



Towards Legged Locomotion on Steep Planetary Terrain

Conference Paper**Author(s):**

Weibel, Cedric; Valsecchi, Giorgio; [Kolvenbach, Hendrik](#) ; [Hutter, Marco](#) 

Publication date:

2023

Permanent link:

<https://doi.org/10.3929/ethz-b-000625001>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/IROS55552.2023.10341665>

Towards Legged Locomotion on Steep Planetary Terrain

Giorgio Valsecchi*, Cedric Weibel*, Hendrik Kolvenbach, and Marco Hutter

Abstract—Scientific exploration of planetary bodies is an activity well-suited for robots. Unfortunately, the regions that are richer in potential discoveries, such as impact craters, caves, and volcanic terraces, are hard to access with wheeled robots. Recent advances in legged-based approaches have shown the potential of the technology to overcome difficult terrains such as slopes and slippery surfaces. In this work, we focus on locomotion for sandy slopes, comparing standard walking policies with a novel crawling-based gait for quadrupedal robots. We fine-tuned a state-of-the-art locomotion framework and introduced hardware modifications to the robot ANYmal, which enables walking on its knees. Moreover, we integrated a novel metric for stability, the stability margin, in the training process to increase robustness in such conditions. We benchmarked the locomotion policies in simulation and in real-world experiments on a martian soil simulant. Our results show a significant improvement in terms of robustness and stability, especially at higher slope angles beyond 15 degrees.

I. INTRODUCTION

Planetary exploration rovers have been successfully deployed on neighboring celestial bodies. Until this day, and except for the recent Ingenuity helicopter [1], solely wheeled locomotion was used to traverse the terrain. Sojourner, Spirit, Opportunity, Curiosity, and Perseverance [2] are well known examples of such rovers. The success of these missions showed the suitability of wheeled locomotion for the most approachable terrains. However, despite considerable development targeting optimal wheel design [3], [4], some systems, such as the Spirit rover, reached the end of the mission due to getting stuck in the soil [5]. In fact, the majority of the rovers experienced excessive sinkage and slip, significantly impacting the operational timeline [6].

In the last decade, the scientific community is increasingly interested in high-risk, high-reward regions such as craters, lava tubes, volcanic regions, and similar morphologies that expose scientifically valuable features, yet lie in hard-to-reach terrain. In fact, several regions of interest (ROI) have been identified for future missions that will require the robot to conquer steep environments of more than 15° [7], [8], [9], [10].

Adding to the issue, a planetary exploration mission might consist of several phases with conflicting requirements on the locomotion system. In addition to the steep slopes of the region of interest, the mission might require traversing long stretches on a flat and relatively easy terrain, for example,

The authors are with the Robotic Systems Laboratory, ETH Zurich, Zurich 8092, Switzerland (e-mail: vgiorgio@ethz.ch; ceweibel@ethz.ch; hendrik@ethz.ch; mahutter@ethz.ch)

This work was supported by ESA Contract Number 4000131516/20/NL/MH/ic and 4000135310/21/NL/PA/pt.

*These authors contributed equally to this work



(a) Front view.

(b) Side view.

Fig. 1: Snapshot from the robot walking up a 25° slope on stone slabs covered with a fine layer of sand. The wheels are rigidly connected to the robot's shin and not driving, used as knee contact surfaces.

in the vicinity of the lander. The same locomotion system should therefore maximize velocity and efficiency in one situation and safety in another.

The climbing task provides different environmental challenges: Unavoidable mid- to large-sized boulders (terrain morphology), steep inclination, and rocky terrain coated with a fine layer of sand drastically aggravating robust movement (lack of grip), as seen in figure 1. Therefore, the robot must emphasize especially on robust and stable locomotion to ensure safety at all times.

Legged systems, both artificial and biological, offer the versatility to adapt their gaits based on the environment. On challenging terrains, such as on slippery surfaces or extreme slopes, a safe gait with the capability to fall back to a stable static stance is preferred whereas, on flat and easy terrains, a dynamic gait achieves higher speeds with relatively low power consumption. Despite their relatively low technology readiness level (TRL) [11] of legged robots, the adaptability of such systems makes them potentially good candidates for future missions. However, the safe locomotion to overcome steep and slippery slopes is still an open research question. This work presents a novel locomotion strategy for steep terrain and testing in comparison to several state-of-the-art approaches on relevant planetary terrain.

