
Unsupervised learning for thermal anomaly detection on the lunar surface

Ben Moseley

Department of Computer Science
University of Oxford

Valentin Bickel

ETH Zurich/ Max Planck Institute
Zurich/ Goettingen

Jérôme Burelbach

NASA Frontier Development Lab
Mountain View, CA

Nicole Relatores

Carnegie Observatories
Pasadena, CA

Daniel Angerhausen

Center for Space and Habitability
Blue Marble Space Institute of Science
University of Bern/ Seattle, WA

Frank Soboczinski

SPHES
King's College London

Dennis Wingo

Skycorp Incorporated
Mountain View, CA

Abstract

Humans will soon be returning to the Moon, with the goal of establishing permanent bases within the next 20 years. A key challenge in this endeavour is to locate resources on the lunar surface, such as metals, which can be used to build such settlements. Although the Moon has been extensively mapped across multiple wavelengths, inferring the presence of metals from this data is challenging because there exists very little ground truth observations to calibrate physical models to. We take a data-driven approach and use unsupervised learning to search for metallic signatures. We train a convolutional variational autoencoder (VAE) on thermal satellite data and find it is able to accurately reconstruct temperature variations over the lunar day. Furthermore we find physically-interpretable correlations between the VAE's latent representation and estimated thermal parameters from physics-based inversion. Using this representation we are able to efficiently generate new types of thermal anomaly maps which potentially indicate the presence of metals on the surface of the Moon.

1 Introduction

Establishing a permanent settlement on the lunar surface is one of the most pioneering endeavours humanity has embarked upon. NASA's Artemis Program plans to land a manned mission on the Moon by 2024, with further plans to deploy a lunar outpost in 2028 (Berger, 2019; NASA, 2018, 2019). In order for humanity to establish a sustainable presence on the Moon, it is critical to understand what resources are available on the lunar surface.

Metal is a key resource and can be deposited on the surface from impactor influx in the form of metal meteorites; it is estimated that 3-4% of impacts are from nickel-iron bodies. Nickel-iron meteoric fragments have also been found embedded in samples from the Apollo missions (Wingo, 2005). However, while large Iron-nickel meteorites have been found on Earth as well as on the surface of Mars, they have not yet been found on the Moon, despite the overwhelming amount of remote sensing data. This presents a challenge as there is very little ground truth to calibrate physical models to.

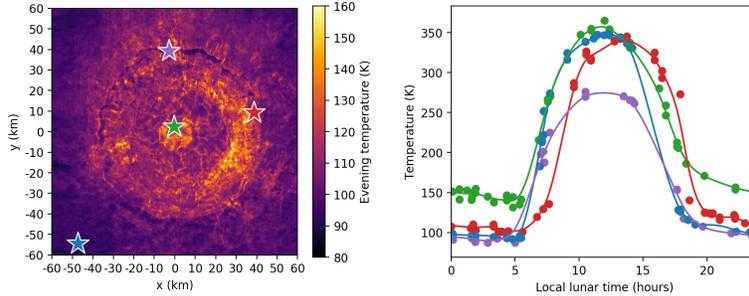


Figure 1: Left: evening temperature map over the Tycho crater, plotting the average temperature over 12 am-4 am in local lunar time. Right: four example temperature profiles extracted over Tycho. Points show temperature measurements, solid lines show the Gaussian process fits of the profiles used for interpolation. Colour coded stars in the left plot indicate the locations of the profiles.

The presence of metal meteorites is known to influence the thermal properties of the lunar surface, providing a potential method for detection (Bandfield et al., 2011, 2015; Zheng et al., 2014; Hayne et al., 2017; Hu et al., 2018; Miller, 2018). In particular, metals are expected to have a higher thermal conductivity than the lunar regolith. However, there exists uncertainty in the accuracy of these thermal models because of the lack of ground truth observations. The Lunar Reconnaissance Orbiter is a satellite which has acquired over 35 TB of point measurements of the surface temperature during 10 years of operation at a resolution of the order of 100 m using its DIVINER instrument, a passive infrared radiometer (Williams et al., 2017, 2018; Sefton-Nash et al., 2017).

We take a data-driven approach and apply unsupervised machine learning to search for thermal anomalies in the DIVINER data. We train a convolutional variational autoencoder (VAE) to reconstruct variations in the surface temperature over the lunar day. We compare the VAE’s reconstructions and latent representation to a 1D physics-based inversion. Finally we use its latent representation to efficiently generate new types of thermal anomaly maps which could indicate the presence of metals.

