Mechanical Engineering
Carnegie Mellon University

Thesis Proposal

submitted for the degree of

Doctor of Philosophy in Mechanical Engineering

entitled

Navigating a Complex World:

Improving Robot Outcomes Through

Social, Regulatory, and Control Theoretic Approaches

Date of Submission: September 29, 2023

Submitted by: James Zhu

Advisor: Dr. Aaron M. Johnson

Department of Mechanical Engineering, Carnegie Mellon University

Committee Members: Dr. Mark Bedillion

Department of Mechanical Engineering, Carnegie Mellon University

Dr. Illah R. Nourbakhsh

The Robotics Institute, Carnegie Mellon University

Dr. Luis Sentis

Department of Aerospace Engineering and Engineering Mechanics,

University of Texas at Austin

Contents

List of Figures

List of Tables

Abstract

In recent years, many wheeled, bipedal, and quadrupedal robots have been released to market. Yet, robots have not had a meaningful presence in public spaces and have instead been mainly restricted to controlled environments like factories and warehouses. What is holding robots back from making wider impacts on our world? This thesis presents an interdisciplinary investigation into the difficulties that robots currently face by synthesizing approaches in control theory, social science, and public policy. Each of these often siloed fields hold insights that will be invaluable to the development of robots that are safe, trustworthy, and equitable.

The first part of this thesis presents methods that analyze the performance of legged robot locomotion and generate trajectories that are safe and robust when operating in uncertain environments. This work is necessary because legged robots are hybrid systems, meaning they undergo discontinuous changes in state and dynamics when their feet touchdown on the ground. These discontinuities violate assumptions key to many traditional control architectures. The presented hybrid systems analysis utilizes the fundamental solution matrix, which characterizes the evolution of initial errors through a trajectory. With this analysis, novel trajectory optimization methods are presented that explicitly reason about the stability and convergence of hybrid trajectories, leading to improved tracking performance for a variety of systems.

This thesis also presents work that investigates legal theory and community attitudes to develop frameworks for equitable robot design. Specifically, I focus on grounding robot design not just in stakeholder and customer preferences, but also in the needs of all community members regardless of their familiarity with robots. This work addresses how current self-defense law can inform robot design as well as how community attitudes toward robots can inform both robot design and future policies. These findings will allow robots to make positive contributions in our vast, human world.

Chapter 1

Introduction

1.1 Motivation

Despite decades of research in mobile robots, large scale deployment is still not a reality outside of tightly controlled factory and warehouse applications. Environmental hazards that humans can traverse such as ledges, ice, and mud can completely immobilize robots, making them unsuited for many human environments both urban and rural. Additionally, robots have encountered resistance from communities who express concerns about safety, privacy, and financial impacts. Policies and regulations addressing robot deployment have been developing rapidly as well, and roboticists must be engaged with how regulation can help equitably distribute benefits and mitigate harms. The obstacles to broad deployment of mobile robots are multi-faceted and require expertise in a variety of traditionally disparate fields. For robots to truly be successful, it is necessary for research to address the totality of social, regulatory, and control theoretic challenges.

1.2 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. Consider, for example, an urban environment like Pittsburgh where sidewalks are not ideally 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 to 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 requires considerable care.

Another contributing factor to the world's complexity is interacting with humans. People are complex creatures that are exceedingly difficult to model and predict. Yet, it is critical for any technology to gain the trust and acceptance of the communities it impacts. This is even more important for robots that intend to operate in public spaces like sidewalks, parks, and malls. Ensuring physical safety of humans is not enough. To achieve broad acceptance from people, robots need to behave in ways that promote trust, even among those who are predisposed to be the most resistant. Robots will need to overcome technical, historical, and cultural barriers to acceptance, which requires engineers to tailor robot design to account for diverse human attitudes. Regulation is an additional avenue that can enable robots to meaningfully contribute to society while preventing misuse, and roboticists have an opportunity to work alongside government partners and take on a major role in shaping these policies.

Although it is not possible to foresee every potential pitfall that robots will encounter, it is crucial for those involved in robotics development to be acutely aware of the consequences of robot failures. Inability to handle environmental uncertainties leads to damaged robots and loss of revenue, while even more crucially, failure to address human factors could lead to inequitable distribution of benefits and costs to society and threaten to reinforce longstanding injustices to historically 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?

Though impossible to answer comprehensively in the scope of a thesis, these questions underpin the approach to the presented topics of safe control for legged robots and community-centered robot design. Not only are these questions important, it is also essential they be addressed with urgency. Mass deployment of robots prior to a thorough understanding of these research questions leaves open the possibility of serious societal-level harms. The goal of this thesis is to push the robotics field forward so that we can provide wide-reaching, equitable benefits to our world.

1.3 Approach

This thesis is presented in two parts. Part I is titled: Navigating a Complex, Uncertain World and presents analysis and control methods that enable legged robots to more safely traverse environments with uncertain terrain and disturbances. Part II is titled: Navigating a Complex, Human World and investigates how social and legal structures are able to inform robot design. Taken all together, this thesis explores several facets of the wide-ranging demands that robots most fulfill in order to find success in the world.

1.3.1 Navigating a Complex, Uncertain World

To enable legged robots to be resilient to environmental uncertainties, Part I presents a line of research that develops a mathematical analysis of stability and convergence of hybrid systems, and from that a set of methods that generate trajectories for legged robots that are resilient to perturbations. The analysis, which draws from prior work extending traditional analysis methods to hybrid systems, derives two important scalar measures: the stability measure and convergence measure. The stability measure describes the asymptotic error behavior of a periodic hybrid trajectory and can be used to predict stability of a legged robot gait. The convergence measure is similar, but can also be applied to aperiodic systems and determines the worst-case evolution of errors across a trajectory. These scalar measures can be utilized in trajectory optimization frameworks to generate stable and robust robot motions. This work formalizes past results that could only be achieved by operator hand tuning and generates novel legged robot behaviors that previously have not been presented. Results have been demonstrated on a range of simulated robot systems and will be shown on quadrupedal robot hardware.

1.3.2 Navigating a Complex, Human World

Large corporations such as Amazon and FedEx have recently scaled back or shut down their robot delivery services, highlighting just how difficult it is for robots to succeed in human environments. One issue in human-aware robot navigation is that robot design requirements are not well defined or understood. What level of human comfort and trust is necessary for robots to be accepted in communities? How does robot design affect human comfort and trust, and how do these effects vary among diverse populations? These are the questions I approach in this second research thread, found in Part II. To accomplish this, I first looked to U.S. self-defense law to establish under what conditions a person may be justified in acting in self-defense against a robot. Of course, robots should not be designed to threaten humans, but a breakdown of trust in a human-robot interaction can still lead to a self-defense scenario. The law can inform roboticists about the minimum standards of comfort and trust needed to ensure safe interactions between robots and humans. Extending on this work, I will examine the impacts of certain robot design parameters on comfort and trust, particularly among diverse communities. The focus on diversity is critical for robots to have equitable impacts on society, since inequitable impacts have already been felt by women [2] and Black communities [3]. This project will draw from interviews with diverse residents of Pittsburgh to generate quantitative and qualitative data on how preferences to delivery robot design vary across gender and race.

1.4 Thesis Outline

Part I: Navigating a Complex, Uncertain World

Chapter 2: Stability and Convergence Analysis of Hybrid Systems (Completed)

Broadly speaking, the goal of robot control is to keep a robot's state close to its intended state. Variational equations describe the evolution of close-together trajectories, and can be used to analyze the success of robot controllers. Chapter 2 defines the mathematical basis of this analysis, presenting two terms, the stability and convergence measures, which will be utilized throughout the rest of this work. This chapter is partially based on the tutorial paper written in conjunction with Dr. Nathan J. Kong and J. Joe Payne, currently under review in the Proceedings of the IEEE [4].

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

Nearly all standard legged robot gaits are periodic, meaning that they repeat themselves after a certain amount of time. Past work has been able to identify behaviors that can improve the stability of periodic gaits, such as swing leg retraction [5]. However, this result has largely been observed only after a significant amount of operator hand tuning, and does not generalize to other possible stabilizing behaviors of legged robots. In Chapter 3, we present hybrid event shaping (HES), a method that can identify stabilizing behaviors based on the stability analysis of Chapter 2. Using HES, we autonomously generate open-loop stable trajectories for legged robots that reproduce known phenomenon like swing leg retraction along with introducing novel behaviors on more complex robot models. This work was published at ICRA 2022 [6].

Chapter 4: Convergent iLQR: Generating Closed-Loop Convergent Hybrid Trajectories (Completed)

Chapter 4 extends HES by reasoning about closed-loop trajectories. In practice, many robotic behaviors are never open-loop stable and require some feedback control. Standard feedback controllers like LQR require significant hand tuning and can struggle with hybrid, underactuated trajectories which are common for legged robots. One possible strategy for improving closed-loop performance is to increase feedback gains, but this may cause undesirable side effects like worsened robustness. Instead, this chapter presents convergent iLQR, a trajectory optimization method that improves the performance and robustness of a closed-loop hybrid trajectory without having to increase feedback gains. This method is demonstrated on a planar quadruped robot in simulation. This paper has been submitted to ICRA 2024 [7].

Chapter 5: Convergence-Based Gait Switching for Perilous Legged Locomotion (Proposed)

One place where convergent iLQR has difficulty is switching between different legged robot gaits. This is due to the local nature of trajectory optimization solvers, which will always struggle to find solutions that are dissimilar to the initially provided guess. Even though methods have been developed that can switch legged gaits [8], it is very difficult for any trajectory optimization method to reliably find a new and improved gait because standard quadrupedal gaits can be quite distant from each other (e.g. finding a walk from a trot). Past work has investigating switching gaits based on energy efficiency, but an alternative reason a robot may want to switch gaits is convergence. Chapter 5 proposes a method to reason about when to switch between predefined gaits based on the difficulty of convergent locomotion for a given terrain map. The results of this project aim to show improved safety navigating narrow beams, corridors, and trails. This work is currently underway and will be submitted to ICRA 2025.

Part II: Navigating a Complex, Human World

Chapter 6: Self-Defense Against Robots: Design and Deployment Considerations for Roboticists (Completed)

No matter how well our robots can maneuver about physical obstacles, their utility will be severely limited if they can not negotiate interactions with humans. If people do not reasonably trust or feel comfortable around robots, the robots will fail to accomplish their goals. Chapter 6 discusses a legal interpretation of exactly 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 these scenarios from occurring, we provide practitioners with actionable steps that will improve the public's trust in robots and work toward understanding the connection between robot design and human attitudes. This work was presented at RO-MAN 2023 [9].

