

Learning manipulation of steep granular slopes for fast Mini Rover turning

Deniz Kerimoglu^{1,*}, Daniel Soto^{1,*}, Malone Lincoln Hemsley², Joseph Brunner¹, Sehoon Ha¹, Tingnan Zhang³
and Daniel I. Goldman¹

Abstract— Future planetary exploration missions will require reaching challenging regions such as craters and steep slopes. Such regions are ubiquitous and present science-rich targets potentially containing information regarding the planet’s internal structure. Steep slopes consisting of low-cohesion regolith are prone to flow downward under small disturbances, making it challenging for autonomous rovers to traverse. Moreover, the navigation trajectories of rovers are heavily limited by the terrain topology and future systems will need to maneuver on flowable surfaces without getting trapped, allowing them to further expand their reach and increase mission efficiency.

In this work, we used a robophysical rover model and performed maneuvering experiments on a steep granular slope of poppy seeds to explore the rover’s turning capabilities. The rover is capable of lifting, sweeping, and spinning its wheels, allowing it to execute leg-like gait patterns. The high-dimensional actuation capabilities of the rover facilitate effective manipulation of the underlying granular surface. We used Bayesian Optimization (BO) to gain insight into successful turning gaits in high dimensional search space and found strategies such as differential wheel spinning and pivoting around a single sweeping wheel. We then used these insights to further fine-tune the turning gait, enabling the rover to turn nearly 90 degrees at just above 4 seconds with minimal downhill slip. Combining gait optimization and human-tuning approaches, we found that fast turning is empowered by creating anisotropic torques with the sweeping wheel.

I. INTRODUCTION

Almost all portions of the celestial bodies and a significant portion of the granular surface on Earth remain inaccessible to the majority of wheeled and tracked vehicles. This is partly due to the limitations of locomoting robots that struggle to effectively interact with their immediate surroundings. Particularly, wheeled systems can get stuck in loose granular media as they sink deeper into the material during the wheel-spinning motion [1]. These difficulties intensify in the vicinity of craters and other destinations characterized by steep, flowable surfaces, where a minimal disturbance could potentially cause yielding and downhill flow. Such regions are scientifically valuable locations, potentially harboring layers of geological history as well as containing shaded

areas where water ice might be preserved [2], [3]. To explore these difficult terrains, the next generation of exploratory robots will need to navigate over steep flowable slopes.

Future missions will also require autonomous systems to negotiate difficult terrain topologies by performing effective maneuvers to increase access to a diverse range of terrain types. This will allow currently-denied extreme terrain to become traversable, hence achieving scientific mission goals would be less time- and energy-consuming due to the greatly shortened paths.

Recent research has addressed some of the mobility limitations to some extent. For loose regolith/soil on level ground, refined terramechanical models [4] and new methods [5] based on granular resistive force theory better predict the rover-terrain interaction. On the robotics side, various wheeled and non-wheel robotic platforms have been developed [6]–[8]. Few have achieved robust mobility in loose granular slopes, [9]–[11]. Various studies focused on finding efficient paths on slopes to avoid failure [12]–[14]. Recently, researchers have been investigating unconventional robots and rovers that depart from traditional designs to navigate across loose granular slopes [15]–[17].” However, robotic turning has not been extensively explored on steep loose granular slopes yet.

Our previous work explored the climbing gaits of a laboratory-scale rover robot, namely Mini Rover, a scaled-down version of a robotic lunar rover prototype called NASA Resource Prospector 15 [18], [19]. The Mini Rover incorporates wheel and leg-like appendages that actively manipulate flowable terrain to climb granular slopes [18]. We extended Mini Rover’s capabilities in [19] using Bayesian optimization, enabling the discovery of fast climbing gaits by exploiting the fluid and solid aspects of granular materials. These investigations exemplify a wider category of interactions involving locomotor and flowable terrain. The locomotor becomes integrated with its surroundings behaving as a unified structure and actively modifies its environment to achieve success.

