

# Terramechanics-Based Mobility Failure Compensation and Soil Manipulation

Catherine A. Pavlov<sup>1</sup>, Arno Rogg<sup>2</sup>,  
and Aaron M. Johnson<sup>3</sup>, Senior Member, IEEE

<sup>1</sup>Mechanical Engineering Department, Carnegie Mellon University, Pittsburgh, PA 15213 USA

<sup>2</sup>NASA Ames Research Center, CA 94035 USA

<sup>3</sup>Mechanical Engineering Department, Carnegie Mellon University, Pittsburgh, PA 15213 USA

Corresponding author: Catherine A. Pavlov (email: cpavlov@alumni.cmu.edu).

This work was supported by a NASA Space Technology Research Fellowship under grant 80NSSC18K1154

**ABSTRACT** In this paper, we enable new mobility and manipulation modes for wheeled planetary exploration rovers through the use of terramechanics modeling and field experiments. Useful modes of wheel-based soil manipulation and examples of rovers driving with degraded mobility systems are first demonstrated in lunar and Martian analog environments. We show a full-scale rover use its wheels to dig trenches up to 10.6 cm deep, dig holes to estimate soil characteristics, and modify terrain to make it accessible to a smaller robot. We also measure the impact of actuator failure on a rover in lunar simulant. Here, we show the slip doubled on moderate slopes for a damaged drive motor, which would exceed the rover’s operational limits for slip, motivating the need for driving strategies that mitigate mobility loss. We then develop an optimization framework which uses a recently developed terramechanics model to automatically generate both open and closed-loop driving strategies for planetary rovers performing terrain manipulation or operating in a degraded state with no need for hand tuning of behaviors. Finally, we demonstrate the generated driving strategies for soil manipulation and mobility compensation on a rover in a controlled lab setting, where we show that 1) mobility is maintained while manipulating soil; and 2) mobility is regained while experiencing failure of steer and drive actuators.

**INDEX TERMS** Failure Detection and Recovery, Motion Control, Space Robotics, Terramechanics, Wheeled Robots

## I. Introduction

Planetary exploration rovers are essential to our exploration of the solar system, and wheeled planetary rovers are the current standard mobility system used due to their stability and ability to traverse rough terrain [1], [2]. However, for missions that require manipulation of the environment, rovers typically require additional payload systems such as arms to perform these functions. Similarly, damage to the mobility systems of rovers is a much higher risk than to terrestrial robots, as they cannot be repaired in-situ. This paper will augment the capabilities of wheeled planetary rovers through the addition of the ability to manipulate terrain and recover mobility in the case of physical system degradation. Both of these aims can be achieved without additional hardware, making them attractive approaches for mass-constrained missions and ongoing missions with rovers that have already been deployed.

There is precedent for the use of rover mobility systems to manipulate terrain; rovers currently use their wheels to “scuff”, or dig trenches, in Martian regolith for the purpose of observing or sampling below the surface layer [3], [4]. Notably, [3] utilized shallow wheel scuffs to infer soil mechanics properties in situ. However, these scuffs are relatively shallow, short, and not conducted while the rover is traversing longer distances. The ability to safely make trenches or scuffs while driving would allow increased subsurface observation without compromising travel distances.

Wheeled rovers have also recovered from degradation of the mobility system in various forms, from rock jams to stuck steering motors. In each instance, the approach was ad-hoc and took significant trial-and-error on representative hardware [1], [5].

In this paper, we present a novel way to generate both open- and closed-loop driving strategies for rovers that are

54 pt  
0.75 in  
19.1 mm

54 pt  
0.75 in  
19.1 mm



FIGURE 1: The KREX-2 rover after digging a trench in unprepared soil in the Atacama Desert. This trench was dug with the rover speed at 20 cm/s, the digging wheel rotated 90° from the direction of travel, and the digging wheel’s speed at 50 cm/s. The white streaks present in the rover’s track are due to halites in the soil.

driving outside of their normal operating modes using a terramechanics-and-optimization-based approach. We show how rovers are able to use their wheels to meaningfully modify terrain through techniques we term “nonprehensile terrain manipulation” (NPTM) as well as recover from mobility loss. We show field tests in Martian and lunar analog outdoor environments of these behaviors, as in Fig. 1, and demonstrate automatically-generated driving strategies successfully manipulate terrain and compensate for failed actuators in a laboratory environment.

The remainder of this paper is organized as follows. Section II provides additional background and motivation from prior work. Section III reports on field demonstrations of the off-nominal rover behavior we are interested in achieving in this paper. Section IV uses a model of wheel-soil interaction from [6], [7] in order to optimize the rover driving conditions to achieve the desired behavior. Section V shows experimental results using this technique on a rover. Finally, Section VI concludes the paper with a summary of its implications and areas for future work.

## II. Background/Motivation

In this section, we detail the current state of off-nominal driving techniques for planetary rovers, with a focus on techniques demonstrated in-situ. We further motivate the development of off-nominal driving strategies by identifying upcoming rover missions for which the addition of manipulation or compensation for degradation would be high-impact.

### A. Manipulation of terrain by rovers

Nonprehensile manipulation is any manipulation that does not rely on grasping, or force closure, to accomplish tasks. It has been used to stably push objects from one place to another [8], reposition objects for further manipulation [9], [10], and simultaneously move objects to clear areas [11].

Much of the work on nonprehensile manipulation has focused on the use of planar surfaces, such as arms [11] or paddles [8], [10]. In [12], nonprehensile rearrangement of boxes was demonstrated in simulation on NASA Ames Research Center’s KREX-2 rover (Fig. 1) augmented with a front-mounted bulldozer blade. This was then demonstrated on hardware by having the KREX-2 use its frame to push cardboard boxes on flat outdoor terrain in [13]. The use of the rover’s frame in [13] and the sides of a robot’s arms in [11] are both examples of using features of a robot not intended for manipulation for that purpose. A notable example of wheel-based manipulation is the Mobipulator, which uses its wheels to move objects on a smooth surface [14]. The Mobipulator is capable of scooting paper around and rolling pencils by placing its wheels on top, and uses only frictional forces from its wheels to do so [15]. In [16] and [17], intentional control of the interaction between robot feet and wheels with granular media is used to improved the substrate’s properties for locomotion.

There are several significant examples of NASA missions utilizing ad-hoc nonprehensile manipulation, to varying degrees of success. The dislodging of the “Potato” rock from Spirit’s wheel, as addressed in Section II.B, is one such example [1], [5]. Spirit and Opportunity have also used their wheels to investigate terrain properties through intentional construction of scuff marks and trenches [3], [4], though this technique was first explored during the Sojourner mission [18]. More recently, Curiosity and Perseverance have incorporated wheel scuffs into their operations, and use the maneuver to sample from lower layers of regolith [19], [20].

### B. Mobility system failure in rover missions

Planetary rovers operating over long periods may encounter mobility system failures of varying impact, ranging from temporary reduction in capabilities to full end-of-mission. Historical mitigation of mobility system failure has been largely ad-hoc and platform-specific. The Mars Exploration Rovers (MER), Spirit and Opportunity, both encountered reduced mobility due to actuator failures during their operational lifetimes.

Spirit temporarily had a rock dubbed “the Potato” jammed in its rear right wheel. Use of a twin rover on Earth allowed operators to test wheel motions for ejecting the rock with the aid of gravity, and then implement them successfully on the real rover [1], [5]. Spirit later had a partial mobility failure in the form of increased current draw in its front right wheel, likely due to poor distribution of lubricant within the actuator [21]. Load on the wheel was decreased by driving the rover backwards and attempts were made to

redistribute lubricant through intentional heating, and careful driving, but Spirit’s front right wheel eventually degraded to a fully stalled condition. This resulted in sideways drifting and yawing during driving maneuvers, which was somewhat mitigated through driving techniques tested in a terrestrial testbed and on Mars. When Spirit later became embedded in soft soil, reduced thrust due to the failed actuator contributed to the inability to extricate itself, and the rear right wheel also stalled, then fully failed. Ultimately, Spirit was unable to escape the soft soil, resulting in end of mission [22].

At one point, Opportunity was temporarily embedded in soft soil but was able to escape through driving strategies tested on earth [23]. Opportunity showed signs of potential lubrication issues in its front right wheel, similar to Spirit, but lessons learned from mitigating Spirit’s mobility failures allowed operators to maintain actuator functionality through driving strategies to reduce load on the actuator, intentional heating, and resting of the actuator when needed [21]. Opportunity later had another partial mobility failure when its front right steer actuator froze at a 7° angle, which resulted in yaw error and side slip while driving, as well as reduced steering capabilities. Ad-hoc driving strategies using three-point turns and approaching science targets backwards had to be developed to enable full operation of the rover, and these methods likely increased strain on the other actuators. The other front steering actuator later failed as well, though Opportunity’s end of mission was ultimately due to loss of power as a result of a powerful dust storm [24].

Previous works on mitigating mobility loss as experienced by Spirit and Opportunity have utilized optimization as a tool for generating controllers for degraded motion. In [25], the authors use a simplified soil model to generate new steering methods for a rover suffering from a variety of mobility failures, and in [26] the authors optimized a strategy for rovers to escape entrapment in soft soil.

As of this writing, the rovers Curiosity and Perseverance have not experienced significant mobility system failures beyond the degradation of Curiosity’s wheels [27]. However, actuator failure has been explicitly cited as a potential threat to Curiosity’s mission [27]. Loss of a steering or drive actuator as seen on Spirit and Opportunity would seriously limit Curiosity’s mobility, and the rover additionally has braking mechanisms which could stall a wheel if they were to fail in a closed position.

The 2022 Planetary Science and Astrobiology Decadal Survey calls for a 1000 km long rover mission, “Endurance-A,” in which a rover would traverse orders of magnitude further than previous missions [28]. With such a long mission timeframe, Endurance-A is almost certain to encounter degraded mobility during operation, and mobility system failure mitigation should be systematically developed and qualified in advance. Additionally, NASA’s Volatiles Investigating Polar Exploration Rover (VIPER), which would search the lunar surface for water ice, has a four-wheeled active suspension [29]. The actuated suspension gives VIPER

flexible extreme terrain mobility at the cost of an increased number of mobility actuators and potentially a higher mobility cost due to actuator loss than might occur on a similar six-wheeled rover. With high-risk wheeled rover missions planned for the near future, intelligent strategies to mitigate mobility failures are essential.

### III. Hardware demonstrations of manipulation and failure compensation

In this section, we demonstrate off-nominal driving on a variety of full-size rovers in field environments. The platforms were chosen for their similarity to planetary rovers and proximity to representative field environments. We begin with demonstrations of NPTM actions in a Martian analog environment with KREX-2, a large, four wheeled rover (Section III.A) [30]. These experiments show that the rover can excavate trenches, dig holes, and clear paths for smaller robots. We then assess the impact of actuator failure on the mobility of NASA Ames’ Volatiles Investigating Polar Exploration Rover (VIPER) in depth through a series of quantitative and qualitative driving tests (Section III.B) [29]. Finally, we demonstrate hand-tuned driving strategies to compensate for two types of failure on KREX-2, which is approximately the same size as VIPER (Section III.C).

#### A. NPTM field demonstrations

In this field experiment, we demonstrate wheel-based soil manipulation on a large rover in a Martian analog environment, illustrating various applications for NPTM. The KREX-2 rover [30] was used to dig trenches as an example of continuous excavation; dig holes while stationary to determine the soil’s angle of repose; and carve a small ramp into terrain to allow a smaller rover to cross an obstacle.

We conducted field demonstrations of nonprehensile terrain manipulation both at NASA Ames Research Center’s Roverscape and in the Atacama Desert in Chile. As the driest place on Earth, the Atacama Desert is a common Martian analog, with almost no annual rainfall, plant, or animal life [31]. The rover is roughly the size of a golf cart and weighs approximately 230 kg in its nominal configuration without sensor payloads. For these tests, KREX-2 was outfitted with inflated rubber tires for all experiments unless noted otherwise. The approximate wheel geometry is described in Table 1. All demonstrations described here are performed using a single rover wheel to manipulate terrain. Additionally, the trenching actions performed are also used to test the soil flow model proposed by [6] on unprepared natural terrain.

##### 1) Wheel-based trenching

We performed four demonstrations of the “trench digging” action in soil in Chile’s Atacama Desert, with the aim of both achieving large amounts of soil displacement and experimentally testing the model proposed in [6]. The soil

54 pt  
0.75 in  
19.1 mm

54 pt  
0.75 in  
19.1 mm



FIGURE 2: Setup of trenching experiments in the Atacama Desert. The KREX-2 rover (left) was driven over unprepared soil, with the FARO<sup>®</sup> LIDAR (center) scanning before and after and video recording of each test (right). Not pictured is the Leica Total Station used to track the rover’s position and speed during testing [33].

was soft, dry, and noncohesive, with a fragile crust layer less than 1 cm thick. The testing site was selected for its level, undisturbed ground and soft soil, and there was no preparation of or traffic over the testing site in advance of experiments. KREX-2 can be seen in Fig. 1 after completing a trenching test.