Chapter 7: Community-Focused Design For Sidewalk Delivery Robots (Proposed)

A key implication of the self-defense against robots work is that the law is grounded in justifying "reasonable" human actions. While there is no strict definition of reasonableness, any attempt to understand this legal standard should be based on the people that are least likely to trust robots. Thus, it is critical to introduce diverse perspectives into our understanding of human-robot interaction. For robots operating in public spaces, however, studies have focused solely on evaluating perspectives of potential customers, which are likely biased toward those with an affinity for robots [10]. Considering most pedestrians that a delivery robots would encounter on a sidewalk will not be customers, exploration of more diverse populations is necessary. In Chapter 7, I propose a two-phased interview methodology inspired by value sensitive design [11] to systematically extract the core values that diverse people have toward sidewalk delivery robots. Phase 1 of this work will consist of extracting qualitative data from semi-structured interviews, and the preliminary results will be presented at a RO-MAN 2024 workshop on equitable human-robot interaction. Phase 2 will analyze quantitative measurements of human perceived safety based on viewing videos of a realistic interaction with a delivery robot simulated in the Unity engine. This conclusive work will be submitted to HRI 2025.

1.5 Thesis Defense

I will complete all proposed work and defend my thesis by December 2024. During my defense, I will present new results for both the gait switching and community-focused design projects. The results of the gait switching project will demonstrate a physical quadrupedal robot navigating narrow beams and other narrow corridors, something only shown so far with the addition of reaction wheel actuation. I will also present the analysis of the delivery robot attitudes interviews, which will include the significance of various robot design parameters on human perception of comfort, trust, and perceived safety, as well as generalized utility functions regarding optimal zones of robot parameters.

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 [12,13]. Following [4], 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. $\mathcal{J} := \{\mathrm{I}, \mathrm{J}, ...\} \subset \mathbb{N}$ 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. $\mathcal{D} := \amalg_{\mathrm{I}\in\mathcal{J}}\, D_\mathrm{I}$ is the collection of domains, where D_I is a C^r manifold and the state $x \in D_\mathrm{I}$ while in mode I.

4. $\mathcal{F} := \amalg_{\mathrm{I}\in\mathcal{J}} F_\mathrm{I}$ is a collection of C^r time-varying vector fields, $F_\mathrm{I} : \mathbb{R} \times D_\mathrm{I} \to \mathcal{T} D_\mathrm{I}$.

5. $\mathcal{G} := \amalg_{(\mathrm{I},\mathrm{J})\in\Gamma}\, G_{(\mathrm{I},\mathrm{J})}(t)$ is the collection of guard sets, where $G_{(\mathrm{I},\mathrm{J})}(t) \subseteq D_\mathrm{I}$ for each $(\mathrm{I},\mathrm{J}) \in \Gamma$ is defined as a regular sublevel set of a C^r guard function, i.e. $G_{(\mathrm{I},\mathrm{J})}(t) = \{x \in D_\mathrm{I} | g_{(\mathrm{I},\mathrm{J})} (t, x) \leq 0\}$ and $\mathrm{D}_x g_{(\mathrm{I},\mathrm{J})}(t, x) \neq 0 \;\forall\; g_{(\mathrm{I},\mathrm{J})}(t, x) = 0$.

6. $\mathcal{R} : \mathbb{R}\times\mathcal{G} \to \mathcal{D}$ is a C^r map called the reset that restricts as $R_{(\mathrm{I},\mathrm{J})} := \mathcal{R}|_{G_{(\mathrm{I},\mathrm{J})(t)}} : G_{(\mathrm{I},\mathrm{J})}(t) \to D_\mathrm{J}$ for each $(\mathrm{I},\mathrm{J}) \in \Gamma$.

An execution of a hybrid system [14] begins at an initial state $x_0 \in D_\mathrm{I}$. With input $u_\mathrm{I}(t, x)$, the system obeys the dynamics F_I on D_I. If the system reaches guard surface $G_{(\mathrm{I},\mathrm{J})}$, the reset map $R_{(\mathrm{I},\mathrm{J})}$ 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 [15], branching [16], and Zeno phenomena where infinite transitions occur in finite time [17]. Following prior literature [18–20], we assume these behaviors do not occur, such that guard surfaces are isolated and intersected transversely [13,14] 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 [21]. In each continuous domain, the linearized variational equation is discretized from timestep i to $i + 1$ and is:

$$\delta x_{i+1} \approx (A_\mathrm{I} - B_\mathrm{I} K_\mathrm{I})\delta x_i \tag{2.1}$$

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 [21]. For hybrid events, the analogous variational term is the saltation matrix $\Xi_{(\mathrm{I},\mathrm{J})}$, 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 $\delta x(t^+)$ [19], such that:

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

The formulation of the saltation matrix is:

$$\Xi_{(\mathrm{I},\mathrm{J})} := \mathrm{D}_x R^- + \frac{\left(F_\mathrm{J}^+ - \mathrm{D}_x R^- F_\mathrm{I}^- - \mathrm{D}_t R^-\right) \mathrm{D}_x g^-}{\mathrm{D}_t g^- + \mathrm{D}_x g^- F_\mathrm{I}^-} \tag{2.3}$$

where

$$F_\mathrm{I}^- := F_\mathrm{I}(t^-, x(t^-)) \tag{2.4}$$

$$F_\mathrm{J}^+ := F_\mathrm{J}(t^+, x(t^+)) \tag{2.5}$$

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

$$\mathrm{D}_x R^- := \mathrm{D}_x R_{(\mathrm{I},\mathrm{J})}(t^-, x(t^-)) \tag{2.7}$$

$$\mathrm{D}_t R^- := \mathrm{D}_t R_{(\mathrm{I},\mathrm{J})}(t^-, x(t^-)) \tag{2.8}$$

$$\mathrm{D}_x g^- := \mathrm{D}_x g_{(\mathrm{I},\mathrm{J})}(t^-, x(t^-)) \tag{2.9}$$

$$\mathrm{D}_t g^- := \mathrm{D}_t g_{(\mathrm{I},\mathrm{J})}(t^-, x(t^-)) \tag{2.10}$$

More information on the saltation matrix and a rigorous derivation can be found in [4].

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 [20], the fundamental solution matrix represents the transformation of error from the initial state to final state:

$$\delta x_f \approx \Phi \delta x_0 \tag{2.11}$$

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 [6]. For a hybrid trajectory with N domains, the fundamental solution matrix can be formulated as:

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

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 [20,22]. 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, $\max(|\lambda|)$ [20]. 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 [20]. 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:

$$\chi = ||\Phi||_2 \tag{2.13}$$

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 [8,23]. One strategy for hybrid control is to cancel out the effects of hybrid events by working with an invariant subsystem [24–26]. 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 [5,27,28]. For example, [27] 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 [26,29] generated open-loop swing leg trajectories that produced deadbeat hopping of a SLIP-like system. Each of these

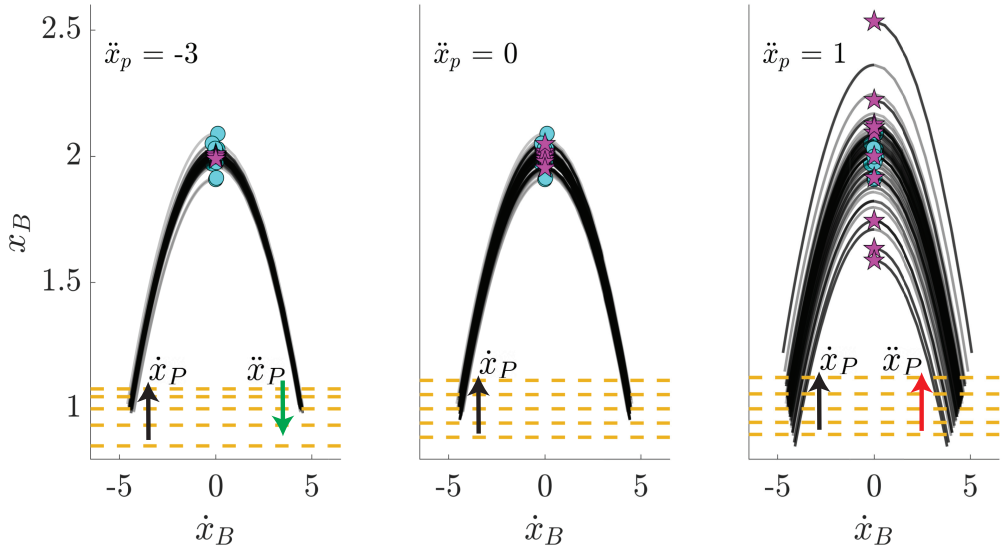

Figure 3.1: The paddle juggler system [27] has no control authority while the ball is in the air. The paddle acceleration at impact determines the convergence/divergence of the system from initial points (cyan dots) to the final states (magenta stars) after 5 cycles. This example underscores how hybrid 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 [5,27] 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 ($\mathrm{D}_x g$ and $\mathrm{D}_t g$) and reset maps ($\mathrm{D}_x R$ and $\mathrm{D}_t R$) 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 touch-down, 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 [5].

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 [30]. 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 [31]. For a diagonalizable Φ, the derivative of the eigenvalues with respect to h_n can be computed in closed form [32]. Assume that matrix $\Phi(h_n)$ has simple (non-repeating) eigenvalues, $\lambda_1, \ldots, \lambda_N$, and let $\boldsymbol{j}_i$ and $\boldsymbol{k}_i$ denote the left and right eigenvectors associated with λ_i. Then the derivative $\frac{d\lambda_i}{dh_n}$ is:

$$\frac{d\lambda_i}{dh_n} = \boldsymbol{k}_i' \frac{d\Phi}{dh_n} \boldsymbol{j}_i \tag{3.1}$$

For matrices with eigenvalues that repeat, the derivatives of the repeated eigenvalues can be calculated similarly with a matrix of associated eigenvectors [32].

$\frac{d\Phi}{dh_n}$ 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:

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

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 [33]. 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 [34] 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,

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

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 [35]. For dynamical systems, these costs are generally functions of state and inputs, with constraints imposed on system dynamics and any other physical limits [36]. For specific problems, other aspects of the system may be added into the cost or constraint functions such as design parameters or minimizing time [37,38]. 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 [39]. 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 [35,36,40].

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

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 [27] and swing leg retraction [5], 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 [27], 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:

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

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: $A_1(T) = A_2(T) = \left[\begin{smallmatrix} 1 & T/2 \\ 0 & 1 \end{smallmatrix}\right]$ 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:

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

This is confirmed in [27], 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.