2 Methods

To search for thermal anomalies we extract all the temperature point measurements recorded by the DIVINER instrument over 2010-19. We use channel 7 of the instrument, which measures radiance in the 25-41 μm band and is most sensitive to evening and night-time temperatures in the range 69-178 K. We bin these temperature measurements onto a 200 x 200 m grid and sort the points in each bin by the local lunar time at which they were recorded, obtaining a 1D profile of the temperature variation over the lunar day at each bin location.

We extract examples of these profiles over 43 areas of interest (AOI). We select AOIs based on existing physical maps; 38 have anomalies in their mean annual temperature or total magnetic field strength, which we believe may indicate the presence of metals; 5 are background locations sampled from the majority of the lunar surface which does not show these anomalies. The AOIs include impact craters of different ages, large pyroclastic deposits and swirls (magnetic anomalies), with areas ranging from 1 - 40,000 km^2 . We extract profiles whose maximum temporal sampling is lower than four lunar hours, which results in 112,214 example profiles.

We search for anomalies in the profiles by using a VAE. We use the VAE to compress each profile into a small set of latent variables and then map these latent variables over the lunar surface. We chose a VAE to encode the profiles because of its ability to separate independent factors of variation from its input distribution (Kingma and Welling, 2013). Before inputting the profiles into the VAE, we interpolate them onto a regular grid by using Gaussian Process (GP) regression (Rasmussen and Williams, 2005). We fit a GP to each profile with a Matern-1.5 kernel, maximum length scale of 6 hours and assume 10 K standard deviation of uncertainty in each temperature measurement. We sample every 0.2 hours, resulting in a 1D input profile with 120 samples.

We use a convolutional encoder-decoder for our VAE architecture, with eight convolutional layers in the encoder and seven transposed convolutional layers in the decoder. All hidden layers use ReLU

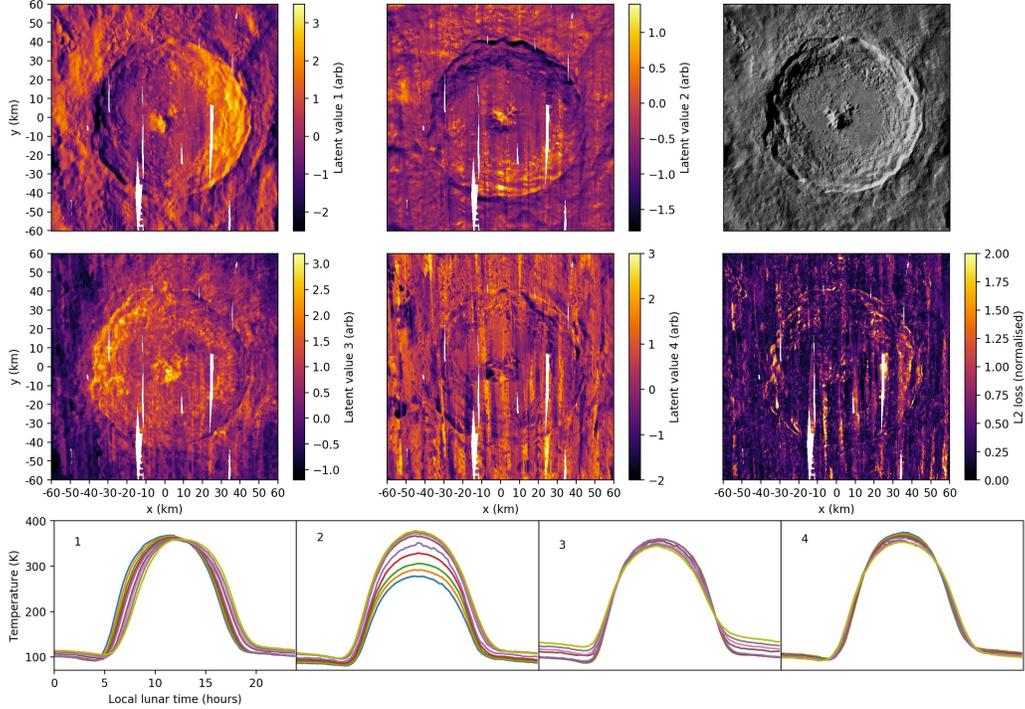


Figure 2: Top left two columns show maps of the four VAE latent variables over the Tycho crater. Top right shows an optical image of Tycho. Middle right shows a map of the VAE L2 reconstruction loss. Bottom plots show the reconstructed profiles generated from the VAE when sampling each latent variable independently and fixing the other latent variables to their mean value.

