

**Navigating a Complex World:
Improving Robot Outcomes Through
Social, Regulatory, and Control Theoretic Approaches**

Submitted in partial fulfillment of the requirements for
the degree of
Doctor of Philosophy
in
Mechanical Engineering

James Zhu

B.E., Mechanical Engineering, Vanderbilt University
M.S., Mechanical Engineering, Carnegie Mellon University

Carnegie Mellon University
Pittsburgh, PA
September 6, 2024

Copyright © 2024 by James Zhu

All Rights Reserved

Acknowledgements

Of all the lessons I've learned during my PhD, the most profound has been understanding what it truly means to dream. To lie awake, envisioning a distant goal, even without knowing the path to get there. To be honest, earning this degree wasn't a long-held dream of mine. It wasn't in 2006, when, during a family vacation to Boston, I bought a postcard of Harvard Medical School, imagining myself as a future heart surgeon. It wasn't in 2016 when I started at Vanderbilt University as an Economics major, only to change course within my first semester. It wasn't in 2019 when I accepted a consulting job in D.C. I wasn't even sure it was in 2022, when I was ready to leave the program entirely.

My dream is to become a mentor, an advisor, and an advocate. I hope to be surrounded by big dreamers and help transform their aspirations into reality. I want to uplift others and inspire them to believe that no dream is too far out of reach. I know the impact such a person can have because I've been fortunate to meet so many who have taught me to dream and to believe I can achieve anything.

My parents, whose lives have been a testament to selfless dedication, offering me every possibility and urging me to reach my fullest potential. My brother, Peter, whom I've always seen as the smarter of us, destined for a brilliance that shines ever brighter.

The friends who have journeyed with me through the seasons of my life: my oldest childhood companions, Austin and Kathy; my Vanderbilt squad: Jason, Elissa, Jeremy, Sam, Tay, and Michaela; the dearly cherished Lisa and Kacie; Chase, Darren, and Jenny, with whom I've explored countless new horizons; and my steadfast comrade Arihant- working alongside you has been a profound honor.

To all my lab mates, with whom I've forged enduring bonds, I am deeply grateful. Especially to the brilliant Joshua Ramos, whose growth I've witnessed firsthand and whose journey fills me with pride. I owe much to Ardalan Tajbakhsh, whose support has been invaluable and who I am certain will achieve greatness in his own right. My heartfelt thanks also extend to my co-authors for their invaluable contributions: Nathan, Joe, George, David, Selvin, and Anoushka.

To my advisor, Aaron Johnson, who urged me to wander through the desert of uncertainty and discover my true passions. To Kate Darling and the RES team– thank you for granting me the chance to explore realms beyond my qualifications; my time at the AI Institute was truly transformative. I extend my deepest gratitude to my thesis committee: Illah Nourbakhsh, whose unwavering support guided me into uncharted and daunting research territories; Mark Bedillion, widely celebrated as the kindest soul in the Mechanical Engineering department; and Luis Sentis, an icon and a figure I hold in profound admiration. My work was supported by the National Science Foundation under grant #CMMI-1943900.

Finally, I'd like to turn to myself for a moment of reflection. This journey has been fraught with challenges, yet you have persevered and I am profoundly proud of what you have achieved. You've overcome such long odds, but not because you're one in a million – because you're one of one.

Abstract

In recent years, many wheeled, bipedal, and quadrupedal robots have been released to the market. However, their presence in public spaces remains limited, with most robots confined to controlled environments like factories and warehouses. What is hindering robots from making broader impacts? This thesis investigates the challenges robots face by synthesizing insights from control theory, social science, and public policy. Each of these often-siloed fields offers valuable perspectives for developing robots that are safe, trustworthy, and equitable.

The first part of this thesis presents methods to analyze legged robot locomotion and generate safe, robust trajectories for challenging environments. Legged robots are hybrid systems that undergo discontinuous changes in state and dynamics upon foot touchdown, violating assumptions of many traditional control architectures. The presented hybrid systems analysis uses the fundamental solution matrix to characterize the evolution of initial errors through a trajectory. This analysis leads to novel trajectory optimization methods that explicitly consider the stability and convergence of hybrid trajectories, improving tracking performance across various systems. Additionally, this analysis enhances our fundamental understanding of how gait parameters like duty factor impact a robot's performance on difficult terrain.

This thesis also explores legal theory and community attitudes to develop frameworks for equitable robot design, ensuring fair distribution of benefits and risks. I emphasize grounding robot design in the needs of all community members, not just customer preferences. By considering self-defense law and the perspectives of food pantry patrons toward autonomous delivery robots, this work aims to improve both robot design and public policies.

Overall, this thesis tackles the diverse factors essential for enabling robots to make wide-ranging, positive impacts on the world. Neglecting any technical, social, or regulatory aspect of robotics risks undermining their lofty potential.

Contents

1	Introduction	1
1.1	Motivation and Problem Statement	1
1.2	Thesis Outline	3
I	Navigating a Complex, Uncertain World	6
2	Stability and Convergence Analysis of Hybrid Systems	7
2.1	Introduction	7
2.2	Hybrid Systems Definition	7
2.3	Saltation Matrix	8
2.4	Fundamental Solution Matrix	10
2.5	Stability and Convergence Measures	10
2.5.1	Periodic Stability Analysis	10
2.5.2	Aperiodic Convergence Analysis	11
3	Hybrid Event Shaping: Generating Open-Loop Stable Hybrid Orbits	12
3.1	Introduction	12
3.2	Methods	13
3.2.1	Hybrid Event Shaping	13
3.2.2	Virtual Hybrid Events	15
3.2.3	Stability Measure Derivative	15
3.2.4	Shape Parameter Stability Optimization	16
3.2.5	Trajectory Optimization with Hybrid Event Shaping	17
3.3	Examples and Results	18
3.3.1	Paddle Juggler	18
3.3.2	2D Hopper	20
3.3.3	Walking Biped Trajectory Optimization	23
3.4	Conclusion	26
4	Convergent iLQR: Generating Closed-Loop Convergent Hybrid Trajectories	27
4.1	Introduction	27
4.2	Related Works	29
4.3	iLQR for Hybrid Systems	30
4.4	Convergent iLQR	32
4.5	Examples and Results	35

4.5.1	Rocket Hopper	35
4.5.2	Planar Quadruped	37
4.6	Conclusion	41
5	The Effect of Gait Parameters on Safe Quadrupedal Robot Locomotion	42
5.1	Introduction	42
5.2	Related Works: Effects of Gait Parameters in Legged Locomotion	45
5.3	Preliminaries	47
5.3.1	Hybrid Systems	47
5.3.2	System Definition	47
5.3.3	Convergence and Cost of Transport Analysis	48
5.4	Methods	49
5.4.1	Investigation 1: Whole-Body Trajectory Optimization with LQR Feedback . .	50
5.4.2	Investigation 2: Centroidal MPC	51
5.5	Experiments and Results	52
5.5.1	Experiments of Investigation 1	52
5.5.2	Results of Investigation 1	53
5.5.3	Experiments of Investigation 2	54
5.5.4	Results of Investigation 2	56
5.6	Conclusion	57
II	Navigating a Complex, Human World	58
6	Self-Defense Against Robots: Design and Deployment Considerations for Roboticists	59
6.1	Introduction	59
6.2	U.S. Self-Defense Law	62
6.2.1	A Reasonable Belief of Imminent Physical Harm	62
6.2.2	A Proportional Response to Threat	63
6.2.3	Self-Defense Against Robots	64
6.3	Reasonable Perspectives Toward Robots	65
6.3.1	Behavioral Norms in Human Interactions	65
6.3.2	Human Attitudes Toward Robots	67
6.4	Human-Aware Planning	69
6.4.1	Explainability	69
6.4.2	Legibility	71
6.4.3	Human-Aware Planning for Non-Expert, Non-Users	72
6.5	Justified Human Self-Defense Against Robots	74
6.6	Discussion and Conclusion	75
7	Improving Equity in Robot Deployment: A Study of Food Pantry Patrons	77
7.1	Introduction	77
7.2	Background and Related Works	80
7.2.1	Food Insecurity and Shopping Patterns	80
7.2.2	Food Pantries	81
7.2.3	Autonomous Delivery Robot Deployment	82

7.2.4	Digital Technologies and Robots for Food Pantries	83
7.3	Methods	84
7.3.1	Participants	84
7.3.2	Positionality Statement	85
7.3.3	Interview Design and Procedure	86
7.3.4	Data Analysis	88
7.4	Results	89
7.4.1	Price Dominates	90
7.4.2	Patrons' Social Connections	92
7.4.3	Language and Technological Literacy Barriers	94
7.4.4	Mixed Views on Robots, But Openness to Delivery Services	95
7.5	Discussion	98
7.5.1	Shaping the Future of ADRs	98
7.5.2	Robots in Public Spaces	99
7.6	Limitations	100
7.7	Future Work	101
7.8	Conclusion	101
7.9	Acknowledgements	102

III Concluding Remarks 103

8 Conclusion 104

8.1	Future Work	105
8.1.1	Stability of Systems Under Stick-Slip Transitions	105
8.1.2	How People Develop Internal Models of Robots	106
8.1.3	Policy Making For Public Robots	106
8.1.4	Uniting Robot Failure Modes	106
	References	107

List of Figures

1.1	The Robotics Deployment Stack	2
3.1	Paddle Juggler Stability Example	13
3.2	2D Hopper System Diagram	20
3.3	2D Hopper Error Results	22
3.4	Biped System Diagram	24
3.5	Biped Error Results	25
4.1	Planar Quadruped Convergence Example Trajectories	28
4.2	Planar Quadruped Error Results	40
4.3	Planar Quadruped Robustness Results	41
5.1	Experimental Hardware Setup for Beam Walking Experiments	44
5.2	Replication of Gait Efficiency Results	45
5.3	Gait Convergence Results	53
5.4	Gait Convergence for Beam Walking Results	54
5.5	Gait Disturbance Rejection Results	55
5.6	Experimental Setup in Simulation	56
7.1	Summary of Participant Responses	89

List of Tables

3.1	2D Hopper Swing Leg Retraction Stability Results	21
3.2	Biped Stability Results	25
4.1	Rocket Hopper Controller Parameters	34
4.2	Rocket Hopper Convergence Measure and Simulation Results	37
4.3	Planar Quadruped Convergence Results	39
5.1	Simulation Experiment Results	56
5.2	Hardware Experiment Results	57

Chapter 1

Introduction

1.1 Motivation and Problem Statement

The world is an abundantly complex place and is rife with dangers that can cause a robot to fail. One source of this complexity is environmental uncertainties such as slick surfaces, tall obstacles, or gaps in terrain. In urban environments like Pittsburgh, sidewalks can be poorly maintained and conditions like rain or snow make traversal even more difficult. Wheeled robots can only go so far before reaching an upper limit on their capabilities, so researchers have turned to the agility that legged robots can provide to handle these scenarios. However, legged robots introduce their own wealth of difficulties, such as maintaining stability and robustness under unforeseen perturbations. One reason why controlling legged robots is difficult is because traditional control methods used for smooth systems do not extend well to hybrid systems where contact is made and broken. Continuity is a key assumption for many of these methods (e.g. Lyapunov analysis [1]), so adapting them to account for the discrete nature of hybrid systems is a challenge.

Interacting with humans adds to the world's complexity. People are difficult to model and predict, yet gaining their trust is crucial for any technology, especially robots operating in public spaces like sidewalks, parks, and malls. Ensuring physical safety isn't enough; robots must behave in ways that foster trust, especially among those who are initially most resistant. Overcom-




Figure 1.1: Successful robot deployments demand an interdisciplinary understanding of many interwoven subjects, encompassing advanced trajectory planning and control techniques, alignment with stakeholder values, and compliance with legal and regulatory standards. The absence of even one of these layers could jeopardize the success of any robotic system.

ing technical, historical, and cultural barriers to acceptance requires engineers to design robots with diverse human perspectives in mind. Regulation also plays a key role in enabling robots to contribute meaningfully to society while preventing misuse, and roboticists have a responsibility to shape these policies.

While not every pitfall can be anticipated, it is crucial for roboticists to understand the consequences of robot failures. Inability to handle environmental uncertainties can result in damaged robots and financial loss, while neglecting human factors could lead to unequal distribution of risk and reinforce longstanding injustices against marginalized groups. This thesis is guided by two sweeping questions:

- What factors cause robots to fail in complex environments?
- How can robots be designed to mitigate the likelihood of failure?

I address these questions through what I call the “Robot Deployment Stack”, shown in Fig. 1.1. Inspired by the concept of a software stack, this model comprises interconnected layers of engineering, social science, and policy design, all of which must be carefully refined to enable successful robot deployment. The Robot Deployment Stack includes three layers of robotic control

theory: safe whole-body feedback control, which depends on stable trajectory generation, which in turn relies on robust high-level gait planning. However, even a perfectly engineered system cannot succeed without a thorough understanding of how stakeholders will perceive and interact with the technology. Deploying robots into real-world environments demands ongoing engagement with community members, particularly marginalized groups who may be at increased risk of harm. Additionally, robot deployment encounters significant regulatory and legal challenges. It will strain existing legal frameworks, necessitating both new interpretations of existing laws and the creation of new ones. As a result, roboticists must become adept at navigating the complex and varied legal landscape, including city, state, federal, and international regulations. This thesis presents my contributions to each layer of the stack and examines the additional work needed to unite these layers into a cohesive ecosystem, enabling robots to reach those who could benefit from them the most.

1.2 Thesis Outline

This thesis is divided into two parts. Part I, *Navigating a Complex, Uncertain World*, presents mathematical analysis and control methods to help legged robots safely traverse uncertain terrain. Part II, *Navigating a Complex, Human World*, examines how social and legal structures shape robot design. Together, this thesis explores the diverse demands that robots must meet to succeed.

Part I: Navigating a Complex, Uncertain World

Chapter 2: Stability and Convergence Analysis of Hybrid Systems

Broadly speaking, the goal of robot control is to keep a robot's state close to its intended state. Variational equations, which describe the evolution of nearby trajectories, can assess the effectiveness of robot controllers. Chapter 2 lays the mathematical foundation for this analysis, introducing the stability and convergence measures that will be used throughout this work. This chapter is partly based on a tutorial paper co-authored with Dr. Nathan J. Kong and J. Joe Payne,

to be published in the Proceedings of the IEEE [2].

Chapter 3: Hybrid Event Shaping: Generating Open-Loop Stable Hybrid Orbits

Standard legged robot gaits are periodic, meaning that they repeat themselves after a certain amount of time. Previous research has identified behaviors like swing leg retraction that enhance the stability of these gaits [3], but these results often rely on extensive heuristic tuning and do not generalize to other stabilizing behaviors. Chapter 3 introduces Hybrid Event Shaping (HES), a method that identifies stabilizing behaviors based on the stability analysis from Chapter 2. HES autonomously generates open-loop stable trajectories for legged robots, reproducing known phenomena like swing leg retraction and introducing novel behaviors in more complex robot models. This work was published at IEEE ICRA 2022 [4].

Chapter 4: Convergent iLQR: Generating Closed-Loop Convergent Hybrid Trajectories

In practice, many robotic behaviors are never open-loop stable and require some feedback control. Chapter 4 extends HES by reasoning about closed-loop trajectories. Standard controllers like LQR often need extensive tuning and can struggle with hybrid, underactuated trajectories typical of legged robots. While increasing feedback gains might improve closed-loop performance, it can also reduce robustness. Instead, this chapter introduces Convergent iLQR (χ -iLQR), a trajectory optimization method that enhances robustness without increasing feedback gains. This approach is demonstrated in simulation on a planar quadruped robot and was presented at IEEE ICRA 2024 [5].

Chapter 5: The Effect of Gait Parameters on Safe Quadrupedal Robot Locomotion

Despite advances in autonomous locomotion through trajectory optimization, many gait parameters—such as speed, gait sequence, and duty factor—must still be set manually, leading to brittleness on unfamiliar terrain. Previous research has explored the relationship between gait parameters and efficiency [6], but their connection to robustness remains unclear. Chapter 5 leverages the

convergence measure to establish a relationship between increased duty factor and improved convergence. This insight was applied to a balance beam walking environment where default gaits with low duty factors failed, but increasing duty factor lead to improved performance. This was shown using multiple control strategies in simulation as well as hardware experiments. This work will be submitted to IEEE RA-L [7].

Part II: Navigating a Complex, Human World

Chapter 6: Self-Defense Against Robots: Design and Deployment Considerations for Robotists

No matter how adept robots are at navigating physical obstacles, their utility is limited if they can't manage human interactions. Without sufficient trust and comfort, robots will fail in their objectives. Chapter 6 discusses a legal interpretation of how robots can fail when humans act in self-defense. Humans that feel threatened by a robot based on reasonable beliefs are justified to defend themselves by damaging or possibly destroying a robot. To prevent such scenarios, this work offers practitioners actionable steps to enhance public trust in robots and explore the link between robot design and human attitudes. This work was presented at IEEE RO-MAN 2023 [8].

Chapter 7: Improving Equity in Robot Deployment: A Study of Food Pantry Patrons

Whereas Chapter 6 focuses on mitigating risks to the public, Chapter 7 explores how to make the benefits of robotic technology more accessible to underrepresented communities. This work consisted of 21 interviews with food pantry patrons about their food shopping experiences and perspectives toward delivery services and autonomous delivery robots. Qualitative analysis of these interviews surfaced insights into factors influencing food shopping decisions, attitudes toward current grocery delivery services, and perspectives on using food delivery robots. This study showed that, despite challenges in deploying robots for food pantries, there is both a need for and interest in such services. This work is in preparation [9].

Part I

Navigating a Complex, Uncertain World

Chapter 2

Stability and Convergence Analysis of Hybrid Systems

2.1 Introduction

This chapter formally defines a hybrid system and the saltation matrix, which is an important tool for analyzing the evolution of errors for a hybrid trajectory. With this analysis enabled by the saltation matrix, we can generate the fundamental solution matrix, which yields two scalar measures that represent aspects of a trajectory's tracking performance. These measures are in turn incorporated into trajectory optimization frameworks in the following chapters.

2.2 Hybrid Systems Definition

Hybrid systems are a class of dynamical systems which exhibit both continuous and discrete dynamics [10,11]. Following [2], we define a C^r hybrid dynamical system for continuity class $r \in \mathbb{N}_{>0} \cup \{\infty, \omega\}$ as a tuple $\mathcal{H} := (\mathcal{J}, \Gamma, \mathcal{D}, \mathcal{F}, \mathcal{G}, \mathcal{R})$ where:

1. $\mathcal{J} := \{I, J, \dots\} \subset \mathbb{N}$ is the finite set of discrete modes.
2. $\Gamma \subseteq \mathcal{J} \times \mathcal{J}$ is the set of discrete transitions forming a directed graph structure over $\mathcal{J}$.

3. $\mathcal{D} := \coprod_{I \in \mathcal{J}} D_I$ is the collection of domains, where D_I is a C^r manifold and the state $x \in D_I$ while in mode I .
4. $\mathcal{F} : \mathbb{R} \times \mathcal{D} \rightarrow \mathcal{TD}$ is a collection of C^r time-varying vector fields, $F_I := \mathcal{F}|_{D_I} : \mathbb{R} \times D_I \rightarrow \mathcal{TD}_I$, for each $I \in \mathcal{J}$.
5. $\mathcal{G} := \coprod_{(I,J) \in \Gamma} G_{(I,J)}(t)$ is the collection of guard sets, where $G_{(I,J)}(t) \subseteq D_I$ for each $(I, J) \in \Gamma$ is defined as a regular sublevel set of a C^r guard function, i.e. $G_{(I,J)}(t) = \{x \in D_I | g_{(I,J)}(t, x) \leq 0\}$ and $D_x g_{(I,J)}(t, x) \neq 0 \ \forall \ g_{(I,J)}(t, x) = 0$.
6. $\mathcal{R} : \mathbb{R} \times \mathcal{G} \rightarrow \mathcal{D}$ is a C^r map called the reset that restricts as $R_{(I,J)} := \mathcal{R}|_{G_{(I,J)}(t)} : G_{(I,J)}(t) \rightarrow D_J$ for each $(I, J) \in \Gamma$.

An execution of a hybrid system [12] begins at an initial state $x_0 \in D_I$. With input $u_I(t, x)$, the system obeys the dynamics F_I on D_I . If the system reaches guard surface $G_{(I,J)}$, the reset map $R_{(I,J)}$ is applied and the system continues in domain D_J under the corresponding dynamics defined by F_J . The flow $\phi(t, t_0, x_0, U)$ describes how the hybrid system evolves from some initial time t_0 and state x_0 until some final time t under input sequence U .

Hybrid systems may exhibit complex behaviors including sliding [13], branching [14], and Zeno phenomena where infinite transitions occur in finite time [15]. Following prior literature [16–18], we assume these behaviors do not occur, such that guard surfaces are isolated and intersected transversely [11,12] and no Zeno executions occur. These assumptions are not generally detrimental to the validity of this theory to applications like legged locomotion.

2.3 Saltation Matrix

For both continuous domains and hybrid transitions, linearized variational equations can be constructed to characterize the evolution of perturbations δx [19]. In each continuous domain, the

linearized variational equation is discretized from timestep i to $i + 1$ and is:

$$\delta x_{i+1} \approx (A_I - B_I K_I) \delta x_i \quad (2.1)$$

with A_I and B_I being the derivatives of the discretized dynamics in mode I w.r.t. state x_i and control inputs u_i , respectively, and K_I are linear feedback gains [19]. For hybrid events, the analogous variational term is the saltation matrix $\Xi_{(I,J)}$, which describes the transition between modes I and J. The saltation matrix is the first-order approximation of the change in state perturbations from before the hybrid event at $\delta x(t^-)$ to perturbations after $\delta x(t^+)$ [17], such that:

$$\delta x(t^+) \approx \Xi_{(I,J)} \delta x(t^-) \quad (2.2)$$

The formulation of the saltation matrix is:

$$\Xi_{(I,J)} := D_x R^- + \frac{(F_J^+ - D_x R^- F_I^- - D_t R^-) D_x g^-}{D_t g^- + D_x g^- F_I^-} \quad (2.3)$$

where

$$F_I^- := F_I(t^-, x(t^-)) \quad (2.4)$$

$$F_J^+ := F_J(t^+, x(t^+)) \quad (2.5)$$

$$x(t^+) := R_{(I,J)}(t^-, x(t^-)) \quad (2.6)$$

$$D_x R^- := D_x R_{(I,J)}(t^-, x(t^-)) \quad (2.7)$$

$$D_t R^- := D_t R_{(I,J)}(t^-, x(t^-)) \quad (2.8)$$

$$D_x g^- := D_x g_{(I,J)}(t^-, x(t^-)) \quad (2.9)$$

$$D_t g^- := D_t g_{(I,J)}(t^-, x(t^-)) \quad (2.10)$$

More information on the saltation matrix and a rigorous derivation can be found in [2].

2.4 Fundamental Solution Matrix

Consider a trajectory that begins at state $x_0 = x(t_0)$ for some initial time t_0 and is executed until time t_f where it arrives at state $x_f = \phi(t_f, t_0, x_0, U)$. Our control objective is to bring any nearby initial state $\bar{x}_0 = x_0 + \delta x_0$ towards the nominal trajectory so that at time t_f , $\bar{x}_f = \phi(t_f, t_0, \bar{x}_0, \bar{U}) = x_f + \delta x_f$ is closer to x_f . To characterize the closeness of $\bar{x}_f$ and x_f , we utilize the fundamental solution matrix, Φ . Following [18], the fundamental solution matrix represents the transformation of error from the initial state to final state:

$$\delta x_f \approx \Phi \delta x_0 \quad (2.11)$$

The fundamental solution matrix can be computed by sequentially composing the linearized variational terms in each continuous domain ($\tilde{A} := A - BK$) and the saltation matrices (Ξ) at each hybrid event [4]. For a hybrid trajectory with N domains, the fundamental solution matrix can be formulated as:

$$\Phi = \tilde{A}_N \Xi_{(N-1,N)} \dots \Xi_{(2,3)} \tilde{A}_2 \Xi_{(1,2)} \tilde{A}_1 \quad (2.12)$$

2.5 Stability and Convergence Measures

2.5.1 Periodic Stability Analysis

If $x_f = x_0$, then the trajectory is periodic, with period $T = t_f - t_0$. In this case, the fundamental solution matrix is also known as the monodromy matrix [18,20]. The monodromy matrix determines local asymptotic orbital stability (which we refer to simply as stability). For nonautonomous systems, stability is determined by the maximum magnitude of the eigenvalues, $\max(|\lambda|)$ [18]. We refer to this as the stability measure, ψ , where a trajectory is stable when $\psi < 1$. Autonomous systems always have an eigenvalue that is equal to 1 since for non-time varying dynamics, perturbations along the flow of the orbit will by definition map back to themselves after period T [18].

Assuming non-convergence in this direction is allowable, ψ for autonomous systems is based on the remaining eigenvalues.

2.5.2 Aperiodic Convergence Analysis

In cases that are not periodic, the fundamental solution matrix captures the change in errors across a trajectory and the singular values of Φ characterize error change along principle axes of state space. The largest singular value, which is equivalent to the induced 2-norm of Φ , describes the evolution of the most divergent direction of initial error δx_0 . We define the convergence measure, χ to be exactly this worst-case value:

$$\chi = \|\Phi\|_2 \quad (2.13)$$

χ is a continuous measure of local convergence, where smaller values of χ indicate stronger reduction of worst-case final errors. A value of $\chi < 1$ indicates errors in all directions will shrink. Therefore, the convergence measure directly correlates with tracking performance of a closed-loop trajectory.

Chapter 3

Hybrid Event Shaping: Generating Open-Loop Stable Hybrid Orbits

3.1 Introduction

In general, the walking and running gaits of legged robots are naturally unstable and challenging to control. Hybrid systems such as these are difficult to work with due to the discontinuities in state and dynamics that occur at hybrid events. These discontinuities violate assumptions of standard controllers designed for purely continuous systems, and work is ongoing to adapt these controllers for hybrid systems [21,22]. One strategy for hybrid control is to cancel out the effects of hybrid events by working with an invariant subsystem [23–25]. We propose instead that the effects of hybrid events are valuable due to rich control properties that can be used to stabilize trajectories of a hybrid system.

