THE UNIVERSITY
of EDINBURGH

Understanding the Fundamentals of Bipedal Locomotion in Humans and Robots

Christopher McGreavy

Doctor of Philosophy

Institute of Perception, Action and Behaviour
School of Informatics
University of Edinburgh

2023

Abstract

Walking is a robust and efficient method of moving around the world, which would greatly enhance the capabilities of humanoid robots, although they cannot match the performance of their biological counterparts. The highly nonlinear dynamics of locomotion create a vast state-action space, which makes model-based control difficult, yet biological humans are highly proficient and robust in their motion while operating under similar constraints. This disparity in performance naturally leads to the question: what can we learn about locomotion control by observing humans, and how can this be used to develop bio-inspired locomotion control in mechatronic humanoids? This thesis investigates bio-inspired locomotion control, but also explores the limitations of this approach and how we can use robotic platforms to move towards a better understanding of locomotion.

We first present a methodology for measuring and analysing human locomotion behaviour, specifically disturbance recovery, and fit models to this complex behaviour to represent it in as simple as possible such that it can be easily translated into a simple controller for reactive motion. A minimum-jerk Model Predictive Control algorithm at the Centre of Mass (CoM) best captured human motion during multiple recovery strategies instead of using one controller for each strategy, which is common in this area. Capturing this simple CoM model of complex human behaviour shows that bio-inspiration can be an important tool for controller development, but behaviour varies between and even within individuals given similar initial conditions, which manifests as stochastic behaviour. Coupled with the ability to only measure expressed behaviours instead of direct control policies, this stochasticity presents a fundamental limit to using bio-inspiration for control purposes, as only indirect inferences can be made about a complex, stochastic system.

To overcome these barriers, we investigate the use of mechatronic humanoid robots as a means to explore invariant aspects of the vast dynamic state-space of locomotion which are described by physical laws, and are therefore not subject to the stochastic behaviour of individual humans, that apply to both biological and mechatronic humanoid forms. We present a pipeline to explore the invariant energetics of humanoid robots during stepping for push recovery, where the most efficient stepping parameters are identified for a given initial CoM velocity and desired step length. Using this to explore the stepping state-space, our

analysis finds a region of attraction between disturbance magnitude and optimal step length surrounded by a region of similarly efficient alternatives which corresponds to the stochastic behavior observed in humans during push recovery, which we would be unable to identify without reproducibility, direct access to internal measurements and known full body dynamics, which is not available in humans.

We expand this paradigm further to investigate the invariant energetics of continuous walking using a full-body humanoid by exploring the state-space of step-length and step-timing to identify the most efficient sub-spaces of these parameters which describes the most efficient way to walk. Through analysis of this state-space, we provide evidence that the humanoid morphology exhibits a passive tendency towards energy-optimal motion and its dynamics follow a region of attraction towards Cost of Transport-optimal motion.

Overall, these findings demonstrate the utility of robotics as a tool with which to explore certain aspects of legged locomotion and the results gained from our methodology suggest that humans do not need to explore a vast state-action space to learn to walk, they need only internalise simple heuristics for the natural dynamics of stepping that are easy to learn and can produce rapid, reactive and efficient stepping without costly decision-making processes.

Lay Summary

Human beings are experts at walking. They can walk over almost all terrain, no matter how unstructured, climb, jump over gaps, and they do these reliably, efficiently and while doing other things like talking and carrying objects.

If robots were able to walk with the same level of performance, they could be very beneficial to society. They could move around human-centred environments for healthcare assistance, work in unstructured and dangerous environments, such as construction and disaster relief, and even be used for planetary exploration with adaptability that is not possible with wheeled robots.

Naturally, if we knew exactly how humans performed their impressive feats of motion, we might be able to make humanoid robots walk and perform the same way. Unfortunately, though robot walking has come a long way in the past few decades and there has been a lot of effort invested into learning from humans through a process called bio-inspiration, robots are still a long way from humans in terms of performance.

One of the major problems in learning motion skills from humans is that we cannot directly access any information about how they are thinking and how they make decisions about moving their limbs to walk in the way they do. We are instead limited to looking at the motions they perform and trying to reverse-engineer how these are produced.

Because of this limitation, we do not understand the underlying 'rules' of locomotion from a human standpoint well enough to appropriately reproduce this motion. The aim of this thesis is search for some of these rules and explore how they might be useful to robot control and expanding our understanding of how humans move.

Since we cannot observe human control processes, we look at human movement data first, then investigate to see if humanoid robots can be used to study humans, as we have full access to the internal processes and measurements of these machines. Though mechanistic robots are not exactly the same as biological humans, they are a platform on which we can study the physics of walking in an accurate and repeatable way.

Through this method, we find that there are some consistencies in how the hu-

manoid body works during locomotion. It tends to be energy-efficient, and we can make simple rules to represent this efficiency that can be useful for robot control and explain how humans are so efficient without thinking about it - they are working with their bodies to walk with as little energy expenditure as possible, and we argue that this is an evolutionary advantage, as early humans would not be required to waste energy on thinking about how they walk, nor waste energy on unnecessary walking styles.

Acknowledgements

First my, thanks and gratitude go to my supervisor, Zhibin (Alex) Li, whose guidance, insight and seemingly limitless knowledge has been an indispensable source of academical, professional and personal support. Thanks also go to my second supervisor Sethu Vijayakumar, for his advice, support and opportunities to expand my horizon in robotics. They took a chance on me as an aspiring, novice roboticist; I wouldn't be where I am today, without the opportunity and support they provided. For that I am truly and forever grateful.

Deep thanks go to the members of my lab and the SLMC lab, who were so kind, patient and knowledgeable. The help and support of every lab member has been indispensable. Wouter, Mohammad, Steve, Quentin, Kai, Vlad, Wenbin, Jordan and Eleftherios in particular gave their time and effort at pivotal moments.

When I first moved to Edinburgh, my fellow inhabitants of office 1.38 were immensely welcoming, kind and took me under their collective wing. They taught me so much of what I know about areas of robotics that I thought I'd never understand, offered sage advice on just about anything and most of all, they have brought so much fun to my life. I am very grateful to them.

Special thanks to Daniel Gordon for his unparalleled patience, gait-lab expertise and for being the enthusiastic co-conspirator in our countless side-projects for which the world is worse off for us not having completed (/started), I'm sure!

I am immensely thankful for my friends, who were there throughout my PhD and make life so enjoyable. My regular D&D group made Wednesdays so fun and gave me something to look forward to. Tatianna and Matthew have been a constant source of advice, fun and escapism. Sam, Sean and Luke(Rick) never fail to make me smile, even when things are difficult. Caitlin, Christina, Dan, Calum and Andre are nerds. Wonderful nerds. I am proud and grateful to be one of them. So many other friends have supported me during my PhD in ways they cannot know, I'm deeply thankful to all of them.

Finally, my most profound gratitude goes to my family. They support and encourage everything I do. They instilled in me a love of engineering, a work ethic and persistence that is the cornerstone of my achievements that would otherwise be impossible, and a belief that no matter how tricky things get, "it be reet".

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(*Christopher McGreavy*)

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Robots that walk as a means of mobility have incredible value to society. Where wheeled robots need continuous, structured surfaces, legged robots can move over discontinuous, cluttered environments and even traverse gaps, making them good candidates to work in human-centric environments as healthcare assistants, in construction, in unstructured disaster areas and as human substitutes for planetary exploration.

Though these benefits are clear, the progress towards usable legged robots has been slow. Only recently have commercial quadruped robots become available to consumers and industry and are beginning to make an impact in inspection, exploration and haulage. Humanoid robots, however, are still largely constrained to highly controlled lab settings, and even those that emerge from the lab in a limited capacity are not yet ready to enter the mainstream.

The progress towards usable legged robots has been slow, however, due largely to the complexity of the subject. Walking robots have complex nonlinear dynamics, making their movements difficult to predict. They need to be careful when placing their feet and regulating their interactions with the environment to ensure stability. For quadrupeds, these problems have been largely mastered, as they are much less susceptible to instability than bipeds. Humanoid bipeds present a more complex task, as they are top-heavy and thus unstable during walking and have difficult to compute dynamics. In combination, these problems create an infinitely large state space in which bipedal robots operate during walking, making reliable bipedal robot control a grand challenge in robotics.

Although bipedal locomotion is more complex, using two legs instead of four has its benefits. Humanoids are better suited to navigate and interact with environments and tools designed for humans [6, 7]. Moreover, in terms of biology, humans have the ability to travel longer distances on foot than their four-legged counterparts, which may result in a theoretically superior operational time [8].

Engineering approaches to this challenge have been successful in producing walking motions that for a long time have been rudimentary and are gradually becoming more impressive, but still fall short of the performance needed for humanoids to be useful tools and the highest performance motions are rarely seen on actual hardware.

Instead of a bottom-up engineering approach, others have turned to biology as a means of understanding locomotion. Biological humans are of particular interest, as they have mastery over bipedal locomotion and are the morphological templates for humanoid robots. Locomotion performed by biological humans is naturally an ideal template for humanoid robots, it is robust, agile, and is performed with such ease that a person can do other things while walking with minimal effect on performance.

This bio-inspired approach to locomotion research has led to a greater understanding of walking, and useful tools for designing robot controllers [9, 10, 11]. Moving into the future, biological humans will continue to provide a vital template for research into locomotion, but this must be done with the recognition of the limits to bioinspiration, both that not all human behaviour can be transplanted directly onto mechatronic humanoids, due to the differences between the biological and manufactured counterparts, and that there are fundamental limits to how we can measure human control due to the lack of observability in in-situ biological processes – in other words, we cannot directly measure how humans think or control their limbs, we are only able to observe from behaviour for the most part.

We can, however, find other ways of using current technology to investigate locomotion despite the challenges of direct measurement. There are fundamental principles that apply to the humanoid morphology – biological or mechatronic – during locomotion and robots may be a good candidate to investigate these, as we can fully observe their internal processes and control them precisely. In

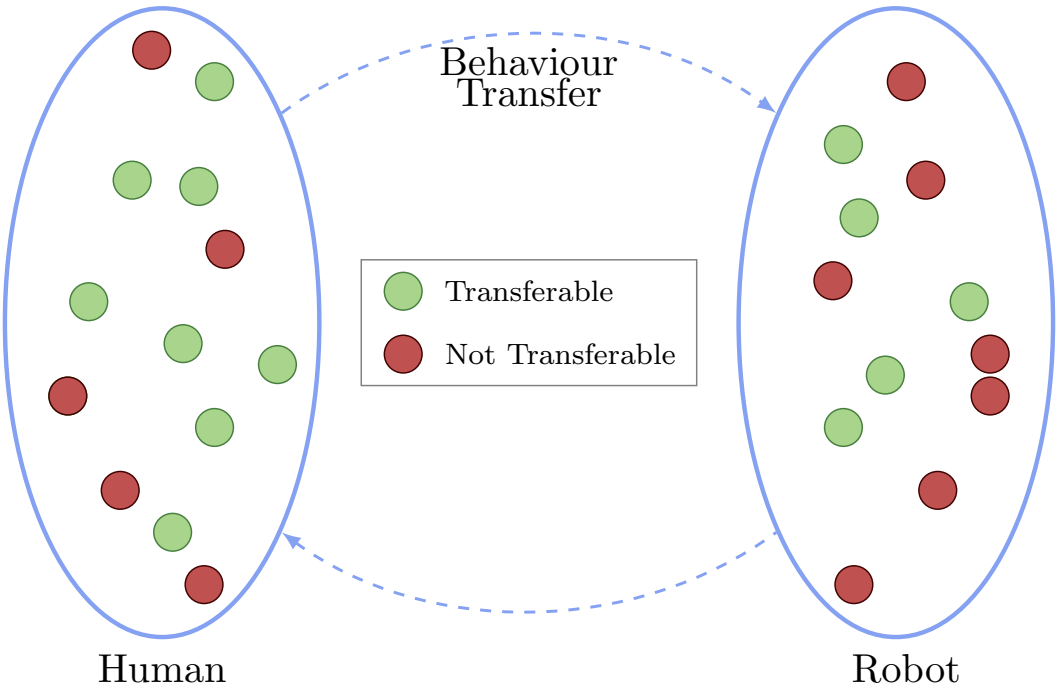

Figure 1.1: The space of potential behaviours and methods that might be transferred between biological and mechatronic humanoids.

this thesis, we investigate the underlying, invariant principles of locomotion in humans based on behavioural measurement, and explore how robots can be used as test-beds to investigate human control decisions and dynamics to enhance our understanding of bipedal locomotion in humanoid morphology.

1.1 Biological and Mechatronic Humanoids as Surrogates

Biological humans are a valuable resource for understanding high-performance locomotion, but obtaining the necessary information from them is a significant challenge. As discussed in detail in Section 2.4, it is not possible to observe internal human control processes during motion using current or even emerging technology, and we are constrained to making inferences about control using behavioural data, which is easy to measure.

Behavioural inference remains a useful tool for robotics, as we explore in Chapter 4, and can give us beneficial insight into the underlying principles of legged locomotion and feeds a large pool of knowledge and data for different control methods

[12, 13]. We also argue that, given the limitations of biological observability, that mechatronic humanoid robots can be used as surrogates to investigate walking behaviour in humans and humanoid morphology in general, which we explore in Chapters 5 and 6.

Both of these approaches must be employed with a recognition of the constraints of transferring knowledge from one to the other, in either direction. It is clear that transferring human behaviour onto mechatronic humanoids is beneficial and desirable for legged robotics research, and there is an implicit assumption that, given sufficient knowledge of biological control policies, all facets of human walking behaviour can be transferred to mechatronic humanoids, and that mechatronic systems have little to offer in terms of learning about biological human motion. To continue along this research track, we must recognise the similarities and differences between the two systems and understand the limits and potential of transferring in each direction.

Biological and mechatronic humanoid systems clearly share a morphology, as the mechatronic structure is based on the biological template. In both systems, the bodies are subject to the same universal forces of physics, with very similar joint configurations. Contact switching is a requirement for walking, so both systems must move their bodies to meet this need, and the aim of the robot is to produce agile, robust and efficient motion similar to that of humans.

The systems differ in their actuation, computation, mass distribution, repeatability, material components and process observability, which are barriers to transferability (Figure 1.2), but some of these differences make each one ideal for studying the other. For example, consider the differences in performance and repeatability; humans perform locomotion more deftly than the best robotic systems, but cannot reproduce conditions and results exactly, where robots can, despite their lower performance. Similarly, ethical considerations are vital for human research, especially for long-duration and dangerous experiments, where this does not exist in robotics.

In terms of transferability, it is not reasonable to assume that all human motion behaviour is transferable to robots and vice-versa. Toe-off motions are a good example of this. In humans, toe-off motions are a vital part of the gait cycle [14], but given current technology and underactuated nature of the motion, toe-

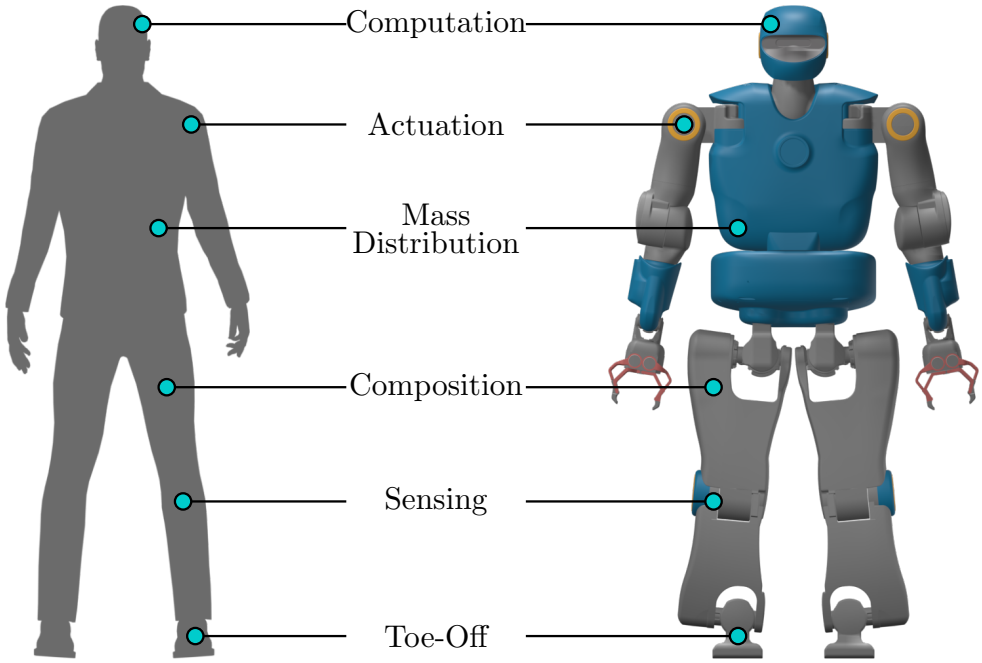

Figure 1.2: Humans and robots share similar morphology, but the difference in processing and actuation means we need to investigate which behaviours are transferable.

off motions are very difficult to perform on mechatronic robots and have only been demonstrated in a few cases [15]. Similarly, there is no value in projecting specific engineered controllers onto humans, as humans can already inherently out-perform them and there is no insight that can be gained from this.

As we see it is not reasonable to assume that all human behaviour can be transferred to robots in a one-to-one fashion and similarly, we must be careful about making generalisations from mechatronic robots back to humans. Biorobotics is a field of robotics that deals with this boundary between biology and mechatronics which has dealt with these questions of applicability between them. When taking this biorobotic approach, as part of the experiment, it must be considered on what level the modelling is being done: on a neuronal level, structural level, adaptation level, etc. and be careful to apply any findings to the biological template whilst acknowledging the differences in between the systems and knowing that models will have imperfections [16, 17].

In this thesis, we investigate the underlying principles of the dynamics of humanoid locomotion on a structural and dynamic level. More specifically, we aim to investigate how biological and mechatronic humanoid morphology perform on

structural and dynamic levels in a range of conditions, and, through analysis, search for common principles that exist between them that might be fundamental rules of locomotion dictated by the interaction between the physical human morphology and invariant rules of dynamics. Given our findings, we then discuss reasonable and logical extensions within this scope, such as how they might have arisen, if and how they are used in human locomotion and how they might benefit robotics research in the longer term.

1.2 Invariant Principles of Locomotion

The need for contact switching is part of the definition of walking and is a required part of the motion, and as such is an invariant, immutable component of walking, whereas the actual contact pattern that is used by a person or robot changes depending on the structural morphology and the context.

In this thesis, we investigate the motions of humans and robots and analyse the results to search for principles that are present in locomotion and are present across contexts and morphologies. We pay particular attention to how walking parameters are selected in different circumstances, the consistencies between these, as well as energy efficiency and inherent morphological optimality.

1.3 Thesis Structure

We can summarise the aims of this thesis in the following research questions:

1. Are there common features between different push recovery strategies in the humanoid morphology?
2. Can we identify subspaces in the dynamics of the humanoid morphology that could be used to improve our understanding of bipedal locomotion?
3. Are there heuristics or simplifications of the humanoid morphology that might explain human reactive speeds that can also be used in robot control?

To begin answering these questions, we first present and discuss historical and modern state-of-the-art approaches to locomotion control in robotics, as well as methods of motion analysis and biorobotics in Chapter 2 to provide context for the research presented in the rest of this thesis. The key concepts of locomotion,

modelling and gait analysis related to the technical aspects of the research are given in Chapter 3. In Chapter 4, our work on human locomotion analysis is presented, which explores the common aspects of human motion and stepping during push recovery.

Open questions from human study lead us to the exploration of robots as human surrogates for studying energy-efficiency during stepping in Chapter 5, where we present a proof-of-concept pipeline for building efficiency maps and studying invariance in locomotion.

Building on this, we then explore robotic locomotion more expansively in Chapter 6, where walking control is analysed to find common and invariant aspects of the relationship between walking parameters and energy-efficiency when using different robots and controllers. Finally, in Chapter 7, the commonalities between the work is presented, the implications of the findings in a wider context are discussed, and we propose future research directions to build on these findings.

1.4 Thesis Contributions

The approaches and results from this thesis contribute to both the body of work in understanding of locomotion, and to robot locomotion control.

Invariant Morphological Optimality: Our findings in biological and mechatronic humanoids showed a tendency towards optimality of humanoid locomotion, perhaps as a consequence of passive optimality of the morphology itself.

Simple Heuristics for Locomotion Control: Optimising for efficiency in either model-based, optimal control or machine learning controllers is difficult and expensive due to the nonlinearity of the energy-cost and the need for a global view of efficiency landscape. This study showed that despite this complexity, there are simple rules which capture the efficiency region and can be incorporated into controllers as a computationally rapid method of improving efficiency going forward.

Insight into Human Motion Heuristics: The vast state-space of actions available to humans is at odds with the ease and robustness by which they move. The heuristics and passive optimality that we propose in this research give us

insight into how this might be achieved, as a simple relationship patterns may be possible for humans to learn or memorise and could produce motion with minimal cognitive processing.

Optimal Regions of Attraction: The consensus of work suggests that human walking is energy-optimal. This may emerge as a result of explicit optimising for Cost of Transport or some metabolic rate; or as a by-product of another consideration, but full explorations of the locomotion landscape to investigate this have not been well studied. Our findings suggest that regions of attraction towards energy-optimal step locations in the humanoid morphology may play a part in this decision making. Further, this suggests that the humanoid morphology itself tends towards these regions of near-optimality and optimality as a consequence of their dynamics, reducing the need for explicit processing. Interestingly, we show that the regions emerge when explicitly minimising energy during optimisation 5 and even when this is not explicitly optimised for (Chapter 6).

Mechatronic Humanoids as Human Analogues: The mastery of robust and agile motion in biological humans cannot be directly observed at a policy level, but mechatronic humanoids share their morphology and we investigate their use as a biorobotic test-bed for research into hypotheses about human motion. We show that robots are a suitable tool for human locomotion research within strict and overt limitations presented by current technology.

Methodologies for Studying Efficiency: We present methodologies for studying the efficiency landscape of given locomotion parameters, processing results and producing efficiency landscape maps to analyse and characterise the patterns of efficiency on a global scale, and the relationship between efficiency and walking parameters.

1.5 Publications

Original contributions by the author are adapted to form part of this thesis:

1. **McGreavy, C.**, Yuan, K., Gordon, D., Tan, K., Wolfslag, W.J., Vijayakumar, S. and Li, Z., 2020, May. Unified push recovery fundamentals: Inspiration from human study. In 2020 IEEE International Conference on Robotics and Automation (ICRA) (pp. 10876-10882). IEEE. DOI:

10.1109/ICRA40945.2020.9196911

2. **McGreavy, C.** and Li, Z., 2022. Reachability Map for Diverse and Energy Efficient Stepping of Humanoids. IEEE/ASME Transactions on Mechatronics. DOI: 10.1109/TMECH.2022.3174961

In particular, Chapter 4 is based on work in Publication 1, and Chapter 5 is adapted publication 2, with Chapter 6 being prepared for journal publication.

The author has contributed to publications based on some of the work presented in this thesis:

Song, D.R., Yang, C., **McGreavy, C.** and Li, Z., 2018, November. Recurrent deterministic policy gradient method for bipedal locomotion on rough terrain challenge. In 2018 15th International Conference on Control, Automation, Robotics and Vision (ICARCV) (pp. 311-318). IEEE. DOI: 10.1109/ICARCV.2018.8581309

Yuan, K., **McGreavy, C.**, Yang, C., Wolfslag, W. and Li, Z., 2020. Decoding motor skills of artificial intelligence and human policies: A study on humanoid and human balance control. IEEE Robotics & Automation Magazine, 27(2), pp.87-101. DOI: 10.1109/MRA.2020.2980547

Wolfslag, W.J., **McGreavy, C.**, Xin, G., Tiseo, C., Vijayakumar, S. and Li, Z., 2020, October. Optimisation of body-ground contact for augmenting the whole-body loco-manipulation of quadruped robots. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 3694-3701). IEEE. DOI: 10.1109/IROS45743.2020.9341498

Triantafyllidis, E., **McGreavy, C.**, Gu, J. and Li, Z., 2020. Study of multimodal interfaces and the improvements on teleoperation. IEEE Access, 8, pp.78213-78227. DOI: 10.1109/ACCESS.2020.2990080

Triantafyllidis, E., Yang, C., **McGreavy, C.**, Hu, W. and Li, Z., 2020. Robot intelligence for real-world applications. IET Computing and Networks, UK.

Triantafyllidis, E., Hu, W., **McGreavy, C.** and Li, A., 2021. Metrics for 3D Object Pointing and Manipulation in Virtual Reality. IEEE Robotics and Automation Magazine. DOI: 10.1109/MRA.2021.3090070

Gordon, D.F., **McGreavy, C.**, Christou, A. and Vijayakumar, S., 2022. Human-

in-the-Loop Optimization of Exoskeleton Assistance Via Online Simulation of Metabolic Cost. IEEE Transactions on Robotics, 38(3), pp.1410-1429. DOI: 10.1109/TRO.2021.3133137

Triantafyllidis, E., Hu, W., **McGreavy, C.** and Li, Z., 2021. Metrics for 3D Object Pointing and Manipulation in Virtual Reality: The Introduction and Validation of a Novel Approach in Measuring Human Performance. IEEE Robotics & Automation Magazine, 29(1), pp.76-91. DOI: 10.1109/MRA.2021.3090070

Chapter 2

Prior Work

2.1 Template Models

The locomotion control problem is complicated by the fundamental imbalance of a bipedal system, such that it is prone to falling, and the computational complexity of the nonlinear forces during motion. Model-based approaches to locomotion are particularly vulnerable to this complexity, as the approach requires accurate simulation of the physical world to forward simulate the motion of the whole body.

To reduce complexity, a battery of models that simplify the dynamics of the humanoid morphology have been developed. These models simplify the computation of this complex motion, but limit how accurately highly dynamic motions can be modelled. Typically, these models condense the complexity of the links and joints of the high-DoF robots into a single point, or a low number of connected points, with known, simple dynamics. The Zero-Moment-Point (ZMP) [18] is the unique point at which the Ground Reaction Forces (GRF) acting at contact points would cause momentum or rotation (pitch or roll) at the contact points themselves. If the ZMP lies within the Support Polygon (SP), any momentum is being counteracted by equal rotational potential elsewhere in the SP, which thus the robot does not fall.

The ZMP in its original form is only applicable to flat-ground walking, and application to other surfaces is non-trivial, though extensions have been made to project 3D ZMP stability and friction constraints onto a 2D surface [19]. As

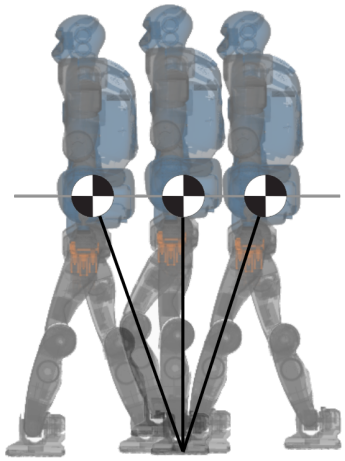

Figure 2.1: The Linear Inverted Pendulum Model approximates the mass of the robot as a single point with a constant height and telescopic leg.

a metric, the ZMP acts as both a binary representation of whether or not the system is stable, but also as a continuous means of measuring how stable it is – the closer the ZMP is to the centre of the SP, the greater the degree of stability and resistance to disturbance. The ZMP has a closed-form formulation and dynamics, so it is relatively simple to use it to compute footstep plans [20, 21] with the constraint that each new contact point is made before the ZMP is moved outside the stance foot, otherwise the ZMP would be outside the SP during the swing-phase, which results in imbalance.

The Inverted Pendulum (IP) [22] and Linear Inverted Pendulum (LIPM) [23] models simplify the whole body dynamics by representing it as a pendulum. Both approximate the entire mass of the body as a single point at the CoM which is connected to a pivot at the CoP, this connection represents the legs of the robot and the distance between the CoM and CoP. The IP models the relationship between the CoP and the CoM as a standard pendulum, accounting for the nonlinear rotation of the pendulum, and accommodates changes to CoM height and has been extended to model angular momentum as a flywheel connected to the CoM [24], which reintroduces some of the nonlinearity of full, whole body dynamics. The dynamics of the IP are rapid to compute and have been successfully integrated into control schema, both online [25, 26] and offline for planning or learning [27], but still have limitations in their applicability to highly dynamic whole-body movements.

The LIPM similarly models the dynamic relationship between CoP and CoM,

whilst making further linearising assumptions about CoM dynamics. Here, the CoM height is fixed, and the robot legs are modelled as a massless and telescopic link between the CoP and CoM. Further linearising the dynamics results in a greater modelling error, but results in computation speed fast enough to compute step plans and CoM trajectories online [28].

LIPM-based planners and controllers have been popular in the locomotion community thanks to rapid computation and reliability and forms the basis some effective stability for walking. The Capture Point (CP) [24], Divergent Component of Motion (DCM) [29] and the Extrapolated CoM [30] are similar methods derived from the LIPM formulation, which quantify the unstable aspect, or orbital energy of LIPM motion. In simple terms, these metrics represent the point in space, usually on the ground, which reduce the CoM velocity to zero if the CoP coincides with them. DCM was formulated for 3D space, such that it is not only applicable to flat ground walking, and the CP formulation has been modified to calculate capture regions [24], which are wider areas on which steps can be taken to bring CoM to a stop, and N-Step capturability, which is used as a method of quantifying the stability of the CoM in terms of the ability to come to a stop within N-steps [31]. CP, similarly to ZMP, is a useful, rapid method of step planning, as it has simple, known dynamics which are rapid to compute, and keeping the CP inside the SP is a sufficient condition for balance, such that this is balance is always maintained if this condition is met.

A plethora of other reduced models have been proposed which bridge the divide between LIPM and full models. Spring-Loaded Inverted Pendulums [32, 33], where the legs are modelled as springs to more accurately represent the energy profile of leg motion. Knees have been added to simple models to more closely represent the humanoid anatomy [34], point masses can be added to represent the swing and stance leg, [35] [36] or even the pelvis and upper and lower leg segments individually [37].

Typically, the output of these methods are time-sequenced CoM trajectories and the contact points, but since the solutions are in task space, they must be combined with Inverse Kinematic low-level controllers to realise the trajectories on the robot by tracking the point-mass trajectories and solve for joint positions.

Reduced models are a successful, simple method of reducing the vast state-action

space of the bipedal walking problem and have led to a great deal of progress in the field in enabling reliable, relatively safe walking methods for a range of robots [31, 38, 39] However, these methods produce highly inefficient and slow locomotion, so cannot be viewed as a sustainable component of the future of locomotion. Continual improvements to computing speed gives us the opportunity to begin moving towards more complete models for model-based control. Model-based controllers still very much have a place in locomotion in the meantime, but model-based feedback solutions which are rapid to compute whilst removing some of the barriers to more agile motion are needed to move the field forward.

2.2 Optimisation

Optimisation approaches can focus on different components of locomotion (discussed in 3.2), focusing on contact planning optimisation, which is then fed to a trajectory optimisation, then time-indexed trajectories are tracked by a whole-body controller, often Inverse Kinematics (IK) for position control or, Inverse Dynamics Quadratic Programming whole body torque control – more appropriate for dynamic motion and interaction.

Footstep planning optimisers solve for step planning sequences, steering and obstacle avoidance [40, 41] of the contact points and whole-body, multi-contact cases [42, 43]. Contact sequence plans are optimised based on some high-level goal, such as a desired walking speed or navigation through a complex environment via discrete search, or in some cases, continuous optimisation [44], but this cannot be used for non-continuous terrain, so is restricted to flat-ground walking.

Motion plans are then generated for the step-sequence, which most often uses CoM transition dynamics to produce time-indexed trajectories of the CoM and the end-effectors using reduced models.

Hybrid approaches combine some or all these phases into a single method. By including additional constraints/variables, the discrete aspects of locomotion can be folded into the optimisation formulation, making them contact invariant; Hybrid Zero-Dynamics (HZD) for cyclic walking [45] takes this approach, as does the contact invariant framework in [46], which adds additional optimisation variables to declare contacts active or inactive given their state.

The Hybrid components of motion can also be optimised by modelling centroidal and swing-foot reachability, representing the robot as a single rigid body [47], which allows motions and step plans to be solved whilst constraining the solution to feasible dynamic and friction bounds for N-contact points in complex terrain and highly dynamic motions. Alternatively, full dynamics models can be used to optimise motions, but at the expense of computational speed [48, 43].

Optimisation-based approaches produce highly agile, often dynamically feasible motions, but in most cases, the computational resource requirements of numerical solutions mean they take too long to converge to the final motion. The consequence of this is that motions are often computed offline and replayed on the robot using joint-angle tracking, making them vulnerable to changes in feedback and modelling errors. This class of solutions would benefit from greater knowledge of the underlying principles of locomotion to further reduce the vast state-action space and therefore decrease the required computational resources.

2.3 Machine Learning

Machine learning is a powerful tool for locomotion based on data-driven algorithms that do not necessarily require detailed dynamic models. Learning methods can be used to augment other approaches, such as learning the residual dynamics of reduced models [49, 50], learning the footstep placement and motion planning as decoupled problems [51, 52], and as an additional augmentation layer to work in parallel with model-based approaches [53, 54].

Machine learning is also uniquely placed to approach locomotion control in an end-to-end fashion; that is, an end-to-end solution receives sensory input from the robot and produces joint-level commands whilst having implicitly solved the contact planning, motion planning and control.

Learning locomotion has evolved quickly from 2D walking [55] to multi-expert approaches which are capable of producing a range of locomotion behaviours from a set of individually trained networks [56]. The most prominent method for learning is Deep Reinforcement Learning (DRL), which is trained over a large amount of trial-and-error episodes, guided by a reward function.

DRL methods can also be augmented with existing real-world data and an imita-

tion term in the reward function to match the motion of the expert and decrease the search space and thus training time [13]. Imitation data can be provided in the form of measured human experimental data [13], or even extracted from existing videos of human and animal motion [57].

A large amount of data, generated by training episodes, is needed to train such policies, so are primarily trained in simulation and the models implemented on the real robot. In DRL in particular, the difficulty in training in the real world is that a vast majority of the early iterations will fail drastically and lead to robot damage, and it is very difficult to reset the robot to a precise configuration given the stochasticity of real hardware. As a result, there is a body of research dedicated to resolving these sim-to-real problems in applying a network trained in simulation to a real world due to the modelling errors [58], although, in some work, sim-to-real is less of a barrier to implementation [56].

The result of the offline training is a real-time locomotion controller which can reactively respond to novel scenarios, and often solve contact planning, motion planning and control simultaneously. But, due to the all-in-one nature of these approaches, it can be difficult to provide ad-hoc goals to each component part, provide long-term goals, or modify the components for a particular task without costly retraining.

Although machine learning is a valuable tool for locomotion, there are some downsides: safety and verification cannot be obtained, as the output of the model is dependent on environmental conditions that cannot be foreseen in advance and there is no way of systematically testing and verifying the long-term performance of the network, which raises reliability concerns. This is more of a problem with bipeds than quadrupeds, as bipeds are inherently less stable and more likely to fall over given a certain set of conditions.

Additionally, in DRL for locomotion, the gaits produced by the training are often unnatural, arrhythmic and inefficient, even when trained on human motion data. This has huge implications for operational time for long-term deployment, even if the safety concerns are mitigated. These methods would benefit from a greater understanding of the relationship between open reward-function parameters that could improve the efficiency of the gait.

2.4 Bioinspiration and Biomimetics

The study of biological systems to inform robotics development has obvious benefits – the biological world is teeming with refined templates of complex, adaptable behaviours that would be invaluable to autonomous mobile systems to accomplish any task that can be similarly performed by biological beings.

Indeed, looking to nature for insights into engineering tasks (known as bioinspiration) has proven fruitful for robotics and can take many forms. The concept of neural networks, the basis of deep learning is based on the neurons in biological systems [59], similarly, reinforcement learning takes inspiration from human early-years development and learning [60]. Human motion, formulated mathematically using Inverse Reinforcement Learning (IRL), can produce walking [11] and running [61] behaviour, motion primitives can be extracted from human behavioural data, learned and linked together to form complex motion graphs [62], or a motion lexicon [63], muscle-model-based controllers for legged locomotion are used for bio-inspired walking [64, 10], and neuromorphic control seeks to perform locomotion using an expanded model of biological motor systems [65, 66].

A prerequisite for using bio-inspiration is access to sensing technology to measure and capture biological behaviour and processes. Stop-motion stills of running horses are an early example of the coupling between technology and bioinspiration 2.2 [67], where trip-wire triggered cameras recorded snapshots of horse locomotion in different gaits. The descendants of this work are seen today in locomotion research, as many quadruped controllers employ the patterns gaits described within [47, 68], and methods exist to specifically select such gait patterns [69, 70]. With sufficiently advanced technology which could reveal maximum resolution, neuron-level activation both spatially and dynamically, sensory processing and muscle activation, it may be possible to directly recreate some complex agile and robust behaviours of biological life in a mechatronic form.

However, current technology is far from this capability. Methods of measuring humans in motion are limited, and most recording takes place in either action space using optical or inertial motion capture, joint level motion capture with joint-angle flex sensors, and muscle activation sensing with Electromyography (EMG). Neuroimaging techniques have a very limited scope of application for in-situ motion research. Function Magnetic Resonance Imaging (fMRI) measures

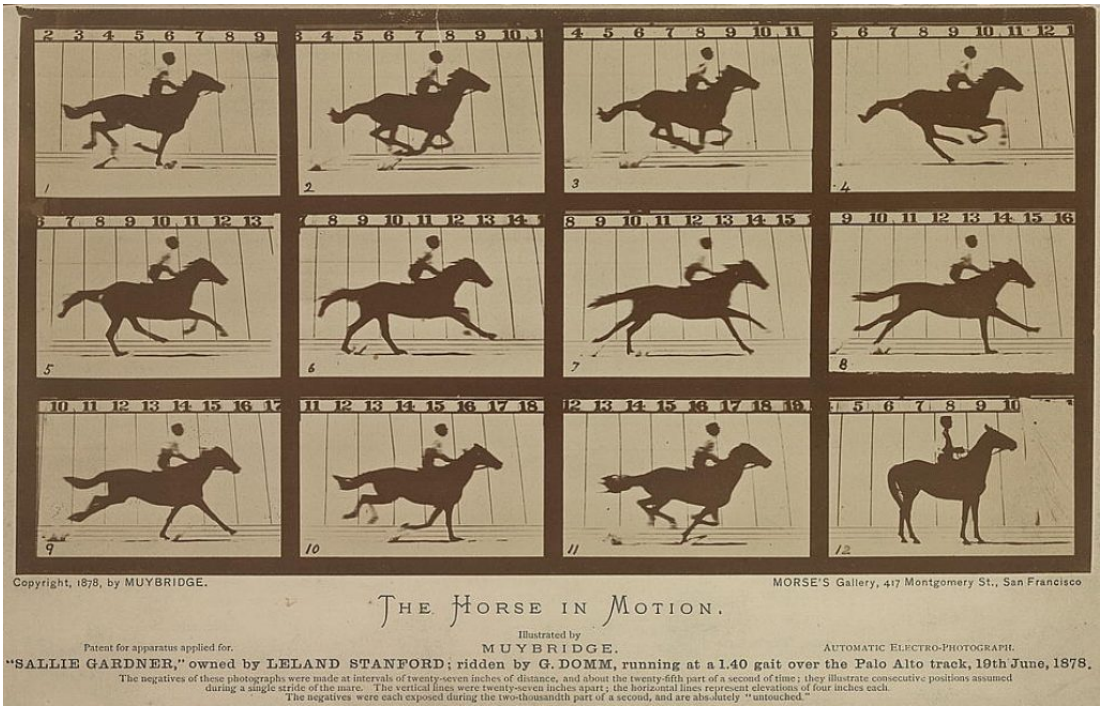

Figure 2.2: Early locomotion analysis used a series of trip-wire cameras at fixed distances to capture gait patterns and limb motion [1].

the dynamic change of cortical activation at good resolution, but due to the size of the machinery, is primarily limited to imagined motion [71] and other methods that are suitable for imaging during motion, such as EEG and PET [72], but have drastically lower resolution and is highly sensitive to individual differences such as age and pathology [73]. As such, high-resolution observation of in-situ control policies made by humans is out of reach in the near future.

As an alternative, we can use model the latest knowledge about biological systems and use the models to test hypotheses about control policies and behaviour. Modelling of complex systems and their use as methods of prediction and verification is widely used in the modern world, from modelling weather systems, planetary mechanics, and quantum mechanics, to computational neuroscience and evolution.

Biorobotics shares similarities with other types of modelling, in that existing processes and mechanisms, are being developed, but in biorobotics, the model itself is not an end-point of the process, it serves as a platform for hypothesis testing about biological systems, which may lead to further questions and development [17] and design iteration it non-standardised and difficult to predict [74].

Biorobotics takes a counterfactual reasoning-like approach to biology, where we can reason about the state of the world based on what effect a hypothesis being true would have on it [75].

The advantage of biorobotics comes from the level of control over experimental variables, the observability of internal processes and the repeatability inherent in mechanical systems. Morphologies and parameters can be precisely altered, we can both test hypotheses about biological processes, and, through the design process, highlight any practical issues or considerations with biological theory that may not be considered during purely mathematical modelling.

Biorobotics has seen great success in both the construction of robots as biomimetic analogues and successfully employing them to test hypotheses without direct access to the biological subject. Modelling invertebrate perception, behaviour [16], and locomotion [76] has been particularly successful. As part of biorobotics projects, hypotheses are posed, and robot models are constructed with the relevant aspects of the target organism. Neural circuits and gait transitions tested on a biomimetic salamander robot [9] showed that Central Pattern Generators (CPG) produced similar gait patterns to the biological host – CPG algorithms have since had a positive impact on legged locomotion [77, 78, 79].

In bipeds, passive dynamic walkers (Figure 2.3a) based on human anatomy are used to study locomotion in humanoids. Passive walkers were built to study the passive dynamics of the humanoid morphology [80, 2, 81], and have demonstrated the dynamics of the humanoid gravitates towards stable locomotion by virtue of their morphology, requiring no actuation to walk, on a sloped surface. Though these are accurate models of human dynamics that produce natural walking, they have not led to improvements in robotics hardware or control development. This is partly caused by their narrow purpose, need for detailed control in the real world [82], and difficulties in scaling up the design to a more general purpose machine, due to the technology not being available.

As a result, passive dynamic walkers which capture the humanoid morphology and dynamics well are almost different species to 'versatile' humanoids [83], which are much more common. The humanoid robots that are common in research today fall into this 'versatile' category (Figure 2.3b), in that they share the human morphology, but a trade-off is made between the accuracy of this similarity, and

the practicality of the robot as a useful platform and tool, given current technology. Though the anatomy, mass distributions, and muscle actuation are very well studied and have been very well refined for decades [84], the technology to replicate this whilst maintaining the generalisability of movement associated with humans is yet to materialise. This disparity makes studying humanoid motion with biorobotics a different problem to smaller beings for which appropriate technology exists to model them.

Even though mechanical humanoids do not match their biological counterparts exactly, there is a great deal of value in exploring bioinspiration and biorobotics with current technology. Within an existing humanoid lies the same morphology as the human, and while the dynamics may not match exactly, they are still a close analogue and provide a great test-bed for hypothesis testing with full observability and control which, as stated earlier cannot be gained with current technology from humans. Further, as stated in [85], to investigate human walking systematically, participants would need to be trained to walk and run differently and repeatability, which is not possible in stochastic biological systems. In this thesis, we investigate both bioinspiration and biorobotics to study efficiency, optimality and dynamic invariance in the humanoid morphology.

2.4.1 Passive Dynamic Walkers, Versatile Humanoids and Human Simulation

It is important to motivate the use of humanoids as a basis for bio-inspiration and as a bio-robot. As stated above, passive walkers are more similar to human anatomy than 'versatile' humanoids [83] that are most widely used today, so why not use passive walkers for studying bipedal locomotion? While passive walkers have successfully demonstrated control-free locomotion, their findings have not been translated to humanoid robot developments in a lasting way and have yet to find broad application, and robots with control are preferred [82]. These mechanisms, though impressive lack any wider application either in robotics or biorobotics, due, in part to the large technology gap between making more biologically accurate robots and making them versatile.

'Versatile' humanoids lie on the other end of the spectrum, where bio-mimicry is a desired outcome of the design, but is only implemented as far as current

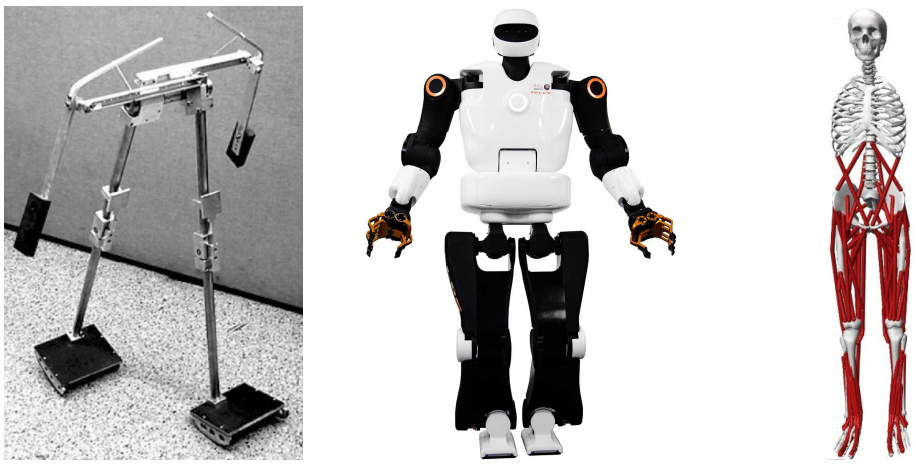

(a) (b) (c)

Figure 2.3: Examples of potential biological analogues for humans: (a) passive walking robots [2], (b) 'Versatile' TALOS humanoid robot [3] and (c) and biomechanical models of humans [4].

technology allows. Trade-offs in terms of biological accuracy are made in favour of versatile motion and control. The level of control offered by versatile humanoids makes them more suitable for many biorobotic scenarios than passive walkers, as changing the gait requires physical reconfiguration of the robot.

Naturally, more biologically similar humanoids would be preferred, but given that it is more likely that developments in versatile humanoids will be geared more towards biological accuracy than passive walkers will be made more versatile, it is valid to focus on versatile humanoids to guide their development towards a more biological standing. One of the outcomes of biorobotics is gaining new insights into the robots themselves [17], using standard humanoids as biorobotics can help prioritise the development of new features over others in terms of their potential gain. For example, showing the gains that can be made from a compact toe-off morphology with a high power-to-weight ratio and an appropriately advanced control system might yield more benefit that putting effort into biological muscle actuation.

Similarly, we also have the option of using simulated human models (Figure 2.3c) as the basis for our study in biorobotics, and even use neuromorphic control to fully mimic human motion and dynamics. This would be an exciting avenue of study, but control of human models in fully dynamic simulation without reference

motion is still a difficult problem [4] and the field is in its infancy. In a recent competition [4], this was achieved using reinforcement learning [86], but with a model with a reduced number of joints, fused torso and no arm segments.

The work presented in this thesis is focused primarily on the dynamic properties of the full, multi-body humanoid morphology, not limited to the subset of morphology which biological inhabit, so we feel confident that 'versatile' humanoid robots fill that role, 'Versatile' humanoid robots are also the dominant platform for research and implementation for tasks involving this morphology and are intended for use in human environments, thus by adopting this focus, we maximise the potential applicability of our findings to both fields.

Chapter 3

Background

Necessary context for the work presented in the remainder of the thesis is explained and discussed in this chapter. The general problem of bipedal locomotion is discussed first, followed by the anatomy of the gait cycle and definitions of the specific aspects of robot control. Legged locomotion is a method of mobility concerned with switching contact locations from one position to another and moving the rest of the body appropriately to achieve some movement goal. Complexity lies in each part of this problem for general legged locomotion, but is particularly pronounced in the case of bipedal morphologies.

3.1 The Bipedal Locomotion Problem

The humanoid form uses two legs, and much of its mass is concentrated in the upper body, which makes them inherently unstable, unlike quadrupeds, whose mass is centred low to the ground between four contact points. Bipeds are top-heavy, so it is much harder to statically balance.

The reason that humanoid dynamics and their forward simulation is difficult in model-based control methods is that humanoid robots are typically made up of at least 32 DoFs, each with their own dimensions, mass, inertia, dynamic effects and mass distribution according to the equations of motion for a rigid body system:

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}) = \boldsymbol{\tau} \tag{3.1}$$

Where $\mathbf{q}, \dot{\mathbf{q}}$ and $\ddot{\mathbf{q}}$ are vectors of joint positions, velocities and accelerations respectively (of around $\mathbb{R}^{32}$ in a full humanoid) $\mathbf{M}$ is the inertia tensor (in $\mathrm{kg}\cdot\mathrm{m}^2$), $\mathbf{C}$ the nonlinear forces acting on the system (1 to 100s Newtons N), such as Coriolis and centrifugal force and $\mathbf{g}$ is the gravity vector ($9.81\,\mathrm{m/s}^2$).

Each joint is a connection between links, which, again have their own dynamic and kinematic properties. Since this is a floating-base system, the motion of each link has an effect on each of the other links and joints. The precise nature of the contact with the ground also has dynamic effects on the links and joints of the robot. Contact dynamics are notoriously difficult to model, and any error in this modelling, or even slight imperfections in the modelling of the robot itself propagates and accumulates into a much larger modelling error, which is then exacerbated further when these models are projected forward in time to compute future states, the further forward the projection, the greater the modelling error, and thus the inaccuracy of the prediction.

Modelling errors are less severe and less impactful for other types of robotics, like fixed-based manipulators, as the joints can be frozen in an emergency, and the robot will remain in place. Humanoids, however, are inherently unstable and the physical consequence of their morphology is for the mass in the upper body to adhere to gravity and move closer to the ground, or to put it in other terms: humanoids want to fall down, and they will if the correct control commands are not executed to maintain balance.

Additionally, the state-action pairing can only be viewed on the trajectory level for natural walking. A single set of joint-level actions has downstream consequences for the whole stepping trajectory, so when planning joint commands in the early stages of the stepping trajectory, a seemingly valid and feasible command may lead to a situation where zero valid commands are available later on in the step, so each command must be computed with a view to its impacts along the time-horizon.

All these factors combine to make humanoid locomotion a very complex problem, which is very difficult to model and control.

3.2 Step Planning, Trajectory Planning and Control

There are two parts to the locomotion problem: deciding where to step and controlling the body to reach the step location, which makes locomotion a hybrid control problem, since the steps are discrete locations in space but the body movement to reach these locations is continuous and planned at the joint Level.

