

Learning and Autonomy for Extraterrestrial Terrain Sampling: An Experience Report from OWLAT Deployment

Pranay Thangeda^{*}, Yifan Zhu[†], Kris Hauser[‡], and Melkior Ornik[§]
University of Illinois Urbana-Champaign, Urbana, Illinois 61801

Ashish Goel[¶], Erica L Tevere^{||}, Adriana Daca^{**}, and Hari D Nayar^{††}
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109

Erik Kramer^{‡‡}
University of California, Los Angeles, Los Angeles, California 90095

Extraterrestrial autonomous lander missions increasingly demand adaptive capabilities to handle the unpredictable and diverse nature of the terrain. This paper discusses the deployment of a Deep Meta-Learning with Controlled Deployment Gaps (CoDeGa) trained model for terrain scooping tasks in Ocean Worlds Lander Autonomy Testbed (OWLAT) at NASA Jet Propulsion Laboratory. The CoDeGa-powered scooping strategy is designed to adapt to novel terrains, selecting scooping actions based on the available RGB-D image data and limited experience. The paper presents our experiences with transferring the CoDeGa model trained on a low-fidelity testbed to the high-fidelity OWLAT evaluating the feasibility of transferring the model to similar systems. In addition, it also validates the model’s performance in novel, realistic environments, and shares the lessons learned from deploying learning-based autonomy algorithms for space exploration. Experimental results from OWLAT substantiate the efficacy of CoDeGa in rapidly adapting to unfamiliar terrains, effectively making autonomous decisions under considerable domain shifts, thereby endorsing its potential utility in future extraterrestrial missions.

I. Introduction

The exploration of ocean worlds stands as a pivotal element in our solar system exploration, encompassing critical research objectives including the quest for potential signs of life and the comprehensive understanding of conditions fostering habitability [1], [2], [3]. Past lander and rover missions including the Mars exploration program [4] and the Perseverance rover mission [5] are human-in-the-loop systems with expert teams on Earth supervising the terrain sampling process and controlling them based on the collected data. However, unlike Mars missions, many of the ocean world missions are anticipated to have short durations, on the order of tens of days, due to the intensity of the radiation environment, adverse thermal conditions, low availability of solar energy, and using battery as the sole power source for cost and planetary protection reasons [6]. The limited mission duration combined with the long communication delays between Earth and the ocean worlds necessitates a high degree of autonomy for the success of these time-critical missions [7].

Autonomy in terrain sampling missions is challenging due to the high degree of uncertainty in the surface topology at the landing site, terrain material properties, composition, and appearance. Constraints on the number of samples that can be analyzed in-situ, coupled with the risk of system failures, further limits the extent of exploration [8]. Any realistic sampling strategy needs to be able to make decisions under uncertainty and rapidly adapt to the environment to maximize the scientific return from the mission. To navigate these challenges, our prior work [9] proposed Deep Meta-Learning with Controlled Deployment Gaps (CoDeGa), an adaptive scooping strategy that uses deep Gaussian

^{*}Graduate Student, Department of Aerospace Engineering, pranayt2@illinois.edu

[†]Graduate Student, Department of Computer Science, yifan16@illinois.edu

[‡]Professor, Department of Computer Science, kkhouser@illinois.edu

[§]Assistant Professor, Department of Aerospace Engineering, AIAA Member, mornik@illinois.edu

[¶]Robotics Technologist, Robotic Surface Mobility Group, ashish.goel@jpl.nasa.gov

^{||}Robotics Technologist, Robotic Surface Mobility Group, erica.l.tevere@jpl.nasa.gov

^{**}Robotics Technologist, Robotic Surface Mobility Group, adriana.daca@jpl.nasa.gov

^{††}Member of Technical Staff, Robotic Surface Mobility Group, hdnayar@jpl.nasa.gov

^{‡‡}Graduate Student, Department of Mechanical and Aerospace Engineering, ehkramer@ucla.edu

process trained with a novel meta-learning approach. CoDeGa learns online from very limited experience on target terrains despite large domain shifts from the training set. These prior experiments were conducted on a low-fidelity testbed at the University of Illinois Urbana-Champaign (UIUC) designed to rapidly collect large-scale data for training and testing the models on a wide range of terrains.