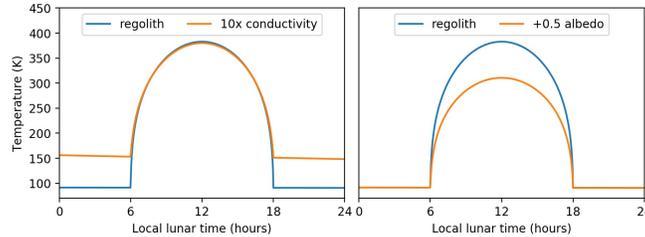


Figure 3: Left: 1D numerical simulations of the lunar surface temperature against local lunar time. Our baseline model uses thermal properties estimated for the lunar regolith. The two plots show the resulting profiles when the thermal conductivity is increased by a factor of 10 and when the albedo is increased by 0.5.

activation functions, batch normalisation, 32 hidden channels, kernel sizes of 2 and strides of 2. The final layers of the encoder and decoder have linear activation functions. The output of the encoder is a latent vector, its length being a hyperparameter, which is reparameterised before being fed into the decoder. The VAE is trained to reconstruct an input profile using a loss function which sums the mean squared error and KL divergence between the input and reconstructed profile, with a β -value of 0.2 (Higgins et al., 2017). We train using the Adam optimiser, a learning rate of 1×10^{-4} and a mini-batch size of 200 samples, reserving 20% of the example dataset for cross validation.

We compare our results to a 1D physics inversion. We build a physics model, as proposed by Miller (2018), simulating the temperature of a surface of material in a vacuum which receives the same solar radiation over time as a point on the equator of the Moon. We run inversion using this model, allowing the inversion algorithm to vary the albedo, thermal conductivity and onset time of solar radiation to best match each input interpolated profile, using a L2 reconstruction loss. Bayesian optimisation with an expected improvement acquisition function is used as the inversion algorithm

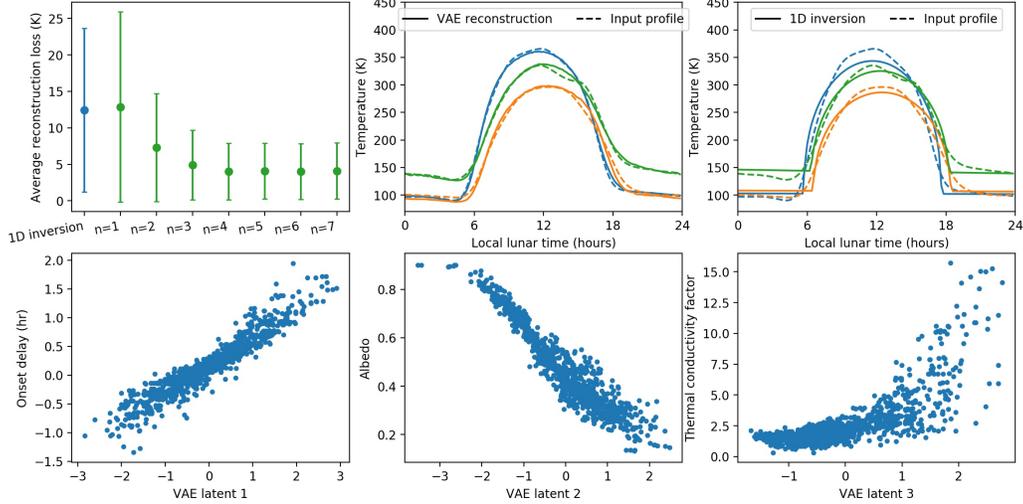


Figure 4: Top left: Average L1 reconstruction loss for VAEs with varying latent vector lengths n and the physics inversion, over 1000 profiles from the validation set. Top middle and right: Example reconstructions from the VAE with $n = 4$ and the 1D physics inversion, for 3 randomly selected profiles in the validation set. Bottom: plots showing the correlation between the first 3 VAE latent values when $n = 4$ and the estimated parameters from physics inversion, over 1000 profiles from the validation set. Thermal conductivity scale is relative to the regolith.