Terrain was mapped with a FARO<sup>®</sup> LIDAR 3D scanner [32] (Fig. 2) before and after each test, with fiducial markers mounted to the rover used to align KREX-2’s initial and final position and orientation to scans. The rover’s position and velocity were tracked using a Leica Total Station and a single rover-mounted reflector prism during each trenching action [33]. The location of the prism relative to the fiducial markers was fixed, and used to correlate the rover’s speed and position information with the terrain maps with sub-centimeter accuracy. Each test was recorded with a camera mounted to the rover chassis and an off-board video camera.

Four trenching actions were tested, with the rover driving in a straight line while trenching with its rear right wheel for all four, as seen in Figs. 1 and 2. The rover’s travel speed for each test was approximately 20 cm/s with the rear right wheel spinning at 50 cm/s, with slip angles  $\beta$  of 0°, 30°, 60°, and 90°. Individual wheel speeds and angles varied for the remaining three non-trenching wheels, but each had the front left wheel turned slightly inward, the rear left wheel turned slightly outward, and the front right wheel spinning slightly faster than the nominal travel speed of 20 cm/s. These speeds, while slightly faster, are in the same regime of typical rover operations, where wheel-soil interaction is dominated by static forces and little material is dynamically ejected. The speeds and angles of the other wheels were hand tuned in the field to allow the rover to drive straight while trenching. Section IV describes how to automate this process.

$r$	Wheel radius [cm]	23
$b$	Wheel width [cm]	15.2
$h_g$	Shearing radius [cm]	25.3

TABLE 1: Wheel geometry parameters for trenching experiments with KREX-2 in the Atacama Desert. The shearing radius is an estimate that incorporates lug height on the wheels.

Each trenching action was run a single time for 20 seconds. Between the rover’s speed and time to reach a steady state trench geometry, each test resulted in a steady-state trench of approximately 2.5 meters long.

Despite being dug with the same wheel speeds and in the same terrain, four very distinct trench geometries were formed, as seen in Figs. 3 and 4. Trenches ranged from 18 cm wide and 1.3 cm deep with the trenching wheel at 0° to 38 cm wide and 10.6 cm deep with the wheel at 90°, with soil piled up to 1/3 of the wheel diameter. The width and depth of each trench are reported in Table 2. Note that trench width here is measured as the width of the region where the soil has been excavated below surface level, and does not include the pile up off to the left side of the trenches.

We compared the observed trenches to those predicted by the soil flow model described in [6]. We present the results of the trench prediction in Fig. 4 and Table 2. For each trench, the wheel sinkage and wheel location relative to the trench were manually extracted from LIDAR scans, along with projections of sections of the trench onto a 2D plane seen in blue in Fig. 4. We estimated the cohesion to be zero, and measured the soil angle of repose on several purpose-constructed soil piles; this process is described in Section III.A.3.

The trench shape results had good qualitative agreement between the model prediction and the observed trench shapes. The average error in predicting the profile’s shape ranged from 0.5 cm for the 0° trench to 3.0 cm to the 90° trench. There is some sideways shift in the shape possibly due to error in aligning the wheel location to the trench from scans, and the model tends to underpredict soil transport from the right side of the wheel to the left, which likely accounts for most of the error in the 90° trench. Note that the wheel geometry used to dig these trenches does not exactly align with the theoretical wheel geometry in the soil flow model, as KREX-2’s rubber tires are rounded and have large treads rather than uniform grousers, and fully accounting for the wheel’s shape and tuning the volume fraction  $\zeta$  would enable better trench shape prediction. Additionally, the 30° model prediction shows the formation of a step in the slope which does not appear distinctly in the scanned trench shape. Despite this, the shapes of the trenches predicted qualitatively describe the trenches observed.

The trenches dug also illustrate how NPTM can be used to aid scientific surveying and sampling. In Figs. 1 and 3



FIGURE 3: Trenches dug by KREX-2 in the Atacama Desert with 1 foot ruler for scale. The wheel angle for each trench, from left to right, was 0°, 30°, 60°, and 90°.



FIGURE 4: LIDAR scans of trenches dug by the KREX-2 rover in the Atacama Desert (blue) overlaid with predicted trench shape (red). The wheel angle for each trench, from left to right, was 0°, 30°, 60°, and 90°.

Trench Type	Trench Width [cm]	Trench Depth [cm]	Avg. Error [cm]	Median Error [cm]	Depth Error [cm]
$\beta = 0^\circ$	18	1.3	0.5	0.5	0.9
$\beta = 30^\circ$	24	4.6	1.2	1.1	2.1
$\beta = 60^\circ$	34	7.8	1.5	1.4	1.8
$\beta = 90^\circ$	38	10.6	3.0	1.9	0.4

TABLE 2: Trench geometry and quality of soil flow model fit for the four trenches dug by the KREX-2 rover in the Atacama.

the color difference between the surface and subsurface soil is clearly visible, and in the excavated tracks faint white streaks can be seen. These streaks are due to the presence of halites in the soil, which blend in with the soil and rocks in the undisturbed terrain. Trenching while driving can be used to survey subsurface soil over long travel distances without stopping to sample.

With these demonstrations, we have shown that the “trench digging” NPTM action can move meaningful amounts of soil, aid scientific sampling, be performed without getting stuck, and can be modeled with the soil flow model described in [6].

## 2) Multipass trenching

Next, we conducted multiple trenching passes over the same terrain, with the aim of seeing how much increase in soil motion can be gained. The rover was run in three passes in one direction over the same terrain, with manual alignment

of the rover’s heading performed at the start of each run. The rear right wheel was used to dig, with the front right wheel running inside the trench on the second and third passes. The trench dimensions are reported in Table 3, and images of the three passes along with LIDAR scans of the trench can be seen in Fig. 5. Multiple passes over the same trench had a moderate increase in trench size from the first to second pass, but a smaller change from the second to third pass, which can be seen in Fig. 5. While each trench varied in width along its length due to deviation in the rover’s heading, the width of the trench following the first pass was 27 cm, which increased to 29 cm on the second and 30 cm on the third pass.

Maintaining a trajectory along trenched terrain is difficult without active control of rover heading, as the trenching wheel tends to slip into the deepest part of the trench, where it moves little soil. Closed-loop control of trenching via either IMU or visual odometry, which we perform in Section V.A.1, would be needed for effective re-excitation of trenches. For an objective like subsurface sampling, a single trenching pass is best, as multiple passes have diminishing returns in excavation area and require more advanced control. Notably, the trench dug by three passes with the trenching wheel at 60° was both shallower and narrower than the trench dug by a wheel at 90° in Table 2, suggesting that the driving primitive used to trench may have larger impact than repeated passes.

## 3) Hole digging

In addition, a hole digging action was demonstrated. Seven holes/soil piles were constructed by moving a single wheel

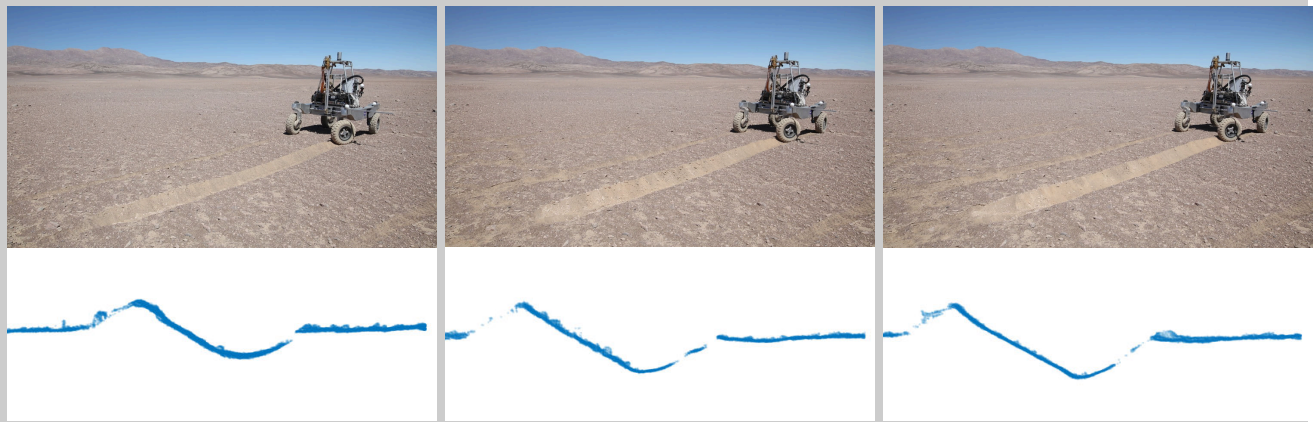


FIGURE 5: KREX-2 performing three trenching passes over the same terrain in the Atacama Desert (top row), and LIDAR scans of all three runs (bottom). All three passes were performed with the wheel at a slip angle of  $60^\circ$  with a rover body velocity of 20cm/s and wheel rim velocity of 50cm/s. The trench had a width of about 27 cm after the first pass (left), a width of 29 cm after the second (center), and a width of 30 cm after the third pass (right).

Pass Number	Trench Width [cm]	Trench Depth [cm]
1	27	6.0
2	29	6.5
3	30	7.7

TABLE 3: Dimensions of a trench dug in multiple passes by KREX-2 in the Atacama. The trench was excavated with the digging wheel inclined at  $60^\circ$  to the direction of travel.

could use nonprehensile terrain manipulation to construct paths for smaller rovers performing monitoring or other tasks over a set area.

In the first demonstration, shown in Fig. 7, a ridge of rocks is assembled in NASA Ames Research Center’s Roverscape on otherwise level and compacted terrain. As seen in the first two frames, MiniRHex is unable to cross the rocks, flipping over and landing on its back when it tries to do so. The third frame shows KREX-2 leveling a rock pile by using a wheel to clear a path through the rocks. In the final frame, we see that MiniRHex is now able to cross the rock ridge without falling over.

The second demonstration was conducted in the Atacama Desert, where a naturally occurring step of about 20 cm in the terrain was located. The sharp step, seen in Fig. 8, was formed by water runoff and has a steep vertical face, but is made of friable soil that crumbles under moderate force. While KREX-2 is able to easily drive over obstacles this size, MiniRHex flipped over backwards and got stuck when attempting to climb it, as seen in the first two frames of Fig. 8. KREX-2 then used its wheel to construct a ramp in the step by crushing the terrain, seen in the third frame, which then allowed MiniRHex to climb up to the higher terrain, as shown in the final image. For this test, KREX-2 had a tensegrital metal wheel design with a radius of 29 cm, which is slightly larger than its usual 23 cm rubber tires. [35], [36].

These two robot teaming scenarios demonstrate how actions for moving rocks and soil can be used by a large rover to increase the mobility of a smaller robot. One can imagine a large rover like KREX-2 systematically clearing paths for a MiniRHex-sized robot to move along performing monitoring or observation tasks, allowing the larger rover to traverse elsewhere, assured of the smaller robot’s safe mobility.




FIGURE 6: The KREX-2 rover with five soil piles dug in the Atacama desert with the purpose of displacing soil to measure its angle of repose (left), and closeup of KREX-2 digging one of the holes (right).

with the rover stationary to estimate soil angle of repose as in [3], as seen in Fig. 6. Maximum slope angle measurements on these soil piles were taken with a handheld inclinometer, with an average value of  $\phi = 33.5^\circ$ . This value was used in the model evaluations performed in Section III.A.1, and thus has clear use in aiding scientific work and mobility planning.

#### 4) Robot teaming

KREX-2 also demonstrated robot teaming scenarios with MiniRHex, an open-source miniature hexapod robot approximately 19 cm long. [34]. In a mission context, a larger rover




FIGURE 7: Images from a video of KREX-2 helping MiniRHex traverse a pile of rocks by moving them. Video recorded at NASA Ames Research Center’s Roverscape. Left to right: MiniRHex tries to cross the rocks; but is unable to and flips over; once KREX-2 clears rocks from the path; MiniRHex is able to climb through.



FIGURE 8: Images from a video of KREX-2 helping MiniRHex traverse a natural step in terrain due to water flow in the Atacama Desert. Left to right: MiniRHex tries to climb the step; but is unable to and flips over; once KREX-2 crushes a ramp into the step; MiniRHex is able to climb the obstacle.




FIGURE 9: Moon Gravity Representative Unit 3, a mobility prototype for VIPER (left) and illustration of the actuators and joints referred to in this section (right). The actuators referred to in this section are the rear (dark blue) and front (green) drive actuators, steering (light blue), and suspension (red) actuators.

prototype Moon Gravity Representative Unit 3 (MGRU3), shown in Fig. 9.

1) Experimental setup for rover mobility testing

All tests were conducted on MGRU3 in GRC-1 lunar simulant [40] at NASA Glenn Research Center’s Simulated Lunar Operations (SLOPE) Lab [41]. The Moon Gravity Representation Unit version 3 (MGRU3) is a prototype of VIPER, designed with representative software, controls, controllers, mechanisms, and wheels, but with a mass approximately 1/6th of the Earth equivalent. There are many potential failure modes for the actuation of VIPER; each of the twelve motors can potentially fail in a “locked” (fixed position) state such as in the case of a rock jam, or an “unpowered” state such as in a power loss or actuator damage event [21]. In addition, the position in which a stuck suspension or steering actuator fails can vary the impact on mobility performance. A subset of potential failure modes were explored due to limited testing time, with a mixture of operationally likely failure states and failure states at the extremes of joint limits to cover the full range. The following failure states were tested individually: free-rolling drive actuator, stuck drive actuator, suspension locked with single wheel raised, and a single steer actuator locked at a fixed nonzero angle. VIPER’s suspension can be set to either maintain a fixed pose or move according to force thresholds for a coarse form of force control; unless otherwise noted, all tests were run with the suspension set to maintain a fixed posture with all wheels held at a neutral (0°) angle to the body.

