

LEARNING TOOL-AWARE ADAPTIVE COMPLIANT CONTROL FOR AUTONOMOUS REGOLITH EXCAVATION

Andrej Orsula¹, Matthieu Geist², Miguel Olivares-Mendez¹, and Carol Martinez¹

¹University of Luxembourg

²Earth Species Project

ABSTRACT

Autonomous regolith excavation is a cornerstone of in-situ resource utilization for a sustained human presence beyond Earth. However, this task is fundamentally hindered by the complex interaction dynamics of granular media and the operational need for robots to use diverse tools. To address these challenges, this work introduces a framework where a model-based reinforcement learning agent learns within a parallelized simulation. This environment leverages high-fidelity particle physics and procedural generation to create a vast distribution of both lunar terrains and excavation tool geometries. To master this diversity, the agent learns an adaptive interaction strategy by dynamically modulating its own stiffness and damping at each control step through operational space control. Our experiments demonstrate that training with a procedural distribution of tools is critical for generalization and enables the development of sophisticated tool-aware behavior. Furthermore, we show that augmenting the agent with visual feedback significantly improves task success. These results represent a validated methodology for developing the robust and versatile autonomous systems required for the foundational tasks of future space missions. *The source code is available at github.com/AndrejOrsula/space_robotics_bench.*

Key words: Reinforcement Learning, Compliant Control, Resource Extraction.

1. INTRODUCTION

The success of ambitious international efforts like the Artemis program depends on establishing a sustained human and robotic presence on the Moon [1]. In-Situ Resource Utilization (ISRU) is a foundational principle for this endeavor [2]. The ability to harness local materials is a critical enabler for long-term missions because it significantly reduces reliance on complex and costly Earth-based supply chains. At the forefront of ISRU is the autonomous excavation of lunar regolith. This single capability underpins a wide range of future operations, including the large-scale construction of infrastructure and the extraction of vital resources like water, ice, and oxygen [3]. The successful execution of these demanding




Figure 1: A generalizable excavation skill is learned within the Space Robotics Bench by training agents on a vast distribution of randomized scenarios with procedurally generated tools, which fosters an adaptive compliance policy for successful tool-aware interaction.

tasks over multi-year missions will depend on the development of robotic systems that can operate reliably and adaptively with minimal human supervision.

Developing an autonomous excavation system presents a formidable set of interconnected challenges. The first major obstacle is the complex and unpredictable physics of lunar environments covered with regolith, which is a fine-grained granular material whose mechanical behavior defies simple analytical models [4, 5]. The properties of this material are known to be highly variable across the lunar surface, yet our current understanding is derived from a small number of past missions [6]. Future operations will target unexplored regions where key characteristics like particle size, compaction, and the presence of ice are largely unknown. This fundamental uncertainty renders it impossible to create a single, accurate, mission-specific model, as a robotic manipulator digging into this medium will encounter highly variable and often large resistance forces [7]. Traditional rigid control strategies are inherently brittle under such unpredictable conditions and can easily lead to mission-critical failures. Such failures include the manipulator stalling, drawing excessive power, or sustaining catastrophic damage to its hardware. A robust system must therefore be able to manage these unmodeled physical interactions through compliant behavior rather than rigid control [8].

The need for operational versatility compounds this physical challenge. A single robotic platform will be expected to perform a variety of excavation tasks over its operational lifetime. These diverse tasks will naturally require different tools, such as wide scoops for bulk material transfer or specialized toothed buckets for breaking compacted ground. This need for adaptation extends beyond intentional tool changes to include unintentional ones as well. The abrasive nature of regolith can cause significant wear and tear over a long-duration mission [9], leading to unforeseen changes in tool geometry, such as blunted edges, broken teeth, or even partial loss of the tool. A practical autonomous system cannot be pre-programmed for a fixed set of tool configurations. It must possess a general excavation skill that allows it to intelligently adapt its strategy to its changing physical form, whether from a deliberate tool swap or gradual degradation. The harsh lunar environment further constrains these operations. The goal of maximizing collected regolith must be carefully balanced against the critical need to minimize abrasive dust generation, which poses a significant and persistent threat to nearby actuators and sensors.