A. Related Work

Previous works already addressed some of the challenges of terrains encountered in planetary exploration [12], [13]. SpaceBok [14], is a highly dynamic quadruped with specially designed feet to locomote on granular terrain [15]. SpaceClimber [16] and Lauron V [17] are hexapod robots optimized to overcome steep inclines and large obstacles. Another example is Capuchin [18], a robot designed to climb

vertical walls using grooves without using active grippers. Lemur 3 [19] uses micro spines to generate gripping force and gave an impressive demonstration of its capability of climbing vertical rock walls. However, all previously mentioned robots lack the ability to adapt to a wide range of environments. Each robot is designed with a specific use case and terrain in mind.

Recent years saw the development and commercialization of general-purpose quadrupedal robots, such as Spot [20], AlienGo [21] and ANYmal [22]. Some of them, like ANYmal, are designed for challenging outdoor environments. They offer a wide range of manipulability and range of motion with their legs and are capable of dynamically stable locomotion. Some of the works using ANYmal integrated sensorized feet and optimized locomotion strategies for challenging terrains [23], [24]. Machine learning techniques are capable of producing controller architecture extremely robust against a variety of terrains [25], [26]. At the same time, the simulators used in the training process are becoming more and more capable of representing complex terrains efficiently, accelerating the design of new locomotion controllers with GPUs to minutes instead of hours [27], [28]. Specific modules can also be integrated into the simulation environment to model increasingly complex terrain, including the terramechanics [29]. The flexibility and performance of such approaches open the possibility to try out many experimental reward terms and design new locomotion strategies.

From a more theoretical point of view, other works explored the aims to theoretically quantify stability at each time instance based on how close the robot is to tumbling over [30]. A subsequent work [31] compared two different leg configurations for quadrupedal robots based on the prior mentioned stability criterion with promising results.

B. Contributions

This work proposes a novel locomotion concept, with both hardware and software features, meant to decrease mission-endangering tumbling and slipping on slopes. We optimize the gait for the aforementioned environments, by training a control policy for the ANYmal robot with a setup similar to that shown by Rudin et al. in [27], using the physics simulator IsaacGym [28]. We use custom rewards for promoting the stability of the robot and introduced hardware modifications based on the insights gained in simulations. We validate our approach experimentally on steep slopes of up to 25° composed of martian soil simulant and rocks, and benchmark against baseline locomotion policies on ANYmal. Moreover, we tested a hybrid wheeled-leg variant in the same conditions. These tests provide an overview of possible failure modalities and challenges of different locomotion concepts.

II. METHODS

We approach the problem of legged locomotion on steep planetary terrain by advancing an existing platform and control architecture to the specific environment. We selected ANYmal and the reinforcement learning pipeline introduced




Fig. 2: Stability polyhedron for a four contact point stance. The contact points are numbered and the lines connecting them are the tumbling axes. Contacts are marked with a dark red circle and the center of gravity as a yellow dot.

in [27] for their versatility, robustness, and abundance of previous results. In particular, we focused on the following points of improvement:

- Introduction of an explicit reward function term to increase the stability margin
- Optimization of the training process for the specific environment
- Introduction of additional contact surfaces in simulation and hardware

The following subsections describe in detail why and how we addressed these points.

A. Stability margin

In our previous works, locomotion failure is penalized with a negative reward in case of a collision with the ground. No explicit terms are used to further quantify and increase the stability margin. If a greater stability margin is to be achieved, a metric for stability has to be introduced. Many definitions of stability are possible, according to the failure mode considered. For the case of the steep slope, we consider tumbling as the worst possible failure, being it a more traumatic and energetic event than other failures such as slippage or base-to-ground collision.

