ARC-LIGHT: Algorithm for Robust Characterization of Lunar surface Imaging for Ground Hazards and Trajectory

University of Michigan Human Lander Challenge 2024 Team

Category: Instrumentation Performance Through the Dust Cloud During Landing

Alexander Cushen (Graduate, Climate and Space Science and Engineering)
Samuel Carrico (Undergraduate, Aerospace Engineering)
Ariana Bueno (Graduate, Applied Physics)
Corrydon Wettstein (Undergraduate, Computer Science Engineering)
Jaykumar Ishvarbhai Adalja (Graduate, Space Engineering)
Mengxiang Shi (Undergraduate, Computer Science)
Naila Garcia (Undergraduate, Mechanical Engineering)
Yuliana Garcia (Undergraduate, Mechanical Engineering)

Faculty Advisors: Mirko Gamba Chris Ruf

University of Michigan: ARC-LIGHT

Major Objectives & Technical Approach

- Enable position and attitude calibration and hazard detection throughout landings impacted by PSI
- Enhance the utility of existing spacecraft sensors at low altitudes to improve landing safety and precision
- Deploy sensor fusion algorithm to combine available data streams into a robust data product
- Utilize real-world, experimental, and synthetic datasets to train a flexible and adaptable algorithm

Image/Graphic:

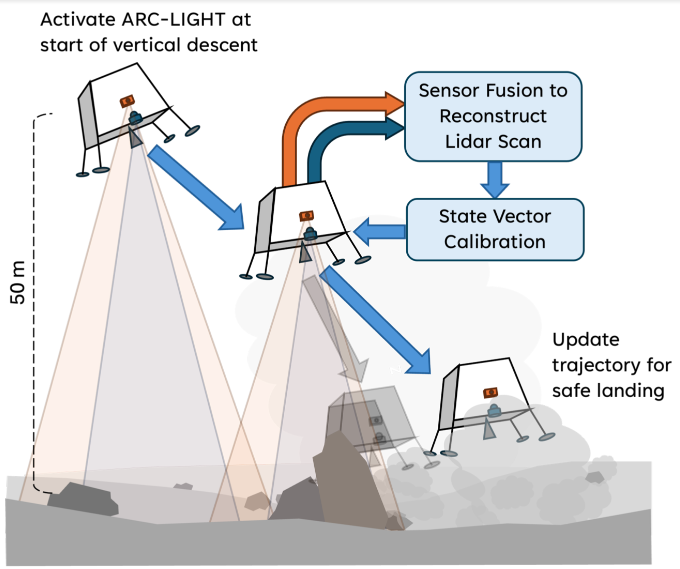

Key Design Details & Innovations of the Concept

- Employs camera + lidar sensor fusion techniques matured by autonomous vehicle hazard detection systems
- Algorithm framework is compatible with many spacecraft designs by making use of common sensors without adding payload mass
- Machine learning-driven system reduces computational cost during use, facilitating real-time use with minimal resources
- Synthetic data combined with limited experimental data used to robustly train algorithm for all landing conditions

Summary of Schedule & Costs for the proposed solution's path to adoption

- 3 ¼ year schedule to finalize algorithm design, gather training data, and complete spacecraft integration and testing
- Multiple subteams working in parallel to reduce project length and ensure coordination between algorithm and data development
- $330 k baseline cost, composed of 1548 FTE weeks and $275 k for equipment and computational resources

Contents

1 Executive Summary

Safe and reliable lunar landings are crucial for future exploration of the Moon. The regolith ejected by a lander's rocket exhaust plume presents a significant obstacle in achieving this goal. It prevents spacecraft from reliably utilizing their navigation sensors to monitor their trajectory and spot emerging surface hazards as they near the surface. We propose to develop and implement a machine learning-based sensor fusion system, ARC-LIGHT, that integrates sensor data from the cameras, lidars, or radars that landers already carry but disable during the final landing phase. Using these data streams, ARC-LIGHT will remove erroneous signals and recover a useful detection of the surface, for use by the spacecraft to adjust its descent profile. It also provides a layer of redundancy for other key sensors, like inertial measurement units. The feasibility of this technology is validated through the development of a prototype algorithm, which is trained on data from a purpose-built testbed facility. Based on these findings, a refined algorithm architecture, development timeline, and budget for ARC-LIGHT are presented.

2 Introduction

2.1 Problem Statement

As international interest in lunar science and exploration grows, the capability for safe and precise landings on the Moon's surface has become more important than ever. Recent missions like SLIM, Chang'e, IM-1, and Chandrayaan have made major strides in this field, but NASA's ambitious Artemis and Moon to Mars initiatives, which aim to provide regular access to the surface for crew and cargo, necessitate enhanced reliability and accuracy of lunar landings. A significant hazard during these landings is the ejection of regolith as a lander's rocket plume impinges on the surface. This Plume Surface Interaction (PSI) results from the engine exhaust of a lander interacting with the surface of a planetary body during the landing and ascending phases. PSI can eject large amounts of regolith particles that can limit visibility, spoof navigation systems, and damage surrounding surface assets. Therefore, understanding and mitigating challenges due to PSI is paramount for the safety and success of upcoming lunar missions.

A significant challenge presented by PSI is dust interference with navigation systems; the lofted clouds of regolith can obstruct the view of the cameras, radars, and lidars used for navigation. During the Apollo landings, astronauts observed various instances where PSI was responsible for visual limitations and radar interference during landing, most notably in Apollo 15 [1]. Current lunar missions typically employ cameras for computer vision-based optical navigation, and lidar and radar systems to model surface features and determine their state vector–their position and attitude relative to surface landmarks. As the landers near the surface, these systems are continuously affected by the regolith lofted by PSI. The PSI-induced dust clouds obscure the view of the surface, change the ambient lighting conditions, and provide a diffuse reflective surface for lidar and radar signals to bounce off. As such, contemporary lunar landers like Chang'e typically do not use these sensors during the final stage of vertical descent due to PSI-induced dust interference [2].

To address this issue, we propose to develop a **machine learning-based sensor fusion system** that **integrates sensor data from standard navigation sensors to be effectively utilized** during the final descent phase of landing. Named Algorithm for Robust Characterization of Lunar surface Imaging for Ground Hazards and Trajectory (ARC-LIGHT), this system will fuse measurements from cameras, lidars, or other employed sensors, to remove signal noise and allow the spacecraft to "see-through" the PSI cloud. This, in turn, will enhance landing accuracy, by allowing for re-calibration of the state vector during vertical descent and improve landing safety by enabling hazard avoidance scans at lower altitudes. Since the system employs data from sensors already on the spacecraft, ARC-LIGHT will not add any hardware mass or complexity to the landers and is readily compatible with multiple lander types. To demonstrate the feasibility of this technology, we have established a testbed facility to gather sensor data affected by a diffuse optically scattering medium. The dataset from this experiment is used to demonstrate a prototype sensor fusion algorithm for reconstructing lidar scans. These results inform the proposed budget, timeline, and project risk management strategy for the development and deployment of ARC-LIGHT.

2.2 Background

Modern lunar landers rely on autonomous systems to safely conduct entry, landing, and descent (EDL) on the Moon. These landers often rely on lidar or radar altimeters and velocimeters, inertial measurement units (IMUs), star trackers, and vision-based methods to properly locate themselves during descent. They also combine two-dimensional image-based methods with 3-dimensional lidar readings to locate and avoid hazards [3]. These sensors are enabled and disabled as needed during different phases of descent.

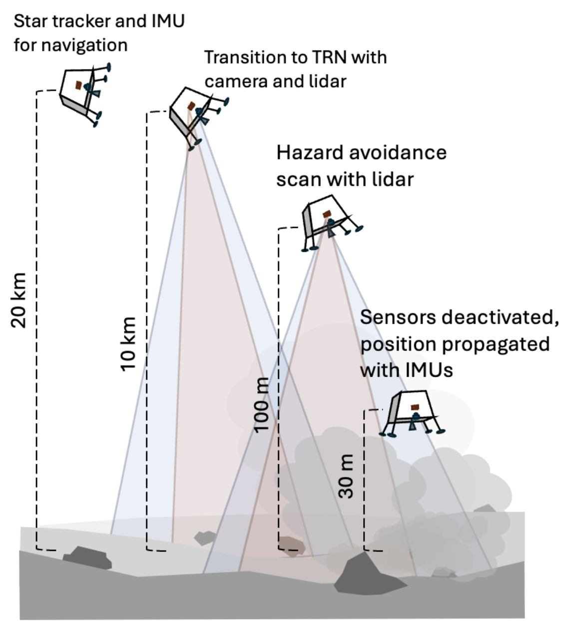

Figure 1: *General framework for sensor usage during EDL, showing navigation methodology at different stages, including deactivation for vertical descent.*

The recent robotic landings of Chang'e 3-6, Chandrayaan-3, IM-1, and SLIM each employ a similar EDL profile, illustrated in fig. 1 [4–7]. Following descent orbit insertion, landers begin powered descent around 20 km above the surface. Powered descent consists of three phases: braking, approach, and descent. The braking phase serves to reduce the lander's velocity from orbital speeds. A star tracker and IMU remain active throughout this phase with lidar altimeters and velocimeters activated around 10 km. These serve to update the state vector as it evolves from the initial orbital parameters. During the approach phase, ranging from around 5 km to 100 m above the lunar surface, the lander establishes visual contact with the landing ellipse. The star tracker, IMU, altimeters, and velocimeters remain active, with an optical camera also activated during this phase. Finally, the vertical descent phase occurs, preceded by a hazard detection and avoidance (HDA) scan at 50-100 m above the surface, which sees a lidar used to determine a safe landing site. Local landmarks are used to provide a local reference frame, called terrain relative navigation (TRN) [6]. The lander continues using sensors as it navigates to a position approximately 50 m above the selected landing site. Below this altitude, it disables every sensor except for its IMU for final vertical descent [5].

