THE UNIVERSITY
of EDINBURGH

A NOVEL APPROACH FOR REPRESENTING, GENERALISING, AND QUANTIFYING PERIODIC GAITS

HSIU-CHIN LIN

Doctor of Philosophy

Institute of Perception, Action, and Behaviour

School of Informatics

University of Edinburgh

2015

ABSTRACT

Our goal is to introduce a novel method for representing, generalising, and comparing gaits; particularly, walking gait. Human walking gaits are a result of complex, interdependent factors that include variations resulting from embodiments, environment and tasks, making techniques that use average template frameworks suboptimal for systematic analysis or corrective interventions. The proposed work aims to devise methodologies for being able to represent gaits and gait transitions such that optimal policies that eliminate the inter-personal variations from tasks and embodiment may be recovered.

Our approach is built upon (i) work in the domain of null-space policy recovery and (ii) previous work in generalisation for point-to-point movements. The problem is formalised using a walking phase model, and the null-space learning method is used to generalise a consistent policy from multiple observations with rich variations. Once recovered, the underlying policies (mapped to different gait phases) can serve as reference guideline to quantify and identify pathological gaits while being robust against interpersonal and task variations.

To validate our methods, we have demonstrated robustness of our method with simulated sagittal 2-link gait data with multiple ground truth constraints and policies. Pathological gait identification was then tested on real-world human gait data with induced gait abnormality, with the proposed method showing significant robustness to variations in speed and embodiment compared to template based methods. Future work will extend this to kinetic features and higher degree-of-freedom.

ACKNOWLEDGMENTS

There are a number of people whom I wish to thank for their support, help and advice when conducting this research and writing this thesis.

First, I would like to thank my first supervisor, Sethu Vijayakumar, for inspiring me into the world of statistical learning for rehabilitation robotics, for offering me the opportunity to conduct this research, and for all of the invaluable support and guidance he has given me over the past four years.

Second, I would like to offer my special thanks to Matthew Howard for his close academic support in my research, for his feedback on the algorithms and experiments, and for his patience with correcting my writing.

Third, I would like to thank my fellows in the SLMC group that have made working and studying such a enjoyable experience. Also, I would like to thank the participants of my experiments for their contributions to my research.

I cannot list, here, all my friends, for their encurangement and emotional support. In particular, their help with my uncountable struggles of academic work and daily life throughout the last phase of my PhD time will not be forgotten.

Last, but very certainly not least, I owe a huge debt to my family in Taiwan. Throughout the years away from home, they give me the freedom and motivation to pursue my interests. Finally, I would like to thank my mother for forgiving my physical absence during her difficult times.

Again, many thanks to everyone mentioned above. I am lucky to have you all, this thesis would certainly not have been possible without any of you.

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.

Hsiu-Chin Lin

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ACRONYMS

DPL	Direct Policy Learning
NPL	Null-space Policy Learning
UPE	Normalised Unconstrained Policy Error
CPE	Normalised Constrained Policy Error
APD	Approximated Policy Difference
CoM	Centre-of-Mass
CoG	Centre-of-Gravity
BoS	Base-of-Support
GRF	Ground-Reaction-Force
CoP	Centre-of-Pressure
NPOE	Normalised Projected Observation Error
NPPE	Normalised Projected Policy Error
RBF	Radial Basis Function
ZMP	Zero-Moment-Point
EMG	Electromyography

LIST OF SYMBOLS

Below is a list of symbols used throughout this thesis. Note that we use bold upper-case letters to denote matrices, bold lower-case letters to denote vectors, and normal letters to denote scalars. Notations of the form f(.) denote an argument should be passed to the function f.

$\mathbf{x}$	Observed state
$\mathbf{u}$	Observed action
$\mathbf{u}^{\text{ns}}$	Null-space component
$\mathbf{u}^{\text{ts}}$	Task-space component
$\boldsymbol{\pi}(.)$	Null-space policy
$\mathbf{b}(.)$	Task-space policy
$\mathbf{N}(.)$	Null-space projection matrix
$\mathbf{P}(.)$	A projection matrix which projects a vector onto the null-space component
$\mathbf{A}(.)$	Pfaffian constraint matrix
$\mathbf{J}(.)$	Jacobian matrix
$\mathbf{I}$	Identity matrix
t	Time
$\mathbf{q},\ \dot{\mathbf{q}},\ \ddot{\mathbf{q}}$	Angle, velocity, and acceleration of the joint-space
$\mathbf{r},\ \dot{\mathbf{r}},\ \ddot{\mathbf{r}}$	Position, velocity, and acceleration of the End-effector space
$\boldsymbol{\rho},\ \dot{\boldsymbol{\rho}},\ \ddot{\boldsymbol{\rho}}$	Position, velocity, and acceleration of the task-space
$\boldsymbol{\tau}$	Joint torques

$\hat{\mathbf{n}}$	Constraint vector
$\boldsymbol{\beta}$	Basis function
$\mathbf{W}$	Weight Matrix
$\mathbf{K}(.)$	Gaussian Kernel
$\mathcal{N}$	Number of data points
$\mathcal{K}$	Number of trajectories
$\mathcal{T}$	Number of time steps
$\mathcal{M}$	Dimensionality of the basis functions
$\mathcal{P}$	Dimensionality of the state space
$\mathcal{Q}$	Dimensionality of the action space
$\mathcal{S}$	Dimensionality of the task-space
$\mathbb{R}$	Real numbers
$\mathbf{R}(.)$	Rotation matrix. Arguments denote the rotation angles
$\mathbf{G}(.)$	Givens rotation matrix. Arguments denote the rotation angles and axes
$\mathrm{E}(.)$	Objective function. Arguments denote the quantity to be optimised.
$\mathrm{U}(.)$	Uniform distribution. Arguments specify the range of the distribution
$\tilde{\mathbf{f}}(.)$	An estimation of $\mathbf{f}(.)$. e.g., $\tilde{\mathbf{f}}(\mathbf{x})$ is the estimate of $\mathbf{f}(\mathbf{x})$ at point $\mathbf{x}$
$\mathbf{f}^*(.)$	The optimal estimation of $\mathbf{f}(.)$
$\hat{\mathbf{a}}$	The direction of vector $\mathbf{a}$
$\|\|\mathbf{a}\|\|$	The magnitute of vector $\mathbf{a}$
$vec(\mathbf{A})$	The vector operation applied to matrix $\mathbf{A}$
$\mathbf{A}^\dagger(.)$	Moore-Penrose pseudo inverse of matrix $\mathbf{A}$
$\mathbf{A}^\top(.)$	Transpose of matrix $\mathbf{A}$

$\int$	Integral
$\prod$	Product
$\otimes$	Tensor product
∂	Partial derivative

1

INTRODUCTION

> *"since man took his first steps, no one has asked himself why he walks, how he walks, if he has ever walked, if he could walk better, what he achieves in walking..."*
>
> *— Honoré de Balzac, Theory of Walking, 1978*[1]

Many everyday human skills can be considered as a form of periodic movement. For example, locomotion can be considered as a periodic motion of the legs, and wiping a table can be a kind of periodic motion of the arms. The form of these movements is influenced by many factors, namely, (i) the embodiment of the subject (e.g., limb lengths, mass properties, etc.), (ii) the environment in which the behaviour is performed (e.g., is the subject walking on flat or uneven terrain?) and (iii) task contextual factors (e.g., is the subject hurrying to a meeting or just taking a walk in the park?). Nevertheless, despite these variations, some consistency appears that causes us to identify behaviours, such as walking, as belonging to the same class.

The fact that such variations exist within a single class of behaviour, such as walking, indicates the presence of redundancy in the system. That is, the presence of additional degrees of freedom allow the constraints induced by these various factors to be satisfied, while at the same time satisfying some underlying consistent behavioural goal. The latter could be to minimise effort, maintain comfort, or other such criteria. This dependence on the various factors makes modelling human gait hard in general (Multon et al., 1999; Boulic et al., 1990), especially given the fact that the precise influence of different factors on the movement can be hard to assess (e.g., how is the foot placement strategy affected by differences in terrain (Pongas et al., 2007; Fukuoka et al., 2003)).

The motivation behind this thesis is to proposed a method in representing, generalising, and comparing periodic movement that are subject to unknown variations; particularly, we are interested in walking gaits. Potential applications include robot-assisted

1 Honoré de Balzac, Théorie de la démarche (Theory of Walking), Paris: Pandora, 1978

rehabilitation and clinical gait analysis. For example, what are the properties and attributes of *normal walking*, so that we can quantify the degree of "pathology" of a mobility-impaired patient and assist the patient through use of robot technology.

Examples in modelling/comparing human gaits can be found in exoskeleton systems such as the Lokomat, the Skywalker, and the Lower Extremity Powered Exoskeleton (LOPES) (Duschau-Wicke et al., 2010; Artemiadis and Krebs, 2011; Veneman et al., 2007). Although these devices are cleverly designed, much work is needed to personalise gait correction. One major challenge of designing wearable devices is when and how much to control the user. Assisting the patients through a predefined reference gait with a feedback controller seems the most prominent approach in gait rehabilitation (Krebs et al., 1998; Lum et al., 1993).

reference gait

This reference gait is normally obtained by taking the average or polynomial fit of some data gathered from healthy subjects. Note that, actively controlling the patients to follow the reference trajectory would restrict the patients to walk at certain speed, slope, or step size, but this seems inadequate. For instance, this approach would consider faster or slower walks as deviations from a normal gait. However, clinical results show that motivating the patients to walk more proactively at a preferred pace promotes the overall results of rehabilitation (Hidler et al., 2005).

Some existing devices have a more flexible approach by incorporating impedance control and/or tolerating small deviation from the reference gait (Jezernik et al., 2004; Veneman et al., 2006; Riener et al., 2005). However, taking the average template framework as the reference might be suboptimal since some joints or dimensions may need no correction. Instead of restricting the patients to follow these predefined rules of training, a more appropriate way is to let the patients walk the way they prefer and correct them only if needed. For us, an ideal reference gait should represent the fundamental components of normal gaits and that is independent from natural variations (from embodiments, environments, and behaviours). On the other hand, "the important aspect of normal walking" has no clear definition.

Previous research in gait analysis have shown very small variations in human kinematics when performing the same gait. For example, (Stokes et al., 1989) evaluated the repeatability of kinematic data of 40 subjects. Although there was variance in the range of motion, the kinematic patterns in the saggital plane were highly predictable. (Ivanenko et al., 2004) applied dimensional reduction analysis to determine whether the pattern of muscle activities can be described by some underlying manifold. Their work showed that while the Electromyography (EMG) signals varied significantly when walk-

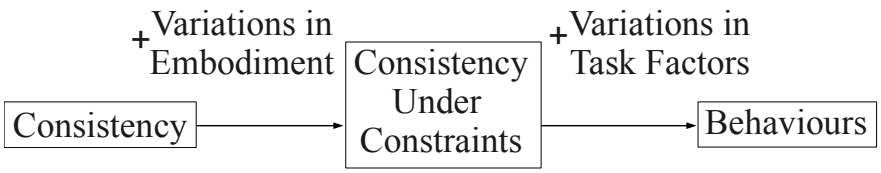

Figure 1.1: We hypothesise that the behaviours that we observe are the combinations of some consistent characteristics and variations in embodiment and task factors. Our approach is based on examining various behaviours to see if such consistency can be found.

ing at different speeds, using Principle Component Analysis (Jolliffe, 2005), just the first five principle components of these EMGs could account for the main features of the signals.

From the results of gait-analysis discussed above, *we hypothesise that the behaviours that we observe are the combinations of some consistent characteristics and variations in embodiment and task factors.* For instance, we normally prefer walking in a comfortable way (the consistency) regardless of our walking speeds (the variations). As mentioned before, we are seeking the important aspect of walking for gait-rehabilitation. Our idea is to take out the consistency of normal walking while eliminating the variations as much as possible.

In recent years, a number of new tools have become available in the learning and robotics community that allow data from constrained and/or redundant systems (Howard et al., 2009; Towell et al., 2010) to be used to uncover underlying consistent behaviours that may be otherwise masked by the constraints. These techniques are proven to be effective in extracting the consistency in robotic task. For example, reaching some specific targets (the variations) while moving to a default position (the consistency).

We recognised an analogy between reaching task of robots and human walking, for both problems aim at extracting the consistent components from some observations. Our approach is based on examining various walking behaviours in the light of such methods, to see if certain underlying characteristics of walking can be found (see Fig. 1.1). These may be intra-personal (e.g., one person walks with a particular style that maintains that person's most comfortable posture) or inter-personal (e.g., people normally walks to minimise some measure of effort).

We narrow down our problem to a single class of human locomotion such that only walking gait on even terrain is considered. In this thesis, we examine various walking behaviours within this class, and we proposed a statistical learning method for

extracting the consistency. If we can model this consistency, we suggest a subsequent technique to measure the abnormality of a walking gait. An extensive set of experiments are reported in order to validate the approach and to evaluate the performance of the framework developed. For these, the proposed techniques are performed on both simulated data and human walking data subject to variations in embodiments and behaviours.

1.1 THESIS OUTLINE

In the following, we give a brief outline of the thesis. For each chapter, we outline the key content and highlight the original contributions made in that chapter.

In **Chapter 2**, we review the state of the art in learning and modelling walking gaits; specifically, we focus on related work in gait rehabilitation, gait analysis, and statistical learning for dealing with walking gaits.

In **Chapter 3**, we describe our approach on the representations of walking gaits. In particular, we assume that the walking gait we observe is the combination of the *consistent characteristics* of walking and *variations* from environment, embodiments, and behaviours. An effective way for representing walking gait is to model the consistent characteristics as an *unconstrained policy*. *representation*

Original Contributions:

- A constrained tracking control scheme such that the observations can be decomposed into the *consistent characteristics* of walking and *variations*.
- Numerous examples of representations and control schemes for the walking phase model provided, including examples from position control, end-effector space control, and balance control.

In **Chapter 4**, we propose a novel method for *generalising the consistent characteristics of walking gaits* subject to variations in environment, embodiments, and behaviours. This is achieved by reconstructing an *unconstrained policy* of the walking phase model (described in Chapter 3) which captures the characteristics of walking. We show that our proposed method can effectively recover the policy *without explicit knowledge of the variations*. *generalisation*

Original Contributions:

- Analysis of how variations in walking affect observed movements from the viewpoint of learning

- Novel method to approximate the *consistent characteristics* of walking gait based on optimising consistency of an *unconstrained policy* with the movements subject to variations. Numerous experiments presented, validating the proposed method with various systems and policies.

- Evidence for the feasibility of learning the *unconstrained policy* based on geometric analysis.

In **Chapter 5**, we propose a procedure to *quantify the distance between walking gaits* which ignores the differences coming from variations. Specifically, we discuss our approach that compares the *constrained policies* between a walking gait and a reference gait without explicitly knowing the variations. (e.g., the walking gait in question and the reference gait can be a mobility-impaired patient and the healthy subjects, respectively.) *quantification*

Original Contributions:

- Novel approach to measure the distance between gaits by comparing the distance between two *constrained policies*, which can be applied to measure the difference between the *characteristics* of two walking gaits. Numerous experiments presented, validating this approach for both stationary and non-stationary systems with various policies

- Novel learning method developed to estimate the *projection matrix* in a generic way, which can be applied to both stationary and non-stationary constraint systems.

In **Chapter 6**, we explore the utility of our approach on human walking data. Our analysis is based on kinematic and kinetic features of subjects walking with various speeds. Our goal is to see whether we can (1) extract consistency across walking behaviours and subjects and (2) use the extracted gait to quantify the difference between normal and pathological gaits. Experimental results will show that our method is more robust than the direct approach such as direct regression.

Original Contributions:

- Numerous experiments presented, validating our approach for generalising and quantifying walking gaits on kinematic and kinetic features of human walking data

Lastly, in **Chapter 7**, we give final conclusions and propose directions for future work.

1.2 PUBLICATION SUMMARY

This thesis provides a presentation of some materials which have been published or submitted. The essential concepts of our work have been published in Lin et al. (2014a). (This includes parts of the walking phase model in Chapter 3, the generalisation algorithms in Chapter 4, the quantification approach in Chapter 5, and the experiment with kinematic features in Section 6.3.) Lin et al. (2014b) is an extension of the previous publication where formal proof and validations discussed in Section 4.3 and Section 4.4 were included. Lin et al. (2015) discusses the quantification approach in higher degree-of-freedom based on Chapter 5. Finally, we aim to submit the experiment on kinetic features of human walking based on Chapter 6.4.

Publication :

- Lin, H., Howard, M., and Vijayakumar, S. (2014). A novel approach for generalising walking gaits across subjects and walking speeds. In *Proceedings of International Conference on Biomedical Robotics and Biomechatronics*, pages 1009-1015.

- Lin, H., Howard, M., and Vijayakumar, S. (2014). A novel approach for representing and generalising periodic gaits. *Robotica*, 32 (08), pages 1225-1244.

- Lin, H., Howard, M., and Vijayakumar, S. (2015). Learning null-space projection. In *Proceedings of International Conference on Robotics and Automation*, pages 2613-2619.

2

BACKGROUND

"Now man, being a biped and making his change of position in the natural way with his two legs, bends them forward for the reasons set forth"

—Aristotle, De Motu Animalium (350 B.C.E) [1]
translated by A. S. L. Farquharson

This chapter provides some relevant work and basic concepts in the fields of gait analysis, bipedal locomotion, robot-assisted rehabilitation, and statistical learning for motor control. In particular, we begin by summarising the definition and clinical results in gait analysis in Section 2.1. We also discuss three different control mechanisms of bipedal walking robots in Section 2.2 followed by examples of lower-limb rehabilitation devices in Section 2.3. Finally, we discuss the related work on learning control policies and operational-space formulation in Section 2.4.

2.1 GAIT ANALYSIS

Although walking is a familiar, everyday activity, it is hard to formally define the walking gait. The field of gait analysis has been attempting to uncover the nature of walking. In this section, we summarise some definitions of walking as well as results from quantitative gait analysis.

2.1.1 *Phase-based Decomposition of Walking*

A common approach in the gait analysis literature (Whittle, 2003), and one that we will follow in this thesis, is to decompose walking into a series of *gait cycles* where *gait cycle*

1 Aristotle (384-322 BC) was the first person who studied human gaits, and the author of the earliest book regarding to human walking, *De Motu Animalium* (On the Gait of Animals).

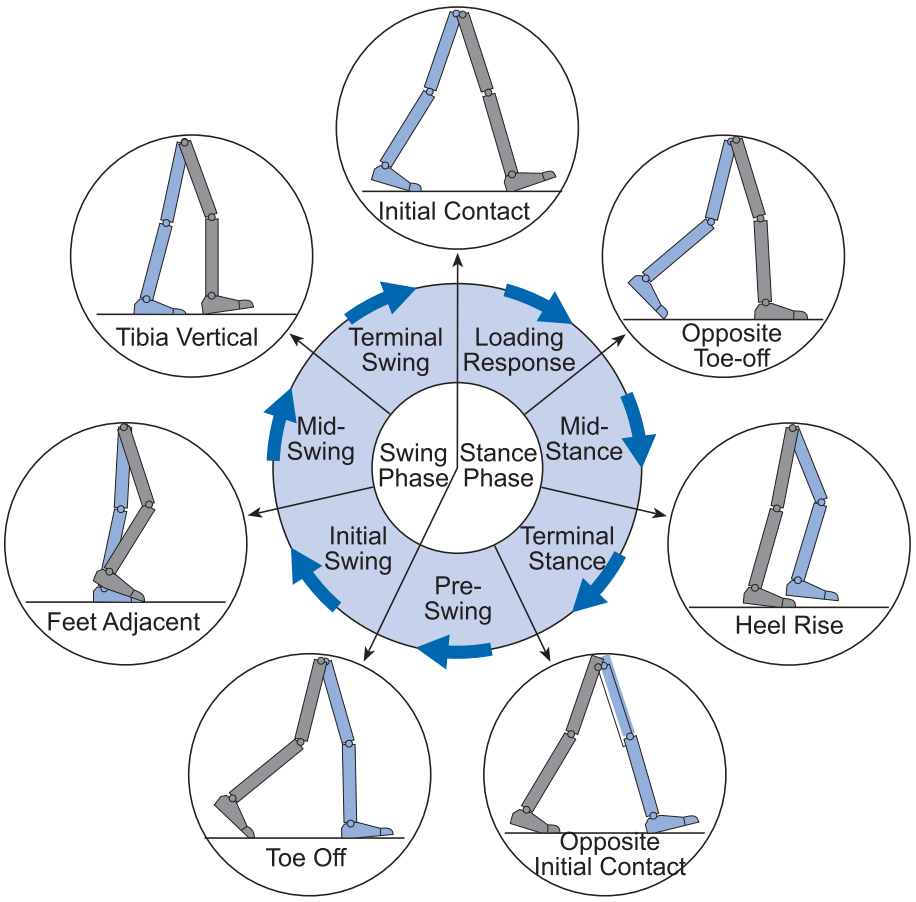

Figure 2.1: Decomposition of a gait cycle by the right leg (gray). The inner circle shows the gait phases and the outer circle shows the gait events. Figures were extracted and modified from Whittle (2007)

one walking cycle is defined as the time between two consecutive occurrences of an event. In gait analysis, the instant at which one heel strikes the ground (Initial Contact) is often demarcated as the beginning of a cycle, that continues until the same heel strikes the ground again.

