LUXEMBOURG SPACE AGENCY

esric

LUNARLAB

# A LUNAR FOUNDATION MODEL

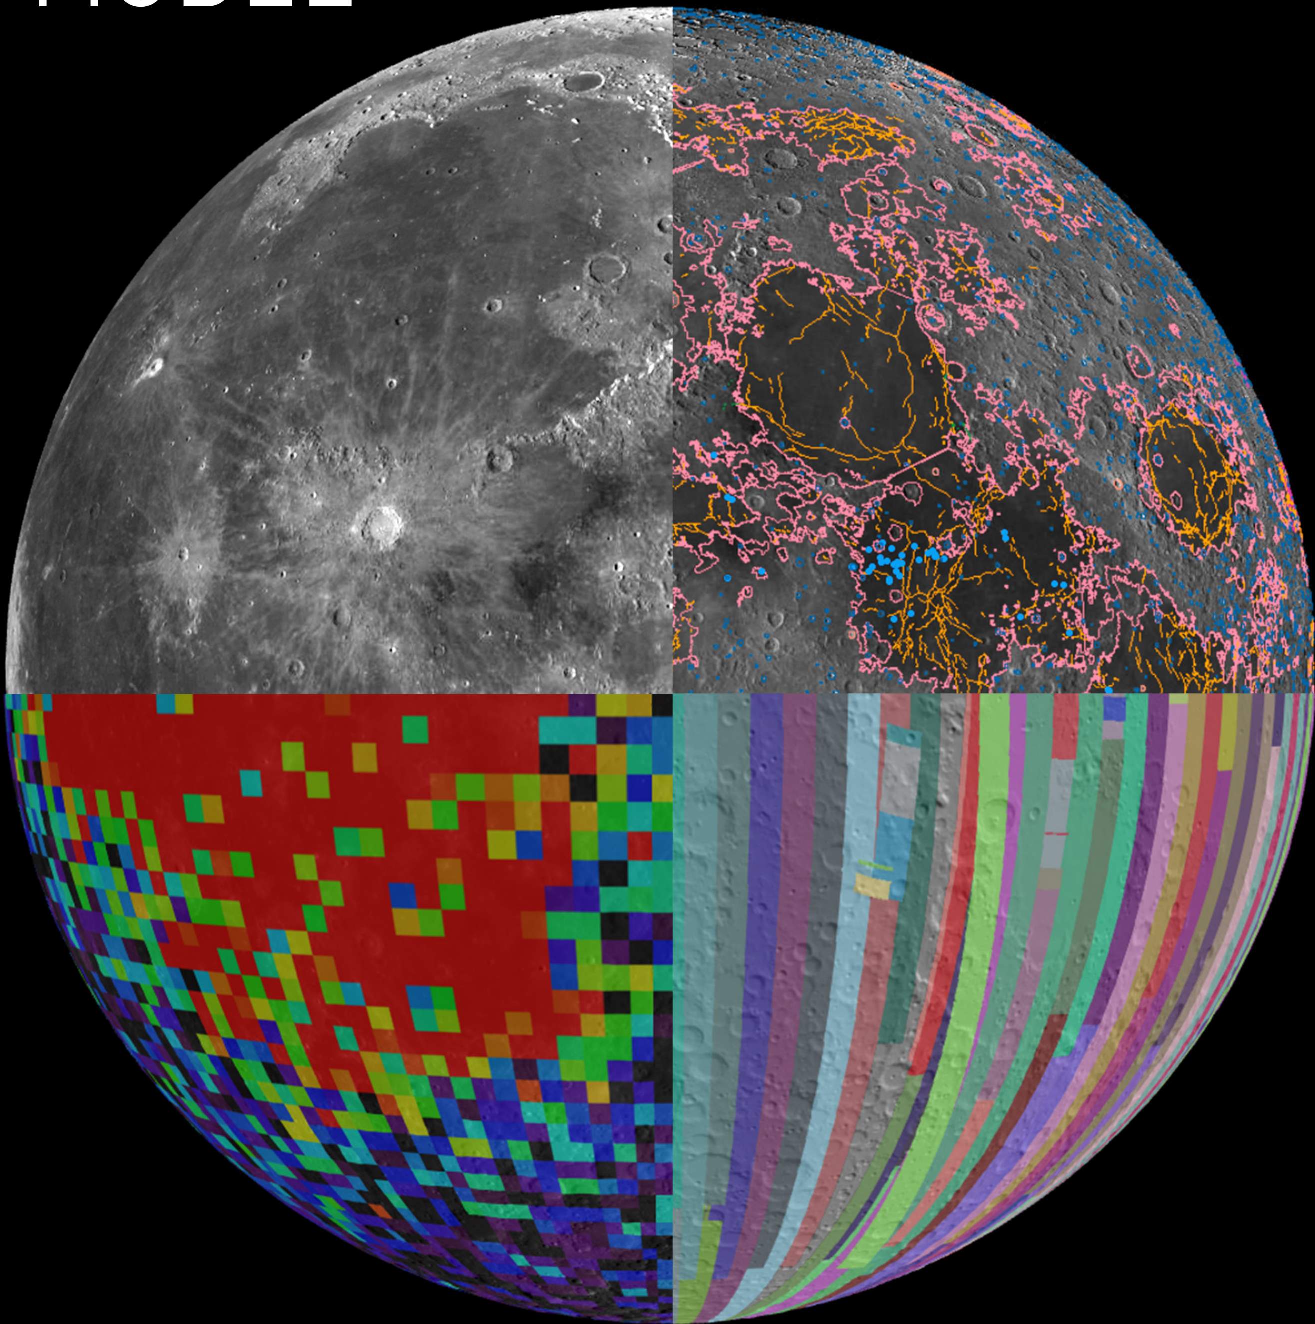

FDL.AI NETWORK

10 YEARS OF APPLIED AI RESEARCH FOR ALL HUMANKIND

Google Cloud NVIDIA SCAN Datarock LUX PROVIDE

AD ASTRA PER ALGORITHMOS

LUNARLAB.AI

TRILLIUM TECHNOLOGIES

# We’ve shrunk the Moon.

# And enabled you (and your robot) to talk to it.

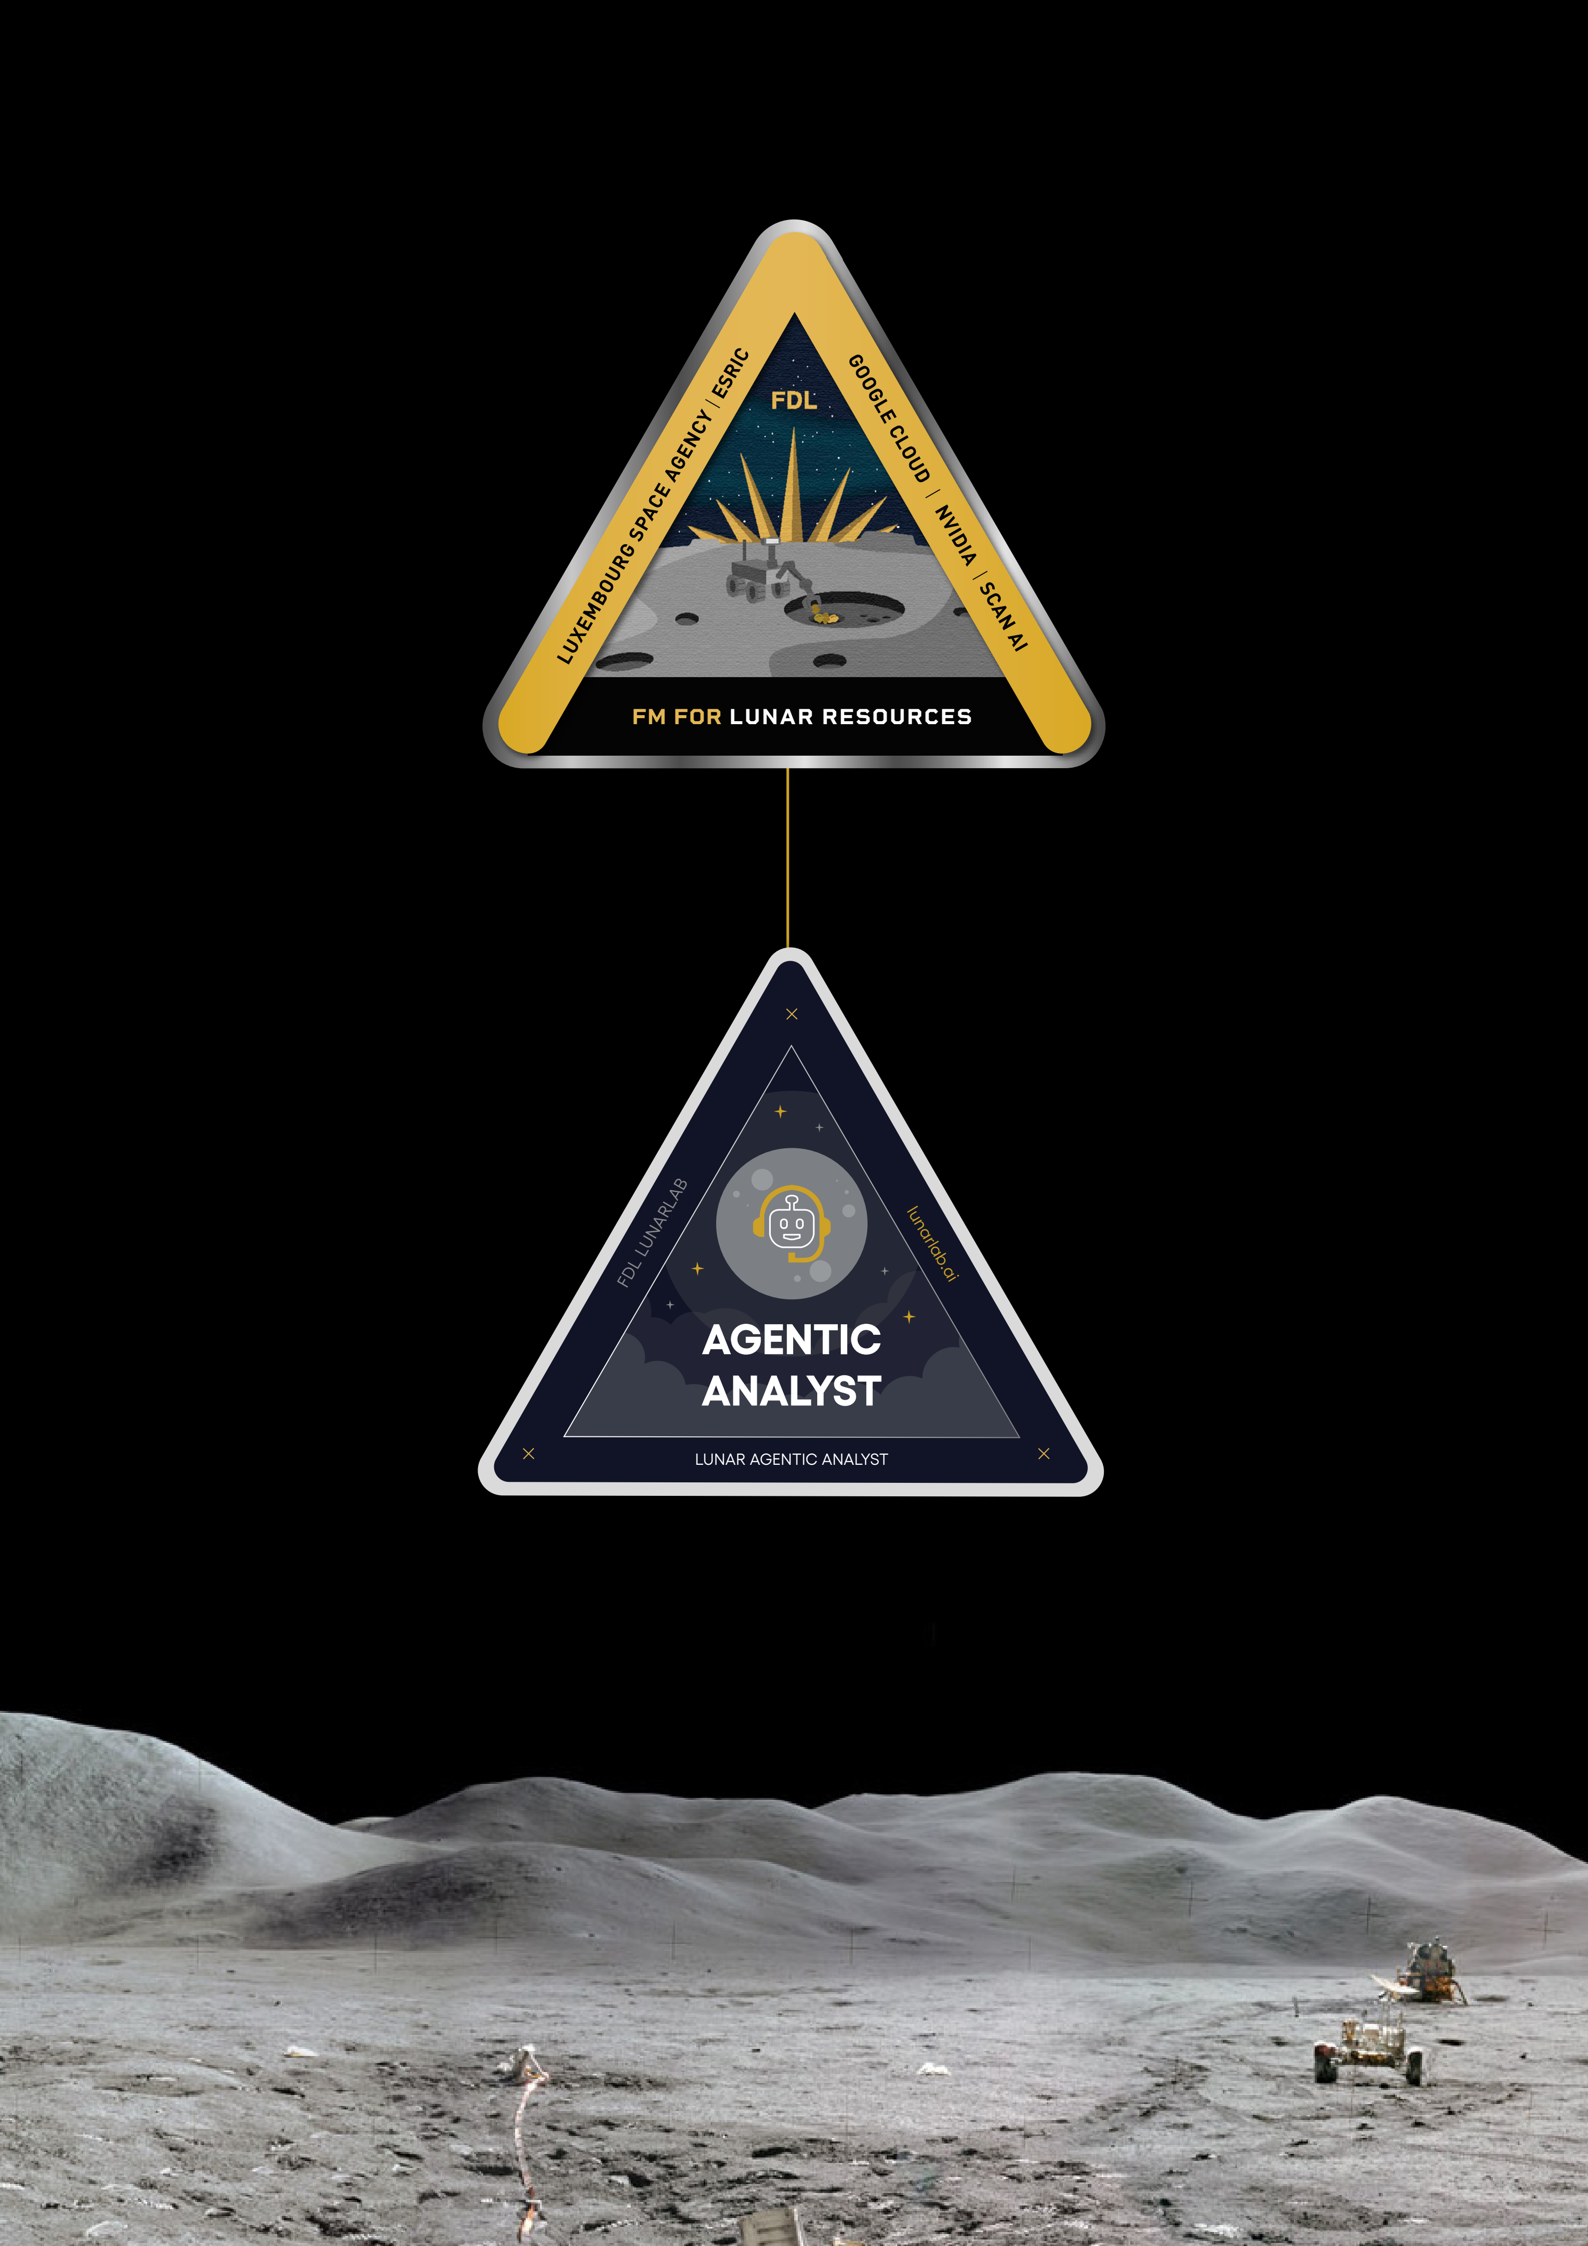
FDL
LUXEMBOURG SPACE AGENCY | ESRIC
GOOGLE CLOUD | NVIDIA | SCAN AI
FM FOR LUNAR RESOURCES
FDL LUNARLAB
lunarlab.ai
AGENTIC
ANALYST
LUNAR AGENTIC ANALYST

# Lunar-FM is humanity's first AI Foundation Model of the Moon.


NASA
esa

# CONTENTS

# THANKS TO OUR PARTNERS

LUXEMBOURG SPACE AGENCY

esric

# ACCELERATING LUNAR EXPLORATION

## Lunar-FM is a 113.3 million parameter multimodal foundation model developed to address the fragmentation and heterogeneity of remote sensing data essential for lunar resource prospecting and scientific analysis.

The model integrates 18 distinct data layers from multiple orbital missions (e.g., LRO, GRAIL, Clementine) spanning modalities: optical imagery, topography, thermal emissions, radar reflectivity, spectroscopy and gravity anomalies.

Lunar-FM is **300x smaller** than the integrated input data, allowing users to run scientific investigations without advanced computing infrastructure.

**Lunar-FM** multi-model data inputs

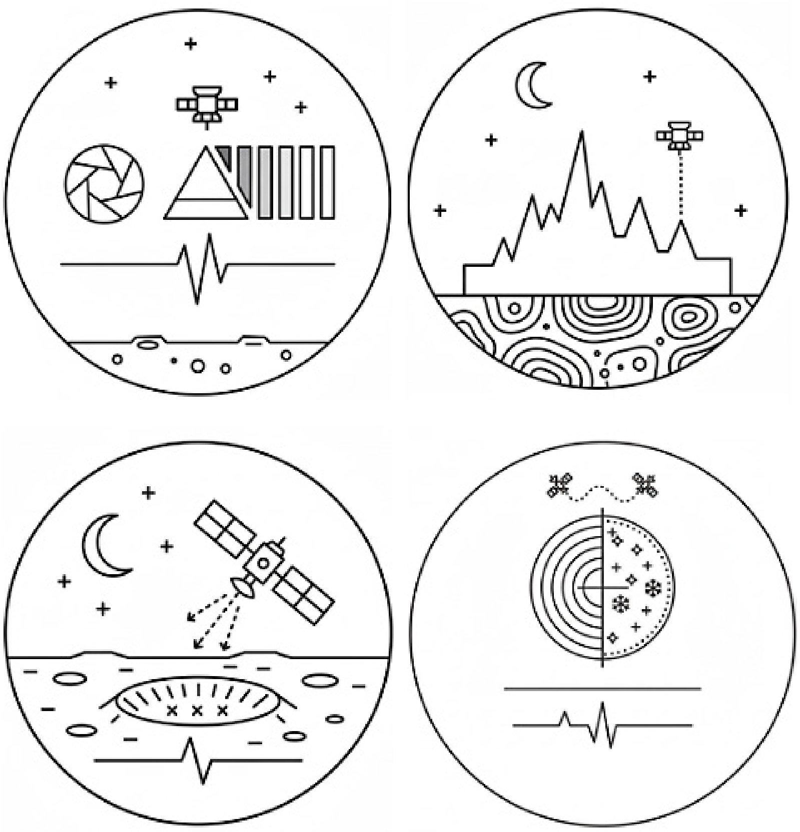

1.Optical and spectroscopic imagery

2.Topograpgy and radar reflectivity

3.Thermal emissions and rock abundance

4.Gravity and volatiles

## Lunar-FM's input data

| Modality | Data Product (Source/ Type) | Resource Identification & Prospecting Utility |
|---|---|---|
| **1. Optical and spectroscopic Imagery** | | Essential for surface composition, elemental abundance, and identifying exposed materials. |
| | **Clementine mineralogy abundance map** (Derived from UVVIS: Ultraviolet-Visible Spectroscopy) | Provides data on the concentration of **Titanium Dioxide,** valuable for extraction of titanium and oxygen. TiO2 is often correlated with Ilmenite, a potential feedstock for oxygen production.<br>Concentration of Iron Oxide (FeO), a primary resource for extracting construction materials and elemental iron. Areas rich in FeO are key targets for basaltic regolith mining.<br>Measures surface brightness. Low albedo can indicate dark, iron/titanium-rich mare basalts, while variations help map soil maturity and exposure of fresh material. |
| | **LROC WAC Global Albedo/Reflectance**<br><br>LROC is of three cameras that make up the Lunar Reconnaissance Orbiter Camera (LROC) instrument suite aboard NASA's Lunar Reconnaissance Orbiter (LRO) spacecraft, which has been orbiting the Moon since 2009.<br><br>Its primary function is to provide the global context and multispectral data of the Moon, complementing the high-resolution, monochrome images captured by the companion Narrow Angle Cameras | Measures surface brightness. Low albedo can indicate dark, iron/titanium-rich mare basalts, while variations help map soil maturity and exposure of fresh material. |
| **2. Topography and Radar Reflectivity** | | Provides structural and environmental context, essential for mission planning, traverse safety, and modeling volatile |
| | **LOLA Global Digital Elevation Model (DEM)**<br><br>The LOLA Global Digital Elevation Model (DEM) is the highest-precision, standardized global map of the Moon's three-dimensional surface shape, or topography.<br><br>It is the foundational data product derived from the Lunar Orbiter Laser Altimeter (LOLA) instrument aboard NASA's Lunar Reconnaissance Orbiter (LRO) spacecraft. | Defines the shape and height of the surface. Crucial for calculating slope, determining line-of-sight for communications, and assessing accessibility of resource sites. |
| | **LOLA Global Slope Map** | Identifies terrain steepness. Prospecting vehicles require areas with low slopes for safe traverse and access to resource deposits. |

| | | |
|---|---|---|
| | **LOLA Global Surface Roughness (RMS Height)** | Measures small-scale variation in elevation. Used for assessing landing site safety and identifying regions where impact gardening is more or less intense. |
| | **Mini-RF S-band (12.6 cm) Circular Polarization Ratio (CPR)** | Key volatile detection product. High CPR in PSRs is a strong proxy signature for the presence of subsurface water ice or blocky scattering within the top few meters of regolith. |
| | **Mini-RF X-band (4.2 cm) Circular Polarization Ratio (CPR)** | Provides a secondary, shorter-wavelength perspective on radar scattering, helping to constrain the depth and distribution of any hypothesized ice deposits or scatterers. |
| | **Mini-RF S-band Backscatter Coefficient** | Measures the overall radar reflection power. Variations are linked to surface roughness and dielectric constant, helping distinguish between different regolith properties. |
| 3. Thermal emission and rock abundance | | Crucial for volatile mapping by indicating surface temperature cycles, which control the stability of water ice, and for mapping |
| | **Diviner Nighttime Minimum Temperature**<br><br>The Diviner Lunar Radiometer Experiment (DLRE), often just called Diviner, is a key instrument aboard NASA's Lunar Reconnaissance Orbiter (LRO). Its primary function is to measure the temperature of the lunar surface, and the Diviner Nighttime Minimum Temperature product is one of its most critical datasets, particularly for lunar resource identification. | **Primary indicator for cold traps.** Permanently Shadowed Regions (PSRs) with extremely low minimum temperatures are the most likely locations for stable water ice deposits (a primary resource).By mapping the absolute minimum temperatures reached in every location on the Moon during the long lunar night, the Diviner data precisely delineates the extent and depth of these cold traps. These cold areas are the highest-priority targets for missions seeking to harvest water ice as a resource for life support and propellant production. |
| | **Diviner Rock Abundance Map** | Measures the fraction of the surface covered by rocks larger than the scale of the soil. High rock abundance can complicate surface operations but is a proxy for material strength and structure |
| | **Diviner Thermal Inertia Map** | Indicates the subsurface composition and grain size. Low thermal inertia suggests fine, dusty regolith, while high inertia can indicate rocks or subsurface compaction. |
| | **Diviner Maximum Subsolar Temperature** | Defines the warmest surface temperature. Used in volatile modeling to determine the maximum thermal stress on any near-surface resources. |
| 4. Gravity & Volatiles | | Provides insight into subsurface structure, crustal thickness, and the global distribution of key volatiles (Hydrogen/Water) |

| | | |
|---|---|---|
| | **GRAIL Bouguer Gravity Anomaly**<br><br>The GRAIL Bouguer Gravity Anomaly map from the Gravity Recovery and Interior Laboratory (GRAIL) mission, which orbited the Moon from 2011 to 2012.<br>This map is essential for understanding the Moon's subsurface structure, which has significant implications for resource identification and geological analysis.<br>The GRAIL mission measured the Moon's gravity field by precisely tracking the minute distance changes between two twin spacecraft as they flew over areas of varying mass.<br>The resulting gravity field data is processed into different types of maps. The Bouguer Gravity Anomaly map is a sophisticated product because it takes the raw gravity measurements and removes the gravitational effects caused by topography (the mountains and valleys measured by LOLA). | Indicates density variations beneath the surface (after removing effects of topography). Used to map subsurface geological features like ancient impact basins and intrusive bodies, which can be associated with mineral-rich terrains. |
| | **GRAIL Crustal Thickness Model** | Constrains the global distribution of the lunar crust. Thickness variations are linked to different geological regimes and may influence the distribution of KREEP-rich (Uranium, Thorium, Potassium, Rare Earth Elements) materials. |
| | **LEND Water Equivalent Hydrogen (WEH) Map** | Direct volatile resource map. Maps the global concentration of Hydrogen (a proxy for water ice) in the upper meter of regolith, particularly at the poles, via neutron spectroscopy. |
| | **(Additional Placeholder Layer - e.g., K, Th, U)** | Elemental maps (e.g., Potassium, Thorium, Uranium) from Lunar Prospector/LRO are essential for identifying regions of elevated radiogenic elements (KREEP), which are resource targets. |

