancing our scientific knowledge of the Moon to readying ourselves for future space exploration. An example of these utilities can be seen in figure I.2, where the South Pole's PSRs are highlighted. The following subsections offers a concise overview of some of the most prominent aspects of lunar research.

1.1.1. Lunar terrain

The research of lunar terrain, this is, its topography and geomorphology, is a prevalent area of lunar science. Missions like the LRO, Kaguya, and Clementine (Spudis et al., 1994) have employed various instruments dedicated to mapping and terrain analysis. For instance, between 2009 and 2016, LRO's Lunar Orbiter Laser Altimeter (LOLA) gathered a remarkable 7 billion measurements of the lunar surface using its laser, supplying data on altitude, surface roughness, slope, and even topographic information about polar craters and permanently shadowed regions. The LRO Camera instrument has also significantly contributed to the geomorphological analysis of the lunar surface, with a dataset currently comprising over 2 million images. Stereo images from different locations have been utilized to create terrain maps (Tran et al., 2010), and multi-view shape-from-shading techniques have been applied to lunar data to enhance and optimize terrain models derived from stereo image pairs (Alexandrov and Beyer, 2018). These efforts showcase the extensive and diverse information that can be acquired by studying lunar terrain. Generating lunar surface roughness maps is another vital aspect of lunar science. The global surface slopes and roughness map obtained from LOLA data has yielded valuable insights into the distinct properties of lunar soil between the highlands and mare regions (Rosenburg et al., 2011). This information is crucial for mission planning, as it offers a better comprehension of the environment in which future lunar missions will operate. Moreover, the Moon's lack of atmosphere, geological activity, and erosional processes (beyond asteroid impacts) means that surface information is critical for understanding the history of the solar system. The lunar surface serves as

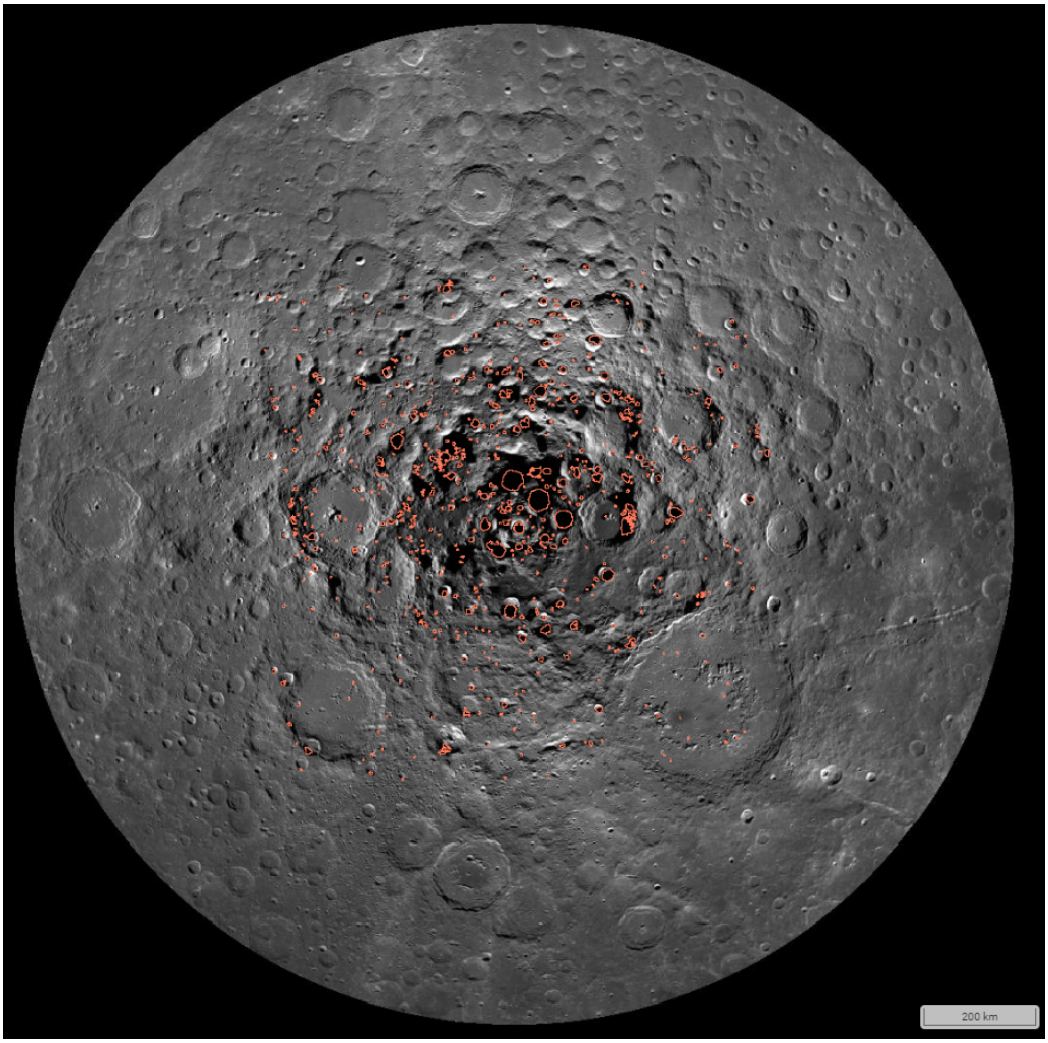

Figure I.2: NASA's Lunar Reconnaissance Orbiter map of the PSRs in the South Pole. Source: NASA Quickmap

an ancient archive of past Solar System activity, making the study of this information essential to understand our origins and the formation of our planet. The surface radiance gathered by LRO's Diviner instrument has also been employed to create roughness maps and analyze various properties of the regolith (Bandfield et al., 2015). These efforts emphasize the importance of understanding the properties of the lunar surface for both mission planning and scientific exploration.

1.1.2. Lunar regolith

Lunar regolith, the layer of loose, fragmented material that covers the solid bedrock of the Moon's surface, presents significant potential for a variety of uses in future lunar exploration. This material, comprised by fragments of rock, breccias, isolated mineral grains and impact melts of varied compositions as well as agglutinates developed by space weathering., is an abundant resource that can be harnessed for various applications. Firstly, lunar regolith can be used as a raw material for construction on the Moon. With advancements in 3D printing technology, it is conceivable to use regolith to construct habitats and other infrastructure, thereby reducing the need to transport heavy materials from Earth. Furthermore, regolith may be processed to extract useful elements, such as oxygen, silicon, and various metals, which can support life support systems and fuel production for lunar bases or deep space missions. Additionally, the scientific study of lunar regolith can offer valuable insights into the Moon's geological history and the broader history of our solar system. As such, lunar regolith holds a pivotal role in enabling sustainable human presence and exploration activities on the Moon. One particularly vital aspect of regolith study involves the analysis of its thermophysical properties (Vasavada et al., 2012a). By examining the radiance emitted by the lunar surface during both day and night cycles, scientists have determined the composition of regolith across the Moon. Diviner data was employed to create a global thermophysical model of the Moon (Hayne et al., 2017), mapping the thermal conductivity, thermal inertia, and heat capacity of the lunar soil. This model was further refined through studying near-surface regolith around the equator (Vasavada et al., 2012b). The surface temperatures of the regolith also offer insight into rock abundance (Bandfield et al., 2011b), which is connected to both the age of craters and the thickness of the regolith. The mini-RF data from the LRO has also been utilized in preliminary investigations, revealing that regolith fines in polar craters contain low metal abundance (Heggy et al., 2020).

1.1.3. Water presence

Lastly, it is important to highlight the significance of water, an essential element for future human missions to the Moon. In extraterrestrial environments, water is crucial for maintaining long-term human presence, as it serves as a life support resource and a fuel

source. Therefore, a considerable portion of lunar research is dedicated to discovering water ice and volatiles on the lunar surface. The availability of an adequate amount of ice would make certain locations favorable for future missions. The Moon's polar regions are the primary focus of this search for water, as they are thought to contain large ice deposits confined in shadowed areas known as Polar Shadowed Regions (PSRs). These regions are among the coldest in the entire solar system, having not received sunlight for millions of years. In the event of a comet impact, water ice from the asteroid could combine with the lunar soil rather than sublimating (Li et al., 2018), due to the absence of an atmosphere.

In the last decade, a variety of remote sensing instruments have been employed to detect water ice on the lunar surface. ISRO's Chandrayaan-1 mission aimed its Mini-SAR instrument at the North pole (Spudis et al., 2010), while the Chandrayaan-2 mission targeted the South pole (Kumar et al., 2022). The Lunar Crater Observation and Sensing Satellite (LCROSS) mission (Colaprete et al., 2012), under the Lunar Reconnaissance Orbiter (LRO), unveiled the presence of water ice on the lunar surface through an impact event (Colaprete et al., 2010). The LRO also carries additional sensors, such as the Mini-RF radar, which analyzed the surface properties of polar craters and found anomalous regions indicating the existence of water ice deposits (Spudis et al., 2013). Data gathered after the LCROSS impact on the lunar surface estimated the hydrogen concentration in the Cabeus crater impact site, corresponding to 0.5-4.0% water ice by weight (Gladstone et al., 2010a; Mitrofanov et al., 2010b), depending on the thickness of the dry regolith layer.

2 Machine Learning

2.1. Super Resolution

Image super-resolution (SR) (Van Ouwerkerk, 2006; Yue et al., 2016) is a crucial category of image processing techniques within computer vision and image processing. It refers to the technique for recovering high-resolution (HR) images from their low-resolution (LR) counterparts. This process finds widespread practical real world applications across various domains, including medical imaging, remote sensing, and face recognition.

In medical imaging (Mahapatra et al., 2019; Zhang et al., 2018), SR techniques aid in enhancing the clarity and level of detail in medical scans, enabling more accurate diagnoses and treatment planning. Remote sensing applications benefit from SR by improving the resolution of satellite or aerial imagery (Dong et al., 2020; Yang et al., 2015), leading to better analyses and informed decision-making in areas like agriculture, environmental monitoring, and disaster management. Moreover, SR also proves invaluable in face recognition techniques (Jiang et al., 2016; Qin and Li, 2020), where it helps sharpen facial features and improve identification accuracy, thereby contributing to security and authentication systems.

As the demand for higher-quality visual data continues to grow across multiple industries, image super-resolution remains a critical field of research, with ongoing advancements in deep learning and AI-powered approaches revolutionizing its capabilities. The ever-expanding real-world applications of SR underline its importance and significance in the realm of computer vision and image processing.

Super-resolution techniques can be broadly classified into two categories based on the number of input LR images: single image super-resolution (SISR) and multi-frame super-resolution (MFSR). In the case of MFSR (Farsiu et al., 2004; Huang et al., 2015), the goal is to reconstruct a clean, sharp, and often higher-resolution output image from multiple degraded and noisy input images of a scene. By recursively fusing the information from different input images, MFSR approaches excel at reconstructing finer details that are not attainable from a single image alone.

In contrast, SISR (Chen et al., 2022; Yang et al., 2019) is more widely popular due to its higher efficiency. However, SISR have significant challenges as it deals with an inherently ill-posed problem. A specific low-resolution (LR) input can correspond to multiple potential high-resolution (HR) images, making the HR space (often representing the natural image space) intractable to map the LR input accurately. SISR methods face two primary drawbacks: firstly, the ambiguity in defining the mapping between the LR and HR space, and secondly, the difficulty in establishing a complex high-dimensional mapping using vast amounts of raw data.

Recently, deep learning-based SR methods (Deudon et al., 2020; Ledig et al., 2017; Lim et al., 2017) have shown remarkable progress by capitalizing on their ability to extract powerful high-level abstractions that bridge the gap between the LR and HR space.

These advanced techniques have achieved significant improvements in both quantitative performance and visual quality of the output images.

2.2. Normalizing Flow

In the domains of statistics and machine learning, a fundamental objective has been to model probability distributions based on samples drawn from those distributions. This task falls under unsupervised learning and is commonly referred to as generative modeling. Its significance lies in the abundance of unlabelled data compared to labelled data. Various applications leverage this approach, including density estimation, outlier detection, prior construction, and dataset summarization.

One notable family of generative models that offers tractable distributions with efficient and precise sampling and density evaluation is known as Normalizing Flows (NF, Ho et al., 2019; Rezende and Mohamed, 2015; Winkler et al., 2019). A Normalizing Flow involves transforming a simple probability distribution (e.g., a standard normal) into a more complex distribution through a sequence of invertible and differentiable mappings. This method offers a powerful mechanism to create new families of distributions. It begins with choosing an initial density and then chaining together a set of parameterized, invertible, and differentiable transformations. This results in a new density that can be efficiently sampled from (by sampling from the initial density and applying the transformations), and the density at a specific sample (i.e., the likelihood) can be computed as described above. The Normalizing Flow approach, allows to effectively construct and work with sophisticated distributions, offering flexibility in representing various complex data patterns. The ability to efficiently sample from and evaluate the density of these distributions makes Normalizing Flows an attractive and practical choice for generative modeling tasks.

3 Objectives and contributions

Lunar remote sensing data serves as a crucial tool for gaining insights across a wide array of scientific fields. Thanks to the collected data, researchers have been able to deepen their understanding of the Solar System's history, pinpoint potential sources of water ice for future human missions, and identify safe landing sites on the Moon. To

effectively process and analyze this data, tools such as the Ames Stereo Pipeline (ASP) and Integrated Software for Imaging Spectrometers (ISIS) play a vital role. They are used to generate valuable products like Digital Elevation Maps (DEMs) and calibrated images. However, despite the abundance and diversity of available data, challenges remain. Data can sometimes be insufficient or inaccurate for specific tasks, or overwhelming and difficult to manage due to its volume. This is where Artificial Intelligence (AI) enters the picture. AI holds the potential to revolutionize the way we process remote sensing data by automating the analysis of massive amounts of information, leading to accelerated progress in scientific research and a deeper understanding of various topics. Moreover, AI can be used to improve and enhance existing data. Considering that the only alternative to obtaining more accurate data through remote sensing is to launch new satellites with more advanced instruments and sensors, AI offers a more cost-effective and efficient approach. AI techniques, such as Machine Learning and Deep Learning, can be employed to identify patterns and correlations in the data, enabling more accurate predictions and analysis. For instance, AI algorithms can be used to automatically classify terrain features, detect changes in the lunar surface over time, and identify potential resources. Furthermore, AI can assist in mitigating issues with data quality by filling gaps or correcting inconsistencies in the data, ultimately improving the overall reliability of the information. By integrating AI into the processing and analysis of lunar remote sensing data, scientists can overcome many of the current challenges, leading to more accurate and efficient research outcomes. As a result, this will enable a better understanding of the Moon, and pave the way for more advanced and successful future missions, both robotic and human.

In this regard, this thesis focused on answering two main research questions: what methods can be employed to obtain scientific insights from the vast quantities of lunar data available and how can we refine and enhance the already collected lunar data while maintaining its scientific validity.

3.1. What methods can be employed to obtain scientific insights from the vast quantities of lunar data available?

Typically, the enormous volume (on the order of terabytes) of lunar datasets makes extracting scientific conclusions a complex task, due to the considerable number of entries

in the data collection and the need for manual analysis of the information for validation. This complexity arises from several factors, such as the high dimensionality of the data, the diverse nature of measurements acquired by various instruments and the scale of the datasets. Moreover, the validation process often involves the expertise of domain specialists, who must carefully study the data to obtain scientific conclusions, which is a rather time-consuming and slow task. As a result, it is essential to develop efficient methods and tools that can effectively handle, process, and analyze these vast quantities of lunar data. Ultimately, the goal is to maximize the scientific value and insights derived from the available data, thereby advancing our understanding of the Moon and its history, geology, and potential for future exploration and utilization.