Several works have examined the utility of controlling hybrid event conditions to improve system stability without any closed-loop continuous-domain control [3,26,27]. For example, [26] found that for the paddle juggler system, paddle acceleration at impact uniquely determines the local stability properties of a periodic trajectory, Fig. 3.1. Other works [25,28] generated open-loop swing leg trajectories that produced deadbeat hopping of a SLIP-like system. Each of these




Figure 3.1: The paddle juggler system [26] has no control authority while the ball is in the air. The paddle acceleration at impact determines the convergence/divergence of the system from initial points (cyan dots) to the final states (magenta stars) after 5 cycles. This example underscores how hybrid event shaping can stabilize a periodic hybrid system.

works found that controlling a hybrid system only at the moment of a hybrid event is sufficient to provide stabilization. So far, however, these results have only been produced for each specific problem structure and a clear connection between these works has yet to be established.

In this work, we propose the concept of hybrid event shaping (HES), which describes how hybrid event parameters can be chosen to affect the stability properties of a periodic orbit. We also propose methods to produce values of these hybrid event parameters to optimize a stability measure of a trajectory. This approach is tested on both existing examples from [3,26] and on a new bipedal robot controller.

3.2 Methods

3.2.1 Hybrid Event Shaping

Hybrid events can greatly affect the stability of an orbit due to the unbounded discontinuous changes that are made to perturbations. The saltation matrix allows for an explicit understanding

of how to perform “hybrid event shaping” (HES), i.e. choosing hybrid event parameters (such as timing, state, input, and higher order “shape parameters”) to improve the stability of a periodic trajectory. The key insight is that hybrid event shaping introduces a generalizable method to stabilize hybrid systems that is independent of continuous-domain control, but that can work in concert with it.

In general, the open-loop continuous variational equations of a hybrid system are functions of initial and final time, initial state, and system dynamics. However, it is challenging to alter any of these parameters because changes will propagate through the rest of the trajectory and periodicity may be violated, though we present a trajectory optimization method below to handle this. The saltation matrix is a function of nominal event time, state and dynamics, but additionally may be a function of higher order shape parameters h that do not influence the dynamics of the system. These parameters arise from the derivatives of the guards ($D_x g$ and $D_t g$) and reset maps ($D_x R$ and $D_t R$) but are not present in the guard, reset map, or vector field definitions themselves. Therefore, shape parameters have absolutely no effect on the nominal trajectory and can be chosen completely freely.

One example of a shape parameter is the angular velocity of a massless leg of a spring-loaded inverted pendulum. Since a massless leg does not induce any torque in the air or forces at touchdown, only the position of the leg at touchdown affects the trajectory of the body. However, leg velocity appears in the saltation matrix and has a significant effect on orbital stability [3].

For more complex models of robots, there may not be any physical shape parameters that can be tuned. For example, legged robots with non-massless legs can not vary leg velocity at impact without also changing their trajectories. These cases can be handled by running a trajectory optimization at the same time as applying HES, as we show in Sec. 3.2.5, or by adding additional virtual hybrid events.

3.2.2 Virtual Hybrid Events

Certain control systems naturally have discontinuities in control inputs, such as bang-bang control, sliding mode control, or systems with actuators that have discretized (i.e. on-off) inputs. These discontinuities in control can cause an instantaneous change in the dynamics of the system, resulting in a virtual (as opposed to physical) hybrid event. Virtual hybrid events act no differently than physical hybrid events and induce saltation matrices with shape parameters to be tuned for stability. Even for systems where discontinuous control inputs are not necessary, the addition of virtual saltation matrices and shape parameters allows for a greater authority to improve stability.

3.2.3 Stability Measure Derivative

Our goal is to determine the optimal choice of hybrid event parameters that minimizes the stability measure of a trajectory. Since directly computing eigenvalues in closed-form is not generally feasible, one solution is to use numerical methods to perform optimization [29]. However, this strategy becomes untenable for high dimensional problems. Instead, by using the saltation matrix formulation (2.3), derivatives of the stability measure can be directly computed, allowing for the use of more efficient optimization methods and making the problem much more tractable.

Assuming that the monodromy matrix Φ depends continuously on each shape parameter h_n , the eigenvalues of Φ are always continuous with respect to h_n [30]. For a diagonalizable Φ , the derivative of the eigenvalues with respect to h_n can be computed in closed form [31]. Assume that matrix $\Phi(h_n)$ has simple (non-repeating) eigenvalues, $\lambda_1, \dots, \lambda_N$, and let $\mathbf{j}_i$ and $\mathbf{k}_i$ denote the left and right eigenvectors associated with λ_i . Then the derivative $\frac{d\lambda_i}{dh_n}$ is:

$$\frac{d\lambda_i}{dh_n} = \mathbf{k}_i' \frac{d\Phi}{dh_n} \mathbf{j}_i \quad (3.1)$$

For matrices with eigenvalues that repeat, the derivatives of the repeated eigenvalues can be calculated similarly with a matrix of associated eigenvectors [31].

$\frac{d\Phi}{dh_n}$ can be found using the derivative product rule, which simplifies if each shape parameter only appears in one saltation matrix. We make this assumption here to improve computational efficiency, but it is not required generically. Without loss of generality, take $\frac{d\Xi_{(1,2)}}{dh_n} \neq 0$, so that:

$$\frac{d\Phi}{dh_n} = \Xi_{(N,1)} A_N \dots A_2 \frac{d\Xi_{(1,2)}}{dh_n} A_1 \quad (3.2)$$

Substituting (3.2) into (3.1) allows us to compute the derivative of the stability measure with respect to each of the shape parameters. Eq. (3.1) is not valid for non-diagonalizable monodromy matrices. However, the guaranteed continuity of the stability measure allows for finite-difference methods to be used in any non-diagonalizable cases.

The derivative computation from (3.2) can be adapted for changes in x and t as well. Without loss of generality, consider again $\Xi_{(1,2)}$. For hybrid event time $t_{(1,2)}$, the derivative $\frac{d\Xi_{(1,2)}}{dt_{(1,2)}}$ can be computed in closed-form. Additionally, the derivatives $\frac{dA_1}{dt_{(1,2)}}$ and $\frac{dA_2}{dt_{(1,2)}}$ are non-zero and can be computed through standard methods [32]. The product rule expansion of $\frac{d\Phi}{dt_{(1,2)}}$ consists of additional terms compared to (3.2) but otherwise can be computed similarly. $\frac{d\Phi}{dx_{(1,2)}}$ for hybrid event state $x_{(1,2)}$ can be computed this same way.

3.2.4 Shape Parameter Stability Optimization

Optimization techniques [33] are able to select optimal hybrid event parameters that minimize the stability measure. Two optimization methods are presented here: the first optimizing the shape parameters without affecting the dynamics of the nominal orbit, and the second optimizing both the hybrid events and periodic orbit simultaneously.

Shape parameters are powerful because they do not appear in the dynamics of the system and have no effect on the nominal trajectory. This means that for a given periodic trajectory, the shape parameters can be chosen freely. We use an optimization framework to choose these

shape parameters with the goal of optimizing the stability measure $\psi(\Phi(h))$ of a trajectory,

$$\underset{h_{1:M}}{\text{minimize}} \quad \psi(\Phi(h)) \quad (3.3)$$

The ability to compute derivatives of the stability measure allows for this optimization to be more computationally efficient. The examples below show how this optimization method is able to reproduce swing leg retraction in a one-legged hopper system by determining optimal inputs of shape parameters to minimize the stability measure.

3.2.5 Trajectory Optimization with Hybrid Event Shaping

Some systems do not have shape parameter terms in their saltation matrices or do not have enough to sufficiently improve stability. In these cases, we can change the trajectory of the system itself so that the timing, state, and input parameters of the continuous variational matrices and saltation matrices improve stability properties. However, it must be ensured that the dynamics, periodicity, and other constraints of the system are obeyed.

Trajectory optimization methods are a class of algorithms that aim to minimize a cost function while satisfying a set of constraints [34]. For dynamical systems, these costs are generally functions of state and inputs, with constraints imposed on system dynamics and any other physical limits [35]. For specific problems, other aspects of the system may be added into the cost or constraint functions such as design parameters or minimizing time [36,37]. Here we propose including the stability measure in the cost function to search for optimally stable trajectories. Eq. (3.4) gives the simplest form of this trajectory optimization problem, with periodicity and dynamics constraints being enforced, where dynamics constraints obey continuous dynamics in each domain and reset mappings at each hybrid event [38]. Additional costs and constraints may be included such as reference tracking costs, input costs, and any physical constraints. We solve this problem using a direct collocation optimization with a multi-phase method to handle

hybrid events [34,35,39].

$$\begin{aligned}
& \underset{x_{1:N}, u_{1:N-1}, t_{1:N-1}, h_{1:M}}{\text{minimize}} && \psi(\Phi(x, t, h)) \\
& \text{subject to} && x_N = x_1 \\
& && x_{i+1} = \phi(t_{i+1}, t_i, x_i, u_i)
\end{aligned} \tag{3.4}$$

3.3 Examples and Results

Here we demonstrate how HES can improve the stability of periodic trajectories for a variety of hybrid systems without any use of continuous-domain feedback control. While continuous-domain feedback could be implemented into any of these systems and should be in practice, these examples emphasize the stabilization capabilities of HES alone.

The first two examples describe how previously discovered results, paddle juggling [26] and swing leg retraction [3], fit into an HES framework. The final example demonstrates how HES can be used even without any shape parameters and how virtual hybrid events can help stabilize a complicated biped walking system.

3.3.1 Paddle Juggler

The paddle juggler system [26], bouncing a ball with a paddle, is known to be stabilized by impacting the ball with a paddle acceleration in a range of negative values, (3.6), Fig. 3.1. The system state consists of the ball's vertical position and velocity such that $x = [x_B, \dot{x}_B]^T$. This periodic hybrid system can be defined with two hybrid domains (descent and ascent) connected by two guards (impact and apex). The domain D_1 represents the ball's aerial descent phase where $\dot{x}_B < 0$ and D_2 represents the ball's aerial ascent phase where $\dot{x}_B > 0$. The guard set $g_{(1,2)} := x_B - x_P \leq 0$ is defined when the ball impacts the paddle, where the paddle follows a twice differentiable trajectory $x_P(t)$. The reset map $R_{(1,2)}$ is defined by a partially elastic impact law, $R_{(1,2)}([x_B, \dot{x}_B]^T) = [x_B, (1 + \alpha)\dot{x}_P - \alpha\dot{x}_B]^T$, with a coefficient of restitution α . The continuous dynamics in both domains follow unactuated ballistic motion: $\mathcal{F}_1 = \mathcal{F}_2 = [\dot{x}_B, -g]^T$, where

g is the acceleration due to gravity.

Using these definitions, the saltation matrix (2.3) between domains 1 and 2 is constructed:

$$\Xi_{(1,2)} = \begin{bmatrix} -\alpha & 0 \\ \frac{(1+\alpha) \cdot (\ddot{x}_P + g)}{\dot{x}_P - \dot{x}_B} & -\alpha \end{bmatrix} \quad (3.5)$$

Observe that $\ddot{x}_P$ appears in the saltation matrix even though it does not appear anywhere in the definition of the guards, reset maps, or vector fields of the system, making it a shape parameter that can be chosen independently of the periodic orbit.

The guard set $g_{(2,1)} := \dot{x}_B \leq 0$ is defined when the ball reaches the apex of its ballistic motion. Its reset map $R_{(2,1)}$ is identity and there is no change in dynamics. Thus, $\Xi_{(2,1)}$ is identity and has no effect on the variations of the system.

The continuous variational matrices of the paddle juggler can be written exactly in closed form: $A_1(T) = A_2(T) = \begin{bmatrix} 1 & T/2 \\ 0 & 1 \end{bmatrix}$ where T is the total time spent in the air and also the period of the system. The periodicity of the system means the ball spends an equal time, $T/2$, ascending as descending.

The monodromy matrix, Φ is constructed by composing together the continuous variational matrices and saltation matrices such that $\Phi = \Xi_{(2,1)} A_2 \Xi_{(1,2)} A_1$.

For a given periodic bouncing trajectory, the monodromy matrix Φ is almost fully defined except for the shape parameter, $\ddot{x}_P$ in $\Xi_{(1,2)}$. Given the 2-dimensional state space of this problem, the eigenvalues for any given $\ddot{x}_P$ value can be solved for explicitly. We can then solve exactly for where $\psi(\ddot{x}_P) < 1$, giving a stable region of:

$$-2g \frac{1 + \alpha^2}{(1 + \alpha)^2} < \ddot{x}_P < 0 \quad (3.6)$$

This is confirmed in [26], where the simple dynamics of the system allowed the return map to be computed explicitly without the saltation matrix. However, that computation is generally not possible for more complex dynamics.




Figure 3.2: SLIP-like system with actuator and damper in parallel.

3.3.2 2D Hopper

The spring-loaded inverted pendulum (SLIP) is a popular model for dynamic legged locomotion [40–42]. This simple hopper model is effective at capturing dynamic properties of animal and robot locomotion [43] and has been used as a test bed for hybrid controllers [44].

2D Hopper Hybrid Model

Consider a point mass body with a massless leg consisting of a spring, damper, and linear actuator all in parallel, Fig. 3.2. This system has two domains (flight and stance) connected by two guards (touchdown and liftoff). The actuator is activated while in the air to preload the spring, but immediately releases at touchdown and provides no forces during stance. For a periodic trajectory to occur, the actuator must preload the same amount of energy that is dissipated by the damper during stance. The only control authority that exists is of the leg angle in the air, which only affects the dynamics of the body at touchdown.

During flight (D_1), the state of the hopper is represented by the horizontal velocity, vertical position, and vertical velocity of the body: $x = [\dot{x}_B, y_B, \dot{y}_B]^T$. Horizontal position is not included because it is not periodic. During stance, the body position x_s and y_s is defined with the toe at the origin. Horizontal position is added back into the state of the hopper such that $x = [x_s, \dot{x}_s, y_s, \dot{y}_s]^T$. In flight, the dynamics of the body follow ballistic motion, while in stance

Table 3.1: Stability measures of 2D hopper trajectories without and with optimized shape parameters.

Shape parameters	K	$\dot{\theta}$	Stability Measure
Zero	0	0	13.756
Optimal (zero seed)	0.129	-0.015	0.948
Optimal (alternate seed)	0.129	-0.589	0.948

there are also forces applied by the spring and damper.

The touchdown guard is defined by the preload length of the leg ℓ_{td} and angle of the leg θ_{td} such that $g_{(1,2)} := y_B - \ell_{td} \cos(\theta_{td})$. The liftoff guard is crossed when the force exerted by the spring-damper, F_{sd} , becomes zero: $g_{(2,1)} := F_{sd}$. There is no change in physical state of the system at the hybrid events and the reset maps only characterize the change in coordinates between domains.

Given a set of model parameters, a trajectory from an initial condition depends only on ℓ_{td} and θ_{td} . ℓ_{td} is held fixed, but θ_{td} is modulated from its nominal position $\bar{\theta}_{td}$ at time $\bar{t}_{td}$ in two ways. A proportional feedback term with gain K is added to stabilize the forward velocity of the system around a nominal $\bar{x}$ and angular velocity $\dot{\theta}$ of the massless leg is also free to be chosen. K and $\dot{\theta}$ are shape parameters that can be used to stabilize this system.

$$\theta_{td} = \bar{\theta}_{td} + K(\dot{x} - \bar{\dot{x}}) + \dot{\theta}(t - \bar{t}_{td}) \quad (3.7)$$

2D Hopper HES Results

For a chosen initial apex height of 2.5 with a forward velocity of 2, ℓ_{td} and $\bar{\theta}_{td}$ were solved for to produce a nominal orbit, though the following results generalize for any choice of feasible values.

With fixed shape parameters $[K, \dot{\theta}] = [0, 0]$, the system is highly unstable. K and $\dot{\theta}$ can be optimized following (3.3) to improve the stability of this orbit. Doing so results in optimal shape parameters $[K, \dot{\theta}] = [0.129, -0.015]$ that stabilize the trajectory, Table 3.1. Setting $\dot{\theta} = -0.015$ rad/s is a slow retraction rate, but there exists an interval of values $\dot{\theta} \in (-0.5892, -0.015)$ that give equivalently minimal stability measures for a fixed K value.




Figure 3.3: Error of perturbed initial states for the 2D hopper asymptotically decrease to zero. Transparent lines represent each of the 20 trials, while the bold line represents average error at each step. Convergence is neither monotonic nor very fast, but this is expected with asymptotic stability.

The results were confirmed in simulation by initializing 20 random perturbed points in a 0.1 radius ball around the nominal initial condition. Each of these trials was simulated for 100 steps and the error (2-norm of the difference in perturbed state x and nominal state $\tilde{x}$) at apex was recorded at each step, Fig. 3.3. The zero shape parameter trajectories are not shown in the figure as every trial diverged within just 5 steps.

2D Hopper Discussion

The feedback term of (3.7) is based on the Raibert stepping controller [45], which was utilized to great success for stabilizing early running robots. Other works have found that this simple controller is effective on more complex models [46].

Another stabilizing property of legged locomotion that has been studied is swing-leg retraction [3,29]. It was noted in [3] that a 2D SLIP was able to run stably if it impacted the ground within a range negative angular leg velocities $\dot{\theta}$.

The results of a negative $\dot{\theta}$ and positive K agree with qualitative expectations from [3] and [45]. While a formal equivalency is yet to be proven, this is significant because the HES shape pa-

parameter optimization has no a priori knowledge that would bias its results to match these works. HES synthesizes two independently generated controllers and produces shape parameter values that stabilize an orbit. This evidence supports the potential for HES to explore other stabilizing shape parameters that are not as well studied.

3.3.3 Walking Biped Trajectory Optimization

For a legged system with non-massless legs, the leg velocity shape parameter disappears as it is no longer independent of the trajectory. Without shape parameters, an HES trajectory optimization can choose timing, state, and input parameters along with injecting virtual hybrid events to discover stable orbits.

Walking Biped Hybrid Model

In this example, we consider a fully-actuated compass walker [47] with knees, Fig. 3.4. This biped model consists of two legs connected by an actuated hip joint. Each leg is separated into two sections, the upper leg (thigh) and lower leg (shank), which are connected by an actuated knee joint that has a hard stop when the thigh and shank are aligned. The ankles are also actuated.

We restrict the gaits to be left-right symmetric and exclusively consist of single stance phases. The stance leg is locked such that its shank and thigh are aligned with each other until liftoff. There are 3 points of actuation at the hip, swing knee, and stance ankle. The state space is defined by three angles relative to vertical: stance leg, swing thigh, and swing shank, denoted $(\theta_1, \theta_2, \theta_3)$.

This system has two domains. D_1 is the unlocked knee domain where the swing leg thigh and shank can swing freely while we enforce that $\theta_3 < \theta_2$. D_2 is the locked knee domain where the thigh and shank are constrained to be aligned with each other ($\theta_2 = \theta_3$). In this domain, there are only two actuators because the swing knee can no longer exert torque. The dynamics of this model are described in [47].

The kneestrike guard set, between the unlocked and locked knee domains, is $g_{(1,2)} := \theta_2 - \theta_3$ and the touchdown or toestrike guard set is $g_{(2,1)} := \theta_1 + \theta_2$. The reset maps at kneestrike and




Figure 3.4: Biped walker system with kneestrike and toestrike hybrid events.

toestrike are also computed in [47].

We add discrete changes in the inputs that induce virtual hybrid events to analyze their utility in stabilizing walking trajectories. Specifically, we choose to include 5 virtual hybrid events in D_1 and 2 more virtual hybrid events in D_2 , where the values of inputs between virtual hybrid events are decision variables for the optimization. The virtual guard functions are chosen such that $g_{v_i} := \theta_2 - \theta_3 + p_i$ for the first 5 virtual hybrid events and $g_{v_i} := \theta_1 + \theta_2 + p_i$ for the last 2 virtual hybrid events for some offset p_i that is also chosen by the optimization.

A direct collocation method was used with the cost consisting of the stability measure and a regularization on the input. Dynamics and periodicity constraints were included along with a ground penetration constraint. The initial conditions of the system, given as the state after touchdown, were allowed to vary within a bounded range.

Walking Biped HES Results

In this experiment, three trajectories were compared to examine how HES can generate stable trajectories and the effect that virtual hybrid events have in further improving stability. A trajectory without HES was produced as a control, with its objective to minimize energy expended by using just an input regularization term in the cost. This minimum energy (ME) trajectory is comparable to how conventional robot locomotion trajectories are generated. Two HES trajectories were generated, one with virtual hybrid events (HES w/ VHE) and one without (HES w/o VHE).

Table 3.2: Stability results for the walking biped optimization.

Trajectory	Stability Measure	Energy Cost
ME	7.8123	0.985
HES w/o VHE	0.4715	1.337
HES w/ VHE	0.3266	3.450




Figure 3.5: Error of perturbed Minimum Energy (ME), Hybrid Event Shaping without virtual hybrid events (HES w/o VHE) and HES w/ VHE trajectories over several steps. Center lines show the average error at each step and shaded regions indicate ± 1 standard deviation. ME trajectories become highly divergent within 4 steps, while both HES trajectories appear convergent after 10 steps. The initial increase in error of the HES trajectories is allowable and is not considered by the stability measure.

The ME trajectory is highly unstable, while the both HES trajectories are stable with the trade off of a higher input cost. Specifically, HES w/ VHE has the lowest stability measure and highest energy cost, whereas HES w/o VHE was in between for both stability and cost, Table 3.2.

The stability properties of the generated trajectories were confirmed through simulation. 50 randomized trials of each trajectory were initialized with perturbations in position and velocity between $(-0.01, 0.01)$. Over a sequence of 10 steps, the state error at each step was tracked for each trial. Fig. 3.5 shows that after 10 steps, the HES trajectories have nearly converged back to the nominal trajectory whereas the ME trajectories diverge quickly. The HES w/ VHE trajectory converges to a smaller error after 10 steps compared to the HES w/o VHE trajectory, which supports the findings of the stability measure.

3.4 Conclusion

While the idea of hybrid event control is not novel, hybrid event shaping provides a generalized method to analyze the stability of hybrid orbits and select hybrid event parameters to optimize stability. HES unifies results of previous simple hybrid event controllers while also being compatible with trajectory optimization techniques to produce stable trajectories for complex systems. HES computes the derivative of the stability measure, improving computational efficiency compared to previous stability optimization methods. Compared to previous work, HES does not rely on human observation and tuning to design stabilizing hybrid event parameters.

In this work, there was no use of continuous-domain feedback that is commonly utilized in hybrid systems control. We believe that hybrid event shaping is one aspect that can be used in conjunction with continuous-domain feedback to improve the success rate of robots performing dynamic behaviors in real-world settings. This is not be prohibitively complex because saltation matrices are not affected by feedback control laws in the continuous domains. The next chapter expands the ideas of HES to trajectories continuous-domain feedback control to produce even more stable closed-loop trajectories.

Chapter 4

Convergent iLQR: Generating Closed-Loop Convergent Hybrid Trajectories

4.1 Introduction

Legged robotics research has increasingly focused on enabling highly dynamic and agile motions such as jumping, leaping, and landing [48–51]. These capabilities are inherently highly unstable and implementing them reliably requires intelligent planning of both feedforward and feedback controllers. This would improve legged robot performance in applications such as extraterrestrial or urban environment navigation where jumping up on ledges or leaping across chasms may be necessary. However, jumping and leaping are dangerous maneuvers, with failure often resulting in catastrophic outcomes for the robot.

What makes these actions challenging is that they induce trajectories that are both hybrid and underactuated, which doubly contribute to the difficulty in controlling legged robots. Broadly speaking, a system is hybrid if it undergoes discrete changes in state and/or dynamics [10,11], and it is underactuated if there exists a direction of acceleration in state space that can not be commanded by any valid input [52, Ch. 1.2]. Even when an underactuated system is controllable, driving the system to a desired target state may require significant time and control effort,




Figure 4.1: Similar quadruped gaits tracked with equivalent LQR controllers display enormous differences in final tracking performance. Results for 50 paired trials of trajectories sampled from an initial error covariance of 10^{-2} in all directions are shown. The top blue trajectory was generated with a standard (vanilla) iLQR algorithm, while the bottom green trajectory was generated with our novel χ -iLQR, which improves the average closed-loop convergence by 92%. The horizontal distance between the displayed frames is exaggerated for visual clarity. Only the convergence of the position states is represented, and does not indicate convergence of the velocity states.

neither of which may be readily available. For instance, a robot jumping in the air can not arbitrarily choose how much time it has until its feet touchdown on the ground. This means that the controller needs to spend a lot of effort to correct tracking errors prior to touchdown, or else discontinuous, unbounded saltation effects [2] can cause arbitrarily large divergence if incoming errors are not sufficiently mitigated, e.g. with grazing impacts. Increasing control gains is one possible solution to improve stability, though that strategy comes at a large drawback of worsening robustness in the face of modelling errors and uncertainties [53, Ch. 13]. Instead, this work leverages nonlinearities in continuous and hybrid dynamics that make some trajectories easier to stabilize than others, even under equivalent feedback controllers.

This chapter presents a novel adaptation of the iLQR trajectory optimization algorithm that improves closed-loop convergence under an equivalent feedback controller (i.e. without changing the LQR controller weights), as demonstrated in Fig. 4.1. Our simulation results show that this convergent iLQR (χ -iLQR) achieves three simultaneous improvements over standard iLQR:

superior tracking performance from initial perturbations, reduced feedback control effort over the trajectory, and improved robustness to large initial errors. Compared to existing methods, χ -iLQR has two additional key strengths. Firstly, it is based on an analysis that is simple to compute compared to methods such as sum-of-squares. Additionally, χ -iLQR captures the local tracking performance of a closed-loop trajectory, which directly predicts experimental results.

4.2 Related Works

A strategy that has been used to enable dynamic yet precarious behaviors for legged robots is leveraging highly accurate, complex models and full-body trajectory optimization to plan precise motions [50,54]. While these methods incorporate feedback controllers to stabilize the generated trajectories, there has been little focus on how these feedback controllers should be designed along with planned trajectories to stabilize closed-loop systems under error and uncertainty.

Robust trajectory planning has been implemented for smooth systems like wheeled robots, with some recent results being adapted to hybrid systems like legged robots. These works have focused on optimizing over uncertainties in system dynamics, such as unknown disturbances and modelling errors [55–57]. These methods plan trajectories under the anticipation of some worst-case or average-case disturbance sequence. For example, [55] designs robust closed-loop trajectories for smooth systems by optimizing the volume reduction of an ellipsoidal disturbance set, but was not applied to hybrid systems. Risk-sensitive planning and control is an alternate method that optimizes over the variance of a cost distribution that evolves through the trajectory [58,59]. Other approaches present trajectory optimization algorithms for legged robots over uncertain terrain [60] and compute a forward reachable set to bound closed-loop errors [61]. Many of these methods require the distribution of errors to be prespecified, which is not always clear how to tune. Additional actuation, such as reaction wheels or tails, also relieves the difficulties of underactuated systems [62,63]. However, this comes with obvious tradeoffs of increased cost, size, and weight.