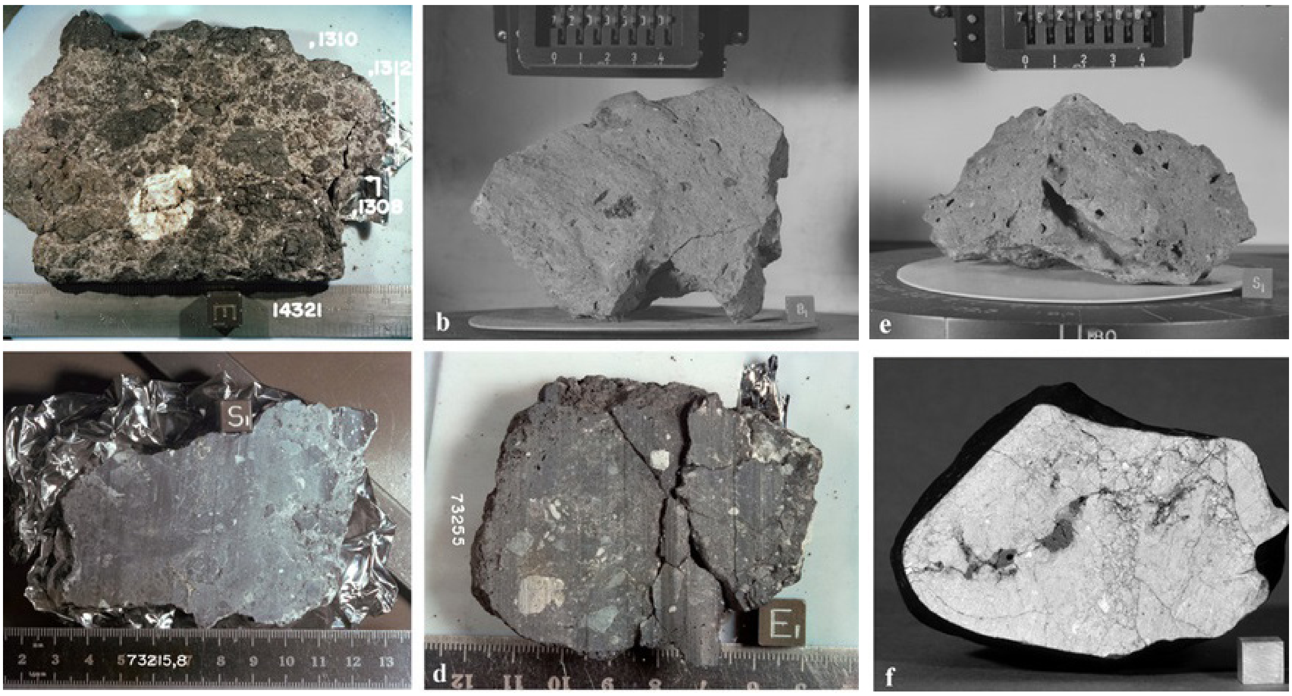

# **LUNAR-FM** TECHNICAL BRIEFING

**Lunar-FM is an AI foundation model** based on a novel application of **self-supervised deep learning** with a primary **focus on lunar resources.**

**Key achievements of Lunar-FM**

1. Multimodal Data Fusion: Lunar-FM is based on a Multi-modal Masked Autoencoder (MultiMAE) architecture with a Vision Transformer (ViT) backbone. This is a significant step in planetary science, as it demonstrates a scalable, self-supervised method for integrating 18 heterogeneous data layers across five modalities (optical, topography, thermal, radar, gravity) into a cohesive model for the first time.
2. Creation of a unified latent feature space: Lunar-FM is a unified, information-dense 768-dimensional latent embedding space that represents global lunar properties at a 0.5$^{\circ}$×0.5$^{\circ}$ resolution. This achievement standardizes data representation, achieving a 300x data compression while retaining rich semantic information required for diverse downstream tasks.
3. Validation of predictive capability: The resulting embeddings were validated as superior features for quantitative analysis, demonstrated by the global identification of titanium dioxide ($TiO_2$) abundance using simple linear models with a low Mean Average Error (MAE of 0.062).
4. Demonstration of few-shot knowledge amplification from scarce samples: Validation of Lunar-FM showed the model's ability to facilitate expert-curated few-shot learning. This demonstrates that high-quality global predictive maps can be generated from an extremely small set of expert-labeled examples (e.g., 8 data points), effectively mitigating the problem of scarce ground-truth data.
5. Integration into an agentic AI interface: Lunar-FM's embedding space was successfully integrated as the core knowledge base within an AI Agent ("Lunar Analyst Copilot") architecture. This establishes a 1.0 system where a Large Language Model (LLM) interprets natural language queries to execute complex data-science tasks via the Lunar-FM feature set.

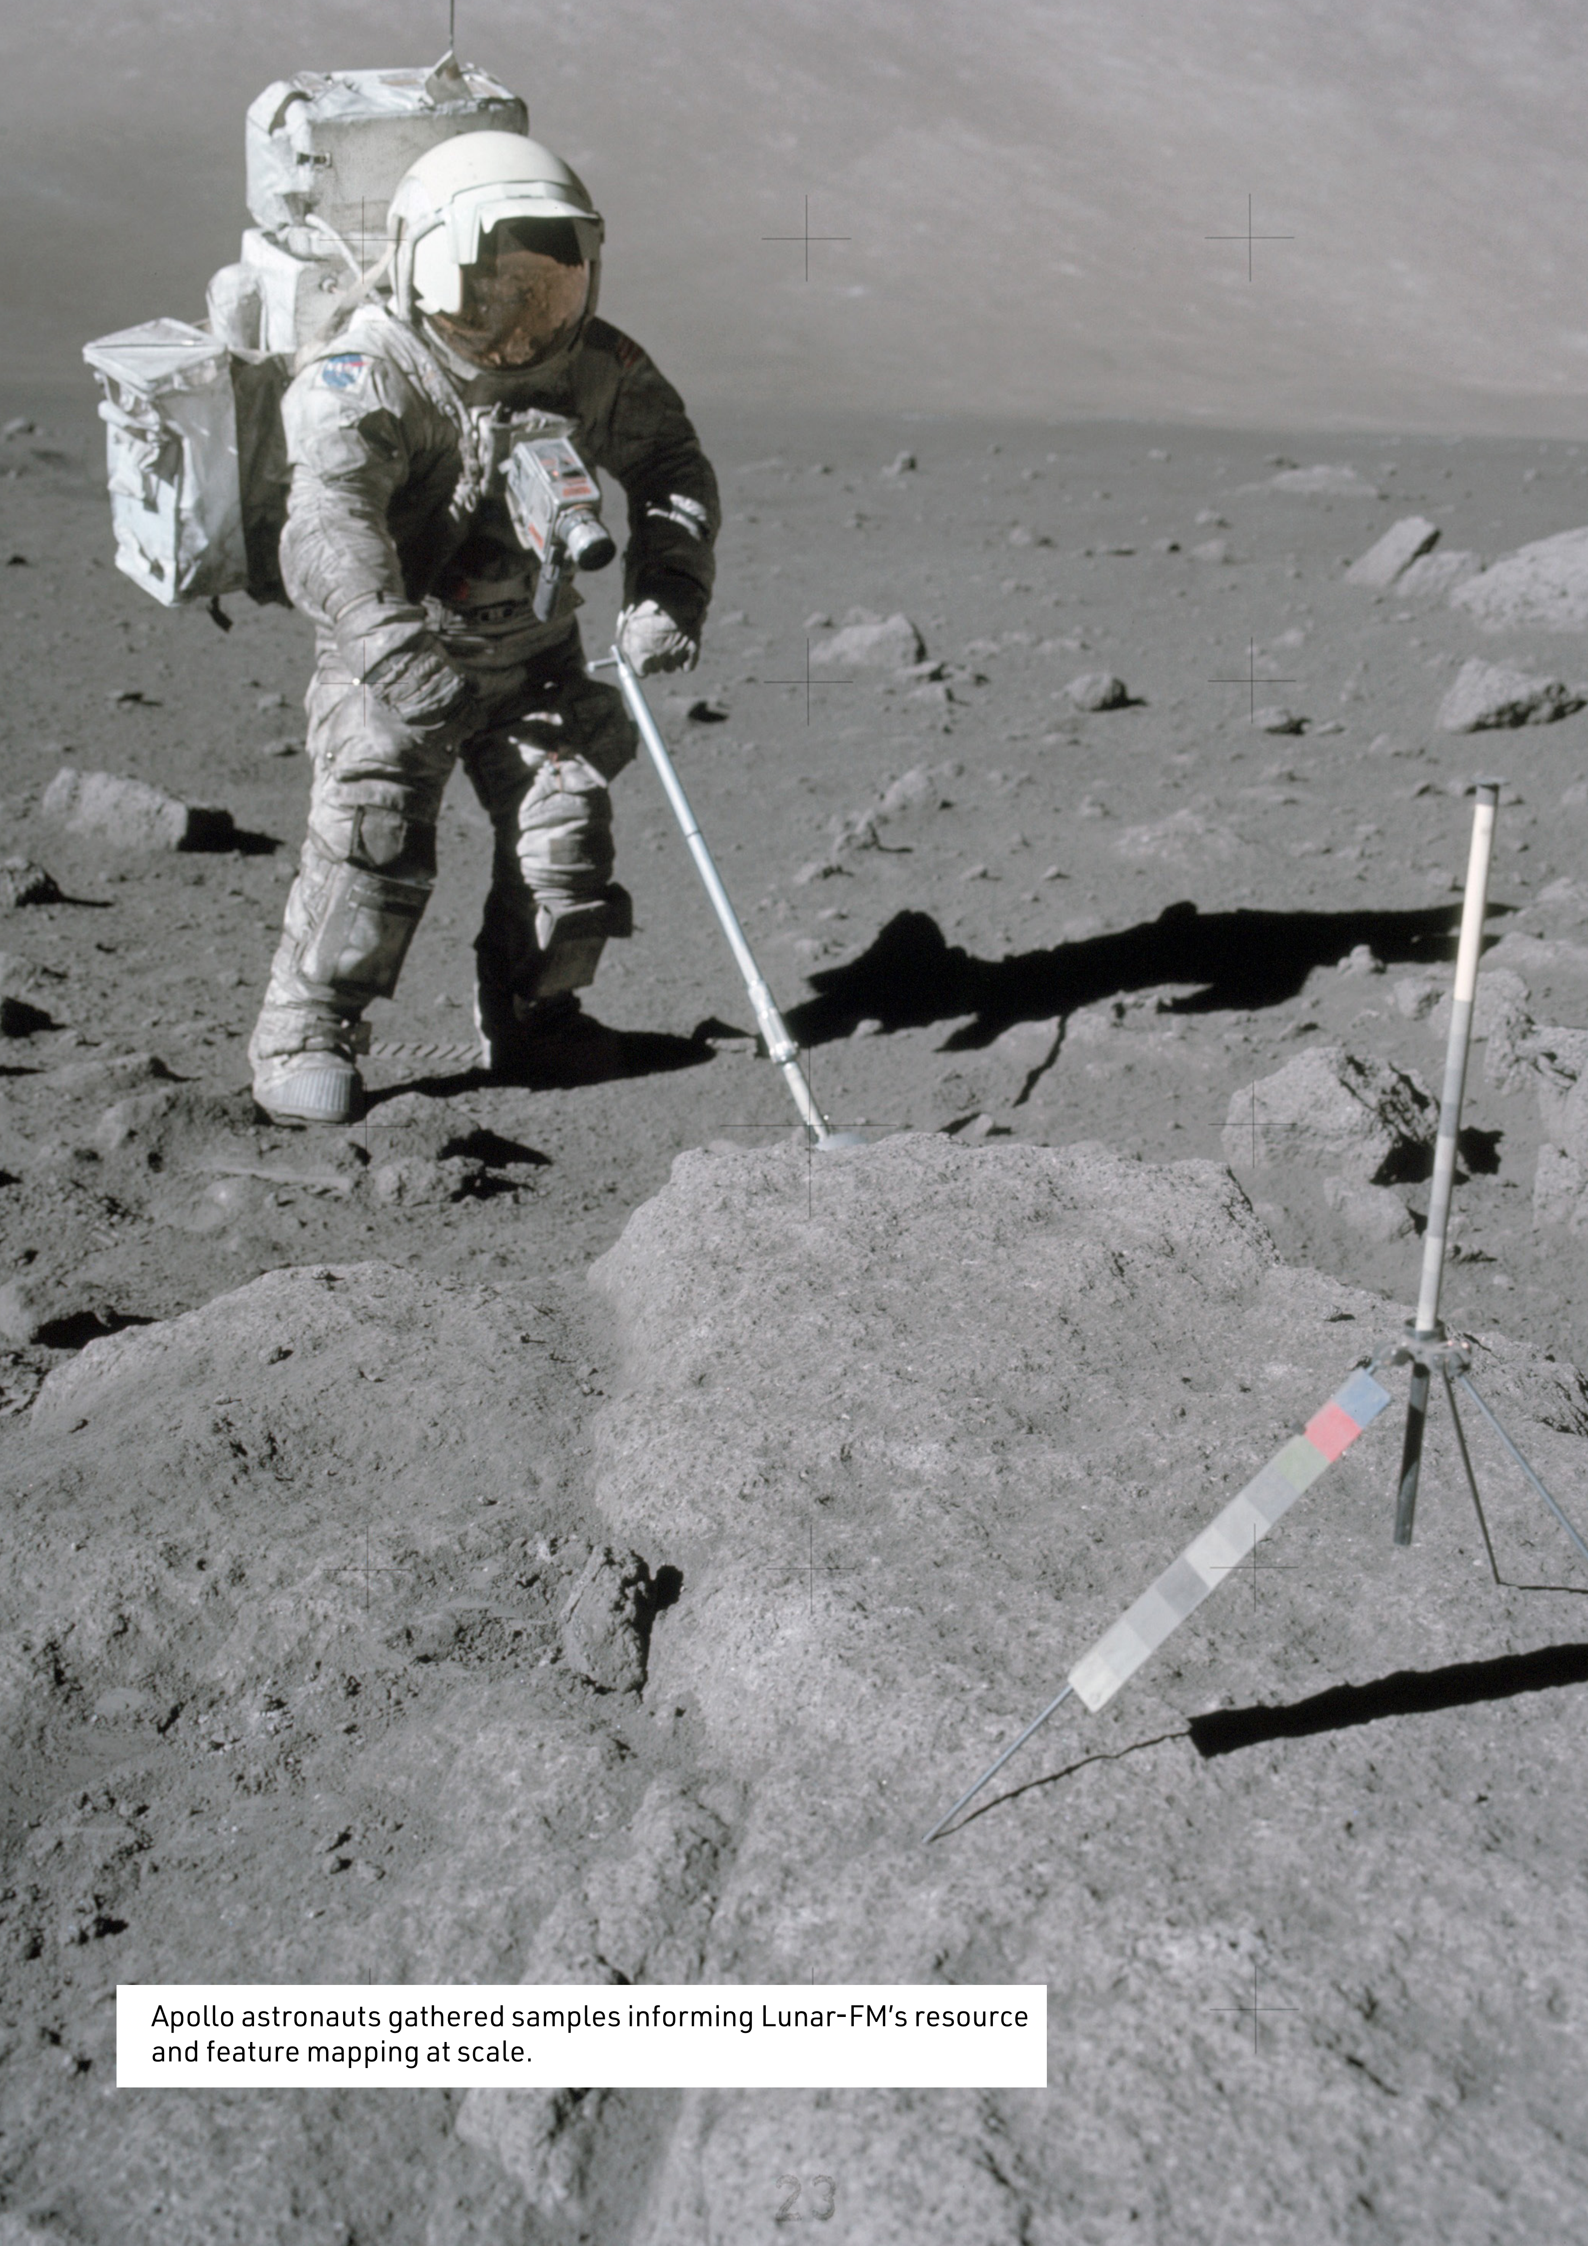

Apollo astronauts gathered samples informing Lunar-FM's resource and feature mapping at scale.

# **LUNAR-FM** ANATOMY

**0.5 ° × 0.5 ° chips from 70 ° S to 70 ° N**

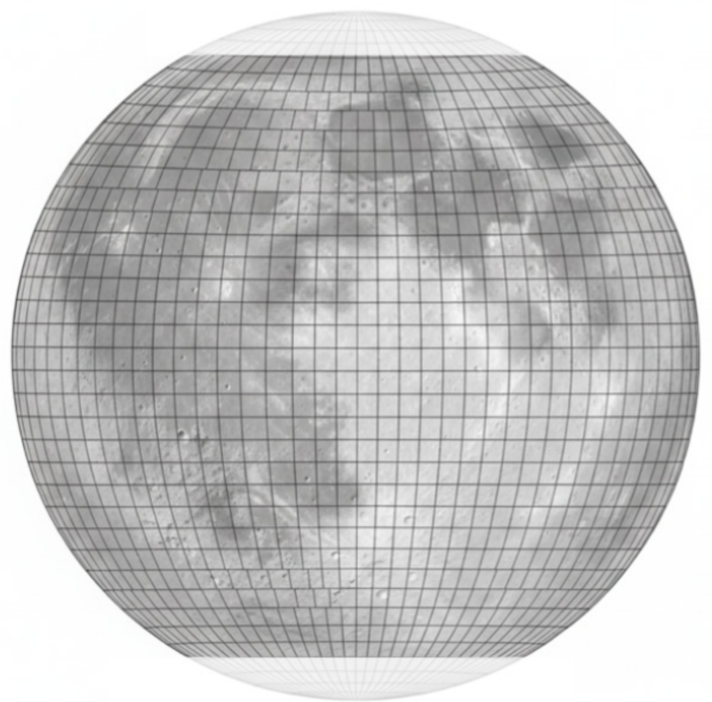

300 x compression

**113.3 Million Parameter Foundation Model**

**Topography / Elevation (DEM)**

**Thermal Data**

**Albedo / Broad Imagery (Visible Light)**

**Spectroscopy / Mineral Data**

**Gravity Anomaly Data**

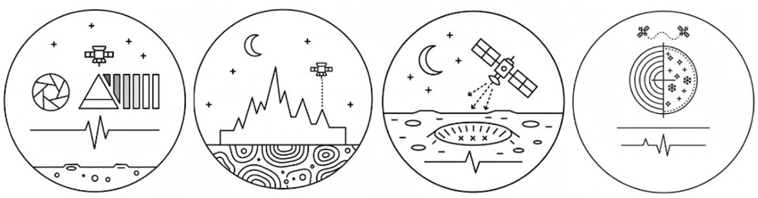

1.Optical and spectroscopic imagery

2.Topograpgy and radar reflectivity

3.Thermal emissions and rock abundance

4.Gravity and volatiles

**768-dimensional unified vector embeddings per chip**

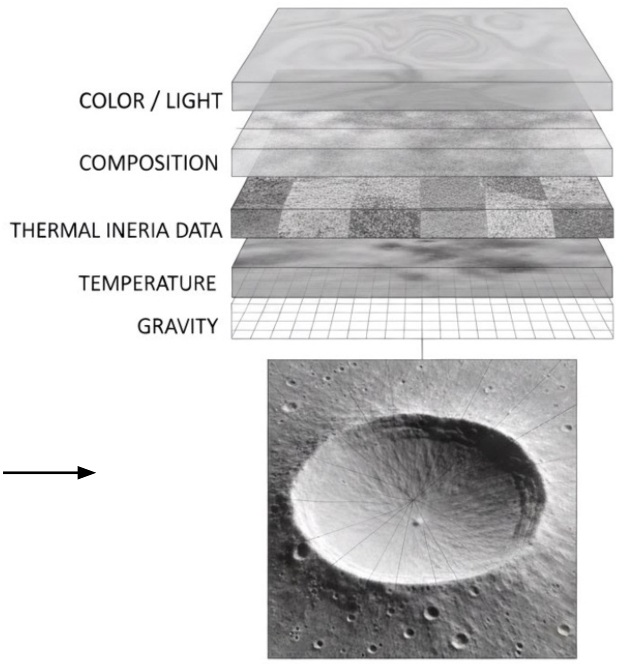


Integrated Data Stack for One Lunar Chip

**Terrain classification**

**Anomaly Detection**

## Generalized downstream use cases

**Rare feature discovery**

**Global mapping based on sparse samples**

**WITHIN EACH CHIP**

1. Information on the elevation of the terrain and the shape of features.
2. How bright or dark the surface is (reflectance) which helps distinguish surface maturity and basic feature boundaries.
3. The chemical and mineral makeup of the surface (e.g., iron, titanium, magnesium content, and specific rock types).
4. How quickly the surface material heats up or cools down, which indicates the size of the surface grains (fine dust vs. rocks).
5. Information about the mass distribution beneath the surface, which relates to subsurface geology and large structures.

# **LUNAR-FM** SAMPLE RESOURCE PROSPECTING PIPELINE

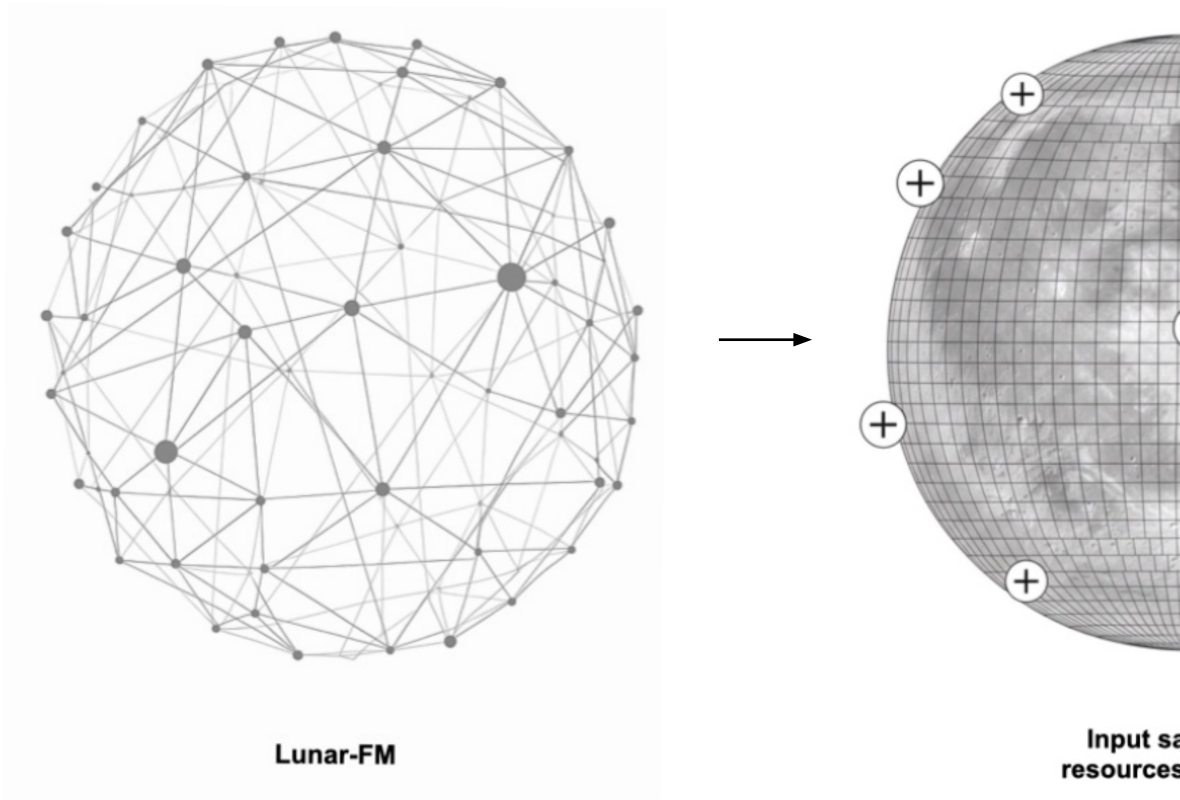

**Foundation models are an important tool for science as they streamline data science investigations from months to days or hours. Precomputed queries can be answered instantly in natural language.**

Foundation models are 'generalizable' - that is, that a single model can be adapted to many different downstream tasks. This capability is known as 'generalizability'. Lunar-FM was validated by showing how it could be adapted to multiple downstream use-cases, unlocking novel capabilities in lunar geology and prospecting.