3.2. How can we refine and enhance already collected lunar data while maintaining its scientific validity?

When it comes to obtaining higher resolution data in extraterrestrial environments, there are two main options. The first alternative is to design, construct, and launch a new satellite equipped with advanced instrumentation. This new data would offer improved resolution, leading to increased precision in mission planning. However, this solution would necessitate substantial time and financial investments, making it a challenging endeavor. Additionally, it would take years to build a remote sensing dataset as extensive as NASA's LRO, for instance. It is worth noting that in-orbit satellites could also obtain more accurate data by approaching closer to the lunar surface, but this would come at the cost of expending fuel and dramatically reducing the mission duration, rendering it non-viable.

3.3. Thesis overview

The aim of this thesis was to emphasize the practical aspects of research, demonstrating how ML research can be applied to actual planetary science issues and the advantages that researchers can gain through its implementation. The primary contributions of this thesis involve exploring the potential of ML architectures in lunar science applications, improving the existing dataset to offer higher resolution data with scientific assurances, and processing vast volumes of data to obtain scientific results in a more efficient way, addressing the two research questions outlined in Section 3. This work has contributed

being one of the first exploring the topic of lunar data enhancement using DL-based architectures and showcasing potential benefits of efficient ML application alternatives to current methods used by planetary science researchers.

The thesis proposes utilizing deep learning-based techniques for direct lunar data enhancement, providing a more practical alternative to deploying satellites with better instrumentation into orbit and allowing for direct application to the vast archives of data already gathered. These methods have the potential to significantly improve data resolution, yielding higher quality and more detailed versions. As data processing tasks, they require only computational resources, which is a considerable reduction in the financial and time investments compared to launching new space missions. This thesis additionally proposes two distinct machine learning-based approaches to derive scientific conclusions from the largest dataset in the LRO collection, the thermal measurements gathered by the Diviner instrument. Through the development of these architectures here introduced, this thesis demonstrates a comprehensive pipeline that goes from raw thermal point measurements to lunar surface analysis results.

The structure of this work follows a cumulative thesis format. After a thorough introduction, which covers the state of the art, research questions, and thesis contributions, the relevant publications that were produced as part of the PhD are presented in chapters 2, 3 and 4. A list of all the publications relevant to this research is provided in the Appendix. Subsequently, the author's individual contributions to the included research papers are discussed.

4 Relevant publications

4.1. Articles in this dissertation

- Research paper [1]: J. I. Delgado Centeno, P. J. Sanchez Cuevas, C. Martinez Luna, and M. A. Olivares Mendez (2021b). "Enhancing lunar reconnaissance orbiter images via multi-frame super resolution for future robotic space missions". In: *IEEE Robotics and Automation Letters* 6.4, pp. 7721–7727.

- Research paper [2]: J. I. Delgado Centeno, P. Harder, B. Moseley, V. Bickel, S. Ganju, M. Olivares-Mendez, and A. Kalaitzis (2021a). "Single Image Super-Resolution with Uncertainty Estimation for Lunar Satellite Images". In: *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*.

- Research paper [3]: S. Bucci, J. I. Delgado Centeno, B. Gaffinet, Z. Liang, V. Bickel, B. Moseley, M. Olivares-Mendez, et al. (2022). "Thermophysical Change Detection on the Moon with the Lunar Reconnaissance Orbiter Diviner sensor". In: *Proceedings of the Neural Information Processing Systems" Machine Learning and the Physical Sciences" Workshop (ML4PS)*. Neural Information Processing Systems Foundation, Inc.

4.2. Other peer-reviewed publications not included in this dissertation

- J. I. Delgado Centeno, P. Harder, B. Moseley, V. Bickel, S. Ganju, M. Olivares-Mendez, and A. Kalaitzis (2021a). "Single Image Super-Resolution with Uncertainty Estimation for Lunar Satellite Images". In: *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*.

4.3. Other non peer-reviewed publications not included in this dissertation

- J. I. Delgado Centeno, P. J. Sanchez Cuevas, C. Martinez Luna, and M. A. Olivares Mendez (2021d). "Lunar Surface Images Enhancement for Space Resources Localization and Extraction". In: Space Resources Week

- J. I. Delgado Centeno, P. J. Sanchez Cuevas, C. Martinez Luna, and M. A. Olivares Mendez (2021c). "Lunar Highres-Net: Super Resolution For Lunar Surface Imagery". In: 72th International Astronautical Congress, IAF Space Exploration Symposium

4.4. Individual contribution to the included research papers

- **RP1 [Enhancing lunar reconnaissance orbiter images via multi-frame super resolution for future robotic space missions]:**
  - Role: Lead co-authorship
  - Description: : Conceptualization, methodology, data preprocessing, implementation, validation, writing and review. Visualization.

- **RP2 [Super-Resolution of Lunar Satellite Images for Enhanced Robotic Traverse Planning]**
  - Role: Lead co-authorship
  - Description: Conceptualization, methodology, data preprocessing, implementation, validation, writing and review.

- **RP3 [Thermophysical Change Detection on the Moon with the Lunar Reconnaissance Orbiter Diviner sensor]:**
  - Role: Lead co-authorship
  - Description:Conceptualization, methodology, data preprocessing, implementation, validation, writing and review.

II | Research Paper 1

Full Title: Enhancing lunar reconnaissance orbiter images via multi-frame super resolution for future robotic space missions

Authors: J. I. Delgado Centeno, P. J. Sanchez Cuevas, C. Martinez Luna and M. A. Olivares Mendez,

Published in: IEEE Robotics and Automation Letters (RAL)

Abstract:
This paper presents a novel application of a Multi- frame Super Resolution (MFSR) method for lunar surface imagery called Lunar HighRes-net (L-HRN). In this work, we adapted and used NASA's Lunar Reconnaissance Orbiter (LRO) image database to train the Deep Learning architecture for image super resolution. Additionally, we also gathered an artificial image dataset from our virtual Moon to improve the amount of input data in the neural network training process. The network's architecture follows a standard MFSR algorithm that was enhanced for this specific use case. The proposed MFSR method has been evaluated using the well-known peak signal-to- noise ratio (PSNR) metric against other generic super-resolution methods of the state of the art. This work aims to improve environmental knowledge about the lunar surface to enhance future autonomous robots capabilities on the surface of the Moon.

Keywords: Space Robotics and Automation: Aerial and Field Robotics

DOI: 10.1109/LRA.2021.3097510

Scopus: 94%

Enhancing Lunar Reconnaissance Orbiter Images via Multi-Frame Super Resolution for Future Robotic Space Missions

J. I. Delgado-Centeno, P. J. Sanchez-Cuevas, C. Martinez, and M.A. Olivares-Mendez

***Abstract*—This paper presents a novel application of a Multi-frame Super Resolution (MFSR) method for lunar surface imagery called *Lunar HighRes-net (L-HRN)*. In this work, we adapted and used NASA's Lunar Reconnaissance Orbiter (LRO) image database to train the Deep Learning architecture for image super resolution. Additionally, we also gathered an artificial image dataset from our virtual Moon to improve the amount of input data in the neural network training process. The network's architecture follows a standard MFSR algorithm that was enhanced for this specific use case. The proposed MFSR method has been evaluated using the well-known peak signal-to-noise ratio (PSNR) metric against other generic super-resolution methods of the state of the art. This work aims to improve environmental knowledge about the lunar surface to enhance future autonomous robots capabilities on the surface of the Moon.**

***Index Terms*—Space robotics and automation: aerial and field robotics.**

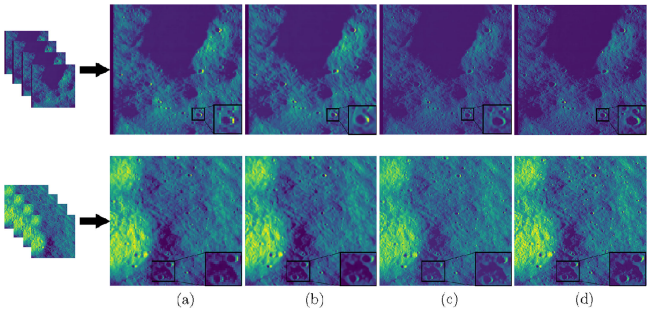

Fig. 1. Super resolution methods applied to NASA's LRO mission lunar surface images. (a) Bicubic, (b) ESDR, (c) Lunar HighRes-net (this work), and (d) Ground Truth. The images are presented in color map to perceive better the difference between the results of each super resolution method.

I. Introduction

SPACE applications are currently attracting the interest of several agencies and companies such as NASA, ESA, SpaceX and Blue Origin. They are investing a lot of resources in new missions [1]–[4] and exploitation plans. These missions present a wide variety of objectives and challenges, such as the study of the geological composition of celestial bodies and the study of life's presence at some point in the history of Mars, among others. In most cases, they involve robotic systems to perform *in-situ* and remote-sensing operations to avoid putting at risk human life or integrity. Two recent examples are the *Perseverance* rover [5] for Mars exploration and the *Volatile Investigating Polar Exploration Rover (VIPER)* [6] which will operate on the lunar surface. Particularly, *Perseverance* rover [5], has a crucial role in the NASA's Mars 2020 mission [2]. It counts with a wide variety of sensors to monitor and study the red planet's environmental conditions, evaluate signs of life and gather data to prepare human exploration in Mars. The *VIPER* rover [6] will explore the South Pole of the Moon and gather data about water-ice concentration in this region.

Nowadays, autonomous space missions heavily rely on the perception of the local environment. Moreover, it is well known that the conditions where rovers navigate are varied and harsh, and the success of the mission is ligated to have solid prior planning. Then, it is crucial to have a good understanding of the issues that a rover can face while performing its mission, such as the type of obstacles, slopes, and craters in the way of the robot.

In this sense, it is clear that having good models of the surface becomes necessary to select the landing site and optimally pre-plan the mission. This is a crucial part while defining a mission because, although the robot has onboard sensors to avoid dealing with unexpected situations, the chances of success are doubtful. Moreover, considering the limited endurance of these rovers due to extreme environmental conditions, it is vital to minimize this kind of reactivity actions. Fig 1 shows one specific result of this work applying image super resolution (SR) to lunar images. The authors identify this technique as a promising tool that can improve decision-making in future missions.

Pre-landing information is mostly obtained through remote sensing missions like NASA's Lunar Reconnaissance Orbiter (LRO) [7] and ISRO's Chandrayaan-2 [8]. The first one is a mission with the sole purpose of mapping the lunar surface by using satellite instrumentation, providing an almost complete lunar map. In contrast, the second one is a lunar exploration mission consisting of a lunar orbiter, a lander and a rover. In this

Manuscript received February 24, 2021; accepted June 20, 2021. Date of publication July 16, 2021; date of current version August 17, 2021. This letter was recommended for publication by Associate Editor J. Cacace and Editor P. Pounds upon evaluation of the reviewers' comments. *(Corresponding author: José Ignacio Delgado-Centeno.)*

The authors are with the SpaceR, Space Robotics Research Group, SnT Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg, Luxembourg 1855, Luxembourg (e-mail: Jose.delgado@uni.lu; psanchez16@us.es; carol.martinezluna@uni.lu; miguel.olivaresmendez@uni.lu).

Digital Object Identifier 10.1109/LRA.2021.3097510

case, the orbiter analyzed the surface area prior to landing the rest of the mission's components with a high-resolution camera.

However, there can be complications due to the lack of resolution in the gathered data used with this purpose because it can potentially omit relevant details that would be vital for the mission's design. For example, NASA's LRO [7] have a resolution of up to 0.5 meters per pixel in the images obtained through the Narrow-Angle Camera (NAC). Furthermore, database statistics show that only a 40 % of the images are this accurate. The rest of the database presents half or worse (1-5 m/px) resolution, which could lead to the issues previously mentioned. Unfortunately, nowadays, the only alternative to improve this data seems to be the launch of new space missions with better instrumentation and sensors, which is costly and challenging to accomplish.

Image enhancement research on the field of remote sensing applied to space bodies has not been widely explored, but in its equivalent used on Earth imagery, there have been some improvements in recent years. In [10], [11] and [12] different methods of Single Image Super-Resolution (SISR) has been proved and applied to already existing databases. This SISR process enhances the quality of an image using either traditional Computer Vision (CV) algorithms or Deep Learning (DL) architecture, such as Convolutional Neural Networks (CNN) or Generative Adversarial Networks (GAN). Depending on the type of input image (RGB, Grayscale, Multispectral), the network varies, but the algorithms follow similar architectures. The improvement in the data quality can be obtained without any new mission and in a much faster manner. It is, therefore, a profitable and useful technique to be used in space robotics applications. As an alternative to SISR, another method can be used to achieve SR using multiple frames from the same image. This method is called Multi-Frame Super Resolution (MFSR), and it fuses the information available in each low resolution frame of a scene to generate a super resolution image. In [18], an example of this method applied to remote sensing is presented. A deep neural network is trained to perform the low-to-high resolution mapping from several remote sensing captures of the same image.

This work[1] aims to solve the lack of image quality problem by using the DL architecture *Lunar HighRes-net*. This network is based on the MFSR architecture of *HighRes-net* [13] and it has been applied to the data collected through NASA's LRO mission. The selection of this methodology was based on the study of its equivalent on Earth [10]–[12] that shows that this type of solutions provides a viable alternative to the launch on new and costly missions. Due to the importance of precision, while performing path planning for the missions, a multi-frame fusion method is employed instead of a single image one. The use of single captures of a particular region can lead to missing details as environmental conditions are decisive when gathering the information. The position of the sun, for example, can lead to shadow a region with rocks or smaller craters that won't appear in the image, potentially causing problems when navigating these areas during the mission. However, the use of multi-frame captures of a specific area of the surface will provide more details that will be taken into account while performing SR, as the images won't be captured under the same conditions.

The rest of the paper is structured as follows: Section II presents the different databases created for this work, including the source and the method used to obtain image datasets suited for the neural network training. Section III describes *Lunar HighRes-net*, the network architecture employed for the lunar surface MFSR. In Section IV, the results of this work are shown, both images and the metric used for the evaluation of the performance of the network. Lastly, in Section V the conclusions of this work are stated, and future lines of work are presented.

II. Moon's Surface Images

A. NASA's Lunar Reconnaissance Orbiter

NASA's LRO [7] mission started in 2009 with the objective of mapping the Moon to properly plan future space missions, providing information about craters, landing sites and regions of interest. Two stereo Narrow-Angle Cameras (NAC) and a Wide Angle Camera (WAC) performed remote sensing, capturing images of every sector of the Moon and providing three-dimensional information of the lunar surface thanks to the stereo image pair. After several years of orbiting the Moon, the mission was able to almost complete (98.2%) the map of the whole lunar surface. The regions that remain unmapped are mostly permanently shadowed areas within deep craters.

B. SpaceR's Virtual Moon

In the computer vision field, synthetic datasets have been becoming popular in the last few years. Artificial image datasets present an alternative to complete image collections for the training of neural networks. In both [23] [24], examples of the use of Unreal engine for the creation of virtual worlds and training scenarios are presented. In fields such as space, real images are hard to obtain, as the instrumentation used for this purpose is not easily accessible, and artificial datasets allow to have well-rounded databases. In [25], a virtual environment with lunar rocks was developed. This dataset presents an image collection suited for lunar surface image segmentation. Another example of these types of space synthetic data collection is introduced in [26], where a photorealistic simulator was created to train deep learning solutions for in orbit spacecraft pose estimation.