An illustration of the experimental setups is shown in Fig. 10. For all tests, soil preparation consisted of manual loosening and raking of simulant between runs as prescribed

**B. Impact of actuator failure on mobility for VIPER**

NASA’s Volatiles Investigating Polar Exploration Rover would launch to the lunar surface to search for water [29]. To test the effect of actuator failure on mobility, we use a test version of VIPER that has a four-wheeled actuated suspension. Each wheel has an in-wheel drive motor and individual actuators for steering and raising/lowering the suspension. The rover is approximately the size of a golf cart but weighs only 86 kg [37]. This mobility system has twelve actuators, enabling advanced locomotion techniques such as wheel-walking and swimming-like gaits [38], [39] but potentially at a higher mobility cost in the case of actuator loss than might occur on a similar six-wheeled rover with a passive suspension. We performed both quantitative and qualitative assessments of rover mobility in the form of drawbar pull tests and motion-tracked driving with VIPER

54 pt  
0.75 in  
19.1 mm

54 pt  
0.75 in  
19.1 mm




FIGURE 10: Experimental setup for testing the impact of joint actuator failure on VIPER. Drawbar pull tests were performed in a tethered configuration (left), and qualitative drive tests were performed untethered (right), with both tests in GRC-1 simulat [40].

in [42]. The rover’s position and velocity were recorded by 3D motion capture using markers mounted to the rover, and the rover’s actuator speeds and positions were recorded internally. Quantitative driving tests were performed by having the rover drive for either 30 seconds or until it left the prepared soil area. For drawbar pull testing, a fixed load was applied to the rover chassis via a tether to induce slippage, with tether load and length measured. The applied load began at 50N and was increased every 20 seconds in 50N increments. Nominal driving performance with all actuators operational was measured with the same experimental setup as a mobility benchmark. For each test, the rover was driven open-loop with a nominal speed of 10 cm/s, with speed and position control on individual actuators but no closed-loop control on the full rover’s state.

2) Qualitative drive testing on a rover

a: *Stuck drive motor*

The MGRU3 rover can be seen after attempting to drive with a nonrotating (stuck) rear left drive actuator in Fig. 11 (green). The affected wheel embedded several centimeters into the soil, putting significant drag on the rover, while the driving wheels excavated a large amount of soil without gaining meaningful traction. The rover dragged the affected wheel behind while pivoting about it, resulting in less than a meter of forward progress when the rover should have driven 3m, as can be seen in Fig. 11.

b: *Unpowered drive motor*

MGRU3 was able to drive normally with a single unpowered motor, as seen in Fig. 11 (yellow) – mobility performance on flat ground was not visibly different from nominal driving, but the reduced thrust available to the rover may be a problem on slopes, which is explored in Section III.B.3.

c: *Stuck steering actuator*

Fixing MGRU3’s rear right steering actuator so that the wheel is pointed 30° outward, as shown in Fig. 12, resulted in very slight drift in heading but did not impede the rover’s ability to make forward progress, as seen in Fig. 11 (red). A steering actuator stuck at a moderate angle would have a

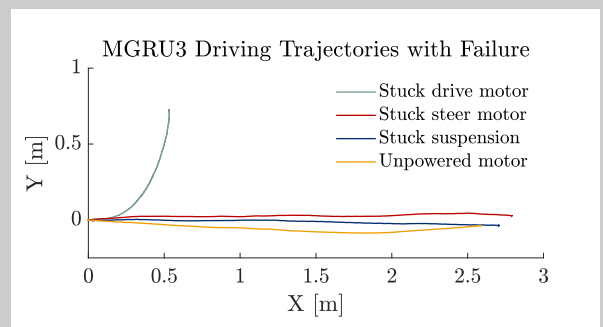


FIGURE 11: Trajectories followed by MGRU3 driving with various failed actuators. The rover barely moved with a non-rotating drive motor (green) but was able to make forward progress with other tested failure modes.




FIGURE 12: MGRU3 driving with its rear right steering actuator stuck at 30° on the rear right wheel.

larger impact on steering performance than normal driving performance, as was observed on Opportunity [21].

d: *Stuck suspension*

The rear right wheel was lifted at a 40° angle to simulate the failure of a suspension actuator. Due to the landing configuration of VIPER’s suspension, where the wheel is stowed with the suspension fully raised, it is most likely that if one of these actuators were to fail it would occur in the fully raised state rather than the fully lowered state. While MGRU3 was able to maintain forward motion with only three wheels touching the ground, it pitched between the two support triangles formed by its wheels, as seen in Fig. 13. This motion would make driving VIPER via either teleoperation or visual odometry extremely difficult and could cause further damage to other parts of the rover through repeated impact. Despite this, the rover was able to drive straight, as shown in Fig. 11 (blue).

3) Drawbar pull testing on VIPER

The same failure states were tested again with a drawbar pull load applied to the rover via a cable to simulate slope climbing. Drawbar pull loads began at 50 N and increased in approximately 50 N increments. Additionally, the unpowered front wheel trial was run a second time with the suspension’s force control method enabled. For a given drawbar pull force




FIGURE 13: MGRU3 driving with a suspension joint stuck in a raised position pictured before moving (left) and immediately after driving started (right), with the rover tipped backwards onto the affected actuator.

$F_{DBP}$  and rover weight  $W_{rover}$ , the equivalent slope angle  $\alpha$  is given by

$$\alpha = \tan^{-1} \left( \frac{F_{DBP}}{W_{rover}} \right) \quad (1)$$

To compute slip, the rover's body velocity was measured at the geometric center of its body using 3D motion capture. In the tests conducted here, all wheel speeds were set to a rim velocity of 10 cm/s, with the exception of the trials where the wheel was not rotating (stuck) or allowed to spin freely (unpowered).

The overall rover slip was computed as an average of the actual measured speed of the three unaffected wheels relative to the rover's body velocity. Given that the disabled drive motor tests result in the affected wheel going slower than the commanded drive speed and all other wheels were traveling at the same velocity, the maximum slip on all wheels is equivalent to the average slip on unaffected wheels for these tested scenarios. As VIPER's nominal operational limit on slope climbing is to maintain wheel slip below 0.4, [43], keeping both the maximum and average unaffected wheel slip below that value is more conservative than looking at the average slip value across all wheels, which can mask slip variation between wheels.

The induced wheel slip for each equivalent slope angle tested can be seen in Fig. 14. Note that the discontinuity and double measurement for the nominal trial at a 10° slope angle is due to that trial being separated into two runs, while all others were recorded in a single continuous run. Nominal (unaffected) driving had the lowest amount of slip on all slopes, ranging from 0.02 on an equivalent slope of 3° to 0.59 on a slope of 19°.

We can draw several conclusions from this data, which shows both the impact of various failure modes and the relative importance of different types of mobility testing. The impact of actuator failure on slope climbing ability depends greatly on which actuator is affected.

*a: Stuck or unpowered drive motor*

Loss of a rear or front drive motor was associated with an increase in slip from 0.4 to 0.8 on slopes of 15°, and an increase from 0.02 to 0.55 on shallow inclines of 3 degrees. While a locked front wheel showed similar traction reductions to an unpowered wheel in these tests, the rover

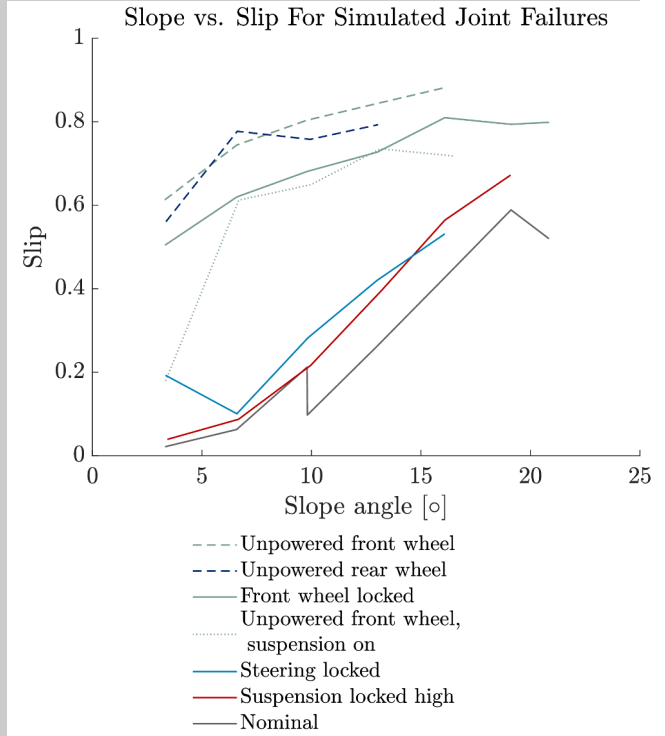


FIGURE 14: Plot of slip vs. slope angle for VIPER driving with failed actuators.

was unable to maintain a straight heading; rover heading had to be manually adjusted during data recording to keep the rover on its path. Use of the suspension to balance forces on the wheels with an unpowered front wheel reduced slip on the 3° slope to 0.2 from 0.6 but had a more modest impact on steeper slopes.

*b: Stuck steering actuator*

Locking a steering joint increased slip from 0.02 to 0.2 on shallow inclines but only from 0.1 to 0.3 on slopes of 10° and from 0.4 to 0.5 on slopes of 15°, though the rover was unable to maintain a straight heading and had to be manually readjusted during trials to keep it in the testbed.

*c: Stuck suspension*

Locking a suspension joint into a raised position only increased slip to 0.1, with minimal impact on low angles. Given that the rover rocked back and forth between the wheels in the free driving test, the tether may have increased the stability of the rover.

4) Experimental comparison

Although the quantity of data presented here is insufficient to draw quantitative trends with certainty, it can be used to illustrate the risk of actuator loss to planetary rover missions and motivate varied testing campaigns. While both the drawbar pull tests presented here and the free driving

54 pt  
0.75 in  
19.1 mm

54 pt  
0.75 in  
19.1 mm





FIGURE 15: KREX-2 driving with a stuck steer motor in the Atacama Desert with hand-tuned compensation. The rear right wheel is stuck at 30° (left), 60° (center), and 90° (right).



FIGURE 16: KREX-2 rover driving with a stuck drive motor without compensation (left) and with open-loop compensation (right).

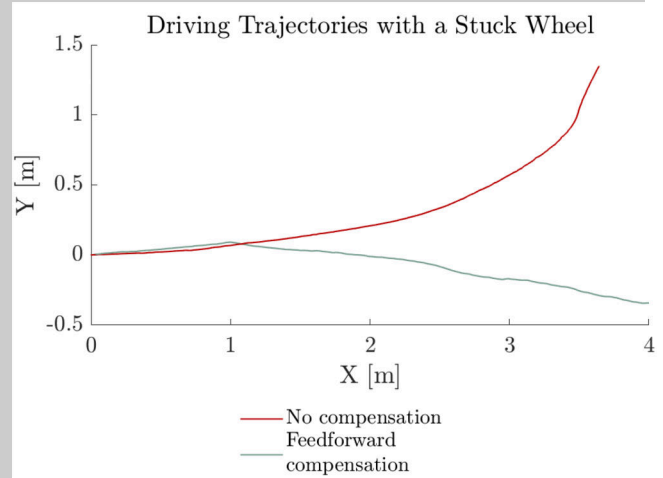


FIGURE 17: KREX-2’s drive trajectories with the rear right wheel stuck (nonrotating). All trajectories are open-loop, with no steering controls used. Note that the positive x axis represents the forward direction and the positive y axis represents left as seen from the view in Figure 16.

When driven with a rear wheel held in a fixed position, as in the case of a rock jam or stuck brake mechanism, the rover yaws about the stuck wheel and veers off track, as seen in the left image in Fig. 16 and in the plot of the rover’s trajectory in Fig. 17. A hand-tuned feedforward compensation strategy was developed, with the front right wheel (same side as the stuck wheel) sped up, and the front two wheels turned slightly left. When this feedforward (open-loop) compensation strategy is applied, changing the speeds and angles of the other three wheels, KREX-2 is able to maintain a relatively straight heading and stay on track, as shown in Figs. 16 and 17.

KREX-2’s forward travel speed with compensation was 6.4 cm/s, with an average slip ratio on the unaffected wheels of 0.66, and a maximum slip ratio of 0.94 on the front left wheel. This represents a very large amount of slip for a rover, and is above the slip threshold of 0.4 set for VIPER’s operations [43]. The loss of thrust from the stuck wheel and the drag it induces are major impediments to KREX-2’s locomotion, though it is able to retain some mobility with the hand-tuned driving strategy.