A significant reason for autonomous landers disabling their sensors during final descent is the effect of PSI. The high density of dust particles ejected during landing can interfere with the sensors causing them to be unreliable. The Chang'e 3 mission observed this interference 20 meters above the lunar surface during its landing [4], but PSI onset at higher altitudes has been observed for higher thrust landers [8, 9]. The lack of active sensor guidance during this period of the flight introduces a margin of error, potentially leading to catastrophic consequences for the spacecraft.

3 Technology Concept and Innovation

3.1 Outline

To improve lunar landing accuracy and safety, we propose to implement a **sensor fusion-based navigation algorithm**, ARC-LIGHT, which takes advantage of the otherwise unused sensor data during the final stage of vertical descent. ARC-LIGHT is motivated by three key objectives:

1. **Improve landing accuracy** within the landing zone, by allowing for re-calibration of the spacecraft state vector using TRN during vertical descent.
2. **Improve hazard detection** by enabling hazard avoidance scans at lower altitudes to spot small surface features that could not be seen from higher altitudes, and pits or rocks moved or formed by PSI.
3. **Provide sensor redundancy** in the case of IMU failure during landing.

ARC-LIGHT fulfills these objectives by deploying a number of sensor fusion signal processing techniques, namely a convolutional neural network (CNN), to integrate multiple data streams and reconstruct a clear lidar scan of the surface. This machine learning-based framework is flexible and lander-agnostic, easily adjusting to take advantage of different sensor configurations, including lidars, cameras, and radars. With ARC-LIGHT, the spacecraft will be able to "see-through" the PSI cloud, providing an additional layer of navigational redundancy during the crucial terminal landing phase.

3.2 Sensor Fusion Algorithm

ARC-LIGHT uses sensor fusion, which is the signals processing framework of combining separate sensor datasets into an integrated data product to provide greater measurement confidence and precision [10]. This technique is most useful where multiple sensor types (e.g. cameras and lidars) observe a scene simultaneously; lidars offer 3D structure information, while cameras provide a broader contextualization of the scene. Since each sensor responds to signal interference differently, the combination of their encoded information allows for a clearer signal to be extracted.

Multiple sensor fusion techniques may be applied to help improve the spacecraft's lidar and camera observations during vertical descent. The choice of which techniques to implement depends on the sensor availability and configuration on the spacecraft, as well as computational resources and landing precision requirements. Figure 2 outlines a proposed architecture for ARC-LIGHT, demonstrating how multiple techniques may be combined. This candidate algorithm has been refined from the example presented in the proposal, based on lessons learned in our research and development. Future refinement of this algorithm will be informed by trade studies and mission requirements for computational cost and performance metrics.

In this proposed algorithm, a LIDROR-type filter [11] is first applied to the lidar input to remove lidar points with low local intensity, corresponding to stochastic cloud backscatter. This form of backscatter is caused by strong reflection from lofted regolith particles. The camera data is fed into a CNN which estimates the optical depth of the lofted regolith in different sectors of the image (e.g. a 16×16 grid) to accommodate spatial inhomogeneity of the PSI cloud. The optical depth τ is a dimensionless measure of light extinction, which is related to the number density of lofted regolith n, regolith cross-section σ, and light path length s:

$$\tau = n\sigma s \tag{1}$$

The image data is also fed into a dehazing module, which increases signal-to-noise for better identification of surface features. Dehazing is an image processing formula that aims to remove haze or blurriness from an image. One common dehazing algorithm is dark channel prior (DCP). DCP uses pixels with low intensity in one of the RGB channels; these pixels are deemed the 'dark pixels'. It should be noted that spacecraft navigation cameras may be monochromatic. However, even if there are no RGB channels, DCP is still feasible, since the algorithm runs by channel; therefore, it would only exercise through the one available channel. It is one of the fastest dehazing algorithms, and its light computational cost is ideal for a resource-constrained lunar lander.

The optical depth estimate, dehazed image, and filtered lidar data are then fed into a second CNN, which integrates these data streams to recreate the lidar scan of the surface as if no PSI

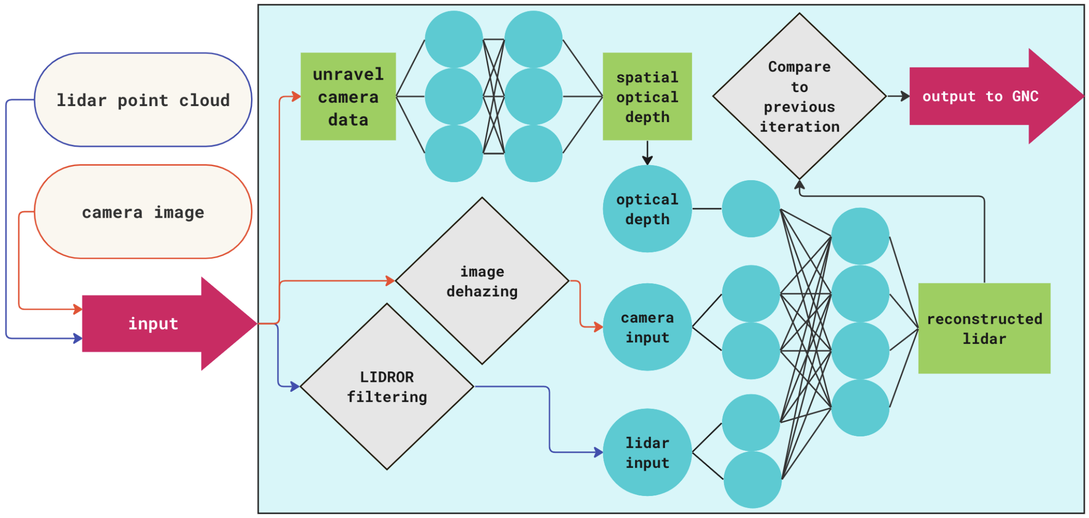

Figure 2: *Prototype algorithm architecture. lidar and camera data are input. Image data is processed in initial CNN to determine optical depth of image in multiple sectors, and dehazed to improve signal-to-noise. lidar data is filtered to remove significant outliers corresponding to strong PSI cloud backscatter. These intermediate quantities are combined in a second CNN to reconstruct lidar scan without interference. Output is sent to spacecraft GNC.*

was present. This reconstructed lidar scan is then compared to previous iterations, to detect any significant errors. If output error exceeds operational bounds due to excessive PSI interference levels, ARC-LIGHT is automatically deactivated. Assuming the output is within bounds, the lidar reconstruction is returned to the lander GNC in the native lidar format for immediate use at a target 1 Hz cadence. This cadence is subject to computational cost and precision requirements of the mission but is feasible to achieve given the low computational cost of using trained CNNs. Heuristics strategies may also help to improve computational efficiency. These are techniques to constrain the decision-making processes and choose optimal solutions without expending unnecessary computational resources trying to find a mathematically perfect one [12].

3.3 Training and Verification

Large amounts of lunar lander data will be needed to train and test the algorithm before operational use. It is infeasible to provide a large enough amount of data solely from available lunar landings given how few have occurred. Ground testing will be employed to supplement these few real-world cases, but cost and time limitations limit the number of experimental runs which can be performed. Additionally, in a lunar testbed, it is infeasible to control all experimental factors such as lighting, vacuum conditions, scale, and gravitational conditions to reproduce flawlessly a lunar landing. Given these limitations, we propose to pre-train our network using synthetic data generated from computational PSI simulations.

PSI models can simulate various scenarios with precise control across a broad parameter space including particle size and density, surface bonding force, and lander descent profiles [13–15]. They have been developed using various methods to recreate the interaction between rocket plumes and regolith particles in the lunar environment [16–19]. For instance, the Direct Simulation Monte Carlo (DSMC) method simulates the behavior of lunar regolith ejecta generated by PSI on a kinetic level by computing individual "macro-particle" trajectories, including stochastic collisions, and resolving time-accurate details like density, velocity, and particle size distribution.

These time-resolved datasets provide the basic data needed to generate the needed training data for ARC-LIGHT. Using both high-fidelity optical simulation and 3D modeling tools like Blender, the regolith dust distribution can be translated into synthetic images simulating what a lander would observe during a landing. Further, Mie scattering calculations allow for accurate lidar backscatter simulations, providing the associated lidar data alongside the camera images. Testbed experiments with regolith simulant will help calibrate these optical simulations. While any given simulation and sensor reconstruction is unlikely to mimic reality perfectly, the broad parameter space captured in this synthetic data encompasses the conceivable range of conditions the CNN may face on the Moon.