In order to improve the amount of data that would be used to train the SR neural network, we created an extra collection of images generated from a virtual model of the Moon designed using the graphic engine *Unreal Engine 4* [19]. As the software provides several tools for the alteration of the environment, the surface of the Virtual Moon can be modified as needed. They allow the possibility of adding and removing details to the scene whenever it is needed, which allows obtaining a complete and well-balanced details-wise dataset. Also, in *Unreal Engine 4* aerial images can be taken from the virtual lunar surface with different resolutions emulating a space remote sensing mission. Therefore, it can be used as ground truth to evaluate how the MFSR approach works. Some examples of both SpaceR's

[1]A video with a summary of our work was published in: https://youtu.be/3SA7qDRDJxU

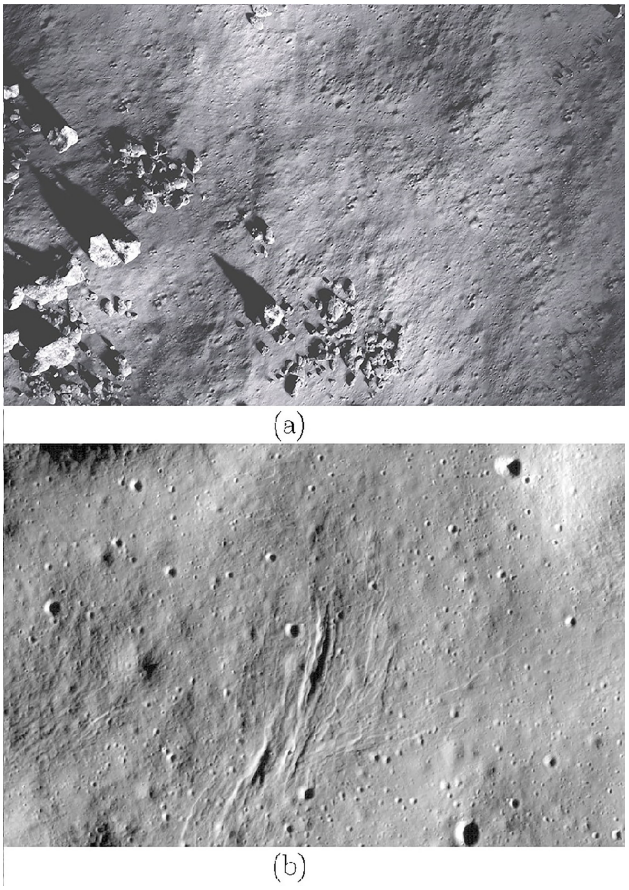

Fig. 2. Examples of the images used to generate the datasets of this work. (a) SpaceR's Virtual Moon, (b) NASA's LRO.

Virtual Moon and NASA's LRO images are shown in Fig. 2. Additionally, a diagram of the database creation process done for this work can be seen in Fig. 3. For the creation of the synthetic image dataset, a script was developed to emulate the flight of a satellite taking captures of the Virtual Moon's surface. After obtaining 100 images, the illumination conditions were changed and the script was launched again. This process was repeated four times and every gathered image was then downgraded and sliced to obtain the required patches and LR-HR sets for the training of the neural network.

C. Esa's Proba-V

The datasets used in this work follow the structure of the one used in [13], where an MFSR was developed to enhance the satellite images of the Earth of ESA's PROBA-V database [14]. PROBA-V was obtained by taking captures of the Earth's surface with an orbiting satellite. To every region of interest present in the dataset correspond at least 9 low resolution and one high resolution images. During this mission, a satellite captured low resolution (300 m/px) and high resolution (100 m/px) images from the surface of the Earth. Thus, each region captured in the database has one high and several low resolution images of the area due to the capture frequency set for the on-board sensors. Those images are also accompanied by metadata which provides

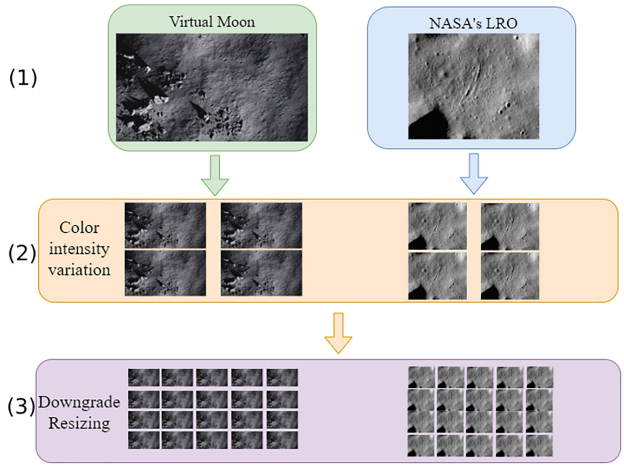

Fig. 3. Database creation diagram. The process can be described as: 1. Image acquisition (from LRO database and direct captures from the Virtual Moon), 2. color intensity variation to emulate different lighting conditions and 3. different downgrade methods applied for the resizing to obtain LR image sets for each scene.

much more details of every region, as the atmospheric and environmental conditions in which the low resolution images are taken are different.

The images from ESA's PROBA-V missions were taken from 74 different regions of the Earth at different points in time by the satellites of the said mission. The database is comprised of 1450 scenes, which are split into 1160 scenes for training and 290 scenes for testing. Each scene is represented by one grayscale HR image, with a 100 meters per pixel resolution and 9 to 20 grayscale LR images, with a resolution of 300 meters per pixel. The size of the images is 384 $\times$ 384 pixels for the HR ones and 128 $\times$ 128 pixels for the rest. Also, each image has associated a clearance map that indicates the area of the image that should not be processed by the neural network, as it is covered by clouds.

D. Datasets Adaptation

In this work, we prepared an equivalent to ESA's PROBA-V dataset using NASA's LRO and Virtual Moon images [2] for MFSR. In the case of the Virtual Moon, the illumination settings of the environment where modified to obtain 5 captures of the same scene with different illumination. In the case of LRO's images, the intensity of the pixel was modified to simulate the illumination change. Then, the same procedure was followed for both types of images. First, the original images were sliced in patches of 384x384 pixels, the considered ground truth. Then, the high resolution images were downscaled to generate low resolution (128 $\times$ 128 pixels) versions of them. By applying different interpolation methods, such as bicubic, bilinear and nearest neighbour, among others, every region of the lunar surface captured selected for the training of the network would have a ground truth and 20 low resolution images. This downgrade process allowed to perform MFSR on both lunar surface dataset.

[2]Both of the datasets utilized in this work are published and available to the public in: https://wwwfr.uni.lu/snt/research/spacer/datasets_tools, in their respective sections *LRO datasets* and *VirtualMoon.*

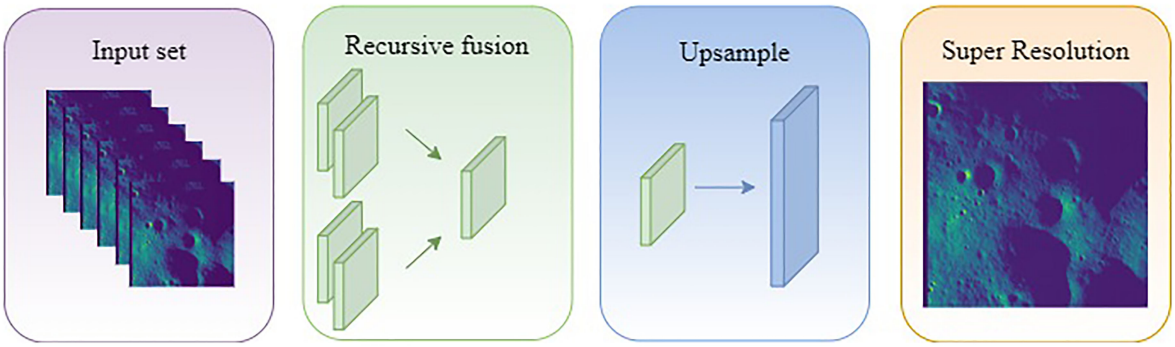

Fig. 4. Lunar HighRes-net MFSR diagram adapted from [13]. A set of multiple images of the same scenes are used as input of the network. These images are fused recursively in pairs in the embedding layer of the network. Finally, upsampling is performed and the SR image is generated.

Following the structure from ESA's PROBA-V, the 147 sets of images gathered from the Virtual Moon and the 60 sets from NASA's LRO are comprised by one HR grayscale image with a size of 384 $\times$ 384 pixels and 20 artificially downgraded LR grayscale images with a size of 128 $\times$ 128 pixels. The sets are divided into 135 for training and 12 for testing for the Virtual Moon database, and 54 for training and 6 for testing for LRO's database. The artificial downgrade was done over the HR images to be able to compare the results of the neural network when enhancing the resolution of the images x3. This downgrade was done rescaling the images using traditional CV interpolations and varying the color intensity of the images by 1%. This process emulates capturing images from the surface of the Moon at different points in time and with a variety of boundary conditions.

III. MFSR Method

A. Lunar HighRes-Net

Image super resolution methods [15] enhance the resolution of a single or multiple images using either traditional CV algorithms or DL techniques. In other words, it increases the perceptual quality and the number of details and features that appear in one captured scene. Even though the most widespread technique is SISR, in some cases, MFSR presents a good alternative to enhancing images with a lesser amount of details in the scene, like for instance, space images such as lunar surface pictures. The main assumption of MFSR is that a set of several views collectively contain more details than any single image of the same scene. The work presented in this paper focuses on the utilization of MFSR with lunar surface satellite imagery, as the enhancement of the images will provide higher perceptual quality and resolution. The current satellite images of the lunar surface or, in general, any celestial body often present a lack of precision in details about the surface characteristics, and therefore it can lead to imprecise planning of the space mission.

The neural network presented in this paper, *Lunar HighRes-net* follows the same architecture of [13]. It introduces *Highres-net*, a model to apply MFSR to remote sensing images of Earth inside a single spectral band (grayscale). In general, the network follows an encoder-decoder architecture that can be trained by using multiple images from the same scene. The input sets consist of a ground truth image with several lower resolution versions of it. This type of architectures works by fusing the information available in the low resolution images of a scene. The embedding layer of the network consists of a convolutional layer followed by two PReLU activations. The embedded states are fused recursively in pairs, reducing by a half its number each time. The final fused state contains the information of every view, and it is encoded into a HR state. Finally, this state is decoded into the resulting SR image. *HighRes-net* is paired with another network called *Shift-net* that would perform sub-pixel translations in order to align the ground truth and the super-resolution generated image. This operation ends up maximizing their similarities thanks to the sub-pixels shifts, thereby improving the quality of the results of the image enhancement. Fig. 4 shows the Lunar HighRes-Net architecture diagram of the process using the proposed network for MFSR.

The architecture of *HighRes-net* [13] included a clearance map to train the network to prevent clouds for interfering with the process. Thus, the network's recursive fusion layers would consider the clouds as a region with no information, while obtaining features and details from any of the others low-resolution captures of the same scene. In this work, we have modified this network architecture to avoid having this clearance map in the process and maximize how the dataset is exploited. Apart from this variation, *Lunar HighRes-net* follows the same pipeline. The input and output images of the network are 128 $\times$ 128 and 384 $\times$ 384 pixels respectively. Thus, the network perform a three times resolution scaling over the low resolution image set to produce the SR image. The datasets employed in the training and testing task have followed the same structure as the one used ESA's PROBA-V database originally, as it was stated in Section II.

B. Validation Metric - PSNR

For the evaluation of this work, the metric Peak Signal-to-Noise Ratio (PSNR) was used. It is the most used metric for SR evaluation. It defines the ratio between the maximum power of a signal and the noise that affects the fidelity of its representation. When applied to images, it can be easily defined via the Mean Squared Error (MSE). It can be described also as an objective metric to measure the quality of the reconstruction of a lossy

transformation. Given a monochrome image I and its noisy approximation N, the PSNR can be stated as:

$$MSE = \frac{1}{mn}\sum_{i=0}^{m-1}\sum_{j=0}^{n-1}[I(i,j) - N(i,j)]^2 \quad (1)$$

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

where i and j are pixel's coordinates and MAX the maximum intensity pixel value of the image I.

This metric has been widely used in the state of the art for SR for a long time. In most cases, the average dB values of PSNR obtained by the latest research works are in the range 30-40 dB.

IV. Results

In this section, we validate the usage of the *Lunar HighRes-net* neural network for obtaining SR images of the lunar surface. To evaluate the performance of the network, the widely used metric Peak Signal-to-Noise Ratio is used, following most of the SR works in the literature. Additionally, three different training input for the network are presented in the section to perform a comparison among them. Lastly, the results produced by the network are introduced along with a comparison to other commonly used image SR methods.

A. Training and Transfer Learning

The training of the network to achieve SR on lunar surface images was divided into two main steps. First, a base training was done with ESA's PROBA-V as input. Then, while having the weights of this base training, a transfer learning approach was used to prepare the network to work with lunar images. Three different trainings were done to test the performance of the network using different inputs. The different trainings done over the network are:

1) PROBA-V: ESA's database was the only one used as an input for the training. The result trained network is equivalent to the original *HighRes-net*.
2) PROBA-V + Virtual Moon: This case includes the previous training, while performing transfer learning and training the network with additional image datasets from the Virtual Moon database.
3) PROBA-V + Virtual Moon + LRO: Finally, this case includes transfer learning from the previous one as well with another training with the extra input from the adapted NASA's LRO database.

The core training of the network with PROBA-V's database was done over 400 epochs, with a batch size of 8. This size was selected to be able to perform the training on a *Nvidia's RTX 2080* graphic card. Default hyperparameters for the ADAM optimizer and the same learning rates showed in [13]. For the training after the transfer learning with both Virtual Moon and LRO's databases, 100 epochs were selected, keeping the batch size and hyperparameters equal to the first part of the training.

The average values of PSNR, in dB, obtained through the training with the three different inputs were analyzed to look for the best results for both Virtual Moon and LRO databases, concluding that the combination of the three presented datasets was key to a multipurpose solution for both image collections.

TABLE I
AVERAGE PSNR TRAINING RESULTS WITH DIFFERENT INPUT DATASETS (IN DB)

Training input	Virtual Moon	LRO
PROBA-V	31	23
PROBA-V + VM	38	28
PROBA-V + VM + LRO	38	31

TABLE II
PSNR VALUES COMPARISON BETWEEN DIFFERENT SR METHODS (IN DB)

Img	Bicubic	FSRCNN	ESPCN	ESDR	L-HRN
L01	23.89	24.28	24.25	24.29	31.61
L02	26.45	26.79	27.03	27.24	30.76
L03	25.08	25.63	25.60	25.91	28.11
L04	27.67	28.38	28.37	28.61	31.42
L05	26.02	26.69	26.68	27.21	28.98
V00	31.80	33.04	32.94	33.29	47.81
V01	26.53	27.84	27.94	28.15	39.81
V02	32.61	35.99	35.23	35.73	39.38
V03	26.51	27.05	27.16	27.19	39.11
V04	30.72	33.36	33.33	34.19	38.31
V05	29.32	32.48	32.48	33.01	41.88
V06	27.06	28.25	28.13	28.69	40.09
V07	31.06	34.08	34.20	35.41	41.98
V08	29.53	31.52	31.52	32.43	39.55
V09	32.31	35.23	35.26	36.34	43.53
V10	28.24	30.33	30.26	31.38	39.76
V11	26.75	27.40	27.35	27.55	36.72
Avg.	28.33	29.90	29.87	30.39	**37.58**

Lunar HighRes-net trained with images from the three databases presented in this work have been proven to provide the best results, as it can be seen in Table I. The addition of LRO's images to the training does not increase the SR performance of the network with Virtual Moon's images. However, it can be seen how the image enhancement of LRO does improve when including every type of dataset. It can also be seen how the *Lunar HighRes-net* network perform better with lunar surface images and a complete training than the base *HighRes-net*.