Building upon this foundational work, the next crucial step is to validate and refine our strategy within a more sophisticated and representative environment. The Ocean Worlds Lander Autonomy Testbed (OWLAT) [10] at NASA Jet Propulsion Laboratory (JPL) provides such an environment. OWLAT is a high-fidelity testbed developed to validate autonomy algorithms for future ocean world missions. It serves as a state-of-the-art platform for simulating various potential future planetary missions over a wide range of dynamic environments, including surface operations on small bodies where recreating the dynamics in low gravity is critical. By integrating CoDeGa with OWLAT, we aim to assess the robustness of our autonomy algorithms in conditions that closely mimic those of actual extraterrestrial landscapes, thereby bridging the gap between preliminary tests and real-world deployment.

In this paper, we report our experiences deploying the CoDeGa-trained adaptive scooping model on the OWLAT testbed. Our contribution is three-fold: (1) we assess the feasibility of transferring the model across systems with similar sensor suites, end-effectors, and primitive actions, (2) we validate the model’s adaptability in novel, realistic environments and (3) we share the lessons learned from designing and deploying learning-based autonomy algorithms in the context of space exploration. Experimental results provide strong evidence that CoDeGa-trained model adapts to the significant domain shifts presented by the OWLAT testbed, reinforcing its applicability and promising role in autonomous terrain sampling for future off-world missions.

II. Preliminaries and Background

This section describes the scooping problem setting in detail and provides a brief overview of the solution approach using the CoDeGa-trained model as proposed in [9]. We also highlight the differences between the UIUC testbed and OWLAT and the importance of deploying and validating the solution approach on the OWLAT testbed.

A. Scooping Problem

We study the problem of scooping in which the goal is to collect high-volume samples from the lander’s workspace with a limited budget of attempts. The problem is formulated as a sequential decision-making problem where the robot observes the terrain RGB-D image $o \in \mathcal{O}$, uses a *scooping policy* to apply action $a \in \mathcal{A}(o)$ where $\mathcal{A}(o)$ is a discrete set of parameterized scooping motions dependent on the RGB-D image, and receives reward $r \in \mathcal{R}$, which is the scooped volume.

Given a target terrain T_* , the robot’s goal is to find a series of scooping motions that maximize the total reward from scoops across the first k attempts. During the n -th attempt, for $n \leq k$, the robot has access to the history of scoops on this terrain $H = \{(o^j, a^j, r^j) \mid j = 1, \dots, n-1\}$. The robot also has access to prior scooping experience, which consists of a set of M terrains $\{T_1, \dots, T_M\}$, and a training dataset $D_i = \{(o^j, a^j, r^j) \mid j = 1, \dots, N_i\}$ of past scoops and their rewards for each terrain $i = 1, \dots, M$.

For a terrain, we suppose a *latent variable* α characterizes its composition, material properties, and topography, which are only indirectly observed. Let α_* characterize T_* and α_i characterize T_i for $i = 1, \dots, M$. Moreover, the observation is dependent on the latent variable and an action’s reward is an unknown function of the action and latent variable. Standard supervised learning applied to model $r \approx f(o, a)$ will work well when α_* is within the distribution of training terrains and α_* is uniquely determined by the observation o or the reward is not strongly related to unobservable latent characteristics. However, when T_* is out of distribution or the observation o leaves ambiguity about latent aspects of the terrain that affect the reward, the performance of the learned model will degrade.

Online learning from H has the potential to help the robot perform better on T_* . We now describe an adaptive online learning approach that can adapt to T_* despite large domain shifts from the training set.

B. Adaptive Scooping using CoDeGa-trained Model

The solution approach proposed in [9], Deep Meta-Learning with Controlled Deployment Gaps (CoDeGa), leverages a deep Gaussian process (GP) model to capture the relationship between the observation-action pair (o, a) and the reward r . The deep Gaussian process model employs deep mean functions and deep kernels where the input to the GP kernel is transformed by a neural network. We note that, in addition to the observation-action pair, the model is also conditioned on the *online support set* that contains the history of previous scoops and their outcomes on the same