In this work, motivated by leveraging terrain manipulation strategies, we study the turning of the Mini Rover robot on a steep granular slope. To this end, simulation environments to investigate the Mini Rover gaits are impractical since simulating interactions between the robot and the millions of granular particles is extremely time-consuming even with the use of high-performance computing tools. Moreover, both the robot and the terrain exhibit complex, nonlinear dynamics,

*Co-first authors

¹Deniz Kerimoglu, Daniel Soto, Joseph Brunner, and Daniel I. Goldman are with School of Physics, Georgia Institute of Technology, Atlanta, USA
dkerimoglu6@gatech.edu

²Malone Lincoln Hemsley is with the Department of Applied Physics and Engineering, Morehouse College, Atlanta, USA

³Sehoon Ha is with the School of Interactive Computing, Georgia Institute of Technology, Atlanta, USA

⁴Tingnan Zhang is with the Google Robotics, Mountain View, California, USA

hence deriving closed-form solutions to represent such interactions is not yet possible, hindering model-based simulation approaches. Therefore we rely on laboratory experiments and optimization methods to find and investigate Rover gait patterns.

In this work, motivated by laboratory experiments and optimization methods to find and investigate robotic gait patterns, we study the turning of the Mini Rover robot on a steep granular slope. We performed robophysical experiments because simulation environments to investigate the Mini Rover gaits are impractical since simulating interactions between the robot and the millions of granular particles is extremely time-consuming even with the use of high-performance computing tools. Moreover, both the robot and the terrain exhibit complex, nonlinear dynamics, hence deriving closed-form solutions to represent such interactions is not yet possible, hindering model-based simulation approaches. Therefore we rely on laboratory experiments and optimization methods to find and investigate Rover gait patterns.

We discovered that the Mini Rover achieves fast turning by combining multiple terrain manipulation strategies. Mini Rover first creates a yaw motion via differential wheel spinning, followed by a single rear wheel sweeping motion that exploits the multiphase feature of the granular media by generating anisotropic torques. Finally, the rover aligns its roll angle to the changing slope condition by extending and retracting its wheels accordingly. Hence, our ML-inspired new gait outperforms gaits found solely by using Bayesian Optimization by a large margin, turning in 4.5 seconds as opposed to 2 minutes, respectively.

II. MATERIALS AND METHODS

A. Robotic Platform: Mini Rover

The Mini Rover consists of 12 Dynamixel AX-12 actuators and 3D-printed linkages and grousers (Fig. 1B). Each wheel-leg complex can lift up and down as well as spin and sweep. The wheel-leg appendages function in a way that resembles legs for propulsion, but they also spin to disrupt the surrounding terrain as wheels operate. The experimental testbed on which we run our trial is a tiltable bed containing 1 mm diameter poppy seeds (Fig. 1A). The bed size is 2.5 m by 1.2 m, which is large enough to avoid boundary effects during turning. The poppy seeds represent a flowable and deformable terrain, a laboratory analog for flowable and deformable surfaces [18]. The granular terrain is reset to a loosely consolidated state after each experiment via air fluidization, creating a controlled uniform initial state for systematic experiments (Fig. 1A). For robot tracking purposes, we placed three blue spherical markers on top of the rover and three webcams were mounted on the testbed to capture the top and side-view videos. Markers constitute a triangle which is used to track the position and orientation of the rover using MATLAB Computer Vision Toolbox. Control of the rover is achieved by creating open-loop gait trajectories in MATLAB which are tracked internally via Dynamixel actuators' motor drivers and PID controllers.




Fig. 1. Experimental bed and the image of the rover. (A) shows a tiltable and air-fluidized testbed of poppy seeds. (B) shows the Mini Rover on loose granular slop. (C) illustrates the actuation parameters of the Mini Rover.

To understand the locomotor-terrain interaction we performed experiments to measure the torque generated by the wheel moving in poppy seeds. The single-wheel sweep and spin experiments were performed on a platform capable of moving vertically via a Firgelli linear actuator. A 6-axis force sensor from ATI Industries is mounted to the linear actuator. The sweeping and spinning wheel system is mounted to the force sensor to capture the resistive torques induced by the poppy seeds. First, the wheel system is immersed 2 cm into the poppy seeds, and then the substrate is air fluidized, resetting the media in the vicinity of the wheel. Then the wheel started to spin and sweep at low and high rates. The resistive torque data is logged to a desktop computer via a NI-DAQ card and Labview software.

