Off-Nominal Rover Driving: Terrain Manipulation and Degraded Mobility Compensation

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Mechanical Engineering

Catherine Pavlov

B.S., Mechanical Engineering, California Institute of Technology

Carnegie Mellon University
Pittsburgh, PA
February, 2023

Acknowledgements

This thesis would not have been possible without incredible support in many forms – financial, academic, and emotional. This work was funded by a NASA Space Technology Research Grant (Grant # 80NSSC18K1154), a NASA STTR/SBIR Grant (Grant # NNX16CA42P), a Carnegie Mellon University Presidential Fellowship, and the Carnegie Mellon University Department of Mechanical Engineering.

First and foremost, I would like to thank my advisor, Aaron Johnson. The path to completing this thesis was not straightforward, and you handled every setback and delay with empathy, understanding, and grace. Thank you for taking a chance on me, it has been both an honor and a pleasure to help grow the Robomechanics Lab from the very beginning.

I was fortunate to have an incredible thesis committee who contributed both wisdom and resources in many forms. David Wettergreen first introduced me to terramechanics and the joys of testing wheels in sand, advised on testing terramechanics models and wheels in general, and generously and repeatedly loaned both equipment and sand. Dimi Apostolopoulos and Terry Fong both guided the early nonprehensile terrain manipulation work that both predated and formed the initial chapters of this work, and it is no overstatement to say that without their vision this thesis would not exist. The passion and enthusiasm you all showed for this work were truly inspirational, and it has been an absolute pleasure to work with you all.

I could not have completed this thesis without the wonderful folks in the Intelligent Robotics Group at NASA Ames, most especially my NSTRF collaborator Arno Rogg. Thank you all for letting me tag along on field trips both near and far, do ridiculous things to your rovers, and for helping me compile decade old code in pursuit of said ridiculous things. Participating in ARADS testing with you all was and will remain one of the highlights of my career.

The journey is more important than the destination, and the journey was made all the better by those I shared the adventure with. The Robomechanics Lab has been a consistent source of positivity and inspiration. To Joe Norby - I couldn't have asked for a better person to start a new lab with. I've immensely enjoyed our adventures both in and out of lab, and I hope we leave a legacy of taking joy in the work but taking care of yourself most of all. Justin, Joe, Sean, Nathan, Paul, James, Stutt, and Ardalan, it's been so fun working with all of you, from the depths of Scaife to the deepest recesses of Wean.

Tess and Yun, thank you for being the best housemates I could ask for, and for making the yellow house a home. Bowen, thanks for literally catching my falls. Eugene, thanks for being my NASA buddy and occasional co-conspirator. Joe Payne, you're an incredible friend and an absolutely delightful neighbor.

To the incomparable Explorers Club of Pittsburgh - you are an incredible community I am so grateful to be a part of, and you truly made Pittsburgh home. I relished having all three types of fun with you, whether we were up to our knees in a bog making up songs about the experience or sitting warm and dry around a fire. Thank you for keeping me grounded by teaching me to access new heights.

To my parents, Liz and George, and my sisters Caroline and Madeline - thank you for believing in me even when I doubted myself, and for always being just a call away. And last, but certainly not least, thank you to my partner, James. Your support has been everything to me. I can't find the words to convey how much you mean to me. I can't wait to see what we build together.

Abstract

This dissertation develops new mobility and manipulation capabilities for planetary exploration rovers through comprehensive wheel-soil interaction modeling and then demonstrates these new techniques on rovers.

While wheeled rovers are the current paradigm for robotic space exploration, wheeled mobile robot performance is only well understood for normal driving conditions. A better understanding of wheel-soil interaction lets us do more with rovers in two key ways: 1) add manipulation capabilities through the use of nonprehensile manipulation, e.g. by using wheels to dig, and 2) retain mobility when experiencing degradation of the mobility system, e.g. loss of a wheel motor. Both of these areas require knowledge of wheel-soil interaction forces beyond the scope of existing methods.

First, we present the concept of Nonprehensile Terrain Manipulation (NPTM) and illustrate the scope of potential applications for planetary exploration rovers. We then select one NPTM action, wheel-based trench excavation in soft soil, as a candidate action that is achievable on current rovers with no hardware addition but that requires new modeling.

Next, we present a closed-form model of soil flow around a wheel driving in regolith that can be used both to model trenches dug by rover wheels and to improve terramechanics models. We then detail a new terramechanics model that covers all slip angles and all ranges of slip and skid and validate it with tests on two wheels operating over a wide range of states.

Then, we demonstrate the feasibility of NPTM, the need to account for mobility system failure, and the viability of recovering from failure on full-scale rovers operating in a variety of environments. This is done through a series of demonstrations on NASA rover prototypes in lunar simulant and a Martian analog environment.

Finally, we implement an optimization framework to automatically generate driving strategies for both digging trenches in soil and recovering from multiple types of mobility system failure. We demonstrate the generated driving strategies on a miniature rover in soft soil and use the optimization to predict overall rover mobility in these modes.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

We have entered a new era of robotic space exploration in the past decade, and are on the cusp of an even greater time for scientific discovery enabled by mobile robots. With the landing of the Curiosity rover on the Martian surface in 2012 [1] and Perseverance in 2021 [2], we have greater scientific power on the Martian surface than ever before. Perseverance recently collected and cached its first sample of material as part of the Mars Sample Return mission, which could see Martian regolith bought back to Earth for study within the next decade [3]. NASA will send its first mobile robot to the lunar surface in 2024 with the launch of the Volatiles Investigating Polar Exploration Rover [4], and several commercial and academic organizations will soon send their on rovers to the moon [5, 6]. Robotic explorers have pushed the boundaries of scientific knowledge of our solar system, and will continue to be at the forefront of its exploration.

The cost of launching planetary robotic platforms is astronomical; gaining additional functionality from existing onboard actuators is therefore of great interest. Wheeled platforms are the design of choice for NASA planetary surface missions [7], due to their ability to passively traverse large objects and evenly distribute weight [8]. The capabilities of planetary rovers can be expanded with no additional hardware through two avenues: 1) the creative use of their wheels and other drive system actuators as manipulators, and 2) novel driving strategies for rovers with degraded mobility systems.

For the first application, we seek to sidestep the typical trade-offs encountered when adding manipulation capabilities to a platform. Energy storage and mass are limited onboard spacecraft, causing competition between mobility, computing, and scientific instrument systems. By minimizing additional actuators, we can avoid adding complexity and volume to the system. This naturally leads to the question "How can we gain greater functionality from existing systems with minimal physical additions?"

In general, field robots seek to avoid contact with the environment except for intentional manipulations performed by an end effector. Interactions between planetary rovers and terrain are incidental; tracks are formed as the rover drives, and small rocks move out of the way, but the rover makes no attempt to reshape the environment to suit its needs. Nonprehensile manipulation seeks to utilize these incidental interactions with the environment to accomplish tasks. By treating the ground itself as a target for manipulation and modification, nonprehensile manipulation can be used to increase mobility or locally terraform for a variety of reasons. For example, a rover might use a wheel to remove topsoil to image below the surface, as shown in Figure 1.1.

As for the second application, we are concerned with increasing the mission duration and preventing premature mission end due to mobility concerns. Planetary rovers have experienced stuck steering actuators, loss of actuator power, and rocks jammed in wheels during mission operations [9]. To date, recovery from mobility failures has been entirely ad-hoc, with adaptations to changing rover kinematics developed as needed through hardware testing on analog rovers. With increasing mission durations up to 1000 km called for in the 2022 Planetary Science and Astrobiology Decadal Survey [10], mobility degradation must be factored into mission planning as an anticipated eventuality rather than an exceptional event. This drives us to ask the question "How can we maintain the functionality of mobility systems when they have degraded beyond their nominal state?"

To ensure mission safety for both applications, the outcome of maneuvers beyond the nominal use of locomotion systems must be reliably predictable. Non-cohesive soil poses one of the primary threats to wheeled robots, which can easily become embedded [11, 12]. This motivates us to ask

"How do we predict the performance of rovers operating in soft soil beyond their intended use?"

Taking all of these questions together, the goal of this research is to gain a deeper fundamental understanding of the mechanics of robot-environment interactions, in particular wheel-soil interaction, and in doing so develop both useful nonprehensile terrain manipulation capabilities and locomotion strategies that mitigate the impact of mobility system failures. Current modeling of wheel-soil interaction for rovers is limited to normal driving maneuvers, so a better understanding of the mechanics involved is required for rovers to manipulate terrain and use new driving primitives without introducing additional risks to mission safety.

Figure 1.1: KREX-2 digging a trench in the Atacama Desert, Chile, using its wheels as manipulators. The white streaks visible in the subsurface material reveal the presence of halites in the soil.

1.1 Dissertation Overview

In this dissertation we develop new terramechanics models and off-nominal driving capabilities for planetary exploration rovers, specifically for the applications of nonprehensile terrain manipulation and degraded mobility performance. We seek to answer several questions:

1. What is the space of possible nonprehensile terrain manipulation actions and how do we classify them?
2. How can we represent soil movement around a rover wheel in a computationally efficient manner?
3. How do wheels behave when experiencing high amounts of lateral slip? How do they behave when lateral slip is combined with longitudinal skid?
4. How do we manipulate soil with a wheel without compromising mobility?
5. Can we recover from mobility system degradation in a systematic manner?

In answering these questions, we define the scope of nonprehensile manipulation capabilities available to rovers, and identify the development required to realize them. We then select a candidate nonprehensile terrain manipulation action and develop a closed-form model of soil flow that we use to represent the terrain modification. Next, we integrate that soil flow model into a new terramechanics model that is capable of predicting the wheel-soil interaction forces experienced both during wheel-based soil manipulation and degraded mobility conditions. Finally, building off of these new models, we use numerical optimization to generate full-rover control strategies for performing soil manipulations and driving with degraded mobility.

The rest of the thesis is organized as follows:

Chapter 2 provides background on the progress in development of nonprehensile manipulation for various platforms and previous instances of rovers adapting to mobility system degradation. It also explores the current state of research into terramechanics, terrain deformation modelling, and terramechanics-based control of rovers.

Chapter 3 delineates potential Non-Prehensile Terrain Manipulation (NPTM) capabilities for planetary exploration rovers into robot “actions", and breaks down the use cases and requirements for each action. The selection of wheel-based trenching as an example for NPTM is justified here. This work was written up in the final grant report for NASA STTR/SBIR No. NNX16CA42P.

Chapter 4 introduces a novel closed-form model for predicting the shape of soil deformed by a wheel during trenching, and demonstrates the model in a single-wheel testbed. This work was presented at ICRA 2019 in [13].

Chapter 5 presents a new terramechanics model capable of predicting the forces on rover wheels during off-nominal driving, as in NPTM or mobility system failure. This model is an extension of existing terramechanics models, which do not currently cover the range of slip angles and ratios needed. Notably, this work is the first to cover wheels that are in a state of both skid and a nonzero amount of lateral slip simultaneously, and validates a lateral slip model beyond the 40° previously covered in the literature. This work is in review [14].

Chapter 6 shows field demonstrations of rovers doing NPTM and driving with degraded mobility systems. Wheel-based trenching, multipass trenching, and teaming with a smaller robot to enhance its mobility are all performed in the Atacama Desert. Additional testing of the model described in Chapter 4 is conducted in unprepared terrain in the field. Experimental analysis on the impact of mobility actuator failure is performed on a lunar rover analog, and a large rover is shown compensating for failed drive and steering actuators. Portions of this work were presented as posters in [15] and [16].

Chapter 7 puts the terramechanics model and soil flow model developed in Chapters 4 and 5 into an optimization framework to generate whole-rover behaviors for using wheels to trench while driving without getting stuck, drive straight with one wheel subjected to a reduced speed, and drive straight with one steer actuator stuck at a high angle. A paper on this work and the demonstrations in Chapter 6 is in preparation [17].

Finally, in Chapter 8 we conclude by summarizing this dissertation's contributions to off-nominal driving and proposing a path for future development. Steep slope mobility is identified as an area where intentional reshaping of the soil surface may dramatically increase performance. The potential to optimize rover design for both off-nominal driving and nominal driving is also discussed, as are future modeling developments needed to further expand rover off-nominal driving capabilities.

Chapter 2

Background

The development of off-nominal driving capabilities requires an understanding of wheel-soil interaction and the scenarios in which it is pushed beyond typical behavior, whether through intentional nonprehensile manipulation or as a result of degradation of the mobility system. In this chapter, we survey works relevant to off-nominal driving in several different areas. In Section 2.1, we look at the history of nonprehensile manipulation and recent work on nonprehensile manipulation with planetary exploration rovers. Section 2.2 goes over the history of mobility system failures in rover missions and steps taken to mitigate them. Section 2.3 goes into the history and current state of wheel-soil interaction models, and shows that there is a gap in currently existing models. Section 2.4 describes the current state of simulations of rovers that incorporate soil motion to motivate the need for simple methods. In Section 2.5, an overview of existing work on controlling rovers using terramechanics knowledge is given, with particular emphasis on those that have been experimentally verified.

2.1 Nonprehensile Manipulation

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 [18], reposition objects for further manipulation [19, 20], and simultaneously move objects to clear areas [21].

Much work on nonprehensile manipulation focuses on the use of planar surfaces, such as arms [21] or paddles [18, 20]. In [22], nonprehensile rearrangement of boxes was demonstrated in simulation on NASA Ames Research Center's KREX-2 rover augmented with a front-mounted bulldozer blade. This was then demonstrated on hardware by having KREX-2 use its frame to push cardboard boxes on flat outdoor terrain in [23]. The use of the rover's frame in [23] and the sides of HEBI's arms in [21] 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 [24]. 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 [25]. While the Mobipulator's wheels were purpose-designed for manipulation tasks, this did not diminish their functionality as drive actuators.

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 2.2, is one such example. Spirit and Opportunity have also used their wheels to investigate terrain properties, through intentional construction of scuff marks and trenches [26, 27], though this technique was first explored during the Sojourner mission [28].

More recently, NASA's InSight Mars lander found that its burrowing sensor, dubbed the "mole", became stuck and unable to progress descending below Mars' surface, due to unexpected soil properties. Attempts were made to use the lander's scoop arm to tamp down soil above the mole [29]. While these efforts were not ultimately successful in aiding the mole's progress, they serve to further illustrate the potential of using actuators beyond their intended purpose in space missions.

2.2 Mobility System Failure in Planetary Rover Missions

Planetary rovers operating over long periods may encounter mobility system degradation 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 anomalies 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 [30, 7]. Spirit later had increased current draw in its front right wheel, likely due to poor distribution of lubricant within the actuator [9]. 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 [12].

At one point, Opportunity was temporarily embedded in soft soil, but was able to escape through driving strategies tested on earth [11]. Opportunity showed signs of potential lubrication issues in its front right wheel, similar to Spirit, but lessons learned from mitigating Spirit's issues allowed operators to maintain actuator functionality through driving strategies to reduce load on the actuator, intentional heating, and resting of the actuator when needed [9]. Opportunity later had its front right steer actuator freeze 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 [31].

As of this writing, the rovers Curiosity and Perseverance have not experienced significant mobility system failures beyond the degradation of Curiosity's wheels [32]. However, actuator failure has been explicitly cited as a potential threat to Curiosity's mission. 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. As Curiosity only uses visual odometry to check drive progress approximately every meter and otherwise relies on IMU data to determine rover state [32], reliable open-loop driving strategies are required for mobility. Thus, if Curiosity or Perseverance were to experience a drive actuator failure during their mission durations, open-loop driving strategies like those developed in an ad hoc manner for Opportunity would be required for continued operation.

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 [10]. 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 upcoming Volatiles Investigating Polar Exploration Rover (VIPER), which will search the lunar surface for water ice, has a four-wheeled active suspension [33, 34]. 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 for mitigating mobility loss are essential.

2.3 Terramechanics

The field of terramechanics is concerned with predicting the forces arising from wheel-terrain interaction, with a particular focus on predicting driving capabilities. In this section we break terramechanics models down into two categories: "classic" terramechanics, in which analytical descriptions of wheel stress based on empirical studies are used, and Discrete Element/Finite Element Methods, in which the wheel and/or terrain are treated as either collections of separate elements or nodes in a mesh.

2.3.1 Classic terramechanics

[35] provided the initial development of terramechanics for rigid wheels based on analysis of the pressure-sinkage relationship. The incorporation of slip-induced shearing of the soil beneath the wheel was provided later by [36, 37] and [38] based on [39], which established the behavior of soil shear beneath tracked vehicles. [40] showed that wheel slip has a significant contribution to sinkage. More recent works have captured the slip- and skid-induced wheel sinkage effect through the introduction of a sinkage exponent scaled by slip, enabling more accurate prediction of wheel sinkage for wheels experiencing nonzero slip [41, 42, 43, 44, 45]. [46] additionally incorporated the slip sinkage induced during dynamic steering maneuvers by breaking the sinkage exponent into a static component and a dynamic component that is a function of wheel angle and applied load. In this document, "classic" terramechanics models refer to those derived from the work pioneered by Bekker, Wong, and Reece. Generally, these techniques take the wheel geometry, slip ratio s, and slip angle β, as defined in Figure 5.1, and use those to predict wheel forces.

Classic terramechanics models were developed for large wheeled vehicles and are scale-dependent, so [47] and [48] added terms to the pressure-sinkage relationship to accommodate small wheel radii as found on planetary exploration rovers. Additionally, the soil shearing behavior assumed in terramechanics modeling is scale-dependent, as noted in [49], necessitating proper calibration of terramechanics models before use on small wheels. Despite that, [50] found overall behavior of stress distributions under small wheels conforms enough to the classic methods used in [35] and [37] to provide useful predictions. While early researchers could only test the accuracy of terramechanics models by measuring the overall force on the wheel, more recent work has enabled the direct measurement of stress distributions [50] and soil failure under the wheel [51].

One major drawback of standard terramechanics models is evaluation time; the central equations do not have a closed form, and are typically solved through an iterative approach. [52] proposes a simplified, closed-version of classic terramechanics models, while [53] applies quadratic fits to key equations to allow for closed-loop evaluation, both for wheels in slip. Similarly, [45] applies linear and quadratic fits to wheels in skid.

Most terramechanics models look at steady-state behavior, with the effect of wheel lugs, or grousers, averaged over the surface of the wheel. [54] established design guidelines for grousered wheels based on terramechanics experiments, which were expanded on in [55]. [56] develops a dynamic terramechanics model which takes into account the time variation of pressure distribution under a wheel due to grouser motion, and replicates the observed oscillations in forces. [43] incorporates the average contribution of grousers to wheel locomotion with a detailed analysis of grouser interaction mechanics, and uses this to help predict wheel sinkage due to slip, which is not accounted for in classic terramechanics models that look only at static pressure. All of these works look at wheels which are being driven straight, and do not consider skid for grousers.

The majority of terramechanics models consider only a wheel driving straight, but some work has gone into predicting the forces on a wheel while it is being driven at an angle, such as during steering maneuvers. In this state, the wheel experiences stresses in two additional directions not experienced during nominal driving; a lateral shear stress on the wheel rim parallel to the wheel's axis of rotation, and a sidewall force due to the face of the wheel plowing into the soil. One of the earliest attempts at modeling both this lateral shear stress and sidewall force was by [57], who calculated the lateral shear stress separately from the tangential but with the same shear modulus. The lateral shear stress beneath the wheel has been further explored by other researchers [58, 59, 53]. Several papers treat the tangential and lateral shear deformations separately, with separate shear deformation moduli for each axis [58, 59, 60, 61]. Both [62] and [53] combined the lateral and tangential shear deformations, and calculated the component of the stress along each axis by scaling the total shear stress by the magnitude of the shear velocity along that axis. [63] also combined the lateral and tangential shear deformations, and scaled each component based on the position on the wheel and the sign of the slip angle. Similarly, [64] models the lateral shear forces experienced by a wheel turning in place, but does not take into account the stress on the sidewall.

Numerous approaches have been taken to compute the sidewall stress, or bulldozing stress, anchored in soil cutting theories derived from classic soil mechanics research, including the work of [65, 35, 66, 67], and [68]. Yoshida and Ishigami [69, 58, 70, 71] used an approach based on the work of [66], though this was only experimentally validated up to a wheel slip angle of 30°. [53] the bulldozing stress from both the formulation of Terzaghi's passive earth pressure presented by [35] as well as by the approach [58] takes, but only validated their lateral wheel force up to a slip angle of 16°. Other approaches have been used but not directly validated through single wheel tests, such as [59] and [72], which calculated bulldozing stress from Terzaghi's soil bearing capacity formulation as given in [65], and both [62] and [60], which computed bulldozing stress from [73]. [57] followed the theoretical methods in [68], but only applied the results to pneumatic tires. [55] tested a variety of wheels at slip angles up to 30°, but does not attempt to model the side forces.

The prediction of wheel stresses is governed by the wheel-soil interaction geometry, which is driven by the sinkage. While the point along the wheel rim that first contacts the terrain (the entrance angle) can generally be computed directly from the sinkage, the corresponding exit angle on the rear of the wheel rim is harder to determine. The exit angle is either assumed to be near zero as in [59, 42], and [43]; given as a pure function of the entrance angle as in [60] and [44]; or given as a function of the exit angle and slip ratio as in [58], [62, 53, 50] (though [50] notes the exit angle is constant in skid). [74] noted large deviation from expected contact geometry at high slip ratios, and empirically modeled the entrance and exit angles as a function of soil and wheel geometry. These approaches rely on either directly measuring this relationship or tuning it in the model without external verification. None of these works have explored the impact of rover slip angle on the determination of exit angle. Other methods involve simulating the terrain in addition to the rover: [41] and [53] treated the terrain surface as static, while [59] and [72] fully represented the terrain as deformable. While [75] showed FEM and DEM methods are capable of capturing the terrain deformation and corresponding wheel-terrain interaction geometry, these methods are computationally expensive and less suitable to real time prediction of wheel-terrain interaction forces. [76] and [13] explored the shape of terrain deformed by a wheel explicitly, but this has not been used to improve classic terramechanics modeling directly.

For the nonprehensile terrain manipulation and degraded mobility work proposed here, we require a terramechanics model capable of simulating relatively small wheels with grousers undergoing a full range of slip and skid conditions [-1, 1) and a full range of slip angles [0°, 90°]. This requires the formulation and testing of a new terramechanics model, which is described in Chapter 5.

2.3.2 DEM and FEM methods

While traditional terramechanics models may suffer from computational time, they are faster than discrete element methods, as shown in [77] and noted in [78]. However, what DEM methods lack in speed they can make up for in predictive power, making them important for general prediction of wheel performance but less suitable for use in real-time modeling. In [79], a DEM model was used to predict wheel performance over non-smooth soil, which classic terramechanics models are unable to do.

Soil Contact Modeling uses meshes to represent the wheel and terrain, but is more computationally efficient than standard FEM techniques, and is designed specifically for rover driving simulations over rough terrain [80].

Resistive Force Theory takes a slightly different approach, treating wheels (or other locomotors) as a superposition of flat plate sections and the soil as a fluid-like substance [81]. Notably, RFT differentiates itself from terramechanics in capturing dynamic phenomena, not just steady-state behavior. [82] found that while in some situations RFT yields similar results to DEM, it struggles to capture grouser interactions correctly. [83] implemented RFT as well as the Material Point Method, a continuum plasticity-based implementation, and found that while RFT and MPM outperformed a terramechanics model significantly in predicting sinkage, all three models had fairly comparable performance in predicting wheel forces.

2.4 Terrain Modeling

There is a wealth of research in simulating the motion of terrain; most techniques fall into either Finite Element or Discrete Element Methods (FEM & DEM), in which the soil is either approximated as a continuum of connected vertices or treated as many small particles, respectively.

[84] and [85] combine FEM methods and DEM methods to take advantage of the faster solve times for FEM in the majority of the material bulk, with detailed DEM simulations performed on material close to the tool-soil interface. Demonstrations have focused on earthmoving applications such as pushing soil with a blade, with no focus on predicting wheel performance. While many, such as [86] use terramechanics theory from [35] or other fundamental soil mechanics to predict soil motion, they are not designed for predicting wheel performance and have not been validated for use in that area.

Other work on terrain modeling focused on driving performance simulations typically uses FEM and DEM to incorporate soil motion into modeling of the forces on the wheel. The DEM methods mentioned in Section 2.3.2 can be used for modeling the shape of terrain after deformation, though some of the FEM techniques mentioned, including some RFT implementations, do not specify soil movement. [75] uses information on Martian regolith to create DEM simulations of the interaction between rover wheels and terrain with high fidelity, while [76] modeled the shape of a rut created by a rolling wheel using FEM.

These methods are all more computationally expensive than classic terramechanics methods, and thus more suitable for evaluation of standard mobility techniques than generation of new strategies.

2.5 Rover Control Incorporating Knowledge of Terramechanics

There is a variety of work on using either terramechanics models or key terramechanics properties to control rovers during rough terrain driving. There are not currently any off-the-shelf simulation methods available which have high fidelity terramechanics modeling or easy incorporation of custom terramechanics models since the discontinuation of ANVEL. [87] uses ANVEL to simulate rover locomotion over various types of rough terrain. Unfortunately, ANVEL is no longer accessible to the public.

It is common for researchers to implement their own multibody simulation of rover motion to test their terramechanics models on. [41] models the full dynamics of a rover, using terramechanics as force inputs to the wheels, using a MATLAB toolbox called SpaceDyn. This was tested against experimental data of the rover on flat terrain, and then used to simulate rover traversal over rough terrain. [58] and [88] use a full kinematic model of the rover and the forces predicted by terramechanics to simulate the rover's motion. [81] uses a simplified terramechanics model and a physics engine to implement a real-time simulation of a rover ascending a slope that is able to predict the forces on the wheels, slip, and sinkage very well. [64] fuses a terramechanics model with a terrain height map that incorporates soil mechanics properties to simulate the locomotion of a six-wheeled rover with good predictive capabilities of the rover's motion and driving forces.

The Soil Contact Model (SCM) developed in [80] is used for a multibody simulation of a rover, though the full simulation is not validated in that publication. [89] shows that uncertainty in terrain properties leads to poor predictive capability of rover simulation in the case of SCM.

Chapter 3

Classification of Nonprehensile Terrain Manipulation Actions

The first step in developing nonprehensile terrain manipulation capabilities is to identify and classify the possible manipulation actions that can be achieved with the actuators present on rovers or with the addition of minimal tools. Identifying these manipulation actions and their relevance to planetary exploration mission scenarios allows us to appropriately focus the work of developing NPTM capabilities. We also pair robot operations with their corresponding NPTM actions, and for each identify possible targets of the action, required hardware and software capabilities, and metrics for completion of the action. This categorizations enables the quick identification of potentially feasible NTPM actions for a given robot platform.

3.1 Scope

This work specifically targets wheeled platforms such as KREX-2 [90], VIPER [34], K10 Mini [91], and Curiosity [92], shown in Figure 3.1. While later sections will focus on four-wheeled rovers, the nonprehensile terrain manipulation mission scenarios and robot operations detailed here are not dependent on the number of wheels and apply to six-wheeled rovers as well. Here it is assumed that any rover has individual drive actuators for each wheel, steering actuators for some wheels (but not necessarily all), and the ability for the wheels to move vertically relative to each other. In most cases this takes the form of a passive rocker or rocker-bogie suspension [90, 91, 92], but some rover designs will have additional actuation built into the suspension [34]. The identified robot operations allow for the addition of simple tools with up to two degrees of freedom beyond what is present on the base rover.