From the literature, two metrics are available to quantify tumble stability: Gravito-Inertial Inclination Margin (GIIM) and Gravito-Inertial Acceleration Margin (GIAM) [30]. Both make use of the so-called Gravito-Inertial Acceleration vector (GIA) \mathbf{a}_{gi} . A vector combining gravitational acceleration \mathbf{g} and total robot acceleration \mathbf{a}_{total} . The total robot acceleration is defined as the sum of all body part impulses divided by the total mass of the robot, as defined in equation 1.

$$\mathbf{a}_{gi} = \mathbf{g} - \mathbf{a}_{total} = \mathbf{g} - \frac{1}{\sum_{i=0}^N m_i} \sum_{i=0}^N m_i \mathbf{a}_i \quad (1)$$

In order to assess imminent tumbling, the metrics make use of a stability polyhedron, defined as a pyramid with the vertices at the robot CoG and all contact points with the ground. A possible tumbling axis is defined as the line between subsequent body parts contacting the ground. Each pyramid's face defines a normal vector \mathbf{n}_{gab} pointing outwards. The vertices $\mathbf{p}_{a/b}$ indicate the location of the robot-to-ground contact. Figure 2 shows the edges of the stability polyhedron as green dotted lines. In contrast to previous work [30], [31] we do not consider gripping forces, external

forces, or external moments. A more in-depth derivation can be found in [30].

The robot is in a stable configuration w.r.t tumbling if the GIA vector points inside the prior defined stability polyhedron.

The first metric, GIIM, defined by equation 2 computes the smallest angle between the GIA vector \mathbf{a}_{gi} and the normal of the polyhedron face \mathbf{n}_{gab} . The value is normalized with $-\frac{\pi}{2}$ to ensure negative angle if the GIA vector points out of the polyhedron.

$$\theta_{margin} = \min_{faces} \left[\arccos\left(\frac{\mathbf{n}_{gab} \cdot \mathbf{a}_{gi}}{|\mathbf{n}_{gab}| \cdot |\mathbf{a}_{gi}|}\right) \right] - \frac{\pi}{2} \quad (2)$$

The second metric, GIAM, defined by equation 3, represents the maximum acceleration the robot could produce without tumbling over.

$$a_{margin} = \min_{faces} \left[-\frac{\mathbf{n}_{gab} \cdot \mathbf{a}_{gi}}{|\mathbf{n}_{gab}|} \right] \quad (3)$$

B. Learning locomotion on sloped terrain

We developed our locomotion policy based on the framework presented in [27]. The following paragraphs explain in detail how we integrated the stability margin in the reward function, and how we modeled the environment and the robot in the simulation.

1) *Reward Design:* Both equations 2 and 3 are not suitable for training, due to only incorporating the most critical angle between a face and the GIA vector. Furthermore, GIAM encourages excessive actuator torques. We utilize a slightly modified version of the proposed GIIM metric in the training. We sum over all angles between the GIA vector and the polyhedron faces, as stated in equation 4.

$$\theta_{reward} = \sum_{faces} \left[\arccos\left(\frac{\mathbf{n}_{gab} \cdot \mathbf{a}_{gi}}{|\mathbf{n}_{gab}| \cdot |\mathbf{a}_{gi}|}\right) - \frac{\pi}{2} \right] \quad (4)$$

With this modification, we get a smoother value progression and maximize all angles and not only the most critical one. This improves the convergence of the solver.

The GIIM reward term alone will not produce satisfactory locomotion of the robot and is not used as termination criteria. Apart from rewards for correct tracking of the controlled linear and angular velocity, contact with critical body parts are penalized, specifically the base and actuator contact body as illustrated in figure 3. Moreover, we use an acceleration penalty in the z-direction of the robot base frame on each foot to discourage the policy from impacting the ground with high velocities, which in return would result in high impact forces at the knee-to-shin interface. In order to ensure an efficient and smooth gait, changes in applied torque and joint acceleration of each actuator are penalized.

2) *Environment Design:* The robot is trained on three different terrains: sloped crater, linear slope, and exponential slope.

Additionally, we chose to use the game-inspired curriculum approach, which has been proven to be effective [27]. The difficulty is increased for an agent if it reaches its




Fig. 3: Modified ANYmal collision bounding boxes. Contacts to the knee and foot are tolerated and contacts to the actuators and the base are penalized.