Separately, consider the problem of quantifying the stability or convergence properties of a system. A very popular method is Lyapunov analysis, where the existence of a positive definite differentiable scalar function with negative definite derivatives, called the Lyapunov function, can guarantee asymptotic stability of the system [32]. While effective, Lyapunov functions can be difficult to compute, particularly for hybrid systems, and can require methods such as sum-of-squares [64] or machine learning [65] to be tractable. A similar strategy known as control barrier functions, which restricts the system from entering some set of undesirable states, has been successfully implemented on legged robot hardware [66], but has the same drawback as Lyapunov functions.

A different strategy to analyze the stability and convergence of trajectories is contraction analysis [67], which tracks the distance between two close trajectories. If this distance monotonically decreases over the trajectory, then the system is contractive and local asymptotic stability can be guaranteed [67]. Contraction analysis has been incorporated into path planning and trajectory optimization algorithms on smooth systems [68,69], but applying contraction analysis to hybrid systems is difficult because many mechanical hybrid systems are not contractive at hybrid events [17]. [4] loosened the contraction criterion and optimized the stability of open-loop periodic orbits using monodromy matrix analysis. Here, we extend that work by generalizing to non-periodic trajectories under feedback control.

4.3 iLQR for Hybrid Systems

The iterative linear quadratic regulator (iLQR) is a trajectory optimization method that also computes LQR feedback gains over the generated trajectory [70]. iLQR is convenient because compared to other trajectory optimization methods like direct collocation [34], it is less computationally intensive and guarantees a feasible trajectory. We draw from recent work that adapts the iLQR algorithm for use on hybrid dynamical systems [22,71].

In brief, iLQR solves the optimal control problem over N discretized timesteps:

$$\min_U \ell_N(x_N) + \sum_{i=0}^{N-1} \ell_i(x_i, u_i) \quad (4.1)$$

$$\text{where } x_0 = x(0) \quad (4.2)$$

$$x_{i+1} = \phi(t_{i+1}, t_i, x_i, u_i) \quad \forall i \quad (4.3)$$

where $\ell_i(x_i, u_i)$ and $\ell_N(x_N)$ represent the nonlinear stage cost and terminal cost, respectively, $X := \{x_0, x_1, \dots, x_N\}$ is a sequence of states with $x_i \in \mathbb{R}^n$ the system state at timestep i and $U := \{u_0, u_1, \dots, u_{N-1}\}$ is a sequence of control inputs with $u_i \in \mathbb{R}^m$ the control input at timestep i . We also record the sequence of domains $M := \{D_0, D_1, \dots, D_N\}$ with D_i the domain at timestep i such that $x_i \in D_i$. ϕ is the aforementioned flow of the trajectory.

iLQR computes gradient and Hessian information of the cost, which results in a quadratic approximation of the cost function. As such, the state and terminal costs can equivalently be simplified as quadratic functions such that the cost function is simplified to:

$$J = x_N^T Q_N x_N + \sum_{i=0}^{N-1} x_i^T Q_i x_i + u_i^T R_i u_i \quad (4.4)$$

with $Q_i, Q_N \in \mathbb{R}^{n \times n}$ and $R_i \in \mathbb{R}^{m \times m}$ all positive definite.

iLQR solves the optimal control problem by alternating between forward passes that simulate the system under a given control input sequence, and backward passes that solve for a new locally optimal control sequence. In the backward pass, the value function, which is the optimal cost to go at any timestep, is propagated through the trajectory in reverse, and gives locally optimal feedforward inputs and feedback gains at each timestep. Computing the value function relies on gradient and Hessian computations of the cost function and Jacobians of the dynamics, which equates to computing the linearized variational equations discussed in Sec. 2.3. For much greater detail of iLQR for hybrid systems, see [22,71].

4.4 Convergent iLQR

Here we present a novel trajectory optimization algorithm called convergent iLQR or χ -iLQR, summarized in Algorithm 1. In this method, the convergence measure (2.13) is added to the cost function from (4.4) such that the algorithm minimizes:

$$J_\chi = Q_\chi \chi + x_N^T Q_N x_N + \sum_{i=0}^{N-1} x_i^T Q_i x_i + u_i^T R_i u_i \quad (4.5)$$

where Q_χ is a scalar weighting parameter.

Typically in iLQR, the cost function J is evaluated after each forward pass, since it is only dependent on the states and inputs of the most recent trajectory. However, in this case the convergence measure portion of the cost function is dependent on the feedback gains generated by the algorithm. This means that the gradient and Hessian terms of the cost function rely on the feedback gains that are being updated at every timestep in the backward pass. Due to this, the cost function derivatives are highly coupled with the gains and become convoluted to compute.

To resolve this, we propose executing two separate backward passes that each compute a different set of gains. First, the tracking backward pass computes the feedback gains that will be used as the LQR tracking controller gains and to compute the convergence measure. It is equivalent to the backward pass in standard iLQR using the cost function J (4.4), which solves the Riccati equation for the most recent trajectory. With the gains generated in the tracking backward pass K_t , the convergent cost function J_χ (4.5) can be computed. The search backward pass takes J_χ from the tracking backward pass and computes the gradients of the convergent cost function with controller gains K_t . The feedforward inputs k_s and the feedback gains K_s from this pass are used to search for an improved trajectory in the forward pass.

Since J_χ is returned by the tracking backward pass, a line search is performed after this function call to guarantee the reduction of the cost function J_χ . If the line search condition is not satisfied, the forward pass and tracking backward pass are looped until the line search condition is passed.

Within the search backward pass, iLQR requires computation of the gradient and Hessian of χ . The derivatives of χ can be computed by leveraging the singular value decomposition of $\Phi = USV^T$ where S is a diagonal matrix of singular values and the columns of U and V are the left and right singular vectors, respectively. This largely mirrors the derivative formulation in the previous chapter. χ is the largest singular value of Φ and let u_χ and v_χ be its corresponding left and right singular vectors. Following [72], the derivative of χ with respect to the state at timestep i is:

$$\frac{\partial \chi}{\partial x_i} = u_\chi^T \frac{\partial \Phi}{\partial x_i} v_\chi \quad (4.6)$$

$\frac{\partial \Phi}{\partial x_i}$ in turn can be computed by using the product rule along with leveraging the fact that only $\tilde{A}_i$ and $\Xi_{(i,i+1)}$ are functions of x_i , and all other $\tilde{A}$ and Ξ terms have zero derivatives with respect to x_i . For notational brevity, let:

$$P_i = \tilde{A}_N \cdots \Xi_{(i+1,i+2)} \tilde{A}_{i+1} \quad (4.7)$$

$$O_i = \Xi_{(i-1,i)} \tilde{A}_{i-1} \cdots \Xi_{(1,2)} \tilde{A}_1 \quad (4.8)$$

such that $\Phi = P_i \Xi_{(i,i+1)} \tilde{A}_i O_i$ and $\frac{\partial P_i}{\partial x_i} = \frac{\partial O_i}{\partial x_i} = 0$. Thus:

$$\frac{\partial \Phi}{\partial x_i} = P_i \frac{\partial \Xi_{(i,i+1)}}{\partial x_i} \tilde{A}_i O_i + P_i \Xi_{(i,i+1)} \frac{\partial \tilde{A}_i}{\partial x_i} O_i \quad (4.9)$$

Derivatives with respect to the input u_i follow equivalently. To improve computational efficiency, O_i at each timestep can be computed recursively during the forward pass and each P_i can be computed recursively in the tracking backward pass. Since the rollout does not yet have feedback gain information, the initial O values must be computed separately.

Because the scalar χ is derived from the norm of the matrix Φ , the gradient of χ relies on computing a 3-dimensional tensor of $\tilde{A}_i$ and Ξ derivatives, and the Hessian of χ is computed from a 4-dimensional tensor of matrix second derivatives. While recent work has enabled faster

Algorithm 1 Convergent iLQR Algorithm

```
Initialize  $U, Q_\chi, Q_N, Q_i, R_i, n_{\text{iterations}}$   
 $X, U, M, J \leftarrow \text{Rollout}(U)$   
 $K_t, J_\chi, P \leftarrow \text{TrackingBP}(X, U, M, J)$   
 $O \leftarrow \text{ComputeO}(X, U, M, K_t)$   
for  $i \leftarrow 1$  to  $n_{\text{iterations}}$  do  
   $k_s, K_s \leftarrow \text{SearchBP}(X, U, M, J_\chi, O)$   
  repeat  
     $X, U, M, J, O \leftarrow \text{ForwardPass}(X, U, M, k_s, K_s)$   
     $K_t, J_\chi, P \leftarrow \text{TrackingBP}(X, U, M, J)$   
  until  $\text{LineSearchIsSatisfied}(J_\chi)$   
return  $X, U, M, K_t$ 
```

Table 4.1: Rocket Hopper Controller Parameters

Trial	LQR Parameters			
	Q_χ	Q_N	$R_{i_{\text{air}}}$	$R_{i_{\text{stance}}}$
1	50	500I	0.01I	0.1I
2	50	800I	0.005I	0.01I
3	50	250I	0.02I	0.05I
4	75	500I	0.01I	0.01I

computation of second derivatives of dynamics [73] which can aid in the computation of the 3-D tensor derivatives, computing 4-D tensor derivatives is generally untenable. Instead, numerical methods like finite differences for gradients and BFGS [74] for Hessians can perform at reasonable speed. In order to approach real-time computation, it is likely that the full Hessian of χ is not necessary to find an appropriate search direction and that a partial computation or even leaving out the Hessian completely is sufficient to compute optimal trajectories. Future work will address this gap. Nonetheless, the algorithm in its current form can still be useful for offline planning for trajectories that are expected to have a high degree of risk, such as leaping across ledges or traversing narrow beams. In real-world applications, it can be acceptable for a robot to pause and plan a safe trajectory before executing these dangerous maneuvers.

4.5 Examples and Results

In this section, we demonstrate the convergence improvements of our method on a spring hopper system and a planar quadruped robot model. Simulation results show that the improved convergence measure correlates with an improvement in average tracking performance, robustness to large disturbances, and feedback control effort. Both examples were implemented in MATLAB, with forward simulations using the `ode113` function. Cost function gradients were computed using (4.9), derivatives of $\tilde{A}_i$ and $\Xi_{(i,i+1)}$ were computed with finite differences, and Hessians were computed with BFGS.

4.5.1 Rocket Hopper

Rocket Hopper Model

This system is made up of a point mass body with a single massless spring leg. The state of the hopper is characterized by the positions x_B, y_B of the body, the angle θ of the leg and their derivatives $\dot{x}_B, \dot{y}_B, \dot{\theta}$ such that the full state is a 6×1 vector. The system has two domains: an aerial phase D_1 and a stance phase D_2 . Taking a constant ground height at zero gives a touchdown guard function $g_{(1,2)}$ that is the height of the foot and a liftoff guard function $g_{(2,1)}$ that is the ground reaction force applied by the spring leg. Both reset maps $R_{(1,2)}$ and $R_{(2,1)}$ are identity since position and velocity are continuous.

The system has two inputs: a hip actuation and an actuation in the direction of the leg. In the air, this allows the hopper to rotate the leg around the body and exert a propulsion in the direction of the leg, somewhat akin to a rocket, though this force can approximate forces from other legs or actuators. A small rotor inertia in the air ensures the dynamics are well-conditioned when controlling the massless leg. In stance, the hip torque and rocket force exert ground reaction forces on the body. The body mass of the hopper was chosen as 1 kg, spring constant as 250 N/m, and resting leg length as 0.75 m.

Rocket Hopper Results

The objective for this system is to begin in the air at rest with a height of 2 m and end in the air at rest with the same height displaced 0.2 m horizontally. The system is given 1.5 s for this trajectory.

We generated four trials of paired trajectories with varied weighting parameters, shown in Table 4.1, and compared the performance of the standard (vanilla) iLQR method (where $Q_\chi = 0$) to χ -iLQR. There is no reference trajectory to track, so Q_i is zero for all trials.

For 3 of these trials, vanilla iLQR generated a trajectory with $\chi > 1$, meaning the worst-case error direction was expansive, see Table 4.2. χ -iLQR decreases every convergence measure to below 1 so that all error directions are reduced over the trajectory. On average, χ -iLQR decreased χ by 28.79% compared to the vanilla method.

To validate these trajectories, each closed-loop trajectory was simulated 100 times with small random initial perturbations in both positions and velocities with covariance matrix $\text{cov}(X_0) = 10^{-4}I$. A small covariance was chosen so that the linearizations assumed in the convergence measure and LQR control are valid. For each simulation run, the initial error δx_0 and the final error δx_f were recorded, along with the sequence of control inputs $V := \{v_0, v_1, \dots, v_{N-1}\}$. Note that these inputs are distinct from the nominal feedforward inputs to the system U because there is additional feedback effort exerted by the actuators. The same set of initial errors was used to test both trajectories.

Two values were recorded during each simulation run to characterize the convergence properties of the trajectories. The first is the error ratio, defined as the ratio of the final error 2-norm to the initial error 2-norm:

$$E = \frac{\|\delta x_f\|_2}{\|\delta x_0\|_2} \quad (4.10)$$

A lower error ratio means better tracking performance, and $E < 1$ indicates a net reduction in error on average. The second value is the feedback effort, which is the sum of squares of the

Table 4.2: Rocket Hopper Convergence Measure and Simulation Results

Trial	Convergence Measure			Mean Simulated Error Ratio			Mean Simulated Feedback Effort		
	Vanilla	χ -iLQR	%Diff.	Vanilla	χ -iLQR	%Diff.	Vanilla	χ -iLQR	%Diff.
1	1.01	0.71	-29.70%	0.42	0.33	-21.72%	$1.74 \cdot 10^{-5}$	$1.6 \cdot 10^{-5}$	-7.16%
2	0.78	0.51	-34.50%	0.32	0.24	-25.74%	$4.66 \cdot 10^{-5}$	$3.25 \cdot 10^{-5}$	-30.31%
3	1.14	0.94	-17.70%	0.49	0.45	-8.00%	$7.12 \cdot 10^{-6}$	$7.05 \cdot 10^{-6}$	-1.01%
4	1.01	0.68	-33.24%	0.41	0.32	-21.73%	$1.89 \cdot 10^{-5}$	$1.62 \cdot 10^{-5}$	-14.20%

difference between V and U .

$$F = \sum_{i=0}^{N-1} (v_i - u_i)^2 \quad (4.11)$$

Table 4.2 shows that the simulation results support our assertion that an improved convergence measure correlates with an improved mean tracking performance and feedback effort. The mean error ratio and feedback effort over the 100 simulations were both lower for trajectories generated with χ -iLQR. The average improvement over the four trials was 19.30% for mean error ratio and 13.17% for feedback effort.

None of the simulated runs had an error ratio greater than one, which is sensible since the worst-case direction occurs with probability zero. However, even if none of the sampled initial errors aligned exactly with the worst-case direction predicted by the fundamental solution matrix, nearby initial error directions still see improvement in convergence, which explains the improvement in mean simulated error ratio.

4.5.2 Planar Quadruped

Here we demonstrate the improvements of χ -iLQR on a more complex robot model akin a standard quadruped robot. The model is simplified as a planar quadruped, meaning that all movement occurs in the sagittal plane and the left-right pairs of legs move identically.

Planar Quadruped Model

In the sagittal plane, we can model the robot with 7 positional states. x_B, y_B, θ_B are the position and orientation of the body. The front and back sets of legs each have two states for the hip angle α_f, α_b and knee angle β_f, β_b . Thus the full state is dimension 14.

This system has four domains: the aerial domain D_1 , front stance domain D_2 , back stance domain D_3 , and full stance domain D_4 . The impact guard function is the height of the foot and the guard function for liftoff is the vertical ground reaction force. The dynamics of the robot body in the aerial phase follow ballistic motion, while the legs are simplified to be massless while including the aforementioned rotor inertia. The impact reset map for each foot consists of a discrete update to the hip and knee velocities, while the body states are unchanged due to the massless legs. The liftoff reset map is identity. The input vector for this system is 4-dimensional to actuate the hip and knee joints.

In this model, parallel torsion springs are added to the knee joints. Parallel joint springs have been utilized to mimic tendons found in animals [75] that increase the energy efficiency of legged locomotion [76,77].

Due to the resonance of the natural spring dynamics, controlling these systems requires special care [78,79]. For example, [80] solved for optimal gait timings to leverage resonant spring frequencies. These spring models of legged robots are good candidates for χ -iLQR because the dynamics of the stance phase depend strongly on the leg configuration at touchdown. Thus, a small error in leg states at touchdown can have a large effect on tracking performance.

The inertial and dimensional properties were chosen to match the Ghost Robotics Spirit 40 quadruped. The added torsional knee spring has a spring constant $75 \text{ N} \cdot \text{m} \cdot \text{rad}^{-1}$ and rest angle 1.2 rad .

Planar Quadruped Results

The trajectory optimization task for the planar quadruped is to generate a gait with a forward velocity of 0.25 m/s . The robot begins in the air with a body height of 0.3 m . The hip joints begin

Table 4.3: Mean convergence results for the quadruped model for vanilla and χ -iLQR with covariance magnitude 10^{-4}

	Vanilla	χ -iLQR	%Difference
Nominal Cost	65.58	74.39	+13.43%
Convergence Measure	60.52	41.35	-31.68%
Mean Simulated Error Ratio	6.66	4.78	-28.23%
Mean Simulated Feedback Effort	0.016	0.013	-16.56%

at an angle of 0.6 rad and the knee joints begin at 1.2 rad. The terminal target state is translated 0.0875 m in the x-direction from the initial state. The trajectory is given 0.35 s to execute. We choose to set a constant input weight of $R_i = 5 \cdot 10^{-4}I$. The terminal weight is $Q_N = 500I$ and the convergence weight for χ -iLQR is $Q_\chi = 1$.

We set up a similar experiment to the prior example, with the addition of simulating over a range of covariance magnitudes. This is done to evaluate the basin of attraction of each trajectory over larger initial errors that introduce greater nonlinear effects. The two trajectories were evaluated with 6 sets of 100 paired simulation runs with random initial error covariance magnitudes of 10^{-4} , $5 \cdot 10^{-4}$, 10^{-3} , $5 \cdot 10^{-3}$, 10^{-2} , and $5 \cdot 10^{-2}$ in each direction. The lowest covariance magnitude of 10^{-4} approximates local linear behavior well, while $5 \cdot 10^{-2}$ is the maximum magnitude before some trials begin with the robot's feet below the ground.

Table 4.3 shows the nominal cost, convergence measure, mean simulated error ratio, and mean simulated feedback effort of the two trajectories at the covariance magnitude 10^{-4} . The nominal cost of the convergent trajectory increases as expected, while the convergence measure and simulation values improve. We believe this trade off between nominal cost and the convergence performance of the trajectory is a valuable tool for roboticists to have access to and tune appropriately. Compared to the vanilla trajectory, the mean simulated error ratio for the convergent trajectory at this small covariance magnitude was 28.23% less and the mean simulated feedback effort was 16.56% less, again showing the simultaneous improvement in both tracking performance and feedback control effort. Fig. 4.2 displays a histogram of the error ratio for each of the trials, and shows the improvement made by the convergent trajectory.




Figure 4.2: Histogram of error ratio for 100 paired simulated trials of the quadruped model with small initial perturbations. Error ratio is the 2-norm of final errors divided by the 2-norm of initial errors.

As the magnitude of initial errors grows, the performance of the LQR tracking controller becomes worse due to the increase in nonlinear effects. Fig. 4.3 shows the simulation results for each trajectory over a range of initial error covariance magnitudes. Each pair of lines indicates the success rate of the respective closed-loop trajectories at maintaining error ratios of less than 50, 10, and 5 respectively. An error ratio of greater than 50 is representative of a catastrophic failure, which the vanilla trajectory encounters at a covariance magnitude of $5 \cdot 10^{-4}$, while the convergent trajectory first experiences a failure at covariance magnitude 10^{-2} . This difference in performance suggests the convergent trajectory is more robust to larger initial errors and nonlinearities.

Even with the convergence improvements in the χ -iLQR trajectory, the controller is still not able to reduce errors in all directions. The simulation results found that most of the time, there was some growth in error, which is reasonable since the body dynamics are fully unactuated in the aerial phase and the system undergoes multiple hybrid events. A combination of higher feedback gains and a more global footstep planner could be able to grant this system full conver-




Figure 4.3: Plots showing success of LQR controllers at tracking vanilla (V) and convergent (χ) trajectories over various initial perturbation covariances. V_{50} and χ_{50} indicate proportion of trials where error ratio was below 50, V_{10} and χ_{10} below 10, and V_5 and χ_5 below 5.

gence. Even so, this work can be valuable to ensure that between iterations of a global footstep planner, the closed-loop system does not diverge too far from its target trajectory and experience a failure, such as falling over.

4.6 Conclusion

In this work, we present a novel trajectory optimization method, χ -iLQR that optimizes over the worst-case error growth of a hybrid trajectory. This method is based on the fundamental solution matrix, which maps the evolution of perturbations through a trajectory. Incorporating the saltation matrix into the fundamental solution matrix allows for straightforward handling of hybrid events. The simulation results presented on two legged robot models demonstrate that this method produces trajectories with improved tracking performance, decreased feedback actuation effort, and improved robustness to large perturbations. Even for a quadrupedal trajectory that was very difficult to track, χ -iLQR produced a trajectory that was superior at avoiding failures.

Chapter 5

The Effect of Gait Parameters on Safe Quadrupedal Robot Locomotion

5.1 Introduction

Robotics has long had the goal of performing tasks in precarious work environments where human laborers are at risk of bodily harm. Recently, quadruped robots have begun to make this goal a reality, being deployed to serve areas such as construction sites [81], offshore electrical substations [82], and mine tunnels [83]. In these situations, robots are required to navigate complex terrain such as steep slopes and narrow pathways. Not only do robots need to be able to plan feasible trajectories through these environments, but they must also do so safely in the presence of disturbances and errors. While many algorithms have been produced to improve the robustness of legged locomotion [5,55,60,84,85], these methods still rely on non-trivial parameter tuning. In trajectory optimization, it is common that some subset of contact sequence, contact timings, gait period, speed, etc. must be pre-specified, and good performance is dependent on intelligent choices of these parameters. Robust control methods tend to have additional parameters that require tuning such as a risk sensitivity parameter [84] or an estimate of disturbance distributions [55]. In order to adequately tune these parameters, operators must have an

understanding of how gait attributes like speed and duty factor affect the performance of the robotic system.

There have been many investigations over several decades on how gait parameters such as step length, gait frequency, and duty factor affect the performance of locomotion in terms of efficiency, speed, and stability. These works have studied many kinds of legged systems, including biological creatures like humans [86] and horses [87], as well as robotic bipeds [88] and quadrupeds [6]. A detailed review of works in this area can be found in Sec. 5.2. One well-known result connecting gaits to performance that has been demonstrated on both horses [87] and quadrupedal robots [6,89] is that the 4-beat walking gait is most efficient at low speeds while the 2-beat trotting gait is advantageous at higher speeds. We replicate this result in our investigation, shown in Fig. 5.2, comparing the perfect walking and trotting gaits, meaning that simultaneous touchdown and liftoff occurs between consecutive feet. With the 12-degree-of-freedom quadruped robot model defined in Sec. 5.3.2 and the cost of transport metric defined in 5.3.3, we find that the perfect walk is energetically advantageous up to about 1.5 m/s and the perfect trot is advantageous beyond this threshold. Based on these results, prior works have concluded that the choice of gait predicts the efficiency of locomotion. However, we present evidence in the rest of this work that indicates the existence of a confounding variable (namely, duty factor) that is a more direct predictor of robot performance.

Another understudied aspect of legged stability is performance on narrow pathways, which we will call beams. Robots navigating through mine tunnels or crowded urban sidewalks may experience circumstances that constrain the stance width of their feet. In fact, one advantage that legged robots have over their wheeled counterparts is that they can adapt their stance width to maneuver across narrow pathways. As a legged robot's stance width narrows, the support polygon, which defines an area where the center of mass is easily stabilized, shrinks. Having a large support polygon is an assumption utilized by many common legged robot controllers, like the zero-moment point method [90]. To counteract this deficiency, prior works have relied on inertial elements to stabilize the robot, using the torso of the robot itself [91] or by augmenting




Figure 5.1: State-of-the-art planning algorithms have difficulty navigating narrow, beam-like environments without significant modifications. In this work, we demonstrate that understanding the effect of gait parameters on locomotion performance can provide straightforward solutions for navigating beam-walking scenarios.

the robot with a tail [92] or reaction wheel [62]. In contrast, we aim to investigate how stability can improved by modulating nominal gait parameters, without need for specialized control strategies or additional hardware.

The objective of this work is to provide insights into the relationship between gait parameters such as duty factor, speed, and stance width and performance attributes like efficiency and stability. We demonstrate these relationships on a 12-degree-of-freedom quadruped robot design that is common in the field today [93,94], but to the authors’ knowledge has not been examined in this type of study. In addition, we show that these gait parameter relationships hold for two distinct planning and control methods: a whole-body trajectory optimization with LQR feedback and a centroidal MPC planning framework. These two investigations provide support that the results of this paper are generalizable to modern state-of-the-art methods [95–98]. These results are crucial for enabling robots to adapt more effectively to unknown environments without the need for additional hardware. Identifying key gait parameters can also enhance data-driven control methods by reducing the training time required to discover these factors.




Figure 5.2: Replication of results from [87,89] indicating that the perfect gait is more efficient at lower speeds and the perfect trot gait at high speeds. Note that these curves are cross sections of Fig. 5.3a at the 0.5 duty factor trot and 0.75 duty factor walk.

In this work, we explore two primary research questions:

- **RQ1: How do the gait parameters of duty factor, speed, and stance width relate to the efficiency and stability performance of two common quadrupedal gaits: the 4-beat walk and trot?**
- **RQ2: Given a stance width constraint, can a combination of gait parameters be generated that can traverse a beam safely, while preserving reasonable speed and efficiency?**

5.2 Related Works: Effects of Gait Parameters in Legged Locomotion

A body of literature has investigated the relationship of gait parameters to locomotion performance of various biological and robotic legged systems. The most common biological system that has been studied is human bipedal locomotion. Experiments found that human walking

speed is highly correlated with both step frequency and step length, and that these correlations are predicted by the optimization of energy consumption [86,99]. This indicates that humans tend to walk in ways that are most energy efficient. Similarly, [100] found that humans prefer a step width that is energy optimal. These works all agree that under normal walking conditions, humans tend to employ gaits that optimize efficiency. On the other hand, studies have found that humans walk most stably at their preferred stride frequency [101] and step width [102], while walking slower than preferred improved stability [103]. These results suggest that there is a complex relationship between walking efficiently and walking stably, where some complementary gait parameters allow for these performance attributes to be optimized simultaneously, while other parameters force a trade-off between efficiency and stability. These trade-offs also extend to bipedal robotics, where researchers discovered two gaits for a passive dynamics walking robot: one with improved efficiency and stability properties, but the other capable of travelling at higher speeds [88].