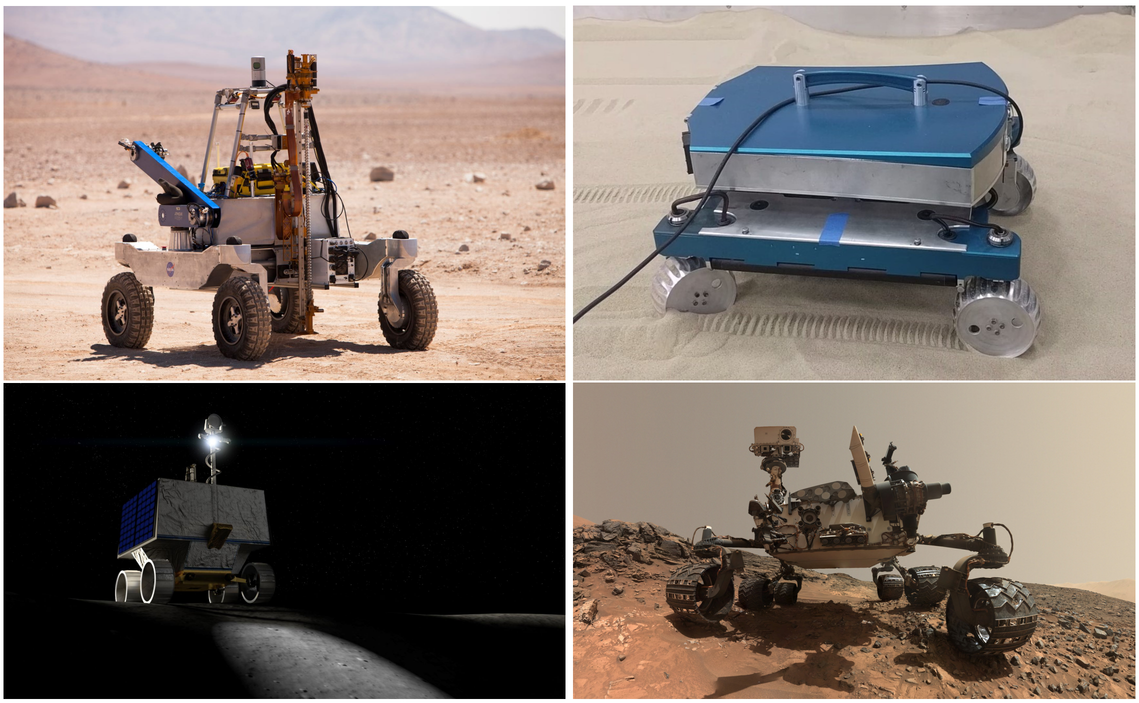

Figure 3.1: KREX-2 rover (top left) [90], K10 Mini rover top right (image credit: author), VIPER rover (bottom left) [34], and Curiosity rover (bottom right) [92].

Additionally, the primary target of this work is planetary exploration missions. While NPTM strategies can be applied to terrestrial missions, the clearest use-case is in the context of weight-limited rover missions. The focus is on missions with some kind of regolith and non-negligible gravity, as that describes the current extent of wheeled rover exploration. The moon and Mars are the most commonly visited extraterrestrial bodies, and this work should be viewed with those destinations in mind. Low gravity bodies in our solar system (such as the Bennu asteroid) may have regolith on their surfaces [93], but limited research has been conducted into the efficacy of wheeled locomotion on their surfaces.

3.2 Robot Actions and Mission Scenarios

Table 3.1 (end of section) identifies 15 robot actions of potential use to a planetary exploration rover. A robot "action" refers to the high-level goal of the robot, such as digging a trench or rolling a rock over. The mission scenarios described here take into account that there may be additional robots or human explorers present in addition to the robot performing NPTM, and many of the tasks identified aim to aid in the missions of other actors. The actions laid out in the first column use rocks as a proxy for movable rigid bodies on a scale comparable to that of the rover. The "purpose" column refers to three distinct categories: mobility, landing site/habitat prep, and science. "Mobility" actions serve to augment the mobility of the rover or other actors, generally by manipulating the environment to make terrain more traversable. "Landing site/habitat prep" focuses on the preparation of the environment in anticipation of future tasks being performed there, such as clearing an area of hazards in advance of the landing of a crewed mission. "Science" focused tasks increase the ability of the rover to collect samples; specific science goals for planetary rover missions are not explored here. These categories are intended to suggest potential applications for NPTM actions, not restrict their uses; while the primary categories for trenching are prep and science, the trenching action may be adapted for mobility purposes.

This table illustrates the variety of manipulation actions that a rover can take and their applicability to mission scenarios. As noted in Section 2.1, several Mars missions have performed ad-hoc terrain manipulation in the form of digging trenches and small holes with wheels [26, 27, 28]. This table places those manipulation actions in context, and shows how only a small subset of what is possible has been explored. For example, selection of a landing site free of obstacles is critical for stationary lander missions, but a rover already on-site could be used to create safe landing sites, opening up a wide array of terrains to potential scientific observation. Similarly, rovers can be used to improve mobility for other, less mobile platforms or for human explorers through the construction of paths and removal of obstacles. This would benefit small robots such as PUFFER [94], which could be used for monitoring of an area while the larger rover travels elsewhere. In Section 6.1.4, two robot-teaming scenarios are demonstrated to this effect, with a larger rover clearing a path for a smaller robot.

3.3 Robot Operations and Characteristics

Table 3.2 takes the actions identified in Table 3.1 and identifies the associated "robot operation" for each, referring to low-level behaviors on the robot, such as rotating a wheel's angle or raising a scoop. A variety of simple tools are considered for addition to the rover, including drills, scoops, plows, and sweeps, as shown in Figure 3.2. As this work aims to avoid increasing the mass and complexity of rovers, only tools with up to two additional degrees of freedom (2DOF) are considered. In particular, this excludes robotic arms. Plows like those seen on bulldozers could be attached to the rover frame to facilitate movement of larger quantities of regolith. A "sweep" mechanism refers to the attachment of a plow-like structure to a wheel below its steering motor, such that steering the yaw of the wheel is used to actuate the sweep.

In breaking down the actuator requirements and sensor requirements for NPTM, Table 3.2 shows that there is much more that can be done with the manipulators and sensors most planetary rovers already have. All mobile rovers NASA has deployed on planetary missions have wheels and a rigid frame, which are the primary systems identified in the majority of robot operations listed. Additionally, many planetary exploration rovers already possess drills and scoops for sampling processes and would not require the additional manipulation listed for actions involving them [92]. While most planetary rovers may not have F/T sensing, all have visual sensing in the form of cameras. Several of the actions listed in Table 3.2 identify a need for depth sensing. In some cases this can be achieved with visual sensors through use of digital image correlation as in [95].

The development of NPTM capabilities is thus not problem of hardware but of a gap in understanding; by building better models of how rovers interact with the environment we can enable new robot actions.

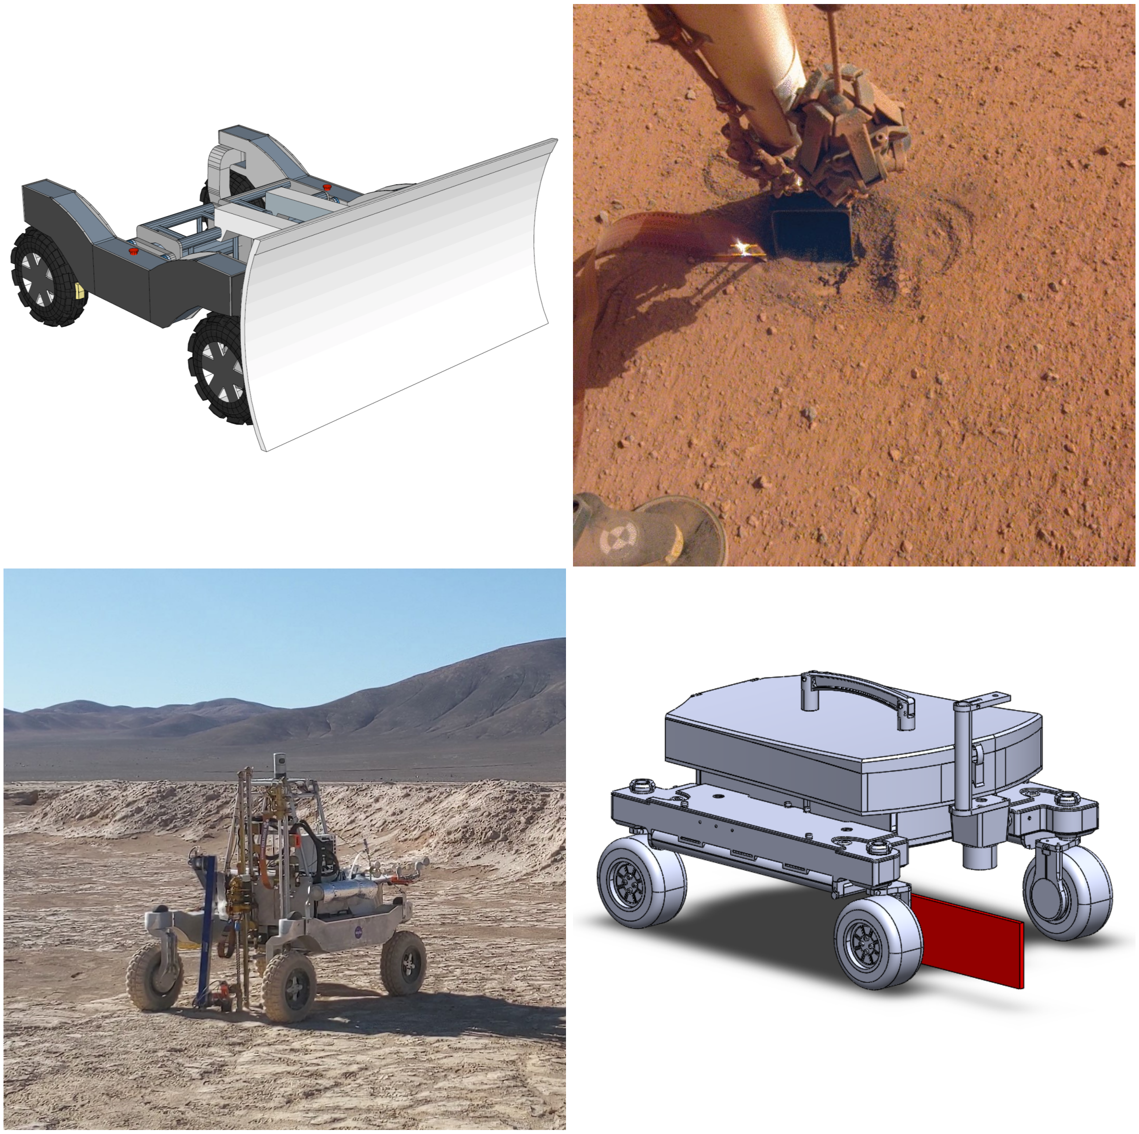

Figure 3.2: Examples of rover-mounted tools: KREX-2 with a front-mounted plow (upper left) [22], NASA's InSight lander's scoop (upper right) [29], KREX-2 drilling (lower left), and illustration of K10 Mini with a sweep bar (red) (lower right).

3.4 Selection of NPTM Actions from Tables

The identified nonprehensile terrain manipulation actions represent a broad array of robot behaviors for interacting with regolith and rigid objects. As seen in Table 3.2, there are a great number of NPTM actions which can be implemented with the addition of no hardware. This work focuses on actions using only the rover wheels, with a particular focus on digging trenches, the first item in both tables. Digging a hole in soft soil with a wheel (the “dig hole” action), moving soil to make a small ramp (“level pile (soil)” and “pile soil” actions), and pushing rocks to clear a path (“level pile (rocks) action”) are all demonstrated in Chapter 6, while wheel-based trenching (“dig trench”) is explored in greater depth.

As discussed in Section 2.1, existing rearrangement planning algorithms can be repurposed for nonprehensile terrain manipulation of rocks and other solid objects. Manipulation of soft soil is less well understood, and using a single wheel to manipulate soil as when trenching can be used for a variety of mission scenarios.

Action	*Purpose*	*Potential Mission Scenarios*
Dig trench	Landing site/habitat prep, science	Access lower strata of surface for sampling. Prep sites for humans/other robots, i.e. by digging a trench to lay wires in.
Dig hole	Landing site/habitat prep, science	Access lower strata of surface for sampling. Prep landing sites, i.e. by digging holes for posts or anchoring
Drill hole	Landing site/habitat prep, science	Access lower strata of surface for sampling. Prep landing sites, i.e. by drilling holes for anchors
Fill ditch (soil)	Mobility, landing site/habitat prep	Make trenches/depressions passable for robot platform, other robots, or humans. Prep landing sites by creating level terrain or burying laid wires.
Level pile (soil)	Mobility, landing site/habitat prep	Prep landing sites for habitat setup, mobility for other robots
Fill ditch (rocks)	Mobility, landing site/habitat prep	Make trenches/depressions passable for robot platform, other robots, or humans. Prep landing sites by creating level terrain or burying laid wires.
Level pile (rocks/discrete objects)	Mobility, landing site/habitat prep	Same as level pile (soil)
Pile rocks	Mobility, landing site/habitat prep	Habitat setup; accumulation of regolith as wall, insulation
Pile soil	Mobility, landing site/habitat prep	Habitat setup; accumulation of regolith as wall, insulation
Roll rocks	Mobility, landing site/habitat prep	Sample collection, site prep, mobility. Used to clear rock out of the way or bring rock somewhere.
Rotate large rocks	Mobility, landing site/habitat prep, science	Move obstacles (for mobility and landing site prep), access area under rocks for sampling (useful for acquisition of samples near surface but shielded from surface conditions)
Tip/push rocks	Mobility, landing site/habitat prep, science	Move obstacles (for mobility and landing site prep), access area under rocks for sampling (useful for acquisition of samples near surface but shielded from surface conditions)
Break rocks	Mobility, science	Eliminate obstacles too large to move, access interior of rocks for sampling
Level slope (soil)	Mobility, landing site/habitat prep	Prepare landing site, make slopes traversable by platform robot, other robots, and humans
Compact soil	Mobility, landing site/habitat prep	Make landing sites more stable for other robots, habitats, and humans. Make untraversable loose soil traversable by platform robot, other robots, and humans

Table 3.1: Table of robot actions possible with NPTM and their applicable mission scenarios.

Robot Operation	*Action*	*Purpose*	*Additional manipulation required*	*Additional manipulator required*	*Type of environment modified*	*Focus of action*	*Sensors required for performance*	*Sensory feedback required for success evaluation*	*Metrics for success*	*Input parameters*	*Expected automation level achievable*
Drag/reverse one wheel while forward driving other three, or use scoop	Dig trench	Landing site prep, science	None/1DOF	None/front scoop	Loose soil/small rocks	Environment	F/T, depth	Basic visual, depth	Alter depth of terrain in desired area	Terrain friability, trench depth	Full automation with minimal perception
Wiggle wheel back and forth	Dig hole	Landing site/habitat prep, science	None	None	Loose soil/small rocks	Environment	F/T, depth	Basic visual, depth	Alter depth of terrain in desired area	Terrain friability, hole depth	Full automation with minimal perception
Drill	Drill hole	Landing site/habitat prep, science	2DOF	Drill	Compact soil/rock	Environment	F/T, depth	Basic visual, depth	Alter depth of terrain in desired area	Terrain friability, hole depth	Full automation
Reverse drive front wheels while fixing back (or vice versa), use plow	Fill ditch (soil)	Mobility, landing site/habitat prep	None/1DOF	None/Plow/ Front scoop	Loose soil/small rocks	Environment	Visual	Basic visual	Alter depth of terrain in desired area	Terrain friability, ditch dimensions	Partial automation, human control likely needed
Reverse front wheels while fixing back (or vice versa), use plow, sweep plow	Level pile (soil)	Mobility, landing site/habitat prep	None/1DOF	None/Plow/ Front scoop	Loose soil	Environment	Visual	Basic visual	Alter depth of terrain in desired area	Terrain friability, pile dimensions	Partial automation, human control likely needed
Push objects with robot frame, reverse drive or fix one wheel while driving others forward	Fill ditch (rocks)	Mobility, landing site/habitat prep	None	None	Mixed/medium-large rocks	Environment	Visual	Basic visual	Alter depth of terrain in desired area	Terrain friability, pile dimensions, nearby rock size	Partial automation, human control likely needed
Push objects with robot frame, reverse drive or fix one wheel while driving others forward	Level pile (rocks)	Mobility, landing site/habitat prep	None/1DOF	None/front scoop	Mixed	Environment	Visual	Basic visual	Alter depth of terrain in desired area	Rock size, orientation	Partial automation, human control likely needed
Push objects with robot frame, reverse drive one wheel, lift with scoop	Pile rocks	Mobility, landing site/habitat prep	None/1/2DOF	None/front scoop	Mixed/medium-large rocks	Object	F/T, visual	Advanced visual	Specific objects in desired orientation	Terrain friability, pile dimensions, nearby rock size	Partial automation, human control likely needed
Push with reverse driven wheels, lift with scoop	Pile soil	Mobility, landing site/habitat prep	None/1DOF	None/front scoop	Loose soil	Object	F/T, visual	Basic visual	Alter depth of terrain in desired area	Terrain friability, pile size	Full automation with minimal perception
Push with reverse driven front wheel(s), push with robot frame, lift and push with scoop	Roll rocks	Mobility, landing site/habitat prep	None/1DOF	None/front scoop	Mixed	Object	F/T, visual	Advanced visual	Specific objects in desired orientation	Terrain friability, rock size	Partial automation, human control likely needed
Push with robot frame, push with sweep	Rotate large rocks	Mobility, landing site/habitat prep, science	None	None	Mixed/medium-large rocks	Object	F/T, visual	Advanced visual	Specific objects in desired orientation	Terrain friability, rock size	Partial automation, human control likely needed
Push with contact wheels fixed, other wheels driving	Tip/push rocks	Mobility, landing site/habitat prep, science	None	None	Mixed/medium-large rocks	Object	F/T, visual	Advanced visual	Specific objects in desired orientation	Terrain friability, rock size	Partial automation, human control likely needed
Drill or other specialized tool	Break rocks	Mobility, science	2DOF/ specialized tool	Drill, chisel	Mixed/medium-large rocks	Object	F/T	Advanced visual	Specific objects in desired orientation	Rock size, hardness	Partial automation, human control likely needed
Reverse front wheels while driving back (or vice versa), use plow, sweep plow	Level slope (soil)	Mobility, landing site/habitat prep	None/1DOF	None/Plow/ Front scoop	Loose soil/small rocks	Object	Visual, angle	Basic visual	Alter depth of terrain in desired area	Terrain friability, slope inclination	Full automation with minimal perception
Drive subset of wheels with high slip rate, tamp down with plow	Compact soil	Mobility, landing site/habitat prep	None/1DOF	None/Plow/ Front scoop	Loose soil/small rocks	Environment	F/T	F/T, visual	Soil too loose to traverse is now traversable	Terrain friability	Partial automation, human control likely needed

Table 3.2: Table of feasible robot operations and their associated actions, purposes, and requirements.

3.5 Conclusions

In this chapter, we have identified the scope of feasible nonprehensile terrain manipulation actions for planetary rovers with minimal or no hardware modification, and classified these actions by their relevant characteristics, such as hardware and environmental requirements. We have identified wheel-based manipulation of soft soil, or trenching, as a NPTM action which is underexplored, useful, and feasible with existing hardware.

In the following chapters, we develop the models needed to deeply understand trenching as a form of NPTM. Chapter 4 develops a model of how terrain is deformed during trenching, Chapter 5 presents a new terramechanics model to predict the forces involved in wheel-based trenching, and Chapter 6 demonstrates the soil flow model in a Mars-analog environment. Chapter 7 uses the models developed in Chapters 4 and 5 to automatically develop full rover driving primitives to achieve a specific trenching task, and Chapter 8 discusses potential extensions of the trenching action for mobility.

Chapter 4

Soil Deformation Modeling

Knowledge of the shape of terrain deformed by a wheel is essential to manipulation of terrain, and a better understanding of wheel-soil interaction geometry can also aid terramechanics modeling. This chapter develops a closed-form model for predicting the shape of a trench or rut formed behind the wheel of a rover and demonstrates it for a small rover wheel in controlled experiments.

4.1 Nomenclature

Table 4.1 lists variables used to develop the soil flow model in this section.

4.2 Soil Flow Model Derivation

The geometry of the wheel-soil interface, Figure 5.1, is driven by the sinkage (h_0) and driving parameters (slip angle, β, and slip ratio, s).

When driving at steady state, all soil displaced by the wheel must flow around and back behind it. Displaced soil is either moved by plowing, where it flows around the sides of the wheel, or by entrapment in the grousers due to rotation of the wheel. The volume of sand in any unit length along the wheel's tracks must be equal to the volume of sand per unit length in front of the wheel (with no

Symbol	Value	Units
h_0	Sinkage	m
s	Slip ratio	
ω	Rotational speed of wheel	rad/s
v_x	Translational velocity of wheel	m/s
β	Slip angle	$^\circ$
r	Radius of wheel	m
b	Width of wheel	m
h_g	Height of grousers	m
μ	Fraction of wheel surface covered by grousers	
ϕ	Soil angle of repose	$^\circ$
A_{tot}	Total area of soil displaced by wheel	m^2
A_{el}	Elliptical portion of displaced soil	m^2
A_{L1}	Soil displaced to the left	m^2
A_{R1}	Soil displaced to the right	m^2
A_{rot}	Soil displaced by grousers	m^2
ζ	Volume fraction of grousers	
h_{L2}	Height of left soil pile on side of wheel	m
h_{R2}	Height of right soil pile on side of wheel	m
h_{L3}	Height of left soil pile after reflow	m
h_{R3}	Height of right soil pole after reflow	m
d	Depth of trench	m
w	Width of wheel at intersection with soil	m
p	Width of trench bottom	m
m	Height of step in trench	m
q	Width of step in trench	m

Table 4.1: Soil flow model variables used in this section.

compaction, an assumption we test in Sec. 4.3). As such, here we consider an infinitesimal slice of soil volume and refer to it as a planar soil area. At steady-state operation, all parameters that determine the planar soil area (e.g. β, s, h_0) are constant. We classify the resulting trench shape into pointy bottomed, flat bottomed, [96], and stepped, as seen in Figure 4.1.

To calculate the resulting profile, we use the constraints imposed by the angle of repose, which is the maximum slope angle of the soil [97], and by conservation of area. The calculation, shown in Figure 4.2, breaks the soil flow into four phases: 1) Soil Separation, 2) Grouser Soil Removal, 3), Inward Flow, and 4) Transported Pile Deposition.

1) Soil Separation: First, the area of soil interaction, A_{tot}, is assumed to be equal to the projected

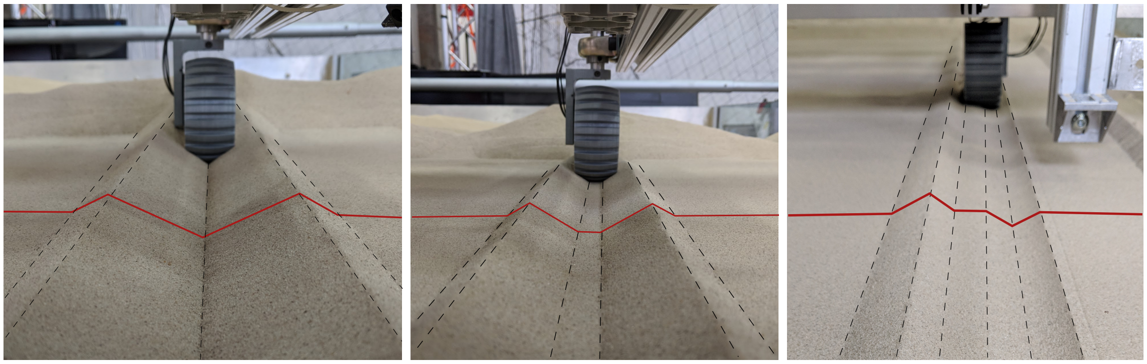

Figure 4.1: Example trenches from the model validation experiments with a pointed bottom (left), flat bottom (center-left), and stepped geometry (center-right).

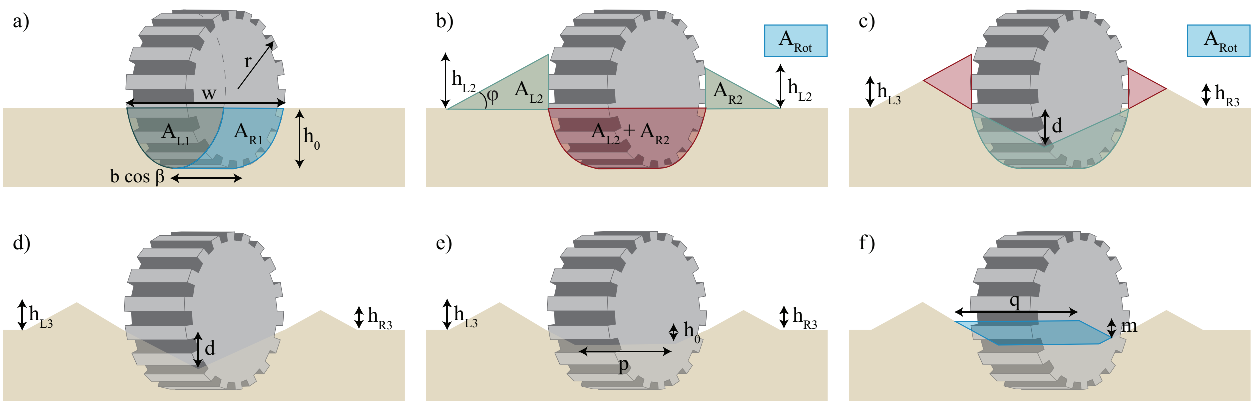

Figure 4.2: Illustration of the soil flow model from behind the wheel. a) The planar soil area (4.1) and division of soil into left and right piles around the wheel, with $A_{L1} = A_{el}$ (in green) (4.2). b) A_{rot} (blue) is removed from the right section (A_{R2}, green) (4.3) and the soil is piled to either side of the wheel (4.4). c) Piles flowing back into the trench, with the red and green areas equal (4.5) - (4.7). d) The resulting trench with a pointed bottom. e) The resulting trench with a flat bottom (4.8) - (4.10). f) The formation of a step from A_{rot} (4.11) - (4.18).

area of the wheel perpendicular to direction of travel, Figure 4.2 a), as determined by the sinkage h_0, slip angle β, wheel radius r, and width b,

$$A_{tot} = A_{el}(r, h_0, \beta) + h_0 b \cos\beta \tag{4.1}$$

where A_{el} is the area of the semi-elliptical cross-section.

The sand is then separated into the piles that flow to the left and right of the wheel, A_{L1} and A_{R1}, respectively. We assume, without loss of generality, that the wheel is turned to the right. For large and small slip angles ($\beta < 10^\circ$ or $\beta > 80^\circ$), the soil is assumed to split evenly around the wheel, creating a symmetric trench. For intermediate angles, the soil flow is assumed to split about the leading edge of the wheel, Figure 4.2 a),

$$A_{L1} = A_{el}, \qquad A_{R1} = h_0 b \cos\beta \tag{4.2}$$

2) Grouser Soil Removal: The soil flows past the wheel via two mechanisms: "plowed" sand, which flows around the left and right of the wheel, and "grouser transported" sand which is moved by the rotation of the wheel. The amount of soil transported by the rotation of the wheel is,

$$A_{rot} = 2\pi r \zeta \mu h_g b \frac{\omega}{v_x}, \tag{4.3}$$

where ζ is the volume fraction of the grousers (fraction each grouser filled with soil), μ is the fraction of the surface area covered by each grouser, h_g the grouser height, ω the angular speed of the wheel in rotations/s, and v_x the forward velocity of the wheel as in (Figure 5.1). This area is bounded by the amount of soil encountered by the grousers on the rotating surface of the wheel. For small angles ($\beta < 10^\circ$), this soil is drawn equally from both the left and right piles, and is bounded by A_{tot}. For larger angles, it is drawn from the pile in the direction of rotation, A_{R1}. The right and left areas after the grouser transported sand is accounted for, in the low angle case, are $A_{R2} = A_{L2} = (A_{tot} - A_{rot})/2$, while for larger angles, $A_{R2} = A_{R1} - A_{rot}, A_{L2} = A_{L1}$.

Determining where sand transported by rotation goes also depends upon the slip angle β of the wheel. For high angles ($\tan\beta > b/(2\sqrt{2rh_0 - h_0^2})$), when the front and rear corners of the wheel align), A_{rot} is added to the opposite volume of plowed soil (so $A_{L2} = A_{L1} + A_{rot}$). For intermediate and low angles, A_{rot} is set aside for now.

Soil piled on either side of the wheel, of area A_{R2} and A_{L2}, is placed initially in triangles with an angle ϕ, the angle of repose of the soil, Figure 4.2 b). The height of the initial soil piles is given by,

$$h_{L2} = \sqrt{2A_{L2}\tan\phi}, \quad h_{R2} = \sqrt{2A_{R2}\tan\phi}. \tag{4.4}$$

3) Inward Flow: The soil that was piled against the wheel then flows back into the trench behind the wheel, Figure 4.2 c). The profile is assumed at first to have a pointy bottom ($p = 0$) with a final height of h_{R3} and h_{L3} on each side and a maximum depth of d, and is determined by solving a set of constraints,

$$\begin{aligned} h_{R3}^2 + h_{L3}^2 &= d^2 - A_{rot} \tan\phi & (4.5) \\ 2h_{R3} + 2h_{L3} + 2d &= h_{R2} + w \tan\phi + h_{L2} & (4.6) \\ 2h_{R3} + d &= h_{R2} + \frac{w - b\cos\beta}{2} \tan\phi & (4.7) \end{aligned}$$

where $w = b\cos\beta + 2\sin\beta\sqrt{2rh_0 - h_0^2}$ is the width of the wheel at its intersection with the soil surface. First, the area of the sand is preserved, (4.5). Next, the soil is assumed to move along the steepest gradient, meaning only inward flow, maintaining the width of disturbed terrain, (4.6). Finally, the deepest part of the resulting trench, of depth d, is assumed to align with the leading corner of the wheel, Figure 4.2 d), (4.7).

