Lunar Excavator Mission Operations using Dynamic Movement Primitives

Joseph M. Cloud[1,2], Minh Q. Tram[1,2], William J. Beksi[2], and Michael A. DuPuis[1]

***Abstract*— To support sustainable infrastructure on the Moon, NASA must leverage robots to extract lunar resources for in-situ processing and construction. As part of this effort, NASA is launching the in-situ resource utilization (ISRU) Pilot Excavator later this decade to validate a robotic regolith excavator based on the Regolith Advanced Surface Systems Operations Robot (RASSOR). RASSOR is designed to extract and transport regolith to meet the needs of ISRU architectures. During its mission, Pilot Excavator will be tasked with driving in test patterns to demonstrate the operational concept. One of these tests is a circular trajectory around the lander while avoiding miscellaneous surface hazards such as lunar rocks. To this end, we utilize dynamic movement primitives to represent navigation sequences as primitive trajectories. Here, we introduce a novel obstacle avoidance parameter, which is configured to avoid rocks throughout testing exercises. We demonstrate the effectiveness our method in a newly developed simulation tool called the Simulated Excavation Environment for Lunar Operations (SEELO) using models based on the NASA RASSOR 2.0 excavator. After making key changes to the obstacle avoidance formulation, our results show that the robot is able to safety and robustly navigate the lunar surface with densely populated rock obstacles while retaining the desired circle pattern behavior.**

***Index Terms*— Space Robotics and Automation; Learning from Demonstration; Mining Robots**

I. Introduction

Establishing an affordable and sustainable human presence on the Moon necessitates using the resources available on the lunar surface. In-situ resource utilization (ISRU) is the practice of making use of these local resources [1]. From the lunar soil, also referred to as *regolith*, many vital resources can be extracted to sustain a habitat such as oxygen for rocket propellant and human breathing [2], [3], or using the soil as a construction material [4]. Advanced technologies such as additive manufacturing [5], [6] have been proposed for the construction of lunar habitat structures including landing pads, berms [7], [8], trenches [9], and radiation shields to name a few. Many of these technologies will utilize robots to improve operational efficiency [10].

Although robots will be extremely useful for lunar operations, there exists the following fundamental difficulties in robot operations on the lunar surface: (i) the presence of rocks [11] poses navigational hazards; (ii) limited communication bandwidth and high latency hinders the practicality of Earth-based teleoperation [12]; (iii) the abrasive properties of lunar dust accelerates wear-and-tear on mechanical systems [13], [14], [15]; and (iv) limited ability for the astronauts to directly interface with robots on the surface means that robots cannot rely on persistent real-time human supervision [16]. Consequently, ISRU robots must be designed to withstand these harsh conditions and continuously operate to fulfill mission goals for mining and construction [9], [10], [17].

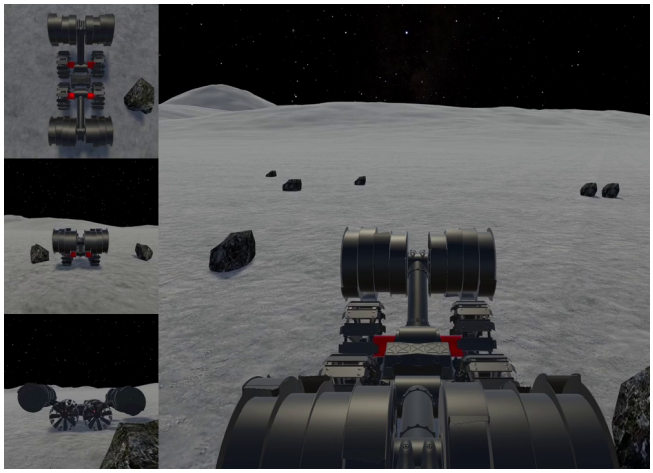

Fig. 1: The SEELO simulation tool depicting a scaled model of RASSOR 2.0 executing a circular dynamic movement primitive with obstacle avoidance. Multiple viewports are depicted.

As part of robotic precursor missions, excavators such as the ISRU Pilot Excavator (IPEx) will need to undergo testing in their relevant environments. This may include driving in specific patterns on the surface, various excavation geometries, and other tests to verify and validate the functionality of the excavator. The representation of these tests must be carefully considered due to limited on-board computational capability and the need to account for uncontrolled variables of the unstructured environment during the execution of the tests. To this end, we use dynamic movement primitives (DMPs) to represent the class of trajectories these robots will perform as part of their testing and performance duties. Specifically, we investigate circle pattern driving since it is part of the IPEx mission objectives.

Complex sequences of muscle movements have long been theorized by neuro-biologists to be composed of 'building block' movements known as motor primitives [18]. DMPs attempt to present this theory of motor primitives within an elegant mathematical framework. Robot learning from demonstration (LfD) is focused on enabling robots to autonomously acquire skills to complete tasks via hu-

[1]Advanced Engineering Development Branch, NE-L6, NASA John F. Kennedy Space Center, Kennedy Space Center, FL 32899, USA. Emails: {joseph.m.cloud, michael.a.dupuis}@nasa.gov

[2]Department of Computer Science and Engineering, The University of Texas at Arlington, Arlington, TX 76013, USA. Emails: minh.tram@mavs.uta.edu, william.beksi@uta.edu

man observation [19], [20]. Yet, one of the challenges in LfD is devising learning methods that can generalize, from demonstrations, and thus enable robots to perform tasks in dynamic environments. DMPs have been used within LfD frameworks with great success as they provide the means of generalizing from a demonstration such that trajectories can be adapted to new starting and goal locations, and even adapted mid-execution to account for secondary goals (e.g., obstacle avoidance).