By "pre-training" the CNN on a broad training dataset, the need for real-world data is drastically reduced. This approach has demonstrated success for CNN object prediction based on synthetic image-based training [20, 21]. The CNN developed from this training can then be "fine-tuned" using the limited high-fidelity testbed data and currently available sensor data from lunar landings [22, 23]. Fine-tuning is a training technique to refine the accuracy of a pre-trained CNN toward the specific use case by performing small training steps on real-world data.

To ensure the safety of the mission and validate the success of ARC-LIGHT, error quantification will be studied. Error quantification will provide a baseline to establish shut-off protocols. ARC-LIGHT will be set to deactivate once PSI becomes significant enough that reliable outputs can no longer be guaranteed. This cut-off point can be determined from an initial hazard scan. Once the standard error is calculated we can program the system to shut off ARC-LIGHT if the distance measured is more than one standard deviation from the initial scan.

Additionally, ARC-LIGHT can be verified as a passive system during an initial lunar lander mission to test out the system, collect data, and build confidence in the pre-training approach. Again since ARC-LIGHT does not add additional hardware, this implementation to verify will not require a lot of additional cost or added risk to the mission since it will be incorporated as a passive system to be tested.

3.4 Operation

Currently, landers accomplish vertical descent by integrating their position based on IMU measurements. This leaves them essentially blind as they are unable to rely on optical sensors like lidar to confirm their location. This results in growing positional uncertainty as to where the lander is located relative to where it believes it is. From the Chang'e 3 and Chang'e 5 missions, positional uncertainty of $\Delta x \approx 2$ meters was observed [2, 5]. This position uncertainty is partially due to the combined effects of error in state vector determination during EDL and the steady impact of IMU drift between state vector calibration. Given the constant error present in any IMU, the integrated error in acceleration and rotation results in a positional uncertainty [24]:

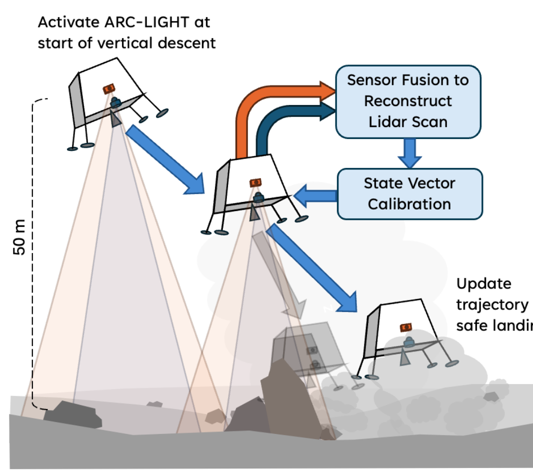

Figure 3: *Illustration of final descent where ARC-LIGHT is used to update spacecraft trajectory.*

$$\Delta x = vt \sin\left(\frac{\text{ARW}}{60}\sqrt{t}\right) + vt \sin\left(\frac{\text{BI}}{3600}t\right) \tag{2}$$

where v is the lander velocity and t is time since calibration. ARW represents the IMU's angle random walk, which is the random noise error of an IMU's gyroscopes. BI represents the IMU's bias instability, which is the drift the IMU's measurements have from the average value of its output rate. For this calculation, we used the parameters from the Chang'e 5 landing, a velocity of 1.5 m/s and a descent time of 66.66 seconds, and a commercial IMU, with a BI of 0.25 degrees/hour and an

ARW of 0.125 degrees per $\sqrt{\text{hour}}$ [5, 25]. In this case, the positional uncertainty caused by IMU drift is only ~ 4 centimeters, which is very small compared to the total uncertainty of ~ 2 m. We can conclude that the vast majority of a lander's positional uncertainty is caused by other forms of error, such as error in the initial HDA scan or environmental effects. ARC-LIGHT can combat these difficult to quantify errors through constant recalibration of the lander's position.

ARC-LIGHT allows the lander to make periodic re-calibrations of the state vector throughout vertical descent to account for any drift, as illustrated in fig. 3. This state vector correction is performed by the GNC with the same software already used during earlier descent stages. This is done using an extended Kalman filter that combines observational data with the propagated position estimate, accounting for the computational delay of processing surface observation data [6, 26].

Beyond its ability to increase the precision of a lunar lander, ARC-LIGHT also allows the lander to detect changes in the surface which include hazards uncovered by PSI, or features that are too small to resolve at the initial HDA scan height. The reconstructed lidar scan can be used to assess the magnitude of these hazards and perform course corrections if needed. ARC-LIGHT also provides redundancy in case the IMU were to fail, as it could help the lander calculate its inertial measurements based off of the images taken. All of these benefits of ARC-LIGHT will help landers execute safer and more accurate landings on the Moon.

4 Feasibility Analysis

4.1 Outline

To fulfill its objectives of improving landing accuracy, safety, and redundancy, ARC-LIGHT calls for the development of machine-learning systems that are robust to the optical interference of the PSI landing environment. To assess the feasibility of this concept, we have undertaken an experimental campaign to study the response of lidars and cameras to a simulated lunar environment and develop a prototype ML-based sensor fusion algorithm for lidar reconstruction. This section reviews the present operational uses of ML-based navigation in spaceflight, introduces our experimental facility and prototype algorithm, and presents our findings from this study.

4.2 Review of Machine Learning in Spacecraft Navigation

Sensor fusion and machine learning-based navigation systems have a demonstrated heritage in autonomous vehicles [27–29], and have been proposed for numerous spaceflight applications [3, 26, 30–32]. Examples of currently operational ML-based space-based navigation can also be seen in two of NASA's Mars rovers, Perseverance and Curiosity.

Perseverance employs the ML-based navigation software Enhanced AutoNav. It sorts a list of potential paths for the rover to traverse and works with the Approximate Clearance Evaluation (ACE) algorithm to evaluate the safety of ranked paths [33]. To reduce the high computational cost of operating ACE, two heuristics are used to accelerate the path selection process: the Gradient Convolution and Learned Heuristics. Gradient Convolution Heuristic assesses the terrain roughness by analyzing the terrain's gradient. Learned Heuristic predicts the ACE value based on the heightmap data of the terrain. Both heuristics evaluate the local terrain to prioritize optimal paths and reduce the computational cost to the rover.

Curiosity utilizes AEGIS (Autonomous Exploration for Gathering Increased Science), another ML-based system used to autonomously select and prioritize targets for analysis [34]. With AEGIS, Curiosity's Chemistry and Camera instrument can autonomously select target rocks for its laser spectrometer and telescopic camera. It analyzes images from Curiosity's stereo Navigation Camera or ChemCam's Remote Micro-Imager to identify potential targets based on adjustable criteria set by scientists [35].

These examples illustrate the practicality and reliability of machine-learning applications in complex space systems. With sufficient development and training, these techniques have helped improve planetary navigation and streamline robotic operations.

4.3 ARC-LIGHT prototype development

4.3.1 Experiment Methodology

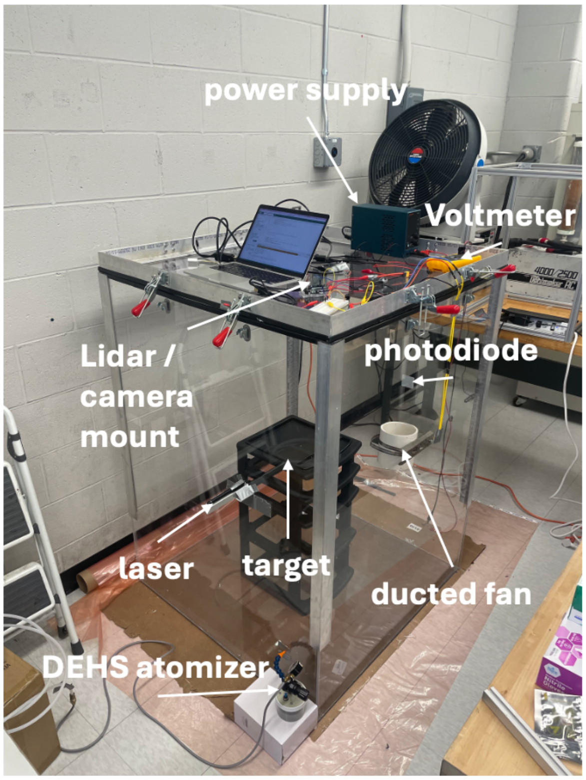

Figure 4: *Annotated image of SELENE. The tank is sealed with latches, allowing for the lid to be removed for interior access.*

We have constructed a physical test chamber to simulate the optical interference encountered during lunar vertical descent and provide us with training data for prototype algorithm development. The construction of this facility, the Sensor Efficacy in the Lunar Environment Experiment (SELENE), was enabled by the University of Michigan Space Institute Power Grant, which this team submitted a proposal for and received in January 2024.

SELENE is composed of a $0.91 \times 0.91 \times 1.2$ m sealed acrylic tank with several mounting points for the sensors, fans, laser, photodiode, and atomizer used to simulate and study this optical environment. Figure 4 shows the layout of the experiment. A 2D scanning lidar and camera are mounted on the underside of the removable lid pointed toward the base of the tank. This is to mimic the nadir-pointing sensors used by landers. A 2D lidar was used for prototype simplicity and budget constraints. A target can be placed at the base of the tank; either a model of the lunar surface for image processing-focused experiments or a series of simple blocks for lidar evaluation. The large scale of the experiment is necessitated by the magnitude of the manufacturer-stated precision of our lidar, a Hokuyo UST-10LX. This error is ± 4 cm, meaning that the tank and lidar targets must be far larger to ensure we receive a large signal-to-noise ratio.