While KREX-2 drove much straighter using a compensated driving strategy than without, variation in the rover’s heading and horizontal position can still be observed throughout the rover’s trajectory, as seen in the curve of

the rover tracks in Figs. 16 and 17. This is likely due to a combination of several factors; variation in soil preparation, wear on the rover’s wheels, and the hand-tuned nature of the selected driving strategy.

This demonstration shows that recovery from failure is possible for planetary exploration rovers. However, it also shows the limits of ad-hoc compensation strategies; even with hours of hand-tuning on the exact hardware and environment of interest, it is extremely difficult to fully balance the forces and moments induced on the rover by a disabled wheel. In the next section, we address this by presenting methods to automatically generate driving strategies to overcome actuator failure.

#### IV. Optimization-based generation of driving strategies

In this section we present an optimization-based approach for generating rover driving strategies, starting by looking at a rover driving straight over flat ground. We utilize the




FIGURE 18: Mini rover used to test off-nominal driving strategies. The rover has two rockers with a passive pivot, with independent steer and drive motors for each wheel. The rover’s center portion bearing the electronics is fixed to the right rocker, and it is tethered for power and communications. A RealSense T265 tracking camera is used to obtain the rover’s position and orientation.

terramechanics model of wheel-soil interaction forces developed in [6], [7]. We then detail how to use the optimization technique to generate driving primitives for trenching objectives or degraded mobility systems. The system analyzed is a small four-wheeled rover about the size of a microwave, pictured in Fig. 18, and is referred to in this section as the “mini rover”.

In setting up the overall optimization strategy, we first define the decision variables the optimization is given to work with and how those get converted into the full state of the rover. Next, we discuss the constraints that force the optimization to choose a kinematically feasible driving strategy. Finally, we go over the objective function. See Table 6 for a full list of variables used in this section along with their meanings and units.

### A. Terramechanics Modeling

In order to predict rover mobility while digging with wheels or experiencing mobility degradation, we must know the forces occurring from wheel-soil interaction, whether the wheel is in slip or skid, and driving straight or steered at an angle. We also need to know how soil flows around the wheel, both to predict the shape of trenches dug by the wheels for soil manipulation and to understand the geometry of the wheel-soil interaction surface for force prediction. In [6], the authors developed a closed-form model of how sandy soil flows around a wheel. In addition, in [7] we created a terramechanics model for predicting wheel-soil interaction forces at any steering angle and wheel slip condition, incorporating the soil flow model to reduce empirical tuning. This terramechanics model, which is used in this work, has two key tuning parameters dependent on the wheel’s state, which define the shape of the normal stress distribution  $\sigma_n$  under the wheel. The normal stress under a wheel is treated as a curve with a peak at an angular location  $\theta_m$  on the wheel’s surface that is driven by the wheel’s slip ratio  $s$  such that

Symbol	Value	Units
Rover and wheel state		
$\beta$	Slip angle	$^\circ$
$\omega$	Wheel angular velocity	rad/s
$h$	Wheel sinkage	m
$W$	Load on each wheel	N
$v_x, v_y$	Translational velocity along $x, y$ axes	m/s
$v$	Net translational velocity along direction of travel	m/s
Rover parameters		
$r$	Radius of wheel	m
$b$	Width of rover wheel	m
$h_g$	Grouser height	m
$l_r$	Rover rocker length	m
$h_r$	Height of rocker pivot from wheel center	m
$w_r$	Width of rover between wheel centers	m
$W_{rover}$	Weight of rover	N
Terramechanics parameters		
$a_0, a_1$	Empirical max stress angle coefficients	
$b_0, b_1$	Empirical exit angle coefficients	
$a_c, a_s$	Constant coefficients within $a_0$ and $a_1$	
$a_{c\beta}, a_{s\beta}$	$\beta$ -dependent coefficients within $a_0$ and $a_1$	$1/^\circ$
$\theta_f, \theta_r$	Front, rear angles where wheel enters, exits soil	$^\circ$
$\theta_m$	Angle of max normal stress under wheel	$^\circ$
$\zeta$	Grouser transport volume fraction	

TABLE 6: Table of variables used in rover driving optimization.

we have

$$\theta_m = \theta_f (a_0 + a_1 s) \quad (2)$$

where  $a_0$  and  $a_1$  are empirical tuning parameters dependent on the wheel and the soil.

The results in [7] showed that for the type of soil and wheels used here, we can define  $a_0$  and  $a_1$  as a linear function of the slip angle  $\beta$ . Using the fits from [7], we calculate  $a_0$  and  $a_1$  for a wheel at an arbitrary slip angle  $\beta$  by

$$a_0 = a_c + a_{c\beta} \quad (3)$$

$$a_1 = a_s + a_{s\beta} \quad (4)$$

As all other parameters in the terramechanics model are not state-dependent, we now have the tools needed to compute the forces experienced by a wheel at any driving state. We next apply these models to a full rover to optimize driving strategies for rovers operating outside of their nominal conditions.

### B. Decision variables

When a rover is driving, the only true free variables are the speed and angle of each wheel; the vehicle will either eventually reach a steady state with constant wheel sinkages, travel velocity, and rover angular velocity, or else become entrapped in the soil. As the terramechanics model developed in [6], [7] is not closed-form and only looks at a single wheel

at a time, we cannot directly solve for those final state values. We can, however, take a given rover state and check if it is at a steady state compatible with forward motion.

We set the rover's desired angular velocity to zero and set the rover's travel velocity to our desired speed, which is 3 cm/s here. We then only have decision variables corresponding to each wheel separately. For each wheel, we have the wheel's rotational speed  $\omega$ , slip angle  $\beta$ , and sinkage  $h$ . We additionally include the individual wheel's vertical load  $W$ . While we could choose either  $h$  or  $W$  and solve for the other iteratively, doing so induces problematic numerical discretization, which is explained in detail in Section IV.I. Thus, rather than including a numerical sinkage finding, as in [7], here we let the optimization routine solve for the driving strategy and sinkage simultaneously.

We therefore have sixteen decision variables, and our solution  $x$  takes the form

$$x = [\omega_{FR}, \omega_{FL}, \omega_{RR}, \omega_{RL}, \beta_{FR}, \beta_{FL}, \beta_{RR}, \beta_{RL},$$

$$W_{FR}, W_{FL}, W_{RR}, W_{RL}, h_{FR}, h_{FL}, h_{RR}, h_{RL}]$$

where the subscripts  $FR, FL, RR,$  and  $RL$  refer to the front right, front left, rear right, and rear left wheels in the rover body frame.

The decision variables are subject to lower and upper bounds that constrain them to feasible values, such that

$$0 \leq \omega \leq \omega_{max} \quad (6)$$

$$-\frac{\pi}{2} \leq \beta \leq \frac{\pi}{2} \quad (7)$$

$$0 \leq W \leq \frac{W_{rover}}{2} \quad (8)$$

$$h_g \leq h \leq \frac{3}{2}r + h_g \quad (9)$$

where  $\omega_{max}$  is the maximum speed of the drive motor,  $W_{rover}$  is the total weight of the rover, and  $h_g$  is the height of the grousers.

Every time the optimization uses  $x$  to check for feasibility of the constraints or evaluate the cost function, it must first go from the decision variables to the full state of the rover by using the terramechanics model to calculate the forces and moments on the rover. First, we find the velocity of each wheel in its own reference frame,  $v_x$  and  $v_y$ , from the rover velocity,  $v_{x_{rover}}$  and  $v_{y_{rover}}$ , so that for each wheel  $i \in \{FR, FL, RR, RL\}$  we have

$$v_{x_i} = v_{x_{rover}} \cos \beta_i + v_{y_{rover}} \sin \beta_i \quad (10)$$

$$v_{y_i} = v_{y_{rover}} \cos \beta_i - v_{x_{rover}} \sin \beta_i \quad (11)$$

The terramechanics model requires a maximum stress angle  $\theta_m$  for wheels in slip, which can be calculated based on  $\beta$  and empirical constants  $a_0$  and  $a_1$ . For wheels in skid, we find the maximum stress angle by finding the maximum stress angle for a slip ratio of zero at our given slip angle  $\beta$ . To do this, we temporarily set the slip ratio to zero and perform a binary search over wheel sinkage until the applied load matches the computed  $F_z$  for the wheel, from which we

can calculate the entrance angle  $\theta_f$  and find  $\theta_m$  for skid, as in [7].

### C. Constraints for driving straight

Our constraints derive from the requirement that the rover is at steady-state for a feasible driving strategy. With our state-dependent soil geometry values set, we compute the forces ( $F_x, F_y,$  and  $F_z$ ) and moments ( $M_x, M_y,$  and  $M_z$ ) on each wheel as in [7]. To compute these, we set the sinkages to the values given in  $x$  and compute the resulting forces. Next, we compute the net forces and moments on the entire rover with the following equations, where individual forces and moments from each wheel have already been transformed into the rover's frame of reference. The free body diagram for each axis can be seen in Fig. 19. The angles of the right and left rockers,  $\phi_R$  and  $\phi_L$ , relative to horizontal, are determined by the sinkages of each wheel and the rover rocker length  $l_r$ , and given by

$$\phi_R = \sin^{-1} \left( \frac{h_{RR} - h_{FR}}{l_r} \right) \quad (12)$$

$$\phi_L = \sin^{-1} \left( \frac{h_{RL} - h_{FL}}{l_r} \right) \quad (13)$$

For a rocker pivot height from wheel center  $h_r$  and width between axles  $w_r$ , we have

$$\Sigma F_x = F_{x_{FR}} + F_{x_{FL}} + F_{x_{RR}} + F_{x_{RL}} = 0 \quad (14)$$

$$\Sigma F_y = F_{y_{FR}} + F_{y_{FL}} + F_{y_{RR}} + F_{y_{RL}} = 0 \quad (15)$$

$$\Sigma F_z = F_{z_{FR}} + F_{z_{FL}} + F_{z_{RR}} + F_{z_{RL}} - W_{rover} = 0 \quad (16)$$

$$\begin{aligned} \Sigma M_x = & M_{x_{FR}} + M_{x_{FL}} + M_{x_{RR}} + M_{x_{RL}} \\ & + (F_{y_{FR}} + F_{y_{RR}}) \cos \phi_R h_r \\ & + (F_{y_{FR}} - F_{y_{RR}}) \sin \phi_R l_r / 2 \\ & + (F_{y_{FL}} + F_{y_{RL}}) \cos \phi_L h_r \\ & + (F_{y_{FL}} - F_{y_{RL}}) \sin \phi_L l_r / 2 \\ & + (-F_{z_{FR}} + F_{z_{FL}} - F_{z_{RR}} + F_{z_{RL}}) w_r / 2 = 0 \end{aligned} \quad (17)$$

$$\begin{aligned} \Sigma M_{y_R} = & (-F_{x_{FR}} - F_{x_{RR}}) \cos \phi_R h_r \\ & + (-F_{x_{FR}} + F_{x_{RR}}) \sin \phi_R l_r / 2 \\ & + (-F_{z_{FR}} + F_{z_{RR}}) \cos \phi_R l_r / 2 \\ & + (F_{z_{FR}} + F_{z_{RR}}) \sin \phi_R h_r = 0 \end{aligned} \quad (18)$$

$$\begin{aligned} \Sigma M_{y_L} = & (-F_{x_{FL}} - F_{x_{RL}}) \cos \phi_L h_r \\ & + (-F_{x_{FL}} + F_{x_{RL}}) \sin \phi_L l_r / 2 \\ & + (-F_{z_{FL}} + F_{z_{RL}}) \cos \phi_L l_r / 2 \\ & + (F_{z_{FL}} + F_{z_{RL}}) \sin \phi_L h_r = 0 \end{aligned} \quad (19)$$

$$\begin{aligned} \Sigma M_z = & M_{z_{FR}} + M_{z_{FL}} + M_{z_{RR}} + M_{z_{RL}} \\ & + (F_{x_{FR}} - F_{x_{FL}} + F_{x_{RR}} - F_{x_{RL}}) w_r / 2 \\ & + (F_{y_{FR}} - F_{y_{RR}}) l_r / 2 \cos \phi_R \\ & + (F_{y_{FL}} - F_{y_{RL}}) l_r / 2 \cos \phi_L = 0 \end{aligned} \quad (20)$$




FIGURE 19: Free body diagram of the mini rover while driving. All labeled forces and moments are in the rover frame. Note that the moment balances on the left and right rockers are performed separately due to the passive pivot between them.

Note that as the rockers are free to pivot relative to each other, the moments about the  $y$ -axis are treated separately for each rocker, designated  $\Sigma M_{y_R}$  for the right side rocker and  $\Sigma M_{y_L}$  for the left side rocker. In the above equations we neglect the roll of the rover for simplicity as it is sufficiently small. At steady state, the sum of all of the forces and moments should be zero, so by setting (14)–(20) equal to zero we have a set of nonlinear constraints.

Because we are allowing the optimization to select both the individual wheel sinkages and load distribution on the wheels, we also need a constraint to ensure that our loads and sinkages match up for each wheel. To accomplish that, we add the constraints (where  $F_z$  is computed from the sinkage, slip angle, and rotational speed through the terramechanics model)

$$F_{z_{FR}} - W_{FR} = 0 \quad (21)$$

$$F_{z_{FL}} - W_{FL} = 0 \quad (22)$$

$$F_{z_{RR}} - W_{RR} = 0 \quad (23)$$

$$F_{z_{RL}} - W_{RL} = 0 \quad (24)$$

If we take these equations and substitute them back into (16), we can restate it as

$$\Sigma F_z = W_{FR} + W_{FL} + W_{RR} + W_{RL} - W_{rover} = 0 \quad (25)$$

This new constraint is a linear function of  $x$ , which is easier for the solver to handle than (16), which relies on the output of the terramechanics model and is therefore highly nonlinear.

