# Learning Physically Informed Maps for Off-world Autonomous Scientific Discovery

Margaret Hansen[1]*and David Wettergreen[1]

***Abstract*— Planetary exploration with mobile robots has relied on scientist selection of sites for scientific observation and sampling. Recent advancements in informative path planning enable adaptive site selection; however, the underlying models of information used in such planners remain limited. Integrating physical process knowledge into these models can improve predictive capability, provide richer estimates of the underlying natural phenomena, and eventually allow for onboard analysis of scientific hypotheses. These capabilities are demonstrated in this work on a computationally synthesized distribution of lunar volatiles by incorporating underlying scientific knowledge into mapping with deep Gaussian process regression.**

## I. Introduction

Remote robotic exploration of non-terrestrial planetary bodies has mainly proceeded with input from human operators on the ground [1]. However, limitations such as travel time and communications bandwidth preclude reliance on ground operators for robotic operations in the outer solar system and beyond [2]. Even on the moon and Mars, awaiting instruction from the ground can prevent efficient adjustment of sampling strategies as a result of incoming information. Intelligent systems are capable of adapting to novel observations, incorporating scientific hypotheses and models, and distilling data into meaningful packets prior to downlink to Earth; autonomy of this sort enables future planetary exploration missions to the moon, Mars, and even more distant destinations.

Informative path planning (IPP), which maximizes information gain based on a map of expected observations and their associated uncertainties, has made progress in this direction. While IPP is applicable for tasks ranging from geology and site surveying to environmental monitoring [3–5], these planners tend to oversimplify the scientific information available at each location by computing a single information-theoretic metric, frequently entropy, from a basic spatial model of observations. Some methods learn correlations or links between variables [6,7], but the inclusion of scientific information should be extended.

This work incorporates physical process models into rover mapping and examines its specific application to lunar volatiles prospecting. By drawing on physics-informed machine learning literature, for which there has been ample development of methods that learn differential equations or sensor models, this mapping approach expands the types of scientifically-relevant information available for planning and analysis [9,10]. The resulting models are more explainable than correlation-based or black-box methods and are capable of providing predictions and uncertainty estimates of underlying natural phenomena in addition to sensor measurements. This technique, termed *physical process-informed mapping*, has been tested on data synthesized from geologic simulation of the lunar surface and representing the distribution of volatiles in lunar south polar regions. These experiments demonstrate its applicability to a planetary exploration scenario for which many upcoming robotic missions are planned: NASA's VIPER and Endurance rovers, Astrobotic's MoonRanger rover, and China's Chang'e-7 [11–14].

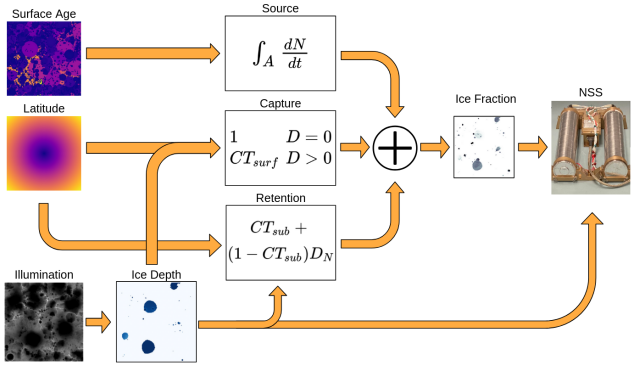


Fig. 1: Diagram depicting the deep Gaussian process regression model applied to lunar volatiles, based on the ice favorability index [8]. Input variables include the illumination, surface age, and latitude, hidden variables are ice depth and ice fraction, along with the source, capture, and retention terms, and outputs are observations from a neutron spectrometer system (NSS). The included equations are used to develop mean functions for the corresponding layers of the deep GP. Variables combined using the additive model are log-transformed before and after combination.