A gait cycle is divided into the *stance phase* and the *swing phase*. The stance phase is the time interval where the foot is in contact with the walking surface, which covers approximately 60 percent of the cycle. The swing phase is the time interval while the leg is swinging in the air, and covers the remaining 40 percent.

stance/swing

A gait cycle can be further divided into smaller *gait phases* according to special *gait events* . Fig. 2.1 shows an example of a gait cycle by the right leg (gray). The inner circle shows the gait phases and the outer circle shows the gait events. (Figures were extracted and modified from Whittle (2007).)

gait phases

gait events

Note that, in gait analysis literature, the descriptions of walking gait are often described in terms of gait events or gait phases, where the descriptions might include kinematic features (e.g., joint-angles, joint-velocities, joint-accelerations), kinetic features (e.g., joint-moments, joint-powers, joint-works), or electrical activity (e.g., EMG). To further understand normal locomotion, in the following, we briefly discuss some interesting properties of each gait event.

kinematics

kinetics

Initial Contact (IC)

Also known as the ”heel-strike”. The knee angle extends rapidly during swing phase and is nearly fully extended at the end of swing phase. At initial contact, the knee angle begins to flex (referred to as ”stance phase knee flexion”), and its magnitude is highly dependent on walking speeds. In the mean while, the hip angle begins to increase and the ankle begins to flex so the foot is lowered to the ground.

Opposite Toe off (OT)

When the opposite leg leaves the ground, the knee reaches the peak of ”stance phase knee flexion” and begins to extend again. The amount of flexion is normally between 10° to 20°, but varies across subjects and speeds. Meanwhile, the foot is normally flat on the ground, the ankle reaches the peak of plantar-flexion and begins to flex.

Heel rise (HR)

or ’Heel off’. The time at which the heel begins to leave the walking surface, the timing varies between individuals and with walking speed. Knee extension reaches its maximum around the same time.

Opposite initial contact (OI)

The heel of the opposite leg contacts the ground. Normally, hip angle reaches its maximum extension at this point. Both hip and knee angle begin to flex which causes the leg to rotate forward about the forefoot.

Toe off (TO)

The toe leaves the ground at which point the stance phase ends and the swing phase begins.

Feet adjacent (FA)

The time at which the leg passes the opposite leg, two feet are about side by side,

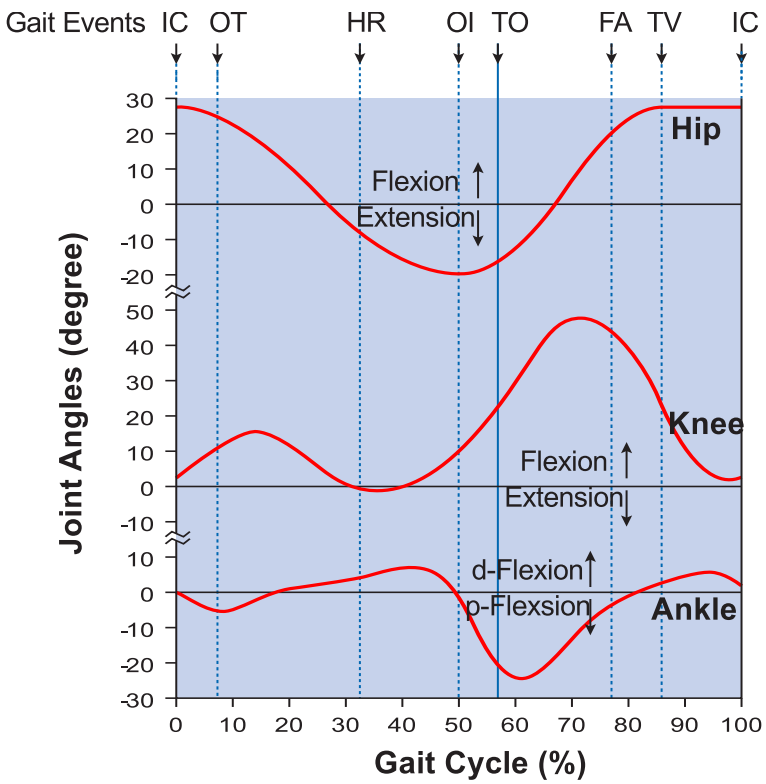

Figure 2.2: An example of joint angles in a gait cycle extracted and modified from (Whittle, 2007). The figures from the top to the bottom show the hip, knee, and ankle trajectories, respectively.

and occurs approximately in the middle of the swing phase or of the gait cycle. The knee angle reaches its peak flexion and starts to extend again. The amount of flexion is affected by walking speeds, and faster walks typically have less flexion and vice versa.

Tibia vertical (TV)

The time at which the tibia is perpendicular to the ground. This event marks the time at which hip angle stops flexion.

Note that, these gait phases and gait events were widely studied in the gait analysis community since they characterise a switch or change in kinematics and/or kinetics. Fig. 2.2 shows an example of hip, knee, and ankle within a gait cycle where the vertical lines are the gait events. Figures were extracted and modified from (Whittle, 2007).

When a person is walking with his/her own preferred speeds, the gait events discussed above occurred approximately at 7%, 15%, 32%, 50%, 60%, 77%, and 86% of the gait cycle; however, this time varies with the walking speed. For instance, when

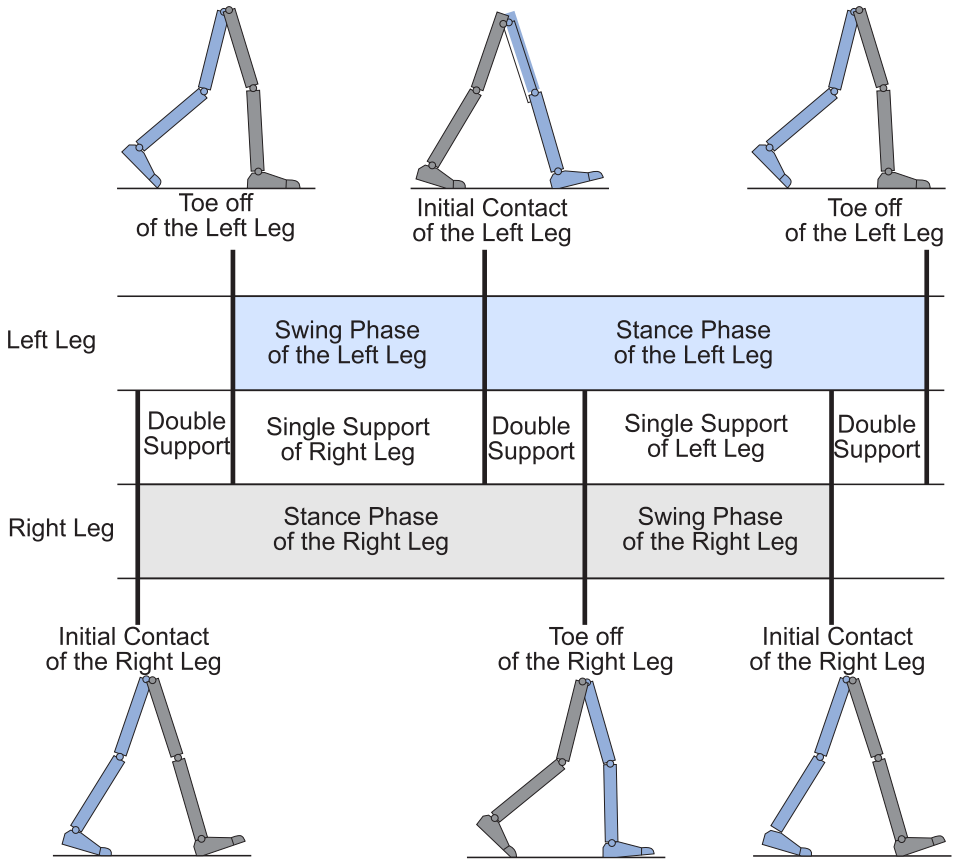

Figure 2.3: Single Support and Double Support Phases

walking speed slows, stance time increases, while the swing time remains constant (Herman, 1976).

Approximately, the first 10% and the last 10% of the stance phase are spent in the *double support phases* (when both feet are in contact with the ground), while the rest is in the *single support phases* (only one foot is in contact with the ground). This is also specified by the gait events, as shown in Fig. 2.3.

double support

single support

2.1.2 *Balance*

Winter (1995) defines *Balance*, as a generic way to describe the posture or motion that is free from falling. Although it is not clear how such posture is maintained, it is assumed that our balance system passively controls the Centre-of-Mass (CoM). By

centre-of-mass

definition, the CoM is a point equivalent to the weighted average of the total body mass and body segments. For an N-segment body, the CoM is calculated by

$$CoM = \sum_{i=1}^{N} \frac{m_i CoM_i}{M}$$

where m_i is the mass of the i^{th} segment, CoM_i is the centre of mass of the i^{th} segment, and M is the total body mass. The segment centre of mass CoM_i is a point where the weighted relative position of the distributed mass sums to zero. Assuming that the i^{th} segment can be divided into J sections such that $m_{i,j}$ is the mass of the j^{th} section of the i^{th} segment, and $x_{i,j}$ is the distance from the j^{th} section to the edge of that segment, the CoM_i is defined as

$$CoM_i = \frac{1}{m_i} \sum_{j=1}^{J} m_{i,j} x_{i,j}$$

centre-of-gravity
base-of-support
CoM displacement

An underlying assumption is, the *dynamical stability* of walking is the ability to control the Centre-of-Gravity (CoG) in relationship to the Base-of-Support (BoS). The CoG refers to the vertical projection of the CoM on the surface of walking. The gait is balanced if the CoG is close enough to the position of BoS, which is the area of the body in contact with the supporting surface. Therefore, the CoM displacement, or the distance between CoG and the BoS, is highly correlated to the balance of the gait (Shumway-Cook and Woollacott, 1995).

2.1.3 *Repeatability in Normal Walking*

For decades, researchers have been questioning whether or not there exists a 'normal' profile in walking. Especially in the field of pathological gait analysis, it is useful to know whether or not there exists an underlying walking pattern across normal population.

intra- or inter-personal

Assuming such a normal profile exists, there are many interesting questions to consider: (i) Is there an inter-personal or intra-personal pattern? (ii) What variables are related to the normal profile? (i.e., speed, step-size) (iii) In what configuration space can we find such a normal profile? Here, we briefly summarise some experimental results in quantitative gait analysis.

Kadaba et al. (1989) studied the repeatability of kinematic and kinetic data in normal adult gait. The author found that when the subjects are walking at their own preferable speed, the gait variables are repeatable.

A similar outcome is also discussed in Winter (1984), when a single person is walking with various cadences, the kinematic pattern is highly repeatable. However, there is a high variation on the joint moment at the knee and hip joints, and presumably, such a variation is the result of adaptation to produce the same kinematic pattern. Winter (1984) also observed variability in the behaviour across the normal population, and found a high degree of covariance in the hip and knee patterns between subjects.

Orendurff et al. (2004) examined the relationship of the centre-of-mass in the vertical and frontal direction, and found that the CoM displacement changes substantially with walking speed. Although centre-of-mass is essential to whole body balance, even normal individuals show significant mediolateral CoM displacement at slow speeds.

Pedotti (1977) studied the variability of EMG signals and how EMG might be related to the variability of joint-moments. The author found that, although the EMG signals are different from one subject to another, the overall *shapes* of the EMG signal are very similar. This result is also confirmed by Arsenault et al. (1986), such that the EMG profiles are different across subjects. On the other hand, apart from these inter-personal variations in the EMG signals, the EMG from each individual can be represented by a specific profile.

Experimental results suggest the evidence of consistency and variations at different space (kinematics, kinetics, etc) and different levels (either intra-personal or inter-personal). In our work, we utilise the above results to hypothesise the decomposition of walking gait and formulate our walking phase model.

2.2 CONTROL MECHANISM FOR BIPEDAL WALKING ROBOT

Over the past fifty years, research into humanoid robots has drawn from various perspectives, and many of them were inspired from human gaits. One essential component is to make the robot *human-like*, such as maintaining balance and/or minimising energy consumption during walking, like is done in the human body.

The early experimental results of bipedal walking robots were reported in Kato (1973). WABOT-1 was the first full-scale anthropomorphic robot, which achieved stable walking while transporting objects with its gripper. The best present-day robots can walk, turn, and climb, including the Asimo (Hirai et al., 1998) and HRP-3 (Kaneko et al., 2008).

Since there are successful stories in humanoid robots, can we adapt the same control paradigm for rehabilitation robots? In this section, some common approaches to

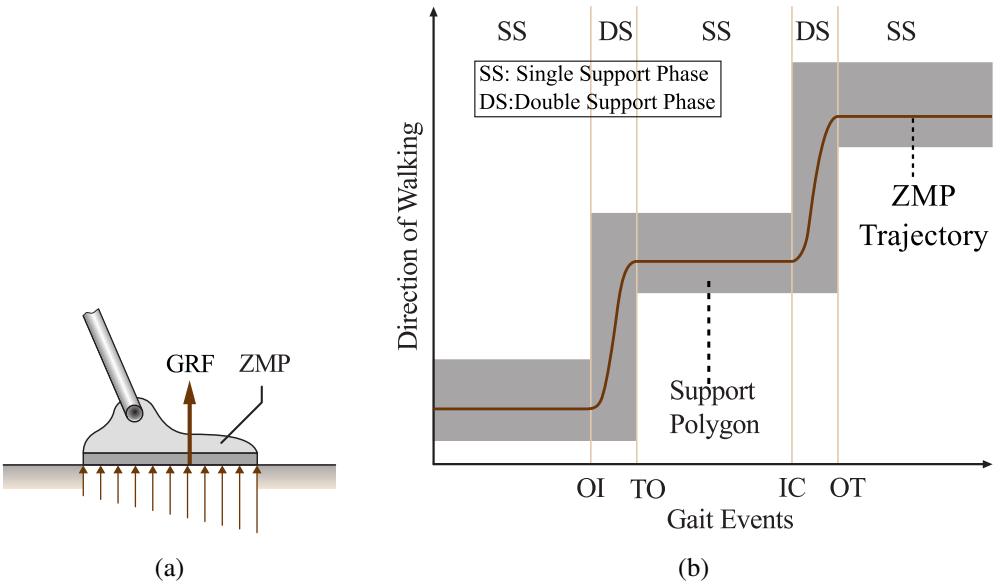

Figure 2.4: Illustration of zero-moment-point: (a) Definition of zero-moment-point (b) Example of support polygon within a gait cycle (Figures extracted and modified from Kajita and Espia (2008))

humanoid robots are examined, including (i) zero-moment-point, (ii) passive dynamic walking, and (iii) learning by imitation. Further details on control of bipedal walking and humanoids robots can be found in the books of Grizzle et al. (2007) and Siciliano and Khatib (2007).

2.2.1 *Zero-Moment Point*

Ever since the Zero-Moment-Point (ZMP) was introduced by Vukobratović and Stepanenko (1972), this concept had become one of the most versatile approaches for controlling the dynamical stability of bipedal walking robots.

zero-moment-point

Before discussing the concept of ZMP, some definitions should be clarified. The Ground-Reaction-Force (GRF) is the sum of all active forces acting on the feet (i.e., inertia, gravitation, and Coriolis) during walking. The Centre-of-Pressure (CoP) is the point on the feet of the stance leg at which the GRF acts on. The ZMP refers to the point on the ground which the ground reaction force is acting on (Fig. 2.4a).

ground-reaction-force
centre-of-pressure

Note that the ZMP coincides with the CoP when a robot is stabilised. A change in the dynamic of walking will change the direction and magnitude of the GRF, and hence,

causing a change in CoP and ZMP. During the stance phase, the CoP moves from the heel toward the toe.

The central idea of ZMP is that, to maintain stability of the robot, the GRF should act at a particular point on the foot to balance out all the forces acting on the robot during walking. To achieve this equilibrium, the moment of GRF with respect to ZMP should be zero. Therefore, by controlling the ZMP location within some range of stability, which is usually referred as the *support polygon* (the shaded area in Fig. 2.4b), the robot may induce forward motion while maintaining balance.

support polygon

Although the ZMP position cannot be controlled directly, it can be ensured by controlling the appropriate movement of the joints. Several previous works have used a controller to minimise the error between the desired ZMP and the output ZMP for humanoid locomotion (Hirai et al., 1998; Kajita et al., 2003; Ogura et al., 2006; Guan et al., 2006).

In spite of the fact the ZMP approach has produced many successful results, maintaining ZMP condition is not sufficient for asymptotic stability of a periodic walking motion. In additional to the standard ZMP, the extension and variations were introduced, including the foot rotation indicator (FRI) (Goswami, 1999) and the centroidal moment pivot (CMP) (Popovic et al., 2004). Additional information on ZMP-based method and comparisons are given in Popovic et al. (2005).

2.2.2 *Passive Dynamic Walking*

Another control paradigm relies on the *passive dynamics walking* principles, which uses gravitational force to gain the energy necessary for walking. A passive dynamic walker consists of a stance leg and a swing leg connected by the hip. Inspired by the passive nature of human walking, the swing leg is described as a double pendulum (Mochon and McMahon, 1980), and the stance leg behaves like an inverted pendulum (Miura and Shimoyama, 1984; Raibert, 1986).

passive dynamics

inverted pendulum

The first robot that is based on passive dynamic principles can walk stably down a shallow slope without actuation (McGeer, 1990a). The first generation of passive walkers can only allow a small set of initial condition, and several variations has been built upon the first generation passive walker. For instance, the simplest walking model assume that the feet are massless, based on the fact that the hip mass is normally much larger than the foot mass, so the motion of the swing does not affect the motion of the hip (Garcia et al., 1998). For a compass-gait model, the distributed inertia is rep-

resented by hip mass and leg mass (Goswami et al., 1998). Some work extend the original passive walker by adding knees and/or feet (McGeer, 1990b; Chen, 2007).

The main limitation of a purely passive walker is that the robot needs to walk down on a slope, and it is unable to climb, pause, turn, or run. Some extension have been proposed to introduce partially actuated systems which were governed primarily by passive dynamics, such as adding hip actuation (Wisse, 2004) or ankle push-off (Tedrake et al., 2004).

Note that, in contrast with the ZMP paradigm discussed in the preceding section, which constantly controls joint-angles to ensure the ZMP stays in the support polygon, the passive dynamic walking approach is more energetically efficient. This concept has been a backbone for much research on bipedal robot locomotion (Spong, 1998; Kuo, 1999; Chevallereau et al., 2005).

2.2.3 *Trajectory Tracking*

Another well-established approach for controlling a humanoid robot is to combine a priori definition of desired motion to follow. This desired motion, or reference gait, can be obtained from either human demonstration or purely based on simulation. Motivated by a human's capability of learning and imitating demonstrated behaviour, tracking demonstrated behaviours has been explored as an efficient method to accomplish desired movements.

Pollard et al. (2002) use human demonstrated gesture to control the upper body of a Sarcos humanoid robot. Nakaoka et al. (2003) capture human dance to control a biped humanoid. Although not directly applied on a humanoid robot, tracking human demonstration has been also used in the field of graphics and animation (Arikan and Forsyth, 2002; Ren et al., 2005; Suleiman et al., 2008).

Biologically inspired approaches based on *central pattern generators* (CPGs) have drawn much attention for movement generation. The CPGs are a set of neural oscillators which are capable of generating rhythmic motion without any external inputs (Ijspeert, 2008). Early work of Taga et al. (1991) has inspired many applications of using CPGs for bipedal locomotion (Okada et al., 2002; Geng et al., 2006).

Ijspeert et al. (2003) proposed the *dynamic motor primitives* approach where the movements are encoded as a set of autonomous non-linear differential equations. This approach was later extended to bipedal locomotion (Nakanishi et al., 2004; Aoi and Tsuchiya, 2005).

In contrast with ZMP controller and the dynamic walking control paradigm, the trajectory tracking approach does not rely on precise modelling or assumptions on human dynamics and allows encoding complex human movements. In the next section, we will explore some statistical methods for generating the reference movements from human demonstration.

2.3 ROBOT-ASSISTED REHABILITATION

In the preceding section, we listed a few examples of bipedal walking and humanoid robots. Since there are many successful stories, can we employ similar control paradigms for a more complex system where a human subject is involved in the control loop? In this section, we discuss some examples and issues in *robot-assisted rehabilitation*.

In 2011, the Department of Work and Pensions estimated that 6.6 million people in the United Kingdom have some type of mobility impairment (e.g., spinal cord injury, cerebral palsy, etc) that necessitates the use of assisting devices (Department for Work and Pensions, 2012). On top of that, according to the World Health Organization, the aging population is increasing remarkably in all nations. The average median age was 22.4 in 1975 and 28.0 in 2005. With such a rapid increase, the median age is expected to be 38.1 by 2050 (World Health Organization, 2011). Disability rates are higher among older people; hence, we can foresee a higher demand for assistance.

During rehabilitation, the patients have to repeatedly perform some movement or exercise according to the instructions of the therapist. The training procedure in rehabilitation is highly labour intensive, which usually requires several trained therapists to assist a single patient. Also, the manual assistance provided by different therapists over multiple training sessions can be quite diverse.

rehabilitation robotics

The goals of rehabilitation robotics are to use wearable physical devices that interact with the patients, assist them to exercise, and correct them when they fail to move in the correct way (Fig. 2.5). Introducing a robotic system to rehabilitation has several benefits (Fasoli et al., 2004). The success of such a system can relieve the manual labour needed from therapists, and the training period can be longer and more repeatable. With the high demand in rehabilitation in line with future demographic shift, it is important to advance robot-assisted rehabilitation.