The depth of the pointy bottom trench is then checked against the wheel sinkage h_0 – the trench cannot be any deeper than the deepest part of the soil removed. If the predicted trench is too deep, then the trench has a flat bottom and the soil profile is recalculated with a known trench depth of h_0, as shown in Figure 4.2 e).

The heights of the soil piles on the side are then recalculated:

$$h_{R3} = -h_0 + \sqrt{h_0^2 - \left(\frac{d^2}{2} - dh_{R2} - \frac{h_{R2}^2}{2} - \frac{w - b\cos\beta}{2}\frac{d}{3}\tan\phi\right)} \quad (4.8)$$

$$h_{L3} = -h_0 + \sqrt{h_0^2 - \left(\frac{d^2}{2} - dh_{L2} - \frac{h_{L2}^2}{2} - \frac{w - b\cos\beta}{2}\frac{d}{3}\tan\phi\right)} \quad (4.9)$$

The final width of the trench is thus given by

$$p = (h_{L3}^2 + h_{R3}^2 - d^2)/(d\tan\phi) \quad (4.10)$$

4) Transported Pile Deposition: We now consider the soil transported by rotation, A_{rot}. The equations above give the final trench profile for the high angle case ($\tan\beta > b/\left(2\sqrt{2rh_0 - h_0^2}\right)$), as we have already taken the rotated soil into account.

In the low angle case, $\beta < 10°$, the area transported by rotation A_{rot} is placed on the profile, centered along the wheel axis. The new profile then has a flat bottom, whose depth is increased until all of A_{rot} is accounted for. If the original trench had a pointy bottom, the new maximum depth is given by,

$$d_2 = d - \sqrt{A_{rot}\tan\phi} \tag{4.11}$$

while if the original trench had a flat bottom, the new maximum depth is given by,

$$d_2 = d + \frac{p\tan\phi}{2} - \sqrt{\left(\frac{p\tan\phi}{2}\right)^2 + A_{rot}\tan\phi} \tag{4.12}$$

In the intermediate angle case ($\beta > 10°$ and $\tan\beta < b/(2\sqrt{2rh_0 - h_0^2})$), the soil is piled from the bottom up until either a stepped or flat-bottomed trench is created. A step with width q and height m is formed with area equal to A_{rot}, as seen in Figure 4.2 f). Soil is piled in a step with rightmost edge hitting where the right edge of each grouser meets the soil. It is piled until its height reaches the leftmost edge of the wheel, and is then expanded further to the right. If A_{rot} is large enough that the step runs into the right side of the trench, a flat bottom forms.

As the furthest edge of the grouser is located at $\frac{b\cos\beta}{2}$, the step must have width q of at least $\frac{p}{2} + \frac{b\cos\beta}{2}$. Note that if our initial trench is pointy, $p = 0$ and $q = b\cos\beta$. If the initial trench profile was flat the height m of the step is then given by,

$$m = \frac{A_{rot}}{q} \tag{4.13}$$

If the initial trench profile was pointed, the height of the step is given by

$$m = \frac{A_{rot}}{q} - \frac{q\tan\phi}{4} \tag{4.14}$$

In both cases, the height of the step is limited by the farthest point on the trench on which the grousers are able to dump sand. If m exceeds this limit,

$$m_{max} = (2r\sin\beta + b\cos\beta - p)\frac{\tan\phi}{2} \tag{4.15}$$

Then we set $m = m_{max}$ and make the step wider, with q given by

$$q = \frac{A_{rot}}{m_{max}} \tag{4.16}$$

for the initially flat trench and

$$q = \frac{2m_{max}}{\tan\phi} - \sqrt{\left(\frac{2m_{max}}{\tan\phi}\right)^2 - \frac{4A_{rot}}{\tan\phi}} \tag{4.17}$$

for the initially pointy trench.

If the step of the initially flat trench becomes wide enough to run into the other side of the trench, the step width is then given by,

$$q = \frac{1}{2}\left(2 + \frac{4m}{\tan\phi} - \sqrt{\left(2p + \frac{4m}{\tan\phi}\right)^2 - 4\left(p^2 + \frac{4A_{rot}}{\tan\phi}\right)}\right) \tag{4.18}$$

If the step is so wide as to cover the entire bottom of the trench, the trench's profile is computed using the same equations as for the low angle case, (4.11) and (4.12).

4.3 Soil Flow Model Validation

We demonstrate the effectiveness of the soil flow model presented in Sec. 4 with a set of trenching experiments conducted in the soft-soil testbed in Carnegie Mellon's Field Robotics Center.

The goal of these experiments is to qualify the model as a whole by showing that the overall predicted soil profile closely matches that seen in experiments. In addition, some specific assumptions

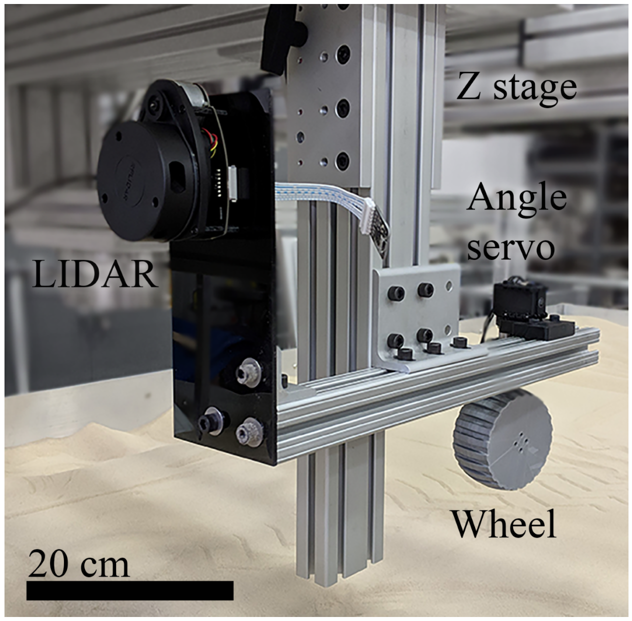

Figure 4.3: The single-wheel soil flow testbed setup. For each trench, the wheel is lowered to a fixed sinkage (h_0) and moves across level sand at a fixed slip ratio (s) and slip angle (β).

are investigated:

- Minimal soil compaction occurs
- The depth of trench follows the model
- The small and large slip angle transition points ($\beta < 10°$ and $\beta > 80°$) for soil division are reasonable.

In each trial, the wheel is set to a fixed sinkage h_0, slip angle β, and slip ratio s, and then driven through a prepared smooth soil bed in a straight line with a forward velocity v_x of 3 cm/s. Consistent soil preparation is achieved via jigs affixed to the testbed for loosening and leveling the sand. The resulting trench is imaged with a LIDAR scanner at 12 locations along its length, 5-10 cm apart, with each sample yielding a profile comprised of 300-400 unique points. The experimental setup is shown in Figure 4.3 and the video attachment, and Table 4.2 contains relevant soil properties and wheel geometries. The wheel geometry was chosen to match the K10 mini rover [91], Figure 3.1, which was used for preliminary trenching experiments.

Sym.	Description	Value
ϕ	Angle of repose [$^\circ$]	29
c	Cohesion [Pa]	0
r	Wheel radius [mm]	48.0
b	Wheel width [mm]	50.0
h_g	Grouser height [mm]	5.0
ζ	Grouser volume fraction	0.1

Table 4.2: Sand parameters (top) and rover parameters (bottom) used in this section. Angle of repose is known for the sand used (Soil Direct #90), other values are typical as given in [53]. Rover parameters match that of the K10 mini rover.

Trench Type	Avg. Error [mm]	Median Error [mm]	Depth Error [mm]	Compaction [%]
All trenches	2.9	2.1	3.5	4.29
$\beta = 0^\circ$	2.1	1.9	5.7	23.1
$\beta = 22.5^\circ$	3.3	2.3	4.0	7.19
$\beta = 45^\circ$	3.1	1.9	2.5	-9.3
$\beta = 67.5^\circ$	2.6	2.0	2.6	-3.02
$\beta = 90^\circ$	3.6	2.2	2.8	3.47
s =-1	3.5	2.3	3.4	5.76
s =-0.5	3.4	2.8	4.0	1.13
$s = 0$	3.3	2.1	4.1	17.12
$s = .8$	2.2	1.3	2.4	-2.48
$s = .9$	2.7	1.5	3.5	-1.0
$h_0 = 5$mm	2.7	1.9	4.1	9.41
$h_0 = 15$mm	2.7	2.1	3.6	1.9
$h_0 = 25$mm	3.5	2.4	2.7	1.56

Table 4.3: Quality of soil flow model fit for all trials with the grousered wheel. Each row is an average over all other test conditions with the listed quantity held constant. Compaction is area change over A_{tot}.

LIDAR scans of the flat prepared surface have a standard deviation of 0.6 mm, which is less than the typical feature scale in these experiments. Calibration scans of alignment jigs are used to correct for rotation and offsets of the LIDAR scanner and testbed mounts. Experiments iterate over a representative range of wheel sinkages ($h_0 \in \{5, 15, 25\}$ mm), slip ratios ($s \in \{-1, -0.5, 0, 0.8, 0.9\}$), and slip angles ($\beta \in \{0^\circ, 22.5^\circ, 45^\circ, 67.5^\circ, 90^\circ\}$), for a total of 69 trenches and 828 sampled profiles[1].

[1]Note that the following trenches at the maximum sinkage were not completed due to excessive deflection of the test rig, indicating conditions where a rover would certainly get stuck: $\beta = 67.5^\circ$ for s =-1 & -0.5, and $\beta = 90^\circ$ for s =-0.5

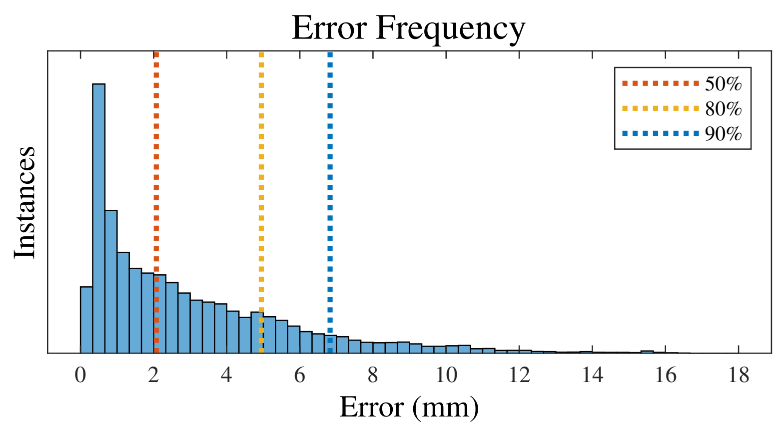

Figure 4.4: Frequency of error in trench profile across all trials. 50% of points had less than 2.1 mm errors, 80% less than 4.9 mm, and 90% less than 6.8 mm

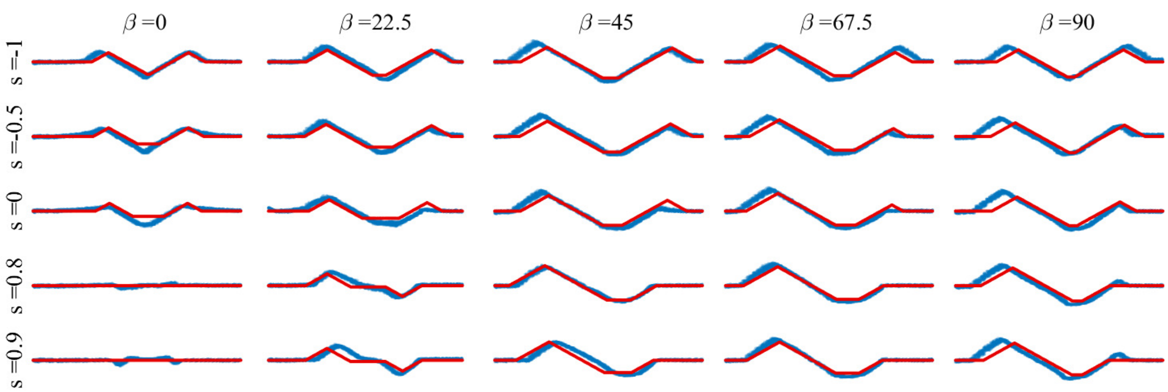

Figure 4.5: Plots of measured (blue) and model-predicted (red overlay) trenches for varied slip angle and slip ratio at fixed sinkage of $h_0 = 15$mm. All 12 measurements of each trench are shown, showing the consistency of the trench. Note that the model captures the steps at the $\beta = 22.5°$ angle, the filled in trench at $\beta = 0°$ with high slip ratios, and the overall width and depth of most trenches.

For each, the average error, δ, of the trench model is evaluated by finding the average Euclidean distance from each measured point along the soil profile to the model profile generated with parameters listed in Table 4.2, and then taking the average error in 0.8mm wide bins. This is the smallest bin size such that no bin is empty. Each bin is equally weighted in determining the overall δ for the trench, to eliminate any bias arising from non-uniform distribution of sampled points. This error is given by,

$$\delta = \frac{1}{n_{bins}} \sum_{i=1}^{n_{bins}} \frac{1}{n_i} \sum_{j=1}^{n_i} |y_{actual,i,j} - y_{expected,i,j}| \tag{4.19}$$

The average error for each trench is reported in Table 4.3, as is the median error between bins

through 0.9.

and the error on the maximum depth of the trench. Finally, these scans are also used to evaluate compaction, by comparing the net area of soil loss post-trenching to the area of soil disturbed by the wheel as predicted by the model, A_{tot}. A slice of the data for $h_0 = 15$ mm is shown in Figure 4.5.

Model fit: The qualitative trench fit was very good, with most qualitative trench shapes (flat vs. pointy, stepped or unstepped) successfully predicted by the model, Figure 4.5. The average error over all trenches was 2.9 mm, with the maximum error at any individual point 18.2 mm. The frequency of error magnitudes can be seen in Fig 4.4.

The model fit was best for low and intermediate sinkages, and worst for high sinkages. Interestingly, the model performed well at high slip. This supports our quasistatic assumption, as the soil does not behave significantly differently at higher rotational speeds and thus any dynamic effects are small.

Compaction: The average compaction for all trenches was 4.29%, but this conceals more interesting behavior. Trench compaction was the most significant for trenches with $\beta = 0°$ and trenches with $s = 0$. In these cases, the wheel does not displace the soil as much as it rolls over, and exerts more force in the vertical direction, compacting it. At intermediate angles ($\beta = 45°, 67.5°$) and very high slip ($s = 0.8, 0.9$), the compaction is negative. This corresponds to the sand getting aerated by the wheel.

Total depth: The overall average error on the deepest part of the trench was comparable to the average error of the trench fit as a whole, as seen in Table 4.3. The error on the maximum depth was about 3mm for most trenching configurations; however, the model was particularly good at predicting trench depth for trenches with intermediate or high β and particularly poor for low β.

Small and large angle soil split rule: The transition points in the soil division rule described in Sec. 4 occur at $\beta = 10°$ and $\beta = 80°$. The angles tested on either side of these cutoffs ($\beta = \{0°, 22.5°\}$, $\beta = \{67.5°, 90°\}$) match the model well. Shifting the transition points to $\beta = 25°$ and $\beta = 75°$ such that the $\beta = 22.5°$ and $\beta = 67.5°$ trials are in the low and high angle regimes, respectively, gives worse fits. This suggests that the transition points do lie within the ranges $[0°, 22.5°]$ and $[67.5°, 90°]$. However, more experiments are needed to refine these points.

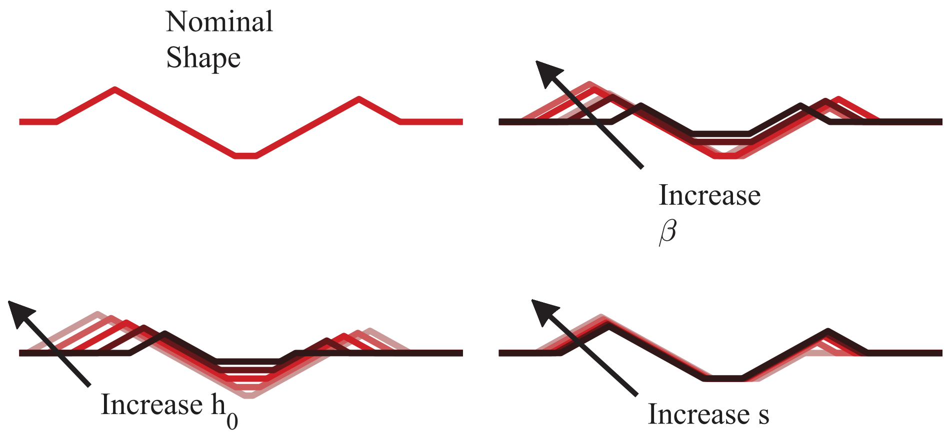

Figure 4.6: Nominal trench shape for $h_0 = 12$ mm, $\beta = 45°$, and $s = 0$, and the effect of varying each parameter individually from this case. Note that sinkage changes the shape of the trench's bottom, while slip angle and ratio affect the shape of the trench's sides.

4.4 Insights into Trenching from Soil Flow Model

Numerical results based on this model were generated by testing different slip ratio and slip angle values using the parameters listed in Table 4.2.

Trench Shape: A surprising variety of qualitatively distinct trench profiles emerge from the variation of slip angle, slip ratio, and sinkage. This model allows us to predict the shape of the trench, which may have a flat bottom or pointed bottom, smooth sides or a step, and piles of sand on both sides or just one side. We are able to select driving parameters (e.g. s, β) to achieve a desired trench geometry – for example, one might choose to dig a trench with all soil piled to one side, so that a cable can then be laid in the trench and reburied with a single pass. Figure 4.6 illustrates the effect of driving parameters on trench shape, and Figure 4.7 shows the range of shapes possible for a single sinkage.

Grouser Effect: In preliminary tests it was observed that even wheels without grousers were capable of transporting soil by rotation. An ungrousered wheel can be used to dig in the same manner as a wheel with grousers, but the amount of soil moved by rotation is more difficult to predict, and likely depends on surface properties of the wheel as much as geometry. Further investigation in this area is needed.

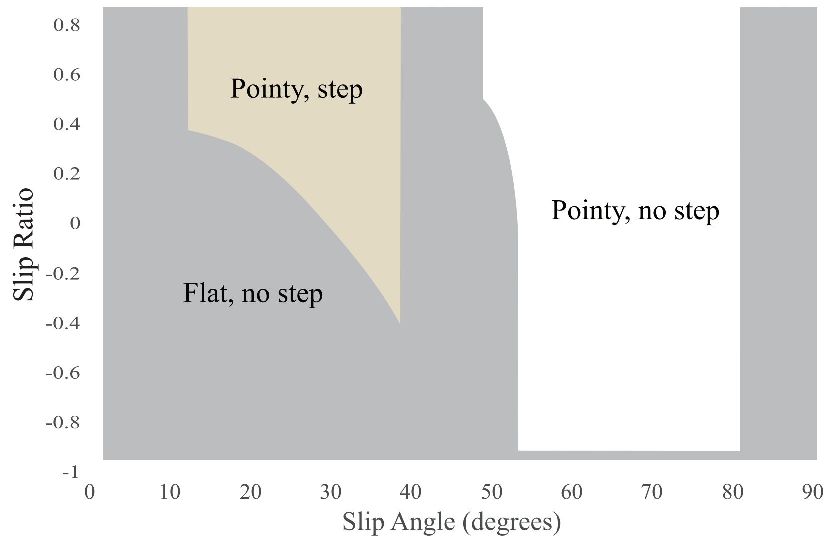

Figure 4.7: Phase plot of trench shape for $h_0 = 10$mm generated from the soil flow model. Flat trenches have $p > 0$, and stepped trenches have $q > 0$.

4.5 Conclusions

In this chapter we have developed a new closed-form model of how terrain is deformed by a driving rover wheel, using only soil mechanics principles and simplifying assumptions. We showed how several distinct types of trench shapes form and how to mathematically predict them. We then experimentally validated this model on a wheel in sand, and showed we can predict the resultant trench shape. This model enables the trenching action described in Chapter 3 and can be used to improve terramechanics models, which will be explored in Chapter 5.

Chapter 5

Terramechanics Modeling

During terrain manipulation, rover wheels are pushed beyond their intended usage and thus beyond the scope of most models concerning them. As detailed in Section 2, there are no existing terramechanics models which cover small wheels with grousers at nonzero slip angle in negative slip (skid). Furthermore, while formulations for wheel forces are valid at high slip angles, none have been experimentally verified beyond 40°. Therefore, in this section we present a new model for predicting the forces on a wheel during terrain manipulations which fall outside of normal driving maneuvers, as well as a set of terramechanics experiments to validate it.

5.1 Nomenclature

See Table 5.1 for a full list of variables used in this section along with their meanings and units.

5.2 Terramechanics Model Development

All values described in this section are given in a wheel-fixed reference frame unless otherwise noted, with the positive x-axis pointed forwards, y-axis out from the face of the wheel, and z-axis pointing up, as shown in Figure 5.1. A wheel traveling with a forward velocity v_x and lateral velocity

Symbol	Value	Units
	Wheel state	
v_x, v_y	Translational velocity along x, y axes	m/s
v	Net translational velocity along direction of travel	m/s
β	Slip angle	$^\circ$
s	Slip ratio	
ω	Wheel angular velocity	rad/s
	Wheel geometry	
r	Radius of wheel	m
r_s	Shearing radius of wheel	m
μ	Grouser area ratio	
h_g	Grouser height	m
b	Width of wheel	m
	Terramechanics variables	
h_0	Sinkage of wheel	m
h	Height of a point relative to the bottom of wheel	m
θ	Angular location of point on wheel	$^\circ$
θ_f, θ_r	Front, rear angles where wheel enters, exits soil	$^\circ$
θ_m	Angle of max normal stress under wheel	$^\circ$
θ_e	Equivalent angle	$^\circ$
θ_0	Shear flow transition angle in skid	$^\circ$
σ_n	Normal stress on wheel surface	kPa
σ_b	Bulldozing stress on side face of wheel	kPa
R_b	Bulldozing resistance on unit width blade	kN/m
q	Soil surcharge pressure	kPa
τ_L, τ_T	Lateral, tangential shear stress under wheel	kPa
j_L, j_T	Lateral, tangential soil deformation	m
v_{jL}, v_{jT}	Lateral, tangential shear rate	m/s
	Tuning parameters	
a_0, a_1	Empirical max stress angle coefficients	
b_0, b_1	Empirical exit angle coefficients	
ζ	Grouser transport volume fraction	
	Soil properties	
c	Cohesion stress	kPa
ϕ	Internal friction angle	$^\circ$
ρ	Soil specific weight	kN/m^3
K	Shear modulus	m
n	Sinkage exponent	
k	Sinkage modulus	kN/m^{n+2}
X_c	Destructive angle	$^\circ$

Table 5.1: Terramechanics variables used in this section.

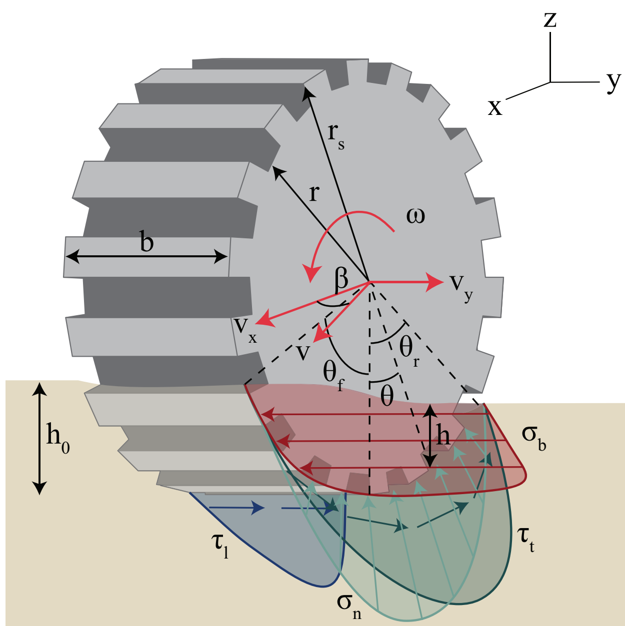

Figure 5.1: Definition of key geometry and state variables for terramechanics models. Descriptions of variables can be found in Table 5.1.

v_y has a net travel velocity $v = \sqrt{v_x^2 + v_y^2}$ and a slip angle given by

$$\beta = -\tan^{-1}\frac{v_y}{v_x} \tag{5.1}$$

For a wheel rotating at an angular velocity of ω, the slip ratio is given by

$$s = \begin{cases} (r_s\omega - v)/(r\omega) & \text{for } |r_s\omega| > |v| \text{ (slip)} \\ (r_s\omega - v)/(v) & \text{for } |r_s\omega| < |v| \text{ (skid)} \end{cases} \tag{5.2}$$

where r_s is the radius at which soil failure occurs along the wheel rim. For a wheel with grousers of height h_g and wheel inner rim radius of r, the shearing radius is assumed to be $r_s = r + h_g$. For a smooth wheel, this radius is assumed to be the wheel's radius.

The wheel-soil contact geometry is primarily defined by two parameters: the wheel's sinkage and soil exit angle. The wheel's sinkage h_0 is the vertical displacement along the z-axis relative to the undisturbed soil surface, as measured at the shearing radius. By convention, angles on the wheel are measured clockwise from vertical in the x-z plane (i.e. about the negative Y axis, from the negative z-axis), as shown in Figure 5.1. In general, a location on the wheel can be defined by its angle as determined by height h from the bottom of the wheel with $\theta = \pm\cos^{-1}(1 - h/r_s)$. The entrance angle θ_f of the wheel is the angle to the point where the wheel first contacts the undisturbed soil on the front, and the exit angle θ_r is the angle to the point where the wheel leaves the soil at the rear. The entrance angle is thus given by $\theta_f = \cos^{-1}(1 - h_0/r_s)$. The sinkage and exit angle are both functions of the wheel state, including a dependence on slip angle and slip ratio. The exit angle can significantly impact the performance of the terramechanics model, but it is typically either empirically tuned or assumed to be uniformly zero for simplicity. In Section 5.2.1, we implement both the standard empirically-tuned method of predicting the exit angle and present a new method to determine θ_r with less tuning.

5.2.1 Prediction of wheel-soil contact geometry

[37] established that the location of the maximum normal stress σ_n under the wheel, θ_m can be approximated as a linear function of slip s, with $\theta_m = \theta_f\,(a_0 + a_1 s)$, where θ_f is the entrance angle and a_0 and a_1 are empirically determined constants for a given wheel and soil, with $0 \leq a_0 \leq 1$, $0 \leq a_1 \leq 1$, and $a_0 + a_1 \leq 1$.

The exit angle θ_r is analogously determined by $\theta_r = \theta_f\,(b_0 + b_1 s)$, with $-1 \leq b_0 \leq 0$, $-1 \leq b_1 \leq 0$, and $b_0 + b_1 \geq -1$. [50] noted that both the maximum stress angle θ_m and exit angle θ_r are constant in skid, calculated based on the entrance angle at zero skid, $\theta_f|_{s=0}$. We refer to these values as θ_{m0} and θ_{r0}, respectively. With this, the maximum stress angle is given by

$$\theta_m = \begin{cases} \theta_f\,(a_0 + a_1 s) & \text{(slip)} \\ \theta_m\big|_{s=0} = \theta_f\big|_{s=0} a_0 = \theta_{m0} & \text{(skid)} \end{cases} \tag{5.3}$$

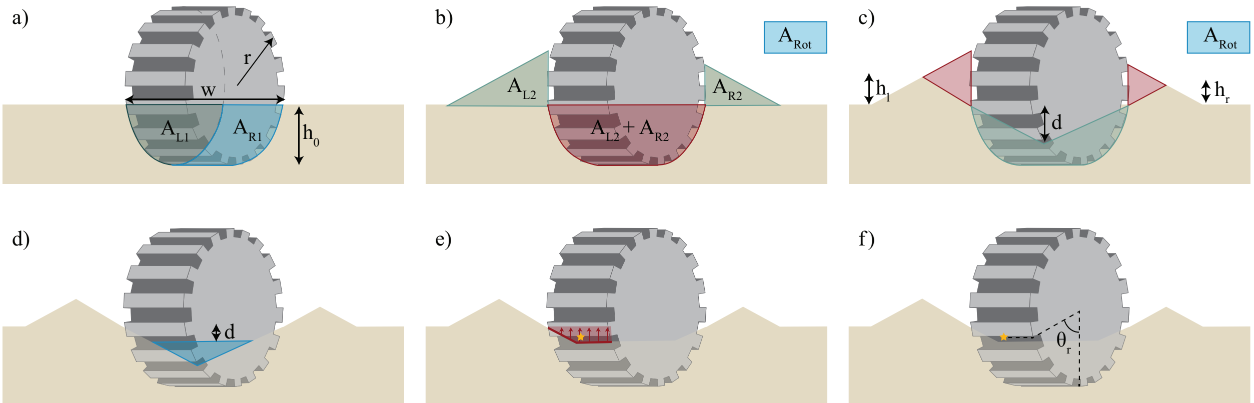

Figure 5.2: Illustration of the soil flow model used to determine exit angle θ_r. (a) The projected area of the wheel is used to separate the terrain into soil contacting the wheel face, A_{L1}, and wheel rim, A_{R1}. (b) The soil caught in the grousers, A_{rot} is separated out, and the remaining soil is piled to the sides of the wheel. (c) The soil is allowed to flow behind the wheel. (d) The soil caught in the grousers is deposited onto the soil profile. (e) The region of the rear rim contacting the soil profile is identified, and the average depth along that portion is identified. (f) The average depth along the contact region is converted to exit angle θ_r. Figure adapted from [13].

and the exit angle is given by

$$\theta_r = \begin{cases} \theta_f\left(b_0 + b_1 s\right) & \text{(slip)} \\ \theta_r\big|_{s=0} = \theta_f\big|_{s=0} b_0 = \theta_{r0} & \text{(skid)} \end{cases} \tag{5.4}$$

Previous works have not explored the possible dependence of a_0, a_1, b_0, and b_1 on slip angle β, so each value is tuned separately for every slip angle. Section 5.4.3 explores the dependence of a_0 and a_1 on β.