**For example a global map of TiO2 was built using only 8 validated samples (Diagram below).**

Lunar-FM has case study notebooks available for the community to replicate and improve.

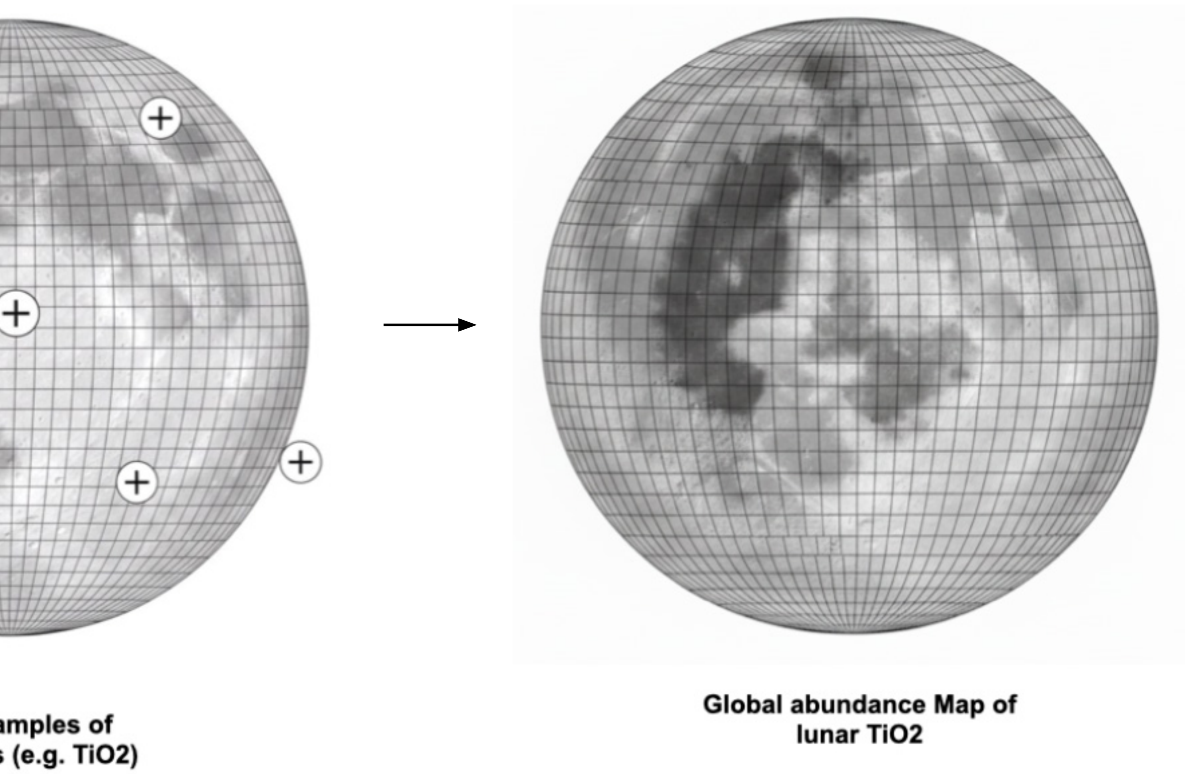


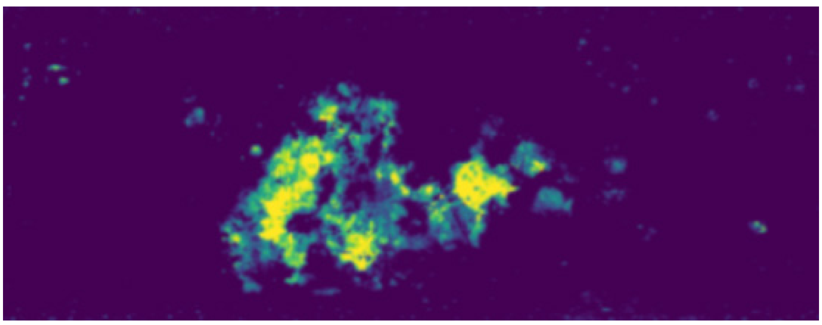


Trace Metallics
Excavated Regolith Metallics
Exogenic Metallics
LUNARLAB
FDL.AI
NETWORK
Giordano Bruno Crater

# Q&A: A LUNAR FOUNDATION MODEL FOR OPERATIONAL SCIENCE WITH LUNAR GEOLOGIST

*Dr. Abigail Calzada Diaz is a leading lunar geologist and R&T Associate at the **European Space Resources Innovation Centre (ESRIC)**, renowned for translating fundamental lunar science into mission-critical applications. Her PhD work established new insights into the **source regions of lunar meteorites**, deepening our understanding of the Moon's crustal composition. Currently, she is a key scientific figure in the burgeoning space resources sector, advising on lunar resource utilization (ISRU) and mission design.*

esric

**DR ABIGAIL CALZADA DIAZ**

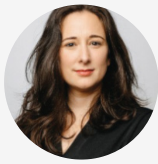

**Question 1:** What is the critical challenge that the lunar foundation model addresses concerning current lunar data synthesis and knowledge extraction?

**ACD:** *The primary issue is the fragmented, multi-source nature of lunar data which until now has made lunar investigations labor-intensive. Lunar-FM provides a standardized, unified data infrastructure enabling knowledge extraction and synthesis from disparate datasets, transcending simple data aggregation because potential insights are now, in a shared embedding space that can be adapted to multiple investigations.*

**Question 2:** What novel capabilities does this tool offer for active scientific inquiry, distinguishing it from general exploration software?

**ACD:** *It moves beyond passive simulation (e.g., Google Moon). Its core value lies in its query-driven capability of real data to facilitate active scientific work, allowing researchers to interrogate the underlying phenomena and derive complex insights not accessible via standard data exploration tools.*

**Question 3:** How does an agentic interface impact the research workflow, particularly for non-computer science experts in planetary geology?

**ACD:** *Being able to talk to the data democratizes advanced analysis by providing easily operable tools for geologists and prospectors bypassing the need for extensive coding expertise. This lowers the barrier to entry for complex data interpretation.*

**Question 4:** In what capacity does Lunar-FM support theoretical lunar science and mission operational planning?

**ACD:** *Lunar-FM is a translational bridge between pure scientific understanding of lunar processes and the operational requirements of missions. It converts theoretical knowledge into science-informed insights and resource-focused targets.*

**Question 5:** Considering future missions, what is the most critical next developmental step for ensuring the Lunar-FM's longevity?

**ACD:** *Continuous inclusion of new observational data from future missions such as Artemis is an exciting possibility. Foundation models shouldn't be static or lone-solutions. The agentic capabilities of Lunar-FM unlock a future where multiple models can work together to support researchers and surface operations.*

# THE LUNAR AGENTIC ANALYST COPILOT

The **Lunar Analyst Copilot** represents the operational and user-facing implementation of the Lunar foundation model. It is an **AI Agent architecture** designed to bridge the gap between complex, multidimensional lunar data and natural language queries by scientists, geologists, and mission planners.

The model's embeddings function as core tools within an AI Agent architecture, which uses a Large Language Model (LLM) to interpret natural language queries. This system routes questions to the appropriate tool (e.g., Similarity Search or Regression Models) and synthesizes a coherent response by integrating the model's numerical outputs with external knowledge from a comprehensive knowledge based of scientific papers, technical reports and mission documents.

**This represents a prototype for future systems allowing scientists and mission planners to interact with complex lunar data through an intuitive conversational interface.**

### Core Architecture and function

The Agentic Analyst's function is centered on providing dynamic, conversational data analysis by leveraging the rich feature space of the Lunar-FM embeddings.

- **Large Language Model (LLM) Interface:** The system utilizes a Large Language Model as its front-end. This allows users to interact with the vast lunar dataset using natural language queries (e.g., *"Where are the best places to land near the South Pole to search for water ice?")*.
- **Query Translation and Routing:** The LLM acts as an interpreter, translating the high-level, human-driven query into a specific sequence of executable **data science tasks**. It understands which underlying data modalities and models are required to synthesize an answer.
- **Tool Composition:** The Agent then autonomously routes the request to the appropriate tool (function) to leverage the Lunar-FM embeddings. .

# LUNAR-FM APPLICATIONS

The capabilities of the Lunar foundation model (Lunar-FM) support a range of critical applications for science, resource utilization, and mission planning:

- **Identify High-Potential Resource Sites** (e.g., ilmenite-rich basalts, FeO-rich regolith, **KREEP** terrains).
    - **KREEP:** Rock or regolith components rich in K (Potassium), R(Rare Earth Elements), and P (Phosphorus). These terrains are of high interest for rare Earth element.
- **Locate and Characterize Polar Volatile Deposits**, including mapping likely water-ice stability zones using thermal data (e.g., Diviner Nighttime Minimum Temperature).
- **Select Optimal Landing Sites** by combining critical topographical and physical constraints: **slope**, surface **roughness** (e.g., RMS height), **rock abundance**, and illumination conditions.
- **Plan Safe and Energetically Efficient Rover Traverses** using integrated topography, roughness, and thermal data (to predict potential component stress).
- **Model Thermal Environments** for lander, rover, and habitat design, considering diurnal (day/night) cycles, **cold traps** (Permanently Shadowed Regions, or **PSRs**), and overall thermal stress on hardware.
- **Assess Regolith Mechanical Properties** (grain size, compaction, blockiness) to support excavation and **ISRU**operations.
    - **ISRU:** In-Situ Resource Utilization—the practice of harvesting and using materials found.
- **Detect Subsurface Geological Structures** (e.g., mascons, buried basins, intrusions) linked to mineral enrichment, primarily using gravity anomaly data (e.g., GRAIL Bouguer maps).
    - **Mascons:** Mass Concentrations—regions beneath the lunar surface with an excess of mass (high density), typically associated with buried impact basin floors.
- **Differentiate Geological Units** (e.g., mare vs. highlands, pyroclastics, fresh vs. mature surfaces) to refine exploration strategies and understand the Moon's geological evolution.
- **Estimate Near-Surface Hydrogen Abundance** (a proxy for water content) for water extraction feasibility studies, primarily using neutron spectrometer data (e.g., LEND).
- **Support Long-Term Infrastructure Planning**, including site selection for power grids, communications relays, mobility corridors, and bulk storage facilities.

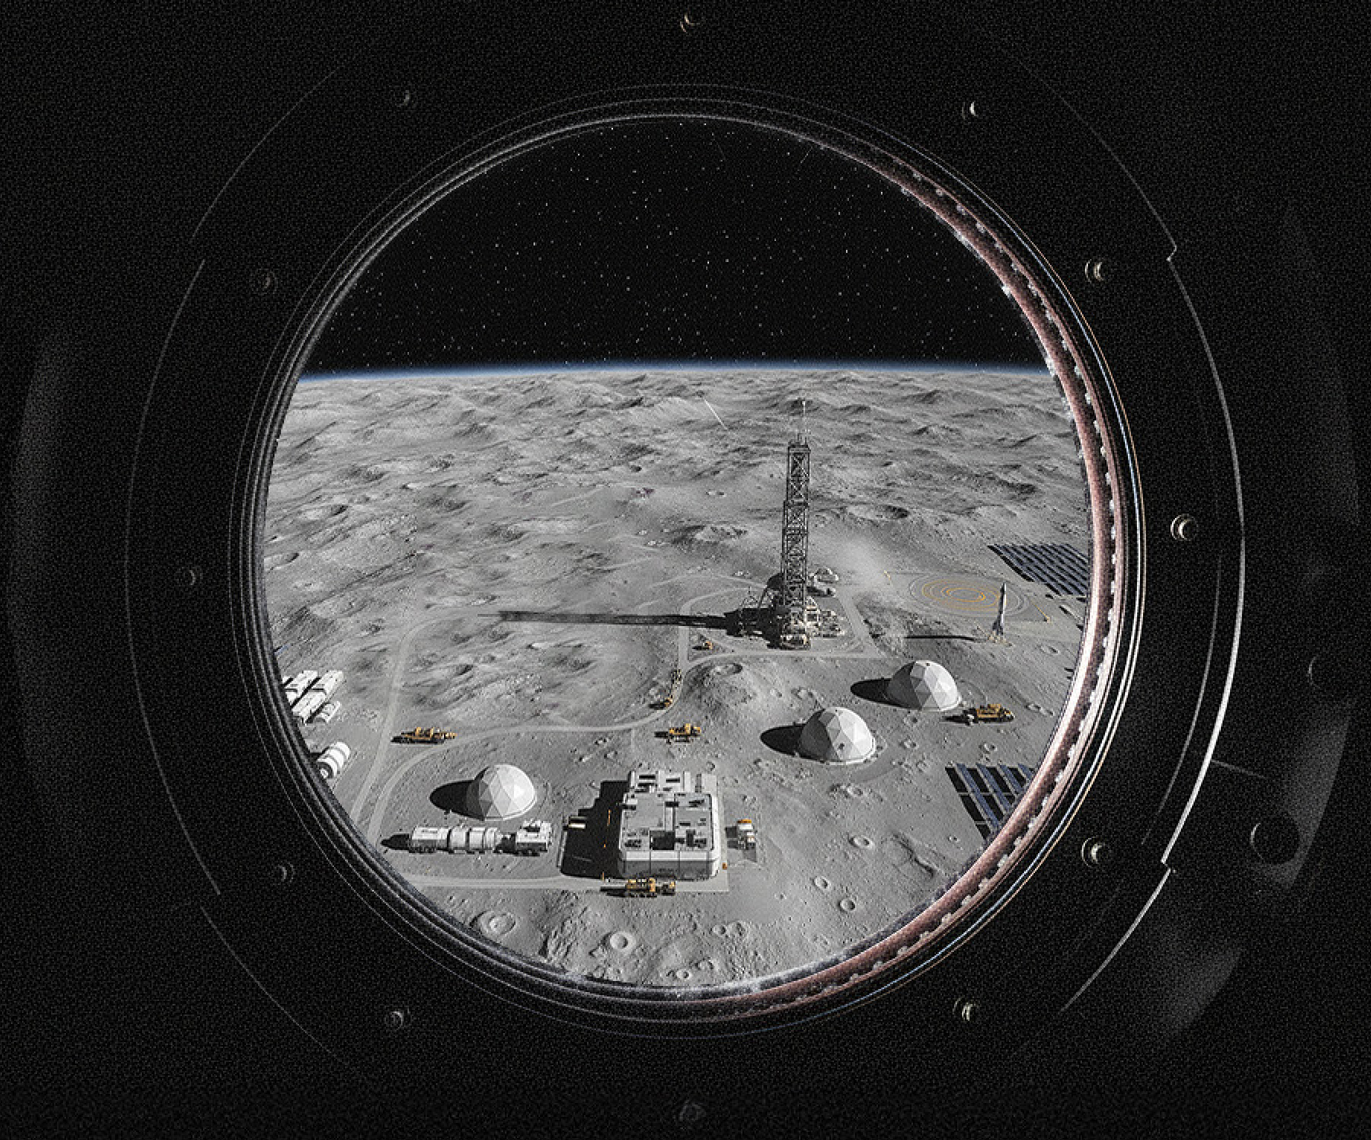

# GENERALIZABILITY TO A WIDE ARRAY OF TASKS...

## Understanding user needs

Lunarlab develops AI-driven tools for lunar science and exploration. Different users engage with Lunarlab in different ways, from scientists validating models to developers integrating data into new tools. By mapping their goals, needs, and measures of success, we can shape a release strategy that ensures Lunarlab is practical, valuable, and inspiring for all.

The ideas that follow explore potential application directions.

**The FDL Lunar-FM team at SpaceR Research Group, the LunaLab, and the Zero-Gravity Lab at the University of Luxembourg's Interdisciplinary Centre for Security, Reliability and Trust (SnT)**

# 1) INTERDISCIPLINARY ELEMENTS FOR A RESPONSIBLE FRONTIER

**LUNAR AGENT PROMPT**
*Show me regions where resource potential, environmental stability and heritage preservation overlap. Which analogies can be drawn from other fields?*

**UNDERLYING (SCIENCE) QUESTION**
Can we maximize resource exploration avoiding or minimizing environmental and heritage risks? What elements are most useful from other disciplines? Is studying the Moon in isolation sufficient, or would a broader understanding colored by other disciplines be beneficial?

**TOOLS NEEDED - FLAG THE NEED FOR EMBEDDINGS**
Impact scoring algorithms combining geological, operational, legal and ethics constraints into a unified index.

**LIMITATIONS & OPPORTUNITIES**
Concept of “sustainability” is yet to be defined. Limited knowledge of heritage and sensitive locations. Framework for sustainability metrics and decision-support layers.

**PERSONA**
Lunar Sustainability Officer, Lunar Policy Officer, Space Lawyer, Resource Exploration Geologist

**CONTRIBUTOR(S)**: Abigail Calzada Diaz, Russell Spiewak

# 2) LUNAR ANALYST FOR SITE OPS AND PLANNING (LASOP)

**LUNAR AGENT PROMPT**
*Help me define top priority sites with abundance of materials X, Y, X while taking into consideration availability of solar light for power needed to set up operations.*

**UNDERLYING (SCIENCE) QUESTION**
Can we identify X,Y,Z materials and specific lunar features (lighting, geographic, etc.) with the model, and then make an assessment of viability of setting up operations.

**TOOLS NEEDED - FLAG THE NEED FOR EMBEDDINGS**
Embeddings and additionally might need to include other data: time series, shadows, lighting cycle, simulation data, etc.

**LIMITATIONS & OPPORTUNITIES**
Current resolution could be useful in a first stage, but would need higher resolution for complete assessment, depending on the application (hazard avoidance might need higher resolution for autonomous vehicles will definitely need higher resolution).

**PERSONA**
Mission Designer, Mission Operations, Control Staff

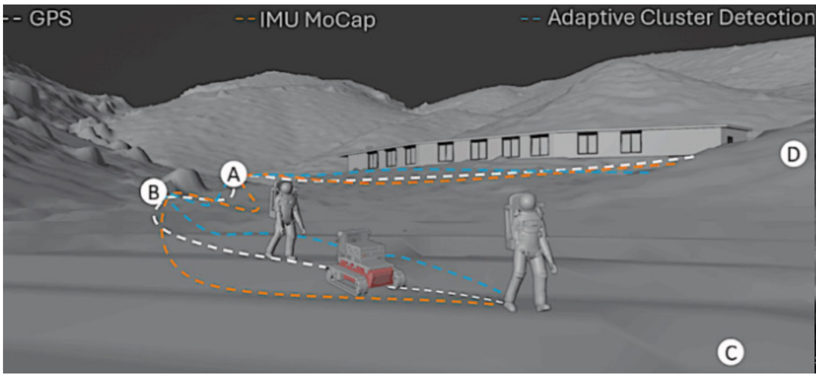


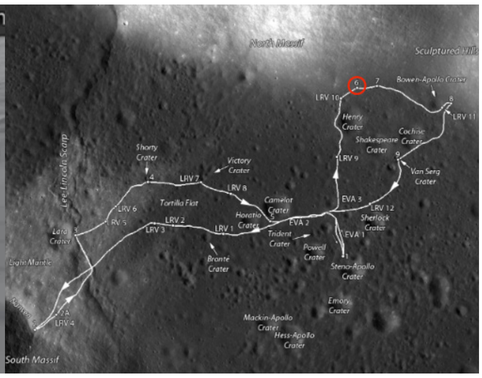


**CONTRIBUTOR(S):** David Cheuvront, Carmen Waters, Raul Ramos, Sumit Goski

# 3) LUNAR CRISIS RESPONSE TABLETOP EXERCISE

**LUNAR AGENT PROMPT**
*A solar flare warning has just been received, identify the nearest / safest lunar pit for shelter.*

**UNDERLYING (SCIENCE) QUESTION**
Disaster response on the lunar surface in response to a solar flare.

**TOOLS NEEDED - FLAG THE NEED FOR EMBEDDINGS**
Position of the sun in relation to the Moon, access high-res NAC, with tooling to identify and quantify shielding offered by lunar pits.

**LIMITATIONS & OPPORTUNITIES**
Opportunity: Enable cross foundation model communication, i.e., solar foundation model. Integrate these outputs with video-based navigation instruction and emergency wayfinding.
Limitation: No GPS on the Moon (need to build out the navigation capability without this.

**PERSONA**
Mission Operations, Astronauts

**CONTRIBUTOR(S):** James Parr, Poliana Santana, Marc Girona-Mata

# 4) LUNAR-FM SIMULATOR

**LUNAR AGENT PROMPT**
*Can you simulate a dataset of typical terrain for traverse planning with realistic illumination.*

**UNDERLYING (SCIENCE) QUESTION**
Integrate Lunar-FM and physics based modeling, including a simulation layer in Lunar-FM. This would allow for relevant regions to be modeled in a generative sense and extend current high resolution imagery.

**TOOLS NEEDED - FLAG THE NEED FOR EMBEDDINGS**
Cesium/Unreal or similar physics engine, with Lunar-FM basis and similarity search to create appropriate mapping.

**LIMITATIONS & OPPORTUNITIES**
Resolution is the major limiting factor, while gamification is a huge opportunity. There is a need to define a baseline simulation for a representative high resolution area then extend this using Lunar-FM.