Initially, SELENE was designed to use regolith simulant as an optical barrier. The simulant would be lofted by the fans mounted on the walls to produce a homogenous cloud of suspended regolith particles. The tank sealant was designed to safely contain the particles, which present a health hazard and contaminant to the people and other experiments in the laboratory. The 650 nm laser mounted on one wall, pointed across the tank to a photodiode on the opposite face, measures the total column density of lofted particles. The use of lunar simulant is postponed until an additional redundant containment system (e.g. a dedicated chamber surrounding the tank) can be established. This is necessary to eliminate all possibilities of contamination, even when the testbed lid is open.

In place of regolith, a Di-Ethyl-Hexl-Sebacate (DEHS) atomizer is used to inject a mist of ~ 1 micron droplets into the chamber. DEHS is a non-reactive, colorless, odorless substance that can persist for hours as a mist. It is also safe for human exposure, although masks are worn at all times during experiments. The fan system is used to circulate the mist around the chamber to establish a homogeneous distribution of particles. By measuring the photodiode voltage before and after DEHS is introduced, the optical depth of the chamber can be inferred.

Although the DEHS droplets are in a similar size range as lunar regolith, their optical properties

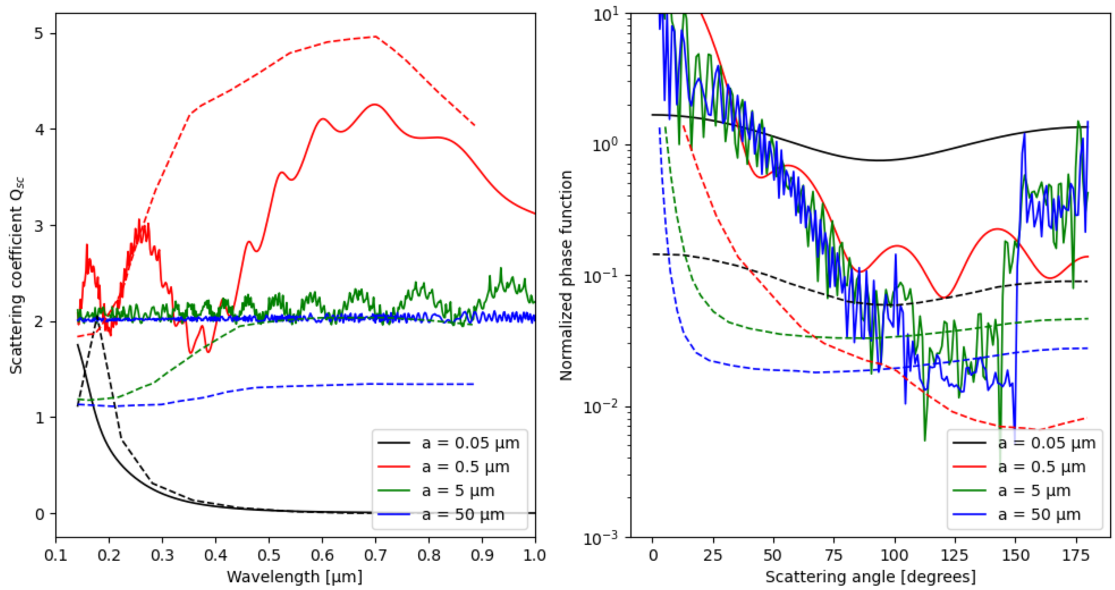

Figure 5: *Left: Dimensionless scattering coefficients for irregular regolith grains (dashed) and spherical DEHS droplets (solid) of different radii. Right: Normalized phase function of the same particles for 0.65* µm *wavelength light. Regolith data from* [36].

differ. Since we are focused on understanding the performance of lidars and cameras, the effective cross-section and scattering phase function (SPF) of these particle types are the key parameters to compare. Optical depth characterizes the amount of light extinction through some medium and depends on particle cross-section σ as shown in eq. 1; as such, knowledge of the relative cross section allows for a scaling between DEHS and regolith densities required to achieve the same optical extinction. The SPF describes the probability of light scattering into a given direction. As such, the value of the SPF at 180° represents the backscattering efficiency of the particle. This is important for the lidar, which infers the distance to an obstacle using backscattered light.

Figure 5 shows the dimensionless scattering coefficient $Q_{sc} = \sigma / \left(\pi a^2\right)$ and SPF on the left and right, respectively, for both particle types and four particle radii a. The DEHS droplets are assumed to be spherical and their properties were calculated using `miepython` [37]. The selected particle sizes span the size distribution of regolith [38]. Our lidar operates at 0.9 microns, and the laser-photodiode system operates at 0.6 microns. At these wavelengths, the left panel shows that the effective cross-section of DEHS is typically within $\sim 20\%$ of the regolith cross-section, especially for smaller particles. By contrast, the right panel shows that the backscattering efficiency (i.e. the SPF at large scattering angles) is typically one or two orders of magnitude larger for the DEHS droplets, especially at the lidar wavelength of 0.9 micron. Together, this means that the DEHS droplets will pose a far more significant obstacle to the lidar than regolith grains would. Their comparable cross sections mean both will scatter a similar amount of light out of the line of sight, which means that their impact on camera images is more comparable.

Experiments with SELENE proceed by placing a target/s of interest in the tank and sealing the lid. The photodiode voltage with the laser incident on it is recorded. Then, DEHS is gradually injected in controlled bursts, with the fan active to thoroughly circulate it. Between each addition of DEHS, a lidar scan, camera image, and photodiode voltage measurement are gathered and recorded. This process is repeated until the DEHS presents a completely opaque barrier to the lidar, which is the most sensitive to the particles as discussed above. This occurs around $\tau \approx 0.6$, which is comparable to the optical depth a lander ~ 10 m above the surface would encounter, assuming a uniform density of $n = 10^9$ particles/m^3 [9]. These experiments allow us to gather sensor data in a

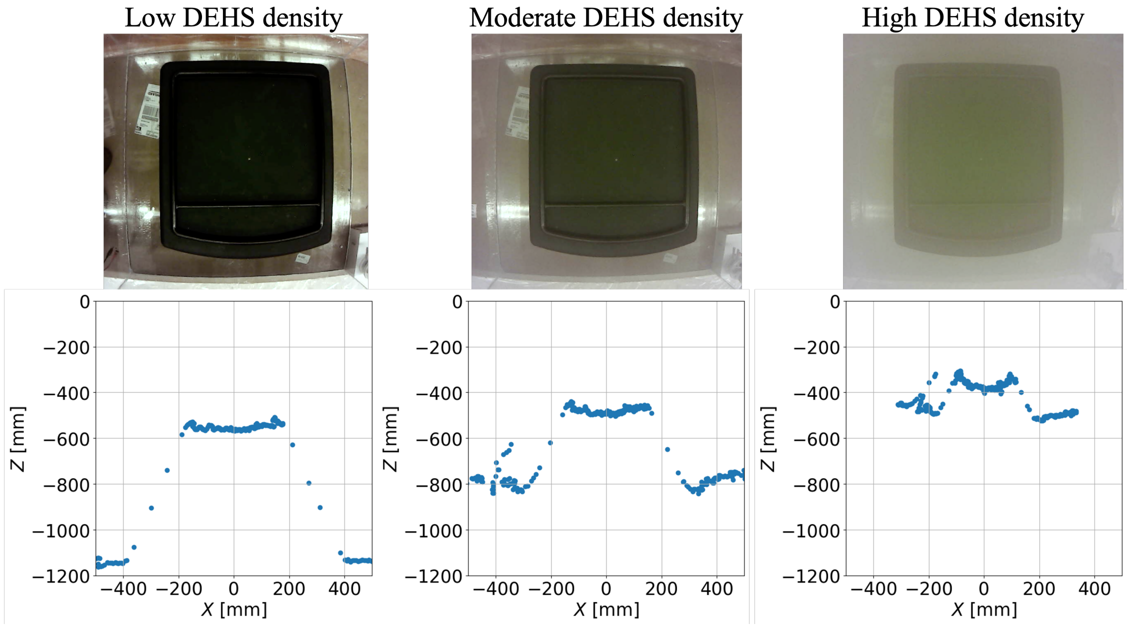

Figure 6: *Example camera images (top) and lidar scans (bottom) from SELENE. DEHS density increases from left to right. The lidar scan crosses the center of the image from top to bottom.*

controlled setting, and independently characterize the amount of interference using the optical depth as a proxy. Figure 6 shows images and lidar scans gathered for a simple square target. The lidar scan maps to the vertical centerline of the image, where X is the position along the tank base and Z is vertical height.

4.3.2 Software Development Methodology

To demonstrate ML-based sensor fusion feasibility and utility in the context of PSI-impacted landings, we have developed a prototype algorithm that fuses data from the camera and lidar to deliver a better data product than either sensor could provide alone. This prototype represents a preliminary version of the optical depth CNN and lidar reconstruction components of the proposed ARC-LIGHT algorithm shown in fig. 2.