## II. Methodology

Physical process-informed mapping (PPIM) incorporates physical knowledge into Gaussian process-based mapping for mobile robot exploration. Data-driven approaches, namely including process variables as inputs or additional outputs, are relatively straightforward and have already been used to model NSS observations in the Mojave desert [15,16] and rainfall in Australia [6]. Architecture-based approaches include hierarchical structures designed to integrating hypothesized information into mapping [3], along with combinations of Gaussian processes and graphical models for learning linkages between relevant variables [7]. The for-

[1] Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA.
* Corresponding author. margareh@cs.cmu.edu.

| Model | NSS MAE (↓) | Depth MAE (↓) | Wt % MAE (↓) |
|---|---|---|---|
| Spatial Baseline | 0.8374 | – | – |
| PV Baseline | 1.8669 | – | – |
| Spatial + PV Baseline | 0.8601 | – | – |
| MOGP Baseline | 0.9448 | – | – |
| Simple Deep | 0.4386 | 1.6260 | 0.0565 |
| Simple Deep + LV Loss | 0.4134 | 2.9709 | 0.0750 |
| Full Deep | 0.4540 | 1.8914 | 0.1080 |
| Full Deep + LV Loss | 0.4584 | 2.5425 | 0.0735 |

TABLE I: Mean absolute error over both NSS detectors, as well as for the depth and weight fraction, for each model. Baselines are listed in the top panel while deep GP models are in the bottom panel.

mer, however, does not simultaneously estimate the entire hierarchical model, while the latter learns structural connections between variables but does not analyze specific models proposed by science. This work uses deep Gaussian processes to estimate a hierarchical model between variables that integrates scientific knowledge, streamlining the model into one larger structure that allows extraction of a variety of key components for planning.

### A. Deep Gaussian Processes

Deep Gaussian processes allow hierarchical modeling by stacking multiple Gaussian processes into layers in a manner similar to a neural network [17,18]. In addition to predicting the observations associated with sampling locations or other covariates, the deep GP hierarchical structure can enable learning "hidden" variables corresponding to natural phenomena. Unlike standard GP regression, deep GPs require approximate inference due to the layering approach. Variational inference is used in this work, with inducing point locations learned during training [17].

### B. Application to Lunar Volatiles Mapping

Physical process-informed mapping can be applied to lunar volatiles mapping by incorporating relevant physical process variables, extending the GP regression architecture to use deep GP models, and relying on a simulated calibration curve linking the depth and weight fraction of ice to NSS observations for a model-driven loss term. Incorporating process variables in the data used to train the model is a fairly straightforward data-driven approach, with the key improvements made in this work stemming from the deep GP structure and loss term.

Two deep GP architectures are tested, each with and without the extra loss term (denoted with "+ LV Loss" when used). The first architecture ("Simple Deep") models the NSS observations as a function of spatial locations with one hidden layer of two nodes. The second architecture ("Full Deep", represented by Figure 1) relies on the ice favorability index model from [8] to link illumination, age, and depth to ice favorability, which is used as a proxy for weight fraction in these tests.

The four combinations of deep architectures and loss terms are tested against four baselines: a basic spatial Gaussian process regression (the "Spatial Baseline"); a basic regression using process variable inputs only ("PV Baseline"); a combination of the two ("Spatial+PV Baseline"); and a multi-output model ("MOGP Baseline")[1] similar to the model from [6]. This last model predicts NSS observations along with the other process variables as covarying functions of location. The baseline models are all trained using exact inference. All models use the Matérn kernel with $\nu = 1.5$ and constant means, except for the full deep GP models, for which the mean functions are defined based on the equations shown in Figure 1.

## III. Simulating Lunar Volatiles Distributions

The lunar south polar region is hypothesized to contain subsurface and surficial volatiles such as water ice in permanently shadowed regions due to consistently low lighting angles, making it a particularly important target for exploration missions for both lunar science and resource prospecting [11]. As such, it is a compelling case study for physical process-informed mapping. This work relies on a simulation of lunar volatiles across geologic time scales (3.8 Gyr) to produce realistic "present day" maps of their depth and weight fraction. This simulation iteratively adds craters using a crater size-frequency distribution for the lunar south polar region, which serves to both move regolith and deliver volatiles, along with also simulating the effects of ballistic hopping and impact gardening. The `synthterrain` software is used to add craters and degrade the terrain over time in a realistic manner [19], while the `MoonPIES` simulation is used to deliver, move, and remove volatiles based on the crater distribution. This latter simulation was edited to allow for modeling an entire region instead of individual craters [20]. Finally, a neutron spectrometer system (NSS), one of the instruments designed to detect near-surface (within 1 m) hydrogen, is also simulated based on a modeled instrument calibration curve [11,13,21].