Reaching step locations can be further broken down into 2 sub-categories: computing mid-level trajectories (CoM, swing foot), then mapping those trajectories into the robot to be tracked by a position, or torque control algorithm to convert the trajectories into joint actuation. Depending on the control method, planning and control can be computed at the same time or decoupled, but in many cases, when they are computed together, a feed-forward model is used for the control, such that control may need to be computed again at runtime.

But it is clear that they are inherently linked - the body reachability must be considered when planning step positions and the step position is the inherent goal of the control, so is vital to control computation. This is part of what makes me bipedal locomotion so difficult, since the already complex control of the morphology must be projected forward in time to calculate where the body will be at the next step.

Part of what we hope to learn from humans is how they manage this complexity and how they incorporate this into their locomotion control policy. Although humans are well known to use feed-forward planning in their movement, it seems unlikely that the speed, reactivity and agility of locomotion would be possible if explicit forward simulation of their whole body and environment was required at a high frequency.

3.3 Static and Dynamic Walking

It is important to make the distinction between static and dynamic walking. In static walking, the projection of the centre of mass always lies within the support polygon of the robot's contact points such that the robot would not fall if it stops mid-motion for any reason. In dynamic walking the CoM can move outside the support polygon, which makes the motion more unstable and difficult to control,

but this type of locomotion has many advantages including faster walking, more efficient walking and greater contact switching reach (stepping further). Static walking is used in many robotic locomotion methods, as it simplifies the dynamics of the problem by minimising the nonlinear forces involved in rapid motion and reduces the computational cost, whereas humans mostly walk dynamically.

3.4 Open Parameters in Locomotion

Part of the complexity of the locomotion problem is due to the number of decisions that need to be made, the size of their respective decision spaces, and the complex interactions between them. Key components of locomotion include:

Gait Pattern: Static Walk, Dynamic Walk, Trotting, Running, Hopping, Jumping. This dictates, or reflects the footstep sequence

Contact/Footstep Selection: What is the translation and orientation of the next contact point?

Step Timing/Speed: How long should a swing-phase take/what is the speed of the swing foot during transition? How long should each foot stay in contact for? The longer contact is maintained, the more the walking speed decelerates, but also becomes more unstable

Swing Motion: Constrained by step timing, how should the swing foot move towards the next contact point? What is the maximum height of the swing foot at its apex? What acceleration profile should be used to reach the next step goal?

Whole-body Transition: Most of the mass of a bipedal robot is concentrated in the upper body, is rigidly connected to the rest of the robot and has a 6-DoF floating base. The motion of this mass must be solved with consideration of its dynamics and its connection/control over the legs and contact points

Control: Given a whole-body motion, how can that be achieved at the joint level?

Reachability/Feasibility: While solving the other issues, it is also necessary to consider the capabilities of the robot and the constraints of the environment, i.e. can the robot produce enough torque to move the whole body to a given footstep

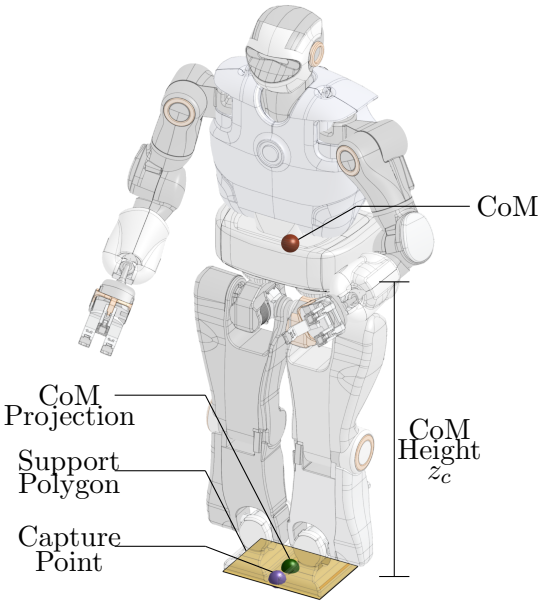

Figure 3.1: Relevant metrics and how they relate to the robot morphology.

location in the given time? Does the robot have enough kinematic reachability to move to a new location?

Even if we pre-emptively calculate and enforce feasibility constraints, the action space for each open parameter is vast, and the knock-on effect of one action on another is not well understood.

3.5 Definitions: Models and Terms

3.5.1 Linear Inverted Pendulum Model

The LIPM [87] is an equation of motion describing CoM motion when controlled by a contact point, but constrains the CoM to a fixed height, ignoring the non-linearity of vertical motion z_c (Figure 3.1):

$$\ddot{x} = \frac{g}{z_c} x \tag{3.2}$$

Where the acceleration of the CoM $\ddot{x}$ is determined by the natural frequency of the pendulum, calculated by dividing the gravity constant g by the fixed height z_c, multiplied by the current CoM position, x.

The primary method of control in the LIPM is switching the contact position p_x represented by the CoP, at the base, altering the formulation:

$$\ddot{x} = \frac{g}{z_c}(x - p_x) \tag{3.3}$$

3.5.2 Capture Point

LIPM equations of motion in 3.3 can be used to calculate the orbital energy of the system and, by extension, calculate the step location p_x required to reduce that orbital energy to zero, thus bring the CoM to a velocity of zero. This step location – the capture point (p_c) – is determined by the equation:

$$p_c = \dot{x}\sqrt{\frac{z_c}{g}} \tag{3.4}$$

3.5.3 Definitions of Terms

Centre of Mass (CoM) is the mean spatial position of the mass distribution across a body, or assembly of adjoined links (Figure 3.1).

Centre of Mass Projection, the CoM as projected onto the current support surface (Figure 3.1).

Support Polygon (SP) is the convex hull formed by the points of the body in contact with the ground (Figure 3.1).

Double Support Phase, where both feet are in contact with the contact surface and contributing to balance.

Single Support/Swing Phase, when only one contact is engaged with contact surface and the other is moving towards a new step position.

Swing Foot/Leg, during single support phase, the foot/leg that is not currently in contact and is moving towards the next contact point.

Stance Foot/Leg, during single support phase, the foot/leg currently in contact and responsible for balance during the swing phase.

Swing Apex, the maximum height of the swing foot during the swing motion.

Chapter 4

Optimality in Human Balance Recovery

We began investigating push recovery as a specific aspect of locomotion because the overall process is incredibly intricate, and we aimed to narrow the scope of our study and to reduce the complexity of analysing human motion. In this chapter, we collect data from human push recovery experiments, process this data and model the behavioural data in order to fit a simple controller to the human observation.

Push recovery is segmented into a set of discrete strategies, each with a different method of control. In robotics, strategies are used in two ways: either as standalone push recovery controllers [88, 24], which are useful against a small range of pushes, or combined into a single controller. Using a single controller which can combine all strategies is effective against a wider range of pushes [89, 90], but this often means there are sudden switches from one strategy to another [91]. Transferring between strategies, especially when each uses a different controller, can lead to failure cases. A second drawback is that users must select how much each action emerges during recovery [89]. Even with careful tuning, failures can occur when the tuned ratios do not match the needs of the recovery.

In contrast, humans combine strategies into continuous motions [30], combine them in many ways [92], and switches between them occur rapidly [93]. Thus, it is unlikely that humans decide how strategies are used ahead of time or control each strategy differently, especially when sensory delays are considered.

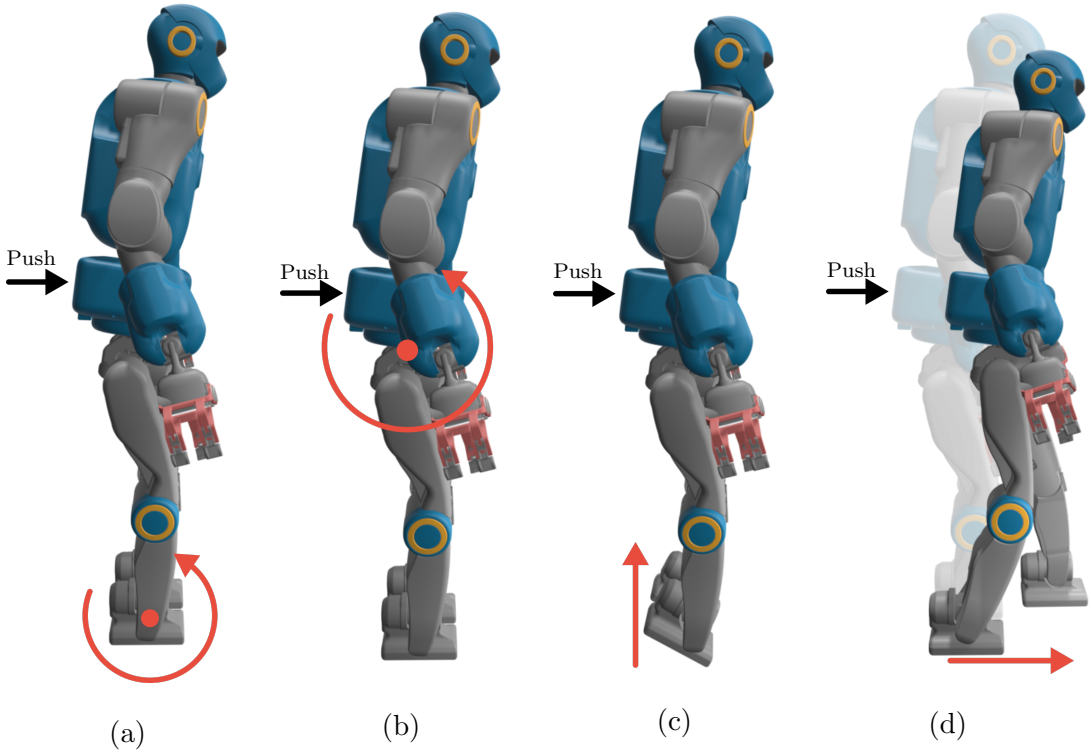

Figure 4.1: Core strategies used in push recovery: (a) Ankle Strategy, (b) Hip Strategy, (c) Toe-Strategy, (d) Step Strategy

Motivated by this human behaviour, we hypothesise that the discrete actions share fundamental principles which can be used as a unified controller across all strategies. For robotics this would mean tuning the emergence of each strategy is no longer required and reduces the need to use different controllers for each strategy.

4.1 Push Recovery Methods

Ankle Strategy (Figure 4.1a)regulates ankle torque to modulate the Centre of Pressure (CoP) to affect CoM motions for small pushes. Observed in humans [92] and used legged robots [88], this strategy is effectively modelled by the Linear Inverted Pendulum Model [94].

Hip Strategy (Figure 4.1b) applies torque at the hip joint to induce angular momentum around the CoM. This is observed in humans [95] and applied to robotics [88]. Though some work shows it has limited effect on recovery [96].

Toe-Lift Strategy (Figure 4.1c)generates upward CoM motions, and is prominent

in humans [97]. The strategy has been applied in Robotics by controlling the under-actuated motion [15].

Step Strategy ((Figure 4.1d)) modulates the Support Polygon (SP) halting the CoM. Its effectiveness especially during large disturbances has been shown both in humans [30] and robots [24].

Discrete strategies can be combined using Model Predictive Control (MPC) [89], but parameter tuning dictates the desired emergence of strategies. Linear Quadratic Regulation [90] can effectively implement ankle and hip strategies by minimising joint deviations but does not include step strategy. Using a Proportional Integral Derivative as a CoM reference for Quadratic Programming (QP) whole-body control shows ankle and step strategies can emerge naturally [98]. Strategies can also emerge during learned locomotion [99].

As push recovery strategies appear simultaneously in humans [100, 93], we investigate this behaviour aiming to find shared principles of push recovery strategies. Applying this principle to robotics could reduce the number of parameters needed to tune recovery behaviour, since strategies would emerge naturally. To investigate these common principles, we first collect data of humans recovering from pushes. Data will be analysed to extract any common principles which may exist between the strategies. Lastly, we investigate how any principles might be adapted to use in robotics.

4.1.1 Contributions

This work presents the following contributions:

1. **Experimental Design (Section 4.2)**: Experimental design extracting useful physical quantities to identify recovery strategies in humans
2. **Strategy Classification (Section 4.3)**: Classification criteria of human push recovery strategies
3. **Invariant Minimum Jerk Principle (Section 4.6.2)**: Evidence of minimum jerk regularisation as core strategy during push recovery.
4. **Minimum Jerk Controller Formulation (Section 4.6.2)**: Control design of a minimum jerk controller that resembles human CoM behaviour

1. R. Acromium
2. L. Acromium
3. Clavicle
4. R. ASIS
5. L. ASIS
6. V. Sacral (Coccyx)
7. R. Thigh Upper
8. R. Thigh Lateral
9. R. Thigh Medial
10. L. Thigh Upper
11. L. Thigh Medial
12. L. Thigh Lateral
13. R. Knee Lateral
14. R. Knee Medial
15. L. Knee Medial
16. L. Knee Lateral
17. R. Shank Upper
18. R. Shank Lateral
19. R. Shank Medial
20. L. Shank Upper
21. L. Shank Medial
22. L. Shank Lateral
23. R. Ankle Lateral
24. R. Ankle Medial
25. R. Heel
26. L. Ankle Medial
27. L. Ankle Lateral
28. L. Heel
29. R. MTP5
30. R. MTP1
31. L. MTP1
32. L. MTP5

Figure 4.2: Motion capture marker placements on subjects on body landmarks

5. **Application of Recovery Controller for Robotics (Section 4.7.3)**: Application of extracted core principles to robotics controller with better performance.

4.2 Experimental Setup

4.2.1 Human Motion Capture

Optical motion tracking was used to recorded the movement of the subjects, using the VICON motion tracking system. Motion tracking cameras were positioned around the experimental space, illustrated in Figure 4.3, to provide full visibility of the subject and minimal obscuration by the experimenter. Reflective markers were placed on body landmarks according to a skeleton template, as shown in Figure 4.2. The template focuses on standard bony landmarks to ensure consistent and accurate tracking across subjects. A skeleton motion tracking model is scaled to each subject, and the 3D spatial data is recorded, shown in Figure 4.4.

A template of human body mass distribution was scaled to each participant in OpenSim [101]. This was used to calculate CoM position, joint angles, positions and body dimensions used to calculate angular momentum [12].

4.2.2 Disturbance Force Application and Sensing

A mounted `Optoforce` 6-axis force/torque sensor (Figure 4.5a) was used to push subjects and measure the applied force.

4.2.3 Experimental Design

Each participant underwent 15 push trials. Instructions were to try to return to the initial position after the push. If a step was required, they were told to come to a stop and not attempt to return to the initial position. Illustrated in Figure 4.6, participants were pushed at the coccyx using the force/torque device. Pushes, with magnitude and timing unknown to the subjects, were applied after 2-3s of quiet standing followed by a recovery action of the subject. To ensure purely reactionary responses without sensory cues, pushes were applied by experimenters. The magnitude of the impulses were varied and ranged from 12.1 N·s to 55.5 N·s, with a mean and standard deviation of 35.9 N·s and 10.3 N·s.

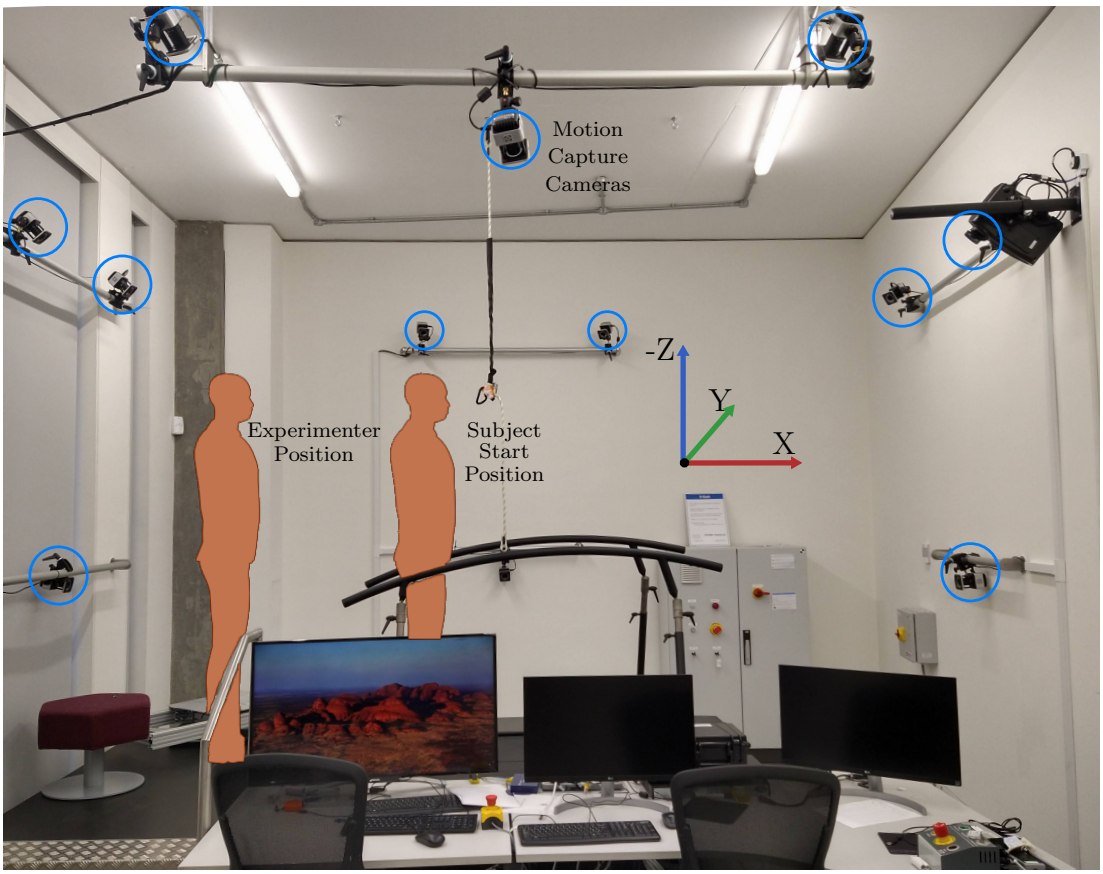

Figure 4.3: Experimental setup, showing instrumented treadmill, motion capture cameras, experimenter and subject positions, and treadmill frame of reference.

4.2.4 Data Consideration and Post Processing

Recovery begins after the push force is removed and ends when stability is regained. Motion capture data is trimmed to remove movement when force sensor readings are below 2 N and after recovery has ended. Only sagittal movement is considered as all pushes were performed in the sagittal plane. Velocity, acceleration, and jerk were obtained through differentiation of marker positions from motion capture. CoM states are normalised by leg length and are unitless. Low-pass filtering was applied to de-noise all data (4th order Butterworth, cut-off frequency at 6 Hz).

4.2.5 Subjects

We collected 60 trials from 4 participants recruited from the University of Edinburgh. Ethics approval was gained. Each subject gave written informed consent before testing. All ethics related forms are included in Appendix D.

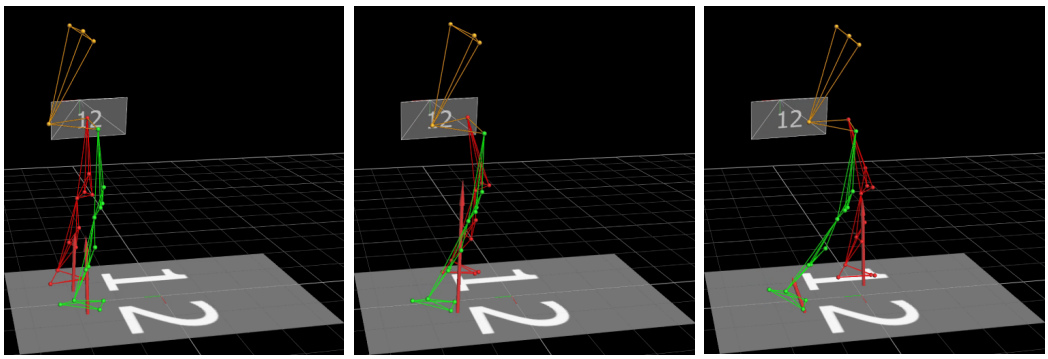

Figure 4.4: Snapshots of captured motion during push recovery.

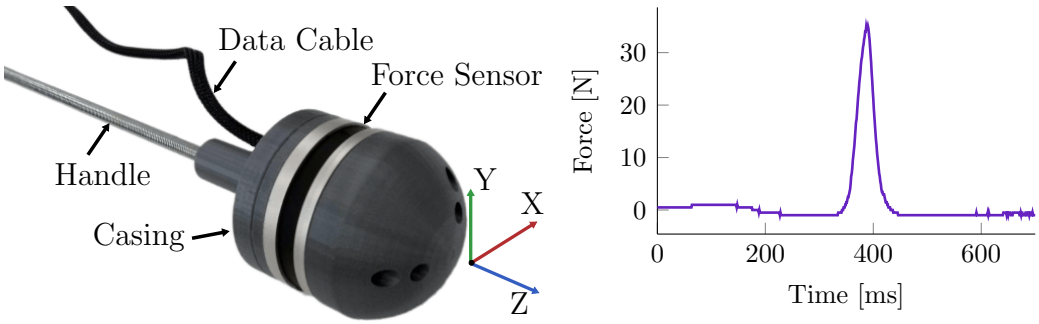

(a) (b)

Figure 4.5: A pole equipped with a force sensor was used to disturb the balance of participants: (a) pole and force sensor assembly, with coordinates shown; (b) example force readings from force sensor Z-axis (flipped to show force as positive)

4.3 Strategy Identification Criteria

We define *Control Actions* as the active component for recovery, e.g., CoP modulation, and use the term *Strategies* as label for an entire recovery trial. *Strategy* labels denote the highest *Control Action* that was used during a trial. To identify Strategies in the data, we determine the Control Actions from a set of criteria and thresholds (Table 4.1).

Since multiple recovery strategies simultaneously, we use the nomenclature of *Control Actions* and recovery *Strategies* to distinguish between discrete methods of recovery and a label for an entire recovery trial respectively. The discrete push recovery actions discussed in Section 4.1 are all based on a single active component. For example, the active component of the ankle strategy is manipulating the CoP. The criteria in Table 4.1 can then be used to identify which control actions are active at each timestep of a recovery trial. *Strategy* labels denote

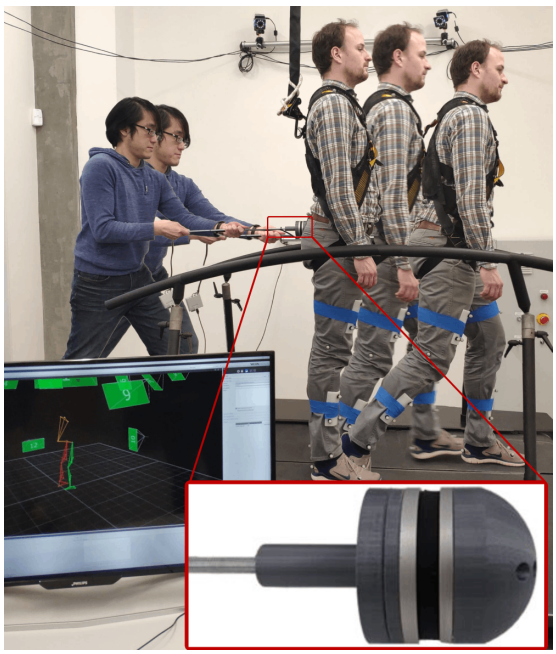

Figure 4.6: Sequence of motion during a single push recovery trial.

the highest *Control Action* that was used during a trial. Control Actions are ranked by the maximum normalised impulse rejected during each trial (Table 4.2). By making this distinction, we can qualitatively investigate how humans perform push recovery. This will add to the evidence that humans are unlikely to pre-select discrete recovery motions or use different policies to control them. The mean CoM trajectories for each Strategy are plotted in Figure 4.10.

4.4 Relationship Between Control Actions and Strategies

The dependence of control actions on the push magnitude are shown in Figure 4.7. Since subjects cannot predict the push force in advance, they gradually employ increasingly more effective techniques: lower ranked actions are used for smaller pushes and more actions are incorporated as the magnitude increases. For large pushes, weaker actions are skipped completely, and subjects immediately use higher ranked actions. The further the CoM diverge from the equilibrium (*i.e.* quiet standing), the stronger the Control Actions (blue boxes) become (Figure 4.8).

Table 4.1: Threshold values for determining active strategies.

Action	Criteria	Threshold	Human Max	Unit	Method
CoP	$p_{\text{heel,z}} \leq \delta_{\text{heel,z}}$	$\delta_{\text{heel,z}} = 0.01$	Variable	m	N/A
Angular Momentum	$L_t > \delta_t$	$\delta_t = 0.1$	0.49	s(Normalised)	Torso sway with feet fixed to ground
CoM Height	$\dot{z} > \delta_{\dot{z}}$	$\delta_{\dot{z}} = 0.02$	0.1269	m/s	Stand on toes as quickly as possible
Support Polygon	$p_{\text{toe,z}} > \delta_{\text{s},v_{\text{L,R}}} > \delta_{\text{s}}$	$\delta_{\text{s}} = 0$	N/A	N/A	N/A

Notation: $p_{\text{heel,z}}$ Heel height, L_t Hip Angular Momentum, $\dot{z}$ CoM Velocity (Z Axis), $p_{\text{toe,z}}$ Toe Height, $v_{\text{L,R}}$ Foot Velocity (X Axis) $\delta_{\text{heel,z}}$ Heel Height Threshold, δ_t Hip Angular Mom. Threshold, $\delta_{\dot{z}}$ CoM Height Threshold, δ_s Step Threshold

Table 4.2: Max normalised impulse observed in each action.

Rank	Action	Impulse [ms]
1	CoP Modulation	52
2	Angular Momentum Modulation	59
3	CoM Height Modulation	72
4	SP Modulation	77

4.5 Stability Regions

Through calculating a set of stability regions, which represent which actions can affect the CoM in a given state, we can identify their effectiveness and when they are used. We calculate these regions for humans and for an example robot with robot values using the specifications of the Valkyrie robot and human values obtained experimentally (Table 4.3), where p_{min} and p_{max} are the limits of the SP, θ_{max} and τ_{max} is the maximum angle and torque of the hip joint, $v_{z,max}$ and z_{max} the maximum vertical movement speed and position during toe-lift, and v_{min} and v_{max} are the minimum and maximum step speeds.

By adding physical constraints to a subject, we sequentially isolate the effect of each action. CoP Modulation parameters are obtained from the mean foot length among participants. The Angular Momentum Modulation constraints τ_{max} are determined by fixing the subject's heels to the ground enabling only hip recovery behaviour and increasing push forces until recovery using the hip was no longer possible. CoM Height Modulation parameters were obtained as in Section 4.3. For identifying the constraints during Support Polygon Modulation increasing pushes were applied until a flight phase was needed to recover.

Stability regions are defined at the CoM initial state $\boldsymbol{x}_0 = [x_0, \dot{x}_0]$ and are calculated for each action. An action can stabilise the system if $\boldsymbol{x}_0$ is in its stability region (Figure 4.9).

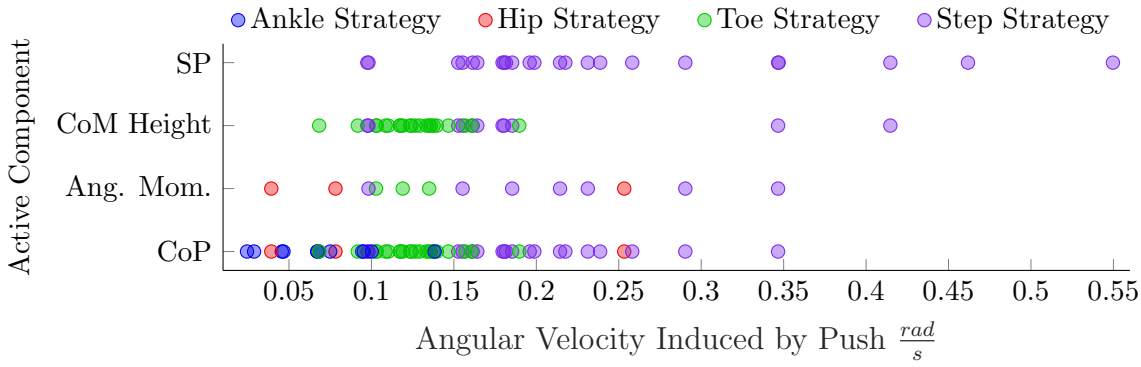

Figure 4.7: Active components of strategies (y-axis, circles) appear at different frequencies depending on the angular velocity (CoM w.r.t the ankle) caused by a push. Each colour shows which main strategy is used in that trial. Vertically aligned circles also represent a single trial, though connecting lines were removed for visibility. Push magnitude is represented by the caused angular momentum around subjects' ankles.

4.5.1 CoP Modulation

To remain stable whilst balancing, the CP must lie within the Support Polygon:

$$p_{\min} \leq x + \dot{x}\omega \leq p_{\max}, \tag{4.1}$$

where $\omega = \sqrt{g/z_c}$, and foot dimension (support polygon bounds) $p_{\min}, p_{\max}$. This bound is more practical than the ZMP criterion [102], as the CoM must be within the SP at the end of the recovery (at $t \to \infty$). The CP can then be controlled by modulating the CoP within the SP using the dynamics of the LIP model [94].

4.5.2 Angular Momentum Modulation

This action allows additional control via torque around the CoM, thus expanding the Capturability Region [24]:

$$p_{\min} - \alpha \leq x + \omega\dot{x} \leq p_{\max} + \alpha, \tag{4.2}$$

with $\alpha = \frac{\tau_{\max}}{\beta^2 mg}(\beta - 1)^2$, $\beta = e^{\omega T_{\max}}$, mass m, gravity constant g, maximal torso actuator torque $\tau_{\max}$, $\omega = \sqrt{g/z_c}$, maximum time $T_{\max} = \sqrt{4I/\tau_{\max}(\theta_{\max}) - \theta_0)}$, inertia I, maximal torso angle $\theta_{\max}$ and starting angle θ_0.

Table 4.3: Physical properties required for calculating balance limits and their associated boundaries.

	CoP Modulation		Angular Momentum Modulation		CoM Height Modulation		Support Polygon Modulation		
	$p_{\min}$ [m]	$p_{\max}$ [m]	$\theta_{\max}$ [rad]	$\tau_{\max}$ [N·m]	$v_{z,\max}$ [m/s]	$z_{\max}$ [m]	Step Length [m]	$v_{\min}$ [m/s]	$v_{\max}$ [m/s]
Human	0.10	0.17	0.27	3.01	0.13	0.13	0.5	0.1	3.95
Robot	0.12	0.19	0.66	150	0.10	0.07	0.25	0.1	3.00

4.5.2.1 CoM Height Modulation

This control action increases virtual leg length, e.g., through toe-tilting, and provides a force f perpendicular to the CoM velocity that reduces the horizontal velocity $\dot{x}$. To find the Stability Region for toe-tilting, we need to analyse whether horizontal CoM velocity can be reduced to zero at the edge of the foot.

The horizontal velocity $\dot{x}_0$ is induced by external pushes, while the vertical velocity $\dot{z}(x_0, \dot{x}_0)$ can be added through CoM Height Modulation. We assume that CoM motion will follow a straight line from initial condition $[-x_0, z_c]$ to endpoint $[0, z_c + \Delta z_{\max}]$. This over-approximates the Stability Region by assuming that the required force f can be generated at all times. For a more conservative estimation, a stability margin for the endpoint can be set instead of letting it be at the edge of the foot.

To achieve a straight line, a vertical velocity $\dot{z}$ needs to be set:

$$\dot{z}(x_0, \dot{x}_0) = \sqrt{\frac{\Delta z_{\max}^2 \dot{x}_0}{(x^2 + \Delta z_{\max}^2)(1 - \frac{\Delta z_{\max}^2}{x_0^2 + \Delta z_{\max}^2})}}, \tag{4.3}$$

while not exceeding the physical capabilities of the robot:

$$\dot{z}(x_0, \dot{x}_0) \leq \dot{z}_{\max}. \tag{4.4}$$

Due to f being perpendicular to the CoM velocity, we can use the energy balance to compute velocity at the foot edge. At this point $E_{\text{kinetic}} \leq \Delta E_{\text{potential}}$ must hold with $E_{\text{kinetic}} = \frac{1}{2} m(\dot{x} + \dot{z}(x_0, \dot{x}_0))^2$, and $\Delta E_{\text{potential}} = mg\Delta z_{\max}$. If this constraint is not met, the robot will pivot around the toe and fall.

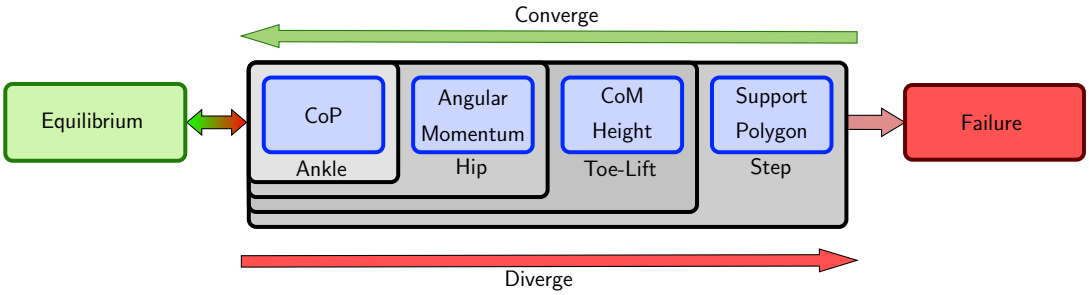

Figure 4.8: Relationship between control actions and strategies. As a system diverges from the equilibrium, a sequence of control actions (blue) will be activated. A strategy is defined by the highest active/dominant control action (grey gradient).

4.5.2.2 Support Polygon Modulation

The stability region is gained using the controller proposed in Section 4.6.2.1 with a given initial CoM state. If for the given initial condition, the robot can halt using the given constraints, then the initial condition is in the stability region.

Figure 4.9 shows these stability regions in relation to human CoM trajectories and Control Actions. The approximate stability regions are a close match to the Actions (coloured dots) shown in human movement. In the non-stepping cases, the CoM remains in the original support polygon, whereas in the stepping case, the CoM diverges to a new SP.

4.6 CoM Behaviour During Recovery Strategies

Following, we present a single control principle matching trajectories across push recovery actions and strategies without a switching mechanism between Control Actions.

4.6.1 CoM Modelling

From our human motion analysis, a dominating effects of the CoM dynamics was observed. Hence, we use a point mass model which is in agreement with past study [103] to analyse human CoM data, and to find the fundamental principles between the strategies:

$$\frac{\mathrm{d}^3}{\mathrm{d}t^3}x = u, \tag{4.5}$$

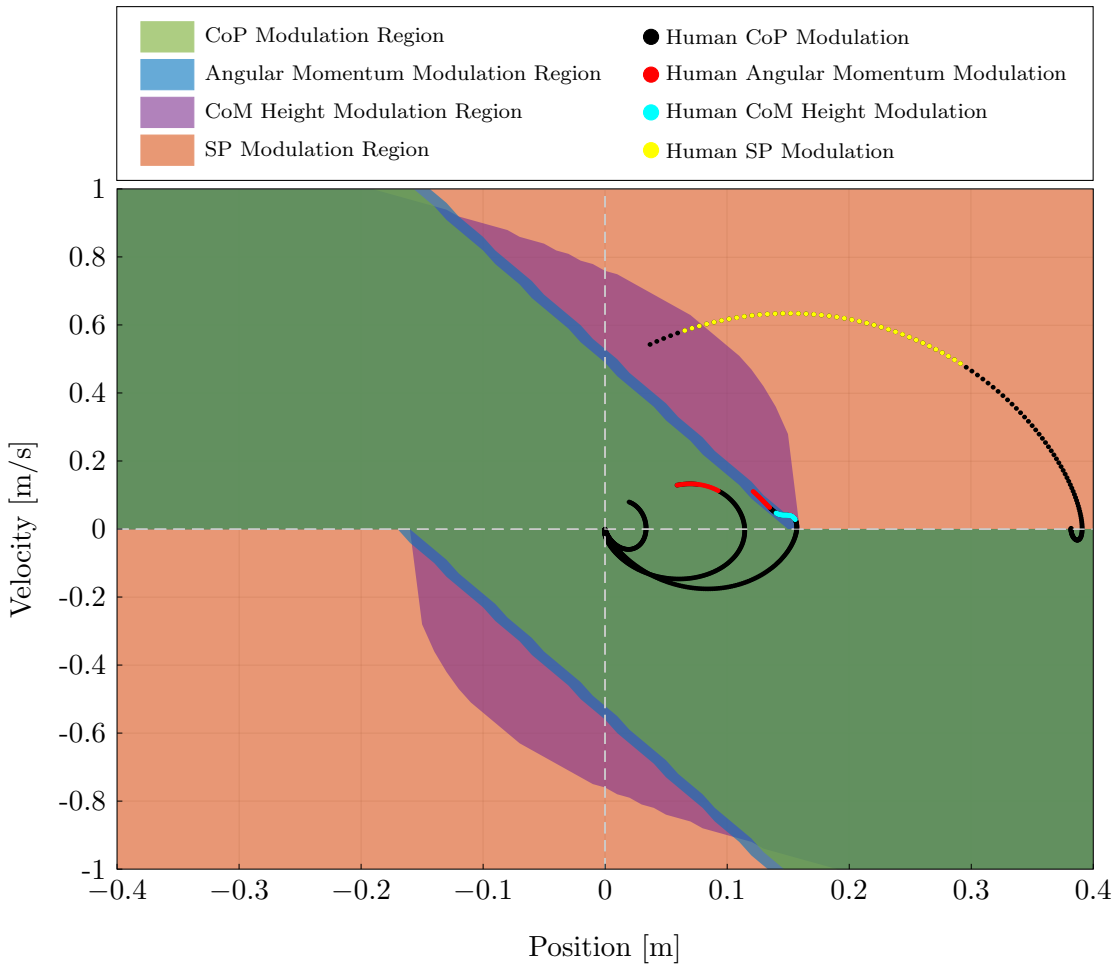

Figure 4.9: Human recovery data plotted against shaded stability regions.

where x is CoM position and the controlled variable u is CoM jerk $\dddot{x}$. This assumption will be ensured by constraining the input in the MPC. An independent integrator model is used for each axis, horizontal and vertical.

4.6.2 Minimum Jerk Model Predictive Control (MJMPC)

Our aim is to find a single control principle which accurately matches measured human trajectories in all strategies.

4.6.2.1 Model-Predictive Control Formulation

A minimum jerk principle is used as the single principle control principle to investigate fit across strategies.

We chose to use minimum jerk control as the control model for push recovery, as it is a widely used model for explaining human motion in other contexts and was considered appropriate for this application. The strong resemblance between the human data and minimum jerk trajectories used to explain human motions in other studies [104] provided further justification for this decision.

A minimum jerk MPC controller is designed, which minimises the objective function:

$$C(u(t)) = \frac{1}{2}\int_0^{t_f} \left(\frac{\mathrm{d}^3 x(t)}{\mathrm{d}t^3}\right)^2 dt = \frac{1}{2}\int_0^{t_f} u(t)^2 \mathrm{d}t. \tag{4.6}$$

Jerk $\frac{\mathrm{d}^3 x}{\mathrm{d}t^3}$ is used as control effort u with final time t_f. The MPC solves the following constrained optimisation problem:

$$\begin{aligned}
&\min_{u(t)} && C(u(t)) \\
&\text{subject to} && \text{Eq. 4.5} \\
& && [x(0), \dot{x}(0), \ddot{x}(0)] = [x_0, \dot{x}_0, \ddot{x}_0] && (4.7)\\
& && [x(t_f), \dot{x}(t_f), \ddot{x}(t_f)] = [x_f, \dot{x}_f, \ddot{x}_f] && (4.8)\\
& && [x_{\min}, \dot{x}_{\min}, \ddot{x}_{\min}] \leq [x, \dot{x}, \ddot{x}] \leq [x_{\max}, \dot{x}_{\max}, \ddot{x}_{\max}], && (4.9)
\end{aligned}$$

with initial condition $[x_0, \dot{x}_0, \ddot{x}_0]$, terminal condition $[x_f, \dot{x}_f, \ddot{x}_f]$, and Eq. 4.9 for physical feasibility.

In the MJMPC scheme (Algorithm 1), first, the desired CoM states X_{des}, the desired CoM lateral position, and Z_{des}, the desired CoM vertical position, are set as terminal conditions in the constrained optimisation problem. While the desired CoM height stays constant, the desired horizontal CoM x_d is computed by the step optimiser (Section 4.6.3.2) if the CP exceeds the SP. Using the current state as the initial condition X_0, Z_0 and the final time t_f (Section 4.6.3.3), the control effort $u_{x,1:t_f}, u_{z,1:t_f}$ is calculated over prediction horizon t_f by solving the constrained optimisation (4.7). Lastly, a whole-body QP controller calculates feasible joint torques τ by executing the first CoM reference position gained by minimising CoM jerk value $X_{\text{ref}} \leftarrow X_{\text{ref},1:t_f}(1)$ from the whole control input trajectory $u_{i,1:t_f}$ in an MPC fashion.

Although Angular Momentum Modulation was present in the human data, we found that by using two linear point models, we were able to recreate CoM motion during hip strategy needing to control Angular Momentum.

Algorithm 1: MJMPC control algorithm

$X_{\text{des}} \leftarrow [0,0,0]^T$, $Z_{\text{des}} \leftarrow [z_c,0,0]^T$
while $X \neq X_{des}$ **do**
  if $CP > SP$ **then**
    $X_{\text{des}} \leftarrow step\_optimiser(X)$
  end
  $X_0 \leftarrow X$, $Z_0 \leftarrow Z$
  $X_{\text{ref}} \leftarrow \min_{u(t)} f(u(t))$
  $X_{\text{ref}} \leftarrow X_{ref,1:t_f}(1)$
  $Z_{\text{ref}} \leftarrow \min_{u(t)} f(u(t))$
  $Z_{\text{ref}} \leftarrow Z_{ref,1:t_f}(1)$
  $\tau \leftarrow$ Whole Body Control$(X_{\text{ref}}, Z_{\text{ref}})$
  $X, Z \leftarrow$ Robot(τ)
end

4.6.3 Optimising Parameters to Fit Human Data

Three components are required before these controllers can be fitted to human trajectories: initial condition $\boldsymbol{x}_0$, terminal condition $\boldsymbol{x}_f$ and hyperparameters of each controller.

4.6.3.1 Initial Conditions

The parameters will be fit across all human trials using the CoM state at the moment the push is removed as the initial condition $\boldsymbol{x}_0$ for each trial.

4.6.3.2 Terminal Conditions

The reference state $\boldsymbol{x}_f = [0,0,0]^T$ as the terminal condition is applied to all non-stepping actions. For stepping actions, a new CoM reference is required, which is provided by the optimisation method for foot placement presented in [105]. This step optimisation considers kinematic and dynamic limits of the physical system as inequality constraints and outputs a new step location which is the final CoM position x_d, yielding the terminal constraint as $\boldsymbol{x}_f = [x_d,0,0]$. Step are only taken when the CP is beyond the SP. Our collected human data only involve a single step, but this can be extended to multiple steps for large pushes.

Table 4.4: Comparison of controller and human trajectories.

Axis	Mean MSE [mm^2]				
	Ankle	Hip	Toe	Step	Total
X	0.022	0.075	0.623	0.560	0.366
Z	0.003	0.005	0.07	0.01	0.0257

4.6.3.3 Final Time

The hyperparameters for each controller are determined via the least square fitting. The fitting minimises the least square error between the CoM trajectory generated by humans Y and controller Y^*, given the controller-specific parameters P at each time step t:

$$\sum_{i=1}^{N}\sum_{t=1}^{T}(Y^*_{t,i}(P) - Y_{t,i})^2. \tag{4.10}$$

Optimisation was performed using the MATLAB `fmincon`, and t_f was fit to both the X and Z axes for all trials.

4.7 Results

We present fitting results to the human data, and analyse the stability regions of the generated motions.

4.7.1 Characterising Human Data

Applying MJMPC with a fitted prediction horizon of 123 ms and 101 ms for the X and Z axis respectively, a close fit between human data and controller was achieved (Table 4.4). The controller has a total Mean Square Error over all trials of 0.366 mm^2 and 0.0257 mm^2 for X and Z axis respectively. The MJMPC scheme's ability to closely characterise the human policy can be seen by the close resemblance between human and robot trajectories (Figure 4.10).

4.7.2 Stability and Feasibility Analysis of MJMPC

Despite the MPC constraints on CoM motion, infeasible motions may still occur if the boundaries are incorrect. Therefore, we analyse the stability regions of the

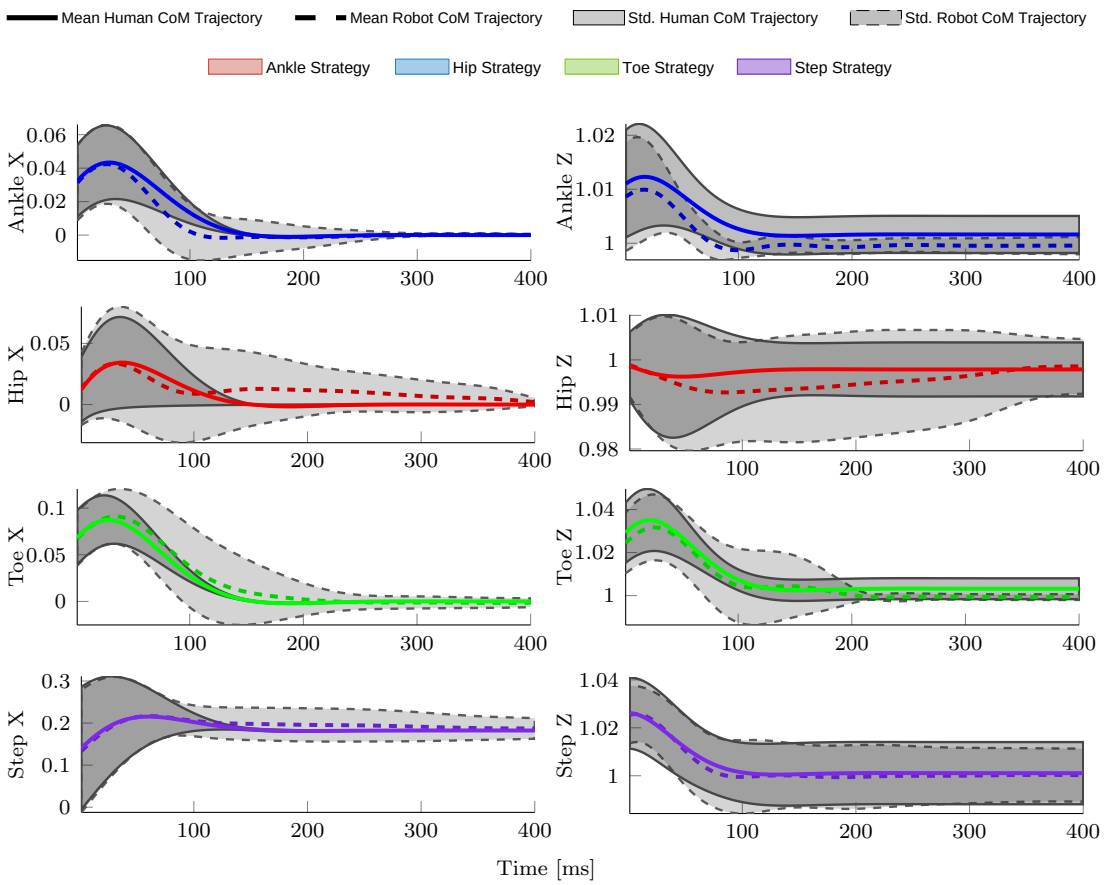

Figure 4.10: Comparison of human (solid line) and robot (dashed line) performance. Rows: ankle, hip, toe, step strategies in X and Z axes (columns). Units: leg length l.

system and consideration of the joint torque limits is conducted.

We validate the feasibility of the calculated CoM trajectory using an Inverted Pendulum Model (IPM), as detailed in Appendix A The max force and torque used during these trajectories is 1384 N and 352 N respectively. This is well within the Valkyrie specifications: $f_{\max}$ =2500 N, the magnitude of force applied to the ground by the robot, and $\tau_{\max}$ =2500 N·m. We show the time elapsed illustration of the performed IPM simulation is shown in Figure 4.12.

4.7.3 Applying Push Recovery Principles to Robotics

Having shown a single control principle can be used to explain human CoM motion across the main push recovery strategies, we show this can be used for robotics. In the earlier section, we used the MJMPC controller to match human

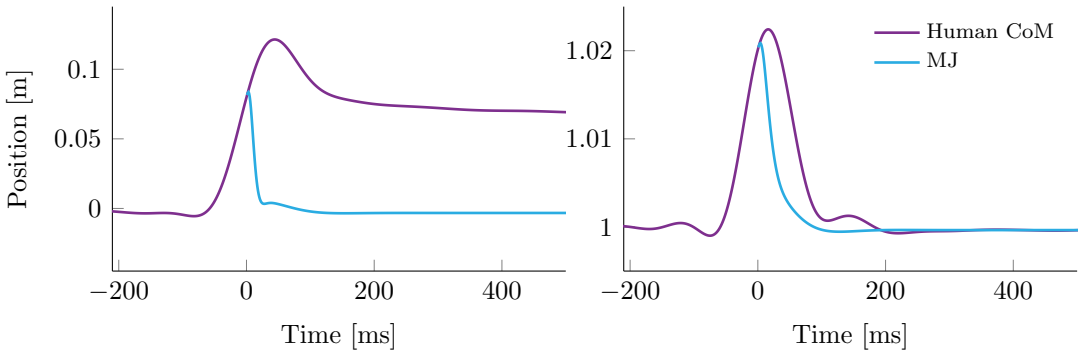

Figure 4.11: CoM trajectories tuned for time optimal performance. Compared to humans MJMPC converges faster and does not step to recover from the push in the (a) X-axis and (b) Z-axis.

performance. Our goal, however, is not to blindly imitate human behaviour, but to extract useful principles from humans such that they can be used in robotics. By tuning the time parameter of the MJMPC such that the CoM reaches the desired state as quickly as possible, we may see better performance while being able to exhibit each of the strategies. The reformulated control objective uses a time-optimal-like control cost, subject to force and torque limits of our example robot in Table 4.3:

$$\min_{x(t),z(t)} \quad \frac{1}{2}\int_0^{t_f} ((x_d(t)-x(t))+(z_d(t)-z(t)))^2 \mathrm{d}t \tag{4.11}$$

$$\text{subject to:} \quad \tau \leq \tau_{\max}, \quad \mathbf{F} \leq \mathbf{F}_{\max} \tag{4.12}$$

where $x(d), z(d)$ are the desired positions for the controller, set to the initial position of the human CoM during quiet standing before the push starts. This output is constrained to ensure the torque τ and force $\mathbf{F}$ produced are within the robot's limits τ_{max}, F_{max}. If the CP moves out of the SP, using the variables stated in Table 4.3 the same step selection criteria as in Section 4.6.3 were used.