Our prototype algorithm uses a camera image to estimate the optical depth of the tank, which allows for the attenuated lidar signal to be re-projected as if the measurement was made without any scattering medium present. This is possible because the lidar scan at large optical depths represents a combined signal of the base of the tank and the backscatter from the DEHS. Figure 6 illustrates this; in the middle panel, the lidar is still able to detect the difference between the target and the base of the tank, even though the entire structure is measured to be closer to the lidar located at (0,0). It is hypothesized that knowledge of the optical depth would allow for the lidar scan to be projected back to the distances observed when no DEHS is present, allowing us to "see-through" the interference.

A training dataset of hundreds of images taken inside SELENE at different DEHS densities, ambient lighting conditions, and arrangement of objects within the camera field of view is used to train the CNN. Critically, the laser-photodiode system allows us to label this dataset according to the measured optical depth. Figure 7 shows the architecture of the CNN; it accepts a 100×100 pixel black and white image as input and outputs an estimate of the optical depth. Three hidden layers of 64 nodes (or "neurons") each are used, with ReLU activation functions between each hidden layer. A sigmoid activation function is used for the last layer, which ensures the output is a value between 0 and 1, spanning the range of optical depths tested. For each training step ("epoch"), the CCN is fed a training image and returns an estimate of the optical depth $\hat{\tau}$. The difference between

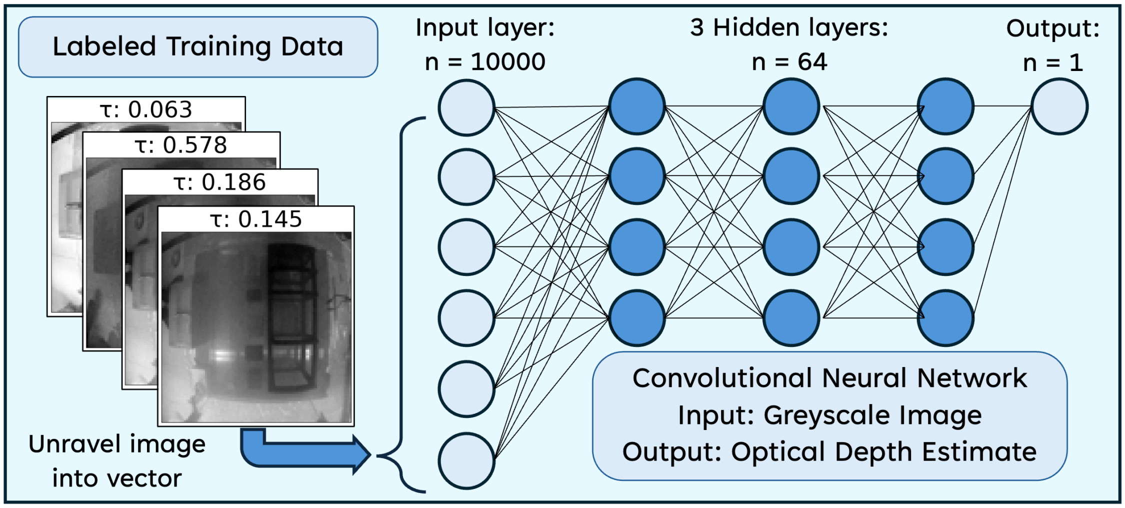

Figure 7: *Prototype CNN architecture. The training images with known optical depths are used to train the network, which is composed of multiple layers of n neurons each. For a given image* (100×100 *pixel values), it outputs an estimate of the optical depth (value between 0 and 1).*

$\hat{\tau}$ and the true optical depth τ, the "loss", is calculated. The network then "learns" by adjusting the weights and biases of each neuron via backpropagation to minimize the loss. As such, the loss of the CNN decreases as it trains, representing convergence towards a solution.

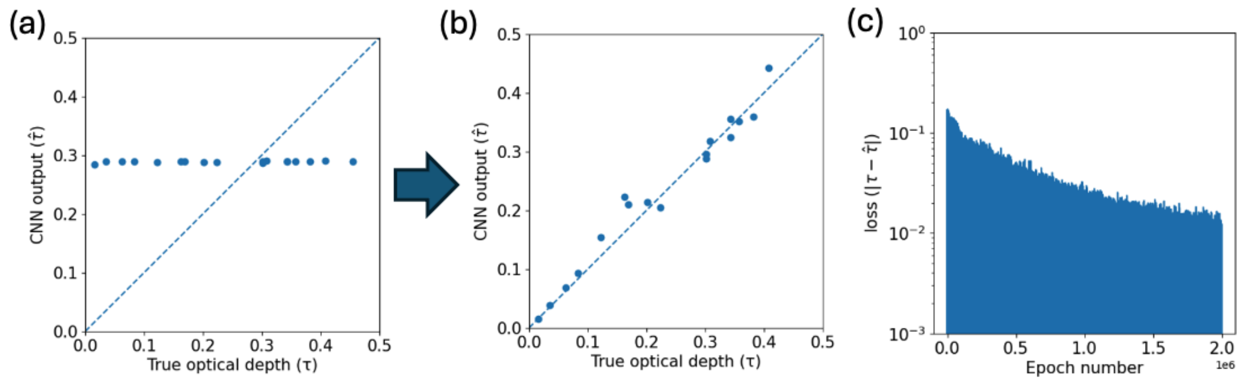

Figure 8: *CNN training metrics. (a),(b): CNN estimate of the optical depth for the test data, plotted against the true value. (a) is before training, and (b) is after training completes. (c): Plot of CNN loss across the training epochs.*

To quantify training success, 20 random images from the SELENE dataset are reserved for testing its prediction accuracy. The CNN is not trained on these images. Figure 8 (a) and (b) shows the CNN output before and after training respectively for these test images, as well as the loss over the training epochs, (c). The center panel shows that $\hat{\tau}$ successfully converges to τ after training, indicating that the CNN can generalize from its training data and estimate τ for new images.

Alongside this, lidar scans at varying optical depths are used to develop an analytic fit for the signal attenuation, $\chi(\theta, \tau)$. χ is defined as the lidar measured distance over the true distance at each scan angle θ and optical depth τ. For this fitting, scans of the flat base of the tank without any obstacles were used for consistency. The solid lines in the left panel of fig. 9 shows χ as a function of θ for multiple optical depths, as indicated by the line color; red represents small optical depth, and blue is large. These curves are each fit to the function:

$$\chi = a\theta + b + \theta^{-c} \tag{3}$$

where a, b, and c are free parameters determined by the fitting. These fits are shown as dashed lines. The right panel of fig. 9 shows the values of the free parameters as a function of τ. A second set of functions are fit to these free parameters. Together, this allows for the determination of χ as a function of θ and τ for any new lidar scan. The algorithm thus deploys the CNN to compute $\hat{\tau}$ and uses this output to compute χ at each scan angle θ. Dividing the raw lidar signal by χ gives the projected lidar scan.

4.3.3 Prototype Results

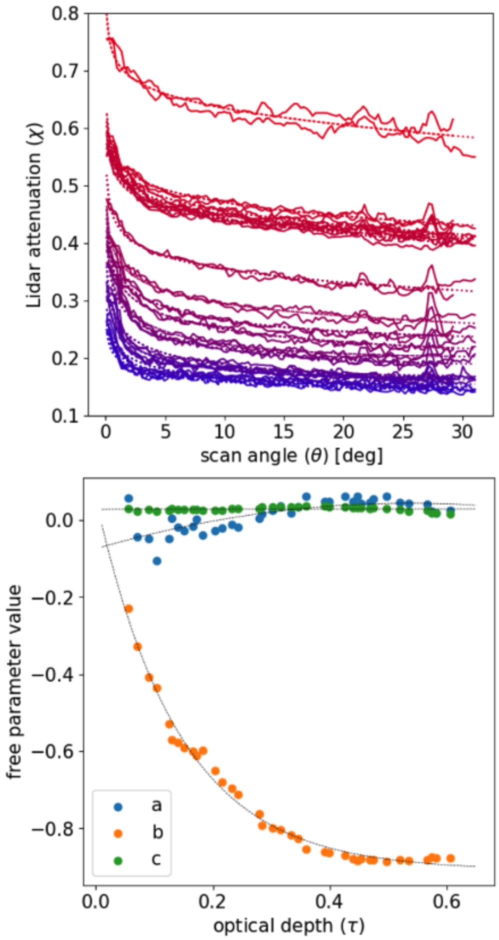

Figure 9: *Top: lidar distance attenuation as function of angle. 0 degrees is looking straight down. Color represents the optical depth of the tank for each scan, ranging from 0 (red) to 0.6 (blue). Solid lines are lidar data, dotted lines are analytic fit. Bottom: eq. (3) free parameters as function of* τ*.*