B. BO-RRP Gait to TRRP Gait

To start inducing turning motion with the Mini Rover we modified the Bayesian Optimized Rear Rotator Pedaling (BO-RRP) gait, a previously studied successful gait pattern in [19] (see Fig. 2A). BO-RRP lifts, spins, and sweeps the rear wheel-leg appendages in an alternating style, exploiting the multi-phase feature of granular media by selectively solidifying/fluidizing the mound around the wheels. In the "solidifying phase" the wheel spin halts and the sweeping wheel exerts high pressure that pushes grains backward, creating high reaction force, as shown with dashed black curves in Fig. 2A. In the "fluidizing phase", the wheel spin is activated to create fluidizing shear when resetting the appendage's position, generating a low reaction force, as shown with dashed red curves. Although this strategy significantly increased the rover's locomotion speed, it was incapable of creating maneuvering gaits and hence would only climb straight.

We modified BO-RRP by disabling the selective solidification/fluidization of the rear wheel, i.e. disabling the spinning motion while sweeping the wheel toward the rear of the rover (sweep-out), on the turning side of the Mini Rover (dashed black arrow in Fig. 2B). In addition, we set RRP of both




Fig. 2. Mini Rover gait diagrams. (A) illustrates the snapshots of the rover executing BO-RRP. The white rear wheel sweeps-out 100° (toward the rear of the rover), while the red wheel sweeps-in, resetting the wheel position (toward the front of the rover). The sweep-out direction is shown in black arrows. The dashed arrow indicates that the wheel spinning is disabled during sweep-out. (B) shows the TRRP gait, where both rear wheels perform sweeping-out at the same time while the right rear wheel spinning is disabled. Once sweep-out is completed, both rear wheels spin and sweep-in.

sides to run simultaneously, similar to breaststroke motion. In doing so the modified gait, called Turning RRP (TRRP), created a higher push force on one side and a lower force on the other, yielding a differential thrust and eventually turning in one direction.

C. Bayesian Optimization of TRRP

We designed the new TRRP gait based on our prior experience, however, these changes may cause the parameters inherited from BO-RRP to be suboptimal. Therefore, we run another round of Bayesian optimization to optimize control parameters further, a technique that has demonstrated effectiveness in handling complex robot-terrain interactions [20], [21]. Operating efficiently on systems with low sampling rates, BO treats the objective function as a stochastic function without making any assumptions by assigning a prior probability distribution using a Gaussian process. Over multiple iterations, the outcomes of function evaluations update the priors to create a posterior distribution, continuing until converging to the maximum of the objective function.

The Bayesian optimization is initiated with TRRP so that the gait parameter search starts near a relatively successful gait. The Bayesian optimization is performed for 30 experiment episodes on a 25° granular slope and we omitted the results ending in failure. The open-loop gaits are set to run for 60 cycles taking about 2 minutes. The turning trials are initiated at a body orientation of 0° and the objective function is set to be 90° in the counterclockwise direction to the world frame. We tracked the body orientation, and after each experiment, we updated the objective function to be the achieved turning angle. BO primarily varied the sweep range of the rear wheels, the spin direction, and the speed of all wheels as well as the extension/retraction range of wheel-leg

appendages. The optimization search space is illustrated in Fig. 1C.

III. RESULTS

A. Baseline Gaits

We start maneuvering experiments with the TRRP gait. TRRP achieved gradual turning, taking over 2 minutes to reach 90° of body orientation (see Fig. 5 blue curve). Moreover, the Mini Rover drifted and slid down moderately during turning and failed 2 times out of 7 trials. This is in part due to its inability to manipulate the underlying terrain to create enough turning thrust as well as limitations in its capacity to adapt to the changes in the slope.