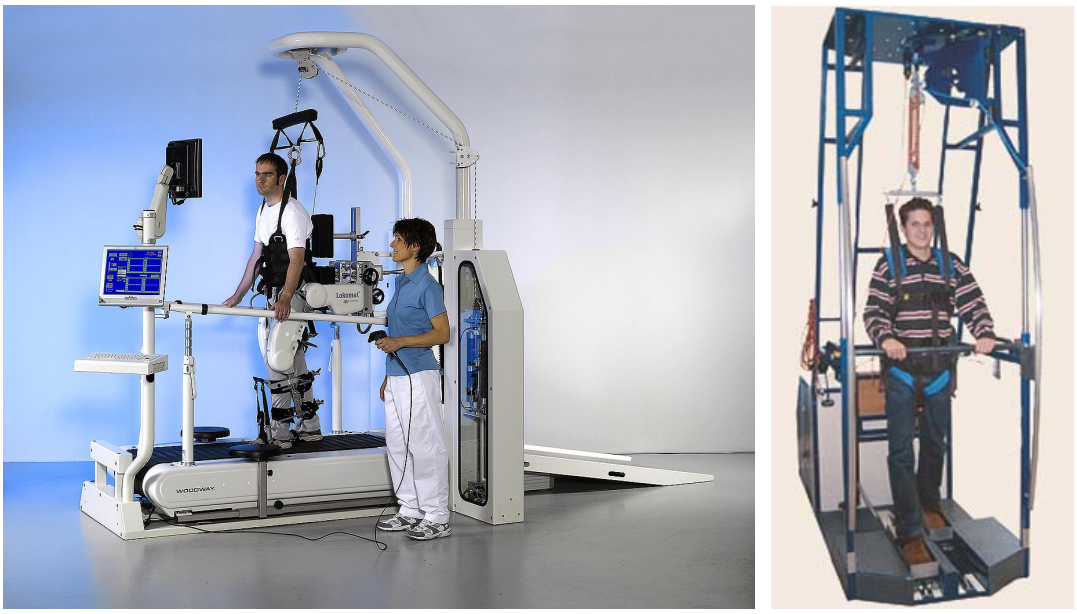

(a) (b)

Figure 2.5: Robot-assisted rehabilitation devices in clinics. (a) Lokomat by Hocoma, Switzerland (Lünenburger et al., 2007) (b) Gait Trainer by RehaStim, Germany (Hesse and Uhlenbrock, 2000)

2.3.1 *Wearable Lower-limb Devices*

In the past few decades, several wearable lower-limb devices were built for various purposes. The most prominent method of gait rehabilitation is body-weight supported treadmill training, which enables motor deficient patients to support their own body weight and experience locomotion. Currently, the devices available in clinics include Lokomat from Hocoma AG, Switzerland (Fig. 2.5a), AutoAmbulator from HealthSouth, USA (Fig. 2.5b), and Gait Trainer from RehaStim, Germany (Duschau-Wicke et al., 2010; Hesse and Uhlenbrock, 2000).

Meanwhile, research institutes and industries are developing and improving gait rehabilitation devices such as the Lower Extremity Powered Exoskeleton (LOPES), Active Leg EXoskeleton (ALEX), Pelvic Assist Manipulator - Pneumatically Operated Gait Orthosis (PAM-POGO), and the MIT Skywalker (Veneman et al., 2007; Banala et al., 2009; Aoyagi et al., 2007; Artemiadis and Krebs, 2011).

Some work considers overground walking instead of a body-weight support system. For example, ReWalk (Zeilig et al., 2012), the powered knee brace, RoboKnee (Pratt

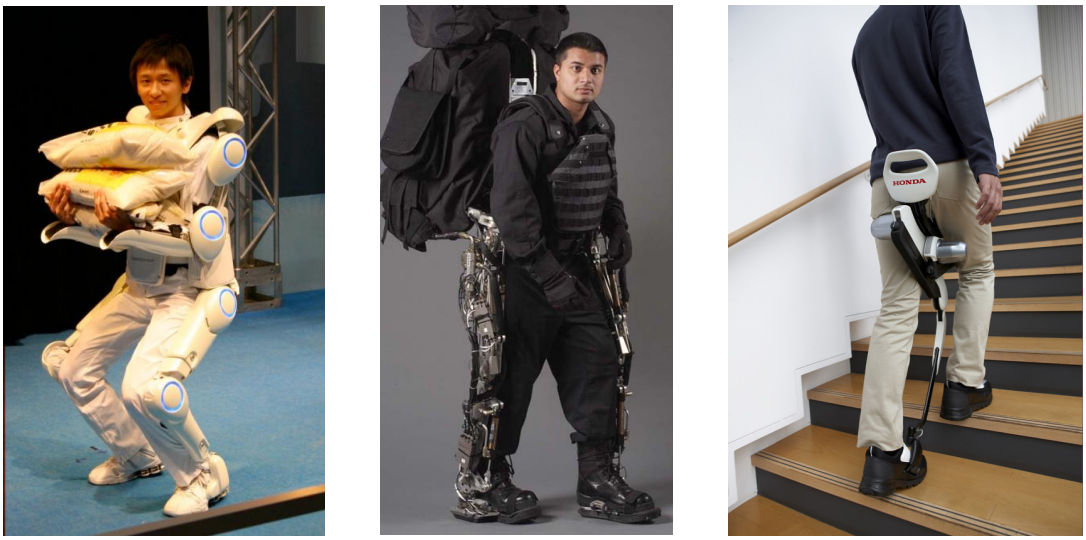

(a) (b) (c)

Figure 2.6: Overground walking devices (a) HAL-3 (Hayashi et al., 2005) (b) BLEEX (Zoss et al., 2005) (c) Body-weight Support Assist (Honda, 2009)

et al., 2004) and Stride Management Assist and Bodyweight Support Assist (Fig. 2.6c) for mobility assistance (Honda, 2009).

Another category is that of exoskeleton devices for general purpose (Fig. 2.6), such as the Hybrid Assistive Legs (HAL-3) (Lee and Sankai, 2002) and Berkeley's Lower Extremity Exoskeleton (Bleex) (Kazerooni et al., 2005). Although the original intention for these systems is to amplify the capabilities of the wearer, there is the potential application of gait rehabilitation.

2.3.2 *Control Mechanism*

One major challenge of designing wearable devices is when and how much to control the user. Assisting the patients through a predefined reference gait with a feedback controller seems the most prominent approach in gait rehabilitation (Krebs et al., 1998; Lum et al., 1993).

An example of such a control scheme is sketched in Fig. 2.7. This reference gait $\mathbf{q}^{\text{ref}}$ is normally obtained by taking the average or polynomial fit of some data gathered from healthy subjects, and the feedback controller calculates the torque τ needed to correct the subject. For example, as the participant moves away from the reference gait,

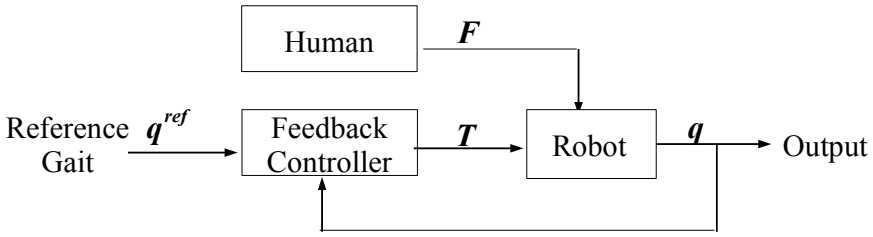

Figure 2.7: An example of the control scheme for robot-assisted rehabilitation.

the feedback controller should increase the torque proportionally. A simple position controller can be formed as

$$\tau = \mathbf{K}(\mathbf{q}^{\text{ref}} - \mathbf{q}) + \mathbf{C}\dot{\mathbf{q}}$$

where τ is the joint-torque, $\mathbf{K}$ is the stiffness factor, and $\mathbf{C}$ is the damping factor.

Most devices mentioned in the prior section either enforce the position of the limbs and/or the force interaction between the device and the limbs. Most commercially available systems take the former approach while the more recent and on-going work has shifted from position-control to force-control.

Variations can also be found in over-ground walking assistance, such as the balance controller approach implemented in BLEEX and STRING-MAN (Steger et al., 2006; Surdilovic et al., 2007), the feet-position controller in the MIT Skywalker (Artemiadis and Krebs, 2011), and the EMG-based control in HAL-3 and ankle-foot orthosis (Hayashi et al., 2005; Ferris et al., 2005). Detailed surveys of rehabilitation devices and exoskeletons can be found in Marchal-Crespo and Reinkensmeyer (2009) and Herr (2009).

2.3.3 *Assist-as-needed*

Early work in robotic rehabilitation mainly applied position control that moved the patient's leg through a prescribed reference trajectory, irrespective of the patient's capabilities or preferences. However, clinical trials suggest that motivating the patient to walk actively promotes the overall results of rehabilitation (Hidler et al., 2005). Recent work adapts the *assist-as-needed* concept, in which the robot assists the patient without externally constraining them to a predefined reference trajectory (Emken et al., 2007). Some examples of research methodology that attempts this are discussed next.

assist-as-needed

One common strategy is to tolerate the movement within some deviation from a given fixed reference trajectory; for instance, the "position control" allows some spatial and temporal deviation from the reference trajectories (Riener et al., 2005), and the "virtual tunnel" approach allows the wearer to have some level of free movement (Kim et al., 2010).

Another solution to aid flexibility is by implementing impedance control and direct adaptation of the gait pattern (Jezernik et al., 2004; Ekkelenkamp et al., 2005), automatically adjusting the level of assistance according to the performance of the patient (Krebs et al., 2003; Banala et al., 2009), or using a spring-damper system to allow some deviation between the actual joint positions and reference positions (Lee and Sankai, 2005).

Some work aims to be flexible with the walking speed of the user. The "path control" of Lokomat allows movements within some temporal error (Duschau-Wicke et al., 2010). Veneman et al. (2006) uses a reference trajectory that is variable in time and constantly synchronises with the users' movements. Another approach is to consider representing the state of the system by the relative position in the gait cycle, so the speed of the movement is flexible. For example, MIT Skywalker measures the relative position of the feet in the horizontal dimension to determine the desired position in the vertical dimension (Artemiadis and Krebs, 2011).

On the other hand, we consider this type of approach is probably suboptimal since some dimensions may need no correction. The main issue is that the reference gait is taken from the average framework of normal walking, which is affected by natural variations such as embodiments, environments, and behaviours.

2.4 STATISTICAL LEARNING FROM DEMONSTRATION

As mentioned in Section 2.3, a common principle in physical therapy has been an attempt to make movements more "normal"; thus, the desired movement is to develop movement patterns that are similar to healthy individuals. For this reason, generalising the walking gait to a variety of subjects is particularly critical in gait rehabilitation.

On the other hand, walking gaits are highly influenced by intra-personal and inter-personal variations (Section 2.1). What defines the normal walking gait has yielded no clear answer. One way to arrive at the best solution is to employ *optimisation* strategies that somehow capture the essence of the walking gait. This optimisation process is normally accomplished though *statistical learning*.

statistical learning

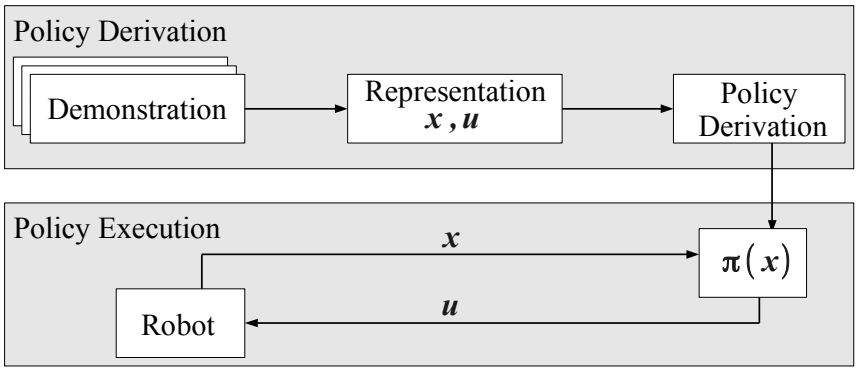

Figure 2.8: A schematic of statistical learning from demonstration. In the policy derivation step, the behaviour is approximated as a policy $\boldsymbol{\pi}$. In the policy execution step, the derived policy $\boldsymbol{\pi}$ enables the robot to select an action $\mathbf{u}$ based on the current state $\mathbf{x}$

Statistical learning from demonstration lies at the heart of many approaches to the motor control problems in robotics. A skill is generalised from demonstration, and the resulting model can be re-applied on novel situations or behaviours. Specifically, given demonstrated data $\mathbf{D}$ as pairs of observed states $\mathbf{x} \in \mathbb{R}^{\mathcal{P}}$ and observed actions $\mathbf{u} \in \mathbb{R}^{\mathcal{Q}}$, the goal of statistical learning is to find the *policy* $\boldsymbol{\pi}$ that maps between the observed states $\mathbf{x}$ and observed actions $\mathbf{u}$ which allows the robot to select an action $\mathbf{u}$ based on current state $\mathbf{x}$ (Argall et al., 2009).

state/action

policy

A schematic is illustrated in Fig. 2.8. The demonstrated data $\mathbf{D}$ is formed as N pairs of observed states $\mathbf{x}$ and observed actions $\mathbf{u}$. During the *policy derivation* step, the behaviour derived from $\mathbf{D}$ is approximated as a policy $\boldsymbol{\pi}$. In the *policy execution* step, the policy enables the robot to select a $\mathbf{u}$ based on the current state $\mathbf{x}$.

2.4.1 *Policy Derivation*

In this section, we survey a few examples of policy derivation from human demonstration. There exist various methods to derive the policy from demonstrated data; here, we briefly discuss the (i) mapping function and (ii) system model. For further details on robot learning by demonstration, extensive survey can be found in Billard et al. (2007) and Peters and Schaal (2008).

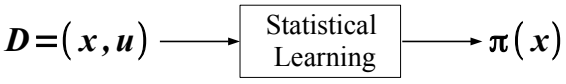

Figure 2.9: Policy derivation by direct mapping

2.4.1.1 *Mapping Function*

One way to approximate the policy $\boldsymbol{\pi}$ is to directly map from the observed state $\mathbf{x}$ to the observed action $\mathbf{u}$ from the demonstrated data $\mathbf{D}$ (Fig. 2.9)

$$\boldsymbol{\pi} : \mathbf{x} \rightarrow \mathbf{u},\ \boldsymbol{\pi} \in \mathbb{R}^{\mathcal{P}} \rightarrow \mathbb{R}^{\mathcal{Q}}$$

Note that, many human behaviours are difficult to describe. This approach is generic and less task-specific, so it does not require much prior knowledge of the domain dynamics.

Classical examples of learning by imitation normally formulate the problem in action-space (Craig et al., 1987), with feedback/feed-forward controller (Kawato, 1990). Locally Weighted Regression (LWR) (Cleveland and Devlin, 1988) has been adapted to learn patterns in joint-trajectories for bipedal robot locomotion (Nakanishi et al., 2004). Locally Weighted Projection Regression (LWPR) is an extension of LWR which operates efficiently even in high dimensional space (Vijayakumar and Schaal, 2000). Successful application can be found in the work of Grollman and Jenkins (2007), which enable the Aibo robot to mimic behaviours such as ball seeking.

Other regression methods can be applied to learn the policy such as Neural Network (Kawato, 1990; Butz et al., 2007; Gomes et al., 2011), and Gaussian process regression (Plagemann et al., 2008; Nguyen-Tuong et al., 2009). Support Vector Regression (SVR) (Vapnik, 2000) leverages the standard regression method by mapping the data-set to higher dimensionality to handle non-linear problems, and has been applied for robotic grasping (Pelossof et al., 2004) and approximating centre of mass (Zhao et al., 2013).

Instead of learning the policy directly, some authors take a probabilistic approach, such as using Gaussian Mixture Regression (Calinon et al., 2010; Ghahramani and Jordan, 1994), Gaussian Process Dynamical Models (Wang et al., 2008), and Hidden Markov Models (Inamura et al., 2004).

The dynamical system approach mentioned in Section 2.2 also falls within this category. This approach includes a network of non-linear differential equations for the encoding of the demonstrations and a policy derivation technique for shaping the

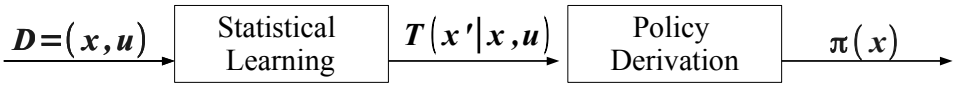

Figure 2.10: Policy derivation by reinforcement learning

landscapes according to the demonstration, which is well-established for humanoid robots (Ijspeert et al., 2003).

2.4.1.2 *System Model*

Instead of learning the mapping directly, another approach is to encode the movement by model-based learning. From the demonstration $\mathbf{D}$, not only the policy $\boldsymbol{\pi}$, but a predictive model of the skill is approximated (Wolpert and Kawato, 1998).

This approach is often formulated as a *Reinforcement Learning (RL)* problem (see Fig. 2.10), where a state-transition function $\mathbf{T}(\mathbf{x}' \,|\, \mathbf{x}, \mathbf{u})$ is approximated. The optimal policy $\boldsymbol{\pi}$ is derived from the state-transition function $\mathbf{T}$ with a reward function $\mathbf{R}(\mathbf{x})$ that maximises the cumulative reward function $\mathbf{V}$

$$\mathbf{V}(\mathbf{x}) = \int_{\mathbf{u}} \boldsymbol{\pi}(\mathbf{x}, \mathbf{u}) \int_{\mathbf{x}'} \mathrm{T}(\mathbf{x}' \,|\, \mathbf{x}, \mathbf{u}) \left[\mathbf{R}(\mathbf{x}) + \gamma \mathbf{V}(\mathbf{x}')\right]$$

where γ is a discounting factor for future reward (Sutton and Barto, 1998).

Note that, the reward function $\mathbf{R}$ is either defined by the user or learnt from demonstrations. Defining a reward function to accurately capture the demonstrated skill is not always trivial, especially in a real world system. For this, some research proposed the *Inverse Reinforcement Learning* to address this issue, where the reward function is learnt rather than explicitly specified (Russell, 1998). Examples of using IRL for motor control can be seen from the work on swing-up of a pendulum for a humanoid robot (Atkeson and Schaal, 1997) and Small legged robot (Kolter et al., 2008).

Instead of IRL, an alternative is to take the gradient approach, which directly optimises between the observed actions and the reward functions (Neu and Szepesvári, 2007). Some others model the probabilistic distribution of the reward functions and use Bayesian Inference to find the most likely reward (Ramachandran and Amir, 2006).

2.4.2 *Constraint Model*

Following up on our discussion on policy derivation, a promising way to provide robots with skills is to take examples of human demonstrations and attempt to learn a mapping

between the state and actions that somehow capture the demonstrated behaviours under real world *constraints*.

constraints

Many everyday human behaviours can be considered in terms of performing some task subject to a set of self-imposed or environmental constraints. For example, when pouring water from a bottle, self-imposed constraints apply to the position and the orientation of the hand so that the water falls within a glass. When wiping a table, the surface of the table acts as an environmental constraint that restricts the hand movements when maintaining contact to the surface. Based on the principles of analytical dynamics, a skill can be formalised as a *constraint model* (Udwadia and Kalaba, 2007).

constraint model

Given demonstration $\mathbf{D}$ as pairs of observed states $\mathbf{x} \in \mathbb{R}^{\mathcal{P}}$ and observed actions $\mathbf{u} \in \mathbb{R}^{\mathcal{Q}}$, the policy can be described as systems of the form $\mathbf{u}(t) \equiv \boldsymbol{\pi}(\mathbf{x}(t))$. In a constraint model, the behaviour is assumed to be subject to a set of k-dimensional Pffaffian constraint

constraint matrix

$$\mathbf{A}(\mathbf{x})\,\mathbf{u} = \mathbf{b}(\mathbf{x})$$

where $\mathbf{A}(\mathbf{x}) \in \mathbb{R}^{\mathcal{S}\times\mathcal{Q}}$ and $\mathbf{b}(\mathbf{x}) \in \mathbb{R}^{\mathcal{S}}$ together describe the constraints. By inverting the constraint model, the observed actions can be described as

$$\mathbf{u} = \mathbf{A}^{\dagger}(\mathbf{x})\,\mathbf{b}(\mathbf{x}) + \mathbf{N}(\mathbf{x})\,\boldsymbol{\pi}(\mathbf{x}) \tag{2.1}$$

where $\mathbf{A}^{\dagger}$ is the pseudo-inverse of $\mathbf{A}$, $\mathbf{N} \equiv \mathbf{I} - \mathbf{A}^{\dagger}\,\mathbf{A} \in \mathbb{R}^{\mathcal{Q}\times\mathcal{Q}}$ is the *projection matrix* that projects the policy $\boldsymbol{\pi}$ onto the null-space of the constraint matrix, and $\mathbf{I} \in \mathbb{R}^{\mathcal{Q}\times\mathcal{Q}}$ is the identity matrix.

pseudo-inverse
projection matrix

The effect of constraint $\mathbf{A}$ is to modify the policy $\boldsymbol{\pi}$ such that $\boldsymbol{\pi}$ can only act on the null-space of $\mathbf{A}$. This model is generic and can be applied to wide variety of problems.