The exit angle can alternately be determined through knowledge of the terrain shape behind the wheel. In [13], we developed a closed-form model of the steady-state shape of the rut behind a wheel for non-cohesive soils. In this soil flow model, which is presented in Chapter 4 and summarized in Figure 5.2, the soil is either caught in the wheel's grousers or pushed to the sides of the wheel, allowed to pile up, and then flowed back as the wheel passes through, with soil coming to rest at its angle of repose ϕ. This model relies on a single tuning parameter ζ describing the fraction of sand caught in the grousers which is ultimately transported within them, and otherwise uses only the soil angle of repose ϕ, wheel geometry, and driving state variables.

The trench shape output by the soil flow model can be used to identify the exit angle. The contact angle can vary along the width of the wheel's rim, so a weighted average of the depth along the rear wheel rim is used to determine the exit angle θ_r. This formulation is able to identify the exit angle for any slip angle β and slip ratio s, eliminating the need for the tuning parameters b_0 and b_1.

The terramechanics model described in this section is evaluated in Section 5.4 with both the classic slip-tuned definition of exit angle, (5.4), as well as with this soil flow model.

5.2.2 Pressure-sinkage relationship & normal forces on wheel

The forces on a wheel are largely governed by the normal stress and the shear stress on the wheel rim and sidewall. The normal stress σ_n is zero at the entrance angle θ_f and at the exit angle θ_r, and it takes a maximum value at an angle of θ_m, the location of which is governed by the wheel's slip. [35] observed that the normal stress at a point on the wheel rim can be described by

$$\sigma_n = kh^n \tag{5.5}$$

for a given depth h, where k and n are the sinkage modulus and sinkage exponent, respectively, and are theoretically intrinsic soil properties based on the pressure-sinkage relationship of a flat plate. k is commonly broken into a constant term and a wheel width dependent term as in [38], though [47] noted that additional corrections are required for small radius wheels as the stress distribution diverges from that of a flat plate. Proper identification of each of these parameters requires testing a wide parameter space of wheel geometries; for simplicity in this section the values of n and k are identified separately for each wheel.

A constant value of n only accounts for static sinkage of the wheel and does not accurately the capture slip sinkage and skid sinkage phenomena. Thus, this model treats n as a linear function of slip ratio in both the slip regime, as in [41] and [43], and the skid regime, as in [44]. The sinkage exponent n for a given wheel and soil is

$$n = \begin{cases} n_0 + n_1 s & \text{(slip)} \\ n_0 - n_2 s & \text{(skid)} \end{cases} \tag{5.6}$$

where n_0, n_1, and n_2 are all positive. For a rotating wheel, the normal stress is zero at the edges of the contact region given by the entrance angle θ_f and exit angle θ_r, and reaches a maximum at some angle θ_m, which is a function of the slip ratio and other factors. The normal stress at a given angle θ along the wheel is given by applying (5.5) with $h = r_s - r_s \cos\theta$ at the grouser tips and $h = r - r\cos\theta$ on the wheel rim, giving

$$\sigma_n(r) = kr^n(\cos\theta_e - \cos\theta_f)^n \tag{5.7}$$

where θ_e is the equivalent angle, which is defined in front of and behind θ_m as

$$\theta_e = \begin{cases} \theta & \text{for } \theta_m < \theta < \theta_f \\ \theta_f - \frac{\theta - \theta_r}{\theta_m - \theta_r}(\theta_f - \theta_m) & \text{for } \theta_r < \theta < \theta_m \end{cases} \tag{5.8}$$

The equivalent angle formulation ensures that the stress reaches a maximum value at $\theta = \theta_m$ and that the effective depth at the exit angle is zero, so that the normal stress vanishes to zero at the edges of the contact region. The total normal stress is a weighted average of the stress at the grouser tip radius, $\sigma_n(r_s)$, and the radius of the surface of the wheel, $\sigma_n(r)$, yielding

$$\sigma_n = \mu\sigma_n(r_s) + (1-\mu)\sigma_n(r) \tag{5.9}$$

where μ is the grouser area ratio, or the fraction of the wheel's surface taken up by the grousers, and is one for a smooth wheel, as in [53].

5.2.3 Shear forces on wheel

The shear stress under the wheel is tangential to the wheel's surface and broken into two components: lateral, τ_L, which is parallel to the wheel's y-axis, and tangential, τ_T, which is perpendicular to the y-axis. Shear forces at each point on the wheel are a function of the normal stress σ_n applied to the same point on the wheel, as originally defined by [39] for two dimensional soil failure, such that

$$\tau = (c + \sigma_n \tan\phi)(1 - e^{-j/K}) \tag{5.10}$$

where c is the soil cohesion, ϕ is the soil friction angle, K is the shear modulus, and j is the shear displacement a unit of soil undergoes before reaching the given angular position on the wheel.

The scaling of the shear stress τ along each component, τ_T and τ_L, is based on a fusion of approaches. As in [58], the tangential and lateral shear deformations are combined into a single net shear deformation, scaled by a single modulus K. The magnitude of each component is scaled by the relative magnitude of the shear velocity as in [53], such that the shear stress is zero when the shear velocity is zero. This yields

$$\tau_i = \frac{v_{ji}}{\sqrt{v_{jT}^2 + v_{jL}^2}}(c + \sigma_n \tan\phi)\left(1 - e^{-\sqrt{j_T^2 + j_L^2}/K}\right) \tag{5.11}$$

where $i \in \{T, L\}$, v_{jT} and v_{jL} are the shear velocity of the soil tangential to the rim in the x-z plane and shear velocity along the y-axis, respectively, and j_T and j_L are the corresponding shear deformations. Each shear deformation is obtained by integrating the shearing velocity of the soil, or the velocity difference between the wheel surface and soil at its interface, with respect to time. For each shear velocity, the corresponding shear deformation is

$$j_i = \int_0^t v_{ji} dt \tag{5.12}$$

The tangential and lateral shear velocities for a wheel in positive slip as defined in (5.2) are given by

$$v_{jT} = r_s\omega - v_x\cos\theta \qquad \text{Tangential, slip} \tag{5.13}$$

$$v_{jL} = -v_y \qquad \text{Lateral, slip} \tag{5.14}$$

When the wheel is in slip, the soil under the wheel is always flowing from front to back in parallel to the wheel's rim in the x-z plane. In skid, a portion of the soil under the wheel's surface is instead pushed along the front of the wheel, flowing opposite to the wheel's rotation, as described in [36]. This results in a transition point within the wheel-soil contact area where the tangential shear velocity is zero, which is identified by the angle θ_0. We approximate θ_0 as a function of slip ratio as in [43], as

$$\theta_0 = \theta_f\left(1 + \frac{\omega r_s - v}{2v}\right) \tag{5.15}$$

In the front region, the shear velocity in skid is given by

$$v_{jT} = r_s\omega - v_x\cos\theta + K_v v_x \tag{5.16}$$

where the coefficient K_v can be determined as in [36] by setting the tangential shear deformation j_T equal to zero at θ_0. The shear velocity in skid for the rear region of the wheel is the same as in slip, as given in (5.13). The lateral shear velocity is equivalent in slip and skid, as in (5.14). Thus the tangential and later shear velocities for a wheel in skid are

$$v_{jT} = v_y \qquad \text{Tangential, skid, } \theta < \theta_0 \tag{5.17}$$

$$v_{jT} = -v_x\cos\theta + \frac{\sin\theta_f - \sin\theta_0}{\theta_f - \theta_0}v_x \qquad \text{Tangential, skid, } \theta > \theta_0 \tag{5.18}$$

$$v_{jL} = v_y \qquad \text{Lateral, skid} \tag{5.19}$$

Using the substitution $\omega = \frac{d\theta}{dt}$, the tangential shear deformation at a point θ can be found by

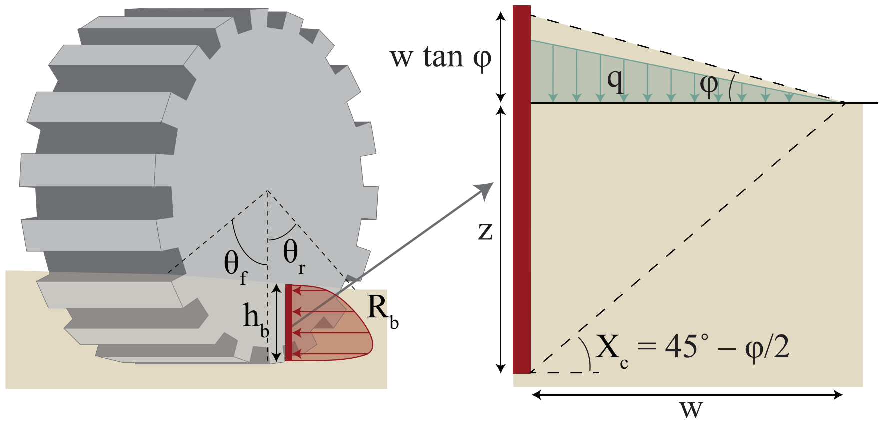

Figure 5.3: Side view of a wheel with detail of the bulldozing stress (left) and closeup of a single unit-width blade for calculating bulldozing stress (right).

integrating (5.12) with either (5.13), (5.17), or (5.18) to get

$$j_T = \begin{cases} r_s(\theta_f - \theta) - (\sin\theta_f - \sin\theta)\frac{v_x}{\omega} & \text{Slip} \\ r_s(\theta_0 - \theta) - (\sin\theta_0 - \sin\theta)\frac{v_x}{\omega} & \text{Skid, } \theta < \theta_0 \\ \left(\frac{\sin\theta_f - \sin\theta_0}{\theta_f - \theta_0}(\theta_f - \theta) - (\sin\theta_f - \sin\theta)\right)\frac{v_x}{\omega} & \text{Skid, } \theta > \theta_0 \end{cases} \tag{5.20}$$

The lateral shear deformation is the distance a particle of soil on the wheel gets displaced along the lateral direction. Using the same $\omega = \frac{d\theta}{dt}$ substitution, we integrate (5.12) with (5.14) or (5.19) to get

$$j_L = \int_{\theta}^{\theta_f} \frac{v_{jL}}{\omega} d\theta = (\theta_f - \theta)\frac{v_y}{\omega} \qquad \text{Slip and skid} \tag{5.21}$$

Note that when $\omega = 0$, j_L is infinite. If the wheel's rotation does not move a particle into and out of contact with the wheel at the edges, it will (in theory) stick to the wheel and continue to move with it indefinitely, representing an infinite shear deformation. This results in a lateral shear deformation of $\tau_L = c + \sigma_n \tan\phi$.

5.2.4 Sidewall forces on wheel

When the wheel is angled relative to its direction of travel there are forces on the side face of the wheel. The most commonly used method of determining the force on the side face was presented by [58] and derived from [66], but we take a different approach here based on soil cutting theory that is more widely used than [66]. In this work, we determine the bulldozing stress σ_b for a point at depth z on the side face of the wheel from [38]'s formulation for passive earth pressure

$$\sigma_b = \rho z \cot^2 X_c + q \cot^2 X_c + 2c \cot X_c \tag{5.22}$$

for soil density ρ, cohesion c, destructive angle $X_c = 45° - \phi/2$ (which defines the line along which shear failure propagates through the soil in front of the blade, as seen in Figure 5.3), angle of internal friction ϕ, and weight of the soil surcharge above the undisturbed soil surface plane q.

The pressure due to this surcharge is given by the weight of the soil piled along the wheel's side surface which reaches the point where the shear line meets the soil surface and has angle of repose ϕ, as shown in Figure 5.3. Based on this, we derive the soil surcharge as

$$q = \frac{1}{2}\rho z \cot X_c \tan\phi \tag{5.23}$$

The force on a unit-wide blade R_b is taken by integrating the passive earth pressure σ_b over the depth of the wheel, where the depth of the bottom of the wheel at a given angular location is h_b, yielding

$$\begin{aligned} R_b &= |\sin\beta| \int_0^{h_b} \rho z \cot^2 X_c + \frac{1}{2}\rho z \cot^3 X_c \tan\phi + 2c \cot X_c dz \\ &= |\sin\beta| \left(\frac{1}{2}\rho h_b^2 \cot^2 X_c (1 + \frac{1}{2}\cot X_c \tan\phi) + 2h_b c \cot X_c \right) \end{aligned} \tag{5.24}$$

The term $|\sin\beta|$ is included to incorporate the inclination of the wheel's face relative to the direction of travel, as this formulation is not dependent on velocity. The profile of the undisturbed soil along the wheel's side surface is assumed to follow a straight line between the entrance and exit angles θ_f

and θ_r. The depth h_b measured to this undisturbed soil surface is given by

$$h_b = \frac{\sin\theta - \sin\theta_r}{\sin\theta_f - \sin\theta_r}(\cos\theta_r - \cos\theta_f)r_s - (\cos\theta_r - \cos\theta)r_s \tag{5.25}$$

[58] uses an alternate method to determine the bulldozing resistance based on [66]. However, the two formulations are slightly different, with each paper's formulation for R_b given by

$$\begin{aligned} R_b &= \quad \rho h_b^2 \cot^2 X_c \left(1 + \cot X_c \tan\phi\right) + 2h_b c \cot X_c \text{ [58]} \\ R_b &= \frac{1}{2}\rho h_b^2 \cot^2 X_c \left(1 + \cot X_c \tan\phi\right) + 2h_b c \cot X_c \text{ [66]} \end{aligned}$$

The equations above have been rearranged for visual similarity, and simplified using the fact that $\tan(X_c+\phi) = \tan(\pi/4-\phi/2+\phi) = \tan(\pi/4+\phi/2) = \tan(\pi/2-(\pi/4-\phi/2)) = \tan(\pi/2-X_c) = \cot X_c$. Compared to [66], [58] is off by a factor of $\frac{1}{2}$ in the first term and therefore has gained an additional term of $\frac{1}{2}\rho h_b^2 \cot^2 X_c \left(1 + \cot X_c \tan\phi\right)$. The formulations by both [58] and [66] differ from the version derived here from [38], with Hegedus having an additional term of $\frac{1}{4}\rho h_b^2 \cot^3 X_c \tan\phi$ compared to Wong.

5.2.5 Net forces and moments on wheel

The net forces on the wheel surface can be found by integrating the stresses given in equations (5.5) and (5.11) over the wheel, giving

$$F_x = r_s b \int_{\theta_r}^{\theta_f} -\sigma_n \sin\theta + \tau_T \cos\theta d\theta \tag{5.26}$$

$$F_y = r_s b \int_{\theta_r}^{\theta_f} \frac{R_b \sin\beta}{b} \cos\theta + \tau_L d\theta \tag{5.27}$$

$$F_z = r_s b \int_{\theta_r}^{\theta_f} \sigma_n \cos\theta + \tau_T \sin\theta d\theta \tag{5.28}$$

Note that [58] integrates the bulldozing stress differently, such that the sidewall force given in (5.27) is instead given by

$$F_y = r_s b \int_{\theta_r}^{\theta_f} \frac{1}{b}\left(1 - \frac{1}{r_s} h(\theta)\cos\theta\right) R_b \sin\beta + \tau_L d\theta \tag{5.29}$$

which does not correspond to integrating horizontally across the wheel's surface as would be consistent with the definition of bulldozing resistance as derived from [66].

The moments on the wheel are similarly given by

$$M_x = r_s^2 b \int_{\theta_r}^{\theta_f} -\tau_L \cos\theta + \frac{R_b \sin\beta \cos\theta h_{COP}}{r_s b} d\theta \tag{5.30}$$

$$M_y = r_s^2 b \int_{\theta_r}^{\theta_f} -\tau_T d\theta \tag{5.31}$$

$$M_z = r_s^2 b \int_{\theta_r}^{\theta_f} -\tau_L \sin\theta + \frac{R_b \sin\beta \sin\theta \cos\theta}{b} d\theta \tag{5.32}$$

where h_{COP} is the height of the center of the bulldozing pressure for an individual unit-width blade as shown in Figure 5.3, and for the bulldozing stress in Equation 5.24 is given by

$$h_{COP} = \frac{\frac{1}{3} h_b^2 \rho \cot^2 X_c (1 + \frac{1}{2}\cot X_c \tan\phi) + \frac{1}{2} h_b c \cot X_c}{\frac{1}{2} h_b \rho \cot^2 X_c (1 + \frac{1}{2}\cot X_c \tan\phi) + c \cot X_c} \tag{5.33}$$

Equations 5.26 – 5.32 assume the wheel's sinkage h is known. If instead, the vertical load $F_{z_{applied}}$ is known, the sinkage can be found by performing a binary search over possible sinkage values until the applied load matches the computed load within a given tolerance ε, such that

$$|F_{z_{applied}} - F_z| \leq \varepsilon \tag{5.34}$$

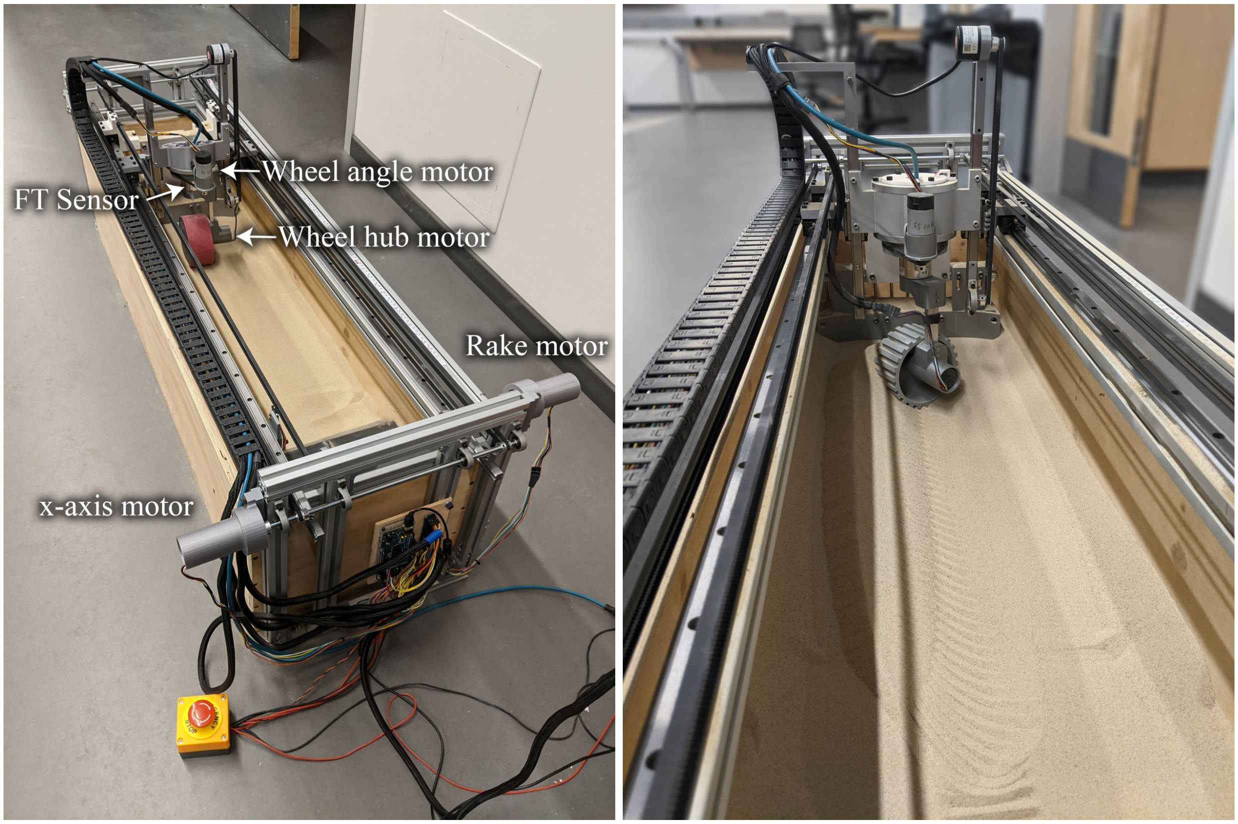

Figure 5.4: Labelled image of the single wheel terramechanics testbed (left) and grousered wheel in the testbed at the end of a single test run (right).

5.3 Experimental Validation of Terramechanics Model

5.3.1 Terramechanics testbed

Figure 5.4 shows the terramechanics testbed used to measure forces and sinkage of a wheel operating at fixed slip angle and slip ratio. The testbed is based on NASA Ames Research Center's terrain manipulation testbed developed by Loic Tissot-Daguette as shown in [98]. It allows a wheel to be driven at a forced slip angle and slip ratio along a prepared sandbox while recording forces and torques on the wheel and the wheel's sinkage.

The testbed consists of a main belt-driven x-axis gantry mounted to long linear rails, with a frame constructed from slotted aluminum profile connected by 3D printed brackets. The gantry has shorter vertical linear rails to allow for free vertical translation, an actuated rotational axis to control wheel angle, and a hub motor for driving the wheel. Thus, the wheel can be driven at a fixed slip ratio and slip angle, with load applied by weights loaded into the main gantry. A Bota Systems Rokubi Serial 6-axis force-torque sensor is mounted between the wheel angle motor and the wheel driving hub motor [99]. The x-axis rails also support a rake and smoothing mechanism, which is used for repeatable soil preparation between experiments. For each trial, the soil preparation mechanism sweeps the full length of the testbed, raking the soil with tines separated by 3 cm and covering the full width of the testbed. The soil preparation mechanism also incorporates a blade at the rear, so that the soil is smoothed and leveled to the same state for every trial.

A full set of tests was run on two wheels, grousered and smooth, shown in Figure 5.5 with the parameters given in Table 5.2. Each wheel was run at all combinations of the slip ratios {-1, -.7, -.5, -.2, 0, .2, .5, .75, .9} and slip angles {0°, 15°, 30°, 45°, 60°, 75°, 90°}, with three trials performed for each run.[1] For both wheels over all test conditions, repeat measurements of the same trial had an average spread of below 1.9N for force measurement and below 3mm for sinkage measurements. All forces are relative to a wheel-fixed frame of reference as shown in Figure 5.1. The sinkage is measured relative to the prepared soil surface, and does not reflect the depth of soil pile buildup in front of the wheel, with a larger sinkage value corresponding to the wheel being embedded deeper into the soil. The testbed contains Soil Direct #90 sand [100]. Forces and torques are measured while the wheel travels 850 mm across prepared sand, with force data extracted from the steady-state portion of the wheel's motion. For most runs, steady-state motion occurred from 500mm to 750mm. In several low slip angle, low slip ratio trials the sinkage was high enough that the wheel drive motor hub contacted the sand, increasing the observed forces. The trials at slip angles of 75° and 90° did not reach steady state sinkage over the full length of the testbed and are thus excluded from analysis here but included in the data set for completeness.

[1]The raw and processed data for all trials can be found at https://github.com/robomechanics/3d-terramechanics

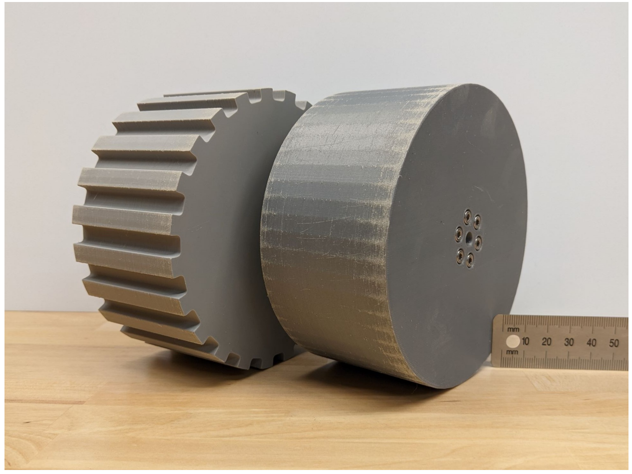

Figure 5.5: Wheels used in terramechanics experiments. Grousered (left) and smooth (right) wheels tested.

	r (mm)	b (mm)	h_g (mm)	μ
Smooth Wheel	62.5	60.0	0	1
Grousered Wheel	57.5	60.0	5	.5

Table 5.2: Wheel geometry parameters for terramechanics experiments.

5.3.2 Soil characterization and model tuning

As only some key terramechanics properties for the Soil Direct #90 sand used were provided in datasheets, the remaining values were obtained though a combination of measurement and tuning, the results of which can be found in Table 5.3. The soil's specific weight ρ was measured at 13.03 kN/m^3 by measuring the volume and density of a sample of soil from the testbed with the soil in a loosely packed state. ϕ was provided in soil datasheets, and the testbed was modified to measure the sinkage modulus k of the soil. For this measurement, the testbed's wheel was replaced with circular plates with 60mm, 80mm, and 100mm diameters. In order to measure the pressure-sinkage relationship of the sand, the same soil preparation was used as in the single wheel tests and each plate was inserted slowly into the soil while measuring sinkage and load on the plate. The measured sinkage modulus k was 8000 kN/m^{n+2}. The classic sinkage exponent n was measured as well, though this value was not used as the model described here formulates the sinkage exponent as a function of slip ratio and slip angle.

Hand tuning was used to select values for the cohesion c and shear deformation modulus K as

Wheel	k (kN/m^{n+2})	ϕ (°)	ρ (kN/m^3)	c (kPa)	K (m)	n_0	n_1	n_2
Smooth	8000	29	13.03	1.0	0.021	1.46	0.01	0.55
Grousered	8000	29	13.03	1.0	0.021	1.46	0.01	0.74

Table 5.3: Soil parameters for terramechanics tests. k, ϕ, and ρ were measured, all other values were tuned.

in [101] and [44], along with the sinkage exponent coefficients n_0, n_1, and n_2 for both wheels, with the values reported in Table 5.3.Note that while the static component of the sinkage exponent n_0 is a soil property, the wheel's sinkage due to slip and skid depend on the wheel's grouser design and other wheel-specific factors, and thus these values differ between the two wheels tested here. Hand tuning was performed by starting from typical values for sand and incrementally varying values to improve the error on the model's prediction of F_x, F_y, and the sinkage h over all tests.

As high slip angles have not been investigated in depth in existing work, the relation between slip angle and a_0, a_1, b_0, and b_1 is not known, and it cannot be assumed that these values are independent of slip angle. MATLAB's *fmincon* was used to determine a_0, a_1, b_0, and b_1 for each slip angle β separately. This was done by minimizing the sum squared error of the forces on the wheel driving at a given slip angle over all measured slip ratios via the following procedure, with other soil parameters set to typical values:

1. Set initial parameter guesses

2. Set wheel sinkage to measured value

3. Compute expected forces F_x, F_y, F_z from terramechanics model

4. Compute sum squared error of forces normalized to measured values of F_x, F_y, F_z

The resultant values are given in Table 5.4, though the number of parameters required to run the model can be greatly reduced from those listed here, as discussed in Section 5.4.3. Notably, the values of b_0 and b_1 are not used in the model implementation that uses soil flow to predict exit angle, and a_0 and a_1 can be fit as linear functions of β.