B. Optimized TRRP Gait (BO-TRRP)

Next, we fed the TRRP gait parameters into the Bayesian Optimization and evaluated the turning performance of the Mini Rover with iterative experiments, with a random exploration number of 4. Illustrated in Fig. 3A, in the first 10 iterations, the rover turned in the clockwise direction, opposite to the desired orientation, peaking its performance at iteration 4. Despite being in the wrong direction, the Mini Rover achieved 82° turning by discovering the tank-turning maneuver with the front wheels spinning in the reverse direction to each other, a conventional turn-in-place strategy. In doing so, the Mini Rover turns in place via the torque generated by the reciprocal spinning of the front wheels. The Mini Rover started to turn in the right direction at iteration 10. Then, at iteration 14 the optimizer learned the differential spinning (DS) of both the front and rear wheels and turned in the desired orientation as illustrated with the time-lapse images in Fig. 3B. The Mini Rover still performed RRP motion with the rear wheels, however, incorporating DS motion enabled moderate turning. Consequently, at iteration 19 the optimizer learned to perform RRP only with the appendage located at the opposing side of the desired orientation (single RRP) as shown in Fig. 3C, while the remaining wheels kept performing DS motion. The cooperation of both strategies, called BO-TRRP, resulted in almost 90° of turning as shown with the gait diagram in Fig. 3D.

Continued experiments showed little progress in discovering novel gait patterns, only yielding faster DS motion at iteration 29. The plateauing of gait performance is an indication of entrapment into a local minimum, a commonly observed phenomenon in high-dimensional systems. Additionally, although BO provided us with important strategies for rover turning, the process was rather slow taking around 2 minutes.

C. ML-inspired Gait Design

Although the high-dimensional actuation empowers the Mini Rover to manipulate the underlying terrain effectively, the search for optimized gaits in the vast gait space by solely relying on experiments is challenging. Additionally, the hardware experiments are prone to local minima issues due to the lack of governing equations of motion for rover-granular media interactions. To facilitate faster search for




Fig. 3. Evaluation of Bayesian Optimization of turning gaits. (A) depicts the evolution of Mini Rover's body orientation as a function optimization iteration. The circles and crosses represent successful and failed experiments, respectively. (B) and (C) illustrate the tracked images of the Mini Rover's CoM (green curves) and orientation (θ) while turning on 25° slope at iterations 14 and 19, respectively. The optimization iterations 14 and 19 discovered DS and single RRP gaits, respectively. (D) The snapshot images of the rover executing the best-performing optimized gait (BO-TRRP). The straight and curved arrows indicate the wheel's spinning and sweeping directions whereas the straight arrows on the wheels depict the direction of the wheel spinning, respectively. The white and red rear wheels indicate sweeps-out (toward the rover's rear) and sweeps-in (toward the front of the rover), respectively. BO-TRRP starts with the DS motion which the single RRP of the right rear wheel follows.

turning and avoid local minima, we analyzed instances of BO-TRRP gaits. In doing so, we aim to understand how the locomotor-terrain interaction induces turning and design more effective gaits.

Closer investigation of the BO-TRRP revealed that one of the main failures was due to the shifting of the Rover's center of mass (CoM) outside the support polygon as its body orientation changed from one slope to another. Additionally, the Mini Rover's ability to generate reaction forces by pushing against the mound decreased a few seconds after the turning maneuver began. The sweeping leg cannot extend deep enough to reach the mound accumulated around the wheel. This is due to Mini Rover's inability to properly align its roll orientation and CoM position, to the changing slope conditions. Furthermore, in the single RRP motion, the wheel kept spinning during sweeping-out which limits the material resistance and creates low turning torque.

We used these observations and gait cues obtained from

BO-TRRP to update the gait scheme for fast turning. We first increased the DS speed of the wheels to further boost the initial downward tilting of the Mini Rover body. Then, we addressed the Mini Rover's roll angle adaptation issue by extending the wheels that are in the turning direction, facing the downward section of the slope, and retracting the wheels, facing the upward section of the slope. In doing so, the rover can adjust its body during and after the completion of its maneuver. This adjustment also enabled the Mini Rover to push against a larger granular mound and create high resistive force during sweep-out.

To further enhance the turning motion, we varied the spin and sweep speed parameters of the single RRP motion. We first, stopped the wheel spinning during the sweep-out motion and decreased the speed of the sweep, performing a solid manipulation on the substrate. In doing so, the wheel benefited more from the solid-like response of the mound surrounding the wheel by creating more sweeping torque in the direction of turning (Fig. 7A 1). Once the wheel started to sweep-in, we initiated the wheel spinning at a high rate and increased the sweep-in speed, generating a fluid manipulation (Fig. 7A 2)). This high spin rate of the wheel created less torque during sweep-in, hence yielding less torque in the opposing direction to the Mini Rover turning. Combining all these strategies we created the ML-inspired turning gait as illustrated in (Fig. 4A). The table below depicts the control parameters and their associated values used for the ML-inspired gait.