This work introduces a learning-based framework where an autonomous agent acquires generalizable excavation skills to address these challenges. Our approach, conceptually illustrated in Figure 1, centers on training the agent in a high-fidelity simulation featuring a rich distribution of procedurally generated scenarios. The agent learns to master complex physical interactions not through rigid commands but by dynamically modulating its own physical compliance. This method enables the development of sophisticated, tool-aware strategies that are robust and efficient. This paper makes several key contributions. We first present a comprehensive simulation framework for learning regolith excavation within the Space Robotics Bench (SRB) [12]. This environment integrates procedural generation of both lunar terrains and diverse excavation tools inside a high-fidelity particle physics simulation. We then demonstrate that training an agent with a procedural distribution of tools is critical for achieving a robust and generalizable policy. We also provide an analysis of the emergent behaviors that result from learned adaptive compliance. Finally, this investigation explores how augmenting the agent with visual feedback from depth perception enhances its ability to coordinate the complex interactions between tools and regolith.


2. RELATED WORK

The field of terramechanics traditionally addresses the challenge of robotic interaction with granular media [10]. It provides analytical models that can predict interaction forces and offer valuable insights for tasks like rover traversal. However, these models are often insufficient for the high-dimensional dynamics of manipulation [11]. The reaction forces during an excavation task depend on specific tool geometry, the precise trajectory through the medium, and highly variable local regolith properties. The behavior of granular material is highly non-linear, and its properties are difficult to model, especially when

combined with the uncertain conditions of space. Creating accurate analytical models for such intricate interactions is therefore exceptionally difficult. This difficulty motivates our data-driven approach, where a control policy can implicitly learn the nuances of excavation through direct experience.

Due to the limited availability of data, the development of such data-driven policies for space robotics is only feasible through high-fidelity simulation. This work is a direct application of the principles behind SRB, our simulation framework designed to facilitate robot learning for diverse space applications [12]. A core philosophy in modern robotics is that policy robustness stems from a massive diversity of experience rather than a single perfect digital twin. Foundational techniques like domain randomization have proven effective in bridging the sim-to-real gap by varying simulation parameters to enforce robustness to uncertainty [13]. We extend this principle with procedural content generation (PCG), which allows us to create a nearly infinite distribution of unique training scenarios. While the majority of prior research has applied PCG to simple 2D environments [14, 15], we have demonstrated the effectiveness of this approach for contact-rich assembly [16] and sim-to-real transfer of a navigation policy [17]. The key distinction in this work is that we apply PCG directly to the robot’s own physical embodiment. We procedurally generate a wide variety of excavation tool geometries, rather than merely creating diverse environments. This encourages the agent to learn a generalizable skill not tied to a specific end-effector.

Reinforcement Learning (RL) has emerged as a powerful paradigm for discovering sophisticated strategies in problems where optimal behaviors are non-intuitive [18]. Its potential for autonomous excavation has been demonstrated in terrestrial applications, such as controlling wheel loaders to perform soil scooping [19] and planning excavation trajectories for rigid objects [20]. The complexity of interaction dynamics has also been explored in related domains such as water scooping [21], where curriculum learning was employed to accelerate the learning process. Our methodology is distinct from prior work in two critical ways. First, we employ a model-based RL algorithm. Methods like DreamerV3 learn a predictive world model that can significantly improve sample efficiency, allowing an agent to plan more effectively in environments with complex physical dynamics [22]. Second, we focus on the control representation that defines how the agent physically executes its decisions. Rigid kinematic control is often brittle in contact-rich scenarios [8]. To address this, Operational Space Control (OSC) provides a principled framework for managing interaction forces through direct control of end-effector compliance [23]. Beyond using manually tuned gains, OSC can be extended with learning-based approaches that dynamically adjust stiffness and damping based on sensory feedback [24]. This adaptability is key to enhancing the performance of autonomous systems in dynamic tasks. While prior research has combined OSC with model-free RL for terrestrial tasks [25], its application to space robotics remains unexplored. The central innovation of



(a) Parametric pipelines generate diverse lunar terrains and excavation tools, creating a near-infinite source of unique training assets. (b) A visualization of 16 parallel training environments with procedural diversity that creates a massively diverse learning experience.

Figure 2: Our procedural paradigm, with the SimForge engine populating many parallel environments in the Space Robotics Bench to facilitate the rich and varied training experience that is essential for learning truly generalizable skills.

our work is the fusion of these two concepts. We leverage a model-based RL agent to dynamically learn and modulate the stiffness and damping parameters of its own OSC controller. This creates a tight synergy between high-level learning and low-level physical control, enabling the agent to discover adaptive compliance strategies tailored to the unique challenges of robotic excavation in space.