## IV. Experiments

Each model was learned using a 20 x 20 grid of observations sampled from one 200 m x 200 m map produced by the lunar volatiles simulation described in Section III. For all models, 1000 training iterations and a learning rate of 0.01 was used. These tests show that the deep Gaussian process models outperform the best baseline model,

[1]Since the NSS produces two values, technically all models are multi-output in this work. The MOGP baseline adds output variables to the set.

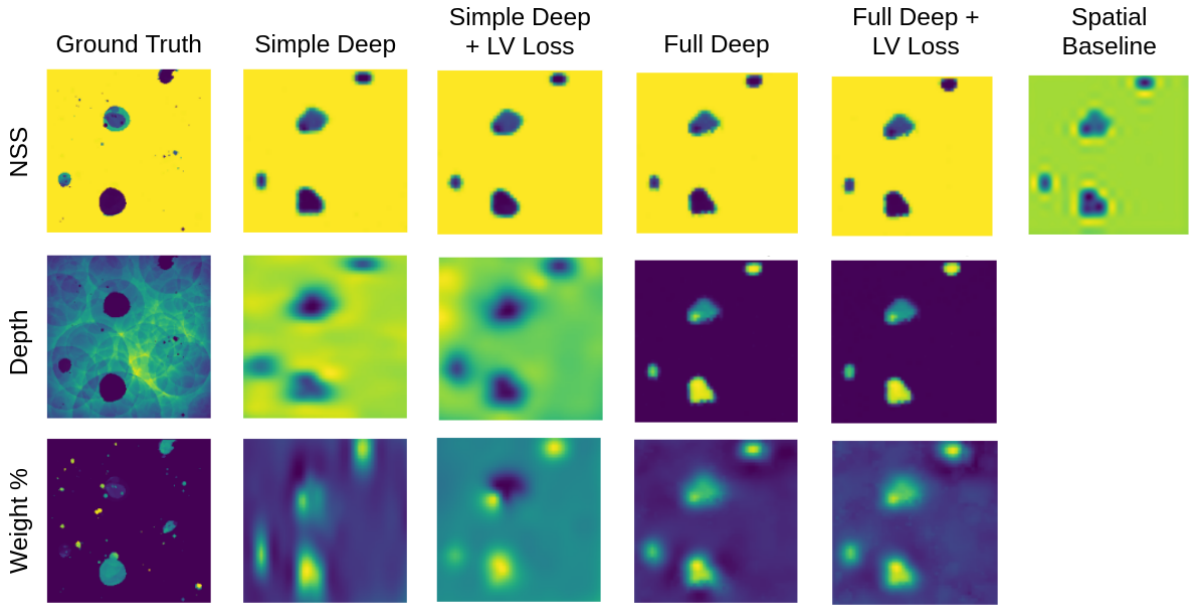


Fig. 2: Map predictions from the four deep GP models models for NSS detector 1, depth, and weight fraction. The ground truth for each value is shown on the left of each row. The predicted NSS map from the spatial baseline is included on the right for comparison. For all maps, the color scale ranges from purple indicating low values to yellow indicating high values.

the spatial baseline, in terms of the mean absolute error over all test data points and both NSS detectors (Table I). Interestingly, the simplest baseline model outperformed the other baselines, indicating that there are likely only weak relationships between the process variables included and the simulated NSS observations. The simple deep model outperforms the full deep model in terms of mean absolute error of the NSS observations, depth, and weight fraction, which is intuitive given that this model is more flexible in what it can learn; the full model has a more specific structure that can be restrictive.

The strength of the deep models is exemplified in Figure 2, which displays predicted values for one NSS detector,[2] along with the hidden layers corresponding to depth and weight fraction. For models without the additional loss term, nothing is explicitly driving the hidden layers to represent the depth and weight fraction; however, the learned latent variables still resemble the depth and weight fraction in terms of pattern. Interestingly, the simple deep model seems to learn a more correct representation for depth and weight fraction than the model with the loss incorporated.

Using deep models in this manner allows extraction of the distributions associated with depth and weight fraction, the two natural quantities of interest. Distributional information of these can allow an informative path planner to plan over the uncertainty associated with the depth or weight fraction rather than over the observational uncertainty. Richer models like the full deep model can likewise provide distributional information on specific aspects of scientific models (i.e. the source, capture, and retention terms in this example). The deep GP models also tend to be better calibrated than the baselines, exhibiting higher correlations between variance and error (around 0.5-0.6 instead of around 0.04-0.1) [22]. This link is important for mobile robotic exploration since uncertainty is frequently used to determine where to sample.