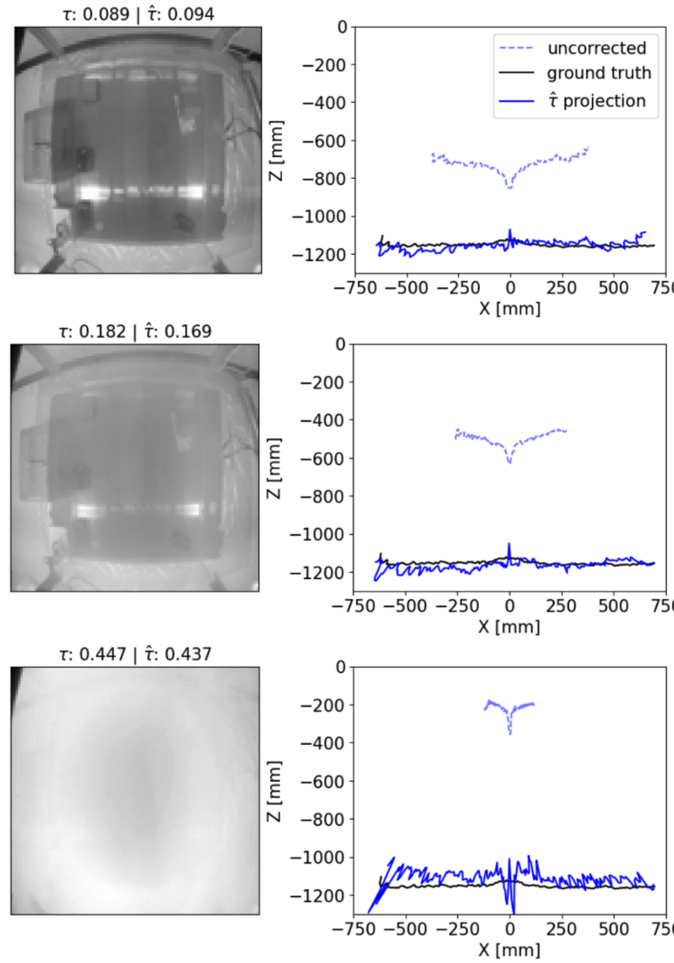

Figure 10: *Projections of lidar scans for three DEHS densities. Left: images from SELENE labeled by the measured optical depth and CNN estimate. Right: Associated lidar scan, showing the raw lidar scan (dashed blue), ground truth geometry (black solid), and projected lidar scan using* $\hat{\tau}$ *(blue solid).*

Figure 10 shows three example lidar reconstructions of SELENE's tank base at different optical depths using this prototype algorithm. The lidar scan corresponds to a slice down the center of the image. In all cases, τ and $\hat{\tau}$ are similar, indicating that the CNN successfully estimates the image optical depth. This is to be expected, given the tight agreement shown in the center panel of fig. 8. The right panels show the associated lidar scans for each image. The raw scan data, shown as the dashed light blue line, incorrectly measures the tank base to be closer to the lidar than it is. The black dashed line shows the true distance, as measured when no DEHS is present. The solid blue lines show the lidar scan projected using the CNN output $\hat{\tau}$ and the lidar attenuation parametrization χ. In all cases, there is good agreement between the ground truth and the projected lidar scan. This indicates that this system can estimate the distance to the tank base even at high optical depths where the lidar scan is significantly impacted by the DEHS.

These promising results suggest that this sensor fusion methodology can provide significant advantages in measuring the distance between the lidar and the target obstacle. Nonetheless, this system relies on multiple limiting assumptions, which could be easily addressed in the near future. For instance, in calculating the true optical depth, the photodiode voltage was assumed to scale linearly with incident intensity. The optical depth between the laser source and photodiode was further assumed to be identical to the optical depth of the camera image. Correcting these assumptions would improve the accuracy of the optical depth ground truth values used for CNN training, and hence the CNN output. At present, due to the simple analytic fit used to determine χ, the lidar projection is only reliable for the same geometry it was fit to. Distance estimation error begins to arise when the height of an obstacle is a significant fraction of the distance between the lidar and the target since the tops of these obstacles are therefore at different optical depths. Using a second CNN to predict χ would remove the limitation of requiring a consistent surface geometry for the lidar projection. Modifying the CNN to independently predict τ in different sectors of the image would also help build resilience to this issue. As identified in sect. 4.3.1, the DEHS droplets are particularly effective in backscattering the lidar signal. To improve the fidelity of this prototype, experiments with lofted regolith or droplets of a smaller refractive index (and hence a lower backscattering efficiency) would be valuable.

This prototype demonstrates how ARC-LIGHT will fulfill its three objectives. By reconstructing a compromised lidar scan using camera imagery, this system allows for an independent measurement of the sensor's position relative to the surface and the identification of its geometry. As outlined in sect. 3.4, this may be used for state vector calibration or IMU redundancy. The proposed addition of a CNN for lidar attenuation estimation will further enhance this system's ability to resolve small surface hazards and help decrease the systematic error near nadir angles. On a technical level, it demonstrates how such a system may operate with the limited computational resources present on the lander; while the CNN training is slow, it is extremely fast to query the trained network to determine optical depth and project the lidar points, taking a fraction of a second on a standard laptop.

5 Technical Management

5.1 Risk Analysis

The development and deployment of ARC-LIGHT carries inherent risks. Here, we review the primary risks related to project timeline, budget, and outcomes, and discuss measures to mitigate them.

5.1.1 Algorithm Architecture Revisions

Developing a robust and accurate ML algorithm for PSI mitigation is a complex task. With limited lunar regolith data available, the algorithm might struggle to learn the nuances of PSI phenomena. This could lead to inaccurate performance and an inability to effectively mitigate PSI risks during lunar landings, necessitating significant revisions to the algorithm design, causing delays and exceeding project budgets. Our synthetic data generation approach, which will be developed in parallel to the algorithm architecture, helps to mitigate this risk. By establishing a flexible simulation framework and data generation pipeline, we can modify the training data type to suit any needed changes in the algorithm. Initial trade studies and algorithm prototyping will also help reveal any significant challenges early in the project.

5.1.2 Algorithm Errors

An inherent challenge stems from the nature of ARC-LIGHT as a convolutional neural network. Unlike traditional, rule-based systems designed by humans, neural networks cannot be easily "debugged" by hand. While extensive training aims to produce the desired outputs, unforeseen behavior or errors might still occur in rare cases. To mitigate this, we will establish safeguards and checks that can override the system's output in such scenarios. Methods that inherently "bound" the network's outputs may also be used, to ensure that problematic outputs are not returned to the GNC. An example of this is our use of a sigmoid activation layer, as described in sect. 4.2, which intrinsically bounds the CNN output between 0 and 1.

5.1.3 Personnel Expertise

Finding qualified personnel with expertise in both machine learning and spacecraft navigation systems might be challenging. Delays in hiring or a lack of qualified candidates could negatively impact project timelines. The specialized skill set required for this project may necessitate higher salaries or recruitment efforts focused on specific academic or professional backgrounds. To address this challenge, the project prioritizes a modular design for ARC-LIGHT with clear and well-defined interfaces between signal processing modules. This modularity allows specialists in machine learning and GNC systems to work more independently on their respective components, simplifying the hiring process by making it easier to find qualified candidates with expertise in each specific domain.

5.1.4 Hardware & Software Integration

Integrating ARC-LIGHT with existing GNC systems could pose compatibility challenges, given the diversity of existing landers developed for CLPS, Artemis, and other initiatives. By designing the algorithm output to match the native data format of the lander's sensors, our approach aims to minimize the need for extensive modifications to either ARC-LIGHT or the existing GNC software. We have also allocated a significant margin in our schedule to accommodate any delays in spacecraft integration and testing.

5.1.5 Testing & Validation

As with all machine learning systems, large quantities of data are needed for ARC-LIGHT training and testing. Given the paucity of lunar landings, alternate means of gathering this data are required. To address this, the project employs a combination of experimental and synthetic data generation. Experimental campaigns using a higher fidelity testbed will help supply realistic scattering data for the cameras and lidars, and extensive synthetic data will cover the broad parameter range of conditions the lander may encounter. Additionally, the performance of ARC-LIGHT integrated into a lander will be thoroughly validated through ground and in-flight testing. Since ARC-LIGHT imposes no additional payload burden, the software can be deployed for its initial flight in a non-active mode, to verify the system without adding risk to the landing. This would involve the system performing its sensor fusion calculations without returning output to the GNC, allowing for post-flight verification of outputs before implementing it as a live system for following landings.

5.1.6 Signal Processing Computational Cost and Reliability

Each signal processing technique employed by the algorithm carries the risk of failing in the lunar environment or demanding a high computational cost. For instance, the two biggest problems with DCP-based dehazing are categorized into obvious and non-obvious failure. Obvious failure is when pixels appear brighter than the ambient light in an image. Non-obvious failure results from pixels being significantly brighter and standing out from their neighboring pixels. Thus, dehazing is difficult in bright environments, or in images that are brightly colored or mostly white, such as the moon.

To rectify obvious failure, pixel minimum can be rescaled to ensure bright surfaces do not occur in the same section of an image as pixel minimum. Non-obvious failure can be mitigated by normalizing the minimum pixel value in a manner that appropriately corresponds to the average haziness observed throughout the scene. This modified DCP also performs significantly faster than the original [39]. Heuristics strategies, such as those discussed in sect. 4.2, may also help to improve computational performance to achieve the target cadence.

5.2 Timeline

The process of ARC-LIGHT development and deployment is divided into three key activities. Development begins by assembling the software and dataset generation teams, defining project requirements, and identifying suitable spaces for software development and testbed operations. Preferably, this location would have an existing "dirty" vacuum chamber which would allow the use of regolith simulant without threatening sensitive hardware. Several compatible facilities investigating PSI exist, including one at NASA Marshall [40]. The testbed experiment should be designed to reflect PSI on the lunar surface as accurately as possible to provide a realistic optical environment for the cameras and lidars. Our work with SELENE provides a basic template of the key features of this facility. The team working on the testbed will be divided into 3 subteams, with 1 for each sensor (camera and lidar), and 1 for the testbed's structure. At the same time, we also study the available PSI models and develop the software pipeline to run large numbers of simulations over a broad parameter range (as detailed in sect. 3.3) and translate them into synthetic sensor data. These tasks are represented in green in fig. 11.

