

Mass-Constrained Robotic Climbing on Irregular Terrain

Paul Nadan

CMU-RI-TR-25-05

February 2025

The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee

Aaron M. Johnson, Chair

William (Red) Whittaker

Zeynep Temel

Spencer Backus (NASA Jet Propulsion Laboratory)

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Robotics.*

Acknowledgments

I want to begin by acknowledging the many funding sources that made this research possible. This work was supported by a NASA Space Technology Graduate Research Opportunities Award. Portions of this work were funded by the Chevron Technical Center (CTC) and the Pennsylvania Infrastructure Technology Alliance (PITA). This material is also based upon work supported by the National Science Foundation under grants #CMMI-1943900 and #ECCS-1924723. A portion of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

When I was applying to graduate school, the most common advice I received was to choose a good advisor, because they have the potential to make or break your Ph.D. experience. I am incredibly grateful to my advisor, Aaron Johnson, for making my experience such a positive one. Aaron, thank you for giving me both the freedom to explore and the support to take risks. It has truly been an honor to work with you.

I would like to thank my thesis committee, Red Whittaker, Zeynep Temel, and Spencer Backus, for sharing their experience and wisdom throughout this process. This document would be in much worse condition if not for their feedback. I am also grateful to Spencer for his mentorship over the last four years as my NSTGRO collaborator. Spencer, thank you for always standing up for me. You are an endless source of good ideas and I have learned so much from you.

Next, I would like to acknowledge the many amazing collaborators who contributed to this work. First, the magnetic wall-climber team: Jai, Nate, Jesse, Nathan, Hairong, Alexander, and Greg – thank you for all of your hard work to make this project a success. To my mentees: it has been a privilege to watch you learn and grow over the years. I would also like to express my gratitude to everyone on the EELS project – thank you all for making me feel like part of the team from day one. Rohan, thank you for always valuing my input and ensuring that I got the most out of the experience. Rudra, although our project did not quite make it into this thesis, I am still grateful I had the opportunity to learn so much working with you.

To my labmates in the Robomechanics Lab: you all are the reason I look forward to coming in to work each day. Catherine, Joe Norby, and Joe Payne: starting a PhD during a pandemic is never a good time, but playing D&D with you all got me through the worst of it. Thank you for welcoming me into the lab with open arms and for being amazing friends and role models throughout this journey. Justin, thank you for being an endless source of enthusiasm and helping me see the beauty in everything from brick patterns to e-waste. Sean, your visionary leadership as lab manager and your wry sense of humor were integral to lab culture. Ardalán, thank you for always having more confidence in me than I do myself. David, thank you for keeping me company for all those late nights in the lab. Steven, Naomi, Wensen, Diana, Aja, and Vivek: thank you for so many fascinating conversations about cool bugs, operating systems, terrifying extinct bugs, space robots, and whether we should all just give up and use RL. I'm going to

really miss you all, but I know the lab is in good hands and I cannot wait to see what you accomplish next.

To my housemate, Sam: thank you for being a great friend and entertaining all of my most ridiculous ideas. Jeremy, Daniel, and Nate: thank you for providing so many welcome distractions from my research and for all your support even from afar. To everyone in the Tartan Wind Ensemble: you all are awesome and talented and I have really enjoyed making music together.

To my brother, Will: it has been so inspiring to watch you turn your passions into a career, and I know you will find success in anything you decide to pursue. Thanks for putting up with me for all these years. To my parents, Michele and Chris: for as long as I can remember you have encouraged and supported my curiosity wherever it led, and that encouragement is what shaped me into the person I am today. Thank you for always believing in me, for always being able to make me smile, and for always being just a phone call away. And finally, to my partner, Adhayana: thank you for your endless support and for bringing so much joy into my life. I cannot wait to start our next chapter together.

Abstract

Climbing robots can operate in steep and unstructured environments that are inaccessible to other ground robots, with applications ranging from the inspection of artificial structures on Earth to the exploration of natural terrain features throughout the solar system. Climbing robots for planetary exploration face many challenges to deployment, including mass restrictions, irregular surface features, and communication delays. We present a hierarchical approach that overcomes these obstacles via underactuated design to comply with surface features, internal force control to maximize adhesion, and hybrid path planning to enable autonomy on complex 3D terrain.

We first apply compliant mechanism design at the scale of individual microspines, which rely on flexible suspensions to load-share and conform to surface features. We present a new suspension design that can be 3D printed from a single material to enable faster and easier fabrication with no reduction in mechanical performance. We then apply underactuation at the gripper scale, presenting a fully passive microspine gripper and wrist that can nevertheless solve the paired problems of conforming to uneven terrain and maintaining adhesion over a wide range of loading angles. These grippers are mounted on the lightweight rock-climbing robot LORIS, which has demonstrated the successful ascent of several different climbing surfaces. Next, we move to the full robot scale with the design of Sally, a magnetic-wheeled inspection robot for steel structures. Sally uses a passive suspension to maintain end effector contact despite surface irregularities and repurposes the steering and sensor deployment actuators to assist in difficult corner transitions.

At the control level, we generalize the bioinspired directed inward gripping technique as an internal force optimization problem that we apply to both LORIS and Sally. We also present a unified shape, contact, force, and motion control scheme that was used by the snake-like robot EELS to successfully descend an ice shaft. At the planning level, we present a motion planning approach for hybrid locomotion that applies the sampling-based AIT* algorithm to rapidly search for paths along arbitrarily shaped terrain manifolds modeled by point clouds. This global planner relies on a local planning algorithm to select high-quality footholds. Lastly, we examine planner solutions both on hardware and in simulation. We end by discussing possible extensions of this work to conventional quadrupeds and the pursuit of dynamic climbing capabilities.

Contents

1	Introduction	1
1.1	Thesis Outline	3
2	Microspine Design for Additive Manufacturing	7
2.1	Introduction	7
2.2	Related Work	8
2.3	Microspine Design Approaches	10
2.3.1	SDM Microspine Design	10
2.3.2	HDM Microspine Design	11
2.3.3	FDM Microspine Design	11
2.4	Experimental Characterization	16
2.5	Discussion	19
3	Design of a Lightweight Rock-Climbing Robot	23
3.1	Introduction	23
3.2	Robot Design	24
3.2.1	Splayed Microspine Gripper	25
3.2.2	Passive Wrist Joint	27
3.2.3	Robot Prototype	28
3.3	Climbing Gait	30
3.4	Experimental Performance	30
3.5	Discussion	31
4	Design of a Magnetic-Wheeled Inspection Robot	35
4.1	Robot Design	37
4.2	Interior Corner Transitions	39
4.2.1	Transition Model	39
4.2.2	Transition Strategies	41
4.2.3	Transition Control	44
4.3	Experiments	45
4.3.1	Interior Corner and Obstacle Traversal	45
4.3.2	Field Deployment	45

4.3.3	Line Scanning	46
4.4	Discussion	47
5	Unified Control of Shape, Contact, Force, and Motion	49
5.1	Introduction	49
5.2	Robot Model	50
5.3	Contact Force Estimation	51
5.4	Internal Force Optimization	52
5.4.1	General Formulation	53
5.4.2	Adhesion Constraints	54
5.4.3	Optimization Results	55
5.5	Unified Control Scheme	56
5.6	Discussion	59
6	Hybrid Motion Planning on Irregular Surfaces	61
6.1	Introduction	61
6.2	Related Work	62
6.3	Problem Setup	63
6.4	Local Planning	64
6.4.1	Foothold Rejection Sampling	64
6.4.2	Foothold Cost Heuristic	65
6.4.3	Body Pose Selection	65
6.4.4	Adhesion Validation	66
6.5	Global Planning	67
6.6	Planning Results	68
6.7	Discussion	69
7	Conclusion	71
7.1	Future Directions	72
	Bibliography	75

List of Figures

1.1	The LORIS, Sally, and EELS 1.5 climbing robots	2
1.2	A conceptual framework for solving the key challenges of robotic climbing	3
2.1	A simple gripper composed of 3D-printed microspines, subject to a 30 N load . .	8
2.2	Microspine geometry and examples made with each manufacturing process	10
2.3	The Hybrid Deposition Manufacturing (HDM) process for microspines.	11
2.4	Failure modes encountered during microspine design	12
2.5	Geometric parameters involved in computing the effective stiffness of a microspine, and empirical load-displacement curves compared against the linear model	14
2.6	Experimental setup for measurement of microspine axial and normal stiffness . .	17
2.7	Overlaid axial and normal load-displacement curves for each microspine design .	18
2.8	Stress-strain curves of SDM, HDM, and FDM microspines loaded until failure . .	19
2.9	SDM and FDM microspine performance measured over 100 load cycles	20
3.1	The LORIS climbing robot	24
3.2	A splayed microspine gripper and analysis of its loading capabilities	26
3.3	Analysis of ball-joint placement for the passive wrist	28
3.4	An annotated visualization of the LORIS robot design	29
3.5	The LORIS robot ascending cinder block, vesicular basalt, slag, and tufa surfaces	32
3.6	The revised LORIS robot design	33
4.1	Images taken of Sally at various locations around an industrial site	36
4.2	An annotated visualization of the Sally robot design	37
4.3	Free-body diagram for an interior corner transition	40
4.4	Minimum tire friction coefficients for interior corner transitions	42
4.5	Minimum tire friction coefficients with assistance from pXRF elevator	43
4.6	Motor torques and velocities during floor-to-wall transition	44
4.7	Key frames taken during interior corner and obstacle transitions	45
4.8	Sensing comparison test locations at the field site	47
5.1	Feasible contact force sets for various end effectors	54
5.2	Optimal contact force distributions for the LORIS robot's typical climbing posture	56
5.3	Optimal contact force distributions for Sally during an interior corner transition	57

6.1	Local planner solutions for LORIS on a slag climbing wall.	69
6.2	Global planner solutions for LORIS in simulated environments	70

List of Tables

2.1	Microspine Model Parameters and Validation	16
2.2	Experimental Comparison of Mechanical Properties	18
3.1	Comparison of LORIS to Prior Work	25
3.2	Climbing Experimental Results	31
6.1	Planning Results on Simulated Terrain	69

Chapter 1

Introduction

Advances in the capabilities of robotic quadrupeds and humanoids have led to a recent explosion of research and commercial interest. In contrast to their wheeled predecessors, these robots offer the potential to operate nearly anywhere a human can go, enabling the automation of a wide range of real-world tasks. However, there remain many environments that even humans struggle to access due to steep terrain features, such as rocky cliffs, cave systems, icy crevasses, or vertical structures. Climbing robots thus provide a solution to explore environments that humans and conventional robots cannot reach, and to automate some of the most difficult and dangerous tasks that still require human labor.

As scientists seek to understand our solar system, the data they can collect is often limited by the capabilities of our current rovers, while climbing robots could open up vast new opportunities for exploration. On Mars, the Opportunity rover observed sedimentary layers in rock outcrops of Victoria crater, but was unable to access them due to the steepness of the incline [1]. Another feature of interest, dark markings called recurring slope lineae that may be caused by briny flows, are only found on 25-40° slopes [2]. Below the surface, lava tubes and caves may provide favorable conditions for the preservation of biosignatures [3]. On the moon, pits formed by collapsed lava tubes expose layers of bedrock stratigraphy, presenting an opportunity to study the evolution of the Moon's secondary crust [4]. On Saturn's moon Enceladus, shafts in the icy crust could grant access to the subsurface ocean to search for signs of life [5]. In microgravity environments, like small moons, asteroids, and comets, locomotion is effectively a form of climbing in the absence of sufficient ground pressure to create traction [1].

Here on Earth, inspection and maintenance tasks often require the construction of costly scaffolding so human workers can access the upper regions of structures. Extensive safety measures are needed to avoid exposing workers to fall risks, and in certain industrial settings additional precautions such as hazmat suits may be necessary to protect workers from hazardous chemicals in the environment. There has been growing interest in climbing robots that can take on these tasks far more cheaply and without putting humans in harms way, with application areas including bridges [6], pipes [7], storage tanks [8], ship hulls [9], and urban buildings [10].

This diverse array of climbing tasks has motivated the development of a similarly wide range




Figure 1.1: From left to right: the limbed rock-climbing robot LORIS, the magnetic-wheeled inspection robot Sally, and the snake-like ice climbing robot EELS 1.5 (courtesy NASA/JPL-Caltech). LORIS and Sally were both developed as part of this thesis.

of different climbing mechanisms (Fig. 1.1). However, climbing robots near universally face a common set of three challenges: mass, adhesion, and autonomy.

A major difficulty of space exploration is that heavy payloads are exceedingly costly to transport to another body. Climbing is an especially high-risk activity, with catastrophic consequences in the event of a fall, and therefore deploying a large climbing robot may pose an unacceptable risk for such an expensive resource investment. In contrast, multiple small, expendable climbing robots could instead be carried alongside a larger conventional rover to be deployed at sites of scientific interest, reducing risk through redundancy and ensuring that their potential failure would not jeopardize the overall mission. Even beyond transportation costs, greater mass requires increased actuator size, power capacity, and structural strength to deal with the larger forces involved, all of which can increase the mass even further. Beyond a certain point, the robot’s weight may surpass the limits of its chosen adhesion mechanism, providing a hard upper bound on the robot’s mass.

This brings us to the second challenge: adhesion. Unable to rely on normal and frictional contact forces alone, climbing robots must generate adhesive forces to support their weight and maintain contact with the climbing surface. On rough rocky surfaces that lack sufficient handholds, this is typically accomplished with microspines – arrays of sharp hooks embedded in flexures that can catch on small asperities (bumps and pits) in the rock [11]. On smooth surfaces like glass, gecko-inspired adhesives can closely conform to the surface to generate adhesive van der Waals forces [12]. On steel structures, magnetic attachment offers a much stronger form of adhesion, enabling the support of larger robots as in [13]. On icy surfaces, robots can drill or cut into the surface using ice screws to anchor themselves [14, 5]. Still other methods include suction and chemical adhesion [15]. Note that the ability to break adhesion when necessary can be equally important, especially for legged robots. In all of these cases, improvements to adhesion can enable increases in payload capacity and reliability.

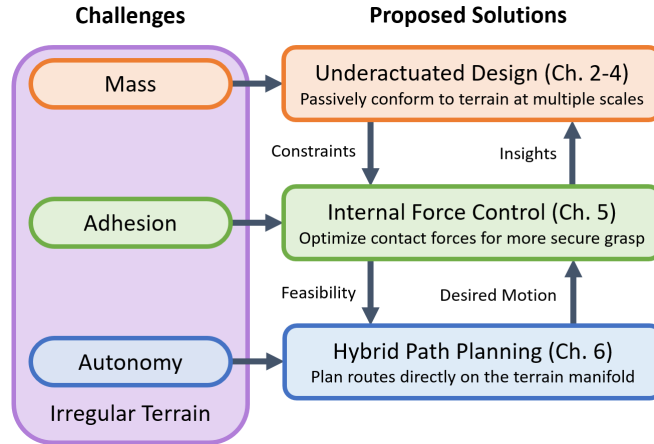


Figure 1.2: A conceptual framework for solving the key challenges of robotic climbing.

The third challenge is autonomy: many potential targets of exploration missions are far enough from Earth to create significant communication delays. For instance, Mars ranges from 4 to 24 light-minutes away while the distance to Enceladus is over a light-hour. These great distances can make human teleoperation inconveniently slow or even entirely infeasible, motivating the need for a partially or fully autonomous system that can respond to quickly changing conditions, such as a gripper slipping, in real-time. For the automation of tasks on Earth, autonomy instead serves the practical purpose of reducing labor requirements, improving task efficiency, and eliminating the risk of human error leading to a fall.

All three of these challenges are intensified in the presence of irregular surface features. While many past robots have been designed to climb flat walls, natural terrain is often highly uneven, and even artificial structures can contain obstacles like beams, rivets, and corners. At large scales these features can prevent the robot’s end effectors from maintaining contact with the surface due to kinematic limits, motivating the use of more complex, high degree-of-freedom systems with increased actuator mass. At smaller scales, these features can complicate the adhesion problem, requiring the generation of adhesive forces in a range of different directions depending on the local orientation of the surface. Autonomy becomes more challenging as well, because the planning space effectively moves from a 2-dimensional plane or 2.5-dimensional height map to a surface in 3 dimensions.

1.1 Thesis Outline

In this dissertation, we propose to overcome the challenges of mass, adhesion, and autonomy on steep, irregular surfaces via a hierarchical approach (Fig. 1.2). We explore this hierarchy primarily through the development of the Lightweight Observation Robot for Irregular Slopes (LORIS), a 3.6 kg quadrupedal rock-climbing robot equipped with microspine grippers to perform in-situ observations on cliff faces or in caves [16]. Along the way we also discuss Sally, a

magnetic-wheeled climbing robot we developed to automate the inspection of steel structures, and the Exobiology Extant Life Surveyor (EELS), a snake-like robot developed by researchers at the NASA Jet Propulsion Laboratory to descend through plume shafts in the icy crust of Enceladus using heated ice screws and search for life in the liquid ocean underneath [5] (Fig. 1.1). This thesis aims to validate three key research hypotheses about robotic climbing:

1. At the design level, underactuation at multiple scales can mitigate the difficulties of irregular terrain while avoiding the mass penalty of additional actuators. In **Chapter 2**, we apply compliant mechanism design at the scale of individual microspines, which rely on flexible suspensions to load-share and conform to surface features. We present a new suspension design that can be 3D printed from a single material, greatly simplifying the manufacturing process compared to previous techniques. This work is adapted from [17]. In **Chapter 3** we apply underactuation at the scale of the microspine gripper and wrist, presenting a fully passive solution that can still solve the paired problems of conforming to uneven terrain and maintaining adhesion over a wide range of loading angles. This work is adapted from [16]. In **Chapter 4** we focus on the full robot scale, discussing Sally’s use of a passive suspension to maintain end effector contact despite surface irregularities and Sally’s repurposing of existing actuators to assist in difficult corner transitions. This work is adapted from [18].
2. At the control level, regulation of the internal forces between different end effectors can improve grasp quality. In nature, both cockroaches and geckos generate opposing lateral forces between their feet while climbing [19], and these distributed inward gripping (DIG) forces can improve adhesion by preloading the directional attachment mechanisms on their feet [20]. In **Chapter 5** we generalize the DIG approach as an internal force optimization problem that can be applied to LORIS, Sally, and EELS despite their varying adhesion mechanisms and morphologies. We also explore the coupling between different layers of our hierarchical framework: the robot design determines the constraints on the optimization problem, and therefore by examining the effect of design parameters on the optimal solution cost, we can gain insights that help to guide key design decisions. In the other direction, desired motions of the body and end effectors produced by the path planner must be integrated with the desired force setpoints in a unified control scheme. This work is drawn from [16], [18], and [21].
3. At the planning level, the terrain can be directly represented as a 2-dimensional manifold embedded in 3-dimensional space, offering a lower-dimensional space in which to search for paths while preserving the higher-dimensional topology. In **Chapter 6**, we present a motion planning approach for hybrid locomotion on such a manifold, represented as a point cloud. We build on the AIT* planning algorithm [22], with a low-fidelity reverse search that computes a cost-to-go heuristic on the terrain manifold to guide the forward search. We again see coupling with the previous layer of the framework: the planning algorithm uses the internal force optimization method to quickly evaluate the feasibility of candidate sets of footholds. This work is adapted from [23].

The result of this thesis is a framework for the design of autonomous, mass-constrained robots to operate on steep, irregular terrain. Using this framework, we have developed LORIS to serve as a proof-of-concept – the first robot of its size to achieve free-climbing on irregular vertical rock faces in full Earth gravity. We also demonstrate the broader value of this approach with the examples of Sally and EELS, two very different robots that nevertheless benefit from the same approach. Finally, in **Chapter 7** we summarize the key takeaways from this solution framework for the development of future climbing robots. We end by discussing possible extensions of this work to conventional quadrupeds and the pursuit of dynamic climbing capabilities.

Chapter 2

Microspine Design for Additive Manufacturing

2.1 Introduction

Humans rely on surface variations at the scale of our hands and feet to provide graspable holds when we rock climb, but the absence of such features on flat regions of a rock face can render a climb difficult to impossible for us. Even so, depending on the type of rock, these flat regions can still feature significant texture at the millimeter to sub-millimeter scale. Thus far we have viewed terrain irregularity as a challenge to overcome, but at this scale it also offers a solution: the small bumps and pits, called asperities, provide potential anchor points where a small, sharp hook called a microspine can catch on the surface. Unlike other adhesion mechanisms such as suction, magnetism, or dry adhesion, microspines are well-suited to the rough rock and cliff surfaces that climbing robots seek to access during planetary exploration [15, 24, 1].

While an individual microspine can only support a small force, a gripper comprised of multiple microspines can support larger forces through load-sharing [11]. The use of multiple microspines in parallel presents two challenges: adapting to local variations in surface height, and distributing the load evenly among the spines. This requires at least two degrees of freedom per spine, normal and tangential to the surface (Fig. 2.1, left), and in an array of a dozen or more spines it quickly becomes infeasible to actuate each axis independently. Instead, microspines rely on the first element of our solution framework from Chapter 1: underactuated design. Compliant suspensions enable microspines to move independently, individually catching on asperities as a gripper is dragged along a surface. These suspensions are typically made using Shape Deposition Manufacturing (SDM), a process in which alternating subtractive machining and casting steps produce a multi-material part containing both rigid and elastic elements, as well as the embedded steel hook [11]. This approach has proven extremely effective, and microspines made with SDM have been used on a wide range of climbing and perching robots [10, 1, 25, 26, 27].

However, the SDM process is time, labor, and equipment intensive due to long curing times




Figure 2.1: A simple gripper composed of nine 3D-printed microspines contained in a carriage. Left: a diagram of the gripper geometry, with the side wall partially cut away. Right: the gripper is shown supporting a 30 N tangential load on a vertical cinder block surface. One can observe that individual microspines have translated relative to each other in order to catch on asperities and distribute the load between them.

and the use of CNC machine tools, with fabrication of new microspine designs taking multiple days of skilled work [26]. Additive manufacturing (AM), or 3D printing, offers potential improvements in all of these areas; parts can be produced quickly and cheaply with minimal human labor. Furthermore, additive manufacturing may enable the use of more durable or long-lasting materials and new designs with more complex geometry. In particular, Fused Deposition Manufacturing (FDM) printers are widely available, used by hobbyists and researchers alike, and offer low part costs and large batch sizes in comparison to other additive methods [28].

In this chapter, we introduce a design model for microspines using two additive manufacturing approaches: Hybrid Deposition Manufacturing (HDM) and Fused Deposition Manufacturing (FDM). We then experimentally compare microspines made via SDM to those made using these new methods. While the HDM method suffers from issues with delamination, we demonstrate that FDM achieves superior mechanical properties to SDM with significant reductions in manufacturing time, labor, and fabrication infrastructure. This makes additive manufacturing a valuable tool for rapid prototyping of microspine designs and increases the broader accessibility of microspine gripper technology.