In this work, we apply DMPs to a mobile robotic platform based on the Regolith Advanced Surface Systems Operations Robot (RASSOR) excavator. Furthermore, we demonstrate the method in our new Simulated Excavation Environment for Lunar Operations (SEELO) ISRU simulation tool, Fig. 1. We focus on circular pattern driving with obstacle avoidance to represent one of the potential test driving procedures for NASA's ISRU Pilot Excavator mission. *To the best of our knowledge, this is the first paper to evaluate the use of DMPs on a mobile ground vehicle in a simulated representative environment*.

The rest of the paper is structured as follows. We provide an overview of related research on obstacle avoidance (including DMPs) and lunar simulation tools in Section II. The specifics of the RASSOR excavator along with the mathematical properties of DMPs are described in Section III. In Section IV, the details of our methodology are presented. The experimental setup and results are discussed in Section V. We conclude and provide directions for future work in Section VI.

II. Related Work

A. Obstacle Avoidance with Movement Primitives

In related work, a bio-inspired approach to obstacle avoidance called the *steering angle* [21] has been used in conjunction with DMPs. The steering angle approach models human behavior of obstacle avoidance with a differential equation that relates the angle between a human's velocity and obstacle heading to a steering velocity to avoid colliding with the obstacle. In prior work, the steering angle method for obstacle avoidance was applied to DMPs for robotic applications [22], [23]. These works focused on obstacle avoidance via robotic manipulators with point-to-point trajectories using distinct start and goal locations. In [24], DMPs were applied with the steering angle obstacle avoidance formulation to a 2D grid-world mobile robotics simulation. In contrast, we apply obstacle avoidance to a more sophisticated 3D environment with circular trajectories, which required key changes not applicable to discrete point-to-point trajectories.

B. Lunar Rock Obstacle Avoidance

Various obstacle avoidance methods have been proposed over the years to address navigation around rocks on the lunar surface. Many of these methods have been tested in Earth-based field environments with vision-based sensing [25], [26], [27] to perform reactive obstacle avoidance and path planning. For conventional vision-based methods, these systems are tuned to work with path planners to assess the traversability of a trajectory. Consequently, these methods do not adhere to a particular pattern or "type" of trajectory as planning is based on gross waypoint navigation. Conversely, our primitive-based method allows us to retain the functional behavior of the trajectory and govern (via tuning the controllers) the obstacle avoidance term's strength in perturbing the primitive trajectory online during navigation.

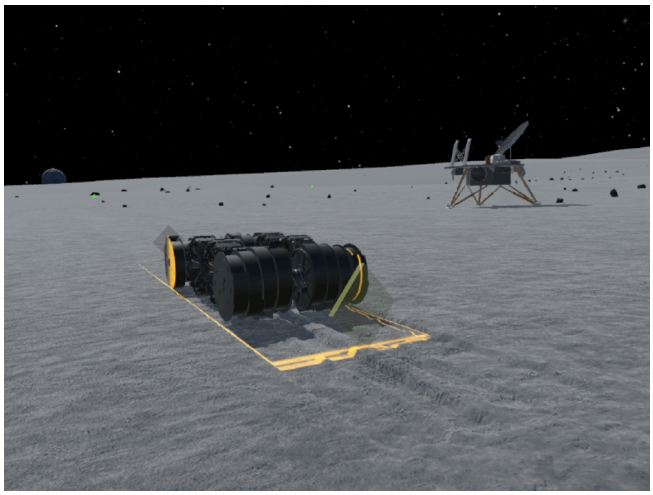

Fig. 2: The SEELO simulation depicting RASSOR excavating the lunar surface. As the robot grates the surface it leaves scoop tracks similar to the physical platform. We are able to replicate physical phenomena such as the depletion of on-board batteries and the increase in mass of the robot when excavating material.

In other work, an end-to-end reinforcement learning technique [28] has been demonstrated on lunar rock obstacle avoidance using vision data as input. End-to-end methods have the capacity to resolve non-trivial relationships between states and actions. However, the lack of explainability at the module-level means that these models are difficult to audit when they fail [29]. In contrast, using movement primitive-based trajectories means we know the structure of the executable behavior and we can modify its parameters in a deterministic fashion.

Developing a comprehensive lunar robotic simulation is a challenging problem since fidelity comes at the cost of performance and both must be carefully balanced. Several robotic lunar simulation tools are currently under active development [30], [31]. Nevertheless, to the best of our knowledge, none support terrain deformation, which is a unique feature of both SEELO (Fig. 2) and our precursor simulation developed in prior work [32]. Previously, RASSOR was simulated using Gazebo [33], an open-source 3D robotics simulation. Gazebo has been used in other works to simulate the lunar environment (e.g., [30], [31]).

Despite widespread use in the robotics community, Gazebo lacks the configuration flexibility offered by some game physics engines such as Unity, thus rendering it unsuitable for higher fidelity work. In particular, Gazebo does not natively support terrain deformation whereas SEELO is designed with this capability in mind. Moreover, Gazebo