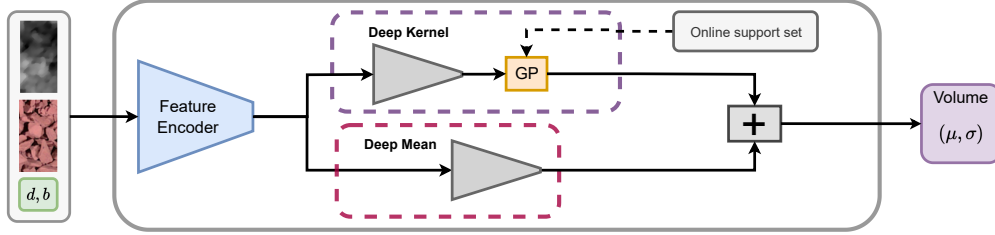


Fig. 1 Overview of the CoDeGa-trained model with deep mean, deep kernel, and feature encoder modules.

terrain. It’s important to differentiate the definition of the term deep GP in [9] from its other usage in the literature where it can also refer to deep belief networks constructed from compositions of GP models [11]. In the context of a scooping task, this model predicts the scooped volume given an input consisting of a local patch of an RGB-D image and action parameters, as illustrated in Figure 1.

In the CoDeGa training procedure, the training terrains are split into mean and kernel training sets containing different materials. Doing so encourages the kernel to encounter residuals representative of those in out-of-distribution tasks. The deep mean is first trained on the mean training set to minimize error, and the GP kernel is then trained on the residuals of the deep mean model applied to the kernel training set. This process is repeated, similar to k -fold cross-validation, with a common kernel trained over aggregated losses across folds. The strength of CoDeGa lies in its ability to generate a model that performs well under deployment gaps, which are common in real-world applications.

Given a model to robustly predict the scooped volume, [9] employs a Bayesian optimization approach for selecting the scooping action. Rather than simply using the mean prediction of the model, it utilizes an acquisition function that also takes uncertainty into account. This function, in essence, serves as a scoring system that guides the selection of actions and encourages the exploration of actions with uncertain outcomes, allowing for a more robust performance under varying conditions.

The CoDeGa-based scooping strategy, while effective, requires a large amount of data across different terrains for training the deep mean and kernel. In the next section, we describe the UIUC testbed designed to collect such data with minimal human supervision and the different terrains and compositions on which the data is collected.

C. UIUC Testbed

The data collection testbed at the University of Illinois at Urbana-Champaign is designed for large-scale data collection and testing of learning-based approaches for scooping tasks. The setup includes a UR5e arm with a scoop mounted on the end-effector, an overhead Intel RealSense L515 RGB-D camera, and a wheeled simulant bin that is approximately 0.9 m x 0.7 m x 0.2 m. A *scoop action* in this setup is a parameterized trajectory for the scoop end-effector tracked by an impedance controller. Figure 2 shows an illustration of the testbed.




Fig. 2 Comparison of the UIUC testbed (left) and the Ocean Worlds Lander Autonomy Testbed (OWLAT) (right).

A *terrain* is defined as a unique composition of one or more materials. The setup uses multiple wheeled simulant bins to parallelly collect data in one bin and set up a new terrain in another that can be quickly swapped. For each new

terrain, data is collected by executing different scooping actions and measuring the volume of scooped material. A detailed description of the different materials considered and the scoop action parameters is provided in [9].

The UIUC testbed enables rapid data collection at the cost of limited representation of the ocean world dynamics. In the next section, we describe the OWLAT testbed designed for validating autonomy algorithms for ocean worlds missions.

D. Ocean Worlds Lander Autonomy Testbed

The Ocean Worlds Lander Autonomy Testbed (OWLAT) [10] is a testbed developed at NASA Jet Propulsion Laboratory (JPL) to test and validate the performance of various autonomy algorithms and architectures for future missions to ocean worlds. As shown in Figure 2, the testbed hardware consists of a 7-DOF Barrett WAM7 robotic arm with a host of interchangeable end-effector tools representing the manipulator, an intel Realsense D415 mounted on a pan-tilt mount for 3D perception, and force-torque sensors located at the interface between the arm and the platform and also at the end of the arm’s wrist. The testbed also has a 6-DOF Stewart platform that is used to simulate the lander and a simulant area that hosts the testbed’s terrains with different simulants and surface features.