**PERSONA**
Mission Planner, Developers


**CONTRIBUTOR(S):** Big Think

# A FOUNDATION MODEL FOR LUNAR RESOURCES

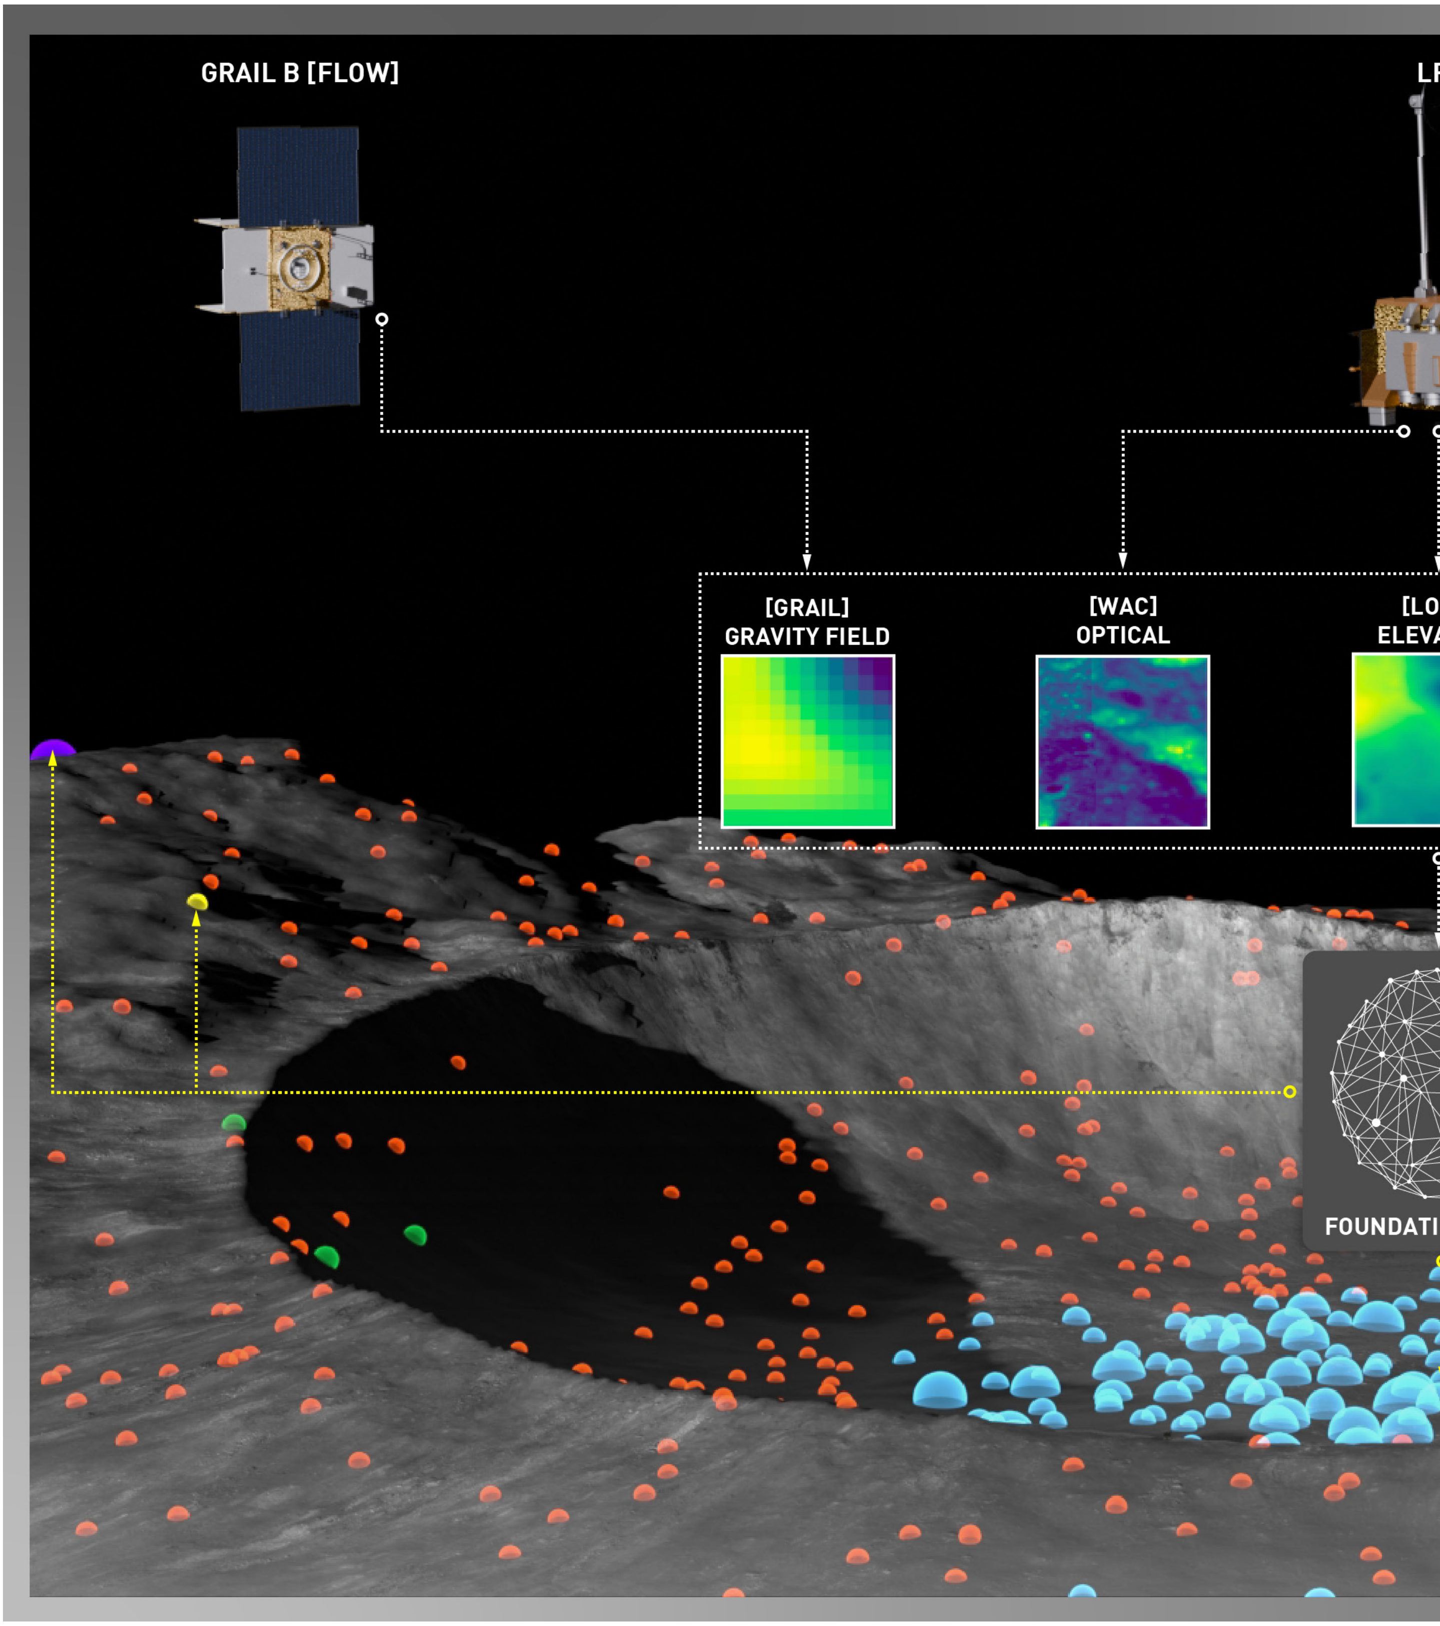

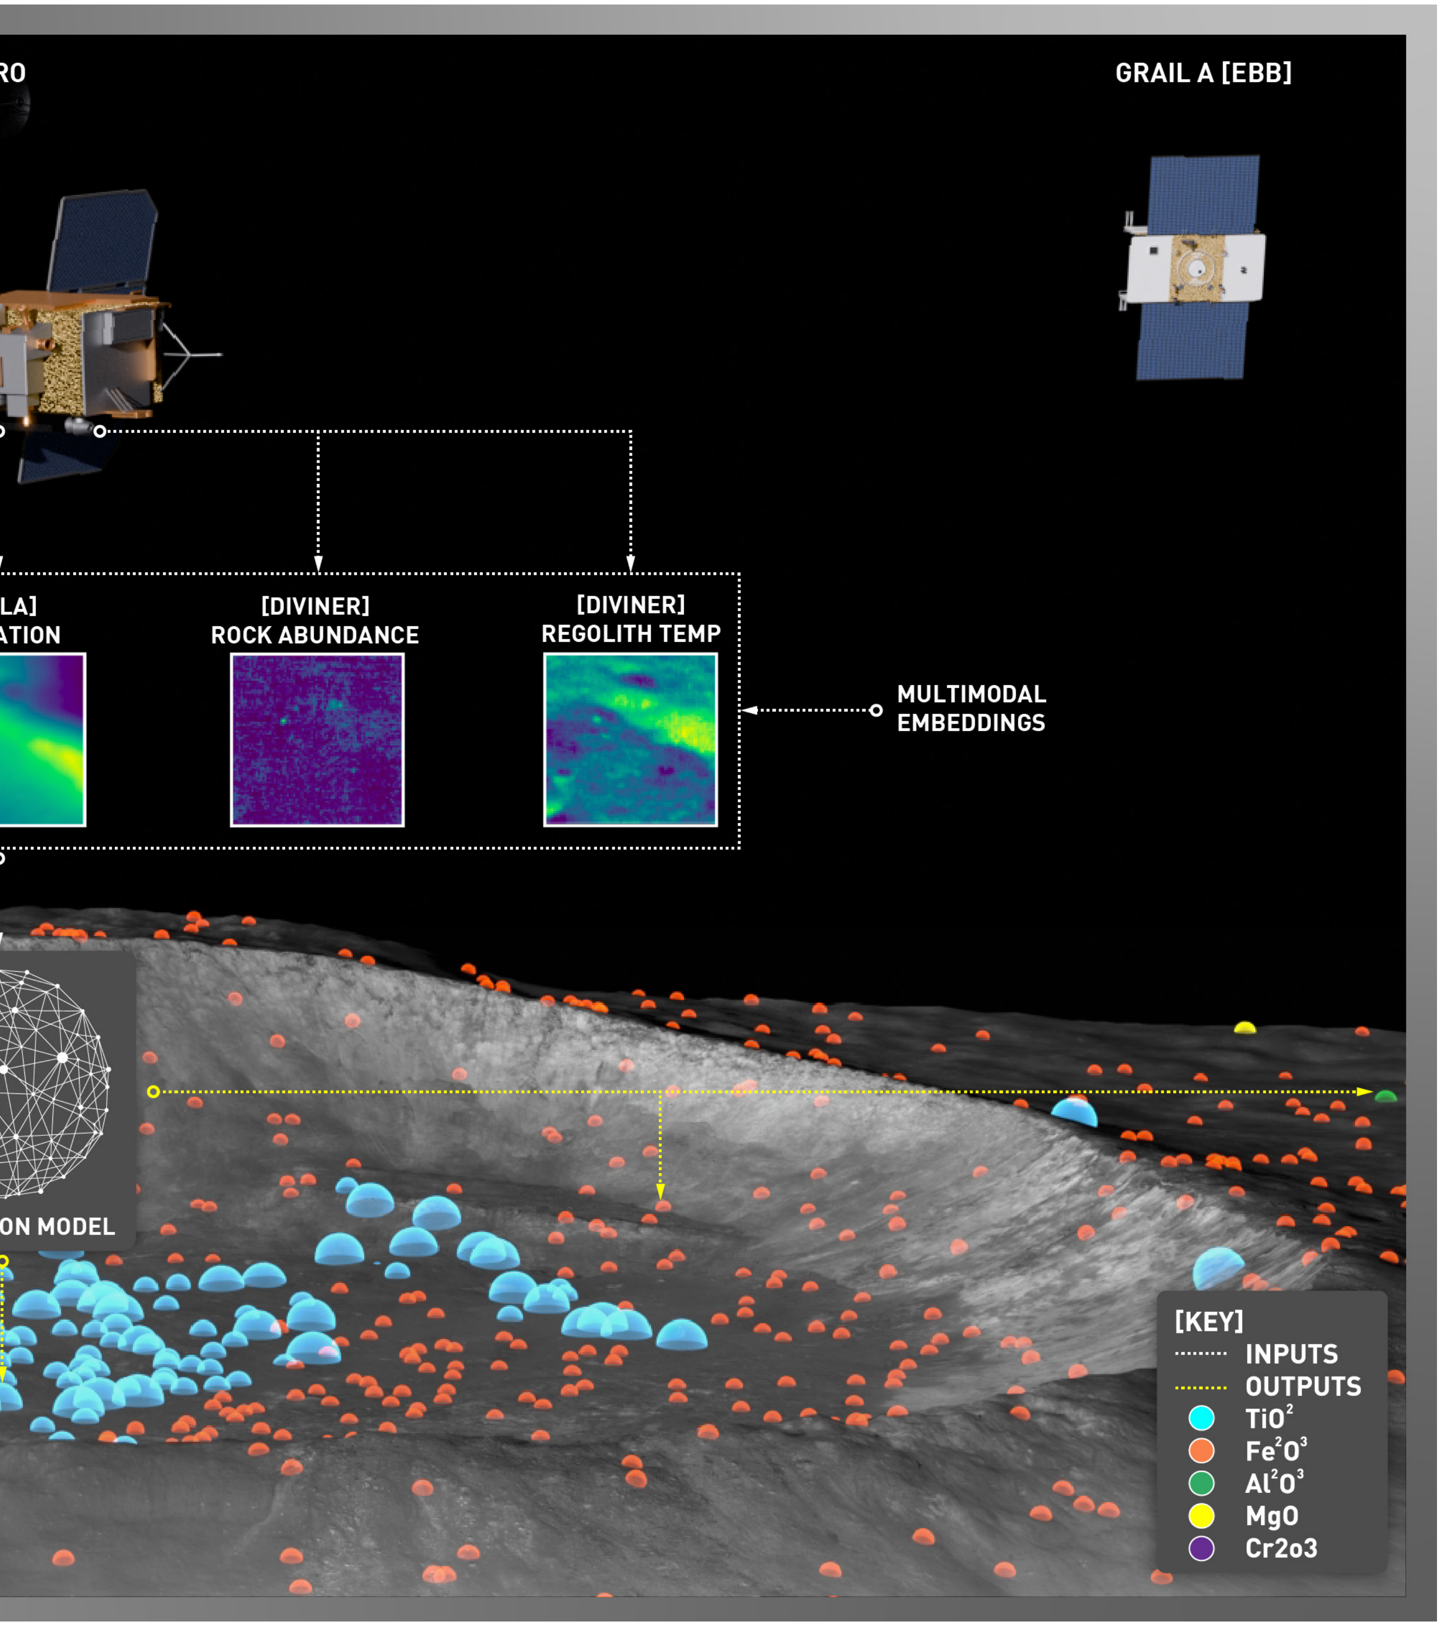

RO
GRAIL A [EBB]
LA]
ATION
[DIVINER]
ROCK ABUNDANCE
[DIVINER]
REGOLITH TEMP
MULTIMODAL
EMBEDDINGS
ON MODEL
[KEY]
INPUTS
OUTPUTS
$TiO^2$
$Fe^2O^3$
$Al^2O^3$
MgO
Cr2o3

# ENGAGING THE SCIENCE COMMUNITY

**Release stages:**

### Evaluation

Establish and third party validation of downstream use cases for Lunar-FM. Initial user testing for Lunar Agent.

### Science

Reproduction of rigorous scientific output using Lunar-FM. Refinement of Lunar Agent, with deployment strategy.

### Public

Lunar-FM hosted and publicly available, with literature and tutorials for onboarding.

### Future

Next development cycle for Lunar-FM 2.0.

There is a need for **clear communication** and **community coordination**.

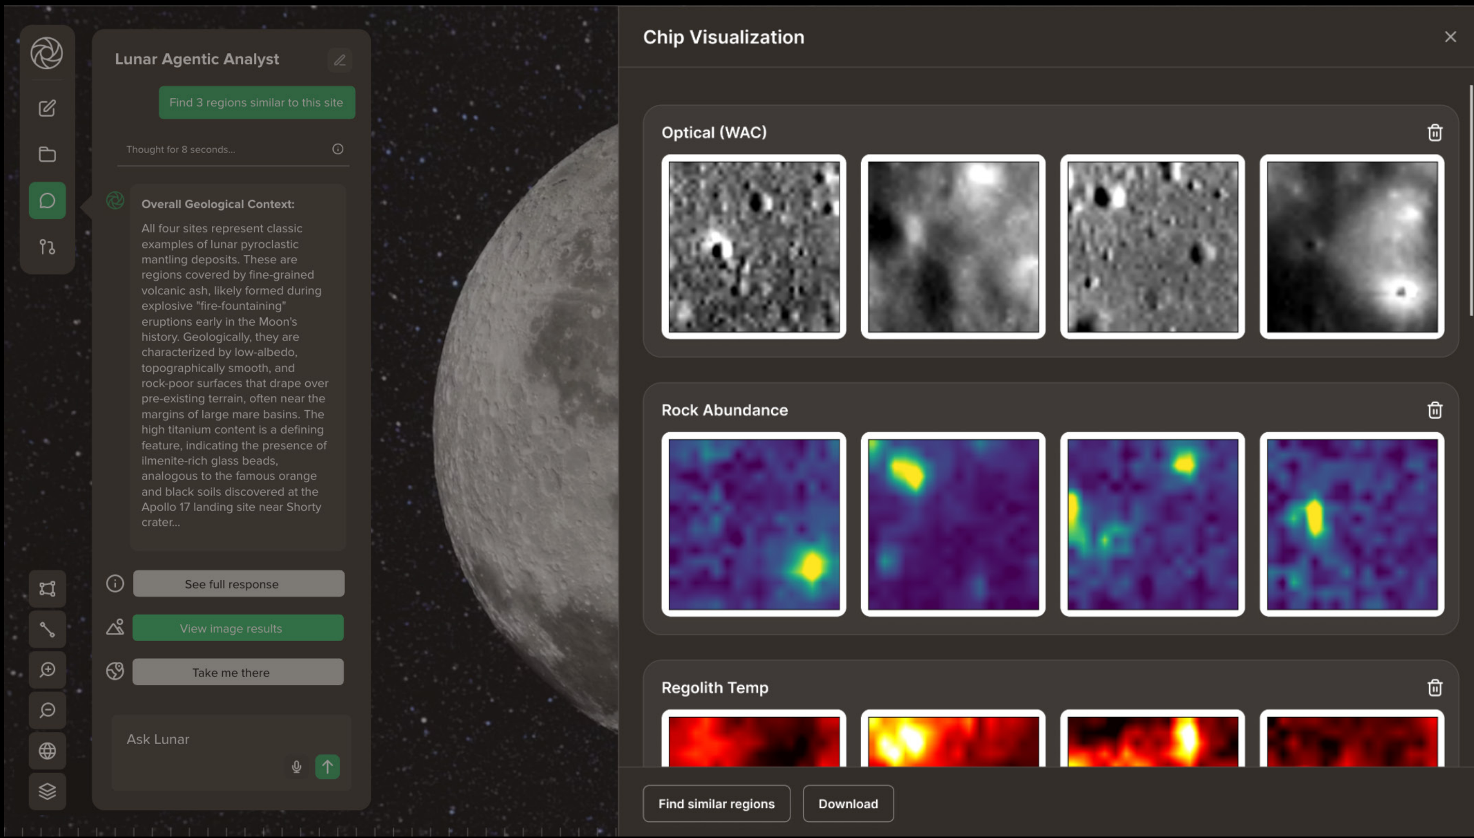


# COME TALK TO THE MOON WITH US!

Email Mike Heyns to register your interest in joining evaluation release testing.

**mike.heyns@trillium.tech**

FDL
LUXEMBOURG SPACE AGENCY | ESRIC
GOOGLE CLOUD | NVIDIA | SCAN AI
FM FOR LUNAR RESOURCES

LUXEMBOURG SPACE AGENCY esric LUNARLAB TRILLIUM TECH

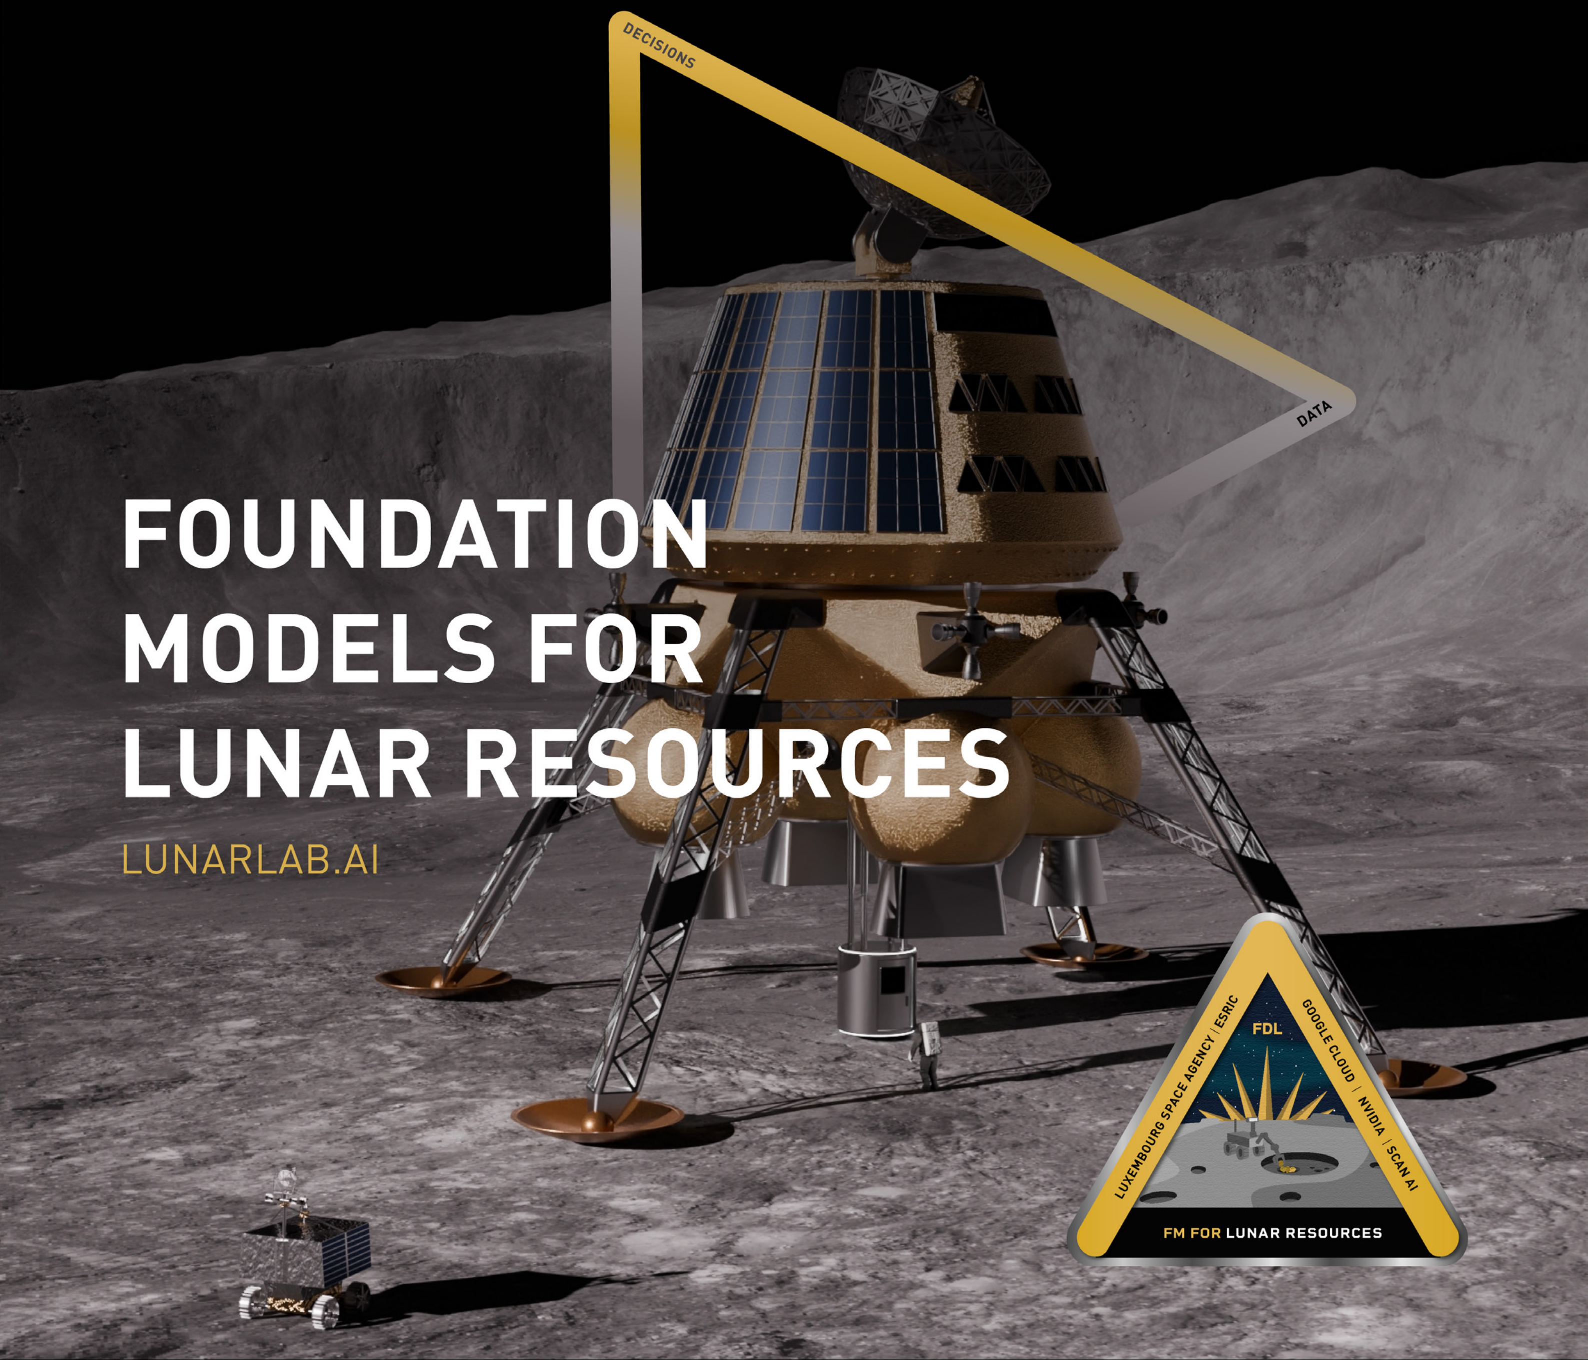


# FOUNDATION MODELS FOR LUNAR RESOURCES

LUNARLAB.AI

**Sumit Goski**
Researcher

**Jakob Gawlikowski**
Researcher

**Sylvester Kaczmarek**
Faculty

**Gautier Bardi de Fourtou**
Researcher

**Marc Girona-Mata**
Researcher

**Raul Ramos-Pollan**
Faculty

# CHALLENGE SUMMARY

**The Lunarlab team has developed the first foundation model for Lunar Resources, called Lunar-FM. This multimodal foundation model for the Moon uses unified multimodal lunar embeddings to combine global datasets, including multispectral, radar, temperature, and gravity sensors, using a multi-MAE architecture.**

Lunar-FM can translate between different data modalities and reconstruct masked-out data across modalities, demonstrating its ability to learn valuable information both within and across modalities. This model can further perform similarity searches using its embeddings to find similar lunar locations and build importance maps to identify relevant areas in each modality, providing interpretability.

By using a few positive and negative examples of titanium dioxide or other mineral concentrations, the model can derive common characteristics and extrapolate over the Moon, resulting in a reasonable static map for potential resources that correlates well with ground truth maximum concentration.

Not only has the team developed lunar embeddings for a variety of use cases, now we can talk to the Moon. The team has a prototype for a Lunar Analyst Agent that allows users to interact with lunar knowledge using natural language to inquire about different regions of the Moon based on location or features.

Taken together, we now have a new capability stack for lunar exploration, including a unified data stack, core multi-modal embeddings, high-confidence prospectivity Maps, and a prototype for an agentic analyst.