We now have all of the constraints necessary to ensure our rover's motion is fully defined and at a steady-state. With this set of constraints, we can generate feasible strategies for a rover to drive straight forward.

### D. Constraints for steering while driving

While rovers such as Curiosity and Perseverance use feed-forward driving strategies and only close the loop on heading with visual odometry every meter [27], many terrestrial

wheeled robots use closed-loop controls to maintain heading. We can generate steering controllers for following a course using the same optimization technique by changing only our constraints. To control for the rover's heading, or yaw angle, we simply offset the moment balance about the vertical  $z$ -axis by a small amount  $M_{z_{offset}}$ , replacing (20) with

$$\begin{aligned} \Sigma M_z = & M_{z_{FR}} + M_{z_{FL}} + M_{z_{RR}} + M_{z_{RL}} \\ & + (F_{x_{FR}} - F_{x_{FL}} + F_{x_{RR}} - F_{x_{RL}})w_r/2 \\ & + (F_{y_{FR}} + F_{y_{FL}} - F_{y_{RR}} - F_{y_{RL}})l_r/2 \\ & - M_{z_{offset}} = 0 \end{aligned} \quad (26)$$

We use the same strategy to generate a steering controller for the rover's side slip by adding a small force offset  $F_{y_{offset}}$  along the  $y$ -axis, replacing (15) with

$$\Sigma F_y = F_{y_{FR}} + F_{y_{FL}} + F_{y_{RR}} + F_{y_{RL}} - F_{y_{offset}} = 0 \quad (27)$$

As we may later not be able to assume symmetry when one wheel is either trenching or degraded, we generate the left and right steering controls for both yaw and  $y$  position separately by solving the optimization problem with both positive and negative offsets for each value. We then take that solution and subtract the original driving primitive to get a steering controller which we rescale with PID control on the rover's yaw or  $y$  error and add to the feedforward driving primitive.

The steering controllers generated here were used as a default steering approach for trenching objectives and degraded driving conditions. The resultant steering controller in yaw has the rover speed up the wheels on one side and slow them down on the other. The controller for  $y$  turns all steering motors in the desired direction with a small speed-up. In both cases the optimization was able to find controllers that match intuition. These methods are the default feedback steering control used in experiments in Section V.

Empirical Terramechanics Parameters			
Parameter	Description	Smooth	Grousered
$k$	Sinkage modulus	$8000 \text{ kN/m}^{n+2}$	$8000 \text{ kN/m}^{n+2}$
$\phi$	Angle of repose	$29^\circ$	$29^\circ$
$\rho$	Density	$13.03 \text{ kN/m}^3$	$13.03 \text{ kN/m}^3$
$c$	Cohesion	1.0 kPa	1.0 kPa
$K$	Shear modulus	0.021	0.021
$n_0$	Sinkage exponent	1.46	1.46
$n_1$	Sinkage exponent (slip)	0.01	0.55
$n_2$	Sinkage exponent (skid)	0.01	0.74

TABLE 7: Soil parameters for terramechanics tests for both smooth and grousered wheels, as defined in [7].  $k$ ,  $\phi$ , and  $\rho$  were measured, all other values were tuned.

### E. Objective function

The next step is to select a suitable objective function. While any value for the cost will allow us to find a feasible driving strategy, it will not necessarily select a sensible one. Here, we minimize the actuator effort required to drive, which also serves as a rough proxy for minimizing the power. By taking the square of the torque on each wheel, we have a cost function that minimizes the overall power needed to drive. The objective function is therefore

$$c(x) = \bar{M}_{yFR}^2 + \bar{M}_{yFL}^2 + \bar{M}_{yRR}^2 + \bar{M}_{yRL}^2 \quad (28)$$

where  $\bar{M}$  is the wheel torque, which is the moment about the wheel's  $y$  axis. In Section IV.I, we discuss rescaling the objective function by a scalar to improve convergence.

### F. Complete optimization problem

Putting this all together into an optimization problem,

$$\min \text{ Objective function, } c(x), \quad (28) \quad (29)$$

$$\text{w.r.t. Decision variables: } x, \quad (5)$$

$$\text{s.t. Bounds: } x_{LB} \leq x \leq x_{UB}, \quad (6) - (9)$$

$$\text{Force Balance: } \Sigma F_x = 0, \Sigma F_z = 0, \quad (14), (25)$$

$$\Sigma F_y = 0 \quad (15) \text{ or } (27),$$

$$\text{Moment balance: } \Sigma M_x = 0, \Sigma M_{yR} = 0,$$

$$\Sigma M_{yL} = 0, \Sigma M_z = 0, \quad (17) - (20) \text{ or } (26)$$

$$\text{Vertical load balance: } F_z - W_z = 0, \quad (21) - (24)$$

We now have the tools needed to generate a feedforward driving strategy and feedback controllers for both the rover's yaw and  $y$  position. We solve these optimization problems with the rover parameters and force/moment steering offsets listed in Table 8 and with the soil parameters listed in Table 7, except where otherwise noted. The soil parameters  $k$ ,  $\phi$ , and  $\rho$  were directly measured, while other values were tuned. Note that while cohesion ( $c$ ) is often approximated as zero, the measured cohesion of Martian simulants such as the one used here can vary from 0.2 kPa to 1.4 kPa [44].

Rover parameters		
$r$	Radius of wheel	5.8 cm
$b$	Width of rover wheel	5 cm
$h_g$	Grouser height	0.5 cm
$l_r$	Rover rocker length	32 cm
$h_r$	Height of rocker pivot from wheel center	11.2 cm
$w_r$	Width of rover between wheel centers	33.1 cm
$W_{rover}$	Weight of rover	34.3 N
$v$	Rover travel velocity	3 cm/s

Terramechanics parameters		
$a_c$	Constant coefficient within $a_0$	0.24
$a_\beta$	$\beta$ -dependent coefficient within $a_0$	$1/129.1^\circ$
$a_s$	Constant coefficient within $a_1$	0.69
$a_{s\beta}$	$\beta$ -dependent coefficient within $a_1$	$-1/133.5^\circ$
$\zeta$	Grouser transport volume fraction	.2

Offset forces		
$M_{zoffset}$	Yaw moment offset for steering control	$\pm 0.2 \text{ Nm}$
$F_{yoffset}$	$y$ force offset for steering control	$\pm 1.0 \text{ N}$

TABLE 8: Rover and soil parameters for nominal driving optimization.

### G. Control for wheel-based trenching

Using the trench shape model developed in [6], we can easily factor trench geometry into both the constraints and the objective function. For example, we could specify that we want a trench with a flat bottom by setting the width of the trench base to be nonzero as an optimization constraint or to maximize the width of the trench bottom by including it in the objective function. For these experiments, we choose to make a trench with a "clean" side, where the soil is entirely scraped away on one side of the wheel and piled onto the opposite side, such that no soil reflows in on the clean side. To do this, we specify an additional constraint that the height of the soil pile to the right of the wheel,  $h_{R2}$ , is zero for the wheel being used to trench. Here we select the rear right wheel to dig, and so we add the constraint

$$h_{R2,RR} = 0 \quad (30)$$

We can also add inequality constraints, such as requiring that the trench depth  $d$  be at least a minimum amount. In these experiments, we require the trench to be at least as deep as half of the wheel's width  $b$  by specifying

$$b/2 - d \leq 0 \quad (31)$$

Here, we use the same objective function as before, (28), seeking to minimize the energy used in digging while we drive. The optimization problem, (29), was solved with these additional constraints and the values outlined in Table 8, with the exception that  $\zeta$  was set to 1. Both a feedforward driving primitive as well as optimized yaw and  $y$  steering controllers were computed. In addition, for testing we compared with the steering controls generated for nominal driving.

**H. Control for driving with degraded mobility**

We also use the same techniques to enable rovers to drive with degraded mobility by incorporating the mobility failure(s) as additional optimization constraints. For example, a wheel jammed by a rock would be represented by setting that wheel’s rotational speed to zero, or a motor that has lost power could be represented by setting the wheel torque to zero. Here, we detail how to represent two failure modes and generate driving strategies for them, which are later tested in Section V.

First, we look at the case of a single wheel that needs to be run at a lower speed than the desired travel velocity for the full rover. This could be due to reduced available power or a need to reduce strain on the wheel while it warms up or recirculates lubricant, such as with drive motor degradation experienced by Spirit and Opportunity [21]. The second degraded mobility case is that of a rover with a steering actuator stuck at a fixed angle, such as what happened to Opportunity’s front right wheel [21]. For both of these scenarios, we keep our objective function set to minimize actuator effort and treat the rear right wheel as the affected actuator.

For the speed-limited wheel, we add an additional inequality constraint that the rim speed of the wheel does not exceed the assigned speed limit,  $v_{lim}$ . As we are looking at the rear right wheel, this gives

$$\omega_{RR} r_s - v_{lim} \leq 0 \tag{32}$$

We re-solve the optimization problem with this added constraint, using the values specified in Table 8, and additionally solve it with  $y$  force and yaw moment offsets to compute the optimal steering controls that do not require higher speeds on the affected wheel.

A rover with a steering actuator stuck at a constant angle  $\theta_{fail}$  needs an additional equality constraint of

$$\beta_{RR} - \theta_{fail} = 0 \tag{33}$$

Again, we solve the optimization problem with this additional constraint, and then separately solve it with the force and moment offsets to get steering controllers that do not require different wheel angles on the affected wheel.

While not explored here, we could additionally set constraints based on the forces and torques imposed on the rover’s structure due to driving, such that we could protect weakened mechanisms from undue stress or ensure that generated driving strategies do not decrease the lifespan of an otherwise healthy rover.

**I. Challenges with solving the optimization problem**

A number of difficulties arise when complex models such as those used in terramechanics are treated with traditional optimization methods. There are several aspects of the terramechanics model used here which make embedding it in an optimization problem difficult:

- Numerical integration

- Mode transitions
- Internal iterative loops
- Decision variable and constraint with different scaling
- A small feasible space

In this section, we break down how each of these aspects impacts the implementation of the terramechanics-based optimization in (29) and discuss how we solved these problems. The solver used for this paper was MATLAB’s `fmincon` with the `interior-point` algorithm in “Feasibility Mode”.

*a: Numerical integration*

As the full terramechanics models presented in [6], [7] are not closed-form, the calculation of  $F_x, F_y, F_z, M_x, M_y,$  and  $M_z$  for each wheel must be numerically integrated to compute the forces on a wheel at a given sinkage and evaluate the constraints and objective function. Additionally, this means that we do not have an analytical gradient for either the constraints or objective function, and must numerically compute the gradient of both for each step of the optimization process using finite differences. As numerical integration is a discrete process, this necessarily makes the gradient of an objective function or constraint function non-smooth. For finite differences to perform properly on models with numerical integration, the step size for differentiation must be large enough that nonsmoothness due to numerical integration does not impact the local shape of the gradient. Furthermore, `fmincon` automatically rescales the finite difference step size in each component based on a typical value of  $x$ . Here, we increase the step size from the typical value of  $10^{-10}$  to  $10^{-6}$  so that the finite differentiation is large enough to not get caught by nonsmoothness from numerical integration but small enough to capture local gradient shape.

*b: Mode transitions*

The terramechanics model in [7] covers both slip and skid but uses different soil shearing behavior in those two regimes. Additionally, the soil flow model in [6] has discrete shape transitions dependent on the wheel’s speed and slip angle (such as the transition in how the soil splits around the wheel as its slip angle crosses  $10^\circ$ , as explained in [6]). As solvers find local optima, discrete transitions in the gradient can lead to the solver getting “stuck” in an infeasible or sub-optimal region. By selecting initial conditions which are close to a feasible solution, we remove the need for the solver to cross regions of the terramechanics model with discontinuities to find a feasible solution, which makes it easier to solve the optimization problem. Non-strategically selected initial conditions, such as starting from  $x = 0$ , do not result in the solver finding a feasible solution.

We employed two strategies to select a good initial condition. The first strategy is selecting a “reasonable” initial point based on intuition and experience. This method was sufficient for the generation of the feedforward trenching strategy and stuck steer motor combination. For the trenching

54 pt  
0.75 in  
19.1 mm

54 pt  
0.75 in  
19.1 mm

strategy, we began with initial conditions that we knew would produce a good-sized trench with the  $y$  forces and  $z$  moments balanced due to symmetry, with the front wheels pointed forward and driving at 4.5 cm/s, and the rear wheels pointed symmetrically outward at  $\pm 60^\circ$  and driving at 12.5 cm/s. For the stuck steer motor, we used initial conditions with the stuck wheel set at its fixed angle of  $50^\circ$  pointed outwards and all four wheel drive speeds set to 3.5 cm/s, which was the wheel drive speed obtained from solving the straight driving optimization problem. We set the wheel loads to be equally distributed among all four wheels and the wheel sinkages to be  $r_s/4$  for all optimization problems solved, as we expect the force balance on the wheels to be relatively even and  $r_s/4$  is close to the sinkage observed on the mini rover when driving in the testbed. With these initial conditions, we were able to solve both the stuck steering and trenching problems.