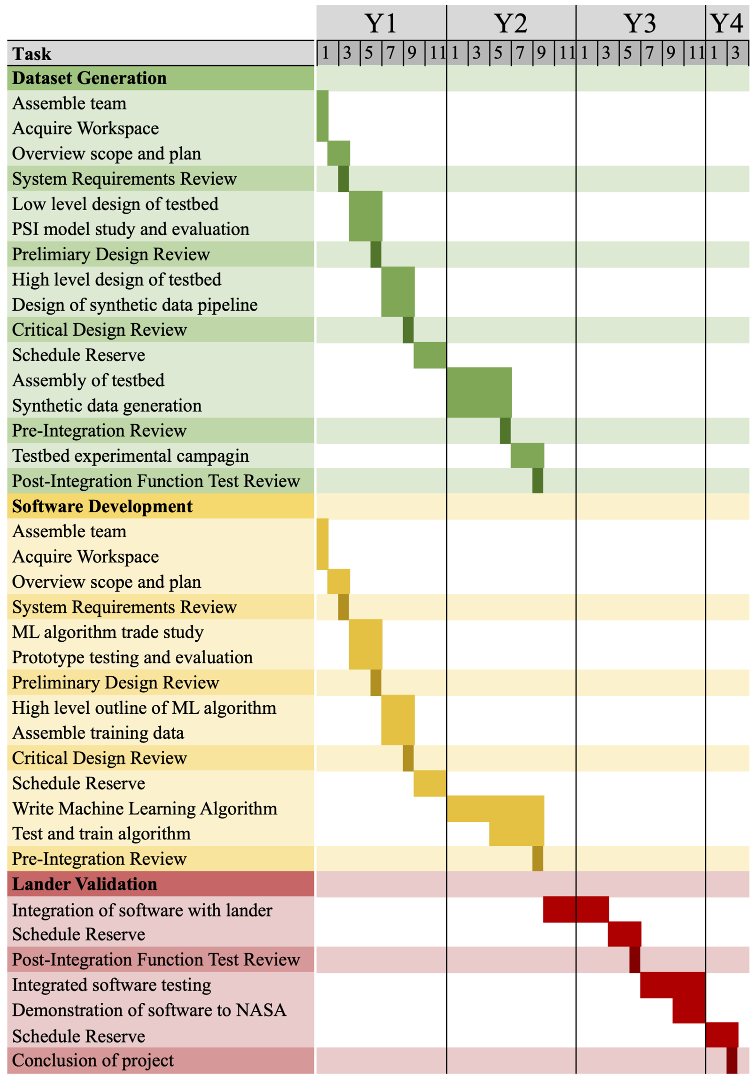

Figure 11: *Timeline organized by year and month.*

The yellow tasks correspond to software development. While the experimental and synthetic datasets are prepared, a team working in parallel begins work on algorithm trade studies, prototyping, and ultimately final version development. The team will be further divided into 3 subteams, 1 assembling training data, 1 developing the code, and 1 verifying the data and the code. This work builds on our prototype development by scrutinizing lander software and hardware constraints to arrive at a more refined algorithm architecture than our candidate design shown in fig. 2. The algorithm team will work closely with the testbed and synthetic data teams to ensure the appropriate data is produced for the selected algorithm design.

Last, the red tasks encompass integration and testing with the lander. As with all "payloads", ARC-LIGHT will be integrated into the spacecraft and tested to ensure successful communication with the spacecraft sensors and GNC. Schedule reserves have been allocated to address any unexpected delays or issues with integration. Once complete, ARC-LIGHT is ready for flight.

5.3 Budget

The project requires a total of 15 employees during its 3 $\frac{1}{4}$ year runtime. The workforce will be divided into 6 subteams, with the 3 testbed subteams having a lead engineer, junior engineer, and technician and the 3 software subteams having a lead engineer and technician. 5 of the lead engineers will be needed for the entire project, with an estimated total workload of 780 full-time equivalent (FTE) weeks. The 6th lead engineer will only be needed for 21 months of the project, with an estimated workload of 84 FTE. The 3 junior engineers are needed for 21 months of the project, with an estimated total workload of 252 FTE. The 6 technicians are only needed for 18 months of the project, with an estimated total workload of 432 FTE. The total estimated cost for the combined workload of the workforce required for this project is 1548 FTE.

To develop the high-fidelity lunar testbed, several hardware items are required. The vacuum chamber structure is estimated to cost \$150,000, the sensors have \$6,900 allocated to them, the small rocket thruster is estimated to cost \$50,000, and the regolith simulant will cost \$3,250. A total of \$10,000 has been allocated for extra hardware that may be required for the project. The total estimated cost for the hardware used for this project is \$220,150. These costs could be reduced if a suitable existing vacuum chamber can be accessed.

High-grade computers will be necessary for the creation of the ARC-LIGHT software. 4 computers would most likely be needed, totaling an estimated \$20,000. 3 years of cloud computing access would be necessary for conducting the PSI simulations for synthetic data generation and CNN training, totaling an estimated \$30,000. A total of \$5,000 has been allocated for any additional software costs. The total estimated cost for the software development is \$55,000. The combined estimated cost of the ARC-LIGHT project is \$275,150 plus 1,548 FTE weeks, approximately amounting to \$3,300 k depending on salary costs.

Category	Cost				Notes
	Number	Unit	Unit Cost	Final Cost	
A. Salaries			(FTE Weeks)	(FTE Weeks)	
Lead Engineers	5	Employees	156	780	Needed for duration of project
Lead Engineer (Training Data subteam)	1	Employees	84	84	Only needed until testbed is finished with assembly
Junior Engineers	3	Employees	84	252	
Technicians	6	Employees	72	432	Hired at the beginning of PDR, and only needed until
Salaries Total				**1,548**	
B. Hardware			(\$)	(\$)	
Lidar	3	N/A	\$2,000	\$6,000	One for prototyping, one for testing, one for backup
Camera	3	N/A	\$300	\$900	
Thruster	1	N/A	\$50,000	\$50,000	Replicates thruster on lander
Testbed & Vacuum Chamber Structure	1	N/A	\$150,000	\$150,000	Contains experiment
Regolith Simulant	50	Kg	\$65	\$3,250	Propelled by thruster
Additional Hardware	1	N/A	\$10,000	\$10,000	Additional equipment needed for assembly and testing
Hardware Total				**\$220,150**	
C. Software					
Computers	4	N/A	\$5,000	\$20,000	Needed to create ML algorithm and run tests
Cloud Space	3	Years	\$10,000	\$30,000	Storage Space and Power
Additional Software	1	N/A	\$5,000	\$5,000	Additional software needed to create algorithm
Software Total				**\$55,000**	

Total Cost	FTE	Dollars
	1,548	**\$275,150**

Figure 12: *Budget estimating the total cost of the development of ARC-LIGHT.*

References

[1] James McDivitt. Apollo 15 Mission Report. Technical report, NASA Manned Spacecraft Center, Houston, 1971.

[2] Shuang Li, Xiuqiang Jiang, and Ting Tao. Guidance summary and assessment of the Chang'e-3 powered descent and landing. *Journal of Spacecraft and Rockets*, 53(2):258–277, 2016. ISSN 15336794. doi: 10.2514/1.A33208.

[3] Luca Ghilardi and Roberto Furfaro. Image-Based Lunar Hazard Detection in Low Illumination Simulated Conditions via Vision Transformers. *Sensors*, 23(18), 2023. ISSN 14248220. doi: 10.3390/s23187844.

[4] Shuang Li, Xiuqiang Jiang, and Ting Tao. Guidance summary and assessment of the Chang'e-3 powered descent and landing. *Journal of Spacecraft and Rockets*, 53(2):258–277, 2016. ISSN 15336794. doi: 10.2514/1.A33208.

[5] Honghua Zhang, Ji Li, Zeguo Wang, and Yifeng Guan. Guidance Navigation and Control for Chang'E-5 Powered Descent. *Space: Science and Technology (United States)*, 2021, 2021. ISSN 26927659. doi: 10.34133/2021/9823609.

[6] Joel Getchius, Devin Renshaw, Daniel Posada, Troy Henderson, Lillian Hong, Shen Ge, and Giovanni Molina. Hazard Detection And Avoidance For The Nova-C Lander. Technical report, 4 2022. URL http://arxiv.org/abs/2204.00660.

[7] JAXA. SLIM Project Press Kit, 2024. URL https://global.jaxa.jp/countdown/slim/SLIM-mediakit-EN_2310.pdf.

[8] James McDivitt. Apollo 12 Mission Report. Technical Report March, NASA Manned Spacecraft Center, Houston, 1970.

[9] Christopher Immer, John Lane, Philip Metzger, and Sandra Clements. Apollo video photogrammetry estimation of plume impingement effects. *Icarus*, 214(1):46–52, 7 2011. ISSN 00191035. doi: 10.1016/j.icarus.2011.04.018.