OWLAT hardware is complemented by a software interface to command the manipulator and the camera along with its pan-tilt mount for carrying out surface operations and collecting data. The force-torque sensors located at the end of the wrist and at the interface between the arm and the Stewart platform play a critical role in replicating the dynamical environment such landers are likely to experience on the low-gravity icy moons of Jupiter and Saturn. As the tool interacts with the simulant in the testbed, the reaction forces measured are fed into a dynamics model of the system. The computed motion is imposed on the Stewart platform in real-time. The use of high bandwidth Ethercat force-torque sensors allows OWLAT to close the loop at 1 kHz and study test cases demonstrating how interaction with the surface on objects with gravity as low as Enceladus ($g = 0.13 \text{ m/s}^2$) can cause the legs of the lander to lift off the ground, thereby achieving Earth gravity compensation without the use of suspension cables and gantry mechanisms.

In the next section, we describe the process of deploying the CoDeGa-trained model on the OWLAT testbed, highlighting the challenges of interfacing a learned model with new hardware configurations vastly different from the training setting.

III. Deployment

This section details the process of deploying the CoDeGa-trained model on OWLAT testbed for selecting scooping actions based on RGB-D information of the robot’s workspace and past experience in the deployed environment.

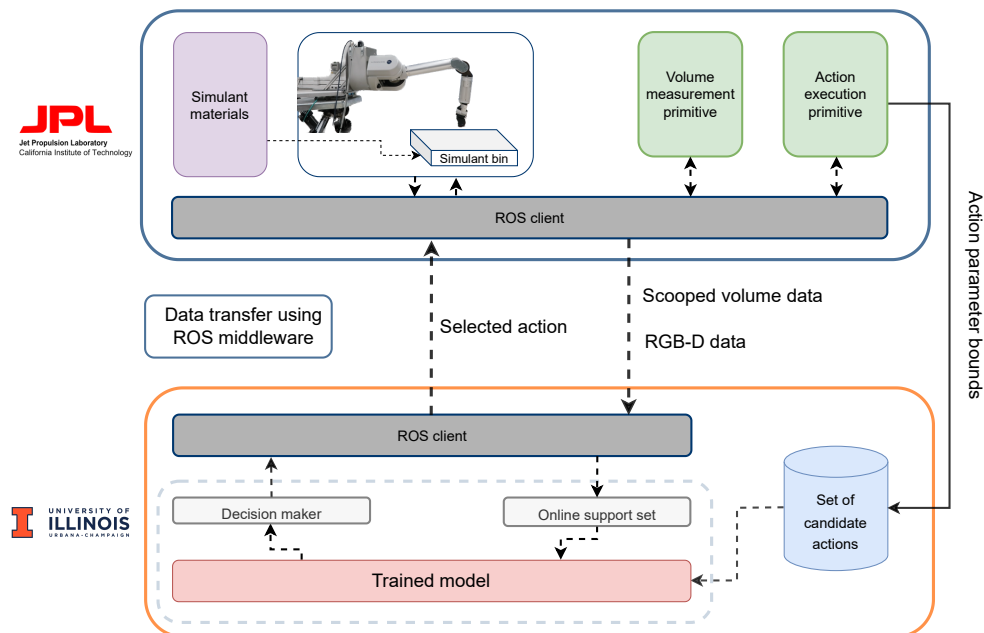


Fig. 3 Overview of remotely deploying the adaptive scooping strategy with UIUC-trained model on OWLAT.

Deploying the adaptive scooping strategy with the CoDeGa-trained model involves taking a sequence of scooping actions where the model adapts based on the reward from the previously executed actions. A single scooping attempt involves (i) collecting RGB-D data capturing the current state of the simulant bin, (ii) passing the data and the set of all candidate scoop actions to the client, (iii) evaluating the candidate actions using the model, taking into account the history of previous actions and their outcomes, (iv) selecting an action from the candidate set based on the model’s evaluations, (v) executing the selected action on the robot, (vi) measuring the scooped volume, and (vii) recording the outcome along with the executed action and corresponding observation in the online support set.

Deploying the scooping strategy, therefore, requires the testbed to provide (i) a simulant bin with testing terrains, (ii) an interface to capture RGB-D data of the simulant bin, (iii) scoop action primitives that are parameterized similar to UIUC testbed, and (iii) a mechanism to measure the scooped volume.