(Frazier, 2018). We then compare the estimated parameters from the inversion to the latent variables generated by the VAE.

3 Results and Discussion

Figure 1 shows example temperature data from the DIVINER instrument over one of our AOIs, the Tycho crater. An optical image of this crater is shown in Figure 2. We observe a wide variation in temperature around this crater. Outside of the crater, the surface temperature increases to around 350 K at midday and decreases to 100 K at night. On the Eastern slope of the crater the profiles are shifted in time by approximately 2.5 hours, which is consistent with the slope’s shadow delaying the incidence time of the Sun’s radiation. On the Northern slope the peak temperature is lower at 270 K, which is consistent with the slope receiving less thermal radiation per unit area from the Sun’s elevation at Tycho’s latitude. In the centre of the crater we observe a profile which has a higher evening temperature around 150 K. Figure 3 (left) plots our physics model simulation when varying thermal conductivity; one explanation is that the material at this location has a higher thermal conductivity than its surroundings, which could potentially indicate the presence of metal.

Figure 4 (top) shows the average reconstruction loss after training the VAE, when varying the number of latent variables, n . We observe that this reconstruction loss asymptotes at $n = 4$. We also plot this loss after the 1D physics inversion, which is significantly higher than the VAE loss. Figure 4 suggests that only 4 latent values are needed to describe the temperature profiles and furthermore the VAE is more accurate at reconstructing the profiles than the physics model.

Figure 2 shows example profiles generated by the VAE when varying each latent variable independently and fixing the others to their mean values, for the VAE with $n = 4$. We observe that latent variable 1 responds to the onset of the profile, latent 2 responds to the peak temperature, latent 3 responds to the evening temperature “sidelobe” and latent 4 appears to learn a residual. We compare the latent values with the estimated parameters from 1D physics inversion in Figure 4 (bottom). We find that latent 1 and 2 strongly correlate to the estimated solar radiation onset time and albedo respectively, whilst latent 3 somewhat correlates with the estimated thermal conductivity.

Given this physical interpretation, we generate maps of the latent variables to search for anomalies. We note that generating these maps is two orders of magnitude faster than running the inversion, which

allows this approach to be scaled over large areas. Maps of the variables and the L2 reconstruction loss over Tycho are shown in Figure 2. Latent 3 is particularly interesting as it shows high values on the Western side and centre of the crater which could indicate high thermal conductivity, potentially related to metals. The L2 loss map can also be used to identify locations with extreme thermal behaviour outside of the VAE training distribution which could relate to high contrasts in chemical or mineralogical composition.

4 Conclusion and Further Work

We have shown that unsupervised learning can aid the detection and understanding of thermal anomalies on the lunar surface. We trained a VAE which was able to accurately encode variations in the lunar surface temperature. When combined with physics-based inversion, its latent representation improved our understanding of the physical factors of variation in the surface temperature. Its latent representation allowed us to generate new types of anomaly maps, two orders of magnitude faster than inversion, which potentially indicate the presence of metals.

Further work includes producing global anomaly maps with our workflows, building more complex physics models to aid our physical interpretation of the VAE and incorporating physics models directly in the training of our VAE to constrain our latent representation further.

Acknowledgements

This work was completed as part of the NASA Frontier Development Lab (FDL). We would like to thank our FDL mentor Daniel Angerhausen (University of Bern) for helping us with the 1D thermal simulations; our FDL mentors Frank Soboczinski (King’s College London) and Dennis Wingo (Skycorp Incorporated) for their general advice; Jean-Pierre Williams (UCLA) for his help in understanding the DIVINER instrument; the FDL and all of its partners (Luxembourg Space Agency, Google Cloud, Intel AI, Hewlett Packard Enterprise, Element AI and USGS); the SETI Institute for hosting the FDL; and our other FDL mentors; Abigail Calzada Diaz (iSpace), Alfredo Kalaitzis (Element AI), Ignacio López-Francos (NASA Ames) and Allison Zuniga (NASA Ames).