3. METHODOLOGY

Our approach is founded on an integrated framework that combines a scalable and diverse simulation testbed with a sophisticated learning methodology. The testbed is designed to generate the rich experience necessary for training a robust policy, while the learning methodology then enables an agent to translate that experience into a generalizable and physically adaptive excavation skill.

3.1. Procedural Testbed for Regolith Excavation

All agent training and evaluation occur within the Space Robotics Bench [12]. This simulation framework for robot learning in space is built on top of NVIDIA Isaac Lab [26], which provides a scalable foundation through its massively parallel GPU-accelerated physics and rendering. This high-throughput architecture is essential for our procedural workflow to facilitate the generation of diverse experiences at scale, which is crucial for generating the vast and varied datasets required by modern RL algorithms. To accurately capture the complex physics of excavation, we model the regolith as a granular medium using high-fidelity particle physics that is based on extended position-based dynamics (XPBD) [27]. This physics-based approach is crucial for representing the complex, non-linear dynamics of the material. It also provides the realistic and highly variable force feedback that the manipulator experiences during interaction, which is a necessary component for learning a physically grounded skill.

The foundation of our approach for developing a robust policy is training with diverse data. We achieve this

through two synergistic mechanisms that are depicted in Figure 2. First, we generate unique structural assets for each parallel training instance through our integrated procedural engine SimForge [28]. The procedural pipelines for both lunar terrains and excavation tools are hand-crafted using the modular Geometry Nodes system within Blender. This node-based system allows for the creation of complex parametric rules that deterministically define an entire family of 3D models. This geometric generation is combined with procedural materials based on Shader Nodes. These materials are then baked into a set of physically-based rendering (PBR) textures before the final asset is spawned within the simulation. To produce a near-infinite variety of unique assets, the engine automatically samples the pipeline parameters from a seeded random distribution for each exported variant. The resulting assets are showcased in Figure 2a.

The procedural generation of the excavation tools themselves is a central innovation of our methodology. The tool geometry pipeline features over forty parameters that control a wide array of morphological features. These parameters modulate the fundamental size and shape of the scoop, including its width, depth, and curvature. They also control features like mounting interface alignment, the number and shape of teeth, and even the overall deterioration of the tool to simulate wear from long-term use. As showcased in Figure 3, this level of control enables us to generate a continuous spectrum of end-effectors. These tools range from wide scoops for bulk material transfer to narrow buckets with sharp serrated edges for breaking compacted regolith with elevated cohesion. By training on this procedural distribution of tools, the agent is forced to learn an abstract and generalizable understanding of excavation. It cannot rely on a single memorized trajectory. It must instead develop an adaptive strategy that can be generalized to the specific geometry of the tool it currently holds.

A complementary layer of diversity is added through extensive domain randomization of the simulation’s physical and visual parameters. Physical randomization targets the underlying dynamics to build robustness against un-

modeled physical effects and environmental uncertainty. This includes randomizing factors such as the density, friction, and cohesion of individual regolith particles. The magnitude of the gravity is also sampled from a distribution and varied at random intervals to simulate operation in different lunar regions. Visual randomization is applied to the ambient illumination to match the harsh lighting of the lunar day. The true power of our methodology is realized through its scalability. Each of the parallel environments in Figure 2b is not just structurally unique due to PCG, but it is also parametrically unique in every episode due to domain randomization. This comprehensive combination of procedural and parametric diversity creates the rich, high-variance distribution of training scenarios required to forge a truly generalizable policy.




Figure 3: A grid of procedurally generated excavation tool geometries. Each row isolates the effect of modulating selected parameters. This demonstrates the fine-grained control over the resulting tool morphology.

3.2. Learning Tool-Aware Adaptive Control

We formulate the excavation task as a Partially Observable Markov Decision Process (POMDP) because the agent cannot access the complete state of the world. Key physical properties, such as the shape of the terrain and the position of every regolith particle, are not directly observable. Similarly, we do not provide the agent with the intrinsic properties of local regolith, such as friction and cohesion, due to the difficulty of measuring these properties in situ. The unobservability is compounded by the fact that the agent does not have access to the geometry of its equipped tool. This creates a challenging setting where the agent must infer hidden states from its history of physical interactions.