β	a_0	a_1	b_0	b_1	ζ
Smooth wheel					
0°	0.20	0.70	-0.48	-0.00	0.10
15°	0.22	0.70	-0.47	-0.01	0.10
30°	0.57	0.20	-0.51	-0.01	0.10
45°	0.63	0.17	-0.46	-0.15	0.10
60°	0.72	0.13	-0.54	-0.13	0.10
Grousered wheel					
0°	0.27	0.67	-0.66	-0.18	0.20
15°	0.29	0.65	-0.50	-0.41	0.20
30°	0.47	0.45	-0.44	-0.56	0.20
45°	0.62	0.30	-0.51	-0.49	0.20
60°	0.71	0.21	-0.53	-0.47	0.20

Table 5.4: Tuning parameters for terramechanics model. a_0, a_1, b_0, and b_1 were tuned with the optional soil flow model disabled.

5.4 Terramechanics Modeling Results and Discussion

5.4.1 Force prediction

The model detailed in Section 5.2 can be seen overlaid on the data collected in Section 5.3 in Figure 5.6.[2] In general, the load on a wheel is known and we wish to determine the wheel's sinkage and tractive forces, so here we take the measured vertical load on the wheel for each trial, perform a binary search over sinkage to match the applied load, and plot the resultant forces along the x and y axes along with the predicted sinkage. The model performance is presented both with classic tuning of the exit angle (green, solid), and with the soil flow model-based exit angle calculation presented in Section 5.2.1 (red, dashed). Additionally, [58]'s implementation of sidewall force is shown for sake of comparison (blue, dotted).

The model captures the trends of wheel-soil interaction forces for both wheels, with better performance in slip than skid for forces along the wheel's x-axis. For the smooth wheel, the average error in predicted force along the x-axis over all slip ratios ranged from 8.3 - 16.1% of the applied vertical load without the soil flow model, and 22.3 - 35.4% with the soil flow model, as reported in Table 5.5 for each slip angle tested. The model performed similarly on the grousered wheel,

[2]The code for generating the plots in this section can be found at
https://github.com/robomechanics/3d-terramechanics

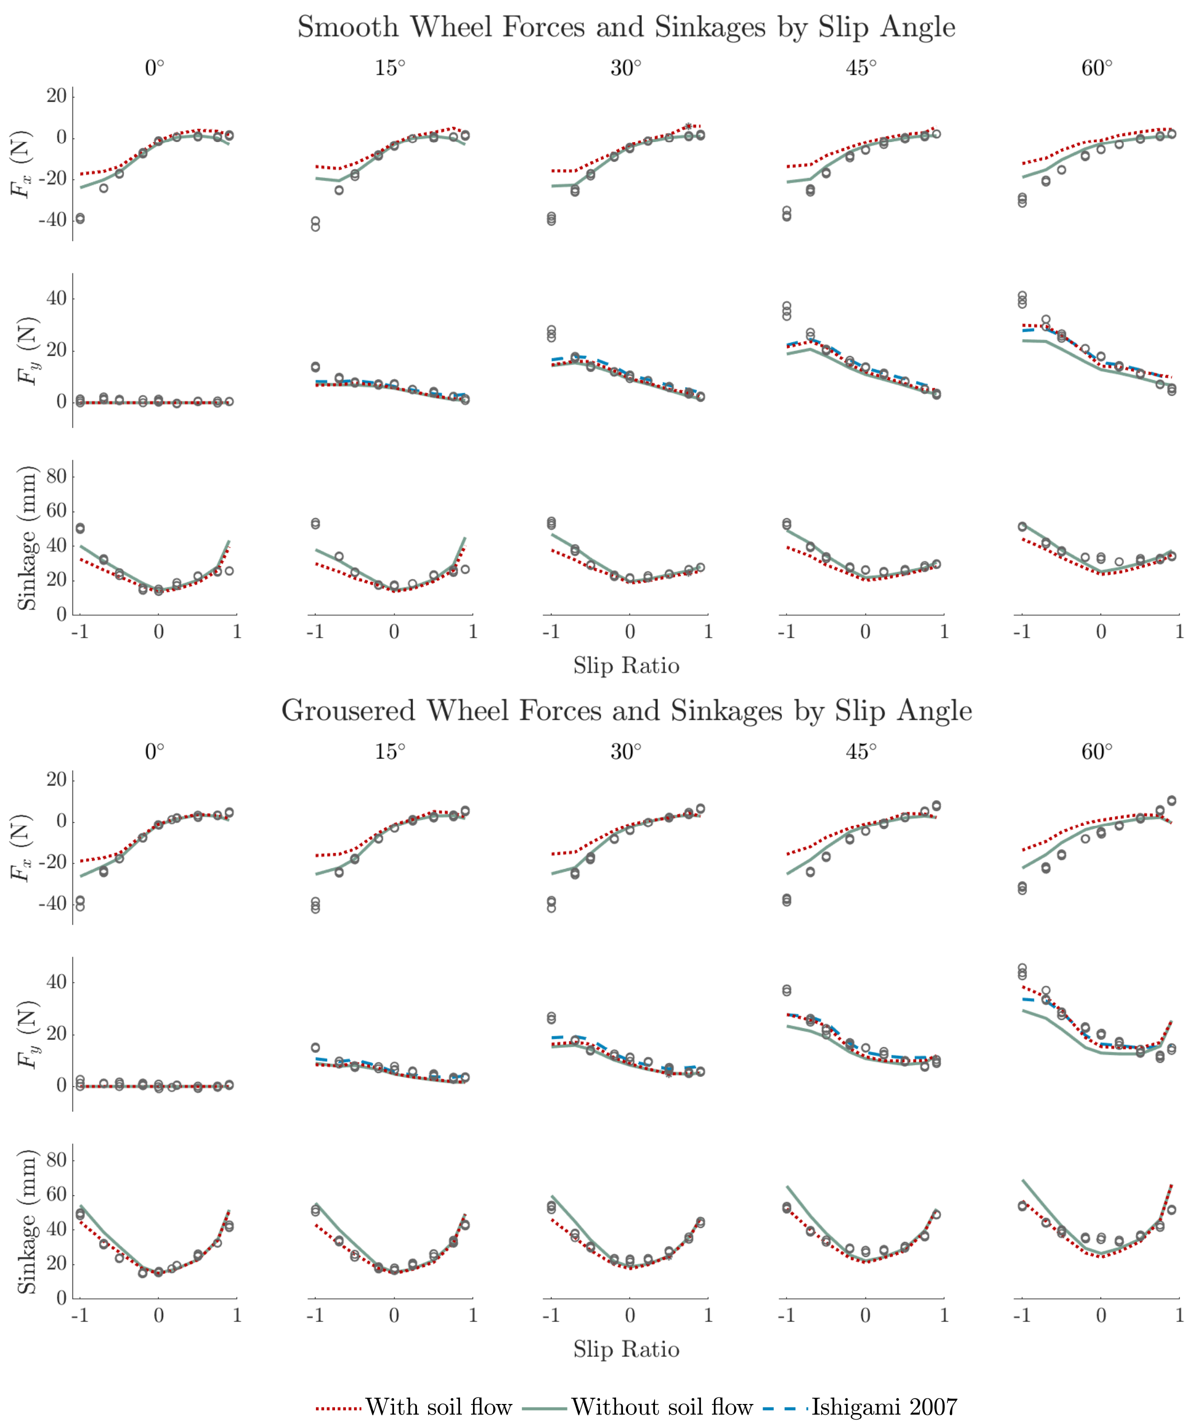

Figure 5.6: Performance of the terramechanics model at predicting forces and sinkages for a given slip angle and applied load on smooth and grousered wheels. The model is shown with the new exit angle formulation derived from soil flow (red, dotted) and with the standard tuned exit angle (green, solid). [58]'s implementation of bulldozing force is presented for comparison (blue, dashed).

with ranges of 7.1 - 12.1% error and 15.6 - 27.2% error for without and with the soil flow model, respectively.

The model performed well at prediction of forces along the wheel's y-axis, which are dominated by the bulldozing force on the sidewall. For the smooth wheel, the average error in force along the y-axis over all slip ratios ranged from 9.0 - 23.2% without the soil flow model and 4.9 - 11.1% with the soil flow model, excluding $\beta = 0°$, which nominally has $F_y = 0$. The model had comparable performance at predicting sidewall force on the grousered wheel, with 2.3 - 19.7% error and 0 - 9.3% error without and with the soil flow model, respectively. The y-axis forces are comparable to those predicted by [58], but not identical.

Overall, the addition of the soil flow model to predict the exit angle moderately reduced predictive quality of the model for forces along the x-axis but increased predictive quality of forces along the y-axis.

5.4.2 Sinkage prediction

The model has excellent prediction of wheel sinkage both with and without the soil flow model, with less than 7.9% average error relative to wheel radius for each slip angle. Sinkage prediction was accurate on both wheels in both slip and skid, as seen in Figure 5.6 and Table 5.5.

5.4.3 Extrapolation of tuning parameters

As noted in Section 2, previous works incorporating nonzero slip angle have not investigated the impact of slip angle on the location of the maximum stress angle θ_m or the exit angle θ_f. The soil flow model presented in Section 5.2.1 and validated in [13] answers this question for exit angle, but does not address the maximum stress angle.

In this work, we have separately tuned values of a_0 and a_1 for each slip angle β tested, with the maximum stress angle defined for a given slip angle and slip ratio in (5.3). While tuning discrete values of a_0 and a_1 is sufficient for model validation, continuous definitions of a_0, a_1, and θ_{m0} are necessary for a full model that is continuous over all possible values of β.

Slip Angle	F_x Error (%)	F_y Error (%)	h_0 Error (%)
Smooth wheel			
Without soil flow model			
0°	8.3 ± 25.7	-2.9 ± 2.3	2.2 ± 11.9
15°	14.7 ± 39.7	-9.0 ± 10.5	0.0 ± 14.2
30°	11.5 ± 26.8	-11.5 ± 20.0	-0.7 ± 4.5
45°	16.1 ± 26.1	-19.2 ± 26.6	-1.7 ± 3.6
60°	16.1 ± 19.2	-23.2 ± 26.9	-1.6 ± 5.8
With soil flow model			
0°	22.3 ± 32.8	-2.9 ± 2.3	-3.1 ± 13.4
15°	31.2 ± 46.1	-8.1 ± 12.0	-5.9 ± 15.6
30°	29.5 ± 36.5	-8.1 ± 20.8	-6.4 ± 7.6
45°	33.6 ± 36.6	-11.1 ± 24.1	-7.9 ± 5.9
60°	35.4 ± 26.3	-4.9 ± 21.4	-7.7 ± 5.0
Grousered wheel			
Slip Angle	F_x Error (%)	F_y Error (%)	h_0 Error (%)
Without soil flow model			
0°	7.1 ± 21.2	-2.3 ± 2.6	5.5 ± 6.8
15°	9.7 ± 25.2	-9.1 ± 8.7	3.3 ± 6.7
30°	11.5 ± 24.9	-12.4 ± 17.2	2.2 ± 6.8
45°	12.1 ± 23.6	-15.5 ± 22.1	3.5 ± 10.0
60°	10.3 ± 25.9	-19.7 ± 33.4	4.8 ± 14.2
With soil flow model			
0°	15.6 ± 32.3	-2.3 ± 2.6	1.6 ± 6.4
15°	23.8 ± 38.5	-8.3 ± 10.2	-2.6 ± 6.4
30°	25.9 ± 41.0	-9.3 ± 17.0	-4.2 ± 4.3
45°	26.9 ± 38.2	-6.0 ± 17.6	-2.5 ± 5.7
60°	27.2 ± 38.4	0.0 ± 22.0	-1.6 ± 12.6

Table 5.5: Force and sinkage percentage errors for the terramechanics model on the grousered and smooth wheels. Subsections of table (sections of rows) for the three different model combinations. Forces are normalized to the applied load and sinkages to the wheel radius.

The values of a_0 and a_1 obtained in Section 5.3.2 for both wheels are shown in Figure 5.7. A linear fit as a function of slip angle β approximates the values of a_0 and a_1 very well for the grousered wheel, but less well for the smooth wheel, as shown in Figure 5.7. For the grousered wheels and soil tested in this work we can then use continuous definitions of a_0 and a_1 to apply the model presented here to arbitrary slip angles. Further experiments which directly measure the stress distribution under wheels are needed to conclusively determine the impact of steering angle on normal stress, but a linear approximation can be used for this implementation.

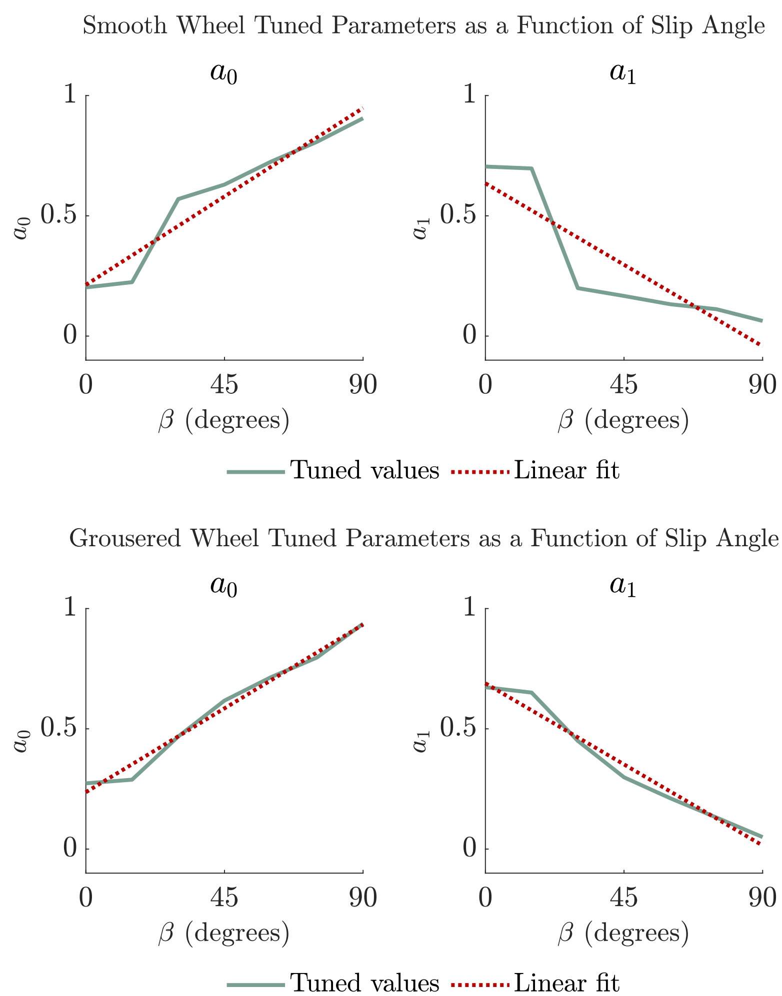

Figure 5.7: Tuned values of a_0, a_1, and θ_{m0} and linear fit to approximate to intermediate slip angles on both a smooth and grousered wheel.

5.5 Conclusions

In this chapter we present a terramechanics model capable of describing wheel-soil interaction forces over all possible slip and skid states, including slip angle, which has been previously unexplored. In addition to fusing existing terramechanics and soil failure models, we introduce a closed-form soil flow model to better determine the wheel-soil contact geometry without extensive measurement and tuning. We also present a terramechanics dataset covering a range of states previously unexplored, particularly high slip angle, on two small wheels in sand. This dataset is used to validate the terramechanics model on all slip ratios and on slip angles up to 60°. In tuning the maximum stress angle relationship, we observed that the parameters describing the maximum stress angle follow a linear dependence on slip angle. Further work directly measuring the stress distributions under a wheel at high slip angle is needed to confirm this relationship. Furthermore, the model proposed here is validated on a limited dataset consisting of two wheels and a single soil type, with a single applied load; additional testing is required for verification of applicability to a broader range of wheel geometries and soils.

The model presented in this chapter enables prediction of wheeled mobile robot mobility during off-nominal driving scenarios. This includes unintended mobility changes, such as the failure of a steering actuator or driving actuator, in which case the rover may experience skid conditions with a high slip angle or a variety of slip ratios simultaneously on different wheels. Off-nominal driving scenarios also include the intentional use of wheels for manipulation of terrain, such as by digging trenches. With the ability to model both slip and skid conditions for a wheel at arbitrary slip angles, it is now possible to automatically generate driving strategies to both recover from mobility failures and enable safe driving during intentional off-nominal driving conditions.

Chapter 6

Hardware Demonstrations of Off-Nominal Driving

Thus far, we have laid out a roadmap for Nonprehensile Terrain Manipulation (NPTM) and created models needed for both trenching and driving with failed actuators, but have not ourselves shown that either are possible on hardware. In this chapter, we show off-nominal driving on full size rovers. We begin with NPTM demonstrations in a Martian analog environment, showing a large rover excavate trenches, dig holes, and clear paths for a smaller robot. We then assess the impact of actuator failure on the mobility of VIPER in depth through a series of quantitative and qualitative driving tests, and then demonstrate hand-tuned driving strategies to compensate for two types of failure on a rover the same size as VIPER.

6.1 Field Demonstrations of Nonprehensile Terrain Manipulation

In this section we demonstrate wheel-based soil manipulation on a large rover in a Martian analog environment, illustrating various applications for NPTM. The KREX-2 rover 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 at NASA Ames Research Center's Roverscape and in the Atacama Desert in Chile. The Atacama Desert is a common Martian analog, with almost no annual rainfall and no plant or animal life in many areas. For these tests, the KREX-2 rover [90] was outfitted with inflated rubber tires for all experiments unless noted otherwise. The approximate wheel geometry is described in Table 6.1. All demonstrations described here are performed using a single rover wheel to manipulate terrain. These demonstrations show the application of several robot actions described in Tables 3.1 and 3.2, and how they can be applied to mission scenarios. Additionally, the trenching actions performed are also used to test the soil flow model developed in Chapter 4 on unprepared natural terrain.

6.1.1 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 motion and providing another data set to compare the model developed in Chapter 4 against. The soil was soft, dry, and relatively 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. The KREX-2 rover can be seen in Figure 6.1 after completing a trenching test.

Terrain was mapped with a FARO® LIDAR scanner [102] (Figure 6.2) before and after each test, with fiducial markers mounted to the rover used to align the rover'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 [103]. 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. Each test was recorded with a GoPro HERO camera [104] mounted to the rover chassis and an off-board video camera.

Figure 6.1: KREX-2 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.

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 Figure 6.2. 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 of 0, 30°, 60°, and 90°. The speeds and angles of the other wheels were hand tuned in the field to allow the rover to drive straight while trenching. Chapter 7 describes how to automate this process.

Each trenching action was run 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.

Despite being dug at the same wheel speeds and in the same terrain, four very distinct trench geometries were formed, as seen in Figures 6.3 and 6.4. Trenches ranged from 18 cm wide and 1.3

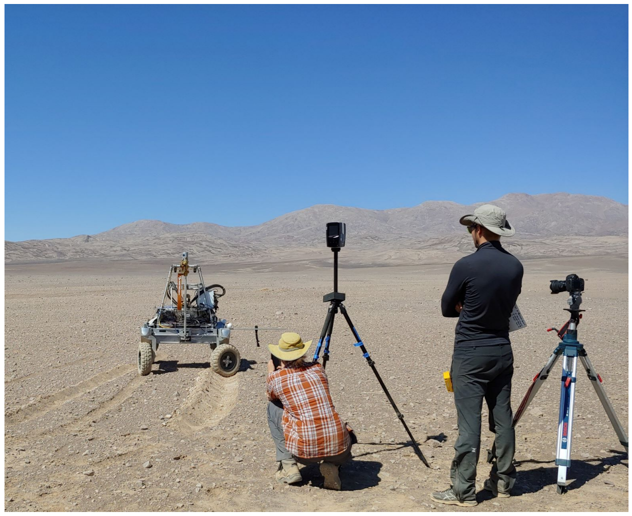

Figure 6.2: Setup of trenching experiments in the Atacama Desert. The KREX-2 rover (left) was driven over unprepared soil, with the FARO® LIDAR (center) scanning before and after and video recording of each test (right). Not pictured is the Leica Total Station used to track the rover's position and speed during testing [103].

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

Table 6.1: 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.

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 6.2. 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 Chapter 4 using the same analysis techniques. We present the results of the trench prediction in Figure 6.4 and Table 6.2. 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 Figure 6.4. 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 6.1.3.

Figure 6.3: Trenches dug by the KREX-2 rover 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 6.4: 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°.

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 ζ 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 Figures 6.1 and 6.3 the color difference between the surface and subsurface soil is clearly visable, 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

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

Table 6.2: Trench geometry and quality of soil flow model fit for the four trenches dug by KREX-2 in the Atacama.

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 modelled with the soil flow model described in Chapter 4.

6.1.2 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 trench dimensions are reported in Table 6.3, and images of the three passes along with LIDAR scans of the trench can be seen in in Figure 6.5. Multiple passes over the same trench had 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 Figure 6.5. 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 visual odometry, which we perform in Section 7.3.1, would be needed for effective re-excavation 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

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

Table 6.3: 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° to the direction of travel.

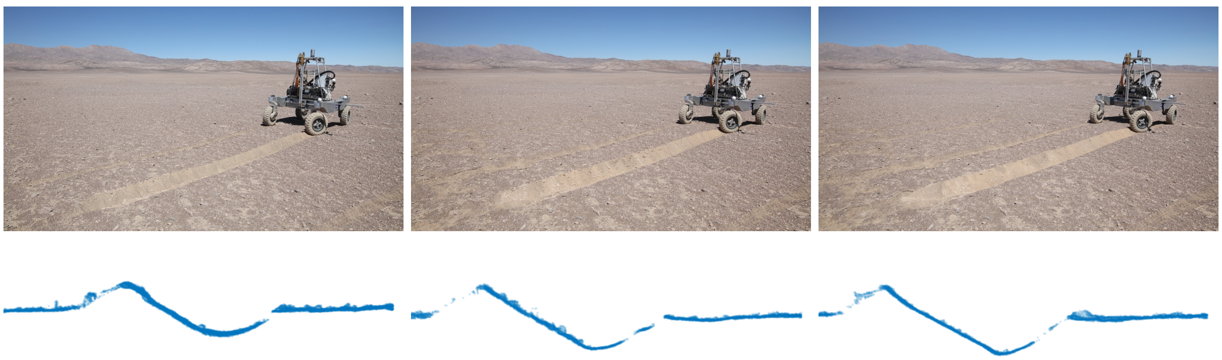

Figure 6.5: KREX-2 performing three trenching passes over the same terrain in the Atacama Desert (top row), and LIDAR scans of all three runs (bottom). All three passes were performed with the wheel at a slip angle of 30° with a rover body velocity of 20cm/s and wheel rim velocity of 50cm/s. The trench 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).

both shallower and narrower than the trench dug by a wheel at 90° in Section 6.1.1, suggesting that the driving primitive used to trench may have larger impact than repeated passes.

6.1.3 Hole digging

In addition, a hole digging action was demonstrated, as described in the second rows of Tables 3.1 and 3.2. Seven holes/soil piles were constructed by moving a single wheel with the rover stationary to estimate soil angle of repose as in [26], as seen in Figure 6.6. Maximum slope angle measurements on these soil piles were taken with a handheld inclinometer, with an average value of $\phi = 33.5°$. This value was used in the model evaluations performed in Section 6.1.1, and thus has clear use in aiding scientific work and mobility planning.

Figure 6.6: KREX-2 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).

6.1.4 Robot teaming

KREX-2 also demonstrated robot teaming scenarios with MiniRHex, an open-source miniature hexapod robot about the size of a brick [105]. In a mission context, a larger rover 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 Figure 6.7, 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 executing the 7th action in Tables 3.1 and 3.2, leveling a rock pile, by using a wheel to clear a path through the rocks. In the final frame, we see that MiniRHex is then 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 Figure 6.8, was formed by water runoff and has a steep vertical face, but is made of relatively friable soil that crumbles under significant 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 Figure 6.8. KREX-2 then used its wheel to construct a ramp in the step by crushing the terrain, seen in the

Figure 6.7: 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. MiniRHex tries to cross the rocks (upper left) but is unable to and flips over (upper right). Once KREX-2 clears rocks from the path (lower left), MiniRHex is able to climb through (lower right).

third frame, which then allowed MiniRHex to climb up to the higher terrain, as shown in the final image. This action was a combination of the “pile soil" and “level pile" actions in Tables 3.1 and 3.2, with soil moved from the top portion of the step to the bottom portion. For this test, KREX-2 had a tensegrital wheel design with a radius of 29 cm, which are slightly larger than its usual 23 cm rubber tires [106, 107].

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 rover to drive along performing monitoring or observation tasks, allowing the larger rover to traverse elsewhere, assured of the smaller robot's safe mobility.

Figure 6.8: Images from a video of KREX-2 helping MiniRHex traverse a natural step in terrain due to water flow in the Atacama Desert. MiniRHex tries to climb the step (upper left) but is unable to and flips over (upper right). Once KREX-2 crushes a ramp into the step (lower left), MiniRHex is able to climb the obstacle (lower right).

6.2 Impact of Actuator Failure on Mobility for VIPER

NASA's Volatiles Investigating Polar Exploration Rover will soon launch to the lunar surface to search for water [34]. VIPER has a four-wheeled actuated suspension, as seen in Figure 6.9. Each wheel has an in-wheel drive motor and individual actuators for steering and raising/lowering the suspension [108]. This mobility system has twelve actuators, enabling advanced locomotion techniques such as wheel-walking and swimming-like gaits [54, 109] 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 prototype Moon Gravity Representative Unit 3 (MGRU3), shown in Figure 6.9.

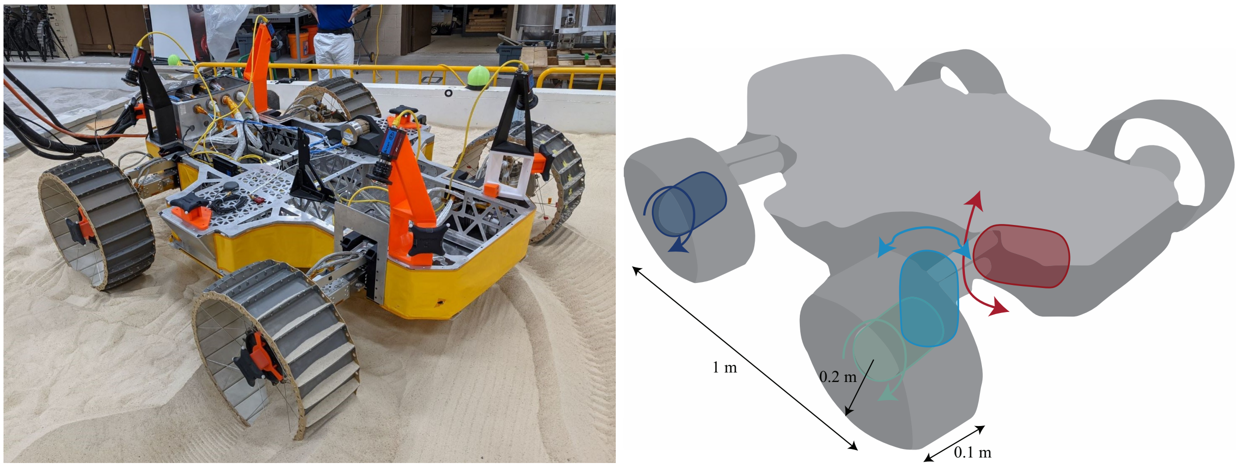

Figure 6.9: Moon Gravity Representative Unit 3, a mobility prototype for VIPER (left) and illustration of the actuators and joints referred to in this section (right). The actuators referred to in this section are the rear (dark blue) and front (green) drive actuators, steering (light blue), and suspension (red) actuators.

6.2.1 Experimental setup for VIPER mobility testing

All tests were conducted on MGRU3 in GRC-1 lunar simulant [110] at NASA Glenn Research Center's Simulated Lunar Operations (SLOPE) Lab. MGRU3 has flight software, motors, gearboxes, and joints, and has a lower mass than VIPER to simulate lunar equivalent weight. There are many potential failure modes for the actuation of an active suspension; 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 [9]. In addition, in the case of a stuck suspension or steering actuator the position at which an actuator fails can massively alter the mobility impact. A subset of potential failure modes were explored due to limited testing time, with a mixture of more operationally likely failure states and an attempt at representative coverage. 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 Figure 6.10. For all tests, soil preparation

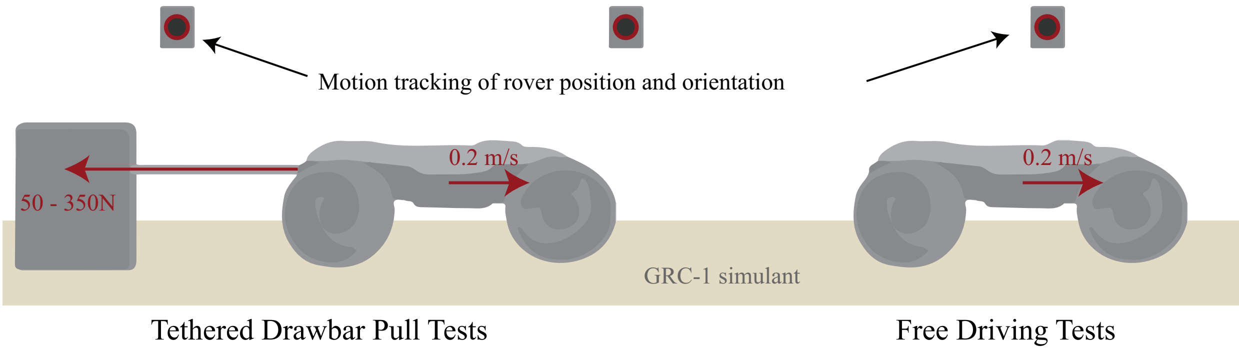

Figure 6.10: Experimental setup for testing the impact of joint actuator failure on VIPER. Drawbar pull tests were performed in a tethered configuration (left), and qualitative drive tests were performed untethered (right), with both tests in GRC-1 simulant [110]

consisted of manual loosening and raking of simulant between runs. 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.