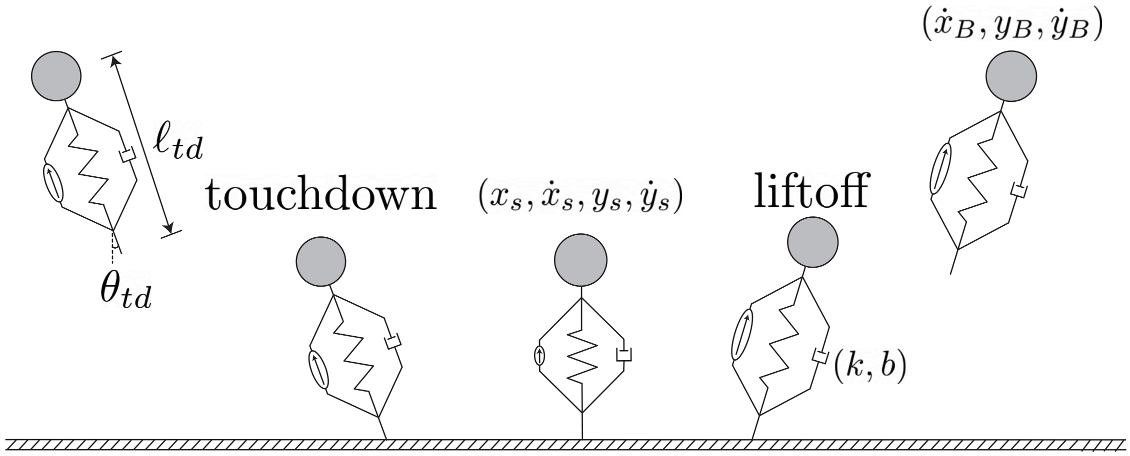

Figure 3.2: SLIP-like system with actuator and damper in parallel.

3.3.2 2D Hopper

The spring-loaded inverted pendulum (SLIP) is a popular model for dynamic legged locomotion [41–43]. This simple hopper model is effective at capturing dynamic properties of animal and robot locomotion [44] and has been used as a test bed for hybrid controllers [45].

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 $\overline{\theta}_{td}$ 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 $\overline{\dot{x}}$ 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.

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

2D Hopper HES Results

For a chosen initial apex height of 2.5 with a forward velocity of 2, ℓ_{td} and $\overline{\theta}_{td}$ 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.

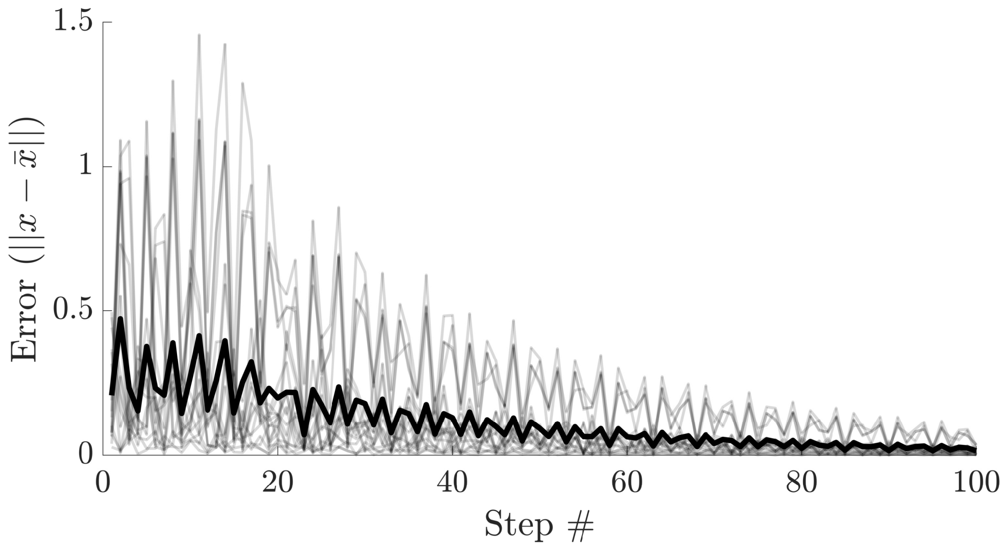

Figure 3.3: Error of perturbed initial states for the 2D hopper asymptotically decrease to zero. Transparent lines represent each of the 20 trials, while the bold line represents average error at each step. Convergence is neither monotonic nor very fast, but this is expected with asymptotic stability.

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 $\bar{x}$) 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 [46], 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 [47].

Another stabilizing property of legged locomotion that has been studied is swing-leg retraction [5,30]. It was noted in [5] 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 [5] and [46]. While a formal equivalency is yet to be proven, this is significant because the HES shape 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 [48] 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 [48].

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

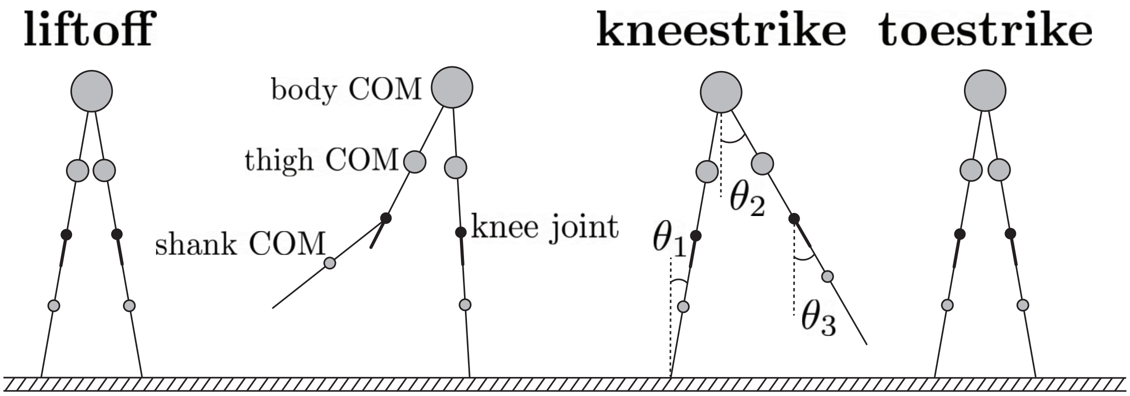

Figure 3.4: Biped walker system with kneestrike and toestrike hybrid events.

toestrike are also computed in [48].

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

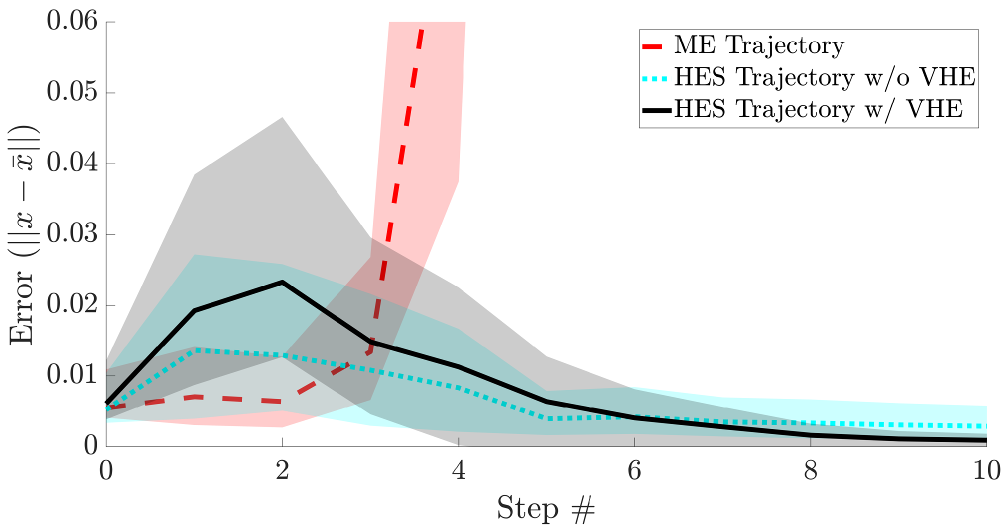

Figure 3.5: Error of perturbed Minimum Energy (ME), Hybrid Event Shaping without virtual hybrid events (HES w/o VHE) and HES w/ VHE trajectories over several steps. Center lines show the average error at each step and shaded regions indicate ± 1 standard deviation. ME trajectories become highly divergent 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 [49–52]. 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 [12,13], and it is underactuated if there exists a direction of acceleration in state space that can not be commanded by any valid input [53, 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,

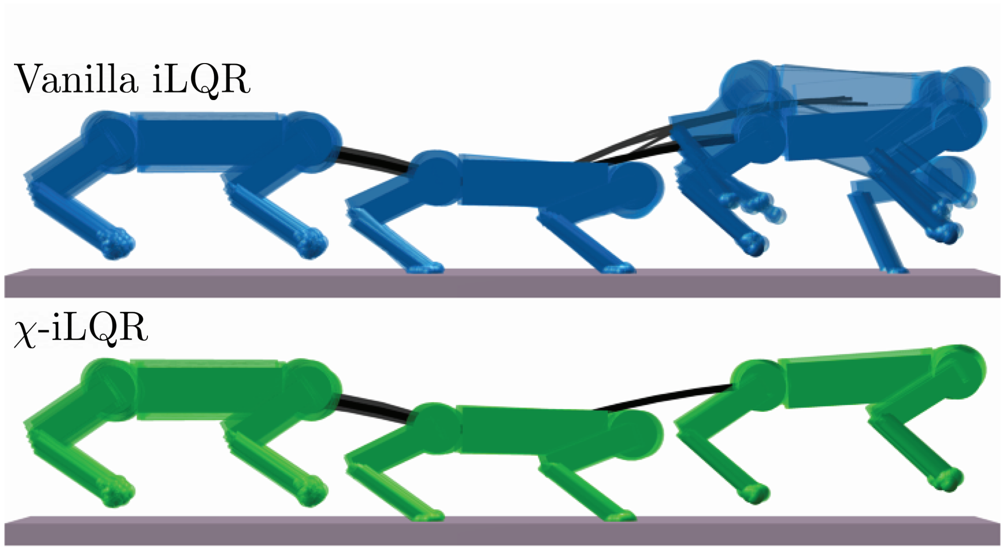

Figure 4.1: Similar quadruped gaits tracked with equivalent LQR controllers display enormous differences in final tracking performance. Results for 50 paired trials of trajectories sampled from an initial error covariance of 10^{-2} in all directions are shown. The top blue trajectory was generated with 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 [4] 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 [54, 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 paper 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 [51,55]. 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 [56–58]. These methods plan trajectories under the anticipation of some worst-case or average-case disturbance sequence. For example, [56] 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 [59,60]. Other approaches present trajectory optimization algorithms for legged robots over uncertain terrain [61] and compute a forward reachable set to bound closed-loop errors [62]. 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 [63,64]. 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 [33]. While effective, Lyapunov functions can be difficult to compute, particularly for hybrid systems, and can require methods such as sum-of-squares [65] or machine learning [66] 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 [67], but has the same drawback as Lyapunov functions.