TABLE I
PARAMETERS USED FOR ML-INSPIRED GAIT

Property	Value
Number of gait cycles	2
$\theta_{SweepRange}^L, \theta_{SweepRange}^R$	[0, 90] (deg)
$\theta_{RearLiftMin}^L, \theta_{RearLiftMax}^L$	[-35, -35] (mm)
$\theta_{RearLiftMin}^R, \theta_{RearLiftMax}^R$	[-35, 0] (mm)
$\theta_{FrontLiftHeight}^L, \theta_{RearLiftHeight}^R$	[-35, 35] (mm)
$\omega_{Wheel}^L, \omega_{Wheel}^R$	[-11, 11] (rad/s)




Fig. 4. Mini Rover executing ML-inspired gait. (A) and (B) illustrate the side and top view of the Mini Rover performing fast turning with ML-inspired gait on a 25° slope. Straight and curved arrows depict spin and sweep directions, respectively. Double-arrows illustrate the difference in the appendage leg extension and retraction.




Fig. 5. Time evolution of Mini Rover's orientation using various turning gaits on a slope of 25° . TRRP gait achieves turning 90° gradually in about 2.5 minutes. Optimized gaits start turning rapidly initially, however, their performance degrades over time (orange and green curves). The ML-inspired gait yielded a body turning of almost 90° in just above 4 seconds, outperforming previous gaits by a large margin. Additionally, the ML-inspired gait requires much fewer gait cycles than others, suggesting a reduced energy consumption.

The DS motion first tilts the rover's body toward the desired orientation. Then, DS is accompanied by the single RRP motion inducing more body turning by pedaling the right rear appendage. The continued DS motion helps the rover to agitate the material, reducing the material's resistance to the robot's turning. Finally, the Mini Rover adapts its roll angle to the changing slope condition.

We analyzed the trajectories of the Mini Rover body orientation under different controllers to assess the turning performance of the gaits introduced in this study, as depicted in Fig. 5. While TRRP achieved gradual and slow turning behavior, the optimized gaits initiated a fast turning, however, their performance degraded over time (orange and green curves). The ML-inspired gait yielded a body turning of almost 90° in just above 4 seconds, outperforming previous gaits by a large margin. Additionally, the ML-inspired gait requires much fewer gait cycles than others, suggesting a reduced energy consumption.

We systematically varied the control parameters of the ML-inspired gait to investigate their effects on the rover turning. Specifically, we varied parameters related to the DS and single RRP gaits. Interestingly, the parameters related to the single RRP wheel turned out to be highly coupled and had less effect on the rover turning. On the contrary, we found that the DS speed directly affects the rover's final turning angle as illustrated in Fig. 6. While the experiments with the Mini Rover suggest various gaits that might be




Fig. 6. Mini Rover orientation as a function of differential spin rate. For zero DS speed, the Mini Rover slightly turns in the negative direction.




Fig. 7. Single wheel spin and sweep experiments. (A) Illustration of solid and fluid manipulation of granular terrain with the single wheel-leg system. In (1) the wheel stops spinning while sweeping at a slow speed, creating more resistive torque. In (2) the wheel spins and sweeps at a high rate, leading to less resistive torque. (B) The apparatus used to perform spin and sweep experiments. (C) The time evaluation of torques obtained from the experiments. The blue curve depicts the net torque of a non-spinning wheel sweeping 90° at 1 rad/s . The torque gradually increases and saturates at 1 second. The orange curve represents the wheel spinning and sweeping 90° at 11 rad/s and 6 rad/s , respectively. The wheel spin yields linear-like torque where the onset of sweep generates a peak in the torque.

deployed on a full-size rover, the implementation of these gaits onto a real-world system presents challenges such as scaling to operate in reduced gravity with larger robot inertia as well as locomoting on the different surface material. To this end, Li et al. have shown that for various dry and friction-dominated granular media, the profiles of the vertical and horizontal stresses per unit depth of an intruding object have significant similarities [5]. This might suggest that the proposed strategies in this work might be scaled, at least to some extent, on large-scale rovers on different surfaces. However, further validation through experiments and simulations is required to make an extensive claim.