Figure 3 shows an overview of the deployment process with the OWLAT testbed. OWLAT provides a ROS interface with access to the RGB-D camera data and ROS actions over the ROS middleware. We implemented ROS actions that allow us to execute scoop actions parameterized similar to the UIUC testbed. The RGB-D data captured by the camera is passed to the client over the ROS middleware. The deployment client preprocesses the RGB-D data and passes it to the CoDeGa-trained model along with the candidate actions. The decision maker then decides the next action to execute based on the model’s evaluations and sends the action to the testbed over the ROS middleware. The testbed executes the action on the robot and measures the scooped volume. Given that the testing on OWLAT involves far fewer overall scooping attempts when compared to the data collection process in the UIUC testbed, we manually measure the scooped volume. We note that a deployment-ready ocean world mission is expected to include the instrumentation to analyze the scientific value of the scooped material, thereby addressing the issue of scooped volume measurement.

We now discuss in detail the steps involved in processing RGB-D data to be compatible with the CoDeGa-trained model and generating the candidate action set that can be easily implemented on any robotic system with a scoop end-effector.

A. Processing RGB-D Data

The CoDeGa-trained model utilizes a CNN-based feature encoder to extract relevant information from the RGB-D data. While effective, these encoders are susceptible to variations in input data, requiring a degree of consistency in feature scale and camera orientation relative to the training set for optimal performance. The UIUC testbed standardizes data capture across materials with varying feature scales by using a static overhead RGB-D camera, thereby streamlining data collection and preprocessing by minimizing environmental inconsistencies.

In contrast, real-world systems rarely offer such controlled conditions. For example, in the OWLAT testbed, the RGB-D camera is mounted on a pan-tilt unit at the base of the robot, capturing data from this unique perspective, as shown in Figure 2. Additionally, real-world data is often further complicated by variations in camera resolution and quality across systems, and it commonly includes noise, occlusions, and other artifacts.




Fig. 4 Overview of the steps involved in processing the RGB-D data captured by the OWLAT testbed.

To align the OWLAT data with the training conditions of the CoDeGa model, we developed a preprocessing pipeline (shown in Figure 4). This pipeline includes (i) reprojecting the point cloud to a top-down view to emulate the UIUC testbed perspective, (ii) inpainting to fill in gaps in the RGB and depth data while accounting for occluded and out-of-range regions, and (iii) using system state and constraints to filter out anomalies. The resulting images are then used to produce localized image patches based on the action parameters.

In the next section, we discuss how our choice of parameterized motion primitives as actions enables us to easily deploy the CoDeGa-trained model on any robotic system with a scoop end-effector. Additionally, we detail our approach of generating valid candidate actions while accounting for system constraints.

B. Generating Action Candidates

In our implementation, the scoop action is a motion primitive characterized by an initial x, y position, the yaw angle θ of the scoop, the scooping depth d , and the stiffness parameter b . The z -coordinate is deduced from the depth map at the outset of the action. This action primitive follows the parameterized trajectory described in [9], which is executed via an impedance controller. By defining the action in the end effector space, we maintain a consistent action interpretation across varied systems. Collaboration with the OWLAT team has allowed for the adoption of the same trajectory for the scoop action primitive as established in [9] for deployment. Moreover, empirical evidence from [9] indicates that higher stiffness values yield superior performance on diverse terrains; consequently, we fix the stiffness parameter b to 'high' value within the OWLAT action primitive to leverage these observed benefits.




Fig. 5 Illustration of the candidate action set generation process using OWLAT testbed data.

To synthesize the set of candidate actions, we use the processed depth data and system constraints to delineate regions in the simulant bin amenable to scooping. Within these identified regions, we establish a uniform grid to represent potential action locations. At each grid point, we evaluate eight possible yaw angles—excluding any that are deemed infeasible—coupled with four distinct scooping depths (d values of 0.2 cm, 0.4 cm, 0.6 cm, and 0.8 cm) and a single 'high' stiffness value (b). We note that the size of the resultant action set is dynamic, as it adjusts with the changes in the environment. For each candidate action, we generate corresponding RGB and depth image patches centered at the action x, y and aligned with the scoop orientation θ , as illustrated in Figure 5. The image patches along with the stiffness parameter and the selected depth value are then used to evaluate the action candidate, as illustrated in Figure 1.

The next section describes the experiments conducted on the OWLAT testbed to evaluate the performance of the CoDeGa-trained model.