2.2 Related Work

The SDM process, developed in [29, 30], can form multi-material parts with embedded components through alternating subtractive and additive steps. In [11], this process is applied to create

microspines using hard and soft urethane polymers with embedded fishhooks. The procedure requires several steps. First, a cavity is cut into a machinable wax block using a CNC milling machine. At this point the fishhook is embedded into the part and the cavity is filled with hard urethane resin. After the urethane has cured, new cavities are cut into the part during a second CNC pass. These are filled with flexible urethane resin and left to cure, after which the finished microspine can be removed from the wax block.

The FDM process is an additive manufacturing technique based on filament extrusion [28, 31]. This technique is simple, rapid, and low-cost, while its drawbacks include the appearance of layers, surface roughness, and minimum feature size. New materials with desired material properties like stiffness, strength, flexibility, conductivity are available [32, 31, 33]. Advances in multi-material printing using dual extruders facilitates fabrication of structures with complex geometries and overhangs [28]. This technique is commonly used in mechanisms such as grippers and compliant joints for a robot [34, 32, 35].

The HDM process combines elements of both SDM and FDM to produce multi-material parts [36]. One or more solid bodies are first 3D-printed and assembled together with a 3D printed mold. The cavities are then filled with a resin such as urethane, which is left to cure. Finally, the mold is discarded to leave the finished part. More complex parts can be built up by repetition of this process. A single monolithic printed part can even use sacrificial mold walls that are removed after casting to eliminate the need for a separate mold. HDM offers similar capability to SDM while reducing manufacturing effort and enabling overhanging features that are not possible with SDM alone.

A few alternative approaches to SDM have previously been used for microspine fabrication. Researchers in [37] developed fully metallic microspines that can function in space environments. These microspines use the bending of thin metal ribbons to provide desired stiffness properties. Subsequent work on metallic microspines in [38, 39] instead used a serpentine ribbon geometry to achieve the desired stiffness. In [26], researchers developed rotary microspines that could be fabricated from a single material. They used cast urethane with a higher stiffness than conventional SDM microspines, which was rigid enough to enable embedding of the hook while also flexible enough to enable load-sharing. In [40, 41], researchers embedded microspines at the ends of compliant, laser-cut acrylic and 3D-printed PLA legs. The higher material stiffness is compensated for by the length of the legs, but the materials used cannot provide sufficient compliance when scaled to the smaller form factor of most grippers. Another approach uses laser-cut sheets of soft and hard material that are then glued together along with a spine [42]. Other grippers have avoided the use of flexures entirely: the JPL-Nautilus gripper in [43] instead used springs and tendons for load sharing, and researchers in [44] used an array of spring-loaded, linearly constrained spines.

In this work we focus on multi-material microspines with elastic flexures as in [11] due to their widespread use and advantages such as multi-directional compliance and compact form factor.

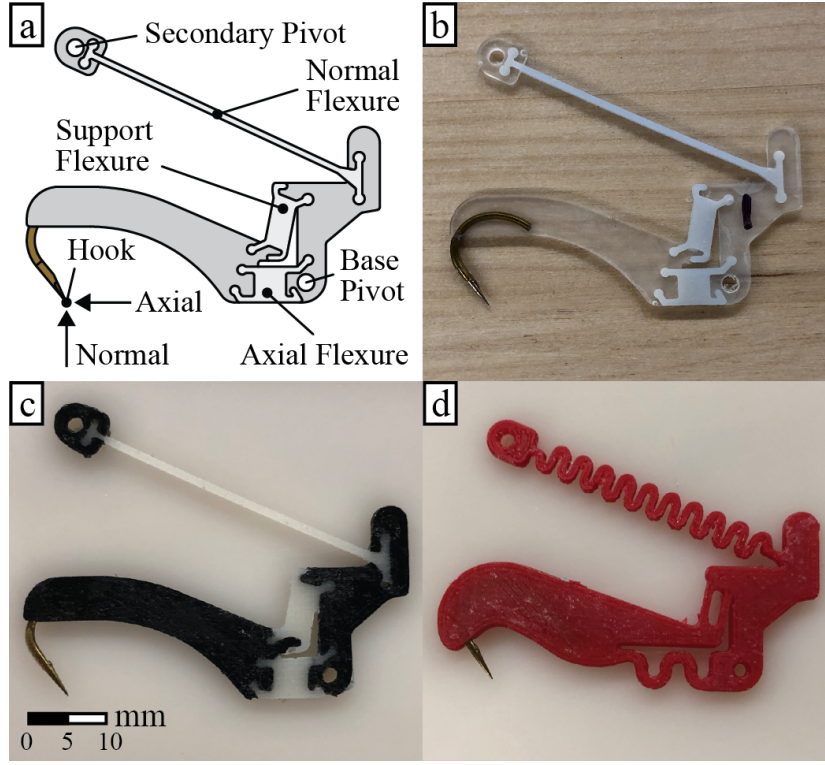


Figure 2.2: a) Diagram of basic microspine geometry. b-d) Images of microspines fabricated using b) Shape Deposition Manufacturing, c) Hybrid Deposition Manufacturing, and d) Fused Deposition Manufacturing techniques.

2.3 Microspine Design Approaches

To perform a controlled comparison among manufacturing approaches, this section considers a typical SDM microspine design (Fig. 2.2b), originally used for the microspine gripper in [25], and details the design adjustments required to apply the HDM and FDM processes.

2.3.1 SDM Microspine Design

The SDM microspines, following the design in [25], are cast from rigid Smooth-On Task 9 and flexible Smooth-On VytaFlex 60 urethane resins, with a Tiemco TMC 300 #2 fishhook for the hook. This SDM microspine design is meant to provide relatively high axial stiffness (parallel to the surface) and low normal stiffness (into the surface), and to minimize rotation of the hook during loading. This is achieved using three separate elastic flexures, labeled in Fig. 2.2a. The axial flexure enables the microspine to stretch axially for load-sharing between microspines as they catch on different asperities. When stretched, the normal flexure applies a torque about the base pivot point that presses the hook against the surface. Lastly, the support flexure provides a vertical constraint that prevents the hook from rotating in plane as the microspine is axially loaded, thereby maintaining a constant contact angle between the hook and the surface [11].




Figure 2.3: The HDM manufacturing process for microspines. a) Diagram of HDM microspine geometry showing the 3D-printed rigid elements (black) and break-away mold walls (red). b) HDM microspine ready for casting. c) HDM microspine after casting and before mold wall removal.

2.3.2 HDM Microspine Design

The standard HDM process can be applied to microspines with minimal modification. In this work, PLA was used for the 3D-printed rigid material, and VytaFlex 60 was used for the cast flexible material. Thin, break-away mold walls were added to the design to contain the urethane during casting (Fig. 2.3a). As in [36], 0.1 mm gaps were left in the design in between the break-away walls and the rest of the part. The gaps are filled in during printing due to the width of the extruder, but the programmed gaps reduce the strength of the connection for easier removal of the mold walls after casting. The 3D printer must be paused prior to the final layer, enabling a fishhook to be press-fit into the exposed cavity as in [40, 41]. Using the PrusaSlicer slicing software it is straightforward to insert a pause into the program. After the print is resumed, the final layer will fully cover the hook.

The finished 3D-printed part is placed on a layer of adhesive tape to prevent leakage during casting (Fig. 2.3b). For best results, the same side of the microspine that was contacting the print bed should contact the tape because of its smoother surface finish. The flexible urethane casting process can then proceed as in the SDM approach (Fig. 2.3c), after which the mold walls can be removed (Fig. 2.2c).

2.3.3 FDM Microspine Design

To further simplify fabrication and eliminate resin casting entirely, here we propose replacing the cast microspines with FDM microspines printed using a flexible thermoplastic polyurethane (TPU) filament. The challenge with this approach is that filaments with comparable Shore hardness to the flexible urethane are uncommon, difficult to work with, and only compatible with 3D printers capable of both multi-material printing and handling very flexible filaments. Instead, we propose the fabrication of microspines using a single material for both rigid bodies and elastic flexures. This approach requires significant changes to the microspine geometry, for which purpose we present a model of microspine stiffness as a function of material and geometric parameters.




Figure 2.4: Examples of failure modes encountered during microspine design. a) Delamination at the material interface of an HDM microspine. b) Insufficient material around hook, c) out-of-plane twisting, and d) hook tearing free of FDM microspines.

One of the key advantages of the single-material approach is that there are fewer interfaces between different materials that can delaminate under load, as seen in Fig. 2.4a. While single-material metallic microspines were designed in [37, 38, 39], the use of a more flexible material enables simpler microspine geometries and less precise manufacturing methods.

For the flexible filament, we selected NinjaFlex Cheetah, a TPU filament with 95A Shore hardness, 26 MPa tensile modulus, and 9 MPa yield strength. The higher stiffness of the Cheetah filament enables the hook to be directly embedded in the flexible material, as in [26]. Using the original SDM microspine geometry, the flexible filament is less effective at constraining the fishhook and it is free to rotate out of plane (Fig. 2.4b). Extra material is added around the fishhook to address this, leaving less of the hook exposed.

Additionally, microspines made of a flexible filament have lower rigidity, which can manifest as out-of-plane twisting of the region between the hook and the axial flexure when the microspine is loaded (Fig. 2.4c). To compensate, a greater length of the fishhook shaft is embedded in the part. Load applied to the hook is transferred along the shaft of the hook to the axial flexure, so that no twisting moment can occur. Due to the shape of the fishhook, this design change requires reducing the size of the cutout between the hook and base pivot, which was introduced in [25] to reduce interference with surface irregularities; however, this cutout was not present in prior microspine designs, so the larger size is not considered essential.

The cavity in which the hook is embedded must also be more carefully sized than in the case of HDM microspines; due to the softer material, any extra clearance around the fishhook may

enable it to twist out of plane and tear through the side of the microspine (Fig. 2.4d).

Most significantly, the spring stiffness of all three flexures increases dramatically due to the increased tensile modulus of the material. This increase can be counteracted by changing the flexure geometry. In [45], the authors demonstrate that a serpentine ribbon geometry can reduce stiffness by multiple orders of magnitude. In our FDM microspine design, both the axial and normal flexures use serpentine geometries to provide a lower stiffness than the TPU material properties could achieve in a rectilinear geometry. The relatively wide axial and supporting flexures of the original design each fully constrain translation and rotation in the flexure plane, while a serpentine flexure provides only a translation constraint along the flexure axis. Making both flexures serpentine would therefore leave one degree of freedom unconstrained. Instead, the axial tendon is made serpentine, while the supporting flexure is replaced with two narrow parallel flexures. The parallel flexures act as a 4-bar linkage to constrain both translation along their axis and rotation, so that in combination with the axial flexure the system remains fully constrained as in the original design. For a more general approach to flexure design, the Freedom and Constraint Topology (FACT) approach [46] provides a comprehensive mapping between desired motions and constraints.

We derive a linear elastic stiffness model for the two serpentine flexures and the support flexure in order to match the stiffness of the original SDM design at small displacements. Though nonlinear constitutive laws for viscoelastic materials [47] would better generalize to large displacements, these laws would require the measurement of additional material parameters and thus increase the difficulty of applying the model to new materials. Similarly, other approaches such as the Beam Constraint Model (BCM) [48] can account for elastokinematic nonlinearities at larger flexure displacements, but a linear approximation is sufficient for a relative comparison. For the axial and normal flexures, we use the serpentine ribbon model from [45] to select the flexure geometry. They define the ribbon geometry by the flexure width w , arc radius R , arm length l , and arc opening angle α , as shown in Fig. 2.5a, from which the stiffness of a given flexure design can be computed. According to their model, the reduction in stiffness is scale invariant and monotonic with smaller w/R , larger l/R , and larger α . We can use their model to compute the stiffness of the axial flexure (k_a) and normal flexure (k_n),

$$k_a = \frac{Ew_a h}{L_a} S \left(\frac{w_a}{R_a}, \frac{l_a}{R_a}, \alpha_a \right) \quad (2.1)$$

$$k_n = \frac{Ew_n h}{L_n} S \left(\frac{w_n}{R_n}, \frac{l_n}{R_n}, \alpha_n \right) \quad (2.2)$$

where subscripts a and n denote parameters of the axial and normal flexure respectively, L is the flexure length, h is the flexure thickness, E is the material tensile modulus, and S is the relative effective stiffness of a serpentine ribbon as derived in [45].

For the support flexure, we derive the stiffness of a flexure loaded at an angle relative to its longitudinal axis. Here, we consider a support flexure of length L_s and width w_s , rotated θ degrees from the horizontal x axis (Fig. 2.5b). With one end of the flexure fixed and the other rotationally constrained but free to translate, we apply a horizontal load F_x at the free end,

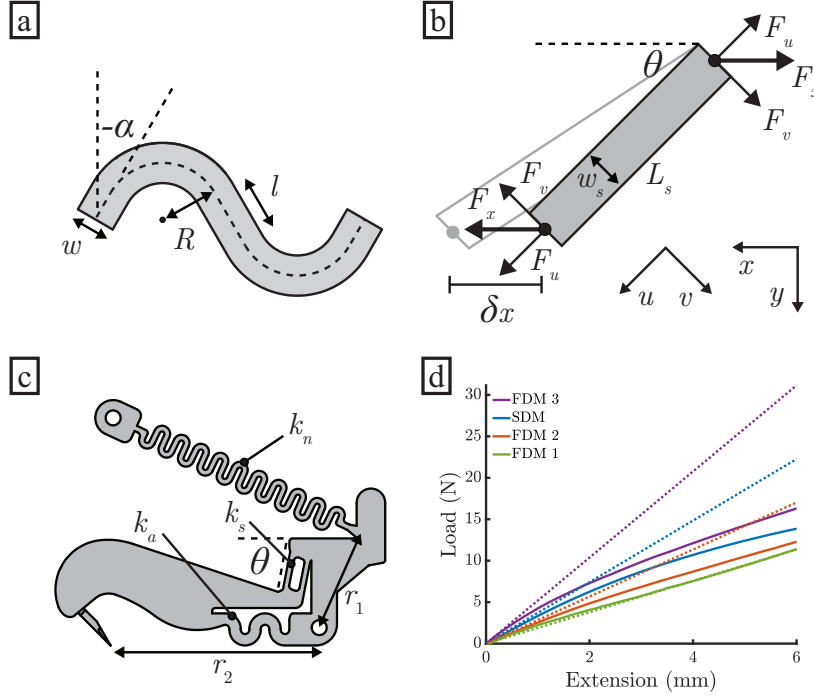


Figure 2.5: a) Geometric parameters for a serpentine ribbon flexure, as defined in [45]. b) Diagram of an individual flexure subject to a tensile load F_u and bending load F_v caused by a lateral displacement δx . c) Labeled stiffness contributions of each flexure, as well as moment arms r_1 and r_2 and support flexure angle θ . d) Load-displacement curves (solid) for the SDM microspine design and three variations of the FDM microspine design, overlaid against linear model predictions (dotted). Legend entries are arranged in the same order as the curves. The SDM and FDM 2 results are averaged over 5 samples, while the FDM 1 and FDM 3 results are each based on a single trial.

such that the flexure experiences a horizontal deflection δx along the x axis. The load can be decomposed into a tensile force along the flexure's length due to axial elongation, F_u , and a shearing force perpendicular to the flexure due to bending, F_v , given by,

$$F_u = \frac{E w_s h}{L_s} \delta x \cos \theta \quad (2.3)$$

$$F_v = \frac{12 E I_s}{L_s^3} \delta x \sin \theta \quad (2.4)$$

where $I_s = \frac{w_s^3 h}{12}$ is the area moment of inertia.

We can then superimpose and project the resulting forces onto the axis of applied load to get the combined applied load F_x and compute the spring stiffness k_s for the support flexure

along the axial direction of the microspine,

$$F_x = F_u \cos(\theta) + F_v \sin(\theta) \quad (2.5)$$

$$k_s = \frac{F_x}{\delta x} = \frac{E w_s h}{L_s} \cos^2 \theta + \frac{12 E I_s}{L_s^3} \sin^2 \theta \quad (2.6)$$

Applying this model to a microspine (Fig. 2.5c), the axial stiffness is determined by both the serpentine axial flexure and the support flexure, and the normal stiffness is determined by the normal flexure adjusted for the lever arm of the applied load. The overall microspine stiffness is then given by,

$$k_{axial} = k_a + n_s k_s \quad (2.7)$$

$$k_{normal} = k_n \left(\frac{r_1}{r_2} \right)^2 \quad (2.8)$$

where n_s is the number of support flexures and r_1 and r_2 are the moment arms about the base pivot of the normal flexure and the normal load, respectively (Fig. 2.5c). Normal compliance of the support and axial tendons is negligible compared to the compliance of the normal tendon, and thus not included for simplicity.

To validate this model, we produced three microspine designs with differing axial flexure geometry. The geometries were chosen to vary the predicted microspine axial stiffness while meeting space constraints. These three designs, along with a standard SDM microspine, were fabricated and axially loaded to measure stiffness as in Section 2.4. Geometric parameters and tensile modulus for each microspine design are given in Table 2.1. The tensile modulus of the urethane resin was not published by the manufacturer, so the 100% modulus was used as a substitute. For the FDM design we treat the support flexure as two thinner individual flexures. When entering dimensions into the model, it is important to account for changes in the part size during printing; we found that the width of each flexure increased by 0.3 mm from the original design, and the height of the part increased by 0.1 mm.

Axial load-displacement curves for each microspine are compared to linear model predictions in Fig. 2.5d, showing agreement at small displacements to within 20%. Table 2.1 provides the percent error between the predicted stiffness k_a and the measured stiffness \hat{k}_a over the first 0.5 mm of extension. While the assumptions of this linear model do not apply to large displacements, we see that the relative stiffness ordering of the four designs remains unchanged and all exhibit similar nonlinearity. Thus by matching the predicted small-displacement stiffness designers can still generate designs that produce similar forces at large displacements. The predicted normal stiffness values were less accurate than the axial predictions. For the normal flexures in particular, small amounts of preload in the test setup could have had a large effect on the measured stiffness values.

Despite these limitations, the model is effective at comparing relative axial stiffness between different SDM and FDM designs despite very different material properties and flexure geometries. Using this model enables designers of new FDM microspines to optimize their suspension

Table 2.1: Microspine Model Parameters and Validation

Parameter	SDM	FDM 1	FDM 2	FDM 3
E , MPa	2	26	26	26
h , mm	1.5	1.6	1.6	1.6
θ	79°	79°	79°	79°
r_1 , mm	14.5	14.5	14.5	14.5
r_2 , mm	30.0	30.0	30.0	30.0
L_s , mm	6.5	4.3	4.3	4.3
w_s , mm	3.0	0.9	0.9	0.9
n_s	1	2	2	2
L_a , mm	4.0	9.0	8.4	8.4
w_a , mm	4.5	1.5	1.9	2.3
R_a , mm	-	1.5	1.4	1.4
l_a , mm	-	0.3	0.1	0.5
α_a	-	0°	0°	-15°
S_a	1	0.076	0.156	0.334
L_n , mm	33.0	30.8	30.8	30.8
w_n , mm	0.8	1.1	1.1	1.1
R_n , mm	-	0.8	0.8	0.8
l_n , mm	-	1.7	1.7	1.7
α_n	-	0°	0°	0°
S_n	1	0.030	0.030	0.030
k_{axial} , N/mm	3.71	1.89	2.84	5.19
\hat{k}_{axial} , N/mm	3.35	2.28	2.52	4.64
Error, %	20.4	-11.2	-10.7	-9.7

geometry appropriately given an expected change in tensile modulus to maintain the desired microspine compliance.

2.4 Experimental Characterization

Relative performance of the three manufacturing methods was evaluated through tests simulating the loading experienced by a microspine during use in a gripper. Mechanical properties of each microspine type were evaluated using an Universal Testing machine (Instron 5969, 1 kN load cell). These included the microspine stiffness in the axial direction, the stiffness in the normal direction, the maximum axial load before the microspine suffers permanent damage, the observed failure mechanism, and the effect of fatigue measured by the decrease in required axial force to cause the same amount of displacement. For the microspine stiffness characterization (both axial and normal direction) the loading velocity was 5 mm/min and for the fatigue test, 100 cycles were performed with a 5 mm/min loading and unloading cycle. Five microspines of each type were tested for axial and tangential stiffness and axially loaded to failure. The FDM microspines that were measured correspond to the FDM 2 design above. In addition, one SDM and one FDM microspine were each subjected to cyclic axial loading. Only two of the five HDM




Figure 2.6: Experimental setup for measurement of microspine axial stiffness (left) and normal stiffness (right). The top diagrams show a section view of each assembly, where the upper fixture is in light blue and the lower fixture is in dark blue. The bottom images show the test assemblies with the microspine unloaded.

microspines were able to reach a displacement of 5 mm before failure, so an HDM microspine was not included in this test.

During load testing, microspines were constrained as they would be in practice when incorporated into a gripper. Separate fixtures were used to support microspines during axial and normal loading, as shown in Fig. 2.6. In each case, the microspine was translationally constrained relative to the lower fixture at the two pivots. A socket head bolt in the upper fixture was used to securely engage the hook tip.

The measured mechanical properties of microspines manufactured with each method are summarized in Table 2.2. The stiffness measurements for both axes were based on measured load at 6 mm of displacement (leaving out any samples that failed early). The axial load-displacement curves for each microspine type are shown in Fig. 2.7, left, averaged across the five samples. A Butterworth filter was used to smooth out the data during post-processing. We see very close agreement between all three microspine types, with HDM microspines 15.5% stiffer than SDM and FDM microspines 11.3% less stiff than SDM. The agreement between the HDM and SDM microspines is to be expected because both microspines use the same flexure geometry and flexible material. However, the agreement between the SDM and FDM microspine results from deliberate tuning of the FDM microspine’s serpentine flexure geometry.

The FDM microspines also exhibit a more linear load-displacement curve, possibly due to the serpentine flexure geometry or the material properties of the filament. The nonlinearity of the SDM curve is slightly beneficial for load-sharing at larger displacements, because differences in displacement between microspines have less effect on their relative forces, but the effect is small enough to be negligible.

Table 2.2: Experimental Comparison of Mechanical Properties

Measurement	SDM	HDM	FDM
Axial stiffness, N/mm	2.31 ± 0.29	2.67 ± 0.43	2.05 ± 0.06
Normal stiffness, N/mm	0.028 ± 0.004	0.033 ± 0.005	0.022 ± 0.005
Maximum load, N	$15.9 \pm 2.1^*$	10.9 ± 8.0	41.1 ± 2.1
Force decrease, 100 cycles, N	2.36 (17.4%)	-	1.30 (12.4%)

*Maximum load for SDM microspines is typically around 30 N.