While there is relatively less work that exists studying biological quadrupeds, horses have been the subject of several gait parameter experiments. The Introduction discussed seminal results where different horse gaits were optimally efficient in different speed regimes [87]. Later experiments also found strong correlations between speed and parameters such as stride frequency and duty factor [104]. Regarding robots, early work on quadruped walkers employed gaits that maximized static stability through the trajectory [105] and found that increasing duty factor improved stability [106,107]. More recent work found that adjusting duty factor was necessary to optimize efficiency at different speeds [108]. However, the quadrupedal robots used in each of these investigations differ drastically to current state-of-the-art robots in terms of design and control methodology.

There are complex trade-offs between gait parameters and performance, as well as significant variation among different legged systems. In this work, we aim to extend concepts from prior work to modern quadrupedal robot systems with state-of-the-art planning and control algorithms.

5.3 Preliminaries

5.3.1 Hybrid Systems

Legged robots are a type of hybrid dynamical system, where domains of continuous dynamics (i.e. stance and flight phases) are linked by discontinuous hybrid events (i.e. foot touchdown and liftoff) [10]. Following [2], we describe a hybrid dynamical system as a set of modes $\{I, J, \dots, K\}$, each equipped with a domain D_I and a time-varying vector field F_I that defines the domain's dynamics. $G_{(I,J)}$ is a guard surface that triggers a transition between mode I and mode J and $R_{(I,J)}$ is the reset map defining that discontinuous transition.

An execution of a hybrid system [12] begins at an initial state $x_0 \in D_I$. Given an input $u_I(t, x)$, the system evolves according to dynamics F_I . If the system reaches guard surface $G_{(I,J)}$, the reset map $R_{(I,J)}$ is applied and the system is mapped into the new domain D_J , where it continues to evolve under dynamics F_J . The flow $\phi(t, t_0, x_0, U)$ describes the systems' evolution from initial time t_0 and state x_0 to final time t under input sequence U .

5.3.2 System Definition

The specific hybrid system we will investigate in this work is a three-dimensional quadruped robot modelled after the Ghost Robotics Spirit 40 with a motor model based on the T-Motor AK60-6 which is designed for legged robot applications. The robot is modelled as a rigid body with four legs that each have three degrees of freedom: two at the hip and one the knee. The configuration space of the robot q consists of seven body states $q_B = [x_B, y_B, z_B, \mathbf{q}_B]^T$ where $\mathbf{q}_B$ is a unit quaternion, and the twelve joint angles $q_J = [\alpha_1, \beta_1, \gamma_1, \dots, \alpha_4, \beta_4, \gamma_4]^T$ where α , β , and γ represent the ab-ad, hip, and knee joint angles of each leg. The velocity space is 18-dimensional since the angular velocities can be represented in three dimensions compared to the four-dimensional quaternion. The full state space is the concatenation of the configuration and velocity spaces, which is size 37 and the input space is size 12.

5.3.3 Convergence and Cost of Transport Analysis

In this subsection, we will introduce the quantitative measures of convergence and cost of transport, respectively. These measures will be used in Sec. 5.4.1 to analyze the performance of different gaits. The convergence measure is equivalent to (2.13) used in the previous chapter. To compute the efficiency of our quadrupedal gaits, we define the cost of transport, which is the energy consumed by the system normalized by mass and distance travelled. Here, we estimate energy consumption as the sum of Ohmic motor losses and positive work done by the motors.

More formally, the instantaneous rate of Ohmic losses, P_Ω can be defined as:

$$P_\Omega = I^2 R \quad (5.1)$$

where I is the motor current and R is the effective motor resistance estimated to be a constant $R = 3023 \, \Omega$. Since commanded motor torque is more easily available than commanded current, we can convert between these units with the motor torque coefficient K_t :

$$I = \frac{u}{K_t} \quad (5.2)$$

In our motor model, we use a value of $K_t = 0.068 \, \text{Nm/A}$ for the ab-ad and hip motors and a doubled amount for the knee motor to account for an additional 2:1 reduction of this motor in the Ghost Spirit 40 design. Therefore our equation for Ohmic loss rate is:

$$P_\Omega = \frac{u^2 R}{K_t^2} \quad (5.3)$$

Given the joint velocity space $\dot{q}_J$ and the joint torque inputs u , the total positive motor work rate P_m is:

$$P_m = \max(\dot{q}_{J,t,i} u_{t,i}, 0) \quad (5.4)$$

Integrating the Ohmic loss rate and positive motor work rate gives us the approximation of the energy consumption:

$$W = \sum_{t=1}^N \sum_{i=1}^{12} (P_{\Omega,t,i} + P_{m,t,i})h \quad (5.5)$$

where h is the length of the timestep and N is the number of timesteps.

The cost of transport is then the energy consumption normalized by distance travelled Δx , mass m , and acceleration due to gravity g :

$$\text{CoT} = \frac{W}{mg\Delta x} \quad (5.6)$$

This analysis simplifies the energy consumption of a physical robot system by ignoring factors like frictional losses. Therefore, it should be viewed as an underestimation and is not directly comparable to empirical tests of transport cost. There are alternate methods to theoretically computing cost of transport [6,109,110], though our investigation did not find significant differences in trends when using any alternate formulation.

5.4 Methods

Based on our research questions, we established the following hypotheses:

- **Hypothesis 1** *Increasing the speed will worsen convergence*
- **Hypothesis 2** *Increasing the duty factor of a gait will improve the convergence*
- **Hypothesis 3** *Narrowing stance width will worsen convergence*
- **Hypothesis 4** *Overall, the walking gait will tend to be more convergent than the trotting gait*

Our goal is to evaluate if these hypotheses hold regardless of the exact planning and control structure employed. Therefore, we conducted two investigations on the relationship between

speed, duty factor, and stance width and convergence. The first investigation employs a whole-body trajectory optimization and LQR feedback using an iLQR framework, while the second utilizes a centroidal model-predictive control (MPC) model.

5.4.1 Investigation 1: Whole-Body Trajectory Optimization with LQR Feedback

In Investigation 1, we employ a whole-body direct collocation trajectory optimization strategy with LQR feedback. This method allows us to directly leverage prior work [5] to characterize closed-loop stability performance of trajectories.

To evaluate these hypotheses, we wrote a custom direct collocation trajectory optimization method using Pinocchio [111] to compute the continuous and hybrid dynamics and the corresponding dynamics derivatives. We chose a direct collocation method because it allows us to explicitly define the contact sequence and timings to control the gait and duty factor of each trajectory. Pinocchio allows for fast computation of these dynamics and their derivatives which are used to compute the fundamental solution matrix. The cost function consists of standard quadratic costs on input torques and shaping costs on the body height and leg joint velocities:

$$\sum_{i=0}^{N-1} u_i R u_i + x_{z,i} Q_z x_{z,i} + x_{lv,i} Q_{lv} x_{lv,i} \quad (5.7)$$

where $X := \{x_0, x_1, \dots, x_N\}$ is a sequence of states with $x_i \in \mathbb{R}^n$ the system state at timestep i , $U := \{u_0, u_1, \dots, u_{N-1}\}$ is a sequence of control inputs with $u_i \in \mathbb{R}^m$ the control input at timestep i , and h is the timestep length. x_z is the element of the state vector corresponding to body height and x_{lv} the vector of elements corresponding to the leg velocities. R , Q_z , and Q_{lv} are positive-definite weighting matrices of appropriate dimension.

A hard constraint is applied to the final state to match the initial state such that the robot travels a specified velocity in the x-direction along with periodicity of the remaining states being enforced. The peak swing height of the feet are also constrained to ensure sufficient ground

clearance. Soft constraints are applied to the initial position and velocity of the robot along with the timestep length so the gait period can be chosen by the optimizer based on a given duty factor and forward speed. The mathematical description of the constraints have been omitted for brevity.

For every trajectory, an LQR controller is generated with constant weights on state and input. On the whole, we aimed for this closed-loop trajectory generation method to be similar to recent methods for whole-body control of legged robots [112,113].

5.4.2 Investigation 2: Centroidal MPC

In this second investigation, we explore a different planning and control stack, Quad-SDK [98], that utilizes a centroidal model-predictive controller (MPC) to plan ground-reaction forces (GRFs) that are in turn converted into joint torque commands. This control method differs from Investigation 1 because trajectories are planned using a simplified centroidal dynamics model and tracked with a receding horizon MPC method. Because of the centroidal dynamics model, there will be modelling errors present in deployment to both simulation and hardware, making this platform useful for investigating robustness properties. Quad-SDK has been deployed on quadrupedal robot hardware and adapted for other robust locomotion applications like walking through entanglements [114]. The hierarchical and modular structure of Quad-SDK allowed for quick modification of gait parameters and simulation environments, while taking advantage of its existing local footstep planner and MPC.

Quad-SDK is composed of three main modules, namely its global body planner, local planner, and robot driver. For the purposes of these experiments, we directly commanded body velocities, using twist inputs bypassing the global body planner entirely. Kinodynamic constraint approximations made during global planning make it difficult to find feasible plans across narrow beams, especially given that the planner uses a reduced order dynamics model that only regards the motion of the body. The local planner determines valid foothold positions using a predetermined traversability map [96], and solves the corresponding MPC problem to determine

the necessary ground reaction forces to track that reference center of mass trajectory. The robot driver then interfaces directly with the robot's motors, converting the GRFs using inverse kinematics to generate joint torques and joint angles for the robot to actuate.

A variety of experiments were run in both Gazebo Simulation and on a Ghost Robotics Spirit-40 quadruped robot to assess the impact of gait type, duty factor and velocity on a quadruped's beam walking performance.

5.5 Experiments and Results

5.5.1 Experiments of Investigation 1

In this work we performed three experiments. In Experiment 1-1, we generated walking and trotting gaits at 6 speeds (0.25, 0.5, 1, 1.5, 2, 2.5) m/s and 6 duty factors (0.75, 0.792, 0.821, 0.844, 0.861, 0.875 for the walk and 0.5, 0.583, 0.643, 0.688, 0.722, 0.75 for the trot). Note that 0.5 and 0.75 are the minimum possible duty factors of a trot and walk, respectively, without a full aerial phase. For each of the 72 trajectories, we computed the energy efficiency (positive work over distance) and convergence measure and plotted the results in Fig. 5.3.

In Experiment 1-2, we selected three gaits that exhibited good efficiency and convergence properties: the 1.5 m/s, 0.75 duty factor trot (Trot 1), 1.5 m/s, 0.75 duty factor walk (Walk 1), and 1 m/s, 0.875 duty factor walk (Walk 2). We then generated trajectories with each of these speed/duty factor combinations while constraining the gaits to narrowing stance widths of (0.34, 0.3, 0.25, 0.2, 0.15, 0.1, 0.05, 0) m. 0.34 m is the default stance width with the feet directly below the hips.

For Experiment 1-3, these same three gaits were evaluated at the same beam widths based on their capacity to reject an impulse disturbance at the beginning of the trajectory. 100 trials of each gait were given a random impulse to the center of mass of magnitude 0.01 Ns. The error between the state immediately after the impulse was applied and the reference state, δx_0 was recorded along with the tracking error at the end of the trajectory, δx_f . To evaluate the




Figure 5.3: The walking and trotting gaits transition quite smoothly between their natural duty factor regimes. Our results indicate the combination of duty factor and speed is a strong predictor of a gait’s energy efficiency, while convergence is primarily correlated with duty factor. Note there are 4 absent data points at the combination of the 2 highest speeds and 2 highest duty factors because the algorithm failed to converge to a solution.

disturbance rejection performance of each trial, we computed the error ratio, E which is the ratio of the magnitude of δx_f and δx_0 .

$$E = \frac{\|\delta x_f\|_2}{\|\delta x_0\|_2} \quad (5.8)$$

For each gait type and beam width combination, we computed the error reduction rate, which is the proportion of trials out of 100 where $E < 1$. This provides a measure of how effective each closed-loop gait was at reducing tracking error over a trajectory.

5.5.2 Results of Investigation 1

The results of Experiment 1-1 are shown in Fig. 5.3, where there is a nearly smooth transition between the two gaits at 0.75 duty factor. While the walk is generally more efficient at lower speeds and the trot at higher speeds, duty factor is a much more direct predictor of efficiency, with lower duty factors being more efficient at higher speeds and vice versa. Fig. 5.3b shows that speed does not have a strong relationship to convergence and that increasing duty factor overall improves convergence.




Figure 5.4: As stance width narrows, convergence generally worsens, while efficiency does not differ significantly. The performance of walking and trotting gaits with equivalent speed and duty factor remain similar. There is also indication of a trade-off between sacrificing efficiency to improve convergence, at least up until a stance width of about 0.05.

Experiment 1–2 found that the efficiency and convergence of the gaits with equivalent speeds and duty factors maintain their similarity as stance width narrows. Fig. 5.4 shows that a trade-off between efficiency and convergence exists up until stance width of about 0.05 m.

The results of Experiment 1–3 are shown in Fig. 5.5. Agreeing with Experiment 1–2, these results show Walk 2 was the most effective at reducing tracking error across all stance widths. For stance widths greater than 0.1 m, Walk 2 was 100% successful at rejecting disturbances, indicating this gait was robust. The other gaits at lower duty factors did not achieve 100% error reduction at any stance width.

Overall, the results from Investigation 1 indicate support for Hypotheses 2 and 3, while providing evidence against Hypotheses 1 and 4.

5.5.3 Experiments of Investigation 2

For Investigation 2, we performed two experiments, aimed at assessing the impact of gait parameters on narrow beam walking in simulation and hardware.

In Experiment 2–1, we designed a terrain map in simulation measuring 20 m in length. The map consists of nine stages, starting with a width of 0.375 meters, which decreases incrementally by 0.05 meters until it reaches a final width of 0.025 meters. The initial width reflects the nominal




Figure 5.5: At each stance width, the 3 gaits were simulated 100 times with random initial impulses. Over all stance widths, the high duty factor walk was significantly more successful in reducing error magnitudes by the end of the trajectory, succeeding 100/100 times at a 0.15 m stance width and higher. Meanwhile, the lower duty factor trot and walk never achieved a perfect success rate. This corroborates with our analysis in Fig. 5.4b.

stance width of the robot, while the width of the last stage is slightly larger than the tip of each foot. Three types of gaits were evaluated across three different speeds, totalling nine trials. The gaits were a 0.5 duty factor trot (which is the default gait of Quad-SDK), a 0.75 duty factor trot, and a 0.75 duty factor walk. The speeds tested were (0.3, 0.4, 0.5) m/s. Fig. 5.6 displays the simulation environment in this experiment. For each trial, the minimum beam width the gait was able to traverse was recorded.

Experiment 2-2 evaluated the effect of duty factor on beam walking performance for a physical quadrupedal robot. In order to fit the size of the experiment environment, a terrain map composed of three stages measuring a total 6 m in length was used. Like the map used in simulation, the initial width is 0.375 m, this time decreasing incrementally by 0.1 m per stage. In practice, this map was outlined on the ground to simulate the environment as perceived by the robot without risking the physical integrity of the robot. Fig. 5.1 shows the experimental setup. The same three gaits as the previous experiment were evaluated with just the 0.5 m/s speed. Like Experiment 2-1, we recorded the minimum beam width traversed.




Figure 5.6: Nine different gaits, varying in gait type, duty factor, and speed, were evaluated in simulation to determine the minimum beam width they could traverse.

Speed (m/s)	Gait Type	Duty Factor	Min Beam Width (m)
0.3	Trot	0.5	0.375
	Trot	0.75	0.025
	Walk	0.75	0.025
0.4	Trot	0.5	0.325
	Trot	0.75	0.025
	Walk	0.75	0.025
0.5	Trot	0.5	0.325
	Trot	0.75	0.025
	Walk	0.75	0.025

Table 5.1: Results of Experiment 2-1 show that, similar to Investigation 1, increasing the duty factor enhances beam traversal performance, while gait type has no discernible effect.

5.5.4 Results of Investigation 2

The results of these experiments aligned closely with the results from Investigation 1. In both simulation and hardware, we found that increasing the duty factor improved beam traversal performance, while different gait types with the same duty factor showed no performance differences. Table 5.1 shows the results of Experiment 2-1 and Table 5.2 displays the results of Experiment 2-2. Again, Hypotheses 2 and 3 were supported, and Hypotheses 1 and 4 were not supported.

Speed (m/s)	Gait Type	Duty Factor	Min Beam Width (m)
0.5	Trot	0.5	0.275
	Trot	0.75	0.175
	Walk	0.75	0.175

Table 5.2: Results with a physical quadrupedal robot align closely with simulation results.

5.6 Conclusion

This work presents an analysis of the impact that duty factor and speed have on gait robustness and beam walking performance. To do this, we conducted two independent investigations, one using a whole-body trajectory optimization method with LQR control and the other using a centroidal MPC method. In both investigations, we found support for the hypotheses that narrow beam widths make traversal more difficult and that increasing duty factor enhances performance in narrow environments. Conversely, speed and gait type did not significantly impact performance. These results advance the field of legged robotics by offering insights into how practitioners can adjust planning and control methods to enhance stability and robustness. Additionally, this work paves the way for future research to further explore the relationships between gait parameters and locomotion performance.

Part II

Navigating a Complex, Human World

Chapter 6

Self-Defense Against Robots: Design and Deployment Considerations for Roboticists

6.1 Introduction

Recent work in the robotics field has focused on deploying robots in public environments around humans for applications such as delivery [115], security [116], and personal assistance [117]. Delivery robotics companies such as Starship, Kiwi, and Serve have been rapidly growing the number of robots operating in public spaces like university campuses and urban sidewalks. In these environments, robots experience many close encounters with a diversity of humans. It is unavoidable that a number of these encounters will result in some kind of danger to the robot and/or human, such as a robot combusting [118], college students vandalizing a robot [119], or a robot crashing into a car [120]. Though uncommon, these edge cases underscore that humans and robots can, and will, come into direct physical conflict with each other.

While the examples above were a result of some combination of accident, negligence, and wanton violence, there are cases where a human using physical force against a robot could be justified under fear of threat to one's own well-being. Consider the following scenario, which is discussed in further detail in Section 6.5:

- A person is tired after a long day of work and slowly walks home at night. A robot, mindful of keeping a safe distance from the person, quietly treads behind them at a constant distance. The person, on alert due to walking home alone at night, becomes fearful of something stalking them from behind. Eventually, the robot decides there is enough room on the sidewalk to accelerate and pass the person on their left.

In this example, both the human and robot behave reasonably, but the context and environment in which this interaction occurs could lead the human to feel threatened and act against the robot in self-defense. Our research is focused on analyzing such situations where use of force in human-robot interaction may be justified under self-defense law. In particular, we focus on a sub-group of people we anticipate to be the most likely to feel threatened by a robot: the non-expert non-user. Compared to experts and/or users of robots, this group may not have an understanding of the capabilities that a robot possesses.

We propose categorizing self-defense scenarios as the tuple: (threat, protector, protectee), where the protector defends the protectee from some kind of perceived threat. In standard self-defense where a human protects themselves from another human, this self-defense tuple is (human, human, self) where the protector and protectee are the same person. There are also many other legally established tuples such as a human defending someone else from harm (human, human, other), a human defending their own property from someone else (human, human, property), and a human defending themselves from a non-human, which could be an animal or a non-living object (property, human, self). Note that even though not all of these tuples represent a human defending themselves, we use the term self-defense in this work due to its broad familiarity. This chapter analyzes in detail the (robot, human, self) tuple, a special case of (property, human, self), while our future work will entail characterizing other self-defense situations involving robots.

We concentrate our analysis on public ground robots that have a primary task other than social interaction with humans. These public robots contrast with industrial or care giving robots, where any person that interacts closely with them can be expected to have been trained or at

least be quite familiar with how the robot operates. Ground robots contrast with aerial robots like drones, which do not generally come into as close proximity to humans.

The objective of this chapter is to first extract elements from prior works that are key to the formation of a (robot, human, self) self-defense scenario, and secondly present four actionable recommendations for roboticists to design systems that mitigate the likelihood and severity of these self-defense situations. These recommendations span research, industry, and policy making among the robots field, and were chosen based on their potential to progress the capability of robots to operate in dense public areas. We aim to provide evidence-based guidance for roboticists across the industry so that the field can make positive, equitable impacts.

This work draws upon the fields of self-defense law, human-robot interaction, and robot path planning to present an unaddressed topic that lies at their intersection. These research areas are synthesized to establish how the current robotics state-of-the-art overlooks the possibility of non-expert, non-user humans acting in justified self-defense against a robot. The rest of the chapter is organized as follows:

- Sec. 6.2 reviews the guiding principles that inform U.S. self-defense law and argues for the importance of considering (robot, human, self) scenarios.
- Sec. 6.3 connects these principles to findings from human-robot interaction literature on human attitudes toward robots and identifies aspects of human attitudes that require additional study.
- Sec. 6.4 discusses planning strategies that have been developed for robots to operate around humans and some of their limitations on handling self-defense situations.
- Sec. 6.5 synthesizes the previous sections to present hypothetical scenarios where a human would be justified in taking self-defense action against a robot.
- Sec. 6.6 summarizes this work and discusses future work that the authors will address next to bolster the connection between robot design and self-defense law.

6.2 U.S. Self-Defense Law

In this section, we outline the key tenets of self-defense law in the United States. These laws vary appreciably by jurisdiction in the extent and environments in which self-defense is justified. For example, most U.S. states oblige no duty to retreat from a threat, while others impose a limited duty to retreat in public spaces [121]. Because of the great variance that exists in this domain, practitioners should use discretion with the specific self-defense statutes in their jurisdiction. However, we believe the principles outlined in this section are broadly applicable to the vast majority of self-defense codes in both the U.S. and other countries, and the conclusions drawn in this work are largely independent of the specific intricacies of individual statutes. The legal discourse regarding self-defense is lengthy and the analysis presented here is quite brief. For a deeper discussion on the theory of self-defense law, see [122,123].

The two primary principles that underpin American self-defense law are:

- a reasonable belief of imminent physical harm
- a proportional response to the threat

While terse, these aspects of self-defense carry centuries of legal nuance that must be interpreted by the courts on a case-by-case basis. Below we analyze some important details of these two principles.

6.2.1 A Reasonable Belief of Imminent Physical Harm

This phrase can be broken down into two parts. Firstly, a “reasonable belief” means that the protector in the (threat, protector, protectee) tuple does not need to have definitive proof that the threat of imminent harm is true, but only that it would be reasonable for a person to view a situation as threatening [124]. For instance, if the protector is acting with limited information about the threat, there can exist a reasonable belief of harm even if the protector is ultimately found to have been incorrect about the nature of the threat. Similarly, the reasonableness of the

belief is subject to prior experiences of the protector. So even if an external third party would not necessarily view the belief as reasonable, context such as prior history of physical abuse [125] or legitimate verbal threats of violence that occurred directly prior [126] can provide justification for acting in self-defense.

The second aspect of this definition is “imminent physical harm” [127]. Imminence indicates that the harm to the protectee must be actively happening or about to happen. This means that force used either proactively or after the harm has subsided is invalid to justify self-defense. There has been debate centered around the soundness of imminence as an indicator of the necessity of self-defense, but in practice, this interpretation has been upheld [127]. Additionally, the law tends to interpret the harm that is incurred as needing to be physical. While some have argued that other types of harm such as invasion of privacy may be tantamount to a self-defense justification [128], American common law has yet to grant this argument [129].

6.2.2 A Proportional Response to Threat

Once the protector reasonably believes that self-defense is warranted, they must act in a manner appropriate to the threat level [130]. For self-defense between humans, this means that relatively minor force, such as a punch or kick, can not be responded to with lethal force. However, certain situations can justify a lethal response to a less than lethal threat, such as when the protector is in their own home [131].

In a (robot, human, self) scenario, it is key to understand that the robot is property and has no intrinsic right to act in self-preservation in the way that humans do. Because of this, a self-defense act that destroys the robot can be justified even with a lesser threat to the human. This concept of proportional response emphasizes the need for engineers to carefully design robots to avoid self-defense situations, because any perception of threat could lead to justified destruction of the robot.

6.2.3 Self-Defense Against Robots

Ground robots exhibit characteristics that can uniquely give rise to self-defense justifications. While previous work has discussed self-defense against aerial drones [128], the argument justifying self-defense against these technologies was weakened by the typically large distance between the drone and any given person. Ground robots, on the other hand, are expected to come into close, immediate contact with humans during normal operation, so designing for self-defense situations is critical.

Since we have established that the proportionality criteria can be broadly satisfied in (robot, human, self), the primary challenge to determining when people may be justified to defend themselves against a robot is defining what a reasonable belief is. Because robots are still such a novel and unfamiliar technology to most people, the standard of reasonableness may be lowered to take into account the misconceptions and misunderstandings non-experts tend to carry as they interact with and react to robots.

To begin codifying self-defense law as a serious consideration for robot engineers and to assuage the public's fears about coming into contact with potentially dangerous robots, we make the following recommendation:

Recommendation 1 *Robotics companies and research organizations should publicly advertise that in situations where a human and robot are in direct physical conflict, the human's physical well-being is always valued more greatly than that of the robot, even if the result is damage or destruction of the robot.*

This recommendation reinforces the principle of proportionality in self-defense against robots, assuring people that they have immediate recourse if they feel physically threatened by a robot. Transparency from the robotics field regarding individuals' rights in interactions with robots can foster greater trust and openness toward the deployment of robots in public spaces." As discussed in Section 6.3, the attitudes that humans have towards robots has a large impact on how human-robot interactions play out.

6.3 Reasonable Perspectives Toward Robots

In this section, we take a step toward understanding standards of reasonableness in human-robot interaction. We start by analyzing how behavioral norms can dictate (human, human, self) self-defense scenarios. We next contrast the role of behavioral norms in human-human and human-robot interactions and present results from prior literature on human attitudes toward robots, which contribute to an understanding of reasonable behavior. There lacks consistently interpretable robot behaviors that could be used to establish human-robot norms, and instead suggest considering diverse attitudes toward robots to inform what constitutes reasonable behavior.