A different strategy to analyze the stability and convergence of trajectories is contraction analysis [68], 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 [68]. Contraction analysis has been incorporated into path planning and trajectory optimization algorithms on smooth systems [69,70], but applying contraction analysis to hybrid systems is difficult because many mechanical hybrid systems are not contractive at hybrid events [19]. [6] 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 [71]. iLQR is convenient because compared to other trajectory optimization methods like direct collocation [35], 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 [8,72].

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

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

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

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

where $\ell_i(x_i, u_i)$ and $\ell_N(x_N)$ represent the nonlinear stage cost and terminal cost, respectively, $X := \{x_0, x_1, ..., x_N\}$ is a sequence of states with $x_i \in \mathbb{R}^n$ the system state at timestep i and $U := \{u_0, u_1, ..., u_{N-1}\}$ 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 $M := \{D_0, D_1, ..., D_N\}$ 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:

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

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 [8,72].

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:

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

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 [73], the derivative of χ with respect to the state at timestep i is:

$$\frac{\partial \chi}{\partial x_i} = u_\chi^T \frac{\partial \Phi}{\partial x_i} v_\chi \tag{4.6}$$

$\frac{\partial \Phi}{\partial x_i}$ 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:

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

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

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:

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

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_χ, Q_N, Q_i, R_i, n_iterations
X, U, M, J ← Rollout(U)
K_t, J_χ, P ← TrackingBP(X, U, M, J)
O ← ComputeO(X, U, M, K_t)
for i ← 1 to n_iterations do
    k_s, K_s ← SearchBP(X, U, M, J_χ, O)
    repeat
        X, U, M, J, O ← ForwardPass(X, U, M, k_s, K_s)
        K_t, J_χ, P ← TrackingBP(X, U, M, J)
    until LineSearchIsSatisfied(J_χ)
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 [74] 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 [75] 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 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 $V := \{v_0, v_1, ..., v_{N-1}\}$. 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:

$$E = \frac{||\delta x_f||_2}{||\delta x_0||_2} \tag{4.10}$$

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.

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

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 are constrained to 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 [76] that increase the energy efficiency of legged locomotion [77,78]. Due to the resonance of the natural spring dynamics, controlling these systems requires special care [79,80]. For example, [81] 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 N · m · 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. 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. Fig. 4.2 displays a histogram of the error ratio for each of the trials, and shows the improvement made by the convergent trajectory.

As the magnitude of initial errors grows, the performance of the LQR tracking controller be-

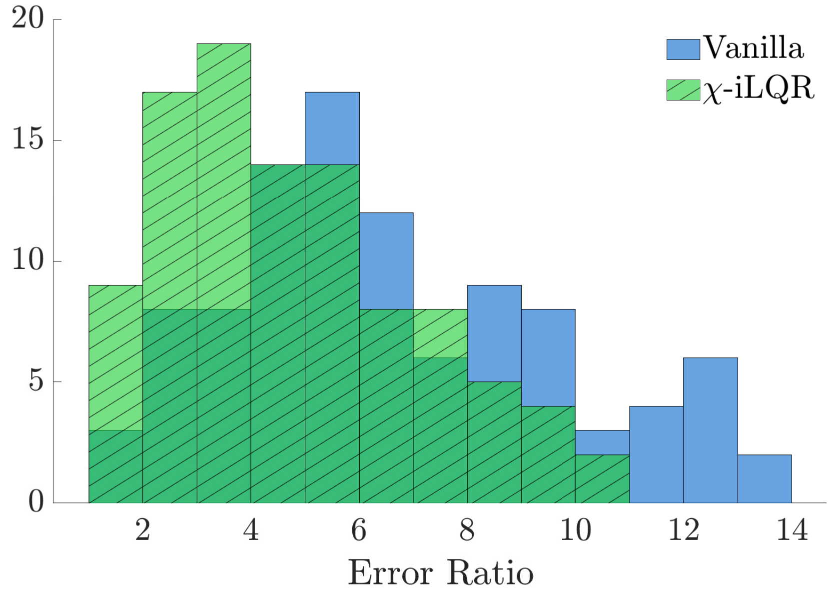

Figure 4.2: Histogram of error ratio for 100 paired simulated trials of the quadruped model with small initial perturbations. Error ratio is the 2-norm of final errors divided by the 2-norm of initial errors.

comes 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 non-linearities.

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 convergence. Even so, this work can be valuable to ensure that between iterations of a global footstep

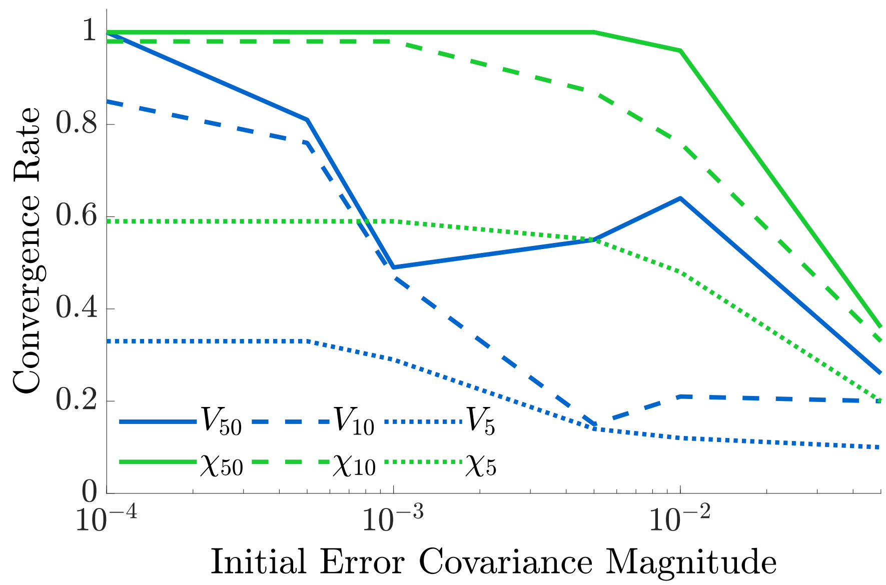

Figure 4.3: Plots showing success of LQR controllers at tracking vanilla (V) and convergent (χ) trajectories over various initial perturbation covariances. V_{50} and χ_{50} indicate proportion of trials where error ratio was below 50, V_{10} and χ_{10} below 10, and V_5 and χ_5 below 5.

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

Convergence-Based Gait Switching for Perilous Legged Locomotion

5.1 Introduction

Legged robots have the potential to outperform their wheeled counterparts in extreme terrain because foot touchdowns allow them to exert forces on the body to recover from any errors or disturbances. However, one major difficulty controlling legged robots is selecting an appropriate sequence of touchdowns even in relatively simple situations. Some legged control algorithms take in a predefined contact sequence from the user [82,83], but this becomes untenable for environments with even minor complexity and uncertainty. Other methods, such as convergent iLQR in the previous chapter, overcome this by using contact-implicit strategies that can alter the contact sequence to search for more optimal trajectories [8,84]. However, even these methods encounter an obstacle that is nigh unavoidable in optimization: local minima. In short, any optimization method is severely limited if it is necessary for performance to get worse before it gets better. This is an issue for legged robots with more than two legs because most standard legged gaits are far away from each other in terms of contact sequences. For example, consider a quadruped robot. In order to continuously transition between a walking gait (where each leg lifts and touches down in series) and a trotting gait (where diagonally opposed legs move in tandem), an optimizer must go through a several suboptimal degenerate gaits with poor performance. This difficulty prevents convergent iLQR from reliably taking advantage of gait changes during its search for convergent trajectories.

To extend the work in Chapters 3 and 4, in this chapter, I propose a high level gait planning scheme that utilizes the convergence measure to intelligently choose quadrupedal gaits based on a given terrain map. This proposed method can have great utility in traversing dangerous environments like walking across narrow beams or squeezing through gaps in obstacles, where the robot's feet are forced to be placed close to the center of the body, limiting its ability to apply lateral recovery forces and yielding a high degree of underactuation. Enabling robots to intelligently adapt their behavior to the complexity of the environment will be very useful in environments that may have sections that are easily traversible, so that the robot can focus on optimizing its gait for speed or energy efficiency, and other sections that are difficult to traverse, where slower and more careful gaits are necessary to achieve safe locomotion.

5.2 Prior Work

Beam walking is a challenging problem because it forces shrinkage of the support polygon, which defines an area where the center of mass is easily stabilized. Having a large support polygon is an assumption utilized by many common legged robot controllers, like the zero-moment point method [85]. Several works have examined this situation such as balancing a quadruped robot with only two feet in contact with the ground [86,87] with [87] also applying this strategy to trotting across a narrow beam. [63] was the first to demonstrate beam walking on hardware by adding a reaction wheel to apply additional stabilizing forces. However, these works did not consider when a robot can reason about the impact of different gaits on performance.

Other works have analyzed strategic gait switching, but were based on velocity to optimize energy efficiency [88]. [89] noted that in order switch between standard gaits, a transient middle gait had to be designed. A method to switch between gaits to improve robustness was demonstrated for a six-legged robot in [90]. These works inspire us to design a method to switch gaits based on convergence for quadruped robots.

5.3 Gaits and Contact Sequences

Legged gaits have been studied extensively on animals, with dogs and horses being common quadrupedal subjects of study [91]. Here, it was observed that quadrupeds tend to employ a few primary gaits, mainly dependent on desired locomotion speed. Naturally, these gaits exhibit some variation across animals, so their definitions are quite broad. For example, a quadrupedal trot consists of diagonally opposed pairs of legs moving in tandem. This gait is also a common default gait for quadruped robots [83,92]. However, there can be variation in contact sequence among just trot gaits, as the timing between leg pairs can lead to one pair lifting off exactly as the other pair touches down (perfect trot), liftoff occurring prior to touchdown (aerial trot), or liftoff occurring after touchdown (stance trot). Due to this inexactness in definition, in this work we will categorize quadrupedal locomotion methods based on contact sequence.