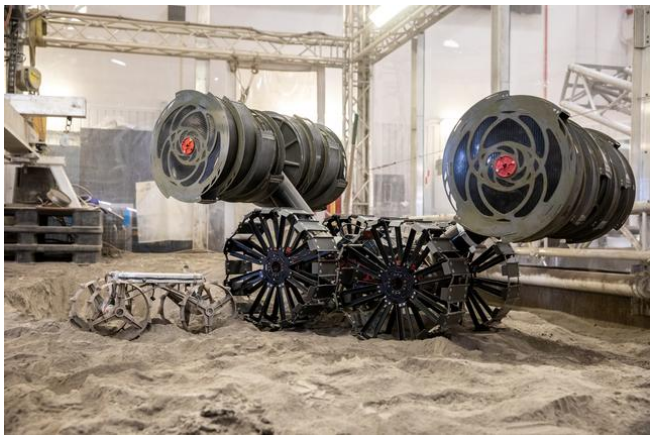

Fig. 3: RASSOR 2.0 in the 'Big Bin' at NASA's Kennedy Space Granular Mechanics & Regolith Operations Laboratory. Here, sensor data was gathered to inform the modeling of our simulation.

is not designed to render photorealistic scenes so it cannot generate dust or other effects that align with our simulation goals. We developed SEELO to address the need for a high-fidelity ISRU excavation simulation. SEELO's capability also enables us to explore the use of movement primitives since our implementation allows us to tightly couple robot control and physical modeling in a more realistic manner.

III. Background

A. RASSOR Excavator

RASSOR 2.0 (also referred to as *RASSOR*) is a Technology Readiness Level 4 (i.e., component/subsystem validation in a laboratory environment) planetary excavation robot approximately 2 m in length and 66 kg of mass, Fig 3. Its counter-rotating bucket drum design enables operation in low gravity environments where it can efficiently perform deep regolith excavation and slot trenching. RASSOR's mechanical design consists of four wheels, two arms, and two sets of bucket drums, which excavate and store regolith material [34]. Additionally, RASSOR has on-board vision and inertial measurement sensors. Each of the robot's actuators includes high-fidelity encoders for precise angular measurements. Our SEELO simulation also possesses similar sensing capabilities. In SEELO, we use a scaled-down model of the RASSOR excavator to approximate the dimensions of the ISRU Pilot Excavator as it is under active development.

B. Dynamic Movement Primitives

A DMP is derived from a non-linear spring-damper system with gains selected to render it critically damped. It is perturbed with additional forcing terms that generate the novel behavior captured in a task demonstration. To do this, the forcing term must be learned from the demonstration represented as a trajectory. For example, the robot could be kinesthetically moved to demonstrate the task as the joint positions are recorded throughout the demonstration. The forcing term is estimated from the demonstration and then learned by performing locally-weighted regression. In our work with mobile robots, we use Cartesian trajectories to represent robot motion.

The major advantage of DMPs is their generalizability. Intuitively, executing a basic motor behavior is localized in space and should be flexible in velocity, along with where the motion starts and ends, while maintaining the fundamental behavior. DMPs possess this property as both the start and goal locations can be adjusted and a separate temporal scalar can be used to adapt the speed of execution [35]. Another important characteristic of DMPs is their ability to perform online adaptation, i.e., the robot can react to environmental changes mid-execution. This includes obstacle avoidance [22] and the adaptation based on force feedback [36]. For instance, this can be seen within exoskeletons where forces imparted by an operator help tune rhythmic DMPs such that the suit learns to minimize human effort [37].

The classical one-dimensional DMP is represented as [38]

$$\begin{aligned} \tau \dot{v} &= K(g - x) - Dv + (g - x_0)f(s), \\ \tau \dot{x} &= v, \end{aligned} \tag{1}$$

where $x, v \in \mathbb{R}$ are the position and velocity, $x_0, g \in \mathbb{R}$ are the initial and goal positions, and $K, D \in \mathbb{R}^+$ are constants for the spring and damping terms. In (1), $D = 2\sqrt{K}$ to render system critically damped, τ is a positive speed scaling factor, and f is the s dependent forcing function to be learned. Given the decay term, α, s abstracts away time using the canonical system and exponentially decays from 1 to 0, i.e.,

$$\tau \dot{s} = -\alpha s. \tag{2}$$

The forcing term is defined as

$$f(s) = \frac{\sum_{i=0}^{N} \omega_i \, \psi_i(s)}{\sum_{i=0}^{N} \psi_i(s)} s, \tag{3}$$

where N Gaussian basis functions, represented by ψ_i, are sum weighted against the learned weights ω. With this, (1) can be rewritten to calculate the target forces. Then, the weights can be computed using locally weighted regression.

This formulation can be extended to vector form to handle multiple degrees of freedom [39]. One of the issues in this formulation is that start and goal locations that are close to each other will limit the influence of the forcing term. Later work [22] resolved this by separating the constant $g - x_0$ from the forcing term and dissipating it using s, i.e.,

$$\tau \dot{\mathbf{v}} = \mathbf{K}(\mathbf{g} - \mathbf{x}) - \mathbf{D}\mathbf{v} - \mathbf{K}(\mathbf{g} - \mathbf{x}_0)s + \mathbf{K}\mathbf{f}(s) \tag{4}$$

In (4), $\mathbf{K}$ and $\mathbf{D}$ are $d \times d$ diagonal matrices for d-dimensional trajectories. The other notable change is that the contribution of the constant $g - x_0$ is phased out with s such that the forcing term is no longer dependent on it.

C. Steering Angle Obstacle Avoidance with DMPs

Building upon the formulation of DMPs in (4), we can append additional forcing terms to further perturb the behavior of the primitive. For instance, in [23] obstacle avoidance functionality was extended to DMPs and demonstrated on a robotic manipulator by incorporating a steering angle term.

This is accomplished by modeling the relationship between the steering velocity corresponding to the angle between the velocity of the end-effector, and a vector pointing in the direction of the obstacle with respect to the current position of the end-effector. Concretely,

$$\dot{\varphi} = \gamma\varphi \exp(-\beta|\varphi|), \tag{5}$$

where γ is the gain that determines the force the obstacle avoidance term on the DMP spring-damper system, φ is the steering angle (described below), and β determines the acceleration of steering.