6.2.2 Qualitative drive testing on VIPER

Stuck drive motor: MGRU3 can be seen after attempting to drive with a nonrotating (stuck) rear left drive actuator in Figure 6.9 (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 Figure 6.11.

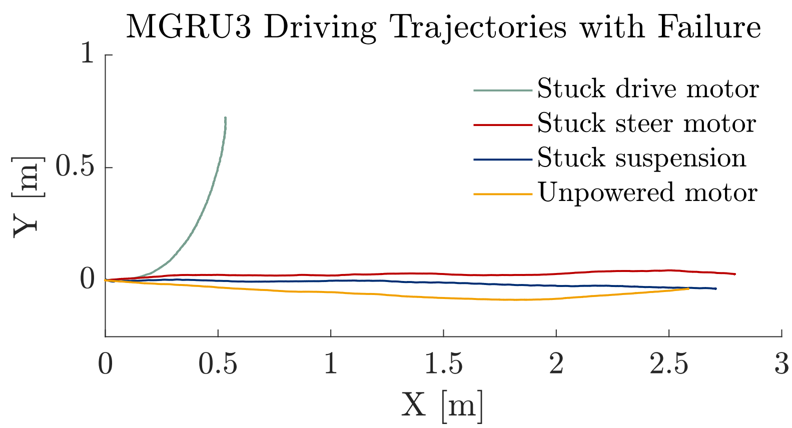

Figure 6.11: Trajectories followed by MGRU3 driving with various failed actuators. The rover barely moved with a nonrotating drive motor (green) but was able to make forward progress with other tested failure modes.

Unpowered drive motor: MGRU3 was able to drive normally with a single unpowered motor, as seen in Figure 6.11 (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 6.2.3.

Stuck steering actuator: Fixing MGRU3's rear right steering actuator so that the wheel is pointed 30° outward, as shown in Figure 6.12, resulted in very slight drift in heading but did not impede the rover's ability to make forward progress, as seen in Figure 6.11 (red). A steering actuator stuck at a moderate angle would have a larger impact on steering performance than normal driving performance, as was observed on Opportunity [9].

Stuck suspension: The rear right wheel was lifted at a 45° angle to simulate the failure of a suspension actuator. Due to the design of VIPER's suspension, 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 Figure 6.13. 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,

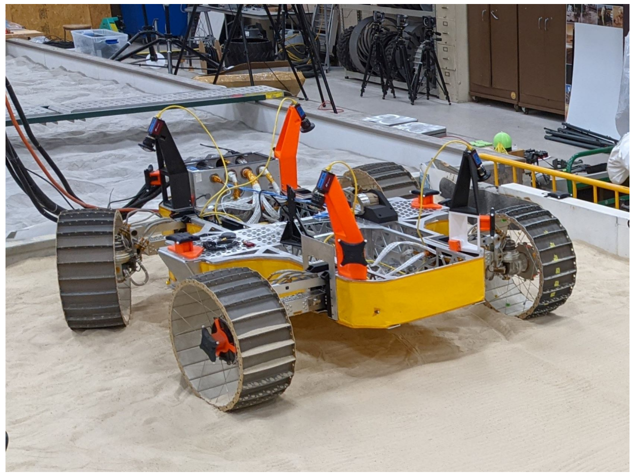

Figure 6.12: MGRU3 driving with its rear right steering actuator stuck at 30° on the rear right wheel.

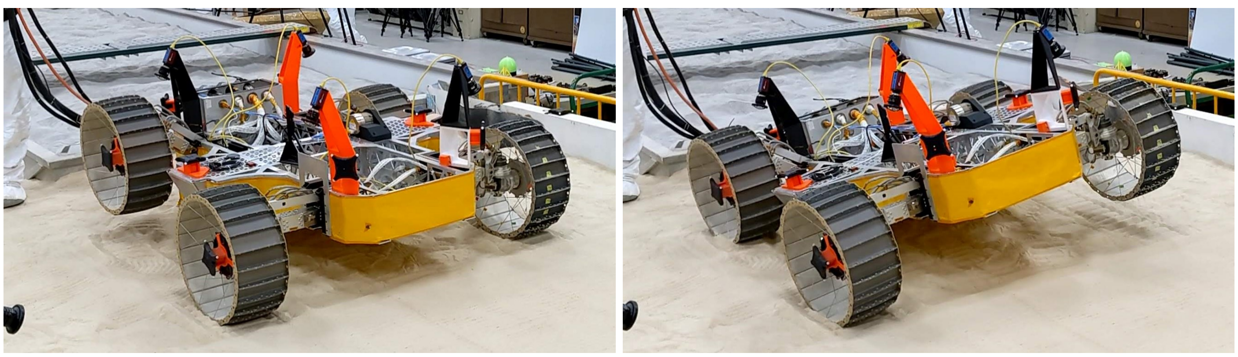

Figure 6.13: MGRU3 driving with a suspension joint stuck in a raised position pictured before moving (left) and immediately after driving started (right), with the rover tipped backwards onto the affected actuator.

the rover was able to drive straight, as shown in Figure 6.11 (blue).

6.2.3 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 F_{DBP} and rover weight W_{rover}, the equivalent slope angle α is given by

$$\alpha = \tan^{-1}\left(\frac{F_{DBP}}{W_{rover}}\right) \tag{6.1}$$

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. 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 operational limit on slope climbing is to maintain wheel slip below 40% [111], 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 Figure 6.14. 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.

Stuck or unpowered drive motor: Loss of a rear or front drive motor was associated with a 200% increase in slip on slopes of 15°, and a 25x increase on shallow inclines of 3 degrees. While a locked front wheel showed similar traction reductions to an unpowered wheel in these tests, the rover 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

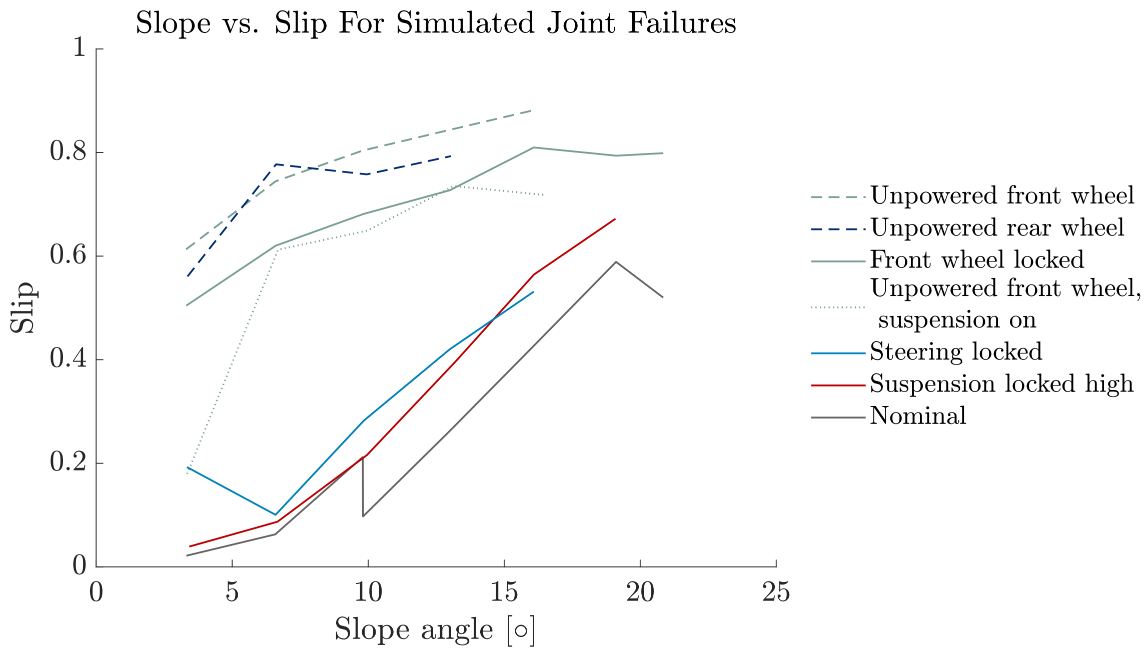

Figure 6.14: Plot of slip vs. slope angle for VIPER rover driving with failed actuators.

impact on steeper slopes.

Stuck steering actuator: Locking a steering joint increased slip 25x on shallow inclines but only 0-50% increase on other slopes, but the rover was unable to maintain a straight heading and had to be manually readjusted during trials to keep it in the testbed.

Stuck suspension: Locking a suspension joint into a raised position only increased slip by 0-50%, 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.

While both the drawbar pull tests presented here and the free driving tests in Section 6.2.2 were performed on the same hardware in the same soil, the impact of joint failure on MGRU3's mobility varied between the two tests for certain affected actuators. In particular, the unpowered wheel had little effect on the free driving tests but a dramatic slip increase in the tethered tests. The rover was nearly unable to progress with a stuck drive actuator in the free driving tests, but had lower slip than the unpowered wheel tests when tethered. Additionally, the rover's heading drift due to a stuck steering motor was much larger in tethered tests than free driving tests. The relative difference in performance between the tethered and untethered tests for the drive motor failure modes suggests that both free driving tests on flat ground and drawbar pull tests to simulate slopes are needed to fully assess rover mobility. Neither testing mode showed lower or higher mobility impact for all tests, so for conservative assessment of rover mobility both free driving and drawbar pull tests should be performed.

Depending on the affected portion of the mobility system, actuator loss could be mission ending for VIPER. Without compensation, the average slip on MGRU3's wheels was above VIPER's operational limit of 40% for both unpowered and stuck drive motors for all drawbar pull tests. Even a stuck steer actuator, which still allowed the rover to make forward progress, would hurt VIPER's mobility by making it yaw unexpectedly. Should VIPER experience mobility actuator loss of any kind, operational constraints in the form of slope limits would need to be reasessed, and new driving strategies would likely be required. In later sections, we develop the techniques needed to perform these tasks and mitigate mobility loss.

6.3 Field Demonstrations of Rover Wheel Failure Compensation

As shown in Section 6.2, the loss of steering, driving, and suspension actuators limits rover mobility, and in extreme cases can fully disable the rover. In this section, we present two examples of a rover using a hand-tuned strategy to compensate for failure of a mobility actuator in order to demonstrate the feasibility of these techniques on rovers on the scale of KREX-2, VIPER, and Curiosity. In Chapter 7, we use the models created in Chapters 4 and 5 to automatically generate mobility degradation compensation strategies for a smaller rover.

6.3.1 Stuck steer motor

The trenching tests described in Section 6.1 also serve as demonstrations of driving straight with a steering actuator that has gotten stuck at 30°, 60°, and 90°. As seen in Figure 6.15, the rover was able to drive straight with the rear right wheel stuck at all three angles using hand-tuned driving strategies. When driving with its rear right wheel pointed straight forwards, the rover's travel velocity was 30 cm/s, and the hand-tuned driving primitives at other steering angles attempted to match this. The rover was able to maintain speed with its steer motor stuck at 30° and 60° with travel speeds of 28 cm/s and 23 cm/s respectively, as seen in Table 6.4, but moved at a much slower 16 cm/s when the wheel was stuck at 90°.

Figure 6.15: 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).

Steer Angle	Rover Speed [cm/s]	Avg Slip
$\beta = 30°$	28	-.05
$\beta = 60°$	23	.08
$\beta = 90°$	16	.39

Table 6.4: Travel speeds and slip for KREX-2 operating with a stuck steer motor. Note that average slip values include the affected wheel.

The rover was able to maintain close to zero slip at the two smaller steer angles, but encountered 40% slip at the 90° steering angle. As these driving primitives were hand tuned we can't guarantee that these driving strategies are optimal, but the impact of a failed steering motor is clearly dependent on the angle at which it gets stuck.

6.3.2 Stuck drive motor

A nonrotating wheel is one of the worst-case mobility system failure in terms of mobility impact, as seen in Section 6.2, so we sought to show that it is possible for a rover on the same scale as

VIPER to recover. In this section, we show KREX-2 recovering forward driving mobility with a stuck drive actuator by compensating with its other three wheels. For this demonstration, KREX-2 was outfitted with cylindrical wheels composed of PVC tubing fit onto the outside of rubber tires, in order to have the wheel-soil interface be more similar to rigid wheels on flight rovers like VIPER and Curiosity. The dimensions of the PVC wheels and rover weight in this configuration are given in Table 6.5. The wheel surface was covered in sandpaper with an adhesive backing to give KREX-2 better traction. The demonstration was conducted in NASA Ames Research Center's Roverscape, in extremely compactible, poorly sorted crushed granite. Between trial runs, the full driving track was loosened with a powered rototiller to a depth of several inches, deeper than the observed wheel sinkage of KREX-2. The soil was then raked to a visually smooth and level state. During each run the rover's position and body velocity were measured with a Leica Total Station as in Section 6.1.1.

r	Wheel radius [cm]	26
b	Wheel width [cm]	20
W	Rover weight [kg]	280

Table 6.5: KREX-2 parameters for wheel failure compensation demos.

When driven with a rear wheel held in a fixed position, as in the case of a rock jam or stuck brake mechanism, KREX-2 yaws about the stuck wheel and veers off track, as seen in the left image in Figure 6.16 and in the plot of the rover's trajectory in Figure 6.17. When a hand-tuned 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 Figures 6.16 and 6.17.

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 40% slip threshold set for VIPER's operations [111]. 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.

Figure 6.16: KREX-2 rover driving with a stuck drive motor without compensation (left) and with open-loop compensation (right).

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 Figures 6.16 and 6.17. 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 on the scale of VIPER and Curiosity. 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 chapter, we address this by presenting methods to automatically generate driving strategies to overcome actuator failure.

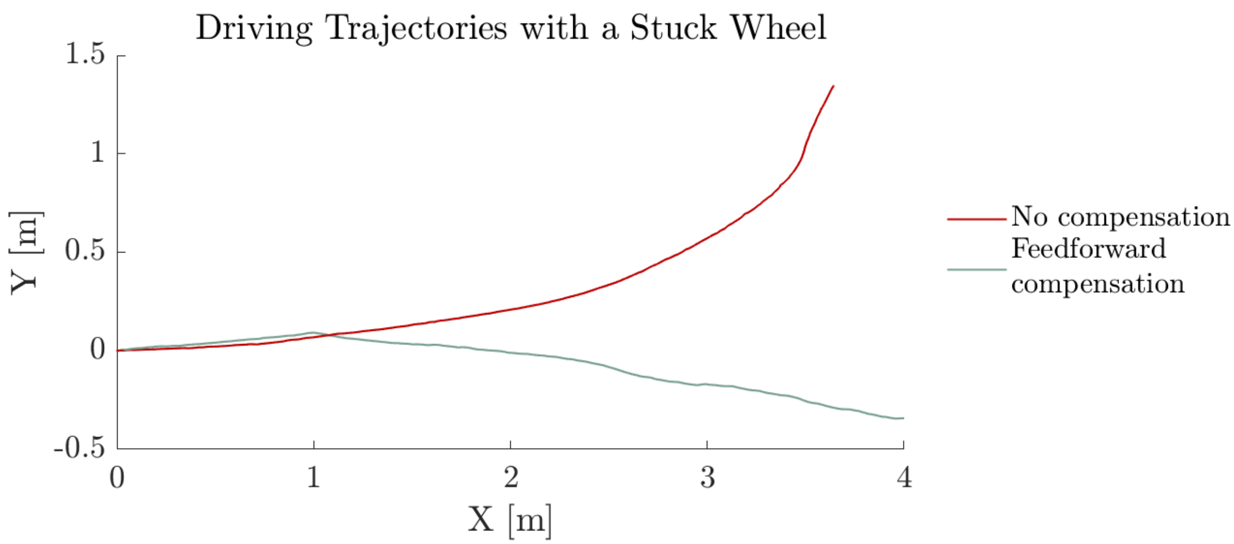

Figure 6.17: KREX-2 drive trajectories with the rear right wheel stuck (nonrotating). All trajectories are open-loop, with no steering controls used.

6.4 Conclusions

In this chapter, 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 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 KREX-2 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.

We have shown here 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, which we detail in Chapter 7.

Chapter 7

Rover Control for Off-Nominal Driving

In Chapter 6 we demonstrated hand-tuned driving primitives for trench digging and recovery from mobility failures. These demonstrations showed it is feasible to drive while trenching or with degraded mobility, but were not generalizable to different contexts and did not allow for closed-loop control of steering. In this chapter, we take the models developed in Chapter 4 and Chapter 5 and embed them in an optimization framework for generating steady-state driving strategies. By carefully selecting constraints for the optimization problem, we can generate driving strategies for nominal driving, driving while trenching, and driving with varied types of degraded mobility.

During nominal operation, wheeled vehicles will drive with all wheels pointed straight forwards and driving at the same velocity. Steered driving is generally accomplished through a set method that is dependent on the kinematics of the vehicle, such as skid steering, Ackermann steering, point turns, or swerve driving. However, when a rover is performing manipulation with a wheel or is experiencing failure of one or more mobility actuators, it will not drive straight when driven normally and may not have the necessary steering capabilities. As the terramechanics model developed in Chapter 5 can be used to determine the forces and wheel-soil contact geometry of a single wheel experiencing arbitrary slip and steering angles, it can in turn be used to predict the motion of a full rover, and thus develop feasible driving strategies. In this chapter, we first propose an optimization-based approach for finding feasible, energy-efficient driving strategies for planetary rovers in off-nominal conditions as well as a steering control approach via the same method. We then show how to use the optimization to achieve soil manipulation objectives and compensate for specific types of mobility failure. Finally, we demonstrate driving primitives generated by the optimization on a miniature rover.

7.1 Nomenclature

See Table 7.1 for a full list of variables used in this section along with their meanings and units.

Symbol	Value	Units
	Rover and wheel state	
β	Slip angle	$\circ$
ω	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	
$a_{c\beta}, a_{s\beta}$	β-dependent coefficients within a_0 and a_1	$1/\circ$
θ_f, θ_r	Front, rear angles where wheel enters, exits soil	$\circ$
θ_m	Angle of max normal stress under wheel	$\circ$
ζ	Grouser transport volume fraction	

Table 7.1: Table of variables used in rover driving optimization.

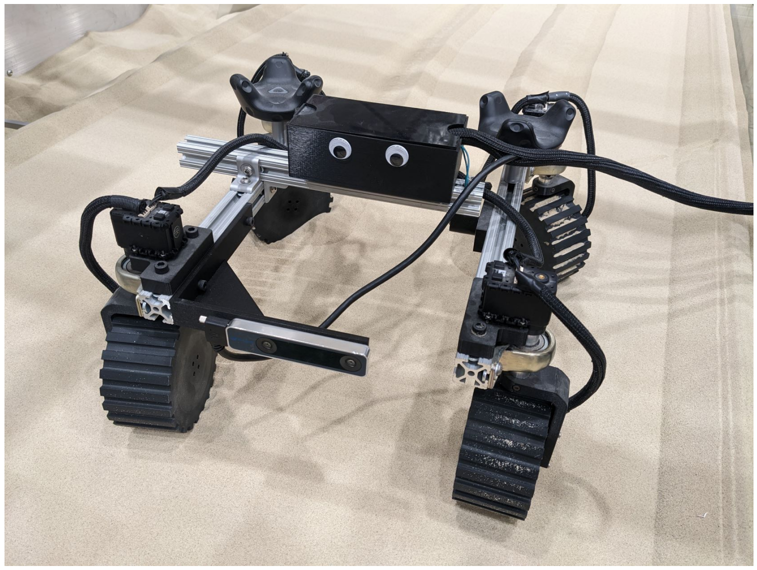

Figure 7.1: Mini rover used to test off-nominal driving strategies. The rover has two rockers with a passive pivot, with independent steer and drive motors for each wheel. The rover's center portion bearing the electronics is fixed to the right rocker, and it is tethered for power and communications. A RealSense T265 tracking camera is used to obtain the rover's position and orientation.

7.2 Optimization for generating off-nominal 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 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 Figure 7.1, 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.

7.2.1 Decision variables

When a rover is driving, the only true free variables are the speed and angle of each wheel; the vehicle will eventually reach a steady state with constant wheel sinkages, travel velocity, and rover angular velocity. As the terramechanics model developed in Chapter 5 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 steady state. To simplify the problem, we set the rover's 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 ω, slip angle β, 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 7.2.7. Thus rather than including a numerical sinkage finding, as we did in Chapter 5, 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

$$x = [\omega_{FR}\ \omega_{FL}\ \omega_{RR}\ \omega_{RL}\ \beta_{FR}\ \beta_{FL}\ \beta_{RR}\ \beta_{RL}\ W_{FR}\ W_{FL}\ W_{RR}\ W_{RL}\ h_{FR}\ h_{FL}\ h_{RR}\ h_{RL}] \tag{7.1}$$

where the subscripts FR, FL, RR, and RL respectively 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

$$0 \leq \omega \leq \omega_{max_{motor}} \tag{7.2}$$

$$-\frac{\pi}{2} \leq \beta \leq \frac{\pi}{2} \tag{7.3}$$

$$0 \leq W \leq \frac{W_{rover}}{2} \tag{7.4}$$

$$h_g \leq h \leq \frac{3}{2}r + h_g \tag{7.5}$$

where $\omega_{max_{motor}}$ 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 as illustrated in Figure 5.1, from the rover velocity, $v_{x_{rover}}$ and $v_{y_{rover}}$, so that for each wheel $i \in \{FR, FL, RR, RL\}$ we have

$$v_{x_i} = v_{x_{rover}} \cos\beta_i + v_{y_{rover}} \sin\beta_i \tag{7.6}$$

$$v_{y_i} = v_{y_{rover}} \cos\beta_i - v_{x_{rover}} \sin\beta_i \tag{7.7}$$

In Section 5.4.3 we showed that for the soil and wheels used here, we can define a_0 and a_1, the parameters for determining the shape of the normal stress, as a linear function of the slip angle β. Using the fits from Section 5.4.3, we calculate a_0 and a_1 for each wheel i at an arbitrary slip angle β by

$$a_{0_i} = a_c + a_{c\beta}\beta_i \tag{7.8}$$

$$a_{1_i} = a_s + a_{s\beta}\beta_i \tag{7.9}$$

With a_0 and a_1 set for each wheel, we find the maximum stress angle θ_m for each wheel that is in slip with Equation 5.3. For wheels in skid, we find the maximum stress angle, as in Equation 5.3, by finding the maximum stress angle for a slip ratio of zero at our given slip angle β. 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 calculate the entrance angle $\theta_f = \cos^{-1}(1 - h/r_s)$ based on the shearing radius of the wheel r_s and then use Equation 5.3 to find θ_m for skid.

7.2.2 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

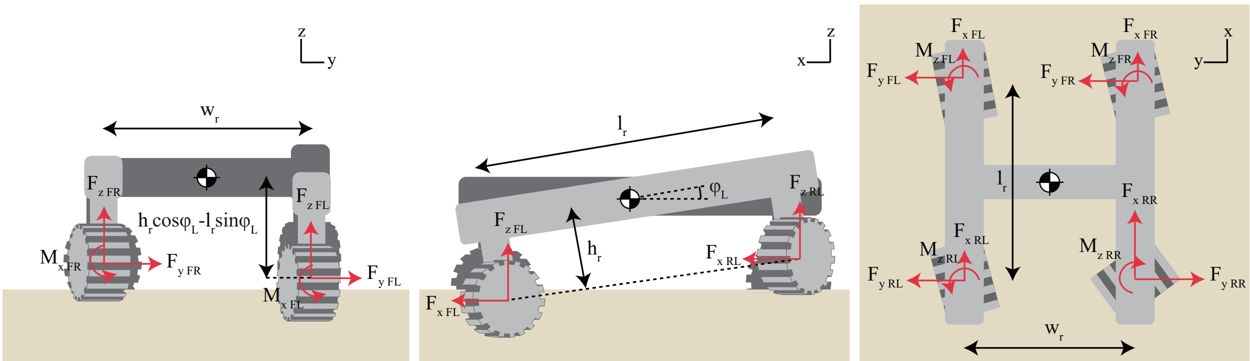

Figure 7.2: Free body diagram of the mini rover while driving. All labeled forces and moments are in the rover frame. Note that the moment balances on the left and right rockers are performed separately due to the passive pivot between them.

F_z) and moments (M_x, M_y, and M_z) on each wheel as in Section 5.2.5. 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 Figure 7.2. The angles of the right and left rockers, ϕ_R and ϕ_L, relative to horizontal, are determined by the sinkages of each wheel and the rover rocker length l_r, and given by

$$\phi_R = \sin^{-1}\left(\frac{h_{RR} - h_{FR}}{l_r}\right) \tag{7.10}$$

$$\phi_L = \sin^{-1}\left(\frac{h_{RL} - h_{FL}}{l_r}\right) \tag{7.11}$$

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

$$\Sigma F_x = F_{xFR} + F_{xFL} + F_{xRR} + F_{xRL} = 0 \tag{7.12}$$

$$\Sigma F_y = F_{yFR} + F_{yFL} + F_{yRR} + F_{yRL} = 0 \tag{7.13}$$

$$\Sigma F_z = F_{zFR} + F_{zFL} + F_{zRR} + F_{zRL} - W_{rover} = 0 \tag{7.14}$$

$$\begin{aligned}\Sigma M_x = {} & M_{xFR} + M_{xFL} + M_{xRR} + M_{xRL} + (F_{yFR} + F_{yRR})\cos\phi_R h_r + (F_{yFR} - F_{yRR})\sin\phi_R l_r/2 \\ & + (F_{yFL} + F_{yRL})\cos\phi_L h_r + (F_{yFL} - F_{yRL})\sin\phi_L l_r/2\end{aligned}$$

$$+(-F_{z_{FR}}+F_{z_{FL}}-F_{z_{RR}}+F_{z_{RL}})w_r/2=0 \tag{7.15}$$

$$\Sigma M_{y_R} = (-F_{x_{FR}}-F_{x_{RR}})\cos\phi_R h_r + (-F_{x_{FR}}+F_{x_{RR}})\sin\phi_R l_r/2 + (-F_{z_{FR}}+F_{z_{RR}})\cos\phi_R l_r/2$$
$$+(F_{z_{FR}}+F_{z_{RR}})\sin\phi_R h_r = 0 \tag{7.16}$$

$$\Sigma M_{y_L} = (-F_{x_{FL}}-F_{x_{RL}})\cos\phi_L h_r + (-F_{x_{FL}}+F_{x_{RL}})\sin\phi_L l_r/2 + (-F_{z_{FL}}+F_{z_{RL}})\cos\phi_L l_r/2$$
$$+(F_{z_{FL}}+F_{z_{RL}})\sin\phi_L h_r = 0 \tag{7.17}$$

$$\Sigma M_z = M_{z_{FR}}+M_{z_{FL}}+M_{z_{RR}}+M_{z_{RL}}+(F_{x_{FR}}-F_{x_{FL}}+F_{x_{RR}}-F_{x_{RL}})w_r/2$$
$$+(F_{y_{FR}}-F_{y_{RR}})l_r/2\cos\phi_R+(F_{y_{FL}}-F_{y_{RL}})l_r/2\cos\phi_L = 0 \tag{7.18}$$

where the subscripts FR, FL, RR, and RL respectively indicate forces or moments on the front right, front left, rear right, and rear left wheels in the rover body frame. 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 ΣM_{y_R} for the right side rocker and ΣM_{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 Equations 7.12–7.18 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)

$$F_{z_{FR}} - W_{FR} = 0 \tag{7.19}$$

$$F_{z_{FL}} - W_{FL} = 0 \tag{7.20}$$

$$F_{z_{RR}} - W_{RR} = 0 \tag{7.21}$$

$$F_{z_{RL}} - W_{RL} = 0 \tag{7.22}$$

If we take these equations and substitute them back into Equation 7.14, we can restate it as

$$\Sigma F_z = W_{FR} + W_{FL} + W_{RR} + W_{RL} - W_{rover} = 0 \tag{7.23}$$

This new constraint is a linear function of x, which is easier for the solver to handle than Equation 7.14, 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.