6.3.1 Behavioral Norms in Human Interactions

In human-human interactions, behavioral norms between humans play a crucial role in determining reasonable behavior. When evaluating (human, human, self) cases, courts judge the reasonableness of a person's behavior based on broad, implicit understandings of how humans typically behave. For example, in the case of *Rowe v. United States*, it was found that even though Rowe kicked a man, which prompted the man to attack him with a knife, Rowe stepping back after the kick revived his right to self-defense [132].

Following the literature [133], we define a norm as a widely adhered to and understood action that helps coordinate behavior. This stepping back is a norm that was found to be a widely understood indication that Rowe was no longer a threat and withdrawing from the fight, which protected him from being retaliated against. Such implicit understandings between humans can be considered a basis of self-defense law because they create standards of reasonableness, allow for generalizations across cases, and ultimately promote fairness in decision-making.

Norms for human-robot interaction, on the other hand, are not well established in the sense that actions taken by the robot are not widely understood and do not facilitate coordinated behavior between the human and robot, which can make robot behavior unpredictable to a non-

expert. Even though human behavior is not completely predictable, years of lived experience interacting with other humans allows for a deeply developed understanding of typical human behavior. Due to the novelty of robots, this understanding between humans and robots is inadequate. Because of the lack of established norms for robots, the standard for justified self-defense in (robot, human, self) should be lower than in (human, human, self) cases.

Frameworks such as COMPANION have attempted to address this gap by encoding human social norms (such as maintaining personal space and moving to the right to avoid colliding with people approaching from the opposite direction) within robot behavior [134]. However, more research is needed to determine whether humans actually expect robots to behave with the same norms as humans. In fact, some studies have indicated otherwise. For example, [135] found that robots were considered more trustworthy when approaching a person quickly (possibly since faster robots were more noticeable than slower robots), whereas humans were more trustworthy when approaching slowly. The study also found that humans performed more corrective reactions (such as stepping back or adjusting eye contact) when a robot invaded their personal space compared to a human. This indicates that humans react differently to violations depending on if the offending party is a human or robot, and simply having robots adopt human norms does not guarantee self-defense situations will be avoided.

Even if robot behavioral norms become standardized among the industry, there is no guarantee that humans will, in their split-second decision making, have enough trust to assume that a robot can reliably follow certain norms. It is important for robots to conform to the preferences and expectations that humans have for their behavior, and we recommend further study of these topics:

Recommendation 2 *Because of the differences in human preferences and expectations when interacting with robots compared to other humans, researchers should explore whether there are robot behaviors that humans react consistently to and if these behaviors can be encoded into a standardized framework of human-robot norms.*

Results that establish norms for even a subset of ground robots (such as wheeled robots,

humanoids, or quadrupeds) could begin to establish a more refined definition of reasonable behavior around robots. Possibilities of behavioral norms for robots could be exhibiting body language or digital facial expressions [136], which could in turn improve the legibility of robots in public environments (as discussed in Sec. 6.4.2). Establishing norms may be difficult due to the variance in attitudes that humans exhibit toward robots. Even if consistent norms cannot be identified for many aspects of robot behavior, it is still important to characterize how attitudes vary among humans and under what circumstances human attitudes can be well modeled.

6.3.2 Human Attitudes Toward Robots

Human attitudes towards robots tend to vary based on several factors, including a person's familiarity with robots and how well a robot's behavior aligns with the human's expectations and preferences. Studies suggest that the more familiar a person is with a robot and the more their expectations align with the robot's behaviors, the more positive their attitude towards that robot [137]. Conversely, when there are gaps and discrepancies in these areas, attitudes tend to shift negatively. The real consequences of negative attitudes toward ground robots and a violation of expectations during their deployment emerge as justified self-defense scenarios. Alleviating the public's negative attitudes and aligning robot design with expectations is essential for safe human-robot interactions. This reinforces the necessity of Recommendation 1, which can help inform the public's expectations of how they can interact with robots.

One highly-documented example of this is the largely negative attitudes communities have expressed toward the deployment of robots by police departments in several U.S. cities [138]. One issue that arose was that communities affected by the police's usage of ground robots were not involved in the development process and expressed frustration over the expensive and possibly dangerous technology [138,139], calling the robot "another danger for Black & Latino residents" [139] and expressing fear toward the futuristic appearance of the robots [140]. Requests to police departments for more information on the purpose of the robots were not always met [138,141].

When impacted communities are not involved in the development of robots, the deployed products can be misaligned with the community's expectations of how these new tools should be used. Research indicates that the level of familiarity people have with robots and the preconceptions they hold influence their attitudes, such as how the fear of sentient robots correlates with negative attitudes [137]. Another study suggests that people may be more likely to support robots doing jobs that require less experience and communication [142]. Therefore, attitudes toward delivery robots, which satisfy both of these conditions, could be more positive than those toward police robots, though further research should test this theory.

Attitudes of marginalized groups toward robots are especially important for developers to consider, since they have been disproportionately affected by harmful uses of novel technologies [143]. Police robots have often been deployed to patrol low-income, Black neighborhoods [138,144], while women have repeatedly been targets of unwelcome surveillance by drones [145,146]. Additionally, [147] found that women tended to be less receptive to the concept of patrolling police robots than men. To combat this inequity, [138] suggests involving marginalized community members in the technology design process. Factoring in the preferences of the stakeholders who interact most closely with robots will help developers align robot design with expectations, reduce negative attitudes toward robots, and promote equity by working for marginalized communities instead of against them. We recommend further investigation into how attitudinal differences manifest among disadvantaged groups:

Recommendation 3 *Due to the variance in human attitudes toward robots and the disparate effects technologies have had, researchers should examine and catalog the attitudinal differences among different groups of people, especially from those that have historically been marginalized.*

Reasonable human behavior varies greatly due to differences in background and past experiences. These differences can be measured in the attitudes, perspectives, and reactions people exhibit toward robots. Ultimately, this variability suggests that even severe human behaviors towards robots can be justified and considered reasonable, at least until the establishment of robot

behavioral norms that are broadly understood by people of many backgrounds. Instead of basing robot behavior on unestablished norms that people must adhere to, it is essential to consider people's diverse attitudes and expectations regarding robots and design robots with this context in mind.

6.4 Human-Aware Planning

While self-defense has so far gone unconsidered in the design and implementation of robots, there has been ample related work in planning robot motions in human environments. Generally, robot path planning is performed by sampling many possible paths a robot can take and selecting the most optimal choice, often based on the shortest path [148,149]. Algorithms are also able to obey specified constraints such as avoiding obstacles. Recent research has adapted these path planning algorithms to predict and react to human obstacles, and to minimize risk of collision with people [150]. Other work has developed robotic behaviors to satisfy desired outcomes such as visibility [151], active communication [152], and following social norms [134,153]. In this section, we examine two primary research thrusts in human-aware planning that have seen significant attention: explainability and legibility. We analyze not only what work has been done, but also the reasons stated in the literature for why these aspects of human-aware planning are important. While aspects of explainability and legibility are useful in mitigating the potential for self-defense situations, current implementations lack the capability to address all environments in which a self-defense scenario may arise.

6.4.1 Explainability

Drawing from [154], we define explainability as the ability of an autonomous agent to produce records of the decisions it has made and understandable reasoning for why those decisions were made. This definition is compatible with how explainability is discussed in prior works, such as generating contrastive explanations (i.e. why A and not B?) [155] and explanations that satisfy

user-defined preferences [156]. Post-hoc explanations are designed to be generated after a robot decision has been executed and in response to some kind of questioning, while some work has examined generating concurrent explanations for behaviors as they are happening [157].

Drawing from the existing discourse on explainability in the field of autonomous vehicles (AVs), [158] discusses post-hoc and concurrent explanations for AVs by analyzing a scenario in which an AV fails to recognize a pedestrian crossing in front of it. Post-hoc explanations to characterize why the AV failed could be useful in a post-accident investigation and for regulators to hold manufacturers accountable. But these post-hoc explanations would not be able to prevent accidents from occurring. Concurrent explanations, such as communication to a passenger that the car will continue through a crosswalk because no pedestrian has been detected, could allow passengers to take emergency actions when they recognize the vehicle has made an error. In this instance, a passenger could activate an emergency brake that stops the car before it collides with the pedestrian.

In the context of ground robots, post-hoc explanations assume that the people who desire explanations for robot behaviors have access to the robot afterward. These explanations could be useful to operators who could recognize errors in their usage of the robot, developers who could better understand errors and implement fixes, and members of the judiciary who could use explanations to assign liability after an accident. However, post-hoc explanations generally exclude members of the general public who interact with the robot for just a fleeting moment, such as passing each other on the sidewalk. Considering that the majority of people interacting with a robot in a public environment will likely not have access to that robot afterward, this exclusion is significant. Concurrent explanations, on the other hand, are able to actively communicate to people in the robot's immediate surroundings. However, concurrent explanations may be difficult to convey to certain people in real-world environments. Explanations announced verbally may not be heard by people on the phone or listening to music, or may be drowned out in loud environments such as construction. Similarly, explanations presented visually may not be suitable for people with visual impairments or in night-time environments. A robot must also

consider that some people may not speak the robot’s default language.

In a survey of 62 papers on explainability, [159] found that the most commonly stated motivation for the work was transparency (i.e allowing people to better understand the inner workings of the robot), followed by trust and collaboration. These motivations go hand in hand, as increased transparency would naturally lead people to trust being around the robot and working with it. Of these surveyed papers, many framed trust around the relationship between robots and their operators or teammates, and the faith these people had that their robots would work reliably [160,161]. With the European Union’s recent General Data Protection Regulation (GDPR) outlining a person’s legal “right to explanation” when encountering autonomous agents [162], explainability has come to even greater relevance.

In this work, we question if explanations are the key to inspiring public trust in robots and allowing for a seamless deployment of robots in public spaces. Explanations can be valuable, but are ineffective in fully mitigating the possibility of self-defense scenarios due to the difficulties of access and communication with typical bystanders. What robots need is the capability to generate implicit methods of communication that can foster an improved understanding of robot behavior. The human-aware planning concept of legibility may be a more suitable method to accomplish this.

6.4.2 Legibility

While explainability focuses on the producing reasons for why a robot behaved a certain way, legibility characterizes how a robot communicates what it is doing or intends to do. Based on the work from [163], we define legibility as the ability of a human to understand a robot’s intentions based on observation. For instance, cars have turn signals to indicate to others what action they are about to perform (e.g. turning right). The turn signal does not explain why the car is turning right, but allows others to understand what it is about to do and react accordingly. Explanations can help robots become more legible, but there are many other factors that can improve legibility such as providing cues [164], mimicking human behavior [165], moving quickly toward

goals [166], and staying in people’s field of view [151].

Legibility is a somewhat vague concept that is difficult to define and measure experimentally. To do this, some authors have evaluated legibility by asking people to predict a robot’s future behavior based on past observed behavior [167,168]. Others evaluated human performance of an unrelated task while a robot navigated around them [169]. Others still used questionnaires to gauge how well subjects felt they understood a robot’s intentions [164], while [166] proposed a numerical measure of a trajectory’s legibility.

Studies have found that legibility is correlated with increased feelings of safety, comfort, and acceptance [169–171]. These goals align very closely with the stated purposes of explainability, but may differ in the groups of people these methods are designed for. While explainability is often framed around expert users, significant work has focused on robot legibility to non-experts, as the lesser amount of information needed may make it easier for non-experts to comprehend. However, as with generating explanations, interacting with a diverse group of humans may make producing legible behaviors much more difficult. In these situations, it may not be possible to stay in everybody’s field of view or to expect all people to notice the robot’s gestures. The complicating factors discussed for explainability can make communication difficult here as well, as visual or verbal signals can breakdown in certain cases. Engineers must also consider how obvious any given message is to non-expert non-users. This relates closely to the concept of human-robot norms, discussed in Sec. 6.3. Norms are not yet established for robots, so developing robots that can be legible to people across diverse backgrounds is an enormous open problem.

6.4.3 Human-Aware Planning for Non-Expert, Non-Users

Based on this section, we conclude that current human-aware planning algorithms have not sufficiently addressed how robots should operate in dense human environments where communication is impeded. In particular, non-expert, non-users have received little attention as to how robots should interact with them. For instance, the IEEE Standard for Transparency of Autonomous Systems lays out guidelines for what information robots should be able to commu-

nicate to people [172]. The transparency standard for users is grounded primarily in providing explanations, but the transparency standard for non-users among the general public is focused exclusively on data privacy and lacks considerations for how robots must communicate with non-users. Other standards associations like ISO and ANSI address robots in industrial, service, or personal care environments, but also do not acknowledge non-expert non-users in public settings. We recommend these standards be revised to acknowledge how non-expert, non-users interact with robots:

Recommendation 4 *Standards, guidelines, and regulations from influential organizations such as IEEE should detail how robots should interact with non-expert, non-users to cultivate perceived safety and trust amongst the public. By drawing from concepts such as legibility, these standards can provide clear direction for how robots should be developed to minimize self-defense occurrences.*

While current standards for transparent and legible robots are unsatisfactory, the literature provides guidance to suggest possible standardization of certain aspects of robot behavior, such as establishing a standard mapping of light color to indicate behavior [128,145] and requiring the deployment of noticeable robots instead of silent, stealthy ones [135]. Another regulation that could be considered is how robots should operate at night. The Federal Aviation Administration currently restricts flying drones at night [128,173] and encouraging similar regulations for ground robots could be productive. However, unlike drones, there is not one government agency that can dictate regulations on ground robots. State and local governments control the use of their roads and sidewalks, so adoption of a consistent set of regulations is unlikely. Organizations like IEEE and leading companies like Boston Dynamics and Agility Robotics may be the vanguards of establishing industry standards for operating robots. Regardless, there is still a long way to go in establishing standards and regulations that fully address self-defense against robots.

6.5 Justified Human Self-Defense Against Robots

Even when explainability and legibility are incorporated into robot planning, there is still potential for self-defense situations to arise between robots and non-expert, non-user humans. Practitioners must understand the dependence a human-robot interaction has on the exact person and environment the interaction takes place in. In the Introduction, we suggested a hypothetical case of a lone person walking home at night, who is followed from behind and then passed by a robot. Consider two additional hypothetical cases where a robot that behaves according to conventional human-aware planning principles makes a human feel uncomfortable and even threatened:

- A robot equipped with gaze tracking technology attempts to stay as close as possible to the center of a person's field of view to maximize legibility as they wait for a bus. This person, however, is attempting to view the numbers of the buses that are passing by, and is unsettled that the robot appears to be blocking them from finding the right bus to leave on. As the robot approaches, they become afraid of the robot's single-minded focus on them.
- Navigating around a blind corner, a legged robot unexpectedly bumps into a person turning the corner in the opposing direction. In this situation, the robot is unable to satisfy the personal space constraint it has been programmed with and reverts to a "safe mode", which is to sit down on the ground. Already flustered by the sudden encounter with the robot, the person finds this behavior particularly unexpected and feels unsafe due to this unpredictability.

In each of these cases, the robots demonstrate some aspects of current human-aware planning methods, which in many circumstances may be appropriate and increase the transparency, trust, and perceived safety that nearby people feel. However, these cases highlight ways that naive implementation of these methods can cause unintended negative effects. The robot in the first case from the introduction takes care to maintain a safe distance and pass according to typical social norms, but fails to account for the context and environment that causes the person to

be fearful of any nearby entity approaching quietly from behind. In the second case, the robot attempts to maximize its visibility, but without the understanding that the human would prefer to not have the robot so central in their field of vision and feels uncomfortable with the intense attention the robot is paying them. Finally, the third case highlights a robot's attempt to embody a norm that indicates a non-threatening disposition. However, this norm is not obvious enough to a person that must make a split-second decision on whether the robot could harm them.

Even though the robots in these examples may not pose an actual threat to the humans, the behaviors of these robots coupled with the people's backgrounds and the unique environment they are in can lead to a perception of threat. This perceived threat could manifest into the humans acting in self-defense against the robots once they are sufficiently close to each other. Self-defense in these (robot-human-self) cases would be justified because there exists a reasonable belief of imminent physical harm. These scenarios could result in damage or destruction of the robot, a potentially appropriate proportional response to the threat.

6.6 Discussion and Conclusion

As roboticists work to rapidly ramp up deployment of their robots in public environments, it is crucial to understand the genuine physical harm these robots could cause human bystanders. Developers must design robots not only to guarantee human safety, but also to maximize the perceived safety of nearby humans. However, as robots are still largely unfamiliar to most of the general population and are often viewed with negative preconceptions, it is likely that some humans will see robots as threats to their physical safety and act in self-defense. In this work, we discuss how self-defense law applies to human encounters with ground robots, the human norms and attitudes that dictate the outcome of human-robot interactions, and the need to expand explainability and legibility to address self-defense cases. Synthesizing these three concepts, we identify scenarios where human self-defense against robots could be justified, even under reasonable robot behavior.

These considerations inform four recommendations to roboticists that aim to reduce the likelihood of justified human self-defense against robots. Recommendation 1 addresses robotics companies and research institutions to provide open communication to the public on the rights that they have when interacting with robots they perceive as dangerous, which will promote public trust and improve human attitudes toward robots. Recommendation 2 suggests researchers examine if any implicit robot behaviors are widely interpretable to humans and if a framework of human-robot norms can begin to be established. Recommendation 3 calls for a more detailed exploration of the attitudes that marginalized groups such as Black communities and women hold toward robot deployment. This will work toward considering previously-excluded people in the development of novel technologies and reinforcing the rights of these marginalized populations. Finally, Recommendation 4 advocates for an overhaul in robot standards, guidelines, and regulations to address legible robot behavior to non-expert, non-users.

We argue that contextualizing robot navigation in self-defense law establishes tangible, relevant outcomes that developers can use to evaluate their algorithms on. We hope that this work will contribute to keeping people of all backgrounds safe and secure as robots are increasingly deployed around them.

Based on Recommendations 2 and 3, one line work we are exploring in more detail is the reasonableness criteria for self-defense against robots. Specifically, we propose experimentally establishing aspects of robot locomotion that people would be more or less likely to perceive as threatening.

Chapter 7

Improving Equity in Robot Deployment: A Study of Food Pantry Patrons

7.1 Introduction

Recent developments in robotics have enabled impressive capabilities in locomotion [174], manipulation [175], and socially-aware navigation [176]. As a result of this progress, it seems that large-scale robot deployment is only a short time away, with applications of manufacturing [177], warehouse management [178], and healthcare [179] receiving significant attention. One robotic technology that has already been deployed for several years is the autonomous delivery robot (ADR). Companies have been deploying ADRs in cities across the world, focusing on delivering food and groceries directly to customers. These companies promote their technologies as improving the convenience, sustainability, and affordability of food shopping [180–182]. Additionally, cities like Boston, USA, have recently launched initiatives in collaboration with private partners to expand the deployment of ADRs [183].

Despite demonstrated interest and commercial viability, there are social, technical, and regulatory challenges surrounding ADRs that are still not resolved. For instance, researchers found that about 40% of YouTube comments relating to ADRs were negative [184]. Another study dis-

covered that public acceptance of ADRs is stratified among different groups [185] and publicized cases of harms caused by AI and robotics [138,143,145] may amplify certain groups' anxiety toward ADRs. Furthermore, ADRs are more reliable traversing wide, well-maintained sidewalks, which are more commonly found in affluent neighborhoods [186] and might exclude underserved communities that lack such resources. Meanwhile, cities and states have started enacting regulations for ADRs, but their approaches have varied significantly. The state of Pennsylvania, USA has a generous policy that limits ADRs to 550 lb and enforces a speed limit of 12 mph [187]. On the other end of the spectrum is the city of Toronto, which in 2021 voted to ban ADRs from operating due to safety concerns for people with disabilities [188]. Social, structural, technical, and policy factors all influence whether people from diverse backgrounds can benefit from ADRs, and there is still no consensus on the strategies needed to achieve this goal.

A relevant population for ADRs that has so far been overlooked is food pantry patrons. Boston has a significant population facing food insecurity, primarily concentrated in its most diverse neighborhoods of Dorchester, Mattapan, and East Boston. In these neighborhoods, over a quarter of residents face food insecurity [189], a rate twice the national average [190]. Food pantries can be a vital resource for these individuals, offering free distributions of food funded partially by the city. However, data from the City of Boston reveals that there are no food pantries in Dorchester or Mattapan, forcing residents to travel considerable distances to access this essential service [189]. ADRs could potentially address this issue, but it is unknown whether current ADR services align with the needs of food pantry patrons. Given Boston's public investment in both ADRs and food pantries, we explore whether these programs can be integrated to deploy ADRs for delivering free food to those in need. This concept is already being tested by the city of Arlington, USA, which has launched a pilot program utilizing robots and drones to make food pantry deliveries [191]. However, the program has not yet publicized any findings on the specific needs of the food pantry community.

To evaluate the potential for ADRs to serve food pantry patrons, we approach the following research question: **How can ADR services be designed and deployed to benefit the food**

pantry community? This question is pertinent because while there is a body of work chronicling food insecurity in the United States, little attention has been given to how technological solutions could address this issue. Conversely, while significant effort has been invested in developing ADR technologies, there has been minimal investigation into their impact on underserved communities like food pantry patrons. Bridging these gaps and addressing this question directly can improve equity in robot deployment, meaning that people from all backgrounds can access and benefit from ADR technologies, regardless of their background.

To this end, we conducted semi-structured interviews with 21 food pantry patrons in the Boston area to gain insight into their food shopping experiences, related technology adoption practices, and perspectives on ADRs. Our data suggests that despite challenges, robot delivery would empower patrons to make food shopping decisions according to their evolving personal needs and preferences. This work offers the following novel contributions:

- an extension of prior literature investigating habits of those experiencing food insecurity by explicitly connecting to a potential technical solution
- identification of patron needs that current ADR services are unable to address and recommendations to tailor ADRs to these needs
- a guide for developers to study how technologies can be deployed equitably, ensuring that benefits are accessible to underrepresented populations.

These contributions will help decision-makers synthesize insights from social science and engineering to deploy services that benefit marginalized communities.

The rest of the paper is organized as follows. Section 7.2 provides an overview of food insecurity in the United States and the role of food pantries in addressing these needs, reviews the development of robots and other digital tools to support food pantry operations, and discusses research on the social, economic, and policy factors that influence ADR deployment. Section 7.3 details the design, execution, and analysis procedure of the qualitative study. Section 7.4 presents four themes that reveal nuanced insights into the experiences and perspectives of food

pantry patrons. Section 7.5 connects these themes to the research question and provides takeaways for decision-makers. Section 7.6 discusses limitations of the study, Section 7.7 introduces extensions to this work that would provide further insight to the research question, and Section 7.8 concludes the work.

7.2 Background and Related Works

This section consists of four parts, summarizing key research on food insecurity, the role of food pantries, the current state of ADR deployment, and the development of digital technologies for food pantries. It underscores two gaps in the existing literature: the limitations of food pantries in fully addressing food insecurity and the lack of focus on how technologies like ADRs could meet the needs of pantry patrons.

7.2.1 Food Insecurity and Shopping Patterns

In 2022, 17 million U.S. households faced food insecurity, unable to provide enough food to meet their needs [190]. A large body of research has examined the food shopping patterns of low-income and food-insecure individuals. Studies reveal that factors such as price, accessibility, food quality, and selection significantly influence their shopping habits. [192,193]. These factors manifest in nuanced ways, such as shoppers evaluating food prices not only based on the sticker price but also considering potential food waste and how quickly the food will be consumed [194]. Additional research reveals that to navigate tight food budgets, low-income shoppers often display resourcefulness by shopping at multiple stores and seeking out sales [195–197]. This means they do not always shop at stores that are closest to them, and may travel farther for better prices rather than shopping at nearby stores [198,199]. However, this is not feasible for many shoppers who do not have a car and live in less affluent areas where public transportation is unreliable [200]. These findings show that dealing with food insecurity is a complex challenge for many households, and that potential solutions must acknowledge the factors that shoppers

consider in order to help them make decisions suited to their needs.

When considering potential solutions, merely increasing the number of grocery stores may not effectively address food insecurity, as shoppers often travel farther for better prices. This has led to the appearance of food mirages– areas with a seemingly adequate number of food providers that remain inaccessible to those facing strict constraints [199]. Researchers have proposed solutions such as expanding nutrition assistance programs like SNAP [201], promoting urban agriculture [202], and enhancing access to affordable food through local food cooperatives and delivery [203]. The last suggestion is the focus of this study, as we investigate how ADR technologies could improve food access while reducing the need for extensive travel.

7.2.2 Food Pantries

Given that up to 89% of food pantry¹ patrons experience food insecurity, food pantries are a crucial resource for many low-income households [205]. Research has shown that they positively impact food security, though they are often constrained by limited resources [206]. Food pantries tend to be independently run organizations connected to food banks and other organizations through partnerships like Feeding America [207]. Many are affiliated with schools, churches, or other community organizations. The frequency at which patrons can receive food varies depending on supply and demand, usually ranging from once a month to once a week. Some pantries also enforce residency requirements, such as neighborhood or zip code restrictions. In summary, food pantries are dynamic and diverse establishments that play a crucial role in their communities.

Research on food pantries and their patrons has involved soliciting direct feedback to enhance the food pantry experience. These studies have identified specific needs of patrons, including a desire for increased food quantity, improved food quality, healthier food options, and the ability to select their own food. [208,209]. Other studies have examined the characteristics

¹A food pantry is a location that distributes food at no cost. In contrast, a food bank is a warehouse that stores large amounts of food that is delivered to food pantries, but usually does not directly interact with patrons [204].

of food pantry patrons, revealing that long-term users often face prolonged unemployment and are likely to depend on government programs such as food stamps (now known as SNAP) [210]. A study conducted in the Pittsburgh, USA area found that ownership of a car was highly correlated with food pantry patronage, indicating patrons may have similar difficulties accessing food pantries as they do with traditional grocery stores [211]. Proposals to improve food pantry services include formalizing food procurement and distribution operations and partnering with health organizations to promote healthy eating [206,212]. This study seeks to enhance the existing literature by exploring the potential of ADRs to address identified needs like increased food quantity and reduced transportation, as well as the challenges that must be overcome to make this solution viable.

7.2.3 Autonomous Delivery Robot Deployment

ADRs aim to solve the last-mile delivery problem, where traditional delivery vans are becoming increasingly costly as demand is increasing, workforce numbers are declining, and sustainability is becoming a more critical component of city planning [213]. ADRs attempt to address each of those challenges, promising financial savings for companies that deploy them [214]. In turn, companies assure customers will reap benefits of cheaper delivery fees and quicker deliveries [180,182]. Research on ADRs has focused on deployment considerations like financial, regulatory, and social factors [215,216].