[2]The performance trends are very similar across both detectors.

## V. Conclusion

These results indicate that incorporating additional structure and knowledge of physical processes into Gaussian process regression can provide a number of added benefits for informative path planning. Such models not only improve the predictive performance of the model, as evidenced by lower mean absolute error of NSS observations, but they are also capable of modeling quantities of interest that are only indirectly sensed. This structure allows for more adaptive choices of objective to be made when planning. Finally, directly including physical model components into the structure enables learning uncertainty associated with these components, providing some sense of how useful they are in analyzing the natural phenomena under study.

Additional experimentation is needed to vet this methodology, including further and more in-depth mapping experiments to analyze the performance across multiple maps and assess statistical significance in performance metrics. In addition, the deep GP models take longer to train and for prediction than the simple models. However, these models remained trainable on a standard laptop CPU within 10 minutes; with the advent of increased compute on planetary exploration robots [23], learning such models onboard is likely possible in near-real-time. Additional improvements are likely possible, including leveraging GPU computation or sparse GP regression [24].

**Acknowledgements** This work was supported by the NASA Space Technology Graduate Research Opportunity (80NSSC22K1206) and NASA LSITP award 80MSFC20C0008 to Astrobotic Technologies. The authors would additionally like to thank Richard Elphic for his help understanding the neutron spectrometer system (NSS).

## REFERENCES

[1] John L. Bresina, Ari K. Jonsson, Paul H. Morris, and Kanna Rajan, "Activity planning for the Mars Exploration Rovers," in *Proceedings of the Fifteenth International Conference on Automated Planning and Scheduling*, 2005.

[2] Masahiro Ono, Richard Reiber, Tony Freeman, Mathieu Choukroun, Michel D. Ingham, Chloe Gentgen, David Murrow, and Daniel Selva, "Is Simplicity Golden? A Survey of Post-Launch Adaptation in Planetary Missions and Lessons Learned," in *2026 IEEE Aerospace Conference*, 2026.

[3] A. Candela, D. Thompson, E. N. Dobrea, and D. Wettergreen, "Planetary robotic exploration driven by science hypotheses for geologic mapping," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 3811–3818, iSSN: 2153-0866.

[4] D. R. Thompson and D. Wettergreen, "Intelligent Maps for Autonomous Kilometer-Scale Science Survey," 2008.

[5] G. Hitz, E. Galceran, M.- Garneau, F. Pomerleau, and R. Siegwart, "Adaptive continuous-space informative path planning for online environmental monitoring," *Journal of Field Robotics*, vol. 34, Dec. 2017.

[6] N. Harrison, N. Wallace, and S. Sukkarieh, "Automated Testing of Spatially-Dependent Environmental Hypotheses through Active Transfer Learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 17 941–17 947.

[7] B. Ayton, "Risk-bounded autonomous information gathering for localization of phenomena in hazardous environments," Ph.D. dissertation, Massachusetts Institute of Technology, 2017.

[8] K. M. Cannon and D. T. Britt, "A geologic model for lunar ice deposits at mining scales," *Icarus*, vol. 347, p. 113778, Sep. 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0019103520301652

[9] J. Pateras, P. Rana, and P. Ghosh, "A Taxonomic Survey of Physics-Informed Machine Learning," *Applied Sciences*, vol. 13, no. 12, p. 6892, Jan. 2023, number: 12 Publisher: Multidisciplinary Digital Publishing Institute.

[10] G. Camps-Valls, J. Verrelst, J. Munoz-Mari, V. Laparra, F. Mateo-Jimenez, and J. Gomez-Dans, "A Survey on Gaussian Processes for Earth-Observation Data Analysis: A Comprehensive Investigation," *IEEE Geoscience and Remote Sensing Magazine*, vol. 4, no. 2, pp. 58–78, Jun. 2016, conference Name: IEEE Geoscience and Remote Sensing Magazine.