When the angle between the current velocity and the direction to the obstacle is close, we want to steer hard in a direction away and eventually around the obstacle. Otherwise, if the robot is moving with a velocity orthogonal to obstacle heading, no perturbation is necessary. We calculate the steering angle, φ as

$$\varphi = \cos^{-1}\left(\frac{(\mathbf{o}-\mathbf{x})^\top \mathbf{v}}{|\mathbf{o}-\mathbf{x}||\mathbf{v}|}\right), \tag{6}$$

where $\mathbf{x}$ is the position of the robot, $\mathbf{o}$ is the position of the obstacle, and $\mathbf{v}$ is the current velocity of the robot. With φ, we can then use (5) to calculate the steering velocity to drive the robot away from the obstacle. We wish to perturb our original transformation system described in (4) with an additional term $\mathbf{p}(\mathbf{x}, \dot{\mathbf{x}})$ representing the obstacle avoidance dynamics, i.e.,

$$\tau\dot{\mathbf{v}} = \mathbf{K}(\mathbf{g}-\mathbf{x}) - \mathbf{D}\mathbf{v} - \mathbf{K}(\mathbf{g}-\mathbf{x}_0)s + \mathbf{K}\mathbf{f}(s) + \mathbf{p}(\mathbf{x}, \dot{\mathbf{x}}), \tag{7}$$

which incorporates the steering angle methodology previously discussed. We use (5) and (6) to write

$$\mathbf{p}(\mathbf{x}, \dot{\mathbf{x}}) = \gamma \boldsymbol{R}\dot{\mathbf{x}}\varphi \exp(-\beta\varphi), \tag{8}$$

where $\boldsymbol{R}$ is a rotation matrix with axis $\left(\frac{(\mathbf{o}-\mathbf{x})\times\dot{\mathbf{x}}}{|\mathbf{o}-\mathbf{x}||\dot{\mathbf{x}}|}\right)$ rotated by $\frac{\pi}{2}$ from [23]. In practice, a threshold is typically used to assess if an obstacle is further away from the goal than the robot. However, this will not work with a circular pattern driving. In the next section, we introduce our method to overcome this drawback.

IV. Methodology

A. Avoiding Obstacles in Circular Movement Primitives

In the previous section, we formalized the DMP structure and described the steering angle approach for obstacle avoidance with discrete point-to-point trajectories. In this section, we present our contributions for an obstacle avoidance method that enables practical use of a circular trajectory. In our work, we consider circular trajectories to have the same start and goal location. Consequently, some of the assumptions made in previous work do not apply. Namely, the distance of the obstacle to the goal location is no longer an appropriate means of filtering out obstacles. Secondly, a priori knowledge that a circular movement primitive is being executed allows us to make another change to the obstacle avoidance term to improve efficiency. Finally, the obstacle avoidance formulation previously introduced assumes global knowledge of obstacles, which may be an infeasible assumption. We address each of these issues separately.

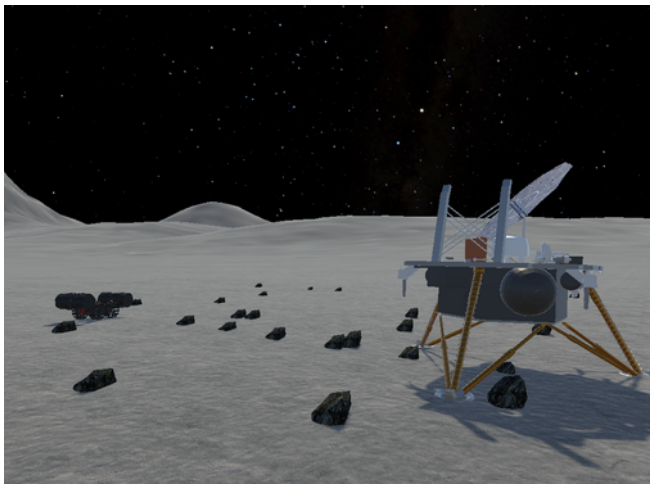

Fig. 4: The SEELO environment used in the experimental evaluation. The rocks are randomly sampled and placed on the surface.

In practice, the original obstacle avoidance formulation [23] used a threshold to determine if an obstacle was further away from the goal location than the robot's current position. This assumption removed the effect of obstacles "behind" the robot. In the case of a circular trajectory, we remove this obstacle filter and instead consider a sensory limit threshold. Using this threshold, we set a radius around the robot's current position to restrict the influence of nearby obstacles. In our implementation, we modified (8) to be

$$\mathbf{p}(\mathbf{x}, \dot{\mathbf{x}}) = \sum_{\mathbf{o}\in\text{obstacles}} \begin{cases} z\gamma \boldsymbol{R}_m \dot{\mathbf{x}}\varphi \exp(-\beta\varphi)\delta^{-1}, & \text{if } z < \delta \\ \mathbf{0}, & \text{otherwise} \end{cases} \tag{9}$$

where $z = |\mathbf{o}-\mathbf{x}|$ and δ is a set threshold representing the maximum detection distance between the robot and the obstacle. All obstacles outside the radius specified by δ are ignored. Without this change, all obstacles (regardless of distance to trajectory) would perturb the robot since they will be tangential to the robot during execution. We show in our experimental results why this change is necessary. We also define $\boldsymbol{R}_m$ to be a matrix that is rotated by δ/r where r is the nominal radius of the executed movement primitive. As the robot heading aligns with the tangent of the circle, obstacles will bias the avoidance term to drive the robot towards the outside of the circle. This parameter allows us to instead adapt the steering of the robot to the interior of the circle such that it tends towards a shorter path when deviating from the canonical DMP. When driving the robot in clockwise circular patterns we multiple δ/r by -1.