Fig. 4: The three terrain environments during training (crater, linear slope, exponential slope) with surface roughness.

domain boundaries, set to a maximum of $\pm 30^\circ$ inclination. Our solution differentiates between a linearly and exponential slope, in order to have more variation on the underground inclination during an episode of training. Moreover, the difficulty is tied to the overall roughness of the terrain to encourage the policy to perform higher steps without slipping or tripping.

Despite previous works [32], [33] extraterrestrial terrains are not well known and hard to model, especially the interaction with dry, granular matter. We, therefore, use a wide range of randomized friction coefficients $\mu_s \in [0.3, 1.25]$ to simulate the interaction of a rigid body of the robot with the wide spectrum of terrain conditions. However, this approach cannot address terrain deformation and more complex terramechanic behaviors such as sliding or triggering an avalanche. In order to take into account disturbances in the form of dimensional uncertainties and external forces acting on the actor, we randomly alter the mass of the 7.5 kg base by $\pm[0, 5]$ kg.

3) *Agent Design:* Each agent interacts with the environment with a simplified geometry for fast collision detection during training of the policy. Contact with the base and the actuators are penalized. The collision model, as seen in figure 3, is based on the geometry after designing the new hardware. Due to requirements given by the physics solver, we model the shape of the knee wheel as a multitude of smaller spheres. This approach proved to be sufficient and fast to simulate.

C. Additional contact surfaces

In order to lower the CoG and therefore increase the stability margin, a new hardware component is required. We designed a knee part, which, coupled with a non-movable




Fig. 5: Modified leg, the red and green shading show which components are rigidly connected. The shin has been tilted seven degrees outwards to enable the full range of motion of the knee joint.

wheel, adds the required contact surface for the robot’s crawling motion. For the sake of simplifying the design process, off-the-shelf and in-house available components, such as the thigh, wheel, and shin with attached foot and shin clamps have been used. The knee, shin, and wheel form a rigid body and are fixed to the knee actuator, which can be seen marked on the left image of the figure 5. The parts are 3D printed out of nylon (SLS technology) and we optimized the topology of the parts in order to maximize the strength-to-weight ratio. The load on the bodies has been recorded from previous runs of the robot walking on its feet. We verified the design strength against multiple corner cases and load scenarios.

III. EXPERIMENTS AND RESULTS

To compare the performance of locomotion concepts (legged, wheeled, crawling), we used both simulations and real-world experiments. The computational infrastructure used for the training process can be leveraged to generate a large amount of data and draw conclusions from a statistical point of view. Moreover, all data are available and experiments are perfectly reproducible. Nevertheless, simulations cannot fully capture the interaction between the real robot and the environment. To complement the result of the simulations, we also performed tests in lab conditions meant to replicate an actual planetary terrain.

A. Simulation Results

1) *Crawling policy*: The emerging policy makes use of the knee to increase its stability and lower its center of gravity. The robot lowers the center of gravity and hence maximizes the possible angle between the GIA vector and the stability polyhedron faces. The movement resembles a crawling motion, which intuitively improves stability in terms of tumbling.

Figure 6 shows a typical walking pattern of a robot using the new policy, crawling up a 20-degree slope. The robot utilizes both feet and knees to contact the ground. The run shows the typical diagonal contact pattern in dynamic walking. During each contact switch, the minimum GIIM value shortly drops but remains positive.

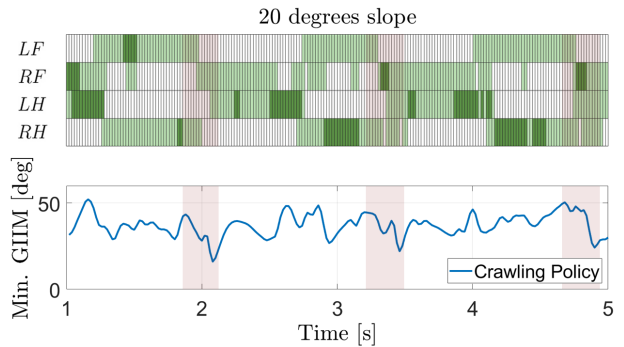


Fig. 6: Contact sequence on a simulated slope. The rows of the contact map show the number of contact points with the ground per leg. Dark green is foot and knee, light green is either foot or knee, and white indicates no contact. The bottom plot run illustrates the GIIM progression on the slope. The areas in shaded red highlight the contact switching phases, during which the stability margin decreases.