LUXEMBOURG SPACE AGENCY esric LUNARLAB

# FOUNDATION MODEL FOR LUNAR RESOURCES: TECH MEMO

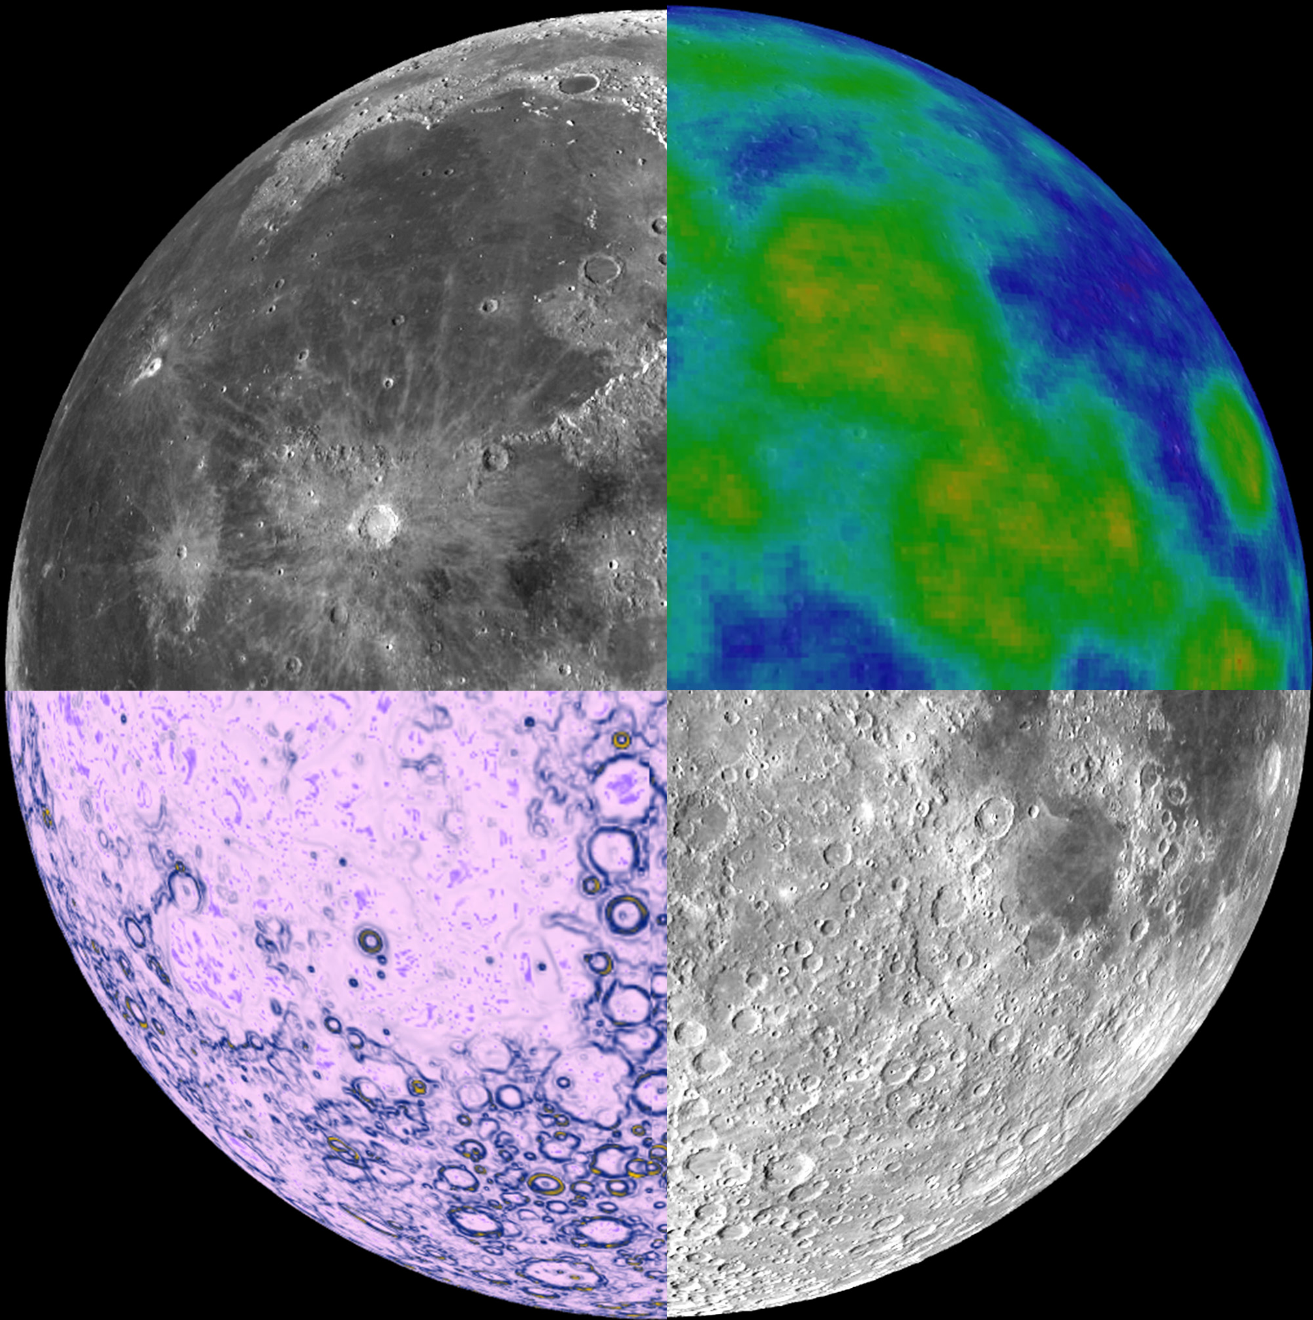

FDL.AI NETWORK | 10 YEARS OF APPLIED AI RESEARCH FOR ALL HUMANKIND | Google Cloud NVIDIA SCAN Datarock LUX PROVIDE

AD ASTRA PER ALGORITHMOS

TRILLIUM TECHNOLOGIES

LUXEMBOURG SPACE AGENCY esric LUNARLAB

TRILLIUM TECHNOLOGIES

*"Let's build an intelligent future"* ™

**Frontier Development Lab (FDL) Technical Memorandum Template**
**Foundation Model for Lunar Resources**


### AUTHORS

Marc Girona-Mata, Jakob Gawlikowski, Sumit Gorski, Gautier Bardi de Fourtou, Sylvester Kaczmarek, Raúl Ramos-Pollan, Kathryn Hadler, Abigail Calzada Diaz, Ben Moseley, Valentin Bickel




### ACKNOWLEDGEMENTS

This work has been enabled by FDL Lunarlab (https://lunarlab.ai/), a public/private partnership between the Luxembourg Space Agency, ESRIC and Trillium Technologies, in partnership with Google Cloud, NVIDIA Corporation, SCAN, and Datarock.

# INTRODUCTION

The Luxembourg Space Agency (LSA) is invested in advancing in-situ resource utilization (ISRU) on the Moon, from water ice to useful minerals, to support future missions and the space economy (*Luxembourg Space Agency, 2024*). Orbital missions (NASA LRO, Chandrayaan-1, etc.) have revealed rich datasets (e.g., lunar regolith composition, temperature, topography) and even water ice deposits in diverse regions of the Moon (*Luxembourg Space Agency, 2024*).

However, these heterogeneous data sources remain largely siloed. The identified need is a multimodal foundation model that can learn from diverse lunar datasets jointly, capturing complex geospatial correlations and enabling new insights for resource prospecting. Such a model would help fuse information from various instruments to better locate and characterize potential resources (e.g., regions high in titanium or hydrogen), addressing the challenge of data overload and knowledge gaps in lunar exploration (*Trillium Technologies, 2024*; *Luxembourg Space Agency, 2024*). By creating a unified representation of the Moon's surface, scientists and mission planners can more easily query and analyze lunar sites of interest using AI. This can be taken a step further, building AI that interacts with lunar data through natural-language queries, enabling "talking to the Moon".

This technical memorandum presents our approach to building a multimodal lunar foundation model (LunarFM), the data and methodology used, and the results on several downstream tasks relevant to lunar science and resource exploration. We also describe how this foundation model can be leveraged in an interactive lunar agent that integrates external knowledge and user queries using the LunarFM embeddings as core tools.

LUXEMBOURG SPACE AGENCY esric LUNARLAB

TRILLIUM TECHNOLOGIES

*"Let's build an intelligent future"* ™

# IDENTIFIED NEED

## Why is this a challenge?

As international space agencies plan for a sustained human presence on the Moon (e.g., Artemis program), the ability to locate and utilize in-situ resources becomes critical. Key lunar resources such as water ice, oxygen bound in oxides (like $TiO_2$ in ilmenite), and other minerals are unevenly distributed and mostly inferred from disparate orbital sensor data. Given the vast number of instruments looking at different aspects of the Moon, several challenges arise when attempting to provide a comprehensive knowledge of all available information.

**Data Fragmentation:** A lot of the recorded data is currently still quite fragmented in the sense that they are available in different places. For example, the Lunar Reconnaissance Orbiter (LRO) alone carries cameras, a laser altimeter, a thermal radiometer, and more, with each instrument yielding separate datasets that are distributed across different archives. Combining these datasets requires significant effort. Thus, the proposed foundation model must address the combined challenge of, first, ingesting all these datasets simultaneously and, second, aligning them in a common space.

**Heterogeneous Modalities:** The Moon's data modalities range from images of the surface (which contain visual-textural information) to geophysical measurements like gravity (which inform about subsurface mass distributions). The data are not naturally compatible as they measure different physical properties, and their spatial resolutions differ. This heterogeneity poses a challenge: many machine learning models excel at single modalities (e.g., convolutional neural networks for images), but struggle when confronted with multimodal input. A key challenge is to design an architecture that can jointly learn from different data types without one dominating or being ignored.

**Generalization**: Certain high-value lunar-resource analyses are currently done in a very specific manner. For instance, mapping titanium content ($TiO_2$) across the Moon has been done by empirically correlating spectral ratios from Clementine or LRO camera data with sample measurements. If we wanted to also factor in mineralogy or extend to other elements following this approach, it would mean developing another bespoke model. Another approach is to develop a foundation model that, once pre-trained, can be easily adapted to multiple such tasks (composition mapping, mineral detection, etc.) by leveraging the same learned representation. This reusability is a strong motivator for the foundation model as a large pre-trained model that can then be fine-tuned for a variety of different specific applications.

# State of the Art

## Lunar Analysis

The mapping of the lunar surface has traditionally been performed by scientific experts interpreting a combination of images and other scientific data (altimeter, radar, spectrometers, etc.) with some single-sensor data processing. The Unified Geologic Map of the Moon (*Fortezzo et al., 2020; v2 2022*) is a result of this approach, representing a major step in the record of planetary exploration. However, the workflow is too time-consuming and necessarily restricted by the cognitive load and bias of human mappers, because this approach relies heavily on human interpretation.

In recent years, more approaches have leveraged machine learning (ML) and deep learning (DL) to automatically analyze some of the datasets available for the Moon, thus allowing for more efficient and accurate analysis. This effort was mainly drawn towards crater detection and characterization (*Jia et al., 2020*; *Silburt et al., 2019*; *Di et al., 2014*). Additional studies have also leveraged these methods for other instruments and analyses, such as thermal analysis (*Moseley et al., 2020*), seismic analysis (*Liu et al., 2024*), or generating an artificial lunar landscape (*Malarkhodi et al., 2025*).
However, all these studies are based on only one modality amongst all the available ones, and very few combine more than one modality: *Scorsoglio et al., 2020*, for autonomous landing, or *Lu et al., 2021*, which highlights the limitations of current SOTA approaches.

## Foundation Models

While a lot of work has been done both on lunar feature analysis and on foundation models, including foundation models for Earth Observation applications, our approach aims at leveraging the gained experiences and adapting them to Moon resources. This application comes with its own challenges, such as the scarcity of data, the high multimodality, and the unique physics of the lunar environment, which requires a tailored approach. Amongst the architectures that showed the best performance for Earth Observation are *TerraMind* (*Jakubik et al., 2025*), the *Galileo* (*Tseng et al., 2025*), or the *AnySat* (*Astruc et al., 2025*), which are highly multimodal, integrating various modalities such as satellite imagery, multispectral and radar data, elevation maps, and natural language descriptions. TerraMind, reaching top performance on the *PANGAEA* (*Marsocci et al., 2024*) benchmark, at the time of this study, significantly outperforms previous geospatial foundation models across different Earth observation tasks. The novelty of our approach is to adapt those architectures to lunar applications, tackling unique challenges including lunar physics particularities, data quality and quantity, and ground truth scarcity.

LUXEMBOURG SPACE AGENCY esric LUNARLAB

TRILLIUM TECHNOLOGIES

*"Let's build an intelligent future"* ™

# DATA DESCRIPTION

To build a foundation model that truly understands the Moon, we curated a comprehensive set of **quasi-global, static lunar maps** from multiple orbital instruments. These datasets provide complementary views of the lunar surface and subsurface, spanning a range of modalities: optical imagery, topography, thermal emission, radar reflectivity, gravity anomalies, and spectral reflectance. Below are all the (18) modalities that were included in our lunar foundation model (also in Table 1):

- **Clementine UVVIS 750 nm (~118 m/px, global):** legacy single-band albedo with broad maturity/composition context. It complements WAC by anchoring FeO/$TiO_2$/OMAT-style proxies and unit mapping at the global scale. Use as a baseline albedo/brightness channel after photometric normalization; regularize to avoid over-weighting albedo vs WAC.
- **7-band Hapke WAC (321-689 nm, ~400 m/px, 70°S–70°N):** photometrically normalized to fixed geometry (i=60°, e=0°, g=60°) to suppress topographic shading. Color/maturity signals (ratios, continuum slopes) are consistent across latitude/lighting. Use for composition/maturity features and as the illumination-invariant multispectral backbone.
- **LOLA LDEM (~118 m/px, global):** laser altimeter DEM capturing elevation, slopes, and relief controlling illumination and permanently shadowed region (PSR) likelihood. Encodes terrain hazards/access and exhumation potential often missed by reflectance alone. Use to derive multi-scale slope/roughness/curvature/horizon and to co-register/normalize all other layers.
- **Diviner Rock Abundance (night-time, ~200-500 m/px, global):** two-component thermal fit (rock+soil) estimating the **areal fraction of warm rocks**. Indicates freshness, blocky ejecta, and landing/trafficability hazards; complements Mini-RF roughness. Use with latitude/slope de-trending and robust per-chip stats (median/p5/p95).
- **Diviner Regolith Temperature (slope-adjusted midnight, ~200-500 m/px, global):** physics-clean **soil-only** night temperature after removing slope and rock contributions. Tracks thermal inertia/regolith thickness and cold-trap propensity relevant to volatile stability. Use alongside PSR/illumination metrics; aggregate extrema/percentiles for robustness.
- **Diviner Bolometric Temperature (day/night, ~200-500 m/px, global):** The **total thermal energy radiated from the surface** (integrated across all wavelengths). It is a direct measure of the surface's thermal state and is essential for modeling the lunar thermal environment, including the stability of volatiles and heat budget analysis.
- **GRAIL Free-air (total) gravity anomaly (10-100 km, global):** observed gravity with topographic signal retained, highlighting basins/mascons and long-wavelength structure. Provides first-order tectonomagmatic context tied to KREEP and mare emplacement. Use with spectral band-splitting (hi/low-pass) and DEM co-interpretation.

- **GRAIL Bouguer anomaly (topography removed, 10-100 km, global):** gravity after subtracting modeled topographic loads to emphasize **density contrasts**. Reveals intrusive bodies, buried flows, and crustal heterogeneity pertinent to subsurface resources. Use multiple density reductions or propagate per-pixel uncertainty to avoid model-lock-in.
- **GRAIL inverted density ("dist", degree-scale, global):** model-derived crustal density field constrained by gravity. Supplies smooth structural priors at long wavelengths that guide regional prospectivity. Use strictly as **context** with uncertainty awareness; don't treat it as truth for supervision.
- **GRAIL measurement uncertainty ("sd", global):** per-pixel error/SD associated with gravity products. Prevents over-trusting weakly constrained regions and improves uncertainty-aware training/evaluation. Use for loss weighting, masking, and to report confidence in downstream maps.
- **Mini-RF CPR (~100–300 m/px, partial strips):** circular polarization ratio sensitive to decimeter-scale roughness and multiple scattering. Useful for blocky ejecta, lava textures, and potential ice flags in PSRs (non-unique). Use with incidence-angle normalization, speckle filtering, and terrain masks for layover/shadow.
- **Mini-RF $\sigma^0$ (S1 total backscatter, ~100–300 m/px, partial strips):** angle-normalized radar backscatter mixing dielectric and roughness responses. Adds sensitivity to regolith packing/buried interfaces that optical/thermal miss. Use multilooked mosaics and per-chip robust stats; explicitly encode coverage to handle striping/gaps.

*Table 1. Input data sources, including instruments and modalities.*

| Instrument (Mission) | Modalities Included | Channels | Description |
|---|---|---|---|
| LROC Hapke WAC (LRO) | Multispectral optical imagery | 7 | UV to visible bands (321–689 nm, ~400 m/px) |
| LOLA Altimeter (LRO) | Topographic elevation | 1 | Surface elevation DEM |
| Diviner Radiometer (LRO) | Thermal emission | 3 | Regolith temperature, rock abundance, bolometric temperature |
| Mini-RF (LRO) | Radar | 2 | Circular Polarization Ratio and related radar measures |

| | | | |
|---|---|---|---|
| Lunar Gravity Ranging System Ka-band ranging (GRAIL) | Gravitational anomaly | 4 | Free-air gravity anomaly, Bouguer anomaly, inverted density, measurement uncertainty |
| UVVIS (Clementine) | Reflectance (near-IR or UVVIS band) | 1 | Spectral reflectance (complementary band) |

# METHODOLOGY

## Data Preprocessing

Each modality was preprocessed to ensure comparability and to suit the model input requirements. This involved scaling each channel (e.g., to be zero mean and unit variance) so that all input layers have similar value ranges. All these data were then reprojected to a common cylindrical projection and co-registered on a 0.5° grid covering the lunar surface between 70°S and 70°N latitude, since polar regions were omitted due to incomplete coverage in some datasets. The Moon's surface was then cut into 201,600 discrete "chips", where each chip covers a 0.5° by 0.5° area (approximately 32×32 km at the equator). We then applied masking to indicate **no-data regions**, such as areas with missing Mini-RF or Diviner values. Rather than attempting to fill these gaps during the preprocessing stage, we relied on the model's ability to handle missing pixels as non-informative pixels and masked out patches with a certain fraction of nan-pixels from the reconstruction loss. We also experimented with different grid splitting strategies for creating training and validation sets. An initial random shuffle of chips into train/val/test sets was used.

However, this (random) split strategy is prone to spatial leakage as adjacent chips, which are very similar, belong to different sets. Therefore, if one pixel is used for training and one of its immediate neighbors is used for testing, the model may effectively "see" part of the test region during training, leading to over-optimistic performance. To avoid this, we employed a **band split** strategy, assigning contiguous latitude bands to each split (so that validation chips are geographically separated from training chips by some buffer). This way, the model is tested on regions of the Moon it has never "seen" during training, providing a more rigorous evaluation of generalization in spatial context. Our final split reserved 70% of the chips for training, 20% for validation, and 10% testing.

# Machine Learning Pipeline

Our multimodal foundation model, LunarFM, is based on the **Multi-modal Masked Autoencoder (MultiMAE)** architecture (*Bachmann et al., 2022*), which is in turn an extension of the **masked autoencoder** (MAE) (*He et al., 2021*), and with a Vision Transformer (ViT) as a backbone.

The challenge with multiple modalities is that they may have different data ranges, dimensions, and even different spatial resolutions. One approach would be to stack all channels into one big "image" and apply a standard MAE. However, that would require resampling everything to a common resolution and might let the model shortcut by focusing only on easier-to-predict modalities. Instead, following the approach of MultiMAE (*Bachmann et al., 2022*), we designed a model that treats each modality *somewhat independently within a shared architecture*. Key aspects of our **Lunar-FM architecture** include:

- **Separate Patch Embedding per Modality:** For each modality (e.g., optical image, elevation map, etc.), the input tile is divided into 8x8 pixel patches. Each modality has its own learned linear projection that turns its patch data into an embedding vector. This results in a set of "tokens" for each modality. **Shared Transformer Encoder:** All modality-specific tokens are then concatenated and fed into a single Transformer encoder (ViT) that treats them as a unified set of input tokens. This encoder has a fixed dimension (we set the latent vector size to 768, a typical size for ViT-Base). In the encoder, self-attention layers allow information to be exchanged across modalities. For instance, a token representing an optical patch can attend to tokens from a corresponding elevation patch or temperature patch in the same location. This encourages the model to learn **cross-modal relationships**. The encoder does not know which tokens came from which modality explicitly (though we do encode modality type and location via separate positional embeddings), so it can mix information freely to build a joint representation.

- **Modality-Specific Decoders:** After the encoder, we use separate lightweight decoder networks for each modality to reconstruct that modality's masked patches. That is, we have one decoder that takes the latent representation and produces optical image pixels for masked optical patches, a different decoder that produces elevation values for masked elevation patches, and so on. This design acknowledges that each data type may have different statistical properties, so a single decoder might not be optimal for all. Each decoder is trained only on the task of reconstructing its own modality.

- **Masking Strategy:** In multimodal masking, we randomly mask out a large fraction of the patches *across all modalities*. We ensure that the overall number of masked patches is about 85% of the available patches across modalities. This means that some modalities can randomly be more masked out than others. This encourages the extraction and embedding of cross-modal information and correlations. This enables the model to perform cross-modal inference (e.g., predict a missing modality's patches) and to fill in spatial gaps within a modality (predict a missing part of an image from other surrounding parts).

The final model has 113.3 million parameters, in line with a ViT-Base. We trained the model from scratch (no pre-training) on our lunar dataset with 0.5-degree chips, using self-supervised learning as described. We masked 85% of available patches and used the Adam optimizer with a learning rate of 1e-4. The model was trained for about 500,000 iterations on three H100 GPUs.

# Specific Algorithms Developed for the Task

## LLM Alignment

We explored the possibility of extending a pre-trained LLM so it can interpret multimodal lunar tokens generated by our foundation model (LunarFM). Amongst the different LLM fine-tuning approaches we considered, we found that leveraging a pre-trained vision-capable LLM (e.g., Gemma 3) was a promising avenue. This involves projecting the LunarFM (multimodal) embeddings into the vision-capable LLM's embedding space so they can be consumed as "image tokens" by the LLM's multimodal API (Figure 1), and training this projector concurrently with LLM fine-tuning. The following is a description of our ongoing work plan to implement this.

A sensible approach is to implement this projection as a linear mapping from the LunarFM embedding dimension to the LLM embedding dimension. Inputs to this mapping can be LayerNormed and rescaled with a learnable scalar to ensure that output norms align with the distribution of text embeddings already used by the LLM. This simple affine mapping is preferred as the baseline since it preserves geometry, minimizes parameters, and makes optimization stable. If linear mapping proves insufficient, it can then be extended (with, e.g., a small residual two-layer MLP), but the intention is to start linearly to avoid hiding too much complexity in the projector. For now, no per-modality type embeddings are introduced into the LLM; instead, modalities are reduced by linear mapping.

Each chip's patch-level tokens are serialized in row-major order, with 2D positional encodings optionally injected to preserve spatial structure. The input sequence to the LLM can, for example, start with the beginning-of-sentence special token [BOS], followed by the special token [IMG_START], then the projected FM (image) tokens, then the special token [IMG_END], and finally the text segment [TEXT_START], followed by the text description. This formatting takes advantage of Gemma 3's ability to distinguish image and text tokens. The order of image and text tokens can be permitted. The pre-existing text vocabulary embeddings are frozen so that the lunar extension does not interfere with general language capability. Training batches interleave chip+text examples with pure text data, while causal LM loss is applied to descriptions generated from chip statistics (as well as other metadata sources, such as lunar surface features databases). Dataset splits are spatial (e.g., by lunar tiles; possibly following the same LunarFM's data split) to avoid leakage, and normalization ensures consistent scaling across modalities.

Fine-tuning is done using LoRA (or QLoRA) adapters inserted in LLM attention layers (and optionally also MLPs), while the projector is trained with full-precision. Preservation of general language skills is encouraged by freezing text embeddings and optionally mixing in standard text-only objectives. Evaluation should cover chip-to-text retrieval (and perhaps vice versa), quality of generated captions, and downstream use cases such as classification of terrain morphology/surface features. Robustness can also be tested under out-of-distribution conditions (e.g., illumination, polar vs equatorial, missing sensor modalities). Standard text benchmarks can be used to ensure language retention / avoid catastrophic forgetting. The overall strategy is therefore: a linear projector combined with LoRA fine-tuning to allow lunar patch embeddings to be integrated into a pre-trained LLM with minimal disruption, aligning the multimodal foundation model with the language model's reasoning space.