In this work, we focus on periodic gaits that consist of each foot lifting off and touching down once per period. This means we can define gait sequences as a series of eight transitions. Let each transition be notated by one letter, L for liftoff and T for touchdown, and one number, 1–4 for the front left, back left, front right, and back right feet, respectively. So the front right foot lifting off is notated as $L3$ and the back left foot touching down is $T2$. A contact sequence is defined as a tuple of minimum length 2 and maximum length 8. Each element of the contact sequence tuple can contain multiple transitions if they occur simultaneously, such as a trot. We define the convention that each sequence tuple begins with $L1$. So the three possible trot sequences are:

$$\text{aerial trot: } (L1L4, T2T3, L2L3, T1T4) \tag{5.1}$$

$$\text{perfect trot: } (L1T2T3L4, T1L2L3T4) \tag{5.2}$$

$$\text{stance trot: } (L1L4, T1T4, L2L3, T2T3) \tag{5.3}$$

5.4 Methods

5.4.1 Contact Sequence Convergence Analysis

During our investigation in Chapter 4, we observed a correlation between system actuation and convergence. For legged robots, actuation is increased when more feet are on the ground. The duty factor of a legged gait describes the percentage of the gait period that a foot is on the ground. Certain gaits naturally have high duty factors, such as the quadrupedal walk which always has a duty factor of at least 0.75 [91]. In this study, we would like to investigate the relationship between duty factor and convergence. If our findings support our hypothesis of a correlation between the two values, we will develop a gait planning method that will detect if a trajectory with a nominal contact sequence is insufficiently convergent, and then switch to a higher duty factor gait that will improve the trajectory convergence. However, it is not just duty factor that might affect the convergence of a trajectory. Consider how a walk and a pronk, $(L1L2L3L4, T1T2T3T4)$, can both have duty factor 0.75, but the entirety of the walk is spent with 3 feet on the ground, whereas the pronk alternates between 4 feet on the ground and a full aerial phase. We hypothesize that the time the robot spends in various underactuated domains will correlate with convergence and that a walk will be more convergent than a pronk due to avoiding the severely underactuated aerial domain.

So far, we have conducted a preliminary experiment to address the second hypothesis with the planar quadruped system from Chapter 4. For this system the front feet are set to move in tandem and the back feet likewise, such that front feet transitions are labelled with 1 and

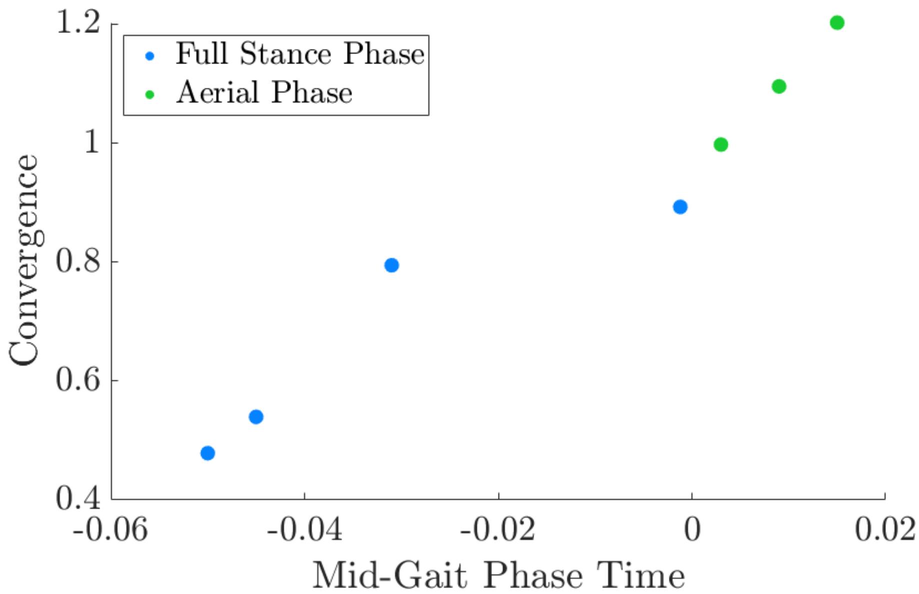

Figure 5.1: Convergence results for planar quadruped bound trajectories with different mid-gait phase times. Positive values indicate a contact sequence $(L1, T2, L2, T1)$ with the x-axis value being the length of the aerial phase (e.g. the point $(0.015, 1.20)$ spends 0.015 seconds in the aerial phase). Negative x-axis values indicate a contact sequence $(L1, L2, T1, T2)$ with the length of the full stance phase the inverse of the x-axis value (e.g. the point $(-0.05, 0.48)$ spends 0.05 seconds in the full stance phase).

the back feet with 2. Figure 5.1 shows the results of 7 different planar quadruped trajectories whose contact sequences were prespecified. 3 of the trajectories were given a contact sequence $(L1, T2, L2, T1)$, with an aerial phase occurring between front liftoff and back touchdown. The other 4 trajectories have contact sequence $(L1, L2, T1, T2)$ with a full stance phase between front touchdown and back touchdown. We compared these trajectories based on the time spent in these mid-gait phases - either the aerial phase or full stance phase. Our hypothesis was that the full stance trajectories would overall be more convergent than the aerial phase trajectories, due to the significant advantage in actuation in this domain. Our results supported this hypothesis and also showed that among the full stance trajectories, the longer the full stance phase was, the better convergence was, with vice versa for the aerial trajectories. These results appear promising, but we also note that the planar bound gaits are close to each other in the sense that it is possible to transition smoothly between an aerial bound and a stance bound. Due to this,

Algorithm 2 Proposed Gait Selection Algorithm

Initialize C_0, χ_{max}, $GaitsExhausted = False$
$X, U, K, \chi \leftarrow$ TrajOpt(C_0)
while $\chi > \chi_{\text{max}}$ $\&$ $GaitsExhausted = False$ **do**
    $C \leftarrow$ GaitStep(C_0)
    $X, U, K, \chi \leftarrow$ TrajOpt(C)
    $GaitsExhausted \leftarrow$ CheckGaitsExhausted
return $X, U, K, \chi, GaitsExhausted$

we aim to explore far apart walking and trotting gaits on a full quadruped model and examine if these patterns hold.

5.4.2 Gait Selection Algorithm

Based on the results of the convergence analysis, we plan to design a library of quadrupedal gaits that our novel algorithm can choose to switch between when improved convergence is necessary. The idea for this algorithm is to begin with some nominal contact sequence C_0, like a perfect trot that is common for quadruped robots, and generate a trajectory along some time horizon with that contact sequence. If that contact sequence satisfies a maximum allowable convergence measure, χ_{max} then the algorithm does not need to make any changes. Otherwise, if the trajectory is too divergent, then the algorithm will switch to a contact sequence that, based on the convergence analysis, will likely be more convergent. The new trajectory will be generated and the convergence will be checked. This process will repeat until a suitable trajectory is found or all gaits are checked without success. In this case, the environment may be too dangerous to traverse. See Algorithm 2 for a broad outline of the proposed method. There will most likely need to be additional intricacies to the method to be effective. One detail is that switching between gaits of different contact sequences is not trivial, and the transition must be executed safely. [93] addressed the problem of gait transitions which may provide useful here too. Prior work has found that velocity is an important factor to the performance of a gait [88], so our algorithm may need to modulate the trajectories desired velocity in order for a new gait to be appropriate.

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 [94], security [95], and personal assistance [96]. 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 [97], college students vandalizing a robot [98], or a robot crashing into a car [99]. 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 paper 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 paper 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 paper 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 paper 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 [100]. 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 [101,102].

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 [103]. 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 [104] or legitimate verbal threats of violence that occurred directly prior [105] can provide justification for acting in self-defense.

The second aspect of this definition is “imminent physical harm” [106]. 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 [106]. 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 [107], American common law has yet to grant this argument [108].

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 [109]. 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 [110].

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 [107], 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 justification of proportionality for self-defense against robots, which will empower people with the understanding that they will have at least some immediate recourse if they feel physically under threat by a robot. In turn, this transparency from the robotics field on the rights that people have when interacting with robots can engender a greater sense of trust and openness to 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 a 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 [111].

Following the literature [112], 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 [113]. 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, [114] 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 [115], 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 [116]. 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 [3]. 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 [3,117], calling the robot "another danger for Black & Latino residents" [117] and expressing fear toward the futuristic appearance of the robots [118]. Requests to police departments for more information on the purpose of the robots were not always met [3,119].

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 [116]. Another study suggests that people may be more likely to support robots doing jobs that require less experience and communication [120]. 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 [121]. Police robots have often been deployed to patrol low-income, Black neighborhoods [3,122], while women have repeatedly been targets of unwelcome surveillance by drones [2,123]. Additionally, [124] found that women tended to be less receptive to the concept of patrolling police robots than men. To combat this inequity, [3] 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 [125,126]. 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 [127]. Other work has developed robotic behaviors to satisfy desired outcomes such as visibility [128], active communication [129], and following social norms [113,130]. 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 [131], 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?) [132] and explanations that satisfy user-defined preferences [133]. 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 [134].

Drawing from the existing discourse on explainability in the field of autonomous vehicles (AVs), [135] 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, [136] 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 [137,138]. With the European Union's recent General Data Protection Regulation (GDPR) outlining a person's legal "right to explanation" when encountering autonomous agents [139], 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 [140], 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 [141], mimicking human behavior [142], moving quickly toward goals [143], and staying in people's field of view [128].

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 [144,145]. Others evaluated human performance of an unrelated task while a robot navigated around them [146]. Others still used questionnaires to gauge how well subjects felt they understood a robot's intentions [141], while [143] proposed a numerical measure of a trajectory's legibility.

Studies have found that legibility is correlated with increased feelings of safety, comfort, and acceptance [146–148]. 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 communicate to people [149]. 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 [107,123] and requiring the deployment of noticeable robots instead of silent, stealthy ones [114]. Another regulation that could be considered is how robots should operate at night. The Federal Aviation Administration currently restricts flying drones at night [107,150] 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

Community-Focused Design For Sidewalk Delivery Robots

7.1 Introduction

The previous chapter discussed how the standard of reasonableness is key to applying self-defense principles to human-robot interactions. It also argued that reasonableness must be defined with respect to people who are least likely to trust robots: non-expert non-users. Additionally, marginalized groups may be particularly distrusting of novel technologies based historic experiences of misuse [151]. Policy makers bear the responsibility of protecting their constituents and will be interested in regulating sidewalk robots on the basis of ensuring both physical human safety as well as perceived safety and comfort. There has already been significant recent movement in the United States and abroad to pass ordinances regulating aspects of sidewalk robot deployment. The state of Pennsylvania has limited robots operating on public sidewalks to be a maximum of 550 pounds unloaded and travel a maximum of 12 miles per hour [152]. In Washington state, those values are 120 pounds and 6 miles per hour, respectively [153]. Elsewhere, major cities such as San Francisco and Toronto have taken a different approach, essentially banning all sidewalk robot operation [154,155], see Table 7.1. Clearly, each of these municipalities

Table 7.1: Sample of current regulations of sidewalk robots in North America demonstrate tremendous variance.