[11] K. Ennico-Smith, A. Colaprete, D. Lim, and D. Andrews, "The VIPER Mission, a Resource-Mapping Mission on Another Celestial Body," in *Space Resources Roundtable XXII Meeting*, Colorado School of Mines, Jun. 2022.

[12] J. D. Baker, J. O. Elliott, J. T. Keane, N. R. Khan, R. P. Kornfeld, H. D. Nayar, and I. A. Nesnas, "The endurance lunar rover sample return mission," in *2024 IEEE Aerospace Conference*, 2024, pp. 1–13.

[13] A. Breitfeld and D. Wettergreen, "Developing local trajectory planning for a lunar micro rover," in *2024 IEEE Aerospace Conference*, 2024, pp. 1–16.

[14] C. Wang, Y. Jia, C. Xue, Y. Lin, J. Liu, X. Fu, L. Xu, Y. Huang, Y. Zhao, Y. Xu, R. Gao, Y. Wei, Y. Tang, D. Yu, and Y. Zou, "Scientific objectives and payload configuration of the chang'e-7 mission," *National Science Review*, vol. 11, no. 2, p. nwad329, 12 2023. [Online]. Available: https://doi.org/10.1093/nsr/nwad329

[15] G. Foil, T. Fong, M. Deans, R. C. Elphic, and D. Wettergreen, "Physical Process Models for Improved Rover Mapping," 2016, p. 8.

[16] M. Hansen, R. Elphic, T. Fong, and D. Wettergreen, "Improving Predictive Modeling for Robotic Resource Mapping with Non-Stationary Process Effects and Sensor Noise Estimation," in *2024 International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Nov 2024.

[17] A. Damianou and N. D. Lawrence, "Deep Gaussian Processes," in *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*. PMLR, Apr. 2013, pp. 207–215, iSSN: 1938-7228. [Online]. Available: https://proceedings.mlr.press/v31/damianou13a.html

[18] A. Sauer, R. B. Gramacy, and D. Higdon, "Active Learning for Deep Gaussian Process Surrogates," *Technometrics*, vol. 65, no. 1, pp. 4–18, Jan. 2023, publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00401706.2021.2008505.

[19] Ross Beyer, Mark Allan, Scott McMichael, Audrow Nash, and Jennifer Nguyen, "synthterrain," 2025. [Online]. Available: https://github.com/NeoGeographyToolkit/synthterrain/tree/main

[20] C. J. Tai Udovicic, K. R. Frizzell, G. R. L. Kodikara, M. Kopp, K. M. Luchsinger, A. Madera, M. L. Meier, T. G. Paladino, R. V. Patterson, F. B. Wroblewski, and D. A. Kring, "Buried ice deposits in lunar polar cold traps were disrupted by ballistic sedimentation," *Journal of Geophysical Research: Planets*, vol. 128, no. 5, p. e2022JE007567, 2023, e2022JE007567 2022JE007567. [Online]. Available: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2022JE007567

[21] P. N. Peplowski, R. C. Elphic, E. L. Fritzler, and J. T. Wilson, "Calibration of NASA's Neutron Spectrometer System (NSS) for landed measurements of hydrogen content of the lunar subsurface," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1049, p. 168063, Apr. 2023.

[22] W. Chen, R. Khardon, and L. Liu, "Adaptive Robotic Information Gathering via Non-Stationary Gaussian Processes," Jun. 2023, arXiv:2306.01263 [cs].

[23] L. Schweitzer, H. Jamal, H. Jones, D. Wettergreen, and W. L. Red Whittaker, "Micro Rover Mission for Measuring Lunar Polar Ice," in *2021 IEEE Aerospace Conference (50100)*, Mar. 2021, pp. 1–13, iSSN: 1095-323X. [Online]. Available: https://ieeexplore.ieee.org/document/9438261

[24] H. Bijl, J.-W. Van Wingerden, T. B. Schön, and M. Verhaegen, "Online sparse Gaussian process regression using FITC and PITC approximations**This research is supported by the Dutch Technology Foundation STW, which is part of the Netherlands Organisation for Scientific Research (NWO), and which is partly funded by the Ministry of Economic Affairs. The work was also supported by the Swedish research Council (VR) via the project Probabilistic modeling of dynamical systems (Contract number: 621-2013-5524)." *IFAC-PapersOnLine*, vol. 48, no. 28, pp. 703–708, 2015. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S2405896315028360