B. MFSR Results

For the experiments, the test image set defined previously in the databases was used. They consist of the comparison between the high resolution ground truth and the super resolution image generated using the presented network. *Lunar HighRes-net* has been evaluated using the images from both the Virtual Moon and NASA's LRO defined in Section II with the PSNR metric. These results have been compared with other SR techniques to validate that *Lunar HighRes-net* is a valid approach to lunar surface image enhancement, and it achieves state of the art scores. The results of enhancing the input dataset resolution x3 with the different solutions are presented in the Table II. We used the well-known open-source computer vision library *OpenCV* as it has built-in methods that use networks for image SR that were evaluated in lunar imagery for reference. For the comparison, image scaling with bicubic interpolation was used to have a reference of a SR traditional computer vision method. Additionally, three machine learning methods were used: *EDSR* [20], *ESPCN* [21] and *FSRCNN* [22]. These networks have been

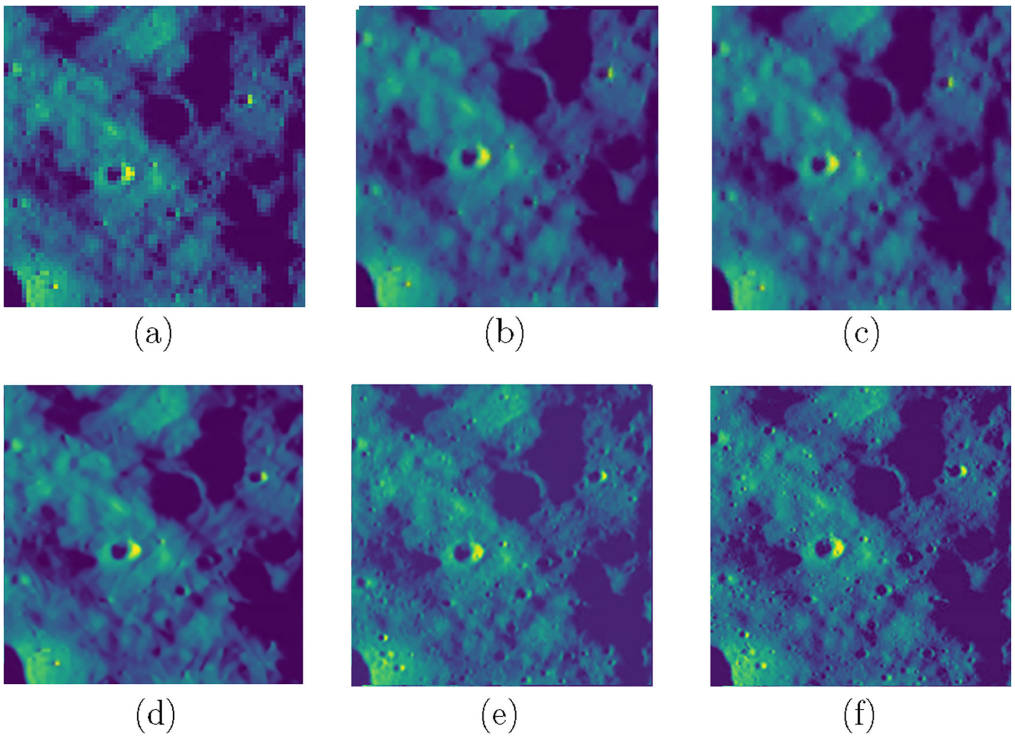

Fig. 5. Full image comparison between the different SR methods used in this paper. (a) Bicubic, (b) ESPCN, (c) FSRCNN, (d) EDSR, (e) Lunar HighRes-net, (f) Ground truth. The size of the figure makes it possible to better appreciate the difference between methods.

proved to provide high detailed SR image solutions for many applications, such as botanic and zoological imagery. Due to the wide variety of SR applications where these architectures can be utilized, they were selected for the comparison in this work. These implementations were used on the lunar surface imagery, and the results were evaluated comparing the SR version of the scene with the ground truth also with the PSNR metric. Fig. 5 presents an example of the results obtained with the different SR algorithms.

The quality of the generated SR images compared with the ground truth of the test dataset is evaluated with the PSNR metric in Table II. The results obtained after performing SR on lunar images from both LRO and the Virtual Moon datasets with the different methods are displayed. Bicubic interpolation is one of the most basic traditional computer vision algorithms of resolution enhancement, and it is compared as reference with the rest of the methods. FSRCNN [22], ESPCN [21] and EDSR [20] are selected as generic methods of deep learning SR for comparison. Finally, the Lunar HighRes-net column shows the results of this work, after the training with the combined remote sensing datasets previously mentioned. The results of Table II show that *Lunar HighRes-net* provides the best results according to the PSNR when compared to the other methods. Furthermore, a perceptual quality increase can be seen in the output of the network. Some examples of the output of the network can be seen in Fig 6 along with one of the input images and the ground truth. High level features present in these images, such as rocks or craters, show the easiest to notice resolution enhancement achieved with *Lunar HighRes-net*. Even though

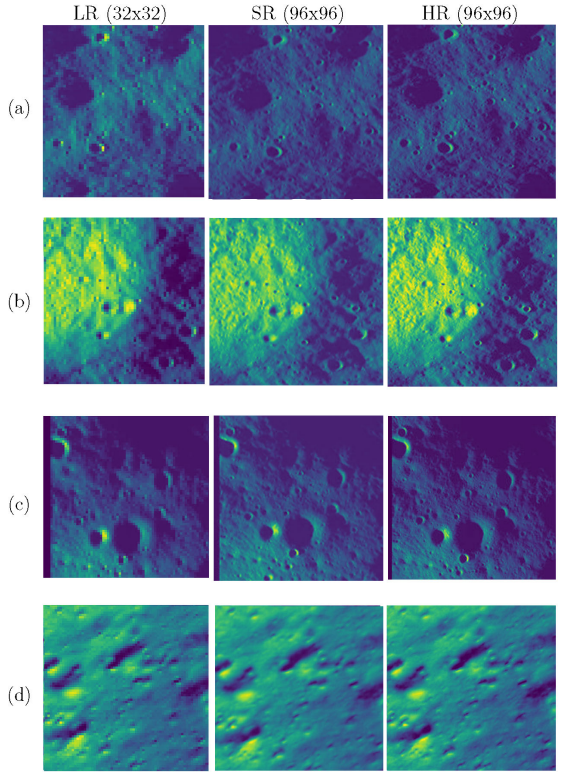

Fig. 6. Results from using Lunar HighRes-net on different images of the databases created for this work. (a), (b), (c) LRO images, (d) Virtual Moon image. The following link provides a folder with real size images: https://wwwfr.uni.lu/snt/research/spacer/datasets_tools, in the corresponding section of *Lunar High-ResNet results.*

some of the SR images from the Virtual Moon display a high PSNR due to the presence of small shadowed regions in the scene, it can also be seen that the results of the x3 enhancement achieve similar results to other SR network architectures in the field of image improvement. As it was mentioned the average values obtained in the process are similar to the state of the art in image SR.

V. Conclusion

In this work, we presented *Lunar HighRes-net*, a deep-learning based Multi-Frame Super Resolution method to enhance images of the Moon's surface. We also introduced and shared two databases that have been created for the training of the presented network, one adapted from NASA's Lunar Reconnaissance Orbiter mission imagery and a second one obtained from a virtual Moon developed in *Unreal Engine 4*. These databases are composed of image sets both artificial and real of lunar surface images. Each set contains one high resolution image (384×384 pixels) and 20 corresponding low resolution images (128×128 pixels) of the same scene. A transfer learning approach was performed from ESA's PROBA-V database training. Training results evidence that the addition of the new datasets created for this work to the training process of the network improve the performance of the network. The results obtained with *Lunar HighRes-net* achieves state of the art performance (35-40 dB) according to the results evaluated with the well-known Peak Signal-to-Noise Ratio for both image datasets. It also presents a significant quality improvement over other well known deep learning super-resolution approaches. Future work in this research line will focus on fine-tuning performed over the network *Lunar HighRes-net*. Additionally, new deep-learning architectures based on Generative Adversarial Networks (GAN) will be explored for space images super-resolution. These type of networks achieve better results in some cases of the state of the art of super-resolution when compared with convolutional architectures. Lastly, it will be studied the estimation of 3-dimensional information from super-resolution stereo image pairs. This will allow to increase the precision of the elevation maps of the lunar surface and therefore aid in preparing new space robotic missions on the Moon.

References

[1] M. Smith *et al.*, "The artemis program: An overview of NASA's activities to return humans to the Moon," in *Proc. IEEE Aerosp. Conf.*, Big Sky, MT, USA, 2020, pp. 1–10.

[2] K. H. Williford *et al.*, "The NASA Mars 2020 rover mission and the search for extraterrestrial life," *From Habitability Life Mars*, Elsevier, 2018, pp. 275–308.

[3] J. Vago *et al.*, "ESA ExoMars program: The next step in exploring Mars." *Sol. System Res.*, vol. 49, no. 7, 2015, pp. 518–528.

[4] T. Hoshino *et al.*, "Lunar polar exploration mission for water prospection-JAXA's current status of joint study with ISRO," *Acta Astronautica*, vol. 176, pp. 52–58, 2020.

[5] J. N. Maki *et al.*, "The mars 2020 engineering cameras and microphone on the perseverance rover: A next-generation imaging system for mars exploration," *Space Sci. Rev.* vol. 216, no. 8, pp. 1–48, 2020.

[6] A. Colaprete *et al.*, "An overview of the volatiles investigating polar exploration rover (VIPER) mission," in *Proc. AGU Fall Meeting Abstr.*, 2019, pp. P34B–03.

[7] G. Chin *et al.*, "Lunar reconnaissance orbiter overview: Theáinstrument suite and mission," *Space Sci. Rev.*, vol. 129, no. 4, pp. 391–419, 2007.

[8] J. N. Goswami, and M. Annadurai, "Chandrayaan-2 mission," in *Proc. Lunar Planet. Sci. Conf.*, no. 1608, p. 2042, 2011.

[9] R. W. Zurek, and S. E. Smrekar, "An overview of the Mars reconnaissance orbiter (MRO) science mission," *J. Geophysical Res.: Planets*, vol. 112, no. E5, 2007.

[10] J. M. Haut, R. Fernandez-Beltran, M. E. Paoletti, J. Plaza, and A. Plaza, "Remote sensing image superresolution using deep residual channel attention," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 11, pp. 9277–9289, Nov. 2019.

[11] K. Jiang, Z. Wang, P. Yi, J. Jiang, J. Xiao, and Y. Yao, "Deep distillation recursive network for remote sensing imagery super-resolution," *Remote Sens.*, vol. 10, no. 11, p. 1700, 2018.

[12] Y. Zhang *et al.*, "Pqa-cnn: Towards perceptual quality assured single-image super-resolution in remote sensing," in *Proc. IEEE/ACM 28th Int. Symp. Qual. Serv.*, 2020, pp. 1–10.

[13] M. Deudon *et al.*, "HighRes-Net: Recursive fusion for multi-frame super-resolution of satellite imagery," 2020, *arXiv:2002.06460*.

[14] M. Francois, S. Santandrea, K. Mellab, D. Vrancken, and J. Versluys, "The PROBA-V mission: The space segment," *Int. J. Remote Sens.*, vol. 35, no. 7, pp. 2548–2564, 2014.

[15] S. C. Park, M. K. Park and M. G. Kang, "Super-resolution image reconstruction: A technical overview," *IEEE Signal Process. Mag.*, vol. 20, no. 3, pp. 21–36, May 2003.

[16] H. Greenspan, "Super-resolution in medical imaging," *Comput. J.*, vol. 52, no. 1, pp. 43–63, 2009.

[17] R. Fernandez-Beltran, P. Latorre-Carmona and F. Pla, "Single-frame super-resolution in remote sensing: A practical overview," *Int. J. Remote Sens.*, vol. 38, no. 1, pp. 314–354, 2017.

[18] M. Kawulok, P. Benecki, S. Piechaczek, K. Hrynczenko, D. Kostrzewa and J. Nalepa, "Deep learning for multiple-image super-resolution," *IEEE Geosci. Remote Sens. Lett.*, 17, no. 6, pp. 1062–1066, Jun. 2020.

[19] Unreal Engine 4, [Online]. Available: https://www.unrealengine.com/

[20] B. Lim, S. Son, H. Kim, S. Nah and K. M. Lee, "Enhanced deep residual networks for single image super-resolution," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2017, pp. 136–144.

[21] W. Shi *et al.*, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 1874–1883.

[22] C. Dong, C. C. Loy, and X. Tang, "Accelerating the super-resolution convolutional neural network," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 391–407.

[23] W. Qiu *et al.*, "Unrealcv: Virtual worlds for computer vision," in *Proc. 25th ACM Int. Conf. Multimedia*, 2017, pp. 1221–1224.

[24] A. Garcia-Garcia *et al.*, "The robotrix: An extremely photorealistic and very-large-scale indoor dataset of sequences with robot trajectories and interactions," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 6790–6797.

[25] R. Pessia and G. Ishigami, "Artificial Lunar Landscape Dataset," [Online]. Available: https://www.kaggle.com/romainpessia/artificial-lunar-rocky-landscape-dataset

[26] P. F. Proença and Y. Gao, "Deep learning for spacecraft pose estimation from photorealistic rendering," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 6007–6013.

III | Research Paper 2

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

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

Published in: IEEE Robotics and Automation Magazine (RAM)

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.

Keywords: Image Super-resolution, Space robotics, Remote sensing, Traverse planning.