Figure 4.11 shows the performance of the MJMPC using the new parameters with mean human CoM behaviour for comparison. We see, that the mean human trajectory after being pushed diverges significantly from the initial position, caused by human subjects taking a step. In contrast, MJMPC can return to the initial position without stepping, since the calculated CP does not move outside the SP (Alg:1, Ln:3).

When we were fitting the controllers to the human CoM trajectories in the previous section, the controllers were implicitly accounting for sensing and muscle

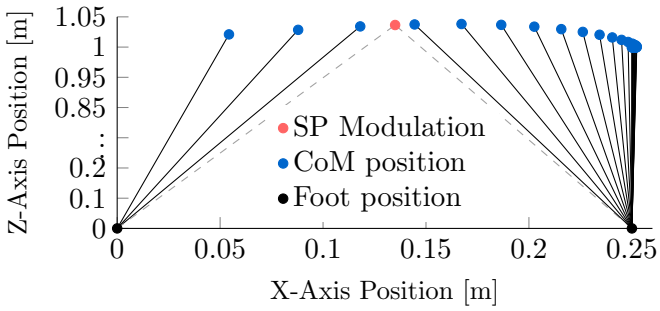

Figure 4.12: Step strategy CoM motion using MJMPC controller.

actuation delays in humans. Since these delays are trivial in robotics, using this time optimal approach, we can achieve higher performance than in the human data.

4.7.4 Comparing Stability Region between Human and Robots

The boundaries delineated in Section 4.5 are used to create stability regions for robots (Figure 4.13) employing the values provided in Table 4.3. These robot stability regions exhibit qualitative similarities to human regions as depicted in Figure 4.9. Interestingly, the limited effectiveness of Angular Momentum Modulation is evidenced by its stability region in both human and robot plots. It substantially overlaps with CoP Modulation, and for the most part, resides within the CoM height modulation stability region. Consequently, Angular Momentum Modulation can be replaced by a combination of CoP and CoM Height modulations, as demonstrated by the reconstruction of human data using this method (see Section 4.7.1). Nonetheless, it should be noted that the influence of arm motion was not factored into the motion analysis or the stability region modeling, and incorporating these effects on balance would likely modify the region.

4.8 Conclusion

This chapter first investigated human push recovery motions with the goal of identifying and studying the invariant aspects of this simple form of locomotion. We quantified a set of criteria which can identify when Control Actions are being used. Once these were identified, we were able to find a common minimum jerk control principle which could accurately match each of these Control Actions.

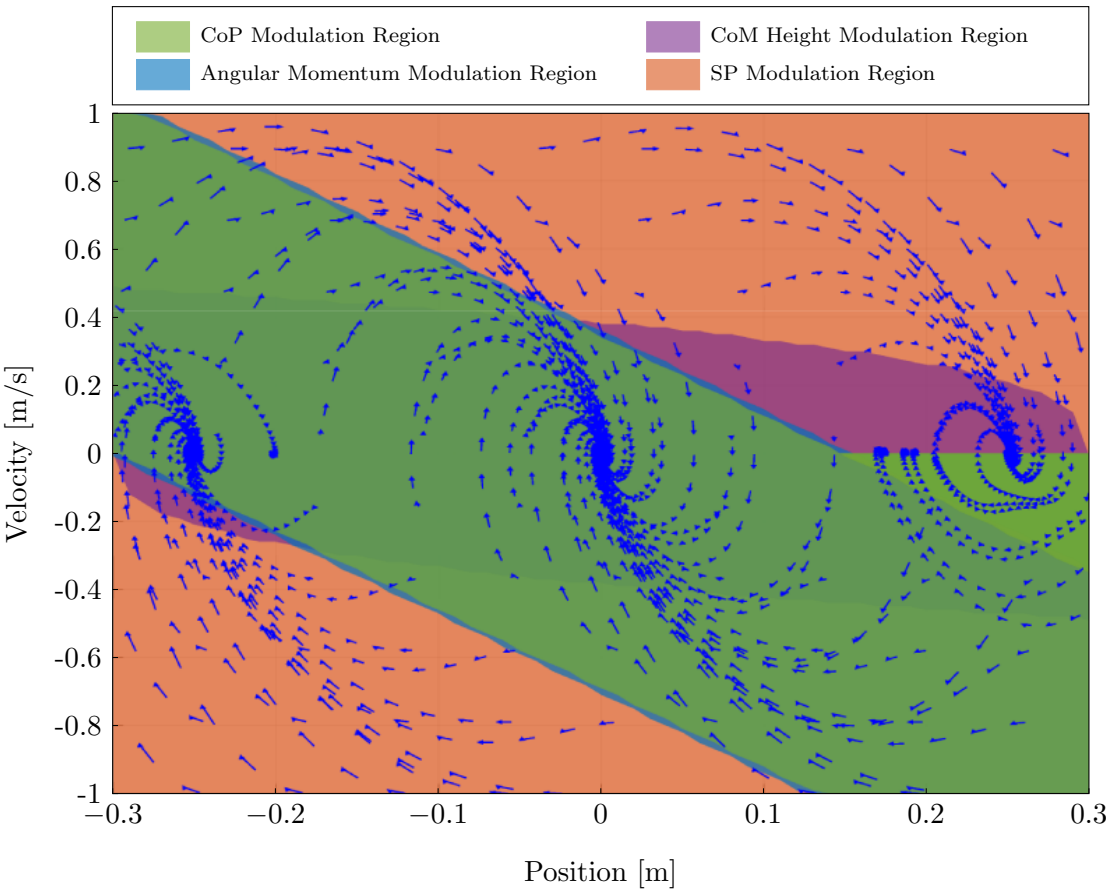

Figure 4.13: Stability regions for each action with coloured curves indicating human CoM motion during recovery trials.

Taking this control principle, we were then able to increase performance by tuning its parameters to perform optimally rather than to fit human data.

From this, we conclude that despite the complexity of human recovery behaviour it is possible to capture the characteristics of the CoM of the human using a single simple rule across all behaviours. Moreover, the key aspect of these findings are that between different means of motion, whether they involve contact switching or not, there is a consistency in their optimality at the CoM level in the sense that they are each modelled by a controller which optimises minimum jerk.

Optimising for minimum jerk implies a level of efficiency, as changes in the derivative acceleration determine energy consumption. Motivated by this, we next continue to study the invariance in human locomotion, but with a focus on one of the unanswered questions from this chapter. Though certain aspects of push

recovery showed a level of invariance, when step positions were taken by participants, there was a level of stochasticity in the choice of step location for similar push magnitudes. In the next chapter this variance is explored in an effort to find an explanation or an underyling invariance that exists in this choice of step position.

Chapter 5

Energy Efficient Stepping for Balance Recovery

5.1 Introduction

The positions where humans choose to step during recovery from pushes exhibit a significant degree of stochasticity that is challenging to capture and explain using current models. In this chapter, we delve into this phenomenon and investigate the underlying principles that govern it, with the aim of gaining a deeper understanding of human motion and exploring potential applications for controlling robot locomotion. Our study takes an energy efficiency perspective, as this is a crucial aspect of human motion and a critical goal for legged robotics in the long term.

To achieve long operation time for legged locomotion, it is essential to study the relationship between energy-efficiency and locomotion behaviours. The energy-efficiency and actuation power of a robot during locomotion are highly related to both its multi-body dynamics and its control, which therefore must be considered together when studying energy consumption.

For high-DoF robots, coupling hardware and software creates high dimensional state-action and solution spaces, such that studying energy consumption for general locomotion is complex and computationally heavy. This chapter investigates stepping behaviour and energy consumption and sheds light on this nonlinear relationship for a narrow aspect of locomotion – stepping and balance recovery –

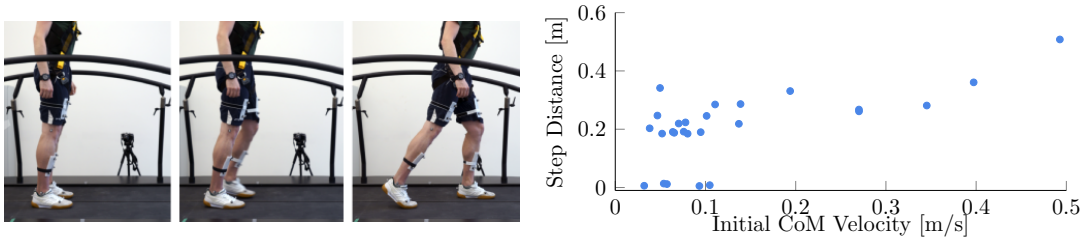

(a) (b)

Figure 5.1: Human stepping and balance recovery which is hard to model as a biological multi-body system: (a) A subject taking a step during push recovery; (b) Distribution of step locations during push recovery [5].

as a proof of concept.

To study energy-efficient stepping, we must consider the robot system and its dynamic behaviour as a whole. Whole Body Motion Planning (WBM) [106] generates motion plans for the full robot system and considers its dynamics throughout the planning stage and therefore can be accurately tracked using Whole Body Control (WBC) [107] and can be formulated to enforce tasks like kinematic reachability [108, 109], collision avoidance [110] and centroidal dynamics [111], but plans are computationally expensive to generate. Mixed Integer Convex Optimisation produces WBM plans for online use [108, 112] using conservative constraints such as static stability, but such motions are unsuitable for efficient walking. Therefore, though WBM methods could be used to optimise for energy-efficiency, it is impractical for the study of a large search space.

Alternatively, reduced models such as point mass models [23] can quickly produce reactive step plans [113] and stabilisation control [114], and produce energy-efficient motion [115] at the planning level. Rapid computation comes at the cost of accuracy, as reduced models simplify the strong nonlinear and coupling effects of dynamics, therefore, optimality of planned motions is lost when projected to the whole system. We propose a proof-of-concept method that uses WBC to evaluate the true energy cost of tracking trajectories generated by reduced models instead of optimising efficiency within the reduced model, resulting in a method of rapid trajectory generation with an accurate associated energy cost.

Cost of Transport (CoT), commonly used to quantify locomotion efficiency, is a dimensionless approximation of energy cost per unit weight of a body during locomotion, expressed as a single numerical value and is used as a unified metric

to compare different bodies moving at different speeds [116]. CoT can aid robot design and control development [117], optimise gait for a given movement speed [118] and, at a high computational cost, for footstep planning [119]. But, since we focus on single-step motions for which it is hard to define walking speed to calculate CoT, it is difficult to apply in this case.

For a robotic system consisting of electric motors, the mechanical power of all joints is ultimately converted from electrical power, meaning the mechanical power consumption can be equally quantified by electrical consumption [120]. Joint torque is a good index for energy dissipation [121, 122, 123], since the torque is directly proportional to current. Therefore, given a constant voltage in a DC motor, torque can be used to reflect energy. In our simulation study, the sum of squared joint torque can be used to quantify whole-body energy use, which is integrated to give torque use during stepping.

We compile reachability maps representing the realistic energy cost of whole body motion into a compressed representation, allowing us to inversely generate energy-optimal trajectories for selecting step locations. Similar reachability maps have been effectively used in previous work for kinematic reachability [124, 125], feasible transition motions [126, 127], obstacle avoidance [110, 128], faster locomotion planning [129, 130], complex end-pose planning [131], dynamic transitions [40], and warm-start learning [132].

Other studies of locomotion have explored efficiency during locomotion, similarly building energy cost landscapes that map onto the existing terrain [74], using reduced models to plan efficient steps based on the passive dynamics of the multi-body system [133], trying to capture the inherent efficiency of passive walking on a slope [134], and analyzing mathematical walking models [135], even using these to produce walking landscapes [136]. Our work builds on these previous studies by using realistic dynamic models of a full humanoid system to analyze a more complete efficiency landscape. Moreover, by analyzing our maps, we demonstrate simple heuristics and a diverse range of efficient stepping locations for planning motions with low computational cost.

Heuristics for energy-efficient step regions provide insights into complex stepping behaviours in humans. Figure 5.1b, from [5], shows step positions for human balance recovery (Figure 5.1a), offset by mean initial velocity of non-stepping

trials (0.1103 m/s) and shows the stochasticity of human stepping [137, 138]. Our heuristics for the underlying optimality help explain the stochasticity of human stepping.

This chapter is motivated to study the relationship between step location and energy optimality for balance recovery from a set of initial conditions by using a whole-body model to evaluate the energy cost of motions generated by simple models. Energy efficiency is not a primary concern in balance recovery, but by studying efficiency we find simple heuristics to rapidly select recovery step positions for full-body, dynamic motions which are limited when using simple models alone. We accurately quantify the energy cost of stepping for the whole stepping space for all joints in a 32-DoF robot in a realistic dynamic physics simulation to fully explore this relationship and develop a method of rapid, energy optimal step planning. We use Bayesian Optimisation (BO), a sample efficient, nonlinear optimisation method suited to whole-body tasks [139] to optimise open parameters for efficient stepping.

5.1.1 Scope

We develop an optimisation pipeline and focus on the following set of stepping motions: taking a single step forward from an initialised standing position with the CoM above stance foot and without toe-off motions. This allows us to search for optimal locomotion parameters given different initial states, and obtain results of the torque usage and energy consumption of a full-body humanoid robot with redundancies.

5.1.2 Contributions

This chapter develops a nonlinear optimisation pipeline and studies the efficiency of power consumption in humanoid stepping. Our contributions are as follows:

1. **Energy-Optimal Stepping (Section 5.4.2.3)** Reachability maps that show energy-optimal step positions based on whole-body dynamics and the use of optimal control;

2. **Optimisation Pipeline (Section 5.4.1, 5.4.2)** A parallelized optimisation pipeline for whole-body control based stepping in full dynamics in a

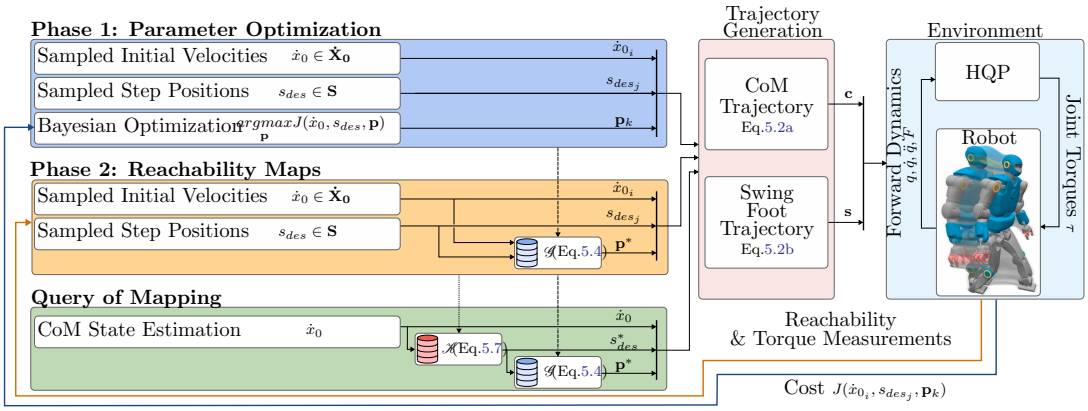

Figure 5.2: A pipeline for optimising trajectory generation and HQP parameters: Phase 1 optimises trajectory generation parameters for sets of motions and outputs a parameter mapping; Phase 2 uses this mapping to build reachability maps and an energy-optimal step selection mapping; both mappings are combined for energy-optimal push recovery motions.

physics simulation;

3. **Reachability Maps (Section 5.4.2.2)** A proposed method for computing precise reachability maps for dynamic motions which can rapidly select energy-efficient step locations;

4. **Finding of Simple Approximation of Optimal Stepping (Section 5.4.2.2)** Finding of simple approximations and the disclosure of a funnel of near-optimal step locations.

Our sampling-based energy cost maps reveal a complex, nonlinear distribution of efficient stepping locations, surrounded by regions of similar efficiency, indicating a simple step selection heuristic for balance recovery – which can only be studied and understood by considering the complexity of whole-body dynamics and optimal control as a whole. We find optimal step regions are different from those predicted by simple models which do not capture whole-body dynamics. Reachability maps can also be used for rapid step selection and give insight into non-uniqueness in human step selection.

Table 5.1: Optimisation variables and descriptions.

Optimization Parameter	Description	Affects	Bounds	Optimized Range	Unit
t_{min}	Minimum swing time	CoM	0.01-0.99	0.01-0.77	s
s_{max}	Maximum step length	CoM	0.01-0.99	0.06-0.98	m
t_{swing_start}	Swing foot start time	Swing Foot	0.01-0.08	0.026-0.078	s
s_{speed}	Swing foot velocity	Swing Foot	0.2-3.0	0.2-1.34	m/s

5.2 Problem Formulation

To study energy-efficiency during stepping, we generate trajectories for the CoM and swing foot using simplified models and evaluate their energy cost by tracking them with a whole-body feedback controller in a full dynamic simulation. Trajectories each have two open parameters and BO is used to optimise the open parameters for energy-efficient stepping, we then produce maps to associate pairs of initial conditions and trajectories to energy cost. We achieve this using our 2-phase pipeline, defined mathematically below.

A Hierarchical Quadratic Programming (HQP) feedback controller [140] is used to map the simple trajectories onto a complex whole-body system with low computational cost due to its quadratic formulation. The HQP, $\mathscr{H}$, tuned for dynamic stepping, solves joint torques, $\boldsymbol{\tau}$, given a set of task constraints:

$$\mathscr{H} : (\mathbf{c}, \mathbf{s}, t, \mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}) \longmapsto \boldsymbol{\tau}, \tag{5.1}$$

where $(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}) \in \mathbb{R}^m$ are the position, velocity and acceleration of the robot's joints, and m is the number of joints and t the current time. $\mathbf{c}$ and $\mathbf{s}$ are reference trajectories for the CoM and the swing foot respectively, additional details are given in Section 5.3.1.

Trajectory planning generates spatial paths which are intended to be followed by a specific part of the robot, in our case we have two trajectories, one for the CoM [105] and one for the swing foot [141, 142]. Similarly, trajectory optimisation seeks to find the best spatial paths between two points according to some task

or cost [47], which in this case is performed by BO. The trajectories are time indexed position references such that $\mathbf{c}(t) \in \mathbb{R}^3$ and $\mathbf{s}(t) \in \mathbb{R}^3$, generated by the parameterized functions:

$$\mathscr{C}_{traj\_gen} : (t_{min}, s_{max}) \longmapsto \mathbf{c} \tag{5.2a}$$

$$\mathscr{S}_{traj\_gen} : (s_{des}, t_{swing\_start}, s_{speed}) \longmapsto \mathbf{s}. \tag{5.2b}$$

CoM trajectory generator $\mathscr{C}_{traj\_gen}$ takes minimum step time, t_{min}, and maximum step length, s_{max}, as arguments. Swing foot trajectory generator $\mathscr{S}_{traj\_gen}$ is a function of desired step position s_{des}, swing start time t_{swing_start} and step speed s_{speed} parameters, which produces 3 dimensional trajectories; note, orientation is not explicitly controlled in this case and the foot is assumed to be parallel to the ground. Step speed s_{speed} is the swing foot target speed, typically reached at swing apex. Open parameters are tuned in Phase 1 of our pipeline, summarised in Table 5.1, and are concatenated in vector $\mathbf{p}$:

$$\mathbf{p} = [t_{min}, s_{max}, t_{swing\_start}, s_{speed}]^T. \tag{5.3}$$

Note that s_des is not an open parameter, it is provided as an initial condition (Phases 1 & 2), or generated from Eq. 5.7. The output of Phase 1 is function $\mathscr{G}$:

$$\mathscr{G} : (\dot{x}_0, s_{des}) \longmapsto \mathbf{p}^* = \arg\max_{\mathbf{p}} \quad J(\dot{x}_0, s_{des}, \mathbf{p}), \tag{5.4}$$

which maps initial CoM velocity, $\dot{x}_0 \in \mathbb{R}^1$ in the lateral plane, and desired step position, s_{des}, to a set of optimal parameters $\mathbf{p}^*$ (optimal values are denoted by $(\cdot)^*$). In Phase 1, the values of s_{des} are set to predetermined intervals increasingly far away from the robot; in Phase 2, optimally efficient step positions s^*_{des} can be generated automatically.

To build this mapping, we use BO to maximise the objective function J (Eq. 5.5) for pairs of initial conditions. Initial conditions and BO parameters are iteratively passed to the trajectory generators (Eq. 5.2a, 5.2b), produced trajectories are then tracked by the HQP (Eq. 5.1), in a full dynamic simulation environment, detailed in Section 5.3.3, where the robot attempts to step towards the desired step position s_{des}. This is highlighted in dark blue in Figure 5.2

The objective function J takes values from the dynamic simulation, which we consider to be a black-box and each term is assigned a manually tuned weight,

each denoted by a subscript of w, defined in Table 5.2:

$$\begin{aligned} J(\dot{x}_0, s_{des}, \mathbf{p}) = -(&w_f(t_{total} - t_{term}) \\ &+w_{swing}(s_{des} - s_{td})^2 + w_{x\_mid}(x_f - s_{mid})^2 \\ &+w_z(z_{nom} - z_f) + w_\tau J_\tau(\dot{x}_0, s_{des}, \mathbf{p})), \end{aligned} \tag{5.5}$$

Weights are tuned to promote successful stepping, low energy cost and scale down large terms. For example, t_{total} is set to 7 s at 1000 Hz, giving 7000 time-steps; termination at 4.5 s would incur a cost of $0.001 \cdot (7000 - 4500) = 2.5$, whereas other terms typically evaluate to less than 1 during success and are nominally high during failure to amplify w_f. Termination conditions are detailed in Section 5.3.3.

The error between swing foot position at touchdown, s_{td}, and desired step location, s_{des}, incurs a cost proportional to the square of the error between them. To increase stability after the step, a cost is applied between final CoM position, x_f, and the midpoint between the stance foot and the desired swing foot position: $s_{mid} = s_{stance} + (s_{stance} - s_{des})$. To encourage straight legs after landing, a cost is applied to the final CoM height z_f to be as close as possible to a nominal height $z_{nom} = 0.925$ m. A term is added to minimise the integral sum of the squared measured torque $\boldsymbol{\tau}$ in all the robot's joints (32 in this case) between the time swing foot motion begins t_{lo} and touchdown time t_{td}:

$$J_\tau : (\dot{x}_0, s_{des}, \mathbf{p}) \longmapsto \int_{t_{lo}}^{t_{td}} \boldsymbol{\tau}^2 \, dt. \tag{5.6}$$

After BO is completed for each pair of initial conditions, we are left with the mapping function $\mathscr{G}$ (Eq. 5.4), which outputs optimised parameters for pairs of initial conditions. In Phase 2, we query this mapping with a larger set of initial velocities and desired step positions to test how the optimised parameters generalise to novel initial condition pairs, illustrated in Figure 5.4. Here, objective function J is no longer used, and we use the optimal parameters in the mapping function; instead, the dynamic simulation returns two values, shown in Figure 5.2: a binary reachability value to encode whether motions are successful, and the integral sum of the square of the measured torque (Eq. 5.6) for every pair of initial conditions.

A forward map then encodes which step positions are reachable from each initial CoM velocity and another map quantifies the integral sum of the square torque

cost during the swing phase of each reachable step position. Using the second map, we create an inverse mapping, $\mathscr{K}$, from an initial CoM velocity $\dot{x}_0$ to an energy-optimal step position s^*_{des} and the optimal parameters used to reach that location.

$$\mathscr{K} : (\dot{x}_0) \mapsto \begin{cases} \mathbf{p}^* = \mathscr{G}(\dot{x}_0, s^*_{des}) \\ \text{s.t.:} s^*_{des} = \underset{s_{des}}{\text{argmin}}\, J_\tau(\dot{x}_0, s_{des}, \mathscr{G}(\dot{x}_0, s_{des})). \end{cases} \tag{5.7}$$

Automatically generated energy-optimal step positions s^*_{des} are passed to $\mathscr{G}$ to reduce the input dimensionality such that only initial CoM velocity $\dot{x}_0$ is required to query optimal trajectory generation parameters $\mathbf{p}^*$.

All maps are built offline and a regression model is used to approximate energy-optimal step positions s^*_{des} for arbitrary initial CoM velocities $\dot{x}_0$. We can query these mappings for stepping or push recovery; we run validation tests using the mappings and previously unseen initial conditions.

5.3 Technical Details

5.3.1 HQP Structure

The HQP feedback controller is a function of forward dynamics and two parameterized trajectories (Eq. 5.1). We tuned the tasks, weights and hierarchy order of an existing HQP Controller [140]. Upper body joints such as the arms and head are not required for the stepping motion, so a posture task is also added to regulate arm motion and keep them near a nominal position with a low weight in the HQP.

5.3.2 Trajectory Generation

CoM and swing foot trajectories are passed to the HQP to determine their energy cost. Each has 2 open parameters which determine their profile and outputs time-indexed position references for the X,Y and Z axes. Parameters determine the length of the trajectory and its gradient and are optimised in Phase 1 (Section 5.4.1).

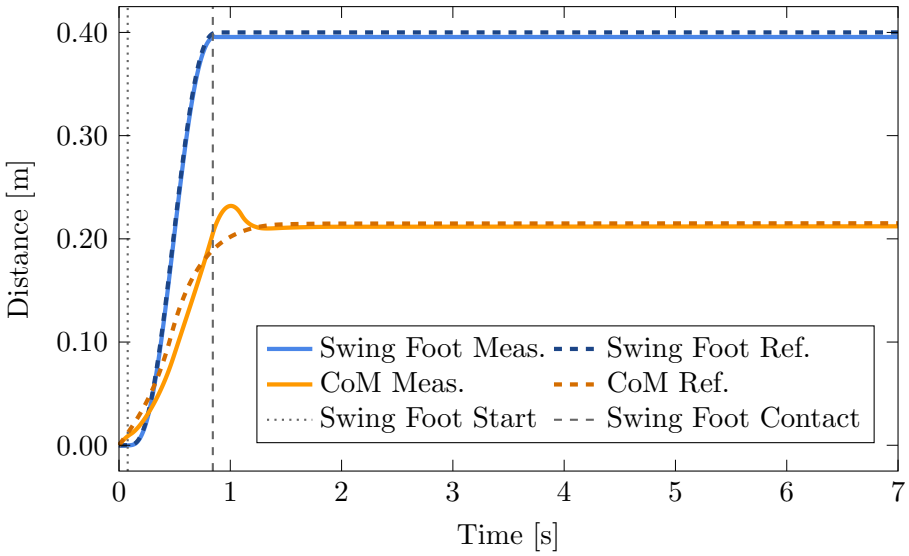

Figure 5.3: Example reference trajectories for the CoM and swing foot, and measured trajectories from the HQP. Initial conditions: $\dot{x}_0$ =0.15 m/s, s_{des} = 0.4 m.

5.3.2.1 CoM Trajectory Generation

An existing LIPM based model [105] produces CoM trajectories as a function of two parameters: minimum step time t_{min} and maximum step length s_{max}. An example trajectory is shown in Figure 5.3. During testing, s_{max} was set to reflect the capabilities of the physical 32-DoF Talos humanoid robot, but this results in falling, so this parameter was added to the optimisation. Since this pipeline is modular, this can be replaced with alternative CoM trajectory generation methods.

5.3.2.2 Swing Foot Trajectory Generation

Swing foot trajectories are 5^{th} degree minimum-jerk polynomials [104] (Figure 5.3), parameterized by the time at which the swing foot starts to move, (t_{swing_start}), and the swing foot speed (s_{speed}). The Z axis consists of two minimum jerk trajectories connected at a via point at the maximum desired swing height z_{max}, for which we found 80 mm to be a reliable value.

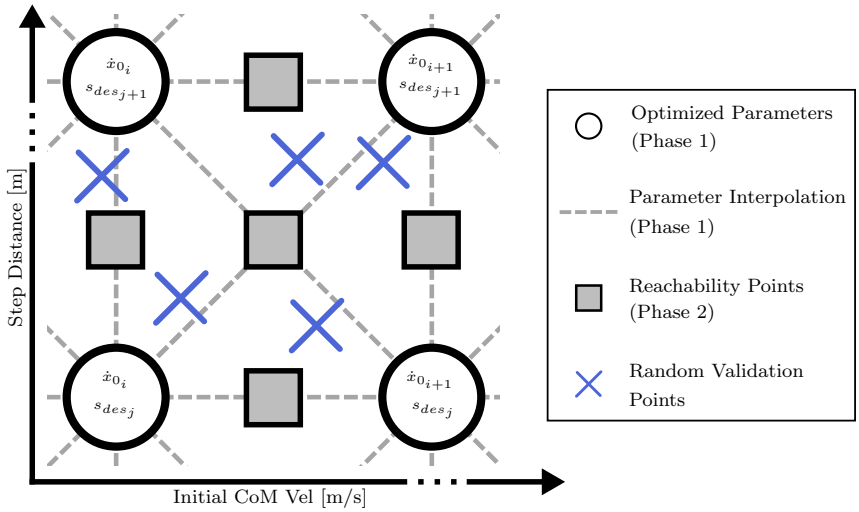

Figure 5.4: Illustration of the distribution of parameter, sampling and validation points.

5.3.3 Simulation Setup

For the dynamic simulation, we used Pinocchio rigid body dynamics library [143], based on the widely used dynamics equations in [144]. In the long term, we aim to move the query of mapping implementation onto the real robot, but was not possible for this chapter. This pipeline is compatible with any humanoid robot, here we used the Talos humanoid robot model (32 DoF), with the complete dynamic and kinematic properties of the real robot including the position, velocity, acceleration and torque limits. The simulation environment was fully dynamic, including friction, torque limits, with both simulation frequency and control frequency of 1000 Hz ($\Delta t = 0.001$ s). The `BayesianOptimisation` package [145] was used for parameter optimisation.

We constrain motions for this proof-of-concept to gait initiation in the X axis (forward), where in each simulation episode the robot starts in a standing configuration with the swing foot 1 cm above the ground to reduce the complexity of optimising weight transfer time while we develop the pipeline.

Initial CoM velocities are induced by directly applying joint torques in simulation, torque values are calculated using the Jacobian from the stance foot to the CoM. On the real robot, we expect CoM velocities to be applied by having the robot lean in a given direction or being pushed. Termination conditions during optimisation are as follows: robot reaches desired foot position and remains standing

Table 5.2: Objective function weights and their values.

Notation	Affects	Value
w_f	Failure to Complete Stepping Motion	0.001
w_{swing}	Swing Foot Position error	50
w_{x_mid}	Final CoM Position	1
w_z	Final CoM Height	1
w_τ	Torque Minimization	0.0002

at $t = t_{total}$ (success) or the norm of joint velocities exceed a threshold ($1e6$) (failure). If $t_f < t_{total}$, the remaining sensor readings are filled with a nominal high value.

5.4 Building Reachability Map: Dynamic Stepping Optimisation Pipeline

This section describes the phases and sub-routines of this pipeline, starting with building an optimal parameter query space and resulting in a rapid, step-selection model for efficient motion.

5.4.1 Parameter Optimisation - Phase 1

Phase 1 of this pipeline deals with the optimisation of stepping parameters and generating a continuous parameter mapping. This is arranged into two sub-steps.

5.4.1.1 Bayesian Optimisation

Phase 1 of our pipeline, shown in the dark blue box in Figure 5.2, optimises a set trajectory generation parameters, $\mathbf{p}$, for pairs of initial conditions. Algorithm 2 shows the nested loops: for every initial CoM velocity $\dot{x}_0 \in \dot{\mathbf{X}}_0$ and every desired step position $s_{des} \in \mathbf{S}_{des}$, a set of parameters $\mathbf{p}$ is optimised via BO, such that for each pair, $[\dot{x}_0, s_{des}]$, is mapped to a set of parameters and stored in the mapping $\mathscr{G}$.

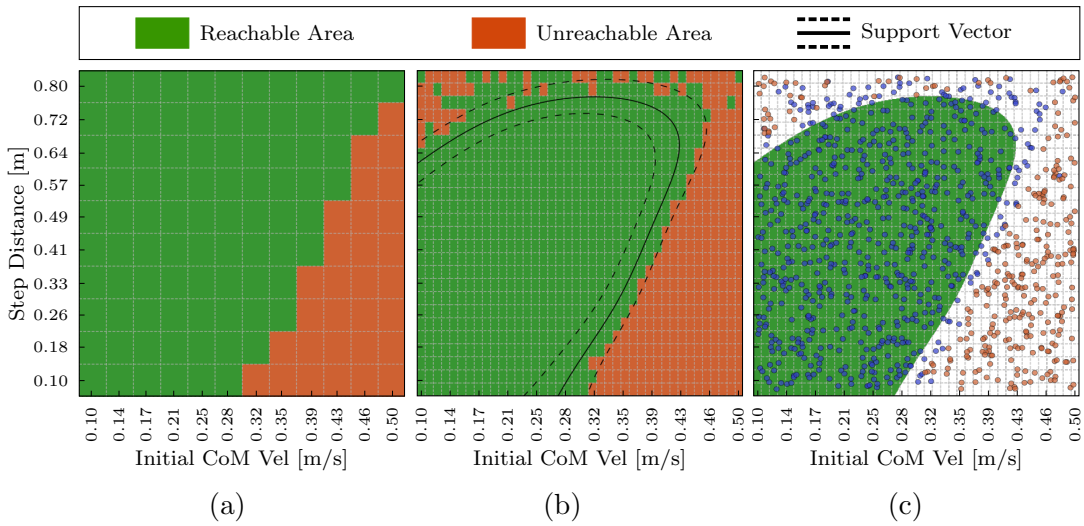

Figure 5.5: Binary reachability maps show successful (green) or unsuccessful (red) steps for pairs of initial conditions: (a) Phase 1: Parameter optimisation map, (b) Phase 2: Dense reachability map, (c) Validation: SVM generated reachability map with validation testing points.

Optimal parameters were found for 150 initial conditions: 15 initial CoM velocities, 10 step positions – sparsely covering the state space, as shown in Figure 5.4 – with 170 BO iterations for each (100 random, 70 Bayesian) using objective Eq. 5.5. Figure 5.5a shows the result where initial conditions are projected onto a 2D space and the area around each is coloured according to a binary classification of whether successful stepping parameters were found. All desired step locations are reachable from low initial CoM velocity and as this increases, steps which are closer to the robot's starting position are no longer reachable, as the robot is moving too quickly to take a single short step in that direction.

5.4.1.2 Parameter Interpolation

Linear, element-wise interpolation between the nearest 4 sets of sparsely optimised optimised points (Figure 5.4) were used to create a continuous mapping $\mathscr{G}$ (Eq. 5.4) from arbitrary initial values to a set of interpolated, optimised parameters $\mathbf{p}^*$.

Algorithm 2: Bayesian Parameter Optimisation

input : List of sampled CoM velocities:$\dot{\mathbf{X}}_0$,
List of desired step positions:$\mathbf{S}_{des}$
output: Mapping function $\mathscr{G}$ (Eq. 5.4)

```
for each CoM initial velocity (ẋ_0) in Ẋ_0 do
    for each step position (s_des) in S_des do
        for each k in BayesOptIterations do
            p(k) ← BayesianOptimisation (Eq. 5.5)
            objective ← DynamicSim(ẋ_0, s_des, p(k))
        end
        G(ẋ_0, s_des) ← arg max_p (objective)
    end
end
```

5.4.2 Reachability Maps - Phase 2

The purpose of Phase 2 (orange box, Figure 5.2) is to use the optimised and interpolated parameters from Phase 1 to create higher resolution maps to denote the reachability and efficiency of stepping parameters, illustrated in Figure 5.4; no further parameter optimisation is performed.

The outputs of this Phase are: a high resolution forward map of reachable step positions, a map of the integral sum of square torque cost to reaching these positions using the simplified model trajectories and an inverse mapping for selecting optimally efficient step positions from any initial condition (Eq. 5.7).

We used 1000 initial condition pairs (40 initial CoM velocities, 25 desired step positions) to sample the mapping function, compared to 150 in Phase 1. Initial CoM velocities were between $\dot{\mathbf{X}}_0 = [0.1\,\text{m/s}, \ldots, 0.5\,\text{m/s}]$ at $0.017\,\text{m/s}$ intervals and desired step positions: $\mathbf{S}_{des} = [0.1\,\text{m}, \ldots, 0.8\,\text{m}]$ at $0.029\,\text{m}$ intervals. A dynamic simulation episode was executed for each initial condition pair using parameters $\mathbf{p}^* = \mathscr{G}(\dot{x}_{0_i}, s_{des_j}) \forall\, \dot{x}_{0_i} \in \dot{\mathbf{X}}_0, s_{des_j} \in \mathbf{S}_{des}$ Eq. 5.4.

Simulation episodes return two values: a binary value to denote if the stepping motion was successful, and the integral sum of square joint torques, $\boldsymbol{\tau}$, using Eq. 5.6. As a result, each pair of initial conditions has an associated reachability value and measured integral sum of square joint torque value that we use to build the reachability maps.

5.4.2.1 Binary Reachability Map

The binary reachability map (Figure 5.5b) projects every sampled pair of initial conditions from Phase 2 onto a 2D space and each pair is coloured according to a binary classification of if the robot successfully reached the desired position and stayed standing or not. Similar to Figure 5.5a, when using the interpolated parameters many steps are successful at lower initial CoM velocities and as this increases the step locations closer to the robot are unreachable. However, the space of parameters that lead to successful motions appears to be much smaller for longer steps, but this is not captured by the interpolation between sparsely sampled points, resulting in noisy regions at the top of the map. In these areas where success is uncertain, the robot may not fall in every case, and it may be possible to take extra steps, but here we consider one-step capturability to be unsuccessful in the noisy region.

A denser set of optimised parameters would remove this noise but would require more optimisation time and since the robot is operating at the limits of its workspace, would still be sensitive to minor changes in modelling or sensor error on the real robot. Instead, we chose a more conservative approach to trim the noisy extremes and leave only the conditions which reliably lead to successful stepping. We train a Support Vector Machine (SVM) model to separate the reliable step locations from the noisy areas, shown in Figure 5.5b. By querying the SVM model, we create a cleaner representation of the safe stepping area (Figure 5.5c). We used a 3^{rd} order SVM with a radial basis function kernel, reachable points had a weight of 1, unreachable points had a weight of 14.

5.4.2.2 Measured Torque Maps

The integral sum of square torque measured during the swing phase of motions from each initial condition pair is shown in Figure 5.6, where each point in the trimmed reachability map is coloured according to the integral sum of the squared joint torques (Eq. 5.6); darker (more purple) colours denote higher measured integral sum of square joint torque and the lowest cost step positions for each sampled initial velocity are marked. The distribution of torque patterns is highly nonlinear, but steps with minimum measured integral sum of square joint torques form a simple trend that can be used for footstep prediction. We consider the energy-optimal stepping positions to be those with the lowest integral sum of

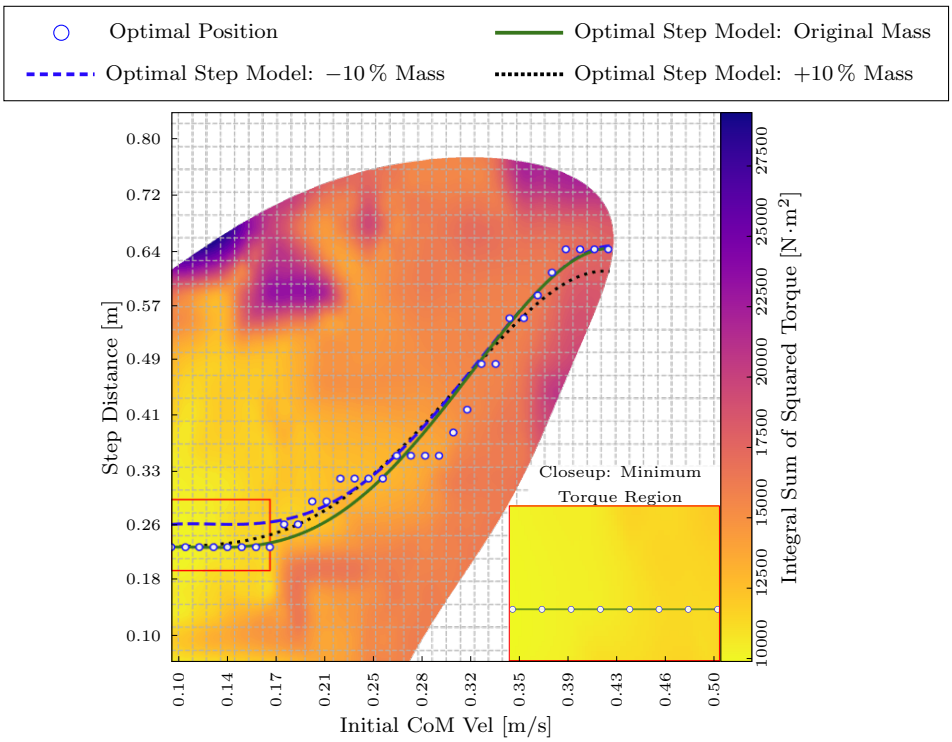

Figure 5.6: Map of the measured integral sum of squared joint torques during the swing phase to reach each position, with optimal step locations marked. A regression model is fitted to optimal step positions, for the original mass and validation tests with mass discrepancies, and a zoomed region is shown in inset with a maximum of 5% deviation from optimal for the original mass.

squared torque integral for all joints.

5.4.2.3 Minimum Energy Step Selection

The relationship between initial CoM Velocity and energy-optimal step positions can be modelled using a simple 4^{th} order polynomial regression, which can quickly approximate energy-optimal step positions given an initial CoM Velocity (Eq. 5.7). Figure 5.6 shows this model captures the minimal energy step positions trends with a mean error of $216.12\,\text{N}\cdot\text{m}^2$ (StD.=$287.84\,\text{N}\cdot\text{m}^2$, Min.=$0\,\text{N}\cdot\text{m}^2$, Max.=$1159.39\,\text{N}\cdot\text{m}^2$).

5.4.3 Query of Mapping

Resultant mappings can be queried to produce trajectories for stepping motions, shown in the green box in Figure 5.2. Therefore, for a given CoM velocity, an energy-optimal step position is output by the mapping $s^*_{des} = \mathscr{K}(\dot{x}_0)$ (Eq. 5.7),

which is then used to generate the stepping parameters to produce this motion using the mapping $\mathbf{p}^* = \mathscr{G}(\dot{x}_0, s^*_{des})$ (Eq. 5.4). This allows us to automatically generate stepping for a given initial CoM velocity during push motions.

Since the parameters in Phase 1 are optimised for a set of discrete points, in the next section we run a series of validation tests to show that motions are still produced reliably given continuous initial conditions. In these validation tests, we test the mapping function $\mathscr{G}$ under a range of initial conditions and the mapping $\mathscr{K}$ to test the energy optimal step selection. In each case, initial conditions passed to one or both of the mapping functions and their outputs are used to run an episode of the dynamic simulation.

5.5 Results

Several sets of robustness tests were carried out to validate the motions of the robot under different forms of model discrepancies. The results of this validation are presented in this section.

5.5.1 Reachability Map Validation

Validation was run for two purposes: first, to verify the interpolation of parameters at random, novel points using the original robot model (Figure 5.4); second is to induce simulated discrepancies to the robot model to validate the robustness of the parameters to modelling disturbances.

In each case, 1000 random pairs of initial conditions, $[\dot{x}_0, s_{des}]$, were generated which were *not* the same as those tested in the previous phases over the entire initial search space. For each pair, an episode of dynamic simulation was run using parameters from the mapping $\mathscr{G}$and episodes were judged successful (coloured blue) or unsuccessful (coloured red) according to the termination conditions in Section 5.3.3.

5.5.1.1 Interpolated Parameter Validation

This validation uses the original robot model and tests the interpolated region between optimised points to verify that this region is in fact reliable for the robot used during optimisation. Figure 5.5c shows the random initial conditions

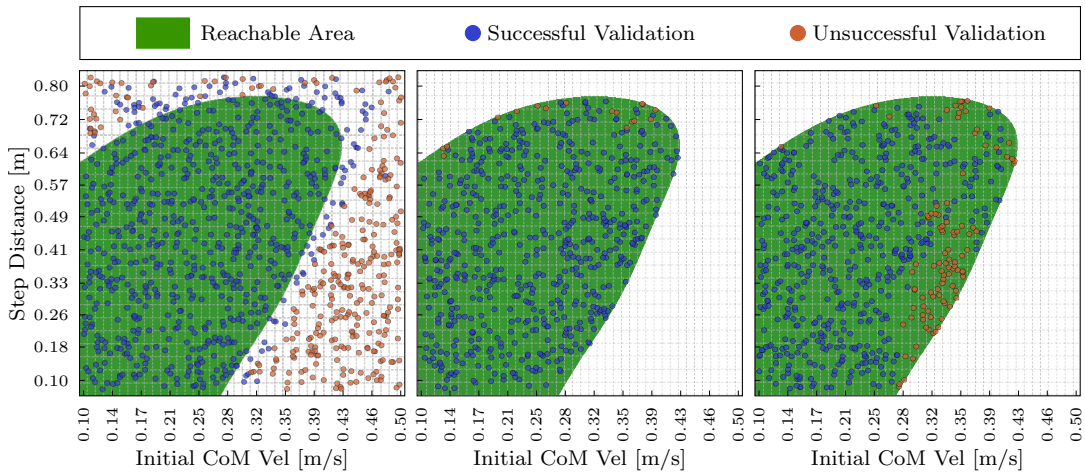

Figure 5.7: Validation testing mass discrepancies, only internal points shown: (a) original mass, (b) +10 % mass, (c) +20 % mass.

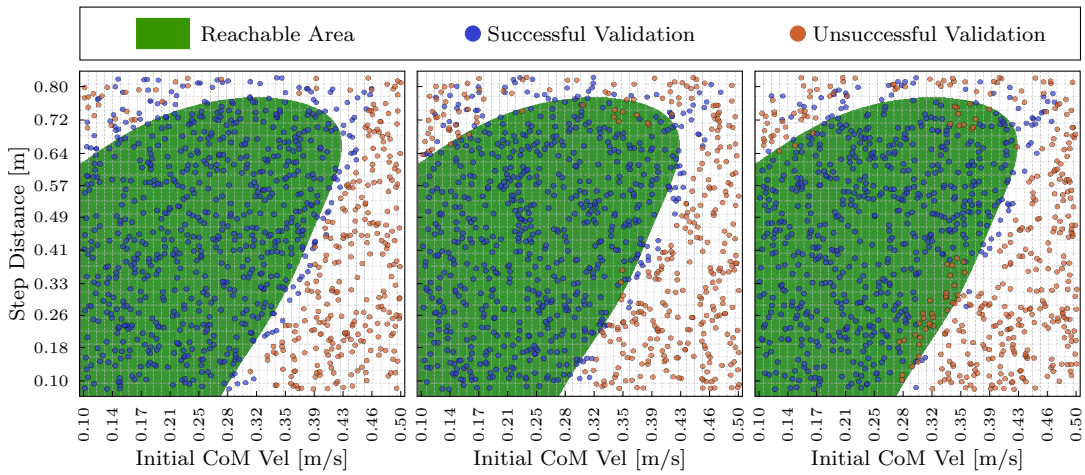

Figure 5.8: Validation testing mass discrepancies at non-uniform intervals, all points shown: (a) −17.3 % mass, (b) +13.1 % mass, (c) +18.5 % mass.

projected onto the trimmed reachability map, highlighting the noisy, uncertain regions for longer step distances and showing that the safe region is conservative given the region of successful trials surrounding this area.

5.5.1.2 Robustness Validation

Next, we introduce modelling errors to validate the robustness of the parameters to dynamic discrepancies.

To enhance the robustness of our analysis, we conducted simulations by modifying

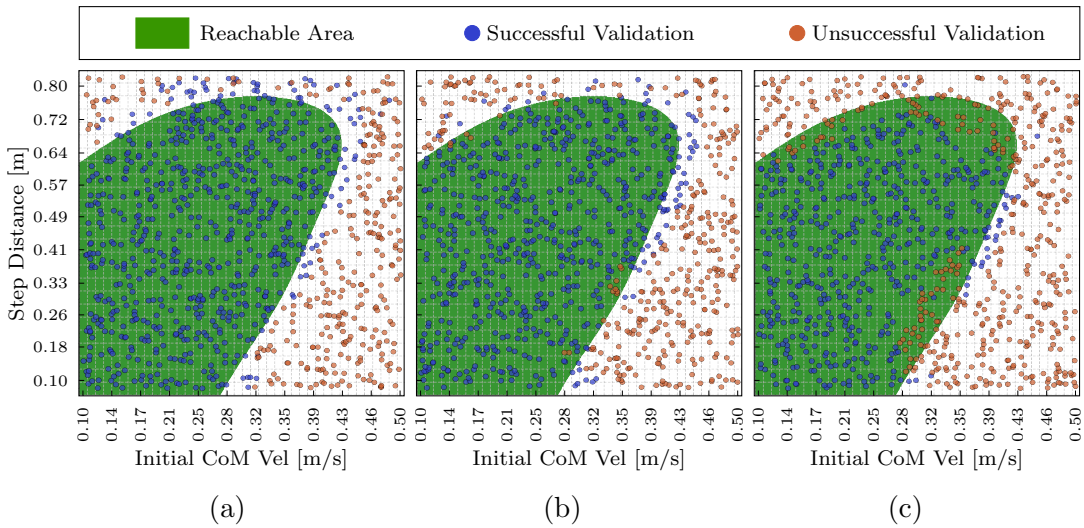

Figure 5.9: Robustness testing with control loop delays. (a) 10 ms, (b) 20 ms, (c) 30 ms.

the robot's mass. Although it is possible to examine alternative CoM heights for robustness analysis, it would require considering non-uniform mass modifications throughout the link, which could complicate the results. Therefore, we opted to alter only the mass of the links as it allowed us to investigate the impact of a single variable without introducing additional complexities.

The robot mass was altered $\pm 10\,\%, \pm 15\,\%, \pm 20\,\%$ of the original mass, with no changes to the HQP feedback gains and generated 1000 new random pairs of initial conditions for each mass variant. Plots show the validation trials projected onto the safe reachability region for the original mass robot. Mass variations were uniform and did not result in alterations of CoM height. Full plots are shown in the Supplementary Material (Figures S1 & S2) and omitted here for clarity. In all test cases where the mass is lower than the original robot ($-10\,\%, -15\,\%, -20\,\%$), an increasingly high number of trials outside the original safe region were successful, making the original safe region appear more conservative.

Figures 5.7b & 5.7c show tests of increased robot mass (10 %,20 %), with failed trials inside the original safe region (trials outside this region have been removed for clarity), the incidence of failure increases proportionally to the mass such that a more conservative safe region would be needed for these variations under the current settings. Given the large mass variation without changing the HQP feedback gains, this shows that even under modelling error, the optimised parameters still lead to successful stepping in a majority of cases.