Munincipality	Max Weight (lb)	Max Speed (mph)
Pennsylvania	550	12
Washington	120	6
San Francisco	Heavily Restricted	
Toronto	Banned	

have arrived at different conclusions to the trade off between encouraging innovation and ensuring pedestrian welfare. In this work, I propose investigating people's attitudes toward delivery robots to understand upper limits to pedestrian comfort around sidewalk robots. Specifically, we address two research questions:

- What aspects of robot design are most salient to pedestrian attitudes?
- Where along these design parameters do attitudes broadly deteriorate?

This work builds upon the value sensitive design framework, which centers technology design on stakeholder values [11]. This method has been used to encode ethical principles into projects such as aerial drones [156] and healthcare robots [157].

7.2 Methods

To examine these two research questions, we propose a two-phased approach. First will be an exploratory phase where we will conduct semi-structured interviews to extract factors of current delivery robot design that participants assess to be beneficial or detrimental to their trust and perceived safety. Phase 2 will consist of showing participants simulated videos of a sidewalk robot passing an avatar human. Design and movement parameters of the robot will be modulated and a modified Godspeed questionnaire [158] will evaluate the effect of each parameter on participant perceived safety. Quantitative analysis drawn from visual conjoint methods [159,160] will allow individualized utility functions to be generated for each participant. Both phases will

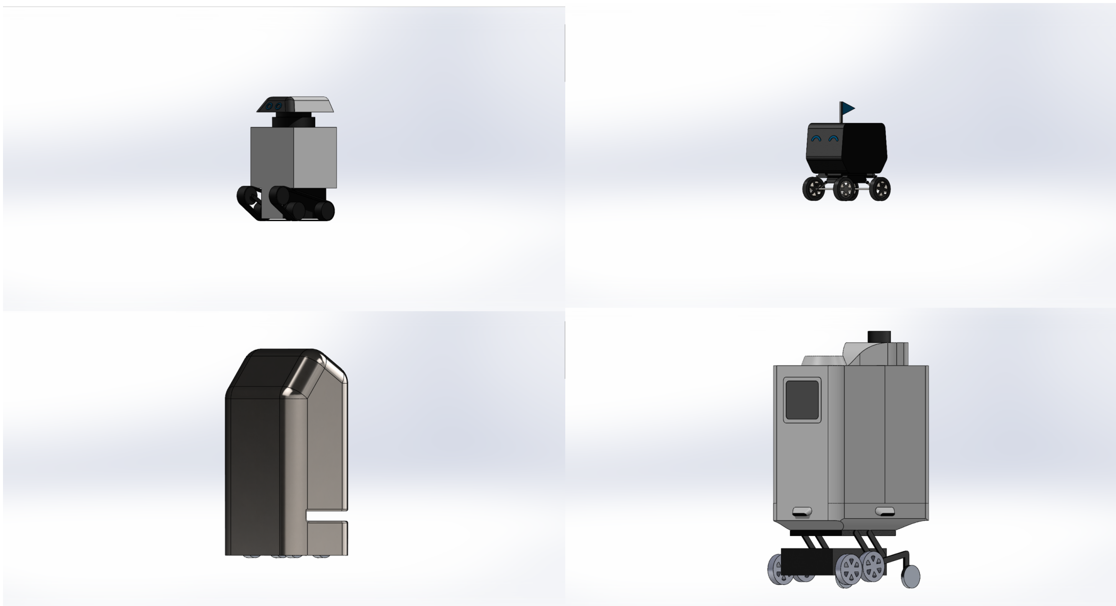

Figure 7.1: Sample images of current delivery robot designs for use in exploratory interviews. Going clockwise from the top left, designs are inspired by Daxbot [161], Kiwibot [162], Keenon W3 [163], and Fedex Roxo [164].

focus on drawing from a diverse sample of participants to glean insight into how diversity may relate to attitudes.

7.2.1 Phase 1: Exploratory Community Interviews

The objective of this phase is to gain insight into how non-expert, non-users from diverse backgrounds perceive current delivery robot designs. To do this, we will engage with members of the Hazelwood, Pittsburgh community. Currently, Carnegie Mellon is developing the Robotics Innovation Center in Hazelwood, so there is an interest in the university's role in community relations. The population of Hazelwood is also split nearly evenly between White and Black residents [165], which will allow for diverse representation among our participant sample.

We will aim to recruit 15 residents of the Hazelwood area, who have limited experience interacting with robots. The interviews will be scheduled for 1 hour, with participants being reimbursed for their time. In these interviews, sample images of current delivery robot designs, Figure 7.1, and questions regarding the designs will be asked in a semi-structured way to allow the participant to give as much of their own insight as possible. Some of the potential questions that may be asked in these interviews include:

- What kind of capabilities do you believe a delivery robot needs to complete its tasks?
- How easily do you believe you could interpret a sidewalk robot's intentions?
- What would be the first aspects of a robot you would notice if you saw one on the sidewalk?

These interviews will be recorded and a transcript will be produced. Afterward, each interview will be analyzed using thematic analysis [166], which focuses on identifying common themes across interviews. In this case, these themes will consist of robot design parameters that participants consistently identify. We hypothesize that some parameters that will be identified are robot size, color, and presence of a digital screen for communication. We also hypothesize that Black participants and participants who are women will express greater levels of distrust of robots, though our sample size will be too small to draw definitive conclusions.

7.2.2 Phase 2: Visual Conjoint for Robot-Pedestrian Interactions

Based on the results of phase 1, we will design an experiment to quantify the effects of changing robot design parameters on human perceived safety. To do this, we will develop a Unity-based simulation environment, using SEAN [167] for integration with the robot control package ROS. In this environment, we will produce videos of a virtual avatar moving along a city sidewalk, with other pedestrians nearby. Additionally, each of these videos will show a robot moving the opposite direction from the avatar as they eventually pass each other. In each video, the design and movement parameters of the robot will be modulated. The exact parameters to be tested in this experiment will be chosen based on the results of phase 1. Each participant will watch a series of videos and after each trial, answer a modified Godspeed questionnaire to assess their perceived safety during the interaction. The Godspeed questionnaire has been used in human-robot interaction experiments to assess perception of social robots [158]. Though it has not been utilized

Figure 7.2: Virtual environment to evaluate pedestrian interaction with sidewalk robots.

specifically for delivery robots, we believe its validity will be maintained in for this application. The questionnaire will appear as follows:

Based on viewing this video, please rate your emotional state on these scales:

- Anxious (1) - (5) Relaxed
- Calm (1) - (5) Agitated
- Still (1) - (5) Surprised

These several questions are asked to establish internal consistency, which can be evaluated using Cronbach's alpha coefficient. To gain a sufficient number of samples, we will crowdsource this study online through a service like Amazon Mechanical Turk. However, it is difficult to assure a diverse sample from crowdsourcing platforms, so we will return to Hazelwood to conduct this experiment with 15 more participants, conducting a post-op semi-structured interview to gain further insight to diverse perspectives.

The questionnaire allows us to generate a numerical score between 1-5 for the perceived safety of each video trial. By varying parameters across trials, this allows us draw from visual conjoint literature to generate utility curves for each participant along each parameter axis. Visual conjoint is a method to draw out user preferences for aspects of product design [159,160]. The utility curves that are generated are then used to inform customer-specific optimal design choices. Here, we leverage these utility curves in a slightly different manner. Our objective is not to necessarily determine optimal parameter values, but rather bounds on design parameters where perceived safety deteriorates. Based on the results of this experiment, we aim to inform policy makers on how regulations might be set on sidewalk robot operations.

Part III

Concluding Remarks

Chapter 8

Conclusion

This thesis focuses on unpacking the myriad of technical and societal factors that impact the success of robot deployment. In particular, I examine two open problems in robotics: navigating complex, uncertain environments and navigating complex, human environments. There are parallels in the approach toward each of these problems. First is the effort to quantify under what circumstances a robot can feel confidently safe in executing a behavior, and second is when it may need to adapt its behavior under heightened risk.

Part I addresses the issue of robots navigating through uncertain terrain where any number of disturbances could cause a robot to become stuck or fall over. 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 shown so far have been implemented in simulation, ranging from simple monopod hoppers to more complex biped and quadruped systems. Proposed work will apply these concepts to hardware experiments, with a quadruped robot navigating various difficult narrow paths.

Part II proposed experiments to understand the relationship between robot design parameters and human attitudes. This part also examined how the results of this proposed investigation connects to both historical and forthcoming legal statutes. 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. In addition to proposing the aforementioned robot design experiments, Chapter 7 also argues for how these studies can inform regulations on sidewalk delivery robots.

8.1 Possible Future Work

There are several possible extensions to this thesis that future researchers could investigate. Below is a non-comprehensive discussion of potential future work.

8.1.1 Control-Initiated Hybrid Events

Chapter 3 introduced the concept of virtual hybrid events, which are state-triggered discrete changes in control effort. These virtual hybrid events induce saltation behavior from the system without requiring a physical impact. As shown in Fig. 3.4, these virtual hybrid events can further improve the stability of a hybrid system, or even be used to inject discrete stability improvements into smooth systems. Considering that in practice, robot controllers tend to discontinuously switch between input commands, this stability improvement need not necessitate a drastically different deployment of the controller. Future work could explore how to design the virtual guard functions that trigger stabilizing input switches within a controller and how the choice of frequency and timing of virtual hybrid events affects this stability.

8.1.2 Stability and Convergence of Systems Under Stick-Slip Transitions

Another extension of the saltation-based analysis of hybrid systems is examining another type of hybrid event, the stick-slip and slip-stick transitions. In our complex world, small slips are common and a strategy to handle these potentially calamitous perturbations is necessary. This is true for robots with legs or wheels, and even autonomous vehicles. These transitions also induce saltation behavior in the system, and similar analysis to this work can be used to analyze them. Because of the difficulty in modelling stick-slip transitions, methods that capture uncertainty in transitions may improve robustness, as currently found in hybrid state estimation literature [168].

8.1.3 Impact of Rights Assurances On Robot Trust

In Chapter 6, I argue that the legal rights of humans will always dominate the right a person has to their property (i.e robot). However, this may not be apparent to non-user, non-experts as they come into close proximity with robots. The understanding of relative rights between a person and robot may have a significant impact on how comfortable people feel around robots. I hypothesize that if humans are reassured that their rights take precedent over a robot's, then their trust and comfort toward that robot will increase. A future experiment could examine the difference in attitudes toward robots based on an a priori rights assurance, manipulated so that either that the human's rights are prioritized, the robot's rights are prioritized, or that there is an equal prioritization of rights. The results of this experiment could inform roboticists, policy makers, and citizen's rights groups the importance of informing the public about the legal structure of human-robot interaction.