DOI: 10.1109/MRA.2023.3276267

Scopus: 89%

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 [10], 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], [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

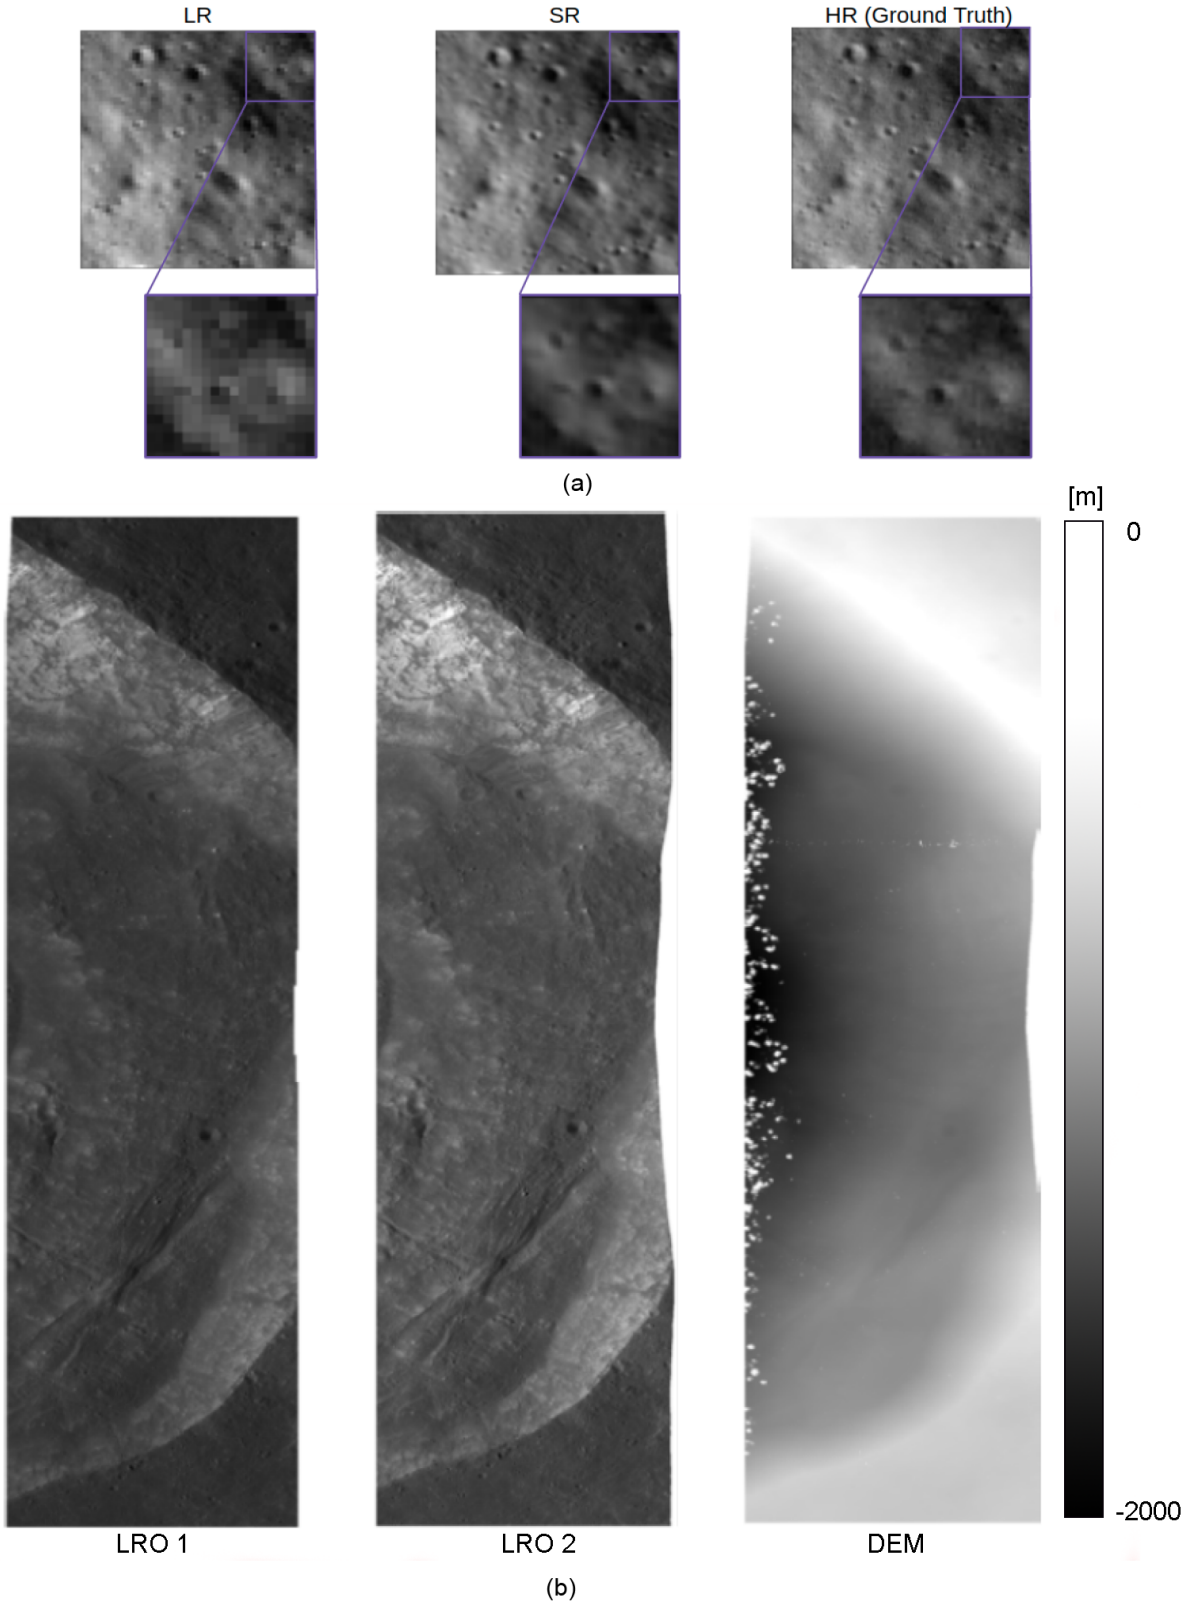

Fig. 1. (a) Direct comparison between lunar images with different resolutions: LR is Low resolution at 1m/px, SR is the product of our image enhancement Super-resolution process at 0.5m/px and HR the ground truth image also at 0.5 m/px. (b) Two map projected LRO images of the Kepler crater (latitude: 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.

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], [11] 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) [12]. 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 [13]. 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.

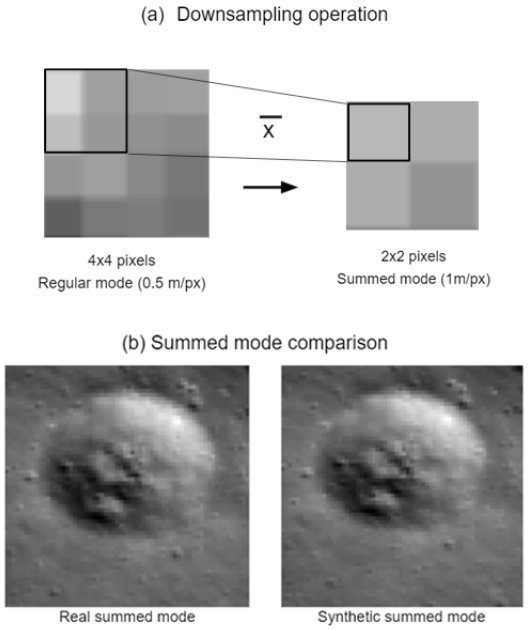

Fig. 2. Summed mode operation. (a) Mean of near neighbour pixels diagram. (b) Perceptual validation of our downsampling approximation. A real Regular-summed mode image pair that cover the same lunar surface area were used to validate our approach.

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 [14] 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 [16]. To do so, 24 instances of the GAN model were trained with different random seeds and initial weights, creating a deep

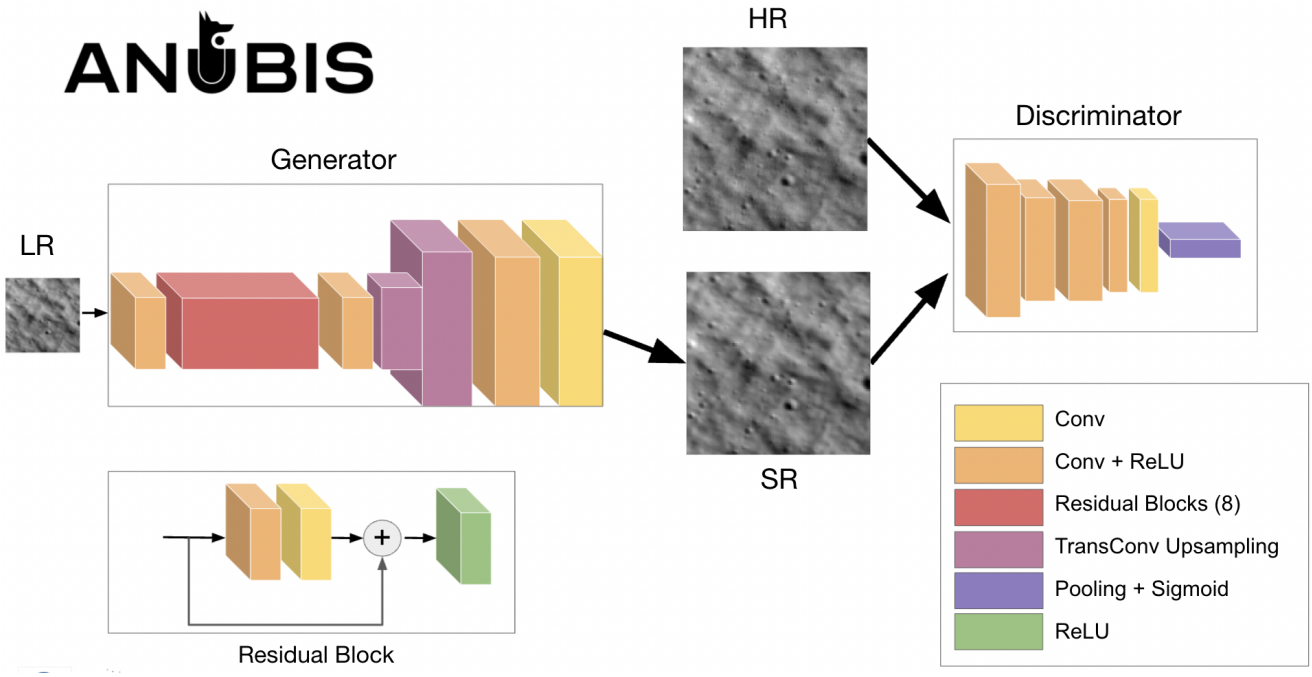

Fig. 3. The ANUBIS architecture consists of a ResNet-based generator and a CNN-based discriminator. Colors indicate the block types.

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.

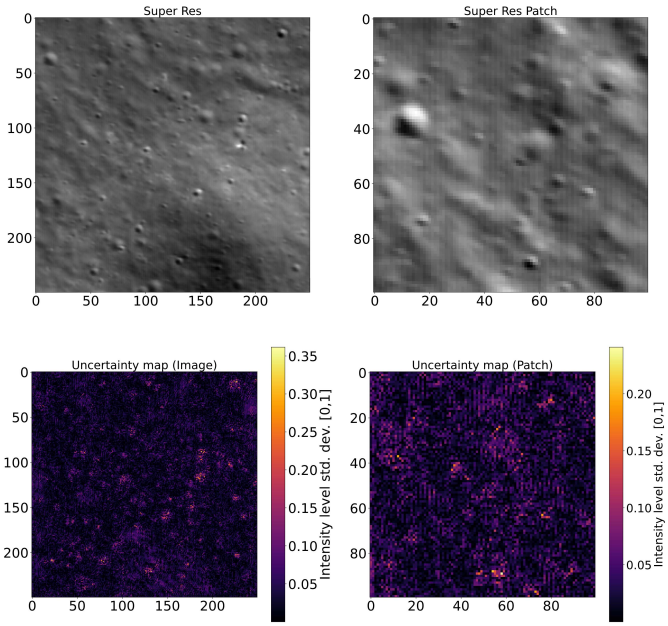

Fig. 4. Estimated uncertainty over the super-resolved image and its upper-left patch. For every pixel (of intensity range [0,1]), we compute the variance over the ensemble set of super-resolved images. The result is a heat map that is associated with the uncertainty of the posterior distribution about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

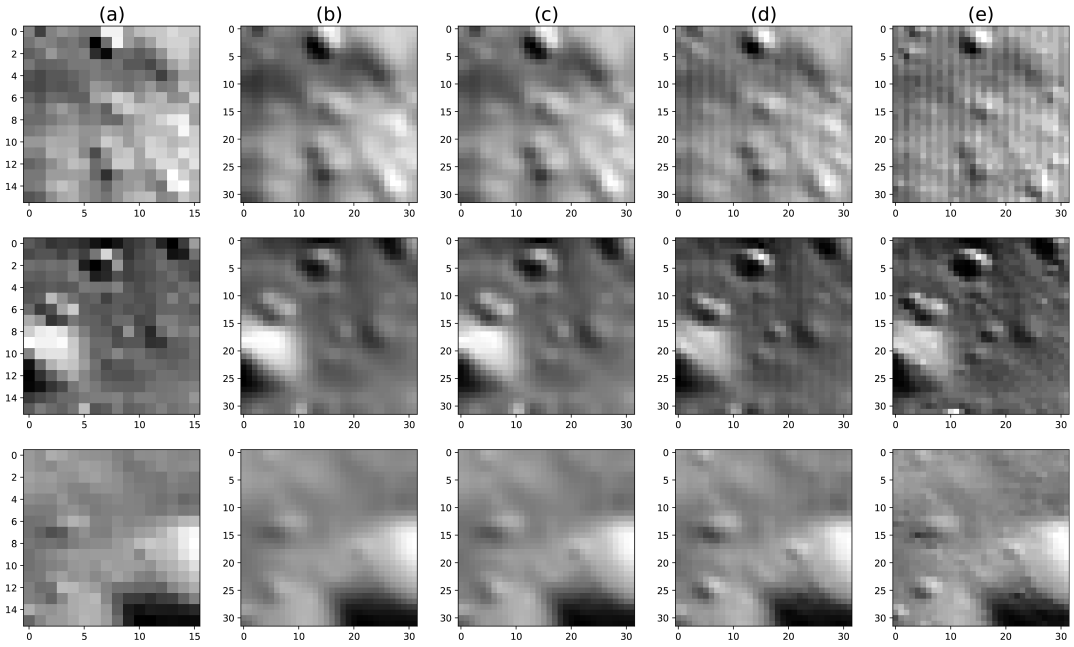

Fig. 5. Three examples of the lunar surface images super resolution comparison: (a) Low Resolution, (b) Bilinear, (c) Bicubic, (d) **ANUBIS (our work)** (e) Ground truth. The selected images were randomly taken out of the test set of images.

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) [18] and the Enhanced Deep Residual Networks for single-image super-resolution (EDSR) [19]. These DL-based techniques are two of the best performing networks metric wise according to one of the latest survey paper [20] 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. 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.

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 ± 0.97**
SSIM	0.785	0.83	0.86	0.86	**0.91 ± 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

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 %

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

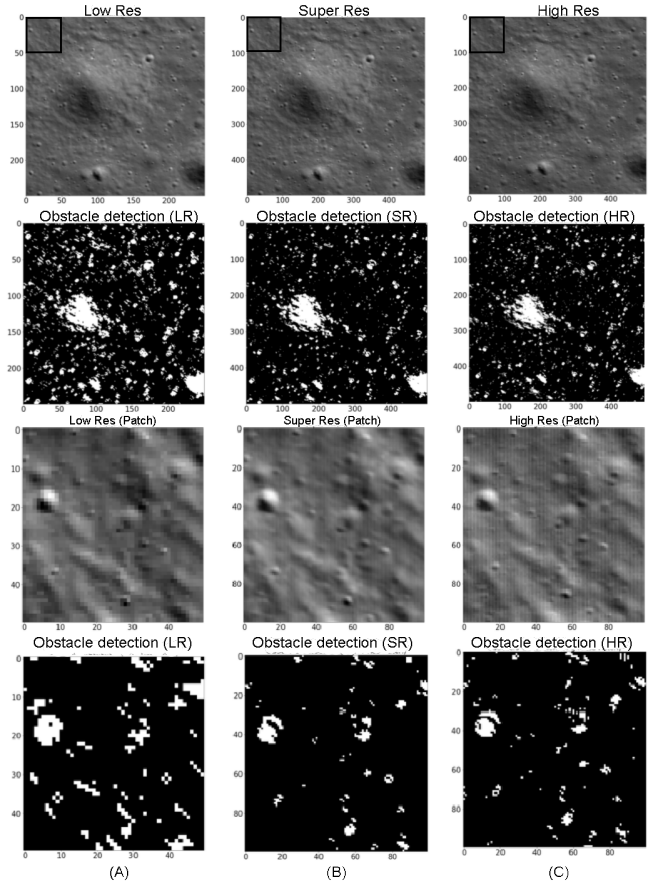

Fig. 6. Obstacle detection image and patch example. (A) Low resolution, (B) ANUBIS, (C) Ground truth. The super-resolved image (B) is the image average of the distribution generated by ANUBIS.