5.5.2 Energy-Optimal Step Selection Validation

Energy optimal step selection was validated with 150 initial CoM velocities, randomly generated within the safe region ($0.1 : 0.43\,\mathrm{m/s}$). A dynamic simulation episode was executed for each point, using an energy-optimal step position s^*_{des}. All trials were successful as shown in Figure 5.11 and the attached video.

Figure 5.13 shows the measured joint positions and torques of the torso (trunk and pelvis), hip pitch, knee and ankle pitch joints in both legs during the validation shown in Figure 5.11c. All joints are included in the motion analysis, but the effect of arm and head joints are negligible so not analysed further. During the swing phase, swing leg joint torque is lower than in the stance leg which serves as a rigid pivot. The landing impact is clear in the measured torque (Figure 5.13 hatched area) and is close to the actuation limits which helps explain noisy regions in Figure 5.5b.

Validation experiments were carried out to investigate how changes in robot mass affect optimal step positions. Phase 2 of the pipeline was run with a range of mass variations using the optimal output from Phase 1 which is based on the original robot mass. Figure 5.6 shows regression models $+10\,\%$ (blue) and $-10\,\%$ (black) variation alongside that of the original robot mass (green). Optimised parameters, interpolation and HQP feedback gains as tuned for the original robot mass were used.

Negative mass variations result in optimal step positions at low initial CoM velocities being further from the robot and being steady for longer than for the original mass, then converging to the same location as initial CoM velocities increase. The variations by increasing the mass result in divergences away from the optimal step locations, which initially match those of the original mass, at around $0.14\,\mathrm{m/s}$, and diverge further from the original optimal as initial CoM velocity increases. These tests demonstrate that parameters are robust to modelling errors such as robot mass, and there is an underlying trend to the dynamics of a multi-body system such that the optimal step position can be calculated given the construction and dynamics of a robot.

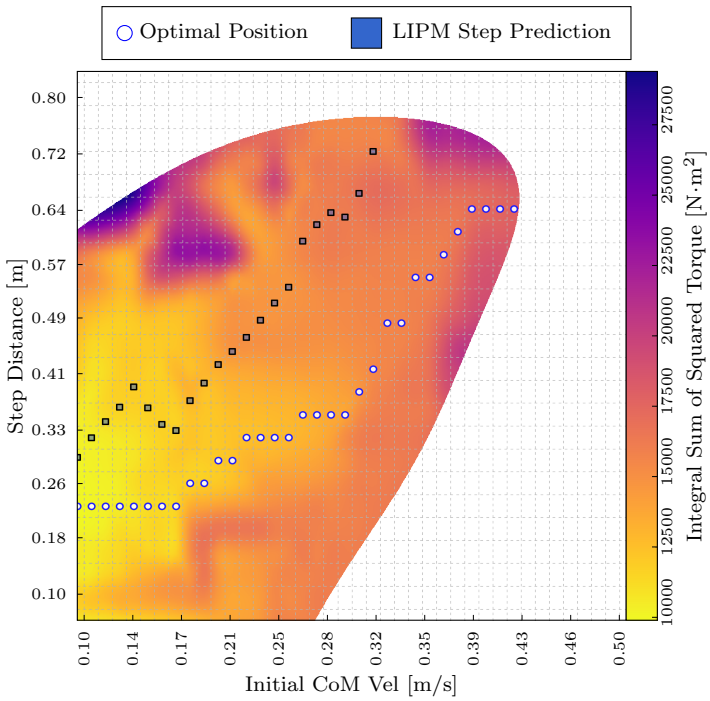

Figure 5.10: Map of the measured integral sum of squared joint torques during the swing phase to reach each position, with optimal step locations marked as circles and compared to predictions from the LIPM model, represented at the square points.

5.5.3 LIPM Comparison

We compared one-step-capturable step positions predicted by the LIPM model given an initial CoM velocity to the energy-optimal step positions from our optimisation, as shown in Figure 5.10. To do this, we queried the mapping $\mathscr{G}$ at the optimal step locations for each sampled initial CoM velocity, forward simulated the LIPM using these parameters and recorded the position of the capture point at the end of the swing time. Final capture points were projected onto the measured integral sum of squared joint torque map to extract the torque cost of moving to that location, we compared this to the measured integral sum of square joint torque of the energy-optimal step location.

After 0.32 m/s, the LIPM predicted step location is beyond the reachable area for the robot in our simulations, so are not included since the torque cannot be measured. We calculated the error in the integral sum of square joint torque between the optimal step positions from our optimisation and those from the LIPM prediction, with a root mean squared error of 1364.5 $\mathrm{N{\cdot}m^2}$(10.9 % increase)

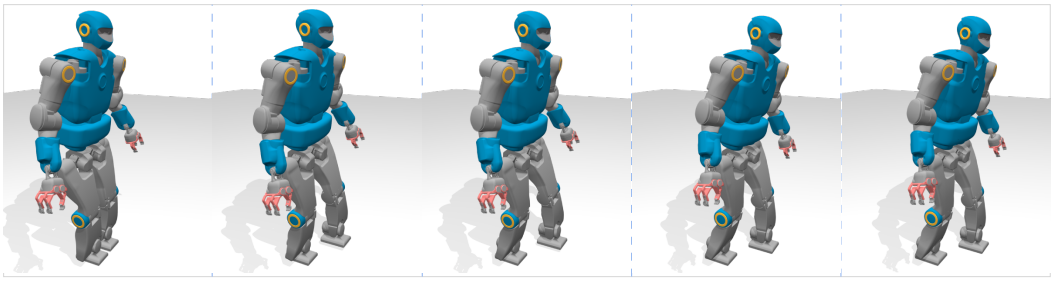

(a)

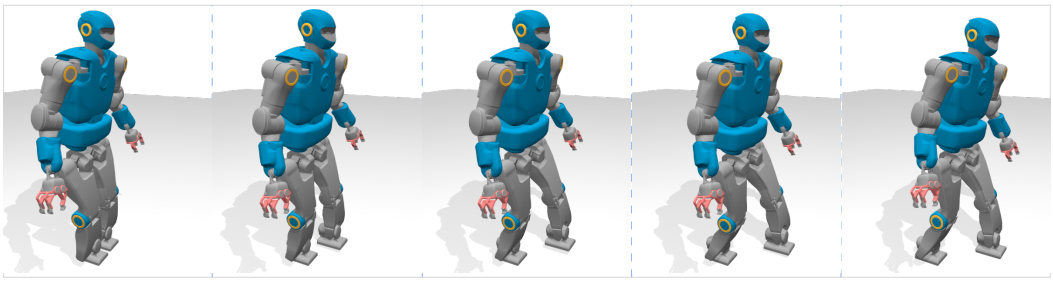

(b)

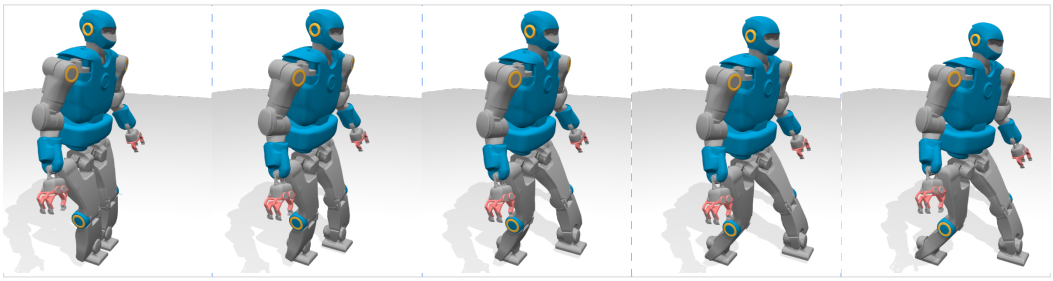

(c)

Figure 5.11: Validation of stepping motions using our automatic step selection validation. Initial CoM velocities 0.164 m/s, 0.292 m/s, 0.398 m/s, were mapped to step positions: 0.248 m, 0.409 m, 0.617 m for (a), (b) and (c) respectively.

(StD.=829.6 $\mathrm{N \cdot m^2}$(6.0 %), Min.=164.3 $\mathrm{N \cdot m^2}$(2.1 % increase), Max. =2531.9 $\mathrm{N \cdot m^2}$ (20.0 % increase)). Comparison of preferred step positions with the LIPM shows the effect of modelling discrepancies between using the full dynamic model and the simplified model. Each LIPM step location in Figure 5.10 uses the same step parameters as the optimised step location below it, which uses our full-body model. This demonstrates the extent to which optimality planned during LIPM motion planning is lost when tracking the trajectories to the full-body model. Moreover, using the LIPM to predict step locations reduces the potential workspace of the

robot, potentially harming the one-step recovery ability of the robot.

5.5.4 Diversity in Efficient Step Selection

In addition to the simple trend of energy-optimal step positions, a diverse range of near-optimal step locations are present on the map. The highlighted area in Figure 5.6 shows step positions where energy costs are a maximum of 5 % above optimal for the same initial condition and Figure 5.12 shows a band of step positions superimposed onto Figure 5.6 which deviates from the optimal to a maximum of 10 %, as a result, any step positions in this range are defined as near-optimal.

Near-optimal regions form simple heuristics: for low initial CoM velocities, stepping between 22 cm - 30 cm will result in minimal impact on energy-efficiency, even if the optimal position is not reached, allowing coarse, yet rapid step selection with trivial changes in energy-efficiency. Additionally, since the regions span a range of initial CoM velocities, inaccurate CoM state estimation can still lead to efficient stepping.

5.5.5 Underlying Energy Optimality

Near-optimal efficiency regions suggest an underlying optimality in dynamic stepping gives us insight into human step selection. If similar regions exist in dynamic multi-body systems in general, humans can learn similar heuristics and use them for simple, rapid, near-optimal step selection. Having near-optimal heuristics for step selection that are robust to sensing delays and inaccuracies would be beneficial to developing biological humans.

Figure 5.1b shows a clustering of selected steps similar to that in the highlighted area in Figure 5.6, but with a higher range of initial CoM velocities. Humans can withstand higher magnitudes using foot tilting behaviour. This is a limitation of our robot control, which does not consider underactuated foot tilting control, hindering the range of feasible step positions.

We also gain insight into energy optimality by looking at the optimised step parameters. Total swing time can be calculated from the optimal parameters ($t_{swing\_start} + s_{des}/s_{speed}$), and plotted in Figure 5.14b, showing all initial condition pairs and their optimised swing time, with optimal points projected into 2D in

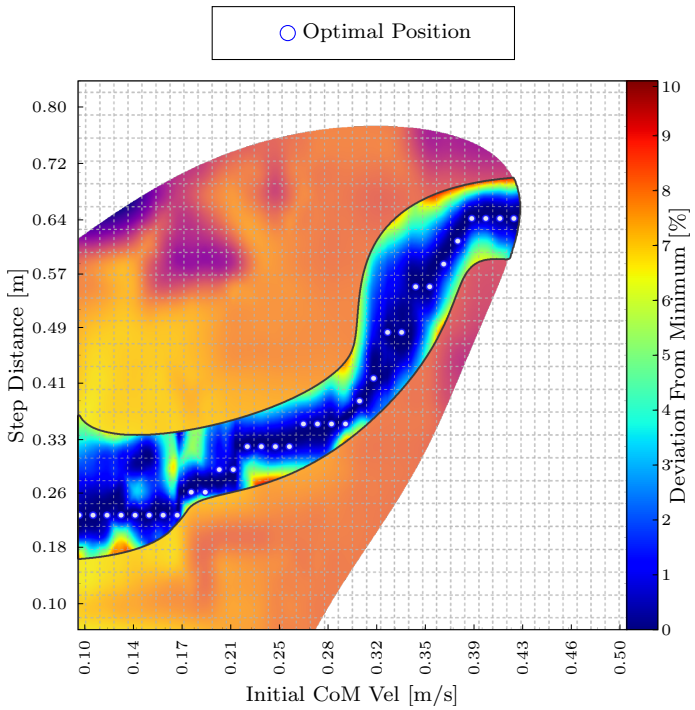

Figure 5.12: Map of the measured integral sum of squared joint torques during the swing phase to reach each position, with optimal step locations marked. The highlighted region shows step locations which deviate from the optimal by a maximum 10% for the same initial CoM velocity superimposed onto 5.6.

Figure 5.14c, from the initial foot position $[0, 0, 1]$cm, swing time is dependent on the desired step position, displayed on the vertical axis. This shows a piece-wise relationship between initial CoM velocity and optimal swing time, where swing time initially drops, levels off, then rises as the initial CoM velocity increases. The majority of the swing time effect is caused by the swing speed parameter, s_{swing_speed}, with only minor changes induced by the swing start time, t_{swing_start}, as shown by the optimal ranges in Table 5.1.

5.5.6 Analysis of Computation Time

Training used an Intel Core i7-8700k with 12 cores (6 physical), 32 GB RAM, Ubuntu 16.04 and Pinocchio 2.5.0. The pipeline is parallelized, with each core optimising all step positions for one initial CoM velocity, and took around 5 hours, reachability map building takes around 40 minutes. Parameters for one swing foot were used for the opposite foot.

Figure 5.14a shows the computation time for parameter optimisation for each

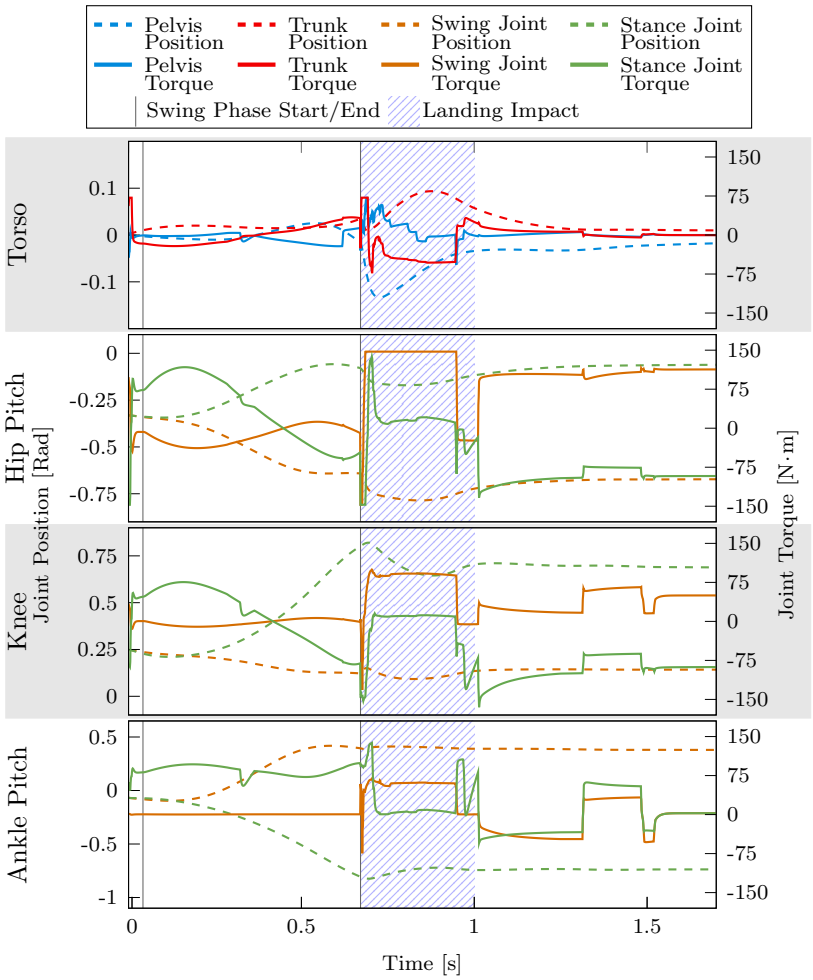

Figure 5.13: Joint profiles in the torso and legs during one validation stepping motion in Figure 5.11c.

pair of initial conditions, with 170 episodes for each pair. For larger initial CoM velocities and larger step distances, the number of early terminations leads to quicker computation. Parallelization scales linearly with the number of cores, where large-scale distribution, with one core per initial condition pair, would lead to 45 min optimisation.

Swing foot trajectories are generated in 0.5 ms, and CoM trajectories in 1.4 s, due to the nonlinear optimisation formation used for generation, but can be reduced to 0.5 ms using the same generation as the swing foot. Querying the energy-optimal step selection takes 0.13 ms.

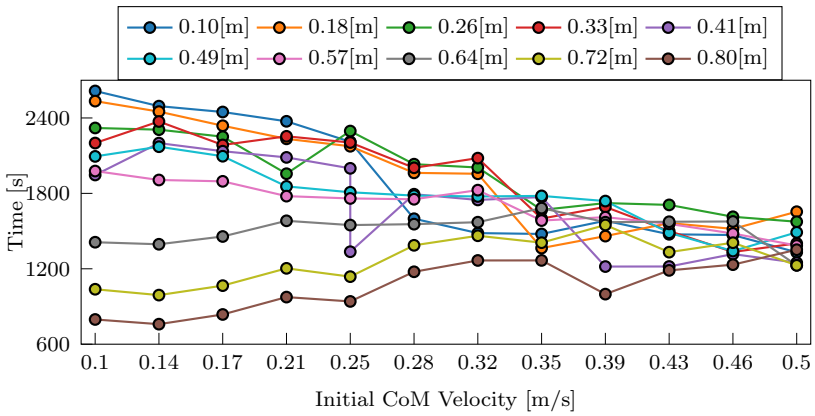

(a)

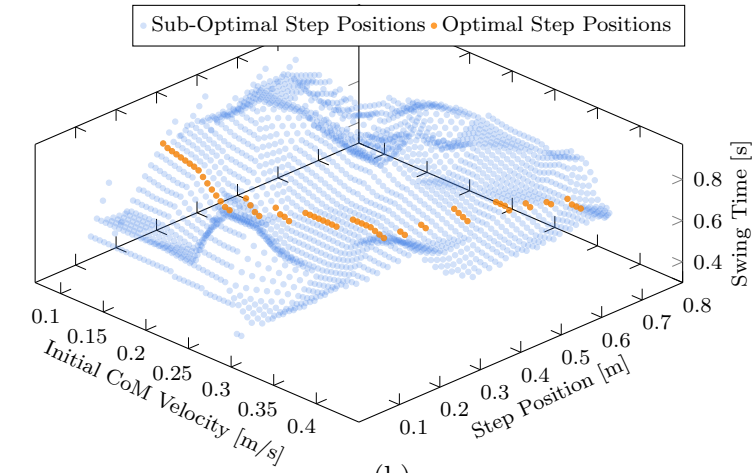

(b)

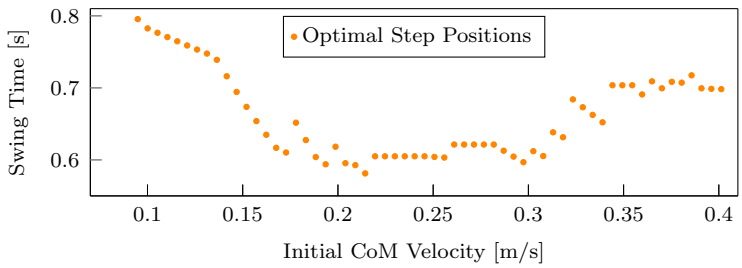

(c)

Figure 5.14: Analysis of pipeline performance and parameter results: (a) Computational time to optimise parameters for each pair of initial conditions, step positions denoted by coloured traces; (b) 3D plot of swing time for optimal and suboptimal points; (c) 2D projection of swing time for optimal points.

5.6 Discussion

The study of energy-efficient locomotion is a complex process that involves the whole-body dynamics and control together as a whole and requires complex op-

timisation, as well as a global search of global energy-optimal step location and timing, which are all highly nonlinear. The complex interplay between multi-body dynamics, control and gait parameters is shown by the optimal gait parameters in Figure 5.14b, where this nonlinear relationship and non-smooth gradient would cause standard gradient search methods to get stuck in the local minima. However, the mapping between initial conditions and key gait parameters, such as step location and swing time, suggests that the gait parameters can be represented by piece-wise approximations, as shown in Figures 5.14b & 5.14c, resulting in a method for rapid step selection. This indicates that despite the complexity of the whole process, this nonlinear relationship in human gait can be possibly learned by humans by prior trials and experience.

Additionally, balance can be recovered using similar step locations and swing times for different initial CoM velocities by trading off energy optimality, which can potentially explain the large variations in step location in human study [5].

5.7 Conclusion

This chapter describes our investigation of energy-efficient step selection using nonlinear optimization to construct offline reachability maps. Our findings revealed that selecting energy-efficient steps during push recovery, or identifying a set of diverse stepping regions, is challenging to characterize using simple models and difficult to compute online, as shown in Figure 5.14a. Therefore, reachability maps can be employed for swift step selection on flat ground. Moreover, the results provide valuable insights into the feasibility of diverse step selection for humanoid robots. In future work, we plan to extend this pipeline to study energy-efficient locomotion in various modes and implement the Query of Mapping for warm-start solutions in online optimization. To further advance this line of research, the next chapter will explore this principle at a more general level, examining continuous walking to determine if this trend is maintained at a broader scale.

Chapter 6

Optimality in Continuous Walking

6.1 Introduction

For a limited locomotion case, where the robot has an initial CoM velocity and takes a single step, there is a region of attraction for walking efficiency, where simple heuristics can map an initial CoM velocity to an energy-optimal next step location, as shown in the previous chapter. This chapter, expands on this work for a continuous walking case, to study the relationship between walking parameters and energy efficiency. Within this work, we again highlight the advantages of using mechatronic humanoids as a tool for investigating walking behaviour in the humanoid morphology, as this study would not be possible on a biological human, though we may gain insights into the underlying principles of locomotion which also apply to the biological case

Continuous walking, much like single steps presents a problem of non-uniqueness and redundancy in the space of possible parameters. Given the infinite possible combinations of walking parameters available to humans, how and why do they decide on those that they do, and how can we apply this knowledge to robots?

In biology, studies consistently show that human step parameter choice tends towards efficiency during walking [146, 147, 148] and even running [149, 150]. The energetic cost, represented by CoT, forms a convex U-shaped curve [151, 152]

when plotting against walking speed, where there is a uniquely energy-optimal set of parameters.

6.1.1 Efficiency in Human Locomotion

Self-selected, or preferred, walking speed is energy-optimal under unconstrained conditions, but is also remarkably consistent under constrained conditions or altered dynamics. When walking with a spring-assisted 'reduced gravity' device, overall energy expenditure decreased, but the energy curve convexity is maintained [153]. Carrying loads to temporarily increase body mass also shows consistency in the energy optimality curve and efficient movement [154], as does the differing internal body mass [155]. Even after lower-limb amputations, [156], the trend is preserved and subjects still walk efficiently even with these constraints. The human ability to walk efficiently even under constrained conditions is surprisingly robust.

Though energy-minimisation is shown to be a feature of human locomotion, it is unlikely to be the only consideration. When walking downhill [157], on sand [158] on smooth, hard, and other inclined surfaces [159], as well as rough terrain [160], other components, such as fatigue minimisation[161], or stability are proposed as additional factors in how biological humans produce motion. In Gast et al's work [160], an optimisation function was proposed to include energy output as one of, but not the sole factor in gait parameter selection:

$$MIN_{\{s\}}[w_1 \times COT(s) + w_2 \times 1/Stability(s) + w_3 \times Time(s)] \tag{6.1}$$

Where s is walking speed, $w_{1,2,3}$ are scalar importance weights, $Time$ is a function for desired completion time, and $Stability$ a function to compute a desired stability margin metric. This formulation is intuitive and represents human walking parameter selection as an optimisation problem, weighted according to their respective importance in the current context to output a walking speed that minimises the weighted terms accordingly. Though it is difficult to determine how weights, stability and completion time are selected, and uncertain if this could be computed quickly enough, or at a high enough frequency to explain the reactive speed and efficient nature of human walking.

6.1.2 Metabolic Regulation

Uncertainty also surrounds the mechanism by which humans sense and regulate their energy expenditure during motion at reactive speeds. Physiologically, humans have blood [162, 163, 164] and muscular [165, 166] metaboreceptors to sense metabolic output, but such sensors are slow-acting, as they are sensitive to metabolic by-products [167] such a lactic acid, which take time to metabolise and reach the receptors, thus cannot account for the speed at which humans regulate their walking efficiency. Humans return to optimally efficient walking within 1-2 seconds of artificial constraints being removed, for example by an exoskeleton [168], thus these sensors are unlikely to be the primary drivers of a rapid energy-optimal feedback loop. This same study suggests that a feed-forward energy-cost prediction is a more suitable explanation for this rapid adaptation.

Based on the results of the previous chapter we proposed that human physiology is such that accurate sensing and regulation are not strictly necessary for locomotion and that the dynamical properties of the morphology are biased towards efficient motion, such that efficient locomotion is the "path of least resistance" and the body falls into this much like a dynamic walker moves forward by virtue of its dynamics, without the need for high-frequency feedback control. Neuromuscular investigation into locomotion behaviour also showed that the musculoskeletal structure of the human leg may contribute to walking economy, instead of this being a purely neural process [169]. A passive efficiency in the humanoid morphology might imply that, at a neuronal level, there is no need to explicitly engage in complex, online muscular planning; instead, all is needed is to learn simple, easy-to-compute heuristics which, when passed to the muscles, produce efficient motion. This is evidenced by 'habituality' in human motion, where motions are not re-planned to account for changes in functionality when muscles are temporarily paralysed using electrical interference during reaching tasks [170], suggesting that muscle-level actuation and optimisation may be learned and not planned in real-time.

6.1.3 Evolutionary Perspective

Evolutionary biology has shown that the motion of the human ancestor *Australopithecus afarensis* was also an optimally efficient walker in terms of energy expenditure [171], and that there is a positive trend towards efficiency through-

out the evolutionary timeline in humans [172] and animals [85, 173, 174], possibly driven, in part, by morphological adaptations during evolution [151]. Motion efficiency would be a key factor in survivability in the calorie-sparse conditions of a hunter-gatherer lifestyle and being able to move for longer and further than a competitor given the same amount of energy would be a distinct advantage.

Evidence of the positive relationship between efficiency and evolution and optimality in modern studies presents a convincing picture that biological motion very likely tends toward energy optimality but questions and issues remain. As we have explored above, the question of *how* humans produce optimally efficient motion is unsolved, as is the question of is efficiency a by-product of another cost function, as suggested in other cases including human running [175].

6.1.4 Global and Local Optimality

Further, there is a tendency of biological-centric research to focus on global optimality in efficiency, whereas, we saw in the Chapter 5 that there is value in analysing the local optimality of motions. Much of the work is focused on how humans perform in different conditions relative to a known global-optima during free walking with preferred gait. For example, walking efficiency on rough terrain is compared to flat ground, and framed as 0.5 % loss of efficiency from the global optimality [160]. But, this framing may be obscuring an underlying trend within the human locomotion efficiency landscape. Chapter 5 showed that for a given CoM velocity there still exists a set of walking parameters (step length and step speed) that is locally optimal for that velocity, but not globally optimal. Hence, if constraints are imposed on preferred step length, this may have a lower global-optimality, but even within these constraints other parameters must be chosen, and we showed previously that locally optimal parameters can still be chosen in the presence of constraints.

The reason that this is an issue for biological study is that in order to assess local optimality in biological studies a global map is required similar to those in Chapter 5. Experiments must sample human walking at a range of different speeds, step lengths and step speeds that are not natural to the human and measure energy cost at each, in a dense sample grid, whilst accounting for fatigue, careful compliance to protocol, ethics, and individual differences. Studies on human step parameter selection have shown a linear relationship between walking

speed and step length [176] over a number of gait cycles. But unfortunately, there is a high level of stochasticity and individual differences, making it difficult to draw conclusions about the general walking case.

6.1.5 Mathematical Models of Efficiency

Mathematical models provide an alternative means of building global efficiency maps. Efficiency maps for walking parameters in locomotion are primarily built on template models of motion and can augment human data. A modified compliant walking model [177] can produce global walking maps [178] and showed a complex energy landscape, but was based on complex Fourier transforms of ground reaction force progression difficult to apply to humans. Using a pendulum template model to optimise efficiency, it is possible to show the optimal gait pattern for a given speed but does not result in a quantified global energy map [179]. Passive walking models can produce detailed energy cost maps with gradient contours showing the difference from optimal [180], where three passive walker models were compared. In this same study, the authors suggested such models are systematic methods of energy landscape estimation and a closer approximation to human dynamics is required. This has been explored using an insect robot model in simulation to build energy landscapes for different scenarios [74], studied in quadrupeds [181], and is explored extensively in this chapter.

6.1.6 Problem Definition and Scope

In this chapter, we expand upon these mathematically based maps and use a full robot model to systematically build a high resolution global energy landscape for walking parameters in a highly accurate simulation. Augmenting the motivation for using robots in the previous chapter, this is a useful tool for this research, as walking aims of the robot are not contaminated by extraneous optimisation goals, we can precisely sample the search space and accurately measure energy cost.

We use three scenarios, each with a different combination of robots and locomotion control to explore the energy landscapes of each and analyse the commonalities between them. By sequentially altering the walking parameters of the robot in each case, we explore the whole state space of the chosen walking behaviour and measure the energy cost of each pair of parameters in each case and study

the characteristics shared between them.

Studying a general walking case is not possible for this project, so the scope of this chapter is constrained to walking forward on flat ground. Full dynamic simulations are used to accurately model the robot dynamics. Gait is initiated from a quiet-standing start in each trial, but this data is discarded, and we focus maintained walking behaviour.

6.1.7 Contributions

1. **Energy-Optimal Walking (Section 6.3.3)** Feasibility maps to show the local and global energy landscapes of walking parameters for three scenarios, using two robots and two control schemes;

2. **Cost of Transport Regions of Attraction as an Invariant Principle of Locomotion (Section 6.5)** Based on our results we propose that the dynamics of the humanoid morphology may have a passive tendency towards producing CoT-optimal motion, evidenced by multiple cases though global search;

3. **Heuristics for Globally and Locally Optimal Stepping (Section 6.5.4)** Analysis of these efficiency maps reveals hierarchical step parameter heuristic for producing efficient motion in unconstrained (globally optimal) or constrained (locally optimal) cases;

4. **Robots as Analogues for Human Locomotion Research (Section 6.5.1)** We demonstrate that mechatronic humanoids are an effective testbed for researching biological humans and can offer deep insight into the humanoid body in the form of invariant energy maps that would be impossible to generate with biological subjects.

6.2 Methodology

To study the energy efficiency landscape of humanoid robots during locomotion in a systematic manner, we used three scenarios with different robot platforms and controllers. Using different models increases the data points we gain about the general humanoid structure and allows us to gain more information about the humanoid morphology as a whole. Different controllers increase this further,

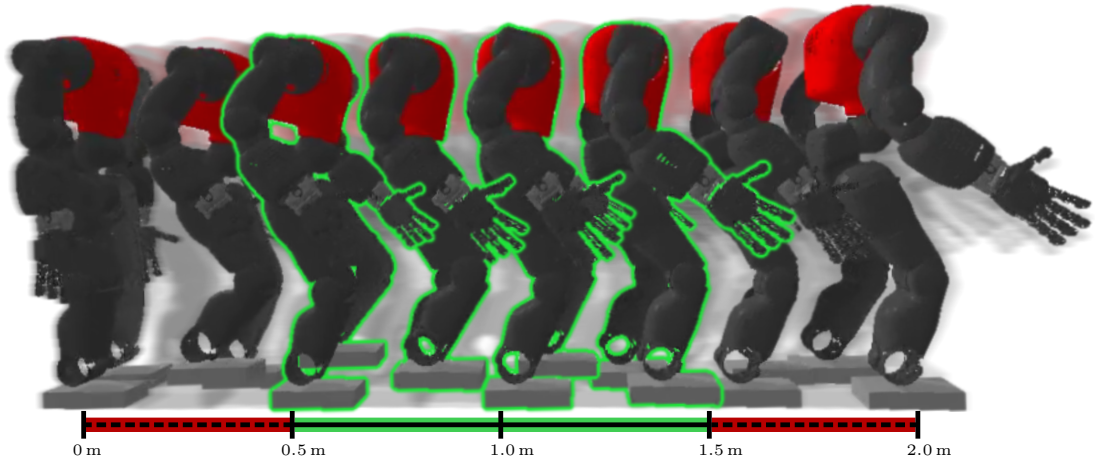

Figure 6.1: Walking trials are run for 2 m and if the robot completes this motion, data from the middle 1 m segment of motion is taken as the trial data.

since, any commonalities between control scenarios is more likely to be representative of the humanoid dynamic system rather than an artefact of a single controller. This section presents our methodology for this systematic study and the conditions under which it was conducted.

6.2.1 Experimental Setup

In this experiment, there are three scenarios in which either the robot platform, or its controller is altered, with the same procedure across all cases. Due to the size of the walking state space, or experiments are constrained to walking in the lateral plane (X-axis in all cases), wherein the robot walks forwards.

The controller in each Scenario had 2 open parameters and all other parameters were constant throughout the trials. Each scenario consists of a number of trials equal to and indexed by the number of open parameters. In each trial, the robot is first initialised in the simulator. During initialisation, the robot joint configuration is set to a quiet standing pose with the legs and arms straight, feet slightly apart and flat on the ground. In world-frame, the CoM position is located at $x, y = 0m$ and the robot's nominal standing height in the Z axis.

After initialisation, the controller is instantiated and the two open parameters are set by the sampling protocol detailed below. In each trial the aim is to walk forwards until the CoM position reached 2 m in the X axis. The two termination conditions were:

1. The CoM_x position reaches 2 m (Success)
2. The CoM_z position drops below a threshold z_{min} indicating the robot has fallen (Failure).
3. Total simulation time exceeds 90 seconds (Failure)

Data is recorded throughout each trial and is indexed by the open parameters used for that trial. We measure the CoM position and velocity in x, y, z axes, the joint positions, velocities, accelerations and torques, step locations and GRFs. After each trial, regardless of success, the data is recorded and the next trial begins until all trials have been completed.

6.2.2 Data Processing

After all trials have been completed, the data is processed to study the characteristics of each walking scenario and its energy profile.

6.2.2.1 Data Trimming

Data is first separated by trial success or failure and only trials in which the robot successfully reached $max(CoM_x) \geq 2\,\text{m}$ and $min(CoM_z) > z_{min}$ are included in the analysis, with all others as failed trials. Each trial is then partitioned to extract a representative 1 m section from the middle of the whole walking motion, between 0.5 m and 1.5 m in the CoM X axis position, as shown in Figure 6.1. The start of the motion is disregarded to remove any artefacts from gait initialisation, as the motion took some time to accelerate to a steady state, especially when rapid step parameters were used. Further, the motion section between 1.5 m and 2 m was present in the trials to ensure stability was maintained after the usable sampling period and that the 1 m section of interest was not simply a stable part of a larger unstable motion. Thus, the 1.5 m and 2 m was removed, leaving a consistent 1 m section of walking data from each trial.

6.2.2.2 Walking Velocity

The walking velocity was not directly controlled by the open parameters and as such was calculated from the data. Walking velocity, $\dot{x}$, was calculated as a mean velocity in the X-axis at the CoM across the representative 1 m sampled area, as follows:

$$\dot{x} = \frac{x_f - x_0}{t_f - t_0} \tag{6.2}$$

where $(x_f - x_0)$ is difference in the CoM_x position at the start and end of the sampled motion (always 1 m), and $t_f - t_0$ is the time taken to travel across the 1 m sampled area.

In the results section, some plots show the relationship between parameters and walking velocity. Here, the term CoM Velocity is used to refer to this mean walking velocity over the sampled distance. Since CoM velocities were not sampled directly, the resultant number of CoM velocities does not necessarily match the number of sampled parameters. As such, when CoM velocities are plotted against the uniformly spaced step parameters, they are assigned to bins, each representing a small range of approximate CoM velocities to maintain the regular grid structure of the data. In cases where more than one parameter was mapped to the same bin, the mean value of the bin is taken, though this was relatively rare, as a high bin resolution was used.

6.2.2.3 Energy Consumption

As in the previous chapter, Eq. 5.6, the squared integral of joint torque is used as an analogue for energy consumption, as it represents the battery-level expenditure to control the joints when both accelerating and decelerating:

$$\int_{t_0}^{t_f} \boldsymbol{\tau}^2 \, dt \tag{6.3}$$

where t_0 and t_f are the start and end times of the sampled distance and $\boldsymbol{\tau}^2$ is the matrix of square joint torques. The squared integral torque is calculated for each successful trial across the sampled distance.

In the previous chapter, this calculation was used as part of the control cost function and for analysis, whereas in this case, this is used as an analytical metric only and the controllers do not explicitly solve for efficiency.

6.2.2.4 Cost of Transport

Cost of Transport is a dimensionless metric used to compare the efficiency of different morphologies and speed during motion and is a function of the energy

(a) (b) (c)

Figure 6.2: Snapshots of robots and environments: (a) CoMAN robot in PyBullet environment, (b) TALOS robot in PyBullet environment, (c) TALOS robot in croccodyll environment.

or power used to move a given distance or velocity, whilst accounting for the mass of the body in motion.

Since, in this case, there is no direct measurement of battery output, the CoT a close approximation is calculated using the squared integral torque as above as an analogue for power consumption as follows:

$$\text{CoT} = \frac{P}{mgv} \equiv \frac{\int_{t_0}^{t_f} \boldsymbol{\tau}^2 \, dt}{mgv} \tag{6.4}$$

where P is the power input into the system, m is the robot mass, as shown in Table 6.2, g is the gravity constant ($g = 9.81\,\text{m/s}^2$) and v is the mean walking velocity as described above.

6.2.3 Controllers

Two control methods were employed for this investigation, using existing implementations, one model-based feedback controller and one optimal control trajectory optimisation method. Both are detailed in this section.

6.2.3.1 Model-Based Feedback Controller

The model-based feedback controller was implemented by Kasaei et al. [50, 77] with minor adjustments for this investigation. The controller uses an architecture built of reduced models for step planning, trajectory planning and low-level control, where a LIPM model with the robot's parameters describes the dynamics

Table 6.1: Physical properties of the CoMAN and TALOS robots.

					Mass Distribution [%]		
Robot	DoF	Height [m]	Leg Length [m]	Mass [kg]	Torso	Arms	Legs
COMAN	29	0.945	0.54	31	35.3	30.8	33.4
TALOS	32	1.75	0.85	89	40.3	20.6	39.1

of the robot during locomotion given the desired parameters. A Partial Fourier Series (PFS) is then generated using the LIPM parameters, the PFS then acts as a Centroidal Pattern Generator. When coupled with a state machine to track current and next state actions, and an Inverse Kinematic solver to resolve joint positions, the architecture can produce omnidirectional locomotion.

For this investigation, the control parameters are fixed to constrain the direction of motion to only move forwards, except for the lateral motion required for weight transfer. The only open parameters were step distance and step timing, which were supplied to the controller as detailed in Section 6.2.1. Supplied parameters were used as high-level controls to determine the formulation of the CPG to achieve the given parameters. In all cases, sampling time: 0.005 ms (200 Hz) and control frequency: 0.005 ms (200 Hz)

The PyBullet [182] environment was coupled with this controller. PyBullet is a popular full dynamic simulation from which dynamic metrics are recorded. Scenarios which show this environment are shown in Figures 6.2a and 6.2b.

6.2.3.2 Optimal Controller

For optimal control, the Crocoddyl [43] library was used. Using Differential Dynamic Programming (DDP) and an auto-differentiation solver, Crocoddyl solves for both optimal trajectories and optimal feedback gains for multi-contact scenarios in the Pinocchio dynamics library [143].

The optimisation problem solves for individual gait cycles, which are specified manually to meet the requirements of our study. For each gait cycle, the following data was required: step length s_x, step height s_z, step time Δ_s (duration of optimisation knots), and step knots s_k (the number of knots per single-support phase). The number of knots for each double support phase was set to 10 in all

cases.

Step length was set to the desired step length for a given trial $s_x = s_{des_i}$, s_z was set to 0.1 m, step knot duration Δ_s was set to 0.03 s and the duration of the swing phase was adjusted by altering the number of step knots

$$s_k = \lfloor \frac{t_{des_j}}{\Delta_s} \rfloor \tag{6.5}$$

such that, after rounding, the number of knots of duration Δ_s matched the desired step time t_{des_j}.

The objective function to be minimized in this case was:

$$J = w_1|y - y_{ref}|^2 + w_2|u|^2 \tag{6.6}$$

where J is the total cost to be minimized, y is the state vector, y_{ref} is the reference state vector, u is the control input vector, and $w_1 = 1$ and $w_2 = 0.1$ are weights for tracking and energy costs respectively.

It is worth noting that in this case, motions are validated using classical equations of motion, while in the other cases, the dynamic PyBullet simulator is used. When validating solely using equations of motion, certain additional interaction and contact forces are not taken into account, whereas these factors are considered in dynamic simulation. This difference may have an impact on the comparison of the two scenarios, which is discussed further in Section 6.5.

6.2.4 Robots

Two robot models were used, the CoMAN [183] robot and the TALOS [3] robot. Both robots have a humanoid morphology with different dimensions and dynamic properties. The properties of both are shown in Table 6.1, which shows that the CoMAN robot is smaller, with a lower mass than the TALOS, which is around twice the height, with almost three times the mass. Mass distributions in both robots are also interesting differences, the CoMAN robot is more top-heavy than the TALOS, but in general is more evenly distributed. TALOS has heavier legs proportional to the rest of its mass, which has an effect on the walking dynamics and potentially a vastly different energy profile.

Table 6.2: Properties and parameters of each experimental scenario

				Parameter Values								
				Step Pos				Step Time				
Scenario	Robot	Controller	Simulator	Min[m]	Max[m]	Diff[m]	N	Min[s]	Max[s]	Diff[s]	N	Total N
1	COMAN	Model-Based	PyBullet	0.01	0.15	0.001	141	0.15	0.645	0.005	100	14100
2	TALOS	Model-Based	PyBullet	0.01	0.5	0.005	100	0.05	1.0	0.01	96	9600
3	TALOS	Optimal	Pinocchio	0.05	0.7	0.0132	50	0.2	1.6	0.0286	50	2500

Mass distribution in the leg segments of CoMAN are relatively close to humans, with a mean of 32.2 % [84], with TALOS being more dis-proportioned. Difference are more apparent between the robots and humans in the concentration of mass in the upper body, where human arms account for 9.8 % [84] on average, there is a marked difference in that of both robots. When combined, human torso and head on average accounts for 58 % of bodily mass [84], from which TALOS is the closer of the two.

These differences from humans are from the means of human cadaver studies [84] and there variance between different humans that we hope to capture by using different robot platforms.

6.2.5 Scenarios

By using different robots, simulator environments and controllers, we hope to gain more insight into the invariant effects of the humanoid morphology between these cases than would be observed from a single case study.

The details of each scenario are presented in Table 6.2. The feedback controller is used for the CoMAN and TALOS robots, and the TALOS robot is also used for the optimal control case.

Table 6.2 details the sampled parameter values for each scenario, with the range and interval of the step positions and step timings shown. Ideally, there would be an equal number of trials, N for each scenario, however, this was limited by computational performance and software instability, particularly in the optimal control case which was prone to failure for sections of the state-space. As a result, the resolution of trials was reduced to ensure the state space could still be covered.

Table 6.3: Descriptive statistics for experimental scenarios.

Scenario	Successful N	Total N	Success [%]	Min Vel.[m/s]	Mean Vel.[m/s]	Max Vel.[m/s]
1	9174	14100	65.06	0.015	0.147	0.391
2	1505	9600	15.68	0.013	0.145	0.557
3	2302	2500	92.08	0.013	0.176	0.641

6.3 Results

In this chapter, results are presented primarily as area plots showing the relationship between two parameters on the X and Y axes and the corresponding values between the pairs are colour coded according to its relative or absolute value as indicated on each plot. Plots for each scenario in Table 6.2 are shown in numerical order in their own sub-figure.

Although the colour gradient used to denote values are consistent within figures, the scale of each colourbar applies only to that subplot, as a consistent figure-wide colour gradient obscures much of the detail. This is also true of the scale of plot axes, as each is adjusted and scaled according to its own parameters. The axis scales are not consistent between plots, as each scenario has a different range of sampled parameters and using an absolute scale would reduce visibility. Non-shaded areas in the parameter-based plots denote failed trials, where one of the termination conditions was met during the trial and therefore was deemed unsuccessful and excluded from the results.

Some plots, especially those showing squared torque integral contain some trials with values dramatically higher than the mean, which affected the colour gradient representation. Removing such outliers was a possibility, but this would mean omitting some important characteristics of the results; instead we decided to scale the colourbars to show this gradient. A power-law was used to scale the colour bar, which clusters higher values towards the top of the colour bar, giving a better impression of the gradient. To avoid misrepresenting the data, 3D plots have also been included to give an alternative view of the gradient; these plots are shown in the Appendix C to improve readability of this chapter.

6.3.1 Descriptive Statistics

Statistics to show the walking performance for each scenario are shown in Table 6.3. The number of successful trials was markedly higher in scenario 3, as the motion itself is optimised for feasibility over multiple passes, whereas in the model-based cases, scenarios 1 & 2, the parameters are pushed to their limits. This is especially noticeable in scenario 2, as only 15 % of the trials were successful, but this gives a misleading picture, as the sampled parameters were set to the limit of the operating region from pilot trials, and looking at Figure 6.3b, we see that after a step distance of around 0.1 m, the band of step times that result in successful trials vastly decreases, leaving only a narrow band of successful trials after this bottleneck.

It is important to note that the bottleneck observed in scenario 2 could be a result of either the controller parameters not being properly adjusted to this particular robot, or that the different dynamic parameters of this robot may render the reduced controller unsuitable for this specific task. Further investigation and experimentation would be necessary to determine the cause of the bottleneck and find a solution to improve the success rate of the trials.

Figure 6.3 gives a graphical representation of the performance of the robot in each scenario, as well as the walking velocity achieved in each parameter pairing. Note that there is some discontinuity in the results for scenario 1, Figure 6.3a for longer step distances. This was caused by instability during the gait cycle due to the walking speed. Rather than removing these as outliers, they instead show the blurred boundary between successful and unsuccessful trials, where the controller is able to complete the motion at the very edge of its capabilities similar to the bottleneck shape of 6.3b, where the range of feasible step times quickly decreases as stride length increases.

6.3.2 Torque Consumption

We explore the relationship between walking parameters and efficiency from both an absolute-torque point of view and from a Cost of Transport angle. This section presents the results and analysis of the torque-consumption landscape and how it is affected by the tested walking parameters.

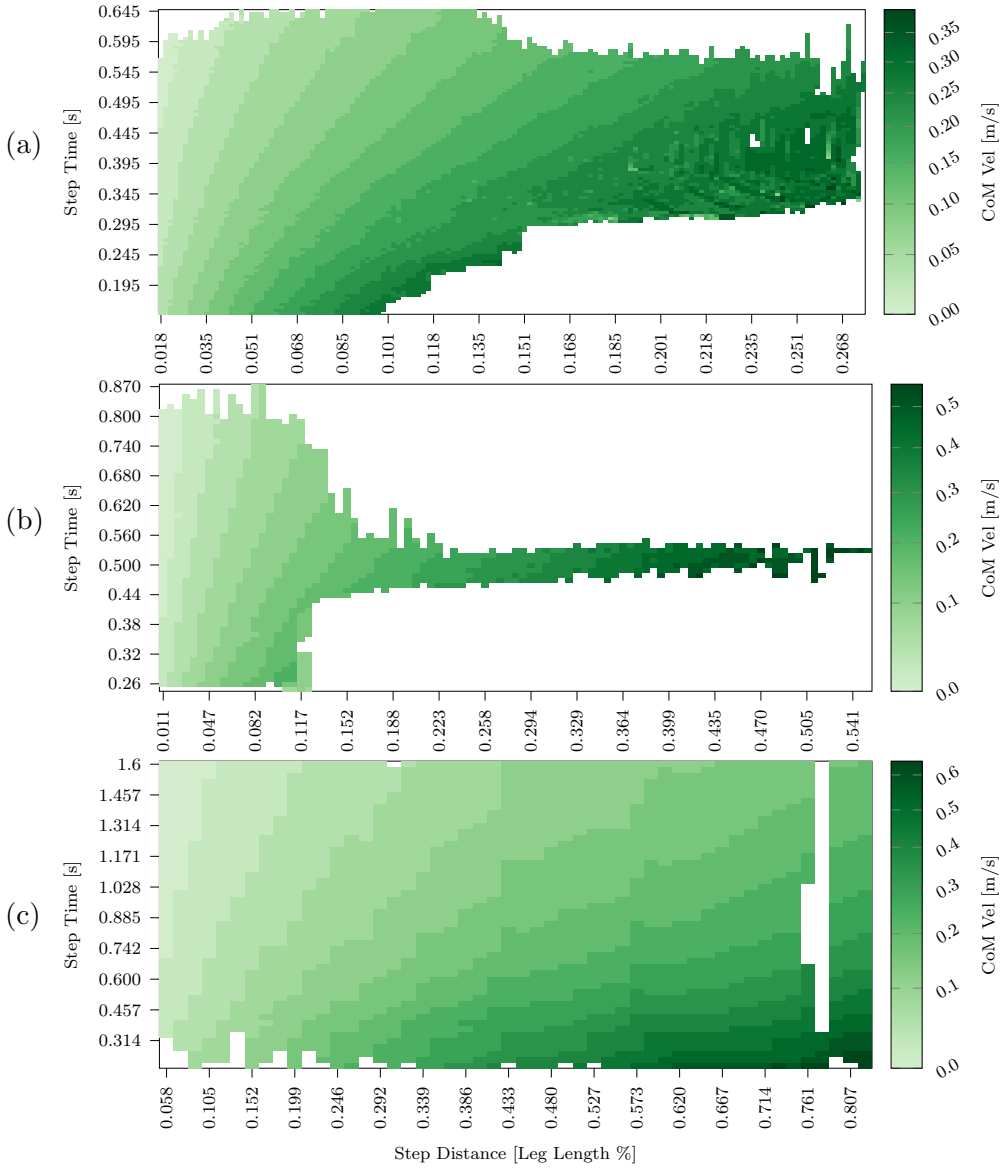

Figure 6.3: Step distance (as a proportion of leg length) vs. step time vs. CoM velocity for Scenarios 1(a), 2(b) & 3(c).

6.3.2.1 Step Distance and Step Time

To analyse the relationship between walking parameters and efficiency, we first look directly at the relationship between parameters and the raw torque consumption across all joints. Figure 6.4 shows how the squared integral torque landscape over the sampled parameters in each scenario, where the more purple the colour, the greater the torque use, and points denote step time with the lowest torque (locally optimal) value for each sampled step distance.