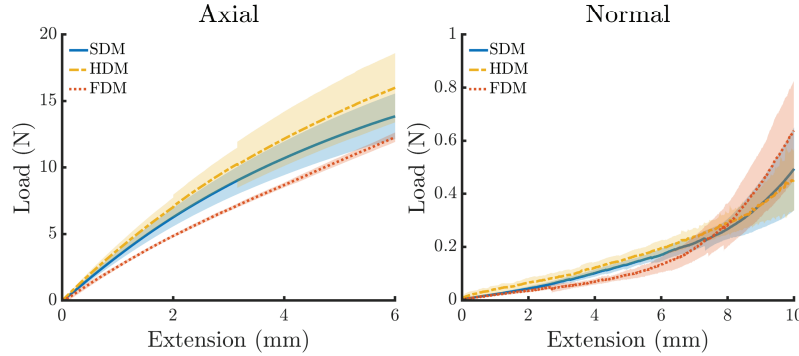


Figure 2.7: Overlaid axial (left) and normal (right) load-displacement curves for each of the measured microspine designs. Shaded regions indicate one standard deviation. Samples are removed from the averaging after a failure occurs, causing small discontinuities in the HDM axial load. Small discontinuities in the FDM normal load plot were caused by slight shifting of the hook relative to the fixture.

Similarly, Fig. 2.7, right, shows agreement between the normal stiffness of all three microspine types for most of the measured range, with HDM microspines 14.9% stiffer than SDM and FDM microspines 21.4% less stiff than SDM. At very large displacements, the normal stiffness of the FDM microspines begins to sharply increase, as the microspine collides with the normal flexure above. While this also occurs for SDM and HDM microspines, the FDM flexures have significantly higher bending stiffness and therefore this interference has a much greater effect. As this interference only occurs near the limit of travel the actual impact on performance is minor in most circumstances.

With further tuning of the serpentine flexure geometry, both the average axial and normal stiffness of the FDM microspines could be made to match the SDM microspines arbitrarily well. In practice, the optimal microspine stiffness is dependent on the distribution of surface asperities, terrain roughness, expected load, and number of microspines in the gripper, so further optimization of the microspine stiffness should be done with a specific application in mind.

In Fig. 2.8, results are shown for each sample being axially loaded to failure. The SDM samples failed at 13-19 N, due to either delamination (3 samples) or a break in the rigid material (2 samples). We later obtained original SDM microspine samples from the authors of [25] and determined that their failure load was significantly higher at around 30 N, although their stiffness properties were identical. This suggests that a defect occurred during our manufacturing process,




Figure 2.8: Stress-strain curves of individual SDM* (left), HDM (center), and FDM (right) samples when axially loaded until failure. One HDM sample evidently delaminated during the mold wall removal process prior to the experiment, and so was not included in the plot. *The maximum load for SDM microspines is typically around 30 N.

and highlights the advantage of easier and more consistent fabrication using the FDM process compared to SDM. HDM microspine performance is much less reliable, with delamination at the urethane-PLA interface occurring anywhere from 0-19 N. We attribute this to poor adhesion between the flexible urethane and PLA materials relative to the adhesion between the flexible and rigid urethane resins. The FDM samples are consistently able to reach around 40 N of applied load, at which point angular deflection of the microspine causes the hook to slip out of the bolt head. The microspine is left intact, so the maximum load before material failure is even greater. This superior performance is due to the use of a single material, which avoids the failure mode of delamination. Additionally, the Cheetah material has a much higher tensile strength than the flexible urethane (39 MPa vs 6 MPa respectively) which enables FDM microspines to support larger loads before rupture.

Fatigue testing results are shown in Fig. 2.9. For each microspine, the peak load reached during each cycle is plotted against the number of elapsed cycles. The peak load measurements are normalized by the peak load of the first cycle for that microspine. The SDM microspine exhibits a 17.4% reduction in force over 100 cycles, compared to a 12.4% reduction for the FDM microspine. These results suggest that FDM microspine performance degrades more gradually with use than that of SDM microspines. For reference, wear of the spine tip is known to occur over a comparable time scale (the tip radius levels off after around 50 cycles) [25].

As a final validation, nine FDM microspines were assembled into a simple gripper (Fig 2.1, right), which was placed against a cinder block surface and tangentially loaded until disengagement. The gripper was able to consistently support a 30 N tangential load before adhesion failure, with the individual microspines visibly displacing to catch on asperities as expected.

2.5 Discussion

Overall, these results demonstrate that FDM microspines can achieve comparable stiffness properties to SDM microspines while providing increased durability in the form of a higher maximum

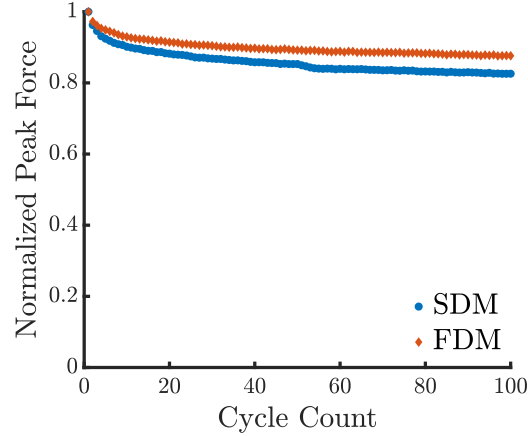


Figure 2.9: Maximum force reached during each load cycle normalized by the force at the first loading cycle, plotted over 100 load cycles with a 5 mm displacement for both SDM and FDM microspines.

load, better maintenance of mechanical properties after cyclic loading, and fewer failure modes. We conclude that microspines fabricated using this new approach are able to match the adhesive performance of existing SDM microspines, as also demonstrated by handheld gripper tests.

FDM microspines offer a number of advantages beyond their mechanical properties. Foremost is the reduction in production time. As a rough estimate, manufacturing a batch of 25 microspines through SDM requires 2-3 hours of machining time, with a similar amount required for setup and cleanup. The hook preparation, casting, and mold removal steps combined required an additional 1-2 hours, but the bulk of the time is the curing process, which requires 1 hour for the rigid polymer and 16 hours for the flexible polymer. The total comes out to around 24 hours or 1 hour per microspine, of which 4 hours or 10 minutes per microspine actively requires human labor. In contrast, a batch of 9 FDM microspines requires 2 hours of printing (on an Original Prusa i3 MK3S+ printer) and 10 minutes of hook preparation, for an average of 15 minutes per microspine, of which only 1 minute of human labor is required. This difference becomes even more drastic for small batches – the machining setup and curing times for SDM are unchanged, so manufacturing a single microspine still requires 2-3 hours of human effort and an overnight casting process. In contrast, manufacturing a single FDM microspine can be done in 15 minutes.

The SDM process also requires users to learn a complicated casting process with many opportunities for error, while setting up the FDM process can be as simple as downloading a G-code file. Other advantages include the shelf life of the materials; once opened, the urethane resins begin to harden and lose efficacy, while the flexible filament is much more stable. Lastly, the SDM process requires access to a CNC mill, fume hood, and vacuum chamber, while the FDM process uses only a basic 3D printer.

More broadly, we have seen that underactuation is fundamental for microspines, enabling them to passively conform to surface irregularities on the scale of millimeters so that they can

be affordably grouped into larger arrays. However, large enough surface variations can exceed the range of motion of the individual spines. In Chapter 3 we examine how hierarchical levels of compliance in the gripper can overcome the challenge of surface irregularity at this larger scale.

Chapter 3

Design of a Lightweight Rock-Climbing Robot

3.1 Introduction

Rock-climbing robots have the potential to explore steep cliff faces, caves, and microgravity environments of scientific interest throughout the solar system [49, 50, 3, 4]. While wheeled rovers are typically limited to shallow slopes, robots equipped with microspine grippers (summarized in Table 3.1) are capable of ascending steep, rough rock faces using arrays of sharp hooks that catch on small asperities in the surface [11, 51]. The earliest microspine grippers were fully passive, relying on the weight of the robot to provide the necessary engagement force, and robots equipped with these grippers were designed to traverse relatively flat surfaces such as the walls of buildings [10, 26, 44].

In contrast, the microspine gripper in [24] used an actuator to generate large, opposing engagement forces between its microspines, enabling the LEMUR series robots [24, 1] to support their greater weight and traverse unstructured natural terrain features. Other researchers have followed a similar approach, creating active grippers that operate on the same principle, but with increased engagement speed and reduced size and mass [52, 53]. Some grippers can close around protruding footholds to provide a robust grasp, but are less effective on a flat surface [54, 55, 56].

In [20], the authors present an insect-inspired strategy called distributed inward gripping (DIG), in which opposed engagement forces are generated between the legs of a climber rather than by each gripper individually. This enables passive grippers to obtain many of the same benefits as active grippers, but without the additional mass and complexity. However, the robot in [20] uses a single hook per foot, limiting it to climbing on a mesh screen.

Despite the demonstrated successes of these robots, climbing remains an inherently high-risk operation, and the premature demise of a primary mission asset during a space mission would be catastrophic. Instead, one or more small, expendable climbing robots could be carried alongside



Figure 3.1: The LORIS robot uses passive microspine grippers and distributed inward gripping to ascend the side of a tufa stone bridge.

a larger conventional rover, to be deployed at sites of scientific interest without endangering the primary mission objectives. Such a system would reduce overall risk, enable division of labor between multiple vehicles, and be more likely to be deployed as the first climbing robot in space. However, this requires a robot that is both lightweight and capable of climbing on unstructured terrain.

In this chapter, we present a novel passive gripper design capable of supporting varying lateral loads along with a robot prototype, the Lightweight Observation Robot for Irregular Slopes (LORIS), that can ascend unstructured steep and vertical terrain (Fig. 3.1) by generating DIG forces between grippers to improve adhesion. We provide climbing experiments on irregular slag and vesicular basalt climbing walls, flat cinder block, and natural rock features. LORIS is to our knowledge the first robot to generate DIG forces using microspines and at 3.2 kg the first robot of its size to demonstrate free-climbing on irregular vertical rock faces in full Earth gravity.

3.2 Robot Design

LORIS is a quadrupedal robot with three joint actuators per limb for the shoulder, “wing” (abduction/adduction), and knee. While wall-climbing robots like RiSE only required two actuators per limb to climb on flat surfaces [10], the third limb joint provides the robot more freedom to select footholds as needed on irregular terrain. At the end of each limb is a passive microspine gripper on a passive 3-degree of freedom (DoF) wrist. This design is analogous to ground quadrupeds that have 3-DoF legs with 3 passive DoF at their point feet, enabling ar-

Robot	Terrain	Active DoF (per limb)	Mass (kg)	Speed (m/min)	Gravity (g)
RiSE [10]	Flat Walls	2	3.8	15	1
HubRobo [54]	Handholds	3	3.0	0.17	0.38
SCALER [55]	Handholds	6	6.3	0.35	1
LEMUR 2B [24]	Irregular	3	10	Unknown	1
LEMUR 3 [1]	Irregular	7	35	0.0027	0.38
RockClimbo [52]	Irregular	4	3.5	0.14	0.67
LORIS	Irregular	3	3.2	0.20	1

Table 3.1: Comparison to Prior Work

bitrary positioning and orienting of the body. Compliant degrees of freedom within both the gripper and the wrist are designed to provide “mechanical intelligence”, the ability to naturally conform to uneven terrain without any actuation or control. A pitch spine actuator in the body aids in traversing larger obstacles and plane transitions, as in [57, 58], and a tail actuator is used to push against the surface with a large moment arm, reducing the adhesive force needed by the front grippers when climbing.

3.2.1 Splayed Microspine Gripper

Microspines rely on a steep angle of attack relative to the contact surface, typically around 70° [26]. Due to this directionality, they are not effective at individually resisting lateral loads. While climbing flat walls rarely requires variation in the load direction, more complicated terrain requires grippers that can support loads over a wider range of directions without repositioning. Active grippers avoid this limitation by placing microspines in a ring, so that a subset of them are always oriented to oppose the external load [24, 54]. However, this approach inherently requires some form of actuation to engage or release the grasp.

Instead, we propose a passive gripper with a single pair of splayed carriages angled $\pm 45^\circ$ to the gripper axis (Fig. 3.2a). These carriages in combination can apply loads with tangential ($+x$) and lateral ($\pm z$) components as force is distributed into one component along each carriage’s axis. The gripper is fully passive and can be disengaged by unloading tangentially, as this direction of motion is not opposed by any microspines. Unlike an active gripper it cannot support negative tangential or purely lateral loads, but in practice the need to resist the pull of gravity or generated DIG forces will create a significant positive tangential component.

There are many possible splayed gripper designs, depending on the number and splay angle of the carriages – at one extreme, the individual microspines could be radially spaced in a half circle. To find the optimum in this design space, we model the in-plane forces each carriage experiences during static loading (Fig. 3.2b). The n microspine carriages are assumed to act as linear springs with stiffnesses k_i oriented at angles α_i . If the gripper translates by a small displacement δu along the angle β , then the resulting tangential carriage force f_i^t and net in-plane

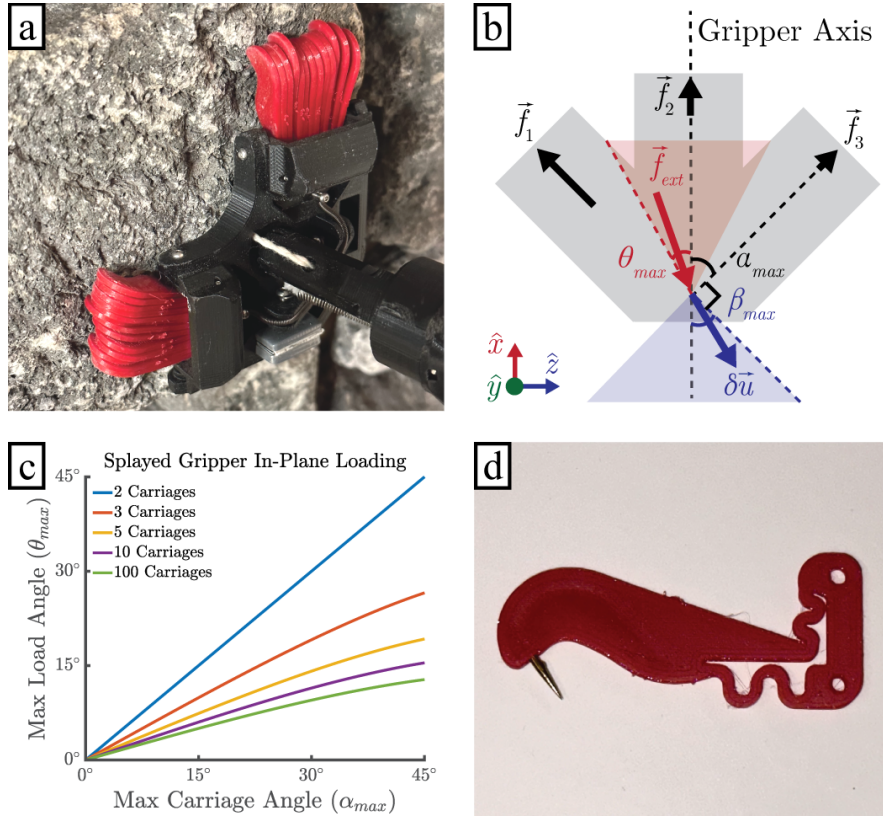


Figure 3.2: a) A splayed microspine gripper engaged with a rock. b) A free-body diagram of a 3-carriage splayed gripper, showing viable displacements in blue and viable loading in red. Angles are defined relative to the gripper axis. c) A plot of the maximum loading angle for varying gripper carriage geometry. d) A 3D-printed microspine with an embedded fishhook and serpentine flexures.

external force \vec{f}_{ext}^{xz} are given by,

$$f_i^t = k_i \cos(\beta - \alpha_i) \delta u \quad (3.1)$$

$$\vec{f}_{ext}^{xz} = - \sum_{i=1}^n f_i^t \begin{bmatrix} \cos(\alpha_i) \\ \sin(\alpha_i) \end{bmatrix} \quad (3.2)$$

If any carriage experiences a displacement in the opposite direction of its microspines, then that carriage loses engagement and the grasp is considered to have failed. This occurs when $\cos(\beta - \alpha_i) < 0$, thus β is bounded by $[\alpha_{max} - 90^\circ, -\alpha_{max} + 90^\circ]$. Computing θ , the angle of \vec{f}_{ext} relative to the gripper axis, at these bounds produces the minimum and maximum loading angles $-\theta_{max}$ and θ_{max} that the gripper can support.

In Fig. 3.2c, the maximum loading angle θ_{max} is plotted for symmetric grippers with different numbers of evenly spaced, equal stiffness carriages as a function of carriage angle α_{max} . Two carriages outperform larger numbers because additional intermediary carriages increase the relative u component of the load without affecting the bounds of β . At α_{max} values greater

than approximately 45° the angle of the outermost spines becomes too steep to effectively catch on asperities during initial engagement.

3.2.2 Passive Wrist Joint

Full control over the gripper's orientation in space would require three additional actuators per limb, but significant mass reduction is possible by letting those degrees of freedom remain passive. We define rotations with pitch about the z axis (i.e. rotating up-down), roll about the x axis (left-right), and yaw about the y axis (counterclockwise-clockwise). To remove the dependence of the limb kinematics on the potentially unknown gripper orientation, we let all three axes of rotation intersect at a common point. Mechanically, $\pm 35^\circ$ of pitch and roll are provided by a ball joint while the yaw degree of freedom is supplemented by a low-friction 360° revolute joint.

The gripper must make contact at 3 or more points to passively align itself with the climbing surface in the pitch and roll axes. For the splayed gripper design above, these are ideally the two sets of microspines and a contact point at the base of the gripper. On highly convex features this may not be the case, potentially causing a grasp failure. Furthermore, when these axes are passive the gripper performance becomes sensitive to the location of the center of rotation.

We model the gripper in terms of the width between the two carriages w , the height between the microspines and the base contact h , the splay angle α , and the center of rotation \vec{p} (Fig. 3.3). Assuming static equilibrium and a frictionless base contact, the system equations,

$$0 = f_{ext}^x + (f_1^t + f_2^t) \cos(\alpha) \quad (3.3)$$

$$0 = f_{ext}^z + (f_2^t - f_1^t) \sin(\alpha) \quad (3.4)$$

$$0 = f_{ext}^y + f_1^n + f_2^n + f_b^y \quad (3.5)$$

$$0 = \tau_{ext}^x + \frac{w}{2}(f_1^n - f_2^n) + p^y f_{ext}^z \quad (3.6)$$

$$0 = \tau_{ext}^z + h(f_1^n + f_2^n) + p^x f_{ext}^y - p^y f_{ext}^x \quad (3.7)$$

can be solved analytically for the forces at the microspines \vec{f}_1 , \vec{f}_2 and the base contact f_b^y in terms of the gripper geometry and the external force \vec{f}_{ext} and torque $\vec{\tau}_{ext}$.

The resulting forces should obey two constraints to prevent the gripper from disengaging: 1) the pull-off force angle on the microspines cannot exceed a threshold ϕ_{slip} , as found by [59], and 2) the required normal force on the base contact cannot be negative (or else the gripper pitches downward, causing the spine angle of attack to become too steep),

$$-f_1^n \leq f_1^t \tan(\phi_{slip}) \quad (3.8)$$

$$-f_2^n \leq f_2^t \tan(\phi_{slip}) \quad (3.9)$$

$$0 \leq f_b^y \quad (3.10)$$

Focusing on the pitch axis and assuming no torque at the joint, let $f_{ext}^z = 0$, $\tau_{ext}^x = 0$, and

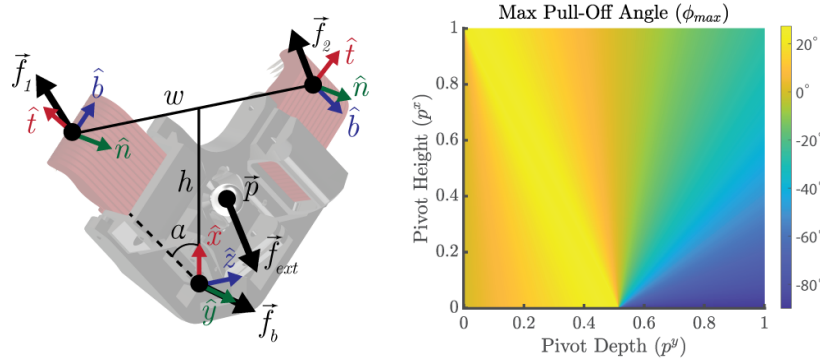


Figure 3.3: Left: an annotated image of a splayed gripper, with the three contact points and the pivot ball joint indicated by circles. Right: a plot of the maximum pull-off force angle relative to the wall as function of the pivot joint location for a 2-carriage gripper with a $\pm 45^\circ$ splay angle. Pivot height and depth are normalized by the gripper height h .

$\tau_{ext}^z = 0$, then solve the constraints numerically to compute the maximum feasible gripper pull-off force angle ϕ_{max} given the pivot location. That maximal angle is plotted in Fig. 3.3, which shows that the optimal pivot location lies roughly along a line. Intuitively, moving the pivot further from the surface results in greater spine pull-off force, while moving the pivot toward the surface causes the gripper to begin rotating at a smaller loading angle. Thus the pivot point should be in line with the spine contact force vector at the gripper's maximum loading angle to minimize both effects. However, due to space constraints the pivot could not be placed as close to the surface as optimal, and so a pair of springs were added to provide a negative preload torque τ_{ext}^z . This shifts the optimum away from the surface, although the exact size of the shift depends on the applied force magnitude.

While the gripper naturally flattens itself against the surface in pitch and roll, the gripper yaw depends on its angle at the moment contact is made. We would like to align the gripper with the nominal in-plane loading direction, which depends on the external gravity vector and any generated internal forces. If the DIG force magnitude is nominally proportional to the gravitational load on each gripper, then the optimal gripper angle is a fixed angular offset from the gravity vector's projection into the surface plane. We can therefore distribute the mass of the gripper such that its minimum-energy yaw angle is equal to this offset, similar to the grippers in [60] (which can be considered to have an offset of zero). A limitation of this approach is the requirement of a significant in-plane gravity component, and thus this strategy does not generalize to fully inverted climbing or micro-gravity environments.

3.2.3 Robot Prototype

The LORIS robot, Fig. 3.4, is designed for ease of manufacturing – all parts are either 3D printed, waterjet from aluminum, or commercially available. The robot body consists of two 2 mm-thick aluminum baseplates separated by standoffs. The frame is 315 mm long by 180 mm wide by 44 mm deep and separated into two equal halves, connected by the body joint. The




Figure 3.4: An annotated visualization of the LORIS robot design, with a splayed gripper enlarged in an inset.

first limb motors project outward from the body to ensure a full range of motion. The limbs are composed of off-the-shelf plastic and custom PLA brackets. The 348 mm tail consists of a carbon fiber tube encased in a rounded PLA shell to easily slide over uneven rocks.

The actuators used are Dynamixel AX18-A servo motors for the wing degree of freedom and tail, and higher torque Dynamixel XM430-W350-T servos for the shoulders, knees, and body joint. LORIS can either be powered by an offboard supply or carry a 3S LiPo battery. The robot can also carry an Intel NUC 11 for onboard computation. An Intel RealSense D435i depth camera mounted on a carbon fiber mast enables the robot to construct a map of the upcoming terrain.