2.4.2.1 *Joint-space Control*

In the case of *kinematic control*, the state and actions are the joint-angle and joint-velocity $\mathbf{x} \equiv \mathbf{q}$ and $\mathbf{u} \equiv \dot{\mathbf{q}}$. The constraint model can be formed as $\mathbf{A}(\mathbf{q})\,\dot{\mathbf{q}} = \mathbf{b}(\mathbf{q})$, and the observed joint-velocity under the constraints are

kinematic control

$$\dot{\mathbf{q}} = \mathbf{A}^{\dagger}(\mathbf{q})\,\mathbf{b}(\mathbf{q}) + \mathbf{N}(\mathbf{q})\,\boldsymbol{\pi}(\mathbf{q})$$

For *force control* , the observed states are represented by the joint-angles and joint-velocities $\mathbf{x} \equiv (\mathbf{q}, \dot{\mathbf{q}})$ and the observed actions are the joint-accelerations $\mathbf{u} \equiv \ddot{\mathbf{q}}$, the dynamics can be expressed in the Lagrangian form

force control

$$\tau = \mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}(\mathbf{q})$$

where $\tau \in \mathbb{R}^{\Omega}$ is the joint torques, $\mathbf{M}(\mathbf{q}) \in \mathbb{R}^{\Omega} \times \Omega$ is the inertia matrix, $\mathbf{C}$ is the matrix of centrifugal, gyroscopic, and Coriolis effects, and $\mathbf{G}$ is the generalised gravitational torque. The constraint model can be formed as

$$\mathbf{A}(\mathbf{q}, \dot{\mathbf{q}})\ddot{\mathbf{q}} = \mathbf{b}(\mathbf{q}, \dot{\mathbf{q}})$$

The desired joint-force can be determined by solving the inverse dynamic of the system (Hollerbach and Suh, 1987)

$$\tau = (\mathbf{A}\,\mathbf{M})^{\dagger}(\mathbf{b} - (\mathbf{A}\,\mathbf{M})^{-1}\mathbf{F}) + (\mathbf{I} - (\mathbf{A}\,\mathbf{M})^{\dagger}(\mathbf{A}\,\mathbf{M}))\,\boldsymbol{\pi}$$

2.4.2.2 *Operational-space Control*

Khatib (1987) introduced the *operational-space* formulation to address the dynamic of task-space movement. In the context of operational-space formulation, a *task* can be any kind of activity that has a one-to-one correspondence between the observed state $\mathbf{x}$ and the state in the task-space $\boldsymbol{\rho}$; i.e., the relationship between them can be described as

task-space

$$\boldsymbol{\rho} = \mathbf{J}(\mathbf{x})$$

where $\mathbf{J}$ is the Jacobian that relates the actuator-space to task-space.

Jacobian

One common application is to control a robot manipulator; that is, given a desired motion in the end-effector space such as reaching a target or fetching an object, we seek the set of joint-velocities for the robot manipulator to accomplish the task. The end-effector position is given by $\boldsymbol{\rho} \equiv \mathbf{r} = \mathbf{J}(\mathbf{q})$ where $\mathbf{J}$ denotes the forward kinematics matrix, the end-effector velocities and accelerations are $\dot{\mathbf{r}} = \mathbf{J}(\mathbf{q})\dot{\mathbf{q}}$ and $\ddot{\mathbf{r}} = \mathbf{J}(\mathbf{q})\ddot{\mathbf{q}} + \dot{\mathbf{J}}(\dot{\mathbf{q}})$.

robot manipulator

This control scheme for robot manipulators commonly appears in scenarios where the robot interact with objects or to avoid joint-limit (Liégeois, 1977), singularities (Nakamura and Hanafusa, 1986), or obstacles (Baillieul, 1986; Khatib, 1986).

Note that, although operational-space control was originally developed for solving redundancy in a robot manipulator, it is not limited to this application (i.e., the constraint can be imposed in various internal or external coordinate reference frames other than the end-effector space). This is particularly useful since many human behaviours may be subject to various constraints that are non-linear in actuator space.

This formalism has been readily applied in humanoid locomotion; for instance, the constraint can be imposed in the centre-of-gravity for balancing robots (Sentis

and Khatib, 2006) or imposed in a heuristic virtual model for bipedal robots (Pratt et al., 2001). A variety of further examples exist, such as controlling compliant legged robots (Hutter et al., 2012), planning foot placements for a humanoid robot (Kanoun et al., 2011), or maintaining ZMP for stabilising humanoid robots (Stephens, 2007; Hofmann et al., 2009; Miura et al., 2009).

In recent years, a hierarchical approach holds promise for robot control with more complex behaviours. For instance, maintaining balance of a robot (higher priority) while accomplishing a end-effector task (lower priority) (Sentis and Khatib, 2005; Gienger et al., 2005) or avoiding obstacles (Stilman and Kuffner, 2008).

2.5 BRIDGE BETWEEN REHABILITATION, GAIT ANALYSIS, AND STATISTICAL LEARNING

A common approach for robot-assisted rehabilitation is to correct the patients through a reference gait with a feedback controller (Section 2.3). From Fig. 2.7, we can see that having a correct reference gait is very crucial in the control loop. Our work is motivated by the need of finding this reference gait. The current issue is that the reference gait is normally defined by taking the average gait from a population of healthy subjects. However, as mentioned in Section 2.1, walking movements vary across behaviours and subjects; and this approach would ignore these natural variations.

While much research in bipedal locomotion has achieved promising results (Section 2.2), the ability to adapt these control mechanisms to robot-assisted rehabilitation remains unsolved. First, humanoid robots maintain stability by controlling ZMP position, yet humans and robots may not follow the same mechanism for balancing. In fact, human locomotion is constantly imbalanced, especially during the swing phase (Shumway-Cook and Woollacott, 1995) (i.e., the ZMP position of human is often at the edge of the support polygon), and this might be the reason that the motions generated with ZMP are not natural.

Second, many bipedal walking robots were based on passive dynamic principles (Section 2.2.2), but there is no direct evidence to prove that the reference gait generated from a passive walker is optimal for human walking. In Section 6.5, we will compare our proposed method and a model based on passive dynamic principle on learning motion-capture data of human walking.

Additionally, although many bipedal walkers adapt the bio-inspired networks (Section 2.2.3), designing this type of system is difficult since one must determine the archi-

tecture of the networks, the number of oscillators, the type of couplings, the waveforms, the type of input, etc, which are not theoretically determined.

In Section 2.4, we reviewed some work in *learning by demonstration* which derives a policy that maps between state and action space. In Section 2.4.2, we discussed a special case of learning by demonstration, which aims at solving the consistent behaviour in robotic task (e.g., reaching different targets). We recognised the similarities between human walking and robot-reaching; in particular, both problems are seeking a consistent policy from observations under various unknown constraints. Hence, we will adapt the technique for robot-reaching task (Section 2.4.2.2) to our problem of walking.

Again, a robust method for learning the reference gait is particularly important for gait rehabilitation. In Section 3.2, we will take the results in clinical gait analysis to outline a model using the constraint formulation discussed in Section 2.4.2. In Chapter 4, we will extend the algorithm in operational space control (Section 2.4.2.2) for extracting the consistency from data generated with various constraints and tasks.

3

REPRESENTATION OF WALKING GAIT BY WALKING PHASE MODEL

> *"...this body like a machine which, having been made by the hand of God, is incomparably better structured than any machine that could be invented by human beings, and contains many more admirable movements. "*
>
> *—René Descartes*

As mentioned in the previous chapter, experimental results in gait analysis suggest the presence of consistency and variations in kinematics and dynamics of walking gait, either intra-personal or inter-personal. Based on this insight, we assume that a walking gait can be described as a combination of the *consistent characteristics* of gait and *variations* coming from environment, embodiment, and behaviours. In order to recover the underlying characteristics, a model such that the *characteristics* and *variations* can be separated is needed.

In this chapter, we describe our approach on representing and modelling locomotion. In Section 3.2, we describe a constrained tracking control scheme in which behaviour is decomposed into task/null-space components. We then describe our interpretation of walking gait and how to determine the control variables in this control scheme (Section 3.2.2). In Section 3.3, we describe some example of representations for this model.

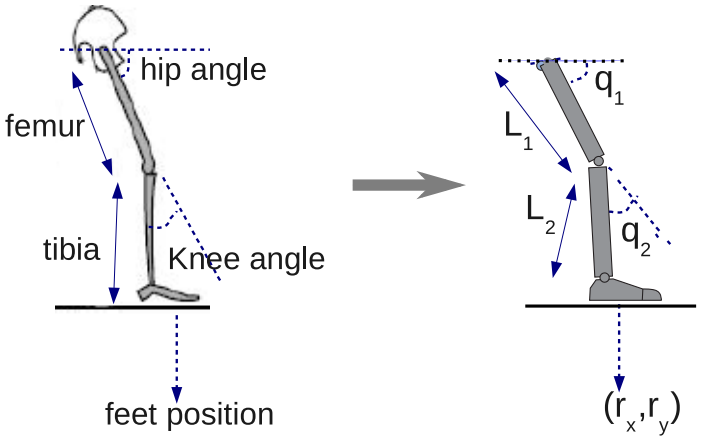

Figure 3.1: Correspondence between leg and 2-Link System

3.1 REPRESENTATIONS OF HUMAN GAIT

A simple model of the human leg can be described by a planar two-link system (Fig. 3.1). In this model, L_1, L_2 are the length of femur and tibia, respectively, the angles $\mathbf{q}_1, \mathbf{q}_2$ represent the hip and knee angles, respectively, and $\mathbf{r}_x, \mathbf{r}_y$ represent the horizontal and vertical position of the foot.

For kinematic control, the state of the system can be represented by the joint angles $\mathbf{x} = (\mathbf{q}_1, \mathbf{q}_2)^\top$, and the controls can be the joint-velocities $\mathbf{u} = (\dot{\mathbf{q}}_1, \dot{\mathbf{q}}_2)^\top$. For dynamical control, the state might be joint positions and velocities $\mathbf{x} \equiv (\mathbf{q}, \dot{\mathbf{q}})$ with actions corresponding to joint torques $\mathbf{u} \equiv \tau$.

Most of the commercially available rehabilitation systems mentioned in Section 2.3, such as Lokomat and GaitTrainer, follow the former representation (Duschau-Wicke et al., 2010; Hesse and Uhlenbrock, 2000). Meanwhile, the majority of rest are force controlled, including LOPES, ALEX, and POGO (Veneman et al., 2007; Agrawal et al., 2007; Aoyagi et al., 2007).

To aid the explanation, we use the two-link system in Fig. 3.1 as an illustrative example throughout this chapter.

3.2 WALKING PHASE MODELS

In this thesis, we follow the gait analysis literature discussed in Section 2.1, and decompose walking gait into a series of *gait cycles* and *gait phases*. As mentioned, these gait phases have been studied since they are diverse in their kinematic and/or kinetic

features. In order to deal with these variations in different phases along with inherent variations in embodiment and task-related factors, a sufficiently flexible model is required that can be used to decompose the movement.

We assume that movements within each walking phase are the result of a composition of components handling the *phase-critical components* of motion (i.e., those that must be controlled for successful completion of a given phase), and redundant components that control consistent aspects of motion. In order to satisfy the phase critical components, it is assumed the set of constraints

phase-critical components

$$\mathbf{A}_k(\mathbf{x}, t)\,\mathbf{u}(\mathbf{x}, t) = \mathbf{b}_k(\mathbf{x}, t) \tag{3.1}$$

is maintained, where $\mathbf{x} \in \mathbb{R}^{\mathcal{P}}$ represents observed state, $\mathbf{u} \in \mathbb{R}^{\mathcal{Q}}$ represents the observed action, t is time, and k indexes the phase. Here, the *task-space policy* $\mathbf{b}_k(\mathbf{x}, t) \in \mathbb{R}^{\mathcal{S}}$ ($\mathcal{S} < \mathcal{Q}$) describes a *task-dependent* control policy. The *constraint matrix* $\mathbf{A}_k(\mathbf{x}, t) \in \mathbb{R}^{\mathcal{S}\times\mathcal{Q}}$ is a matrix projecting the task-space policy onto the relevant part of the control space. Inverting the constraint model in (3.1), results in the relation

task-space policy

constraint matrix

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{A}_k(\mathbf{x}, t)^{\dagger}\,\mathbf{b}_k(\mathbf{x}, t) + \mathbf{N}_k(\mathbf{x}, t)\,\boldsymbol{\pi}(\mathbf{x}) \tag{3.2}$$

where $\mathbf{A}^{\dagger}$ is the Moore-Penrose pseudo-inverse of $\mathbf{A}$, and we define

pseudo-inverse

$$\mathbf{N}_k(\mathbf{x}, t) := \mathbf{I} - \mathbf{A}_k(\mathbf{x}, t)^{\dagger}\,\mathbf{A}_k(\mathbf{x}, t) \tag{3.3}$$

where $\mathbf{N} \in \mathbb{R}^{\mathcal{Q}\times\mathcal{Q}}$ is a projection matrix which projects $\boldsymbol{\pi}$ onto the null-space of $\mathbf{A}$, and $\mathbf{I} \in \mathbb{R}^{\mathcal{Q}\times\mathcal{Q}}$ is the identity matrix. Note that, in (3.2), the second term arises due to the redundancy in the system (since $\mathcal{S} < \mathcal{Q}$), and allows secondary control objectives to be realised through the *null-space policy* $\boldsymbol{\pi}(\mathbf{x}) \in \mathbb{R}^{\mathcal{Q}}$. [1]

projection matrix

null-space policy

We assume that $\mathbf{A}_k$ and $\mathbf{b}_k$ are not explicitly known, but the quantities vary across walking phases to handle different phase-critical components. We assume that the null-space policy $\boldsymbol{\pi}$ is *consistent* and *independent* of the phase, however, the *observed effects* of control toward these objectives (i.e., the null-space component of motion $\mathbf{N}_k\,\boldsymbol{\pi}$) may be influenced by the phase. This is because the null-space policy $\boldsymbol{\pi}$ is subject to the higher priority constraints imposed by $\mathbf{A}_k$ and $\mathbf{b}_k$.

null-space component

1 Note that, *constraint model* is described by a variety of terms within the published literature. To clarify the terminologies, in this thesis, we term $\mathbf{b}_k$ as the *task-space policy* since the constraint matrix $\mathbf{A}$ projectes $\mathbf{b}_k$ onto to the relevant task-space, and we term $\boldsymbol{\pi}$ as the *null-space policy* since $\boldsymbol{\pi}$ can only act on the null-space of $\mathbf{A}$.

3.2.1 *Decomposition of Walking Gait*

We assume that the walking gait is a combination of *characteristics* of the gait and *variations* resulting from different embodiments and behaviours. In this section, we describe how to formulate the characteristics and variations with our walking-phase model.

3.2.1.1 *Consistent Characteristics of Gaits*

As previously stated, we hypothesise that there exists some consistent characteristics within a single class of human locomotion. In walking, such consistency may include energy minimisation, or maintenance of a comfort posture.

The consistent characteristics of walking is captured in our model as the underlying null-space policy $\boldsymbol{\pi}$. An example of such a policy can be a limit cycle policy $\dot{r} = r(\rho - r^2)$ where r and θ are the polar representation of the state such that $\mathbf{x} = (r\cos(\theta), r\sin(\theta))$, ρ is the radius of the attractor, and $\dot{\theta}$ is the angular velocity.

3.2.1.2 *Variations in Embodiments*

The variations in embodiments arise from inter-personal differences in factors such as body size, body type, and physical limits. In our model, we assume that such variations will result in modifications of the *constraint matrix* $\mathbf{A}$. (Hence, $\mathbf{A}$ varies across phases and subjects.)

In (3.1), we define the constraint matrix $\mathbf{A} \in \mathbb{R}^{\mathcal{S}\times\mathcal{Q}}$ as a set of $\mathcal{S}$ *task-constraints*, and each task constraint refers to restrictions on the freedom of some subspace of the system. The constraints can be imposed in wide variety of representations, and a simple case is to restrict some sub-space of the joint-space. For example, if the state of the system is defined as $\mathbf{x} \equiv \mathbf{q} = (\mathbf{q}_1, \mathbf{q}_2)^\top$, and the constraint matrix is set as $\mathbf{A} = [0, 1]$, the knee angle is restricted to follow various task-space policy $\mathbf{b}$ while the hip angle is free to move with the null-space policy $\boldsymbol{\pi}$.

3.2.1.3 *Variations in Environments*

Another variation of walking gait come from factors such walking terrain and slope, collectively aggregated in our model as variation in the environment. In our model, they are also captured by variation of the constraint matrix $\mathbf{A}$.

Controlled Variables		Affected Variables	
Variation	Example	**b**	**A**
behaviour	speed, cadence	✓	
embodiments	leg-length		✓
environment	slopes		✓
phase			✓

Table 3.1: Correspondence between variations in walking gaits and the variables in our proposed model

3.2.1.4 *Variations in Behaviours*

The variations in behaviours are caused by other contextual factors such as the need to hurry for a meeting, and result in variations in, for example, step sizes and speeds. In our model, these variations are captured by changes to the *task-space policy* $\mathbf{b}$, and we assume that $\mathbf{b}$ may vary across gait cycles to handle different behaviours.

In (3.1), we define the task-space policy $\mathbf{b} \in \mathbb{R}^{\mathcal{S}}$ outputs the task-space velocity, in order to accomplish some operations. Depending on which dimension is constrained (i.e., defined by $\mathbf{A}$), the constrained dimensions are restricted to move to a specific target along that dimension.

For instance, the knee joint starts flexion at the end of the stance phase and reaches its peak flexion in the middle of swing phase (i.e., this is the interval between the *toe-off* event and the *feet-adjacent* event mentioned in Section 2.1). The amount of flexion is affected by the walking speeds such that faster walks normally have smaller flexion. An example of $\mathbf{b}$ could be a point-attractor in joint-space $\mathbf{b}(\mathbf{x}) = \omega(\mathbf{x}^* - \mathbf{x})$ where ω is a scaling factor and $\mathbf{x}^*$ is the task-space target (i.e., the angle of peak flexion). In this case, ω can very to control the speed and $\mathbf{x}^*$ can vary to specify the peak flexion.

Table 3.1 summarises how the variations in embodiments and behaviours correspond to the control parameters, and how these parameters affect the variables in the model.

Fig. 3.2 shows the correspondence between walking gaits and the walking-phase model. The observed behaviours $\mathbf{u}$ are the result of some consistent policy $\boldsymbol{\pi}$ modulated by various $\mathbf{A}$, and $\mathbf{b}$. With this formulation, we examine various $\mathbf{u}$ and see if $\boldsymbol{\pi}$ can be recovered.

In Fig. 3.3, we illustrate examples of behaviours produced from three different task-constraints. We set up a planar 2-link system (as shown in Fig. 3.1) consisting of a null-

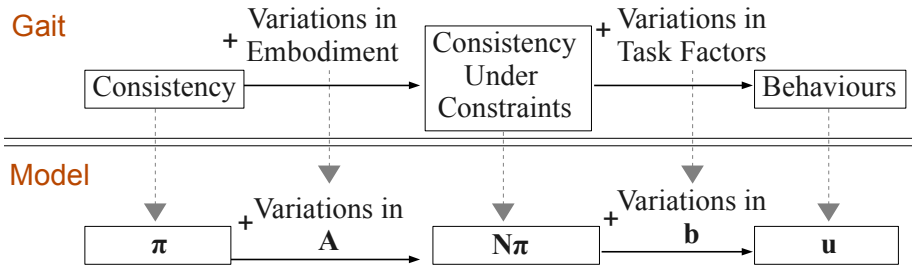

Figure 3.2: Correspondence between walking gaits and walking-phase model

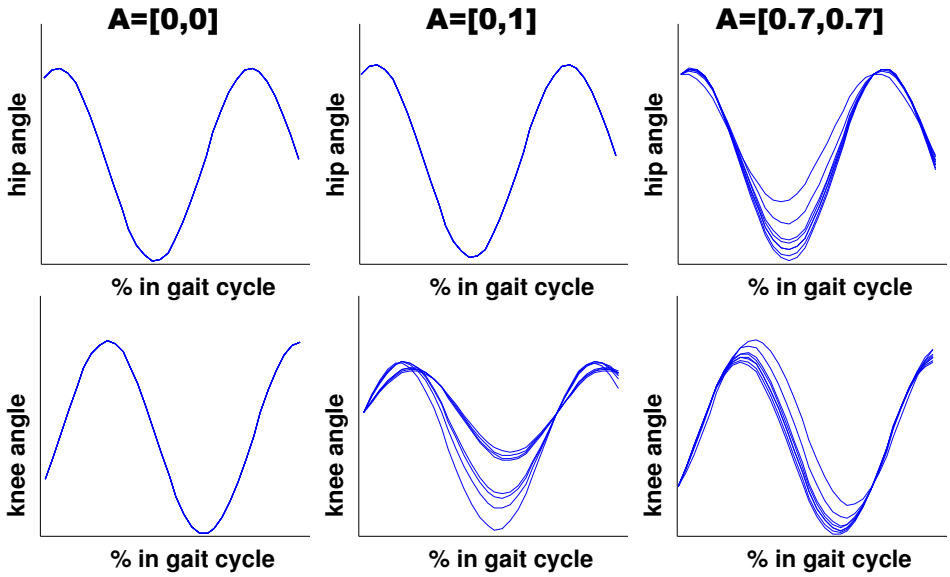

Figure 3.3: Simulated examples of three different task-constraints. (a) $\mathbf{A} = [0, 0]$, (b) $\mathbf{A} = [0, 1]$, (c) $\mathbf{A} \approx [0.7, 0.7]$

space policy $\boldsymbol{\pi}(t) = [\cos(t), \sin(t)]$ and a task-space policy $\mathbf{b}(t) = \alpha * \sin(0.25t + \alpha)$ where α was drawn from $\alpha \sim U[0.4, 1]$.

If the constraint is set to $\mathbf{A} = [0, 0]$ (fully unconstrained, Fig. 3.3a), the movements are simply the output of the null-space policy $\boldsymbol{\pi}$. If the constraint is set to $\mathbf{A} = [0, 1]$ (Fig. 3.3b), the hip angle is unconstrained, and the knee angle is constrained to follow the task-space policy $\mathbf{b}$, resulting in variations driven by the task. If the constraint is set to $\mathbf{A} \approx [0.7, 0.7]$ (Fig. 3.3c), the observations are the combination of the null-space policy $\boldsymbol{\pi}$ and the task- space policy $\mathbf{b}$, and both hip and knee show variance in behaviours consistent with the task.