B. Simulation Environment

In previous work [32], we developed a 3D simulation environment using the Unity Game Engine [40] to represent lunar excavation for reinforcement learning tasks. Building upon this, SEELO enhanced the fidelity of the

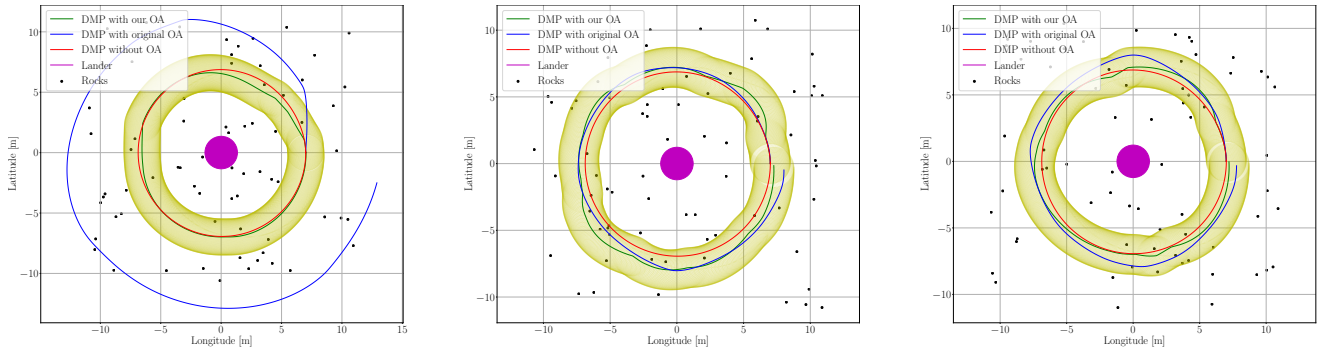

(a) Execution A (b) Execution B (c) Execution C

Fig. 5: In each of the three executions, we see a unique uniformly sampled distribution of rocks (represented by black dots) traversed by the same circular DMP. The 2D representation of the original circular DMP without any obstacle avoidance (OA) is in red, our method is in green, and the original obstacle avoidance method is in blue. We use a yellow boundary to indicate the sensing limits of the DMP with our obstacle avoidance method over the history of the trajectory execution.

terrain deformation and physical modeling of the robot in addition to consolidating other capabilities. SEELO seeks to provide life-accurate mission environments and regolith interaction mechanics while remaining lightweight enough to run faster than real-time, Fig. 4. These combined features allow SEELO to be used for mission planning purposes and offer an environment for lunar robots to perform machine learning for mobility, excavation, and various other activities. SEELO uses lunar surface data collected by NASA's Lunar Reconnaissance Orbiter to realistically represent proposed landing sites [41]. The physical and empirical modeling of the RASSOR excavator at Kennedy Space Center has informed the attributes of the simulated analogue. This includes calculating the baseline power usage, empirically-derived equations relating excavated material to effects on actuator load, and mobility power calculations.

V. Experiments

A. Experimental Setup

We conducted our experiments using the SEELO simulation. In addition to the scaled model of the RASSOR excavator, a lunar lander was placed as a marker in which the robot drove around similar to the IPEx mission objectives. For our results depicted in Figs. 5-7, we model a challenging environment [42] by uniformly sampling 10 rock distributions consisting of 70 rocks each with a volume of 67.8 dm^3 per rock. We selected a rock large enough to be considered non-traversable by the robot. The rocks are distributed within a 22 m $\times$ 22 m worksite with a density of 0.15 rocks/m^2. For each distribution of rocks, we evaluated our modified obstacle avoidance method (referred to as OA in the figures) against both the original formulation presented in [23] and an unmodified DMP without obstacle avoidance. We only trained the DMP once using a synthetically generated circular trajectory with a radius of 7.0 m. We also sampled 100 additional rock distributions and quantitatively evaluated the performance of the obstacle avoidance methods separately.

Within the simulation, we were able to measure the distance driven and power consumed by the robot throughout the operation. These measurements are provided as part of

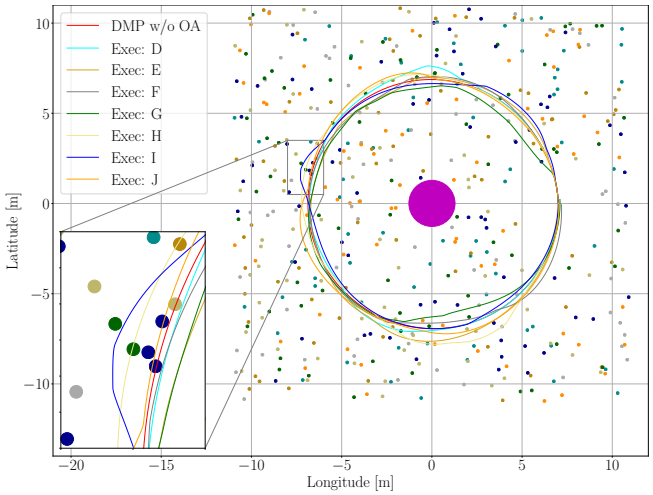

Fig. 6: A simplified representation of the seven remaining executions with novel rock distributions. A darker shade of the trajectory color denotes the associated rock distribution.

the quantitative discussion. We only considered one size of rock and we exaggerated the rock size to improve visibility and more clearly demonstrate the capability of our method. Furthermore, we assessed each DMP-generated trajectory for traversability with obstacles. We recorded collisions for any trajectory that causes the robot to go within a 1.5x safety factor of its radius to the obstacle. The obstacle avoidance gains can be tuned to increase the boundary between the robot and obstacle as needed. In all of our experiments, we set γ to 1500000, β to 10, and the value of δ is 1.5 m. For the DMP, we set α to 30 and included 1000 Gaussian basis-functions. We separately tuned the obstacle avoidance term for the DMP with the original formulation and used $\gamma = 100000$ for all experiments with the exception of Execution A.