The passive grippers are composed of 3D printed PLA. 26 microspines are evenly divided among two carriages forming a 90° angle. The carriages are free to independently rotate toward the surface by up to 30° . This hierarchical compliance aids the gripper in conforming to convex features at a larger scale than an individual spine. Low-stiffness restoring springs ensure that the carriages are pitched downward while making initial contact, but can easily be pressed back when placed on a flatter surface.

The microspines are 3D printed in TPU using the additive manufacturing process from Chapter 2 with embedded fishhooks (Fig. 3.2d). The primary improvement in the design is the removal of the long flexure that rotates the hook into the surface. Instead, a serpentine diagonal flexure is sufficient to provide a similar level of normal compliance, enabling a significant reduction in footprint.

3.3 Climbing Gait

LORIS relies on unified force control strategy to generate optimal DIG forces while simultaneously executing desired motions of the end effectors – this controller is covered in detail in Chapter 5. For now we limit our discussion to the higher-level climbing gait. The gait is defined by a simple state machine and relies on proprioception to detect gripper contact and engagement with the surface, providing robustness to uneven terrain in the absence of external sensing.

1. **Unload** – constrain the contact force on the gripper \vec{f}_i to be zero in the force optimization and wait for the contact force to converge to the new setpoint.
2. **Disengage** – free the microspines from the surface and raise the gripper until it reaches the maximum height above the surface.
3. **Swing** – translate the gripper to the next foothold provided by the path planner or teleoperator.
4. **Engage** – lower the gripper toward the surface until a force threshold is reached, then translate gripper tangentially until a second force threshold is reached.
5. **Load** – remove the constraint on the gripper contact force and wait for convergence.

The disengage state is of particular importance, especially on surfaces with large vesicles (cavities formed by trapped gas during cooling). In this state, the robot attempts to disengage the gripper carriages by translating in the opposite direction to the gripper axis. However, often one carriage comes free before the other, resulting in a lateral load that causes the gripper to yaw. This is addressed by translating the gripper in the direction of the lateral load to reduce it. The second issue is that the grippers sometimes struggle to free their microspines because they are unable to apply a torque with the wrist, so the microspines simply rotate as the gripper tries to pull away. A solution is to lower the gripper back toward the surface whenever a large force into the surface is detected. This reverses the rotation, giving the gripper another chance to work its way free. Additionally, LORIS can react to a gripper snag (unexpected force during swing) or engagement failure (excessive displacement during engagement) by returning to the disengage state, which results in a pawing motion.

3.4 Experimental Performance

The vertical climbing performance of the LORIS robot was evaluated on several different substrates, including cinder block, vesicular basalt, and slag (Fig. 3.5). In each case, the robot was given 10 attempts to climb the wall, and stopped each time it reached the top or fell. Slips that were recoverable were not counted as failures, although any recovery steps did not count towards the step total. A human teleoperator assisted with foothold selection during the swing phase and when necessary helped the robot maneuver free if it became caught on a rock. The vesicular basalt and slag surfaces were rearranged after 5 trials to increase the variety of terrain features.

Surface	DIG	Payload	Failure Rate	Speed (m/min)
cinder block	No	TX2	6.4%	0.29
cinder block	Yes	TX2	2.3%	0.22
Vesicular Basalt	No	None	6.4%	0.24
Slag	Yes	NUC	3.5%	0.20
Tufa	Yes	NUC	-	0.19

Table 3.2: Climbing Experimental Results

A summary of the results can be seen in Table 3.2. The reported failure rate is the ratio of falls to steps taken, where each step equates to approximately 5 cm of travel. The reported speed uses the fastest successful trial.

On cinder block the robot was unable to maintain adhesion with the Intel NUC onboard, which reduces the adhesion margin by shifting the robot’s center of mass away from the wall, so a smaller NVIDIA TX2 was carried instead as a placeholder. While flat surfaces like cinder block are not the target environment for this robot, these trials provided a controlled validation of the DIG strategy’s efficacy. With this strategy, the robot saw a reduction in step failure rate from 6.4% to 2.3%. Qualitatively, the robot traversed the full 1 m wall in 6/10 trials vs 1/10 without the DIG forces.

The vesicular basalt trials were conducted with an older version of the robot with weaker AX-18A knee motors, so DIG forces could not be applied. Although the robot was able to climb with the NUC onboard, the failure rate was high and the trials were instead conducted without an onboard computer or battery. The robot traversed the full 0.5 m in 4/10 trials, with a 6.4% step failure rate.

For the slag trials, the robot carried the NUC onboard and applied the DIG strategy. The robot traversed the full 0.5 m in 7/10 trials, with a 3.5% step failure rate. The ability to carry the full NUC is attributed to large-scale features of the slag surface which enable more robust grasps.

Finally, an outdoor field test was conducted on a bridge made of tufa stone (a variety of limestone) in a local park (Fig. 3.1). Although the highly irregular surface made suitable footholds hard to identify, LORIS managed to travel for 8 steps before failure. However, we were unable to conduct sufficient experiments at this location to produce an accurate failure rate measure.

3.5 Discussion

Although LORIS successfully traversed a variety of terrain types, its reliability is not yet sufficient for real-world applications in which a single fall could be catastrophic. In both sets of cinder block trials, the majority of failures were due to an individual gripper losing purchase. LORIS only has four grippers, one of which is typically in swing phase, so a single unexpected disengagement is often irrecoverable. In contrast, hexapedal robots like RiSE can keep five grippers engaged at all times – no single gripper failure can cause a fall. This redundancy com-




Figure 3.5: From left to right: the LORIS robot prototype ascending cinder block, vesicular basalt, slag, and tufa surfaces.

pounds with the decreased load per gripper to explain much of the reliability discrepancy (RiSE traveled 9.6 m before failure on a stone aggregate surface [10], 4.4 times further than LORIS averages on cinder block).

On the remaining surfaces, spine engagement became significantly easier due to protruding features and the presence of large vesicles. Instead, the primary challenge came from footholds that were too highly curved for the gripper to conform easily. Other challenges included unexpected body collisions with the surface, and grippers snagging on rocks during both disengagement and the swing phase. In terms of climbing speed, a significant amount of time was spent selecting footholds and retrying steps if the gripper did not engage. Without these teleoperator interventions LORIS travels at 0.4 m/min, roughly double the current speed on slag. For these highly irregular environments, an effective foothold selection and path planning algorithm could dramatically improve both speed and reliability.

Despite these challenges, LORIS provides a valuable proof of concept as a lightweight robot that can nonetheless traverse unstructured vertical terrain. In particular, the distributed inward gripping strategy has been shown to produce a significant increase in adhesion, reducing the




Figure 3.6: Left: the revised LORIS robot design with modifications annotated. Right: the revised LORIS prototype posed on a ledge, with the camera placement providing visibility of the terrain ahead.

failure rate on cinder block by 64%. This comes at a small price in power consumption, raising the power draw during stance by 9.3% from roughly 51 to 56 W (of which the NUC draws 20 W).

A later redesign of the LORIS prototype addressed several shortcomings of the initial design (Fig. 3.6 left). The first was insufficient torque of the tail and wing motors: overheating of the tail motor limited the duration for which the robot could operate, and the wing motors similarly struggled to support the robot’s weight while walking on level ground. These five motors were replaced with more powerful XM430-W350-T servos to match the rest of the robot’s actuators. The battery was also resized from 2200 to 5000 mAh to further extend operational life. Additionally, the onboard computer was moved inside the front body cavity, bringing the robot’s center of mass closer to the climbing surface and thereby reducing the moment that the tail must counteract. Another issue occurred when climbing downwards: as the tail slid over the surface its end tended to jam against rocks despite the rounded shell because the tail is generally angled downwards into the surface. The end of the redesigned tail flares upwards such that any jamming occurs at a much shallower angle, alleviating the issue. Finally, the camera is now mounted on the front body segment instead of the back to provide greater visibility of the terrain ahead when the spine joint is at a large angle, as illustrated in Fig. 3.6 right. The price of these improvements is an increase in mass from 3.2 to 3.6 kg, mostly driven by the larger battery, wing motors, and tail motor.

Overall, the design of LORIS is motivated by two key research objectives: applying under-actuation at the gripper scale to reduce mass, and eliminating constraints on the generation of internal DIG forces to improve adhesion. LORIS achieves both of these competing aims with the splayed gripper concept, a passive design that can nonetheless resist loads with a significant lateral component. Compliance within the gripper and wrist enable passive adaptation to the shape of the terrain, so that LORIS can go beyond flat walls and traverse irregular natural rock faces. However, at the full robot scale LORIS is fully actuated, with three DoF per limb for precise foothold placement. In Chapter 4, we shift our focus to underactuation at this larger scale by examining a second robot, Sally.

Chapter 4

Design of a Magnetic-Wheeled Inspection Robot

In recent years, magnetic climbing robots have been developed for inspection and maintenance tasks on a wide variety of steel structures including bridges [6], pipes [7], storage tanks [8], and ship hulls [9]. Robotic inspection can eliminate the need to construct costly scaffolding and send human inspectors into potentially hazardous environments. However, robotic locomotion on these structures can be challenging; for instance, the interior of an oil tanker contains several small obstacles, such as rivets, cables, and pipes, as well as interior corners between the floor, walls, and ceiling of the tank. These corner transitions are particularly difficult for wheeled and tracked robots because the robot must overcome the magnetic adhesion with the previous surface. Inspection robots may also need to carry bulky sensor payloads. In particular, portable X-ray fluorescence (pXRF) sensors are often used by human inspectors to determine the elemental composition of a surface. While automated pXRF inspection has been explored for soil remediation [61, 62], this sensing technology also holds great promise for detecting components of interest such as lead or mercury on steel structures.

Aerial robots are often used for visual inspection and researchers have developed perching solutions to enable close-range measurements [63]. However, aerial robots struggle with large payloads, limited runtimes, and confined environments. Tracked climbers can support larger payloads due to the greater surface area of the treads, as in [64], but typically have difficulty traversing larger obstacles. Legged climbers [65, 66] or hybrid solutions [67, 68, 69] can potentially traverse the widest range of obstacles using switchable adhesion and dexterous limbs, but at the cost of increased complexity and reduced payload. A legged morphology would also make continuous line scanning, discussed below, more challenging. In contrast, wheeled robots can use passive suspensions [6], active suspensions [70], or swiveling magnets [71] to traverse corners and obstacles while maintaining wheel contact, but have thus far been too small to carry a bulky payload like a pXRF sensor and unable to transition to non-ferrous surfaces. A recent exception is found in [72], but this design requires dedicated mechanisms to break contact during




Figure 4.1: Images taken of Sally at various locations around an industrial site. a) Water storage tank, b) I-beam, c) truck bed, d) 1.8 m metal pipe interior, e) 1.8 m metal pipe exterior, f) shipping container.

transitions.

Beyond the capabilities of sensor and robotic technology, the sampling method significantly affects the accuracy and efficiency of inspection tasks. Traditionally, this process involves human inspectors physically climbing inside the structure to perform grid point measurements in several locations within a designated area, the mean of these samples representing the average level of a component of interest in that area [73, 74, 75]. However, with many sensors, including pXRF, a moving-while-sensing sampling method (line scanning) facilitated by robotic movement presents a more accurate and efficient alternative. This concept has been applied in pipe inspection tasks such as erosion detection and thickness measurement using ultrasonic, PEC, or IR laser probe sensors [76, 77, 78]. Line scanning, compared to grid point sampling, has the advantage of covering a larger area in the same amount of time, potentially offering a more accurate estimation of the area mean of the component of interest.

In this work we present Sally (Fig. 4.1), the first magnetic-wheeled robot to carry a pXRF sensor for steel structure inspection. Sally can overcome all interior corner transition types, traverse small obstacles up to 5 cm in height, maneuver in tight spaces, and conduct both point samples and line scans. We analyze the feasibility of various interior corner transitions, determining a set of design constraints that apply to all magnetic-wheeled robots with similar morphology. We find that no single set of parameters can enable a robot to accomplish all transition types unaided, but by re-purposing steering and sensor deployment mechanisms Sally is able to achieve all transition types without adding dedicated mechanisms. We evaluate Sally’s capabilities in both a controlled laboratory setting and a realistic field environment.




Figure 4.2: An annotated visualization of the Sally robot design.

4.1 Robot Design

Sally features an aluminum chassis, measuring 46 cm long by 35 cm wide without including the rotating wheel assemblies. The robot weighs 8.3 kg and drives on level ground at 0.2 m/s (0.44 body lengths per second). An onboard Intel NUC 11 computer and 3S 5000 mAh LiPo battery enable the robot to operate untethered. Dynamixel XM430-W210-T servos are used for all actuators on the robot. Sally’s capabilities derive from several key features of the design, listed below.

Magnetic Wheels

Sally’s wheels each consist of a single axially-polarized NdFeB ring magnet (manufactured by SDM Magnetics Co.), with a keyed central hole for mounting to an axle. Each magnet is 5 cm in diameter, weighs 350 g, and provides up to 30 kg (294 N) of adhesive force on steel surfaces. In general, wheeled robots can either rely on magnetic wheels or magnets fixed to the base of the robot for adhesion. While fixed magnets offer greater flexibility for the magnet form factor, magnetic wheels can remain in contact with the surface to provide consistent adhesion as the robot traverses obstacles. Magnetic wheels can be constructed by mounting several magnets around the rim of a wheel, as in [79], but this results in varying magnetic strength as the wheel rotates, especially for low numbers of magnets. Other robots, like Sally, rely on ring magnets which exert a constant force as the wheel rotates, as in [70].

Crowned Tires

The tires each consist of a 3D printed shell that is press-fit around the magnet to provide a crowned profile with a 40 mm radius of curvature, and then wrapped in rubber grip tape to improve traction and dampen impacts. The shell is 2 mm thick at its widest point and reduces the adhesive force of the wheels on steel to 60 N each. Sally's total mass is 8.3 kg, so this gives an adhesion-to-weight ratio of 2.9. Empirical measurements place the friction coefficient of the grip tape at approximately 0.8 on steel and 1.0 on paint, although the friction coefficient gradually decreases over time as the tape experiences wear and accumulates dirt. The crowned profile of the tires provides multiple benefits. First, the curvature reduces the size of the contact patch, making it easier to turn the wheel in place when steering. Secondly, the curvature reduces the variation in magnetic force as the wheel is tilted, ensuring consistent adhesion when traversing curved surfaces or obstacles. The adhesive force of an axially-polarized magnetic wheel decreases with distance from the surface and tilt angle relative to the surface [80], so the crowned profile reduces the distance at larger angles to partially cancel out these two effects.

Rocker Suspension

Attached to each side of the frame is a freely rotating, 46 cm long aluminum rocker plate. A pair of 0.58 N/mm springs connect each rocker plate to the central body to keep the frame roughly aligned with the rockers. The frame has 7.6 cm of ground clearance when the rockers are level, and the rockers can rotate 20° or more to raise any given wheel above the base of the frame when going over an obstacle. This passive suspension ensures that all four wheels can remain contact with the terrain at all times without any active effort by the robot. While an active suspension could also be used, the extra weight and control complexity make it less desirable than the passive suspension which achieves the goal of enabling obstacle traversal without needing precise sensing and control.

Independent Drive and Steering

The four steering motors are mounted at the ends of the two rocker plates, directly connected to the wheel brackets. Each wheel is supported by a U-shaped aluminum bracket with a drive motor mounted on one side. The ability to independently steer all four wheels enables the robot to move with any desired translation and rotation on the surface, including turning in place. This maneuverability helps Sally to navigate tight spaces and closely spaced obstacles; in contrast, a robot with only two independently steered wheels would have a non-zero turning radius and be unable to move laterally. In addition, when the robot is holding position on a vertical surface the wheels can be turned perpendicular to gravity, enabling the robot to hold position indefinitely without exerting torque. This is especially useful for inspection tasks because pXRF measurements take 30-90 seconds, so reducing power usage during these measurements greatly extends the robot's battery life. Finally, during interior corner transitions, the rear wheels can be twisted to help break contact with the previous surface (Section 4.2.2).

Portable X-Ray Fluorescence Sensor and Elevator

The robot carries a Vanta pXRF sensor to analyze the surface elemental composition. The pXRF sensor is mounted on a moving carriage, which connects to the robot’s frame through a pair of sliders. The sliders are constructed with ball bearings that contact three faces of the rails, constraining the carriage’s motion to one axis. A pair of independently driven belts translate the sensor up and down, so that the sensor can be raised clear of any obstacles or lowered to make contact with the surface. The carriage’s range of motion goes from 2 cm below the driving surface to 7 cm above it, enabling Sally to scan areas that are not level with the wheels. The ability to extend below the driving surface also enables the carriage to press against the ground, aiding with interior corner transitions (Section 4.2.2). Four omni wheels are mounted around the base of the sensor to reduce friction during line scanning and overcome small obstacles.

4.2 Interior Corner Transitions

4.2.1 Transition Model

In [81], Fischer et al. present a model for a magnetic-wheeled robot undergoing an interior corner transition and estimate the required wheel friction coefficient. Their analysis focuses on relatively small robots, for which the wheels are large relative to the robot body and the magnet adhesion force is many times greater than the robot weight. A similar model for interior corner transitions is presented in [82], but focuses only on one robot configuration and one surface friction condition. In contrast, here we examine how the required friction coefficient scales with robot weight relative to the magnetic adhesion force, and additionally consider transitions between non-ferrous and ferrous surfaces.

We consider a wheeled robot characterized by the following dimensionless parameters,

$$M = \frac{mg}{2f_m}, \quad A = \frac{h}{\ell} \quad (4.1)$$

Where m is the robot mass, g is the gravitational acceleration, f_m is the adhesive force of a single wheel pair (e.g. front or back), ℓ is the distance between the wheel axes, and h is the distance from the center of mass to the bottom of the wheels. We refer to M as the weight ratio (weight divided by total adhesive force) and A as the aspect ratio. The following constraints are imposed on the robot parameters,

$$M < 1 \quad (4.2)$$

$$MA < \frac{1}{2} \quad (4.3)$$

The constraint in (4.2) ensures that the magnetic force from the wheels is strong enough for the robot to hang inverted from the ceiling, and (4.3) is the moment balance about the rear wheels which ensures that the robot can drive vertically on a wall without losing adhesion.




Figure 4.3: A free-body diagram of a magnetic-wheeled robot at the moment of front wheel (left) and rear wheel (right) floor-to-wall transition.

In order to simplify the analysis, we assume that the wheel radius is small relative to the size of the robot and that the center of mass is horizontally centered between the two wheels. We further assume that the transition is performed quasi-statically at low speed.

To determine the friction conditions necessary to perform an interior corner transition, we examine the transition in two stages: the front wheel transition and the rear wheel transition (Fig. 4.3). In each case, at the moment of transition the contact force between the specified wheel and the surface it is leaving is zero. We can express the condition of static equilibrium as the vector equations,

$$0 = \mathbf{f}_1 + \mathbf{f}_{m1} + \mathbf{f}_2 + \mathbf{f}_{m2} + m\mathbf{g} \quad (4.4)$$

$$0 = \mathbf{r}_1 \times (\mathbf{f}_1 + \mathbf{f}_{m1}) + \mathbf{r}_2 \times (\mathbf{f}_2 + \mathbf{f}_{m2}) \quad (4.5)$$

Where \mathbf{f}_1 and \mathbf{f}_2 are the front and rear contact force vectors, \mathbf{r}_1 and \mathbf{r}_2 are the position vectors of the front and rear wheels relative to the center of mass, \mathbf{f}_{m1} and \mathbf{f}_{m2} are the magnetic force vectors acting on the front and rear wheels (combining the magnetic force on each wheel to multiple surfaces as needed), and \mathbf{g} is the gravitational acceleration vector.

We decompose each contact force \mathbf{f}_i into tangential and normal components T_i and N_i and expand (4.4)–(4.5) for the front wheel transition,

$$0 = -N_1 + f_m + T_2 + mg_x \quad (4.6)$$

$$0 = T_1 + N_2 - 2f_m + mg_y \quad (4.7)$$

$$0 = \frac{\ell}{2}(T_1 - N_2) + h(N_1 + T_2 + f_m) \quad (4.8)$$

and rear wheel transition,

$$0 = -N_1 - N_2 + 2f_m + mg_x \quad (4.9)$$

$$0 = T_1 + T_2 - f_m + mg_y \quad (4.10)$$

$$0 = h(T_1 + T_2 - f_m) + \frac{\ell}{2}(N_1 - N_2) \quad (4.11)$$

Where all normal forces $N_i \geq 0$.

Our goal is to find the minimum wheel friction coefficient μ for which neither wheel will slip during the transition. We assume that this will occur when both wheels are at the friction limit,

$$\mu = \frac{T_1}{N_1} = \frac{T_2}{N_2} \quad (4.12)$$

We analytically solve (4.6)–(4.8) or (4.9)–(4.11) together with (4.12) for \mathbf{f}_1 and \mathbf{f}_2 in terms of M for varying values of A and \mathbf{g} using the MATLAB Symbolic Math Toolbox. Equations (4.2) and (4.3) are treated as assumptions. Note that (4.12) is nonlinear in the force components, preventing the problem from being solved as a linear system of equations.

We also consider two special cases. First, transitions involving a non-ferrous floor are possible since no magnetic force is required to keep the robot on the driving surface. For non-ferrous-floor-to-wall and wall-to-non-ferrous-floor transitions, the appropriate \mathbf{f}_m terms are set to zero.

The other case we consider is a wall-to-wall transition. In a wall-to-wall transition, gravity is parallel to both surfaces, and therefore does not act in the direction of either the normal or the tangential contact forces. Since our model only considers in-plane forces, the wall-to-wall solutions effectively assume no out-of-plane forces or moments, which is equivalent to assuming that the aspect ratio (or height) of the robot is zero.

The results of this analysis are shown in Fig. 4.4 for each type of transition. We observe that the wall-to-non-ferrous-floor transition is the most difficult transition at low values of M , while the floor-to-wall and wall-to-ceiling transitions are the most difficult at high M values. At the crossover point $M = \frac{1}{2}$, an infeasibly high friction coefficient of 1.0 is required, revealing that no robot of this morphology can achieve all transition types unaided. We also observe that a friction coefficient of $\frac{1-\sqrt{5}}{2} \approx 0.62$ is the minimum required to complete all transitions between ferrous surfaces as the weight ratio approaches zero.

4.2.2 Transition Strategies

Examining the approximate location of Sally on the plot in Fig. 4.4 ($M = 0.345$, $\mu = 0.9$), we predict that wall-to-non-ferrous-floor transitions will be infeasible, while floor-to-wall and wall-to-ceiling transitions are borderline depending on the freshness of the wheel grip tape. However, the presence of additional actuators on the robot for sensor deployment and steering can enable Sally to circumvent these limits.