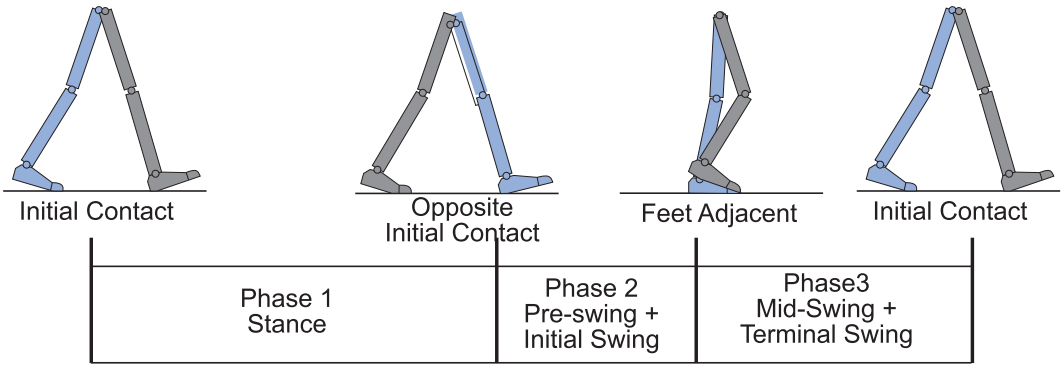

Figure 3.4: An example of three different gait-phases

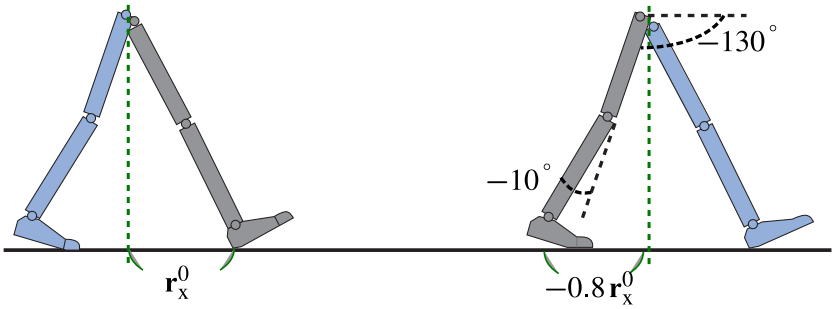

Figure 3.5: An example of stance phase. At the end of the phase, on average, the hip and knee angle are approximately $(-130^\circ, -10^\circ)$, and the horizontal displacement between the heel and the torso is approximately 80% of the step-size.

3.2.2 *An Example of Modelling Walking Phases*

In the preceding section, we presented the formulation of our walking phase model; in the following, we illustrate an example of representing the phases shown in Fig. 3.4.

Please note that, in human walking, we have no prior knowledge about the gait-phases (e.g., what is the best way to divide a gait-cycle?) and the gait parameters (e.g., the constraints $\mathbf{A}$ and the task-space policy $\mathbf{b}$). Our aim in this section is to help the reader understand our model via an example. The choices of phase-divisions and parameters are based on (i) the literature in Winter (1984) and (ii) our interpretation of motion capture data; these selections are not conventional and not necessary optimal.

3.2.2.1 *Phase 1: stance phase*

During the stance phase (Fig. 3.5), the hip flexes to rotate backward and then reverses the rotation right before the end of stance phase. The leg is almost straight so that the stance leg is able support the body weight.

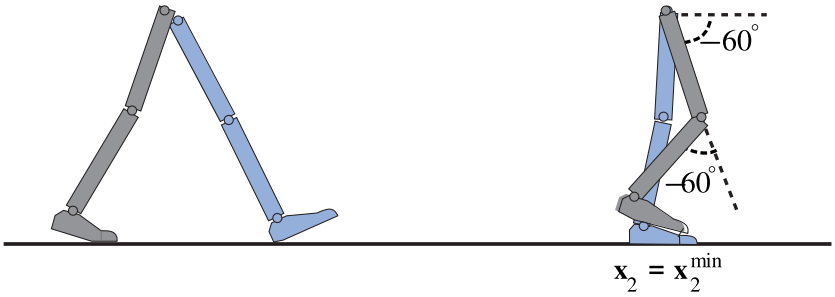

Figure 3.6: From pre-swing phase, the foot lifts off from the ground as the hip joint extends and the knee joint flexes. At the end of the phase, the knee angle $\mathbf{x}_2$ reaches its minimum flexion $\mathbf{x}_2 \approx \mathbf{x}_2^{min}$

This sequence of transitions seem to be a prerequisite for all kind of walking behaviours. We assume that there are some level of redundancy and consistency in both hip and knee angle, yet we have no information about these two quantities. We set the constraint as $\mathbf{A} = [-0.6, 0.8]$ since it is a reasonable choice of unit vector that reconstructs similar movements with our walking data obtained from motion capture. The task-space target is set to $\rho^* = (-130^\circ, -10^\circ)$ since this is a reasonable approximation of the posture at the end of stance phase.

This phase ends when the horizontal displacement between the heel and the torso is approximately 80% of the step-size. We utilise the initial horizontal position of the feet (e.g., this is $\mathbf{r}_x^0$ in Fig. 3.5) and terminate the phase when $\mathbf{r}_x < -0.8\,\mathbf{r}_x^0$.

3.2.2.2 *Phase 2: pre-swing to mid-swing point*

Our second phase combines the pre-swing and the initial-swing phase (Fig. 3.6). During these two phases, the foot lifts off from the ground as the hip joint extends and the knee joint flexes. How fast the knee angle rises is dependent on the speed of walking. For this, we assume that the knee angle is influenced more by the walking speed (e.g., the variations in $\mathbf{x}_2$ is higher than the previous phase). The constraint is defined as $\mathbf{A} = [-0.4, 0.9]$ because this would give more variations to the knee angle and the output behaviour is similar to the motion capture data.

At the end of this phase, the hip joint is almost at its maximum extension and the knee joint is at its maximum flexion. For this, we set the task-space target as $\rho^* = (-60^\circ, -60^\circ)$ since it is a fair posture of maximum hip extension and maximum knee flexion. The knee angle should reach its maximum flexion at the end of this phase, so the phase terminates when $\mathbf{x}_2 < \mathbf{x}_2^{min}$.

Figure 3.7: At the end of swing phase, the foot is in contact on the walking surface $\mathbf{r}_y = 0$

Phase	Constraint $\mathbf{A}$	Task-space Target ρ^*	End Condition
1	$[-0.6, 0.8]$	$(-130^\circ, -10^\circ)$	$\mathbf{r}_1 < -0.8\mathbf{r}_1^0$
2	$[-0.4, 0.9]$	$(-60^\circ, -60^\circ)$	$\mathbf{x}_2 < \mathbf{x}_2^{min}$
3	$[0.1, 0.99]$	$(-40^\circ, 0^\circ)$	$\mathbf{r}_2 < \mathbf{r}_2^0$

Table 3.2: An example of modelling phase critical components by defining parameters $\mathbf{A}$, and ρ^* for each phase

3.2.2.3 *Phase 3: mid-swing point to end of cycle*

In the final phase, the thigh rotates forward, and the knee is once again straightened (Fig. 3.7). The movement of the knee is highly depended on both speed and step-sizes (The position where the heel of the swing leg lands on the surface of walking depends on the step-sizes, and this variation seems to be compensated by the knee angles). We assume that the consistency in $\mathbf{x}_2$ is relatively smaller than the previous two phases, so the constraint in this dimension should be higher. For this reason, the constraint is set to $\mathbf{A} = [0.1, 0.99]$, and this $\mathbf{A}$ seems to resemble the motion capture data as well.

At the end of this phase, the foot is set back on the surface of walking. We denote the height of the left feet as $\mathbf{r}_y$ and the height of the left feet at initial-contact as $\mathbf{r}_y^0$. The terminal condition of this phase is set to $\mathbf{r}_y \leqslant \mathbf{r}_y^0$.

In addition, the joint angles are approximately at the maximum extension at the time of initial-contact. We assume that maximum hip and knee extension are approximately $(-40^\circ,\ 0^\circ)$. For this, the task-space target is set to $\rho^* = (-40^\circ, 0^\circ)$ as a reasonable attractor point.

Under our constraint model, the constraint matrix $\mathbf{A}$, the task-space target ρ^*, and the terminal condition for each phase are summarised in Table 3.2. In Section 4.4.2, we used the example above to validate our method for learning the null-space policy.

3.3 CHOICE OF CONSTRAINED SPACE

As described in Section 2.3, rehabilitation devices differ from their design and control mechanism, such as how the movements are constrained (e.g., kinematic control, force control, foot-position control, etc). The choice of constraints might be a crucial factor in facilitating the improvement after training, but it is not clear what is the optimal way to enforce the constraints. The walking phase model proposed in this chapter can be used to handle a wide variety of representations and constraints. In the following we outline some examples from the kinematic control and balance control.

3.3.1 *Constraint on the Joint-Space*

A simple control scheme that corresponds to the examples discussed in the preceding section is to impose restrictions directly on the joint-space. Namely, the state and action of the system can be represented by the joint angles $\mathbf{x} \equiv \mathbf{q}$, and the joint-velocities $\mathbf{u} \equiv \dot{\mathbf{q}}$, respectively.

Note that, this is similar to the *position control* paradigm in many existing devices for gait rehabilitation discussed in Section 2.3. The subject is instructed to follow the reference trajectories, which is defined in joint space. For this control scheme, the constraint $\mathbf{A} \in \mathbb{R}^{\mathcal{S} \times \mathcal{Q}}$ can be defined as a set of of $\mathcal{S}$ unit vector $\hat{\mathbf{n}}$ where each $\hat{\mathbf{n}} \in \mathbb{R}^{1 \times \mathcal{Q}}$ specifies the constrained joints.

3.3.2 *Constraint on End-effector Space*

An alternative of a constraint-based control scheme is to restrict the end-effector position $\mathbf{r}$. Note that, this is a popular scheme for velocity-based control of rigid-body manipulators discuss in Section 2.4. For example, when grasping an object, the end-effector of the manipulator must be controlled to reach the object.

In this control scheme, we can identify the state $\mathbf{x} \equiv \mathbf{q}$, the action $\mathbf{u} \equiv \dot{\mathbf{q}}$, the task $\mathbf{b}(\mathbf{x}) \equiv \dot{\mathbf{r}}(\mathbf{x})$ specifies the desired end-effector velocities to reach the object, and the constraint $\mathbf{A}(\mathbf{x}) \equiv \mathbf{J}(\mathbf{x})$ where $\mathbf{J}$ is the Jacobian matrix relates the joint space to the end-effector space. This places a constraint on the joint-space velocity such that the manipulator can reach the object. *Jacobian*

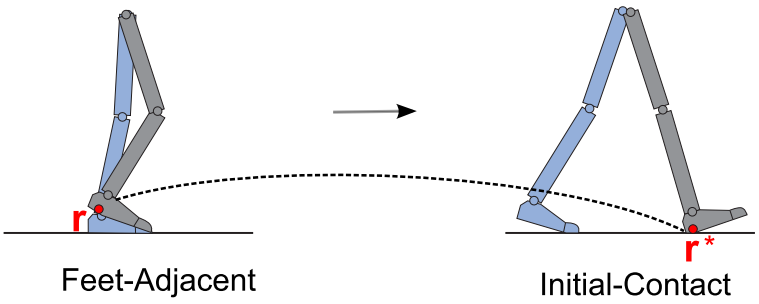

Figure 3.8: An example of using this control scheme is for tracking the trajectories of foot positions so that the heel of the swing leg strikes on the surface of walking at the end of swing phase.

For locomotion or gait rehabilitation, an example of using this control scheme is for tracking the trajectories of foot positions (Fig. 3.8). For instance, at the time of initial-contact event, the heel of the swing leg should strike on the surface of walking. During the swing phase, the foot position of the swing leg should be constrained to ensure the foot is placed on the ground at the end of the terminal-swing phase.

In terms of the 2-link planar system shown in Fig. 3.1, the task can be defined in a way that drives the foot position $\mathbf{r} = [\mathbf{r}_x, \mathbf{r}_y]$ to a desired placement. An example of task-space policy $\mathbf{b}$ could be a point-attractor in the foot position $\mathbf{b}(\mathbf{x}) = \beta^{ts}(\mathbf{r}(\mathbf{x}) - \mathbf{r}^*)$ where $\mathbf{r}^*$ is the target position and β^{ts} is a scaling factor. In this case, β^{ts} and $\mathbf{r}^*$ can very depends on the speed and the step-length. To place a constraint on the joint-velocity, the constraint matrix $\mathbf{A}$ takes the form

$$\mathbf{A} \equiv \mathbf{J} = \begin{bmatrix} \frac{\partial \mathbf{r}_x}{\partial \mathbf{q}_1} & \frac{\partial \mathbf{r}_x}{\partial \mathbf{q}_2} \\ \frac{\partial \mathbf{r}_y}{\partial \mathbf{q}_1} & \frac{\partial \mathbf{r}_y}{\partial \mathbf{q}_2} \end{bmatrix}$$

where $\mathbf{J}$ is the Jacobian that relates the hip and knee angles $\mathbf{q} = [\mathbf{q}_1, \mathbf{q}_2]$ to the foot position.

Another option is to restrict some subspace of the end-effector position. Note that, in the above example, the foot position is restricted to land at a specific point $\mathbf{r}^*$ on the ground (i.e., this would put hard restrictions on the step-size), but this is probably not necessary. We need to ensure that the vertical position of the foot ($\mathbf{r}_y$) is in contact with the surface at the time of initial-contact event, but the foot position in the horizontal direction ($\mathbf{r}_x$) is less critical.

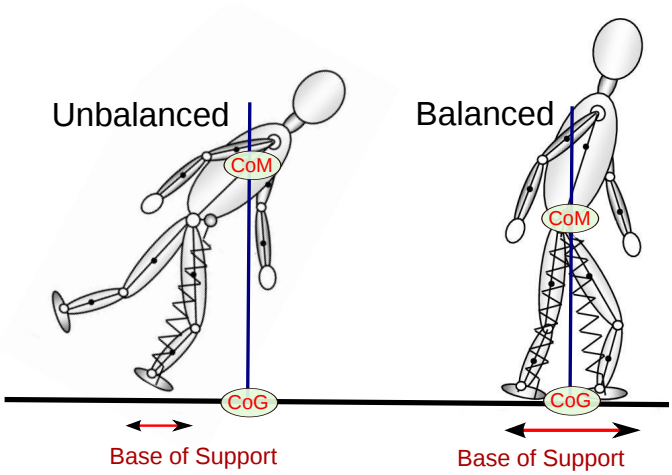

Figure 3.9: If the CoG is closed enough to the base of support, the gait is balanced

One way of characterising such motion is by setting the constraint as $\mathbf{A} \equiv \hat{\mathbf{n}}\,\mathbf{J}$ where $\hat{\mathbf{n}}$ is a vector specifies which dimensions are constrained. Specifically, the constraint matrix can be defined as

$$\hat{\mathbf{n}} = [0, 1] \text{ and } \mathbf{A} \equiv \hat{\mathbf{n}}\,\mathbf{J} = \begin{bmatrix} \frac{\partial\,\mathbf{r}_y}{\partial\,\mathbf{q}_1} & \frac{\partial\,\mathbf{r}_y}{\partial\,\mathbf{q}_2} \end{bmatrix}$$

This would allow motion in the horizontal space provided that the vertical dimension of the foot is tracking the desired trajectory.

3.3.3 *Constraint on Centre-of-Mass*

As described in Section 2.1, maintaining balance is essential during walking. Although the condition for stability is unknown, a common assumption is that the whole body centre-of-mass (CoM) should be controlled within the base-of-support (BoS) (Winter, 1995).

centre-of-mass

The idea is illustrated in Fig. 3.9. The whole-body CoM is a point at which the body is balanced in all directions. In gait analysis, this is a passive variable controlled by the balance control system. The *centre-of-gravity* (CoG) is the vertical projection of the CoM onto the ground. A common assumption in gait analysis community is that the body is balanced if the position of CoG is close enough to BoS. Therefore, the *CoM displacement* (e.g., the distance between CoG and the BoS) is highly correlated to the stability of the subject.

centre-of-gravity

CoM displacement

However, it is not clear what causes the variations and consistency in the whole-body centre-of-mass. As mentioned in Section 2.1, clinicians who use observational

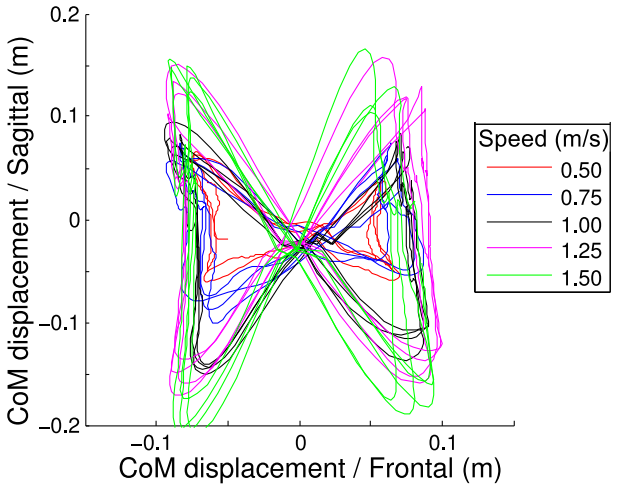

Figure 3.10: Trajectories of CoM displacement taken from a subject walking with 5 different speeds

gait analysis discover that even normal individuals show significant variations in CoM displacement (Orendurff et al., 2004). Note that, this is consistent with the CoM data we collected using force plateforms. Fig. 3.10 shows trajectories of CoM displacement from a subject walking with five different speeds, where the x-axis is the CoM displacement in the frontal plane, and the y-axis is the CoM displacement in the sagittal plane. We can see that, the overall shapes of the trajectories are very similar, but variations appear across different behaviours (especially in the sagittal plane).

Note that we cannot measure the mass and centre of mass of each segment for human subjects, so we cannot calculate CoM directly. Several approaches have been suggested to deal with the CoM estimations, including (i) Approximation from Kinematics (Kingma et al., 1995), (ii) Low Pass Filter (Breniere et al., 1987), and (iii) Gravitational-Line-Projection (Zatsiorsky and King, 1997).

gravitational-line-projection

One way to enforce balance is to set the constraint similar to the ZMP controller discussed in Section 2.2. A simple example is to relate the kinematics and CoM trajectories:

$$\mathbf{x} = \begin{bmatrix} \mathbf{q}_1 \\ \mathbf{q}_2 \end{bmatrix}, \ \mathbf{u} = \begin{bmatrix} \dot{\mathbf{q}}_1 \\ \dot{\mathbf{q}}_2 \end{bmatrix}, \ \mathbf{A} = \hat{\mathbf{n}}\,\mathbf{J}$$

where $\mathbf{x}, \mathbf{u}$ are the joint-angle and joint-velocity, and $\mathbf{J}$ is the Jacobian that relates the joint-velocity to the CoM displacement velocity. However, this is not trivial since the relationship between joint space and the CoM displacement space is undefined (CoM

is approximated). Another way to enforce the balance is to set the state and action space as

$$\mathbf{x} = \begin{bmatrix} \rho_x \\ \rho_z \end{bmatrix}, \ \mathbf{u} = \begin{bmatrix} \dot{\rho}_x \\ \dot{\rho}_z \end{bmatrix}, \ \mathbf{A} = \hat{\mathbf{n}}$$

where ρ_x, ρ_z are the position of CoM-displacement, $\dot{\rho}_x, \dot{\rho}_z$ are the velocity of CoM-displacement in frontal and sagittal plane, respectively, and $\hat{\mathbf{n}}$ specifies the dimension in CoM-displacement to be constrained. In Section 6.4, we will test our approach with these constraints on human walking data.

3.4 DISCUSSION

In this thesis, we assume that the observed walking gait are influenced by natural variations coming from environment, embodiment, and behaviours. In order to recover the essence of the gait, we proposed a formal model for the analysis of walking gait such that the characteristics and variations can be separated.

In this chapter, we describe (i) a *walking phase model* which was built upon previous work in constraint movement analysis (ii) our interpretation of walking gait and how to determine the control variables in this walking phase model, and (iii) example of representations for this model.

Note that, the proposed walking phase model is generic and can be applied to different representations. In Section 3.3, we discussed a few possible variations of constraints; we will use them for our experiments in the later chapters (Section 4.4, Section 5.4, and Section 6). Also, in this thesis, we focus on kinematics representations; however, with some modifications, it can be extended to other representations such as those discussed in Section 3.1.

We will use this walking phase model throughout this thesis. In particular, in Chapter 4, we focus on learning the *characteristics* of walking using this decomposition. In Chapter 5, we apply our walking phase model to compare the difference between walking gaits.

4

GENERALISATION OF A POLICY ACROSS WALKING BEHAVIOURS AND SUBJECTS

A centipede was happy - quite! Until a toad in fun. Said, "Pray, which leg comes after which?" Which threw her mind in such a pitch. She laid bewildered in the ditch. Considering how to run.

—Katherine Craster, Pinafore Poems, 1871 [1]

What defines a normal walking gait is an overloaded question with contradictory interpretations. One way to arrive at a solution is to employ optimisation strategies that somehow capture the essence of walking gait. However, this is not trivial since the walking gait that we observe are masked by intra-personal (e.g., speed, cadence) and inter-personal (e.g., body size, personal preferences) variations.

For instance, Fig. 4.1 shows the observed trajectories of hip joint-angle using motion capture. In Fig. 4.1a, the data was collected from five subjects walking at one meter-per-second, and in Fig. 4.1b, the data was collected from a subject walking with five different speeds ($0.5, 0.75, 1, 1.25, 1.5$ meter-per-second). We can see that, although the subjects were asked to walk in a normal way, there are natural variations that arise from different behaviours.