We design the task as a finite-horizon problem. A single episode requires the agent to learn a sequence of actions

that successfully lifts and stabilizes a significant volume of regolith. Each episode is truncated after a maximum duration of 15 seconds, and we use no other state-based termination conditions. This fixed-duration formulation is a deliberate design choice. It ensures the learned policy represents a modular, reusable skill. The resulting policy can then be integrated as a low-level primitive within a larger hierarchical control system, which could execute long-horizon missions by repeatedly invoking this foundational skill.

The objective is to learn a policy that maps a sequence of observations to actions that maximize the expected cumulative reward. We employ the DreamerV3 [22] model-based RL algorithm to learn the control policy. Its ability to learn a predictive world model is exceptionally well-suited for mastering the complex dynamics of granular media [17]. Its recurrent world model architecture serves as a memory that allows the agent to integrate information over time and infer the hidden states of the environment. Furthermore, its high sample efficiency is crucial for training with the computationally intensive simulation of granular media.

The agent’s observation space is based on proprioception, which includes the manipulator’s normalized joint positions and joint torque measurements. The end-effector pose from Forward Kinematics is also provided, with orientation encoded using a 6D representation [29]. To investigate the impact of augmenting the agent with sensory feedback, we also evaluate a variant that incorporates 128×128 px depth maps of its immediate workspace from two perspectives. We deliberately omit any privileged information regarding the tool geometry. This decision reflects a more realistic scenario and prevents the agent from overfitting to our specific procedural pipeline. Instead, the agent must infer the tool’s geometry through physical interactions, which is a more challenging but ultimately more generalizable approach.

The actions of the agent directly integrate with OSC, as illustrated in Figure 4. The core concept involves using OSC to provide software-defined compliance through adjustable stiffness and damping gains, which can be dynamically modulated by the agent at every time step based on task requirements. In our experiments, we investigate three distinct control strategies that primarily differ in how compliance is handled and learned.

IK [$\dim(\mathcal{A}) = 6$]: A non-compliant baseline using standard differential Inverse Kinematics.

OSC-PART [$\dim(\mathcal{A}) = 12$]: OSC agent that learns to modulate only the stiffness (K_p) gains.

ADAPTIVE [$\dim(\mathcal{A}) = 18$]: OSC agent that learns to modulate both stiffness (K_p) and damping (K_d) gains.

Across all strategies, a core component of the action space $\mathcal{A}$ includes the desired end-effector translational and rotational displacements in $SE(3)$. For the adaptive strategies, the action space is augmented accordingly, where OSC-PART adds six dimensions for the stiffness gains K_p , while ADAPTIVE adds another six dimensions for the damping gains K_d . This allows the agent to ac-




Figure 4: The action space of our agent combines SE(3) motion commands with stiffness and damping gains to achieve learned adaptive compliance through OSC.

tively control the robot’s impedance based on its learned policy. For instance, the agent can choose to make the end-effector rigid when moving through free space and then switch to a more compliant mode when interacting with the regolith. This dynamic modulation of compliance is a key innovation that enables the agent to adapt its physical interaction style to the specific tool geometry and the current state of the excavation task.

We employ a composite reward function to guide the agent towards the desired excavation strategy. The total reward provided at each time step is defined as a sum of several components.

$$R_{\text{total}} = R_{\text{approach}} + R_{\text{lift}} + R_{\text{stabilize}} - P_{\text{create_dust}} - P_{\text{jerk}}$$

The primary objective is captured by the two positive terms R_{lift} and $R_{\text{stabilize}}$. R_{lift} provides a reward proportional to the number of regolith particles successfully lifted to a target height. $R_{\text{stabilize}}$ then provides a substantial bonus for achieving a stable lift, specifically rewarding particles that are within the target height range and have near-zero velocity. This structure heavily incentivizes a smooth, controlled final state where the excavated material is held stationary. Early-stage learning is encouraged by the R_{approach} term, which facilitates approaching the regolith pile and maintaining a correct top-down scoop orientation. These objectives are balanced by the penalty term $P_{\text{create_dust}}$, which penalizes dust creation proportional to regolith particle velocity. Lastly, the term P_{jerk} encourages smoother motion by penalizing action rate and excessive joint torques. This comprehensive reward structure provides a dense learning signal that guides the agent toward mastering the nuanced skill of autonomous excavation. The particle-related reward components are computationally expensive, as the data transfer alone for reading the positions and velocities of particles would require a minimum bandwidth of 1.2 GB/s for every 1M particles. Therefore, these terms are computed only once every 2 seconds and cached for intermediate steps. This design improves simulation performance but introduces reward delays, which makes the notion of credit assignment more challenging.