7.2.3 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

$$c(x) = \bar{M}_{y_{FR}}^2 + \bar{M}_{y_{FL}}^2 + \bar{M}_{yRR}^2 + \bar{M}_{yRL}^2 \tag{7.24}$$

where $\bar{M}$ is the wheel torque, which is the moment about the wheel's y axis. In Section 7.2.7, we discuss rescaling the objective function by a scalar to improve convergence.

Putting this all together into an optimization problem,

$$\begin{aligned} \text{min} \quad & \text{Objective function, } c(x),\ (7.24) \\ \text{w.r.t.} \quad & \text{Decision variables, } x,\ (7.1) \\ \text{s.t.} \quad & \text{Bounds, } x_{LB} \leq x \leq x_{UB}, (7.2) - (7.5) \\ & \text{Force Balance, } \Sigma F_x = 0,\ \Sigma F_y = 0, \Sigma F_z = 0,\ (7.12), (7.13), (7.23) \end{aligned} \tag{7.25}$$

Moment balance, $\Sigma M_x = 0,\ \Sigma M_{y_R} = 0,\ \Sigma M_{y_L} = 0,\ \Sigma M_z = 0,\ (7.15) - (7.18)$

Vertical load balance, $F_z - W_z = 0,\ (7.19) - (7.22)$

The resultant set of wheel angles and speeds is an open-loop driving primitive that will make the rover drive straight.

7.2.4 Constraints for steering while driving

While rovers such as Curiosity and Perseverance use feedforward driving strategies and only close the loop on heading with visual odometry every meter [32], 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 Equation 7.18 with

$$\begin{aligned}\Sigma M_z =& M_{z_{FR}} + M_{z_{FL}} + M_{z_{RR}} + M_{z_{RL}} + (F_{x_{FR}} - F_{x_{FL}} + F_{x_{RR}} - F_{x_{RL}})w_r/2 \\ &+ (F_{y_{FR}} + F_{y_{FFL}} - F_{y_{RR}} - F_{y_{RL}})l_r/2 - M_{z_{offset}} = 0\end{aligned} \tag{7.26}$$

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 Equation 7.13 with

$$\Sigma F_y = F_{y_{FR}} + F_{y_{FL}} + F_{y_{RR}} + F_{y_{RL}} - F_{y_{offset}} = 0 \tag{7.27}$$

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.

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_β	β-dependent coefficient within a_0	1/129.1°
a_s	Constant coefficient within a_1	0.69
$a_{s\beta}$	β-dependent coefficient within a_1	-1/133.5°
ζ	Grouser transport volume fraction	.2
Offset forces		
$M_{z_{offset}}$	Yaw moment offset for steering control	±0.2 Nm
$F_{y_{offset}}$	y force offset for steering control	±1.0 N

Table 7.2: Rover and soil parameters for nominal driving optimization.

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 7.2 and with the soil parameters listed in Table 5.3 except where otherwise noted.

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 7.3.

7.2.5 Control for wheel-based trenching

With the trench shape model developed in Chapter 4, we can easily factor trench geometry into both our constraints and our 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 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}, as defined in Equation 4.4, is zero for the wheel being used to trench. Here we select the rear right wheel to dig, and so we add the constraint

$$h_{R2,RR} = 0 \tag{7.28}$$

We can also add inequality constraints, such as requiring that the trench depth d, as defined in Equations 4.5–4.12, 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

$$b/2 - d \leq 0 \tag{7.29}$$

Here, we use the same objective function as before, Equation 7.24, seeking to minimize the energy used in digging while we drive. The optimization problem, Equation 7.25, was solved with these additional constraints and the values outlined in Table 7.2, with the exception that ζ was set to 1. Both a feedforward driving primitive as well as yaw and y position steering controllers were computed and tested along with the steering controls generated for nominal driving. The steering controls generated relative to the feedforward trenching-while-driving primitive are denoted as the "optimal" steering controls in Section 7.3.1, as they were designed specifically to minimize actuator effort about that set point.

7.2.6 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 7.3.

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 [9]. 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 [9]. 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 use the same constraints as in our nominal straight-driving optimization and add an additional inequality constraint such 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

$$\omega_{RR} r_s - v_{lim} \leq 0 \tag{7.30}$$

We re-solve the optimization problem with this added constraint, using the values specified in Table 7.2, 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 θ_{fail} needs an additional equality constraint of

$$\beta_{RR} - \theta_{fail} = 0 \tag{7.31}$$

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. The feedforward and feedback controls generated are tested in Section 7.3.

7.2.7 Nuances of optimization for complex problems

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 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 Equation 7.25, and discuss how we solved these problems.The solver used for this problem was MATLAB's *fmincon* with the *interior-point* algorithm in Feasibility Mode.

Numerical integration: As the full terramechanics models presented in Chapter 5 are not closed-form, Equations 5.26 – 5.28 and 5.30 – 5.32 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 built on terramechanics nonsmooth. 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.

Mode transitions: The terramechanics model in Chapter 5 covers both slip and skid but uses different soil shearing behavior in those two regimes. Additionally, the soil flow model in Chapter 4 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°, as explained in Section 4.2). 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 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 $\pm 60°$ 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° 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.

Internal iterative loops: In Chapter 5, 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 ε, as in Equation 5.34. When included in an optimization, the numerical accuracy ε set for the binary search creates steps in the sinkage, which makes the constraint and gradient functions nonsmooth, as shown in Figure 7.3. Additionally, having an internal loop slows down the solver. To avoid this, 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, Equations 7.19–7.22.

However, there is a single internal iterative loop that cannot be eliminated by this method; the definition of θ_{m0} (and θ_{r0}, when the trench model is not used to determine soil exit angle) is dependent on the sinkage of the wheel at zero slip, $\theta_m\big|_{s=0}$, as given in Equation 5.3. As the sinkage at zero slip is dependent on the applied wheel load and the wheel's slip angle, θ_{m0} must be evaluated for each iteration of the optimization. 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.

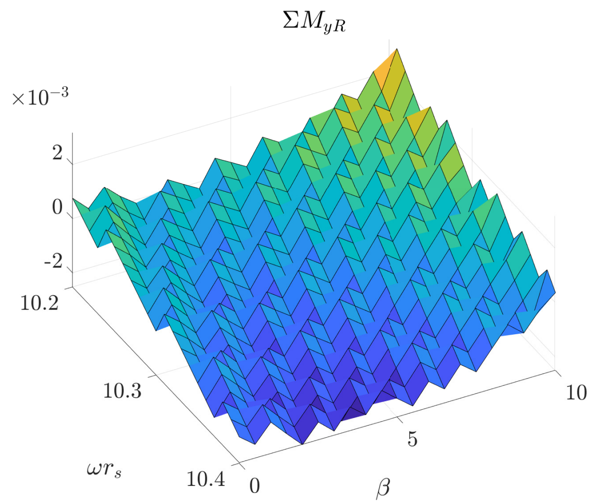

Figure 7.3: Discretization visible in one of the constraint functions due to the tolerance of binary search for finding wheel sinkage. Values shown here are for constraint (7.16) looking at a 2d slice of possible values of x, as the full dimension of decision variables is too high to visualize.

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. 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. These values 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 they 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.

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.

7.3 Experimental Testing of Optimization-Based Control

Testing of the optimization-based control strategies was performed with the mini rover in the same soft soil testbed used in Section 4.3. Both the soil and grousered wheels used in this experiment are the same as those used in the terramechanics model validation described in Chapter 5. Key geometry parameters for the wheel can be found in Table 5.2, and key soil parameters can be found in Table 5.3. The mini rover pictured in Figure 7.1 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. 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. The

Figure 7.4: Full rover control test setup in CMU's soft soil testbed. Realtime state estimation is performed by a RealSense® T265 tracking camera with IMU mounted to the front of the rover.

experimental setup can be seen in Figure 7.4.

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.

7.3.1 Trenching control experiments

Figure 7.5: 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.

The trenching strategy generated in Section 7.2.5 was run on the mini rover using three different control methods: feedforward only, nominal feedback steering controls, and optimized steering controls. The tracks left by each trial can be seen in Figure 7.5, 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 Figure 7.6. The rover drove very straight while trenching in all three control schemes, with less than 3° 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 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 without visual odometry such as Curiosity and Perseverance.

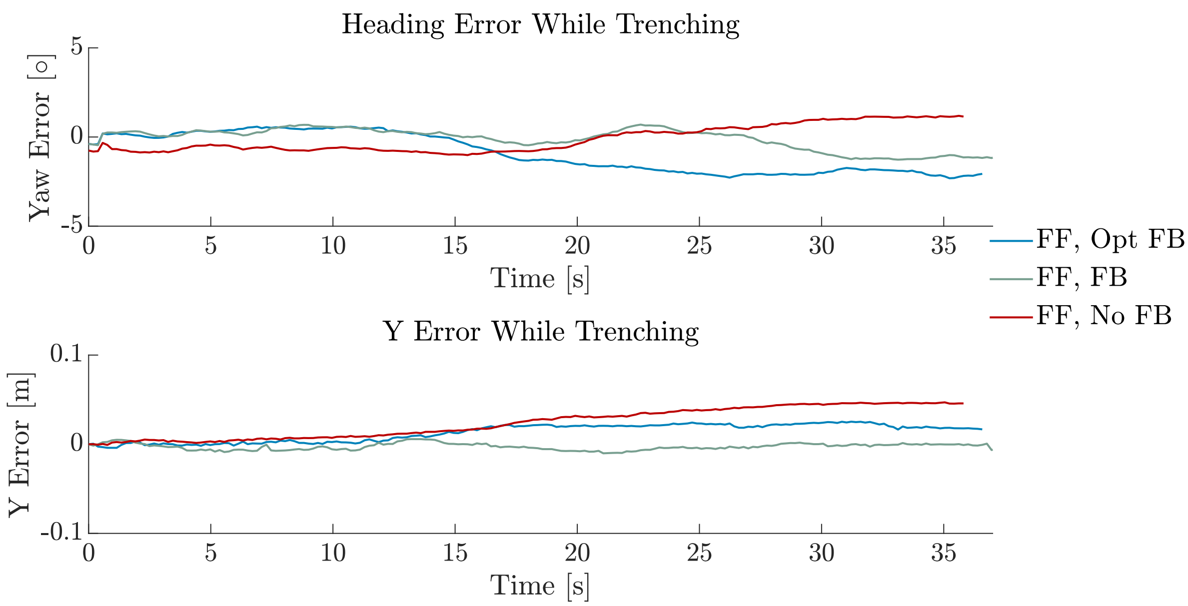

Figure 7.6: Comparison of yaw and y position errors for open and closed loop trajectory following while using the rear right wheel to dig a clean-sided trench. FF: uses a feedforward driving strategy generated by optimization. FB: feedback control enabled, with default steering unless noted and "optimal" steering relative to feedforward strategy for Opt FB. All three methods are able to maintain heading while digging a trench.

A detailed image of the trench dug by the wheel (blue) along with its predicted shape (red) can be seen in Figure 7.7. 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 2.6 cm deep trench and and observed depth of 1.5 cm. This discrepancy is likely due to an under prediction of the wheel's sinkage, possibly due to either a difference in predicted applied load from the actual load, or from the terramechanics model in Chapter 5 being tuned on a slightly higher rover weight than that of the mini rover.

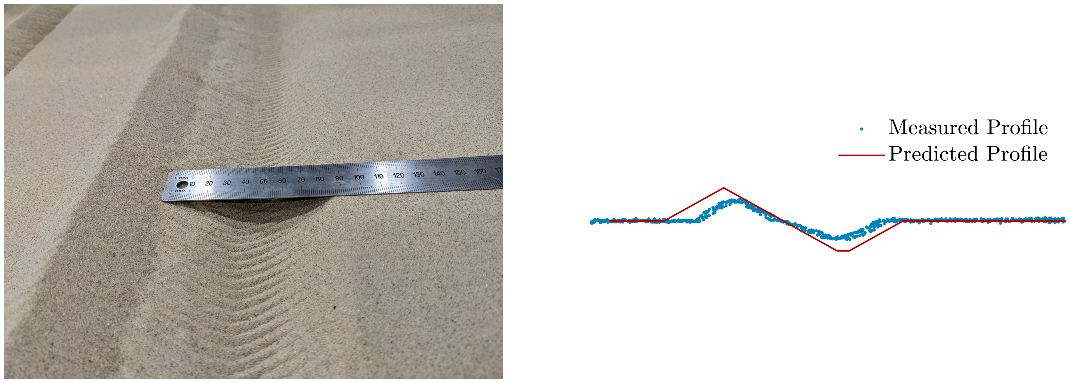

Figure 7.7: Photograph of trench dug by the mini rover while driving with closed-loop steering controls (left), alongside LIDAR scan of the same trench (blue, right) overlaid with its predicted shape (red, right). The rover successfully produced a clean-sided trench, but with a lower depth than expected.

7.3.2 Speed-limited wheel experiments

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 30 cm/s. The tracks left by the rover with no compensation, feedforward compensation only, and feedback compensation optimized for use with the feedforward compensation can be seen in Figure 7.8, while the yaw and y error are plotted for these and additional feedback strategies in Figure 7.9.

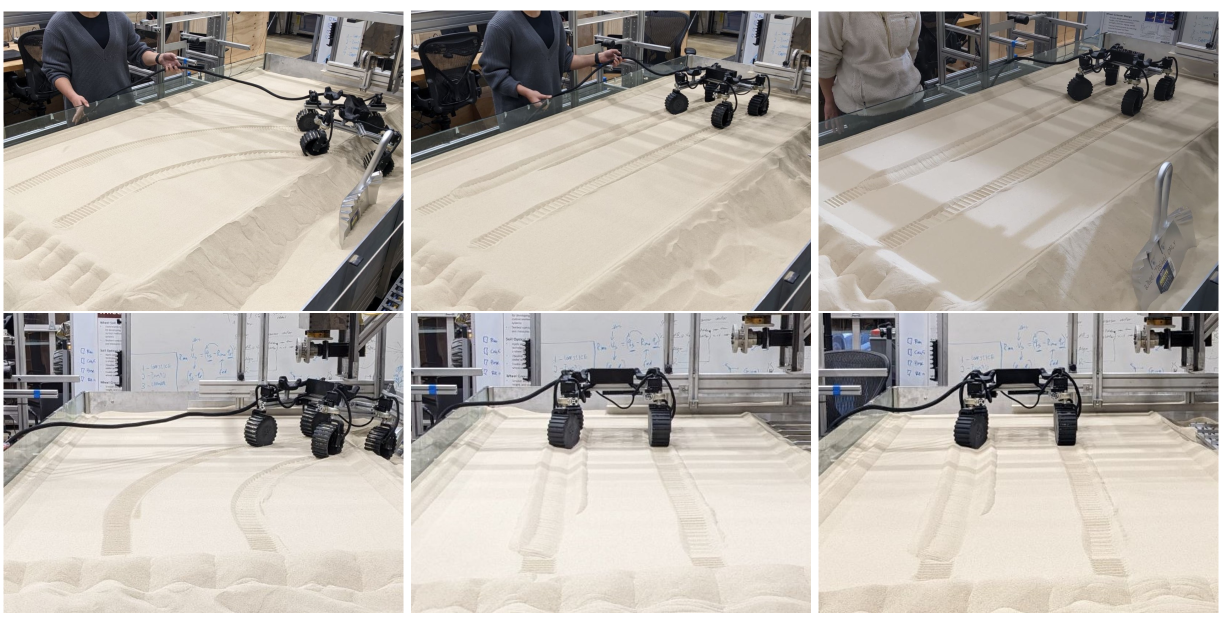

Figure 7.8: Mini rover driving with a speed-limited wheel. From left to right: driving without compensation, feedforward with no feedback control, feedforward with optimal feedback control.

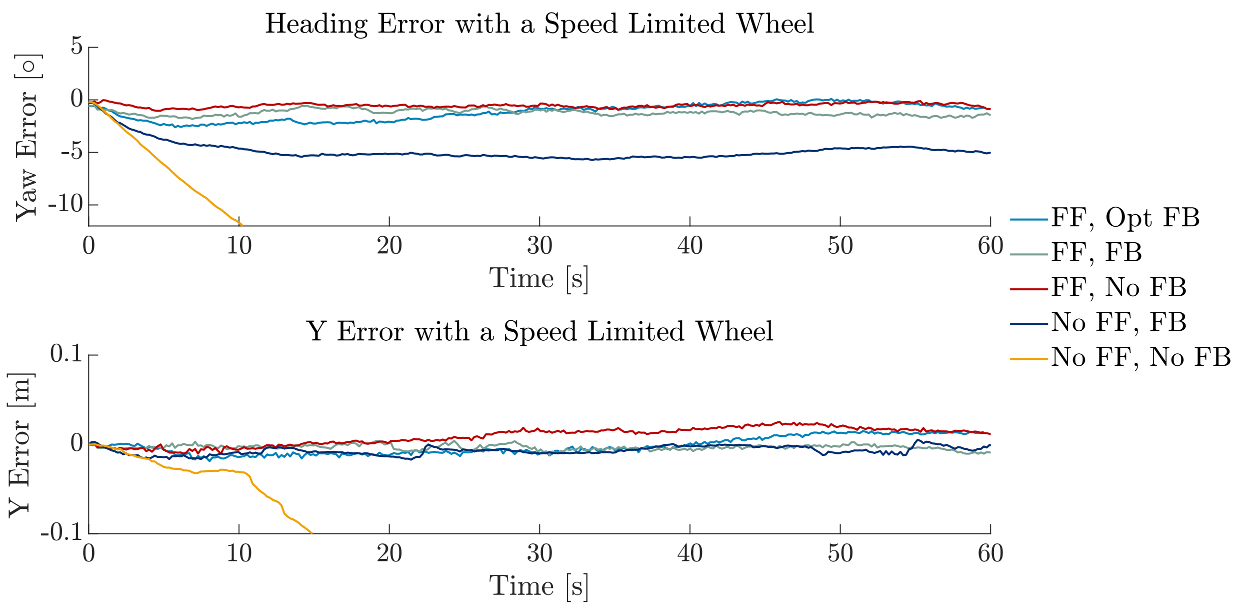

Figure 7.9: Comparison of yaw and y position errors for open and closed loop trajectory following with the rear right wheel limited to 15 cm/s. FF: uses a feedforward driving strategy generated by optimization. FB: feedback control enabled, with default steering unless noted and "optimal" steering relative 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.

With no compensation, the rover immediately began to pull right, veering off the prepared soil track within 1 m of driving. All compensation strategies greatly improved the rover's ability to maintain heading (yaw) without side slip (y error), as seen in Figure 7.9. Without compensation (yellow), the rover's yaw and y position continuously increased. With all combinations of feedforward and feedback compensation (blue, green, red, and dark blue), the rover was able to maintain heading within 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, Figure 7.9) 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).

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 without 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.

7.3.3 Stuck steering motor experiments

As seen in the tracks in Figure 7.10 and the error magnitudes in Figure 7.11, 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.

Figure 7.10: 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.

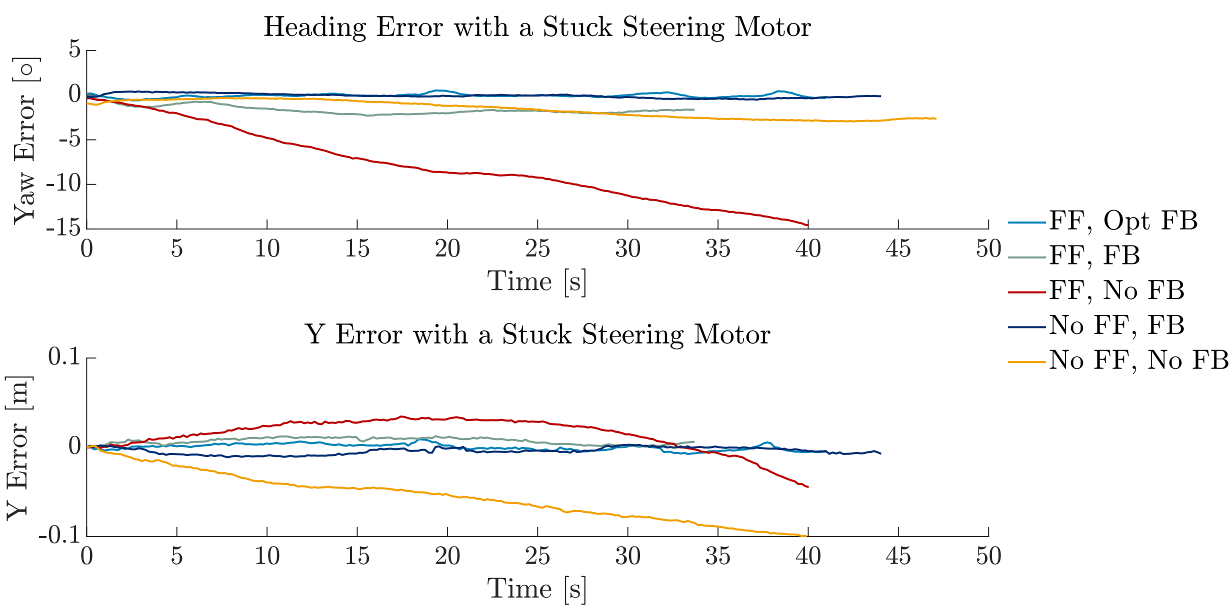

Figure 7.11: Comparison of yaw and y position errors for open and closed loop trajectory following with the rear right steering actuator stuck at -50°. FF: uses a feedforward driving strategy generated by optimization. FB: feedback control enabled, with default steering unless noted and "optimal" steering relative to feedforward strategy for Opt FB.

7.3.4 Mobility system failure insights from optimization

In Section 6.3 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 6.2.3 experimentally simulated slope by applying a load resisting the rover's motion. To simulate various slope angles, the optimization problem described in Section 7.2 was solved with two of the force balance constraints changed to account for the change in the direction of gravity relative to the rover's frame of reference, with equations 7.12 and 7.23 replaced respectively by

$$\Sigma F_x = F_{x_{FR}} + F_{x_{FL}} + F_{x_{RR}} + F_{x_{RL}} - W_{rover} \sin\alpha = 0 \tag{7.32}$$

$$\Sigma F_z = W_{FR} + W_{FL} + W_{RR} + W_{RL} - W_{rover} \cos\alpha = 0 \tag{7.33}$$

for a given slope angle α. 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 6.2.1.

The optimization problem was solved repeatedly with slope constraints for slope angles of 0° to 25° 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 6.3, with the speed-limited wheel driving at half the rover's travel speed and the steer motor stuck at 50°. Figure 7.12 shows both the 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 ω for each wheel, $P = \Sigma M_y \omega$.

Two values are given for the slip ratio in Figure 7.12: the maximum slip on any wheel and the average slip across all wheels not experiencing failure. Wheels experiencing failure are excluded

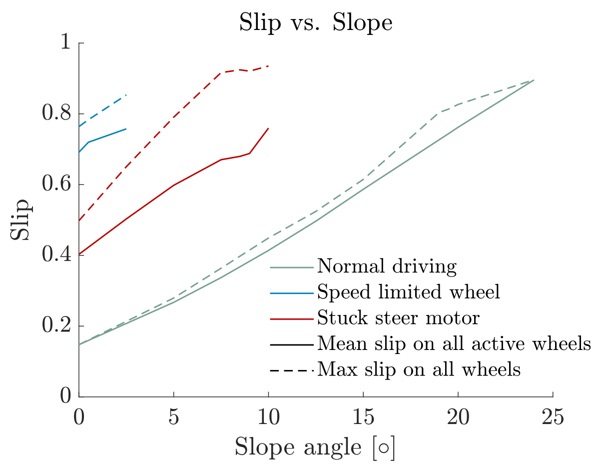

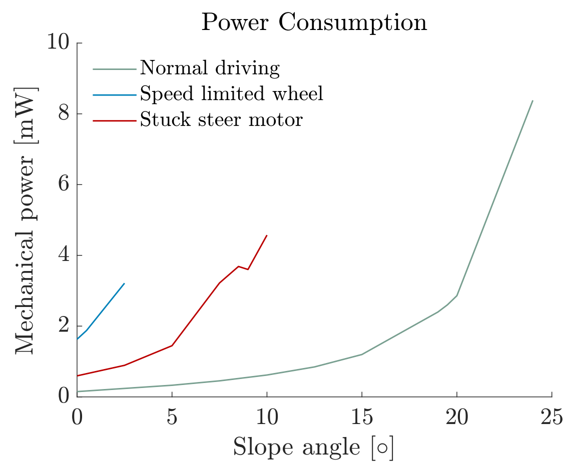

Figure 7.12: Plot of slip vs. equivalent slope angle (left) and power usage vs. slope angle for driving with failed actuators. Slip ratio is plotted as an average of the slip ratios on all non-impacted wheels as well as the maximum value of all wheels. Power is given by the sum of mechanical power (wheel torque × angular velocity) from each individual wheel.

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 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°; 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°, with 40% average slip on flat ground compared to only 15% slip on flat ground for nominal driving. The speed-limited case fares even worse, with 69% average slip and a maximum wheel slip of 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 50% in order to minimize entrapment risk, we expect the rover to be able to safely ascend slopes up to 12.5° with normal driving conditions. We would limit the rover to slopes up to 2.5° 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°, driving up a slope of 5° 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 [112].

7.4 Conclusions

In this chapter we developed 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 then 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. We then 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.

We then 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.

We explored only three types of constrained driving problems here, though 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.

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.

Chapter 8

Conclusion

8.1 Summary and Contributions

In this document we have expanded the capabilities of current and future planetary exploration rovers by directly enabling use of their wheels as manipulators and mitigating the impact of hardware degradation on mobility through the development and implementation of new soil models.

First, we proposed Non-Prehensile Terrain Manipulation (NPTM) as a strategy for augmenting the capabilities of planetary exploration rovers. We classified the different NPTM actions a rover can perform with minimal or no hardware additions, and identified the necessary metrics for performing and evaluating the success of these actions. The tables developed in this work were used to select wheel-based trenching as a candidate NPTM action to explore in-depth and serve as a path forward for the further development of rover NPTM capabilities.

We then created new soil models to fill identified gaps in the literature and enable both NPTM and recovery from degraded mobility. First, we addressed the motion of soil flow around a wheel, which is necessary both for wheel-based trench digging and for improving terramechanics modeling by predicting wheel-soil contact geometry. We created a closed-form solution to terrain shape following the passage of a trenching or driving wheel that can quantitatively and qualitatively describe the geometry of the resulting trench. The model was demonstrated both in a controlled lab environment and on undisturbed soil in the Atacama Desert. We then built a new terramechanics model on the foundations of classic terramechanics methods, with the new model covering a range of slip angles and slip ratios previously unexplored in the literature. The terramechanics model incorporates the soil flow model to determine the wheel-soil interaction geometry based on soil mechanics, eliminating the need for two standard tuning variables in the process. This model was validated on two wheels driving in soft soil in a purpose-built terramechanics testbed.

Trenches up to 1/3 of wheel diameter were dug by the KREX-2 rover's wheels in naturally occurring soil in the Atacama Desert, showing the potential for wheel-based trenching to move meaningful amounts of soil. Multiple passes over the same trench were shown to have diminishing returns. Wheel-based hole digging was used to facilitate the measurement of the soil angle of repose, illustrating how NPTM can be used to aid scientific sampling. Additionally, first-of-their-kind demonstrations of NPTM for rover teaming scenarios were performed in the Atacama Desert and NASA Ames Research Center's Roverscape, in which a large rover was able to directly enable passage through difficult terrain by a smaller robot.

We quantitatively and qualitatively assessed the impact of mobility system degradation on driving for the VIPER lunar rover in lunar simulant at NASA Glenn SLOPE Lab through drawbar pull and untethered driving tests. Ad hoc methods of compensating for actuator failure were demonstrated on the KREX-2 rover in the Atacama Desert and NASA Ames Roverscape, with the rover able to both drive straight with a steering actuator pointed outwards at several angles and with a wheel nonrotating. These tests show the feasibility of compensating for degraded mobility and motivated the need for tools to automatically generate safe driving strategies.

Finally, we created a numerical optimization framework to automatically generate motion primitives for off-nominal driving. Utilizing the terramechanics and soil flow models previously developed, the optimization framework was used to generate both open-loop driving primitives and closed-loop position controllers for multiple modes of off-nominal driving. First, wheel-based trenching was performed, with the optimization successfully determining how to dig a clean-sided trench useful for subsurface soil observation while maintaining mobility. The optimization framework was then used to enable a rover to recover from both a wheel unable to drive at full speed and one stuck at a high steering angle. In both cases the rover was able to recover mobility; notably, even feedforward-only driving strategies can mitigate mobility failure, allowing for direct implementation of this technique on rovers without visual odometry or other means of closed-loop position control. Enabling rovers to continue operations following a wheel actuator failure could easily mean the difference between mission success and failure.