1 *The Centipede* is a short poem written by Mrs Katherine Craster in *Pinafore Poems*, 1871. The centipede walks easily until it begins to think about walking. Later, this is referred as *The Centipede's Dilemma*. When a simple and unconscious activity is disrupted by consciousness, one may find their performance of the task impaired.

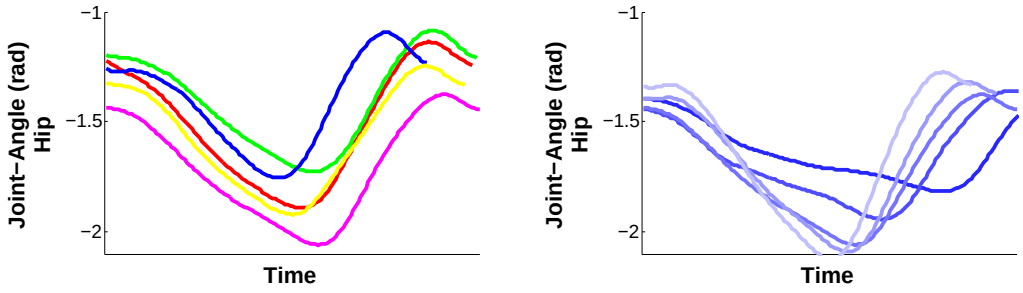

(a) (b)

Figure 4.1: Examples of variations across subjects and walking speeds: (a) Trajectories of five subjects (denoted by colours) walking at 1 meter-per-second. (b) Trajectories of one subject walking with five different speeds ($0.5, 0.75, 1, 1.25, 1.5$ meter-per-second). Although the subjects were asked to walk in a normal way, natural variations arise from different embodiments and behaviours.

As discussed in Section 2.3, a common approach in gait rehabilitation is to assist the patient through a reference gait. This reference gait is normally taken from the average or polynomial fit of walking data from healthy subjects. From Fig. 4.1, we can see that the average framework can be affected the natural variations.

In the preceding chapter, we introduced our walking phase model in which the observed movements are decomposed into the *consistent characteristics* of walking and variations from embodiments and behaviours. In this chapter, we focus on the method to generalise the *consistent characteristics* from various walking demonstrations.

In Section 4.1, we introduce the formulation of learning walking gait and outline the issues of using the standard regression for this problem. In Section 4.2, we propose a method to generalise the *characteristics* of walking by reconstructing an *unconstrained policy*. In Section 4.3, we discuss the necessary condition of our method. In Section 4.4, we validate our method with simulated data.

4.1 PROBLEM FORMULATION: GENERALISING WALKING GAIT

In the following we formulate the problem of generalising the *consistent characteristics* of walking gait across various embodiments and behaviours. We assume that the demonstrated walking movements are given in the form of trajectories; i.e., as $\mathcal{N}$ pairs of observed states $\mathbf{x} \in \mathbb{R}^{\mathcal{P}}$ and observed actions $\mathbf{u} \in \mathbb{R}^{\mathcal{Q}}$. For kinematic control, for example, $\mathbf{x}, \mathbf{u}$ could be joint-angles $\mathbf{x} \equiv \mathbf{q}$ and joint-velocities $\mathbf{u} \equiv \dot{\mathbf{q}}$. In dynami-

cal control, the state might be joint positions and velocities $\mathbf{x} \equiv (\mathbf{q}, \dot{\mathbf{q}})$ with actions corresponding to joint torques $\mathbf{u} \equiv \tau$.

In Section 3.2, we described our working phase model such that the observed actions can be decomposed as

$$\mathbf{u} = \mathbf{A}^{\dagger}\,\mathbf{b} + \mathbf{N}\,\boldsymbol{\pi}$$

where $\boldsymbol{\pi}$ is the *null-space policy*, $\mathbf{b}$ is the task-space policy, $\mathbf{A}$ describes the constraints, and $\mathbf{N}$ is a matrix which projects $\boldsymbol{\pi}$ onto the null-space of $\mathbf{A}$. In this walking phase model, the consistent characteristics of walking corresponds to $\boldsymbol{\pi}$, the variations in embodiments and environment are captured by $\mathbf{A}$, and the variations in behaviours (such as speed) are represented by $\mathbf{b}$.

We assume that (i) $\mathbf{u}$ are generated using the same null-space policy $\boldsymbol{\pi}$, (ii) each observation might have been constrained to accomplish some tasks (that is, $\mathbf{A}\,\mathbf{u} = \mathbf{b}$ for some constraint matrix $\mathbf{A} \neq 0$ and task-space policy $\mathbf{b} \neq 0$), and (iii) $\mathbf{b}$ and $\mathbf{A}$ (and $\mathbf{N}$) are not explicitly known for any given observation.

The objective is to approximate $\boldsymbol{\pi}$, as a mapping from the state space to the action space, and is independent of task and embodiment. By assumption, the task-space policy $\mathbf{b}$ varies across observations, and therefore it could be highly non-linear. The method proposed in Section 4.2 approximates the policy $\boldsymbol{\pi}$ without explicit knowledge of $\mathbf{b}$. For this reason, details of $\mathbf{b}$ are not discussed in this thesis.

4.1.1 *Direct Approach*

Given the observations $\mathbf{x}, \mathbf{u}$, a simple, but naive approach is to apply Direct Policy Learning (DPL) which estimates the policy function $\boldsymbol{\pi}(.)$ using direct regression. Specifically, one seek a direct mapping between state and action space $\mathbf{u} = \boldsymbol{\pi}(\mathbf{x})$ by minimise the error

direct policy learning

$$\mathsf{E}_{\mathrm{DPL}} = \sum_{n=1}^{N} \| \mathbf{u}_n - \tilde{\boldsymbol{\pi}}_n \|^2 \tag{4.1}$$

where $\tilde{\boldsymbol{\pi}}$ is some suitable estimator of the policy function. An effective way to represent $\tilde{\boldsymbol{\pi}}$ is by a linear combination of basis functions $\tilde{\boldsymbol{\pi}}_n(\mathbf{x}) = \mathbf{W}\,\boldsymbol{\beta}(\mathbf{x})$ where $\mathbf{W} \in \mathbb{R}^{M \times M}$ is a matrix of weights, and $\boldsymbol{\beta}(\mathbf{x}) \in \mathbb{R}^{M}$ is a vector of fixed basis functions.

This representation is flexible to represent a class of functions, provided a suitable set of basis function, which can be a linear model where $\boldsymbol{\beta}(\mathbf{x}) = \left[\mathbf{x}^{\top}, 1\right]^{\top}$ or the non-linear model such as Radial Basis Function (RBF)

radial-basis-functions

$$\beta(\mathbf{x}) = \frac{\mathbf{K}(\mathbf{x}-\mathbf{c}_i)}{\sum_{j=1}^{M} \mathbf{K}(\mathbf{x}-\mathbf{c}_j)}$$

where $\mathbf{K}$ is the Gaussian kernels $\mathbf{K} = \mathbf{e}^{-\frac{1}{2\sigma^2}\|\mathbf{x}-\mathbf{c}\|^2}$, M specifies the number of basis functions, and $\mathbf{c}$ are the centres. *Gaussian kernel*

As mentioned earlier, walking gait can be influenced by factors such as embodiments and behaviours, and direct approach results in model-averaging of the walking gait. In terms of our walking phase model, the direct approach ignores the variations in constraints $\mathbf{A}$ (corresponds to embodiments) and task-space policy $\mathbf{b}$ (corresponds to behaviours), and yields the average motion from different trajectories. However, this is unrealistic in everyday behaviour, so minimising (4.1) is unlikely to result in a good model of the gait.

Fig. 4.2 illustrates the effect of variations on the walking gaits and the averaging effect of standard regression. In Fig. 4.2a, the x- and the y-axis are the hip and knee angles, respectively. The blue trajectories are motion capture data collected from a subject walking with five different speeds. We can see that the trajectories vary with the speeds. Direct regression results in averaging of the five trajectories (red) in a way that cannot explain the demonstrated behaviours. Fig. 4.2b illustrates this problem with our walking phase model. The observations $\mathbf{u}_1$ and $\mathbf{u}_2$ (blue) can be considered as the results of a null-space component $\mathbf{u}^{ns}$ (purple) and two different task-space components $\mathbf{u}_1^{ts}$ and $\mathbf{u}_2^{ts}$ (green) from state $\mathbf{x}$. If we apply the direct method, we end up with $\mathbf{u}_{DPL}$ (red), an average between $\mathbf{u}_1$ and $\mathbf{u}_2$

Note that, this corresponds to the default *average template* solution used in many gait and rehabilitation analysis, but this is probably sub-optimal. For example, the demonstrations from healthy subjects could be taken from various walking speeds. If we use direct approach, the resulting model would consider faster or slower walks as deviations from a normal gait. Clearly, a method that can uncover the characteristics and ignore the variations is needed. In the next section, we propose to resolve this issue by modifying the functional in (4.1).

4.2 NULL-SPACE POLICY LEARNING FOR WALKING GAIT

Considering the analysis in the previous section, we take an alternative approach, in which the constraint and task-space variations are explicitly considered. The proposed approach builds on previous research on Null-space Policy Learning (NPL) for point- *null-space policy learning*

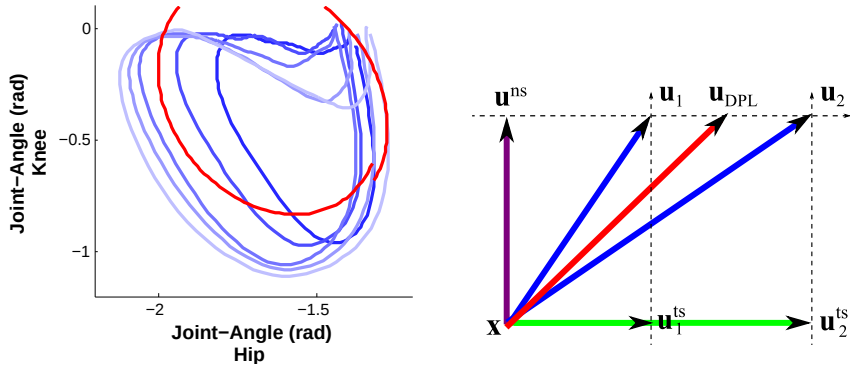

(a) (b)

Figure 4.2: Illustration of the effect of variations on the walking gaits and the averaging effect of standard regression. (a) When a subject walks with different speeds, the observed hip angles (blue) vary. Direct regression results in averaging of the five trajectories in a way that cannot explain the demonstrations (red). (b) The observations $\mathbf{u}_1$ and $\mathbf{u}_2$ (blue) can be considered as the results of two different task-space components $\mathbf{u}_1^{ts}$ and $\mathbf{u}_2^{ts}$ (green) from state $\mathbf{x}$. If we directly optimising E_{DPL} (4.1), we end up with $\mathbf{u}_{DPL}$ (red), an average between $\mathbf{u}_1$ and $\mathbf{u}_2$.

to-point movements (Towell et al., 2010), and adapts it to generalise the characteristics of gaits. The key additional challenges in walking tasks are (i) the differences across subjects due to interpersonal variations and (ii) the temporal switching of constraints between phases.

In the proposed approach, the policy $\boldsymbol{\pi}$ is estimated using two separate steps. The first step is to decompose the observations $\mathbf{u}$ into two orthogonal components: the task-space component $\mathbf{u}^{ts} \equiv \mathbf{A}^{\dagger}\mathbf{b}$ and the null-space component $\mathbf{u}^{ns} \equiv \mathbf{N}\boldsymbol{\pi}$ such that $\mathbf{u} = \mathbf{u}^{ts} + \mathbf{u}^{ns}$. The second step is to reconstruct the null-space policy $\boldsymbol{\pi}$ from the estimated $\mathbf{u}^{ns}$. The following describes the two steps in detail.

task-space component

null-space component

4.2.1 *Step-1: Learning Null-space Components*

The first step is to extract the null-space component $\mathbf{u}^{ns}$ from the raw observations $\mathbf{x}, \mathbf{u}$. As discussed in (Towell et al., 2010), a requirement on this step is that the data are grouped into multiple subsets such that the constraint matrix $\mathbf{A}$ is consistent within each subset. In the present setting, this separation arises naturally from consideration of the different phases of the gait.

Specifically, since the task-constraints might be different across different persons and phases, each phase of each subject is considered an independent subset for the purposes of Step-1. Given input data as pairs of states and actions $\mathbf{D} = \{\mathbf{x}, \mathbf{u}\}$, the data set is divided into i subsets such that subset $\mathbf{D}_i$ is the set of observations from subject i. Then, each $\mathbf{D}_i$ is further divided into k subsets $\mathbf{D}_{i,k} = \{\mathbf{x}_{i,k}, \mathbf{u}_{i,k}\}$ where $k = \{1, 2, 3, ...\}$ denotes the phase number.

For the $(i, k)^{th}$ data subset, we seek a model that minimises the inconsistency between the true null-space component $\mathbf{u}^{ns}_{i,k}$ and the estimated null-space component $\tilde{\mathbf{u}}^{ns}_{i,k}$.

$$E[\mathbf{u}^{ns}_{i,k}, \tilde{\mathbf{u}}^{ns}_{i,k}] = \sum_{n=1}^{\mathcal{N}_{i,k}} \| \mathbf{u}^{ns}_{i,k,n} - \tilde{\mathbf{u}}^{ns}_{i,k,n} \|^2 \tag{4.2}$$

Since we do not have access to the true null-space component $\mathbf{u}^{ns}$, (4.2) cannot be directly optimised. Instead, we attempt to eliminate the components of motion that are due to the task constraints, and learn a model that is consistent with the observations. To achieve this, we seek a projection matrix $\mathbf{P}_{i,k} = \frac{\tilde{\mathbf{u}}^{ns}_{i,k} \tilde{\mathbf{u}}^{ns\top}_{i,k}}{\| \tilde{\mathbf{u}}^{ns}_{i,k} \|^2}$ which projects $\mathbf{u}_{i,k}$ onto the learnt null-space component and satisfies $\mathbf{P}_{i,k}\, \mathbf{u}_{i,k} \equiv \mathbf{P}_{i,k}(\mathbf{u}^{ts}_{i,k} + \mathbf{u}^{ns}_{i,k}) = \tilde{\mathbf{u}}^{ns}_{i,k}$. The objective function (4.2) can be rewritten in terms of this projection as:

$$E_1[\tilde{\mathbf{u}}^{ns}_{i,k}] = \sum_{n=1}^{\mathcal{N}_{i,k}} \| \frac{\tilde{\mathbf{u}}^{ns}_{i,k,n} (\tilde{\mathbf{u}}^{ns}_{i,k,n})^T}{\| \tilde{\mathbf{u}}^{ns}_{i,k,n} \|^2} \mathbf{u}_{i,k,n} - \tilde{\mathbf{u}}^{ns}_{i,k,n} \|^2 \tag{4.3}$$

Fig. 4.3a illustrates an example of this idea. Assuming there are two observation $\mathbf{u}_1$, $\mathbf{u}_2$ from the same $\mathbf{x}$ within a subset, since the constraint is consistent, $\mathbf{u}_1$, $\mathbf{u}_2$ must have the same $\mathbf{u}^{ns}$. We seek a $\mathbf{u}^{ns}$ such that when $\mathbf{u}_1$ and $\mathbf{u}_2$ are projected onto $\mathbf{u}^{ns}$, the error is minimised.

In this thesis, each $\mathbf{u}^{ns}_{i,k}$ is modelled through iterative optimisation of (4.3) using Gaussian radial basis functions. More precisely, $\tilde{\mathbf{u}}^{ns}_{i,k} = \mathbf{W}_{i,k}\, \boldsymbol{\beta}(\mathbf{x}_{i,k})$ where $\mathbf{W}_{i,k} \in \mathbb{R}^{d \times M}$ is a matrix of weights, and $\boldsymbol{\beta}(\mathbf{x}_{i,k}) = \frac{K(\mathbf{x}_{i,k} - \mathbf{c}_m)}{\sum_{m=1}^{M} K(\mathbf{x} - \mathbf{c}_m)} \in \mathbb{R}^M$ is a vector of basis functions, M is the number of basis functions, and $\mathbf{c}_m$ for $m = 1, ..., M$ are the centres. The optimisation is initialised using direct regression to find the initial approximation $\mathbf{W}^0_{i,k} = \arg\min \sum \| \mathbf{u}_{i,k} - \tilde{\mathbf{u}}^{ns}_{i,k} \|^2$.

4.2.2 *Step-2: Learning Null-space Policies*

The output of Step-1, is a set of $i \times k$ intermediate models for the null-space component $\tilde{\mathbf{u}}^{ns}_{i,k} \approx \mathbf{N}_{i,k}\, \boldsymbol{\pi}$. The goal of the second step is to generalise from these to find an

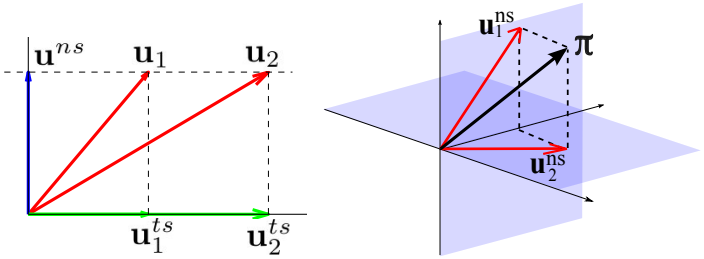

(a) (b)

Figure 4.3: A schematic of the proposed learning algorithm. (a) The goal of *Step-1* is to decompose the observations $\mathbf{u}_1$ and $\mathbf{u}_2$ into two orthogonal components: the null-space component $\mathbf{u}^{ns}$ and the task-space component $\mathbf{u}_1^{ts}$, $\mathbf{u}_2^{ts}$. (b) The goal of *Step-2* is to find an approximate null-space policy $\boldsymbol{\pi}$ that is consistent with all of the null-space component $\mathbf{u}_1^{ns}$ and $\mathbf{u}_2^{ns}$.

approximate policy $\tilde{\boldsymbol{\pi}}$ that is consistent with all of the estimated null-space component $\tilde{\mathbf{u}}_{i,k}^{ns}$. Ideally, the approximation should minimise the error between the true policy and the learnt policy

$$E[\boldsymbol{\pi}, \tilde{\boldsymbol{\pi}}] = \sum_{n=1}^{N} \| \boldsymbol{\pi}_n - \tilde{\boldsymbol{\pi}}_n \|^2. \tag{4.4}$$

Unfortunately, since the true policy $\boldsymbol{\pi}$ is not observed, (4.4) cannot be minimised directly.

Instead, we proceed by noting that, on completion of Step-1 we have the equivalent to a set of $i \times k$ systems that satisfy $\mathbf{A}_{i,k}\, \tilde{\mathbf{u}}_{i,k}^{ns} = 0$. As a result we can adapt the work in Howard et al. (2009) to find a policy that is maximally consistent with the observations.

More precisely, the $i \times k$ intermediate models are combined into a single dataset $(\mathbf{x}, \tilde{\mathbf{u}}^{ns})$, and the approximation is made by minimising the *inconsistency error*

$$E_2[\tilde{\boldsymbol{\pi}}] = \sum_{n=1}^{N} \| \frac{\tilde{\mathbf{u}}_n^{ns}\, \tilde{\mathbf{u}}_n^{ns\,\top}}{\| \tilde{\mathbf{u}}_n^{ns} \|^2} \tilde{\boldsymbol{\pi}}(\mathbf{x}_n) - \tilde{\mathbf{u}}_n^{ns} \|^2 \tag{4.5}$$

An example is illustrated in Fig. 4.3b. Given two (or more) null-space components $\mathbf{u}_1^{ns}$ and $\mathbf{u}_2^{ns}$, the inconsistency error favours models for which there is minimal discrepancy between $\mathbf{u}_1^{ns}$ and $\mathbf{u}_2^{ns}$ and the model, projected onto these observations.

In this thesis, the null-space policy $\tilde{\boldsymbol{\pi}}$ is modelled with Gaussian radial basis functions, through minimisation of (4.5). The entire process of estimating $\boldsymbol{\pi}$ for periodic gaits is summarised in Algorithm 1.