During a front wheel transition, Sally can press the pXRF elevator into the surface to produce an additional vertical force near the center of the robot, helping to pry the front wheels free and




Figure 4.4: Minimum required tire friction coefficient as a function of normalized robot weight for each type of interior corner transition. For each transition, the maximum of the front wheel and rear wheel transitions at each weight ratio is taken as the minimum required friction. Transitions are plotted at an aspect ratio of 0, with shaded regions indicating how the friction conditions change up to an aspect ratio of 0.5.

reducing the required friction coefficient. In this case, Equations (4.4) and (4.7) are modified to include the force \mathbf{f}_p from the pXRF elevator.

$$0 = \mathbf{f}_1 + \mathbf{f}_{m_1} + \mathbf{f}_2 + \mathbf{f}_{m_2} + \mathbf{f}_p + m\mathbf{g} \quad (4.13)$$

$$0 = T_1 + N_2 - 2f_m + f_p + mg_y \quad (4.14)$$

Similarly, during a rear wheel transition, Sally can use the independent steering motors to twist the rear wheels sideways. This has two effects: first, the magnetic adhesion force diminishes at larger angles relative to the surface, making it easier to break free. Second, as the wheel twists the side of the wheel bracket comes into contact with the surface, so the torque applied by the steering motor generates an additional contact force pushing the robot away from the wall.

For each front wheel transition, using the pXRF to push off of the wall reduces the friction coefficient required as shown in Fig. 4.5. The curves for certain transitions terminate early at the point where additional pXRF force would cause Sally to fall. One possible failure mode is due to the rear wheels losing contact with the initial surface; this occurs when the solution would require a negative normal force on the rear wheel. Solutions are filtered out if the normal force corresponding to the real wheel contact is negative. The wall-to-wall and wall-to-ceiling transitions can also fail if the robot starts to slide down the wall. The pXRF force decreases the normal force on the wheels, which reduces the frictional force they can apply and potentially causes the robot to slip. For wall-to-wall transitions, the weight of the robot is evenly distributed




Figure 4.5: The required friction coefficient for each front-wheel transition as force is applied with the pXRF elevator, calculated for Sally’s values of $M = 0.345$ and $A = 0.27$. The coefficient decreases as pXRF force increases, up until the lines cut off at the point where the pXRF force would cause the robot’s rear wheel to slip or lose contact.

between the front and back wheels to balance the net moment, yielding the following condition,

$$N_1, N_2 \geq \frac{mg}{2\mu} \quad (4.15)$$

For wall-to-ceiling transitions, the limiting condition for sliding occurs when the front wheels experience no contact force. In this condition, the rear wheels experience a normal force of $2f_m - f_p$, so the maximum tangential force is $\mu(2f_m - f_p)$. This produces the following force balance equation in the vertical direction,

$$f_m + \mu(2f_m - f_p) - mg = 0 \quad (4.16)$$

Solving for the normalized pXRF force, $P = \frac{f_p}{2f_m}$, we get the following dimensionless no-slip condition,

$$P \leq 1 - \frac{1}{\mu}(W - \frac{1}{2}) \quad (4.17)$$

In cases where $W < \frac{1}{2}$ (such as Sally), the robot is light enough that the magnetic force between the front wheels and the ceiling is sufficient to hold the robot up, so the sliding failure mode can be safely ignored.




Figure 4.6: Average front and rear wheel motor torques and velocities during a floor-to-wall transition. The robot’s wheels are normally in velocity mode, but switch to torque for the duration between the front and rear wheel transitions.

4.2.3 Transition Control

During certain transitions, an additional challenge arises that is not captured by the above static analysis. As the robot moves across a transition, the relative speeds of the wheels relative to the surfaces change due to the corner geometry, with the rear wheels stationary at the start of the transition, the front wheels stationary at the end, and the wheel speeds equal at the halfway point. This presents a problem if the wheels are velocity controlled, because the front wheels will tend to pull the rear wheels free from the previous surface near the end of the transition but before the rear wheels reach the corner. For floor-to-wall and wall-to-floor transitions the robot will naturally pivot around the front wheels until the rear wheels regain contact, but for wall-to-wall, wall-to-ceiling, and ceiling-to-wall transitions this can result in the robot pivoting away from the wall and falling.

With accurate state estimation it would be possible to vary the wheel speeds over the course of the transition and avoid this problem. However, a more robust solution is to use torque control on the wheels between the front wheels breaking contact and the rear wheels making contact, with a relatively low torque setpoint for the front wheels. This ensures that the force exerted by the front wheels will never be sufficient to prematurely pull the robot free of the previous surface. Notably, this solution only requires knowledge of the robot’s contact state, rather than the exact angle of progress through the transition (as encoder-based measurements may become inaccurate due to wheel slip). Fig. 4.6 demonstrates the velocity and torque seen by the front and back motors during a floor-to-wall transition.

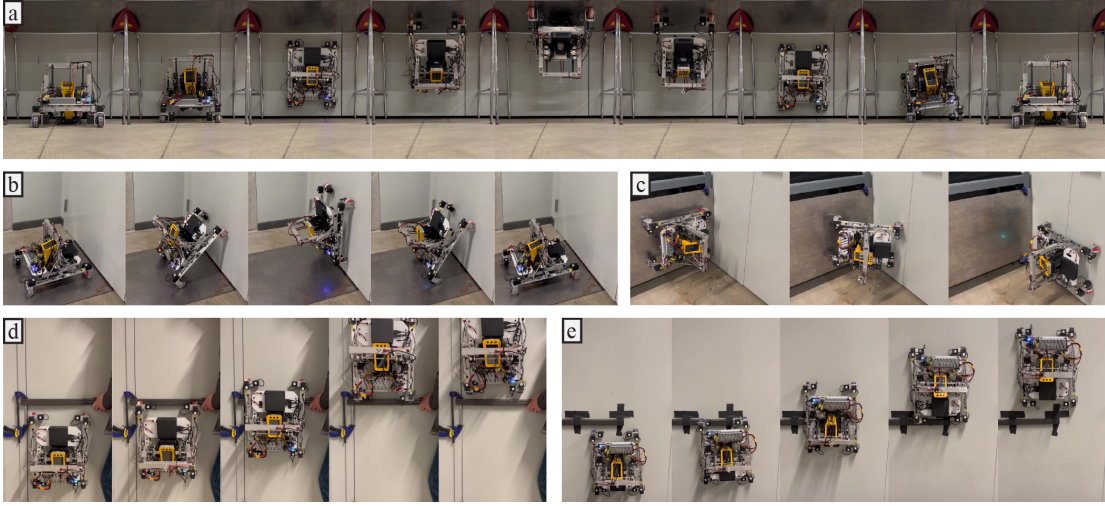


Figure 4.7: Key frames taken during a) non-ferrous-floor-to-wall-to-ceiling-to-wall-to-non-ferrous-floor, b) wall-to-floor-to-wall, c) wall-to-wall, d) two-wheel obstacle (3.8 cm square steel channel) and e) one-wheel obstacle transitions.

4.3 Experiments

4.3.1 Interior Corner and Obstacle Traversal

We attempted all 7 interior corner transition types with fresh grip tape on the wheels, finding that Sally is capable of completing all but the wall-to-non-ferrous-floor transition without assistance from the steering motors or pXRF elevator. However, with the transition strategies described above, Sally was able to complete this transition as well. Key frames from each transition are shown in Fig. 4.7a-c. After repeated trials, the floor-to-wall transition began to require the use of wheel twisting as the grip tape lost traction due to wear.

We also evaluated obstacle traversal while driving on a vertical wall. Two types of obstacle were examined for each test: a segment of 3.8 cm square steel channel and a 5 cm diameter steel pipe. In the first test, each obstacle was placed to obstruct only one side of the drivetrain as a test of the rocker suspension. Sally successfully traversed both obstacles unaided, although the pXRF had to be raised to clear the obstacle. In the second test, each obstacle was placed to obstruct both wheels. Transitioning the back wheels over the obstacle proved challenging, because unlike in a wall-to-ceiling transition the moment due to gravity opposes the wheel breaking free from the wall. Sally was ultimately able to traverse these obstacles as well by applying the pXRF pushing and wheel twisting strategies (Fig. 4.7d-e).

4.3.2 Field Deployment

Sally was also tested at an industrial site which possessed several decommissioned structures available for robot testing, see Fig. 4.1. In particular, the experiments included climbing: 1) the

exterior of a large water storage tank, where Sally was able to drive along the curved surface in any direction, over small bumps and welds, and take point and line scans using the pXRF (Figs. 4.1a and 4.8); 2) a steel I-beam, on which Sally demonstrated a non-ferrous-floor-to-wall transition, point measurement, and wall-to-non-ferrous-floor transition (Fig. 4.1b); 3) a steel truck bed, on which Sally demonstrated a ferrous-floor-to-wall and wall-to-ferrous-floor transition (Fig. 4.1c); 4) a 1.8 m diameter steel pipe, on which Sally demonstrated the ability to drive on the interior and exterior of curved surfaces (Fig. 4.1d and e); and 5) the exterior of a shipping container, on which Sally demonstrated the ability to drive on uneven surfaces (Fig. 4.1f).

Overall, Sally was able to successfully navigate all of these obstacles. Despite these successful demonstrations, several challenges became evident when Sally was taken to a field deployment environment. The most significant is the build up of loose magnetized dirt and rust on the outside of the wheels. This causes a reduction in adhesion, and large enough chunks of material can prevent the wheels from spinning entirely. Another challenge comes from surface patches that have rusted through to the point that they become brittle and the wheel loses adhesion. Finally, Sally struggles to traverse closely spaced rivets because a wheel can pass between two neighboring rivets while the attached drive motor collides with one of them. This issue likely extends to any set of obstacles with spacing wider than the wheel but narrower than the full wheel assembly.

4.3.3 Line Scanning

The pXRF surface sampling can be operated in two different modes: point measurement or line scanning. In traditional grid point measurement, the robot lowers the sensor to the ground and holds its position for the duration of the scan, resulting in a high-precision measurement of the elemental composition of the surface at a single location. In contrast, during line scanning, the robot drives at a constant velocity throughout the sensing, resulting in an average measurement of the surface composition along the line it travels. A typical inspection task is to find any remaining hot spots of target elements after cleaning a surface. Thus, by taking line scans, the robot can relatively quickly cover large areas of the surface and only revisit locations to perform more precise point measurements when non-negligible concentrations of target elements (e.g. heavy metals) are detected. Line scanning is also effective if the study objective is to determine an overall average concentration level for target elements, as it will give a more accurate average for the length of the line than a single scan.

Sally performed a sensing comparison test on the exterior of a water storage tank at an industrial site, where eight grid point measurements and one line scan were conducted to measure the concentration of Zn (Fig. 4.8). On average, each individual point measurement differed by 58.50% from the mean of all eight point measurements ($0.58 \mu\text{g}/\text{cm}^2$). In contrast, the line scan measurement differed by only 36.48% from the mean, despite having the same duration as a single point measurement. Assuming the mean of the eight point measurements accurately represents the true concentration of the estimate area, the line scan's measurement error is




Figure 4.8: Sensing comparison test locations at the field site. Point measurements are marked with blue dots and the line scan is marked with a red line.

37.60% less than the error of a single point measurement. This supports the claim that line scanning gives a better estimate of the area mean than point measurement for the same amount of sampling time.

4.4 Discussion

As predicted by the model, with only the drive actuators Sally is able to complete all transition types except for the wall-to-non-ferrous-floor transition. However, one of the keys to Sally’s effectiveness lies in the dual-use of the robot’s mechanisms, such as the pXRF elevator and steering actuators, in order to complete even these more difficult transition types successfully. In actual deployment, a lower friction coefficient is a safer and more realistic assumption, and thus the two supplemental techniques are critical for ensuring reliable mobility even when performance is degraded. Additionally, the rocker suspension and crowned tire profiles provide passive adaptability to obstacles and other terrain irregularities, maintaining contact and adhesion at the scale of both the robot and the individual wheels.

Sally was successfully deployed at an industrial site where we were able to demonstrate many of the key features of the design. These tests also helped to identify areas for further research and development. The accumulation of magnetized dirt on the wheels, although of less concern in a relatively clean oil tanker setting, is a critical issue in dirty or corroded environments and a self-cleaning solution that scrapes off excess material would be extremely valuable. Using larger wheels or offsetting the drive motors upwards would improve performance when driving over closely spaced rivets.

The line scanning technique, enabled by Sally’s wheeled design, was also shown to provide

a more accurate estimate of the area mean than traditional point measurements over the same period of time. However, scans occasionally fail because of unevenness on the surface, triggering a safety setting within the sensor. This failure mode may be partially addressed by carefully controlling the elevator, but ultimately limits the line scanning technique to fairly smooth surfaces (with the robot’s obstacle and corner maneuverability valuable for transitioning between these surfaces). Despite this limitation, line scanning offers significant time savings across a wide array of tasks as robotic solutions become more prevalent.

Moving forward, an important task will be analyzing the scalability of Sally’s design, particularly the jump to a 6-wheeled rocker-bogey suspension to support a larger payload. Furthermore, the current interior corner transition model assumes quasi-static motion and a straight-on approach, but by using the robot’s momentum or approaching at an angle, one wheel at a time, it may be possible to achieve transitions with a significantly lower friction coefficient. Finally, traversing interior corners and obstacles with Sally currently requires a teleoperator to remember a sequence of operations (e.g. push with pXRF, switch to torque mode, switch back to velocity mode, twist rear wheels), where an incorrect command can cause the robot to lose adhesion and fall. It would be valuable to develop driver-assist features that automate this logic, reducing the burden on the teleoperator and paving the way for a fully autonomous solution. In particular, one potential approach is to deliberately regulate the internal forces between the front and back wheels, in order to maximize the margin of safety – an analogous problem to the optimization of distributed inward grasping forces on LORIS. This internal force control strategy, and its applicability to both robots, will be discussed in Chapter 5.

Chapter 5

Unified Control of Shape, Contact, Force, and Motion

5.1 Introduction

In the preceding chapters, we discussed the design of climbing robots that use underactuated and multi-functional mechanisms to achieve high performance while minimizing actuator mass. We now move on to the second layer of the proposed hierarchy: control of the actuated joints to optimize the internal contact forces while climbing for a more secure grasp. We divide this control task into three components: contact force estimation, internal contact force optimization, and unified Shape-Contact-Force-Motion (SCFM) control.

While it is possible to directly measure contact forces via dedicated sensors, for many robots the additional mass, cost, and complexity of sensor integration makes this approach impractical. In contrast, many modern controllers feature built-in current sensing, which can be used to approximate the torque output. Prior work has used a generalized momentum observer to estimate contact forces from these motor torques [83], but for a quasi-static robot it is possible to obtain a reasonable estimate just by looking at the present time step. We describe our implementation of such an estimation procedure, including in cases of over or underactuation. We also extend this procedure to cases where the robot's orientation relative to gravity is unknown.

Once the contact force estimates are obtained, the next step is to determine the optimal set of contact forces to apply that minimize the risk of falling. A robot's pose is considered statically stable if its center of mass lies inside its support polygon when projected onto the ground plane. For a robot walking on flat ground the support polygon is simply the convex hull of its feet, but on uneven terrain this is more difficult to compute. The authors of [84] discuss how linear programming can quickly determine if a single point lies in the polygon while linear projection can compute the full polygon, albeit more slowly. They present an incremental projection approach to bridge the two methods, which is further improved upon in [85]. However,

these works focus on fully actuated robots with simple friction cone constraints. An alternative approach is the tumble stability margin, which considers gripping forces but requires a no-slip assumption [86].

In this work we extend the linear programming approach to a wider range of robot designs, including underactuated robot morphologies and a variety of different adhesion mechanisms. Optimal solutions are shown for both LORIS (Chapter 3) and Sally (Chapter 4) to demonstrate how this optimization approach can automatically generate distributed inward gripping forces as in [20] and prevent accidental loss of adhesion.

The final step is regulating the robot’s contact forces to achieve the optimal solution. Hybrid force/motion control strategies [87] enable a robotic manipulator to simultaneously regulate contact forces and velocities along different axes. We extend this strategy to self-manipulation [88] with SCFM control, introducing three additional control objectives: maintaining contact with the terrain, controlling the robot’s shape using any redundant degrees of freedom, and obeying safety limits on the maximum output velocity.

This overall control scheme was implemented with minor simplifications on LORIS to achieve the successful climbing results of Chapter 3. A variation of the SCFM control scheme described here was also implemented on the EELS 1.5 robot [21], enabling the controlled descent of a moulin (a vertical shaft in the ice) at Athabasca Glacier [89].

5.2 Robot Model

We begin by defining a model of a climbing robot, which is represented as a tree of rigid links connected by joints where some number of branches terminate in an end effector. We assign reference frames for the world frame W , the robot body frame B (i.e. the root of the tree), the robot link frames L_l for each link l , and the contact frames C_i for each end effector i . Contact frames are defined with their z axis normal to the surface. We define the state space q of a robot in d dimensions by the joint angles $\theta \in \mathbb{R}^n$ and the body transform in the world frame $x_b \in SE(d)$, and with abuse of notation define the state derivative \dot{q} by the joint velocities $\dot{\theta}$ and body twist $v_b \in \mathbb{R}^p$.

$$q = \begin{bmatrix} \theta \\ x_b \end{bmatrix}, \quad \dot{q} = \begin{bmatrix} \dot{\theta} \\ v_b \end{bmatrix} \quad (5.1)$$

We define the equation of motion for the robot during self-manipulation similarly to [88],

$$\bar{M}(q)\ddot{q} + \bar{C}(q, \dot{q})\dot{q} + \bar{N}(q) + A^T(q)f_c = \Upsilon \quad (5.2)$$

where \bar{M} is the mass matrix, \bar{C} is the Coriolis matrix, \bar{N} is the nonlinear (gravitational) forces, A is the velocity constraint matrix, f_c is the contact forces, and Υ is the actuator forces. We assume quasi-static motion, $\ddot{q} = \dot{q} = 0$, eliminating the \bar{M} and \bar{C} matrices. Expanding the

remaining terms,

$$\begin{bmatrix} N \\ -f_{ext} \end{bmatrix} + \begin{bmatrix} -J_h^T \\ G_s \end{bmatrix} f_c = \begin{bmatrix} \tau \\ 0 \end{bmatrix} \quad (5.3)$$

where f_{ext} is the gravitational wrench on the body (in the body frame), N is the gravitational force on each joint, J_h is the hand Jacobian (the mapping from the joint velocities $\dot{\theta} \in \mathbb{R}^n$ to the contact velocities $v_c \in \mathbb{R}^m$), $G_s = -G$ is the self-manipulation grasp map (the mapping from the contact wrenches $f_c \in \mathbb{R}^m$ to the body wrench $f_b \in \mathbb{R}^p$), and $\tau \in \mathbb{R}^n$ is the actuator joint torques.

To compute N , the potential energy V is differentiated with respect to the joint angles θ ,

$$N = \frac{\partial}{\partial \theta} V = \frac{\partial}{\partial \theta} \sum_{l=1}^n -m_l r_l^T g = - \sum_{l=1}^n m_l J_l^T g \quad (5.4)$$

where m_l is the mass of link l , r_l is the position of each link center of mass in the body frame, g is the gravitational acceleration vector, and J_l is the link Jacobian (the mapping from the joint velocities $\dot{\theta}$ to the link velocities \dot{r}_l). For a robot modeled with massless limbs, $N = 0$. The gravitational body wrench f_{ext} can be computed similarly as $-\frac{\partial}{\partial x_b} V$.

We assume a linear stiffness model as in [90], in which contact forces are generated by displacements δx_c of the end effectors relative to the terrain in the contact frames,

$$\delta f_c = -K S_C \delta x_c = K S_C A \delta q \quad (5.5)$$

where K is a symmetric positive definite stiffness matrix capturing both passive and active compliance in the system and S_C is a diagonal selection matrix for the end effectors in contact with the terrain. From this equation, it is evident that $\delta f_c \in \mathcal{R}(K S_C A)$, where $\mathcal{R}(\cdot)$ denotes the range space (column space) of a matrix. This is equivalent to the condition $\mathcal{N}(A^T S_C^T K^T)^T \delta f_c = 0$, where $\mathcal{N}(\cdot)$ denotes the null space of a matrix. In general, the null space is empty when $n + p \geq m$. While the force in this null space remains constant, it may have a nonzero initial value due to preload when making contact (e.g. jamming effects). If K is set to the identity matrix, then this null space is equivalent to the preloaded internal grasp force defined in [91].

5.3 Contact Force Estimation

A robot equipped with 6-axis force torque sensors at each end effector can produce extremely accurate contact force measurements, but such sensors add additional weight, complicate the robot design, and require time-consuming calibration. An alternative solution is to generate lower-accuracy force estimates by measuring the joint torques at each of the robot's actuators via built-in current sensing. A moving average filter can be used to reduce the noise in the raw current measurements. Additionally, many robots are equipped with an inertial measurement unit (IMU) that can determine the robot's orientation. Given the estimated gravity vector \hat{g} it is then possible to compute the external body wrench \hat{f}_{ext} and joint torques \hat{N} induced by

gravity.

Given measurements of the actuator torques $\hat{\tau}$, external body wrench \hat{f}_{ext} , and gravitational joint torques \hat{N} , we can use (5.3) to solve for the estimated contact force $\hat{f}_c \in \mathbb{R}^m$,

$$\begin{bmatrix} -J_h^T \\ G_s \end{bmatrix} \hat{f}_c = \begin{bmatrix} \hat{\tau} - \hat{N} \\ \hat{f}_{ext} \end{bmatrix} \quad (5.6)$$

This equation has a unique solution when $n + p = m$ and there are no redundant joints. If the equation is overconstrained ($n + p > m$) then a least squares solution can be computed, inversely weighted by the measurement covariance Σ to account for the relative accuracy of the IMU and current sensing. This solution minimizes $(A^T \hat{f}_c - b)^T \Sigma^{-1} (A^T \hat{f}_c - b)$, where b is the right hand side vector in (5.6) and A is the constraint matrix as defined above (i.e. the transpose of the left hand side matrix in (5.6)). In this overconstrained case, the solution is,

$$\hat{f}_c = (A \Sigma^{-1} A^T)^{-1} A \Sigma^{-1} b \quad (5.7)$$

If the equation is underconstrained ($n + p < m$) then the least-norm solution can be found, inversely weighted by the stiffness K to find the energy-minimizing solution. This solution minimizes $\frac{1}{2} \hat{f}_c^T K^{-1} \hat{f}_c$, which physically corresponds to the energy stored in the virtual springs. In this underconstrained case, the solution is,

$$\hat{f}_c = K A (A^T K A)^{-1} b \quad (5.8)$$

If a robot does not possess an IMU, then the second row of (5.6) can be removed and the equation instead has a unique solution when $n = m$. Without an IMU measurement of g it is no longer possible to precompute N . Fortunately, (5.4) shows that N is a linear function of g , which can be computed as the linear components of f_{ext}/m_{net} , where m_{net} is the total mass of the robot and $f_{ext} = G_s f_c$. Substituting (5.4) into (5.6) and dropping the second row gives the following linear equation which can be solved for \hat{f}_c as above,

$$\left(-J_h^T - \frac{1}{m_{net}} \sum_{l=1}^n m_l \begin{bmatrix} J_l \\ 0 \end{bmatrix}^T G_s \right) \hat{f}_c = \hat{\tau} \quad (5.9)$$

Finally, we can estimate which end effectors are in contact with the terrain by thresholding the individual contact force estimates $\hat{f}_{c_i} \in \mathbb{R}^{m_i}$ to determine the contact set \mathcal{C} . The contact force estimate \hat{f}_{c_i} for any $i \notin \mathcal{C}$ is rounded to zero.