Taken together, the methods and results presented in this thesis have provided deeper insight into the interaction between rover wheels and granular media, and in doing so have increased the mobility and manipulation capabilities of rovers without any hardware changes.

8.2 Future Work

In this work we have performed foundational work in two key areas of off-nominal driving for rovers, and in doing so have identified new avenues for exploration. Here, we detail the next steps needed to realize the full potential of both NPTM and degraded mobility compensation.

Nonprehensile terrain manipulation as a whole is ripe for future development – of fifteen identified robot actions in Table 3.1, we have touched on a few and developed only trenching in depth. Autonomous reshaping of large swaths of terrain with wheels or rover-mounted tools is relevant to in-situ resource utilization [113], and discrete object manipulation methods as developed by [22] should be merged with soft soil modeling and manipulation methods to work in a wider variety of terrains.

Additionally, wheel-based manipulation of soft soil may be able to increase rover mobility in non-degraded modes via improved methods of climbing slopes. Uncontrolled lateral and longitudinal slip are serious threats to rover mobility on steep slopes of soft soil. The risk of entrapment, rollover, and inability to ascend or descend a slope to reach a destination can easily result in mission failure. Using trenching techniques to intentionally control wheel sinkage could allow rovers to prevent slip while traversing, ascending, and descending slopes. By extending the models developed in Chapters 4 and 5 and applying the same techniques described in Chapter 7, new mobility techniques can be developed. These techniques would effectively use the “dig trench" action described in Table 3.2, but with the purpose of controlling sinkage and similar mission scenarios to the “level slope" action in Table 3.1.

Rovers may be able to descend extremely steep soft slopes in a controlled manner by intentionally entering a high sinkage state with one or multiple wheels, effectively “plowing” down the slope as in [114]. Updated terramechanics and soil flow models would be required to fully simulate driving on a slope, and the same type of numerical optimization used in Chapter 7 could be used to control rover orientation on slopes, or allow some wheels to provide anchoring while others control driving. In [62], a reconfigurable suspension was used to reduce a rover’s roll while traversing slopes, which decreased downhill slip during driving. Additionally, in the same work and in several other works, yaw control was also used to decrease downhill slip. Intentional sinkage of uphill wheels may be able to gain some of the benefits of roll control without the need for a specialized, actuated suspension.

In this work we have focused on optimizing rover motion control for given platforms; we also now have the tools to optimize platform design for specific tasks. By adding in rover geometry such as wheel radius, grouser size, or suspension dimensions as decision variables in the optimization, we can redesign rovers to be better suited to both trenching and driving, or design them to minimize the impact of degradation of mobility systems. One challenge of this is the semi-empirical nature of the terramechanics model used; the relation of sinkage moduli and pressure distribution shape variables to wheel geometry would need to be tested at a wider range of wheel scales. Use of the terramechanics and soil flow models developed here would also require validation for wheels at the scale of VIPER or Perseverance before use on space hardware, and the model would need additional testing to be fully validated up to slip angles of 90°.

Thus far, we have considered only steady-state motion; a dynamic understanding of wheel-soil interaction would enable automatic generation of a broader range of locomotion and manipulation techniques, such as those developed in [109]. As classic terramechanics methods are not well-suited to dynamic modeling, a different approach such as Resistive Force Theory [81, 115] may be more appropriate when combined with full three dimensional modeling of soil flow [113, 116]. However, these methods are far more computationally expensive than the closed-form soil model and classic terramechanics model presented here and would be difficult to use in an optimization for that reason. Further work on efficient modeling of soil movement and wheel-soil interaction forces is necessary to enable novel dynamic motion primitive generation.

Rovers may someday slalom down sandy slopes or ride down on avalanching sand, or autonomously construct landing pads and lunar habitats from regolith. Future planetary rovers may have four wheels, six wheels, legs, neither legs or wheels, or something in between, but the exploration of the solar system will continue to be led by mobile robots pushing the bounds to go where no one has before.

Bibliography

[1] NASA press release, "NASA lands car-size rover beside Martian mountain," August 6th 2012. [Online]. Available: https://mars.nasa.gov/news/1288/nasa-lands-car-size-rover-beside-martian-mountain/

[2] ——, "Touchdown! NASA's Mars Perseverance rover safely lands on red planet," February 18th 2021. [Online]. Available: https://mars.nasa.gov/news/8865/touchdown-nasas-mars-perseverance-rover-safely-lands-on-red-planet/

[3] ——, "NASA's perseverance rover deposits first sample on Mars surface," December 1st 2022. [Online]. Available: https://www.jpl.nasa.gov/news/nasas-perseverance-rover-deposits-first-sample-on-mars-surface

[4] ——, "VIPER mission overview," September 9th 2022. [Online]. Available: https://www.nasa.gov/viper/overview

[5] B. Spice, "CMU's MoonRanger will search for water at Moon's south pole," September 23rd 2020. [Online]. Available: https://www.scs.cmu.edu/news/2020/cmu-moonranger-will-search-water-moons-south-pole

[6] A. Aupperlee, "CMU's Iris rover secured to lunar lander," December 3rd 2021. [Online]. Available: https://www.cmu.edu/news/stories/archives/2021/december/iris-moon-rover.html

[7] P. Leger, A. Trebi-Ollennu, J. Wright *et al.*, "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.

[8] M. Heverly, J. Matthews, J. Lin *et al.*, "Traverse performance characterization for the Mars Science Laboratory rover," *Journal of Field Robotics*, vol. 30, no. 6, pp. 835–846, 2013.

[9] J. Townsend, M. Seibert, P. Bellutta *et al.*, "Mars exploration rovers 2004-2013: Evolving operational tactics driven by aging robotic systems," *International Conference on Space Operations*, no. May, pp. 1–22, 2014.

[10] 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.

[11] NASA press release, "NASA's Opportunity rover rolls free on Mars," June 6th 2005. [Online]. Available: https://mars.nasa.gov/mer/newsroom/pressreleases/20050606a.html

[12] ——, "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

[13] 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.

[14] ——, "A terramechanics model for high slip angle and skid with prediction of wheel-soil interaction geometry," *Journal of Terramechanics*, 2023, under review.

[15] ——, "Field experiments in nonprehensile terrain manipulation with planetary exploration rovers," in *International Symposium on Artificial Intelligence, Robotics and Automation in Space*, October 2020, workshop poster.

[16] 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.

[17] C. Pavlov and A. M. Johnson, “Terramechanics-based joint failure compensation and soil manipulation,” 2023, in prep.

[18] K. M. Lynch and M. T. Mason, “Controllability of pushing,” in *IEEE International Conference on Robotics and Automation*, vol. 1, 1995, pp. 112–119.

[19] J. E. King, M. Klingensmith, C. M. Dellin *et al.*, “Pregrasp manipulation as trajectory optimization,” in *Robotics: Science and Systems*, 2013.

[20] 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.

[21] 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.

[22] 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.

[23] J. King, “Robust rearrangement planning using nonprehensile interaction,” Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, USA, 2016.

[24] M. Mason, D. Pai, D. Rus *et al.*, “A mobile manipulator,” in *IEEE International Conference on Robotics and Automation*, vol. 3, 1999, pp. 2322–2327.

[25] S. S. Srinivasa, C. R. Baker, E. Sacks *et al.*, “Experiments with nonholonomic manipulation,” in *IEEE International Conference on Robotics and Automation*, vol. 2, 2002, pp. 2042–2047.

[26] R. Sullivan, R. Anderson, J. Biesiadecki *et al.*, “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.

[27] J. Biesiadecki, E. T. Baumgartner, R. G. Bonitz *et al.*, “Mars Exploration Rover surface operations,” *IEEE Robotics and Automation Magazine*, vol. 13, no. 2, pp. 63–71, 2006.

[28] 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

[29] NASA press release, “NASA InSight’s ’mole’ is out of sight,” October 16th 2020. [Online]. Available: https://www.jpl.nasa.gov/news/nasa-insights-mole-is-out-of-sight

[30] J. Matijevic and E. Dewell, “Anomaly Recovery and the Mars Exploration Rovers,” in *SpaceOps 2006 Conference*. Reston, Virigina: American Institute of Aeronautics and Astronautics, June 2006.

[31] NASA press release, “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/

[32] A. Rankin, M. Maimone, J. Biesiadecki *et al.*, “Mars curiosity rover mobility trends during the first 7 years,” *Journal of Field Robotics*, vol. 38, no. 5, pp. 759–800, 2021.

[33] NASA press release, “VIPER: The rover and its onboard toolkit,” February 16th 2022. [Online]. Available: https://www.nasa.gov/viper/rover

[34] ——, “New VIPER lunar rover to map water ice on the Moon,” October 25th 2019. [Online]. Available: https://www.nasa.gov/feature/new-viper-lunar-rover-to-map-water-ice-on-the-moon

[35] M. G. Bekker, *Theory of Land Locomotion*. University of Michigan Press, 1956.

[36] J. Y. Wong and A. Reece, "Prediction of rigid wheel performance based on the analysis of soil-wheel stresses. Part II. Towed rigid wheels." *Journal of Terramechanics*, vol. 4, no. 2, pp. 7–25, 1967.

[37] J. Y. Wong and A. R. Reece, "Prediction of rigid wheel performance based on the analysis of soil-wheel stresses. Part I. Performance of driven rigid wheels," *Journal of Terramechanics*, vol. 4, no. 1, pp. 81–98, 1967.

[38] J. Y. Wong, *Theory of Ground Vehicles*. John Wiley & Sons, 2008.

[39] Z. Janosi and B. Hanamoto, *An Analysis of the Drawbar Pull vs Slip Relationship for Track Laying Vehicles*. Land Locomotion Laboratory, 1961.

[40] M. Lyasko, "Slip sinkage effect in soil-vehicle mechanics," *Journal of Terramechanics*, vol. 47, no. 1, pp. 21–31, 2010.

[41] L. Ding, K. Nagatani, K. Sato *et al.*, "Terramechanics-based high-fidelity dynamics simulation for wheeled mobile robot on deformable rough terrain," *IEEE International Conference on Robotics and Automation*, 2010.

[42] H. Gao, J. Guo, L. Ding *et al.*, "Longitudinal skid model for wheels of planetary exploration rovers based on terramechanics," *Journal of Terramechanics*, vol. 50, no. 5-6, pp. 327–343, 2013.

[43] L. Ding, Z. Deng, H. Gao *et al.*, "Interaction mechanics model for rigid driving wheels of planetary rovers moving on sandy terrain with consideration of multiple physical effects," *Journal of Field Robotics*, vol. 32, no. 6, pp. 827–859, 2015.

[44] J. Guo, L. Ding, H. Gao *et al.*, "Longitudinal skid model for wheels of planetary rovers based on improved wheel sinkage considering soil bulldozing effect," *Journal of Terramechanics*, vol. 74, pp. 45–56, 2017.

[45] J. Guo, H. Gao, L. Ding *et al.*, “Linear normal stress under a wheel in skid for wheeled mobile robots running on sandy terrain,” *Journal of Terramechanics*, vol. 70, pp. 49–57, 2017.

[46] L. Ding, H. Yang, H. Gao *et al.*, “Terramechanics-based modeling of sinkage and moment for in-situ steering wheels of mobile robots on deformable terrain,” *Mechanism and Machine Theory*, vol. 116, no. 2017, pp. 14–33, 2017.

[47] G. Meirion-Griffith and M. Spenko, “A modified pressure-sinkage model for small, rigid wheels on deformable terrains,” *Journal of Terramechanics*, vol. 48, no. 2, pp. 149–155, 2011.

[48] G. Meirion-Griffith, C. Nie, and M. Spenko, “Development and experimental validation of an improved pressure-sinkage model for small-wheeled vehicles on dilative, deformable terrain,” *Journal of Terramechanics*, vol. 51, no. 1, pp. 19–29, 2014.

[49] C. Senatore and K. D. Iagnemma, “Direct shear behaviour of dry, granular soils for low normal stress with application to lightweight robotic vehicle modelling,” *International Conference of the International Society for Terrain Vehicle Systems*, pp. 1–11, 2011.

[50] C. Senatore and K. Iagnemma, “Analysis of stress distributions under lightweight wheeled vehicles,” *Journal of Terramechanics*, vol. 51, no. 1, 2014.

[51] K. Skonieczny, S. Moreland, V. Asnani *et al.*, “Visualizing and Analyzing Machine-soil Interactions using Computer Vision,” *Journal of Field Robotics*, vol. 31, no. 5, pp. 820–836, 2014.

[52] 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.

[53] Z. Jia, W. Smith, and H. Peng, “Fast analytical models of wheeled locomotion in deformable terrain for mobile robots,” *Robotica*, vol. 31, no. 1, pp. 35–53, 2013.

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

[55] H. Inotsume, S. Moreland, K. Skonieczny, and D. Wettergreen, “Parametric study and design guidelines for rigid wheels for planetary rovers,” *Journal of Terramechanics*, 2019.

[56] R. A. Irani, R. J. Bauer, and A. Warkentin, “A dynamic terramechanic model for small lightweight vehicles with rigid wheels and grousers operating in sandy soil,” *Journal of Terramechanics*, vol. 48, no. 4, pp. 307–318, 2011.

[57] H. Schwanghart, “Lateral forces on steered tyres in loose soil,” *Journal of Terramechanics*, vol. 5, no. 1, pp. 9–29, 1968.

[58] G. Ishigami, A. Miwa, K. Nagatani, and K. Yoshida, “Terramechanics-based model for steering maneuver of planetary exploration rovers on loose soil,” *Journal of Field Robotics*, vol. 24, no. 3, pp. 233–250, 2007.

[59] C. Senatore and C. Sandu, “Off-road tire modeling and the multi-pass effect for vehicle dynamics simulation,” *Journal of Terramechanics*, vol. 48, no. 4, pp. 265–276, 2011.

[60] A. Azimi, D. Holz, J. Kövecses *et al.*, “A multibody dynamics framework for simulation of rovers on soft terrain,” *Journal of Computational and Nonlinear Dynamics*, vol. 10, no. 3, pp. 1–12, 2015.

[61] Y. Li, L. Ding, and G. Liu, “Attitude-based dynamic and kinematic models for wheels of mobile robot on deformable slope,” *Robotics and Autonomous Systems*, vol. 75, pp. 161–175, 2016. [Online]. Available: http://dx.doi.org/10.1016/j.robot.2015.10.006

[62] H. Inotsume, M. Sutoh, K. Nagaoka *et al.*, “Modeling, analysis, and control of an actively reconfigurable planetary rover for traversing slopes covered with loose soil,” *Journal of Field Robotics*, vol. 30, no. 6, 2013.

[63] L. Ding, Z. Deng, H. Gao *et al.*, “Experimental study and analysis of the wheels’ steering mechanics for planetary exploration wheeled mobile robots moving on deformable terrain,” *The International Journal of Robotics Research*, vol. 32, no. 6, pp. 712–743, 2013.

[64] H. Yang, L. Ding, H. Gao *et al.*, “High-fidelity dynamic modeling and simulation of planetary rovers using single-input-multi-output joints with terrain property mapping,” *IEEE Transactions on Robotics*, pp. 1–21, 2022.

[65] K. Terzaghi, *Theoretical Soil Mechanics*. John Wiley & Sons, Ltd, 1943.

[66] E. Hegedus, *A Simplified Model for the Determination of Bulldozing Resistance*. Center Line, Michigan: Land Locomotion Laboratory Research Division, Army Tank Automotive Command Report, 1960.

[67] A. R. Reece, “The fundamental equation of mechanics,” *Proceedings of the Institution of Mechanical Engineers*, vol. 179, no. 6, 1964.

[68] D. R. P. Hettiaratchi, B. D. Witney, and A. R. Reece, “The calculation of passive pressure in two-dimensional soil failure,” *Journal of Agricultural Engineering Research*, vol. 11, no. 2, pp. 89–107, 1966.

[69] K. Yoshida and G. Ishigami, “Steering characteristics of a rigid wheel for exploration on loose soil,” in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 4, 2004, pp. 3995–4000.

[70] G. Ishigami, “Terramechanics-based analysis and control for lunar/planetary exploration robots,” Ph.D. dissertation, Tohoku University, March 2008.

[71] K. Yoshida, D. Nenchev, G. Ishigami, and Y. Tsumaki, “Space robotics,” in *The International Handbook of Space Technology*, M. Macdonald and V. Badescu, Eds. Berlin, Heidelberg: Springer, 2014, pp. 541–573.

[72] B. Trease, R. Arvidson, R. Lindemann *et al.*, “Dynamic modeling and soil mechanics for path planning of mars exploration rovers,” in *MSC Software 2011 Users Conference, Americas*, 2011, pp. 755–765.

[73] E. McKyes, *Soil Cutting and Tillage*, ser. Developments in Agricultural Engineering 7. Amsterdam: Elsevier, 1985.

[74] J. Guo, W. Li, L. Ding *et al.*, “High-slip wheel-terrain contact modelling for grouser-wheeled planetary rovers traversing on sandy terrains,” *Mechanism and Machine Theory*, vol. 153, p. 104032, 2020.

[75] M. A. Knuth, J. B. Johnson, M. A. Hopkins *et al.*, “Discrete element modeling of a Mars Exploration Rover wheel in granular material,” *Journal of Terramechanics*, vol. 49, no. 1, pp. 27–36, 2012.

[76] J. Hambleton and A. Drescher, “Modeling wheel-induced rutting in soils: Rolling,” *Journal of Terramechanics*, vol. 46, no. 2, pp. 35–47, 2009.

[77] W. Smith, D. Melanz, C. Senatore *et al.*, “Comparison of discrete element method and traditional modeling methods for steady-state wheel-terrain interaction of small vehicles,” *Journal of Terramechanics*, vol. 56, pp. 61–75, 2014.

[78] S. Taheri, C. Sandu, S. Taheri *et al.*, “A technical survey on terramechanics models for tire-terrain interaction used in modeling and simulation of wheeled vehicles,” *Journal of Terramechanics*, vol. 57, pp. 1–22, 2015.

[79] W. Smith and H. Peng, “Modeling of wheel-soil interaction over rough terrain using the discrete element method,” *Journal of Terramechanics*, vol. 50, pp. 277–287, 2013.

[80] R. Krenn and A. Gibbesch, “Soft soil contact modeling technique for multi-body system simulation,” *Trends in Computational Contact Mechanics*, vol. 58 LNACM, pp. 135–155, 2011.

[81] W. Li, L. Ding, H. Gao *et al.*, “ROSTDyn: Rover simulation based on terramechanics and dynamics,” *Journal of Terramechanics*, vol. 50, no. 3, pp. 199–210, 2013.

[82] H. Suzuki, K. Katsushima, and S. Ozaki, “Study on applicability of RFT to traveling analysis of wheel with grousers: Comparison with DEM analysis as a virtual test,” *Journal of Terramechanics*, vol. 83, pp. 15–24, 2019. [Online]. Available: https://doi.org/10.1016/j.jterra.2019.01.001

[83] S. Agarwal, C. Senatore, T. Zhang *et al.*, “Modeling of the interaction of rigid wheels with dry granular media,” *Journal of Terramechanics*, vol. 85, pp. 1–14, 2019.

[84] M. Servin, T. Berglund, and S. Nystedt, “A multiscale model of terrain dynamics for real-time earthmoving simulation,” *Advanced Modeling and Simulation in Engineering Sciences*, vol. 8, 2021. [Online]. Available: https://doi.org/10.1186/s40323-021-00196-3

[85] D. Holz, T. Beer, and T. Kuhlen, “Soil deformation models for real-time simulation: A hybrid approach,” *Workshop on Virtual Reality Interaction and Physical Simulation VRIPHYS*, 2009.

[86] Y. Zhu, X. Chen, and G. S. Owen, “Terramechanics based terrain deformation for real-time off-road vehicle simulation,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 6938 LNCS, no. PART 1, pp. 431–440, 2011.

[87] R. Gonzalez, D. Apostolopoulos, and K. Iagnemma, “Improving rover mobility through traction control: simulating rovers on the Moon,” *Autonomous Robots*, vol. 43, no. 8, pp. 1977–1988, 2019.

[88] H. Inotsume, M. Sutoh, K. Nagaoka *et al.*, “Slope traversability analysis of reconfigurable planetary rovers,” in *IEEE International Conference on Intelligent Robots and Systems*, 2012, pp. 4470–4476.

[89] A. Gallina, A. Gibbesch, R. Krenn *et al.*, “Multibody simulation of planetary rover mobility in condition of uncertain soft terrain,” *Procedia IUTAM*, vol. 13, pp. 118–126, 2015.

[90] NASA press release, “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

[91] ProtoInnovations LLC, “K10-mini : Modular robotic sentry rover.” [Online]. Available: http://www.protoinnovations.com/

[92] NASA Science: Mars Exploration Program, “Mars Curiosity Rover.” [Online]. Available: https://mars.nasa.gov/msl/home/

[93] D. S. Lauretta, D. N. DellaGiustina, and e. a. Bennett, “The unexpected surface of asteroid (101955) Bennu,” *Nature*, vol. 568, no. 7750, pp. 55–60, 2019.

[94] NASA press release, “Origami-inspired robot can hitch a ride with a rover,” March 20th 2017. [Online]. Available: https://www.nasa.gov/feature/jpl/origami-inspired-robot-can-hitch-a-ride-with-a-rover

[95] T. Botha, D. Johnson, S. Els, and S. Shoop, “Real time rut profile measurement in varying terrain types using digital image correlation,” *Journal of Terramechanics*, vol. 82, pp. 53–61, 2019.

[96] E. T. Vincent, “Pressure distribution on and flow of sand past a rigid wheel: technical report,” University of Michigan Research Institute, Tech. Rep., 1960.

[97] R. A. Bagnold, *The Physics of Blown Sand and Desert Dunes*. Dover Publications, Inc., 1941.

[98] D. Catanoso, T. Stucky, J. Case, and A. Rogg, “Analysis of sample acquisition dynamics using discrete element method,” *IEEE Aerospace Conference Proceedings*, 2020.

[99] Bota Systems, “Rokubi: Small 6-axis force torque sensor,” 2022. [Online]. Available: https://www.botasys.com/rokubi

[100] Soil Direct, “Soil Direct,” 2022. [Online]. Available: https://www.soildirect.com/

[101] L. Ding, Haibo Gao, Yuankai Li *et al.*, “Improved explicit-form equations for estimating dynamic wheel sinkage and compaction resistance on deformable terrain,” *Mechanism and Machine Theory*, vol. 86, pp. 235–264, 2015.

[102] FARO, “FARO®Focus Laser Scanners.” [Online]. Available: https://www.faro.com/en/Products/Hardware/Focus-Laser-Scanners

[103] Leica Geosystems, “Total stations.” [Online]. Available: https://leica-geosystems.com/en-us/products/total-stations

[104] GoPro, “GoPro Action Cameras.” [Online]. Available: https://gopro.com/en/us/

[105] 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.

[106] 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

[107] ——, “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

[108] 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

[109] S. Shrivastava, A. Karsai, Y. O. Aydin *et al.*, “Material remodeling and unconventional gaits facilitate locomotion of a robophysical rover over granular terrain,” *Science Robotics*, vol. 5, no. 42, 2020.

[110] 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.

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

[112] 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.

[113] W. Kim, C. Pavlov, and A. M. Johnson, “Developing a simple model for sand-tool interaction and autonomously shaping sand,” *arXiv*, vol. abs/1908.02745 [cs.RO], 2019. [Online]. Available: http://arxiv.org/abs/1908.02745

[114] D. Kohanbash, S. Moreland, and D. Wettergreen, “Plowing for rover control on extreme slopes,” *Springer Tracts in Advanced Robotics*, vol. 92, pp. 399–413, 2014.

[115] S. Agarwal, A. Karsai, D. I. Goldman, and K. Kamrin, “Surprising simplicity in the modeling of dynamic granular intrusion,” *Science Advances*, vol. 7, no. 17, 2021.

[116] Q. Yu, C. Pavlov, W. Kim, and A. M. Johnson, “Wheeled locomotion in granular media using 3D-RFT and a sand deformation model,” 2023, in prep.

[117] R. Krenn and G. Hirzinger, "Simulation of rover locomotion on sandy terrain - modeling, verification and validation," *Proceedings of the 10th Workshop on ASTRA, Noordwijk, The Netherlands*, 2008.

[118] N. Flowers, M. Barragan, and A. M. Johnson, "MiniRHex: An open-source walking hexapod," in *Dynamic Walking*, Pensacola, FL, May 2018, workshop abstract.

[119] NASA press release, "Rover takes on steepest slope ever tried on Mars," March 31st 2016. [Online]. Available: https://www.nasa.gov/feature/jpl/rover-takes-on-steepest-slope-ever-tried-on-mars

[120] NASA Science: Earth's Moon, "Aristarchus crater." [Online]. Available: https://moon.nasa.gov/resources/347/aristarchus-crater/

[121] "Lunar all-terrain utility vehicle for eva, phase ii project." [Online]. Available: techport.nasa.gov:443/file/13692

[122] C. Atwood-Stone and A. S. McEwen, "Avalanche slope angles in low-gravity environments from active Martian sand dunes," *Geophysical Research Letters*, vol. 40, no. 12, pp. 2929–2934, 2013.

[123] C. Pavlov and A. M. Johnson, "Wheel-based trenching: Terramechanics of nonprehensile manipulation for planetary rovers," in *Robotics: Science and Systems Workshop on "Autonomous Space Robotics"*, Pittsburgh, PA, June 2018, Workshop abstract.

[124] ——, "Wheel-based trenching: Terramechanics of nonprehensile manipulation on planetary rovers," in *Robotics: Science and Systems Workshop on "Women In Robotics"*, Pittsburgh, PA, June 2018, Workshop poster.

[125] K. Skonieczny, D. Wettergreen, and W. Whittaker, "Advantages of continuous excavation in lightweight planetary robotic operations," *International Journal of Robotics Research*, vol. 35, no. 9, pp. 1121–1139, 2016.

[126] L. L. Johnson and R. H. King, “Measurement of force to excavate extraterrestrial regolith with a small bucket-wheel device,” *Journal of Terramechanics*, vol. 47, no. 2, pp. 87–95, 2010.

[127] Brendan Juin-Yih Chan, “Development of an off-road capable tire model for vehicle dynamics simulations,” *Virginia Polytechnic Institute*, no. January, 2008.

[128] G. Ishigami, K. Nagatani, and K. Yoshida, “Slope traversal controls for planetary exploration rover on sandy terrain,” *Journal of Field Robotics*, vol. 26, no. 3, pp. 264–286, 2008.

[129] D. M. Helmick, S. I. Roumeliotis, Y. Cheng *et al.*, “Slip-compensated path following for planetary exploration rovers,” *Advanced Robotics*, vol. 20, no. 11, pp. 1257– 1280, 2006.

[130] R. E. Arvidson, P. Bellutta, F. Calef *et al.*, “Terrain physical properties derived from orbital data and the first 360 sols of Mars Science Laboratory Curiosity rover observations in Gale Crater,” *Journal of Geophysical Research: Planets*, vol. 119, no. 1322-1344, 2014.

[131] L. Ding, H. Gao, Z. Deng, and J. Tao, “Wheel slip-sinkage and its prediction model of lunar rover,” *J. Cent. South Univ. Technol.*, vol. 17, pp. 129–135, 2010.

[132] L. Ding, H. Gao, Z. Deng *et al.*, “Longitudinal slip versus skid of planetary rovers’ wheels traversing on deformable slopes,” in *IEEE International Conference on Intelligent Robots and Systems*, 2013.

[133] D. Fujiwara and K. Iizuka, “Study on bearing performance for inching worm locomotion using characteristics of wheel subsidence on loose soil,” *IEEE/ASME International Conference on Advanced Intelligent Mechatronics, AIM*, vol. 2018-July, pp. 930–935, 2018.

[134] K. Terzaghi, R. Peck, and G. Mesri, *Soil Mechanics in Engineering Practice*, ser. Wiley-Interscience publication. Wiley, 1996. [Online]. Available: https://books.google.com/books?id=bAwVvO71FXoC

[135] R. E. Arvidson, K. D. Iagnemma, M. Maimone *et al.*, “Mars Science Laboratory Curiosity rover megaripple crossings up to sol 710 in Gale Crater,” *Journal of Field Robotics*, vol. 34, no. 3, pp. 495–518, 2017.

[136] W. Reid, B. Emanuel, B. Chamberlain-Simon *et al.*, “Mobility mode evaluation of a wheel-on-limb rover on glacial ice analogous to Europa terrain,” *2020 IEEE Aerospace Conference*, pp. 1–9, 2020.