Fig. 7: Plot showing the average GIIM reward for 128 simulated robots per fixed slope. The histogram represents the ratio of robots that contacted the ground with the base contact body during the simulated test.

2) *Stability*: We compared the new policy to two baseline policies. One is our in-house baseline policy, named “General Walking Policy”, as described in [27]. The second one is identical to the “General Walking Policy”, but trained on our new terrain model introduced in II-B, named “Optimized Walking Policy”, in order to have a more objective comparison and assess the influence of the terrain environment.

We compute the metric by averaging over four seconds of recording time and 128 robots simulated in parallel. Each individual robot is placed on a linear slope with a fixed inclination, given a random initial pose and random linear and angular control velocity between 0 and 0.75 m/s.

The mean GIIM per second and robot for the new policy is initially around 12° for low inclination and steadily decreases for steeper slopes, as seen in figure 7.

There is a positive correlation between the GIIM and the success rate of each robot not critically failing, especially on steeper slopes. We defined failure as a collision of the base of the robot with the ground during the recorded time period, represented by bars in figure 7. The robots running the baseline policies nearly all fail from 35° inclination by

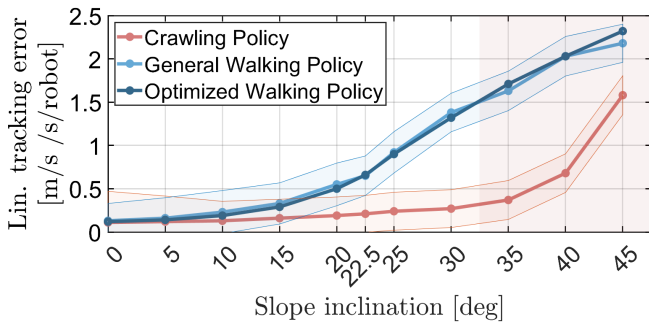


Fig. 8: Plot showing the average linear tracking error reward as a function of the slope angle in simulation.



(a) The hybrid wheeled-legged version of ANYmal, with its wheels slipping. (b) The standard legged ANYmal critically slipping and involuntarily contacting the ground with the knees.

Fig. 9: Two robot configurations on a 20.0° slope tested at Beyond Gravity.

tumbling and stumbling over grooves in the terrain. The new policy does not completely inhibit stumbling, but it reduces the failure rate.

3) *Tracking performance*: An important metric for comparing the policies is the average linear tracking error of the controlled direction per robot and second. As seen in the plot in figure 8, the robots with both walking policies have to significantly deviate from the controlled path to prevent falling over.

However, the crawling policy experience more slippage. Using the knee reduces the range of motion and the manipulability, due to the smaller distance from the knee actuator to the respective contact point. This, combined with the fact that the policy can use both contact points simultaneously per leg, can result in an over-constrained pose resulting in a more pronounced slippage for the crawling policy.

B. Hardware results

In order to validate our contribution in the real world, we tested the crawling gait with the new knee at a locomotion test facility at Beyond Gravity, located in Zurich (Figure 1). The testbed consists of an inclinable, 6m by 6m container which is filled with ES-3 soil simulant and flat stone slabs. ES-3 is a fine-grained sand, analog to that found on the planet Mars, which forms landscape structures such as dunes and sand ripples [15]. The testbed can be inclined up to 25° , a value determined by the internal friction of the soil.

		Inclination			
		0°	10°	20°	25°
ANYmal	Rocks	●	●	◐	◑
	Sand	●	●	◐	◑
ANYmal on wheels	Rocks	●	●	◐	◑
	Sand	●	●	◐	n/a
ANYmal on knees	Rocks	●	●	●	◐
	Sand	●	●	●	◐

● = High safety, ◐ = Non-traversable, n/a = not tested

TABLE I: Qualitative evaluation of robots' locomotion performance. The pie chart indicates the stability of the walking and crawling gaits. The assessment was made based on observation and the number of successful trials.