As ADRs are being deployed in progressively greater numbers, regulation has become a priority for many municipalities. As referenced in the Introduction, states in the USA have begun limiting the size and speed of robots that can travel on public sidewalks [187,217]. Additionally, there is still uncertainty with how ADRs interact with existing laws and regulations such as data privacy, tort liability, and self-defense law [8,214]. These questions have so far gone unanswered in the legal realm and remain an obstacle to widespread ADR adoption. Although this work does not address specific laws and regulations, it documents the experiences of a marginalized community to better equip policymakers in protecting vulnerable populations. In particular, Sec. 7.5 dis-

cusses policy considerations for engineers and governments as they work toward public robotics initiatives.

Recently, increased attention has been paid to public reception of ADRs. As noted earlier, research has shown public support of ADR deployment is mixed [184,185]. However, much of this research has been done without participants actually interacting with a robot. Other studies examining real world human-delivery robot interactions have found mostly positive reactions to ADRs after observing cases of pedestrians greeting or assisting robots [218–220]. This tension suggests that interacting with a robot could alter how a person perceives ADRs, underscoring the importance of ensuring broad ADR access to prevent further stratification between privileged and underprivileged communities. Considering these studies drew from a convenience sample and did not consider underrepresented populations, this risk has so far been overlooked. We aim to address this gap by conducting a qualitative study with a diverse sample of potential ADR users who have not been involved in ADR development. This study did not involve direct interaction with robots, but future research could incorporate this element.

7.2.4 Digital Technologies and Robots for Food Pantries

In recent years, there has been some movement to incorporate digital technologies into food pantry services. Research has found that food pantry staff desire digital tools for staff and volunteer scheduling, inventory management, communicating with volunteers and staff, and more [221]. These technologies can facilitate positive outcomes like encouraging patrons to consume more fresh vegetables [222]. On the other hand, in a review of 39 digital technologies that are currently being developed for food pantries, researchers found that most are lacking in user-centered design and have had limited real-world impact so far, indicating there is still much work to be done [223]. This work contributes to filling this gap by centering food pantry users in technology design.

Robots have been utilized in a few ways to support food pantries and food banks. One example is the “Picking with Purpose” program by warehouse automation company Berkshire Grey,

where their pick-and-place robots packed meals for food banks to distribute [224]. Less directly, robots have also been deployed to disinfect food bank warehouses [225]. Notably, these applications lacked any interaction with patrons.

In our review, we identified only two instances where ADRs have been used in conjunction with food pantry services. One example is at Arizona State University, USA, where ADRs delivered food from a student-led food pantry to students on its Tempe campus [226]. The other case is a pilot program in Arlington, USA, testing ADRs and drones for food bank deliveries [191]. Although the program is still in its early stages, the city has expressed plans to conduct community outreach to better serve patrons. The use of ADRs for food pantries remains an underexplored possibility, and this study examines key considerations for practitioners to facilitate successful ADR deployments. Namely, the needs and preferences of food pantry patrons likely differ from those of typical ADR users. Adapting ADRs to better serve the food pantry community can promote greater equity in deployment by expanding access to this technology.

7.3 Methods

While quantitative analysis is critical to test theories, qualitative analysis is equally necessary when building theories [227]. For a research problem that is poorly understood, like ADRs for food pantries, qualitative studies systematically reveal the holistic and emergent data that can then be expanded on in future studies. In this section, we describe our methodology for this qualitative study, including participant recruitment, the authors' positionality regarding this study, the semi-structured interview procedure, and the data analysis process.

7.3.1 Participants

In order to recruit participants for the study, we connected with two food pantries: the Margaret Fuller Neighborhood House [228] and the East End House [229]. While initial communication with the pantries in part consisted of email conversations, the first author also visited both pantries

in person to volunteer, meet the staff, and assist patrons. These visits allowed the researchers to tailor the focus of the study based on this initial information finding. We then held conversations with the directors of these two food pantries to gauge how this study could be conducted to maximize interest among patrons and minimize disturbance to food pantry operations.

Both partner food pantries have patrons representing a myriad of backgrounds, cultures, and demographics. In particular, there are many different languages spoken at the food pantry beyond English, including Cantonese, Haitian Creole, Spanish, Portuguese, Amharic, and more. Because Cantonese and Haitian Creole were the two most commonly spoken non-English languages, we hired interpreters to conduct interviews in these languages. Ultimately, we recruited 21 participants and conducted 8 interviews in English and 13 in Cantonese. We were not able to conduct any interviews in Haitian Creole as we had planned, which we discuss further in the following Sec. 7.3.2 and Sec. 7.6. In an attempt to preemptively address participant concerns around privacy and use of data, we did not collect any demographic or otherwise identifying data of the participants, including home location, age, race, gender, income, etc.

7.3.2 Positionality Statement

None of the authors have been patrons of food pantries and approach the data with a level of privilege that influenced the project. The first author, who conducted all of the recruitment and interviews, is a mid-20 year old man of Chinese descent, and this may have had some impact on who was willing to participate in the study. It appeared that Cantonese speakers (which presumably were of Chinese descent, though this data was not collected) were significantly more eager to participate in the study compared to the Haitian Creole speaking population that we also desired to recruit. Because of this, the majority of our participants were Cantonese speakers and we did not conduct any interviews in Haitian Creole.

7.3.3 Interview Design and Procedure

We conducted a 30–60 minute long semi-structured interview with each participant that was live interpreted if requested. English interviews lasted 35 minutes on average, while interpreted interviews lasted close to an hour on average due to the additional time needed for live interpretation. Interviews were held at a private room in one of the partner pantries and at public spaces like food halls when requested by participants. At the conclusion of the interview, participants received a grocery store gift card. All participants were completely anonymized and the study was approved by an external institutional review board (IRB).

While the semi-structured nature of the interviews meant that questions were not all predetermined or following a fixed order, the interviews generally consisted of three portions:

1. In the first portion of the interview, participants discussed their grocery shopping and food pantry experience. This consisted of information on where and how often participants shopped for groceries and their method of transportation to the grocery store. On the food pantry side, participants also shared their thoughts on the quality, convenience, and availability of foods along with how they felt the food pantry experience could be improved. Example questions asked in this portion are:

- “What does your average week of grocery shopping look like?”
- “How did you first hear about the food pantry?”
- “What aspects of your food shopping experience would you change?”

2. Next, participants discussed any familiarity they had with grocery delivery services such as Instacart [230]. If the participant had utilized this type of service, they were asked to share why they chose to use the service and their perspective on the experience. If they had not used this service before, they were asked to discuss what changes might lead them to considering adopting the service. Questions in this portion include:

- “Are you familiar with existing food delivery services?”

- “What factors are preventing you from using these services?”
- “How would delivery services change your food shopping habits?”

3. Finally, we broached the subject of ADRs. For participants that were not familiar with these robot technologies, we showed images of different ADRs that have been deployed such as Starship [180] and Serve [182]. Then, with this foundational understanding developed, the participants discussed their initial thoughts on using these robotic delivery services and the factors, both positive and negative, that differentiated the ADRs from human couriers. Types of questions in this portion are:

- “Do you believe a robot could successfully deliver groceries to your house?”
- “If you encountered a robot on the sidewalk, how would you react?”
- “Would you prefer having food delivered by a human or a robot?”

Each portion of the interview played a key role in answering our research question. The first portion allowed us to contextualize the needs, preferences, and practices of this sample in relation to prior studies on similar communities, laying the groundwork for discussions on how delivery services and robots could address these needs. The second portion was essential for evaluating participants’ current relationships with technology, which helped us assess the feasibility of ADRs for this group. For example, if participants were opposed to using digital tools, an ADR solution requiring a mobile app would likely be impractical. However, we found that while participants were not comfortable with navigating apps, they were willing to learn from friends or family members, suggesting potential for ADR adoption. The final portion built on the previous ones by providing insights into how robots might uniquely influence participant behaviors and decisions. This portion addressed the core objective of the study: identifying the design and deployment factors that would make ADRs valuable to patrons.

7.3.4 Data Analysis

Each of the 21 interviews were audio recorded and then transcribed into English. For the Cantonese interviews, the spoken Cantonese was transcribed and then translated to English, giving us two versions of translation: one done by the live interpreter and one by the post-hoc translator. While this means that the translation was not strictly necessary, this was done to ensure the precise meaning of the participants' words was captured. As will be discussed in Sec. 7.6, there sometimes were minor discrepancies between the live-interpreted English and translated English. In cases of discrepancy, we deferred to the translated English for analysis due to the time constraints that make live interpretation difficult.

To analyze the transcripts, we followed the method of thematic analysis which "is a method for systematically identifying, organizing, and offering insight into patterns of meaning (themes) across a data set" [231]. Thematic analysis ensures that takeaways from this study are backed by evidence from multiple participants. We also employed an inductive method of analysis, where codes and themes were not prepared ahead of time, but rather converged upon after several rounds of transcript review [232]. The purpose of this inductive method is to best allow the data to speak for itself, reducing the effect of researchers' assumptions and biases into the analysis. Considering that this population has never been interviewed before about this topic, we found this approach most appropriate.

To conduct the analysis, two formal rounds of data review were conducted. First, an initial round of inductive, semantic coding was conducted. The first co-author coded all the transcripts, while the other five co-authors divided the transcripts among themselves, ensuring each was reviewed by two coders. Each coder was instructed to distill the core, "semantic" meaning of each participant's words into a set of codes.

Next, the coders gathered to critically review the codes, focusing particularly on areas where discrepancies arose between coders. Since not every coder reviewed the data in its entirety, this meeting was also an opportunity to analyze whether the patterns that each coder observed in their portion of the data were representative of the data as a whole. From this, a set of initial




Figure 7.1: Summary of participant responses. Green shows the number of participants expressing agreement, orange shows disagreement, and grey indicates the topic was either not mentioned or no strong opinion was expressed.

proto-themes were generated that encapsulated the salient patterns the coders observed in the data.

In a second coding round, each coder revisited their data and revised, added, and removed codes with respect to the proto-themes that were generated. Then, the first author collected all of the codes to develop a conceptualization of the themes presented in Sec. 7.4, which were refined after further review of the data and codes. Quantitative data was compiled to summarize key features of these themes, which are presented throughout this Sec. 7.4 and consolidated in Fig. 7.1.

7.4 Results

Based on the analysis of the data, we found four themes that are representative of participants' perspectives toward grocery shopping and ADRs. These themes highlight both challenges and opportunities for the successful adoption of ADRs within this community. By addressing these takeaways, technologists and decision-makers can enhance patrons' agency and quality of life,

leading to more effective deployment of ADRs. The implications of these results are discussed further in Sec. 7.5.

7.4.1 Price dominates, but many factors are considered when patrons shop for food.

The first portion of the interviews shed light on how participants made decisions about food shopping and visiting the food pantry. Consistent with prior research [192,193,201], price was a primary consideration for most participants. We found the food pantry played a crucial role in providing flexibility and freedom, allowing participants to balance other important factors in their food shopping experience such as selecting foods they prefer and choosing their own fresh produce. This presents an opportunity for ADRs to further enhance this freedom by offering additional ways for patrons to access food.

Price as a driving force

The most frequently mentioned factor participants considered when grocery shopping was price. This factor influenced where participants shopped and the types of foods they sought out. As P21 expressed: *"If I see an ad, like the ad and see something on sale that I want, then I'm like, okay, I'll go there."* Price also played a crucial role in participants' decisions to visit the food pantry. As P6 noted, *"Because now, ah, food prices are high. Food pantry is free, so we hope to get some food that will help reduce the food budget."* Overall, 85.7% (18/21) of participants cited price as a factor they considered. This suggests that in order for ADR services to be viable for this community, they will likely need to be offered free of charge, subsidized by external funds.

Prioritizing shopping factors based on preferences

However, price was far from the only factor that participants considered when food shopping. They also weighed a variety of other factors, such as the freedom to select the types of foods they

want and the ability to choose their own items to ensure freshness. As P8 and P17 expressed, respectively, *"I like that we can pick out stuffs that we don't need ... the food we like we will take, those we don't like we won't take. Then there won't be a waste."* and *"I go to the Chinese market to buy vegetables. They are so fresh. I want to choose the bag that I like the most."*

As a result of having to carefully balance this complex set of factors, participants often faced inconveniences they were willing to tolerate, such as lengthy travel and long wait times at the food pantry. Some, like P5, spent multiple hours on public transportation to reach the food pantry: *"I take the bus, I think one hour and a half or sometimes two hours [each way]."*

Despite these challenges, participants expressed appreciation for the food pantry's services. They did not view their food pantry experiences as acts of desperation, but rather as opportunities that provided them with agency in their food shopping. The food pantry allowed participants to prioritize essential factors like price, even if it meant sacrificing others, such as time.

P11 – "It's free. So I don't think we can request anything more in my heart. It's impossible. They give away things for free, how can you have too much expectation?"

The results from this theme largely align with prior studies, which highlight price as a dominant shopping factor [192,193,201] and the challenge of traveling long distances [198,199]. Synthesizing these findings allows us to conclude that food pantries empower participants by offering them choices in their food shopping. Unaddressed in prior studies is the potential for ADRs to further enhance this agency by providing patrons with additional options to access the food they need.

Evolving needs

Many participants noted that while their current situations led them to a certain optimal balance of factors, they expect that their needs may change in the near future as they age. As a result, they saw ADRs as a way to maintain flexibility and preserve agency as they navigate the shifting priorities in their lives. Many participants, like P16, currently preferred shopping in person but

recognized that delivery services might become necessary in the near future: *"Because I can buy it myself, and I don't need it for now. When I get older ... maybe I will need it."* For others, these changing circumstances were already a reality:

P17 – "But since I am old, it is too heavy and too many things to carry, so I feel difficult ... But of course, if the delivery service is good, then it's good. We don't have to go there. If we go to there, we'll have to take the public transportation and transfer, and then we have to bring the food back. It is also hard."

ADR deployment for food pantries should focus on increasing the agency of patrons to food shop in ways that align with their preferred balance of factors like price, selecting preferred food types, picking the freshest produce, convenience, etc., especially as their life situations change. While ADRs can only address some of these factors, such as convenience, time, and the ability to select food types, enabling patrons to make choices that best suit their needs makes ADRs worthwhile.

7.4.2 Patrons' social connections are key to their food pantry experiences.

Participants found great value in a network of like-minded individuals with shared cultural backgrounds. This network included friends and family who informed them about the food pantry, as well as connections made with other patrons at the food pantry. Understanding and leveraging this network will be important for ADR adoption since formal methods of advertising may not be as effective. It is also necessary to acknowledge that ADRs may not support the preservation of connections made at the food pantry.

Reliance on social connections

Participants depended on friends and family for various aspects of their lives, including navigating technology and managing schedules. Social connections also influenced their food shopping habits, such as by learning about the food pantry through word of mouth:

P13 – “I talked to my friend about my situation and they just told me that ‘Oh I can bring you somewhere you can get free food.’ That’s how I know.”

While both partner food pantries run substantial awareness campaigns through social media, mailing flyers, and other methods, 66.7% of participants reported learning about the food pantry through word of mouth. Previous research indicates that social connections support SNAP recipients in managing their food needs, with shoppers favoring information from trustworthy acquaintances over unfamiliar sources [233]. Similarly, food pantry patrons may rely on their social networks because of established trust. Therefore, fostering trust with this community will be essential for effectively promoting ADR adoption.

Forming social connections at the food pantry

At the food pantry, participants had the chance to expand their social networks by interacting with others who spoke their native language and shared similar backgrounds. P4 enjoyed the social aspect of visiting the pantry, saying: *“That’s the reason I go to pantry because many people to talk ... the supermarket, I cannot talk too much. So that’s good.”* Similarly, P6 mentioned offering help to those they met at the pantry: *“I will chat with some people. Oftentimes, they don’t speak English, so I will help them a bit, the easy stuff. Sometimes when they need to sign up, they don’t understand.”*

Participants benefited from meeting people at the food pantry, gaining both social interaction and assistance that would otherwise be unavailable. While it may be challenging for ADR technologies to replicate the in-person community building that occurs at food pantries, it’s crucial to consider this factor in ADR development. The ineffectiveness of existing digital tools that automate communication and interaction between patrons might be due to this oversight [223]. ADR development should consider how the preferences and needs of users can be best served by technology while allowing patrons to determine themselves what trade-offs they are willing to make.

7.4.3 Language barriers and gaps in technological literacy are major obstacles to patrons adopting food delivery services, but they are open to learning how to use technology.

Only 9.5% of participants reported having used existing human food delivery services, with most having never considered trying them. The two main reasons cited were English language barriers (62%) and lack of technological literacy (62%). These factors pose challenges for adopting ADRs within this community. However, participants saw these issues as addressable with solutions such as offering mobile apps in multiple languages and providing assistance in navigating these apps, with many expressing a willingness to learn if given proper guidance.

Language barriers

Almost all participants were non-native English speakers, and most were uncomfortable with English in conversation and when operating mobile devices. They expressed a lack of confidence in their ability to navigate an app available only in English. This is consistent with previous findings on the prohibitive nature of language barriers at food pantries [234]. Although a straightforward solution to this problem is to offer the app in Chinese and other languages, identifying the most commonly spoken languages at the food pantry requires deliberate effort by technologists. Engaging with patrons to promote multi-language options will significantly impact the adoption of ADR services.

Challenges with technology, but a willingness to learn

Technological literacy also poses a challenge for using delivery services. Participants expressed low confidence in their ability to navigate apps. As P1 put it: *"But for me, no... Because you know, I don't know how to go to the app."* On the other hand, participants expressed a willingness to learn, with P9 stating, *"I would like to learn as long as someone can teach me."* While the Boston Public Library offers free computer and technology classes (even in languages like Chinese) [235], it did

not appear that participants were familiar with these services. Instead, participants envisioned relying on children and younger family members for guidance. As P18 shared: “ *I will ask my daughter to teach me. She teaches me, and I will know how to use.*” This ties into the earlier theme of patrons’ reliance on social connections, emphasizing the importance of understanding how to leverage these networks to promote ADR adoption.

Previous studies have provided recommendations for designers to make apps more accessible to older adults, who often struggle with smartphone interfaces [236,237] and these lessons can also be incorporated into ADR app design. ADRs introduce the additional challenge of requiring human-robot interaction during drop-off, which participants also expressed anxiety about. As P16 noted: “*Because the real person delivering the goods will talk with you, but I don’t know if the robot knows these things.*” Although studies have explored how visual and auditory cues can enhance communication between robots and humans [238–240], the samples have lacked diversity and need further investigation to be applicable to this community.

7.4.4 Despite mixed views on robot capabilities and risks, patrons are open to using robot delivery services.

Participants had mixed perspectives toward robots in terms of both their capabilities and the risks they posed. Some believed that robots could perform almost any task, while others were concerned about their reliability in safely navigating around people. Participants also expressed concern about surveillance and job loss. These views tie into broader, ongoing discussions about assigning responsibility when robots fail and establishing regulations to prevent negative outcomes. It is crucial for governments to closely consider the perspectives of marginalized groups to ensure that communities most at risk are protected from potential harms caused by robots and other technologies.

Overall, participants had more confidence in humans making deliveries compared to robots. Despite this, an overwhelming majority expressed interest in using ADRs now or in the future. This apparent contradiction may stem from participants’ belief that engineers will eventually re-

solve the current issues with robots. Consequently, roboticists have a responsibility to clearly communicate the capabilities and risks of robots and ensure they meet stakeholder expectations.

Confidence and concern about robots

When asked about situations that ADRs may struggle, some participants expressed full belief in robots. P12 stated, *"I feel like everything's so advanced now, that robots can pretty much do anything..."* Yet, others like P9 were concerned whether robots could safely navigate human environments: *"Yeah, I guess I'd be worried. Can they detect pedestrians? Yeah, I would be concerned about if they could sense a person. I don't want to see a crash or anything like that..."*

Participants also expressed concern about surveillance and job loss, though attitudes towards these issues varied widely. Regarding surveillance, some participants were uneasy about the possibility of a robot using a camera to record them, while others were indifferent.

P12 – "I will feel that they are spying on me when I am walking."

P3 – "I doesn't matter to me. I am an old lady."

Similarly, while some worried about job losses due to automation, others viewed the potential increase in efficiency favorably.

P1 – "Maybe a lot of people, you know, lost jobs. Right?"

P14 – "The good part is that one robot can replace four or five persons; it will save costs that can be used on development."

While the primary goal of this work is to improve equity in robot deployments by expanding access to the benefits of ADRs, these responses underscore the complementary need to equitably distribute the risks associated with ADRs. It is well-documented that certain groups, such as women and elderly adults, tend to have less positive attitudes toward robots [241,242] and face

increased risks from surveillance [145] and job displacement [243]. Regulatory agencies have a responsibility to protect vulnerable populations from harm, especially as interactions with robots will extend beyond the scope of current laws [244]. To effectively mitigate the risks posed to the food pantry community, it will be crucial to deepen our understanding of how robots will impact their lives and develop appropriate regulations.

Openness to using ADRs

Despite mixed views toward robots overall, participants were overwhelmingly consistent in their answers to the following question: “Would you be interested in using a delivery robot for the food pantry?” 81% of participants stated they would be interested in utilizing ADRs at some point if they were available. This result suggests that ADRs for food pantries would address participants’ self-identified needs and would likely be readily adopted by food pantry patrons.

This comes despite most participants believing that robots would not outperform humans and might even perform worse in areas like speed and navigating obstacles such as stairs. Overall, only 14.3% of participants preferred robot delivery over human delivery, while 47.6% favored human delivery, and 38.1% had no strong preference. This apparent contradiction between participants’ preference for human delivery and their interest in using ADR services might be explained by the novelty of encountering a new technology. Based on our data, we also propose that participants had high confidence in engineers’ ability to resolve any lingering issues with ADRs.

P3 – “If [roboticists] make this happen, you must have a good solution, right?”

It is unclear why participants had such high trust in engineers’ ability to preempt any challenges that ADRs might face. This trust could stem from a lack of awareness about the negative impacts robots and AI have had on underrepresented groups [245–247] or from a belief that the potential benefits of ADRs outweigh their concerns. Regardless, ongoing engagement with this community will be essential to ensure that the technologies deployed align with their expectations.

Participants' preference for human delivery raises the question of whether ADRs are a relevant solution at all. However, our discussions with partner pantries revealed that their ability to deliver food is limited by funding and volunteer capacity. Discussed further in Sec. 7.5.1, we hypothesize that recent city initiatives in public AI and robotics suggest governments may be more inclined to invest in technological solutions for food pantries, making ADRs a more feasible option. Further research is needed to explore how government incentives impact funding for various food pantry programs.

7.5 Discussion

Our results in Sec. 7.4 offer insights directly to the research question at hand, but are also relevant to ongoing discourse of how robots should be deployed in public spaces. This section discusses how our findings can contribute to understanding recent government initiatives on AI and robotics in public settings, as well as the regulatory responsibilities to address concerns.

7.5.1 Shaping the Future of ADRs

This work builds on existing literature about food insecurity and food pantry patrons by exploring a potential technological solution. ADRs can offer significant value to food pantry patrons by providing them with additional flexibility to manage the complex food shopping needs in their lives. However, both our findings and previous research indicate that price is a major constraint for this community. To ensure that patrons can fully access ADR services, offering these services at no cost will likely be essential. Current ADR deployments are focused on profit-driven businesses, so implementing this solution will require government intervention and a redefinition of design requirements. But what incentives do governments have to invest in a food pantry ADR program? One motivation is the desire for cities to be at the forefront of technology deployment [248]. Many cities have already made significant investments in drones, robots, and AI chatbots for applications in policing and education [249–252]. Another incentive is the interest in

promoting environmentally friendly alternatives to automobile travel, which was the motivation behind the US Department of Energy funding the Arlington food pantry ADR pilot [191]. We encourage cities to explore whether deploying ADRs for food pantries could simultaneously achieve their goals of becoming technology leaders, promoting environmental sustainability, and serving marginalized communities.

Engineers will also play a crucial role in bringing ADRs to underrepresented communities. Much of the existing human-robot interaction research has focused on privileged populations, potentially alienating those with different needs and preferences [253]. Developers can address this by identifying communities that could benefit from their technologies and forming lasting partnerships to involve diverse groups in the technology development process. Engineers must also recognize that technological solutions are not always the most appropriate response to societal challenges, as highlighted by participants' preference for human delivery over robots. Understanding factors such as market forces and government funding incentives is crucial for identifying where technology can make the most meaningful and positive impact.

7.5.2 Robots in Public Spaces

This work closely ties into ongoing discussions about governing robots in public spaces. Scholars have argued that aspects of existing law can apply to robot interactions, like tort liability [254], privacy rights [128,244], and self-defense [8]. However, legal precedents for these issues have not yet been established, leaving these questions unresolved in court. If technologies do not tightly adhere to regulations, they could face consequences like Cruise, whose autonomous vehicles were banned indefinitely after multiple accidents and traffic violations [255].

However, adhering to existing laws does not necessarily mean technologies are guaranteed to be accepted. For instance, concerns of safety led Toronto, Canada to ban ADRs from operating in the city [188]. While significant robotics research has focused on operating safely in human environment [150,151], there is clearly work to be done to align robot behavior with pedestrians' perceptions of safety. Developers must engage closely with the public and be sensitive to

expressed concerns. This work contributes to a better understanding of community perspectives on robots. We found concerns regarding privacy and job security, which have been echoed broadly by members of the public. As previously mentioned, this may be a consequence of participants not yet interacting with a robot and these perspectives may shift in the future. Continued effort to monitor public opinion of robots will be essential for robots to live up to their potential.

7.6 Limitations

A particularly challenging aspect of this work was navigating the language barrier between the interviewer and participants. As mentioned, this was handled by providing live interpretation services along with translating all the spoken Cantonese into English during transcription. We chose to translate each Cantonese speaker into English twice (once during the interview and once afterwards from the recordings) in order to ensure participants' words were reported as faithfully as possible. We believed that the post-hoc translation would be more accurate because of the time constraints associated with live interpretation. There were occasions where there was minor discrepancy in the transcripts between the interpreted English and translated English, and in all of these cases we deferred to the translated English for coding and reporting. However, because none of the researchers spoke Cantonese, there was no reasonable way to prove which version was more faithful in these cases.