The second method of improving the solution likelihood is to use initial conditions that are the output of a similar optimization problem. This method was employed to solve the speed-limited wheel problem, which was much harder to find a feasible solution for than the other two problems. This was done by starting with a speed limit on the affected wheel that was closer to the rover travel velocity of 3 cm/s and then iteratively using the output as the initial conditions to a new optimization problem with a slightly lower speed limit. By this method, the affected wheel's speed was slowly reduced to 1.5 cm/s. It was unable to be reduced further, as the mini rover is not able to both balance the yaw moments and produce enough traction at lower speeds. This method was also employed to speed up the solution of the yaw and  $y$  position steering problems for each driving mode, by taking the output of the feedforward driving problem as the initial condition to each of the steering problems.

*c: Internal iterative loops*

In [7], the sinkage for a wheel with a given load  $W$  is found using a binary search over possible sinkages until the vertical load  $F_z$  is matched within a given numerical tolerance  $\epsilon$ . When included in an optimization, the numerical accuracy  $\epsilon$  set for the binary search creates steps in the sinkage, which makes the constraint and gradient functions nonsmooth, as shown in Fig. 20. Additionally, having an internal loop slows down the solver. To avoid this, here we allow the optimization to perform sinkage finding by including the wheel sinkage as a decision variable along with the vertical load and incorporating the force balance on each wheel as a constraint, (21)–(24).

However, there is a single internal iterative loop that cannot be eliminated by this method; the definition of  $\theta_{m0}$  (and  $\theta_{r0}$ , when the trench model is not used to determine soil exit angle) is dependent on the sinkage of the wheel at zero slip,  $\theta_m|_{s=0}$ . As the sinkage at zero slip is dependent on the applied wheel load and the wheel's slip angle,  $\theta_{m0}$  must be evaluated for each iteration of the optimization.




FIGURE 20: Discretization visible in one of the constraint functions due to the tolerance of binary search for finding wheel sinkage. Values shown here are for constraint (18) looking at a 2d slice of possible values of  $x$ , as the full dimension of decision variables is too high to visualize.

As this internal loop is only used to determine a portion of the wheel-soil contact geometry, the overall value of the objective and constraint functions are less sensitive to the numerical accuracy of this sinkage-finding loop than they would be for an iterative loop to find sinkage of the entire wheel.

*d: Decision variable and constraint scaling*

For the four-wheeled rover used in this work, we have sixteen variables – the rotational speed, steering angle, vertical load, and sinkage of each individual wheel. Each of these values has different units and they do not scale proportionally to each other with rover size. Without performing any rescaling, the optimization routine is less sensitive to decision variables that are typically smaller, as the step sizes in those directions will be closer to zero. Without rescaling, the optimization would stop based on the small step size without fully optimizing these variables, which left it unable to find feasible solutions.

*fmincon* allows the user to specify typical values of the decision variables to automatically rescale finite difference step sizes, which we used here. For our rover, we have the following typical values for the elements of  $x$ :  $\omega = v/r_s = 0.4762$  radians/second;  $\beta = 10^\circ = 0.1745$  radians;  $W = W_{rover}/4 = 8.6$  N; and  $h = r_s/4 = 0.0158$  m. The range of these values also span multiple orders of magnitude, and so for the trenching problem and stuck steer motor problem we rescaled the  $W$  and  $h$  values by a factor of 100 so their range would be on the same scale as the other variables.

While *fmincon* accepts typical values of  $x$  to rescale the finite difference step sizes, it does not use those values to scale any other solver tolerances. This means that the constraint tolerance, which is a single scalar, is applied to all constraints equally. As our constraints contain forces, moments, and sometimes wheel speeds or angles, we again have a wide variation in scale; a constraint tolerance of  $10^{-6}$  on the vertical force balance for a wheel with a 8.6 N load represents allowable error of 0.01%, while the same tolerance applied to a speed limit of 1.5 cm/s = 0.2381 rad/s

would represent an allowable error of just 0.0004%. Therefore, we rescale the constraints by individually multiplying them by scalars within the constraint function so that they allow the same relative error, or so that we can change the tolerance for various values to reflect how important each equality constraint is. Here, we were able to solve most of the optimization problems with uniform constraint tolerances of  $10^{-6}$ , but for the trench problem we set the constraint tolerance on trench dimensions to be within 1 mm, to reflect that it does not need to be precisely accurate, unlike the force and moment balances on the rover.

*e: A small feasible space*

As we are attempting to find steady-state conditions on a complex mechanism, there are many equality constraints that must be concurrently satisfied. As such, the feasible space is small, and even “feasibility” modes built into solvers do not always find feasible solutions that exist.

As noted above, using initial conditions close to a feasible solution helps the solver but does not guarantee finding a solution. In problems with narrow feasibility regions where the objective function and constraint function have opposing gradients, the solver can find a feasible solution and then later lose feasibility while searching for an optimal solution. *fmincon* encountered this issue for the trench problem, as it struggled to maintain a clean-sided edge and balance the forces at the same time. It would find feasible solutions only to lose feasibility later in the solving process and never regain it.

To solve this issue, we needed to make *fmincon* place a greater weight on feasibility than on optimality. By decreasing the scale of the objective function we force the solver to weight constraints more heavily, leading to an increased chance of finding a feasible solution at the cost of it being less likely that the solution is optimal. For the trench problem, we weighted the objective function by  $10^{-2}$ , which allowed it to find and maintain feasibility.

**V. Optimization test results**

In this section we evaluate the optimization approach to generating driving strategies for different mobility objectives and challenges. Section V.A evaluates the generated feedforward and feedback control strategies on a hardware platform, while Section V.B uses the optimization to gain insight into different mobility failure modalities in simulation.

**A. Laboratory testing of generated driving strategies**

Testing of the optimization-based control strategies was performed with the mini rover shown in Fig. 18 due to constraints on platform and testing facility access. The mini rover was constructed to have similar suspension kinematics and ground pressure to KREX-2, the heavier of the two large rovers employed in the demonstrations presented in Section III. Tests were conducted in Carnegie Mellon University’s soft soil testbed, the same testbed used in [6]. Both the soil

and grousers wheels used in this experiment are the same as those used in the terramechanics model validation performed in [7]. Key geometry parameters for the wheel can be found in Table 8, and key soil parameters can be found in Table 7. The mini rover has four wheels with independent steering and drive actuation and a passive pivot between the two rockers. The drive wheels are actuated with Dynamixel XM-430 servo motors, and steering is done with Dynamixel AX-18A motors. Live tracking of rover position and orientation is done by a RealSense® T265 tracking camera, which is mounted to the rover’s right rocker. The T265 uses a combination of feature tracking through the depth camera and measurements from an internal IMU to measure 6-DOF position and orientation. When enabled, feedback control of the rover’s heading and horizontal position was implemented as PID control on the measurements from the T265, with each component scaled with its own PID controller and the rover’s wheel speeds and angles set as the superposition of the two components and the nominal feedforward motion. The rover is tethered for power and communication. Rover communications and controls operate at 5Hz, with speed and position control of individual motors handled within each Dynamixel at a much higher rate. An offboard 2d LIDAR scanner was used to measure trench profiles as in [6].

For each driving scenario tested, the sand surface is prepared with an automated loosening and smoothing mechanism, and the rover is driven straight forward at 3 cm/s for five seconds before beginning the driving scenario. Rover position and orientation are initialized relative to the rover’s starting location.

1) Trenching

The trenching strategy generated in Section IV.G was run on the mini rover using three different control methods: feedforward only, nominal feedback steering controls, and optimized steering controls. The nominal steering controllers were generated in Section V.D, and have all wheels turn in one direction with a slight increase in speed to correct for  $y$  error or have the wheels speed up on one side and slow down on the other to correct  $M_y$ , subject to any speed or steering constraints imposed by the scenario. The optimized steering controllers were developed via the same method, but were optimized under the constraints imposed by the test case, such as a stuck steering actuator or a slowed wheel.

The tracks left by each trial can be seen in Fig. 21, with all three cases driving straight and digging a clean-sided trench with the rear right wheel. The error on the rover’s heading (yaw) and horizontal ( $y$ ) position can be seen in Fig. 22. The rover drove very straight while trenching in all three control schemes, with less than  $3^\circ$  deviation in yaw and less than 5 cm of side slip over a 1.5m travel distance. Using feedback controls (blue and green) marginally improved  $y$  error, though the rover did not need feedback control to stay on course for this short trial, as seen from the low errors on

54 pt  
0.75 in  
19.1 mm

54 pt  
0.75 in  
19.1 mm



FIGURE 21: Mini rover digging a trench with several driving strategies. From left to right: feedforward trenching primitive only, feedforward with nominal feedback control, feedforward with optimal feedback control.

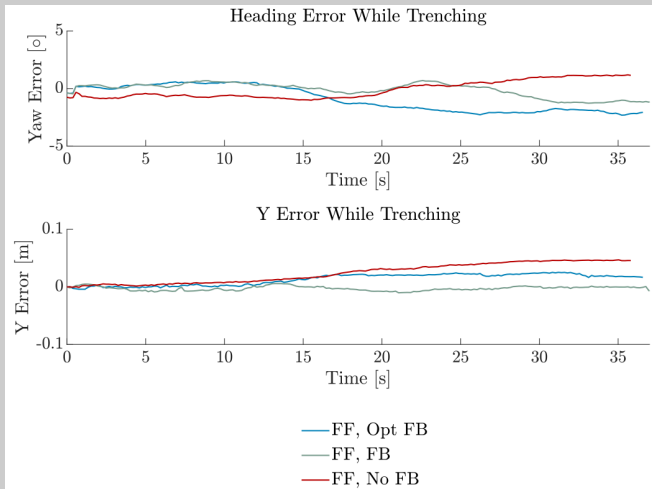


FIGURE 22: Comparison of yaw and  $y$  position errors for open and closed loop trajectory following while using the rear right wheel to dig a clean-sided trench. All use a feedforward (FF) driving strategy generated by optimization. FB: feedback control enabled, with default steering. Opt FB: “optimal” steering relative to feedforward strategy. All three methods are able to maintain heading while digging a trench.

the feedforward-only trial (red). This shows that trenching can be safely achieved without feedback control by using model-generated feedforward trenching primitives, which enables wheel-based trenching on rovers with infrequent visual odometry such as Curiosity and Perseverance.

A detailed image of the trench dug by the wheel (blue) along with its predicted shape (red) can be seen in Fig. 23. The optimization was set to give a trench depth of at least half the wheel’s width, or 2.5 cm, and a clean edge on the right side of the rear right wheel. The qualitative shape of the predicted trench matches that of the observed trench, with a clean-sided edge on the right side. However, the overall depth of the trench was off, with the model predicting a depth of 2.6 cm and an observed depth of 1.5 cm. This discrepancy is likely due to an under-prediction of the wheel’s sinkage, possibly due to some combination of small errors in the terramechanics model at high slip angles, a difference in predicted applied load from the actual load, or from the terramechanics model in [7] being tuned on a slightly higher rover weight than that of the mini rover.

2) Wheel failure: slowed wheel

For the case of the speed-limited wheel, the rover’s rear right wheel was set to a maximum rim speed of 1.5 cm/s, which is half of the rover’s target travel speed of 3.0 cm/s. The tracks left by the rover with no compensation, feedforward compensation only, and feedforward with feedback compensation can be seen in Fig. 24, while the yaw and  $y$  error are plotted for these and additional feedback strategies in Fig. 25.

As seen in Fig. 25, the rover immediately began to pull right with no compensation (yellow), veering off the prepared soil track within 1 m of driving. With all combinations of feedforward and feedback compensation (blue, green, red, and dark blue), the rover was able to maintain heading within





FIGURE 23: Photograph of trench dug by the mini rover while driving with closed-loop steering controls (top), alongside LIDAR scan of the same trench (blue, bottom) overlaid with its predicted shape (red, bottom). The rover successfully produced a clean-sided trench, but with a lower depth than expected.

6° and horizontal position within 3 cm over a 1.5 m run. For this failure scenario, feedback control alone was worst at maintaining a heading; the error in yaw when using only feedback control (dark blue, Fig. 25) was about 6°, while with a feedforward compensation strategy the yaw error stayed below 3°, with or without feedback control. There was less difference in control strategies when looking at the rover’s side slip ( $y$  position), though the error was slightly higher with only a feedforward term (red). The difference between the error resulting from optimized and non-optimized feedback terms was quite small, so while optimized feedback terms can be used for minimizing actuator effort they are not necessary for functional performance in this scenario.

However, the rover did not perfectly maintain its target velocity; while both feedforward and feedback control methods were able to keep the rover moving faster than the skidding wheel’s speed of 1.5 cm/s with travel velocities of 2.1 to 2.4 cm/s, this is considerably slower than the target speed of 3 cm/s. The rover was not actively controlling for forward velocity, but the travel velocities observed here are slower than those predicted by the terramechanics model, which expected that the feedforward term would be able to maintain 3 cm/s.