4. EXPERIMENTS AND RESULTS

Our empirical validation systematically evaluates the core principles of our methodology. We conducted targeted experiments within SRB to deconstruct the complex problem of autonomous excavation. These experiments first establish the foundational role of adaptive

compliance, then demonstrate the critical importance of procedural diversity for generalization, and finally quantify the benefits of visual perception.

4.1. Experimental Setup

Our investigation is structured as a direct comparison of distinct agents that differed in their training experience, compliance strategy, and available sensory observations. All experiments use a simulated Franka Emika robotic manipulator mounted on a stationary pedestal. The simulation is executed with a physics update rate of 250 Hz, while the learning agent operates at a control frequency of 50 Hz. To balance physical fidelity with computational cost, we use 16 parallel environment instances, each populated with 100k regolith particles. For each instance, a particle system is spawned in front of the robot in a pyramid formation above a unique procedural terrain. The initial velocity of each particle is sampled from a uniform distribution $U(-0.5, 0.5)$ m/s for the XY axes and $U(-0.5, 0.0)$ m/s for the Z axis. The particles are then allowed to settle under gravity for up to 300 s or until their median velocity drops below a stable threshold of 0.01 m/s. This settled state is then cached as the starting condition for all subsequent training episodes. During training, the physical properties of the regolith and the environment are varied through domain randomization. Each particle has a radius of 0.004 m, a density sampled from $\mathcal{N}(1500, 200)$ kg/m³, a friction coefficient from $\mathcal{N}(0.9, 0.1)$, and a cohesion value from $\mathcal{N}(0.1, 0.01)$. The global gravity magnitude is also sampled from $U(1.6123, 1.6376)$ m/s² to simulate operations in different lunar regions. The robot begins each episode from a randomized joint configuration with its rigidly attached excavation tool positioned above the regolith pile.

All agents are trained using the DreamerV3 algorithm, which distinguishes between a large world model used during training and a more compact inference policy. The proprioceptive agents utilize an architecture with 86M learnable parameters, with over 90% dedicated to the training-time world model, and the remaining 8M are shared during inference by the encoder and the policy. The agent variant with visual observations incorporates additional convolutional layers to process depth map inputs, which results in a larger architecture with 130M parameters (10M during inference). To achieve satisfactory sample efficiency with the computationally intensive particle simulation, we employ a large replay buffer with a capacity of 2.5M transitions combined with a high training ratio of 256. Other key hyperparameters are kept consistent with the original implementation to ensure a robust baseline [22]. The proprioceptive agents were trained for a total of 5M environment steps with a mean wall-clock time of 71 h on a single NVIDIA RTX 4090 GPU. Due to the increased computational load of processing high-dimensional image data, the VISUAL agent took 218 h to complete training. The resulting inference latency for the final policies was 1.56 ± 0.1 ms for the proprioceptive agents and 2.08 ± 0.2 ms for the VISUAL agent. While we acknowledge that the high-performance

Table 1: Zero-shot generalization performance of all evaluated agents on the held-out set of 8 novel tool geometries.

Agent Configuration			Excavated Volume [unit: L]	Dust Generated (normalized)	Motion Jerk (normalized)
IK	IK	+ PCG Tools	0.13 ± 0.03	1.00 ± 0.03	1.00 ± 0.68
OSC-PART	OSC (K_p)	+ PCG Tools	0.09 ± 0.06	0.97 ± 0.04	0.41 ± 0.27
SPECIALIST	OSC (K_p & K_d)	+ Static Tool	0.02 ± 0.04	0.91 ± 0.03	0.39 ± 0.18
ADAPTIVE	OSC (K_p & K_d)	+ PCG Tools	0.14 ± 0.07	0.94 ± 0.04	0.44 ± 0.17
VISUAL	OSC (K_p & K_d)	+ PCG Tools + Vision	0.27 ± 0.12	0.89 ± 0.09	0.20 ± 0.15

hardware acceleration is not yet space-qualified, this limitation is being addressed by rapid advancements in on-board computing [30]. Additionally, the policy could be further optimized using established techniques like quantization and pruning.