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.

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.

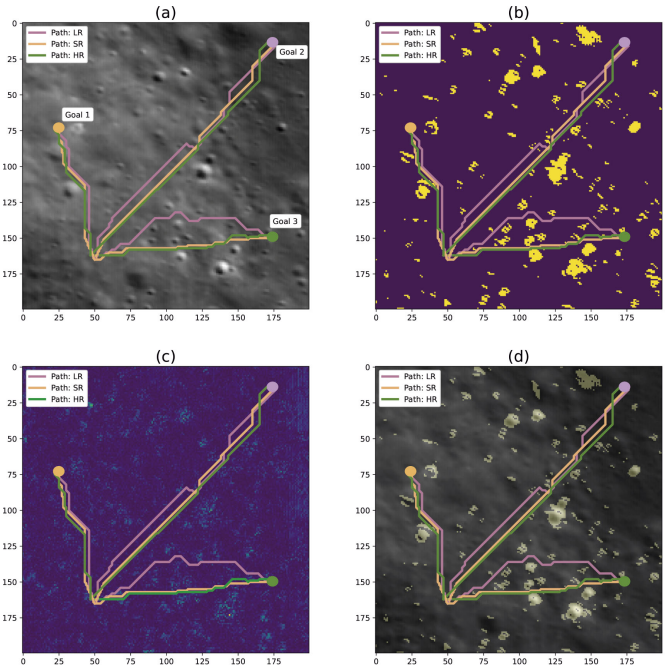

Fig. 7. Path planning downstream task example. 3 different trajectories are overlayed in (a) ANUBIS super-resolved image, (b) Obstacle mask (uncertainty obstacles included), (c) Uncertainty map, (d) ANUBIS super-resolved image + obstacle mask.

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 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

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

contained in it were too small (32×32 px) to generate and plan a relevant trajectory.

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) [17] 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 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

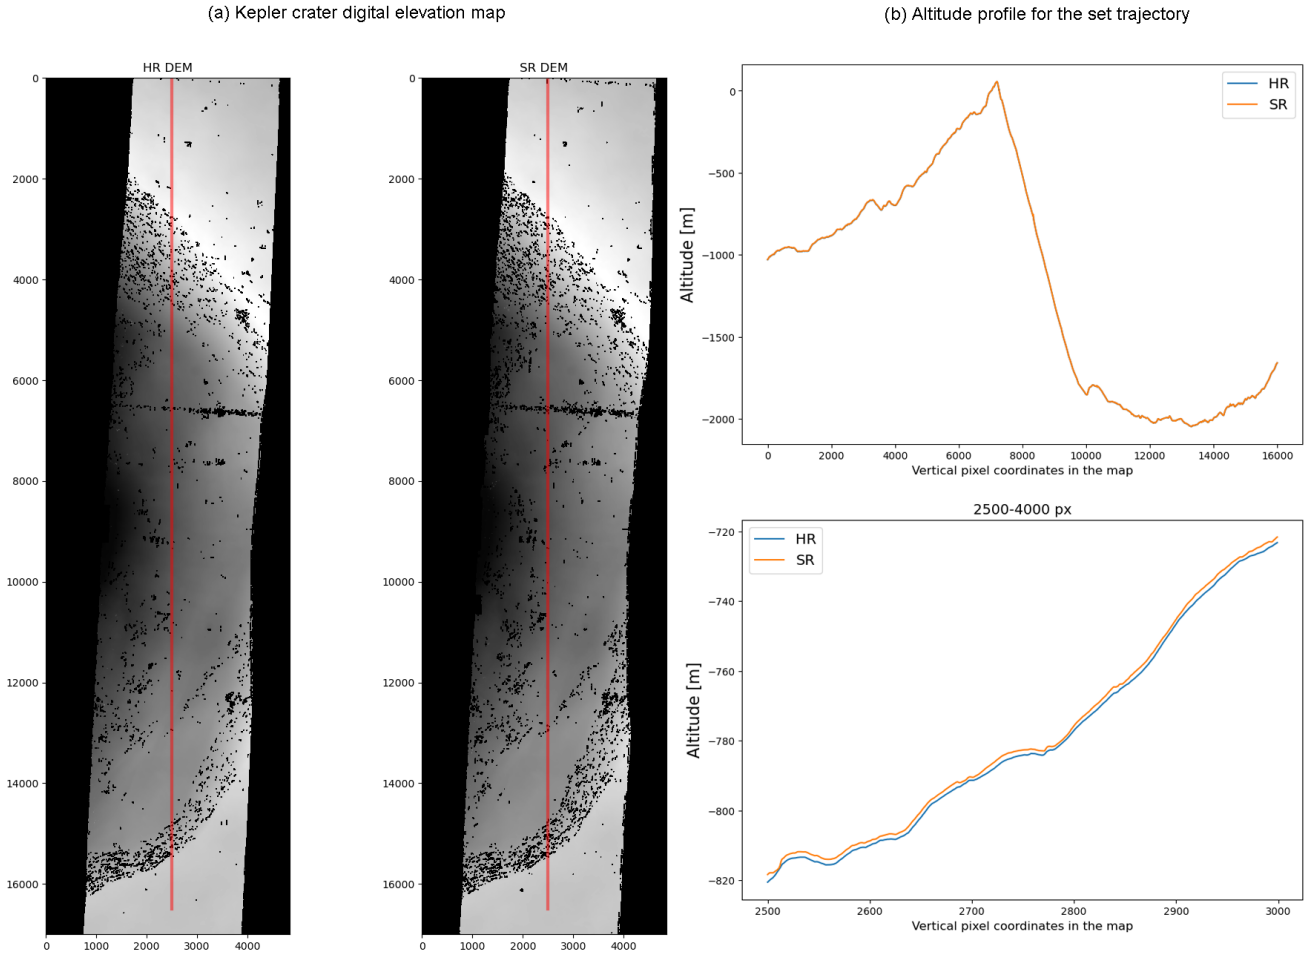

Fig. 8. (a) Digital elevation maps comparison of the Kepler crater between the high resolution ground truth (HR, left) and the super-resolved (SR, right) maps. The red line indicates the selected section to estimate the altitude profile. (b) Altitude profile of the selected region of the Kepler crater.

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

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.

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] R. Vondrak, J. Keller, G. Chin, J. Garvin, "Lunar Reconnaissance Orbiter (LRO): Observations for lunar exploration and science". In Space science reviews, 2010, pp. 7-22.

[11] 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.

[12] 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.

[13] 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.

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

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

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

[17] 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.

[18] 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.

[19] 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.

[20] 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.

IV | Research Paper 3

Full Title: Thermophysical Change Detection on the Moon with the Lunar Reconnaissance Orbiter Diviner sensor

Authors: S. Bucci, J.I. Delgado Centeno, B. Gaffinet, Z. Liang, V. Bickel, B. Moseley and M.A. Olivares-Mendez

Published in: Proceedings of the Neural Information Processing Systems: "Machine Learning and the Physical Sciences" (ML4PS)

Abstract:
The Moon is an archive of the history of the Solar System, as it has recorded and preserved physical events that have occurred over billions of years. NASA's Lunar Reconnaissance Orbiter (LRO) has been studying the lunar surface for more than 13 years, and its datasets contain valuable information about the evolution of the Moon. However, the vast amount and heterogeneous nature of data collected by LRO make the extraction of scientific insights very challenging - in the past most analyses relied on human review. Here, we present NEPHTHYS, an automated solution for discovering thermophysical changes on the surface using one of LRO's largest datasets: the thermal data collected by its Diviner instrument. Specifically, NEPHTHYS is able to perform systematic, efficient, and large-scale change detection of present-day impact craters on the surface. Further work could enable more comprehensive studies of lunar surface impact flux rates and surface evolution rates, providing critical new information for future missions.

Keywords: Thermal Images, Lunar Science

DOI: XXX
GGS: A++

Thermophysical Change Detection on the Moon with the Lunar Reconnaissance Orbiter Diviner sensor

Silvia Bucci*
Polytechnic of Turin, Italy
silvia.bucci@polito.it

Jose Ignacio Delgado Centeno*
University of Luxembourg
jose.delgado@uni.lu

Ben Gaffinet*
RSS-Hydro, Luxembourg
bgaffinet@rss-hydro.lu

Ziyi Liang*
University of Southern California
ziyilian@usc.edu

Valentin Bickel
ETH Zürich VAW
bickel@vaw.baug.ethz.ch

Ben Moseley
ETH Zürich AI Center
benjamin.moseley@ai.ethz.ch

Miguel Olivares-Mendez
University of Luxembourg
miguel.olivaresmendez@uni.lu

Abstract

The Moon is an archive of the history of the Solar System, as it has recorded and preserved physical events that have occurred over billions of years. NASA's Lunar Reconnaissance Orbiter (LRO) has been studying the lunar surface for more than 13 years, and its datasets contain valuable information about the evolution of the Moon. However, the vast amount and heterogeneous nature of data collected by LRO make the extraction of scientific insights very challenging - in the past most analyses relied on human review. Here, we present NEPHTHYS, an automated solution for discovering thermophysical changes on the surface using one of LRO's largest datasets: the thermal data collected by its Diviner instrument. Specifically, NEPHTHYS is able to perform systematic, efficient, and large-scale change detection of present-day impact craters on the surface. Further work could enable more comprehensive studies of lunar surface impact flux rates and surface evolution rates, providing critical new information for future missions.

1 Introduction

The nature, magnitude, and frequency of present-day changes of the lunar surface are of key importance to understand and reconstruct the Moon's and the Solar System's history, to analyze the surface evolution of airless planetary bodies and to identify potential environmental hazards for future robotic and crewed For the Moon, the most extensive datasets that allow for the study of the surface are provided by NASA's Lunar Reconnaissance Orbiter (LRO) mission, which has been operating since 2009. Two of LRO's most important instruments are its NAC (Narrow Angle Camera, Robinson et al. (2010)), which acquires high-resolution ($\sim$0.5 m footprint) optical images, and Diviner (Paige et al. (2010)), which acquires intermediate-resolution ($\sim$200 m footprint) point measurements of the surface temperature.

*Equal contribution

Machine Learning and the Physical Sciences workshop, NeurIPS 2022.

Despite the abundance of these global datasets with long temporal baselines, our knowledge about present-day changes is limited. Past work has been able to identify a small number ($\lesssim 1000$) of such changes - predominantly fresh, small impact craters (Speyerer et al. (2016), Watters et al. (2022), Williams et al. (2018)), but has faced three major challenges: 1) the massive size of the existing datasets (hundreds of terabytes) which makes human mapping impossible, 2) highly variable illumination conditions, and 3) complex noise in the dataset which makes faint and small change signals hard to detect. For example, the most extensive change detection study so far (Speyerer et al. (2016)) relied on the matching of multi-temporal optical NAC images acquired at similar local times, which limited their spatial coverage to only 6.6% of the lunar surface.

To overcome such limitations, we turn to machine learning to provide more automated, scalable and efficient change detection, which accurately detects low signal-to-noise events. In particular, we focus on detecting thermophysical changes within the Diviner dataset; such changes are usually caused by physical events such as small impacts but can have significant spatial thermal expressions and remain visible for hundreds of thousands of years (Williams et al. (2018)). Specifically, we focus on detecting fresh impact craters and their surrounding ejecta blankets, which have different night-time cooling properties than the regular lunar surface, producing so-called cold spots (Williams et al. (2018)) (Fig. 1). Fresh impact craters are believed to represent the largest source of lunar surface change.

We introduce NEPHTHYS (New Event Perceiving High-Trust High-Yield System), which searches for cold spots by using two different neural networks (a CNN-based binary classifier and a normalizing flow based anomaly detector) to analyse preprocessed thermal images constructed from the Diviner thermal point measurements. The output of the networks includes a detection confidence level to increase their reliability. Our workflow is validated by scanning areas with known fresh impacts and then verifying candidates by hand using NAC optical images (Williams et al. (2018)).

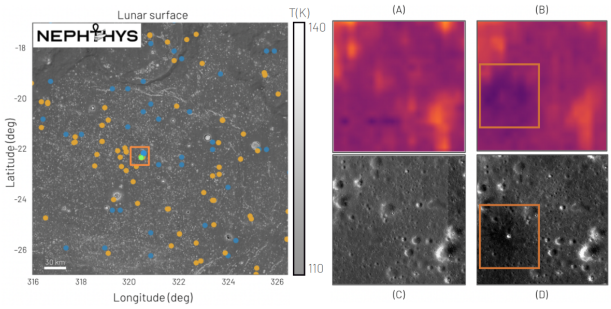

Figure 1: An example of impact detection with NEPHTHYS. Left: thermal Diviner mosaic (showing average night temperature) of a scanned region (280x280 km): orange and blue dots indicate potential candidates for fresh craters found by the binary classifier and anomaly detector, respectively. Green dot inside the bounding box indicates a known recent impact, the target of the search. Validation of the detection: (A),(B) our preprocessed thermal image *before* and *after* the impact, (C),(D) NAC optical image *before* and *after* the impact.

2 Methodology

Our high-level detection workflow consists of two main steps. First, we carry out preprocessing of the raw Diviner point measurements (See Appendix A for a detailed description of the Diviner dataset). Specifically, given an area of interest on the surface, we bin and aggregate the point measurements into two images, representing a *before* and *after* image of the surface temperature spanning the entire operational window of the instrument (Fig. 2). Secondly, we feed these images into two different detection networks, which independently predict whether an impact has occurred by searching for an associated cold spot in the images. Finally, these networks can be scanned over the entire lunar surface to provide an efficient and large-scale detection method.

2.1 Data Preprocessing Workflow

Following Moseley et al. (2020), we first sort and store all of the available Diviner point measurements into three-dimensional "data cubes", indexed by location (in 0.5° latitude and 0.5° longitude intervals) and local time (6-hour intervals), for computational retrieval efficiency. We then produce one *before* and one *after* image for each region of interest: *before* images are populated using the oldest available Diviner measurements over the respective location; *after* images are populated using the most recent Diviner observations over the respective location. Importantly, we currently only use night time point measurements acquired between 22:00 and 06:00 local time to avoid the complex changes in temperature associated with daytime illumination. All data over 2009-2022 is used. Populating the *before* and *after* images consists of two steps: a) *Data normalization* and b) *Binning and aggregation*. *Gaussian smoothing* is also required as third step before creating a *difference* image, used as the input of one of our ML solutions (Section 2.3).