B. Experimental Results

In Fig. 5, we show three of the ten executions. The last seven executions are displayed in Fig. 6. For the trajectories depicted in Fig. 5, we see that our method (in green) only deviates from the canonical DMP (without obstacle

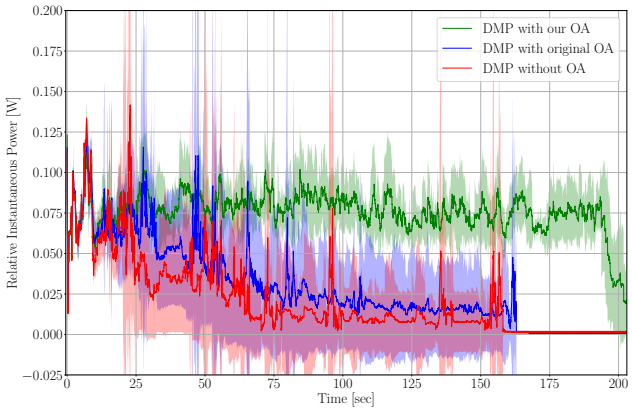

Fig. 7: A power draw comparison between the three DMP approaches (with/without OA) across ten trials. In the simulation, we record the instantaneous power draw from the motors and present the average (dark lines) and standard deviation (shaded area).

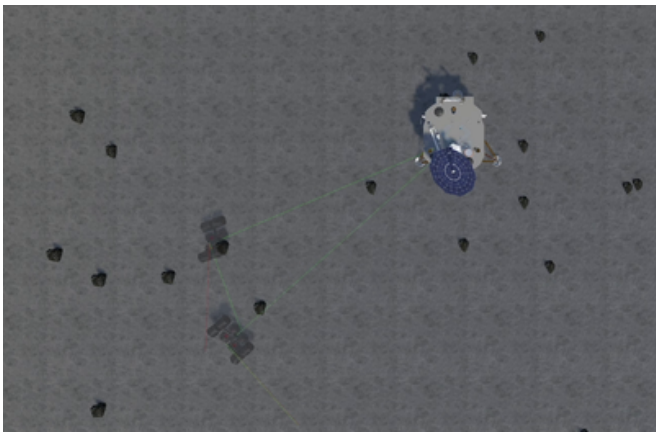

Fig. 8: A top-down view of the SEELO execution of overlayed DMPs. The left-most robot has hit an obstacle and is stuck on an incline whereas the second robot using our DMP obstacle avoidance method can successfully navigate between the rocks.

avoidance in red) when an obstacle is directly in its path. Obstacle avoidance is performed online and the robot is only aware of the obstacle when within 1.5 m as defined by δ.

The original obstacle avoidance formulation introduced in [23] (shown in blue) demonstrates worse tracking performance for the circular DMP. In particular, the original method does not consider proximity nor does it apply a bound on the distance between the obstacles and the robot as we have done in our approach. As a result, since most obstacles are present within the radius of the circular trajectory, the obstacle avoidance term continues to push the blue trajectory "outward" from the lander in a spiral. This is not desirable behavior and we tuned the obstacle avoidance parameters to reduce the effect.

In Fig. 5, the original obstacle avoidance technique in Execution A used the same parameters as our method, whereas in Execution B (and the rest of the tests) we used new parameters specifically tuned for the original approach. Due to the influence of the totality of the obstacles, the individual weight of a single obstacle is diminished, which means that the trajectory can inadvertently steer into a collision such as the case in Execution C at position $(0.0, -7.5)$.

In our quantitative evaluation, we found our method required the robot to drive 44.11 ± 0.65 m on average. In contrast, the original obstacle avoidance formulation traversed 46.83 ± 0.49 m across the terrain. The nominal circular trajectory for the unmodified DMP is 43.45 m, but this was never achieved in the simulation since a collision with an obstacle would halt the robot mid-execution. In Fig. 7, we compare the power consumption of the ten executions across all three DMP modalities. The "spikes" on the original and unmodified DMP methods are due to rock collisions mid-execution. In most cases, these impacts stop the execution of the robot. Our approach was the only method to reach the goal state and it successfully completed every test. Consequently, the average power draw for our method is higher compared to the other methods. When a robot becomes stuck and the wheels freely spin the power consumption is lowered.

We did not collide with any obstacles with our methodology during the ten trials conducted. Therefore, we conducted an additional 100 tests with novel rock positions. We found that our method again avoided all obstacles across the 100 tests, whereas the unmodified DMP would have collided with 426 obstacles in the environment. Similarly, the original obstacle avoidance approach collided with 425 obstacles. Fig. 8 illustrates one of these cases where a robot employing our method successfully navigates between obstacles and another robot collides with a rock (shown pitched upward over the rock) using the unmodified DMP trajectory.