5.4 Internal Force Optimization

Prior climbing robots such as [10] have relied on the straightforward strategy of evenly distributing the gravitational load among the grippers. This can suffice on flat surfaces or for grippers with much greater holding force than the task requires, but may not generalize for smaller

grippers on unstructured terrain. In contrast, the researchers in [92] formulated a combined locomotion and grasping optimization problem for the SCALER robot, but this approach is slow, requiring multiple minutes to find a new solution, and is thus unable to react quickly if a gripper begins to slip. Instead, we present an optimization problem that can determine the ideal force distribution in a matter of milliseconds for any robot posture, not just vertical climbing on a flat wall.

5.4.1 General Formulation

We define an optimization problem to select the desired contact forces f_c at each end effector. We focus on quasi-static climbing, so the optimal forces can be computed at each timestep without considering their effects on future timesteps. The optimization is formulated as a linear program, for which many fast solvers exist, which aims to maximize the adhesion margin c given several linear constraints,

$$\arg \max_{f_c, c} c \quad (5.10)$$

$$\text{s.t. } 0 = G_s f_c - f_{ext} \quad (5.11)$$

$$A_{c_i} f_{c_i} + c \leq b_{c_i} \quad i \in \mathcal{C} \quad (5.12)$$

$$\tau_{min} \leq -J_h^T f_c + N \leq \tau_{max} \quad (5.13)$$

$$0 = \mathcal{N}(A^T S_c^T K^T)(f_c - \hat{f}_c) \quad (5.14)$$

The objective being maximized (5.10) is the adhesion margin $c \in \mathbb{R}$, defined as the minimum disturbance force in any direction required to dislodge an end effector. The adhesion margin is used as a slack variable to represent a minimum safety margin or maximum constraint violation across all grippers depending on its sign. The first constraint (5.11) represents the static equilibrium condition. As described above, the external wrench can be estimated using an IMU or computed from the current contact force estimates as $f_{ext} = G_s \hat{f}_c$.

A_{c_i} and b_{c_i} in (5.12) represent linearized adhesion constraints of the form $A_{c_i} f_{c_i} \leq b_{c_i}$ specific to each type of end effector that is discussed in more detail below. Constraint (5.13) ensures that the joint torques $\tau = -J^T f_c + N$ respect motor torque limits. The final constraint (5.14) ensures that any forces generated are physically realizable according to the stiffness model in (5.5). For a fully actuated robot where A is full rank ($n + p \geq m$), this constraint is equivalent to $f_{c_i} = 0$ for $i \notin \mathcal{C}$.

Although not shown here for simplicity, (5.12) can be replaced with a guard condition equality constraint for end effectors that are desired to change contact modes. For instance, the microspine grippers on LORIS have both disengaged and engaged contact states. To reach the engaged state after making contact, a gripper must generate sufficient force components along the gripper axis and into the surface. Similarly, before an end effector breaks contact its contact force should be set to zero to prevent a sudden jerk when contact is lost.

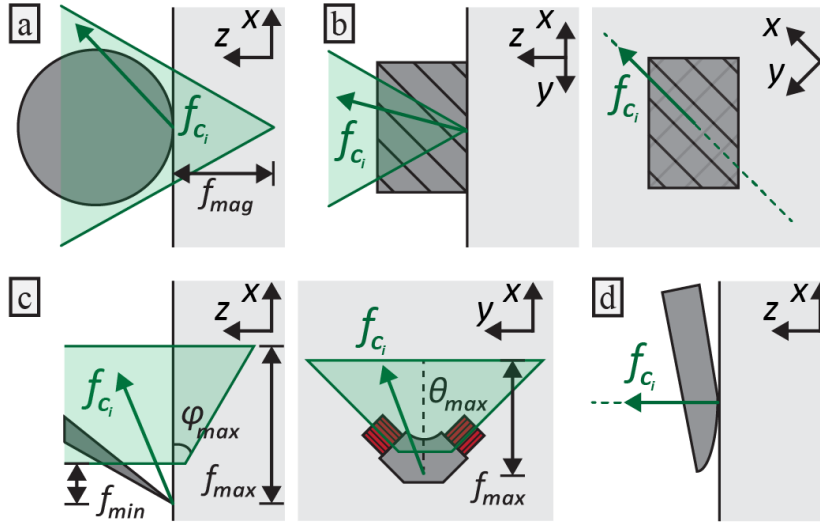


Figure 5.1: Sets of feasible contact forces (in green) that can be supported by each type of end effector without losing engagement with the climbing surface: a) magnetic wheel, b) ice screw, c) splayed microspine gripper, and d) sliding tail. Note that the scaling between force and position is arbitrarily selected for ease of viewing.

5.4.2 Adhesion Constraints

We apply this optimization approach for several end effector types, with adhesion constraints as shown in Fig. 5.1. The first is a magnetic wheel as found on Sally, which generalizes frictional point contact ($m_i = 3$) to include a normal adhesive force f_{mag} ,

$$\begin{bmatrix} -1 & 0 & -\mu \\ 1 & 0 & -\mu \\ 0 & -1 & -\mu \\ 0 & 1 & -\mu \end{bmatrix} f_{c_i} + c \leq \begin{bmatrix} \mu f_{mag} \\ \mu f_{mag} \\ \mu f_{mag} \\ \mu f_{mag} \end{bmatrix} \quad (5.15)$$

These linear constraints define a 4-sided friction cone, but additional sides can be added to improve the approximation. The friction coefficient is given by μ and we assume $\mu \leq 1$, as typical for rubber on steel.

Next, we examine an ice screw as found on the EELS robot [21], which constrains motion normal to the surface and perpendicular to the screw threads, but slides freely parallel to the threads. Here we define the x axis perpendicular to the threads, y axis parallel to the threads, and z axis normal to the surface, but the contact force f_{c_i} only has x and z components ($m_i = 2$), resulting in the constraint,

$$\begin{bmatrix} -1 & -\mu \\ 1 & -\mu \\ 0 & -1 \end{bmatrix} f_{c_i} + c \leq \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad (5.16)$$

The passive splayed microspine gripper, as found on LORIS [16], constrains all linear motion ($m_i = 3$). Here we define the x axis tangential to the gripper, y axis lateral to the gripper, and z axis normal to the surface. The first two rows ensure that the microspines are neither disengaged nor overloaded tangentially. The third and fourth rows ensure that the in-plane loading angle does not exceed θ_{max} , and the final row ensures that the pull-off angle does not exceed ϕ_{max} out-of-plane,

$$\begin{bmatrix} -1 & 0 & 0 \\ 1 & 0 & 0 \\ -\tan(\theta_{max}) & -1 & 0 \\ -\tan(\theta_{max}) & 1 & 0 \\ -\tan(\phi_{max}) & 0 & -1 \end{bmatrix} f_{c_i} + c \leq \begin{bmatrix} -f_{min} \\ f_{max} \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (5.17)$$

Finally, the LORIS robot also possesses a tail, represented as a frictionless point contact that constrains motion normal to the surface ($m_i = 1$). As written, (5.12) assigns an equal weight to the adhesion margin c for each end effector. However, there is no risk of falling if the tail loses contact with the surface, so the adhesion margin can be left out of this constraint,

$$\begin{bmatrix} 0 & 0 & -1 \end{bmatrix} f_{c_i} \leq 0 \quad (5.18)$$

5.4.3 Optimization Results

In Fig. 5.2, representative optimal force distributions are shown for the LORIS robot in a nominal climbing gait with four variations of the gripper design. Non-splayed grippers, able to exert force only along the tangential axis, produce significantly lower adhesion margins even when the grippers are oriented 45° inwards to enable distributed inward gripping (DIG). With splayed grippers, however, this inward angle offset enables the generation of large DIG forces to double the adhesion margin. These optimization results make clear the value added by the splayed gripper design, relative to the directional grippers used by wall-climbing robots like RiSE [10].

Similarly, Figure 5.3 shows the optimal force distributions for Sally during a wall-to-ceiling transition. This is one of the more challenging transition types for a teleoperator: if both wheels are driven at equal velocity then the front wheel will pull the rear wheel off of the wall near the end of the transition (around 80° progress). In contrast, the optimal force distributions shown here actually maintain a positive normal force on the rear wheel, ensuring that a very large (32.5 N) disturbance would be needed to cause early detachment. We also see the difficulty of the rear wheel transition, as evidenced by the relatively small adhesion margin (5.7 N). This aligns with the static analysis from Chapter 4 which showed that Sally's friction coefficient was near the minimum required to complete this transition type (without wheel twisting).




Figure 5.2: Optimal contact force distributions for LORIS's typical climbing posture with the front left limb in swing phase. Adhesion margin is compared between non-splayed (axial) vs splayed grippers, both with and without a 45° inward angle offset. Contact forces are represented by blue arrows, and the 2D projections of the viable loading directions onto the surface plane for each gripper type are annotated in red.

5.5 Unified Control Scheme

A climbing robot has five objectives to pursue. First and foremost, the robot must maintain a stable set of contact forces that lie within the adhesion constraints of its end effectors. Second, the robot should attempt to regain contact with the terrain in the event an end effector comes free unexpectedly. Third, the robot should move along a desired path provided by a teleoperator or high-level planning algorithm. If the robot has any degrees of freedom remaining after meeting the previous objectives then shape constraints can be applied to improve the robot's mobility over longer time horizons. Finally, the robot should comply with any safety constraints on the output velocity. In this section we present a unified control scheme, SCFM control [89], to simultaneously satisfy all five of these objectives.

1) *Force Control*: We apply a linear force feedback control policy,

$$\dot{\tilde{f}}_c = k_p P_{\mathcal{N}}(f_c^* - \hat{f}_c) \quad (5.19)$$

where k_p is a proportional gain, \tilde{f}_c is the desired rate of change of the contact forces, f_c^* is the optimal contact forces computed by the linear program, and \hat{f}_c is the estimated forces. We use $P_{\mathcal{N}}$ to abbreviate $P_{\mathcal{N}(G_s S_C)}$, the null space projection of the product $G_s S_C$,

$$P_{\mathcal{N}(G_s S_C)} = I - (G_s S_C)^\dagger (G_s S_C) \quad (5.20)$$

where † indicates the Moore-Penrose pseudoinverse. Projecting into $\mathcal{N}(G_s S_C)$ ensures that we only regulate internal forces, preventing accumulated drift in the body pose. If $P_{\mathcal{N}} \dot{\tilde{f}}_c = P_{\mathcal{N}} \dot{\hat{f}}_c$ and $P_{\mathcal{N}} \hat{f}_c = P_{\mathcal{N}} f_c$, then the control law (5.19) will ensure exponential convergence of the internal




Figure 5.3: Optimal contact force distributions for Sally during a floor-to-wall and wall-to-floor transition. Forces at each contact are divided in magnetic (red arrows into surface) and contact (blue arrows away from surface) components. The force of gravity is the green originating from the robot's center of mass.

force $P_{\mathcal{N}}f_c$ to $P_{\mathcal{N}}f_c^*$ with no steady-state error.

2) *Contact Control*: The desired end effector velocity is chosen depending on the nominal contact state \mathcal{C}^* and observed contact state \mathcal{C} . We select each desired contact velocity $\tilde{v}_{c_i}^*$ according to three cases,

$$\tilde{v}_{c_i} = \begin{cases} 0 & i \in \mathcal{C} \\ v_{in} & i \notin \mathcal{C} \wedge i \in \mathcal{C}^* \\ v_{c_i}^* & i \notin \mathcal{C} \wedge i \notin \mathcal{C}^* \end{cases} \quad (5.21)$$

Due to the complementarity of force and position, an end effector velocity's cannot be controlled while producing a contact force along the same axis (case 1). However, if the end effector loses contact unexpectedly, the robot can no longer apply forces and should instead move the contact towards the surface with a preset velocity v_{in} until contact is regained (case 2). Finally, when an end effector deliberately breaks contact with the terrain (case 3), it can instead follow a desired velocity $v_{c_i}^*$ provided by the planner (e.g. a swing leg trajectory).

3) *Motion Control*: In addition to the contact force and contact velocity constraints above, we aim to track a desired body twist v_b^* . We select the following combined control policy,

$$J_h \tilde{\theta} = -K^{-1} \tilde{f}_c + G_s^T v_b^* + \tilde{v}_c \quad (5.22)$$

which has terms for the desired force derivative, body twist and contact velocity. To evaluate the effects of this policy, we substitute (5.22) into (5.5) to obtain the resulting force derivative,

$$\begin{aligned} \dot{f}_c &= K S_{\mathcal{C}} A \dot{q} \\ &= K S_{\mathcal{C}} (-J_h \dot{\theta} + G_s^T v_b) \\ &= K S_{\mathcal{C}} (K^{-1} \tilde{f}_c - G_s^T v_b^* - \tilde{v}_c + G_s^T v_b) \end{aligned} \quad (5.23)$$

We assume that the stiffnesses of the end effectors are decoupled, so K has a block diagonal structure to match S_C and $K S_C = S_C K$. We also observe that $S_C \tilde{v}_c = 0$, because $S_{C_i} = 0$ for $i \notin \mathcal{C}$ and $\tilde{v}_{c_i} = 0$ for $i \in \mathcal{C}$. These properties simplify (5.23),

$$\dot{f}_c = S_C \tilde{f}_c - K S_C G_s^T (v_b^* - v_b) \quad (5.24)$$

The derivative of the body wrench is given by,

$$\begin{aligned} \dot{f}_b &= G_s \dot{f}_c \\ &= G_s S_C \tilde{f}_c - G_s K S_C G_s^T (v_b^* - v_b) \\ &= -G_s K S_C G_s^T (v_b^* - v_b) \end{aligned} \quad (5.25)$$

where we substitute in (5.24) and recall that $\tilde{f}_c \in \mathcal{N}(G_s S_C)$. We assume that $\dot{f}_b = 0$ because $f_b = f_{ext}$ is constant, so (5.25) will be satisfied when $v_b = v_b^*$. Finally, substituting this value of v_b into (5.24), we find,

$$\dot{f}_c = S_C \tilde{f}_c \quad (5.26)$$

Thus we see that (5.22) will correctly track the body twist while following the force control policy in the null space $\mathcal{N}(G_s S_C)$.

4) *Shape Control*: If the Jacobian matrix is square and full rank, then (5.22) can be solved directly for the desired joint velocities. However, for an overactuated robot with $n > m$ there are many possible solutions that produce the same end effector velocities. The ambiguity can be resolved by introducing linear shape constraints of the form, $A_s \theta = b_s$. The simplest constraint is fixing an angle to a constant value, but more complicated constraints like symmetry are also supported. We apply a proportional feedback law to correct for drift in the shape joints, giving us the full SCFM control law which can be solved for $\tilde{\theta}$ provided the leading matrix is full rank,

$$\begin{bmatrix} J_h \\ A_s \end{bmatrix} \tilde{\theta} = \begin{bmatrix} -K^{-1} k_p P_{\mathcal{N}}(f_c^* - \hat{f}_c) + G_s^T v_b^* + \tilde{v}_c \\ -k_s (A_s \theta - b_s) \end{bmatrix} \quad (5.27)$$

5) *Safety Constraints*: After computing desired joint velocities $\tilde{\theta}$, we may wish to limit the maximum velocity output of the robot, for instance to respect motor speed limits, avoid dynamic effects, or provide ice screws sufficient time to melt into the surrounding ice. When the robot is operating in a quasi-static regime, we can simply scale the joint velocity outputs by a constant α to reduce our speed without otherwise affecting the robot's trajectory. We consider safety constraints of the form $\|A_{lim} \dot{\theta}\|_{\infty} \leq v_{max}$ and scale the output joint velocities $\dot{\theta}$ to ensure that the constraints are never exceeded,

$$\alpha = \frac{1}{\max \left(1, \frac{v_{max}}{\|A_{lim} \tilde{\theta}\|_{\infty}} \right)} \quad (5.28)$$

$$\dot{\theta} = \alpha \tilde{\theta} \quad (5.29)$$

5.6 Discussion

The estimation, optimization, and control approach described in this chapter has several advantages. Solving the two linear equations and linear program can be done in milliseconds on modern computers, enabling high-frequency control updates that can rapidly respond to changing conditions. Additionally, this approach can be generalized to a wide variety of robot morphologies and adhesion mechanisms. Underactuated robots optimize contact forces as best they can without control of the null space, while overactuated robots can implement shape constraints to achieve additional objectives. Finally, this approach relies primarily on motor current sensing, making it widely applicable especially for smaller robots.

The contact force optimization procedure can also be used as a tool when developing the design and gait of a new robot, enabling a designer to quickly estimate the expected motor torques and adhesion margin for a given morphology and climbing posture, or evaluate trade-offs between potential end effector designs. Another application is in path planning, to quickly estimate the adhesion margin after placing a gripper at a certain location. For this chapter, we have assumed that the desired body and contact velocities are provided by some external source, but to achieve full autonomy the robot must be able to select a route on its own. In Chapter 6 we explore the highest level of the proposed hierarchy – the path planner.

Chapter 6

Hybrid Motion Planning on Irregular Surfaces

6.1 Introduction

To climb autonomously, a robot must find a trajectory that is kinematically feasible and dynamically stable at each time step. As most climbing robots avoid dynamic motions, the problem can often be reduced to finding a path that is kinematically feasible and statically stable at every point. An added complication is that climbing robots are typically hybrid systems that make and break contact with the surface as they climb. This hybrid property divides the robot's state space into discrete contact modes, preventing the robot from simply planning a continuous path through state space.

Many researchers have developed motion planning approaches for legged robots on relatively level terrain. The simplest approach is to select each foothold independently based on a traversability function evaluated at each point in a small neighborhood of a nominal foothold location, as in [93]. Other researchers use an optimization-based approach, to simultaneously ensure kinematic feasibility and dynamic stability along an entire trajectory. In [94], researchers fit convex planar regions to the terrain and optimize each foothold location within the closest planar region. However, this approach can only find footholds within that small region, and for highly irregular surfaces like natural cliff faces the planar regions would have to become even smaller or cease to accurately approximate the surface. Another approach is to find a sequence of statically stable keyframes and interpolate a series of steps between each pair of neighboring keyframes, as in [95]. This enables the planner to quickly find an optimal sequence of keyframes while only connecting keyframes once a feasible path between them is found.

Robots traversing level ground typically approximate the terrain with a 2.5-dimensional (2.5D) height map, but this representation cannot capture the 3D features that are encountered by climbing robots. Researchers in [96] present an algorithm that instead plans directly on the 3D point cloud acquired by their robot's sensors. This enables the planner to reason in full 3D

without requiring any preprocessing of the sensor data.

In this chapter, we present a hybrid path planning approach for the LORIS robot to autonomously climb on irregular terrain. We present a local planning procedure to select footholds that are simultaneously reachable and graspable while maximizing forward progress and static stability. We also present problem-specific modifications to the AIT* algorithm to plan in the reduced state space of the constraint manifold while delegating foothold selection to the local planner. The proposed planning approach is demonstrated both in simulation and on hardware with the LORIS climbing robot.

6.2 Related Work

Informed graph-search planning algorithms such as A* [97] search for a path between two nodes on a graph, using a problem-specific heuristic to explore more promising routes first and quickly find the optimal solution. Incremental sampling-based planning algorithms, like RRT* [98], instead randomly sample points in the state space and connect them with nearby points to gradually build up a graph structure. The AIT* algorithm [22] combines these approaches by constructing a graph from random samples and exploring it using graph-search techniques.

A second area of research is the sampling of points on a constraint manifold embedded in a higher dimensional space. This is applicable for finding foothold locations that lie on the terrain surface. In [99], researchers enumerate several possible approaches to this task: relaxation of the constraints, projection of sampled points to the nearest point on the manifold, sampling points from the tangent space of the manifold at a nearby point, sampling points from an atlas that piece-wise linearly approximates the manifold, and reparameterizing the configuration space to eliminate the constraints. The researchers in [96] use the projection approach to map 3D sampled points to the nearest point in the point cloud.

Another area of related research is in motion planning for hybrid systems. In [100], the climbing problem is explored for a discrete set of possible footholds. Each subset of footholds defines a stance manifold, representing all possible robot states that maintain the same contact mode. The problem is then broken down into finding feasible paths between adjacent stance manifolds, and finding a sequence of connected stance manifolds between the start and goal configurations. In [101], researchers present an integrated task and motion planning (TAMP) approach for manipulation problems. They examine a class of algorithms which alternately sample new points within an orbit (an equivalent term for a stance manifold) and sample new points on the boundary of an orbit to define new orbits, gradually building up a roadmap that connects orbits through transition states with guaranteed asymptotic optimality. In contrast, the TMIT* algorithm [102] samples points within each mode (an equivalent term for an orbit), then projects any samples within a tunable distance threshold of a mode boundary onto the boundary to define a new mode. This projection-based approach is more computationally efficient than directly sampling points on the mode boundary.

6.3 Problem Setup

We divide the planning problem into two subtasks: local and global planning. The local planner is responsible for selecting individual footholds as the robot moves towards a nearby node on the constraint manifold. In contrast, the global planner searches for a sequence of these nodes connecting the start to the goal pose. This division of responsibility is intended to reduce the computation required: the global planner only needs to explore the 2D reduced space of the constraint manifold, while the local planner can make relatively greedy selections while planning over short distances in the full configuration space.

We first define the configuration space q consisting of the body pose $x_b \in SE(3)$ and the foothold positions $r_{c_i} \in \mathbb{R}^3$ for $i \in \mathcal{C}$. The body pose can be decomposed into the position r_b and rotation R_{wb} , or represented by the transform g_{wb} . Overactuated robots have additional degrees of freedom not captured by this representation, so we define these extra shape joint angles as θ_s . The full state is,

$$q = \begin{bmatrix} x_b \\ r_{c_1} \\ \vdots \\ r_{c_n} \\ \theta_s \end{bmatrix} \quad (6.1)$$

The climbing surface is defined as a 2D constraint manifold \mathcal{T} embedded in \mathbb{R}^3 . In practice this is approximated by a point cloud $P \subset \mathcal{T}$. Unlike a heightmap, this representation can capture 3D features such as overhangs. Another alternative would be a volumetric representation such as a voxel grid or octree [103] which can similarly capture these features. For our purposes, we can use any representation which supports two operations: projection of a point onto the surface, and sampling of points within a radius of a query point. These can be computed efficiently for a point cloud using a K-D tree; we use the implementation in PCL [104].