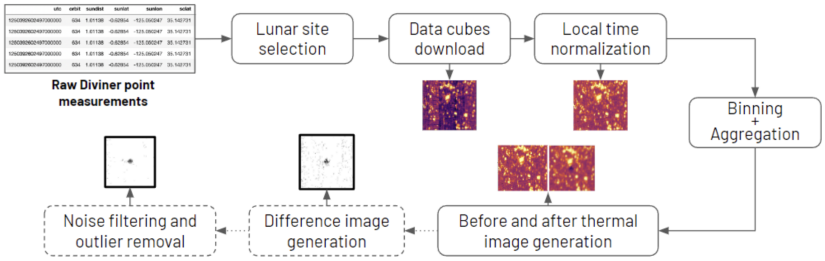

Figure 2: Data preprocessing workflow used by NEPHTHYS.

In Fig. 2 we show a raw temperature image (after Data cubes download). We note that heterogeneous stripes of different temperatures can be observed. This is because some point measurements are taken at an earlier local time (within the 22:00 - 06:00 window used) when the surface has not cooled down completely (brighter color). The darker regions instead correspond to measurements taken at a later local time. To fix this issue, we normalize the measurements to a reference time (precisely, 00:00 local lunar time) computed using Hayne's model (Hayne et al. (2017)), which simulates the background temperature variation with local time for a given latitude on the Moon. After this step, we obtain a more uniform image (Fig. 2, after Local time normalization).

In order to produce *before* and *after* images from the normalized point measurements, we need to go through the second preprocessing step: *Binning and aggregation*. For each image, we bin all the available point measurements onto 100×100 m grids. To create a *before* image we take the measurements from the earliest year available in each bin and for the *after* image, we take the measurements from the most recent year. Lastly we average the selected measurements in each bin to get the final image pairs. To create a *difference* image, we subtract the *after* image from the *before* image. We note that Diviner's pointing accuracy includes slight errors, occasionally resulting in a spatial misalignment of measurements, resulting in spatial noise in the *difference* images. Gaussian filtering (Chung (2020)) is employed to smooth out this noise after the binning step in both images. The last two pairs of images in Fig 2 show the effect of the Gaussian filtering.

2.2 Binary Classification

A straightforward solution to detect the presence of a fresh cold spot by analyzing the *before* and *after* images separately is to check for the presence of a cold spot in both - if the *before* image does not contain a cold spot but the *after* image does, then there was a recent change. We note there is a difference between *old* cold spots and *fresh* cold spots: the former refers to impacts that occurred before 2009 (the year the LRO diviner mission launched); the latter refers to impacts that happened after 2009 - those are actual present-day changes.

We use a Deep Neural Network (DNN) (He et al. (2016)) as a binary classifier to determine whether an image contains a cold spot or not. Our training data is a catalog of $\sim$ 2000 known, *old* cold spots as positive samples (taken from Williams et al. (2018)) and $\sim$ 2000 random points as negative samples taken over the entire lunar surface. We can safely use random points as negative samples given that the probability of picking a point with an old or fresh cold spot is negligibly small (Williams et al.

(2018)). The trained classifier can be deployed on the entire lunar surface: for each point we produce the *before* and *after* images, use the model to classify both, and evaluate for changes.

2.3 Anomaly Detection

Ideally, if there were no changes, the *before* and *after* images of the same location should be identical. However, minor measurement localization inaccuracies add noise (*background noise*). This is the reason why applying simple thresholding techniques on the *difference* images is often challenging for change detection. As an alternative approach, we propose to formulate the problem as an *anomaly detection* task (Hendrycks et al. (2018), Yang et al. (2021)), shown in Fig. 3. Given a *difference* image, we want to assess whether the image lies within the distribution of difference images where no fresh cold spots are present. Our solution is based on *DifferNet* Rudolph et al. (2021), a recent approach in which a Normalizing Flow (NF) Rezende and Mohamed (2015) is applied to the features extracted from a DNN to compute such an *anomaly score*.

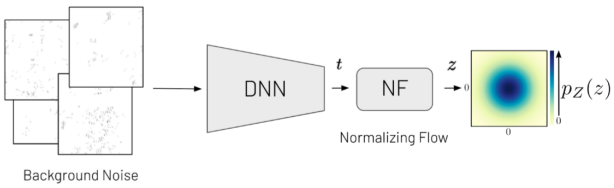

Figure 3: Overview of the architecture used for *anomaly detection* Rudolph et al. (2021)

Let X_T be the set of training data composed of the *difference* images of locations without fresh cold spots (the 2000 random locations used for the BC). For each $x_i \in X_T$, let t_i be the features extracted from the DNN in Fig. 3, and let z_i be the features extracted from the NF. Following Rudolph et al. (2021), we maximize the likelihood of the extracted features t which are quantifiable in the latent space Z. After a change-of-variables and the use of the negative log-likelihood (Rudolph et al. (2021)) obtain the following loss function to train the NF:

$$L(t_i) = \frac{\|z_i\|^2}{2} - \log \left| \det \frac{\partial z}{\partial t} \right| \tag{1}$$

Intuitively, we want that NF maps t as close as possible to $z = 0$ for all of the training examples. During evaluation we can compute the *anomaly score* through the negative log-likelihoods for each test sample $x_j \in X_{eval}$:

$$\tau(x_j) = E[-\log p_Z(z_j)] \tag{2}$$

An image is classified as anomalous if the anomaly score is above some threshold value, Θ.

3 Results and validation

To evaluate the quantitative performance of our binary classifier, we consider a validation set made of 400 images (half positive, half negative): it is able to achieve about 95% of accuracy on them. For the anomaly detector we consider the known fresh cold spots (enhanced with augmented versions of them) and random points: it assigns a score 4-times higher on average for the difference images that contain a cold spot on them. A downstream task for the validation of our approach can be seen in Fig. 1. Here we chose an area of 10x10 degrees in latitude and longitude that contains one known fresh impact that occurred in 2013 (Williams et al. (2018)). The two detectors identify a number of fresh cold spot candidates (orange dots: binary classifier, blue dots: anomaly detector). Importantly, both approaches were able to find the target of the experiment, the known recent impact (green dot). We repeated the same experiment for 4 additional sites with a-priori known fresh impacts: NEPHTHYS was able to consistently find every known fresh impact. Unfortunately, the locations of most existing known fresh impacts have not been made publicly available by LROC, and only these 5 were available from the Diviner team (Williams et al. (2018)) and recognizable in Diviner data.

4 Conclusion, limitations and future steps

We have implemented NEPHTHYS: an ML-based approach to detect present-day thermophysical changes on the lunar surface. A pre-processing pipeline was implemented to convert point measurements from the Diviner instrument into images suitable for ML-driven change detection. Two independent ML solutions were developed, which were able to efficiently scan a region of the Moon and produce a list of potential surface changes. Verifying the detected sites was one of the biggest limitations of our approach due to the lack of known impacts and significant human effort required to check the candidates with NAC imagery. In future work, we would like to submit target acquisition requests to the LROC team to further verify NEPHTHYS' candidates. We would also like to assess whether NEPHTHYS can detect other thermophysical changes, such as moonquakes and landslides. Newly discovered surface changes will provide key insights into the history of the Moon and the Solar System and contribute essential information to the planning of upcoming lunar missions.

Acknowledgement

This work has been enabled by the Frontier Development Lab Program (FDL USA). FDL is a collaboration between SETI Institute and Trillium Technologies Inc., in partnership with NASA, USGS, Google Cloud, NVIDIA, PLANET and the Luxembourg Space Agency (LSA).

5 Broader impact

The main goal of our work is enabling planetary science researchers to process and analyze NASA's Lunar Reconnaissance Orbiter (LRO) data in an automated manner. We believe that our work could have a great impact in the society. Helping in having successful lunar missions, our work contributes to the exploration of the universe whose goal is also to find other planets suitable for life. Considering the climate change and all the other problems that our planet is facing, a backup plan becomes urgent.

References

Chung, M. K. (2020). Gaussian kernel smoothing.

Hayne, P. O., Bandfield, J. L., Siegler, M. A., Vasavada, A. R., Ghent, R. R., Williams, J.-P., Greenhagen, B. T., Aharonson, O., Elder, C. M., Lucey, P. G., and Paige, D. A. (2017). Global regolith thermophysical properties of the moon from the diviner lunar radiometer experiment. *Journal of Geophysical Research: Planets*, 122(12):2371–2400.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.

Hendrycks, D., Mazeika, M., and Dietterich, T. (2018). Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606*.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2017). Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90.

Moseley, B., Bickel, V., Burelbach, J., and Relatores, N. (2020). Unsupervised learning for thermophysical analysis on the lunar surface. *The Planetary Science Journal*, 1(2):32.

Paige, D., Foote, M., Greenhagen, B., Schofield, J., Calcutt, S., Vasavada, A., Preston, D., Taylor, F., Allen, C., Snook, K., et al. (2010). The lunar reconnaissance orbiter diviner lunar radiometer experiment. *Space Science Reviews*, 150(1):125–160.

Rezende, D. and Mohamed, S. (2015). Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR.

Robinson, M., Brylow, S., Tschimmel, M., Humm, D., Lawrence, S., Thomas, P., Denevi, B., Bowman-Cisneros, E., Zerr, J., Ravine, M., et al. (2010). Lunar reconnaissance orbiter camera (lroc) instrument overview. *Space science reviews*, 150(1):81–124.

Rudolph, M., Wandt, B., and Rosenhahn, B. (2021). Same same but differnet: Semi-supervised defect detection with normalizing flows. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1907–1916.

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. (2015). Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252.

Speyerer, E. J., Povilaitis, R. Z., Robinson, M. S., Thomas, P. C., and Wagner, R. V. (2016). Quantifying crater production and regolith overturn on the moon with temporal imaging. *Nature*, 538(7624):215–218.

Watters, T., Speyerer, E., and Robinson, M. (2022). Recent landslides and their relation to young thrust fault scarps on the moon. *LPI Contributions*, 2678:1626.

Williams, J.-P., Bandfield, J., Paige, D., Powell, T., Greenhagen, B., Taylor, S., Hayne, P., Speyerer, E., Ghent, R., and Costello, E. (2018). Lunar cold spots and crater production on the moon. *Journal of Geophysical Research: Planets*, 123(9):2380–2392.

Yang, J., Wang, H., Feng, L., Yan, X., Zheng, H., Zhang, W., and Liu, Z. (2021). Semantically coherent out-of-distribution detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8301–8309.

Appendix

A LRO Diviner dataset

The LRO Diviner instrument (Paige et al. (2010)) is a passive radiometer on board NASA's LRO satellite which has been gathering point measurements of the lunar surface temperature along the ground track of the satellite over the past 13 years. Measurements are collected across nine different wavelength channels from optical to thermal wavelengths. Over multiple orbits, LRO Diviner has covered the entire surface of the Moon, providing multiple measurements at each surface location taken over the full range of local times (noon over midnight to noon). All collected data points are publicly available on NASA's PDS [2]. We use a processed version of the raw Diviner data that was pointing-corrected and calibrated (Level 1), specifically including channels 7, 8, and 9 (covering 12.5 to 200 μm), provided by the Diviner team (Paige et al. (2010)).

B Implementation details

Binary Classification. We build our approach over a standard ResNet-18 (He et al. (2016)) pretrained on ImageNet (Russakovsky et al. (2015)). We train our model for 2000 epochs using batch size 32, learning rate 0.001 and Adam optimizer. We use the standard cross-entropy loss to finetune only the last block of the network keeping the rest frozen with the ImageNet weights. As data preprocessing we first resize the images to 256, then we crop every image into 50x50 crops covering with 121 windows the entire image. Lastly, we resize every crop to 224 and feed the images into the network.

Anomaly Detection. We used as feature extractor AlexNet (Krizhevsky et al. (2017)) followed by the Normalizing Flow block as proposed in DifferNet (Rudolph et al. (2021)). We update the weights only for the Normalizing Flow block keeping the rest frozen with the ImageNet pretrained weights. We first resize the images to 448 and then crop the image to four 224 sub images with no overlap, then resize each sub image to 448. Considering the performances on the validation set, we set the threshold value Θ to 0.8.

For the training of both of the architectures described in the previous section, a NVIDIA Tesla V100 was used. The lightweight nature of the networks used made it possible to run multiple iterations of the training for optimization purposes in a short amount of time. For inference, a similar hardware setup will make it possible to scan specified lunar surface regions with the change detection algorithms.

[2] `https://pds-geosciences.wustl.edu/missions/lro/diviner.htm`

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
   (e) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (f) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (g) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
   (h) If your work uses existing assets, did you cite the creators? [Yes]
   (i) Did you mention the license of the assets? [N/A]
   (j) Did you include any new assets either in the supplemental material or as a URL? [N/A] No new assets
   (k) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] The NASA dataset is publicly available
   (l) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

V | Conclusions

1 Contributions

This thesis presents advancements in the use of machine learning to enhance and process lunar data with the goal of improving and extending the available information for upcoming lunar mission planning. Through the development of Lunar HighResnet, a deep learning-based Multi-Frame Super Resolution method, enhancement of images of the Moon's surface was done for the first time. The creation and utilization of two databases, one composed of imagery from NASA's Lunar Reconnaissance Orbiter mission and another obtained from a virtual Moon developed in Unreal Engine 4, were essential to training this network. The performance improvement seen through the incorporation of these databases in the network's training demonstrates the value of data-rich environments for such endeavors. ANUBIS, an innovative application of Generative Adversarial Networks (GANs) for single image super-resolution and uncertainty estimation of lunar satellite images was also introduced. The application is able to double the input lunar image resolution while providing uncertainty estimation of the process. This method outperforms all existing approaches and shows how Super Resolution can enhance lunar mission planning with current lunar remote sensing data. Notably, this development provides a pathway to improve the low-resolution imagery of the lunar south pole, a primary focus area of upcoming lunar missions. Lastly, the implementation of NEPHTHYS, an ML-based approach for detecting current thermophysical changes on the lunar surface, signifies an efficient and valuable tool for lunar surface monitoring. While verification of detected sites remains a challenge due to the lack of known impacts and the significant human effort required, the potential of this approach to contribute essential information to future lunar mission planning is clear.

Overall, these developments offer substantial contributions to the field of lunar research and exploration, specifically in the areas of lunar imagery enhancement, precision in lunar robotic mission planning, and the detection of lunar surface changes.

2 Limitations of the work

Despite the promising results demonstrated in this thesis, the path towards fully incorporating machine learning techniques into lunar science and future mission planning involves overcoming a series of challenges. One major limitation is the inherent complexity of working with lunar data. The raw data collected from lunar missions, although rich in information, requires a significant amount of preprocessing to make it suitable for ML applications. This preprocessing step involves various tasks such as denoising, data normalization, interpolation, and more, which can be time-consuming and technically challenging. The successful deployment of ML algorithms heavily depends on the quality of the input data, and hence, robust preprocessing mechanisms are essential. In future work, the development of more efficient preprocessing pipelines or even automated data cleansing approaches could further streamline this process.

Another key limitation lies in the demand for scientific validity and accuracy in the machine learning outcomes. While ML can provide insightful and efficient analyses, its results must be reliable and precise enough to be translated into actionable scientific conclusions. In other words, the ML models must not only enhance the data resolution but also preserve the scientific value of the data for it to be useful for upcoming lunar missions. Hence, the challenge here is striking the right balance between technical performance and scientific relevance. This is a nontrivial task given that typical ML performance metrics may not always align with the requirements of scientific validity.