D. Locomotor-Terrain Interaction

To understand and validate the existence of anisotropic torques created during single RRP, we performed single-

wheel spin and sweep experiments to capture the resistive torques as illustrated with the images in Fig. 7B. The blue curve in Fig. 7C illustrates the experiment where the wheel spins and sweeps at a higher rate, while the orange curve shows a non-spinning wheel and low-rate sweeping. The slow sweeping and non-spinning wheel induces a torque profile that is 2 times larger and 3 times longer than the fast sweeping and spinning wheel. Therefore, controlling the spin and sweep speed of the wheel enables the creation of different torque profiles. The findings in Fig. 7C support our experimental observations with Mini Rover that the wheel creates anisotropic torques by controlling the spin/sweep speeds. Hence, the directional torques facilitate rapid Mini Rover turning in the desired direction.

IV. CONCLUSION

In forthcoming extraterrestrial exploration missions, the ability to reach areas characterized by steep granular slopes and perform various maneuvers will be advantageous, serving both scientific data acquisition and mission navigation efficiency objectives. Nevertheless, current exploratory rover designs favor conventional wheels and passive suspension mechanisms that primarily allow navigation on flat granular surfaces with limited mobility over shallow granular slopes. To this end, by experimenting with a laboratory-scale rover robot capable of applying separate forces to each wheel system, we exploited the idea of manipulating granular terrain to achieve turning at steep granular slopes. This paper combines both machine learning and systematic investigation to obtain effective and robust turning gaits on granular materials. We first performed Bayesian Optimization to discover maneuvering gaits starting with a manually tuned turning gait. The optimization found terrain manipulation strategies giving us deeper insight into how the Mini Rover achieved faster turning. We performed systematic investigations and experiments to fine-tune these strategies further. Hence, combining multiple techniques allowed the rover to turn almost 90° in about 4 seconds on a 25° granular slope. The Mini Rover achieved such fast turning by first differentially spinning its left and right-hand side wheels. Secondly, it creates isotropic torques with a single sweeping wheel that halts spinning during sweeping out, creating a pivoting point for the rover. Consequently, the sweeping wheel spins again during sweeping back in, resetting the appendage position. Finally, the rover adapts to the change of slope that occurs during turning by adjusting its body posture using extendable legs.

The interaction between robots and terrain is crucial, particularly in flowable granular slopes where the terrain dynamics cannot be modeled or simulated accurately. Enhanced comprehension of this interaction through optimization and human intuition using experiments can contribute to the development of improved gaits and robots. The strategies shown in this study could prove valuable in designing maneuvering gaits for large-scale rovers and other exploration robots that are capable of terrain manipulation.

ACKNOWLEDGMENT

This research was funded by a Google LLC PRIME Grant.