[10] Huazan Zhong, Hao Wang, Zhengrong Wu, Chen Zhang, Yongwei Zheng, and Tao Tang. A survey of LiDAR and camera fusion enhancement. *Procedia Computer Science*, 183:579–588, 2021. ISSN 18770509. doi: 10.1016/j.procs.2021.02.100. URL https://doi.org/10.1016/j.procs.2021.02.100.

[11] Ali Afzalaghaeinaeini, Jaho Seo, Dongwook Lee, and Hanmin Lee. Design of Dust-Filtering Algorithms for LiDAR Sensors Using Intensity and Range Information in Off-Road Vehicles. *Sensors*, 22(4051):1–13, 2022.

[12] Sadegh Mirshekarian and Dusan Sormaz. Machine Learning Approaches to Learning Heuristics for Combinatorial Optimization Problems. In *Procedia Manufacturing*, volume 17, pages 102–109. Elsevier B.V., 2018. doi: 10.1016/j.promfg.2018.10.019.

[13] Douglas Fontes, James G. Mantovani, and Philip Metzger. Numerical estimations of lunar regolith trajectories and damage potential due to rocket plumes. *Acta Astronautica*, 195 (June 2021):169–182, 2022. ISSN 00945765. doi: 10.1016/j.actaastro.2022.02.016. URL https://doi.org/10.1016/j.actaastro.2022.02.016.

[14] J.E. Lane and P.T. Metzger. Ballistics Model for Particles on a Horizontal Plane in a Vacuum Propelled by a Vertically Impinging Gas Jet. *Particulate Science and Technology*, 30(2):196–208, 2012. doi: 10.1080/02726351.2011.554968.

[15] Peter Liever, Abhijit Tosh, Robert Arslanbekov, and Sami Habchi. Modeling of Rocket Plume Impingement Flow and Debris Transport in Lunar Environment. In *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, Nashville, 2012.

[16] Aaron B. Morris, David B. Goldstein, Philip L. Varghese, and Laurence M. Trafton. Approach for Modeling Rocket Plume Impingement and Dust Dispersal on the Moon. *Journal of Spacecraft and Rockets*, 52(2):362–374, 2015. URL https://doi.org/10.1016/j.actaastro.2018.12.024.

[17] Xiaoying He, Bijiao He, and Guobiao Cai. Simulation of rocket plume and lunar dust using DSMC method. *Acta Astronautica*, 70:100–111, 2012. ISSN 00945765. doi: 10.1016/j.actaastro.2011.07.014. URL http://dx.doi.org/10.1016/j.actaastro.2011.07.014.

[18] Yang Li, Depeng Ren, Zhigang Bo, Wei Huang, Qing Ye, and Yuhong Cui. Gas-particle two-way coupled method for simulating the interaction between a rocket plume and lunar dust. *Acta Astronautica*, 157(September 2018):123–133, 2019. ISSN 00945765. doi: 10.1016/j.actaastro.2018.12.024. URL https://doi.org/10.1016/j.actaastro.2018.12.024.

[19] Arun K. Chinnappan, Rakesh Kumar, and Vaibhav K. Arghode. Modeling of dusty gas flows due to plume impingement on a lunar surface. *Physics of Fluids*, 33(5), 2021. ISSN 10897666. doi: 10.1063/5.0047925.

[20] Hui Tang and Kui Jia. A New Benchmark: On the Utility of Synthetic Data with Blender for Bare Supervised Learning and Downstream Domain Adaptation. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2023-June, pages 15954–15964. IEEE Computer Society, 2023. ISBN 9798350301298. doi: 10.1109/CVPR52729.2023.01531.

[21] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind Apple. Hypersim: A Photorealistic Synthetic Dataset for Holistic Indoor Scene Understanding. In *International Conference on Computer Vision*, pages 10892–10902, 2021. URL http://github.com/apple/ml-hypersim.

[22] Ibrahem Kandel and Mauro Castelli. How deeply to fine-tune a convolutional neural network: A case study using a histopathology dataset. *Applied Sciences (Switzerland)*, 10(10), 5 2020. ISSN 20763417. doi: 10.3390/APP10103359.

[23] Manuel Eugenio Morocho Cayamcela and Wansu Lim. Fine-tuning a pre-trained Convolutional Neural Network Model to translate American Sign Language in Real-time. In *2019 International Conference on Computing, Networking and Communications, ICNC 2019*, pages 100–104. Institute of Electrical and Electronics Engineers Inc., 4 2019. ISBN 9781538692233. doi: 10.1109/ICCNC.2019.8685536.

[24] Mike Horton. How to Easily Estimate Vehicle Localization Errors from IMU Specs, 3 2023. URL https://medium.com/anello-photonics/how-to-easily-estimate-vehicle-localization-errors-from-imu-specs-31e6174ff065.

[25] Honeywell. Compare Our Inertial Measurement Units. URL https://aerospace.honeywell.com/us/en/products-and-services/product/hardware-and-systems/sensors/inertial-measurement-units.

[26] Jianing Song, Duarte Rondao, and Nabil Aouf. Deep learning-based spacecraft relative navigation methods: A survey. *Acta Astronautica*, 191(July 2021):22–40, 2022. ISSN 00945765. doi: 10.1016/j.actaastro.2021.10.025. URL https://doi.org/10.1016/j.actaastro.2021.10.025.

[27] Haibin Liu, Chao Wu, and Huanjie Wang. Real time object detection using LiDAR and camera fusion for autonomous driving. *Scientific Reports*, 13(1):1–12, 2023. ISSN 20452322. doi: 10.1038/s41598-023-35170-z. URL https://doi.org/10.1038/s41598-023-35170-z.

[28] Guotao Xie, Jing Zhang, Junfeng Tang, Hongfei Zhao, Ning Sun, and Manjiang Hu. Obstacle detection based on depth fusion of lidar and radar in challenging conditions. *Industrial Robot*, 48(6):792–802, 2021. ISSN 0143991X. doi: 10.1108/IR-12-2020-0271.

[29] Lili Huang and Matthew Barth. Tightly-coupled LIDAR and computer vision integration for vehicle detection. In *IEEE Intelligent Vehicles Symposium, Proceedings*, pages 604–609, 2009. ISBN 9781424435043. doi: 10.1109/IVS.2009.5164346. URL https://www.researchgate.net/publication/224562486_Tightly-Coupled_LIDAR_and_Computer_Vision_Integration_for_Vehicle_Detection.

[30] Zakaria Chekakta, Abdelhafid Zenati, Nabil Aouf, and Olivier Dubois-Matra. Robust deep learning LiDAR-based pose estimation for autonomous space landers. *Acta Astronautica*, 201:59–74, 12 2022. ISSN 00945765. doi: 10.1016/j.actaastro.2022.08.049. URL https://www.sciencedirect.com/science/article/pii/S0094576522004635.

[31] Genki Ohira, Shuya Kashioka, Yuki Takao, Taketoshi Iyota, and Yuichi Tsuda. Autonomous image-based navigation using vector code correlation algorithm for distant small body exploration. *Acta Astronautica*, 196:400–413, 7 2022. ISSN 00945765. doi: 10.1016/j.actaastro.2020.10.013.

[32] Richard Linares, Tanner Campbell, Roberto Furfaro, and David Edward Gaylor. A Deep Learning Approach For Optical Autonomous Planetary Relative Terrain Navigation. Technical report, 2017. URL https://www.researchgate.net/publication/321534448.

[33] Neil Abcouwer, Shreyansh Daftry, Siddarth Venkatraman, Tyler del Sesto, Olivier Toupet, Ravi Lanka, Jialin Song, Yisong Yue, and Masahiro Ono. Machine Learning Based Path Planning for Improved Rover Navigation. In *2021 IEEE Aerospace Conference*, 11 2021. URL https://hdl.handle.net/2014/54255.

[34] JPL. Under their own AEGIS, 12 2009. URL https://scienceandtechnology.jpl.nasa.gov/under-their-own-aegis.

[35] Guy Webster, Dwayne Brown, and Laurie Cantillo. NASA Mars Rover Can Choose Laser Targets on Its Own, 7 2016. URL https://www.jpl.nasa.gov/news/nasa-mars-rover-can-choose-laser-targets-on-its-own.

[36] R. Lolachi, D. A. Glenar, and T. J. Stubbs. Optical monitoring of the dust environment at lunar surface exploration sites. *Planetary and Space Science*, 234(March 2022):105709, 2023. ISSN 00320633. doi: 10.1016/j.pss.2023.105709. URL https://doi.org/10.1016/j.pss.2023.105709.

[37] Scott Prahl. miepython: Pure python calculation of Mie scattering, 2024.

[38] NASA. Cross-Program Design Specification for Natural Environments. Technical report, 2019. URL https://ntrs.nasa.gov/citations/20200000867.

[39] Irfan Riaz, Xue Fan, and Hyunchul Shin. Single image dehazing with bright object handling. *IET Computer Vision*, 10(8):817–827, 2016. ISSN 17519640. doi: 10.1049/iet-cvi.2015.0451.

[40] Miguel X. Diaz-Lopez, Matthew Gorman, Juan S. Rubio, and Rui Ni. Plume-surface Interaction Physics Focused Ground Test 1: Diagnostics and Preliminary Results. *AIAA*, 1810, 2022. doi: https://doi.org/10.2514/6.2022-1810.