While this study successfully captured the perspectives of a previously unrepresented community toward robot deployment, there are potential sources of sample bias that may have made it difficult for certain groups to participate in the study. First is the fact that by recruiting participants at the food pantry, people who have either chosen to not go to the food pantry or who are unable to were not recruited. This population deserves future attention, as enabling more people to access food pantry services should be a key goal. Additionally, as discussed in Sec 7.3.2, only English and Cantonese speakers participated in this study, excluding speakers of many other languages that may have diverse perspectives toward ADRs. For this study, forming further con-

nections with the food pantry community and related community groups may have informed us of how to better recruit Haitian Creole speakers. These sample limitations are also opportunities for future research to further engage groups that were not fully captured in this work.

7.7 Future Work

One potential extension of this work is to investigate whether ADRs could benefit groups not represented in this study, such as those who do not access food pantry services due to personal choice or mobility issues. Another future direction involves conducting large-scale quantitative studies to identify trends across the broader food pantry patron population, which would be especially valuable for informing policy decisions on ADR deployment. Observing participants directly interact with a robot would also provide valuable insights into how this experience might change their perceptions of robots.

As AI and robotics increasingly shape society, it is crucial for technologists to understand how various populations perceive and interact with these technologies. Future studies could explore the roots of patrons' views on robots, providing roboticists with insights into effectively engaging with these communities and encouraging the adoption of ADRs and similar innovations. Similarly, to understand how advocates can promote public investment in equitable technology applications, future research could examine the decision-making process behind government investments in experimental AI and robotics pilot programs, as well as the main stakeholders shaping these decisions.

7.8 Conclusion

Despite the growing presence of ADRs, research on their impact on underrepresented communities remains minimal. This study extends prior work on food insecurity and automated food delivery systems by exploring how engineers and policymakers can effectively deploy ADRs for food

pantry patrons. We interviewed 21 patrons from the Boston area, who shared their food shopping experiences and perspectives on delivery services and ADRs. Participants emphasized that ADRs could enhance their ability to navigate various food shopping factors. Although challenges such as language barriers, technological literacy, and privacy concerns exist, participants showed strong interest in using ADR services if available. Amid ongoing public discussions about regulating robotics technologies like ADRs, this research provides valuable insights into the potential risks and benefits for food pantry patrons and contributes to improving technology accessibility for underrepresented communities.

7.9 Acknowledgements

We extend our deepest gratitude to Phoebe Heyman of the Margaret Fuller Neighborhood House, Ryan Montblaeu of the East End House, Ingrid Gamez and Atika Kouhail of the Cambridge Health Alliance, and all the interpreters and translators whose support was vital to this study. We also sincerely thank the participants for their kindness, intelligence, and thoughtfulness throughout the research.

Part III

Concluding Remarks

Chapter 8

Conclusion

This thesis examines the technical and societal factors influencing the success of robot deployment. In particular, I examine two open problems in robotics: navigating complex, uncertain environments and navigating complex, human environments. My approach to both of these problems involves identifying robot failure modes, analyzing their causes, and then proposing new designs, algorithms, or policy recommendations to improve robot performance.

Part I addressed the issue of robots navigating through uncertain terrain where disturbances could cause a robot to become unstable. Chapter 2 introduced two local measures of stability and convergence, which are scalar values that reflect the physical behavior of perturbed trajectories. The advantage of these measures is that they can be incorporated into various optimization frameworks to improve robot performance. Chapters 3, 4, and 5 leveraged these measures in varying ways to generate robot trajectories with improved tracking performance. Results have been implemented in simulation, ranging from simple monopod hoppers to more complex biped and quadruped systems, as well as hardware which validated the success of these novel insights.

Part II presented two studies aimed at understanding where large-scale public robot deployments may face social and legal challenges and how engineers can proactively address these difficulties. Chapter 6 discussed how well-established self-defense law extends to human defense against robots and steps that can be taken to mitigate the likelihood of self-defense scenarios

occurring. Chapter 7 revealed insights into how food pantry patrons could benefit from access to autonomous delivery robots and the unique considerations necessary to bring this service to this population.

This thesis advances the robotics field by identifying key failure modes—technical, social, and regulatory—that must be addressed to ensure robots can effectively benefit society. To effectively navigate the complexities of robot deployment, roboticists must position themselves at the intersection of these diverse fields.

8.1 Future Work

While many unanswered questions persist in the individual robotics areas of robust control, human-centered design, and regulation, an even greater challenge lies in the gap between these topics, which are often explored in isolation. Bridging these divides is crucial to ensure that, as robots are deployed in greater numbers in public spaces, decision-makers are prepared to navigate the immense societal impacts. This section outlines several potential extensions to this thesis aimed at working toward this outcome.

8.1.1 Stability of Systems Under Stick-Slip Transitions

An extension of saltation-based analysis for hybrid systems involves examining stick-slip and slip-stick transitions. Small, instantaneous slips are common in both legged and wheeled locomotion, requiring strategies to manage these potentially disruptive perturbations. These transitions also induce saltation behavior, meaning the analysis from this thesis can be applied. Given the challenges in modeling stick-slip transitions, incorporating methods that account for uncertainty may enhance robustness, as suggested by current hybrid state estimation literature [256].

8.1.2 How People Develop Internal Models of Robots

In Chapter 7, Theme 4 highlighted that participants' perceptions of robots varied independently of their prior experience. Many participants expressed beliefs about how robots functioned, their confidence in robot performance, and concerns about risks, often without supporting evidence. To guide developers in educating the public and gathering feedback, it's important to understand how people form their views on robots. This could involve both qualitative studies and large-scale sentiment analyses on platforms like X or Reddit.

8.1.3 Policy Making For Public Robots

Part II highlights the large role that government will play in robot deployment, both in terms of developing regulation as well as investing in public robot initiatives like ADRs for food pantries. Recent anecdotes have shown that engineers, governments, and the public have opposing views on several robot applications [138,143,257]. Yet, municipal governments seem eager to invest public funds into robots, particularly for policing [139,250]. What incentivizes governments to deploy robots and how can engineers and the public influence the decision-making process? Answering this question would enable more democratic participation in shaping society's technological future.

8.1.4 Uniting Robot Failure Modes

Many factors can cause a mobile robot could fail, ranging from environmental issues like using an unstable gait or missing a foothold, to social reasons like inducing a perception of physical danger. While existing work has typically addressed these risks in isolation, they must be considered together in practice. Future work should focus on developing a unified approach to assess these risks. For example, a robot navigating the beam environment from Chapter 5 should weigh failure modes differently than one in the crowded sidewalk environment from Chapter 6. Autonomous evaluation of these risks will be crucial for robots to perform effectively in diverse environments.

Bibliography

- [1] M. S. Branicky, "Multiple Lyapunov functions and other analysis tools for switched and hybrid systems," *IEEE Transactions on Automatic Control*, 1998.
- [2] N. J. Kong, J. J. Payne, J. Zhu, and A. M. Johnson, "Saltation matrices: The essential tool for linearizing hybrid dynamical systems," *Proceedings of the IEEE*, 2024.
- [3] A. Seyfarth, H. Geyer, and H. Herr, "Swing-leg retraction: a simple control model for stable running," *Journal of Experimental Biology*, 2003.
- [4] J. Zhu, N. Kong, G. Council, and A. Johnson, "Hybrid event shaping to stabilize periodic hybrid orbits," *IEEE International Conference on Robotics and Automation*, 2022.
- [5] J. Zhu, J. J. Payne, and A. M. Johnson, "Convergent iLQR for safe trajectory planning and control of legged robots," *IEEE International Conference on Robotics and Automation*, 2024.
- [6] W. Xi, Y. Yesilevskiy, and C. D. Remy, "Selecting gaits for economical locomotion of legged robots," *The International Journal of Robotics Research*, 2016.
- [7] J. Zhu, D. Ologan, S. G. Gonzalez, A. Tajbakhsh, and A. M. Johnson, "Duty factor predicts safe quadrupedal robot locomotion," *In preparation*, 2024.
- [8] J. Zhu, A. Shrivastava, and A. M. Johnson, "Grounding robot navigation in self-defense law," *IEEE International Conference on Robot and Human Interactive Communication*, 2023.

- [9] J. Zhu, P. Reynolds-Cuellar, N. Nakajima, A. Yadav, E. Jain, A. M. Johnson, and K. Darling, "Improving equity in robot deployment: A study of food pantry patrons," *In preparation*, 2024.
- [10] A. van der Schaft and J. Schumacher, "Complementarity modeling of hybrid systems," *IEEE Transactions on Automatic Control*, 1998.
- [11] J. Lygeros, K. H. Johansson, S. Simic, J. Zhang, and S. S. Sastry, "Dynamical properties of hybrid automata," *IEEE Transactions on Automatic Control*, 2003.
- [12] A. M. Johnson, S. A. Burden, and D. E. Koditschek, "A hybrid systems model for simple manipulation and self-manipulation systems," *The International Journal of Robotics Research*, 2016.
- [13] M. Jeffrey, "Dynamics at a switching intersection: Hierarchy, isonomy, and multiple sliding," *SIAM Journal on Applied Dynamical Systems*, 2014.
- [14] S. N. Simić, K. H. Johansson, S. Sastry, and J. Lygeros, "Towards a Geometric Theory of Hybrid Systems," in *Hybrid Systems: Computation and Control*, ser. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2000.
- [15] J. Zhang, K. H. Johansson, J. Lygeros, and S. Sastry, "Dynamical systems revisited: Hybrid systems with zeno executions," in *Hybrid Systems: Computation and Control*. Springer Berlin Heidelberg, 2000.
- [16] M. Aizerman. and F. Gantmacher, "Determination of Stability by Linear Approximation of a Periodic Solution of a System of Differential Equations With Discontinuous Right-hand Sides," *The Quarterly Journal of Mechanics and Applied Mathematics*, 1958.
- [17] S. A. Burden and S. D. Coogan, "On infinitesimal contraction analysis for hybrid systems," *IEEE Conference on Decision and Control*, 2022.

- [18] R. Leine and H. Nijmeijer, *Dynamics and Bifurcations of Non-Smooth Mechanical Systems*. Springer-Verlag Berlin Heidelberg, 2004.
- [19] M. Hirsch, S. Smale, and R. Devaney, *Differential Equations, Dynamical Systems, and an Introduction to Chaos*. Academic Press, 2012.
- [20] H. Asahara and T. Kousaka, "Stability analysis using monodromy matrix for impacting systems," *IEICE Transactions of Fundamentals of Electronics, Communications, and Computer Science*, 2018.
- [21] B. Altın, P. Ojaghi, and R. G. Sanfelice, "A model predictive control framework for hybrid dynamical systems," *IFAC-PapersOnLine*, 2018.
- [22] N. J. Kong, G. Council, and A. M. Johnson, "iLQR for piecewise-smooth hybrid dynamical systems," *IEEE Conference on Decision and Control*, 2021.
- [23] W. Yang and M. Posa, "Impact invariant control with applications to bipedal locomotion," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021.
- [24] S. Burden, S. Revzen, and S. Sastry, "Dimension reduction near periodic orbits of hybrid systems," in *IEEE Conference on Decision and Control*, 2011.
- [25] G. Council, S. Yang, and S. Revzen, "Deadbeat control with (almost) no sensing in a hybrid model of legged locomotion," *International Conference on Advanced Mechatronic Systems*, 2014.
- [26] D. Sternad, M. Duarte, H. Katsumata, and S. Schaal, "Bouncing a ball: Tuning into dynamic stability," *Journal of Experimental Psychology: Human Perception and Performance*, 2001.
- [27] M. M. Ankaralı, U. Saranlı, and A. Saranlı, "Control of underactuated planar hexapedal pronking through a dynamically embedded slip monopod," *IEEE International Conference on Robotics and Automation*, 2010.

- [28] K. Green, R. Hatton, and J. Hurst, "Planning for the unexpected: Explicitly optimizing motions for ground uncertainty in running," *IEEE International Conference on Robotics and Automation*, 2020.
- [29] J. G. D. Karssen, M. Haberland, M. Wisse, and S. Kim, "The optimal swing-leg retraction rate for running," *IEEE International Conference on Robotics and Automation*, 2011.
- [30] E. Tyrtysnikov, *A Brief Introduction to Numerical Analysis*. Birkhäuser Basel, 1997.
- [31] P. Lancaster, "On eigenvalues of matrices dependent on a parameter," *Numerische Mathematik*, 1964.
- [32] H. Khalil, *Nonlinear Systems*. Pearson, 2002.
- [33] J. Nocedal and S. Wright, *Numerical Optimization*. Springer, 2006.
- [34] M. Kelly, "An introduction to trajectory optimization: How to do your own direct collocation," *SIAM Review*, 2017.
- [35] O. Von Stryk and R. Bulirsch, "Direct and indirect methods for trajectory optimization," *Annals of Operations Research*, 1992.
- [36] C. R. Hargraves and S. W. Paris, "Direct trajectory optimization using nonlinear programming and collocation," *Journal of Guidance, Control, and Dynamics*, 1987.
- [37] K. Mombaur, R. Longman, H. Bock, and J. Schlöder, "Open-loop stable running," *Robotica*, 2005.
- [38] O. V. Stryk, "User's guide for DIRCOL (version 2.1): a direct collocation method for the numerical solution of optimal control problems," in *Technische Universität Darmstadt*, 1999.
- [39] J. T. Betts, *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming, Second Edition*, 2nd ed. SIAM, 2010.

- [40] T. A. McMahon and G. C. Cheng, "The mechanics of running: How does stiffness couple with speed?" *Journal of Biomechanics*, 1990.
- [41] P. Holmes, R. J. Full, D. E. Koditschek, and J. Guckenheimer, "The dynamics of legged locomotion: Models, analyses, and challenges," *SIAM Review*, 2006.
- [42] A. Seyfarth, H. Geyer, M. Günther, and R. Blickhan, "A movement criterion for running," *Journal of Biomechanics*, 2002.
- [43] R. J. Full and D. E. Koditschek, "Templates and anchors: neuromechanical hypotheses of legged locomotion on land," *Journal of Experimental Biology*, 1999.
- [44] J. Seipel and P. Holmes, "A simple model for clock-actuated legged locomotion," *Regular and Chaotic Dynamics*, 2007.
- [45] M. Raibert, H. Brown, and M. Chepponis, "Experiments in balance with a 3D one-legged hopping machine," *The International Journal of Robotics Research*, 1984.
- [46] I. Poulakakis and J. W. Grizzle, "The spring loaded inverted pendulum as the hybrid zero dynamics of an asymmetric hopper," *IEEE Transactions on Automatic Control*, 2009.
- [47] V. Chen, "Passive dynamic walking with knees: A point foot model," Master's thesis, Massachusetts Institute of Technology, 2007.
- [48] A. M. Johnson and D. E. Koditschek, "Toward a vocabulary of legged leaping," in *IEEE Intl. Conference on Robotics and Automation*, 2013.
- [49] H. Kolvenbach, E. Hampp, P. Barton, R. Zenkl, and M. Hutter, "Towards jumping locomotion for quadruped robots on the moon," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019.
- [50] C. Nguyen and Q. Nguyen, "Contact-timing and trajectory optimization for 3d jumping on quadruped robots," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022.

- [51] Z. Li, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, "Robust and versatile bipedal jumping control through multi-task reinforcement learning," *Robotics: Science and Systems*, 2023.
- [52] R. Tedrake, *Underactuated Robotics*, 2022. [Online]. Available: <http://underactuated.mit.edu>
- [53] M. A. Haidekker, *Linear Feedback Controls*, M. A. Haidekker, Ed. Elsevier, 2020.
- [54] J. Norby, A. Tajbakhsh, Y. Yang, and A. M. Johnson, "Adaptive complexity model predictive control," *IEEE Transactions on Robotics*, 2024.
- [55] Z. Manchester and S. Kuindersma, "Robust direct trajectory optimization using approximate invariant funnels," *Auton. Robots*, 2019.
- [56] J. Morimoto, G. Zeglin, and C. Atkeson, "Minimax differential dynamic programming: application to a biped walking robot," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2003.
- [57] T. Lew, R. Bonalli, and M. Pavone, "Chance-constrained sequential convex programming for robust trajectory optimization," *European Control Conference*, 2020.
- [58] B. Hammoud, M. Khadiv, and L. Righetti, "Impedance optimization for uncertain contact interactions through risk sensitive optimal control," *IEEE Robotics and Automation Letters*, 2021.
- [59] M. Ahmadi, X. Xiong, and A. D. Ames, "Risk-sensitive path planning via CVaR barrier functions: Application to bipedal locomotion," *IEEE Control Systems Letters*, 2021.
- [60] L. Drnach and Y. Zhao, "Robust trajectory optimization over uncertain terrain with stochastic complementarity," *IEEE Robotics and Automation Letters*, 2021.
- [61] S. Kousik, S. Vaskov, F. Bu, M. Johnson-Roberson, and R. Vasudevan, "Bridging the gap between safety and real-time performance in receding-horizon trajectory design for mobile robots," *The International Journal of Robotics Research*, 2020.

- [62] C.-Y. Lee, S. Yang, B. Bokser, and Z. Manchester, "Enhanced balance for legged robots using reaction wheels," *IEEE International Conference on Robotics and Automation*, 2023.
- [63] Y. Yang, J. Norby, J. K. Yim, and A. M. Johnson, "Proprioception and tail control enable extreme terrain traversal by quadruped robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023.
- [64] A. Papachristodoulou and S. Prajna, "On the construction of Lyapunov functions using the sum of squares decomposition," *IEEE Conference on Decision and Control*, 2002.
- [65] S. Chen, M. Fazlyab, M. Morari, G. J. Pappas, and V. M. Preciado, "Learning Lyapunov functions for hybrid systems," *International Conference on Hybrid Systems: Computation and Control*, 2021.
- [66] R. Grandia, A. J. Taylor, A. D. Ames, and M. Hutter, "Multi-layered safety for legged robots via control barrier functions and model predictive control," *IEEE International Conference on Robotics and Automation*, 2021.
- [67] W. Lohmiller and J. Slotine, "On contraction analysis for non-linear systems," *Automatica*, 1998.
- [68] A. M. Johnson, J. E. King, and S. Srinivasa, "Convergent planning," *IEEE Robotics and Automation Letters*, 2016.
- [69] N. Kong and A. M. Johnson, "Optimally convergent trajectories for navigation," in *International Symposium on Robotics Research*, 2019.
- [70] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems," *International Conference on Informatics in Control, Automation and Robotics*, 2004.
- [71] N. Kong, C. Li, G. Council, and A. M. Johnson, "Hybrid iLQR model predictive control for contact implicit stabilization on legged robots," *IEEE Transactions on Robotics*, 2023.

- [72] J. Townsend, "Differentiating the singular value decomposition," Technical Report 2016, [https://j-towns.github.io/papers/svd-derivative ...](https://j-towns.github.io/papers/svd-derivative...), Tech. Rep., 2016.
- [73] S. Singh, R. P. Russell, and P. M. Wensing, "On second-order derivatives of rigid-body dynamics: Theory & implementation," *IEEE Transactions on Robotics*, 2024.
- [74] C. G. Broyden, "The convergence of a class of double-rank minimization algorithms 1. General considerations," *IMA Journal of Applied Mathematics*, 1970.
- [75] K. Endo and H. Herr, "A model of muscle-tendon function in human walking at self-selected speed," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2014.
- [76] M. Scheint, M. Sobotka, and M. Buss, "Optimized parallel joint springs in dynamic motion: Comparison of simulation and experiment," *IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics*, 2010.
- [77] K. Mombaur, "Using optimization to create self-stable human-like running," *Robotica*, 2009.
- [78] D. Lakatos, K. Ploeger, F. Loeffl, D. Seidel, F. Schmidt, T. Gumpert, F. John, T. Bertram, and A. Albu-Schäffer, "Dynamic locomotion gaits of a compliantly actuated quadruped with slip-like articulated legs embodied in the mechanical design," *IEEE Robotics and Automation Letters*, 2018.
- [79] J. Ackerman and J. Seipel, "Energy efficiency of legged robot locomotion with elastically suspended loads," *IEEE Transactions on Robotics*, 2013.
- [80] M. J. Pollayil, C. D. Santina, G. Mesesan, J. Engelsberger, D. Seidel, M. Garabini, C. Ott, A. Bicchi, and A. Albu-Schaffer, "Planning natural locomotion for articulated soft quadrupeds," *IEEE International Conference on Robotics and Automation*, 2022.
- [81] S. Halder, K. Afsari, E. Chiou, R. Patrick, and K. A. Hamed, "Construction inspection & monitoring with quadruped robots in future human-robot teaming: A preliminary study," *Journal of Building Engineering*, 2023.

- [82] C. Gehring, P. Fankhauser, L. Isler, R. Diethelm, S. Bachmann, M. Potz, L. Gerstenberg, and M. Hutter, "ANYmal in the field: Solving industrial inspection of an offshore HVDC platform with a quadrupedal robot," in *Field and Service Robotics*. Springer Singapore, 2021.
- [83] I. D. Miller, F. Cladera, A. Cowley, S. S. Shivakumar, E. S. Lee, L. Jarin-Lipschitz, A. Bhat, N. Rodrigues, A. Zhou, A. Cohen, A. Kulkarni, J. Laney, C. J. Taylor, and V. Kumar, "Mine tunnel exploration using multiple quadrupedal robots," *IEEE Robotics and Automation Letters*, 2020.
- [84] B. Hammoud, M. Khadiv, and L. Righetti, "Impedance optimization for uncertain contact interactions through risk sensitive optimal control," *IEEE Robotics and Automation Letters*, 2021.
- [85] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, 2022.
- [86] A. D. Kuo, "A simple model of bipedal walking predicts the preferred speed–step length relationship," *Journal of Biomechanical Engineering*, 2001.
- [87] D. F. Hoyt and C. R. Taylor, "Gait and the energetics of locomotion in horses," *Nature*, 1981.
- [88] M. Garcia, A. Chatterjee, and A. Ruina, "Efficiency, speed, and scaling of two-dimensional passive-dynamic walking," *Dynamics and Stability of Systems*, 2000.
- [89] Z. Fu, A. Kumar, J. Malik, and D. Pathak, "Minimizing energy consumption leads to the emergence of gaits in legged robots," *Conference on Robot Learning*, 2022.
- [90] M. Vukobratović and B. Borovac, "Zero-moment point—thirty five years of its life," *International journal of humanoid robotics*, 2004.
- [91] C. Gonzalez, V. Barasuol, M. Frigerio, R. Featherstone, D. G. Caldwell, and C. Semini, "Line walking and balancing for legged robots with point feet," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.

- [92] Y. Yang, J. Norby, J. K. Yim, and A. M. Johnson, "Proprioception and tail control enable extreme terrain traversal by quadruped robots," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023.
- [93] G. Bledt, M. J. Powell, B. Katz, J. Di Carlo, P. M. Wensing, and S. Kim, "Mit cheetah 3: Design and control of a robust, dynamic quadruped robot," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018.
- [94] B. Ugurlu, I. Havoutis, C. Semini, and D. G. Caldwell, "Dynamic trot-walking with the hydraulic quadruped robot — HyQ: Analytical trajectory generation and active compliance control," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013.
- [95] S. Fahmi, C. Mastalli, M. Focchi, and C. Semini, "Passive whole-body control for quadruped robots: Experimental validation over challenging terrain," *IEEE Robotics and Automation Letters*, 2019.
- [96] D. Kim, J. Di Carlo, B. Katz, G. Bledt, and S. Kim, "Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control," *arXiv preprint arXiv:1909.06586*, 2019.
- [97] W. Chi, X. Jiang, and Y. Zheng, "A linearization of centroidal dynamics for the model-predictive control of quadruped robots," *International Conference on Robotics and Automation*, 2022.
- [98] J. Norby, Y. Yang, A. Tajbakhsh, J. Ren, J. K. Yim, A. Stutt, Q. Yu, N. Flowers, and A. M. Johnson, "Quad-SDK: Full stack software framework for agile quadrupedal locomotion," in *ICRA Workshop on Legged Robots*, May 2022, workshop abstract.
- [99] J. E. Bertram and A. Ruina, "Multiple walking speed–frequency relations are predicted by constrained optimization," *Journal of Theoretical Biology*, 2001.

- [100] J. Donelan, R. Kram, and A. Kuo, "Mechanical and metabolic determinants of the preferred step width in human walking," *Proceedings of the Royal Society B: Biological Sciences*, 2001.
- [101] K. G. Holt, S. F. Jeng, R. Ratcliffe, and J. Hamill, "Energetic cost and stability during human walking at the preferred stride frequency," *Journal of Motor Behavior*, 2010.
- [102] P. M. M. Young and J. B. Dingwell, "Voluntarily changing step length or step width affects dynamic stability of human walking," *Gait & Posture*, 2012.
- [103] S. A. England and K. P. Granata, "The influence of gait speed on local dynamic stability of walking," *Gait & Posture*, 2007.
- [104] D. F. Hoyt, S. J. Wickler, D. J. Dutto, G. E. Catterfeld, and D. Johnsen, "What are the relations between mechanics, gait parameters, and energetics in terrestrial locomotion?" *Journal of Experimental Zoology Part A: Comparative Experimental Biology*, 2006.
- [105] R. McGhee and A. Frank, "On the stability properties of quadruped creeping gaits," *Mathematical Biosciences*, 1968.
- [106] T.-T. Lee and C.-L. Shih, "A study of the gait control of a quadruped walking vehicle," *IEEE Journal on Robotics and Automation*, 1986.
- [107] T.-W. Koo and Y. Yong-San, "Dynamic instant gait stability measure for quadruped walking robot," *Robotica*, 1999.
- [108] E. W. McClain and S. Meek, "Determining optimal gait parameters for a statically stable walking human assistive quadruped robot," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018.
- [109] M. Posa, C. Cantu, and R. Tedrake, "A direct method for trajectory optimization of rigid bodies through contact," *The International Journal of Robotics Research*, 2014.
- [110] K. Tsujita, M. Kawakami, and K. Tsuchiya, "A study on optimal gait pattern of a quadruped locomotion robot," *IEEE International Conference on Systems, Man and Cybernetics*, 2004.