REFERENCES

- [1] B. Schäfer, A. Gibbesch, R. Krenn, and B. Rebele, "Planetary rover mobility simulation on soft and uneven terrain," *Vehicle System Dynamics*, vol. 48, no. 1, pp. 149–169, jan 2010.
- [2] V. L. Sharpton, "Outcrops on lunar crater rims: Implications for rim construction mechanisms, ejecta volumes and excavation depths," *Journal of Geophysical Research: Planets*, vol. 119, no. 1, pp. 154–168, 2014.
- [3] P. Mahanti, M. Robinson, T. Thompson, and M. Henriksen, "Small lunar craters at the apollo 16 and 17 landing sites-morphology and degradation," *Icarus*, vol. 299, pp. 475–501, 2018.
- [4] G. Meirion-Griffith and M. Spenko, "A modified pressure–sinkage model for small, rigid wheels on deformable terrains," *Journal of Terramechanics*, vol. 48, no. 2, pp. 149–155, 2011.
- [5] C. Li, T. Zhang, and D. I. Goldman, "A terradynamics of legged locomotion on granular media," *Science*, vol. 339, no. 6126, pp. 1408–1412, 2013.
- [6] A. Thoesen and H. Marvi, "Planetary surface mobility and exploration: A review," *Current Robotics Reports*, vol. 2, no. 3, pp. 239–249, 2021.
- [7] T. d. J. M. Sanguino, "50 years of rovers for planetary exploration: A retrospective review for future directions," *Robotics and Autonomous Systems*, vol. 94, pp. 172–185, 2017.
- [8] C. Weibel, G. Valsecchi, H. Kolenbach, and M. Hutter, "Towards legged locomotion on steep planetary terrain," in *36th IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2023)*, 2023.
- [9] D. M. Helmick, Y. Cheng, D. S. Clouse, L. H. Matthies, and S. I. Roumeliotis, "Path following using visual odometry for a mars rover in high-slip environments," in *2004 IEEE Aerospace Conference Proceedings (IEEE Cat. No. 04TH8720)*, vol. 2. IEEE, 2004, pp. 772–789.
- [10] F. Zhou, R. E. Arvidson, K. Bennett, B. Trease, R. Lindemann, P. Bellutta, K. Iagnemma, and C. Senatore, "Simulations of mars rover traverses," *Journal of Field Robotics*, vol. 31, no. 1, pp. 141–160, 2014.
- [11] P. Arm, G. Waibel, J. Preisig, T. Tuna, R. Zhou, V. Bickel, G. Ligeza, T. Miki, F. Kehl, H. Kolenbach *et al.*, "Scientific exploration of challenging planetary analog environments with a team of legged robots," *Science robotics*, vol. 8, no. 80, p. eade9548, 2023.
- [12] H. Inotsume, C. Creager, D. Wettergreen, and W. Whittaker, "Finding routes for efficient and successful slope ascent for exploration rovers," in *The international symposium on artificial intelligence, robotics and automation in space (i-SAIRAS)*, 2016.
- [13] G. Sakayori and G. Ishigami, "Energy efficient slope traversability planning for mobile robot in loose soil," in *2017 IEEE International Conference on Mechatronics (ICM)*. IEEE, 2017, pp. 99–104.
- [14] J. R. Sanchez-Ibanez, C. J. Perez-del Pulgar, and A. Garcia-Cerezo, "Path planning for autonomous mobile robots: A review," *Sensors*, vol. 21, no. 23, p. 7898, 2021.
- [15] S. Lyu, W. Zhang, C. Yao, Z. Zhu, and Z. Jia, "Modeling and analysis of a reconfigurable rover for improved traversing over soft sloped terrains," *Biomimetics*, vol. 8, no. 1, p. 131, 2023.
- [16] C. Cao, D. Moon, C. Creager, D. K. Lieu, and H. S. Stuart, "Push-pull locomotion: Increasing travel velocity in loose regolith via induced wheel slip," *Journal of Terramechanics*, vol. 110, pp. 87–99, 2023.
- [17] N. Phornpimonchoke, S. Koosermit, A. Tanakijchumroon, and R. Chaichaowarat, "Alternative locomotion modalities for lunar rover," in *2023 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*. IEEE, 2023, pp. 287–292.
- [18] S. Shrivastava, A. Karsai, Y. O. Aydin, R. Pettinger, W. Bluethmann, R. O. Ambrose, and D. I. Goldman, "Material remodeling and unconventional gaits facilitate locomotion of a robophysical rover over granular terrain," *Science Robotics*, vol. 5, no. 42, p. eaba3499, 2020.
- [19] A. Karsai, D. Kerimoglu, D. Soto, S. Ha, T. Zhang, and D. I. Goldman, "Real-time remodeling of granular terrain for robot locomotion," *Advanced Intelligent Systems*, vol. 4, no. 12, p. 2200119, 2022.
- [20] S. Ha, P. Xu, Z. Tan, S. Levine, and J. Tan, "Learning to walk in the real world with minimal human effort," *arXiv preprint arXiv:2002.08550*, 2020.
- [21] W. Yu, J. Tan, Y. Bai, E. Coumans, and S. Ha, "Learning fast adaptation with meta strategy optimization," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2950–2957, 2020.