8.1.4 Policy Making For Public Robots

Beyond strictly academic studies, this thesis also provides implications for policy and standards setting for public robots. One proposed objective of Chapter 7 is to establish constraints on robot design where human trust deteriorates. These results can be valuable to enact well-informed policies from institutions such as IEEE or local and state governments. Other future works such as rights assurances can also be used to form policies that ensure equitable and beneficial deployment of robots into our communities.

8.2 Acknowledgements

I would like to thank many people including my advisor, Dr. Aaron M. Johnson and all of my committee members. I would also like the thank all of my close collaborators, Dr. George Council, Dr. Nathan Kong, and J. Joe Payne, who have all contributed valuable assistance in this thesis and my terrific mentees, Joshua Ramos, Anoushka Shrivastava, and Nikhil Chinnalapatti Gopinath. Thank you to all of the many people who have given me guidance, both research-related and personal, during my time here. That list is far too long to include here. I could not have accomplished any of this without the support of my family, friends, and loving dog Percy.

Bibliography

[1] M. S. Branicky, "Multiple Lyapunov functions and other analysis tools for switched and hybrid systems," *IEEE Transactions on Automatic Control*, 1998.

[2] D. Tourjée, "Several women in this rural town say they were stalked by drones," *Vice Media Group*, 2017.

[3] 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.

[4] N. J. Kong, J. J. Payne, J. Zhu, and A. M. Johnson, "Saltation matrices: The essential tool for linearizing hybrid dynamical systems," *arXiv preprint arXiv:2306.06862*, 2023.

[5] A. Seyfarth, H. Geyer, and H. Herr, "Swing-leg retraction: a simple control model for stable running," *Journal of Experimental Biology*, vol. 206, no. 15, pp. 2547–2555, 2003.

[6] 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.

[7] J. Zhu, J. J. Payne, and A. M. Johnson, "Convergent ilqr for safe trajectory planning and control of legged robots," *arxiv preprint arxiv:2304.00346*, 2023.

[8] N. J. Kong, G. Council, and A. M. Johnson, "iLQR for piecewise-smooth hybrid dynamical systems," *IEEE Conference on Decision and Control*, 2021.

[9] J. Zhu, A. Shrivastava, and A. M. Johnson, "Grounding robot navigation in self-defense law," *arXiv preprint arXiv:2304.00347*, 2023.

[10] 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.

[11] B. Friedman, D. G. Hendry, and A. Borning, *A Survey of Value Sensitive Design Methods*. Now Foundations and Trends, 2017.

[12] A. van der Schaft and J. Schumacher, “Complementarity modeling of hybrid systems,” *IEEE Transactions on Automatic Control*, vol. 43, no. 4, 1998.

[13] J. Lygeros, K. H. Johansson, S. Simic, J. Zhang, and S. S. Sastry, “Dynamical properties of hybrid automata,” *IEEE Transactions on Automatic Control*, 2003.

[14] 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*, vol. 35, no. 11, pp. 1354–1392, 2016.

[15] M. Jeffrey, “Dynamics at a switching intersection: Hierarchy, isonomy, and multiple sliding,” *SIAM Journal on Applied Dynamical Systems*, vol. 13, pp. 1082–1105, 07 2014.

[16] 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, N. Lynch and B. H. Krogh, Eds. Berlin, Heidelberg: Springer, 2000, pp. 421–436.

[17] 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.

[18] 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*, vol. 11, no. 4, pp. 385–398, 01 1958. [Online]. Available: https://doi.org/10.1093/qjmam/11.4.385

[19] S. A. Burden, T. Libby, and S. D. Coogan, “On contraction analysis for hybrid systems,” *arXiv preprint arXiv:1811.03956*, 2018.

[20] R. Leine and H. Nijmeijer, *Dynamics and Bifurcations of Non-Smooth Mechanical Systems*. Springer-Verlag Berlin Heidelberg, 2004.

[21] M. Hirsch, S. Smale, and R. Devaney, *Differential Equations, Dynamical Systems, and an Introduction to Chaos*. Academic Press, 2012.

[22] H. Asahara and T. Kousaka, “Stability analysis using monodromy matrix for impacting systems,” *IEICE Transactions of Fundamentals of Electronics, Communications, and Computer Science*, vol. E101-A, no. 6, 2018.

[23] B. Altın, P. Ojaghi, and R. G. Sanfelice, “A model predictive control framework for hybrid dynamical systems,” *IFAC-PapersOnLine*, vol. 51, no. 20, pp. 128–133, 2018.

[24] W. Yang and M. Posa, “Impact invariant control with applications to bipedal locomotion,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021.

[25] S. Burden, S. Revzen, and S. Sastry, “Dimension reduction near periodic orbits of hybrid systems,” in *IEEE Conference on Decision and Control*, September 2011.

[26] G. Council, S. Yang, and S. Revzen, “Deadbeat control with (almost) no sensing in a hybrid model of legged locomotion,” in *International Conference on Advanced Mechatronic Systems*, 2014, pp. 475–480.

[27] D. Sternad, M. Duarte, H. Katsumata, and S. Schaal, “Bouncing a ball: Tuning into dynamic stability,” *Journal of Experimental Psychology: Human Perception and Performance*, vol. 27, pp. 1163–84, 11 2001.

[28] M. M. Ankaralı, U. Saranlı, and A. Saranlı, “Control of underactuated planar hexapedal pronking through a dynamically embedded slip monopod,” in *IEEE International Conference on Robotics and Automation*, 2010, pp. 4721–4727.

[29] K. Green, R. Hatton, and J. Hurst, "Planning for the unexpected: Explicitly optimizing motions for ground uncertainty in running," in *IEEE International Conference on Robotics and Automation*, May 2020.

[30] J. G. D. Karssen, M. Haberland, M. Wisse, and S. Kim, "The optimal swing-leg retraction rate for running," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 4000–4006.

[31] E. Tyrtyshnikov, *A Brief Introduction to Numerical Analysis*. Birkhäuser Basel, 1997.

[32] P. Lancaster, "On eigenvalues of matrices dependent on a parameter," *Numerische Mathematik*, 1964.

[33] H. Khalil, *Nonlinear Systems*. Pearson, 2002.

[34] J. Nocedal and S. Wright, *Numerical Optimization*. Springer, 2006.

[35] M. Kelly, "An introduction to trajectory optimization: How to do your own direct collocation," *SIAM Review*, vol. 59, no. 4, pp. 849–904, 2017.

[36] O. Von Stryk and R. Bulirsch, "Direct and indirect methods for trajectory optimization," *Annals of Operations Research*, vol. 37, pp. 357–373, 12 1992.

[37] C. R. Hargraves and S. W. Paris, "Direct trajectory optimization using nonlinear programming and collocation," *Journal of Guidance, Control, and Dynamics*, vol. 10, no. 4, pp. 338–342, 1987.

[38] K. Mombaur, R. Longman, H. Bock, and J. Schlöder, "Open-loop stable running," *Robotica*, vol. 23, pp. 21–33, 01 2005.

[39] O. V. Stryk, "User's guide for DIRCOL (version 2.1): a direct collacation method for the numerical solution of optimal control problems," in *Technische Universitat Darmstadt*, November 1999.

[40] J. T. Betts, *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming, Second Edition*, 2nd ed. SIAM, 2010.

[41] T. A. McMahon and G. C. Cheng, "The mechanics of running: How does stiffness couple with speed?" *Journal of Biomechanics*, vol. 23, pp. 65–78, 1990.

[42] P. Holmes, R. J. Full, D. E. Koditschek, and J. Guckenheimer, "The dynamics of legged locomotion: Models, analyses, and challenges," *SIAM Review*, vol. 48, no. 2, pp. 207–304, 2006.

[43] A. Seyfarth, H. Geyer, M. Günther, and R. Blickhan, "A movement criterion for running," *Journal of Biomechanics*, vol. 35, no. 5, pp. 649–655, 2002.

[44] R. J. Full and D. E. Koditschek, "Templates and anchors: neuromechanical hypotheses of legged locomotion on land," *Journal of Experimental Biology*, vol. 202, no. 23, pp. 3325–3332, 12 1999.

[45] J. Seipel and P. Holmes, "A simple model for clock-actuated legged locomotion," *Regular and Chaotic Dynamics*, vol. 12, no. 5, Oct. 2007.

[46] M. Raibert, H. Brown, and M. Chepponis, "Experiments in balance with a 3D one-legged hopping machine," *The International Journal of Robotics Research*, vol. 3, no. 2, pp. 75–92, 1984.

[47] 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*, vol. 54, no. 8, pp. 1779–1793, 2009.

[48] V. Chen, "Passive dynamic walking with knees: A point foot model," Master's thesis, Massachusetts Institute of Technology, 2007.

[49] A. M. Johnson and D. E. Koditschek, "Toward a vocabulary of legged leaping," in *IEEE Intl. Conference on Robotics and Automation*, 2013.

[50] 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.

[51] 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.

[52] 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," *arXiv preprint arXiv:2302.09450*, 2023.

[53] R. Tedrake, *Underactuated Robotics*, 2022. [Online]. Available: http://underactuated.mit.edu

[54] M. A. Haidekker, *Linear Feedback Controls*, M. A. Haidekker, Ed. Elsevier, 2020.

[55] J. Norby, A. Tajbakhsh, Y. Yang, and A. M. Johnson, "Adaptive complexity model predictive control," *arXiv preprint arXiv:2209.02849*, 2022.

[56] Z. Manchester and S. Kuindersma, "Robust direct trajectory optimization using approximate invariant funnels," *Auton. Robots*, 2019.

[57] 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.

[58] T. Lew, R. Bonalli, and M. Pavone, "Chance-constrained sequential convex programming for robust trajectory optimization," *European Control Conference*, 2020.

[59] B. Hammoud, M. Khadiv, and L. Righetti, "Impedance optimization for uncertain contact interactions through risk sensitive optimal control," *IEEE Robotics and Automation Letters*, 2021.

[60] 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.

[61] L. Drnach and Y. Zhao, "Robust trajectory optimization over uncertain terrain with stochastic complementarity," *IEEE Robotics and Automation Letters*, 2021.

[62] 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.

[63] C.-Y. Lee, S. Yang, B. Bokser, and Z. Manchester, "Enhanced balance for legged robots using reaction wheels," *2023 IEEE International Conference on Robotics and Automation*, 2023.

[64] 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.

[65] A. Papachristodoulou and S. Prajna, "On the construction of Lyapunov functions using the sum of squares decomposition," *IEEE Conference on Decision and Control*, 2002.

[66] 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.

[67] 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.

[68] W. Lohmiller and J. Slotine, "On contraction analysis for non-linear systems," *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.

[69] A. M. Johnson, J. E. King, and S. Srinivasa, "Convergent planning," *IEEE Robotics and Automation Letters*, 2016.