- [111] J. Carpentier, G. Saurel, G. Buondonno, J. Mirabel, F. Lamiroux, O. Stasse, and N. Mansard, "The pinocchio c++ library – a fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives," in *IEEE International Symposium on System Integrations*, 2019.
- [112] N. J. Kong, C. Li, G. Council, and A. M. Johnson, "Hybrid ilqr model predictive control for contact implicit stabilization on legged robots," *IEEE Transactions on Robotics*, 2023.
- [113] J. N. Nganga, H. Li, and P. M. Wensing, "Second-order differential dynamic programming for whole-body mpc of legged robots," *IFAC-PapersOnLine*, 2023.
- [114] J. K. Yim, J. Ren, D. Ologan, S. G. Gonzalez, and A. M. Johnson, "Proprioception and reaction for walking among entanglements," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023.
- [115] D. Lee, G. Kang, B. Kim, and D. H. Shim, "Assistive delivery robot application for real-world postal services," *IEEE Access*, 2021.
- [116] A. Lopez, R. Paredes, D. Quiroz, G. Trovato, and F. Cuellar, "Robotman: A security robot for human-robot interaction," *International Conference on Advanced Robotics*, 2017.
- [117] A. Vega-Magro, L. V. Calderita, P. Bustos, and P. Núñez, "Human-aware robot navigation based on time-dependent social interaction spaces: a use case for assistive robotics," *IEEE International Conference on Autonomous Robot Systems and Competitions*, 2020.
- [118] S. Brown-Heidenreich, "KiwiBot catches fire outside Martin Luther King Jr. student union," *The Daily Californian*, 2018.
- [119] L. O'Haver, "UT students arrested after damaging food delivery robot," *WATE*, 2022.
- [120] E. Parks and D. Zoga, "NBC 5 responds: What happens if you're in a crash with a robot?" *NBCDFW*, 2020.

- [121] R. M. Brown, *No duty to retreat: Violence and values in American history and society*. Oxford University Press, 1991.
- [122] R. A. Epstein, "The theory and practice of self-help," *Journal of Law, Economics, and Policy*, 2005.
- [123] D. I. Brandon, M. L. Cooper, J. H. Greshin, A. L. Harris, and J. M. Head, "Self-help: Extrajudicial rights, privileges and remedies in contemporary American society," *Vanderbilt Law Review*, 1984.
- [124] K. W. Simons, "Self-defense: Reasonable beliefs or reasonable self-control?" *New Criminal Law Review*, 2008.
- [125] L. S. Scheff, "People v. Humphrey: Justice for battered women or a licence to kill," *University of San Francisco Law Review*, 1997.
- [126] H. P. Thompson, "The State v. Sutton," *Supreme Court of Georgia*, 2015.
- [127] R. A. Rosen, "On self-defense, imminence, and women who kill their batterers," *North Carolina Law Review*, 1993.
- [128] A. M. Froomkin and Z. Colangelo, "Self-defense against robots and drones," *Connecticut Law Review*, 2015.
- [129] D. Schneider, "Open season on drones?" *IEEE Spectrum*, 2014.
- [130] S. Uniacke, "Proportionality and self-defense," *Law and Philosophy*, 2011.
- [131] B. Levin, "A defensible defense?: Reexamining castle doctrine statutes," *Harvard Journal on Legislation*, 2010.
- [132] D. B. Kopel, "How the United States Supreme Court confronted a hanging judge in the nineteenth century and taught some lessons for jurisprudence in the twenty-first," *American Journal of Criminal Law*, 2000.

- [133] J. E. Anderson and D. Dunning, "Behavioral norms: Variants and their identification," *Social and Personality Psychology Compass*, 2014.
- [134] R. Kirby, R. Simmons, and J. Forlizzi, "COMPANION: A constraint-optimizing method for person-acceptable navigation," *IEEE International Symposium on Robot and Human Interactive Communication*, 2009.
- [135] A. Sardar, M. Joosse, A. Weiss, and V. Evers, "Don't stand so close to me: Users' attitudinal and behavioral responses to personal space invasion by robots," *ACM/IEEE International Conference on Human-Robot Interaction*, 2012.
- [136] I. Cohen, R. Looije, and M. A. Neerincx, "Child's recognition of emotions in robot's face and body," *ACM/IEEE International Conference on Human-Robot Interaction*, 2011.
- [137] M. K. Strait, C. Aguilon, V. Contreras, and N. Garcia, "The public's perception of humanlike robots: Online social commentary reflects an appearance-based uncanny valley, a general fear of a "technology takeover", and the unabashed sexualization of female-gendered robots," *IEEE International Symposium on Robot and Human Interactive Communication*, 2017.
- [138] A. Yunus and S. A. Doore, "Responsible use of agile robots in public spaces," *IEEE International Symposium on Ethics in Engineering, Science and Technology*, 2021.
- [139] M. Zaveri, "N.Y.P.D.'s robot dog returns to work, touching off a backlash," *The New York Times*, 2021.
- [140] S. A. Periwai, "Boston Dynamics' robot dog strikes fear online as fans recall Black Mirror episode 'Metalhead'," *Sportskeeda*, 2020.
- [141] A. Jarmanning, "Mass. state police tested out Boston Dynamics' Spot the robot dog. civil liberties advocates want to know more," *WBUR*, 2019.

- [142] L. Takayama, W. Ju, and C. Nass, "Beyond dirty, dangerous and dull: What everyday people think robots should do," *ACM/IEEE International Conference on Human-Robot Interaction*, 2008.
- [143] J. Buolamwini and T. Gebru, "Gender shades: Intersectional accuracy disparities in commercial gender classification," *Conference on Fairness, Accountability and Transparency*, 2018.
- [144] C. Gidaris, "The rise of the robots: Technocapitalism and the policing of race," *Fast Capitalism*, 2021.
- [145] K. M. Thomasen, "Beyond airspace safety: A feminist perspective on drone privacy regulation," *Canadian Journal of Law and Technology*, 2018.
- [146] D. Tourjée, "Several women in this rural town say they were stalked by drones," *Vice Media Group*, 2017.
- [147] A. Gurinskaya, "Support for robocops: Measuring effects of attitudes towards police and policing technologies," *Digital Transformation and Global Society*, 2022.
- [148] 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*, 1968.
- [149] S. M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," *Tech. Rep. TR 98-11*, 1998.
- [150] J. Chu, F. Zhao, S. Bakshi, Z. Yan, and D. Chen, "Risk-aware path planning with uncertain human interactions," *American Control Conference*, 2021.
- [151] A. V. Taylor, E. Mamantov, and H. Admoni, "Observer-aware legibility for social navigation," *IEEE International Conference on Robot and Human Interactive Communication*, 2022.
- [152] E. Sanoubari, B. David, C. Kew, C. Cunningham, and K. Caluwaerts, "From message to expression: Exploring non-verbal communication for appearance-constrained robots," *IEEE International Conference on Robot and Human Interactive Communication*, 2022.

- [153] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017.
- [154] P. Langley, B. Meadows, M. Sridharan, and D. Choi, "Explainable agency for intelligent autonomous systems," *AAAI Conference on Artificial Intelligence*, 2017.
- [155] M. Cashmore, A. Collins, B. Krarup, S. Krivic, D. Magazzeni, and D. Smith, "Towards explainable ai planning as a service," *ICAPS Workshop on Explainable Planning*, 2019.
- [156] S. Sohrabi, J. A. Baier, and S. A. McIlraith, "Preferred explanations: Theory and generation via planning," *AAAI Conference on Artificial Intelligence*, 2011.
- [157] L. Zhu and T. Williams, "Effects of proactive explanations by robots on human-robot trust," *Social Robotics*, 2020.
- [158] S. Atakishiyev, M. Salameh, H. Yao, and R. Goebel, "Explainable artificial intelligence for autonomous driving: An overview and guide for future research directions," *IEEE Access*, 2024.
- [159] S. Anjomshoei, A. Najjar, D. Calvaresi, and K. Främling, "Explainable agents and robots: Results from a systematic literature review," *International Conference on Autonomous Agents and Multiagent Systems*, 2019.
- [160] M. W. Boyce, J. Y. Chen, A. R. Selkowitz, and S. G. Lakhmani, "Effects of agent transparency on operator trust," *ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts*, 2015.
- [161] N. Wang, D. V. Pynadath, and S. G. Hill, "The impact of POMDP-generated explanations on trust and performance in human-robot teams," *International Conference on Autonomous Agents & Multiagent Systems*, 2016.

- [162] General Data Protection Regulation, "Regulation EU 2016/679 of the European parliament and of the council of 27 April 2016," *Official Journal of the European Union*, 2016.
- [163] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon, "A human aware mobile robot motion planner," *IEEE Transactions on Robotics*, 2007.
- [164] N. J. Hetherington, E. A. Croft, and H. M. Van der Loos, "Hey robot, which way are you going? Nonverbal motion legibility cues for human-robot spatial interaction," *IEEE Robotics and Automation Letters*, 2021.
- [165] J. Guzzi, A. Giusti, L. M. Gambardella, G. Theraulaz, and G. A. Di Caro, "Human-friendly robot navigation in dynamic environments," *IEEE International Conference on Robotics and Automation*, 2013.
- [166] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," *ACM/IEEE International Conference on Human-Robot Interaction*, 2013.
- [167] B. Busch, J. Grizou, M. Lopes, and F. Stulp, "Learning legible motion from human-robot interactions," *International Journal of Social Robotics*, 2017.
- [168] S. Wallkötter, M. Chetouani, and G. Castellano, "A new approach to evaluating legibility: Comparing legibility frameworks using framework-independent robot motion trajectories," *arXiv preprint arXiv:2201.05765*, 2022.
- [169] D. Bortot, M. Born, and K. Bengler, "Directly or on detours? How should industrial robots approximate humans?" *ACM/IEEE International Conference on Human-Robot Interaction*, 2013.
- [170] F. Eyssel, D. Kuchenbrandt, and S. Bobinger, "Effects of anticipated human-robot interaction and predictability of robot behavior on perceptions of anthropomorphism," *ACM/IEEE International Conference on Human-Robot Interaction*, 2011.

- [171] F. Dehais, E. A. Sisbot, R. Alami, and M. Causse, "Physiological and subjective evaluation of a human-robot object hand-over task," *Applied Ergonomics*, 2011.
- [172] I. S. 7001-2021, "IEEE standard for transparency of autonomous systems," 2022.
- [173] Federal Aviation Administration, "Sec. 107.29 operation at night," *Code of Federal Regulations*, 2021.
- [174] J. Carpentier and P.-B. Wieber, "Recent progress in legged robots locomotion control," *Current Robotics Reports*, 2021.
- [175] D. Han, B. Mulyana, V. Stankovic, and S. Cheng, "A survey on deep reinforcement learning algorithms for robotic manipulation," *Sensors*, 2023.
- [176] P. T. Singamaneni, P. Bachiller-Burgos, L. J. Manso, A. Garrell, A. Sanfeliu, A. Spalanzani, and R. Alami, "A survey on socially aware robot navigation: Taxonomy and future challenges," *The International Journal of Robotics Research*, 2024.
- [177] Z. Liu, Q. Liu, W. Xu, L. Wang, and Z. Zhou, "Robot learning towards smart robotic manufacturing: A review," *Robotics and Computer-Integrated Manufacturing*, 2022.
- [178] Ítalo Renan da Costa Barros and T. P. Nascimento, "Robotic mobile fulfillment systems: A survey on recent developments and research opportunities," *Robotics and Autonomous Systems*, 2021.
- [179] M. Kyrarini, F. Lygerakis, A. Rajavenkatanarayanan, C. Sevastopoulos, H. R. Nambiappan, K. K. Chaitanya, A. R. Babu, J. Mathew, and F. Makedon, "A survey of robots in healthcare," *Technologies*, 2021.
- [180] Starship Technologies, "Over 5 million autonomous deliveries now completed! Thousands of Starship delivery robots are operating globally, every day." accessed: 2023-09-10. [Online]. Available: <https://www.starship.xyz/>

- [181] Kiwibot, "Adorable autonomous delivery robots," 2024, accessed: 2024-06-11. [Online]. Available: <https://kiwibot.com/>
- [182] Serve Robotics, "Why deliver 2 pound burritos in 2 ton cars?" 2024, accessed: 2024-06-11. [Online]. Available: <https://www.serverobotics.com/>
- [183] City of Boston, "Sidewalk delivery robots," 2024, accessed: 2024-06-11. [Online]. Available: <https://www.boston.gov/departments/new-urban-mechanics/sidewalk-delivery-robots>
- [184] Y. Du and T. V. Le, "Understanding social perception, interactions, and safety aspects of sidewalk delivery robots using sentiment analysis," *2405.00688*, 2024.
- [185] A. Pani, S. Mishra, M. Golias, and M. Figliozzi, "Evaluating public acceptance of autonomous delivery robots during covid-19 pandemic," *Transportation Research Part D: Transport and Environment*, 2020.
- [186] M. Rajaei, B. Echeverri, Z. Zuchowicz, K. Wiltfang, and J. F. Lucarelli, "Socioeconomic and racial disparities of sidewalk quality in a traditional rust belt city," *SSM - Population Health*, 2021.
- [187] The General Assembly of Pennsylvania, "Senate bill 1199," 2020.
- [188] City of Toronto, "Automated micro-utility devices - accessibility feedback," 2021.
- [189] City of Boston, "Food access and insecurity," 2024, accessed: 2024-06-11. [Online]. Available: <https://storymaps.arcgis.com/stories/956debbdf80c0492bbceeedff9f6a4bac>
- [190] M. P. Rabbitt, L. J. Hales, M. P. Burke, and A. Coleman-Jensen, "Statistical supplement to household food security in the united states in 2022," *Economic Research Service Administrative Publication*, 2022.
- [191] L. Murray, "Arlington gets \$780k grant to research drone & robot delivery from food banks," 2024, accessed: 2024-06-11. [Online]. Available: <https://dallasinnovates.com/arlington-gets-780k-grant-to-research-drone-robot-delivery-from-food-banks/>

- [192] L. Tach and M. Amorim, "Constrained, convenient, and symbolic consumption: Neighborhood food environments and economic coping strategies among the urban poor," *Journal of Urban Health*, 2015.
- [193] C. Antrum, A. Atoloye, O. Ajayi, D. Holter, D. Singerman, and K. C. Stowers, "Black and latina women's lived experiences with navigating neighborhood food swamps to find healthy food: A photovoice approach," *Journal of Human Behavior in the Social Environment*, 2024.
- [194] C. Daniel, "Is healthy eating too expensive?: How low-income parents evaluate the cost of food," *Social Science & Medicine*, 2020.
- [195] E. Pitt, D. Gallegos, T. Comans, C. Cameron, and L. Thornton, "Exploring the influence of local food environments on food behaviours: a systematic review of qualitative literature," *Public Health Nutrition*, 2017.
- [196] B. Caliskan, "Three different studies of the complexity of food access," Ph.D. dissertation, The Ohio State University, 2022.
- [197] S. N. Zenk, A. M. Odoms-Young, C. Dallas, E. Hardy, A. Watkins, J. Hoskins-Wroten, and L. Holland, "'you have to hunt for the fruits, the vegetables': environmental barriers and adaptive strategies to acquire food in a low-income african american neighborhood," *Health Education & Behavior*, 2011.
- [198] L. H. MacNell, "Exploring race, class, and food access across different geographic scales," Ph.D. dissertation, North Carolina State University, 2016.
- [199] S. Santarossa, A. B. Hill, A. R. Sitarik, M. Taylor, S. Hawkins, K. Scher, A. Sohaski, M. Baseer, R. Dombrowski, A. Plum, and C. L. Joseph, "Food insecurity in detroit: exploring the relationship between patient-reported food insecurity and proximity to healthful grocery stores," *Public Health Nutrition*, 2022.

- [200] C. Burns and A. Inglis, "Measuring food access in melbourne: Access to healthy and fast foods by car, bus and foot in an urban municipality in melbourne," *Health & Place*, 2007.
- [201] S. S. Madlala, J. Hill, E. Kunneke, T. Lopes, and M. Faber, "Adult food choices in association with the local retail food environment and food access in resource-poor communities: a scoping review," *BMC Public Health*, 2023.
- [202] N. Haletky, O. Taylor, J. Weidner, and S. Gerbing, "Urban agriculture as a solution to food insecurity: West oakland and people's grocery," *Urban Action*, 2006.
- [203] C. M. Pollard and S. Booth, "Food insecurity and hunger in rich countries—it is time for action against inequality," *International Journal of Environmental Research and Public Health*, 2019.
- [204] T. Waite, "What is the difference between a food bank and food pantry?" 2024, accessed: 2024-07-24. [Online]. Available: <https://www.feedingamerica.org/hunger-blog/what-difference-between-food-bank-and-food-pantry>
- [205] H. A. Eicher-Miller, "A review of the food security, diet and health outcomes of food pantry clients and the potential for their improvement through food pantry interventions in the united states," *Physiology & Behavior*, 2020.
- [206] R. An, J. Wang, J. Liu, J. Shen, E. Loehmer, and J. McCaffrey, "A systematic review of food pantry-based interventions in the usa," *Public Health Nutrition*, 2019.
- [207] Feeding America, "Our vision: An america without hunger," 2024, accessed: 2024-07-24. [Online]. Available: <https://www.feedingamerica.org/>
- [208] C. R. Long, M. M. Bailey, D. C. Cascante, R. S. Purvis, B. Rowland, B. M. S. Faitak, J. Gittelsohn, C. E. Caspi, E. S. English, and P. A. McElfish, "Food pantry clients' needs, preferences, and recommendations for food pantries: A qualitative study," *Journal of Hunger & Environmental Nutrition*, 2023.

- [209] C. E. Caspi, C. Davey, C. B. Barsness, N. Gordon, L. Bohen, M. Canterbury, H. Peterson, and R. Pratt, "Needs and preferences among food pantry clients," *Preventing Chronic Disease*, 2021.
- [210] L. R. Kicinski, "Characteristics of short and long-term food pantry users," *Michigan Sociological Review*, 2012.
- [211] B. O. Daponte, G. H. Lewis, S. Sanders, and L. Taylor, "Food pantry use among low-income households in allegheny county, pennsylvania," *Journal of Nutrition Education*, 1998.
- [212] J. S. Rochester, M. S. Nanney, and M. Story, "Assessing foodshelves' ability to distribute healthy foods to foodshelf clients," *Journal of hunger & environmental nutrition*, 2011.
- [213] N. Boysen, S. Fedtke, and S. Schwerdfeger, "Last-mile delivery concepts: a survey from an operational research perspective," *OR Spectrum*, 2021.
- [214] T. Hoffmann and G. Prause, "On the regulatory framework for last-mile delivery robots," *Machines*, 2018.
- [215] M. Hossain, "Autonomous delivery robots: A literature review," *IEEE Engineering Management Review*, 2023.
- [216] E. Alverhed, S. Hellgren, H. Isaksson, L. Olsson, H. Palmqvist, and J. Flodén, "Autonomous last-mile delivery robots: a literature review," *European Transport Research Review*, 2024.
- [217] State of Washington 66th Legislature, "Engrossed substitute house bill 1325," 2019.
- [218] J. E. Martinez, D. VanLeeuwen, B. B. Stringam, and M. R. Fraune, "Hey?! what did you think about that robot? groups polarize users' acceptance and trust of food delivery robots," *ACM/IEEE International Conference on Human-Robot Interaction*, 2023.
- [219] D. Weinberg, H. Dwyer, S. E. Fox, and N. Martelaro, "Sharing the sidewalk: Observing delivery robot interactions with pedestrians during a pilot in pittsburgh, pa," *Multimodal Technologies and Interaction*, 2023.

- [220] H. R. M. Pelikan, S. Reeves, and M. N. Cantarutti, "Encountering autonomous robots on public streets," *ACM/IEEE International Conference on Human-Robot Interaction*, 2024.
- [221] S. Sundermeir, N. Martin, L. Poirier, D. Reznar, Melissa nad Barnett, A. Uriarte, J. Stephenson, E. Lewis, M. Matsuzaki, and J. Gittelsohn, "Current use and demand for digital tools to enhance food pantry management: Findings from a nationwide survey," *Journal of Hunger & Environmental Nutrition*, 2023.
- [222] S. H. Evans and P. Clarke, "Training volunteers to run information technologies: A case study of effectiveness at community food pantries," *Nonprofit and Voluntary Sector Quarterly*, 2010.
- [223] N. M. Martin, D. J. Barnett, L. Poirier, S. M. Sundermeir, M. M. Reznar, and J. Gittelsohn, "Moving food assistance into the digital age: A scoping review," *International Journal of Environmental Research and Public Health*, 2022.
- [224] Berkshire Grey, "Picking with purpose is our commitment to helping those in need," 2024, accessed: 2024-07-24. [Online]. Available: <https://www.berkshiregrey.com/about-us/picking-with-purposes/>
- [225] R. Gordon, "Csail robot disinfects greater boston food bank," 2020, accessed: 2024-06-11. [Online]. Available: <https://news.mit.edu/2020/csail-robot-disinfects-greater-boston-food-bank-covid-19-0629>
- [226] M. Lok, "Devil's depot partners with DoorDash and Starship to delivery free food to ASU students," 2022, accessed: 2024-06-11. [Online]. Available: <https://www.statepress.com/article/2022/03/asu-devils-depot-partner-doordash-starship-food-pantry>
- [227] S. K. Shah and K. G. Corley, "Building better theory by bridging the quantitative–qualitative divide," *Journal of Management Studies*, 2006.

- [228] Margaret Fuller Neighborhood House, "Building our community together since 1902," 2024, accessed: 2024-06-14. [Online]. Available: <https://margaretfullerhouse.org/>
- [229] East End House, "Welcome to east end house," 2024, accessed: 2024-06-14. [Online]. Available: <https://eastendhouse.org/>
- [230] Instacart, "Order groceries for delivery or pickup today," 2024, accessed: 2024-06-14. [Online]. Available: <https://www.instacart.com/>
- [231] V. Braun and V. Clarke, *APA handbook of research methods in psychology, Vol. 2. Research designs: Quantitative, qualitative, neuropsychological, and biological*. American Psychological Association, 2012, ch. 4.
- [232] D. R. Thomas, "A general inductive approach for qualitative data analysis," *American Journal of Evaluation*, 2003.
- [233] M. F. Monteban, K. D. Bess, C. C. Walsh, H. Baily, S. A. Flocke, E. A. Borawski, and D. A. Freedman, "People and places shaping food procurement among recipients of supplemental nutrition assistance program (snap)," *Health & Place*, 2018.
- [234] S. J. Algert, M. Reibel, and M. J. Renvall, "Assessing foodshelves' ability to distribute healthy foods to foodshelf clients," *American Journal of Public Health*, 2006.
- [235] B. P. Library, "Computer & technology classes," 2024, accessed: 2024-07-21. [Online]. Available: <https://www.bpl.org/computer-classes/>
- [236] M. Awan, S. Ali, M. Ali, M. F. Abrar, H. Ullah, and D. Khan, "Usability barriers for elderly users in smartphone app usage: An analytical hierarchical process-based prioritization," *Scientific Programming*, 2021.
- [237] B. Juliá-Nehme and J. Rosell, "Interaction and design barriers for older adults in food delivery apps: A usability study," *International Journal of Human-Computer Interaction*, 2024.

- [238] C.-M. Chang, K. Toda, X. Gui, S. H. Seo, and T. Igarashi, "Can eyes on a car reduce traffic accidents?" *International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, 2022.
- [239] H. Knight and S. Buchmeier, "The video game to robot driver pipeline: Sociability with humans-in-the-loop," *2023 IEEE International Conference on Robot and Human Interactive Communication*, 2023.
- [240] S. S. Kannan, A. Lee, and B.-C. Min, "External human-machine interface on delivery robots: Expression of navigation intent of the robot," *IEEE International Conference on Robot and Human Interactive Communication*, 2021.
- [241] D. C. May, K. J. Holler, C. L. Bethel, L. Strawderman, D. W. Carruth, and J. M. Usher, "Survey of factors for the prediction of human comfort with a non-anthropomorphic robot in public spaces," *International Journal of Social Robotics*, 2017.
- [242] T. Nam, "Citizen attitudes about job replacement by robotic automation," *Futures*, 2019.
- [243] C.-M. Alcover, D. Guglielmi, M. Depolo, and G. Mazzetti, ""aging-and-tech job vulnerability"": A proposed framework on the dual impact of aging and ai, robotics, and automation among older workers," *Organizational Psychology Review*, 2021.
- [244] M. R. Calo, "Robots and privacy," in *Machine Ethics and Robot Ethics*. Routledge, 2020.
- [245] A. Howard and J. Borenstein, "The ugly truth about ourselves and our robot creations: The problem of bias and social inequity," *Science and Engineering Ethics*, 2018.
- [246] E. Rommes, "Gendered user-representations," *Women, Work and Computerization*, 2000.
- [247] M. A. DeVito, A. M. Walker, and J. R. Fernandez, "Values (mis)alignment: Exploring tensions between platform and lgbtq+ community design values," *Proceedings of the ACM on Human-Computer Interaction*, 2021.

- [248] Bloomberg Philanthropies, "State of cities: Generative ai in local government," 2023.
- [249] Electronic Frontier Foundation, "Atlas of surveillance: Documenting police tech in our communities with open source research," 2024, accessed: 2024-08-22. [Online]. Available: <https://atlasofsurveillance.org/>
- [250] A. Jarmanning, "Mass. state police tested out boston dynamics' spot the robot dog. civil liberties advocates want to know more," 2024, accessed: 2024-07-21. [Online]. Available: <https://www.wbur.org/news/2019/11/25/boston-dynamics-robot-dog-massachusetts-state-police>
- [251] D. Goldstein, "A.i. 'friend' for public school students falls flat," 2024, accessed: 2024-08-22. [Online]. Available: <https://www.nytimes.com/2024/07/01/us/ai-chatbot-los-angeles-schools.html>
- [252] J. Siegel, "'technology is not something we can hide from students': How nyc public schools invited ai into its classrooms," 2024, accessed: 2024-07-21. [Online]. Available: <https://news.microsoft.com/source/features/digital-transformation/how-nyc-public-schools-invited-ai-into-its-classrooms>
- [253] K. Seaborn, G. Barbareschi, and S. Chandra, "Not only WEIRD but "uncanny"? a systematic review of diversity in human-robot interaction research," *International Journal of Social Robotics*, 2023.
- [254] A. Guerra, F. Parisi, and D. Pi, "Liability for robots i: legal challenges," *Journal of Institutional Economics*, 2022.
- [255] D. Kerr, "California orders cruise driverless cars off the roads because of safety concerns," 2023, accessed: 2024-08-22. [Online]. Available: <https://www.npr.org/2023/10/24/1208287502/california-orders-cruise-driverless-cars-off-the-roads-because-of-safety-concern>

- [256] J. J. Payne, N. J. Kong, and A. M. Johnson, "The uncertainty aware salted kalman filter: State estimation for hybrid systems with uncertain guards," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022.
- [257] Future of Life Institute, "Autonomous weapons open letter: Ai & robotics researchers," 2016, accessed: 2024-08-12. [Online]. Available: <https://futureoflife.org/open-letter/open-letter-autonomous-weapons-ai-robotics/>