From these trials, we can see that in the case of a speed-limited wheel, such as in the case of limited available motor power, not only can we find a feasible driving strategy for what would otherwise prevent the rover from following a course, feedforward-only driving strategies are sufficient.

This means that rovers with infrequently updated visual odometry such as Perseverance and Curiosity can recover from a mobility failure in the form of a speed-limited wheel by using feedforward-only driving primitives, provided the soil is sufficiently uniform.

### 3) Wheel failure: stuck steering

As seen in the tracks in Fig. 26 and the error magnitudes in Fig. 27, failure of a steering actuator at 50° has less of an impact on mobility than a driving wheel with limited speed. Without compensation (yellow), the rover drifted 13 cm to the right over a 1.6 m run, with a roughly linear increase in yaw error ending at 3°. While feedforward compensation (red) greatly reduced the  $y$  error from 13 cm to 4 cm, it increased the yaw error to 15° by the end of the trajectory. For this scenario, feedback control with the default steering controller was able to maintain the rover’s heading without using the feedforward term (dark blue). When the feedforward compensated driving strategy was used, the default steering controller showed comparable yaw error to the uncompensated driving (green), while the optimal steering controllers generated for the feedforward driving strategy (blue) performed the best.

Overall, we see that without any compensation a stuck steering actuator results in considerable side slip on the rover’s trajectory and that while a feedforward-only strategy is able to mitigate this side slip it may result in an increased error in yaw. The use of feedback control to maintain rover heading improves the outcome, though using a feedforward-only strategy may be sufficient for short drive distances. Similarly to the slowed wheel scenario, with a stuck steering actuator we do not see path-following improvements from the optimized feedback terms relative to the naive feedback approach, suggesting that optimization of the feedback terms is not required for path following though it can be used to minimize actuator effort.

### B. Mobility system failure insights from optimization

In Section III.C we showed that on flat ground a rover can compensate for the impact of failed mobility actuators through novel control strategies; here, we investigate the impact of the same failures on the rover’s ability to ascend slopes. We simulate slope climbing by changing the constraints on the rover’s force balance, similar to how the drawbar pull tests in Section III.B.3 experimentally simulated slope by applying a load resisting the rover’s motion.

To simulate various slope angles, the optimization problem described in Section IV was solved with two of the force balance constraints, (14) and (25), changed to account for the change in the direction of gravity relative to the rover’s frame of reference,

$$\sum F_x = F_{x_{FR}} + F_{x_{FL}} + F_{x_{RR}} + F_{x_{RL}} - W_{rover} \sin \alpha = 0 \quad (34)$$




FIGURE 24: Mini rover driving with a speed-limited wheel. From left to right: driving without compensation, feedforward with no feedback control, feedforward with optimal feedback control.

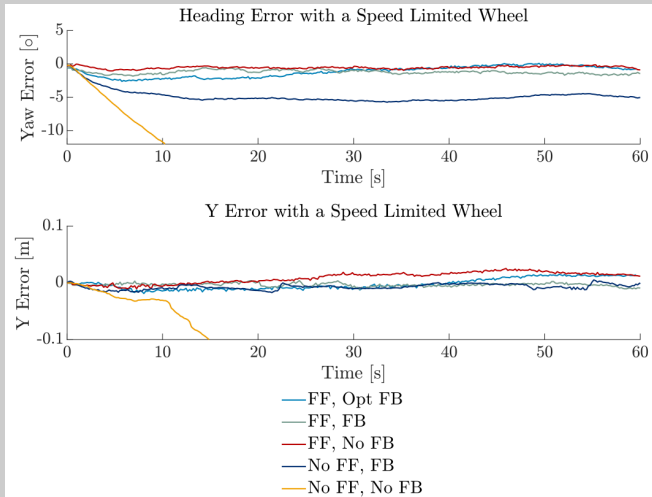


FIGURE 25: Comparison of yaw and  $y$  position errors for open and closed loop trajectory following with the rear right wheel limited to 15 cm/s. FF: uses a feedforward driving strategy generated by optimization. FB: feedback control enabled, with default steering unless noted and “optimal” steering relative to feedforward strategy for Opt FB. Note that the rover is able to drive straight with any form of compensation; feedforward compensation alone is sufficient.

$$\Sigma F_z = W_{FR} + W_{FL} + W_{RR} + W_{RL} - W_{rover} \cos \alpha = 0 \quad (35)$$

for a given slope angle  $\alpha$ .

With these restructured constraints, we re-solve the optimization problem for different modes of impacted or nominal mobility and find the driving strategies for successfully ascending slopes of a given angle. In addition, we can gain insight into three areas of interest: the maximum slope angle a rover can possibly ascend, the amount of slip induced on a given slope, and the relative power needed to climb a slope. With this information, it is possible to set safe operational limits on a rover driving with an impacted mobility system without the need to run a large number of drawbar pull tests such as those conducted in Section III.B.1.

The optimization problem was solved repeatedly with slope constraints for slope angles of  $0^\circ$  to  $25^\circ$  for nominal driving, driving with a speed-limited wheel, and driving with a stuck steering motor. For the speed-limited wheel and stuck steering actuator, the conditions are the same as those explored in Section III.C, with the speed-limited wheel driving at half the rover’s travel speed and the steer motor stuck at  $50^\circ$ . Fig. 28 shows both the simulated slip induced while driving as well as the predicted mechanical power needed to drive at each slope angle for each driving mode, with mechanical driving power given as the sum of the products of the wheel torque  $M_y$  and angular velocity  $\omega$  for each wheel,  $P = \Sigma M_y \omega$ .

Two values are given for the slip ratio in Fig. 28: the maximum slip on any wheel and the average slip across all wheels not experiencing failure. Wheels experiencing failure are excluded from the slip average because they are often in skid, which would make the average slip seem lower; looking at the slip on only unaffected wheels is therefore



FIGURE 26: Mini rover driving with a stuck steering motor. From left to right: driving without compensation, feedforward with no feedback control, feedforward with optimal feedback control.

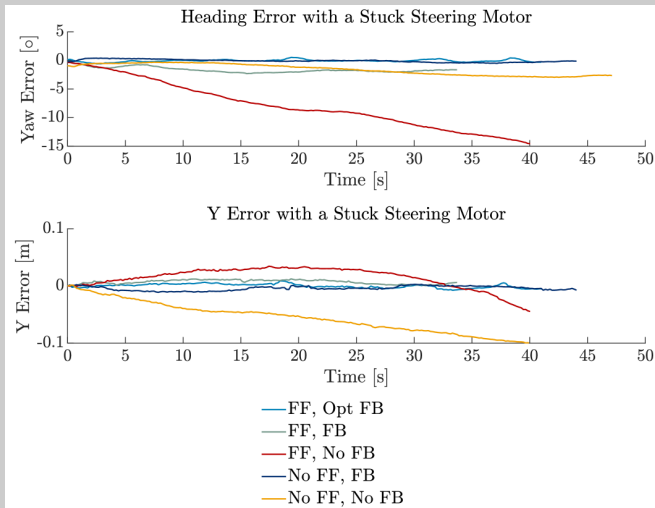


FIGURE 27: Comparison of yaw and  $y$  position errors for open and closed loop trajectory following with the rear right steering actuator stuck at  $-50^\circ$ . FF: uses a feedforward driving strategy generated by optimization. FB: feedback control enabled, with default steering unless noted and “optimal” steering relative to feedforward strategy for Opt FB.

more conservative when looking at induced slip on the full rover. We see that the mini rover is able to ascend slopes up to  $24^\circ$ ; beyond that, the optimization is unable to find a feasible solution. The difference in slip ratio between the average and maximum slip ratios is because the front and rear wheels were run at different speeds, which is necessary

to balance the moments about the  $y$ -axis on each rocker when moving uphill. The induced slip increases sharply for a wheel with a steer motor stuck at  $50^\circ$ , with average slip of 0.4 on flat ground compared to only slip of 0.15 on flat ground for nominal driving. The speed-limited case fares even worse, with 0.69 average slip and a maximum wheel slip of 0.76 on flat ground.

The generated slip values can be used to set operational limits for slopes; for example, if we limit average slip on the mini rover to 0.5 in order to minimize entrapment risk, we expect the rover to be able to safely ascend slopes up to  $12.5^\circ$  with normal driving conditions. We would limit the rover to slopes up to  $2.5^\circ$  with a stuck steering actuator, and would not be able to maintain drive speed with a speed-limited wheel – it would be necessary to slow the entire rover to reduce slip to an acceptable level.

Similarly, we can set terrain limits based on power consumption. The mini rover consumes about the same amount of power driving normally up a slope of  $17^\circ$ , driving up a slope of  $5^\circ$  with a stuck steering motor, and driving on flat ground with a speed-limited wheel. The nonlinear increase in power consumption with slope angle also suggests that it may be more efficient to climb slopes at a nonzero attack angle, such that the rover effectively covers a longer but shallower slope, as explored in [45].

## VI. Conclusions

In this paper, we showed that rovers can retain mobility while manipulating the soil and after experiencing failure, but that hand-tuned driving strategies are limited. First, we demonstrated multiple nonprehensile terrain manipulation

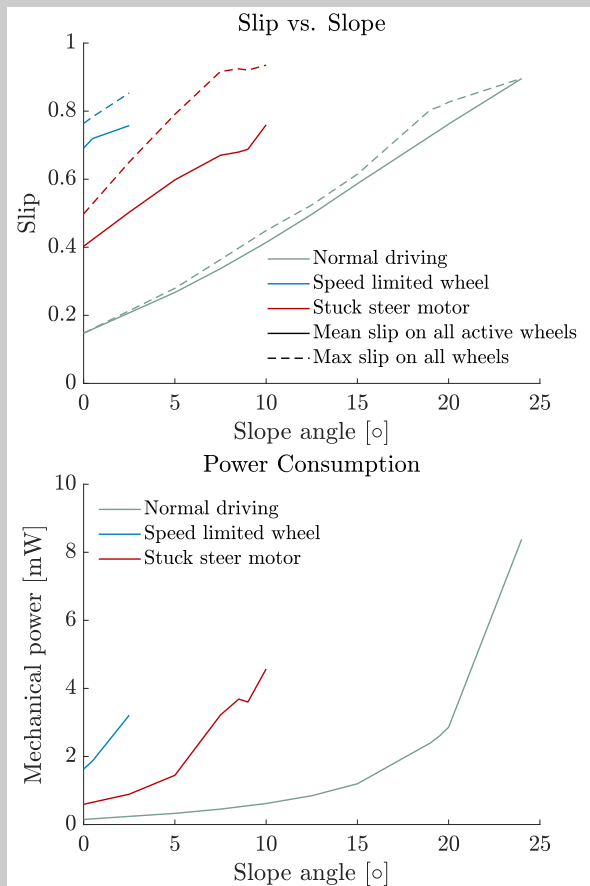


FIGURE 28: Plot of slip vs. equivalent slope angle (top) and power usage vs. slope angle for driving with failed actuators (bottom). Slip ratio is plotted as an average of the slip ratios on all non-impacted wheels as well as the maximum value of all wheels. Power is given by the sum of mechanical power (wheel torque  $\times$  angular velocity) from each individual wheel.

(NPTM) actions on a full-size rover in a Martian analog environment, showing that driving while trenching is both feasible and can move meaningful amounts of soil, with the exposure of subsurface halite deposits as a direct example of NPTM’s potential for scientific sampling. We also used NPTM to determine a key soil property by digging holes to measure the angle of repose and showed path clearing for a small robot in terrain with loose rocks and terrain with friable but steep terrain steps.

We then performed a series of tests on VIPER to measure the mobility loss due to drive system degradation and found that the loss of any actuator negatively impacted tractive ability, but the loss of a drive motor was the most catastrophic. We used hand-tuned driving strategies on the KREX-2 rover to compensate for the failure of a steering actuator and a drive motor, which resulted in straight driving trajectories but at very high slip, and took significant time testing on hardware in-situ to properly tune. While the quantity of these

tests was too small to draw quantitative conclusions on the magnitude of degradation due to various failure modes, they were sufficient to illustrate the impact on driving. Taken together, these field demonstrations show that off-nominal driving for intentional manipulation and degraded mobility are both feasible, but are practical only with a way to automatically develop new driving strategies.

We then created an optimization framework for automatically generating driving strategies for rovers under off-nominal conditions, enabling both nonprehensile terrain manipulation via trenching and compensation for mobility system degradation. We generated example driving strategies and demonstrated them on hardware. For trenching, we were able to automatically generate a motion primitive that excavates a trench while driving straight. The trenching primitive performed well, with the rover able to drive straight both with and without feedback control. Next, we considered a rover with a reduced drive motor speed, again automatically generating both feedforward and feedback compensation strategies. The rover, which veered off track without compensation, was able to drive straight with both feedforward-only control and feedback control. Finally, we considered a stuck steering motor and automatically generated driving primitives for it. The feedforward-only strategy reduced side slip error but had some yaw error compared to uncompensated driving, and all methods of feedback control were able to keep the rover driving straight. There was minimal difference in driving between naive feedback control terms and those output by the optimization, but the optimization could be used to minimize energy consumption from steering when working with a degraded mobility system.