Scenarios 1 and 2 (Figures 6.4a and 6.4b), which use the feedback controller,

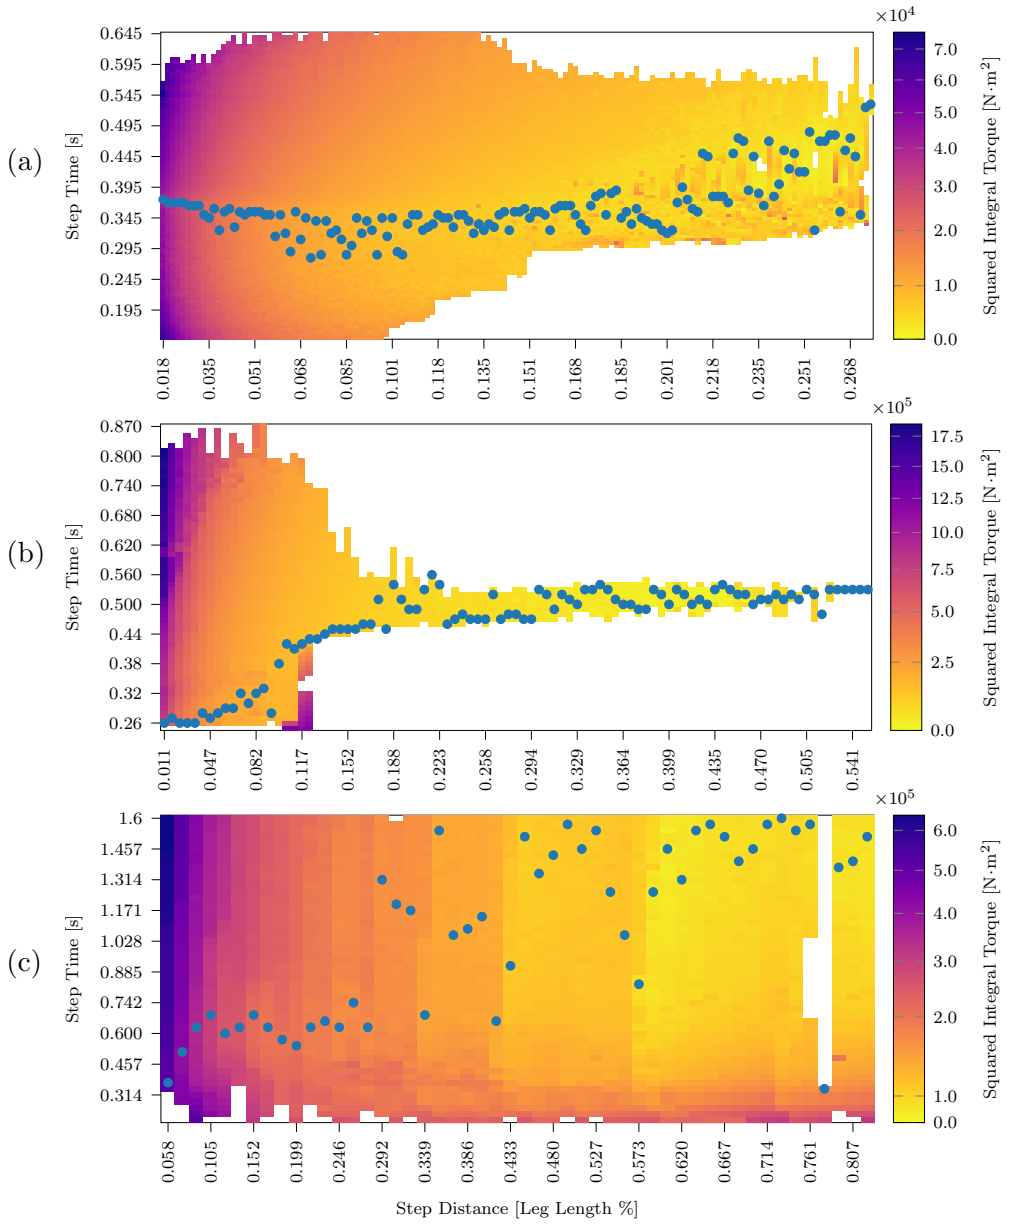

Figure 6.4: Step distance vs. step time vs. Squared Integral Torque for Scenarios 1(a), 2(b) & 3(c). Torque-optimal step times for each step length are marked.

exhibit a simple relationship between step length and torque-optimal swing time. Model-based feedback controller use invariant, continuous models in which the motion produced using a given set of parameters has a differentiable association with the motion produced using marginally different parameters, so perhaps this result is not surprising, but these controllers also form a particular feasible solution space, as such there is no obvious reason why this controller would produce this relatively regular pattern.

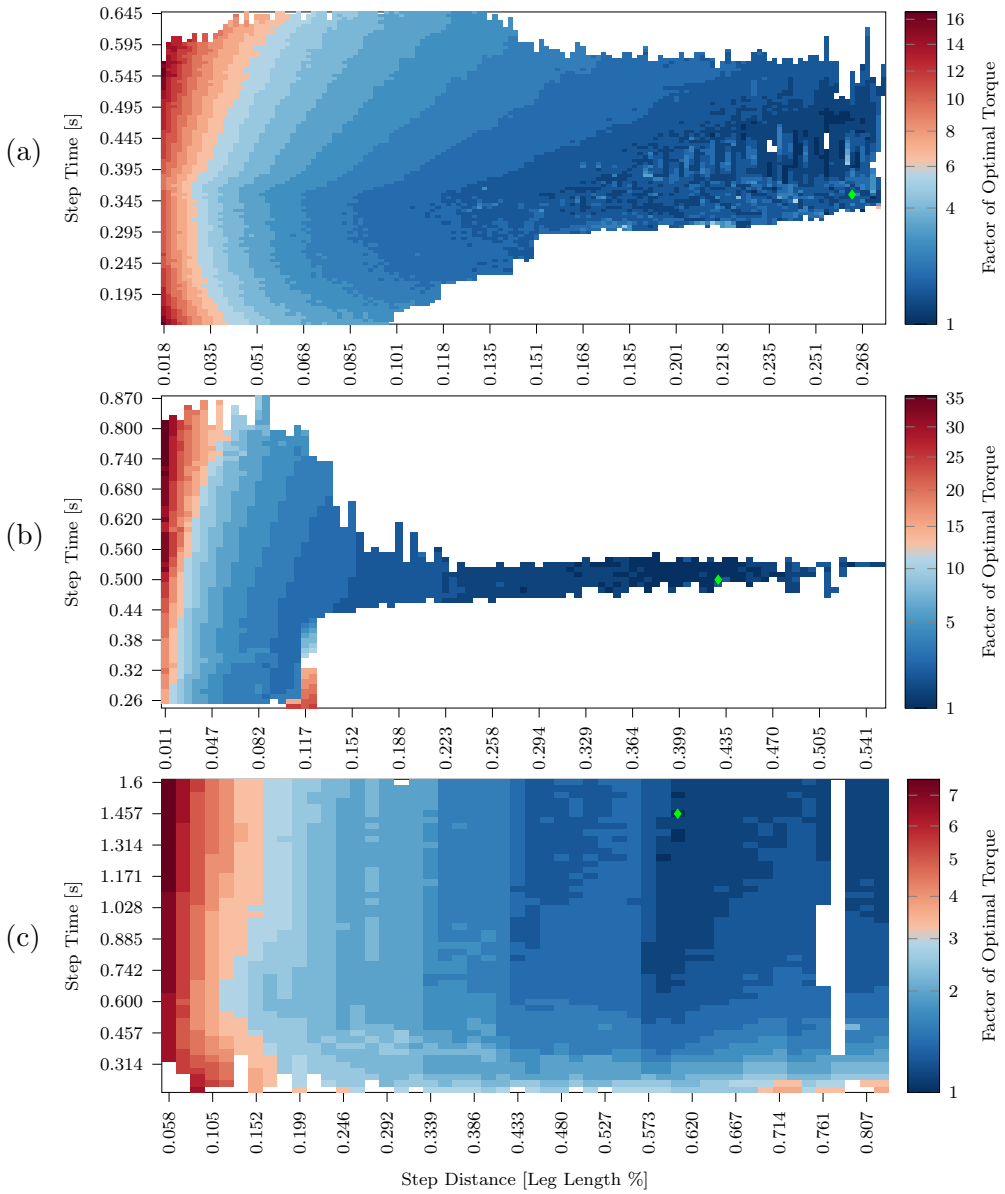

Figure 6.5: Step distance vs. step time vs. difference factor from global torque minimum for Scenarios 1(a), 2(b) & 3(c). The global optimal is marked in green.

For the optimal control case in Figure 6.4c there is a slight upward trend, but there is no discernible relationship between the points due to the one-shot nature of numerical control, where two trials with only slightly different input parameters may converge differently.

Though the local relationship between walking parameters torque consumption seems variable across scenarios, the global relationship appears to be more consistent. In figure 6.5, the globally optimal stepping parameters are denoted with a green marker and the colour gradient represents the factor of difference from this

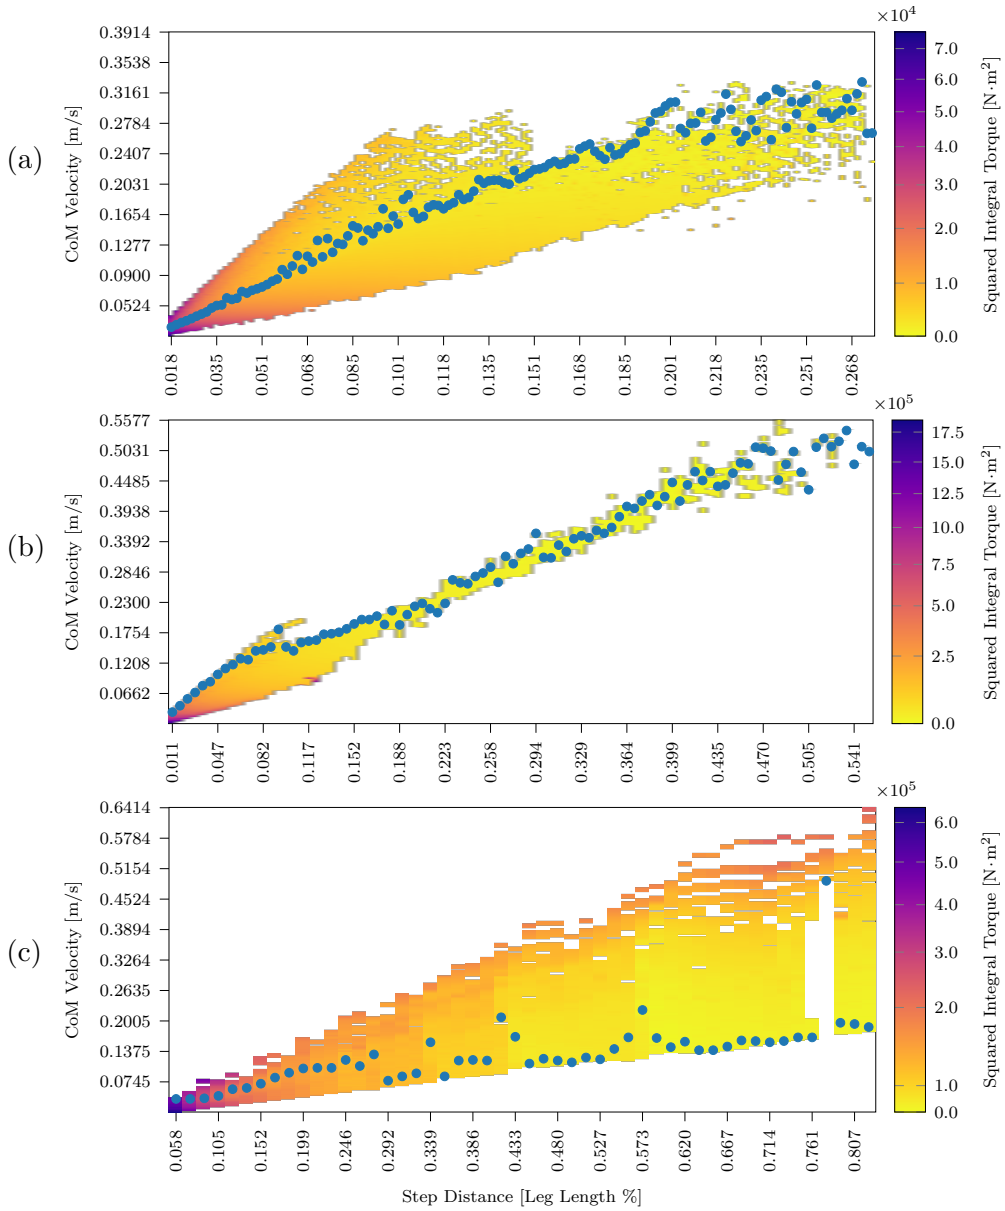

Figure 6.6: Step distance vs. CoM velocity vs. Squared Integral Torque for Scenarios 1(a), 2(b) & 3(c). Optimal walking velocity for each step length is marked.

optima. For all three scenarios, there is a clear radiation from the optimal point, where the optimal point lies in a valley surrounded by a gentle gradient, which increases and becomes very steep towards the least efficient parameters. These regions of attraction are even more pronounced in 3D plots shown in C.4.

In the case of the feedback controllers from scenario 1 & 2, if we follow the gradient of the data towards the optima, it is possible that the limit of the controller is reached before the gradient can continue, indicating that the true optima of the

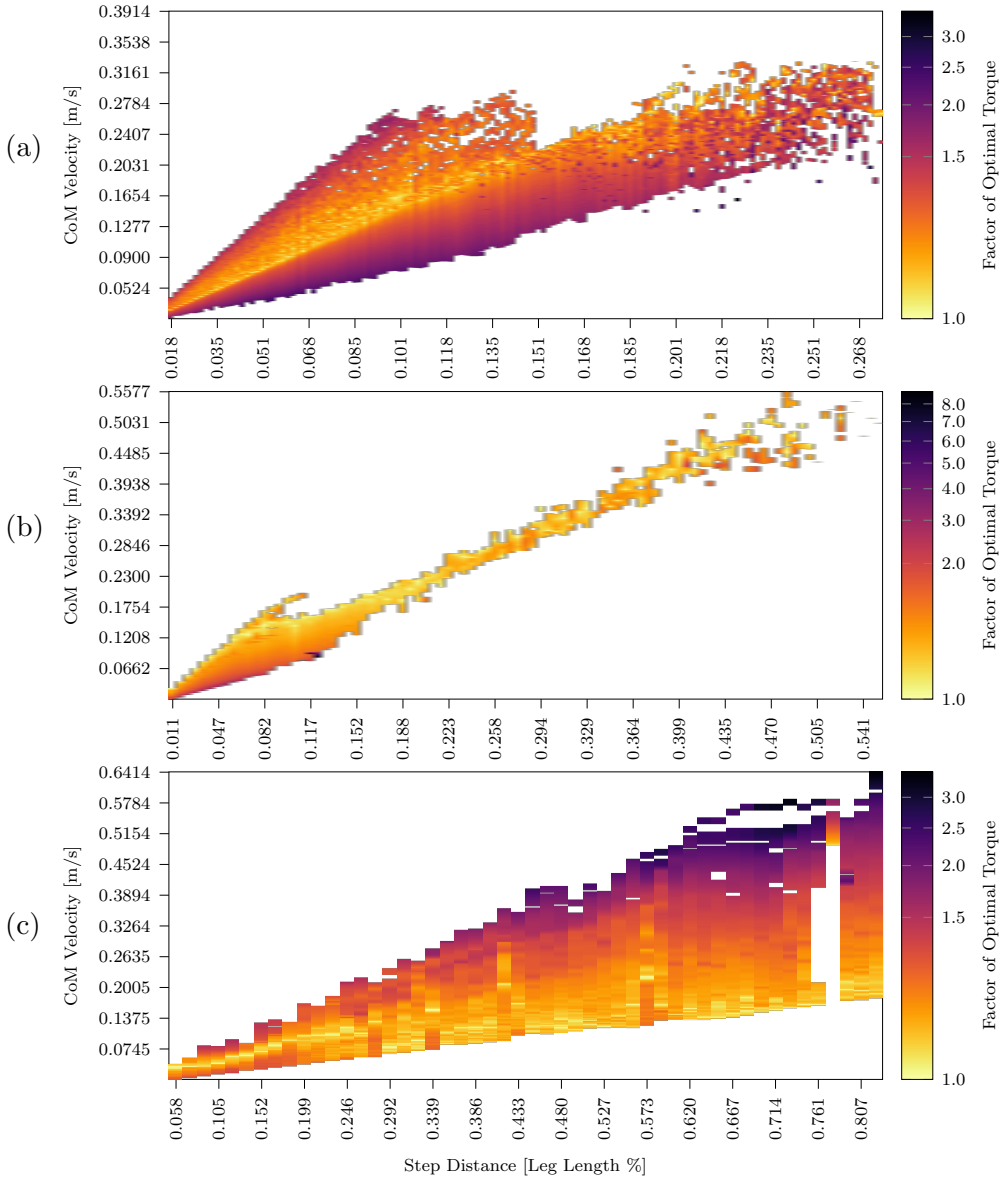

Figure 6.7: Step distance vs. CoM velocity vs. difference from local minimal torque usage for each step distance for Scenarios 1(a), 2(b) & 3(c).

morphology during static walking cannot be reached given the available controller. The optimal control case, scenario 3, (Figure 6.5c) demonstrates this in that the state space is explored further due to the increased feasibility of optimal control in this case, the globally-optimal torque consumption parameters sit in a shallow valley surrounded on all sides, rather than at the limit of the controller.

Torque consumption landscapes seen in the 3D cases (Figure C.4) have a consistently regular, smooth profile across robots and control if we account for the bottleneck in scenario 2. All cases show that short steps are highly inefficient

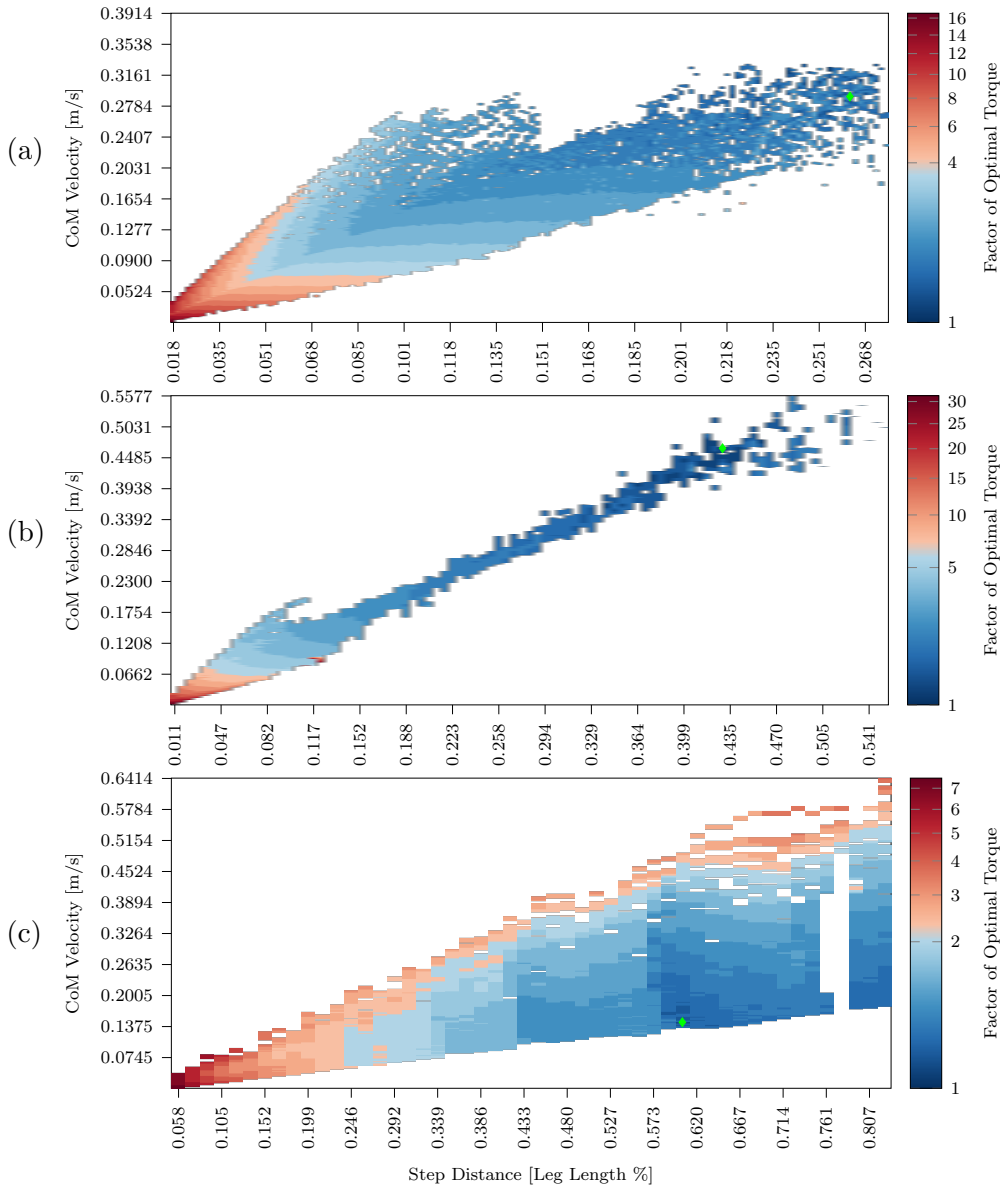

Figure 6.8: Step distance vs. CoM velocity vs. difference factor from global torque minimum for Scenarios 1(a), 2(b) & 3(c). The global optimal is marked in green.

from a torque consumption perspective, and, to a point, the longer the step the better, which highlights the inefficiency of the tendency in locomotion towards slow walking.

6.3.2.2 Step Distance and Walking Speed

Figure 6.6 shows the step length plotted against walking velocity and squared integral torque, with points marking the torque-optimal (locally optimal) walking

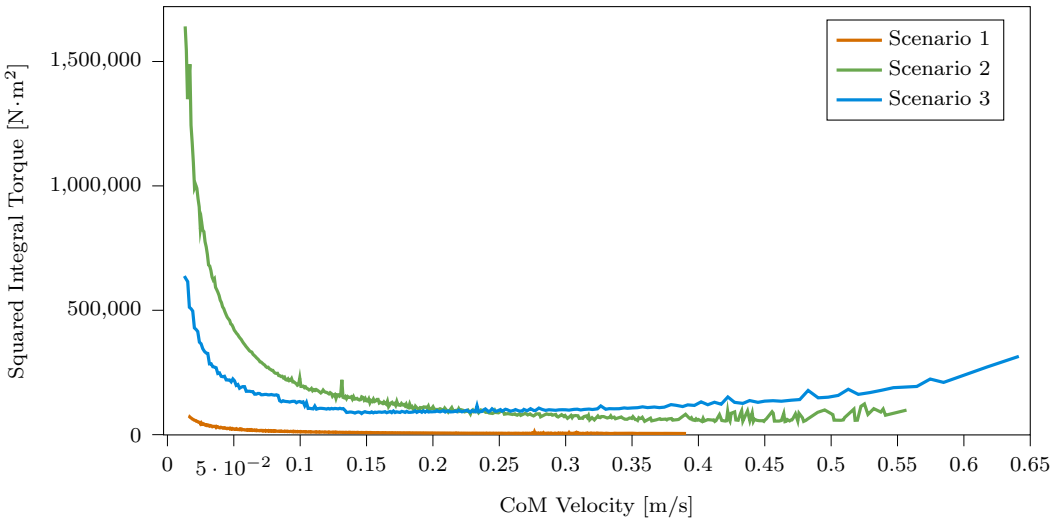

Figure 6.9: Progression of Squared Integral Torque consumption relative to CoM velocity for all scenarios.

velocity for each step length. Note that in scenarios 1 there is an abrupt vertical ledge at around 0.15 % leg length, which is a result of unsuccessful trials for faster step times at the point, as seen below 0.295 m/s in Figure 6.4a.

In each case, the torque-optimal relationship between step length and walking velocity is linear, which is unexpected for the optimal control case, given the discontinuous optimality in Figure 6.5c.

In Figure 6.7, we show the relative difference in energy from the locally optimal walking velocity for each step length to show the change in torque consumption between the optimal and its neighbouring velocities. Darker colors represent greater differences from the optimal, and lighter colors represent smaller differences for that step length.

The trail of optimality in scenario 1 is very distinct and shows that, for a given step length, if the optimal walking velocity is not achievable, it is better to walk faster than slower. A similar pattern is present but much less distinct in scenario 2 (Figure 6.7b). In the scenario 3, this band is clear and runs through the middle for short step lengths, but becomes wider and moves towards lower walking velocities as step length increases, with frequent, noisy exceptions.

Global-optima difference plots in 6.8 show a similar pattern to that in Figure 6.5, where the optimum sits in a shallow valley of near-efficient conditions. Again, in this case we see that in the feedback case, the optima appears in the noisy regions

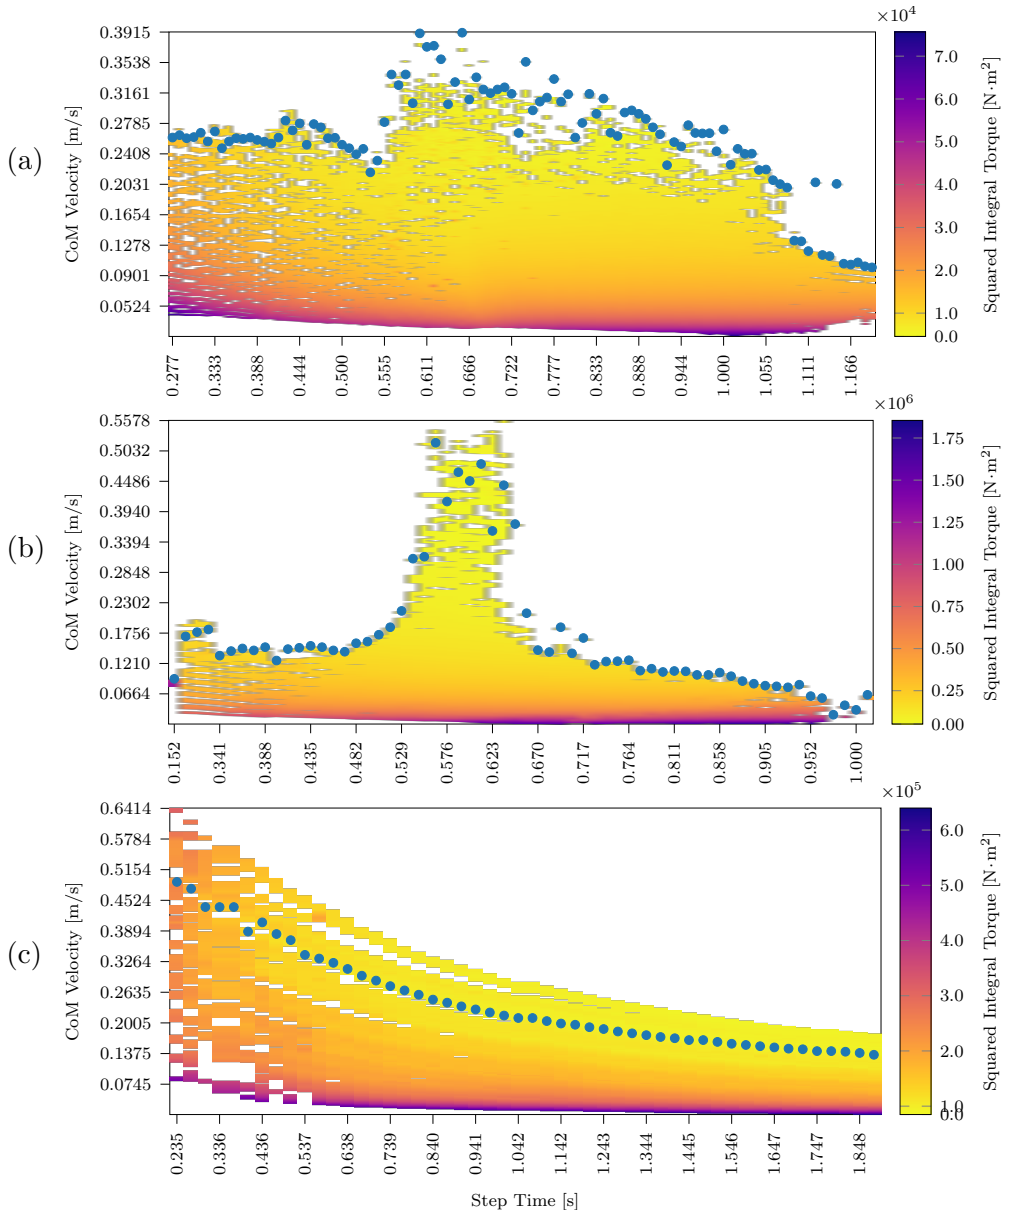

Figure 6.10: Step time vs. CoM velocity vs. Squared Integral Torque for Scenarios 1(a), 2(b) & 3(c). Torque optimal walking velocity for each step time is marked.

at the edge of the controller's operational limits, but in scenario 3, the optima is well within the feasible region.

Walking velocity is plotted in 2D against torque consumption in Figure 6.9. All cases exhibit very high consumption values at low walking velocities and quickly drop steeply. Interestingly, there is a notable intercept between scenarios 2 and 3 at 0.23 m/s, after which point, the optimal control scenario becomes more torque-intensive, in absolute terms, than the feedback control case at higher velocities

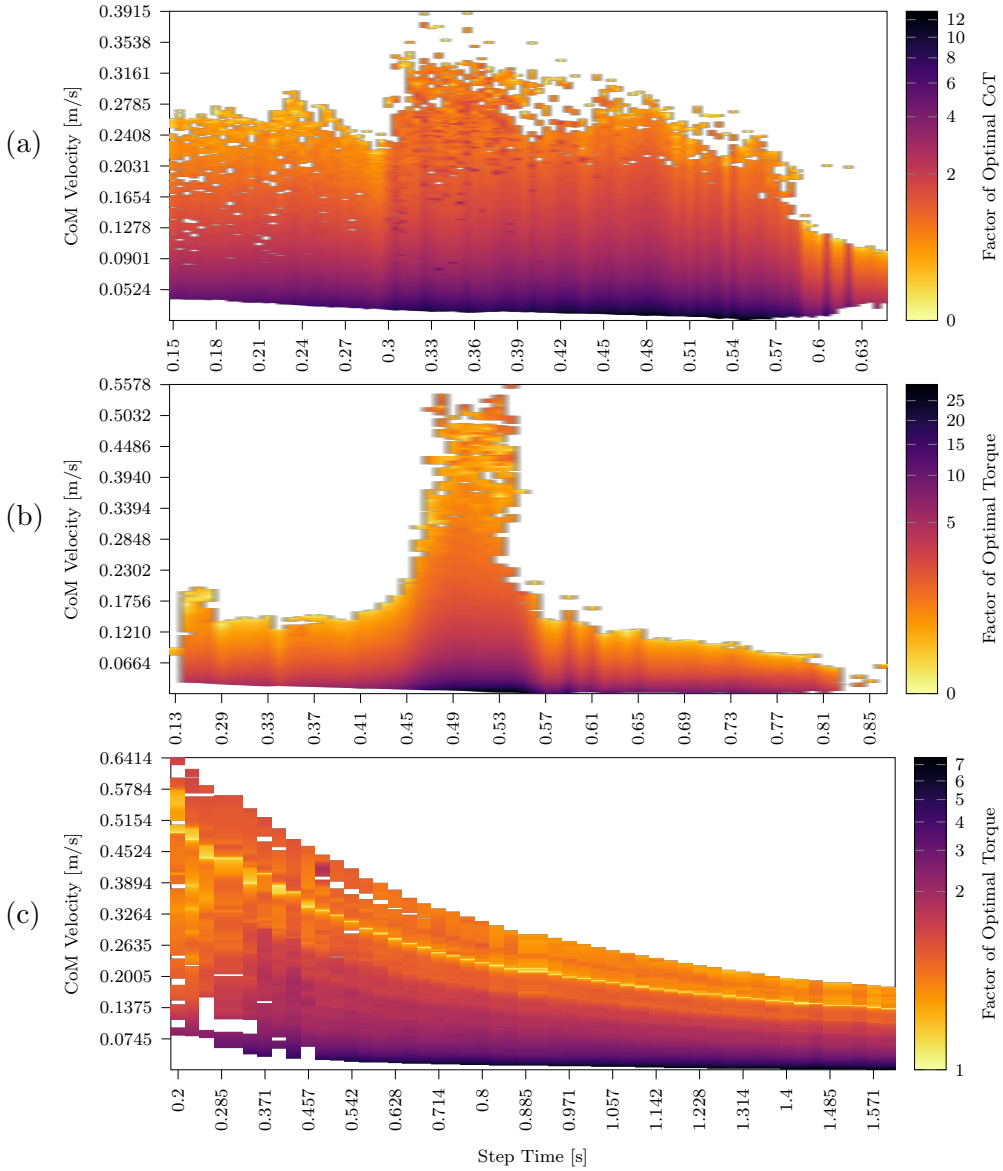

Figure 6.11: Step time vs. CoM velocity vs. difference from local minimal torque usage for each step distance for Scenarios 1(a), 2(b) & 3(c).

using the same robot, which would not be expected given these methods of control, and it is unclear why this is the case.

6.3.2.3 Step Time and Walking Speed

Step time appears to have a more consistent trend of local torque optimality. Scenarios 1 & 2, in Figures 6.10a and 6.10b show that for any given step time parameter, the torque-optimal CoM is as high a possible, which would require as long a step as possible to maintain a high walking speed, but again it appears

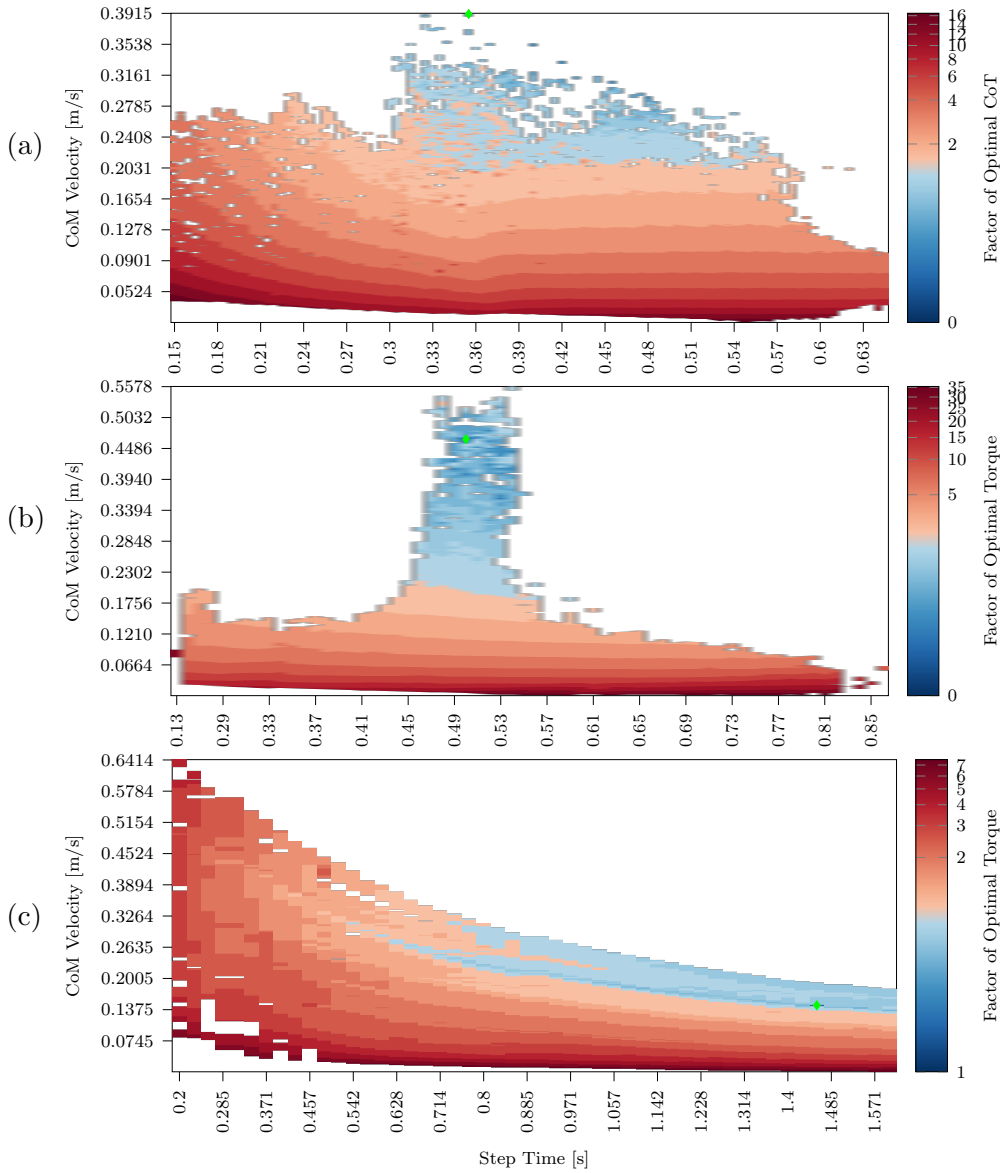

Figure 6.12: Step time vs. CoM velocity vs. difference factor from global torque minimum for Scenarios 1(a), 2(b) & 3(c). The global optimal is marked in green.

that this effect is stifled by the limitations of the controller, as in both cases the optimal walking velocity is at the edge of the set of successful trials. This is mirrored in the local optimality maps in Figures 6.11a and 6.11b, where we see a steep gradient from the locally optimal walking velocity at the top of the map for every sampled step time.

Scenario 3 does not display this issue, and the locally optimal region is much clearer. Figure 6.10c shows a distinct and consistent relationship between step time and walking velocity, with a tendency towards higher walking velocities, but,

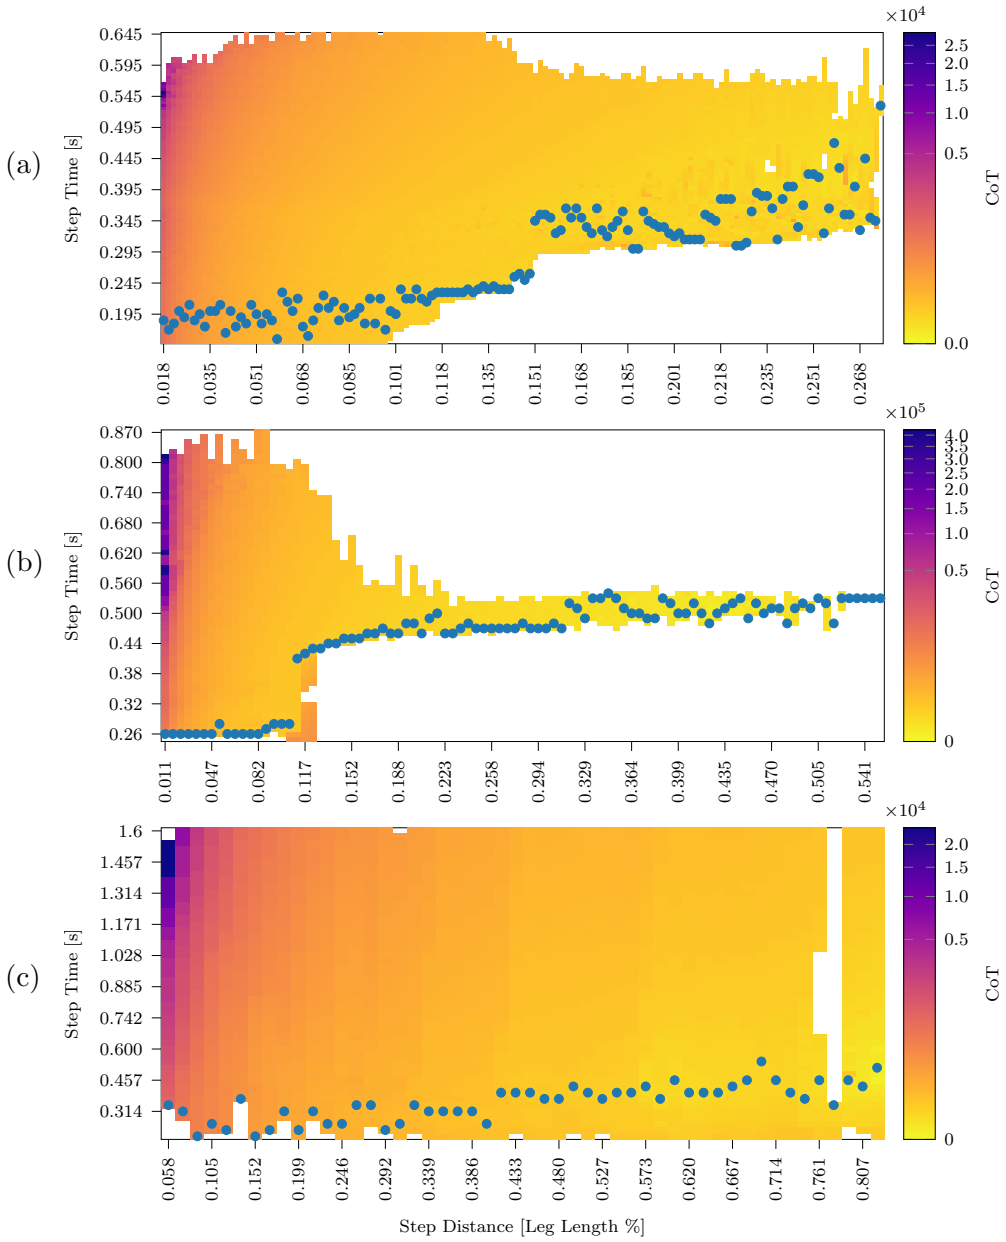

Figure 6.13: Step distance vs. step time vs. Cost of Transport for Scenarios 1(a), 2(b) & 3(c). Optimal step times for each step length are marked.

as shown in Figure 6.11c, this region sits well within the region of the successful trials and shows a clean band of near-optimality. This gives the most pronounced demonstration of the benefit of optimal control for exhaustive exploration of this landscape.

The globally optimal case shows inconsistency between the control cases. Figures 6.12a and 6.12b show that global optimality sits in the noisy regions at the edge of controller capability, with high walking velocity and mid-range step timing,

whereas scenario 3 (Figure 6.12c), torque-optimality unexpectedly appears at longer step times and slower speeds. For the torque-optimal case, this could be due to a higher torque-minimisation term in the optimal control case, but, as shown later in the chapter, the CoT case shows more agreement between the scenarios.

6.3.3 Cost of Transport

6.3.3.1 Step Distance and Step Time

Where analysing the absolute torque usage during walking gives us a view of the raw view of the efficiency during walking, Cost of Transport is more suitable for comparing the energy cost of walking while accounting for the differences between the robots. Since CoT accounts for the differences in mass and for walking speed, it provides a more consistent method of comparison that was not applicable in the previous chapter.

Analysing CoT immediately shows more consistent results across all scenarios than in the torque consumption maps above. Figure 6.13 shows the relationship between step distance, step time and CoT over the 1 m sampled distance. The optimal step time for each step length show a remarkably consistent pattern across all scenarios, where for each case, the CoT-optimality is associated with low step times There is a gentle upward trend as the step length increases in scenario 3 (Figure 6.13c) and towards the noisy regions in scenario 1 (Figure 6.13a).

Local optimality also seems to hug the lower step times as far as is feasible in scenario 1 & 2, (Figures 6.13a & 6.13b), suggesting again that the locally optimal parameters are outside the feasible region of the controller. This also seems to be true of smaller step lengths in scenario 3 (Figure 6.13c), where local optimality lies at the lowest possible step times; unfortunately, the controller in this case became unstable at lower step times, so was not possible to test further.

With respect to the global optima, we again see a shallow region of attraction around the CoT-optimal point in Figure 6.14, where in the feedback case, the optima sits at the limit of the controller capability, but in the optimal control case, the CoT-optimal point is located at a much higher step length and lower step time than the torque-optimal case.

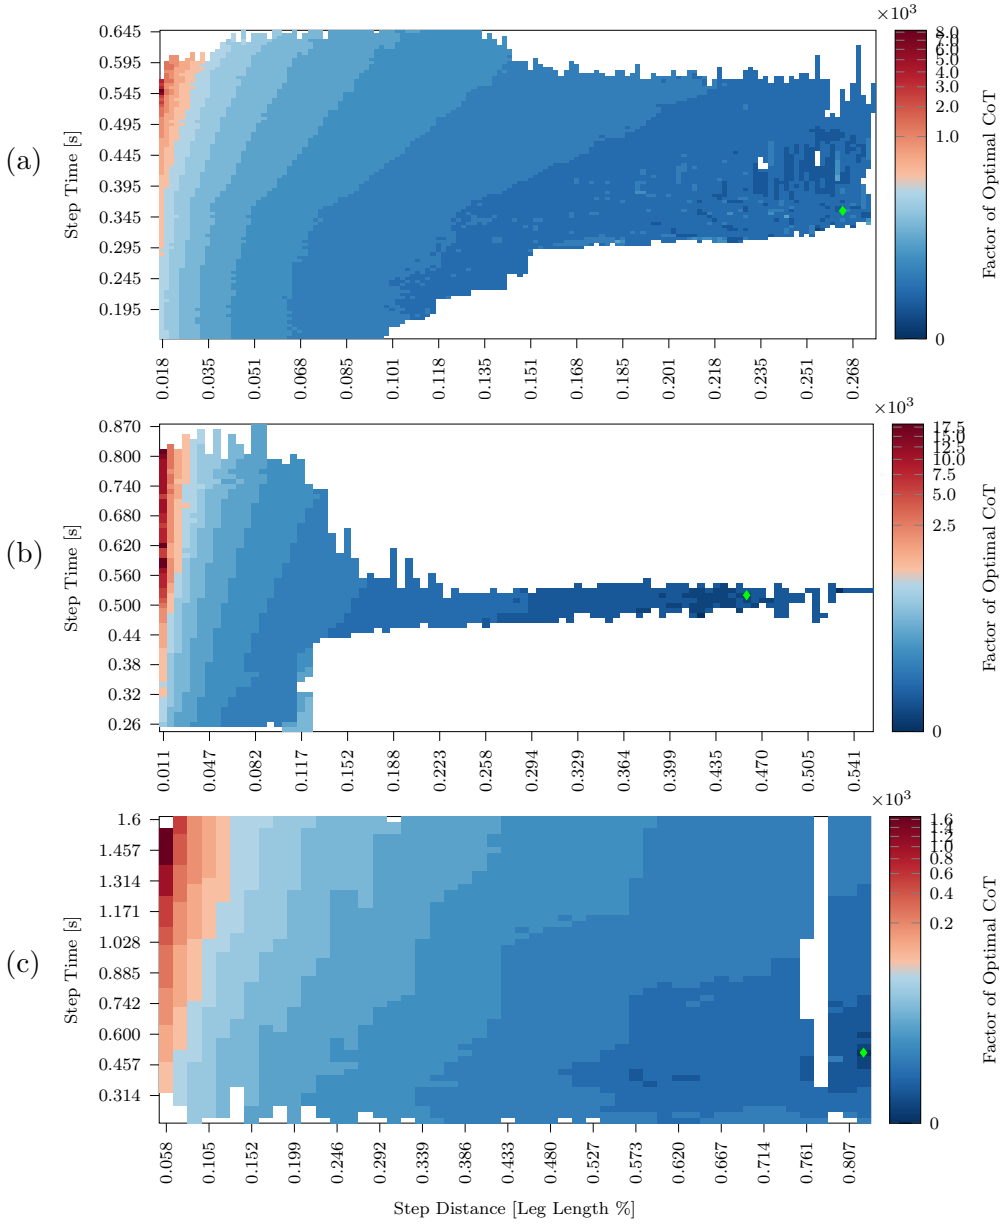

Figure 6.14: Step distance vs. Step Time vs. difference factor from global CoT minimum for Scenarios 1(a), 2(b) & 3(c). The global optimal is marked in green.

We again see this smooth, regular gradient in the 3D representation (Figure C.3). From these plots, we see that the CoT is a much more reliable method of comparison, showing that regardless of the control method or robot, the general trend in the energy efficiency landscape is very consistent.

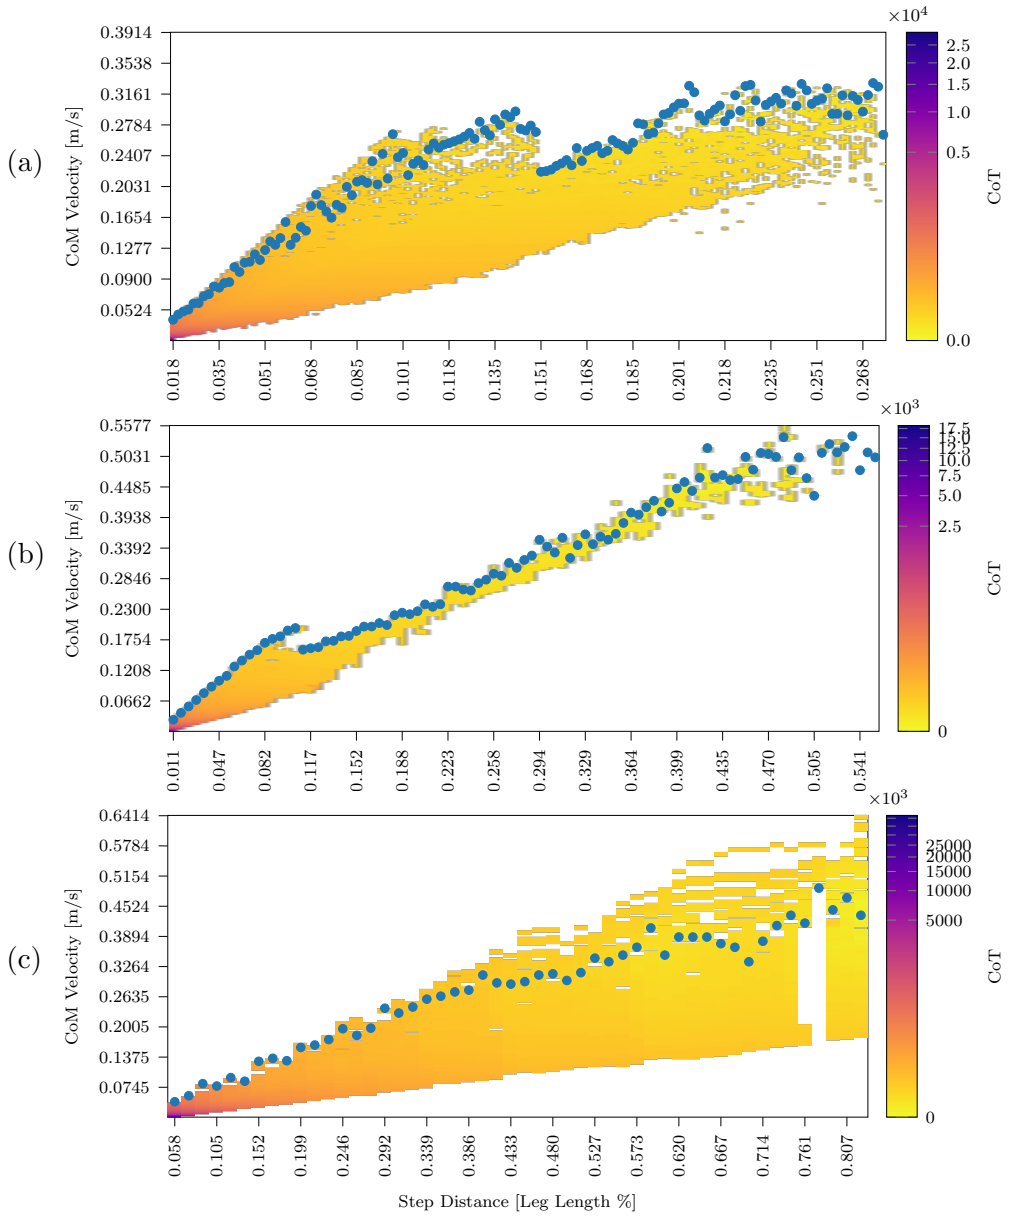

Figure 6.15: Step distance vs. CoM velocity vs. Cost of Transport for Scenarios 1(a), 2(b) & 3(c). The locally optimal CoM velocity for each step length is marked.