We also define the workspace $\mathcal{W}_i(\theta_s)$ of each end effector as a convex set of reachable points in the body frame, parameterized by the current shape joint angles. One potential representation for such a workspace is a linearly constrained set, or equivalently a polytope in \mathcal{H} -representation [105], of the following form,

$$\mathcal{W}_i = \{r \in \mathbb{R}^3 \mid Q_i r \leq q_i\} \quad (6.2)$$

Workspaces can be combined with the Minkowski addition operator \oplus . Several algorithms exist for the Minkowski addition of polytopes [105],

$$S_1 \oplus S_2 = \{p + q \mid p \in S_1, q \in S_2\} \quad (6.3)$$

6.4 Local Planning

The local planner’s task is to select a new foothold given the current stance that brings the robot closer to the goal node; this can be repeated recursively to generate a sequence of steps. Our approach is to 1) sample a set of nearby candidate footholds and reject any that are unreachable or ungraspable, 2) rank the candidates according to a cost heuristic, 3) select the corresponding body pose and shape joint angles for the best candidate, and 4) evaluate the risk of adhesion failure. If a sample is rejected during step 3 or 4, the algorithm tries again with the next best remaining candidate. The steps are ordered by increasing computational complexity, acting as a sieve to filter out more footholds at each step and save the expensive operations for the most promising foothold candidates.

6.4.1 Foothold Rejection Sampling

We begin by defining the reachability criterion. Given an initial state q , goal pose x_{goal} , and candidate foothold $\tilde{r}_{c_i} \in \mathcal{T}$, there must be some new body position \tilde{r}_b that satisfies the following conditions,

$$\tilde{r}_{c_i} \in R_{wb}\mathcal{W}_i + \tilde{r}_b \quad (6.4)$$

$$\tilde{r}_b \in r_{c_j} - R_{wb}\mathcal{W}_j, \quad j \in \mathcal{C}, j \neq i \quad (6.5)$$

(6.4) requires that the candidate foothold is reachable from the new body pose, while (6.5) ensures that the new body pose can still reach the remaining stance footholds. (6.4)–(6.5) can be combined to eliminate \tilde{x}_b ,

$$\tilde{r}_{c_i} \in R_{wb}(\mathcal{W}_i \oplus -\mathcal{W}_j) + r_{c_j}, \quad j \in \mathcal{C}, j \neq i \quad (6.6)$$

These Minkowski sums can be precomputed for a discretized set of shape joint angles to save time during planning. We further require that the centroid of the footholds \bar{r}_c moves closer to the goal by some minimum distance δ_{min} to ensure that the planner is making forward progress,

$$\|r_{goal} - \tilde{r}_c\| \leq \max(\|r_{goal} - \bar{r}_c\| - \delta_{min}, 0) \quad (6.7)$$

We also filter out ungraspable footholds at this stage by constraining the mean curvature H of the terrain in a local neighborhood of each foothold candidate (computed via principle component analysis of the nearby point normals),

$$H_{min} \leq H(\tilde{r}_{c_i}) \leq H_{max} \quad (6.8)$$

A symmetric bound on H ensures the surface is sufficiently flat for the gripper to fully engage. Alternatively, a negatively biased bound will accept convex features while rejecting concave features.

At this point, we can sample a set of points $p_k \sim \mathcal{T}$ within some radius of the swing foot and reject any that violate constraints (6.6)–(6.8). If no points remain, then the local planner reports a failure and the global planner attempts to find a different route. Otherwise, we are typically left with several viable foothold candidates.

6.4.2 Foothold Cost Heuristic

We can now rank the remaining foothold candidates according to the following cost heuristic,

$$J_{foot}(r_{c_i}) = -\text{dist}_{gwb\mathcal{W}_i}(\tilde{r}_{c_i}, r_b - r_{goal}) + \alpha(\hat{n}_{c_i} \cdot \hat{g}) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2) \quad (6.9)$$

where we define the distance operator, $\text{dist}_{\mathcal{W}}(r, v)$, as the distance from a point r inside a polytope \mathcal{W} to the boundary of the polytope along a given direction vector v ,

$$\text{dist}_{\mathcal{W}}(r, v) = \max\{s \in \mathbb{R} \mid r + \frac{sv}{\|v\|} \in \mathcal{W}\} \quad (6.10)$$

The first term in (6.9) represents the distance from the foothold candidate to the edge of its workspace along the opposite direction from the goal. This is effectively the distance that the body can travel in a straight line towards the goal before the same foot needs to take another step. The second term is the dot product of the foothold’s local normal vector \hat{n}_{c_i} and the gravity vector \hat{g} , which represents the steepness of the incline. Lastly, the ϵ term is a random disturbance that adds stochasticity to the selection process, so that repeated attempts do not always select the same foothold. The incline weight α and the variance σ^2 are tuning parameters to trade off between forward progress, foothold quality, and selection variety.

6.4.3 Body Pose Selection

Given the selected footholds, we now determine the corresponding body pose and shape joint angles. To simplify this section we drop the tilde notation, so all variables refer to their values at the new state. We choose the body normal vector \hat{n}_b to align with the best-fit plane for the set of footholds, computed by singular value decomposition (SVD). Specifically, we compute the SVD of the mean-centered matrix of footholds, $r_c - \bar{r}_c$, and align the normal vector with the singular vector that has the smallest singular value, $V_{:,3}$, which represents the direction of least variance in the foothold positions,

$$r_c = \begin{bmatrix} r_{c_1} & \dots & r_{c_n} \end{bmatrix} \quad (6.11)$$

$$(r_c - \bar{r}_c)(r_c - \bar{r}_c)^T = U\Sigma V^T \quad (6.12)$$

$$\hat{n}_b = V_{:,3} \quad (6.13)$$

This still leaves one free degree of freedom for the heading ψ , which can be set according to the following feedback policy to converge toward a desired heading ψ^* with a maximum change

per step of ψ_{step} ,

$$\psi' = \psi + \min(\max(\psi^* - \psi, -\psi_{step}), \psi_{step}) \quad (6.14)$$

One option is to set ψ^* to the direction of travel, for robots that prefer to move in their local forward direction. Another alternative is to align ψ^* with the negative gravity vector. This option is ideal for LORIS, ensuring that the tail remains oriented downward while climbing.

Next, we set the shape joint angles to match the incline of the terrain. For LORIS, which features only a spine pitch joint, we average the normal vectors of the front and rear foot pairs and compute the angle between them in the normal plane to the spine pitch joint axis \hat{y} using (6.15),

$$\theta_s = \sin^{-1}((\hat{n}_{front} \times \hat{n}_{rear}) \cdot \hat{y}) \quad (6.15)$$

We use the centroid of the footholds \bar{r}_c as our initial guess for the body position x_b . We then project the centroid upwards along \hat{n}_b until it is clear of the terrain and lies in the body workspace \mathcal{W}_{body} , defined as the set of body positions from which all of the footholds remain reachable,

$$\mathcal{W}_{body} = \bigcap_i (r_{c_i} - R_{wb}\mathcal{W}_i), \quad i \in \mathcal{C} \quad (6.16)$$

While the initial rejection sampling aims to eliminate any unreachable footholds, it neglects terrain clearance and changes in body orientation and shape joint angle. These changes are typically small, but can potentially result in an empty body workspace. Similarly, for overextended steps the body workspace may not lie above the foothold centroid even if it is non-empty. If the projection along the normal vector fails for these reasons, then the local planner simply tries again with the next best foothold. At this point we also evaluate the joint angles via inverse kinematics to ensure that the angle between each end effector and the surface lies within an acceptable range. Limb collision checking can be performed here as well.

6.4.4 Adhesion Validation

Lastly, given the body position and stance we evaluate the adhesion margin c using the same force optimization setup from Chapter 5. In the optimization, we assume that the swing foot is being pressed into the surface to engage the gripper. This tends to be the most failure-prone stage of the step because the engagement force pushes the robot away from the surface, and the remaining grippers must provide greater adhesive forces to compensate.

If the margin is below a minimum threshold, then we assume that the step is invalid for any foothold with the same swing foot, not just the current candidate. While the position of the body and swing end effector can have a minor influence on the solution, it is unlikely to raise the adhesion margin above the threshold except in very borderline cases. If all of the potential swing feet have resulted in a failed solution, then we flag this state as a dead end and return

failure.

6.5 Global Planning

The local planning approach described above works well over short distances, but its greedy, one-step-at-a-time approach leads to suboptimal routes in the long run. The global planner instead leverages sampling-based planning to find optimal routes in the reduced space of points $p \in \mathbb{R}^3$ on the constraint manifold \mathcal{T} . Traditional sampling-based algorithms like RRT* rely on a distance metric to select the nearest point on the search tree to a given sample. A significant challenge of planning on a manifold, rather than a flat surface, is that Euclidean distance is no longer an accurate measure of the cost to move between two points; the shortest path along a manifold is instead given by a geodesic curve. The AIT* algorithm [22] in particular presents an alternative: computing a cost-to-go heuristic from each sample point to the goal along a reverse search tree.

In more detail, the AIT* algorithm alternates between three steps: sampling, reverse search, and forward search. In the sampling stage, the algorithm samples a batch of new points to form an implicit graph, with neighboring nodes sharing an edge. The reverse search grows a search tree along this graph from the goal to each sample point using a variant of A*. The computationally expensive edge checks are not yet evaluated; edge costs are instead estimated with an admissible heuristic, so the cost-to-go estimate from a point to the goal along this graph is a lower bound on the true cost. In contrast, the forward search uses the cost-to-go estimate from the reverse search as a heuristic to grow a tree from the start to the goal using A*. Unlike the reverse search, the forward search actually evaluates each edge of the tree, ensuring that the resulting path is indeed feasible. To summarize, the reverse search computes a cost-to-go heuristic that the forward search uses to find an optimal path.

We can use this same approach to compute geodesic distance on the constraint manifold: if we bound the Euclidean distance between any two neighbors to at most d , then the entire search tree is constrained to remain within distance $\frac{d}{2}$ of \mathcal{T} . As d approaches zero this approximate cost-to-go converges to an exact estimate of the geodesic distance to the goal. We can also define a cost function over the surface of the manifold, such that the cost-to-go represents the integral of the cost along the optimal path to the goal. We select the following cost function, which contains terms for both incline (weighted by a tuning parameter β and offset by 1 to be non-negative) and Euclidean distance,

$$J_{path} = \int (\beta(\hat{n}(s) \cdot \hat{g} + 1) + 1) ds \quad (6.17)$$

We use the OMPL [106] implementation of AIT* as the foundation for our global planner to optimize the cost function defined in (6.17), with the following modifications.

1. We sample points directly from the point cloud, rather than the full 3D state space. This reduces the dimensionality of the sample space and ensures that all samples lie on the

constraint manifold.

2. When the forward search explores an edge using the local planner, we record the full robot state at the destination node to use as the initial state for future edge checks originating at that node.
3. Multiple subsequent steps with the same foot are combined into a single step to eliminate wasted movement.

6.6 Planning Results

The planning algorithm can be executed in real time by LORIS’s onboard Intel NUC 11 computer. The robot’s current pose and a point cloud representation of the terrain are computed by the onboard RealSense depth camera using the RTAB-Map library [107]. The robot can then follow the planned route using the gait described in Chapter 3.

Representative local planner solutions are shown for LORIS on a slag climbing wall in Fig. 6.1. Note that the terrain was mapped prior to placing LORIS on the wall to include the terrain near the rear feet of the robot. We observe that the footholds are evenly spread out, minimizing the number of steps required to complete the path. We also note the consistency of the two solutions in selecting similar footholds (we have set $\sigma^2 = 0$ for this demonstration).

We observe challenges faced by the planner as well. Predictably, the planner tries to place footholds on the smooth wood above the rock face as it has no knowledge of the surface material. More concerning, in Fig. 6.1 right, the second-to-last foothold of the rear right foot lies on a tiny strip of purple (low cost) surrounded by a large red region (high cost). This is an artifact of occlusion caused by the rocks, which resulted in a very sparse region of the cloud where the curvature estimates were very inaccurate. A potential solution to this issue would be to interpolate missing data to fill in such gaps in the cloud.

The slag climbing wall is too small to demonstrate the effectiveness of the global planner on hardware, but Fig. 6.2 shows several example solutions in simulation. The planner was run for 30 seconds in each instance, not including precomputation of the point cloud normal vectors and curvature. Table 6.1 provides several summary statistics for each solution, including the number of points sampled from the point cloud and the number of edges explored by the local planner. The quality of the path itself is evaluated by the number of steps, the length of the body path, and the mean body incline (computed by averaging the incline cost term for the body normal vector over the length of the path and converting back to an equivalent angle).

In the uneven and cliff environments (Fig. 6.2a and 6.2d), the planner shows its ability to handle irregular terrain features at both the foothold and robot scale. The tube environment (Fig. 6.2b) demonstrates the planner choosing a longer, shallower path, rather than the steeper direct route. Examining Table 6.1, we observe that the planner explored only one edge for this environment, because the local planner was able to solve the problem in a single attempt. In the gap environment (Fig. 6.2c), the planner successfully routes around a large gap that lies




Figure 6.1: Local planner solutions for LORIS on a slag climbing wall. Circles represent footholds, small axes represent planned robot poses, and large axes represent the goal pose. The terrain in the right figure is color-coded by the foothold cost heuristic, on a gradient from purple (ideal) to green to red (ungraspable).

Environment	Points Sampled	Edges Explored	Step Count	Path Length	Mean Incline
Uneven	900	44	38	1.10 m	16.8°
Tube	10600	1	106	2.53 m	45.6°
Gap	200	37	87	3.60 m	90.0°
Cliff	400	39	69	2.54 m	70.0°
Column	2100	3	54	1.37 m	90.6°

Table 6.1: Planning Results on Simulated Terrain

between LORIS and the goal pose, something the greedier local planner cannot manage on its own. We also see that the solution closely hugs the obstacle to minimize total path length.

6.7 Discussion

Stepping blindly, the LORIS climbing robot is capable of adapting to minor terrain variations through its gait, but this purely proprioceptive approach does not scale to larger obstacles or enable the robot to navigate through a complex environment. In this chapter, we present a motion planning algorithm that uses exteroceptive data from an onboard depth camera to reason about both local foothold placement and global path-finding. At the local level we use a sieve approach to progressively filter out unreachable, ungraspable, and low-quality footholds while saving computationally intensive checks for only the most promising candidates. At the global level, we build upon the sampling-based planning algorithm AIT* to integrate our constraint

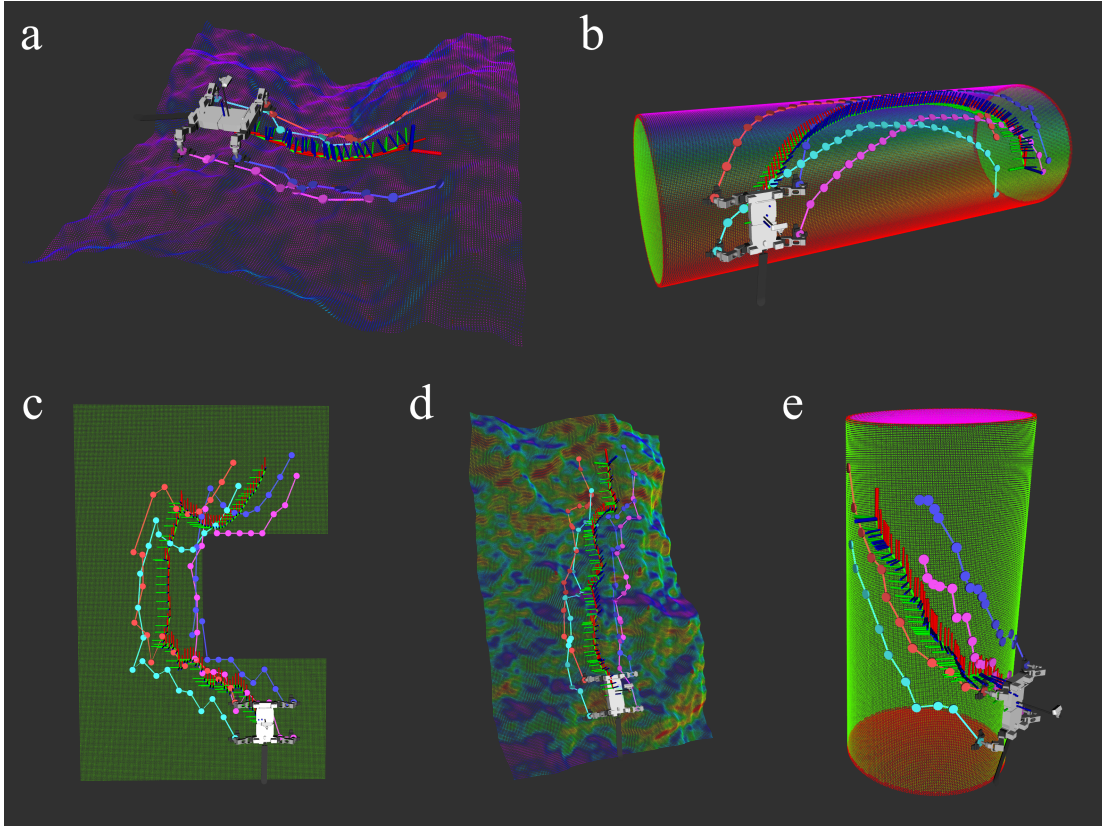


Figure 6.2: Global planner solutions for LORIS in simulated environments: a) uneven, b) tube, c) gap, d) cliff, and e) column. Circles represent footholds, small axes represent planned robot poses, and large axes represent the goal pose. The terrain is color-coded by incline, from purple (0°) to green (90°) to red (135°).

manifold terrain representation and our local footstep planning approach.

In practice, we still see limitations of our planning approach that present opportunities for future development. When applying the planner to real-world sensor data, a major challenge is occlusions; interpolating over gaps in the point cloud could enable the robot to predict the terrain that it cannot yet see. Another area for further study is terrain classification using the color images from the camera. The variety of potential features to identify makes this an ideal case for a learning-based approach that would enable LORIS to identify high quality footholds based on material, not just geometry.

Chapter 7

Conclusion

This dissertation presents new techniques to make climbing robots more lightweight, reliable, and autonomous through innovations in underactuated design, internal force control, and hybrid path planning.

We begin with an additive manufacturing approach for microspine fabrication to replace existing Shape Deposition Manufacturing techniques. Our approach eliminates the need for expensive equipment and detailed expertise, making microspine technology readily available to anyone with a 3D printer. We also reduce manufacturing time by an order of magnitude, enabling designers to rapidly prototype new gripper designs. By lowering the barriers to microspine fabrication, we hope to encourage broader use of this relatively niche technology.

Next, we bridge the gap between small wall-climbing robots and massive fully-actuated climbing robots with LORIS, a lightweight rock-climbing robot that can traverse the irregular terrain encountered in real-world environments. We present a splayed gripper design that overcomes a key limitation of passive grippers: the inability to support lateral force components. Compliant elements in the gripper and wrist are capable of adapting to the local terrain geometry without the need for additional actuation. Climbing experiments with LORIS demonstrate the robot's versatility across both flat walls and uneven rock faces, and the robot serves as a proof of concept for a new paradigm of expendable climbers that could accompany larger rovers on their missions.

We also present Sally, a magnetic-wheeled inspection robot equipped with an X-ray fluorescence sensor to analyze surface composition and a passive suspension to negotiate obstacles. We analyze the difficulty of traversing interior corners with Sally, and find creative ways to reuse Sally's existing steering and sampling mechanisms to aid in the transition process. Sally was also tested in the field to identify the challenges involved in actual deployment, such as the accumulation of magnetized dirt.

In Chapter 5, we present a technique to optimize a robot's internal force distribution and minimize the risk of falling. We first apply this optimization approach to LORIS, discovering the automatic generation of distributed inward gripping forces like we see in natural climbers such as cockroaches and geckos. We then apply the same approach to Sally, revealing a promising

method to automate the complicated procedure that an operator must execute during an corner transition. We also provide a unified control strategy to generate these contact forces while simultaneously recovering from contact loss, tracking the desired body trajectory, and enforcing additional shape constraints and velocity limits. A variation of this unified control strategy was used to great success by NASA Jet Propulsion Laboratory’s EELS 1.5 robot to descend a vertical ice shaft.

Finally, we develop a motion planning approach that overcomes the coupled challenges of hybrid contact and three-dimensional terrain features faced by legged climbing robots such as LORIS. We present a high-level global planner that applies the sampling-based AIT* algorithm to rapidly search for paths along arbitrarily shaped terrain manifolds modeled by point clouds. The global planner relies on a local footstep planner to select footholds that are reachable, graspable, and nearer to the robot’s destination. The local planner also reuses the contact force optimization approach from above to evaluate the robot’s stability. We analyze planned paths both in simulation and on hardware to evaluate the planner’s performance.

Looking at the above work, these contributions fit neatly into the hierarchical framework introduced in Chapter 1. We see the advantages of underactuation at the microspine, gripper, and robot scales to enable passive terrain adaptation while reducing weight. On the control level, optimization of internal forces produces improvements in adhesion across wildly different adhesion methods and robot designs, increasing reliability and payload size. Lastly, a planning approach adapted for the 3D environments encountered by climbing robots provides autonomy, freeing space exploration robots from speed-of-light delays and enabling the automation of difficult and dangerous tasks on Earth. We believe these lessons are broadly applicable, even beyond the specific robots and examples used in this document.

7.1 Future Directions

In fact, one promising avenue for future research is to apply the techniques discussed here to conventional quadrupeds, enabling them to operate on steep and irregular terrain. We see impressive feats of climbing from animals such as goats with a similar morphology to today’s robot dogs, which suggests that new developments like underactuated hoof-inspired grippers could enable these robots to follow suit. Similarly, quadrupeds using internal force control could potentially accomplish tasks like tree climbing or two-wall climbing with only minimal modification to their end effectors, purely by generating internal preload forces with their limbs. A hybrid planning algorithm for 3D environments could also be just as applicable for searching a damaged multi-story structure or underground mine as it is for exploring a lava tube on Mars.

By the same token, can climbing robots take advantage of the dynamic control strategies used by their ground-bound counterparts? Climbing robots tend to be slow and cautious due to the difficulty of their task, but another reason is often the time required to engage and disengage their grippers with the surface. A key benefit of passive grippers like those of LORIS is that no such engagement time is required; the gripper engages just by being pressed against the surface at the right angle. This opens the door to a wide range of useful dynamic maneuvers –

recovering from a fall, lunging for an out-of-reach foothold, or even using a more energy-efficient dynamic gait. As both legged ground robots and climbing robots grow more capable, it will be fascinating to see how they can learn from each other.

We conclude that small, expendable climbing robots present the most realistic path forward for the exploration of steep terrain features in extraterrestrial environments. In this work, we endeavored to expand the capabilities of such robots to navigate irregular terrain, but many difficulties remain. Several of these are typical challenges that must be addressed to operate in space environments, including dust, thermal, power, radiation, and vacuum considerations. But more fundamentally, these small climbing robots will need much greater reliability than research prototypes such as LORIS have been able to demonstrate. We conclude that the most important next step is to increase the reliability of these robots, whether by adding more end effectors for redundancy, implementing better control strategies for dynamic slip recovery, or using less conventional techniques like grappling hooks or grippers on extendable booms. Even so, our work has provided a proof-of-concept that this paradigm of lightweight climbing robots can ultimately enable steep terrain exploration throughout the solar system.