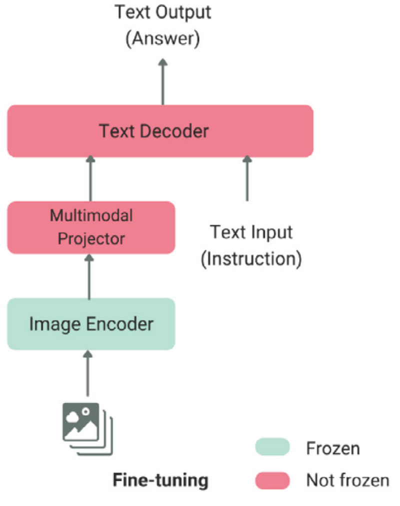


*Figure 1. Schematic representation of vision-enabled LLM fine-tuning (sourced from https://huggingface.co/blog/vlms). For our use case, the 'image encoder' is the MultiMAE pre-trained on lunar multi-instrument data; the 'text decoder' is a pre-trained vision-enabled LLM (e.g., Gemma 3), and the 'multimodal projector' is a linear map or a small MLP. During fine-tuning, the multimodal projector and the LLM (attention weights and optionally MLP weights, but not existing vocabulary embeddings) are jointly trained.*

## Agentic Capability

A core objective of our lunar foundation model (LunarFM) is to make its rich, multimodal embeddings accessible for intuitive exploration. To achieve this, we developed a sophisticated AI agent that serves as a natural language interface, enabling users to "ask questions to the Moon" and receive synthesized answers without needing to write code or understand the complexities of the underlying datasets. This "lunar analyst copilot" is not a monolithic program but a stateful, multi-step reasoning system built using the LangGraph framework. This graph-based architecture defines a clear, robust, and extensible workflow, allowing the agent to dynamically plan and execute actions to fulfill a user's request.

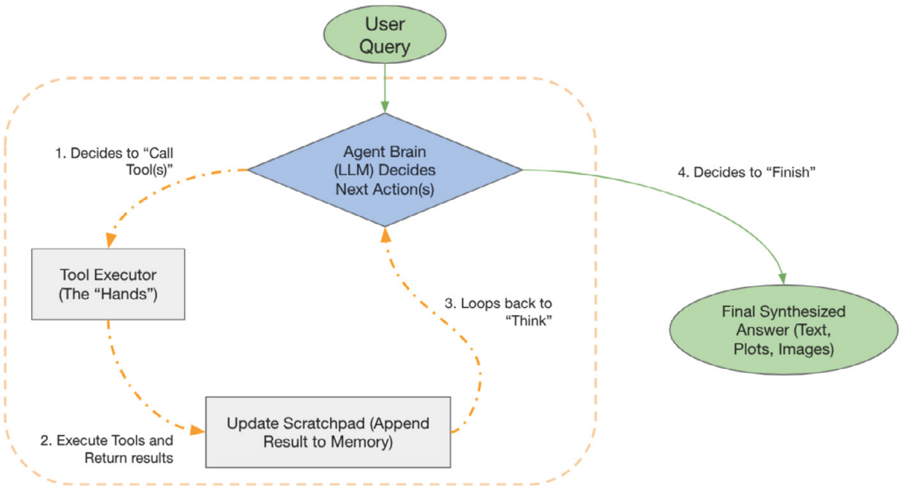


*Figure 2. The architecture of the Lunar Agent, visualized as a state graph. The nodes represent processing steps (intent classification, tool execution), and the edges show the flow of logic as the agent interprets a user's query and generates a response.*

The agent's operational flow is structured as a directed acyclic graph (DAG), where each node represents a specific processing function and the edges represent the conditional logic that routes the state between them. A visualization of this architecture is shown in Figure 2. The entire process can be broken down into four distinct stages:

### 1. Intent Classification (The "Router" or The "Brain")

The process begins when a user submits a query, such as "Find me places that look like the Apollo 11 landing site" or "What data do you have for 89.5°S, 145.0°E?". The initial query is passed to a "router" node, which uses a powerful Large Language Model (Google's Gemini-2.5-Pro) to perform intent classification. The LLM is prompted to categorize the request into one of several predefined intents that map directly to the agent's core capabilities:

a. info_from_coords: The user wants to know about the data available for a specific latitude and longitude.
b. info_from_feature: The user wants to know about the data available for a named lunar feature.
c. similarity_from_coords: The user wants to find regions geologically similar to a specific latitude and longitude.
d. similarity_from_feature: The user wants to find regions similar to a named lunar feature.
e. geology_from_coords: The user is asking a geological question about a specific location.

This initial classification is critical as it determines which downstream tools and logic paths the agent will follow.

**2. Entity Extraction and Disambiguation (The "ID Finder")**

Once the intent is known, the agent must ground the query to a specific location on our lunar grid. This is handled by the get_chip_id node, which intelligently extracts the key geographic entity from the original query and resolves it to a single, valid chip_id-the unique identifier for a data-rich tile in our dataset. This step involves a suite of specialized, purpose-built tools:

a. For coordinate-based intents, an LLM extracts the latitude and longitude from the text. These coordinates are then passed to the `CoordinateLookupTool`, which performs a geospatial query against our chip geometry database to find the corresponding chip_id.
b. For feature-based intents, an LLM extracts the proper name of the feature (e.g., "Shackleton Crater"). This name is passed to the `FeatureNameLookupTool`, which searches a pre-built index of known lunar features. If multiple chips intersect the feature, the tool uses a nearest-neighbor heuristic, selecting the chip whose centroid is closest to the feature's geographic center. This disambiguation step is crucial for providing a single, relevant starting point for analysis.

If a valid chip_id cannot be determined, the agent can gracefully exit and inform the user of the failure.

**3. Tool Execution (The "Action" Nodes)**

With a specific chip_id identified, the agent proceeds to the main action nodes. The conditional logic from the previous step routes the workflow to the appropriate tool based on the user's original intent:

a. For similarity-related intents, the chip_id is passed to the `SimilaritySearchTool`. This tool retrieves the 768-dimensional LunarFM embedding for the query chip and performs an efficient vector similarity search (using squared L2 distance) against the pre-computed embeddings for all other chips. It returns a ranked list of the most geologically similar regions on the Moon. The number of results to return can also be dynamically extracted from the user's query.

b. For information or geology-related intents, the chip_id is passed to the `ChipInfoTool`. This tool queries the datamodule to retrieve metadata about the chip, such as the list of all available data modalities.

**4. Response Synthesis (The "Generator")**

In the final stage, the outputs from the executed tools are collected and passed to a "generator" node. This node uses the LLM to synthesize a coherent, human-readable answer that directly addresses the user's original question. This step transforms raw data into actionable insight:

a. If a similarity search is performed, the LLM receives the list of similar chips and their locations. It then generates a concise geological summary, often highlighting spatial patterns in the results (e.g., "The most similar regions are also located within young mare basalt deposits, suggesting similar composition and surface texture.").

b. If an information query was performed, the LLM formats the list of available data modalities into a clear, easy-to-read summary.

c. If a geological question is asked, the LLM combines the available data modalities with its expert knowledge to provide a preliminary geological assessment.

This modular, tool-based architecture makes the agent highly extensible. New capabilities, such as novel prediction models (e.g., for water ice concentration) or new data sources, can be seamlessly integrated by simply adding new tools to the agent's toolbox. The LLM planner can immediately leverage these new tools by reading their descriptions, paving the way for an ever-more capable and versatile lunar analyst copilot.

# TESTS AND DOWNSTREAM TASKS

After training the foundation model, we designed a series of tests and evaluations to assess its performance and utility on practical tasks. These tests correspond to downstream applications that demonstrate the model's knowledge about the lunar surface and its potential for resource exploration. We briefly describe each test and how it was conducted:

**Reconstruction quality:** As part of the training process, we measured how well the model reconstructs each modality from partial inputs. Visually, reconstructed images appeared plausible, capturing large-scale features but sometimes smoothing out high-frequency details, a known tendency of autoencoders (Figure 3). Quantitatively, the computed mean square errors (pixel-wise computed) were low, indicating the model had indeed learned the joint distribution of the multimodal data. This test confirms that the encoder-decoder has captured meaningful features and learned cross-modal relationships. Though reconstruction alone doesn't guarantee usefulness, it's a necessary first step.

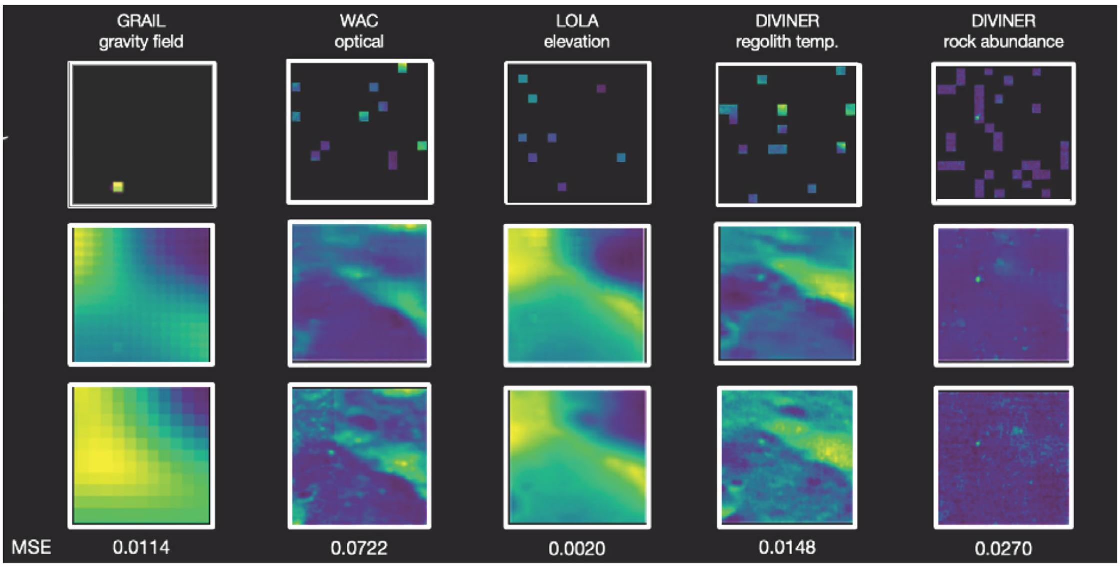


*Figure 3. Visual and quantitative validation of the MultiMAE on lunar data. This figure displays example original multimodal patches (bottom row) and their reconstructions (middle row) from heavily-masked input (top row), demonstrating the model's ability to infill missing pixels and learn cross-modal relationships.*

**Embedding Similarity Search:** we tested whether the model's learned **embeddings** encode semantic information about lunar locations. To do so, we computed the 768-dim embedding for every chip on the Moon using the trained encoder. Then, performed similarity searches in this embedding space. For a given query location (chip) of interest, we retrieve the top 10 most similar other locations by cosine similarity in embedding space. An example of the results showing meaningful similarities is displayed in Figure 4, with only 4 modalities being plotted, even though the distances are computed across the 18 modalities. This also shows that the foundation model can be used as a discovery tool that can find analogues of any given site quickly by embedding comparisons.

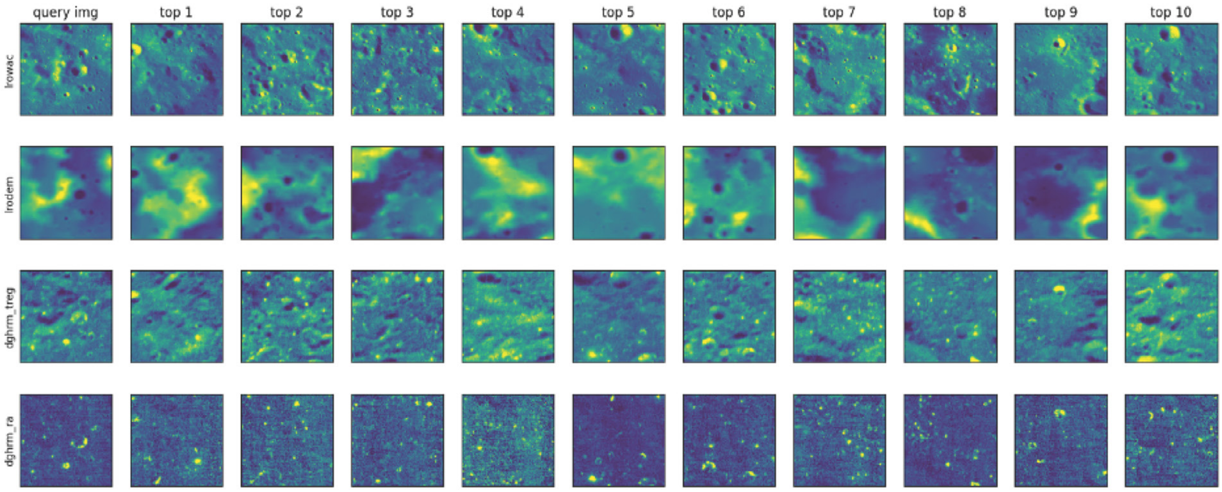


*Figure 4. Embedding similarity search results, illustrating the retrieval of the top 10 most similar chips to a query chip (first column) using cosine similarity on the 768-dimensional LunarFM embeddings.*

**Geologic Unit Classification:** We evaluated whether the embeddings could be used to classify the **geologic unit** of each location, using the USGS Unified Geologic Map of the Moon as ground truth. In this downstream task, each chip is labeled with the geologic map unit that dominates that area. We then trained a simple MLP to predict the geologic unit label of a chip from its representation in the embedding.

We trained this on a subset of the globe and validated and tested it on a separate region following the band split. The performance was measured by accuracy and visual inspection of predicted maps. We found that the model's embeddings do carry significant information about geologic classes, reaching an accuracy of 0.628 on the training set, 0.447 on the validation set, and 0.455 on the test set using only a simple 3-layer MLP on the 1-degree grid. We plotted the model-predicted geologic map versus the actual map in Figure 5. The comparison showed that major boundaries (e.g., the edge of a mare basin, or the extent of a crater's ejecta blanket) were often correctly identified by the embedding-based classification.

Some confusion occurred in transitional areas or where the classes in the geologic map were very detailed. Nonetheless, this experiment proved that **the foundation model's representation encodes both compositional and textural information that aligns with human-defined geologic units**.

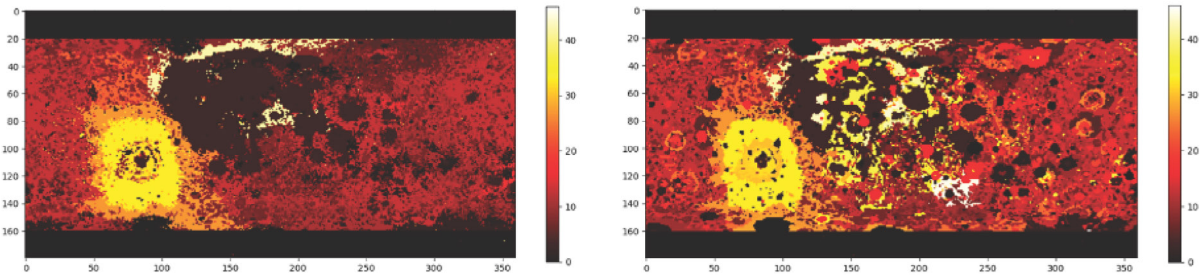

*Figure 5. Geologic unit classification performance. Comparison of the true USGS Unified Geologic Map of the Moon (right) with the model-predicted geologic unit map (left) generated by a 3-layer MLP on the LunarFM embeddings. The visualization confirms that the model's representation aligns with human-defined geologic boundaries and compositional units.*

**$TiO_2$ Abundance Prediction (Regression):** Using LunarFM's learned features to predict the weight percent of titanium dioxide in lunar regolith for a given region. This is a proxy for identifying ilmenite-rich resources. This is a regression problem as we took published global maps of $TiO_2$ abundance and sampled values for each chip. We then fit a simple 3-layer MLP from the embedding to the value. The evaluation is performed by computing the Mean Average Error (MAE) between predicted and actual $TiO_2$ values over a test set of sampled locations. Figure 6 displays the results of the predictions with an MAE of 0.062, proving that the embeddings are rich enough to perform $TiO_2$ predictions.

We took this further by exploring ultra-low-shot learning: **could we predict $TiO_2$ with only a handful of examples, by leveraging expert knowledge?** This simulates the scenario of using an expert's identification of a few promising spots to guide the model. We selected 4 example locations known to have high $TiO_2$ based on expert assessment (*Diaz et al., 2025*) and labeled those as positive; and likewise 4 locations with low $TiO_2$, which we labelled as negative. Training a linear regression on just these 8 points, we still achieved nearly **78% correlation** with the true $TiO_2$ map on evaluation (Figure 7). In contrast, choosing 8 random locations as our training dataset , the correlation dropped to ~0.65. This demonstrates that:

- Lunar-FM embeddings are informative enough so that even a limited number of labeled examples can produce a decent predictive model.
- Combining the embedding with **expert knowledge** (smart selection of examples) can markedly improve results. In other words, an expert geologist could guide the regression or classification model by saying, "these few spots are definitely rich in X", and the predictive model can generalize that insight across the globe effectively.

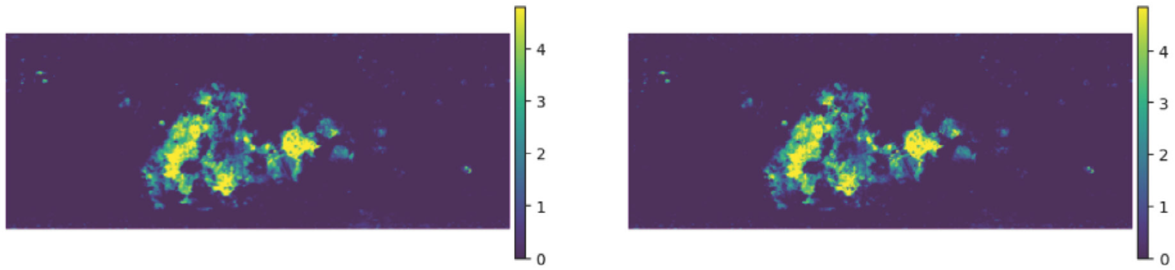


*Figure 6. Global $TiO_2$ abundance prediction map, showing the predicted weight percent of titanium dioxide ($TiO_2$) across the lunar surface (left), derived from a 3-layer MLP trained on the LunarFM embeddings, versus the WAC-derived $TiO_2$ map (right).*

**LLM-Aided Evaluation (RAG and Q&A):** In addition to quantitative tests, we conducted an innovative evaluation using a **Language Model (LLM) agent**. We configured an agent that can answer questions by using our model's capabilities (such as similarity search or regression) along with external knowledge (like a database of lunar science papers). We posed a variety of questions to this agent to see if the foundation model truly contributes to meaningful answers. For example: "Find three regions similar to the Chang'e-5 landing site", the agent would use the similarity tool on the embedding and return locations, then perhaps cross-reference known info (like it might identify that those regions are also young mare basalt areas with high titanium (Diaz et al., 2025)). We found that the agent was generally successful in using the model to answer these complex queries, though it sometimes required careful prompting. This served as a qualitative test of how our foundation model can be used in an interactive scenario. The fact that the LLM agent could integrate our model's output with textual resources to answer such questions is a promising sign for future **systems for smart aid-to-decision support,** where human experts can query an AI assistant for insights about lunar sites.

# RESULTS

Our results confirm the efficacy of the lunar foundation model and illustrate its benefits for both scientific understanding and practical applications:

**Unified Embeddings:** the primary result of our project is the successful training of a unified multimodal embedding for the Moon, and the verification that it contains rich information about lunar surface properties. We obtained a 768-dimensional embedding vector for every 0.5° chip on the Moon, encapsulating information from 18 data layers. The full set of embeddings (~200k vectors) is only ~0.57 GB in size, a drastic reduction from the 160 GB required to store all input mosaics. This compression (about 300× smaller) comes with minimal loss of relevant information. As a result, working with these embeddings is computationally lightweight (small memory footprint, fast to run models on), lowering the barrier to leveraging lunar data.

**Downstream Task Models:** our foundation model's embeddings, when used to train simple models, showed strong performance on multiple tasks, including geologic map classification and $TiO_2$ prediction (performances described in the Test section). Both models trained in a supervised way will be available to be used soon. Additional information on the approach taken to train the models will also be made available in the code shared on GitHub to enable testing other models or different training pipelines to improve accuracy and overall performance. Additional models are also available to predict the concentration of other minerals on the Moon, such as FeO, MgO, but also to predict crater density. The flexible architecture of the available code enables replacing available trained models with new ones to easily test new architectures and improve performance, but also the addition of new downstream tasks.

**Expert-curated Few-shot Learning:** using only 5 to 10 data points (given by an expert), we produced a high-quality global $TiO_2$ prediction map. This suggests that an expert could very rapidly train a custom predictor for a specific resource or feature of interest by just providing a handful of examples. In resource exploration, where labeled examples are scarce (we only have a few ground truth sample locations on the Moon), this capability can be very valuable as it could **amplify expert knowledge through AI**.

As an example, Figure 7 below shows a Linear Regression model prediction using our embeddings trained on expert-curated positive and negative data points for $TiO_2$ (top) and randomly selected ones (bottom).

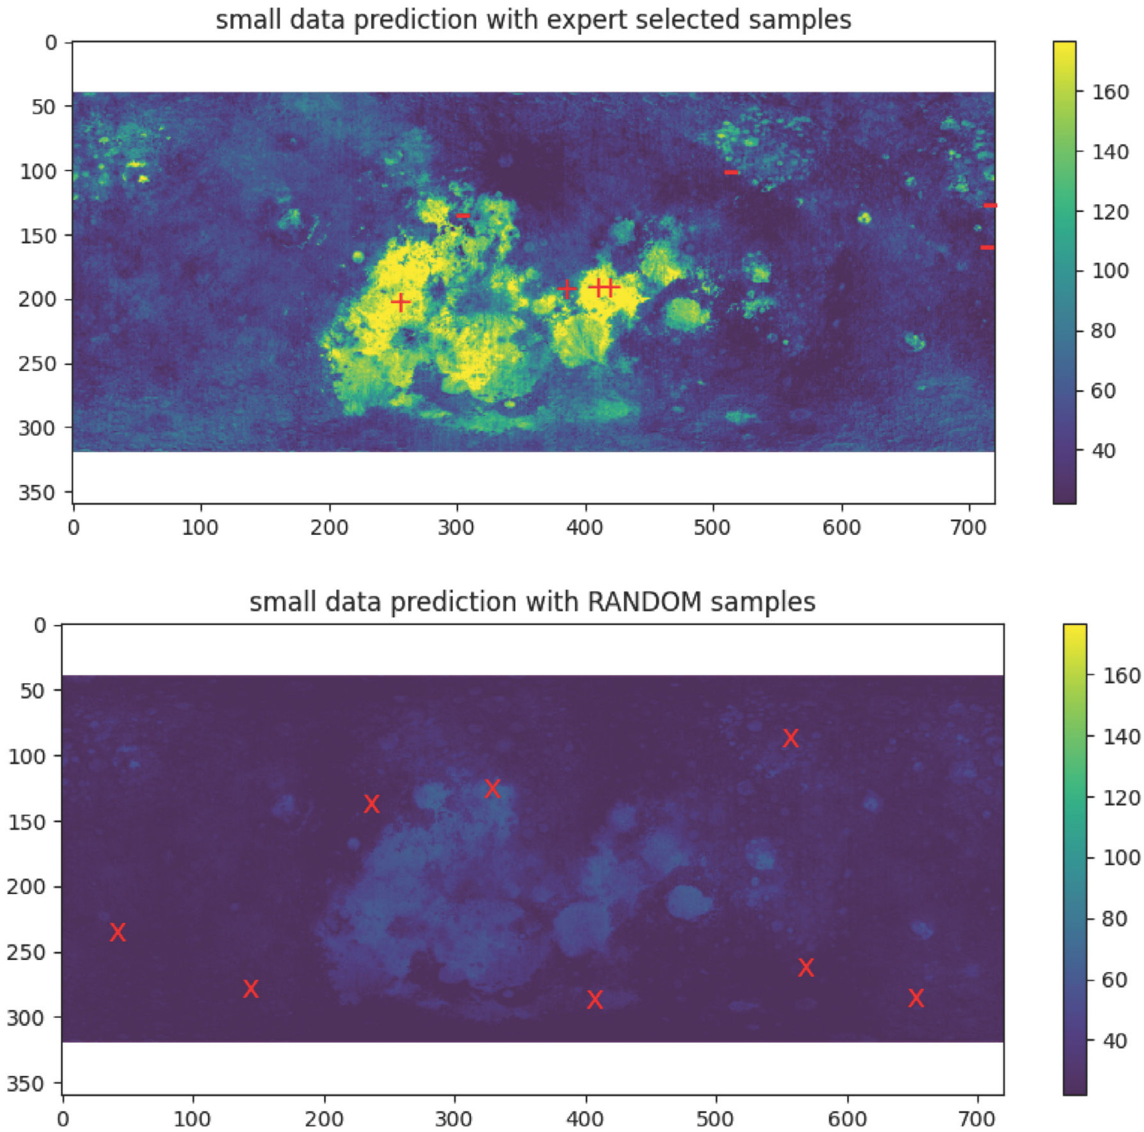


*Figure 7. Few-shot $TiO_2$ prediction: expert vs. random training data points. A comparison of global $TiO_2$ prediction maps generated by a Linear Regression model. The top map is trained on expert-curated positive and negative data points (4 of each), and the bottom map is trained on a randomly selected set of points, highlighting the significant improvement when leveraging expert knowledge with the LunarFM embeddings.*

In fact, when we replicate several times both experiments (expert and random eight data points selection), the correlation coefficient of the linear regression model predictions trained on expert data is not only significantly larger, but also much more stable. See the error bars in the figure below.