The final evaluation of all trained agents is conducted on a held-out test set of 8 novel tool geometries that were not seen during training. As shown in Figure 7, this set includes 4 procedural variants and 4 manual designs to ensure a comprehensive assessment of generalization. We measure performance using three key metrics. The first is the mean volume of successfully excavated regolith, quantified as the normalized count of lifted and stabilized particles at the end of an episode. The second metric is the amount of generated dust, computed as the cumulative fraction of particles exceeding a velocity threshold of 0.02 m/s throughout the episode. The third metric is control smoothness, measured by the time-averaged mean squared jerk of the robot’s joint motion. Each reported result is an average over 50 independent evaluation episodes for each of the 8 distinct tool geometries.

4.2. Results and Discussion

Our experiments were designed to systematically deconstruct the challenge of autonomous excavation by isolating the key factors that contribute to a robust policy. First, we establish the role of adaptive compliance for managing complex physical interactions. We then demonstrate the critical role of procedural diversity in training for generalization. Finally, we quantify the performance benefits of augmenting the agent with visual perception and analyze the implications of our results for robotic tool design. The learning curves of all evaluated agents are presented in Figure 5, and their final performance on the held-out set of novel tools is summarized in Table 1.

Impact of Adaptive Compliance To manage the unpredictable forces of excavation, a robot must interact compliantly with its environment. Our first experiment confirms this necessity by comparing a rigid kinematic controller with agents that could learn to modulate their own compliance. As shown in Table 1, all three agents ultimately learned to excavate a similar volume of regolith. However, the critical difference is in the quality and safety of the interaction. While successful in excavating material, the rigid IK agent achieved its goal through aggressive and uncontrolled interactions. This resulted in significantly jerkier motions and the generation of more




Figure 5: Learning curves of evaluated agents during training, averaged over three seeds (two for VISUAL), with shaded regions representing the standard deviation.

dust. In contrast, the ADAPTIVE agent, which learned to control both stiffness and damping, demonstrated a markedly smoother and more controlled interaction style. It produced the least amount of dust, which indicates a safer and more stable approach. The OSC-PART agent offered slightly smoother motion but was outperformed in the amount of excavated regolith. This result demonstrates that while rigid control can succeed, learned adaptive compliance is essential for performing the task safely and efficiently, which is a prerequisite for long-term hardware reliability in space.

Generalization through Procedural Diversity Having established the need for compliance, our central experiment investigated how to create a truly generalizable excavation skill. We compared a SPECIALIST agent trained with a single, static tool geometry against a generalist (ADAPTIVE) agent trained on the full procedural distribution of tools. The learning curves in Figure 5 show that the SPECIALIST learned more quickly since it only needed to master one condition. The generalist required a longer training period to master the diverse tool geometries. However, the evaluation on the held-out set of novel tools, shown in Table 1, revealed a critical flaw of the SPECIALIST. Its performance collapsed when presented with any tool different from the one it was trained on. The generalist agent, in contrast, maintained a high level of performance across all unseen tools. For completeness, we also evaluated both agents on the specific tool geometry used during the SPECIALIST’s training. The SPE-

CIALIST unsurprisingly maintained its performance from training with a mean excavated volume of 0.15 ± 0.06 L, while the generalist agent performed slightly worse but still competently with 0.12 ± 0.04 L. This outcome provides definitive evidence for our core hypothesis. Training with procedural diversity forces the agent to learn the underlying principles of excavation rather than memorizing a single tool-specific strategy, which is crucial for forging a truly robust and adaptable policy for the unpredictable conditions of space.

Effect of Visual Perception While the proprioceptive agent can infer some information from contact forces, it is effectively blind to the terrain before it. Our final study explores the benefit of augmenting the agent with vision. The visual feedback, shown in Figure 6, provides the agent with a rich, real-time understanding of its workspace, including the shape of the regolith pile and the geometry of its own tool. This added situational awareness proves to be a significant advantage. The learning curves in Figure 5 show that the VISUAL agent learned faster and converged to a higher final reward. The quantitative results in Table 1 confirm this superiority across all metrics. The agent with vision not only excavated a larger volume of material but also did so more safely with less generated dust. This confirms that direct perception enables more effective planning and reactive control, which is crucial for navigating the complex and dynamic environment of an excavation site.




Figure 6: Visual feedback from two views over the course of a single excavation episode. Equalized for clarity.