Algorithm 1 Null-space Policy Learning

Input: $\mathbf{D} = \{\mathbf{x}, \mathbf{u}\}$: data-set of states $\mathbf{x}$ and action $\mathbf{u}$

Output: $\tilde{\boldsymbol{\pi}}$: learnt null-space policy

1: Split $\mathbf{D}$ into $\mathbf{D}_i$ where $\mathbf{D}_i$ is the input from subject i
2: Split $\mathbf{D}_i$ into $\mathbf{D}_{i,k}$ where k denotes the phase number
3: **for all** $\mathbf{D}_{i,k}$ **do**
4: Learn $\tilde{\mathbf{u}}^{ns}_{i,k}$ by minimising (4.3)
5: **end for**
6: Combine $\{\mathbf{x}_{i,k}, \tilde{\mathbf{u}}^{ns}_{i,k}\}$ into a single data-set $\{\mathbf{x}, \tilde{\mathbf{u}}^{ns}\}$
7: Learn $\tilde{\boldsymbol{\pi}}$ by minimising (4.5)

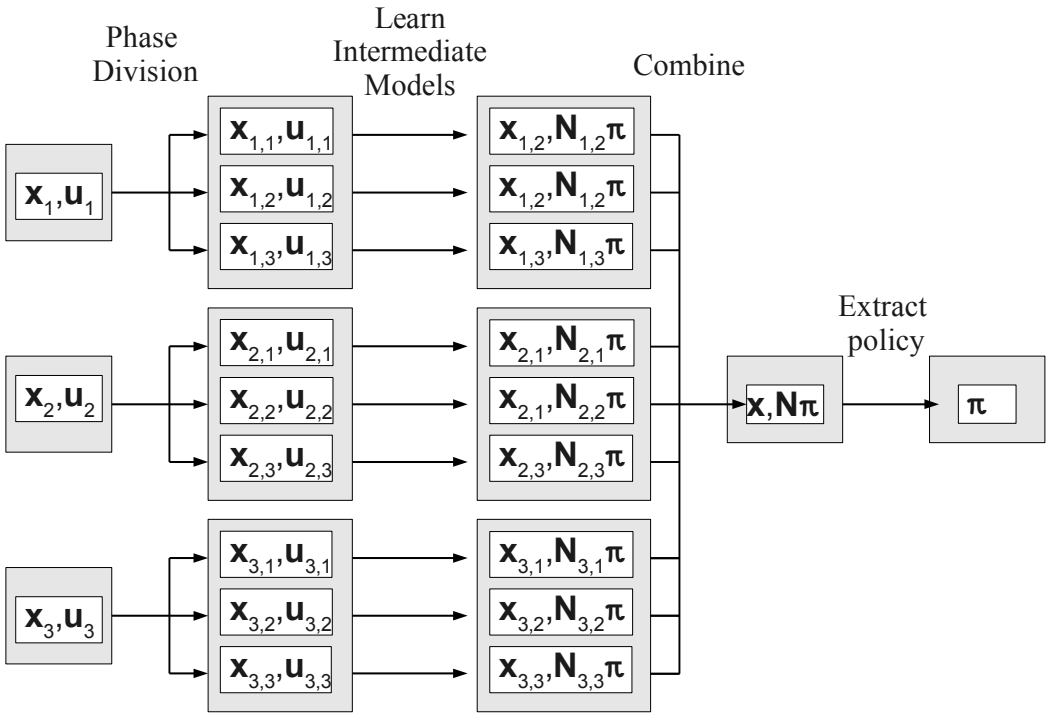

Figure 4.4: An example of learning null-space policy from 3 subjects

Fig. 4.4 shows an example of learning null-space policy from 3 subjects and 3 walking phases. Given data as pairs of $(\mathbf{x}_i, \mathbf{u}_i)$ for $i = 1, 2, 3$, the data from each subject is divided into 3 subsets $(\mathbf{x}_{i,j}, \mathbf{u}_{i,j})$ where $j = 1, 2, 3$ denotes the phase number. We learn an intermediate model by minimising (4.3). The results were combined into a single dataset $(\mathbf{x}, \mathbf{u}^{ns})$, and the null-space policy $\boldsymbol{\pi}$ is extracted by optimising (4.5).

4.3 REQUIREMENTS OF LEARNING CONSISTENT POLICY

In this section, we discuss the requirements and feasibility of learning a consistent policy across subjects.

4.3.1 *Consistent Projection Matrix in Step-1*

The goal of Step-1 is to learn the null-space components $\mathbf{u}^{\text{ns}}$ for estimation of $\boldsymbol{\pi}$ in Step-2. For this to be effective, $\mathbf{u}^{\text{ns}}$ must be consistent within the data set used for estimation. This implies that $\boldsymbol{\pi}$ and the projection matrix $(\mathbf{I} - \mathbf{A}^{\dagger}\,\mathbf{A})$ should both be consistent. Since $\boldsymbol{\pi}$ is expected to capture the characteristics of walking that are consistent across tasks and embodiments, the main variation in observations is expected to come from variations in the constraints $\mathbf{A}$.

Variations in $\mathbf{A}$ may arise due to variations in embodiments and phases. Consistency of $\mathbf{A}$ is ensured, then by breaking the data into independent subsets according to subject and phase.

4.3.2 *Variation of Tasks in Step-1*

The second requirement of Step-1 is that multiple $\mathbf{u}$ are observed at (or near) the same $\mathbf{x}$. This is because, if only one state-action pair $(\mathbf{x}, \mathbf{u})$ is observed, there is more than one way to decompose $\mathbf{u}$ into orthogonal components. $\tilde{\mathbf{u}}^{\text{ns}}$ can be anything that satisfies $\tilde{\mathbf{u}}^{\text{ns}} \approx \frac{\tilde{\mathbf{u}}^{\text{ns}}\,\tilde{\mathbf{u}}^{\text{ns}\,\top}}{\|\,\tilde{\mathbf{u}}^{\text{ns}}\,\|}\,\mathbf{u}$.

For example, consider the situation in Fig. 4.5. In Fig. 4.5a, since only $\mathbf{u} = [1, 1]^{\top}$ is observed, $\mathbf{u}^{\text{ns}}$ can be either $[0, 1]^{\top}$ or $[1, 0]^{\top}$. In Fig. 4.5b, if both observations $\mathbf{u}_1$ and $\mathbf{u}_2$ are observed, $\mathbf{u}^{\text{ns}}$ can be determined.

If the constraint $\mathbf{A}$ and the policy $\boldsymbol{\pi}$ are consistent within each subset, the variation of $\mathbf{u}$ depends on task-space policy $\mathbf{b}$. This suggests that, in order to determine a consistent $\boldsymbol{\pi}$, data should be collected containing a sufficient variation in tasks, for example by asking subjects to walk at different speeds or with different step-sizes.

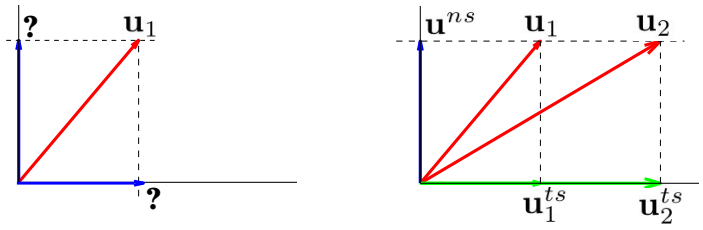

(a) (b)

Figure 4.5: Example of different variations in tasks: (a) if only $\mathbf{u}_1$ is observed, there are more than one way to decompose $\mathbf{u}^{\text{ts}}$ and $\mathbf{u}^{\text{ns}}$. (b) By observing $\mathbf{u}_1$ and $\mathbf{u}_2$, $\mathbf{u}^{\text{ns}}$ can be determined

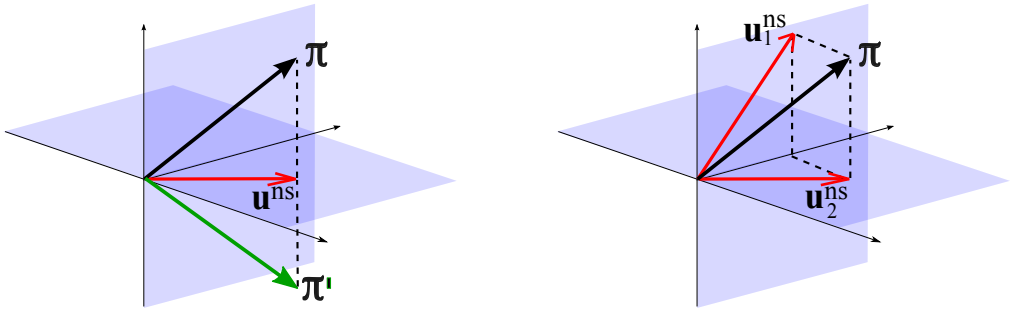

(a) (b)

Figure 4.6: In (a), if only one $\mathbf{u}^{\text{ns}}$ is observed, there are more than one solution for $\boldsymbol{\pi}$. In (b), observing multiple $\mathbf{u}^{\text{ns}}$ determines an unique $\boldsymbol{\pi}$. (Figures were modified from Howard et al. (2009))

4.3.3 *Variation of Projection Matrix in Step-2*

The goal of Step-2 is to learn a consistent $\boldsymbol{\pi}$ given $\mathbf{u}^{\text{ns}}$. A requirement for this is that multiple $\mathbf{u}^{\text{ns}}$ are observed at (or near) the same $\mathbf{x}$.

For example, consider the scenarios depicted in Fig. 4.6. In Fig. 4.6a, if we only have one observation $\mathbf{u}^{\text{ns}}$, there is more than one way to estimate $\boldsymbol{\pi}$. In fact, any vector orthogonal to $\mathbf{u}^{\text{ns}}$ can be a solution (e.g., $\boldsymbol{\pi}'$). In contrast, if multiple $\mathbf{u}^{\text{ns}}$ are observed (Fig. 4.6b), $\boldsymbol{\pi}$ can be determined (Howard and Vijayakumar, 2007).

In the setting considered in this work, variations in the projection matrix $\mathbf{N}$ arise both from variations in the constraint matrix $\mathbf{A}$ across phases, as well as differences in the embodiment of subjects.

Step	Requirement	Parameters	Gait
1	Consistent $\mathbf{N}$	Consistent $\mathbf{A}$	Consistent subject Consistent phase
1	Various $\mathbf{b}$	Various $\mathbf{x}_0$ Various ω	Various step-sizes Various speed
2	Various $\mathbf{u}^{ns}$	Various $\mathbf{A}$	Various subjects Various phases

Table 4.1: Requirements of null-space policy learning

The requirements of null-space policy learning and the ways in which they are fulfilled are summarised in Table 4.1. If the requirements are met, a proof of convergence is given in Appendix A.

4.4 VALIDATION

To illustrate the application of the null-space policy learning approach just described to gait analysis, here we briefly present some numerical results validating its performance. The parameters for these tests (i.e., embodiments, speed, step-size) are chosen so that they are similar to our human data (details will be provided in Section 6.1.1).

4.4.1 *Evaluation Criteria*

The quality of a policy is evaluated by the following criteria:

4.4.1.1 *Unconstrained policy error*

Our primary evaluation criteria is the Normalised Unconstrained Policy Error (UPE), which directly compares the true and the learnt policy

normalised unconstrained policy error

$$\text{UPE} = \frac{1}{\mathcal{N}\sigma_\pi^2} \sum_{n=1}^{\mathcal{N}} \| \boldsymbol{\pi}_n - \tilde{\boldsymbol{\pi}}_n \|^2 \tag{4.6}$$

where σ_π is the standard deviation of the true policy.

4.4.1.2 *Constrained Policy Error*

In some cases, it may be that the variation in constraints is insufficient to fully uncover the true policy $\boldsymbol{\pi}$. However, in such cases, where the constraints exhibit little variation, there may be no need to uncover the hidden components of the fully unconstrained policy (since those components are anyway eliminated by the constraints in normal circumstances). In such circumstances, an alternative quality measure is the Normalised Constrained Policy Error (CPE)

normalised constrained policy error

$$\mathrm{CPE} = \frac{1}{\mathcal{N}\sigma_{\boldsymbol{\pi}}^2} \sum_{n=1}^{\mathcal{N}} \| \mathbf{u}_n^{ns} - \mathbf{N}_n \tilde{\boldsymbol{\pi}}_n \|^2 \tag{4.7}$$

that measures the difference between the data and the estimated policy, when the latter is projected by the same constraints as in the training data.

4.4.1.3 *Normalised Mean-Squared Error*

In many everyday behaviours, we will not have access to the true projections nor the true policy. An alternative is to evaluate how well our learnt projection can reproduce the demonstrated motion without any prior knowledge. Assuming that the policy has been estimated in some way, we measure the distance between the observations and the estimated policy subject to the estimated constraints, namely,

$$\mathrm{E}_{\mathrm{MSE}} = \frac{1}{\mathcal{N}\sigma_{\boldsymbol{\pi}}^2} \sum_{n=1}^{\mathcal{N}} \|\mathbf{u}_n - \tilde{\mathbf{N}}\,\tilde{\boldsymbol{\pi}}\|^2. \tag{4.8}$$

4.4.2 *Experiment on Simulated Data*

To illustrate the application of the null-space policy learning approach just described to gait analysis, here we briefly present some numerical results validating its performance. The parameters for these tests (i.e., embodiments, speed, step-size) are chosen so that they are similar to our human data (details will be provided in Section 6.1.1).

4.4.2.1 *Data*

To simulate recordings of walking data from multiple, healthy subjects, data was collected from a set of simulated 2-link systems. The subjects were assumed to (i) have different embodiments and (ii) perform different tasks.

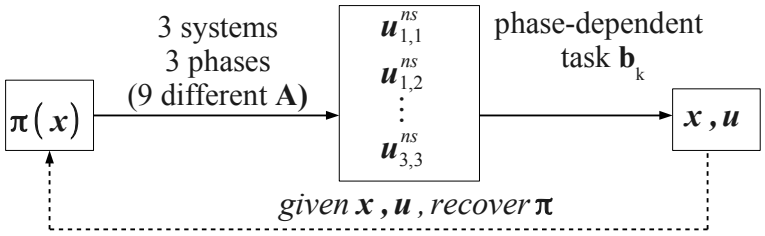

Figure 4.7: Decomposition of simulation data

Specifically, three 2-link systems with different link-lengths were employed, to represent three subjects with different embodiments and each subject (2-link system), was recorded walking at three different speeds ω to represent slow, normal and fast walking. The phase-divisions and task-constraints were chosen based on the example described in Section 3.2.2. The initial joint-angle $\mathbf{x}_0$ of each cycle was chosen to match our human data (ref. Section 6.1.1) hence, each gait cycle had a slightly different task $\mathbf{b}$. As ground truth null-space policies, we considered:

(a) a linear policy: $\boldsymbol{\pi} = \beta^{ns}(\mathbf{x} - \mathbf{x}^*)$ where $\mathbf{x}^* = (-90^\circ, -25^\circ)$ was chosen as a 'comfort' position to which the system tends to track.

(b) a limit-cycle policy: $\dot{r} = r(\rho - r^2)$ with radius $\rho = 0.1$, angular velocity $\dot{\theta} = -2$ rad/s, where r and θ were the polar representation of the state (i.e., $\mathbf{x} = (r\cos(\theta), r\sin(\theta))$).

(c) the same linear policy as (a), but the constraints were imposed in the end-effector space. The constraints were slightly modified to simulate walking on various slopes.

4.4.2.2 *Learning*

Fig. 4.7 shows the generative model of the data. The inputs are pairs of $\mathbf{x}, \mathbf{u}$, and the goal is to recover the policy $\boldsymbol{\pi}$.

The null-space component $\mathbf{u}^{ns}_{i,k}$ was learnt for each phase of each 2-link system by minimising (4.3), which yields 9 different models. Each $\mathbf{u}^{ns}_{i,k}$ was modelled using a set of Gaussian radius basis functions (RBFs), where the number of RBFs N was obtained through cross-validation for $N \in [10, 50]$. The centres were chosen according to k-means and the widths σ^2 were the mean distance between centres.

After Step-1, all $(\mathbf{x}_{i,k}, \tilde{\mathbf{u}}^{ns}_{i,k})$ were combined as the input for Step-2. The null-space policy was modelled separately with another 50 Gaussian RBFs.

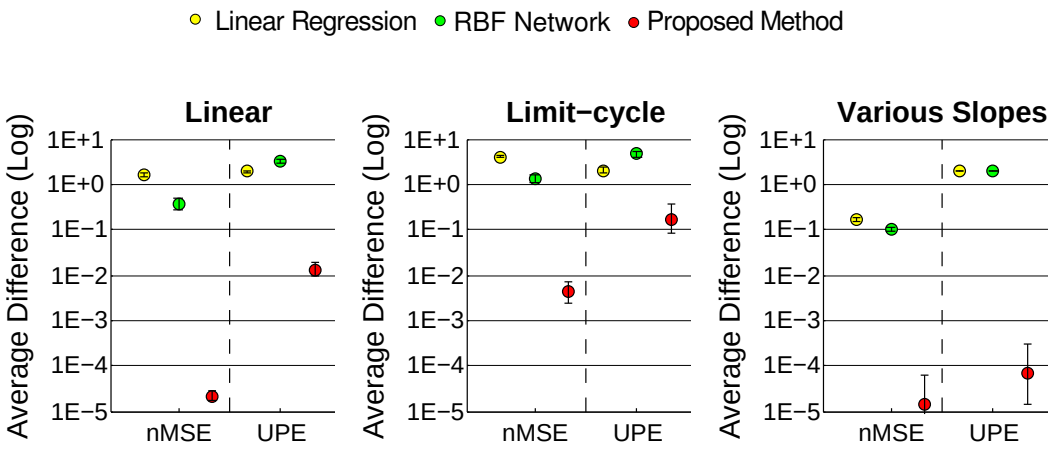

(a) (b) (c)

Figure 4.8: Errors (mean ± 3 std.dev.) of recovering (a) a linear policy, (b) a limit-cycle policy, and (c) a linear policy under various environments.

In the following, the 10-fold cross-validation results were reported when using 90% of the data-set for training the models and reserving 10% for testing. The performance of our proposed method was evaluated based on the criteria discussed in Section 4.4.1. We compared our approach with (i) linear regression and (ii) RBF network.

4.4.2.3 *Results*

Fig. 4.8 summarises the results of recovering (a) the linear policy, (b) the limit-cycle policy, and (c) the linear policy with various slopes.

In all sub-figures, the measurements on the left are the root-mean squared error in joint space (nMSE) and the measurements on the right are the unconstrained policy error (UPE). The error bars are mean±std.dev. in log scale over ten experiments on a hold-out data-set. (Note that standard deviation looks larger at the points lower in the plot due to the log scale.)

From this evaluation, we can see that the proposed method is more accurate in predicting the policy both in terms of the nMSE and UPE. Furthermore, since UPE is a direct comparison between the true and the learnt policy, the proposed method is expected to be more accurate under different task-constraints and behaviours.

Fig. 4.9 is a visualisation of the true and recovered limit-cycle behaviours over multiple gait cycles. Each gait cycle takes a different speed and step-size. (In the following figures, higher speeds are represented by darker colours and vice versa.) These two rows show the hip angle and the knee angle, respectively.

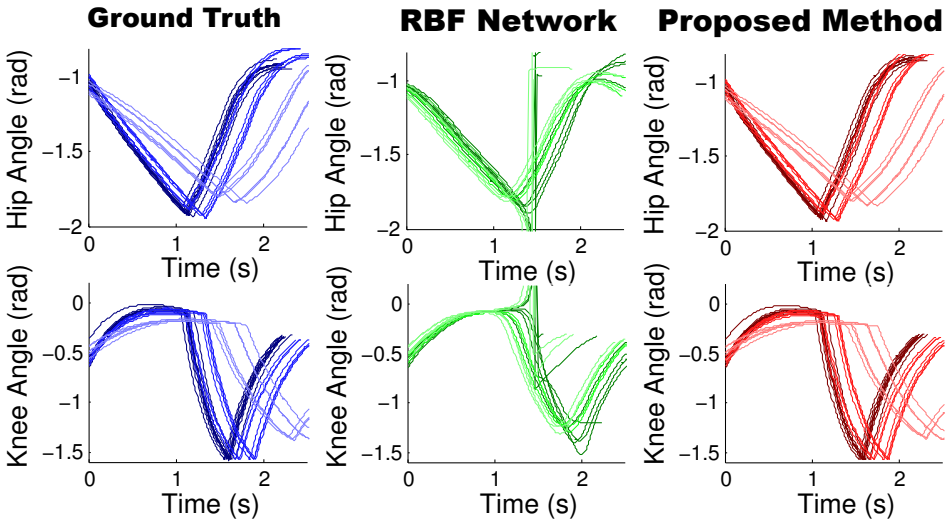

Figure 4.9: Motion generated from (a) the true limit-cycle policy, (b) the learnt policy using RBF network, and (c) the learnt policy using our proposed method

Fig. 4.9a are the motion using the true limit-cycle policy. RBF network (Fig. 4.9b) fails to generate smooth motion, while our method (Fig. 4.9c) reproduces motion that has excellent consistency with the ground truth. Linear regression is worse even than RBF network, so it is omitted from this visualisation. Note that, for RBF network, some initial conditions lead to wrong motion around heel-off. This is probably caused by the difference in these two phases.

In Fig. 4.10, we show an example trajectory. Fig. 4.10a is the motion generated from the limit-cycle policy. Fig. 4.10b is the result of our method, and Fig. 4.10c is the result of RBF network. We picked an example such that RBF network can make decent prediction around heel-off. However, the resulting walker has much smaller knee flexion that the truth data.

4.5 SIMULATED DATA USING PASSIVE DYNAMIC PRINCIPLES

In the last experiment, we tested our method on simulated data which were generated from a constraint model outline in Section 2.4.2 (i.e., the data can be decomposed into a consistent policy, constraints, and tasks). To explore the limitation of the proposed method, we would like to test data that does NOT meet the prerequisites of our method (listed in Section 4.3). In particular, we chose a model based on passive dynamic principles, and see how well the proposed method can generalise the dataset.

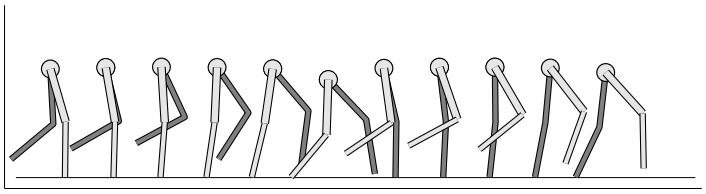

(a)

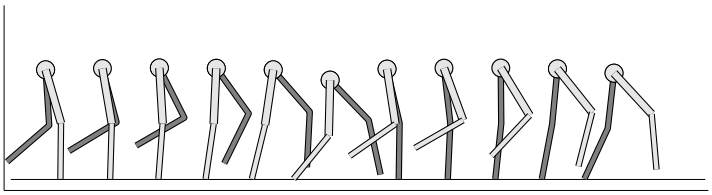

(b)

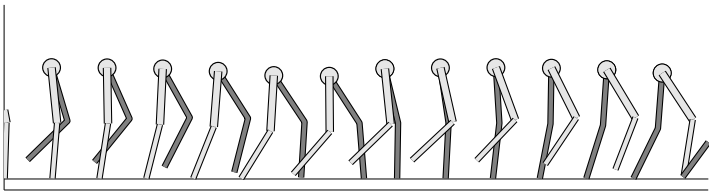

(c)

Figure 4.10: An example of a walking cycle generated from (a) the true limit-cycle policy, (b) the learnt policy using RBF network, and (c) the learnt policy using our proposed method

4.5.1 *Compass-Gait Walker*

The simulated data was based on a compass-gait walker (Goswami et al., 1998). The first compass-gait walker consists of only a swing leg and a stance leg (no knee); to achieve a more anthropomorphic gait, previous work extended the original compass-gait walker by adding two knees (Chen, 2007). The structure of the model is sketched in Fig. 4.11, where m_h, m_t, m_s denote the mass at hip, thigh, and shin, l_t, l_s denote the length of thigh and shin, and b_1, b_2 specify the distance from the beginning of the segment to the centre-of-mass of that segment.