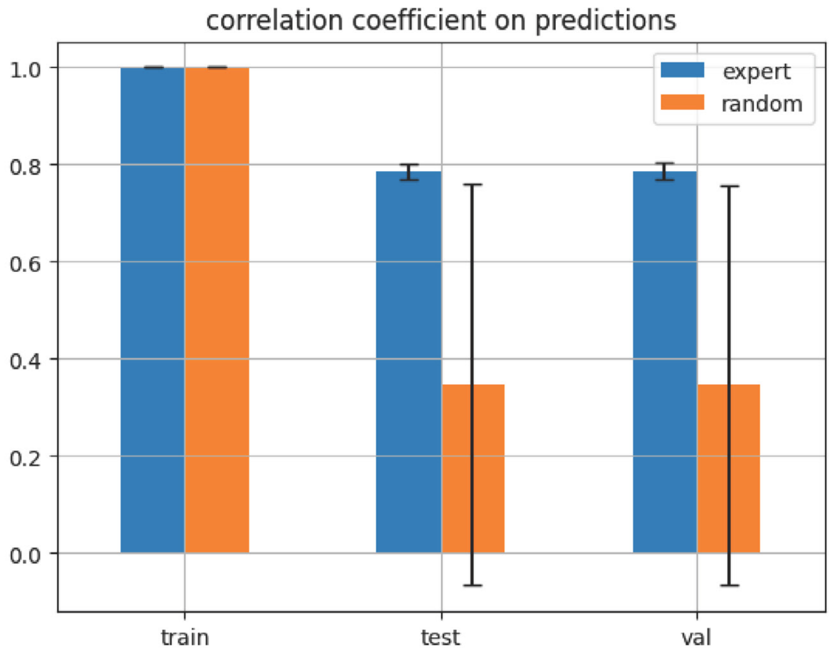


*Figure 8. Stability of few-shot learning with expert curation, comparing the correlation coefficient between predicted and true global $TiO_2$ maps across multiple random trials. The graph shows the higher and more stable correlation achieved when training the linear regression model with expert-curated data points versus randomly selected data points (with error bars).*

**Lunar agent:** one of the showcase results of our project is the development of an agentic interface to the foundation model (Figure 8). We created an AI agent that uses a large language model to interpret natural language questions and then calls relevant tools (functions) that utilize our lunar data and embedding. This agent is also modular, as if, for instance, if a new model or dataset comes along (say, a specialized water-ice prediction model), it can be added as a tool, and the agent can start using it for relevant queries. While this part of the project is more on the prototype side, it illustrates a future where scientists might converse with lunar data through a chat interface (using Gradio), asking high-level questions and getting answers that combine numeric predictions, similarity searches, and literature context in one.

LUXEMBOURG SPACE AGENCY esric LUNARLAB

TRILLIUM TECHNOLOGIES

*"Let's build an intelligent future"* ™

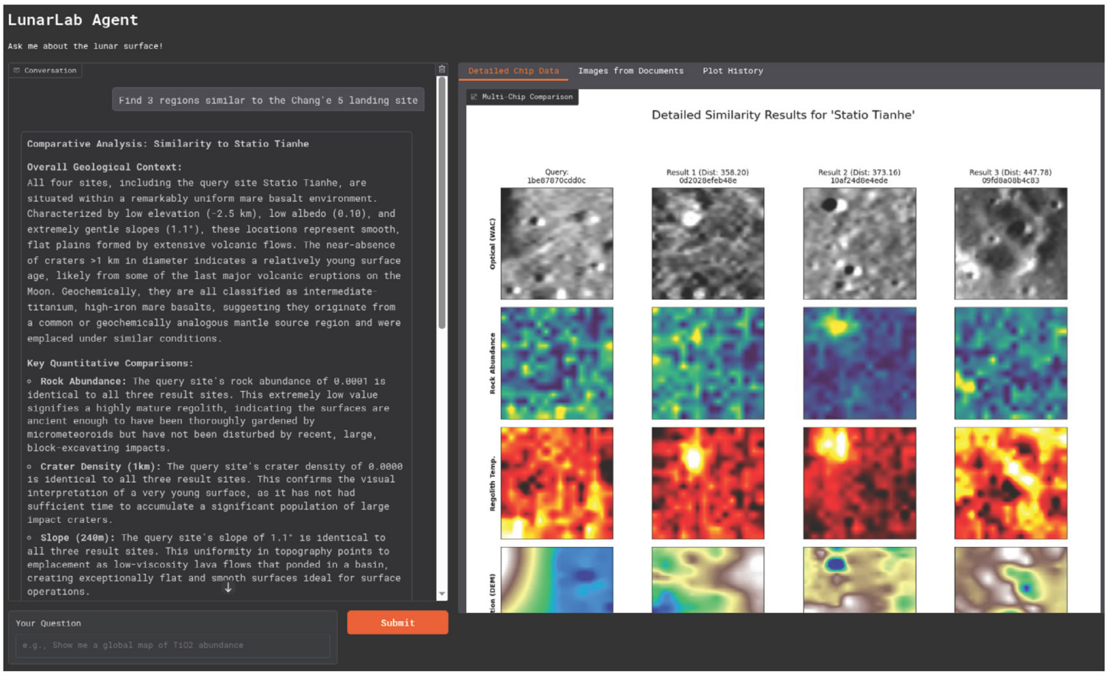


*Figure 9a. Lunar agent's preliminary graphical user interface - Showing similarity search*

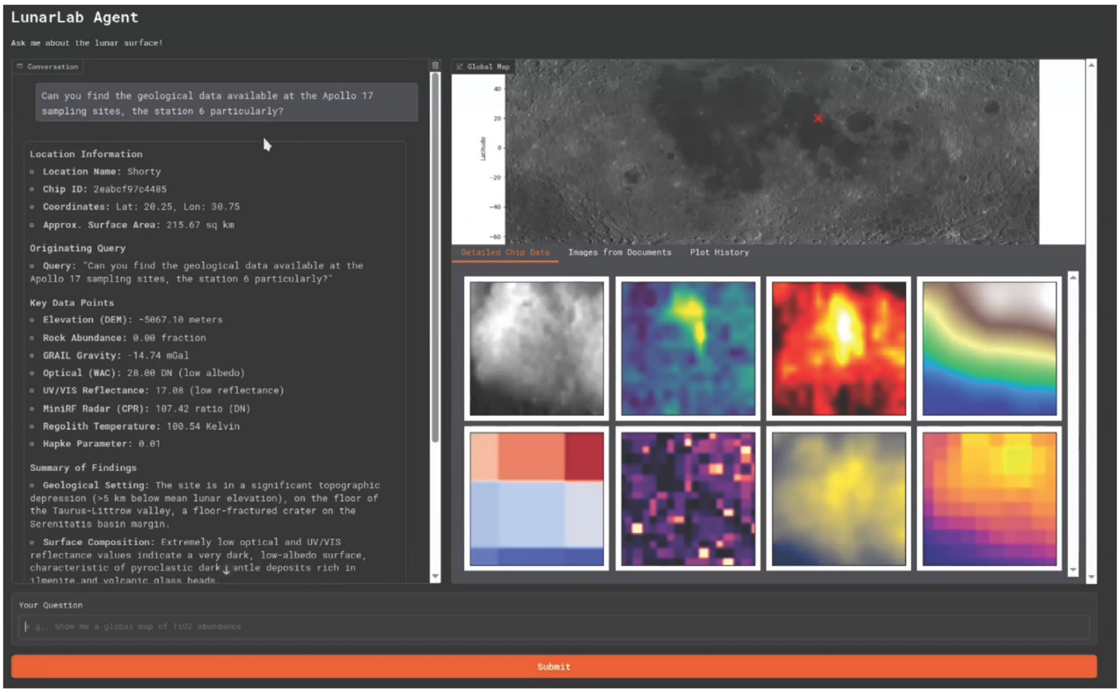


*Figure 9b. Lunar agent's preliminary graphical user interface - Showing data query*

LUXEMBOURG SPACE AGENCY esric LUNARLAB

TRILLIUM TECHNOLOGIES

*"Let's build an intelligent future"* ™

# REPLICABILITY

All source datasets used are publicly available online, mainly through NASA and USGS archives. Namely, LOLA DEM, Diviner GDR (Global Data Records), Mini-RF mosaics, GRAIL gravity models, and Clementine UVVIS data can be downloaded from the Planetary Data System (PDS) and the USGS map through the USGS's Astrogeology portal. Our processed dataset of 0.5° multimodal chips will be available online soon, and researchers will also be able to regenerate the embeddings by following the steps in our code repository, which will also be available online.

Also, the code for model training (based on PyTorch) and for running downstream tasks will also be available and open-sourced in our GitHub repository. This includes configuration files specifying the model hyperparameters (patch size, layer dimensions, etc.) and training routines. Pretrained model weights for the foundation model will also be released, so that the community can directly use the learned embeddings.

For each downstream task (classification, regression, etc.), we include scripts/notebooks that load the pretrained embeddings, train a simple model (or run inference), and produce the reported metrics and figures.

All code for the Lunar Agent (tool definitions, the prompt logic, and example queries) is also included in our repository.

# TOOLS, COMPUTE, AND SOFTWARE ENVIRONMENT

Training a transformer-based foundation model on high-dimensional lunar data required substantial computing resources. We utilized the computing provided by our partners, Google Cloud and ScanAI. In particular, our model was trained on 3 NVIDIA H100 GPUs with 40 GB VRAM. Disk storage of ~1 TB was used to hold the input datasets (global maps and intermediate files).

For the agent component, we utilized the LangChain framework to manage the LLM interactions and tool usage. LangChain allows us to define tools (Python functions) that the LLM can call in a controlled way. The large language model (LLM) itself (Google Gemini) was accessed via API through Google Cloud's Vertex AI.

# CONCLUSION

In this technical memorandum, we have presented how we developed a **foundation model for Lunar Resources**, which, to our knowledge, is the first attempt to fuse such a broad array of lunar remote sensing datasets into a unified machine learning model. By leveraging a multi-masked autoencoder (MultiMAE) architecture, our model learned latent representations that capture the relationships between diverse modalities (imagery, topography, thermal, radar, and gravity data).

The resulting embeddings enable a range of applications from predicting resource-related quantities like titanium to classifying geologic terrain, or performing analog searches for sites of interest. We demonstrated these capabilities and even integrated them into an interactive agent that can answer questions by composing our model's outputs with external knowledge.

This work lays a groundwork for **AI-driven lunar data analysis**: future teams can fine-tune the foundation model with new data, improve the performance of the current simple models for downstream tasks, or even use the embeddings as inputs for new downstream tasks.

In conclusion, our project demonstrates that a foundation model approach, i.e., training once a large model and then fine-tuning smaller models for specific tasks, is feasible and effective for the Moon. It opens up new ways for scientists and mission planners to exploit lunar data. By releasing this model and associated tools, we hope to enable further research and application development in lunar exploration, aligning with LSA's vision of using AI to improve lunar resources identification and utilization.

LUXEMBOURG SPACE AGENCY esric LUNARLAB

TRILLIUM TECHNOLOGIES

"Let's build an intelligent future" ™

# REFERENCES


Astruc, G., Gonthier, N., Mallet, C., & Landrieu, L. (2024). *AnySat: An Earth Observation Model for Many Resolutions, Scales, and Modalities*. arXiv preprint arXiv:2412.14123. Available at: https://arxiv.org/abs/2412.14123

Bachmann, R., Mizrahi, D., Atanov, A., & Zamir, A. (2022). *MultiMAE: Multi-modal Multi-task Masked Autoencoders.* In Proceedings of the European Conference on Computer Vision (ECCV 2022). Available at: https://www.ecva.net/papers/eccv_2022/papers_ECCV/papers/136970341.pdf

Di, K., Li, W., Yue, Z., Sun, Y., & Liu, Y. (2014). A machine learning approach to crater detection from topographic data. Advances in Space Research, 54(11), 2419–2429. https://doi.org/10.1016/j.asr.2014.08.018

Diaz, A. C. & Keszthelyi, L. (2025). *Descriptive models for lunar High-Ti deposits.* Acta Astronautica, 226, 375–384. Elsevier. https://doi.org/10.1016/S0094576524006052

Fortezzo, C. M., P. D. Spudis & S. L. Harrel (2020). *Unified Geologic Map of the Moon, 1:5 000 000-scale*. U.S. Geological Survey Astrogeology Science Center. Available at: https://www.usgs.gov/media/images/fortezzo-et-al-2020-unified-map-moon

He, K., Chen, X., Xie, S., Li, Y., Dollár, P., & Girshick, R. (2022). *Masked Autoencoders Are Scalable Vision Learners.* In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 16000-16009. Available at: https://doi.org/10.1109/CVPR52688.2022.01553

Hoskins, S. (2022). I*nformation-Communication Strategy in Special Library: The Main Activities of a Specialized Information Service in the Information Society.* In P. Czaban & M. Jastrzębski (Eds.), New Perspectives on Library & Information Science (pp. 525-540). Singapore: Springer. https://doi.org/10.1007/978-981-96-1267-3_36

Jakubik, J., Yang, F., Blumenstiel, B., Scheurer, E., Sedona, R., Maurogiovanni, S., Bosmans, J., Dionelis, N., Marsocci, V., Kopp, N., Ramachandran, R., Fraccaro, P., Brunschwiler, T., Cavallaro, G., Bernabe-Moreno, J., & Longépé, N. (2025). T*erraMind: Large-Scale Generative Multimodality for Earth Observation.* arXiv preprint arXiv:2504.11171. Available at: https://arxiv.org/pdf/2504.11171

Jia, Y., Wan, G., Liu, L., Wu, Y., & Zhang, C. (2020). *Automated Detection of Lunar Craters Using Deep*

*Learning.* In Proceedings of the 2020 IEEE 9th Joint International Information Technology and Artificial Intelligence Conference (ITAIC 2020). https://doi.org/10.1109/ITAIC49862.2020.9339179

Liu, X., Xiao, Z., Li, J., & Nakamura, Y. (2024). *Newly Discovered Temperature-Related Long-Period Signals in Lunar Seismic Data by Deep Learning.* Earth and Space Science, 2024. https://doi.org/10.1029/2024EA003676

Lu, Y.; Hu, Y.; Xiao, J.; Liu, L.; Zhang, L.; Wang, Y. (2021). *Three-Dimensional Model of the Moon with Semantic Information of Craters Based on Chang'e Data.* Sensors, 21 (3), 959. https://doi.org/10.3390/s21030959

Luxembourg Space Agency (2024). Resources in Space. Luxembourg Space Agency. Available at: https://space-agency.public.lu/en/space-resources/ressources-in-space.html

Marsocci, V., Jia, Y., Le Bellier, G., Kerekes, D., Zeng, L., Hafner, S., Gerard, S., Brune, E., Yadav, R., Shibli, A., Fang, H., Ban, Y., Vergauwen, M., Audebert, N., & Nascetti, A. (2024). *PANGAEA: A Global and Inclusive Benchmark for Geospatial Foundation Models.* CoRR. Available at: https://openreview.net/forum?id=TgkyGW8jt7

Moseley, B., Bickel, V. T., Burelbach, J. & Relatores, N. (2020). *Unsupervised Learning for Thermophysical Analysis on the Lunar Surface.* The Planetary Science Journal, 1(2): 32. https://doi.org/10.3847/PSJ/ab9a52

Scorsoglio, A., Furfaro, R., Linares, R., & Gaudet, B. (2020). I*mage-based Deep Reinforcement Learning for Autonomous Lunar Landing.* In Proceedings of the AIAA Scitech 2020 Forum (Paper No. AIAA-2020-1910). Available online: https://doi.org/10.2514/6.2020-1910

Silburt, A.; Ali-Dib, M.; Zhu, C.; Jackson, A.; Valencia, D.; Kissin, Y.; Tamayo, D.; Menou, K. (2019). *Lunar crater identification via deep learning.* Icarus, 317, 27–38. https://doi.org/10.1016/j.icarus.2018.06.022

Trillium Technologies (2024). *How FDL-AI and LSA Are Helping Us Go Back to the Moon for Good. Trillium Tech.* Available at: https://trillium.tech/impact-section/how-fdlai-and-lsa-are-helping-us-go-back-to-the-moon-for-good

Tseng, G., Fuller, A., Reil, M., Herzog, H., Beukema, P., Bastani, F., Green, J. R., Shelhamer, E., Kerner, H., & Rolnick, D. (2025). *Galileo: Learning Global & Local Features of Many Remote Sensing Modalities.* In Proceedings of ICML 2025 Poster Track. Available at: https://openreview.net/pdf?id=ggZO3eSZRy

# FOUNDATION MODELS FOR LUNAR RESOUR

***Towards the first multimodal foundation model for the Moo***

**Researchers:** Marc Girona-Mata**,** Gautier Bardi, Jakob Gawlikowski, Sumit Goski
**Faculty:** Raúl Ramos-Pollán, Sylvester Kaczmarek

## CHALLENGE AND APPROACH

The Moon is relatively well mapped but poorly integrated: dozens of orbital datasets (e.g., WAC, LOLA, Diviner, GRAIL, Mini-RF) arrive with incompatible resolutions, geometries, and noise. Thus, key signals remain fragmented across modalities and scales.

We address this by training a **multimodal foundation model (LunarFM)** that aligns all instruments in a unified embedding space, learning shared, physically-meaningful representations without hand-tuned features.

This common latent space supports efficient transfer to downstream tasks, such as resource prospectivity, mineral mapping, and geologic boundary detection; helping convert a patchwork of maps into an integrated, queryable model of the lunar surface.

## METHODS

We leverage the self supervised training paradigm via **cross-modality masked autoencoders**. Specifically, we use a MultiMAE architecture [1], which can deal with different spatial resolutions; and we extend it to better handle no-data pixels/regions.

Our model is able to reconstruct 90% masked inputs in different locations in each modality, and we evaluate the importance of each location within a chip in the unified embedded representation.

INPUT RECONSTRUCTION (model learns to reconstruct from different sparsity across modalities)

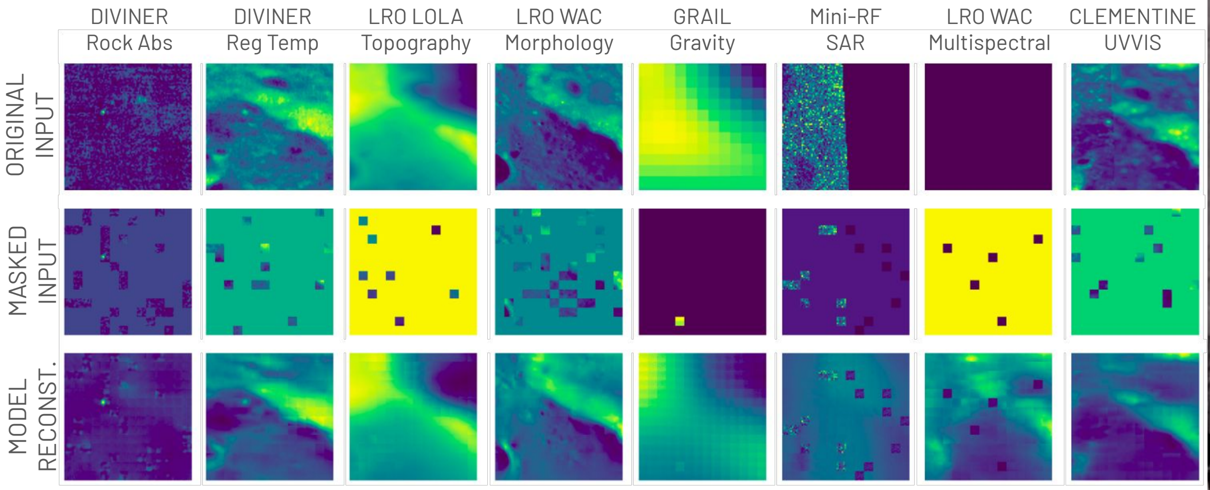


IMPORTANCE MAPS (by masking convolutions and measuring the change in embeddings vectors)

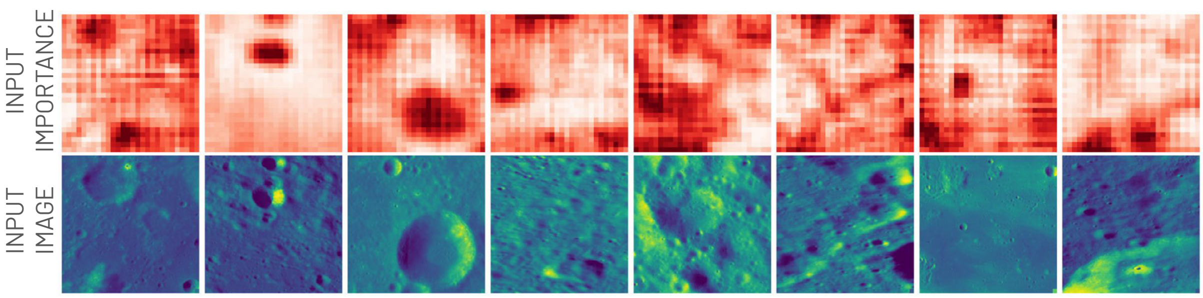

## EXAMPLE DOWNSTREAM TASK: $TiO_2$ PROSPECTIVITY MAPPING

Taking **only 50 known $TiO_2$ examples** (25 positives and 25 negatives), we first reduce the embeddings' dimensionality using principal component analysis (PCA), and then use linear discriminatory analysis (LCA) to learn to map high and low $TiO_2$ concentrations.

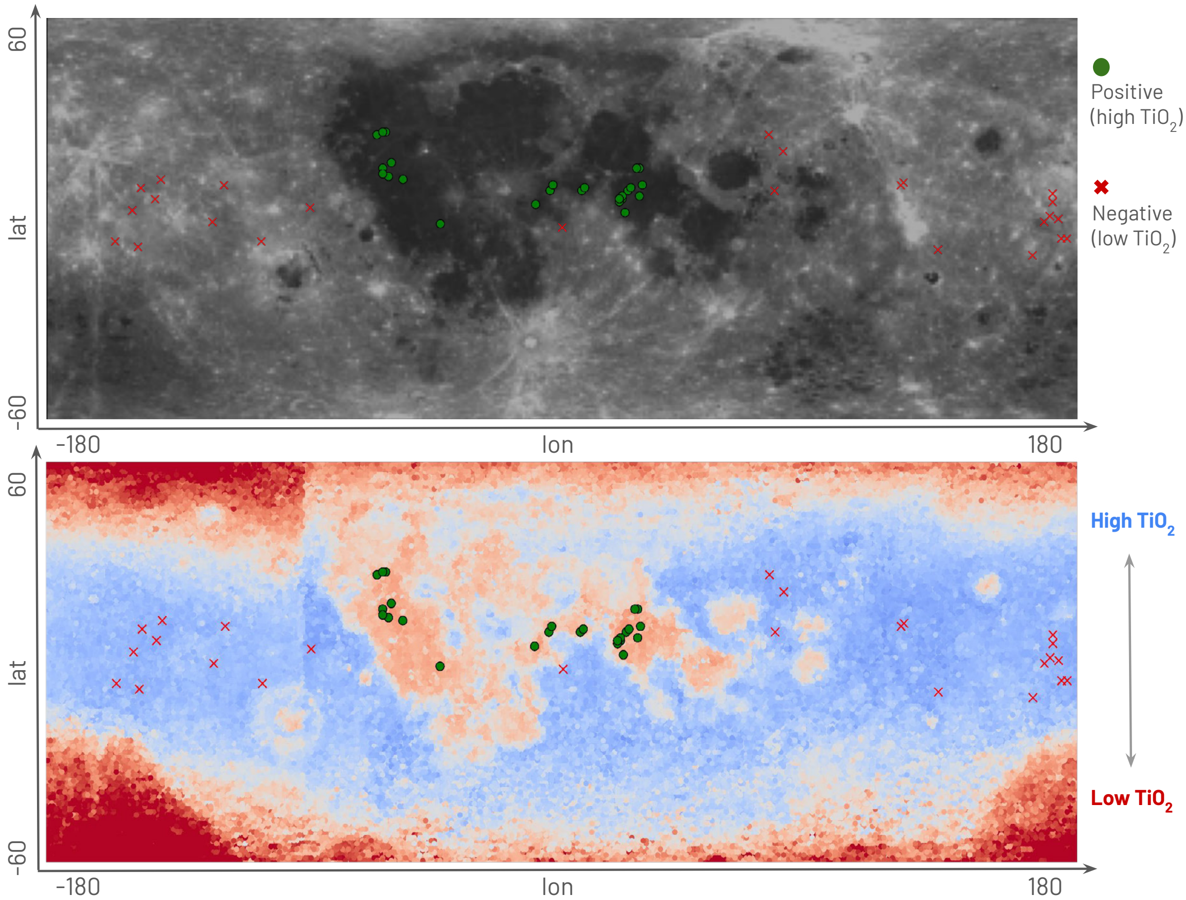


## NEXT STEPS & RELEASE SCHEDULE

- ML-ready input data stack for LunarFM (Sept 2025)
- LunarFM pre-trained embeddings and model weights (Jan 2026)
- Prospectivity maps and additional downstream tasks (Oct 2025)
- Lunar multimodal LLM (early 2026)
- Lunar agentic system (early 2026)

## REFERENCES

[1] *Bachmann, R., Mizrahi, D., Atanov, A., & Zamir, A. (2022, October). Multimae: Multi-modal multi-task masked autoencoders. In European Conference on Computer Vision (pp. 348-367). Cham: Springer Nature Switzerland.*

# **LUNAR SYSTEMS INTELLIGENCE**: TALKING TO THE MOON

Artificial Intelligence (AI) is helping us understand and assist in humanity's return to the Moon. As we enter this new era of sustained lunar activity, the need for intelligent systems that can interpret, predict, and coordinate across diverse data sources in the Lunar environment will be a keystone capability.