Implications for Tool Design Our methodology also presents a novel opportunity to inform hardware development. Predicting the most effective tool design is not straightforward, as each geometry presents a unique interaction challenge, and its performance is difficult to quantify because every tool may require a completely different control strategy. Because our learned agent is generalizable, it serves as a consistent evaluation metric for different mechanical designs. We leverage this capability by analyzing the performance of the VISUAL agent across the held-out set of novel tools, with each tool evaluated over 50 independent episodes. The results, presented in Figure 7, reveal a significant variance in performance attributable to tool geometry alone. Some designs consistently enabled the agent to excavate more material with lower dust generation, indicating a superior suitability for the task with a learned controller. For instance, tools with deeper cavities at their leading edge outperformed their shallower counterparts, while rounded designs generally led to less dust generation. Although this analysis is preliminary and hinges on the specific learned policy,




Figure 7: The performance of the VISUAL agent across the held-out set of novel tool geometries [vol. unit: L].

it nonetheless highlights the potential of using a learned agent as a tool for evaluating and optimizing mechanical designs. This data-driven approach could enable a synergistic co-design process where hardware and control policies are developed in parallel to create tools that are inherently more compatible with autonomous systems.

The results from our experiments provide strong and cohesive evidence for the effectiveness of our integrated methodology. The primary finding is that a generalizable excavation skill is forged from the powerful synergy of procedural diversity and learned adaptive compliance. By training with a wide distribution of tools, the agents discovered the underlying physical principles of excavation. They then expressed this knowledge by dynamically modulating their physical interaction style to suit the specific tool at hand. Furthermore, this work also suggests a new paradigm for co-designing robotic hardware and software. By systematically analyzing which tool geometries lead to superior performance with a generalist agent, we can identify design features that are inherently more compatible with autonomous control. This creates a powerful opportunity to optimize hardware and policies in a unified, data-driven loop. These findings are promising and grounded in simulation, highlighting clear directions for future research. The critical next step is a rigorous sim-to-real validation on a physical robotic testbed. Furthermore, the VISUAL agent faces the additional challenge of the perceptual sim-to-real gap, which remains a significant hurdle in robotics. Finally, the high computational cost of particle physics is a current limitation. Future research could explore more efficient simulation techniques to enable even larger and more complex training scenarios.

5. CONCLUSION

The successful establishment of a sustainable off-world presence requires autonomous systems capable of foundational ISRU tasks like regolith excavation. This paper presented an integrated solution to address the complex physics and hardware variability inherent to this task. This solution is built on the principle that robust autonomy is forged through diversity. By training a model-based RL agent within a procedurally diverse simulation

of both terrains and tools, we have shown it is possible to learn a truly generalizable excavation skill. The agent masters this challenge by learning to modulate its own physical compliance dynamically. This mechanism allows it to develop sophisticated, tool-aware strategies that are robust to novelty and safer in execution. The findings presented here offer more than a successful policy. They provide a validated methodology for overcoming the uncertainties of extraterrestrial environments. They also provide a novel framework for the data-driven co-design of robotic tools and their controllers. While sim-to-real transfer remains a critical next step, this research contributes the foundational knowledge necessary to advance the operational readiness of the autonomous systems that will build our multiplanetary future.