IV. Experimental Evaluation and Discussion

We perform the experiments using an active remote connection with the OWLAT testbed operating at the Jet Propulsion Laboratory (JPL), Pasadena, CA, and the CoDeGa-trained model deployed for inference at the University of Illinois at Urbana-Champaign (UIUC), Champaign, Illinois. Table 1 lists all the materials used for training the model along with their corresponding physical properties. For evaluation, we use a terrain designed by the OWLAT team with two out-of-distribution materials, Comet and Regolith. Comet is an unscoopable composition of grey comet simulant material [12] surrounded by 3D printed PLA features with rugged terrain features from a 3D scan of devil’s golf course in Death Valley National Park, California painted to match the Regolith’s color. Regolith is a fine sand-like material that is visually distinct from the sand used in training. The two materials are composed together to create a hypothetical representation of the ocean world terrain, as depicted in Figure 4.

Following the experimental procedure in [9], we compare the proposed method to the Non-adaptive baseline (Non-Adaptive), i.e. only the deep mean component of the CoDeGa-trained model, and a volume maximizing (Vol-Max) policy, where the action is chosen to maximize the intersection between the scoop’s swept volume and the terrain following a strategy proposed in the excavation literature [13].

We conduct three iterations of the experiment using each method across three different scenarios, each featuring varying terrain topologies. Scenarios 1 and 2 consist of flat Regolith regions emphasizing the unscoopable Comet areas, with minor terrain features in the Regolith region that occur naturally during resets. In Scenario 3, the scoopable Regolith region is designed with terrain features that have heights comparable to those of the unscoopable Comet regions. The depth map in Figure 6 illustrates the terrain topology for Scenario 3, where the three circular green areas represent mounds created in the regolith material.

The objective is to maximize the total volume of material scooped in $k = 5$ attempts. Given that only one of the two materials in the testing terrain is scoopable, we measure mass in each attempt rather than volume to enhance

Table 1 Materials used in experiments, categorized by training and testing sets, with U.S. quarter coin for scale.





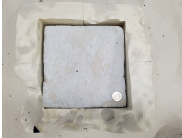





| Training | | | | Testing |
|---|--|---|--|--|
| Sand fine play sand, $\ll 1$ mm  | Pebbles rocks, 0.8 – 1.0 cm  | Slate flat rocks, 2.0–4.0 cm  | Gravel rocks, 1.5–3.0 cm  | Comet comet simulant, unscoopable  |
| Paper Balls crumpled paper, 4.0 – 6.0 cm  | Corn dry corn kernels, 0.3–0.7 cm  | Shredded Cardboard cardboard, 1.0 – 8.0 cm  | Mulch red wood landscape mulch  | Regolith fine sand, 0.1 mm – 0.5 mm  |

Table 2 Results comparing CoDeGa with two baselines in OWLAT testbed. Higher scooped mass is better.

| | Vol-Max (gm) | Non-Adaptive (gm) | CoDeGa (gm) |
|-------------------|--------------|-------------------|-------------|
| Scenario 1 | 0.0 | 3.5 | 52.2 |
| Scenario 2 | 0.0 | 18.8 | 64.2 |
| Scenario 3 | 5.6 | 43.6 | 75.4 |
| Overall | 1.9 | 22.0 | 63.9 |

measurement accuracy and ease. This change is justified given that mass and volume are equivalent due to the direct proportionality between the scooped mass and the scooped volume in the context of our testing environment. We use the average scooped mass over a run as the metric for comparing the efficacy of different methods. Within a run, should robot trajectory planning for a scooping action fail, the subsequent highest scoring action is selected until planning succeeds. We use the candidate action set outlined in Section III.B.




Fig. 6 Example scooping attempts on the OWLAT testbed using three different methods. The cyan-to-yellow colors represent regions with higher terrain elevations.

The average scooped mass for all methods across the three different scenarios is reported in Table 2. The CoDeGa-trained model-based approach significantly outperforms the baseline methods in all scenarios. Figure 6 depicts the scooping actions executed during a representative trial for each method. The Vol-Max method consistently selects scoop locations near the Comet regions due to its inherent preference for areas with steep terrain gradients. While the Non-Adaptive baseline initially targets the Regolith mounds, it fails to modify its policy in response to the data observed online, eventually resorting to ineffective scooping attempts in the Comet region, akin to Vol-Max. On the other hand,

our approach initially engages the Comet region but rapidly adapts its strategy in response to the low mass of scooped material, shifting focus to the 'Regolith' mounds and thereby maximizing the total scooped mass.