For more than a decade, the **Frontier Development Lab** (FDL) has operated at the intersection of AI, space science, and planetary intelligence. In collaboration with **Luxembourg Space Agency** (LSA), our applied research has delivered a series of firsts: from the first-ever view into the Lunar Permanently Shadowed Regions (PSRs), to lunar localization without GPS, thermal anomaly maps and many others. Each milestone has advanced the broader goal of building resilient, reasoning systems capable of operating where latency, bandwidth, and uncertainty are inherent.

Lunar-FM integrates inputs from **multiple spacecraft** into a single multimodal architecture. This fusion allows the model to aggregate a representation of the lunar environment that can be queried in natural language for the first time. As on Earth, the most transformative aspect of the **emergence of foundation models** is multi-modal reasoning. By adapting transformer-based architectures originally developed for natural language understanding, we are enabling models that can converse with complex data, effectively turning lunar datasets into queryable knowledge systems. A scientist or engineer could ask, "Which candidate sites show the most stable illumination for power generation?" and receive an interpretable, physics-informed response supported by visualization and uncertainty estimates. This first model is designed for resource related queries. Future versions of Lunar-FM will unlock more granular capabilities, supporting surface operations and the science goals of NASA's Artemis mission.

Together, hybrid observation, onboard learning, and multi-model reasoning in natural language define a new discipline: Lunar Systems Intelligence - an intelligence architecture for continuous, adaptive understanding of the Moon as a complex, evolving system.
The implications for stakeholders **Luxembourg Space Agency** and **ESRIC** are significant. Lunar-FM can directly support lunar resource characterisation, ISRU (In-Situ Resource Utilization) feasibility assessment, and infrastructure planning, from water ice mining to power distribution, by providing probabilistic, explainable insights at the speed of mission operations.

Lunar-FM represents the beginnings of a **Lunar Operating System** for exploration and utilization; an AI substrate capable of linking simulation, sensing, and decision-making in a single adaptive loop. In doing so, it lays the groundwork for the next phase of humanity's multi-planetary future where intelligent systems don't just observe the Moon, but help us stay there for good.

**LUNARLAB.AI** IS A RESEARCH PARTNERSHIP BETWEEN **LSA, ESRIC** AND **FDL.AI (TRILLIUM TECHNOLOGIES).**

**WE'D LIKE THANK OUR COMPUTE PARTNERS, GOOGLE CLOUD** AND **SCAN COMPUTERS** AND **DATAROCK FOR THEIR EXPERT GEOLOGY INSIGHT.**

# FAQs

- **How does LunarFM consolidate the fragmented and mission-specific lunar remote sensing data?**

*It is designed to fuse 18 disparate data layers from multiple missions into a single, unified feature vector (the 768-dimensional embeddings). This process creates a standardized, global representation of surface and subsurface properties, ensuring consistent data analysis across different geological contexts.*

- **What is the utility of the embedding space for geological tasks beyond resource prediction (regression)?**

*The embeddings are highly effective for qualitative tasks: Geological Classification (allowing accurate delineation of major compositional and textural unit boundaries) and analogue discovery (facilitating efficient similarity searches to find geologically identical regions across the Moon).*

- **Given the scarcity of ground truth data on the Moon, how does LunarFM leverage limited expert knowledge?**

*It enables expert-curated few-shot learning. By allowing an expert to label a very small set of positive/negative examples (e.g., 8 data points for a specific mineral), the model's high-quality embeddings allow a simple linear model to rapidly generate a stable and accurate global predictor map, effectively amplifying expert insight and in situ sampling.*

- **What is the "Lunar Analyst Copilot," and how is it intended to change scientific interaction with the data?**

*The Copilot is an AI Agent architecture that utilizes a Large Language Model (LLM) as a front-end. It allows scientists and mission planners to submit natural language queries (e.g., "Show me areas similar to the Orientale Basin"), which the LLM routes to the appropriate LunarFM tool (Similarity Search or Regression models) to produce a synthesized, conversational answer.*

# THANKS TO OUR ADVISORS AND STEERING COMMITTEE

## Kathryn Hadler (ESRIC)

ADVISOR & STEERING COMMITTEE

## Brenton Crawford (Datarock)

ADVISOR & STEERING COMMITTEE

## Abigail Calzada Diaz (ESRIC)

ADVISOR & STEERING COMMITTEE

## Dave Briguglio (Datarock)

ADVISOR & STEERING COMMITTEE

## Ben Moseley

ADVISOR & STEERING COMMITTEE

## Valentin Bickel

ADVISOR & STEERING COMMITTEE

## STEERING COMMITTEE

**Frank Soboczenski**

**Massimo Mascaro**

**Jonathan Stock**

**David Cheuvront**

**Brian H Day**

**J.L. Galache**

**Jose Ignacio Delgado Centeno**

**Ignacio López-Francos**

**Francesco Civilini**

**Miguel Olivares-Mendez**

*Faculty and researchers are also part of steering committee.*


LUXEMBOURG SPACE AGENCY | ESRIC
GOOGLE CLOUD | NVIDIA | SCAN AI
LUNARLAB
LUNARLAB.AI

# THANKS TO OUR TECHNICAL INTERVIEWERS

Filip Svoboda

Anna Jungbluth

Giacomo Acciarini

Joseph Alejandro Gallego Mejia

Guy Schumann

Raúl Ramos Pollán

Mike Heyns

Giuseppe Castiglione

Josh Veitch-Michaelis

Brianna Isola

Emiliano Díaz Salas Porras

Chetraj Pandey

Karol Bot Concalves

India Jackson

Rutuja Gurav

Angelos Vourlidas

Will Fawcett

James Walsh

Stella Girtsou

Laura Martinez-Ferrer

Iván Venzor-Cárdenas

Marcella Scoczynski

Margarita Bintsi

Daniele de Martini

Grace Colverd

Martin Sanner

LUXEMBOURG SPACE AGENCY

esric

Datarock

## AI FOR SCIENCE IS A TEAM SPORT

**We're deeply grateful to our expert partners, who bring that "seen-it-before" experience as teams dive into the unknown, and to our commercial partners, who provide state-of-the-art hardware (CPUs, GPUs, TPUs), vast RAM, dynamic services, and other innovations that reduce the system-operations overhead of performing AI for science at scale on the critical problems that matter.**

Thank you for being in our corner. We're excited to see what lies on the horizon.

Ad astra per algorithmos.

# THANKS TO OUR PARTNERS

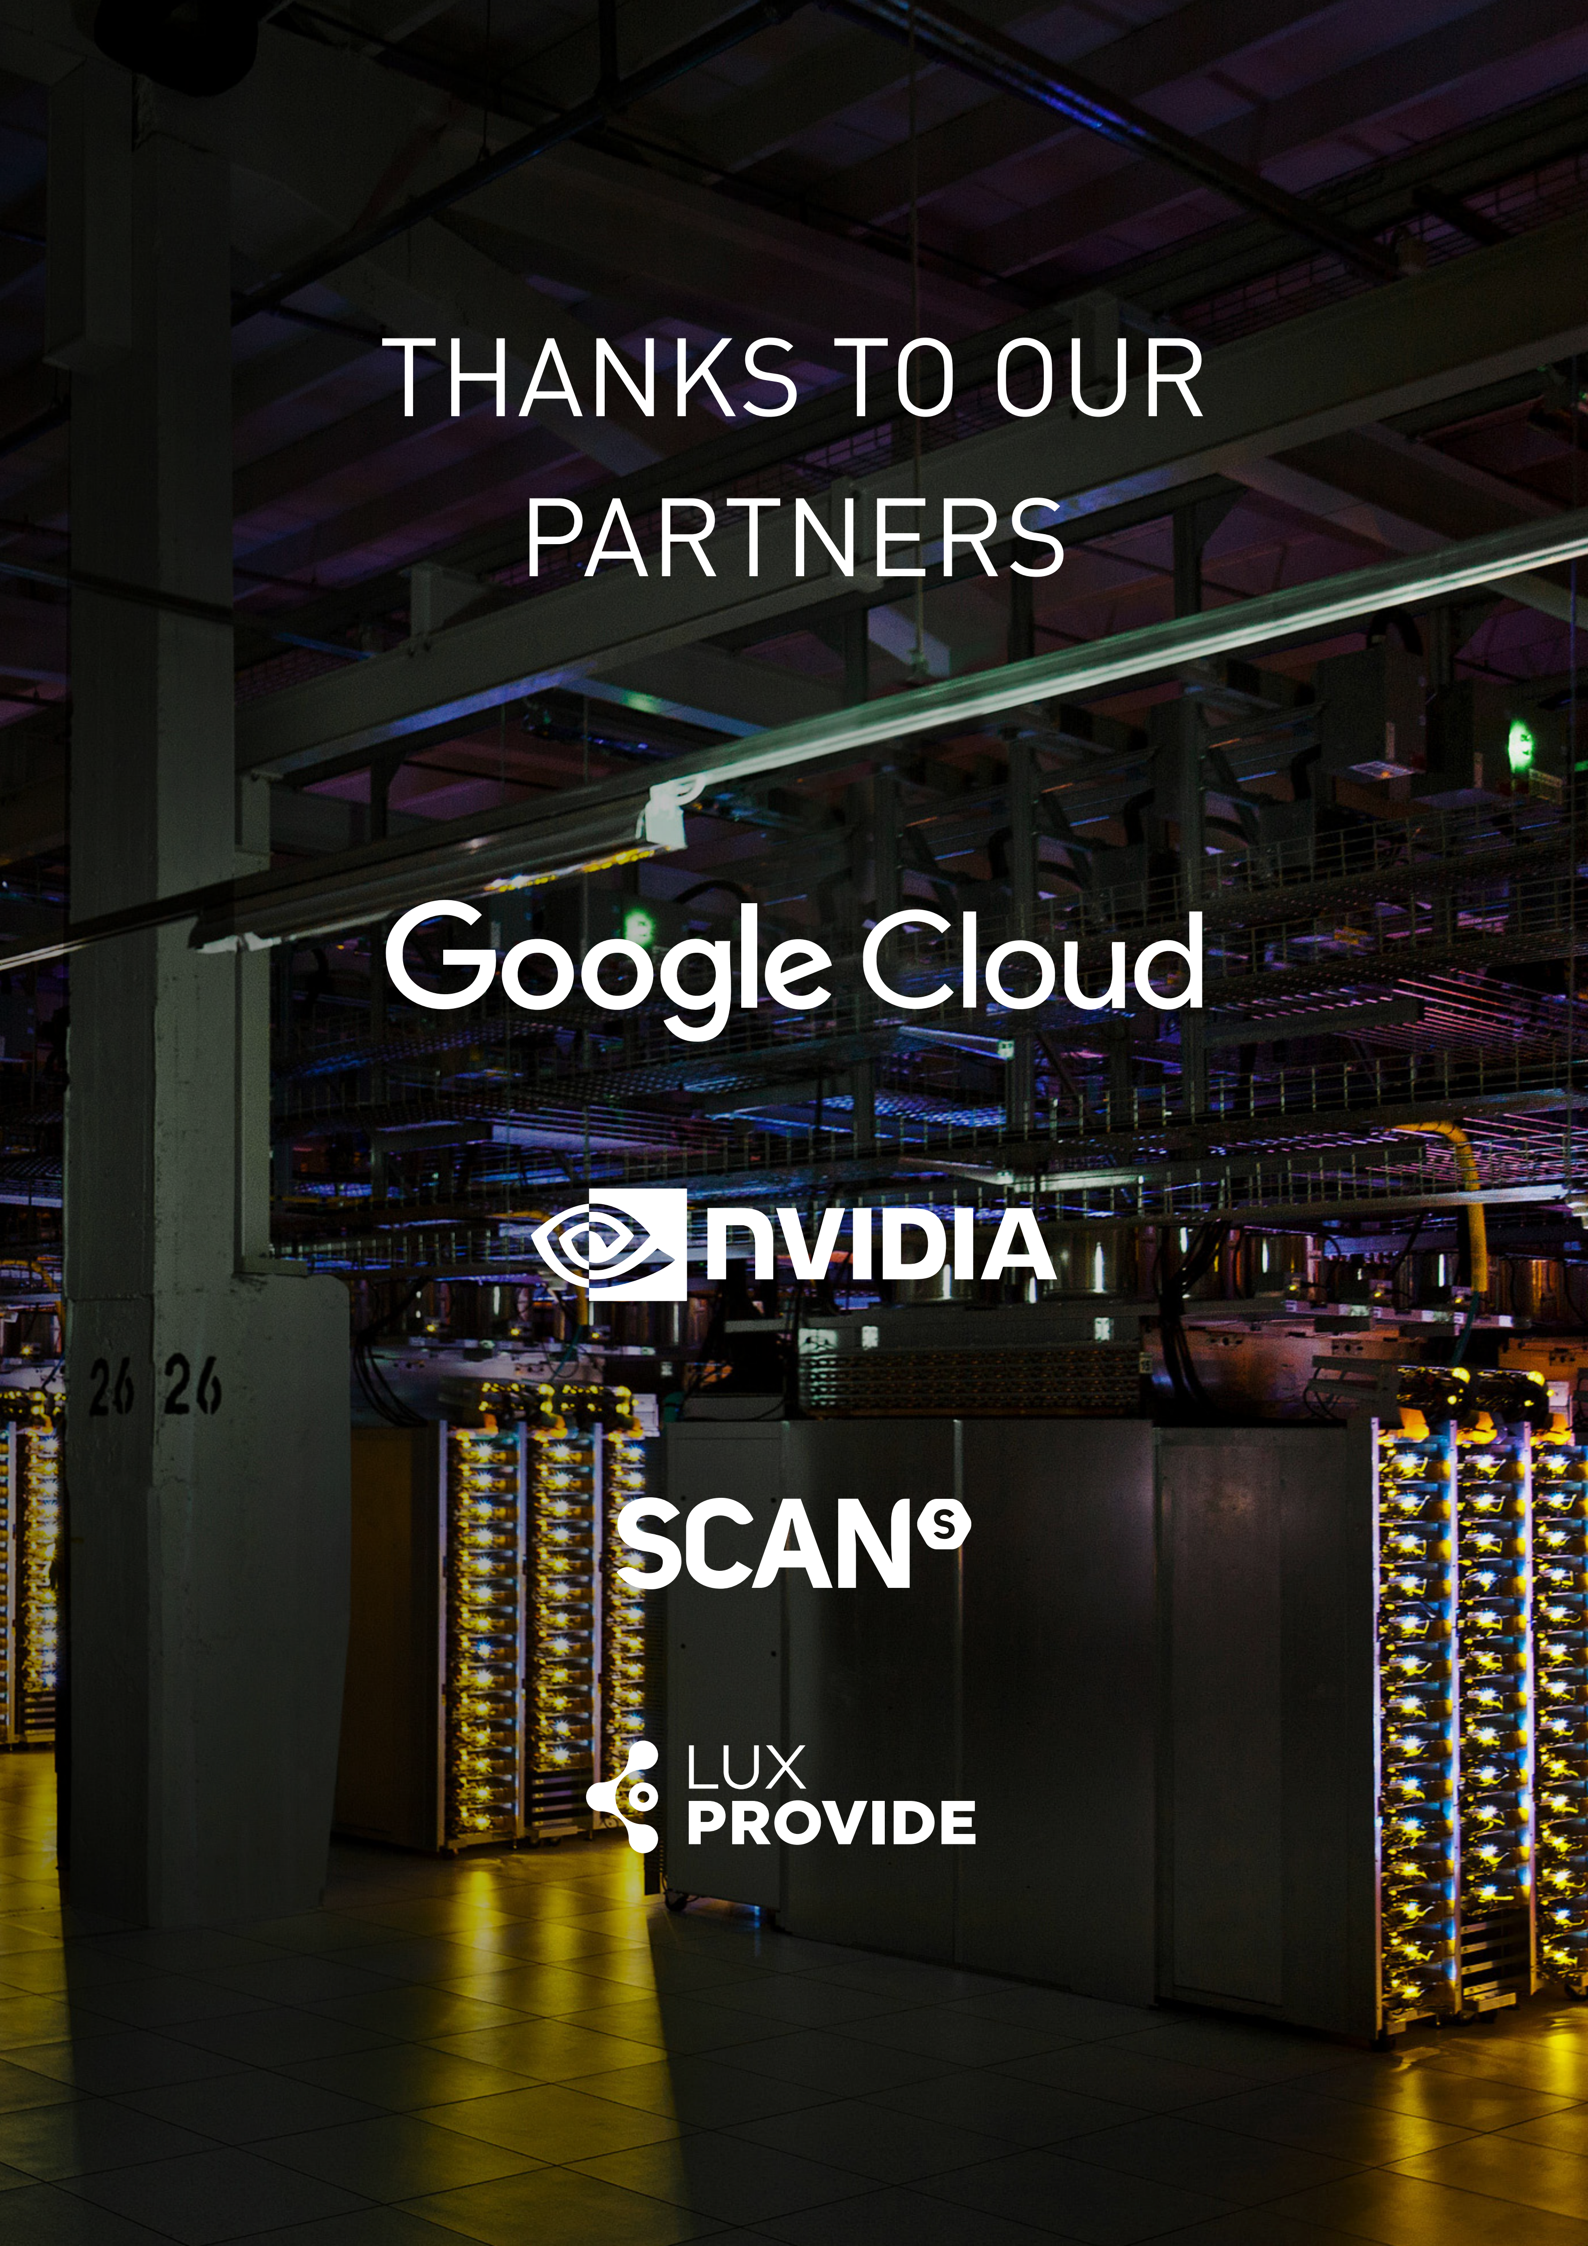

# OVER 50 APPLIED AI FIRSTS

## EARTH SCIENCE

- 2024 3D CLOUDS USING MULTI-SENSORS
- 2024 3D SAR FOR FOREST BIOMASS

### EARTH OBSERVATION

- 2023 GENERALIZABLE MODELS FOR SYNTHETIC APERTURE RADAR
- 2022 LIVE TWIN: HYDROLOGICAL MODELS
- 2022 LIVE TWIN: AEROSOLS
- 2022 SMD KNOWLEDGE GRAPH DISCOVERY
- 2022 SAR DATA FOR CHANGE DETECTION
- 2022 GEOMECHANICS FOR CO2 SEQUESTRATION
- 2022 FORTIFYING THE GRID

### MISSION OPERATIONS

- 2021 SOLAR DRAG
- 2021 WORLD FOOD EMBEDDINGS
- 2021 LIGHTNING UPGRADE
- 2021 HOT SPOT DETECTION
- 2021 AUTO IMAGE ANNOTATION
- 2021 DIGITAL TWIN COASTS

### CLIMATE

- 2020 KNOWLEDGE DISCOVERY FRAMEWORK
- 2020 DIGITAL TWIN EARTH
- 2020 EARTH INTELLIGENCE ENGINE
- 2020 LIGHTNING & EXTREME WEATHER
- 2020 WATERS OF THE US
- 2020 CLOUD BEHAVIOUR
- 2020 CONSTELLATION MGMT

### ON BOARD INTELLIGENCE

- 2019 GROUND STATION PLANNING
- 2019 CLOUD CLASSIFICATION
- 2019 DISASTER RESPONSE: FLOODS
- 2019 EDGE INFERENCE: FLOODS
- 2018 INFORMAL SETTLEMENTS

## PLANETARY DEFENCE

- 2017 LONG-PERIOD COMET DETECTION
- 2016 ASTEROID DEFLECTION DECISION SUPPORT
- 2016 RADAR SHAPE-MODELING
- 2016 METEORITE HUNTING DRONE

## ASTRONAUT HEALTH

- 2021 CAUSAL TRANSFER LEARNING
- 2020 CANCER CAUSAL INFERENCE
- 2019 SYNTHETIC MEDICAL DATA

## ASTROBIOLOGY

- 2022 INFORMATION SIGNATURES OF BIOTIC VS ABIOTIC PROCESSES
- 2018 EXOPLANET DETECTION
- 2018 ATMOSPHERIC RETRIEVAL FROM SPECTRA
- 2018 BIOSIGNATURES

FDL

.AI

## ENERGY FUTURES

- 2022 **CONCENTRATED SOLAR POWER CONTROLLER**
- 2022 **THE H2 DISCOVERY ENGINE**

## LUNAR EXPLORATION

- 2022 **CHANGE DETECTION WITH LRO DIVINIER** — **LUNAR EXPLORATION**
- 2021 **UP SCALING LUNAR RESOURCES**
- 2020 **LUNAR RESOURCE MAPPING II**
- 2019 **LUNAR RESOURCE MAPPING (METALS)**
- 2018 **ROVER LOCALISATION** — **LUNAR AUTONOMY**
- 2018 **AUTONOMOUS ROUTE PLANNING**
- 2017 **CRATER IDENTIFICATION** — **LUNAR MAPPING**

## DISASTER RESPONSE

- 2024 **SAR-FM: A FOUNDATION MODEL FOR SAR**
- 2023 **FOUNDATION MODELS ADAPTERS FOR DISASTERS**
- 2022 **URBAN RAD HUNTER**
- 2022 **WILDFIRE: MULTISPECTRAL ESTIMATION OF FUEL LOADS**
- 2020 **EARLY DETECTION** — **2020 DATAQUEST**
- 2020 **FIRE BEHAVIOUR**
- 2020 **FUEL ASSESSMENT**
- 2018 **DISASTER RESPONSE: DAMAGE ASSESSMENT**

## HELIOPHYSICS & SPACE WEATHER

- 2024 **EARLY WARNING FOR SOLAR ERUPTIONS WITH VIGIL**
- 2024 **SPECTRAL IRRADIANCE OF THE 3D SUN ON MARS**
- 2024 **FORECASTING RADIATION EXPOSURE FOR HUMAN SPACE FLIGHT**
- 2024 **GEOEFFECTIVENESS CL**
- 2024 **THERMOSPHERIC DRAG CL**
- 2023 **SPACE WEATHER AND EARLY DETECTION**
- 2023 **THERMOSPHERIC DRAG**
- 2023 **ARD EUV**
- 2023 **MULTISCALE GEOEFFECTIVENESS**
- 2022 **SEISMIC INSIGHTS WITHIN GEOMAGNETIC AND IONOSPHERIC DATA**
- 2022 **4π EUV IRRADIANCE**
- 2021 **SOLAR WIND FINGERPRINTING**
- 2020 **GEOEFFECTIVNESS OF SOLAR STORMS**
- 2020 **STARSPOTS**
- 2019 **AUTO CALIBRATION**
- 2019 **VIRTUAL TELESCOPES**
- 2019 **GNSS DISTURBANCES**
- 2019 **SUPER RES. MAG. FIELDS**
- 2018 **SDO UV IRRADIANCE**
- 2018 **IONOSPHERIC SCINTILLATIONS**
- 2017 **PREDICT KP INDEX** — **SOLAR PREDICTION**
- 2017 **C-CLASS FLARES**

Live on SPACEML · Published Research · Deployed AI · Data Product / Tool · Commercial Partnership

# Let's build an intelligent future.

**Trillium Technologies is a small international research and development company dedicated to creating intelligent systems for public benefit, planetary stewardship, space exploration and human health. With headquarters in London and virtual offices across the US and Australia, we are at the forefront of addressing some of the world's most pressing challenges using AI.**

As a specialist technology consultancy, Trillium tackles grand challenges, such as obesity, planetary defence (asteroids, comets), space situational awareness (SSA), lunar exploration, sustainability, climate change, disaster response and wildfires. Over the past decade, Trillium has established itself as a leader in applying AI to challenges on the 'to-do' list of NASA, The US DOE, USGS, ESA, CSA, LSA and the Australian Space Agency.  This work is showcased through initiatives like Heliolab (heliolab.ai), ESL (eslab.ai) and Lunarlab (lunarlab.ai)

Our team consists of researchers, scientists, designers, developers and AI specialists, all driven by the vision of achieving impactful solutions for humanity through the accelerated application of trusted intelligent technology.

Learn more and get in touch at trillium.tech or team@trillium.tech

# TRILLIUM TECHNOLOGIES


[Celebrating 10 years of Applied AI]

*In memory of David Cheuvrount,*

*thank you for helping us ta*

*lk to the moon.*

TRILLIUM TECHNOLOGIES

LUXEMBOURG SPACE AGENCY

esric

# FDL LUNARLAB

FOR ALL HUMANKIND

FDL.AI NETWORK

10 YEARS OF APPLIED AI RESEARCH FOR ALL HUMANKIND

Google Cloud NVIDIA SCAN

Datarock LUX PROVIDE

AD ASTRA PER ALGORITHMOS

LUNARLAB.AI