REFERENCES

- [1] National Aeronautics and Space Administration. NASA's Lunar Exploration Program Overview, 2020.
- [2] I. A. Crawford. Lunar Resources: A Review. *Progress in Physical Geography*, 39:137–167, 2015.
- [3] T. Zhang, K. Xu, Z. Yao, X. Ding, Z. Zhao, X. Hou, Y. Pang, X. Lai, W. Zhang, S. Liu, and J. Deng. The Progress of Extraterrestrial Regolith-Sampling Robots. *Nature Astronomy*, 3:487–497, 2019.
- [4] J. Mitchell, W. Houston, R. Scott, N. Costes, W. Carrier III, and L. Bromwell. Mechanical Properties of Lunar Soil: Density, Porosity, Cohesion and Angle of Internal Friction. In *Proceedings of the Lunar Science Conference*, 3:3235, 1972.
- [5] E. Slyuta. Physical and Mechanical Properties of the Lunar Soil. *Solar System Research*, 48(5):330–353, 2014.
- [6] G. Heiken, D. Vaniman, and B. M. French. *Lunar Sourcebook: A User's Guide to the Moon*. Number 1259. Cambridge University Press, 1991.
- [7] J. Connolly and W. D. Carrier. An Engineering Guide to Lunar Geotechnical Properties. In *IEEE Aerospace Conference*, 1–9, 2023.
- [8] N. Hogan. Impedance Control: An Approach to Manipulation: Part I—Theory. *Journal of Dynamic Systems, Measurement, and Control*, 107(1):1–7, 1985.
- [9] J. R. Gaier. The Effects of Lunar Dust on EVA Systems During the Apollo Missions, 2007.
- [10] G. Ishigami, A. Miwa, K. Nagatani, and K. Yoshida. Terramechanics-Based Model for Steering Maneuver of Planetary Exploration Rovers on Loose Soil. *Journal of Field Robotics*, 24(3):233–250, 2007.
- [11] F. E. Sotiropoulos. *Methods for Control in Robotic Excavation*. Massachusetts Institute of Technology, 2021.
- [12] A. Orsula, A. Richard, M. Geist, M. Olivares-Mendez, and C. Martinez. Towards Benchmarking Robotic Manipulation in Space. In *Conference on Robot Learning Workshop on Mastering Robot Manipulation in a World of Abundant Data*, 2024.
- [13] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 23–30. IEEE, 2017.
- [14] K. Cobbe, C. Hesse, J. Hilton, and J. Schulman. Leveraging Procedural Generation to Benchmark Reinforcement Learning. In *International Conference on Machine Learning*, 2048–2056, 2020.
- [15] D. I. Koutras, A. C. Kapoutsis, A. A. Amanatiadis, and E. B. Kosmatopoulos. MarsExplorer: Exploration of Unknown Terrains via Deep Reinforcement Learning and Procedurally Generated Environments. *Electronics*, 10(22), 2021.
- [16] A. Orsula, M. Geist, M. Olivares-Mendez, and C. Martinez. Leveraging Procedural Generation for Learning Autonomous Peg-in-Hole Assembly in Space. In *International Conference on Space Robotics*, 357–364, 2024.
- [17] A. Orsula, M. Geist, M. Olivares-Mendez, and C. Martinez. Sim2Dust: Mastering Dynamic Waypoint Tracking on Granular Media. In *International Conference on Space Robotics*, 2025.
- [18] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, 2018.
- [19] O. Azulay and A. Shapiro. Wheel Loader Scooping Controller Using Deep Reinforcement Learning. *IEEE Access*, 9:24145–24154, 2021.
- [20] Q. Lu, Y. Zhu, and L. Zhang. Excavation Reinforcement Learning Using Geometric Representation. *IEEE Robotics and Automation Letters*, 7(2):4472–4479, 2022.
- [21] Y. Niu, S. Jin, Z. Zhang, J. Zhu, D. Zhao, and L. Zhang. Goats: Goal Sampling Adaptation for Scooping with Curriculum Reinforcement Learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1023–1030. IEEE, 2023.
- [22] D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap. Mastering Diverse Control Tasks Through World Models. *Nature*, 640:647–653, 2025.
- [23] O. Khatib. A Unified Approach for Motion and Force Control of Robot Manipulators: The Operational Space Formulation. *IEEE Journal on Robotics and Automation*, 3(1):43–53, 1987.
- [24] J. Peters and S. Schaal. Learning Operational Space Control. In *Proceedings of Robotics: Science and Systems*, 2006.
- [25] J. Wong, V. Makoviychuk, A. Anandkumar, and Y. Zhu. OSCAR: Data-Driven Operational Space Control for Adaptive and Robust Robot Manipulation. In *International Conference on Robotics and Automation*, 10519–10526, 2022.
- [26] M. Mittal, C. Yu, Q. Yu, J. Liu, N. Rudin, D. Hoeller, J. L. Yuan, R. Singh, Y. Guo, H. Mazhar, A. Mandlekar, B. Babich, G. State, M. Hutter, and A. Garg. Orbit: A Unified Simulation Framework for Interactive Robot Learning Environments. *IEEE Robotics and Automation Letters*, 8(6):3740–3747, 2023.
- [27] M. Macklin, K. Storey, M. Lu, P. Terdiman, N. Chentanez, S. Jeschke, and M. Müller. Small Steps in Physics Simulation. In *Proceedings of the 18th Annual ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 1–7, 2019.
- [28] A. Orsula, M. Geist, M. Olivares-Mendez, and C. Martinez. Advancing Adaptive Autonomy Through Procedural Space Environments. In *International Astronautical Congress*, 2025.
- [29] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li. On the Continuity of Rotation Representations in Neural Networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [30] M. A. Felix, W. S. Slater, D. C. Landauer, R. E. Pinson, and B. B. W. Rutherford. Total Ionizing Dose Radiation Testing of NVIDIA Jetson Orin NX System on Module. In *IEEE Space Computing Conference*, 116–121. IEEE, 2024.