We tested with inclination of 0° , 10° , 20° and 25° . Each policy has been tested on both, pure sand and stone slabs covered with a fine layer of sand. Each run was performed at a fixed slope angle, with the robot moving in the direction of maximum inclination in upward and downward direction.

We conducted the test with a stock version of ANYmal, with the two baseline policies as described, and the one with the newly developed knees and the crawling gait. To have a more comprehensive comparison, we also tested ANYmal's hybrid wheeled-legged variant, ANYmal on wheels (AoW). AoW is capable of driving on wheels and simultaneously changing the leg pose to adapt to the terrain conditions [34]. Figure 9 shows both ANYmal and AoW during critical moments of the tests.

From a qualitative point of view, all policies were only partially capable of walking or crawling on the slopes, differing only in their failure modalities and success rates. The crawling policy behaves similarly to the simulations. It performed a stable gait on slopes up to 25° , the maximum of the test bed. The two baseline policies could be tested up to 20° , with the optimized one failing without producing a single successful run. In contrast, the crawling robot could rescue itself to a stable stance and not tumble over, even when slipping, nevertheless partially failing when the robot showed critical movement and leg positioning, which resulted in necessary intervention from the operator to prevent catastrophic tumbling. Table I summarizes the qualitative outcome of the tests performed in terms of success rate and overall stability.

The three plots of figure 10 show the minimum GIIM value calculated during a time frame of 2.7 seconds, a fraction of the time the robot took to crawl or walk up the slope. The respective sample sections have been shifted in time to align with the most critical GIIM value for comparison, highlighted by a gray band. A GIIM value of below zero indicates a high probability of tumbling, highlighted as a red area on the plot. The crawling policy incorporates most of the time a minimum GIIM above zero degrees for all inclined slopes. Hence, the robot walks with a stable gait in terms of tumbling. In comparison, both walking policies have brief spikes below zero degrees, which could result in tumbling. Figure 11 gives a more comprehensive representation of the tests.




Fig. 10: Comparison between the tested policies at different slope angles. The lowest values of the GIIM have been aligned.




Fig. 11: Box plot with outliers of each run performed on the slope with fixed inclinations of 0, 10, and 20 degrees.

The box plot synthesizes the outcome of all runs, showing the average GIIM margin for the different locomotion policies and slope angles. Overall, it is visible how the increase in the slope angle correlates with a decrease in the stability margin. Most importantly, a trend of the minima and of the outliers is also visible. The crawling policy demonstrates only a few values below the zero-degree mark compared to the general and optimized walking policy. Only at a 20° slope do all walking patterns show critical instances, most notably the general walking policy with nine values below the -60° minimum GIIM, as seen in figure 10.

As mentioned previously, different policies exhibited different failure modalities over different terrains. The stock ANYmal version with the two baseline policies was particularly affected by slipping on rock slabs, digging itself in the sand, jamming its feet (to the point of overloading its actuators and battery, which resulted in a complete shut-down), and stumbling on the rock edges, as illustrated in figure 12b. AoW is also affected by slipping on the rock slabs, resulting in wasted energy. On the sand, AoW is affected by the same issues affecting normal wheeled rovers. Our new locomotion concept is less affected by slippage, but the knee contact results in a lower degree of manipulability and range of motion, making it more sensitive to rock edges, which can block the motion. Figure 12a shows the effects of slippage and figure 12b shows how the sand-covered edge constitutes a very difficult and hard-to-detect obstacle even when using perception.

Root causes of the failures could be identified in the following points, which are partially shared between the



(a) ANYmal feet heavily damaged by slippage after a few minutes walking on the slopes. (b) Robot stumbles at the boundary between sand and stone slabs highlighted in yellow.

Fig. 12: Typical damage and failure modes observed in the tested environment.

different locomotion concepts.