VI. Conclusion

In this paper, we used DMPs to represent lunar excavator trajectories enabling us to define a compact, generalized representation of the motion plans on the surface. We deployed our method in a comprehensive 3D lunar ISRU simulation (SEELO). During the execution of the method within the simulation, we adapted the behavior of the DMP to rock obstacles in the environment. To do this, we devised a modified approach to obstacle avoidance with DMPs that enabled us to traverse obstacle-laden environments in circular patterns. We also compared our technique against the canonical obstacle avoidance method for DMPs and addressed the quantitative and qualitative differences in our results. As part of our future work, we will leverage the movement primitives framework we implemented to investigate rapid-retasking of robots, implement additional forcing terms to address other excavator mission criteria, the generation of other types of primitives relevant to lunar robotic mission operations, and testing on physical mobile robots.

Acknowledgments

We thank the many NASA interns that have contributed to the development of the SEELO simulation. We also thank members of the Granular Mechanics & Regolith Operations Laboratory, Advanced Engineering Development Branch, and the Pathways Office at the NASA Kennedy Space Center for their support.

REFERENCES

[1] G. B. Sanders, W. E. Larson, K. R. Sacksteder, and C. A. Mclemore, "NASA in-situ resource utilization (ISRU) project: Development and implementation," in *Proceedings of the AIAA SPACE Conference & Exposition*, 2008, p. 7853.

[2] M. Anand, I. A. Crawford, M. Balat-Pichelin, S. Abanades, W. Van Westrenen, G. Péraudeau, R. Jaumann, and W. Seboldt, "A brief review of chemical and mineralogical resources on the Moon and likely initial in situ resource utilization (ISRU) applications," *Planetary and Space Science*, vol. 74, no. 1, pp. 42–48, 2012.

[3] P. van Susante and L. Gertsch, "Minerals from space: Terrestrial and extra-terrestrial perspectives," in *Proceedings of ASCE Earth and Space: Engineering for Extreme Environments*, 2018, pp. 390–400.

[4] G. B. Sanders and W. E. Larson, "Integration of in-situ resource utilization into lunar/mars exploration through field analogs," *Advances in Space Research*, vol. 47, no. 1, pp. 20–29, 2011.

[5] B. C. Buckles, R. P. Mueller, and N. J. Gelino, "Additive construction technology for lunar infrastructure," *LPI Contributions*, vol. 2152, p. 5077, 2019.

[6] R. P. Mueller, N. J. Gelino, J. D. Smith, B. C. Buckles, T. Lippitt, J. M. Schuler, A. J. Nick, M. W. Nugent, and I. I. Townsend, "Zero launch mass three-dimensional print head," in *Proceedings of ASCE Earth and Space: Engineering for Extreme Environments*, 2018, pp. 219–232.

[7] R. M. Kelso, R. Romo, C. Andersen, R. P. Mueller, T. Lippitt, N. J. Gelino, J. D. Smith, I. I. Townsend, J. M. Schuler, M. W. Nugent, A. J. Nick, K. Zacny, and M. Hedlund, "Planetary basalt field project: Construction of a lunar launch/landing pad, PISCES and NASA Kennedy Space Center project update," in *Proceedings of ASCE Earth and Space: Engineering for Extreme Environments*, 2016, pp. 653–667.

[8] N. J. Gelino, R. P. Mueller, R. W. Moses, J. G. Mantovani, P. T. Metzger, B. C. Buckles, and L. Sibille, "Off earth landing and launch pad construction—a critical technology for establishing a long-term presence on extraterrestrial surfaces," in *Proceedings of ASCE Earth and Space*, 2021, pp. 855–869.

[9] R. P. Mueller, J. D. Smith, J. M. Schuler, A. J. Nick, N. J. Gelino, K. W. Leucht, I. I. Townsend, and A. G. Dokos, "Design of an excavation robot: regolith advanced surface systems operations robot (RASSOR) 2.0," in *Proceedings of ASCE Earth and Space: Engineering for Extreme Environments*, 2016, pp. 163–174.

[10] G. B. Sanders and W. E. Larson, "Progress made in lunar in situ resource utilization under NASA's exploration technology and development program," *Proceedings of ASCE Earth and Space: Engineering, Science, Construction, and Operations in Challenging Environments*, pp. 457–478, 2012.

[11] G. J. Taylor, P. Warren, G. Ryder, J. Delano, C. Pieters, and G. Lofgren, "Lunar rocks," *Lunar Sourcebook*, pp. 183–284, 1991.

[12] J. O. Burns, B. Mellinkoff, M. Spydell, T. Fong, D. A. Kring, W. D. Pratt, T. Cichan, and C. M. Edwards, "Science on the lunar surface facilitated by low latency telerobotics from a lunar orbital platform-gateway," *Acta Astronautica*, vol. 154, pp. 195–203, 2019.

[13] 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 and Planetary Science Conference*, vol. 3, 1972, p. 3235.

[14] L. Taylor, H. Schmitt, W. Carrier, and M. Nakagawa, "Lunar dust problem: From liability to asset," in *Proceedings of the Space Exploration Conference: Continuing the Voyage of Discovery*, 2005, p. 2510.

[15] C. I. Calle, "The electrostatic environments of Mars and the Moon," *Journal of Physics: Conference Series*, vol. 301, no. 1, p. 012006, 2011.

[16] T. Fong, J. Rochlis Zumbado, N. Currie, A. Mishkin, and D. L. Akin, "Space telerobotics: unique challenges to human–robot collaboration in space," *Reviews of Human Factors and Ergonomics*, vol. 9, no. 1, pp. 6–56, 2013.

[17] K. Zacny, "Lunar drilling, excavation and mining in support of science, exploration, construction, and in situ resource utilization (ISRU)," in *Moon*. Springer, 2012, pp. 235–265.