References

- Eric Berger. Nasa’s full artemis plan revealed: 37 launches and a lunar outpost. May 2019. URL <https://arstechnica.com/science/2019/05/nasas-full-artemis-plan-revealed-37-launches-and-a-lunar-outpost/>.
- NASA. National space exploration campaign report, pursuant to section 432(b) of the nasa transition authorization act of 2017 (p.l. 115-10). September 2018. URL <https://www.nasa.gov/sites/default/files/atoms/files/nationalspaceexplorationcampaign.pdf>.
- NASA. Forward to the moon: Nasa’s strategic plan for human exploration. May 2019. URL https://www.nasa.gov/sites/default/files/atoms/files/america_to_the_moon_2024_artemis_20190523.pdf.
- Dennis Wingo. *Moonrush; Improving Life on Earth with the Moon’s Resources*. Apogee Books, 2005.
- Joshua L. Bandfield, Rebecca R. Ghent, Ashwin R. Vasavada, David A. Paige, Samuel J. Lawrence, and Mark S. Robinson. Lunar surface rock abundance and regolith fines temperatures derived from lro diviner radiometer data. *Journal of Geophysical Research: Planets*, 116(E12), 2011. doi: 10.1029/2011JE003866. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2011JE003866>.
- Joshua L. Bandfield, Paul O. Hayne, Jean-Pierre Williams, Benjamin T. Greenhagen, and David A. Paige. Lunar surface roughness derived from lro diviner radiometer observations. *Icarus*, 248: 357 – 372, 2015. ISSN 0019-1035. doi: <https://doi.org/10.1016/j.icarus.2014.11.009>. URL <http://www.sciencedirect.com/science/article/pii/S0019103514006368>.
- Y.C. Zheng, K.L. Chan, K.T. Tsang, Y.C. Zhu, G.P. Hu, Y.L. Zou, and Z.Y. Ouyang. Catalogue of lunar thermal anomalies. In *45th Lunar and Planetary Science Conference*, 2014.

- Paul O. Hayne, Joshua L. Bandfield, Matthew A. Siegler, Ashwin R. Vasavada, Rebecca R. Ghent, Jean-Pierre Williams, Benjamin T. Greenhagen, Oded Aharonson, Catherine M. Elder, Paul G. Lucey, and David A. Paige. Global regolith thermophysical properties of the moon from the diviner lunar radiometer experiment. *Journal of Geophysical Research: Planets*, 122(12):2371–2400, 2017. doi: 10.1002/2017JE005387. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2017JE005387>.
- Guo-ping Hu, Kwing Chan, Yong-Chun Zheng, and Ao-Ao Xu. A rock model for the cold and hot spots in the chang’e microwave brightness temperature map. *IEEE Transactions on Geoscience and Remote Sensing*, 56:1–10, 04 2018. doi: 10.1109/TGRS.2018.2817654.
- Mark A. Miller. *Modeling Planetary Climates with Python*. 2018. URL <https://github.com/markmillr/whitebark>.
- J.-P. Williams, D.A. Paige, B.T. Greenhagen, and E. Sefton-Nash. The global surface temperatures of the moon as measured by the diviner lunar radiometer experiment. *Icarus*, 283:300 – 325, 2017. ISSN 0019-1035. doi: <https://doi.org/10.1016/j.icarus.2016.08.012>. URL <http://www.sciencedirect.com/science/article/pii/S0019103516304869>. Lunar Reconnaissance Orbiter - Part II.
- J.-P. Williams, J. L. Bandfield, D. A. Paige, T. M. Powell, B. T. Greenhagen, S. Taylor, P. O. Hayne, E. J. Speyerer, R. R. Ghent, and E. S. Costello. Lunar cold spots and crater production on the moon. *Journal of Geophysical Research: Planets*, 123(9):2380–2392, 2018. doi: 10.1029/2018JE005652. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018JE005652>.
- E. Sefton-Nash, J.-P. Williams, B.T. Greenhagen, K.-M. Aye, and D.A. Paige. Diviner lunar radiometer gridded brightness temperatures from geodesic binning of modeled fields of view. *Icarus*, 298:98 – 110, 2017. ISSN 0019-1035. doi: <https://doi.org/10.1016/j.icarus.2017.04.007>. URL <http://www.sciencedirect.com/science/article/pii/S0019103517302737>. Lunar Reconnaissance Orbiter - Part III.
- Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv e-prints*, art. arXiv:1312.6114, Dec 2013.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005. ISBN 026218253X.
- Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew M Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017.
- P. I. Frazier. A Tutorial on Bayesian Optimization. *ArXiv e-prints*, July 2018.