- The lack of perception and visual information makes obstacles such as those of figure 12b particularly challenging.
- The lack of accurate terrain models in simulation results in challenges on the sand.
- The contact surfaces are not optimized for neither the rock nor the sand.
- The kinematic of the legs is sub-optimal for the climbing task, particularly for the crawling gait.

IV. CONCLUSIONS

In this work, we presented a crawling gait for legged robots, meant to investigate locomotion on steep, granular slopes relevant to planetary exploration. We trained the new locomotion policy with a reinforcement-learning-based algorithm in simulation, introducing a special stability reward

term and allowing parts of the robot different from the feet to contact the ground, most notably the knees. The new walking policy was deployed on the quadrupedal robot ANYmal, which has been upgraded to provide contact surfaces on its knees. Finally, we compared the performance of the new concept with those of the standard ANYmal robot with baseline policies in both simulation and real-world experiments in a relevant testing set-up. Moreover, we also tested ANYmal's hybrid wheeled-legged variant.

The tests show that the crawling locomotion successfully improves the robustness of the robot on steep and slippery slopes, although marginally. The robot's probability of tumbling and slipping on slopes up to 25° is lower than for standard solutions, however, it does not perform equally well on all possible terrain conditions and obstacles.

Future works should address the failure points identified in Section III. Integration of sensors and perception pipelines would give the robot a better understanding of the terrain conditions and morphology. Integration of representative terrain models in the simulation, as in [29], would potentially improve performance on the sand. Dedicated terrain simulators could be used to benchmark the impact of terrain models used in training on the final policy performance. The development of specialized contact surfaces for different terrain would influence the overall performance of the robot, as shown in [15]. Finally, the kinematic structure of the leg could be optimized for the set of tasks relevant to the mission, particularly climbing [31]. Moreover, additional work could investigate further the impact of the GIIM on overall performance.

Considering those points, it becomes apparent that robust locomotion for planetary exploration requires additional work. In particular, two challenges should be addressed with research at the system level. Firstly, a long-distance and high-efficiency locomotion strategy should be combined with a strategy for maximizing stability and safety on challenging slopes. Secondly, the robot should be equipped with a multitude of contact surfaces, specialized for different terrains, and a kinematic structure capable of controlling which surfaces dominate the interaction with the ground at any specific moment. This work showed on a conceptual level how a single system could be reconfigured to address this broad range of tasks, however, further research is needed to develop a single solution specifically meant for multi-modal locomotion and optimized for the terrains of a specific mission.

ACKNOWLEDGMENT

The authors want to thank Philipp Oettershagen and his team at Beyond Gravity for supporting and providing the testing facility. This work has been conducted as part of ANYmal Research, a community to advance legged robotics.