Bibliography

- [1] A. Parness, N. Abcouwer, C. Fuller *et al.*, “LEMUR 3: A limbed climbing robot for extreme terrain mobility in space,” in *IEEE International Conference on Robotics and Automation*, 2017, pp. 5467–5473.
- [2] A. S. McEwen, L. Ojha, C. M. Dundas *et al.*, “Seasonal flows on warm martian slopes,” *Science*, vol. 333, no. 6043, pp. 740–743, 2011.
- [3] R. J. L  veill   and S. Datta, “Lava tubes and basaltic caves as astrobiological targets on Earth and Mars: A review,” *Planetary and Space Science*, vol. 58, no. 4, pp. 592–598, 2010.
- [4] I. A. Nesnas, L. Kerber, A. Parness *et al.*, “Moon Diver: a discovery mission concept for understanding the history of secondary crusts through the exploration of a lunar mare pit,” in *IEEE Aerospace Conference*, 2019, pp. 1–23.
- [5] T. S. Vaquero, G. Daddi, R. Thakker *et al.*, “EELS: Autonomous snake-like robot with task and motion planning capabilities for ice world exploration,” *Science Robotics*, vol. 9, no. 88, p. eadh8332, 2024.
- [6] A. Sirken, G. Knizhnik, J. McWilliams, and S. Bergbreiter, “Bridge risk investigation diagnostic grouped exploratory (BRIDGE) bot,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 6526–6532.
- [7] P. A. Bogdan, J. Wheadon, F. B. Klein, and M. Gianni, “Magnetic tracked robot for internal pipe inspection,” in *European Conference on Mobile Robots*, 2021, pp. 1–6.
- [8] A. Shukla and H. Karki, “Application of robotics in onshore oil and gas industry—a review part I,” *Robotics and Autonomous Systems*, vol. 75, pp. 490–507, 2016.
- [9] F. Bonnin-Pascual and A. Ortiz, “On the use of robots and vision technologies for the inspection of vessels: A survey on recent advances,” *Ocean Engineering*, vol. 190, p. 106420, 2019.
- [10] M. J. Spenko, G. C. Haynes, J. A. Saunders *et al.*, “Biologically inspired climbing with a hexapedal robot,” *Journal of Field Robotics*, vol. 25, no. 4-5, pp. 223–242, 2008.

- [11] A. T. Asbeck, S. Kim, W. R. Provancher *et al.*, “Scaling hard vertical surfaces with compliant microspine arrays,” *Robotics Science and Systems Conference*, 2005.
- [12] S. Kim, M. Spenko, S. Trujillo *et al.*, “Smooth vertical surface climbing with directional adhesion,” *IEEE Transactions on Robotics*, vol. 24, no. 1, pp. 65–74, 2008.
- [13] O. Kermorgant, “A magnetic climbing robot to perform autonomous welding in the ship-building industry,” *Robotics and Computer-Integrated Manufacturing*, vol. 53, pp. 178–186, 2018.
- [14] J. Duncombe, “Meet IceWorm: NASA’s new ice-climbing robot,” *Eos*, 2018.
- [15] M. Spenko, “Making contact: A review of robotic attachment mechanisms for extraterrestrial applications,” *Advanced Intelligent Systems*, p. 2100063, 2022.
- [16] P. Nadan, S. Backus, and A. M. Johnson, “LORIS: A lightweight free-climbing robot for extreme terrain exploration,” in *IEEE International Conference on Robotics and Automation*, 2024, accepted for publication.
- [17] P. Nadan, D. K. Patel, C. Pavlov *et al.*, “Microspine design for additive manufacturing,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 5640–5647.
- [18] P. Nadan, J. Kumar, N. Klein *et al.*, “A magnetic-wheeled inspection robot for interior corner traversal,” in prep.
- [19] D. I. Goldman, T. S. Chen, D. M. Dudek, and R. J. Full, “Dynamics of rapid vertical climbing in cockroaches reveals a template,” *Journal of Experimental Biology*, vol. 209, no. 15, pp. 2990–3000, 2006.
- [20] L. R. Palmer III, E. Diller, and R. D. Quinn, “Toward gravity-independent climbing using a biologically inspired distributed inward gripping strategy,” *IEEE/ASME Transactions on Mechatronics*, vol. 20, no. 2, pp. 631–640, 2015.
- [21] R. Thakker, M. Paton, M. P. Strub *et al.*, “EELS: Towards autonomous mobility in extreme terrain with a versatile snake robot with resilience to exteroception failures,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023, pp. 9886–9893.
- [22] M. P. Strub and J. D. Gammell, “Adaptively informed trees (AIT*) and effort informed trees (EIT*): Asymmetric bidirectional sampling-based path planning,” *The International Journal of Robotics Research*, vol. 41, no. 4, pp. 390–417, 2022.
- [23] P. Nadan, S. Backus, and A. M. Johnson, “Autonomous climbing on irregular terrain with passive microspine grippers,” in prep.

- [24] A. Parness, M. Frost, N. Thatte *et al.*, “Gravity-independent rock-climbing robot and a sample acquisition tool with microspine grippers,” *Journal of Field Robotics*, vol. 30, no. 6, pp. 897–915, 2013.
- [25] S. Backus, J. Izraelevitz, J. Quan *et al.*, “Design and testing of an ultra-light weight perching system for sloped or vertical rough surfaces on Mars,” in *IEEE Aerospace Conference*, 2020, pp. 1–12.
- [26] K. Carpenter, N. Wiltsie, and A. Parness, “Rotary microspine rough surface mobility,” *IEEE/ASME Transactions on Mechatronics*, vol. 21, no. 5, pp. 2378–2390, 2016.
- [27] A. Asbeck, S. Kim, A. McClung, and A. Parness, “Climbing walls with microspines,” in *IEEE International Conference on Robotics and Automation*, 2006.
- [28] T. Wallin, J. Pikul, and R. Shepherd, “3D printing of soft robotic systems,” *Nature Reviews Materials*, vol. 3, no. 6, p. 84, 2018.
- [29] R. Merz, F. Prinz, K. Ramaswami *et al.*, “Shape deposition manufacturing,” in *International Solid Freeform Fabrication Symposium*, 1994.
- [30] L. Weiss, R. Merz, F. Prinz *et al.*, “Shape deposition manufacturing of heterogeneous structures,” *Journal of Manufacturing Systems*, vol. 16, no. 4, pp. 239–248, 1997.
- [31] S. Y. Wu, C. Yang, W. Hsu, and L. Lin, “3D-printed microelectronics for integrated circuitry and passive wireless sensors,” *Microsystems and Nanoengineering*, vol. 1, 2015.
- [32] T. D. Ngo, A. Kashani, G. Imbalzano *et al.*, “Additive manufacturing (3D printing): A review of materials, methods, applications and challenges,” *Composites Part B: Engineering*, vol. 143, pp. 172–196, 2018.
- [33] R. H. Sanatgar, A. Cayla, C. Campagne, and V. Nierstrasz, “Morphological and electrical characterization of conductive polylactic acid based nanocomposite before and after FDM 3D printing,” *Journal of Applied Polymer Science*, vol. 136, no. 6, p. 47040, 2019.
- [34] C. Tawk, Y. Gao, R. Mutlu, and G. Alici, “Fully 3D printed monolithic soft gripper with high conformal grasping capability,” in *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, 2019, pp. 1139–1144.
- [35] P. R. Kuppens, M. A. Bessa, J. L. Herder, and J. B. Hopkins, “Compliant mechanisms that use static balancing to achieve dramatically different states of stiffness,” *Journal of Mechanisms and Robotics*, vol. 13, no. 2, 01 2021.
- [36] R. R. Ma, J. T. Belter, and A. M. Dollar, “Hybrid deposition manufacturing: Design strategies for multimaterial mechanisms via three-dimensional printing and material deposition,” *Journal of Mechanisms and Robotics*, vol. 7, no. 2, 05 2015.
- [37] E. G. Merriam, A. B. Berg, A. Willig *et al.*, “Microspine gripping mechanism for asteroid capture,” in *Aerospace Mechanisms Symposium*, 2016.

- [38] A. Parness, A. Willig, A. Berg *et al.*, “A microspine tool: Grabbing and anchoring to boulders on the asteroid redirect mission,” in *IEEE Aerospace Conference*, 2017, pp. 1–10.
- [39] A. Parness, T. Evans, W. Raff *et al.*, “Maturing microspine grippers for space applications through test campaigns,” *AIAA SPACE and Astronautics Forum and Exposition*, no. 203999, pp. 1–17, 2017.
- [40] M. Martone, C. Pavlov, A. Zeloof *et al.*, “Enhancing the vertical mobility of a robot hexapod using microspines,” arXiv, Tech. Rep. arXiv:1906.04811 [cs.RO], 2019.
- [41] R. J. Wallace, C. Pavlov, and A. Johnson, “Design of microspine-enhanced spring legs for robotic running and climbing,” in *Dynamic Walking*, May 2020.
- [42] B. D. Miller, P. R. Rivera, J. D. Dickson, and J. E. Clark, “Running up a wall: the role and challenges of dynamic climbing in enhancing multi-modal legged systems,” *Bioinspiration & biomimetics*, vol. 10, no. 2, p. 025005, 2015.
- [43] S. B. Backus, R. Onishi, A. Bocklund *et al.*, “Design and testing of the JPL-Nautilus gripper for deep-ocean geological sampling,” *Journal of Field Robotics*, vol. 37, no. 6, pp. 972–986, 2020.
- [44] S. Wang, H. Jiang, and M. R. Cutkosky, “Design and modeling of linearly-constrained compliant spines for human-scale locomotion on rocky surfaces,” *The International Journal of Robotics Research*, vol. 36, no. 9, pp. 985–999, 2017.
- [45] S. Yang, S. Qiao, and N. Lu, “Elasticity solutions to nonbuckling serpentine ribbons,” *Journal of Applied Mechanics*, vol. 84, no. 2, 11 2016.
- [46] J. B. Hopkins and M. L. Culpepper, “Synthesis of multi-degree of freedom, parallel flexure system concepts via freedom and constraint topology (FACT)–Part I: Principles,” *Precision Engineering*, vol. 34, no. 2, pp. 259–270, 2010.
- [47] H. Somarathna, S. Raman, D. Mohotti *et al.*, “Hyper-viscoelastic constitutive models for predicting the material behavior of polyurethane under varying strain rates and uniaxial tensile loading,” *Construction and Building Materials*, vol. 236, p. 117417, 2020.
- [48] S. Awtar, A. H. Slocum, and E. Sevincer, “Characteristics of beam-based flexure modules,” *Journal of Mechanical Design*, vol. 129, no. 6, pp. 625–639, 05 2006.
- [49] I. Baldwin, S. Kenig, A. Nicholas *et al.*, “Extreme science: Exploring the use of extreme-terrain rovers in Mars sample return,” in *IEEE Aerospace Conference*, 2016, pp. 1–11.
- [50] S. W. Squyres, A. H. Knoll, R. E. Arvidson *et al.*, “Exploration of Victoria Crater by the Mars rover Opportunity,” *Science*, vol. 324, no. 5930, pp. 1058–1061, 2009.
- [51] D. Schmidt and K. Berns, “Climbing robots for maintenance and inspections of vertical structures—a survey of design aspects and technologies,” *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1288–1305, 2013.

- [52] P. Zi, K. Xu, Y. Tian, and X. Ding, “A mechanical adhesive gripper inspired by beetle claw for a rock climbing robot,” *Mechanism and Machine Theory*, vol. 181, p. 105168, 2023.
- [53] X. Li, W. Chen, X. Li *et al.*, “An underactuated adaptive microspines gripper for rough wall,” *Biomimetics*, vol. 8, no. 1, 2023.
- [54] K. Uno, N. Takada, T. Okawara *et al.*, “HubRobo: A lightweight multi-limbed climbing robot for exploration in challenging terrain,” in *IEEE-RAS International Conference on Humanoid Robots*, 2021, pp. 209–215.
- [55] Y. Tanaka, Y. Shirai, X. Lin *et al.*, “SCALER: A tough versatile quadruped free-climber robot,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 5632–5639.
- [56] W. Ruotolo, F. S. Roig, and M. R. Cutkosky, “Load-sharing in soft and spiny paws for a large climbing robot,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1439–1446, 2019.
- [57] G. C. Haynes, A. Khripin, G. Lynch *et al.*, “Rapid pole climbing with a quadrupedal robot,” in *IEEE International Conference on Robotics and Automation*, 2009, pp. 2767–2772.
- [58] K. A. Daltorio, T. C. Witushynsky, G. D. Wile *et al.*, “A body joint improves vertical to horizontal transitions of a wall-climbing robot,” in *IEEE International Conference on Robotics and Automation*, 2008, pp. 3046–3051.
- [59] A. T. Asbeck and M. R. Cutkosky, “Designing compliant spine mechanisms for climbing,” *Journal of Mechanisms and Robotics*, vol. 4, p. 031007, 2012.
- [60] A. Sintov, T. Avramovich, and A. Shapiro, “Design and motion planning of an autonomous climbing robot with claws,” *Robotics and Autonomous Systems*, vol. 59, no. 11, pp. 1008–1019, 2011.
- [61] V. Nava, N. Sihota, T. Hoelen *et al.*, “Portable X-ray fluorescence for autonomous in-situ characterization of chloride in oil and gas waste,” *Environmental Pollution*, vol. 316, p. 120558, 2023.
- [62] J. Norby, S. Wang, H. Wang *et al.*, “Path to autonomous soil sampling and analysis by ground-based robots,” *Journal of Environmental Management*, vol. 360, p. 121130, 2024.
- [63] W. Myeong, S. Song, and H. Myung, “Development of a wall-climbing drone with a rotary arm for climbing various-shaped surfaces,” in *International Conference on Ubiquitous Robots*, 2018, pp. 687–692.
- [64] J. Hu, X. Han, Y. Tao, and S. Feng, “A magnetic crawler wall-climbing robot with capacity of high payload on the convex surface,” *Robotics and Autonomous Systems*, vol. 148, p. 103907, 2022.

- [65] S. Hong, Y. Um, J. Park, and H.-W. Park, “Agile and versatile climbing on ferromagnetic surfaces with a quadrupedal robot,” *Science Robotics*, vol. 7, no. 73, p. eadd1017, 2022.
- [66] T. Bandyopadhyay, R. Steindl, F. Talbot *et al.*, “Magneto: A versatile multi-limbed inspection robot,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 2253–2260.
- [67] S. T. Nguyen, A. Q. Pham, C. Motley, and H. M. La, “A practical climbing robot for steel bridge inspection,” *IEEE International Conference on Robotics and Automation*, May 2020.
- [68] Y. Bu, Y. Bu, H. Li *et al.*, “Development of wheel-leg hybrid climbing robot with switchable permanent magnetic omni-wheels as feet,” in *IEEE International Conference on Robotics and Biomimetics*, 2022, pp. 1–6.
- [69] J. Quan, M. Zhu, and D. Hong, “A lightweight mobile robot for climbing steel structures with an extending and bending tape spring limb,” in *ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, vol. 8: 47th Mechanisms and Robotics Conference, 2023, p. V008T08A084.
- [70] R. Wang and Y. Kawamura, “Development of climbing robot for steel bridge inspection,” *Industrial Robot*, vol. 43, no. 4, p. 429–447, 2016.
- [71] H. Eto and H. H. Asada, “Development of a wheeled wall-climbing robot with a shape-adaptive magnetic adhesion mechanism,” in *IEEE International Conference on Robotics and Automation*, 2020, pp. 9329–9335.
- [72] B. Wang, P. Li, P. Li *et al.*, “Development of a wheeled wall-climbing robot with an internal corner wall adaptive magnetic adhesion mechanism,” *Journal of Field Robotics*, pp. 1–18, 2024.
- [73] US EPA, OLEM. Tank inspections. [Online]. Available: <https://www.epa.gov/oil-spills-prevention-and-preparedness-regulations/tank-inspections>
- [74] Z. Y. Soon, J.-H. Jung, C. Yoon *et al.*, “Characterization of hazards and environmental risks of wastewater effluents from ship hull cleaning by hydroblasting,” *Journal of Hazardous Materials*, vol. 403, p. 123708, 2021.
- [75] E. Akyuz and M. Celik, “A methodological extension to human reliability analysis for cargo tank cleaning operation on board chemical tanker ships,” *Safety Science*, vol. 75, pp. 146–155, 2015.
- [76] S. Wickramanayake, K. Thiyagarajan, S. Kodagoda, and L. Piyathilaka, “Ultrasonic thickness measuring in-pipe robot for real-time non-destructive evaluation of polymeric spray linings in drinking water pipe infrastructure,” *Mechatronics*, vol. 88, p. 102913, 2022.

- [77] A. Gunatilake, L. Piyathilaka, A. Tran *et al.*, “Stereo vision combined with laser profiling for mapping of pipeline internal defects,” *IEEE Sensors Journal*, vol. 21, no. 10, pp. 11 926–11 934, 2021.
- [78] J. Valls Miro, N. Ulapane, L. Shi *et al.*, “Robotic pipeline wall thickness evaluation for dense nondestructive testing inspection,” *Journal of Field Robotics*, vol. 35, no. 8, pp. 1293–1310, 2018.
- [79] M. Tavakoli, C. Viegas, L. Marques *et al.*, “Omniclimbers: Omni-directional magnetic wheeled climbing robots for inspection of ferromagnetic structures,” *Robotics and Autonomous Systems*, vol. 61, no. 9, pp. 997–1007, 2013.
- [80] R. S. Bisht, P. M. Pathak, and S. K. Panigrahi, “Experimental investigations on permanent magnet based wheel mechanism for safe navigation of climbing robot,” *Procedia Computer Science*, vol. 133, pp. 377–384, 2018, International Conference on Robotics and Smart Manufacturing.
- [81] W. Fischer, F. Tâche, G. Caprari, and R. Siegwart, “Magnetic wheeled robot with high mobility but only 2 dof to control,” in *Advances In Mobile Robotics*. World Scientific, 2008, pp. 319–328.
- [82] C. Park, J. Bae, S. Ryu *et al.*, “R-Track: Separable modular climbing robot design for wall-to-wall transition,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1036–1042, 2021.
- [83] A. Wahrburg, J. Bös, K. D. Listmann *et al.*, “Motor-current-based estimation of cartesian contact forces and torques for robotic manipulators and its application to force control,” *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 2, pp. 879–886, 2017.
- [84] T. Bretl and S. Lall, “A fast and adaptive test of static equilibrium for legged robots,” in *IEEE International Conference on Robotics and Automation*, 2006, pp. 1109–1116.
- [85] A. Del Prete, S. Tonneau, and N. Mansard, “Fast algorithms to test robust static equilibrium for legged robots,” in *IEEE International Conference on Robotics and Automation*, 2016, pp. 1601–1607.
- [86] W. F. Ribeiro, K. Uno, K. Nagaoka, and K. Yoshida, “Dynamic equilibrium of climbing robots based on stability polyhedron for gravito-inertial acceleration,” in *International Conference on Climbing and Walking Robots and the Support Technologies for Mobile Machines*, 2020, pp. 297–304.
- [87] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics Modelling Planning and Control*. Springer, 2009.
- [88] A. M. Johnson and D. E. Koditschek, “Legged self-manipulation,” *IEEE Access*, vol. 1, pp. 310–334, May 2013.

- [89] R. Thakker, B. Jones, Y. K. Nakka *et al.*, “Loco-manipulation for snake robot climbing through simultaneous Shape-Contact-Force-Motion control,” in prep.
- [90] D. Prattichizzo, J. K. Salisbury, and A. Bicchi, “Contact and grasp robustness measures: Analysis and experiments,” in *Experimental Robotics*. Springer, 1997, pp. 83–90.
- [91] J. W. Burdick and E. Rimon, “Wrench resistant multi-finger hand mechanisms,” in *IEEE International Conference on Robotics and Automation*, 2016, pp. 2030–2037.
- [92] Y. Shirai, X. Lin, A. Schperberg *et al.*, “Simultaneous contact-rich grasping and locomotion via distributed optimization enabling free-climbing for multi-limbed robots,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 13 563–13 570.
- [93] D. Kim, D. Carballo, J. Di Carlo *et al.*, “Vision aided dynamic exploration of unstructured terrain with a small-scale quadruped robot,” in *IEEE International Conference on Robotics and Automation*, 2020, pp. 2464–2470.
- [94] R. Grandia, F. Jenelten, S. Yang *et al.*, “Perceptive locomotion through nonlinear model-predictive control,” *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 3402–3421, 2023.
- [95] E. Jelavic, K. Qu, F. Farshidian, and M. Hutter, “LSTP: Long short-term motion planning for legged and legged-wheeled systems,” *IEEE Transactions on Robotics*, vol. 39, no. 6, pp. 4190–4210, 2023.
- [96] P. Krüsi, P. Furgale, M. Bosse, and R. Siegwart, “Driving on point clouds: Motion planning, trajectory optimization, and terrain assessment in generic nonplanar environments,” *Journal of Field Robotics*, vol. 34, no. 5, pp. 940–984, 2017.
- [97] P. E. Hart, N. J. Nilsson, and B. Raphael, “A formal basis for the heuristic determination of minimum cost paths,” *IEEE transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [98] S. Karaman and E. Frazzoli, “Incremental sampling-based algorithms for optimal motion planning,” *Robotics Science and Systems VI*, vol. 104, no. 2, pp. 267–274, 2010.
- [99] Z. Kingston, M. Moll, and L. E. Kavraki, “Sampling-based methods for motion planning with constraints,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, no. 1, pp. 159–185, 2018.
- [100] T. Bretl, “Motion planning of multi-limbed robots subject to equilibrium constraints: The free-climbing robot problem,” *The International Journal of Robotics Research*, vol. 25, no. 4, pp. 317–342, 2006.
- [101] R. Shome, D. Nakhimovich, and K. E. Bekris, “Pushing the boundaries of asymptotic optimality in integrated task and motion planning,” in *Algorithmic Foundations of Robotics*. Springer, 2021, pp. 467–484.

- [102] W. Thomason, M. P. Strub, and J. D. Gammell, “Task and motion informed trees (TMIT*): Almost-surely asymptotically optimal integrated task and motion planning,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 370–11 377, 2022.
- [103] A. Hornung, K. M. Wurm, M. Bennewitz *et al.*, “OctoMap: An efficient probabilistic 3D mapping framework based on octrees,” *Autonomous robots*, vol. 34, pp. 189–206, 2013.
- [104] R. B. Rusu and S. Cousins, “3D is here: Point Cloud Library (PCL),” in *IEEE International Conference on Robotics and Automation*, Shanghai, China, 2011.
- [105] M. Kvasnica, “Minkowski addition of convex polytopes,” 2005.
- [106] I. A. Şucan, M. Moll, and L. E. Kavraki, “The Open Motion Planning Library,” *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [107] M. Labbé and F. Michaud, “RTAB-Map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation,” *Journal of Field Robotics*, vol. 36, no. 2, p. 416–446, 2018.