We finally generated curves showing rover wheel slip and power consumption for increasing slope angles for both failure modes explored, showing that our new driving primitives have reduced locomotion capabilities compared to an unaffected rover.

While we explored only three types of constrained driving problems within the optimization framework, this method can be used for any type of off-nominal driving that can be represented through a combination of rover state and forces. For example, we could optimize to minimize load on a wheel that is losing tread, use constraints to limit the forces on a weakened steering joint, or remove a single wheel from the force balance entirely to represent a suspension like VIPER’s getting stuck in a raised position.

The optimization methodology and terramechanics model employed here were developed to run on a desktop workstation, and while it runs faster than an uplink/downlink cycle to mars (approximately 40 minutes), it has not been explicitly refined for computational efficiency or implementation on flight grade computing hardware. Linearization of terramechanics models could be employed to reduce computation time [46] to be executable onboard a planetary rover. This would enable rapid, in-situ adaptation of novel control strategies, and could be supplemented with more

54 pt  
 0.75 in  
 19.1 mm

54 pt  
 0.75 in  
 19.1 mm

accurate simulations run on an offboard machine and then uplinked. Additionally, in-situ estimations of terramechanics properties as in [26] and [3] can be used to adapt the model in response to variations in the terrain.

With the methodology presented here, we can not only develop strategies to compensate for failure within mobility systems but explore the limits of the generated driving primitives and set operational constraints for our new driving techniques without needing a full test campaign, enabling the ability to adapt to degrading hardware without significant mission delays.

Overall, we have demonstrated through both field tests and controlled laboratory experiments that off-nominal driving strategies are useful both for increasing functionality of existing rovers as well as recovering functionality of degraded rovers, that these off-nominal driving strategies are difficult to implement in an ad-hoc manner, and that new terramechanics models directly facilitate the automatic generation of off-nominal driving strategies.

### Acknowledgment

The authors would like to thank David Wettergreen for the use of the soft soil testbed, along with Antoine Tardy, Georgia Crowther, and Chiara Rissola for their assistance with experiments.

### REFERENCES

- [1] P. Leger, A. Trebi-Ollennu, J. Wright, S. Maxwell, R. Bonitz, J. Biesiadecki, F. Hartman, B. Cooper, E. Baumgartner, and M. Maimone, "Mars Exploration Rover surface operations: Driving Spirit at Gusev Crater," in *IEEE International Conference on Systems, Man and Cybernetics*, vol. 2, 2005, pp. 1815–1822.
- [2] M. Heverly, J. Matthews, J. Lin, D. Fuller, M. Maimone, J. Biesiadecki, and J. Leichty, "Traverse performance characterization for the Mars Science Laboratory rover," *Journal of Field Robotics*, vol. 30, no. 6, pp. 835–846, 2013.
- [3] R. Sullivan, R. Anderson, J. Biesiadecki, T. Bond, and H. Stewart, "Cohesions, friction angles, and other physical properties of Martian regolith from Mars Exploration Rover wheel trenches and wheel scuffs," *Journal of Geophysical Research E: Planets*, vol. 116, no. 2, 2011.
- [4] J. Biesiadecki, E. T. Baumgartner, R. G. Bonitz, B. K. Cooper, F. R. Hartman, P. C. Leger, M. W. Maimone, S. A. Maxwell, A. Trebi-Ollennu, E. W. Tunstel, and J. R. Wright, "Mars Exploration Rover surface operations," *IEEE Robotics and Automation Magazine*, vol. 13, no. 2, pp. 63–71, 2006.
- [5] J. Matijevic and E. Dwell, "Anomaly Recovery and the Mars Exploration Rovers," in *SpaceOps 2006 Conference*. Reston, Virginia: American Institute of Aeronautics and Astronautics, June 2006.
- [6] C. Pavlov and A. M. Johnson, "Soil displacement terramechanics for wheel-based trenching with a planetary rover," in *IEEE Intl. Conference on Robotics and Automation*, Montreal, Canada, May 2019, pp. 4760–4766.
- [7] —, "A terramechanics model for high slip angle and skid with prediction of wheel-soil interaction geometry," *Journal of Terramechanics*, vol. 111, pp. 9–19, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0022489823000691>
- [8] K. M. Lynch and M. T. Mason, "Controllability of pushing," in *IEEE International Conference on Robotics and Automation*, vol. 1, 1995, pp. 112–119.
- [9] J. E. King, M. Klingensmith, C. M. Dellin, M. R. Dogar, P. Velagapudi, N. S. Pollard, and S. S. Srinivasa, "Pregrasp manipulation as trajectory optimization," in *Robotics: Science and Systems*, 2013.
- [10] N. B. Zumel and M. A. Erdmann, "Nonprehensile two palm manipulation with non-equilibrium transitions between stable states," in *IEEE International Conference on Robotics and Automation*, vol. 4, 1996, pp. 3317–3323.
- [11] J. E. King, J. A. Haustein, S. S. Srinivasa, and T. Asfour, "Nonprehensile whole arm rearrangement planning on physics manifolds," in *IEEE International Conference on Robotics and Automation*, 2015, pp. 2508–2515.
- [12] J. E. King, M. Cognetti, and S. S. Srinivasa, "Rearrangement planning using object-centric and robot-centric action spaces," in *IEEE International Conference on Robotics and Automation*, 2016, pp. 3940–3947.
- [13] J. King, "Robust rearrangement planning using nonprehensile interaction," Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, USA, 2016.
- [14] M. Mason, D. Pai, D. Rus, L. Taylor, and M. Erdmann, "A mobile manipulator," in *IEEE International Conference on Robotics and Automation*, vol. 3, 1999, pp. 2322–2327.
- [15] S. S. Srinivasa, C. R. Baker, E. Sacks, G. B. Reshko, M. T. Mason, and M. A. Erdmann, "Experiments with nonholonomic manipulation," in *IEEE International Conference on Robotics and Automation*, vol. 2, 2002, pp. 2042–2047.
- [16] 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. [Online]. Available: <https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/aisy.202200119>
- [17] D. Kerimoglu, D. Soto, M. L. Hemsley, J. Brunner, S. Ha, T. Zhang, and D. I. Goldman, "Learning manipulation of steep granular slopes for fast mini rover turning," *CoRR*, vol. abs/2310.01273, 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2310.01273>
- [18] Rover Team, "Characterization of the martian surface deposits by the mars pathfinder rover, sojourner," *Science*, vol. 278, no. 5344, pp. 1765–1768, 1997. [Online]. Available: <https://science.sciencemag.org/content/278/5344/1765>
- [19] NASA press release, "Scuff stuff," October 5th 2012. [Online]. Available: <https://mars.nasa.gov/news/1369/scuff-stuff/>
- [20] D. Buckner, "Reading the ripples at Observation Mountain," November 15th 2022. [Online]. Available: <https://mars.nasa.gov/mars2020/mission/status/418/reading-the-ripples-at-observation-mountain/>
- [21] J. Townsend, M. Seibert, P. Bellutta, E. Ferguson, D. Forgette, J. Herman, H. Justice, M. Keuneke, R. Sosland, A. Stroupe, and J. Wright, "Mars exploration rovers 2004–2013: Evolving operational tactics driven by aging robotic systems," in *13th International Conference on Space Operations*. American Institute of Aeronautics and Astronautics Inc., 2014.
- [22] NASA press release, "Now a stationary research platform, NASA's Mars rover Spirit starts a new chapter in red planet scientific studies," January 26th 2010. [Online]. Available: [https://www.nasa.gov/mission\\_pages/mer/news/mer20100126.html](https://www.nasa.gov/mission_pages/mer/news/mer20100126.html)
- [23] —, "NASA's Opportunity rover rolls free on Mars," June 6th 2005. [Online]. Available: <https://mars.nasa.gov/mer/newsroom/pressreleases/20050606a.html>
- [24] —, "NASA's Opportunity rover mission on Mars comes to end," February 13th 2019. [Online]. Available: <https://mars.nasa.gov/news/8413/nasas-opportunity-rover-mission-on-mars-comes-to-end/>
- [25] A. Carvalho Leite, B. Schäfer, and M. L. de OLiveira e Souza, "Fault-tolerant control strategy for steering failures in wheeled planetary rovers," *Journal of Robotics*, vol. 2012, no. 1, p. 694673, 2012. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1155/2012/694673>
- [26] J. Guo, Y. Li, B. Huang, L. Ding, H. Gao, and M. Zhong, "An online optimization escape entrapment strategy for planetary rovers based on bayesian optimization," *Journal of Field Robotics*, vol. 41, no. 8, pp. 2518–2529, 2024. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.22361>
- [27] A. Rankin, M. Maimone, J. Biesiadecki, N. Patel, D. Levine, and O. Toupet, "Mars curiosity rover mobility trends during the first seven years," *Journal of Field Robotics*, vol. 38, no. 5, pp. 759–800, 2021.
- [28] National Academies of Sciences, Engineering, and Medicine, *Origins, Worlds, and Life: A Decadal Strategy for Planetary Science and Astrobiology 2023-2032*. Washington, DC: The National Academies Press, 2022.
- [29] NASA press release, "VIPER mission overview," 2024. [Online]. Available: <https://www.nasa.gov/viper>
- [30] —, "Mars rover tests driving, drilling and detecting life in Chile's high desert," March 13th 2017.

[Online]. Available: <https://www.nasa.gov/feature/ames/mars-rover-tests-driving-drilling-and-detecting-life-in-chile-s-high-desert>

[31] A. Azua-Bustos, C. González-Silva, and A. G. Fairén, "The atacama desert in northern chile as an analog model of mars," *Frontiers in Astronomy and Space Sciences*, vol. 8, 2022. [Online]. Available: <https://www.frontiersin.org/journals/astronomy-and-space-sciences/articles/10.3389/fspas.2021.810426>

[32] FARO, "FARO@Focus Laser Scanners," 2024. [Online]. Available: <https://www.faro.com/en/Products/Hardware/Focus-Laser-Scanners>

[33] Leica Geosystems, "Total stations," 2024. [Online]. Available: <https://leica-geosystems.com/en-us/products/total-stations>

[34] M. Barragan, N. Flowers, and A. M. Johnson, "MiniRHex: A small, open-source, fully programmable walking hexapod," in *Robotics: Science and Systems Workshop on "Design and Control of Small Legged Robots"*, Pittsburgh, PA, June 2018, workshop abstract.

[35] G. Crowther, D. Apostolopoulos, and S. Heys, "Tensegrital wheel for enhanced planetary surface mobility: Part 1 – Design and evolution," *Journal of Terramechanics*, vol. 100, pp. 11–24, 2022. [Online]. Available: <https://doi.org/10.1016/j.jterra.2021.11.008>

[36] —, "Tensegrital wheel for enhanced planetary surface mobility: Part 2 - Performance assessment of a ruggedized, double-layer tensegrital wheel," *Journal of Terramechanics*, vol. 100, pp. 87–101, 2022. [Online]. Available: <https://doi.org/10.1016/j.jterra.2021.11.007>

[37] NASA press release, "A VIPER in the sand," January 26th 2022. [Online]. Available: <https://www.nasa.gov/glenn/image-feature/2022/latest-VIPER-prototype-navigates-lunar-surface-of-SLOPE>

[38] S. Moreland, "Traction processes of wheels in loose, granular soil," Ph.D. dissertation, Carnegie Mellon University, 2013.

[39] 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, 2020.

[40] H. A. Oravec, X. Zeng, and V. M. Asnani, "Design and characterization of GRC-1: A soil for lunar terramechanics testing in Earth-ambient conditions," *Journal of Terramechanics*, vol. 47, no. 6, pp. 361–377, 2010.

[41] C. A. Pavlov, A. Rogg, and A. M. Johnson, "Assessing impact of joint actuator failure on lunar rover mobility," in *Spring Meeting of the Lunar Surface Innovation Consortium (LSIC)*, Johns Hopkins University Applied Physics Lab, May 4-5 2022, abstract.

[42] C. M. Creager, V. M. Asnani, H. A. Oravec, and A. C. Woodward, "Drawbar pull (dp) procedures for off-road vehicle testing," 2017. [Online]. Available: <https://api.semanticscholar.org/CorpusID:117248723>

[43] E. Rezich, A. Schepelmann, H. Oravec, and C. Creager, "VIPER 2020 single wheel TREC testing," October 23rd 2020.

[44] B. Dotson, D. Sanchez Valencia, C. Millwater, P. Easter, J. Long-Fox, D. Britt, and P. Metzger, "Cohesion and shear strength of compacted lunar and martian regolith simulants," *Icarus*, vol. 411, p. 115943, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0019103524000010>

[45] H. Inotsume, C. M. Creager, D. S. Wettergreen, and W. R. Whittaker, "Finding routes for efficient and successful slope ascent for exploration rovers," *International Symposium on Artificial Intelligence, Automation and Robotics in Space*, 2016.

[46] H. Shibly and K. Iagnemma, "An equivalent soil mechanics formulation for rigid wheels in deformable terrain , with application to planetary exploration rovers," *Journal of Terramechanics*, vol. 42, pp. 1–13, 2005.

54 pt  
0.75 in  
19.1 mm

54 pt  
0.75 in  
19.1 mm