[70] N. Kong and A. M. Johnson, "Optimally convergent trajectories for navigation," in *International Symposium on Robotics Research*, 2019.

[71] 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.

[72] 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.

[73] J. Townsend, "Differentiating the singular value decomposition," Technical Report 2016, https://j-towns. github. io/papers/svd-derivative . . . , Tech. Rep., 2016.

[74] S. Singh, R. P. Russell, and P. M. Wensing, "On second-order derivatives of rigid-body dynamics: Theory & implementation," *arXiv preprint arXiv:2302.06001*, 2023.

[75] C. G. Broyden, "The convergence of a class of double-rank minimization algorithms 1. General considerations," *IMA Journal of Applied Mathematics*, 1970.

[76] 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.

[77] 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.

[78] K. Mombaur, "Using optimization to create self-stable human-like running," *Robotica*, 2009.

[79] 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.

[80] J. Ackerman and J. Seipel, "Energy efficiency of legged robot locomotion with elastically suspended loads," *IEEE Transactions on Robotics*, 2013.

[81] M. J. Pollayil, C. D. Santina, G. Mesesan, J. Englsberger, 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.

[82] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, "Crocoddyl: An Efficient and Versatile Framework for Multi-Contact Optimal Control," in *IEEE International Conference on Robotics and Automation*, 2020.

[83] 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.

[84] 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.

[85] M. Vukobratović and B. Borovac, "Zero-moment point—thirty five years of its life," *International journal of humanoid robotics*, vol. 1, no. 01, pp. 157–173, 2004.

[86] M. Chignoli and P. M. Wensing, "Variational-based optimal control of underactuated balancing for dynamic quadrupeds," *IEEE Access*, vol. 8, pp. 49 785–49 797, 2020.

[87] 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.

[88] W. Xi, Y. Yesilevskiy, and C. D. Remy, "Selecting gaits for economical locomotion of legged robots," *The International Journal of Robotics Research*, 2016.

[89] Q. Li, L. Qian, S. Wang, P. Sun, and X. Luo, "Towards generation and transition of diverse gaits for quadrupedal robots based on trajectory optimization and whole-body impedance control," *IEEE Robotics and Automation Letters*, 2023.

[90] G. Haynes, A. Rizzi, and D. Koditschek, “Multistable phase regulation for robust steady and transitional legged gaits,” *The International Journal of Robotics Research*, 2012.

[91] R. Alexander, *Principles of Animal Locomotion*. Princeton University Press, 2003.

[92] 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.

[93] D. Owaki and A. Ishiguro, “A Quadruped Robot Exhibiting Spontaneous Gait Transitions from Walking to Trotting to Galloping,” *Scientific Reports*, 2017.

[94] D. Lee, G. Kang, B. Kim, and D. H. Shim, “Assistive delivery robot application for real-world postal services,” *IEEE Access*, 2021.

[95] 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.

[96] 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.

[97] S. Brown-Heidenreich, “KiwiBot catches fire outside Martin Luther King Jr. student union,” *The Daily Californian*, 2018.

[98] L. O’Haver, “UT students arrested after damaging food delivery robot,” *WATE*, 2022.

[99] E. Parks and D. Zoga, “NBC 5 responds: What happens if you’re in a crash with a robot?” *NBCDFW*, 2020.

[100] R. M. Brown, *No duty to retreat: Violence and values in American history and society*. Oxford University Press, 1991.

[101] R. A. Epstein, "The theory and practice of self-help," *Journal of Law, Economics, and Policy*, 2005.

[102] 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.

[103] K. W. Simons, "Self-defense: Reasonable beliefs or reasonable self-control?" *New Criminal Law Review*, 2008.

[104] L. S. Scheff, "People v. Humphrey: Justice for battered women or a licence to kill," *University of San Francisco Law Review*, 1997.

[105] H. P. Thompson, "The State v. Sutton," *Supreme Court of Georgia*, 2015.

[106] R. A. Rosen, "On self-defense, imminence, and women who kill their batterers," *North Carolina Law Review*, 1993.

[107] A. M. Froomkin and Z. Colangelo, "Self-defense against robots and drones," *Connecticut Law Review*, vol. 48, no. 1, 2015.

[108] D. Schneider, "Open season on drones?" *IEEE Spectrum*, 2014.

[109] S. Uniacke, "Proportionality and self-defense," *Law and Philosophy*, 2011.

[110] B. Levin, "A defensible defense?: Reexamining castle doctrine statutes," *Harvard Journal on Legislation*, 2010.

[111] 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.

[112] J. E. Anderson and D. Dunning, "Behavioral norms: Variants and their identification," *Social and Personality Psychology Compass*, 2014.

[113] 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.

[114] 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.

[115] 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.

[116] M. K. Strait, C. Aguillon, 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.

[117] M. Zaveri, “N.Y.P.D.’s robot dog returns to work, touching off a backlash,” *The New York Times*, 2021.

[118] S. A. Periwal, “Boston Dynamics’ robot dog strikes fear online as fans recall Black Mirror episode ‘Metalhead’,” *Sportskeeda*, 2020.

[119] A. Jarmanning, “Mass. state police tested out Boston Dynamics’ Spot the robot dog. civil liberties advocates want to know more,” *WBUR*, 2019.

[120] 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.

[121] J. Buolamwini and T. Gebru, “Gender shades: Intersectional accuracy disparities in commercial gender classification,” *Conference on Fairness, Accountability and Transparency*, 2018.

[122] C. Gidaris, “The rise of the robots: Technocapitalism and the policing of race,” *Fast Capitalism*, 2021.

[123] K. M. Thomasen, “Beyond airspace safety: A feminist perspective on drone privacy regulation,” *Canadian Journal of Law and Technology*, 2018.

[124] A. Gurinskaya, “Support for robocops: Measuring effects of attitudes towards police and policing technologies,” *Digital Transformation and Global Society*, 2022.

[125] 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.

[126] S. M. LaValle, “Rapidly-exploring random trees: A new tool for path planning,” *Tech. Rep. TR 98-11*, 1998.

[127] J. Chu, F. Zhao, S. Bakshi, Z. Yan, and D. Chen, “Risk-aware path planning with uncertain human interactions,” *American Control Conference*, 2021.

[128] A. V. Taylor, E. Mamantov, and H. Admoni, “Observer-aware legibility for social navigation,” *IEEE International Conference on Robot and Human Interactive Communication*, 2022.

[129] 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.

[130] 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.

[131] P. Langley, B. Meadows, M. Sridharan, and D. Choi, “Explainable agency for intelligent autonomous systems,” *AAAI Conference on Artificial Intelligence*, 2017.

[132] 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.

[133] S. Sohrabi, J. A. Baier, and S. A. McIlraith, "Preferred explanations: Theory and generation via planning," *AAAI Conference on Artificial Intelligence*, 2011.

[134] L. Zhu and T. Williams, "Effects of proactive explanations by robots on human-robot trust," *Social Robotics*, 2020.

[135] S. Atakishiyev, M. Salameh, H. Yao, and R. Goebel, "Explainable artificial intelligence for autonomous driving: An overview and guide for future research directions," *arXiv preprint arXiv:2112.11561*, 2021.

[136] S. Anjomshoae, 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.

[137] 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.

[138] 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.

[139] 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.

[140] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon, "A human aware mobile robot motion planner," *IEEE Transactions on Robotics*, 2007.

[141] 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.

[142] 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.

[143] 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.

[144] B. Busch, J. Grizou, M. Lopes, and F. Stulp, “Learning legible motion from human-robot interactions,” *International Journal of Social Robotics*, 2017.

[145] S. Wallkotter, 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.

[146] 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.

[147] 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.

[148] 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.

[149] I. S. 7001-2021, “IEEE standard for transparency of autonomous systems,” 2022.

[150] Federal Aviation Administration, “Sec. 107.29 operation at night,” *Code of Federal Regulations*, 2021.

[151] P. W. Groeneveld, S. S. Sonnad, A. K. Lee, D. A. Asch, and J. E. Shea, “Racial differences in attitudes toward innovative medical technology,” *Journal of general internal medicine*, 2006.

[152] The General Assembly of Pennsylvania, “Senate bill 1199,” 2020.

[153] State of Washington 66th Legislature, “Engrossed substitute house bill 1325,” 2019.

[154] The City and County of San Francisco, “Sec. 794 autonomous delivery devices on sidewalks - permit required,” 2018.

[155] City of Toronto, “Automated micro-utility devices - accessibility feedback,” 2021.

[156] T. Andersen and D. Cawthorne, “Value-sensitive design of unmanned aerial systems: Using action research to bridge the theory-practice gap,” *International Journal of Technoethics*, 2021.

[157] A. Van Wynsberghe, “Designing robots for care: Care centered value-sensitive design,” in *Machine ethics and robot ethics*. Routledge, 2020, pp. 185–211.

[158] C. Bartneck, E. Croft, and D. Kulic, “Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots,” *International Journal of Social Robotics*, 2009.

[159] N. Tovares, P. Boatwright, and J. Cagan, “Experiential conjoint analysis: An experience-based method for eliciting, capturing, and modeling consumer preference,” *Journal of Mechanical Design*, 2014.

[160] S. Orsborn, J. Cagan, and P. Boatwright, “Quantifying aesthetic form preference in a utility function,” *Journal of Mechanical Design*, 2009.

[161] Daxbot, “Delivery for business daxbot: Robots that we like,” accessed: 2023-09-24. [Online]. Available: https://daxbot.com/

[162] Kiwibot, “Adorable autonomous delivery robots,” accessed: 2023-09-24. [Online]. Available: https://kiwibot.com/

[163] RobotShop, "Keenon w3 hotel delivery robot waiter fully automatic welcome and delivery hotel robot, automatic elevator ride," accessed: 2023-09-24. [Online]. Available: https://www.robotshop.com/products/keenon-w3-hotel-delivery-robot-waiter-fully-automatic-welcome-and-delivery-hotel-robot-automatic-elevator-ride

[164] Fedex, "Meet roxo, the fedex sameday bot," accessed: 2023-09-24. [Online]. Available: https://www.fedex.com/en-ge/about/sustainability/our-approach/roxo-delivery-robot.html

[165] Program in Urban and Regional Analysis, University Center for Social and Urban Research, University of Pittsburgh, "The hazelwood neighborhood, 2010," 2012.

[166] S. J. Tracy, *Qualitative Research Methods*. Wiley-Blackwell, 2013.

[167] N. Tsoi, A. Xiang, P. Yu, S. S. Sohn, G. Schwartz, S. Ramesh, M. Hussein, A. W. Gupta, M. Kapadia, and M. Vázquez, "Sean 2.0: Formalizing and generating social situations for robot navigation," *IEEE Robotics and Automation Letters*, 2022.

[168] 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.