V. Conclusion

This paper presents the deployment of a CoDeGa-trained adaptive scooping model on the high-fidelity OWLAT testbed. The model, originally developed on a low-fidelity UIUC testbed, was evaluated on novel terrains in OWLAT testbed with out-of-distribution materials. Experimental results demonstrate the model's capability to rapidly adapt to unfamiliar environments and make effective decisions despite significant domain shifts. Specifically, the model achieved 2-3x higher average scooped mass compared to non-adaptive baselines across varying terrain topologies by adapting its strategy online based on limited experience.

The successful deployment substantiates the potential utility of learning-based autonomy for maximizing scientific return under uncertainty in ocean world missions. Moreover, it provides valuable insights into real-world integration challenges, including data preprocessing and motion planning. Overall, this work endorses the feasibility of transferring learning-based systems from idealized training settings to realistic deployment environments.

Acknowledgments

This work was supported by NASA Grant 80NSSC21K1030. Part of this work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

References

- [1] Hendrix, A. R., Hurford, T. A., Barge, L. M., Bland, M. T., Bowman, J. S., Brinckerhoff, W., Buratti, B. J., Cable, M. L., Castillo-Rogez, J., Collins, G. C., et al., "The NASA roadmap to ocean worlds," *Astrobiology*, Vol. 19, No. 1, 2019, pp. 1–27.
- [2] Howell, S., Stone, W. C., Craft, K., German, C., Murray, A., Rhoden, A., and Arrigo, K., "Ocean worlds exploration and the search for life," *Bulletin of the American Astronomical Society*, Vol. 53, No. 4, 2021.
- [3] Sherwood, B., Lunine, J., Sotin, C., Cwik, T., and Naderi, F., "Program options to explore ocean worlds," *Acta Astronautica*, Vol. 143, 2018, pp. 285–296.
- [4] Shirley, D., and McCleese, D., "Mars exploration program strategy-1995-2020," *Aerospace Sciences Meeting and Exhibit*, 1996.
- [5] Ehrenfried, M., "Surface operations and science," *Perseverance and the Mars 2020 Mission*, Springer, 2022, pp. 91–110.
- [6] Hand, K., Murray, A., Garvin, J., Brinckerhoff, W., Christner, B., Edgett, K., and Hoehler, T., "Report of the Europa lander science definition team," Tech. Rep. JPL D-97667, NASA, 2017.
- [7] Reeves, G., Kennedy, B. A., Tan-Wang, G. H., Backes, P. G., Chien, S. A., Verma, V., Hand, K. P., and Phillips, C. B., "Development of autonomous actions to enable the next decade of ocean world exploration," *Bulletin of the American Astronomical Society*, Vol. 53, No. 4, 2021.
- [8] Thangeda, P., and Ornik, M., "Adaptive sampling site selection for robotic exploration in unknown environments," *International Conference on Intelligent Robots and Systems*, 2022, pp. 4120–4125.
- [9] Zhu, Y., Thangeda, P., Ornik, M., and Hauser, K., "Few-shot adaptation for manipulating granular materials under domain shift," *Robotics: Science and Systems*, 2023.
- [10] Nayar, H., Goel, A., Boettcher, A., Hans, M., Sawoniewicz, J., Gaut, A., Higa, S., Ono, H., Jain, A., Lim, C., Nesnas, I., Ma, R., and Thomsen, M., "Development of a lander autonomy testbed for ocean worlds missions," *Earth and Space*, 2021, pp. 531–540.
- [11] Damianou, A., and Lawrence, N. D., "Deep Gaussian processes," *International Conference on Artificial Intelligence and Statistics*, Vol. 31, 2013, pp. 207–215.
- [12] Carey, E. M., Peters, G. H., Choukroun, M., Chu, L., Carpenter, E., Cohen, B., Panossian, L., Zhou, Y. M., Sarkissian, A., Moreland, S., et al., "Development and characteristics of mechanical porous ambient comet simulants as comet surface analogs," *Planetary and Space Science*, Vol. 147, 2017, pp. 6–13.
- [13] Yang, Y., Long, P., Song, X., Pan, J., and Zhang, L., "Optimization-based framework for excavation trajectory generation," *Robotics and Automation Letters*, Vol. 6, No. 2, 2021, pp. 1479–1486.