Lastly, the success of this interdisciplinary research heavily depends on the seamless collaboration between planetary scientists and ML researchers. To ensure the viability and usefulness of developed ML applications in this field, researchers from both domains need to work closely together. This cross-disciplinary collaboration poses its own set of challenges, primarily rooted in differences in terminologies, methodologies, and objectives across the two fields. However, such collaborations are also vital to align the technical capabilities of ML with the scientific goals of lunar research.

3 Acknowledgements

This work was developed in the Space Robotics Research group (SpaceR) of the University of Luxembourg's Interdisciplinary Centre for Security, Reliability and Trust (SnT). I coauthor the different papers presented in this dissertation along with my supervisors from the group. For RP2 and RP3, a collaboration with researchers of different universities (University of Oxford, ETH Zurich, Polytechnic of Turin, University of Southern California) and companies (NVIDIA, RSS-Hydro) was also done. These two works were developed under the research program Frontier Development Lab (FDL), created by NASA, Google, NVIDIA, Trillium Technologies and the Luxembourg Space Agency).

VI | References

Alexandrov, O. and R. A. Beyer (2018). "Multiview Shape-From-Shading for Planetary Images". In: *Earth and Space Science* 5.10, pp. 652–666.

Bandfield, J. L., R. R. Ghent, A. R. Vasavada, D. A. Paige, S. J. Lawrence, and M. S. Robinson (2011a). "Lunar surface rock abundance and regolith fines temperatures derived from LRO Diviner Radiometer data". In: *Journal of Geophysical Research: Planets* 116.E12.

Bandfield, J. L., R. R. Ghent, A. R. Vasavada, D. A. Paige, S. J. Lawrence, and M. S. Robinson (2011b). "Lunar surface rock abundance and regolith fines temperatures derived from LRO Diviner Radiometer data". In: *Journal of Geophysical Research: Planets* 116.E12.

Bandfield, J. L., P. O. Hayne, J.-P. Williams, B. T. Greenhagen, and D. A. Paige (2015). "Lunar surface roughness derived from LRO Diviner Radiometer observations". In: *Icarus* 248, pp. 357–372.

Bucci, S., J. I. Delgado Centeno, B. Gaffinet, Z. Liang, V. Bickel, B. Moseley, M. Olivares-Mendez, et al. (2022). "Thermophysical Change Detection on the Moon with the Lunar Reconnaissance Orbiter Diviner sensor". In: *Proceedings of the Neural Information Processing Systems" Machine Learning and the Physical Sciences" Workshop (ML4PS)*. Neural Information Processing Systems Foundation, Inc.

Chen, H., X. He, L. Qing, Y. Wu, C. Ren, R. E. Sheriff, and C. Zhu (2022). "Real-world single image super-resolution: A brief review". In: *Information Fusion* 79, pp. 124–145.

Colaprete, A., R. C. Elphic, J. Heldmann, and K. Ennico (2012). "An overview of the lunar crater observation and sensing satellite (LCROSS)". In: *Space Science Reviews* 167, pp. 3–22.

Colaprete, A., P. Schultz, J. Heldmann, D. Wooden, M. Shirley, K. Ennico, B. Hermalyn, W. Marshall, A. Ricco, R. C. Elphic, et al. (2010). "Detection of water in the LCROSS ejecta plume". In: *science* 330.6003, pp. 463–468.

Delgado Centeno, J. I., P. Harder, B. Moseley, V. Bickel, S. Ganju, M. Olivares-Mendez, and A. Kalaitzis (2021a). "Single Image Super-Resolution with Uncertainty Estimation for Lunar Satellite Images". In: *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*.

Delgado Centeno, J. I., P. J. Sanchez Cuevas, C. Martinez Luna, and M. A. Olivares Mendez (2021b). "Enhancing lunar reconnaissance orbiter images via multi-frame super resolution for future robotic space missions". In: *IEEE Robotics and Automation Letters* 6.4, pp. 7721–7727.

Delgado Centeno, J. I., P. J. Sanchez Cuevas, C. Martinez Luna, and M. A. Olivares Mendez (2021c). "Lunar Highres-Net: Super Resolution For Lunar Surface Imagery". In: 72th International Astronautical Congress, IAF Space Exploration Symposium.

Delgado Centeno, J. I., P. J. Sanchez Cuevas, C. Martinez Luna, and M. A. Olivares Mendez (2021d). "Lunar Surface Images Enhancement for Space Resources Localization and Extraction". In: Space Resources Week.

Deudon, M., A. Kalaitzis, I. Goytom, M. R. Arefin, Z. Lin, K. Sankaran, V. Michalski, S. E. Kahou, J. Cornebise, and Y. Bengio (2020). "Highres-net: Recursive fusion for multi-frame super-resolution of satellite imagery". In: *arXiv preprint arXiv:2002.06460*.

Dong, X., X. Sun, X. Jia, Z. Xi, L. Gao, and B. Zhang (2020). "Remote sensing image super-resolution using novel dense-sampling networks". In: *IEEE Transactions on Geoscience and Remote Sensing* 59.2, pp. 1618–1633.

Farsiu, S., M. D. Robinson, M. Elad, and P. Milanfar (2004). "Fast and robust multiframe super resolution". In: *IEEE transactions on image processing* 13.10, pp. 1327–1344.

Gladstone, G. R., D. M. Hurley, K. D. Retherford, P. D. Feldman, W. R. Pryor, J.-Y. Chaufray, M. Versteeg, T. K. Greathouse, A. J. Steffl, H. Throop, et al. (2010a). "LRO-LAMP observations of the LCROSS impact plume". In: *Science* 330.6003, pp. 472–476.

Gladstone, G. R., S. A. Stern, K. D. Retherford, R. K. Black, D. C. Slater, M. W. Davis, M. H. Versteeg, K. B. Persson, J. W. Parker, D. E. Kaufmann, et al. (2010b). "LAMP: the Lyman alpha mapping project on NASA's Lunar reconnaissance orbiter mission". In: *Space Science Reviews* 150, pp. 161–181.

Goswami, J. and M Annadurai (2009). "Chandrayaan-1: India's first planetary science mission to the moon". In: *Current science*, pp. 486–491.

Hayne, P. O., J. L. Bandfield, M. A. Siegler, A. R. Vasavada, R. R. Ghent, J.-P. Williams, B. T. Greenhagen, O. Aharonson, C. M. Elder, P. G. Lucey, et al. (2017). "Global regolith thermophysical properties of the Moon from the Diviner Lunar Radiometer Experiment". In: *Journal of Geophysical Research: Planets* 122.12, pp. 2371–2400.

Heggy, E., E. Palmer, T. W. Thompson, B. Thomson, and G. W. Patterson (2020). "Bulk composition of regolith fines on lunar crater floors: Initial investigation by LRO/Mini-RF". In: *Earth and Planetary Science Letters* 541, p. 116274.

Ho, J., X. Chen, A. Srinivas, Y. Duan, and P. Abbeel (2019). "Flow++: Improving flow-based generative models with variational dequantization and architecture design". In: *International Conference on Machine Learning*. PMLR, pp. 2722–2730.

Huang, Y., W. Wang, and L. Wang (2015). "Bidirectional recurrent convolutional networks for multi-frame super-resolution". In: *Advances in neural information processing systems* 28.

Jiang, J., C. Chen, J. Ma, Z. Wang, Z. Wang, and R. Hu (2016). "SRLSP: A face image super-resolution algorithm using smooth regression with local structure prior". In: *IEEE Transactions on Multimedia* 19.1, pp. 27–40.

Kato, M., S. Sasaki, and Y. Takizawa (2010). "The Kaguya mission overview". In: *Space science reviews* 154.1, pp. 3–19.

Kumar, S., A. Singh, A. Sharma, V. Chaudhary, A. Joshi, S. Agrawal, and P. Chauhan (2022). "Polarimetric analysis of L-band DFSAR data of Chandrayaan-2 mission for ice detection in permanently shadowed regions (PSRs) of lunar South polar craters". In: *Advances in Space Research* 70.12, pp. 4000–4029.

Ledig, C., L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. (2017). "Photo-realistic single image super-resolution using a generative adversarial network". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4681–4690.

Li, S., P. G. Lucey, R. E. Milliken, P. O. Hayne, E. Fisher, J.-P. Williams, D. M. Hurley, and R. C. Elphic (2018). "Direct evidence of surface exposed water ice in the lunar polar regions". In: *Proceedings of the National Academy of Sciences* 115.36, pp. 8907–8912.

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

Mahapatra, D., B. Bozorgtabar, and R. Garnavi (2019). "Image super-resolution using progressive generative adversarial networks for medical image analysis". In: *Computerized Medical Imaging and Graphics* 71, pp. 30–39.

Mitrofanov, I. G., A. B. Sanin, W. Boynton, G Chin, J. Garvin, D Golovin, L. Evans, K Harshman, A. Kozyrev, M. Litvak, et al. (2010a). "Hydrogen mapping of the lunar south pole using the LRO neutron detector experiment LEND". In: *science* 330.6003, pp. 483–486.

Mitrofanov, I. G., A. B. Sanin, W. Boynton, G Chin, J. Garvin, D Golovin, L. Evans, K Harshman, A. Kozyrev, M. Litvak, et al. (2010b). "Hydrogen mapping of the lunar south pole using the LRO neutron detector experiment LEND". In: *science* 330.6003, pp. 483–486.

Nozette, S., P. Spudis, B. Bussey, R. Jensen, K. Raney, H. Winters, C. L. Lichtenberg, W. Marinelli, J. Crusan, M. Gates, et al. (2010). "The Lunar Reconnaissance Orbiter miniature radio frequency (Mini-RF) technology demonstration". In: *Space Science Reviews* 150, pp. 285–302.

Ouyang, Z., C. Li, Y. Zou, H. Zhang, C. Lü, J. Liu, J. Liu, W. Zuo, Y. Su, W. Wen, et al. (2010). "Primary scientific results of Chang'E-1 lunar mission". In: *Science China Earth Sciences* 53, pp. 1565–1581.

Qin, B. and D. Li (2020). "Identifying facemask-wearing condition using image super-resolution with classification network to prevent COVID-19". In: *Sensors* 20.18, p. 5236.

Rezende, D. and S. Mohamed (2015). "Variational inference with normalizing flows". In: *International conference on machine learning*. PMLR, pp. 1530–1538.

Robinson, M. S., S. Brylow, M Tschimmel, D Humm, S. Lawrence, P. Thomas, B. W. Denevi, E Bowman-Cisneros, J Zerr, M. Ravine, et al. (2010). "Lunar reconnaissance orbiter camera (LROC) instrument overview". In: *Space science reviews* 150, pp. 81–124.

Rosenburg, M., O. Aharonson, J. Head, M. Kreslavsky, E Mazarico, G. A. Neumann, D. E. Smith, M. H. Torrence, and M. T. Zuber (2011). "Global surface slopes and

roughness of the Moon from the Lunar Orbiter Laser Altimeter". In: *Journal of Geophysical Research: Planets* 116.E2.

Schwadron, N. A., T Baker, B Blake, A. Case, J. Cooper, M Golightly, A Jordan, C Joyce, J Kasper, K Kozarev, et al. (2012). "Lunar radiation environment and space weathering from the Cosmic Ray Telescope for the Effects of Radiation (CRaTER)". In: *Journal of Geophysical Research: Planets* 117.E12.

Smith, D. E., M. T. Zuber, G. A. Neumann, E. Mazarico, F. G. Lemoine, J. W. Head III, P. G. Lucey, O. Aharonson, M. S. Robinson, X. Sun, et al. (2017). "Summary of the results from the lunar orbiter laser altimeter after seven years in lunar orbit". In: *Icarus* 283, pp. 70–91.

Smith, M., D. Craig, N. Herrmann, E. Mahoney, J. Krezel, N. McIntyre, and K. Goodliff (2020). "The Artemis program: an overview of NASA's activities to return humans to the moon". In: *2020 IEEE Aerospace Conference*. IEEE, pp. 1–10.

Spudis, P. D., E Shoemaker, C Acton, B Burratti, T Duxbury, D Baker, D Smith, J Blamont, M Davies, and E Eliason (1994). *The Clementine mission: initial results from lunar mapping*. Tech. rep.

Spudis, P., D. Bussey, S. Baloga, B. Butler, D Carl, L. M. Carter, M Chakraborty, R. Elphic, J. Gillis-Davis, J. Goswami, et al. (2010). "Initial results for the north pole of the Moon from Mini-SAR, Chandrayaan-1 mission". In: *Geophysical Research Letters* 37.6.

Spudis, P., D. Bussey, S. Baloga, J. Cahill, L. Glaze, G. Patterson, R. Raney, T. Thompson, B. Thomson, and E. Ustinov (2013). "Evidence for water ice on the Moon: Results for anomalous polar craters from the LRO Mini-RF imaging radar". In: *Journal of Geophysical Research: Planets* 118.10, pp. 2016–2029.

Sundararajan, V. (2018). "Overview and technical architecture of India's Chandrayaan-2 mission to the Moon". In: *2018 AIAA Aerospace Sciences Meeting*, p. 2178.

Tran, T, M. Rosiek, R. A. Beyer, S Mattson, E Howington-Kraus, M. Robinson, B. Archinal, K Edmundson, D Harbour, E Anderson, et al. (2010). "Generating digital terrain models using LROC NAC images". In: *At ASPRS/CaGIS 2010 Fall Specialty Conference*.

Van Ouwerkerk, J. (2006). "Image super-resolution survey". In: *Image and vision Computing* 24.10, pp. 1039–1052.

Vasavada, A. R., J. L. Bandfield, B. T. Greenhagen, P. O. Hayne, M. A. Siegler, J.-P. Williams, and D. A. Paige (2012a). "Lunar equatorial surface temperatures and re-

golith properties from the Diviner Lunar Radiometer Experiment". In: *Journal of Geophysical Research: Planets* 117.E12.

Vasavada, A. R., J. L. Bandfield, B. T. Greenhagen, P. O. Hayne, M. A. Siegler, J.-P. Williams, and D. A. Paige (2012b). "Lunar equatorial surface temperatures and regolith properties from the Diviner Lunar Radiometer Experiment". In: *Journal of Geophysical Research: Planets* 117.E12.

Winkler, C., D. Worrall, E. Hoogeboom, and M. Welling (2019). "Learning likelihoods with conditional normalizing flows". In: *arXiv preprint arXiv:1912.00042*.

Yang, D., Z. Li, Y. Xia, and Z. Chen (2015). "Remote sensing image super-resolution: Challenges and approaches". In: *2015 IEEE international conference on digital signal processing (DSP)*. IEEE, pp. 196–200.

Yang, W., X. Zhang, Y. Tian, W. Wang, J.-H. Xue, and Q. Liao (2019). "Deep learning for single image super-resolution: A brief review". In: *IEEE Transactions on Multimedia* 21.12, pp. 3106–3121.

Yue, L., H. Shen, J. Li, Q. Yuan, H. Zhang, and L. Zhang (2016). "Image super-resolution: The techniques, applications, and future". In: *Signal processing* 128, pp. 389–408.

Zhang, S., G. Liang, S. Pan, and L. Zheng (2018). "A fast medical image super resolution method based on deep learning network". In: *IEEE Access* 7, pp. 12319–12327.

Zhao, B., J. Yang, D. Wen, W. Gao, L. Chang, Z. Song, B. Xue, and W. Zhao (2011). "Overall scheme and on-orbit images of Chang'E-2 lunar satellite CCD stereo camera". In: *Science China Technological Sciences* 54, pp. 2237–2242.