REFERENCES

- [1] J. Balaram *et al.*, "The ingenuity helicopter on the perseverance rover," *Space Science Reviews*, 2021.
- [2] N. S. Place. (2021) The mars rovers. [Online]. Available: <https://spaceplace.nasa.gov/mars-rovers/en/>
- [3] J. B. Johnson *et al.*, "Discrete element method simulations of mars exploration rover wheel performance," *Journal of Terramechanics*, vol. 62, pp. 31–40, 2015.
- [4] J. Zheng *et al.*, "Design and terramechanics analysis of a mars rover utilising active suspension," *Mechanism and Machine Theory*, vol. 128, pp. 125–149, 2018.
- [5] G. Webster *et al.* (2009, Dec.) NASA's Mars Rover has Uncertain Future as Sixth Anniversary Nears.
- [6] R. Gonzalez *et al.*, "Slippage estimation and compensation for planetary exploration rovers. state of the art and future challenges," *Journal of Field Robotics*, vol. 35, no. 4, pp. 564–577, 2018.
- [7] A. Seeni *et al.*, "Robot mobility systems for planetary surface exploration: state-of-the-art and future outlook: a literature survey," *Aerospace Technologies Advancements*, 2010.
- [8] N. Potts *et al.*, "Robotic traverse and sample return strategies for a lunar farside mission to the Schrödinger basin," *Adv. Space Res.*, vol. 55, 2015.
- [9] E. S. Steenstra *et al.*, "Analyses of robotic traverses and sample sites in the schrödinger basin for the heracles human-assisted sample return mission concept," *Advances in Space Research*, 2016.
- [10] E. Czaplinski *et al.*, "Human-assisted sample return mission at the schrödinger basin, lunar far side, using a new geologic map and rover traverses," *The Planetary Science Journal*, vol. 2, no. 2, 03 2021.
- [11] H. Kolvenbach *et al.*, "Scalability Analysis of Legged Robots for Space Exploration," in *International Astronautical Congress (IAC)*. IAF, Sept. 2017.
- [12] C. D. Remy *et al.*, "Walking and crawling with alof: a robot for autonomous locomotion on four legs," *Industrial Robot: An International Journal*, 2011.
- [13] H. Kolvenbach, "Quadrupedal robots for planetary exploration," Ph.D. dissertation, ETH Zurich, IRIS, 2021.
- [14] P. Arm *et al.*, "Spacebok: A dynamic legged robot for space exploration," in *ICRA*, 2019.
- [15] H. Kolvenbach *et al.*, "Traversing steep and granular martian analog slopes with a dynamic quadrupedal robot," *Field Robotics*, vol. 2, 2022.
- [16] S. Bartsch *et al.*, "Development of the six-legged walking and climbing robot spaceclimber," *Journal of Field Robotics*, 2012.
- [17] A. Roennau *et al.*, "Lauron v: A versatile six-legged walking robot with advanced maneuverability," in *2014 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, 2014.
- [18] R. Zhang and J.-C. Latombe, "Capuchin: A free-climbing robot," *International Journal of Advanced Robotic Systems*, 2013.
- [19] A. Parness *et al.*, "Lemur 3: A limbed climbing robot for extreme terrain mobility in space," in *2017 IEEE international conference on robotics and automation (ICRA)*, 2017.
- [20] "Spot." [Online]. Available: <https://www.bostondynamics.com/spot>
- [21] "Aliengo." [Online]. Available: <https://www.unitree.com/>
- [22] M. Hutter *et al.*, "ANYmal - toward legged robots for harsh environments," *Advanced Robotics*, vol. 31, no. 17, pp. 918–931, 2017.
- [23] G. Valsecchi *et al.*, "Quadrupedal locomotion on uneven terrain with sensorized feet," *IEEE RAL*, vol. 5, no. 2, pp. 1548–1555, 2020.
- [24] M. G. Catalano *et al.*, "Adaptive feet for quadrupedal walkers," *IEEE Transactions on Robotics*, vol. 38, no. 1, pp. 302–316, 2022.
- [25] J. Lee *et al.*, "Learning quadrupedal locomotion over challenging terrain," *Science Robotics*, vol. 5, no. 47, p. eabc5986, 2020.
- [26] T. Miki *et al.*, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, vol. 7, no. 62, p. eabk2822, 2022.
- [27] N. Rudin *et al.*, "Learning to walk in minutes using massively parallel deep reinforcement learning," 2021.
- [28] V. Makoviychuk *et al.*, "Isaac gym: High performance gpu-based physics simulation for robot learning," *CoRR*, 2021.
- [29] S. Choi *et al.*, "Learning quadrupedal locomotion on deformable terrain," *Science Robotics*, vol. 8, no. 74, p. eade2256, 2023.
- [30] W. Ribeiro *et al.*, "Dynamic equilibrium of climbing robots based on stability polyhedron for gravito-inertial acceleration," in *CLAWAR*, 2020.
- [31] K. Uno *et al.*, "Simulation-based climbing capability analysis for quadrupedal robots," in *Robotics for Sustainable Future*. Springer International Publishing, 2022, pp. 179–191.
- [32] F. Zhou *et al.*, "Simulations of mars rover traverses," *Journal of Field Robotics*, vol. 31, no. 1, pp. 141–160, 2014.
- [33] D. Bickler. "Roving Over Mars," *Mechanical Engineering*, 1998.
- [34] M. Bjelonic *et al.*, "Whole-body mpc and online gait sequence generation for wheeled-legged robots," in *IROS*, 2021, pp. 8388–8395.