[18] S. F. Giszter, F. A. Mussa-Ivaldi, and E. Bizzi, "Convergent force fields organized in the frog's spinal cord," *Journal of Neuroscience*, vol. 13, no. 2, pp. 467–491, 1993.

[19] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009.

[20] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, "Learning and generalization of motor skills by learning from demonstration," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2009, pp. 763–768.

[21] B. Fajen and W. Warren, "Behavioral dynamics of steering, obstacle avoidance, and route selection," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 29, no. 2, pp. 343–362, 2003.

[22] D.-H. Park, H. Hoffmann, P. Pastor, and S. Schaal, "Movement reproduction and obstacle avoidance with dynamic movement primitives and potential fields," in *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, 2008, pp. 91–98.

[23] H. Hoffmann, P. Pastor, D.-H. Park, and S. Schaal, "Biologically-inspired dynamical systems for movement generation: Automatic real-time goal adaptation and obstacle avoidance," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2009, pp. 2587–2592.

[24] Z. Mei, Y. Chen, M. Jiang, H. Wu, and L. Cheng, "Mobile robots path planning based on dynamic movement primitives library," in *Proceedings of the IEEE Chinese Control Conference*, 2017, pp. 6906–6911.

[25] R. Simmons, L. Henriksen, L. Chrisman, and G. Whelan, "Obstacle avoidance and safeguarding for a lunar rover," in *Proceedings of the AIAA Forum on Advanced Developments in Space Robotics*, 1996, pp. 1–9.

[26] L. M. Lorigo, R. A. Brooks, and W. Grimsou, "Visually-guided obstacle avoidance in unstructured environments," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robot and Systems*, vol. 1, 1997, pp. 373–379.

[27] H. Utz and T. Ruland, "Reactive, safe navigation for lunar and planetary robots," in *Proceedings of the AIAA SPACE Conference & Exposition*, 2008, p. 7848.

[28] T. Blum, G. Paillet, W. Masawat, and K. Yoshida, "SegVisRL: development of a robot's neural visuomotor and planning system for lunar exploration," *Advanced Robotics*, vol. 35, no. 21-22, pp. 1359–1373, 2021.

[29] T. Glasmachers, "Limits of end-to-end learning," in *Proceedings of the Asian Conference on Machine Learning*. PMLR, 2017, pp. 17–32.

[30] M. Allan, U. Wong, P. M. Furlong, A. Rogg, S. McMichael, T. Welsh, I. Chen, S. Peters, B. Gerkey, M. Quigley, M. Shirley, M. Deans, H. Cannon, and T. Fong, "Planetary rover simulation for lunar exploration missions," in *Proceedings of the IEEE Aerospace Conference*, 2019, pp. 1–19.

[31] EZRASSOR Team, "EZRASSOR: EZ Regolith Advanced Surface Systems Operations Robot," https://github.com/FlaSpaceInst/EZ-RASSOR, 2021, commit: f034dc9.

[32] J. M. Cloud, R. J. Nieves, A. K. Duke, T. J. Muller, N. A. Janmohamed, B. C. Buckles, and M. A. DuPuis, "Towards autonomous lunar resource excavation via deep reinforcement learning," in *Proceedings of the AIAA Accelerating Space, Commerce, Exploration, and New Discovery Conference*, 2021, p. 4217.

[33] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, 2004, pp. 2149–2154.

[34] R. P. Mueller, R. E. Cox, T. Ebert, J. D. Smith, J. M. Schuler, and A. J. Nick, "Regolith advanced surface systems operations robot (RASSOR)," in *Proceedings of the IEEE Aerospace Conference*, 2013, pp. 1–12.

[35] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, "Dynamical movement primitives: learning attractor models for motor behaviors," *Neural Computation*, vol. 25, no. 2, pp. 328–373, 2013.

[36] A. Kramberger, A. Gams, B. Nemec, and A. Ude, "Generalization of orientational motion in unit quaternion space," in *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, 2016, pp. 808–813.

[37] L. Peternel, T. Noda, T. Petrič, A. Ude, J. Morimoto, and J. Babič, "Adaptive control of exoskeleton robots for periodic assistive behaviours based on emg feedback minimisation," *PLOS One*, vol. 11, no. 2, 2016.

[38] A. Ijspeert, J. Nakanishi, and S. Schaal, "Learning attractor landscapes for learning motor primitives," in *Proceedings of the Advances in Neural Information Processing Systems*, vol. 15, 2002.

[39] M. Saveriano, F. J. Abu-Dakka, A. Kramberger, and L. Peternel, "Dynamic movement primitives in robotics: A tutorial survey," *arXiv preprint arXiv:2102.03861*, 2021.
[40] "Unity 3d game engine," 2022. [Online]. Available: https://unity.com/
[41] D. E. Smith, M. T. Zuber, G. A. Neumann, F. G. Lemoine, E. Mazarico, M. H. Torrence, J. F. McGarry, D. D. Rowlands, J. W. Head, T. H. Duxbury, O. Aharonson, P. G. Lucey, M. S. Robinson, O. S. Barnouin, J. F. Cavanaugh, X. Sun, P. Liiva, D.-d. Mao, J. C. Smith, and A. E. Bartels, "Initial observations from the lunar orbiter laser altimeter (LOLA)," *Geophysical Research Letters*, vol. 37, no. 18, 2010.
[42] J. L. Bandfield, R. R. Ghent, A. R. Vasavada, D. A. Paige, S. J. Lawrence, and M. S. Robinson, "Lunar surface rock abundance and regolith fines temperatures derived from lro diviner radiometer data," *Journal of Geophysical Research: Planets*, vol. 116, no. E12, 2011.