6.3.3.2 Step Distance and Walking Velocity

The CoT profile of the parameters are shown to be more consistent when viewed from a CoT perspective rather than from a raw torque usage perspective. Figure 6.15 shows the CoT value for each pair of parameters, with CoT-optimal walking velocities marked with blue points. In contrast to the torque-centric view, which revealed a linear relationship in each case, but different gradients in each robot,

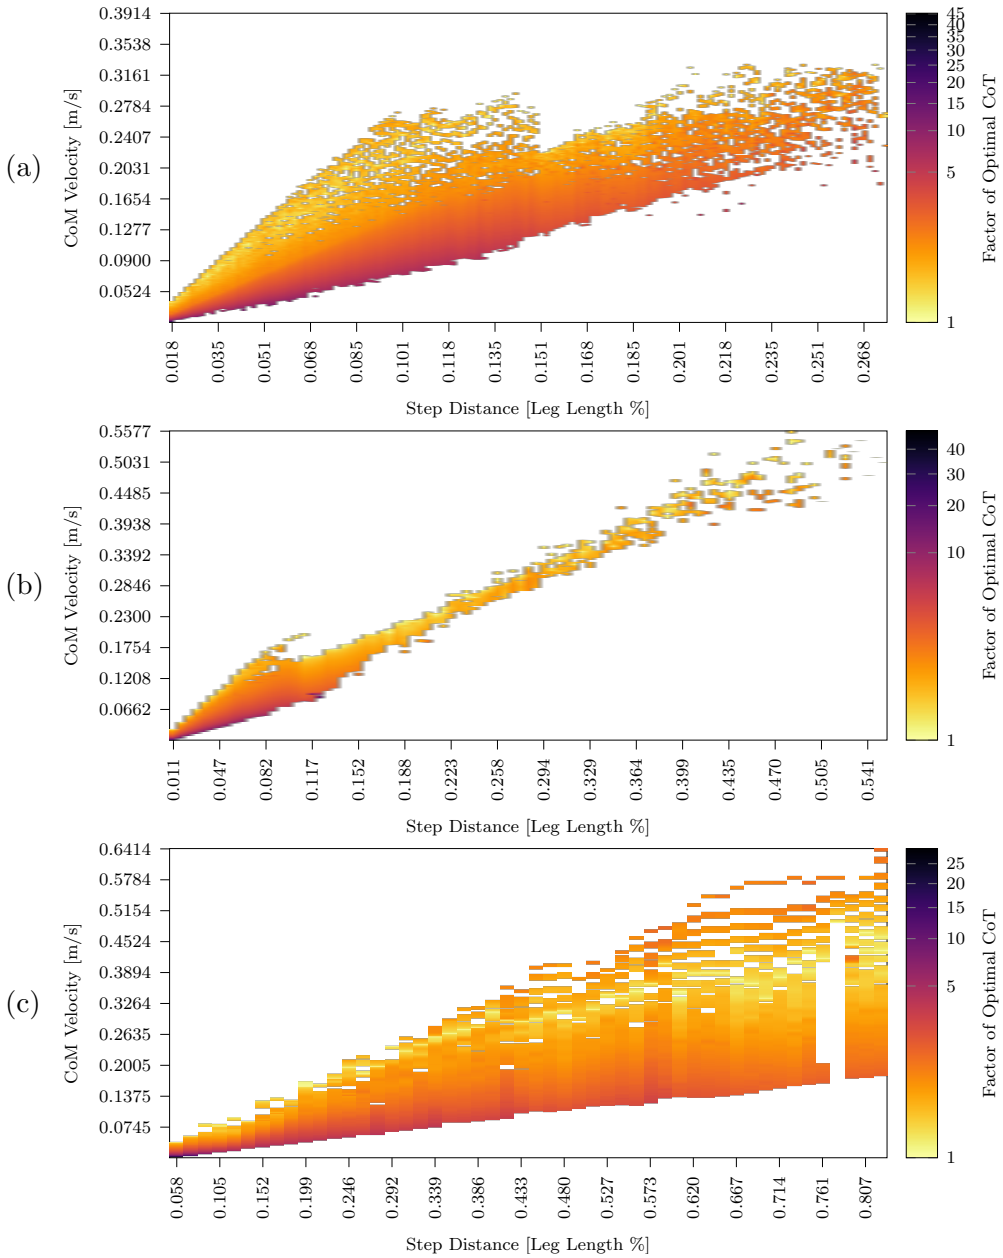

Figure 6.16: Step distance vs. CoM velocity vs. CoT difference from local minima for Scenarios 1(a), 2(b) & 3(c). Darker colours represent a higher differences in CoT from the locally CoT-optimal CoM velocity.

Figure 6.15 shows linear relationships in all cases, but, crucially, trending towards higher walking velocities in all cases, as opposed to the variance seen in the torque consumption landscape in Figure 6.7.

Again, we see that surrounding these locally optimal points is a band of near-optimality in Figure 6.16. In scenarios 1 & 2 (Figures 6.16a,6.16b), the trend is

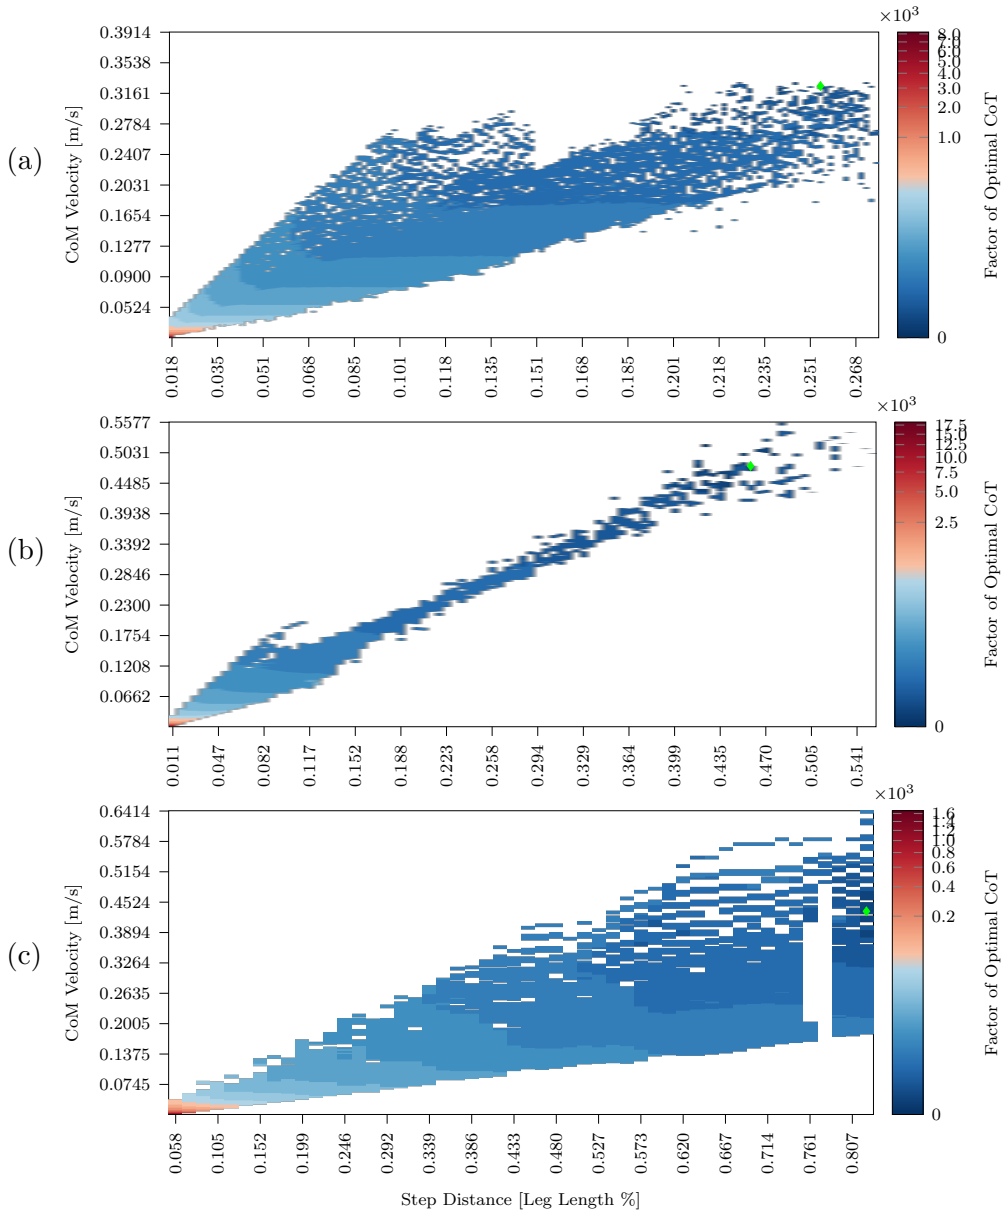

Figure 6.17: Step distance vs. CoM velocity vs. difference factor from global CoT minimum for Scenarios 1(a), 2(b) & 3(c). The global optimal is marked in green.

similar to that in the torque-based maps in Figures 6.7a and 6.7b, hugging the limits of the controller's feasibility. But, crucially, the region of locally optimal walking speeds in scenario 3 (Figure 6.16c) matches this pattern from the other scenarios, showing that from a CoT point of view, there is consistency between the different conditions.

Scenario 3, shown in Figure 6.15c is notable in that until around 0.33 % of leg length, the CoT-optimal walking velocity is at the limit of the achievable walking

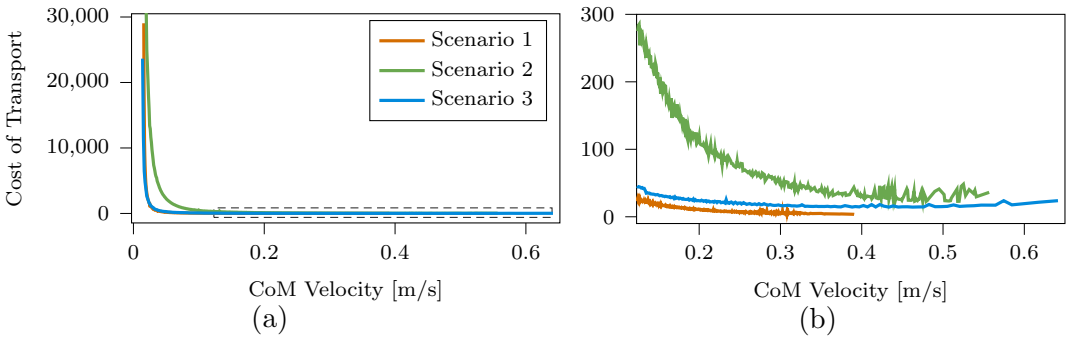

Figure 6.18: Cost of Transport values for each CoM walking velocity both (a) overall and (b) zoomed to the dashed area in (a).

velocities, the trend continues after this point, but does not reach the limit of the achieved walking speeds. This again raises the possibility that in the feedback control cases in scenarios 1 & 2, optimality may be limited by the capabilities of the controller.

Figure 6.17 shows that this may also be true of the global optima. In all three scenarios, we see that the global optima is at the limit of the reliably achievable step length. What is notable from this point is that as globally optimal efficiency tends towards longer step lengths corresponds to the known limitations of existing humanoid robots, in that toe-off motions are a barrier to higher efficiency. This appears to be an obvious point, but we see this across three different scenarios of different robots and controllers, suggesting that toe-off is not just a desirable component to increase the capabilities of the machines, but a lack of this capability is a drastic barrier to efficiency.

Walking velocity/CoT curves in humans show a characteristic U-shape, where the CoT-optimal and preferred walking speed is at the lowest point of this curve. When we plot this curve for our scenarios in Figure 6.18, we see that the very slowest of the trials have astronomically high CoT values (Figure 6.18a), but when we zoom on the lower values in Figure 6.18b, we see the curve is converging asymptotically to form the valley of the U-shape. Given the available robot platforms, it is unclear how the curve would form for higher walking velocities, but given the data from humans, it is a possibility that the curve would reach a global optima then begin to rise again. This again may suggest that not being able to toe-off during locomotion is a physical barrier to increased efficiency, though this could also be a result of the controller, or an alternative issue with the robot

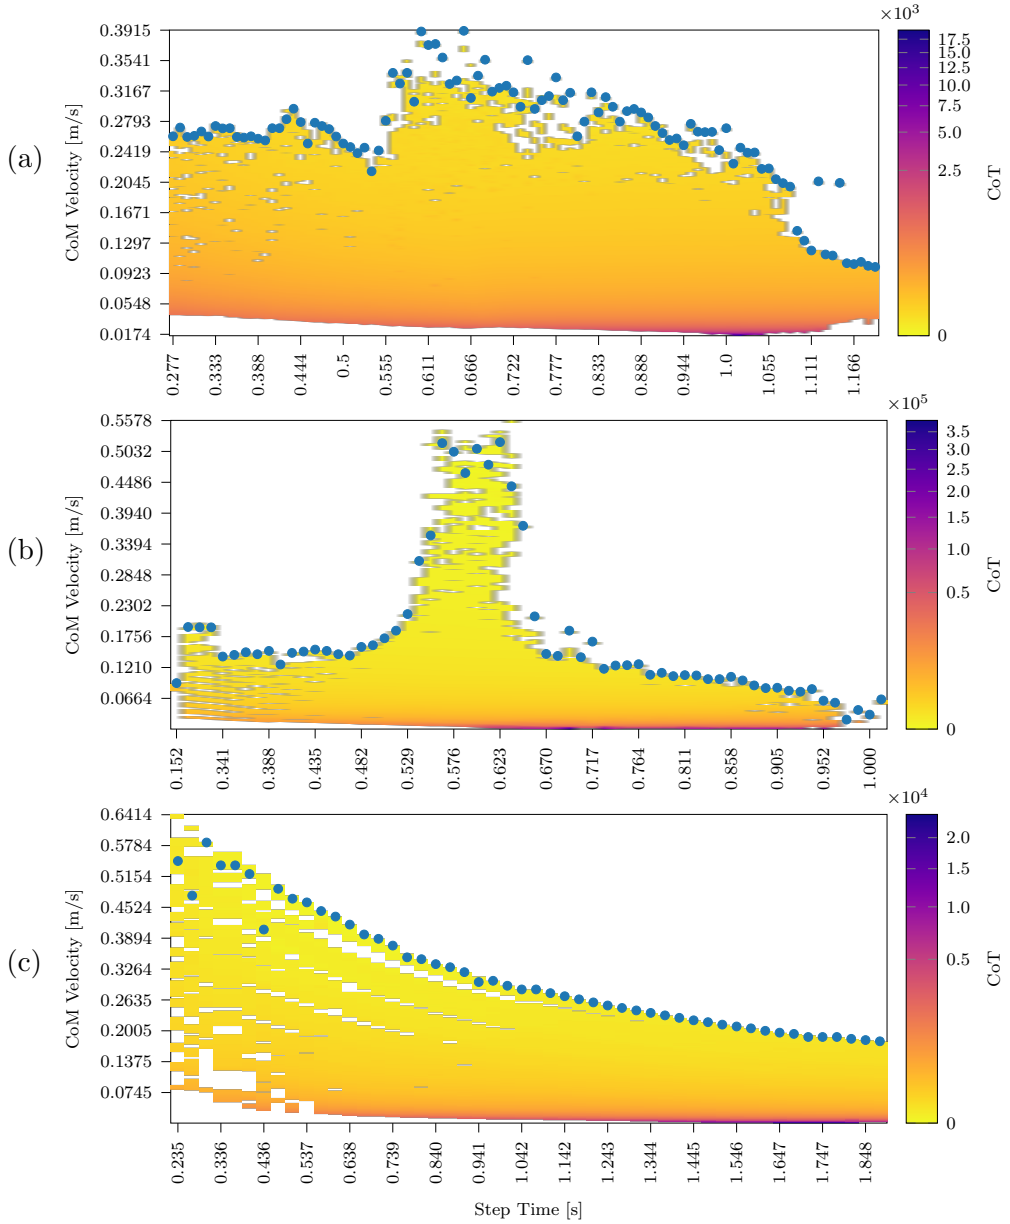

Figure 6.19: Step time vs. CoM velocity vs. Cost of Transport for Scenarios 1(a), 2(b) & 3(c). Optimal walking velocity for each step time is marked.

construction.

6.3.3.3 Step Time and Walking Velocity

During analysis to this point, the characteristics of the torque-optimality landscape for the optimal control case (scenario 3) tend to diverge from those of the feedback controller cases, but when viewed from a CoT point of view, they display a more similar picture; this is also true of the CoT-wise relationship between

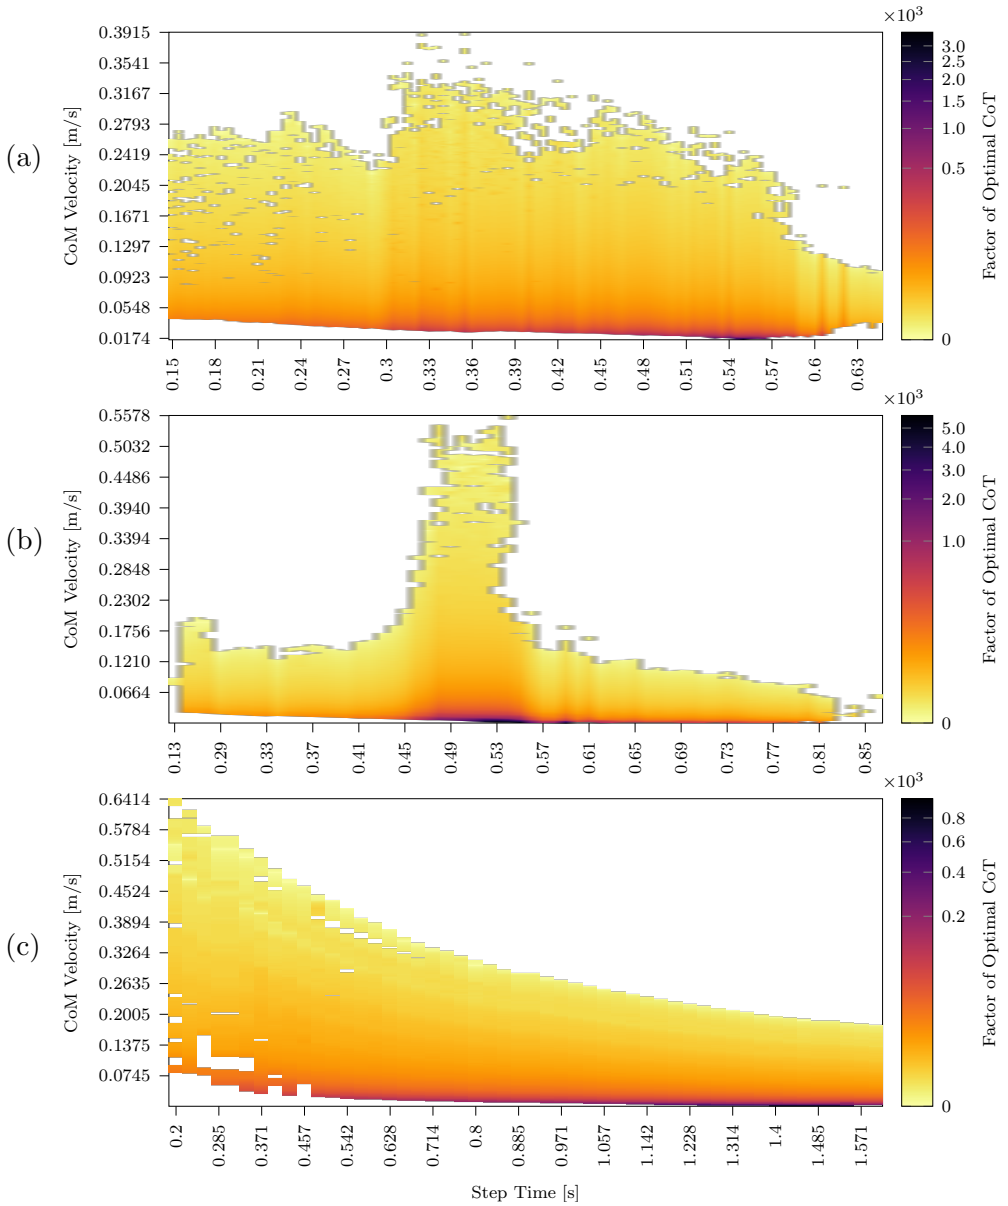

Figure 6.20: Step time vs. CoM velocity vs. difference from local minimal CoT for each step distance for Scenarios 1(a), 2(b) & 3(c).

step time and walking velocity.

All scenarios exhibit a direct invariant, smooth gradient towards higher walking velocity for all sampled step times in both their absolute CoT values (Figure 6.19) and in the local optimality gradient (Figure 6.20).

In the global optimality plots, the three scenarios again converge towards a similar pattern, in that a basin of attraction appears at higher walking velocities and mid-range step times. Notably, the optimal parameters for scenarios 2 & 3, which both

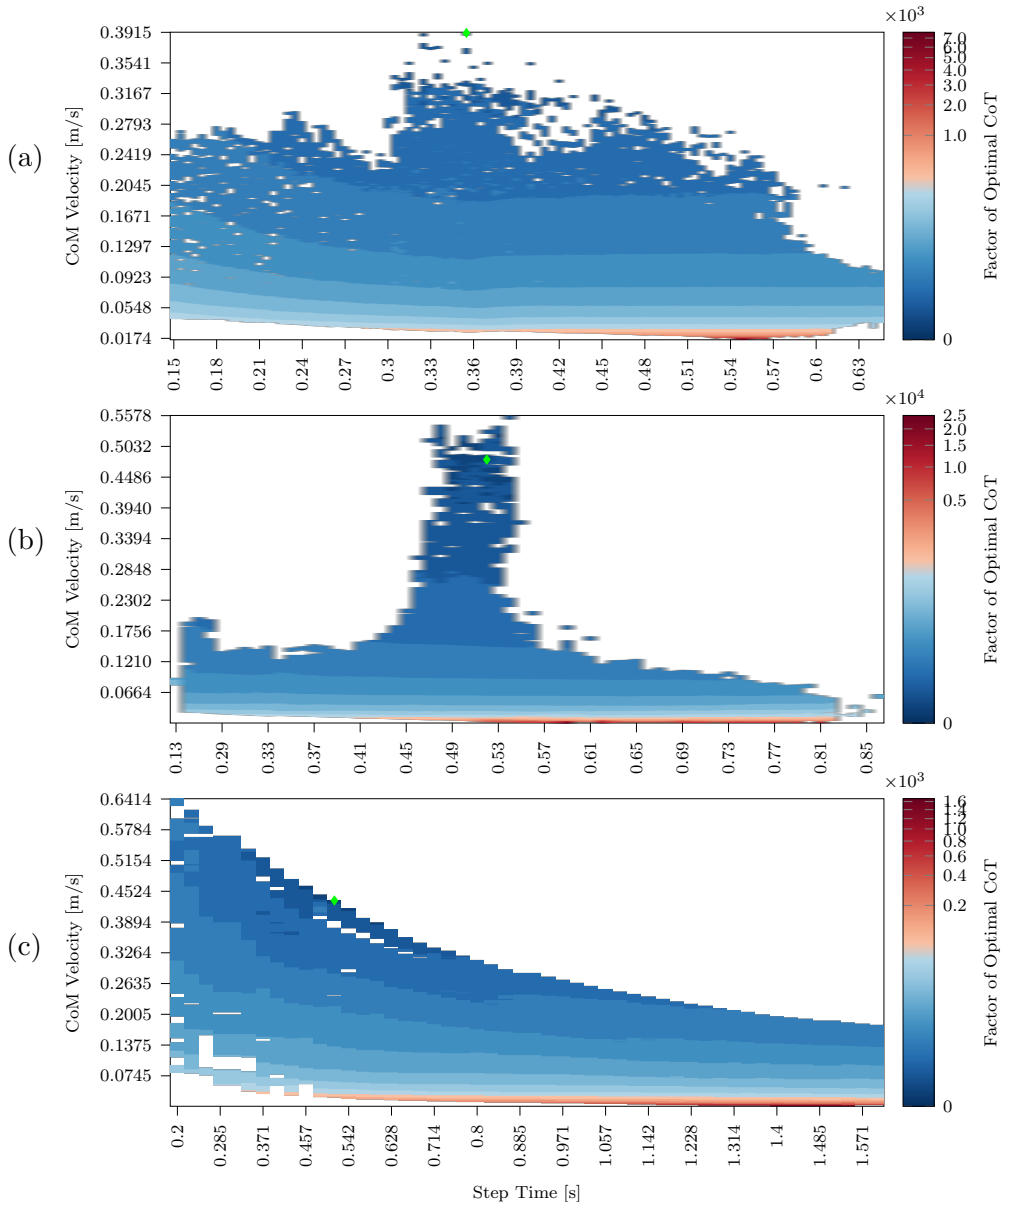

Figure 6.21: Step time vs. CoM velocity vs. difference factor from global CoT minimum for Scenarios 1(a), 2(b) & 3(c). The global optimal is marked in green.

use the same robot, but different controllers are very similar. The optimal walking velocity is 0.48 m/s in scenario 2, and 0.436 m/s in scenario 3 and the optimal step time of both is 0.51 s, which could be explained as a dynamical property of the robot.

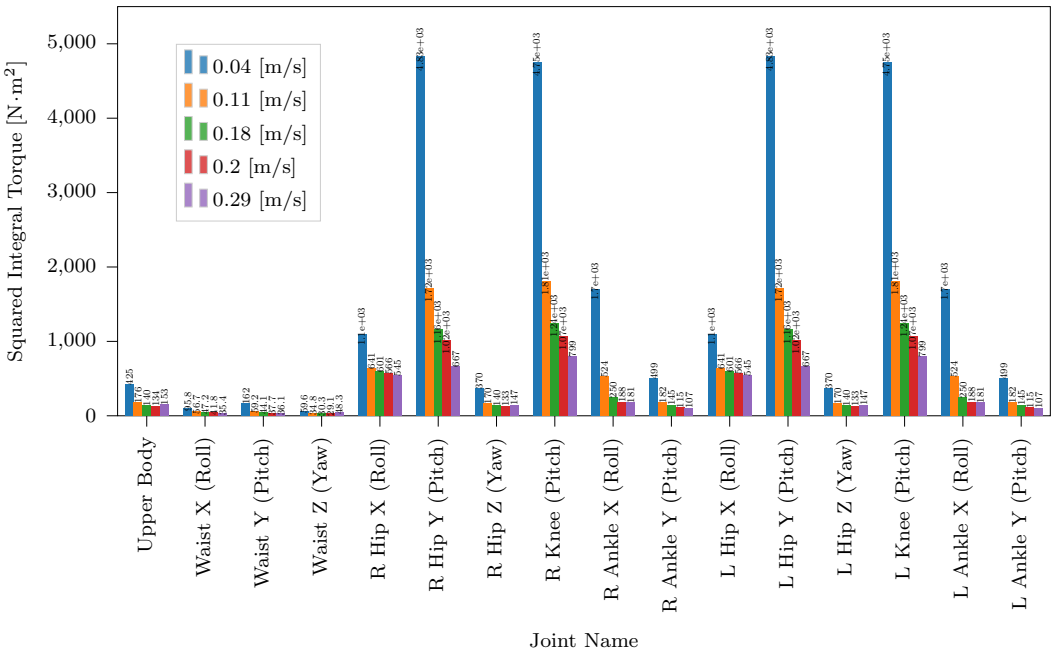

Figure 6.22: Absolute joint torque consumption at each joint during walking at different speeds in the CoMAN robot in scenario 1. Joints above the waist are represented as a single Upper Body value.

6.4 Joint Level Energy Consumption

Figure 6.22 shows the joint-level torque consumption the CoMAN robot in scenario 1 at five increasing walking velocities. The upper body joints torque is negligible and so are combined to give a single value for visibility. From this plot, we see that the absolute joint torque use drops considerably as walking velocity increases, and there is a marked difference in the joints that are the most costly from a torque point of view. From this we see that the joints responsible for sagittal motion are the biggest consumers of joint torque, but this decreases as the walking velocity increases. The second biggest torque consumer are the joints associated with lateral movement, which are heavily used for CoM projection shifting from one contact point to the other.

Figure 6.23 shows the same trials, but displays the torque consumption of each joint as a percentage of the overall torque consumption over the sampled distance. We see that, as walking velocity increases, a greater proportion of torque at joints responsible for mass-shifting side-to-side rather than sagittal forward motion. A redistribution of torque in this way suggests that as walking velocity increases, the preserved body velocity assists the forward motion, lowering the need for

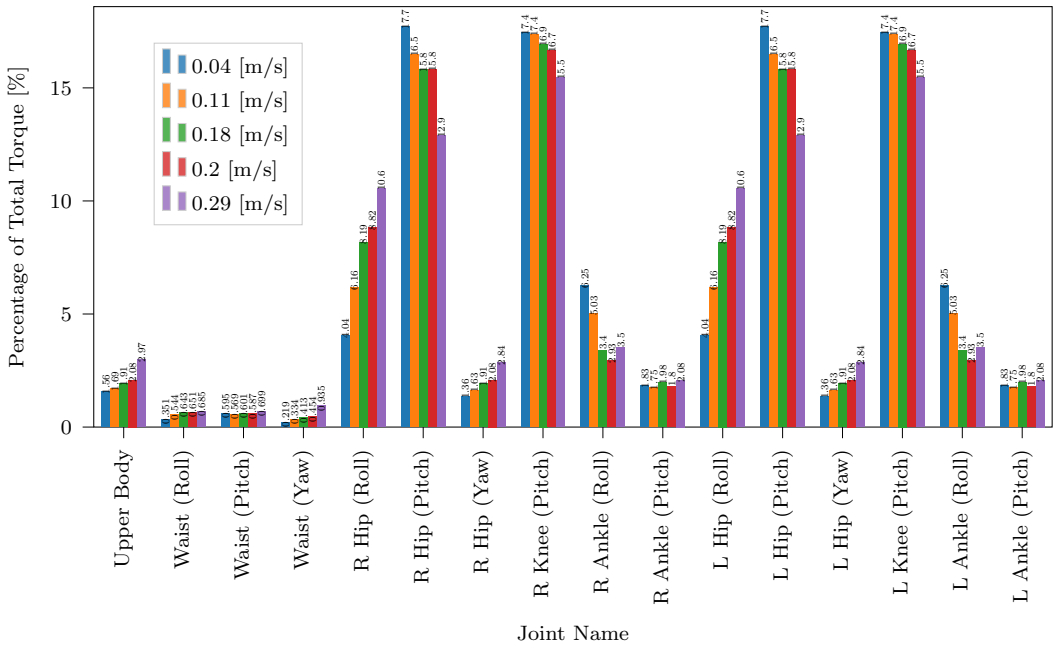

Figure 6.23: Percentage of overall joint torque consumption at each joint during walking at different speeds in the CoMAN robot in scenario 1. Joints above the waist are represented as a single Upper Body value.

torque at the pitch joints. Instead, actuation is needed at joints responsible for shifting the mass projection between the feet to keep the robot stable during for the next step.

As a direction for further research, it would be interesting to see effect of faster, dynamic walking on the torque consumption distribution. To what extent can the torque consumption of the pitch joints be minimised as a result of dynamic energy conservation, and is less requirement for mass-shifting during shorter slower single-support phases, or does this increase due to higher impulse required?

6.5 Discussion

Studying the principles and characteristics of energy-efficient motion is a complex task, which requires global exploration of the walking search space for a given walking type. The highly complex, nonlinear dynamics of the whole body system has the potential to map to similarly nonlinear energy costs. As we saw in the single step case in Chapter 5, this complexity did not emerge, instead we saw a surprisingly simple mapping to energy-optimal step parameters. In this

chapter, we tested parameters for continuous walking and how they were related to efficiency.

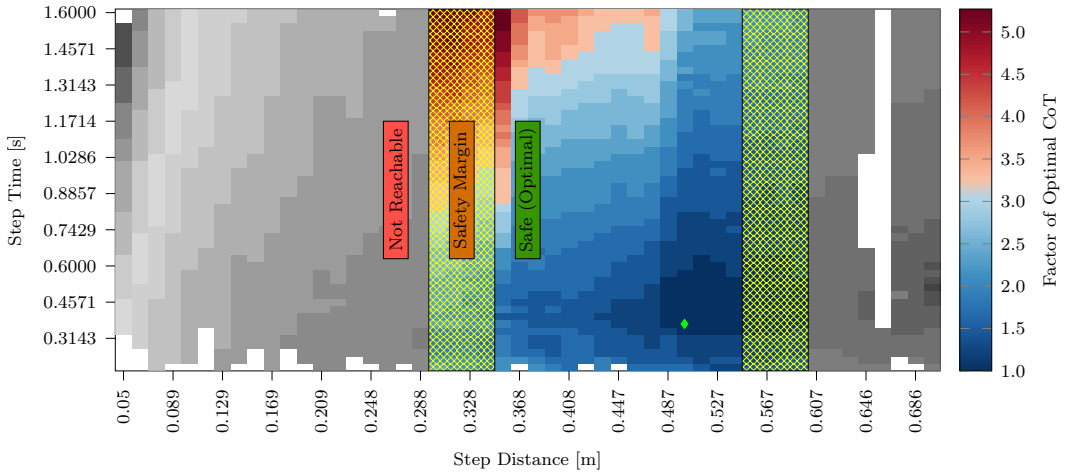

Figure 6.24: Energy-landscape maps can be used as inverse look-up tables to select energy-optimal walking parameters even under spatial constraints and safety margins and select the constrained optimal, marked in green.

Cost of Transport emerged as the key candidate to be an invariant component of the humanoid morphology. When analysing the effect of walking parameters on torque consumption, there was a consistent disparity between the feedback control cases (scenarios 1 & 2) and the optimal control case (scenario 3), but when the same parameters were viewed on the CoT landscape, there was a consistent convergence towards similar gradients and basins of attraction. Figures 6.13, 6.15 and 6.21 in particular show this convergence for different walking parameters on the CoT landscape. Additionally, simple, linear regions of optimality emerged between step length, step time and Cost of Transport in all tested scenarios, providing a computationally rapid method of selecting energy-optimal step parameters.

These findings are encouraging, as they offer insight into the mechanisms determining human locomotion behaviour, as discussed in the next section, and provide a means of more efficient, faster to compute locomotion behaviour for robot locomotion control.

6.5.1 Efficiency in Human Locomotion

In this chapter and the previous one, we observed that the complexity of whole body dynamics did not necessarily translate to complexity in the energy landscape. Notably, the feedback cases exhibited a linear valley of near-optimal parameters surrounding the optimal ones. Interestingly, human step distance selection has a linear relationship with walking speed, but with highly stochastic results around this linear trend over multiple gait cycles [176]. This pattern is consistent with the valley of near-optimality observed in Figures 6.7 and 6.16.

Note that the arms of the robot were commanded to follow trajectories that kept them close to the side of the body with relatively high gain and did not explicitly contribute towards the locomotion performance. If they had been used to explicitly affect the locomotion, the resulting energy landscape may have been different.

Maps such as these are simple enough to be internalised as heuristics for humans, and are also effective in providing a wealth of easy-to-reach near optimal step positions that are easy to compute, with little cognitive load and match existing human data well. Since the results are consistent among our varied robot experimentation scenarios and capture the linearity and stochasticity of the results in [176], these results represent good candidates for being invariant, underlying principles of locomotion for the humanoid morphology.

Simple mappings with regions of near-optimal attraction are an attractive explanation for explaining human locomotive behaviour. They account for the need for human motion to be efficient, cognitively cheap to generate, and the sensor delays and imprecise motion inherent to human motion, as one could spend valuable cognitive resources computing the exact energy-optimal step position for a given context, only for motion execution to be imperfect, wasting that effort. Instead, patterns such as these allow humans to produce 'near-enough', near-optimal locomotion with minimal cognition and, even if footfall is imprecise for any reason, there is negligible energetic consequence.

6.5.2 Efficiency in Humanoid Evolution

These results also suggest that these invariances in the humanoid morphology act as a passive aid to locomotion, which much like a simple machine, or passive

walker, produces complex motion with minimal cognitive input, lending mechanical assistance to any control by virtue of its construction and its dynamics.

This is supported by the evolutionary trend towards increased locomotion efficiency and the optimality in other animals. Because of the shape of the biological body, it naturally aids walking and reduces the amount of cognitive effort needed to perform locomotion, freeing up these resources for other tasks. This is perhaps not a result of conscious effort on the part of the human to produce natural locomotion, but a consequence of the physiology, morphology and biology of the system, such that it would be energetically and cognitively harder to produce a gait that works against the dynamics of the body. In their study on human efficiency, Selinger et al. [168] observed that "energetic cost is not just an outcome of movement, but also plays a central role in continuously shaping it", which seems to be in line with our results.

Bipedal locomotion likely appeared as a result of using arms for manipulation rather than locomotion, and in turn, those who used this strategy were more likely to survive since they could use tools and carry loads. A possible by-product of this is a pressure towards a morphological composition that passively assists movement to free up cognitive resources in the same way as bipedalism freed up the arms for tool use, allowing early bipeds to spend these cognitive resources on other aspects of survival.

6.5.3 Efficiency Maps for Locomotion Planning

A major implementation benefit from having access to these analytical efficiency maps is that they can be queried in a similar manner to those in Chapter 5, such that when walking unimpeded on flat ground, the energy-optimal parameters can be easily extracted, so that an extensive online search is not needed.

Furthermore, the map can also be used to perform locomotion in constrained environments. Given an environment where the full range of motion is not available, such a stepping stones, the restricted, non-reachable part of the feasible region can be blocked off. Figure 6.24, for example shows the CoT global optimality map, where it is only possible to take steps between 30-60 cm. Within this 30 cm range, a ± 5 cm safety margin might be enforced to prevent stepping too close to the edge, leaving a 20 cm 'safe' stepping region. As in Figure 6.24, we

can query the map for globally CoT-optimal step parameters within this reduced workspace.

Benefits such as this make online computation of walking parameters more rapid, as the mapping has already been computed offline. This is especially true of selecting energy-optimal parameters. As illustrated by the amount of previous work in human subjects, obtaining an accurate representation of globally optimal step parameters is difficult without exploring the whole state-space, which is not feasible to do online.

6.5.4 Heuristics for Efficient Walking

From the results in Section 6.3.3.1 see simple, cascading heuristics for walking parameter selection globally and locally. If there are no constraints on stepping distance, the global case suggests the longer the step length the better in all three cases. There is, naturally a limit to this rule in humans, but in the robot case that is limited by toe-off motion, this seems to hold given current technology.

In the local case, after the step length has been selected, there appears be a simple rule that quicker steps are more efficient that slower steps. If the globally optimal step length is feasible and combined with its optimal step time, this will produce globally optimal locomotion for the given platform. But if a sub-optimal step length is chosen or enforced, these maps still provide a method of local optimality that is not explored in human studies. This demonstrates the value of such efficiency study for robot locomotion control, and the value of robots as bio-robotic stand ins, as a study of this resolution and consistency would not be possible in biological subjects.

In this study, we attempted to make these findings as general as possible by using different controllers and robots to study the invariance of the dynamics of the humanoid morphology. Still, the sample size is still too limited to apply to all cases. Given more study on a wider range of robots and control, especially with human-like biological models, enough research could produce a reliable and general method for generating global energy landscapes for a body given its dimensions and actuation parameters.

6.5.5 Future Work

The work carried out in this chapter again displayed the regions of attraction in optimality within sets of walking parameters in different scenarios. However, there were many constraints, such as single direction walking, orientation was not considered, and the arms were not part of the locomotive motions. To investigate this further, some avenues of research that can be explored are:

- **Including arm motion**: Humans use their arms during locomotion and this may contribute to the walking motion and efficiency. Explicitly modelling arm motion in these scenarios may have altered the landscape in terms of the size of the reachable area, the regions of efficiency, or in some other way. Further study into this variable would provide key insight into both the effect of the arms on efficiency and a more accurate view of the humanoid efficiency landscape;

- **Objective function alterations**: The choice and tuning of the objective function has a significant impact on the resulting motion in optimal control. In this case, minimizing the control effort, as shown in Equation 6.6, has an effect on the energy cost of the motion. However, the inconsistent scattering of energy optimality in the torque cost landscapes shown in Figure 6.4 may have influenced the CoT maps shown in Figure 6.13. To further investigate the impact of the objective function on the landscape, it would be interesting to adjust the weight or replace the control cost term with a term that minimizes CoT.

- **Optimal motion transfer**: Transferring optimal motions to a robot in a different context, simulator, or on real hardware might give a more realistic impression of the efficiency of the motions produced within the optimisation. If an optimised motion is replayed in a different context with the same robot, the efficiency may or may not be maintained and studying this will give us a fuller picture of the underlying dynamics of the morphology as a whole;

- **Improved controller tuning**: It is unclear whether the bottleneck in Figure 6.3b is due to the dynamics of the robot, or non-ideally tuned controller. Investigating the parameters of this controller in the scenario in future work would clarify this point and potential changes in landscapes would show the effect of control parameters on the efficiency landscape;

6.6 Conclusion

In this chapter, we investigated the relationship between continuous walking parameters and energy-efficiency over different scenarios. Each scenario altered either the robot or the controller, with feedback-based and optimal-control based methods.

We identified that there is a consistent parameter-efficiency relationship across these scenarios when CoT is used as the efficiency metric. This is consistent with existing work in biological human studies, which show that humans produce energy-optimal behaviour on the CoT level. From this, we propose that the humanoid morphology may be passively efficient and that its structure in itself aids the efficiency of the motion, which reduces the cognitive load for generating motion. This also supports the findings from Chapter 5, in that these complex landscapes can be captured in a simple heuristic that can be learned by developing humans and quickly queried to produce rapid motion

From a robotic point of view, these results give us insight into a simple heuristic for calculating energy-optimal stepping parameters, both globally and locally under spatial constraints. Such heuristics can be used as warm-starts for selecting optimally efficient step parameters for controllers quickly and accurately, as there is no need for computationally expensive exploration.

Chapter 7

Conclusion

The aim of this thesis was to investigate underlying, invariant principles of bipedal locomotion, with a focus on bioinspiration and biorobotics. Secondarily, we explored the concept of using mechatronic humanoid robots as a bio-robotic analogue for biological humans to overcome the difficulty in observing the internal control policies in humans. Here, we explore the findings of this thesis in relation to our initial aims.

7.1 Humanoid Optimality

The goal for this thesis was to study both bio-inspired and bio-robotic approaches to humanoid locomotion, including push recovery, but through the course of our investigations, the major theme that emerged was of optimality in human locomotion.

Through studies on biological humans, we found a trend towards optimality in various motions observed in different subjects specifically in minimum jerk control at the CoM. These results prompted extended investigation into optimality of humanoid locomotion, and we used mechatronic humanoids to study patterns of energy optimality. Viewed in the context of the previous work, this allows us to suggest that the composition of the humanoid morphology itself plays a part in promoting optimality during locomotion and that this is an invariant part of the underlying dynamics of locomotion, regardless of whether is system is biological or mechanical.

In Chapter 4, we observed that in push recovery, a narrow aspect of locomotion, decisions must be made reactively within a limited timeframe. It would therefore be reasonable to assume that optimality is not a high priority in this context. However, the experimental results showed that these motions could be modelled with a minimum-jerk optimal controller, which might suggest humans use a similar mechanism during push recovery. Further, this suggests that there are common features between push recovery strategies in the human morphology, addressing the first research question in Section 1.3.

Subsequently, we used humanoid robots as test-beds for biological humans and showed that in both push recovery (Chapter 5) and continuous walking (Chapter 6) to show that for these mechanisms, with similar morphology, a basin of efficient attraction exists, such that for a given movement speed, not only does a narrow set of optimal movement parameters exists, but this is surrounded by a band of parameters which produce motions that are only marginally less efficient than optimal, shown repeatedly in Figures 5.12, 6.7 and 6.15.

When combined, these findings give us insight that suggests not only have humans developed a control process to make efficient, often near-optimal motion decisions, but they are also working with their bodies and using the advantages of their body structure.

As humans, we do not inhabit and develop in a randomly assigned a morphology which hinders our ability to move and that we then must learn to pilot through extensive cognitive effort; we do, however embody a highly-tuned morphology, developed and refined through evolution over countless generations, which enables and promotes are walking behaviour, almost a plug-and-play morphology which promotes simple motions for us to live, thrive and survive. This perhaps provides a rationale to explain why passive walking robots which mimic the humanoid morphology are able to perform legged locomotion without any control input [2, 81, 80] and only by virtue of its similarity to the biological musculoskeletal system of humans.

7.1.1 Energy-Efficiency as a Constrained Optimisation

When considered in the context of the current knowledge in biological and robotic locomotion, the findings of this thesis reveal an interesting feature of humanoid

locomotion efficiency. As discussed in Section 6.1, the consensus in the field of locomotion is that human walking is globally optimally efficient during self-selected, unimpeded walking [146, 147, 148]. Yet, it is still unknown how humans regulate their efficiency, as metabolic sensing pathways [162, 163, 164, 165, 166] act too slowly to explain in real-time convergence to global optimality after being perturbed [168].

Furthermore, biological studies into scenarios where stability or completion time are higher priorities than efficiency, such as on rough terrain, cluttered or restricted environments show a reduction in energy efficiency relative to the global optimal during preferred, flat ground walking [160]. When formulated as an optimisation problem in [160], motion cost is included as a decision variable, as in Equation 6.1 (repeated below for convenience):

$$MIN_{\{s\}}[w_1 \times COT(s) + w_2 \times 1/Stability(s) + w_3 \times Time(s)]$$

This suggests that a decision is made based on some optimality trade-off or negotiation that is affordable and achievable whilst also maintaining stability. This implies that humans are able to monitor their energy output in real-time to constantly re-evaluate this function to achieve reactive speeds, which, as discussed in Section 6.1.2 is not feasible. Not only does this require humans to sense their efficiency in real-time, it also requires a high-dimensional forward model to map the relationship between stability, completion time, and energy cost. Such a mapping would have inaccuracies caused by sensing delays and would likely be too cognitively demanding, which is at odds with rapid speed at which optimality is converged upon [168].

An alternative view presented by our results is: energy efficiency may not be a weighted term within the formulation along with other considerations (as in Equation 6.1), but instead, factors such as stability and completion time can be represented as constraints, within which the optimally efficient parameters can be selected, as shown in Figure 6.24. With this view, the energy-optimal step parameters do not further complicate the decision making process, but are instead selected as the global optimal of the *constrained* space and could be more accurately represented by a formulation similar to Equation 5.7

$$\mathscr{L} : (E) \mapsto \begin{cases} \mathbf{p}^* = f(E) \\ \text{s.t.:} v_{stability} \geq c_{stability} \\ \text{s.t.:} v_{time} = c_{time} \end{cases} \tag{7.1}$$

Where $\mathscr{L}$ is a mapping function from the current environment conditions and context E to a set of energy-optimal walking parameters $\mathbf{p}^*$. Function f outputs the walking conditions using the parameter relationships shown in Chapter 6 given the environment, subject to some values v of stability and time satisfy the constraints c imposed by the context in which locomotion is being performed. This reduces the demands on memory and computation, as only simple operations and patterns need to be internalised.

On its surface this appears to be a subtle distinction, but would have significant implications, namely that, given the cognitive load for selecting walking parameters would be significantly reduced by at least part of the function f results from the passive optimality of the morphology and heuristics of motion can be learned. Given that the stability and time constraints are applied to an existing mapping, they simply reduce the size of the global manifold such that the computation of optimal walking parameters consists of a simple mapping similar to Figure 6.24.

The results in this thesis can be used to support this view of invariant and passive optimality, but are not definitive, and more research is needed to explore this area and investigate the applicability to human biology.

7.2 Heuristics in Humans

Our results uncovered valleys of optimality and region of attraction for energy-optimal step parameters. These can be approximated to represent complex parameter selection in a compact form that does not require expensive computation to resolve. The patterns of optimality that were revealed are simple enough to be learned by humans, even at an early age; they can then be refined, memorised and form the basis of habitual motions [170].

Heuristics are widely used in human cognition, and are generally advantageous for rapid cognition and decision making. Humans have been shown to employ a

wealth of heuristics to make fast, approximate decisions; paradolia is a heuristic for pattern recognition in faces, we perceive expected words rather than actual words when reading as a heuristic shortcut; heuristics are so widespread that learning recognise them is a large part of critical thinking education [184]. This is widespread in human cognition and decision making, as rapid cognition is often vital for decision making, especially for survival in an evolutionary sense.

Heuristics for locomotion found in this thesis are could be evolutionarily advantageous for humans. Hunter gatherers would have to move across large distances whilst conserving precious energy, performing other tasks, such as tracking prey, observing for predators, it would be wasteful, perhaps even fatal, to spend cognitive resources on careful individual stepping motions over tens, or hundreds of thousands of steps and research has shown that humans perform step planning, both perception and motion planning, reactively for a very limited time-horizon, around 1.5 seconds in advance, even on rough terrain [185]. The heuristics observed in this work that could be a means of rapid human decision making address the final research question in Section 1.3 and provide insight that can be explored further in future work.

Using the observed regions of attraction in the humanoid morphology would be both physically and cognitively efficient, maximising movement distance whilst minimising scarce energy and cognitive load, which addresses the second research question in Section 1.3 Further investigation is required before we can definitively conclude these energy basins exists in humans, but they are strong counterfactual arguments which fit in to the human tendency towards heuristics, and is complementary to the morphology of the human. More research in this area would be hugely interesting for both biology and robotics and could lead to advancement in both fields.

7.3 Biorobotic Humanoids

We explored the arguments surrounding the use of robots as test-beds for hypotheses about biological human locomotion and found that the differences between humans and robots in terms of structure, actuation and control were outweighed by the advantages of the observability, control and repeatability of robots, which offers otherwise unattainable insights into the dynamics of locomo-

tion.

A key component of biorobotics is the bidirectional nature of the research [17], in that, mechatronic hardware platforms are primarily used to understand biological phenomena, but through this process, we also learn more about the platforms themselves. Through these insights, platforms can be iteratively improved to act as even more accurate test beds in future. This phenomenon was illustrated repeatedly throughout this work, particularly in Chapter 6, where gradients towards optimality were abruptly cut off due to the capabilities of the platforms or the controllers.

Adding toe-off motion and reducing the mass of links and joints is likely to improve the capabilities of the platforms, which has been a desire of the robotics community as long as it has existed and is unlikely to be realised anytime soon, but these results give added incentive to this development, as not only will these improvements lead to improved motion performance in general, but will also provide a more powerful tool for biological research.

7.4 Implications for Locomotion Control

The consistent theme throughout this thesis is the optimality of bipedal locomotion as a means of managing the highly redundant state-action space that is inherent in the humanoid morphology. Handling this redundancy is the primary challenge in robotic locomotion control, regardless of the style of control. Moreover, achieving energy efficiency in these controllers is particularly challenging due to the highly nonlinear landscape, making it difficult to compute. However, energy efficiency is desirable for prolonged operation in the real world.

Our findings have significant implications for control methods. We have presented methods to quantify the energy landscape of humanoid robots in an offline manner and shown that simple heuristics can capture the relationship between energy optimality and walking parameters that need to be selected as part of the control process. By warm-starting controllers using these simple relationships, expensive energy-efficiency calculations for selecting walking parameters can be avoided. This increases the speed of solution computation and the efficiency of the solutions themselves.

The relationship between step parameters and energy-optimal step positions can also be useful for hybrid planning methods. Mixed integer optimization techniques, such as [108, 68], can generate globally optimal step and motion plans for bipedal locomotion on rough 3D terrain, while accounting for dynamic and kinematic constraints. The simple relationships discovered in this study can be approximated as convex functions and incorporated into the optimization process to model the robot's reachability and efficiency during planning.

By studying additional robots and control methods in this way, it may be possible to produce a closed-form or simple numerical mapping from robot models to the parameter-efficiency relationship, such that offline studies do not need to be performed on every robot beforehand. A simple approximation of the relationship can be parametrically produced using the robot model.

7.5 Future Work

The results presented in this thesis represent incremental progress in the study of dynamic invariance in humanoid morphology, particularly with respect to optimality and efficiency. Furthermore, the study has raised several additional questions regarding the underlying principles of locomotion.

Further research could expand upon this work by verifying these effects on a wider range of robots and controllers, as well as through human-body simulation when this becomes feasible. Additionally, questions remain about how humans balance passive optimality against active control in various scenarios, how these properties emerged from an evolutionary perspective, and whether this effect persists when considering toe-off and running.

The questions raised regarding the underlying principles of locomotion present an exciting opportunity for cross-disciplinary investigation, which could lead to a deeper understanding of biological motion and the role of embodiment in movement.

To this end, we propose a potential structure for cross-disciplinary data collection and research aimed at achieving this goal:

1. **Baseline human data collection**: Human walking data under both constrained and unconstrained conditions is collected, including metabolic, dy-

namic, and kinematic data. Joint-torque and CoT are calculated in post-processing and use the data as a baseline for computational studies.

Next, we will analyze the data to gain insights into the underlying principles of human locomotion, including how different factors such as terrain, incline, and speed impact joint-torque and CoT. Our goal is to use these findings to inform the development of computational models that accurately capture the dynamics of human walking under varying conditions.

2. **Further robot implementation**: Building upon the work done on continuous locomotion in this thesis, the methods presented in this study can be extended to include more varied morphologies, control methods, and scenarios, such as turning and incline walking. To achieve this goal, we plan to implement and measure these methods on real hardware, collecting data that can be compared to the human baseline established in our previous work.

   If the findings are consistent, we will use the existing results as a foundation for developing more sophisticated computational models that accurately capture the dynamics of locomotion under diverse conditions. This will require additional research and data collection, but we believe that the potential benefits of such a comprehensive approach could be significant, paving the way for new advances in robotics and biomechanics.

3. **Simulated human motion**:As discussed in Section 2.4.1, non-reference-based control methods for human models in simulation are still in their early stages. However, given the rapid advancement in machine learning, it is likely that these experiments can be implemented soon on human models in OpenSim [101] or other software tools [186].

   Building upon the methodology presented in Chapter 6, further investigations can be conducted to determine whether the findings are consistent in human-model control in simulation. Such investigations could shed light on the mechanisms of efficient locomotion in humans and could provide new insights into the development of humanoid robots.

4. **Computational evolutionary biology**: By leveraging existing techniques in computational evolutionary biology, we can explore the trend towards passive optimality across the evolutionary timeline and determine whether

it was a factor in the success of certain species. Such investigations could deepen our understanding of the role of efficient movement in the evolution of life and inform the development of more effective robotic systems.

A multi-disciplinary approach involving roboticists, biologists, physiologists, anatomists, and evolutionary biologists would be necessary to conduct a large-scale study to investigate the role of passive optimality in evolutionary success. If such a study were successful and consistent with the results presented here, the effect size could be significant, potentially impacting a variety of fields, including rehabilitation, prosthetics, sport and exercise, and more.

Bibliography

[1] J. Muybridge, "The horse in motion, illus. by muybridge. "sallie gardner," owned by leland stanford, running at a 1:40 gait over the palo alto track, 19 june 1878: 2 frames showing diagram of foot movements." 1878. [Online]. Available: https://hdl.loc.gov/loc.pnp/ppmsca.06607

[2] S. H. Collins, M. Wisse, and A. Ruina, "A three-dimensional passive-dynamic walking robot with two legs and knees," *The International Journal of Robotics Research*, vol. 20, no. 7, pp. 607–615, 2001.

[3] O. Stasse, T. Flayols, R. Budhiraja, K. Giraud-Esclasse, J. Carpentier, J. Mirabel, A. Del Prete, P. Souères, N. Mansard, F. Lamiraux *et al.*, "Talos: A new humanoid research platform targeted for industrial applications," in *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*. IEEE, 2017, pp. 689–695.

[4] S. Song, Ł. Kidziński, X. B. Peng, C. Ong, J. Hicks, S. Levine, C. G. Atkeson, and S. L. Delp, "Deep reinforcement learning for modeling human locomotion control in neuromechanical simulation," *Journal of neuroengineering and rehabilitation*, vol. 18, no. 1, pp. 1–17, 2021.

[5] C. McGreavy, K. Yuan, D. Gordon, K. Tan, W. J. Wolfslag, S. Vijayakumar, and Z. Li, "Unified push recovery fundamentals: Inspiration from human study," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 10 876–10 882.

[6] C. O. Lovejoy, "Evolution of human walking," *Scientific American*, vol. 259, no. 5, pp. 118–125, 1988.

[7] C. V. Ward, "Interpreting the posture and locomotion of australopithecus afarensis: where do we stand?" *American Journal of Physical Anthro-*

pology: The Official Publication of the American Association of Physical Anthropologists, vol. 119, no. S35, pp. 185–215, 2002.

[8] D. M. Bramble and D. E. Lieberman, "Endurance running and the evolution of homo," *nature*, vol. 432, no. 7015, pp. 345–352, 2004.

[9] A. J. Ijspeert, "A connectionist central pattern generator for the aquatic and terrestrial gaits of a simulated salamander," *Biological Cybernetics*, vol. 84, no. 5, pp. 331–348, 2001.

[10] N. Van der Noot, A. J. Ijspeert, and R. Ronsse, "Biped gait controller for large speed variations, combining reflexes and a central pattern generator in a neuromuscular model," in *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2015, pp. 6267–6274.

[11] K. Mombaur, A. Truong, and J.-P. Laumond, "From human to humanoid locomotion—an inverse optimal control approach," *Autonomous robots*, vol. 28, no. 3, pp. 369–383, 2010.

[12] D. Gordon, G. Henderson, and S. Vijayakumar, "Effectively quantifying the performance of lower-limb exoskeletons over a range of walking conditions," *Frontiers in Robotics and AI*, 2018.

[13] C. Yang, K. Yuan, S. Heng, T. Komura, and Z. Li, "Learning natural locomotion behaviors for humanoid robots using human bias," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2610–2617, 2020.

[14] D. J. Farris, L. A. Kelly, A. G. Cresswell, and G. A. Lichtwark, "The functional importance of human foot muscles for bipedal locomotion," *Proceedings of the National Academy of Sciences*, vol. 116, no. 5, pp. 1645–1650, 2019.

[15] Z. Li, C. Zhou, Q. Zhu, and R. Xiong, "Humanoid balancing behavior featured by underactuated foot motion," *IEEE Transactions on Robotics*, 2017.

[16] B. Webb, "Robots in invertebrate neuroscience," *Nature*, vol. 417, no. 6886, pp. 359–363, 2002.

[17] B. Webb and T. Consilvio, *Biorobotics*. Mit Press, 2001.

[18] M. VUKOBRATOVIĆ and B. BOROVAC, "Zero-Moment Point — Thirty

Five Years of Its Life," *Int. Journal of Humanoid Robotics*, vol. 01, no. 01, pp. 157–173, 2004.

[19] S. Caron, Q.-C. Pham, and Y. Nakamura, "Zmp support areas for multicontact mobility under frictional constraints," *IEEE Transactions on Robotics*, vol. 33, no. 1, pp. 67–80, 2016.

[20] K. Nishiwaki and S. Kagami, "Simultaneous planning of com and zmp based on the preview control method for online walking control," in *2011 11th ieee-ras international conference on humanoid robots*. IEEE, 2011, pp. 745–751.

[21] C. Goldbeck, L. Kaul, N. Vahrenkamp, F. Worgotter, T. Asfour, and J.-M. Braun, "Two ways of walking: Contrasting a reflexive neuro-controller and a lip-based zmp-controller on the humanoid robot armar-4," in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2016, pp. 966–972.

[22] H. Hemami and C. Golliday Jr, "The inverted pendulum and biped stability," *Mathematical Biosciences*, vol. 34, no. 1-2, pp. 95–110, 1977.

[23] S. Kajita, F. Kanehiro, K. Kaneko, K. Yokoi, and H. Hirukawa, "The 3d linear inverted pendulum mode: A simple modeling for a biped walking pattern generation," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1, 2001, pp. 239–246.

[24] J. Pratt, J. Carff, S. Drakunov, and A. Goswami, "Capture point: A step toward humanoid push recovery," in *IEEE Int. Conf. on Humanoid Robots*, 2006.

[25] S. Kajita, M. Morisawa, K. Miura, S. Nakaoka, K. Harada, K. Kaneko, F. Kanehiro, and K. Yokoi, "Biped walking stabilization based on linear inverted pendulum tracking," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 4489–4496.

[26] M. Kelly and A. Ruina, "Non-linear robust control for inverted-pendulum 2d walking," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 4353–4358.

[27] A. Meduri, M. Khadiv, and L. Righetti, "Deepq stepper: A framework for

reactive dynamic walking on uneven terrain," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 2099–2105.

[28] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Yokoi, and H. Hirukawa, "A realtime pattern generator for biped walking," in *Proceedings 2002 IEEE international conference on robotics and automation (Cat. no. 02ch37292)*, vol. 1. IEEE, 2002, pp. 31–37.

[29] J. Englsberger, C. Ott, and A. Albu-Schäffer, "Three-Dimensional Bipedal Walking Control Based on Divergent Component of Motion," *IEEE Transactions on Robotics*, vol. 31, no. 2, pp. 355–368, 2015.

[30] H. Hoffmann, P. Pastor, D.-H. Park, and S. Schaal, "Biologically-inspired dynamical systems for movement generation: Automatic real-time goal adaptation and obstacle avoidance," *IEEE Int. Conf. on Robotics and Automation 2009*, 2009.

[31] J. Pratt, T. Koolen, T. De Boer, J. Rebula, S. Cotton, J. Carff, M. Johnson, and P. Neuhaus, "Capturability-based analysis and control of legged locomotion, part 2: Application to m2v2, a lower-body humanoid," *The int. jrnl of robotics research*, 2012.

[32] M. H. Raibert, *Legged robots that balance*. MIT press, 1986.

[33] F. Iida, Y. Minekawa, J. Rummel, and A. Seyfarth, "Toward a human-like biped robot with compliant legs," *Robotics and Autonomous Systems*, vol. 57, no. 2, pp. 139–144, 2009.

[34] F. Asano, M. Yamakita, N. Kamamichi, and Z.-W. Luo, "A novel gait generation for biped walking robots based on mechanical energy constraint," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 3, pp. 565–573, 2004.

[35] I. R. Manchester and J. Umenberger, "Real-time planning with primitives for dynamic walking over uneven terrain," in *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2014, pp. 4639–4646.

[36] S. Faraji and A. J. Ijspeert, "3lp: A linear 3d-walking model including torso and swing dynamics," *the international journal of robotics research*, 2017.

[37] R. D. Gregg, A. K. Tilton, S. Candido, T. Bretl, and M. W. Spong, "Control and planning of 3-d dynamic walking with asymptotically stable gait primitives," *IEEE Transactions on Robotics*, vol. 28, no. 6, pp. 1415–1423, 2012.

[38] S. Feng, E. Whitman, X. Xinjilefu, and C. G. Atkeson, "Optimization-based full body control for the darpa robotics challenge," *Journal of field robotics*, vol. 32, no. 2, pp. 293–312, 2015.

[39] G. Wiedebach, S. Bertrand, T. Wu, L. Fiorio, S. McCrory, R. Griffin, F. Nori, and J. Pratt, "Walking on partial footholds including line contacts with the humanoid robot atlas," in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2016, pp. 1312–1319.

[40] P. Fernbach, S. Tonneau, and M. Taïx, "Croc: Convex resolution of centroidal dynamics trajectories to provide a feasibility criterion for the multi contact planning problem," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2018.

[41] J. Carpentier, S. Tonneau, M. Naveau, O. Stasse, and N. Mansard, "A versatile and efficient pattern generator for generalized legged locomotion," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 3555–3561.

[42] S. Tonneau, A. Del Prete, J. Pettré, C. Park, D. Manocha, and N. Mansard, "An efficient acyclic contact planner for multiped robots," *IEEE Transactions on Robotics*, 2018.

[43] 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 *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2536–2542.

[44] A. Herdt, N. Perrin, and P.-B. Wieber, "Walking without thinking about it," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2010.

[45] A. D. Ames, "Human-inspired control of Bipedal walking robots," *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1115–1130, 2014.

[46] I. Mordatch, E. Todorov, and Z. Popović, "Discovery of complex behaviors through contact-invariant optimization," *ACM Transactions on Graphics (TOG)*, vol. 31, no. 4, pp. 1–8, 2012.

[47] A. W. Winkler, C. D. Bellicoso, M. Hutter, and J. Buchli, "Gait and trajectory optimization for legged systems through phase-based end-effector parameterization," *IEEE Robotics and Automation Letters*, 2018.

[48] M. Neunert, M. Stäuble, M. Giftthaler, C. D. Bellicoso, J. Carius, C. Gehring, M. Hutter, and J. Buchli, "Whole-body nonlinear model predictive control through contacts for quadrupeds," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1458–1465, 2018.

[49] J. Ahn, J. Lee, and L. Sentis, "Data-efficient and safe learning for humanoid locomotion aided by a dynamic balancing model," *IEEE Robotics and Automation Letters*, 2020.

[50] S. M. Kasaei, "A robust, agile and versatile humanoid locomotion based on analytical control and residual physics," Ph.D. dissertation, Universidade de Aveiro, 2022.

[51] V. Tsounis, M. Alge, J. Lee, F. Farshidian, and M. Hutter, "Deepgait: Planning and control of quadrupedal gaits using deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3699–3706, 2020.

[52] X. B. Peng, G. Berseth, K. Yin, and M. Van De Panne, "Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, pp. 1–13, 2017.

[53] J. Chemin, P. Fernbach, D. Song, G. Saurel, N. Mansard, and S. Tonneau, "Learning to steer a locomotion contact planner," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4430–4437.

[54] S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, and I. Havoutis, "Rloc: Terrain-aware legged locomotion using reinforcement learning and optimal control," *IEEE Transactions on Robotics*, 2022.

[55] X. B. Peng, G. Berseth, and M. Van de Panne, "Dynamic terrain traver-

sal skills using reinforcement learning," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, pp. 1–11, 2015.

[56] C. Yang, K. Yuan, Q. Zhu, W. Yu, and Z. Li, "Multi-expert learning of adaptive legged locomotion," *Science Robotics*, vol. 5, no. 49, 2020.

[57] X. B. Peng, A. Kanazawa, J. Malik, P. Abbeel, and S. Levine, "Sfv: Reinforcement learning of physical skills from videos," *ACM Transactions On Graphics (TOG)*, vol. 37, no. 6, pp. 1–14, 2018.

[58] Z. Li, X. Cheng, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, "Reinforcement learning for robust parameterized locomotion control of bipedal robots," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 2811–2817.

[59] W. S. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *The bulletin of mathematical biophysics*, vol. 5, no. 4, pp. 115–133, 1943.

[60] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[61] K. Mombaur, A.-H. Olivier, and A. Crétual, "Forward and inverse optimal control of bipedal running," in *Modeling, simulation and optimization of bipedal walking*. Springer, 2013, pp. 165–179.

[62] D. Kulić, C. Ott, D. Lee, J. Ishikawa, and Y. Nakamura, "Incremental learning of full body motion primitives and their sequencing through human motion observation," *The International Journal of Robotics Research*, vol. 31, no. 3, pp. 330–345, 2012.

[63] W. Takano, D. Kulic, and Y. Nakamura, "Interactive topology formation of linguistic space and motion space," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2007, pp. 1416–1422.

[64] H. Geyer and H. Herr, "A muscle-reflex model that encodes principles of legged mechanics produces human walking dynamics and muscle activities," *IEEE Transactions on neural systems and rehabilitation engineering*, vol. 18, no. 3, pp. 263–273, 2010.

[65] S. Dutta, A. Parihar, A. Khanna, J. Gomez, W. Chakraborty, M. Jerry, B. Grisafe, A. Raychowdhury, and S. Datta, "Programmable coupled oscillators for synchronized locomotion," *Nature communications*, vol. 10, no. 1, pp. 1–10, 2019.

[66] J. H. Barron-Zambrano and C. Torres-Huitzil, "Fpga implementation of a configurable neuromorphic cpg-based locomotion controller," *Neural Networks*, vol. 45, pp. 50–61, 2013.

[67] J. Muybridge, "The horse in motion," *Nature*, vol. 25, no. 652, pp. 605–605, 1882.

[68] B. Aceituno-Cabezas, C. Mastalli, H. Dai, M. Focchi, A. Radulescu, D. G. Caldwell, J. Cappelletto, J. C. Grieco, G. Fernández-López, and C. Semini, "Simultaneous contact, gait, and motion planning for robust multilegged locomotion via mixed-integer convex optimization," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2531–2538, 2017.

[69] J. Wang, I. Chatzinikolaidis, C. Mastalli, W. Wolfslag, G. Xin, S. Tonneau, and S. Vijayakumar, "Automatic gait pattern selection for legged robots," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 3990–3997.

[70] L. Amatucci, J.-H. Kim, J. Hwangbo, and H.-W. Park, "Monte carlo tree search gait planner for non-gaited legged system control," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 4701–4707.

[71] C. Wang, Y. Wai, B. Kuo, Y.-Y. Yeh, and J. Wang, "Cortical control of gait in healthy humans: an fmri study," *Journal of neural transmission*, vol. 115, no. 8, pp. 1149–1158, 2008.

[72] K. Gramann, J. T. Gwin, D. P. Ferris, K. Oie, T.-P. Jung, C.-T. Lin, L.-D. Liao, and S. Makeig, "Cognition in action: imaging brain/body dynamics in mobile humans," *Reviews in the Neurosciences*, 2011.

[73] D. Hamacher, F. Herold, P. Wiegel, D. Hamacher, and L. Schega, "Brain activity during walking: a systematic review," *Neuroscience & Biobehavioral Reviews*, vol. 57, pp. 310–327, 2015.

[74] R. Othayoth, Q. Xuan, Y. Wang, and C. Li, "Locomotor transitions in the potential energy landscape-dominated regime," *Proceedings of the Royal Society B*, vol. 288, no. 1949, p. 20202734, 2021.

[75] D. Lewis, *Counterfactuals*. John Wiley & Sons, 2013.

[76] H. Cruse, T. Kindermann, M. Schumm, J. Dean, and J. Schmitz, "Walknet—a biologically inspired network to control six-legged walking," *Neural networks*, vol. 11, no. 7-8, pp. 1435–1447, 1998.

[77] S. M. Kasaei, D. Simões, N. Lau, and A. Pereira, "A hybrid zmp-cpg based walk engine for biped robots," in *Iberian Robotics conference*. Springer, 2017, pp. 743–755.

[78] N. Sugimoto and J. Morimoto, "Phase-dependent trajectory optimization for cpg-based biped walking using path integral reinforcement learning," in *2011 11th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2011, pp. 255–260.

[79] M. Mokhtari, M. Taghizadeh, and M. Mazare, "Hybrid adaptive robust control based on cpg and zmp for a lower limb exoskeleton," *Robotica*, vol. 39, no. 2, pp. 181–199, 2021.

[80] T. McGeer *et al.*, "Passive dynamic walking," *Int. J. Robotics Res.*, vol. 9, no. 2, pp. 62–82, 1990.

[81] S. Collins, A. Ruina, R. Tedrake, and M. Wisse, "Efficient bipedal robots based on passive-dynamic walkers," *Science*, vol. 307, no. 5712, pp. 1082–1085, 2005.

[82] A. Ruina, "Passive dynamics is a good basis for robot design and control, not!" in *APS March Meeting Abstracts*, vol. 2017, 2017, pp. X12–013.

[83] A. J. Ijspeert, "Biorobotics: Using robots to emulate and investigate agile locomotion," *science*, vol. 346, no. 6206, pp. 196–203, 2014.

[84] C. E. Clauser, J. T. McConville, and J. W. Young, "Weight, volume, and center of mass of segments of the human body," Antioch Coll Yellow Springs OH, Tech. Rep., 1969.

[85] R. M. Alexander, *Optima for animals*. Princeton University Press, 1996.

[86] B. Zhou, H. Zeng, F. Wang, Y. Li, and H. Tian, "Efficient and robust reinforcement learning with uncertainty-based value expansion," *arXiv preprint arXiv:1912.05328*, 2019.

[87] S. Kajita, F. Kanehiro, K. Kaneko, K. Yokoi, and H. Hirukawa, "The 3D linear inverted pendulum mode: a simple modeling for a biped walking pattern generation," *Proceedings 2001 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No.01CH37180)*, vol. 1, pp. 239–246, 2001.

[88] B. Stephens, "Humanoid push recovery," in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2007.

[89] Z. Aftab, T. Robert, and P.-B. Wieber, "Ankle, hip and stepping strategies for humanoid balance recovery with a single model predictive control scheme," in *IEE Int. Conf. on Humanoid Robots 2012*, 2012.

[90] C. G. Atkeson and B. Stephens, "Multiple balance strategies from one optimization criterion," in *IEEE Int. Conf. on Humanoid Robots 2007*, 2007.

[91] L. Kaul and T. Asfour, "Human Push-Recovery:Strategy Selection Based on Push Intensity Estimation," *Int. Symp. on Robotics*, 2016.

[92] L. M. Nashner and G. McCollum, "The organization of human postural movements: a formal basis and experimental synthesis," *Behavioral and brain sciences*, 1985.

[93] Y.-C. Pai and J. Patton, "Center of mass velocity-position predictions for balance control," *Journal of biomechanics*, 1997.

[94] S. Kajita and K. Tani, "Study of dynamic biped locomotion on rugged terrain-derivation and application of the linear inverted pendulum model," *IEEE Int. Conf. on Robotics and Automation*, 1991.

[95] J. Allum and F. Honegger, "Interactions between vestibular and proprioceptive inputs triggering and modulating human balance-correcting responses differ across muscles," *Experimental brain research*, 1998.

[96] P. Zaytsev, W. Wolfslag, and A. Ruina, "The boundaries of walking stability: viability and controllability of simple models," *IEEE Trans. on Robotics*, 2018.

[97] K. B. Cheng, H. Tanabe, W. C. Chen, and H. T. Chiu, "Role of heel lifting in standing balance recovery: A simulation study," *Journal of Biomechanics*, 2018.

[98] K. Seo, J. Kim, and K. Roh, "Towards natural bipedal walking: Virtual gravity compensation and capture point control," in *IEEE Int. Conf. on Intelligent Robots and Systems*, 2012.

[99] C. Yang, K. Yuan, W. Merkt, T. Komura, S. Vijayakumar, and Z. Li, "Learning whole-body motor skills for humanoids," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2018.

[100] A. Hofmann, "Robust Execution of Bipedal Walking Tasks From Biomechanical Principles," *Computer Science and Artificial Intelligence Laboratory Technical Report*, 2006.

[101] S. L. Delp, F. C. Anderson, A. S. Arnold, P. Loan, A. Habib, C. T. John, E. Guendelman, and D. G. Thelen, "Opensim: open-source software to create and analyze dynamic simulations of movement," *IEEE transactions on biomedical engineering*, 2007.

[102] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, and H. Hirukawa, "Biped walking pattern generation by using preview control of zero-moment point," *IEEE Int. Conf. on Robotics and Automation*, 2003.

[103] Z. Li, C. Zhou, H. Dallali, N. G. Tsagarakis, and D. G. Caldwell, "Comparison study of two inverted pendulum models for balance recovery," in *2014 IEEE-RAS International Conference on Humanoid Robots*, 2014, pp. 67–72.

[104] T. Flash and N. Hogan, "The coordination of arm movements: an experimentally confirmed mathematical model," *Journal of Neuroscience*, vol. 5, no. 7, pp. 1688–1703, 1985.

[105] W. Hu, I. Chatzinikolaidis, K. Yuan, and Z. Li, "Comparison study of nonlinear optimization of step durations and foot placement for dynamic walking," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 433–439.

[106] H. Dai, A. Valenzuela, and R. Tedrake, "Whole-body motion planning with

centroidal dynamics and full kinematics," in *IEEE/RAS International Conference on Humanoid Robots*, 2014, pp. 295–302.

[107] O. Khatib, L. Sentis, J. Park, and J. Warren, "Whole-body dynamic behavior and control of human-like robots," *International Journal of Humanoid Robotics*, vol. 1, no. 01, pp. 29–43, 2004.

[108] R. Deits and R. Tedrake, "Footstep planning on uneven terrain with mixed-integer convex optimization," in *Proc. IEEE/RAS International Conference on Humanoid Robots*, 2014, pp. 279–286.

[109] Y.-D. Hong and B. Lee, "Real-time feasible footstep planning for bipedal robots in three-dimensional environments using particle swarm optimization," *IEEE/ASME Transactions on Mechatronics*, vol. 25, no. 1, pp. 429–437, 2019.

[110] L. Baudouin, N. Perrin, T. Moulard, F. Lamiraux, O. Stasse, and E. Yoshida, "Real-time replanning using 3d environment for humanoid robot," in *IEEE/RAS International Conference on Humanoid Robots*, 2011, pp. 584–589.

[111] B. Ponton, A. Herzog, S. Schaal, and L. Righetti, "A convex model of humanoid momentum dynamics for multi-contact motion generation," in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2016, pp. 842–849.

[112] S. Tonneau, D. Song, P. Fernbach, N. Mansard, M. Taïx, and A. Del Prete, "Sl1m: Sparse l1-norm minimization for contact planning on uneven terrain," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 6604–6610.

[113] J. A. Castano, Z. Li, C. Zhou, N. Tsagarakis, and D. Caldwell, "Dynamic and reactive walking for humanoid robots based on foot placement control," *International Journal of Humanoid Robotics*, vol. 13, no. 02, p. 1550041, 2016.

[114] Z. Li, C. Zhou, N. Tsagarakis, and D. Caldwell, "Compliance control for stabilizing the humanoid on the changing slope based on terrain inclination estimation," *Autonomous Robots*, vol. 40, no. 6, pp. 955–971, 2016.

[115] J. Ding, C. Zhou, and X. Xiao, "Energy-efficient bipedal gait pattern generation via com acceleration optimization," in *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2018, pp. 238–244.

[116] V. A. Tucker, "The energetic cost of moving about: walking and running are extremely inefficient forms of locomotion. much greater efficiency is achieved by birds, fish—and bicyclists," *American Scientist*, vol. 63, no. 4, pp. 413–419, 1975.

[117] E. Ambrose, W.-L. Ma, C. Hubicki, and A. D. Ames, "Toward benchmarking locomotion economy across design configurations on the modular robot: Amber-3m," in *2017 IEEE Conference on Control Technology and Applications (CCTA)*. IEEE, 2017, pp. 1270–1276.

[118] W. Xi, Y. Yesilevskiy, and C. D. Remy, "Selecting gaits for economical locomotion of legged robots," *The International Journal of Robotics Research*, vol. 35, no. 9, pp. 1140–1154, 2016.

[119] M. Brandao, K. Hashimoto, J. Santos-Victor, and A. Takanishi, "Optimizing energy consumption and preventing slips at the footstep planning level," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 1–7.

[120] Z. Li, N. G. Tsagarikis, D. G. Caldwell, and B. Vanderborght, "Trajectory generation of straightened knee walking for humanoid robot icub," in *2010 11th International Conference on Control Automation Robotics & Vision*. IEEE, 2010, pp. 2355–2360.

[121] M. S. Erden and K. Leblebicioglu, "Torque distribution in a six-legged robot," *IEEE Transactions on Robotics*, vol. 23, no. 1, pp. 179–186, 2007.

[122] J. E. Bobrow, B. Martin, G. Sohl, E. Wang, F. C. Park, and J. Kim, "Optimal robot motions for physical criteria," *Journal of Robotic systems*, vol. 18, no. 12, pp. 785–795, 2001.

[123] D. P. Garg and M. Kumar, "Optimization techniques applied to multiple manipulators for path planning and torque minimization," *Engineering applications of artificial intelligence*, vol. 15, no. 3-4, pp. 241–252, 2002.

[124] N. Perrin, O. Stasse, L. Baudouin, F. Lamiraux, and E. Yoshida, "Fast humanoid robot collision-free footstep planning using swept volume approximations," *IEEE Transactions on Robotics*, vol. 28, no. 2, pp. 427–439, 2011.

[125] M. Fallon, S. Kuindersma, S. Karumanchi, M. Antone, T. Schneider, H. Dai, C. P. D'Arpino, R. Deits, M. DiCicco, D. Fourie *et al.*, "An architecture for online affordance-based perception and whole-body planning," *Journal of Field Robotics*, vol. 32, no. 2, pp. 229–254, 2015.

[126] M. Murooka, I. Kumagai, M. Morisawa, F. Kanehiro, and A. Kheddar, "Humanoid loco-manipulation planning based on graph search and reachability maps," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1840–1847, 2021.

[127] J. J. Kuffner, K. Nishiwaki, S. Kagami, M. Inaba, and H. Inoue, "Footstep planning among obstacles for biped robots," in *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No. 01CH37180)*, vol. 1. IEEE, 2001, pp. 500–505.

[128] M. Kallman and M. Mataric, "Motion planning using dynamic roadmaps," in *IEEE International Conference on Robotics and Automation*, vol. 5, 2004, pp. 4399–4404.

[129] S. J. Jorgensen, M. Vedantam, R. Gupta, H. Cappel, and L. Sentis, "Finding locomanipulation plans quickly in the locomotion constrained manifold," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 6611–6617.

[130] F. Burget and M. Bennewitz, "Stance selection for humanoid grasping tasks by inverse reachability maps," in *2015 IEEE International conference on robotics and automation (ICRA)*. IEEE, 2015, pp. 5669–5674.

[131] Y. Yang, W. Merkt, H. Ferrolho, V. Ivan, and S. Vijayakumar, "Efficient humanoid motion planning on uneven terrain using paired forward-inverse dynamic reachability maps," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2279–2286, 2017.

[132] Y.-C. Lin, L. Righetti, and D. Berenson, "Robust humanoid contact

planning with learned zero-and one-step capturability prediction," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2451–2458, 2020.

[133] S. Crews and M. Travers, "Energy management through footstep selection for bipedal robots," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5485–5493, 2020.

[134] H. Dong, M. Zhao, and N. Zhang, "High-speed and energy-efficient biped locomotion based on virtual slope walking," *Autonomous Robots*, vol. 30, no. 2, pp. 199–216, 2011.

[135] F. M. Silva and J. T. Machado, "Energy analysis during biped walking," in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No. 99CH36288C)*, vol. 1. IEEE, 1999, pp. 59–64.

[136] M. F. Silva, J. T. Machado, and A. M. Lopes, "Energy analysis of multi-legged locomotion systems," in *Proc. CLAWAR'2001–4th International Symposium on Climbing and Walking Robots*, 2001, pp. 143–150.

[137] L. Kaul and T. Asfour, "Human push-recovery: Strategy selection based on push intensity estimation," in *Proceedings of 47st International Symposium on Robotics*, 2016, pp. 1–8.

[138] R. Kopitzsch, D. Clever, and K. Mombaur, "Optimization-based analysis of push recovery during walking motions to support the design of rigid and compliant lower limb exoskeletons," *Advanced Robotics*, vol. 31, no. 22, pp. 1238–1252, 2017.

[139] K. Yuan, I. Chatzinikolaidis, and Z. Li, "Bayesian optimization for whole-body control of high-degree-of-freedom robots through reduction of dimensionality," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2268–2275, 2019.

[140] A. Del Prete, N. Mansard, O. E. Ramos, O. Stasse, and F. Nori, "Implementing torque control with high-ratio gear boxes and without joint-torque sensors," *International Journal of Humanoid Robotics*, vol. 13, no. 01, p. 1550044, 2016.

[141] J. Englsberger, A. Werner, C. Ott, B. Henze, M. A. Roa, G. Garofalo, R. Burger, A. Beyer, O. Eiberger, K. Schmid *et al.*, "Overview of the

torque-controlled humanoid robot toro," in *2014 IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2014, pp. 916–923.

[142] C. Zhou, X. Wang, Z. Li, and N. Tsagarakis, "Overview of gait synthesis for the humanoid coman," *Journal of Bionic Engineering*, vol. 14, no. 1, pp. 15–25, 2017.

[143] J. Carpentier, G. Saurel, G. Buondonno, J. Mirabel, F. Lamiraux, O. Stasse, and N. Mansard, "The pinocchio c++ library: A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives," in *Proc. IEEE/SICE International Symposium on System Integration*, 2019, pp. 614–619.

[144] R. Featherstone, *Rigid body dynamics algorithms*. Springer, 2014.

[145] F. Nogueira, "Bayesian Optimization: Open source constrained global optimization tool for Python," 2014–.

[146] M. Y. Zarrugh, F. N. Todd, and H. J. Ralston, "Optimization of energy expenditure during level walking," *European Journal of Applied Physiology and Occupational Physiology*, vol. 33, no. 4, pp. 293–306, 1974.

[147] H. J. Ralston, "Energy-speed relation and optimal speed during level walking," *Internationale Zeitschrift für Angewandte Physiologie Einschliesslich Arbeitsphysiologie*, vol. 17, no. 4, pp. 277–283, 1958.

[148] B. R. Umberger and P. E. Martin, "Mechanical power and efficiency of level walking with different stride rates," *Journal of Experimental Biology*, vol. 210, no. 18, pp. 3255–3265, 2007.

[149] J. C. Selinger, J. L. Hicks, R. W. Jackson, C. M. Wall-Scheffler, D. Chang, and S. L. Delp, "Running in the wild: Energetics explain ecological running speeds," *Current Biology*, vol. 32, no. 10, pp. 2309–2315, 2022.

[150] J. K. Rathkey and C. M. Wall-Scheffler, "People choose to run at their optimal speed," *American Journal of Physical Anthropology*, vol. 163, no. 1, pp. 85–93, 2017.

[151] T. S. Carey and R. H. Crompton, "The metabolic costs of 'bent-hip, bent-knee'walking in humans," *Journal of human evolution*, vol. 48, no. 1, pp. 25–44, 2005.

[152] W. A. Sparrow and K. M. Newell, "Metabolic energy expenditure and the regulation of movement economy," *Psychonomic Bulletin & Review*, vol. 5, no. 2, pp. 173–196, 1998.

[153] D. Abe, S. Sakata, K. Motoyama, N. Toyota, H. Nishizono, and M. Horiuchi, "Economical and preferred walking speed using body weight support apparatus with a spring-like characteristics," *BMC Sports Science, Medicine and Rehabilitation*, vol. 13, no. 1, pp. 1–7, 2021.

[154] G. J. Bastien, P. A. Willems, B. Schepens, and N. C. Heglund, "Effect of load and speed on the energetic cost of human walking," *European journal of applied physiology*, vol. 94, no. 1, pp. 76–83, 2005.

[155] R. C. Browning and R. Kram, "Energetic cost and preferred speed of walking in obese vs. normal weight women," *Obesity research*, vol. 13, no. 5, pp. 891–899, 2005.

[156] M.-J. Hsu, D. H. Nielsen, H. J. Yack, and D. G. Shurr, "Physiological measurements of walking and running in people with transtibial amputations with 3 different prostheses," *Journal of Orthopaedic & Sports Physical Therapy*, vol. 29, no. 9, pp. 526–533, 1999.

[157] E. Monsch, C. Franz, and J. Dean, "The effects of gait strategy on metabolic rate and indicators of stability during downhill walking," *Journal of biomechanics*, vol. 45, no. 11, pp. 1928–1933, 2012.

[158] T. M. Lejeune, P. A. Willems, and N. C. Heglund, "Mechanics and energetics of human locomotion on sand." *The Journal of Experimental Biology*, vol. 201, no. 13, pp. 2071–2080, 1998.

[159] D. B. Kowalsky, J. R. Rebula, L. V. Ojeda, P. G. Adamczyk, and A. D. Kuo, "Human walking in the real world: Interactions between terrain type, gait parameters, and energy expenditure," *PLoS one*, vol. 16, no. 1, p. e0228682, 2021.

[160] K. Gast, R. Kram, and R. Riemer, "Preferred walking speed on rough terrain: is it all about energetics?" *Journal of experimental biology*, vol. 222, no. 9, p. jeb185447, 2019.

[161] M. Ackermann and A. J. Van den Bogert, "Optimality principles for model-

based prediction of human gait," *Journal of biomechanics*, vol. 43, no. 6, pp. 1055–1060, 2010.

[162] J. W. Bellville, B. J. Whipp, R. D. Kaufman, G. D. Swanson, K. A. Aqleh, and D. M. Wiberg, "Central and peripheral chemoreflex loop gain in normal and carotid body-resected subjects," *Journal of Applied Physiology*, vol. 46, no. 4, pp. 843–853, 1979.

[163] M. Fatemian, D. J. Nieuwenhuijs, L. J. Teppema, S. Meinesz, A. G. van der Mey, A. Dahan, and P. A. Robbins, "The respiratory response to carbon dioxide in humans with unilateral and bilateral resections of the carotid bodies," *The Journal of physiology*, vol. 549, no. 3, pp. 965–973, 2003.

[164] M. E. Pedersen, M. Fatemian, and P. A. Robbins, "Identification of fast and slow ventilatory responses to carbon dioxide under hypoxic and hyperoxic conditions in humans," *The Journal of Physiology*, vol. 521, no. 1, pp. 273–287, 1999.

[165] C. M. Adreani, J. M. Hill, and M. P. Kaufman, "Responses of group iii and iv muscle afferents to dynamic exercise," *Journal of applied physiology*, vol. 82, no. 6, pp. 1811–1817, 1997.

[166] P. G. Guyenet and D. A. Bayliss, "Neural control of breathing and co2 homeostasis," *Neuron*, vol. 87, no. 5, pp. 946–961, 2015.

[167] L. R. Webster and S. Karan, "The physiology and maintenance of respiration: a narrative review," *Pain and Therapy*, vol. 9, no. 2, pp. 467–486, 2020.

[168] J. C. Selinger, S. M. O'Connor, J. D. Wong, and J. M. Donelan, "Humans can continuously optimize energetic cost during walking," *Current Biology*, vol. 25, no. 18, pp. 2452–2456, 2015.

[169] P. Krishnaswamy, E. N. Brown, and H. M. Herr, "Human leg model predicts ankle muscle-tendon morphology, state, roles and energetics in walking," *PLoS computational biology*, vol. 7, no. 3, p. e1001107, 2011.

[170] A. De Rugy, G. E. Loeb, and T. J. Carroll, "Muscle coordination is habitual rather than optimal," *Journal of Neuroscience*, vol. 32, no. 21, pp. 7384–7391, 2012.

[171] W. I. Sellers, G. M. Cain, W. Wang, and R. H. Crompton, "Stride lengths, speed and energy costs in walking of australopithecus afarensis: using evolutionary robotics to predict locomotion of early human ancestors," *Journal of the Royal Society Interface*, vol. 2, no. 5, pp. 431–441, 2005.

[172] R. H. Crompton, E. E. Vereecke, and S. K. Thorpe, "Locomotion and posture from the common hominoid ancestor to fully modern hominins, with special reference to the last common panin/hominin ancestor," *Journal of anatomy*, vol. 212, no. 4, pp. 501–543, 2008.

[173] R. M. Alexander, "Models and the scaling of energy costs for locomotion," *Journal of Experimental Biology*, vol. 208, no. 9, pp. 1645–1652, 2005.

[174] G. A. Cavagna, H. Thys, and A. Zamboni, "The sources of external work in level walking and running." *The Journal of physiology*, vol. 262, no. 3, pp. 639–657, 1976.

[175] R. H. Miller, B. R. Umberger, J. Hamill, and G. E. Caldwell, "Evaluation of the minimum energy hypothesis and other potential optimality criteria for human running," *Proceedings of the Royal Society B: Biological Sciences*, vol. 279, no. 1733, pp. 1498–1505, 2012.

[176] S. H. Collins and A. D. Kuo, "Two independent contributions to step variability during over-ground human walking," *PloS one*, vol. 8, no. 8, p. e73597, 2013.

[177] R. M. Alexander, "A model of bipedal locomotion on compliant legs," *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 338, no. 1284, pp. 189–198, 1992.

[178] A. Minetti and R. M. Alexander, "A theory of metabolic costs for bipedal gaits," *Journal of Theoretical Biology*, vol. 186, no. 4, pp. 467–476, 1997.

[179] M. Srinivasan and A. Ruina, "Computer optimization of a minimal biped model discovers walking and running," *Nature*, vol. 439, no. 7072, pp. 72–75, 2006.

[180] A. D. Kuo, "A simple model of bipedal walking predicts the preferred speed–step length relationship," *J. Biomech. Eng.*, vol. 123, no. 3, pp. 264–269, 2001.

[181] S. Seok, A. Wang, M. Y. Chuah, D. J. Hyun, J. Lee, D. M. Otten, J. H. Lang, and S. Kim, "Design principles for energy-efficient legged locomotion and implementation on the mit cheetah robot," *Ieee/asme transactions on mechatronics*, vol. 20, no. 3, pp. 1117–1129, 2014.

[182] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," http://pybullet.org, 2016–2021.

[183] N. G. Tsagarakis, S. Morfey, G. M. Cerda, L. Zhibin, and D. G. Caldwell, "Compliant humanoid coman: Optimal joint stiffness tuning for modal frequency control," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 673–678.

[184] S. Novella, B. Novella, C. Santa Maria, J. Novella, and E. Bernstein, *The Skeptics' Guide to the Universe: How to Know What's Really Real in a World Increasingly Full of Fake*. Grand Central Publishing, 2018.

[185] J. S. Matthis, J. L. Yates, and M. M. Hayhoe, "Gaze and the control of foot placement when walking in natural terrain," *Current Biology*, vol. 28, no. 8, pp. 1224–1233, 2018.

[186] V. Caggiano, H. Wang, G. Durandau, M. Sartori, and V. Kumar, "Myosuite–a contact-rich simulation suite for musculoskeletal motor control," *arXiv preprint arXiv:2205.13600*, 2022.

Appendix A

Stability and Feasibility Analysis of MJMPC

The Inverted Pendulum Model simplifies the dynamics of a robot where the force can be applied along the pendulum and a torque can be used around the pivot. The dynamic equations are given as:

$$r^2\ddot{\theta} + 2r\dot{r}\dot{\theta} - gr\sin\theta = \tau/M \tag{A.1}$$

$$\ddot{r} - r\dot{\theta}^2 + g\cos\theta = f/M, \tag{A.2}$$

with $r = f(x, z), \theta = g(x, z)$, where

$$r = \sqrt{x^2 + z^2} \tag{A.3}$$

$$\theta = \text{asin}(\frac{x}{\sqrt{x^2 + z^2}}). \tag{A.4}$$

To validate the feasibility of the CoM motion, the required force and torque of the ground contact wrench need to be calculated by inputting the CoM motion $\boldsymbol{\xi} = [x, \dot{x}, \ddot{x}, z, \dot{z}, \ddot{z}]$ into (A.2). Max possible force $f_{\max}$ and torques $\tau_{\max}$ are calculated considering the joint torque limits $\boldsymbol{\tau}_{\lim}$. For every joint state $Q = [q, \dot{q}, \ddot{q}]^T$, the maximal admissible ground reaction wrench $\lambda_{\lim} = [f_{\max}, \tau_{\max}]^T$ are calculated via inverse dynamics as:

$$M(q)\ddot{q} + c(q, \dot{q}) + g(q) = \begin{bmatrix} 0_{6\times n} \\ I_{n\times n} \end{bmatrix} \boldsymbol{\tau}_{\lim} \tag{A.5}$$

$$\Rightarrow \lambda_{i,lim} = J_i(q)^{T\#}(M(q)\ddot{q} + c(q,\dot{q}) + g(q) - \begin{bmatrix} 0_{6\times n} \\ I_{n\times n} \end{bmatrix} \boldsymbol{\tau}_{\text{lim}}), \tag{A.6}$$

where $J(q)^{T\#}$ is the pseudo inverse of the Jacobian for ith foot, $M(q)$ is the inertia matrix, $c(q,\dot{q})$ non-linear effects, and $g(q)$ the gravity component.

Appendix B

Robustness Validation Plots

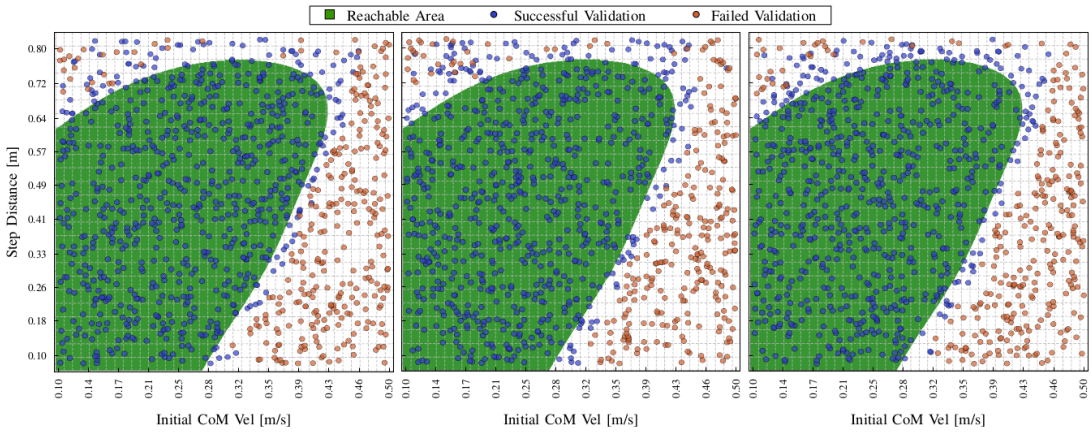

Figure B.1: Full validation testing maps with 1000 randomly sampled initial conditions for robot models with reduced mass compared to the original robot the optimal parameters and feedback controller were tuned for (a) $-17.3\,\%$ mass, (b) $-12.9\,\%$ mass, (c) $-5.8\,\%$ mass.

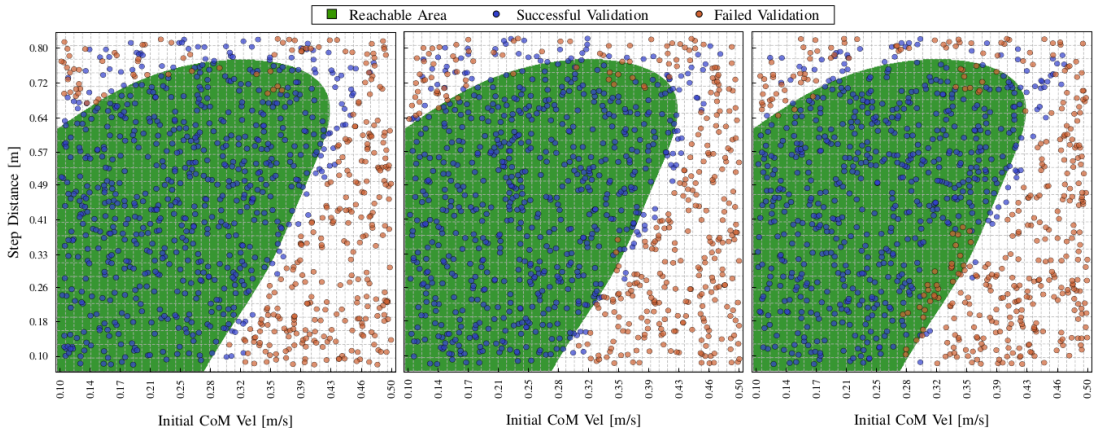

Figure B.2: Full validation testing maps with 1000 randomly sampled initial conditions for robot models with increased mass compared to the original robot the optimal parameters and feedback controller were tuned for (a) +6.2 % mass, (b) +13.1 % mass, (c) +18.5 % mass.

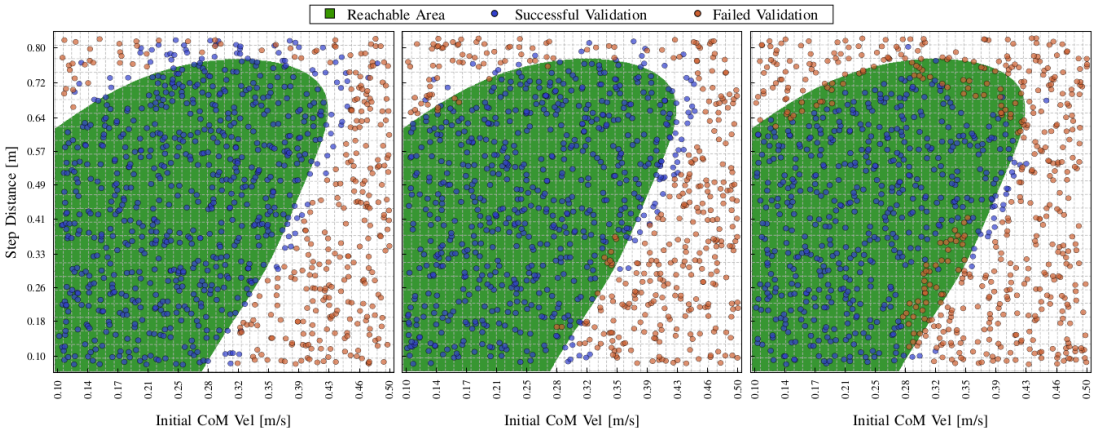

Figure B.3: Full validation testing maps with 1000 randomly sampled initial conditions with added feedback delays to simulate operation on the real robot, tests are shown with delays of (a) 10 ms, (b) 15 ms, (c) 20 ms.

Appendix C

3D Optimality Plots for Continuous Walking

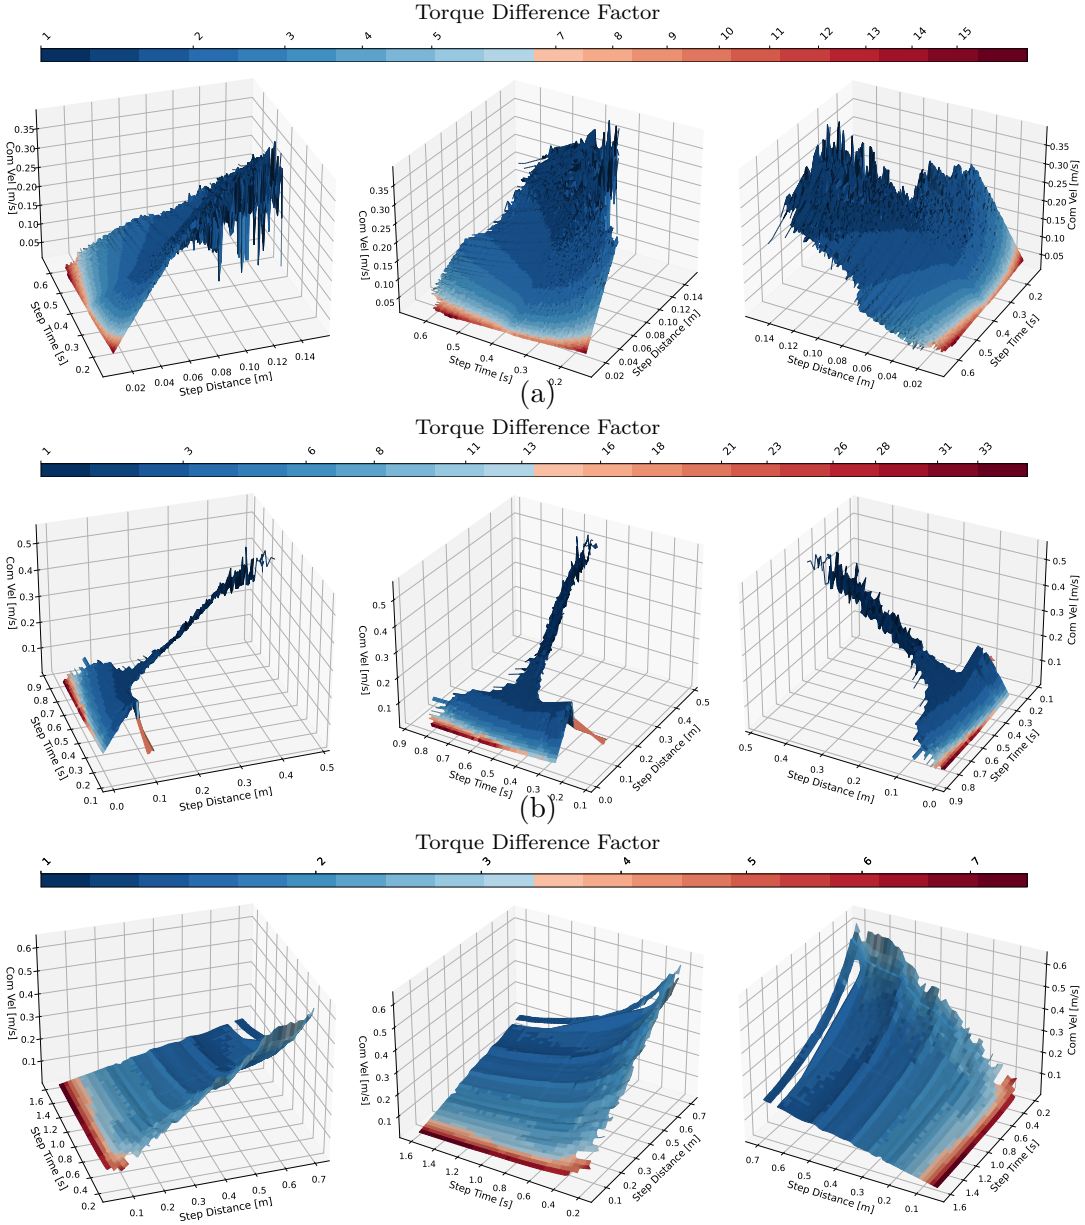

Figure C.1: Step Distance vs. Step Time vs. CoM Velocity, coloured by torque difference factor for scenario 1(a), 2(b) and 3(c).

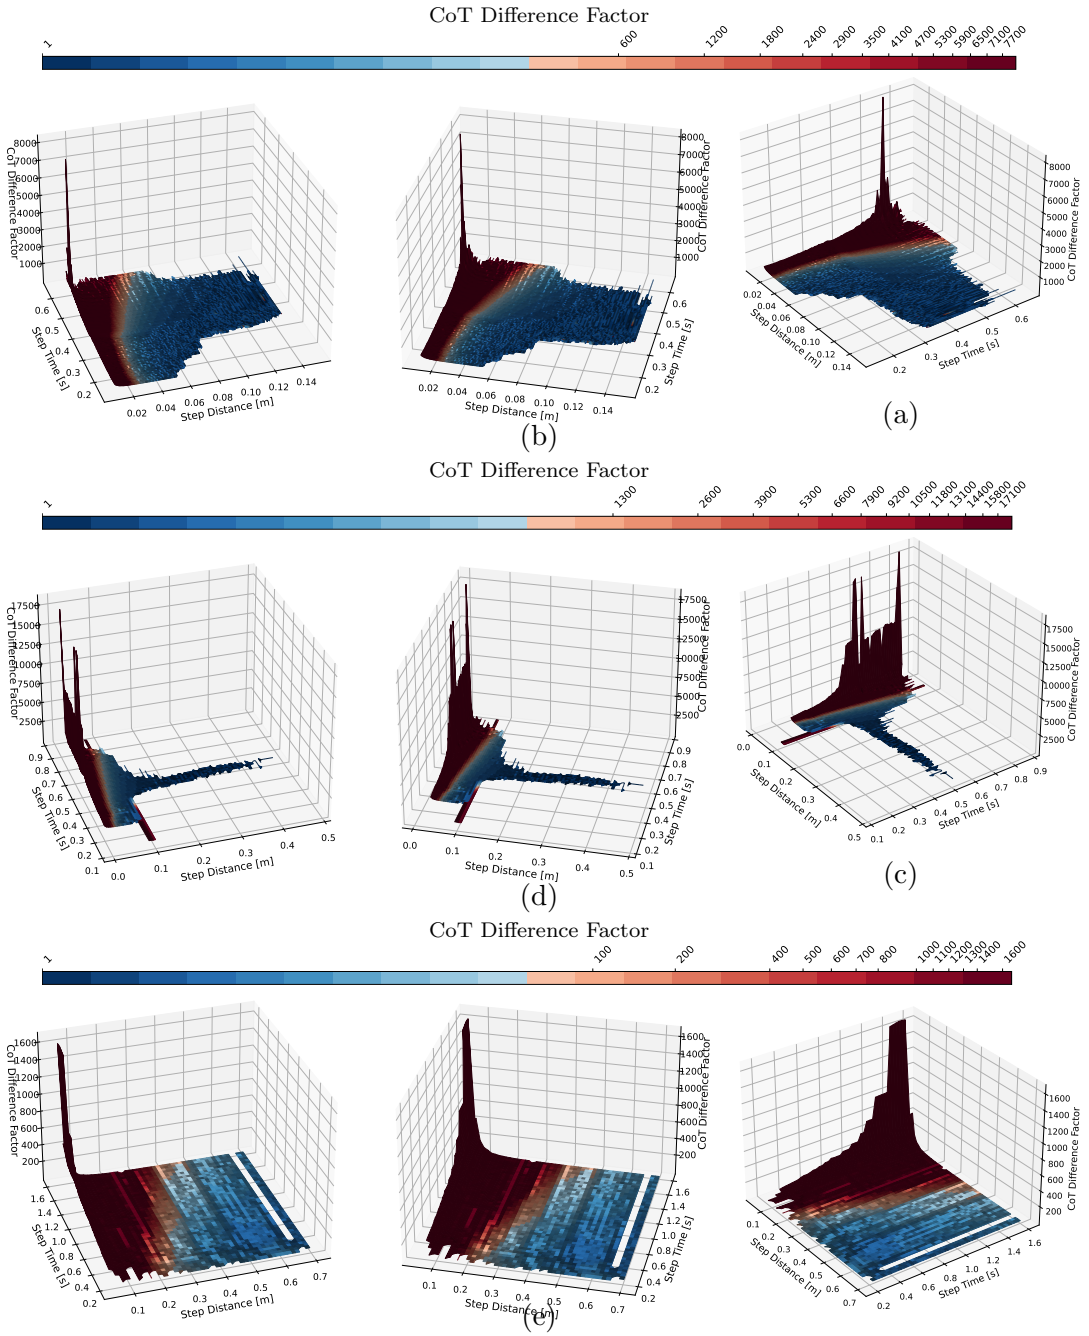

Figure C.2: Step Distance vs. Step Time vs. difference from global CoT minima for scenarios 1(a), 2(b) and 3(c).

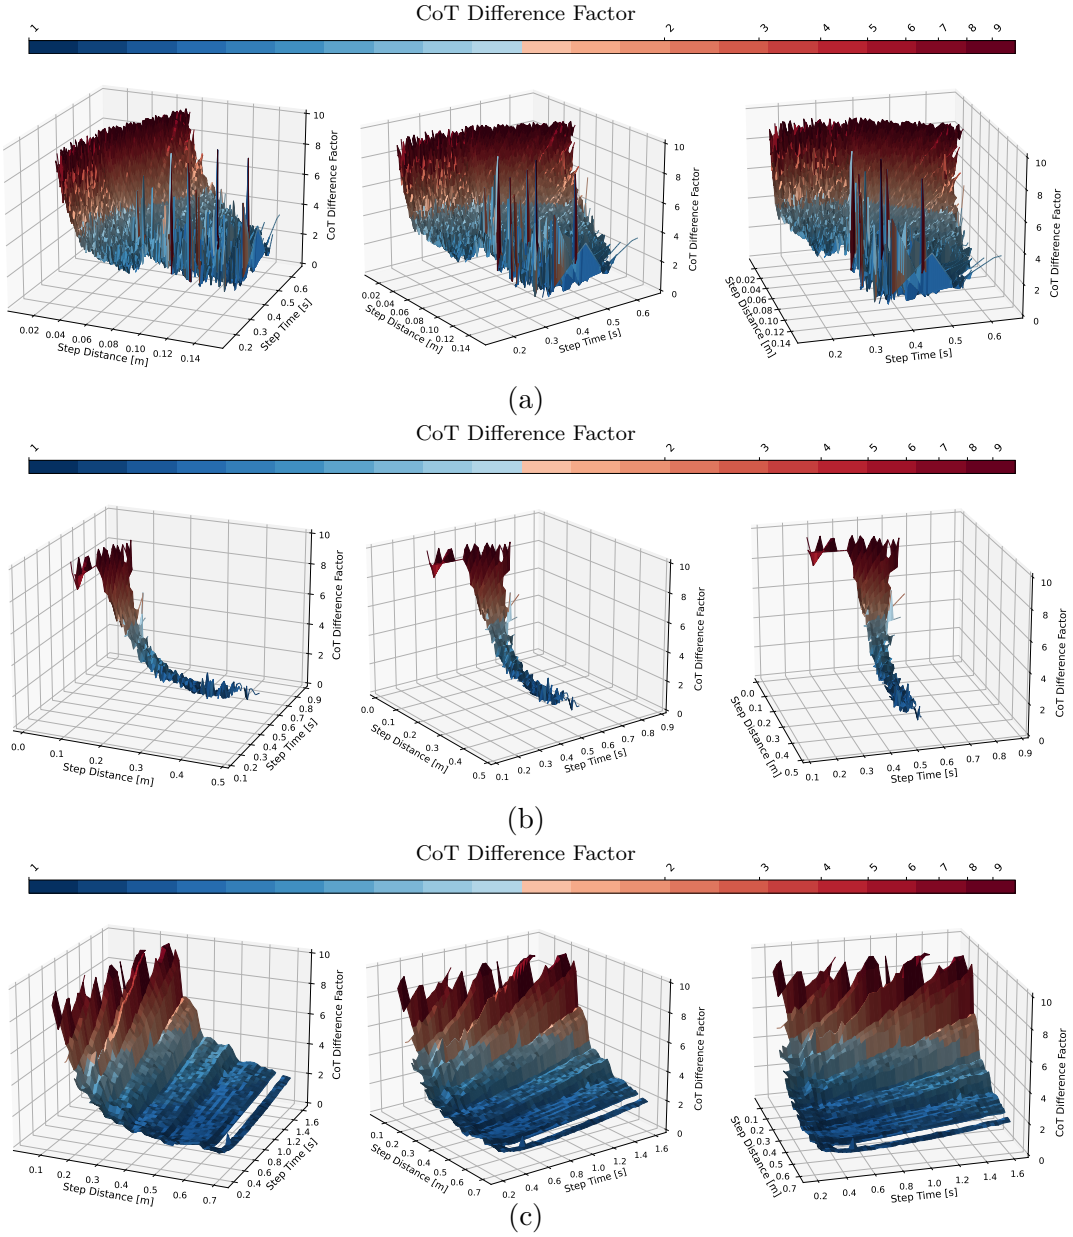

Figure C.3: Step Distance vs. Step Time vs. difference from global CoT minima. Shows the data in Figure C.2, zoomed in on CoT values up to 10 to show curvature for scenario 1(a), 2(b) and 3(c).

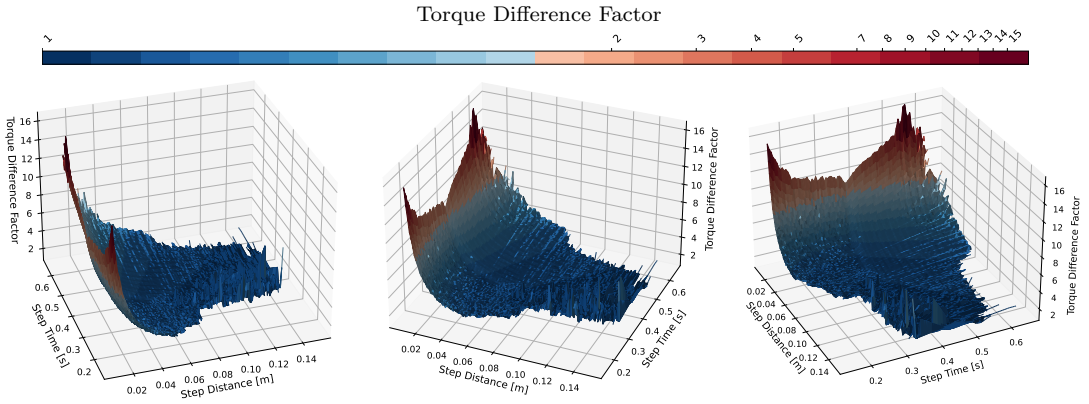

(a)

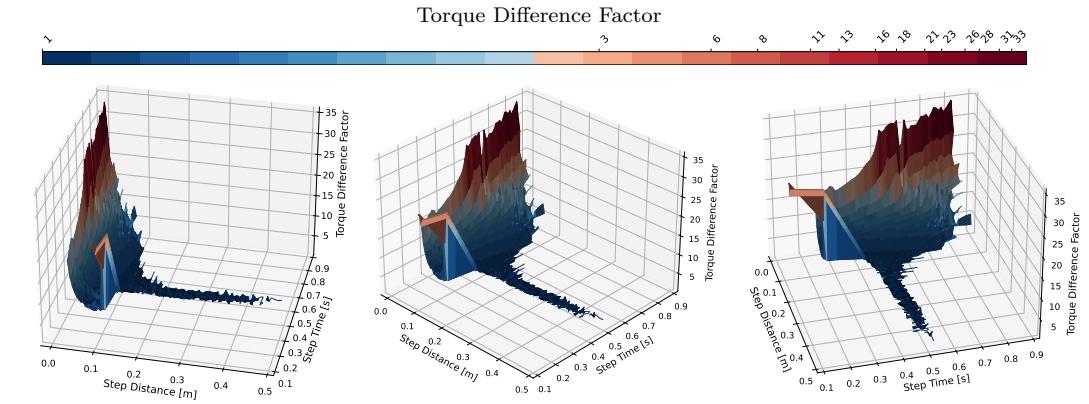

(b)

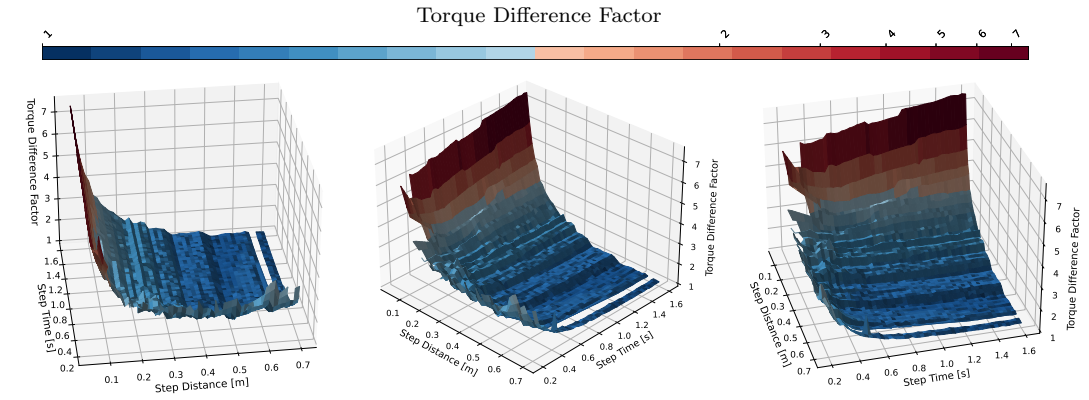

(c)

Figure C.4: Step Distance vs. Step Time vs. difference factor from global torque minimum for scenario 1(a), 2(b) and 3(c).

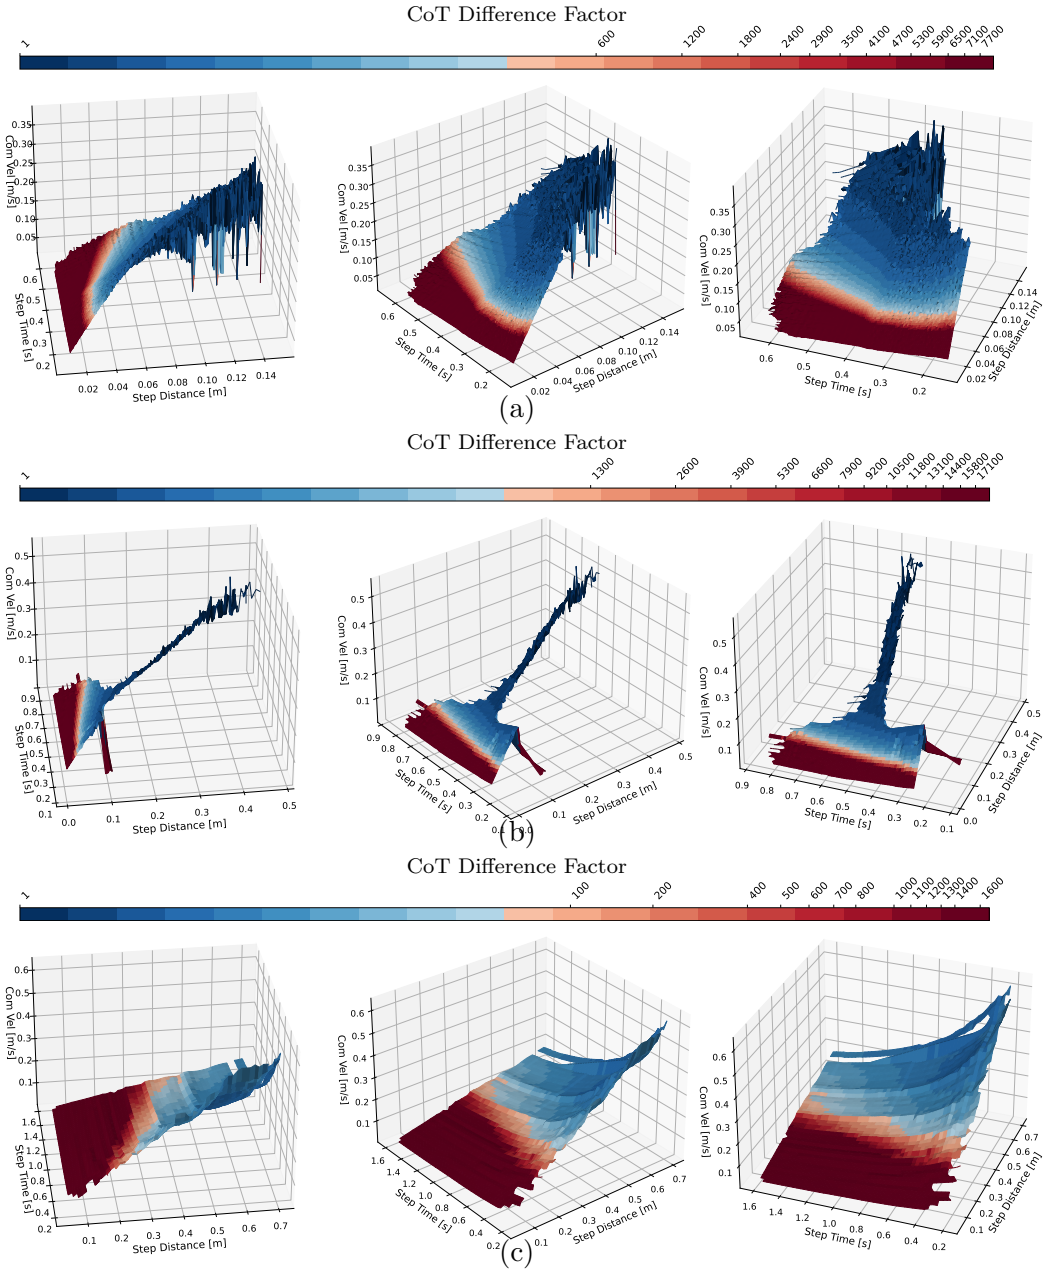

Figure C.5: Step distance vs. step time vs. CoM velocity, coloured by CoT values for scenarios 1(a), 2(b) and 3(c).

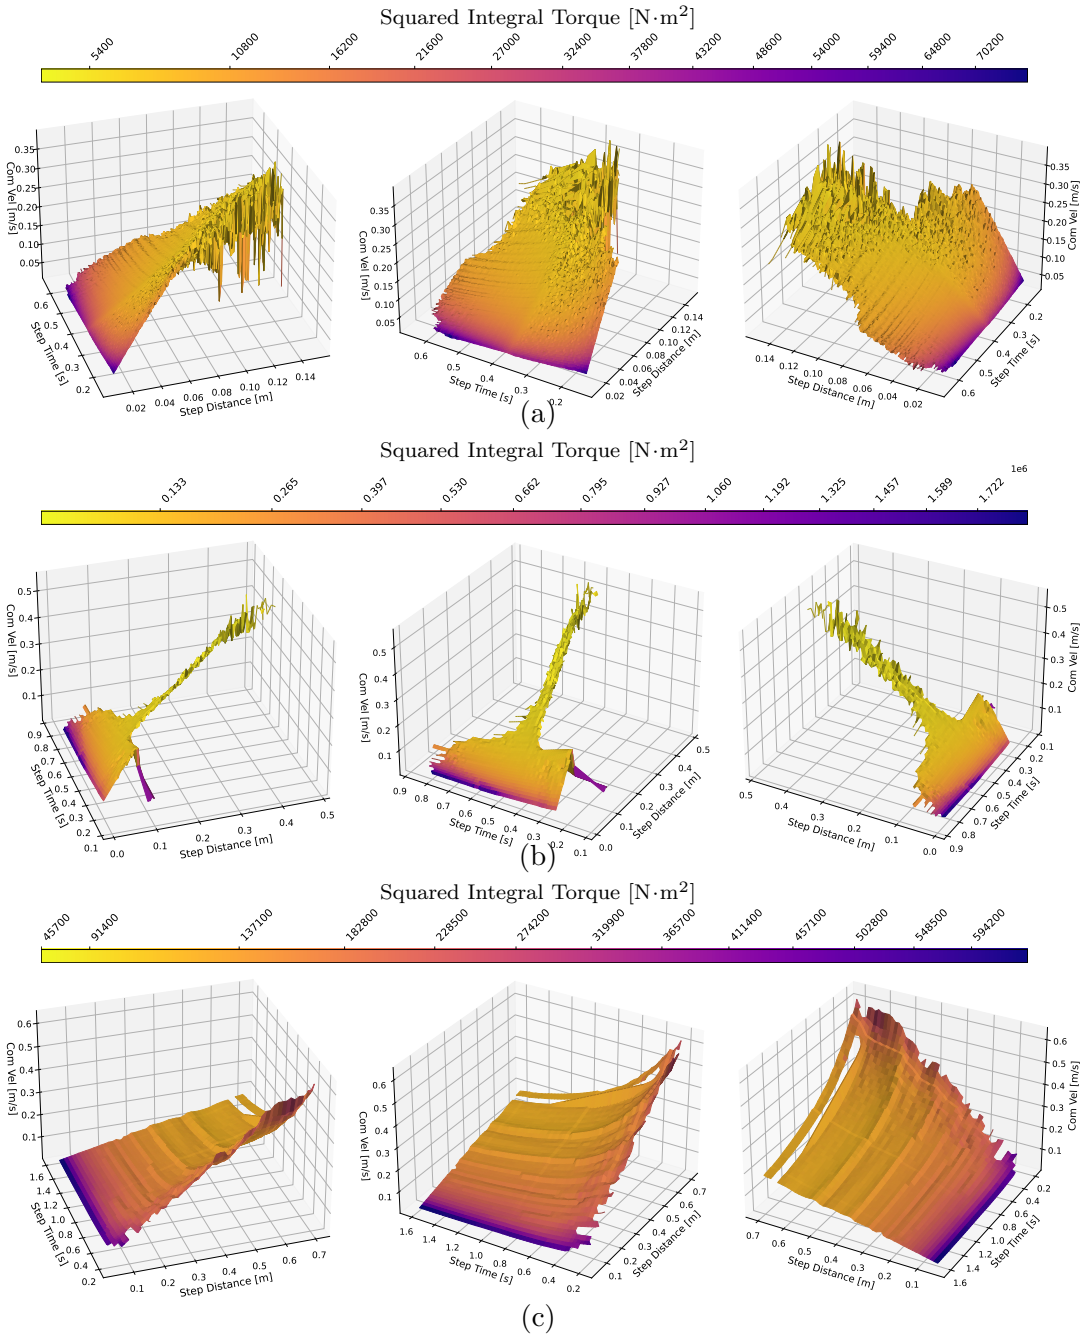

Figure C.6: Step distance vs. step time vs. CoM velocity, coloured by torque consumption values for scenarios 1(a), 2(b) and 3(c).

Appendix D

Human Participant Ethics Forms

Ethical Review Procedures: Level 1

Project Details & Self-assessment

This document is closely modelled on documents used in School of Philosophy, Psychology and Language Sciences provided by Ellen Bard and Cedric MacMartin.

This form is to be filled in and submitted at the same time as the project proposal or the funding application it applies to. The form should be submitted by the Principal Investigator, except in the following cases:

- Post-doctoral fellowships – the proposed postdoc mentor.
- UG, MSc, and PhD research projects – the supervisor.
- Visiting researcher – the staff hosting the visitor.

Please submit the completed form by email to: **infkm+ethics@inf.ed.ac.uk**

This address, with appropriate RT number once issued, should be used for all correspondence (including forms and attached documents). This is essential to ensure proper record keeping. No signature is required if the form is sent from a valid University email address.

Project Details

1 Type Of Project:

□ Research grant proposal	□ UG final year project	☐ MSc project
□ Post-doctoral fellowship	✓ PhD project	□ Research performed by visiting researcher
□ Personal research	□ Other: _____	

2 Is there a sponsor/ funding body? **YES** / NO

3 Does the sponsor/funder require formal prior ethical review? YES / **NO**
If yes, by what date is a response required ?

4 Is any other institution and/or ethics committee involved? YES / **NO**

5 Title of Project Transferring Human Push Recovery Behaviour to Robotics

6 Researchers' names, affiliations, emails

Christopher McGreavy – MSc Student, University of Edinburgh: c.mcgreavy@ed.ac.uk

Dr. Zhibin Li – supervisor, University of Edinburgh: zhibin.li@ed.ac.uk

7 State which professional organisation guidelines you are using:

X School of Informatics research ethics code: http://www.inf.ed.ac.uk/research/ethics/

□ Other ethics code as required by funding body or professional organization:
Title: _____ URL: _____

Part C 2

Self-assessment

Refer to Level 2 form for details on any of the following points.

1. **Protection of research participants' confidentiality**
   Are there any issues of CONFIDENTIALITY which are NOT ADEQUATELY HANDLED by normal tenets of academic confidentiality? YES / **NO**

   These include well-established sets of procedures that may be agreed more or less explicitly with collaborating individuals/organisations, for example, regarding:
   (a) Non-attribution of individual responses;
   (b) Individuals and organisations anonymised in publications and presentation;
   (c) Specific agreement with respondents regarding feedback to collaborators and publication.

2. **Data protection and consent**
   Are there any issues of DATA HANDLING AND CONSENT which are NOT ADEQUATELY DEALT WITH, and compliant with established procedures? YES / **NO**

   These include well-established sets of procedures, for example regarding:
   (a) Compliance with the University of Edinburgh's Data Protection procedures (see http://www.recordsmanagement.ed.ac.uk);
   (b) Respondents giving consent regarding the collection of personal data (via consent form).

3. **Significant potential for physical or psychological harm, discomfort or stress**
   Are there any risks of :
   (a) psychological harm or stress for the participants? YES / **NO**
   (b) physical harm or discomfort for the participants? **YES** / NO
   (c) any kind to the researcher? YES / **NO**

4. **Vulnerable participants**
   Are any of the participants in the research vulnerable, e.g., children, patients, disabled participants? YES / **NO**

5. **Moral issues and researcher/institutional conflicts of interest**
   Are there any SPECIAL MORAL ISSUES/CONFLICTS OF INTEREST? These include:
   (a) Conflict of interest: potential benefit to the researcher, friends or family of a particular research outcome which might compromise the researcher's objectivity or independence;
   (b) The need to keep the purposes of research concealed;
   (c) Use of participants who are unable to provide informed consent (e.g., children);
   (d) Situations where research findings would impinge negatively/differentially upon the interests of participants.

   YES / **NO**

6. **Bringing the University into disrepute**
   Is there any aspect of the proposed research which might bring the University into disrepute? For example, could any aspect of the research be considered controversial or prejudiced?
   YES / **NO**

7. **Use of animals**
   Does the research involve animals? YES / **NO**

8. **Developing countries**
   Does the research involve developing countries? YES / **NO**

9. Dual use
Is the research classified or does it have specific adversarial military applications? YES / **NO**

10. Terrorist or extremist groups
Does your research concern groups which may be construed as terrorist or extremist? YES / **NO**

Can you stop now?

You may want to assure yourself that your 'NO' answers are correct by checking the detailed form in the next section.

If all the YES / NO answers are NO, the self assessment has been conducted and confirms the ABSENCE OF REASONABLY FORESEEABLE ETHICAL RISKS. This form should be signed by the researchers and submitted. The researchers may retain a copy for their own records.

If any answer is YES, please complete the relevant section in the Level 2 form below.

Ethical Review Procedures: Level 2

Detailed Assessment

This material should help you answer the questions in the self-assessment form.
If any difficulties arise, you should fill in the relevant parts of this form in consultation with a near colleague who is not directly involved with the research. You can also seek advice from members of the School Ethics Panel, or from relevant Ethics Committees of other schools.
You should file a new form if you receive advice on changes from the School or College Ethics Committees. For accountability, the School will view the most recent submission as accurate.

1. Protection of research participants' confidentiality

Refer to the University Data Protection Policy to ensure that the relevant conditions relating to the processing of personal data under Schedule 2 and Schedule 3 are satisfied. Details are available at: http://www.recordsmanagement.ed.ac.uk.

1. If the research requires the collection of personal information from e.g., universities, schools, employers, or other agencies about individuals without their direct consent, what information will be sought and why will written consent for access to this information not be obtained from the participants themselves?

2. If any part of the research involving participants will be recorded using any electronic medium, what medium is to be used and how will the recordings be used?

   Data will be recorded using motion capture and force capture systems. Recorded data will be in the form of limb positions over the course of the trials as well as the forces exerted by the person onto the ground. Motion capture data is collected using optical markers attached to the subject. All data will be stored electronically in anonymous files on password protected computers with no traceability to the subject's name or other details, which will be held separately. Descriptive and analytical data will be extracted from the electronic data to be used to analyse statistically.

3. Who will have access to the raw data?

   Only the researchers named above will have access to the raw data from this study. No access will be given to any person not named as a researcher above. All data will be stored on password protected computers with only researchers having access to the passwords.

4. If participants will be identified in your records, how will their consent to quotations/identifications be sought?

   Participant data will be stored and used anonymously.

5. If they will not be identifiable, how will anonymity be preserved?

   Anonymity will be preserved by removing all identifiable information from subjects' trial data. Identifiable information will only be sought during data capture to keep a record of trial order and will be kept completely separate from the data. Subjects cannot be identified from their raw data due to the nature of the data collection being based on motion capture – limb motions during trials

cannot be traced back to an individual.

6. Will the datafiles/audio/video tapes, etc. be disposed of after the study?

   Datafiles will form part of a longer term study beyond this MSc project, into wider PhD research so will not be disposed of directly after this study is completed.

7. If not, how long they will be retained and how will they eventually be disposed of?

   One central copy of the datafiles will be held and once the studies that make use of them are completed, the files will be electronically deleted and removed completely from the device on which they are stored.

8. How do you intend the results of the research to be used?

   The results from this research will be used to as the basis of an MSc project. This MSc project will seek to transfer the data collected in this study into control software for humanoid robots. Potentially, publication will be sought for this study, at which point the results would be published in an aggregated form with no identifiable information included about any participant or their performance during trials.

9. If feedback of findings will be given to participants, how and when will this feedback be provided?

   Collated results can be presented to participants if requested. If this is required it will be provided in the form of a short report which will show the results overall and the implications of results for robotic systems. This can be given after results have been processsed and analysed - approximately 2-3 months after data collection.

2. Data protection and consent

Participants have the following rights over observations and records of their own behaviour:

- If they are engaging in any activity outside their normal daily routine (for example answering a questionnaire, listening for a particular syllable), they must be given some account of what they will be asked to do before they start, and must formally consent to participation;
- In any event, if they will be observed or recorded, they must be informed of and consent to the kinds of record taken;
- They must be assured of anonymity in any publication or dissemination;
- They must consent to how the data will be used;
- They must be free to withdraw from participation at any time.

1. Explain how and when written consent will be obtained from participants or from those responsible for participants unable to consent meaningfully on their own behalf. (If further discussion of this form is needed, please attach a copy of any information sheets and consent forms.)

   A participant information sheet and informed consent form will be supplied to the participants before the study commences. The participant information sheet will give details of the study and what will take place over the course of the trials. Trials will not commence until the informed consent form has been completed to indicate that the subject understands the procedure of the trials and what is expected of them. Only participants who can provide meaningful informed consent

will be provided will take part in this study.

2. If participants cannot meaningfully provide formal consent in this way, normally someone who is legally able to act on their behalf, for example a parent or legal guardian, must do so. If any of the following cases apply, explain how you will obtain the necessary consent and if you will not, how you can proceed ethically without doing so.

    - administrative consent in lieu of participants' consent

        (Administrative consent may be deemed sufficient:

        i. where the data collection involves aggregated statistical information and where the collection of data presents no invasion of privacy and no potential social or emotional risks:

        ii. where studies focus on the development and evaluation of curriculum materials, resources, guidelines, test items, or programme evaluations rather than the study, observation, and evaluation of individuals.)

    - the consent of parents on behalf of minors,
    - the consent or assent (at least verbal) of minors,
    - the consent of participants who do not share a language with the researcher,
    - the consent of participants with special educational needs.

Only participants who can provide meaningful informed consent will be provided will take part in this study.

3. Significant potential for physical or psychological harm, discomfort or stress

If the research could induce any psychological stress or discomfort, state the nature of the risk and what measures will be taken to deal with such problems.

If the research requires any physically invasive or potentially physically harmful procedures, give details and outline procedures to be put in place to deal with potential problems.

Experimental Procedure and Risks

Only one condition of the experiment carries a risk of discomfort, in which the researcher applies a small force to the participant at the back of the hip using a dome-shaped plastic end stop at the end of a rod. This is functional to administer a push to the subject to test how they recover. Pushes will not exceed a force which is able to make the subject take a single step to recover. Two elements of discomfort arise: the force of the push on the pelvis and the risk of falling.

Steps Taken to Minimise Risk

Falls can occur if the person fails to step forwards; the only feasible ways in which the subject would not make a step is if they were not anticipating the disturbance or if there was not enough room to step into. To reduce the possibility of unexpected pushes, a warning will be given to subjects that the trial is about to start within 10 seconds; this warning will be given before the start of each trial. Training trials will also be given at low forces to allow subjects to become used to the structure of the experiment and when to expect the disturbance.

To minimise the likelihood of falls, the participant will be fitted with a safety harness. The harness will be connected to a safety rope, which is also attached to the ceiling above the treadmill to arrest falls. The safety rope has been recently checked for its quality during routine treadmill maintenance by the supplier. The length of the rope is regulated by an adjustable PETZL GRILLON lanyard positioning system developed for safe working in high ropes industries (https://www.petzl.com/US/en/Professional/Lanyards-and-energy-absorbers/GRILLON) connected to a high strength carabiner which is then connected to the harness . The rope will be adjusted using the GRILLON to allow safe movement, but to limit movement outside this allowable trajectory. The rope will be attached to the subjects from the rear of the harness, so there will be no possibility of indirect injury from the rope by swinging around or becoming suddenly taught. Two sizes of harness will be available depending on the participant's height: small and large. The harness will be checked for fit before the participant steps onto the treadmill. Standard fitting procedures will be used, where all fastenings are closed tightly, but still allow space for the index and middle finger to fit between the fastening and the body. This ensures a good fit without limiting manoeuvrability and where sudden tightening of the harness does not affect blood flow.

Push forces will be tested gradually with the participant, where gradual forces will be administered from very light to the maximal push force for the experiment, which will still be relatively small. The discomfort from the device is almost imperceivable, but if participants experience discomfort, padding is made available to cover the lower back. Trials will only commence if participants are comfortable with the push force being applied. The experimenter will administer pushes only within the prearranged comfortable limits pre-arranged with the subject. This experiment is only interested in pushes of small magnitude, so contains an inbuilt tendency to administer small pushes to subjects.

These steps ensure that the participants will no encounter unnecessary risks during this study and will be as free of risk as possible during their trials.

If the research involves the investigation of any illegal behaviour, give details.

If there is a real risk of disclosure of activities which should be reported to the authorities, a warning to this effect must be included in the Informed Consent documents. Please provide the wording of this warning.

If there is any purpose to which the research findings could be put that could adversely affect participants, describe the potential risk for participants of this use of the data. Outline any steps that will be taken to protect participants.

If the research could adversely affect participants in any other way, give details and outline procedures to be put in place to deal with such problems.

If the research could adversely affect particular groups of people, describe these possible adverse effects and the protection to be put in place against them.

If the research is expected to benefit the participants, directly or indirectly, give details.

If the true purpose of the research will be concealed from the participants, explain what information

will be concealed and why.

If participants will NOT be debriefed at the conclusion of the study, explain why not.

4. Vulnerable participants

What criteria will be used in deciding on the inclusion and exclusion of participants in the study?

If any of the participants are likely to be in any of the following vulnerable categories, indicate the category and describe the measures that will be used to recruit, protect and/or inform participants:

under 16 years of age	in the care of a Local Authority
known to have special educational needs	physically or mentally ill
vulnerable in other ways	members of a vulnerable or stigmatized minority
unlikely to share a language with the researcher	in a student-teacher relationship with the researchers
in any other dependent relationship with the researchers	likely to have difficulty in reading and/or comprehending any printed material distributed as part of the study

If participants will receive any financial or other material benefits because of participation, what benefits will be offered to participants and why?

No vulnerable participants will be used in this study and no financial or material benefits will be provided as a result of participation.

5. Moral issues and researcher/institutional conflicts of interest

The University has a draft 'Policy on the Conflict of Interest'. Regarding research the draft states that a conflict of interest would arise in cases where an employee of the University might be

> " . . . compromising research objectivity or independence in return for financial or non-financial benefit for him/herself or for a relative or friend. . ."

The draft policy also states that the responsibility for avoiding a conflict of interest, in the first instance, lies with the individual, but that potential conflicts of interest should always be disclosed, normally to the line manager or Head of Department. Failure to disclose a conflict of interest or to cease involvement until the conflict has been resolved may result in disciplinary action and in serious cases could result in dismissal.

If your research involves a conflict of interest or any situation which could be construed as a conflict of interest, please give details.

No conflicts of interest will arise as a result of this study.

6. Bringing the University into disrepute

If on the level 1 form you have answered that some aspect of the proposed research "might bring the University into disrepute", please elaborate alongside how this might arise, and what steps will be taken by the researcher to mitigate and manage this, to minimise adverse consequences to the University.

7. Use of animals [based on EU FP7 guidelines]

If the proposed research will use animals, please provide the following information:

1. Describe how you have applied the 3Rs: Reduction, Replacement, Refinement.
2. Describe and justify:
   - species and numbers of animals used;
   - humane end points and pain and suffering;
3. Describe how you have explored alternatives to using animals.
4. Answer the following questions:
   - Are those animals transgenic small laboratory animals?
   - Are those animals transgenic farm animals?
   - Are those animals cloning farm animals?
   - Are those animals non-human primates?

8. Developing countries [based on EU FP7 guidelines]

Questions to consider include:

1. Does the research project provide benefit to the local community (in terms of access to healthcare, education, allocation of property rights, capacity to assess and use modern technologies while respecting the population 's own choices and needs, etc.)?
2. Does the research project use local resources (genetic resources, animals, and plants)?

How to deal with research involving developing countries

The categories of issues requiring special attention include:

- A disproportionately heavy burden of diseases (particularly infectious diseases); the breadth and depth of poverty; and high levels of illiteracy
- Wide disparities in health systems and in access to health care; and imbalance between the often-ample resources available for research and the meagre resources available for even basic health care
- Inadequate scientific and ethics infrastructures for the required reviewing process
- The extent of disempowerment of the poor in their personal and communal lives
- Knowledge of the ways in which people of other cultures traditionally view themselves as individuals embedded in communities with respect to the changing boundaries between perceptions of the self that differ from the classical western notion
- The need to understand what it means to be ill in contexts very different from those known to researchers and what can be expected from those one consults for help under those circumstances

9. Dual use [based on EU FP7 guidelines]

1) What is considered as potential dual use

Generally speaking, dual use is a term often used in politics and diplomacy to refer to technology which can be used for both peaceful aims and adversarial military aims. Ethical issues of dual use might arise in cases where:

(d) Classified information, materials or techniques are used in research;
(e) Dangerous or restricted materials e.g. explosives are used in research;
(f) The specific results of the research could present a danger to participants, or to society as a whole, if they were improperly disseminated.

2) How to deal with potential dual use

Regarding implications for the use of and misuse of the research and products, the following measures

and strategies can be applied:

(c) The researcher should show awareness of potential risks to participants and society as a whole from inappropriate dissemination of their results;
(d) Appropriate measures to deal with dangerous or restricted materials should be detailed, where applicable;
(e) An appropriate strategy to deal with issues of informed consent and risk management for participants and for society where classified information, materials or techniques are concerned should be demonstrated;
(f) An advisory board should be included in the project, which should identify risks to participants from particular research activities and devise a strategy for minimising and dealing with these risks;
(g) The dissemination and communication strategy of the study results to a wider audience should be controlled by the advisory board, which should report to the relevant funding body on a regular basis.

EU FP7 ethical guidelines can be found at http://cordis.europa.eu/fp7/ethics_en.html.

10. Terrorist or extremist groups

If your research concerns groups which may be construed as terrorist or extremist, please fill in the following form and submit it with your ethics form.

Prevent Duty supplementary form

The Terrorism Act (2006) outlaws the dissemination of records, statements and other documents that can be interpreted as promoting or endorsing terrorist acts.

1. Does your research involve the storage on a computer of any such records, statements or other documents? Yes / No
2. Might your research involve the electronic transmission (eg as an email attachment) of such records or statements? Yes / No
3. If you answered 'Yes' to questions 1 or 2, you are advised to store the relevant records or statements electronically on a secure university file store. The same applies to paper documents with the same sort of content. These should be scanned and uploaded. Access to this file store will be protected by a password unique to you and your School Research Ethics Officer. Please indicate below that you agree to store all documents relevant to questions 1 and 2 on that file store: Yes

3a. Please indicate below that you agree not to transmit electronically to any third party documents in the document store: Yes

4. Will your research involve visits to websites that might be associated with extreme, or terrorist, organisations? Yes / No
5. If you answer 'Yes' to question 4, you are advised that such sites may be subject to surveillance by the police. Accessing those sites from university IP addresses might lead to police enquiries. Please acknowledge that you understand this risk by putting an 'X' in the 'Yes' box.

   Yes
6. By submitting to the ethics process, you accept that your School Research Ethics Officer and the convenor of the University's Compliance Group will have access to a list of titles of documents (but not the contents of documents) in your document store. Please acknowledge that you accept this by putting an 'X' in the 'Yes' box. Yes

Countersigned by supervisor/manager

1/5/18

Figure D.1: Ethics self-assessment form used for human study - approved by the ethics group at University of Edinburgh School of Informatics

THE UNIVERSITY of EDINBURGH
informatics

Clinical and Experimental Form

Study title: Transferring Human Push Recovery Behaviour to Robotics

To be completed by the researcher:

Subject Number	
DoB	
Mass	
Height	
Date	
Time	

To be completed by the participant:

Given name	
Family name	
Phone number	
E-mail address	

School of Informatics, 10 Crichton Street, Edinburgh, (City of) Edinburgh, EH8 9AB

Figure D.2: Ethics clinical and experimental form for collecting participant data. Data was not linked to participant data, was used only for applicable purposes

THE UNIVERSITY of EDINBURGH informatics

Participant Information Sheet

Study title: Transferring Human Push Recovery Behaviour to Robotics

1. Aim of the study

This study aims to improve the ability of robots to recover from disturbances by external forces using data collected from humans. By applying forces to humans to see how they react, it may be possible to extract their behaviour in order to transfer onto robotic systems. The data which will be collected as part of the study will be used to develop a controller for a robot to recover from disturbances when pushed so as to not fall over.

2. Execution of the study

The data collection sessions will be conducted by researchers trained for this purpose. You will be asked to complete 3 conditions:

Standing: stand perfectly still to measure a baseline stand position

Swaying: voluntary backwards and forwards movements without taking a step

Pushing: You will receive a small push at the back of the pelvis with a force measuring device

In all conditions you will be asked to stand perfectly still at the start, you will then receive instructions as to which condition you will carry out. The movement, whether voluntary or involuntary will be carried out. You will recover your balance and return to a standing position. The trial will the end. This process is repeated until all trials have been completed.

To measure your joint movements accurately, you will be required to wear small reflective markers (see Figure 1) which will be attached to the skin at preset points using medical tape. Before testing begins, an acclimitisation period will take place to find push forces which you are comfortable receiving. The contact by the pushing device is not painful, but padding is available if required. No forces outside the agreed comfortable limits will be applied.

Figure 1: Participant wearing reflective markers3. Your participation

The data acquisition session will last for approximately 30 minutes and will take place in room G.06 of the Bayes Centre, University of Edinburgh. Participation in this study is entirely voluntary. You can refuse to take part or withdraw from the study at any time without having to give a reason. Such a decision has no adverse implications for you.

4. Risk assessment

Your participation to this study involves risk that are as low as reasonably possible. You will be wearing a harness at all times, this will ensure that you cannot fall.

5. Privacy

All data acquired will be treated confidentially. The data might be disclosed anonymously to third parties for the purposes of the study. Your personal information will be stored separately to ensure data protection.

6. Contact

If you have any questions or require further information, please do not hesitate to contact Mr. Christopher McGreavy (c.mcgreavy@ed.ac.uk).

Figure D.3: Ethics Participant Information Sheet

Informed Consent Form

Study title: Transferring Human Push Recovery Behaviour to Robotics

1. I confirm that I have read and understood the Participant Information Sheet for the above study and there is no reason I should not take part. I have had the opportunity to consider the information and ask questions, and have had these answered satisfactorily.
2. I understand that my participation is entirely voluntary and I am free to withdraw at any time without giving a reason.
3. I certify that I have been informed that the data collected during the study will be shared with the scientific community in respect of anonymity, only researchers directly involved with the data acquisition and storage will have direct knowledge of my identity, and they will be bound by professional secrecy.
4. I understand that data collected in this study will be used for the project named above and will also form a part of longer term studies into human movement. Data will be disposed of once these studies are complete.
5. I have been made aware of the risk assessment carried out and the safety measure that are in place to minimize risk.
6. I agree to take part in this study.

.....	.....	.....
Name of participant	Date	Signature

.....	.....	.....
Name of researcher	Date	Signature

School of Informatics, 10 Crichton Street, Edinburgh, (City of) Edinburgh, EH8 9AB

Figure D.4: Ethics consent form signed by all participants after reading Participant Information Sheet (Figure D.3)