The dynamics of an unactuated compass-gait walker can be expressed as

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}(\mathbf{q}) = 0 \tag{4.9}$$

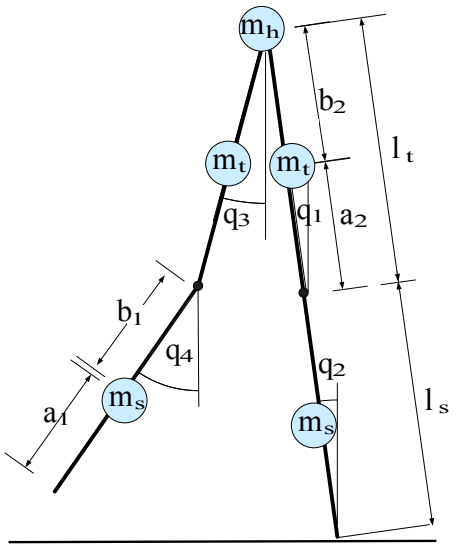

Figure 4.11: Structure of a compass-gait walker with two point masses in each leg

where $\mathbf{M}(\mathbf{q})$ is the inertia matrix, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$ is the Coriolis/centripetal vector, and $\mathbf{G}(\mathbf{q})$ is the gravity vector.

At the beginning of the cycle, the stance leg flexes like an inverse pendulum with two point masses, and the swing leg is acting like a double pendulum. When the swing thigh and shin are aligned, the knee of the swing leg is locked to avoid over extension (this is also referred to as the 'knee-strike' event). When the swing leg hits the ground, the swing and stance leg switch instantaneously.

We collected data from three compass-gait models to represent three persons. Each model has a different body mass and leg lengths, and the values are taken from three human subjects (in Table 6.2). The mass distribution m_h, m_t, m_s and centre-of-mass location b_1, b_2 cannot be measured directly, so these values were approximated using the anthropomorphic table (Winter, 2009). For each compass-gait model, we generated 200 gait cycles. To create different step-sizes and speeds, the initial positions and velocities for were randomly generated within the basin of attractions of the models.

4.5.2 *Learning Reference Policies*

For our method to work well, one requirement is to divide the gait cycle into smaller subsets where the constraints switch. We first divided a gait cycle into stance and

swing phase. For the stance phase, the state of the system $\mathbf{x}$ are joint-angles and joint-velocities, and action $\mathbf{u}$ are joint-accelerations

$$\mathbf{x} \in \mathbb{R}^4 = \begin{bmatrix} q_1 \\ q_2 \\ \dot{q}_1 \\ \dot{q}_2 \end{bmatrix}, \quad \mathbf{u} \in \mathbb{R}^2 = \begin{bmatrix} \ddot{q}_1 \\ \ddot{q}_2 \end{bmatrix}$$

and for the swing phase,

$$\mathbf{x} \in \mathbb{R}^4 = \begin{bmatrix} q_3 \\ q_4 \\ \dot{q}_3 \\ \dot{q}_4 \end{bmatrix}, \quad \mathbf{u} \in \mathbb{R}^2 = \begin{bmatrix} \ddot{q}_3 \\ \ddot{q}_4 \end{bmatrix}$$

We recognised that the dynamics are quite different before/after knee-strike, so we further divided the swing phase at this point. We repeated the same experiment (Section 4.4.2.2) on this dataset. Since this dataset was not drawn from a constraint model, the true null-space policy is not defined. For this reason, we cannot evaluate the model based on UPE nor CPE, and the evaluation is based on behaviour level.

Fig. 4.12 shows an example trajectory generated, and Fig. 4.12a is the true data generated from the compass-gait walker. Fig. 4.12c is the result of direct regression. We can see that, direct policy learning (Fig. 4.12c) failed to imitate the compass-gait walker.

Fig. 4.12b is the result of our method, the error in joint-space is relatively higher than the last experiment by direct observations. On average, our proposed method estimate the null-space policy with the optimisation error $E_2 = 0.0858 \pm 0.0689$ (4.5). In terms of E_2, the optimisation method cannot fit a perfect model for compass-gait data. It is possible that the constraints are not consistent within a single phase or the policy is not consistent across observations. Given that the requirements for our algorithm are not met, we consider this is reasonable approximation.

4.6 DISCUSSION

In our work, we hypothesises that walking can be described as a combination of some consistent characteristics and variations across different embodiments and behaviours.

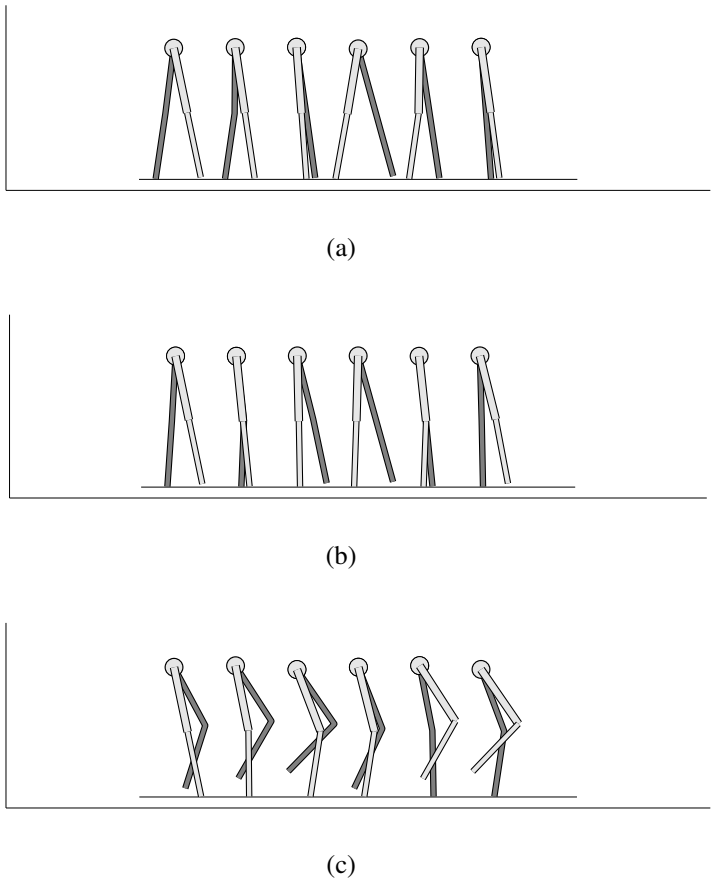

Figure 4.12: An example of walking cycle generated from (a) compass-gait model, (b) model learnt using the proposed method, and (c) model learnt using direct regression

In the preceding chapter, we introduced a walking phase model such that the characteristics and variations can be separated. In this chapter, we explored the problem of generalising the underlying consistent characteristics.

Our approach is built upon previous work in null-space policy recovering; in particular, we learn the characteristics of walking by recovering an *unconstrained policy* (Section 4.2). We outline the requirements for the proposed method to be effective and the verify the theoretical aspect of our approach (Section 4.3).

In Section 4.4, simulation results have shown that our method is effective in reconstructing the policy, even if the true policy, the constraints, and the variations in behaviours are unknown. In Section 4.5, we also shown that the proposed method can approximate well even if the input data does not meet the requirements listed in Section 4.3. Having proven then, in Chapter 6, we will demonstrated our approach on human data.

5

QUANTIFICATION OF DIFFERENCES BETWEEN WALKING GAITS

Good walking leaves no track behind it.
—Lao Tsu, Tao Te Ching

One of our objectives in this thesis is to measure the difference between gaits. A potential application in gait rehabilitation, for instance, is to determine how much the device should correct the subject. The principle is to compare the subject's gait with an appropriate reference gait which is expected to be normal and use their difference in a feedback controller (Fig.2.7). For this, we need a way to compare the difference between two gaits.

We hypothesise that the observed walking gait is a combination of *characteristics* of walking and *variations* from embodiments (e.g., body types), environment (e.g., where the subject is walking), and behaviours (e.g., speed and cadence). This assumption is consistent with the observed motion we collected using motion capture facilities.

In Fig. 5.1, we show an example of how variations in embodiments affect the walking gait. The trajectories are hip (Fig. 5.1a) and knee (Fig. 5.1b) angles from five healthy subjects walking on a treadmill with a fixed speed of one meter-per-second. Although the subjects were asked to walk normally, the presence of inter-personal variations yield no clear definition for a *normal walking* gait.

For this reason, it is sub-optimal to compare the distance between two sets of observations directly, since the observations contain natural variations. A more appropriate approach is to quantify the distance between the *characteristics* of two walking gaits and ignore the differences coming from these variations.

In Chapter 3, we discussed our walking phase model such that the *characteristics* and *variations* can be separated. In Chapter 4, we proposed a method to learn the

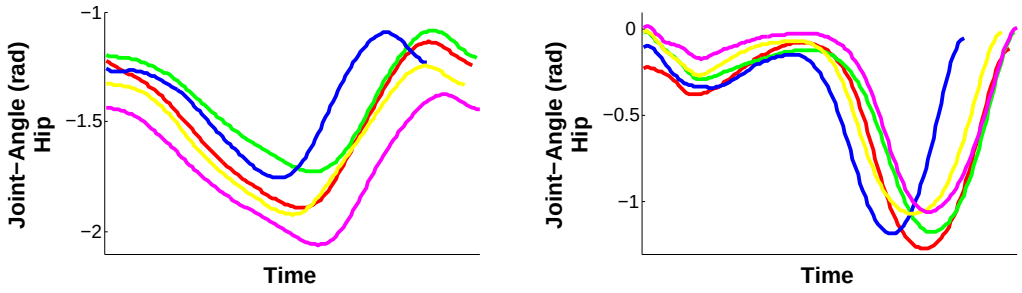

(a) (b)

Figure 5.1: Examples of how variations in embodiments affect the walking gait. The trajectories are taken from five subjects (denoted by colours) walking on a treadmill with a fixed speed of one meter-per-second. The figures are (a) hip angles and (b) knee angles. Although the subjects were asked to walk normally, the presence of interpersonal variations yield no clear definition for a *normal walking* gait.

characteristics of walking *without the knowledge of the variations*. In this chapter, we focus on the method to quantify the distance between a walking gait and a reference gait, which is learnt using the method in Chapter 4.

However, in order to quantify the difference between the two walking gaits, we need to reproduce the motion *under the variation from embodiment*, which is unknown by assumption. Note that, the *variation from embodiment* corresponds to the *constraint* and the *null-space projection matrix* in our walking phase model model (detail see Section 3.2). In this chapter, we also proposed a method to estimate these two quantities.

In Section 5.1, we formalise the problem and the solution of quantifying the difference between gaits. To aid the exposition, we begin by explaining our approach on a simpler problem domain in Section 5.2, and the method for measuring gait abnormality is described in Section 5.3. Finally, in Section 5.4, we summarised some numerical results that validate our method.

5.1 PROBLEM FORMULATION: COMPARING WALKING GAITS

In Section 3.2, we introduced a walking phase model, such that the observed actions can be described as

$$\mathbf{u} = \mathbf{A}^{\dagger}\,\mathbf{b} + \mathbf{N}\,\boldsymbol{\pi}$$

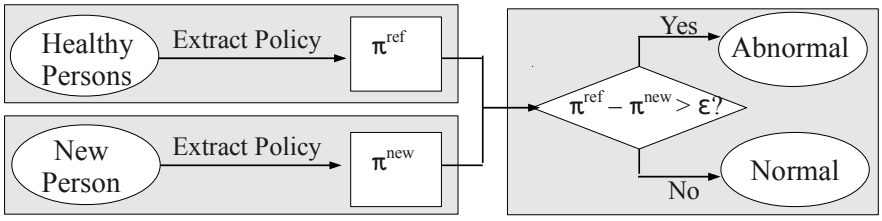

Figure 5.2: Gait abnormality detection by measuring the difference in null-space policies. Ideally, we want to extract the null-space policy from the gait of healthy subjects (as the reference policy $\boldsymbol{\pi}^{ref}$) and compare to the null-space policy from a new subject ($\boldsymbol{\pi}^{new}$). The distance between these two policies ($\boldsymbol{\pi}^{ref}$- $\boldsymbol{\pi}^{new}$) is a quantification of their difference.

where the consistent characteristics of walking is corresponding to $\boldsymbol{\pi}$, the variations in embodiments and environment are captured by $\mathbf{A}$, and the variations in behaviours (such as speed) are represented by $\mathbf{b}$. Note that we don't know the true values of these variables, but we assume that such decomposition exists.

As mentioned in the prior section, to compare the difference between two walking gaits, we only need to compare the *characteristics* of the gait and ignore the difference arise from the variations. With this walking-phase model, comparing two gaits is equivalent to comparing two null-space policies $\boldsymbol{\pi}$.

A general framework is illustrated in Fig 5.2; ideally, we want to extract the null-space policy from the gait of healthy subjects (as the reference policy $\boldsymbol{\pi}^{ref}$) and compare to the null-space policy from a new subject ($\boldsymbol{\pi}^{new}$). A difference above a certain threshold would signify a pathological gait.

reference policy

In Section 4.2, we proposed a method to generalise the characteristics of walking from demonstrations containing variations from embodiments and behaviours. In terms of our walking phase model, we are able to learn $\boldsymbol{\pi}$ without the knowledge of the variations from $\mathbf{A}$, $\mathbf{b}$, and $\mathbf{N}$. We can use this approach to generalise the reference gait $\boldsymbol{\pi}^{ref}$.

However, in order to estimate $\boldsymbol{\pi}^{new}$ and compute this difference, observations from various embodiments are required (Section 4.3), which is infeasible for the new subject. Nevertheless, it is still possible to detect abnormalities in gait by comparing the policies *under variations of embodiments*. Specifically, instead of measuring the difference between $\boldsymbol{\pi}^{new}$ and $\boldsymbol{\pi}^{ref}$, we can instead evaluate the difference between $\mathbf{N}^{new}\boldsymbol{\pi}^{new}$ and $\mathbf{N}^{new}\boldsymbol{\pi}^{ref}$, where $\mathbf{N}^{new}$ is a projection matrix derived from the constraints of the new subject.

On the other hand, the *variations from embodiment* is also unknown by assumption, i.e., we have no knowledge about the constraint and the projection matrix of the new person $\mathbf{A}^{new}, \mathbf{N}^{new}$. In order to measure the constrained difference, we also proposed a method to estimate the $\mathbf{N}^{new}$.

To aid the explanation, in the next section, we will describe our approach on learning null-space projection in a simpler and more generic setting. In Section 5.3, we will outline the method for learning $\mathbf{N}^{new}$ and compare $\mathbf{N}^{new}\boldsymbol{\pi}^{new}$ and $\mathbf{N}^{ref}\boldsymbol{\pi}^{ref}$ in details.

5.2 PROJECTION MATRIX ESTIMATION FOR STATIONARY CONSTRAINT SYSTEMS

We first consider the simpler problem in which the underlying policy is subjected to a set of $\mathcal{S}$-dimensional *stationary constraints*. These constraints act as hard restrictions to the movement of the policy but do not enforce the system to accomplish any tasks.

stationary constraints

Note that this is a simpler problem than the one described in Section 2.4.2 and Section 3.2. The stationary constraint system is a special case where $\mathbf{b} = 0$ and satisfy the relation

$$\mathbf{A}(\mathbf{x}, t)\,\mathbf{u}(\mathbf{x}, t) = O \tag{5.1}$$

This means that the observed action $\mathbf{u}$ is the projection of the policy $\boldsymbol{\pi}$ into the null-space of constraint $\mathbf{A}$; namely, inverting (5.1) results in the relation

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{N}(\mathbf{x}, t)\,\boldsymbol{\pi}(\mathbf{x}) \tag{5.2}$$

where $\mathbf{N}(\mathbf{x}, t) := (\mathbf{I} - \mathbf{A}(\mathbf{x}, t)^{\dagger}\,\mathbf{A}(\mathbf{x}, t)) \in \mathbb{R}^{\mathcal{Q}\times\mathcal{Q}}$ is the projection matrix which projects the null-space policy $\boldsymbol{\pi}$ onto the null-space of $\mathbf{A}$, which in general, has non-linear dependence on both time and state. (Since we are interested in the special case where $\mathbf{b} = 0$, the first term $\mathbf{A}^{\dagger}\,\mathbf{b}$ in Equ. 2.1 disappears in the above equation.)

It would be useful to know the decomposition of $\mathbf{A}$, $\mathbf{N}$, and $\boldsymbol{\pi}$; however, the true quantities of those variables are unavailable by assumption. Several studies mentioned in Section 2.4 have been devoted to learning the null-space policy $\boldsymbol{\pi}$, but, to our knowledge, none have been able to explicitly estimate $\mathbf{A}$ or $\mathbf{N}$.

5.2.1 *Searching the Optimal Constraints*

The proposed method works on data given as $\mathcal{N}$ pairs of observed states $\mathbf{x}_n$ and observed actions $\mathbf{u}_n$. It is assumed that

1. the observations can be decomposed as $\mathbf{u} = \mathbf{N}\boldsymbol{\pi}$
2. $\mathbf{u}$ are generated using the same null-space policy $\boldsymbol{\pi}$,
3. each observation might have been constrained for some $\mathbf{A} \neq 0$, and
4. $\mathbf{A}$ (and $\mathbf{N}$) are not explicitly known for any given observation.

The key to the proposed approach is to use properties of the projection matrix $\mathbf{N}$ in order to find $\mathbf{A}$. By definition $\mathbf{u} = \mathbf{N}\boldsymbol{\pi}$, so $\mathbf{u}$ is the vector $\boldsymbol{\pi}$ projected onto the image space of $\mathbf{N}$. However, it is also the case that *the projection of* $\mathbf{u}$ *also lies in this image space*, i.e.,

$$\mathbf{N}\mathbf{u} = \mathbf{u}. \tag{5.3}$$

Based on this insight, $\mathbf{N}$ can be approximated by seeking an estimate *such that the condition in (5.3) holds*. Specifically, given samples $\{\mathbf{x}_n, \mathbf{u}_n\}_{n=1}^{\mathcal{N}}$ it is proposed to form an estimate $\tilde{\mathbf{N}}$ that minimises the difference between $\tilde{\mathbf{N}}\mathbf{u}$ and the raw observations $\mathbf{u}$, i.e.,

$$E[\tilde{\mathbf{N}}] = \sum_{n=1}^{\mathcal{N}} \|\mathbf{u}_n - \tilde{\mathbf{N}}\mathbf{u}_n\|^2. \tag{5.4}$$

Using (3.3), the n^{th} term of (5.4) can be written

$$\begin{aligned} \|\mathbf{u}_n - \tilde{\mathbf{N}}\mathbf{u}_n\|^2 &= \|\mathbf{u}_n - (\mathbf{I} - \tilde{\mathbf{A}}^\dagger\tilde{\mathbf{A}})\mathbf{u}_n\|^2 \\ &= \|\mathbf{u}_n - \mathbf{u}_n + \tilde{\mathbf{A}}^\dagger\tilde{\mathbf{A}}\mathbf{u}_n\|^2 \\ &= \|\tilde{\mathbf{A}}^\dagger\tilde{\mathbf{A}}\mathbf{u}_n\|^2 \end{aligned} \tag{5.5}$$

where $\tilde{\mathbf{A}} \in \mathbb{R}^{\mathcal{S}\times\mathcal{Q}}$ is an estimate of the constraint matrix $\mathbf{A}$.

Expanding the norm $\|\tilde{\mathbf{A}}^\dagger\tilde{\mathbf{A}}\mathbf{u}_n\|^2 = \mathbf{u}_n^\top(\tilde{\mathbf{A}}^\dagger\tilde{\mathbf{A}})^\top\tilde{\mathbf{A}}^\dagger\tilde{\mathbf{A}}\mathbf{u}_n$, and using the identities of pseudo-inverse $(\mathbf{A}^\dagger\mathbf{A})^\top = \mathbf{A}^\dagger\mathbf{A}$ and $\mathbf{A}^\dagger\mathbf{A}\mathbf{A}^\dagger = \mathbf{A}^\dagger$, (5.4) can be expressed in simplified form

$$E[\tilde{\mathbf{N}}] = \sum_{n=1}^{\mathcal{N}} \mathbf{u}_n^\top\tilde{\mathbf{A}}^\dagger\tilde{\mathbf{A}}\mathbf{u}_n. \tag{5.6}$$

The constraint matrix can therefore be estimated by seeking an estimate $\tilde{\mathbf{A}}$ that minimises (5.6). Note that, through use of the latter, *no prior knowledge of the underlying* $\boldsymbol{\pi}$*, nor of the true projection matrix* $\mathbf{N}$ *is required*.

5.2.2 *Representation of* $\tilde{\mathbf{A}}$

Until this point, no specific assumptions have been made on the form of the estimated matrix $\tilde{\mathbf{A}}$. Here, an appropriate structure of this matrix is outlined, that ensures that the estimate is interpretable in terms of the constraints.

Specifically, $\tilde{\mathbf{A}}$ is assumed to be formed from a set of S linearly independent vectors

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \vdots \\ \mathbf{a}_S \end{bmatrix} \tag{5.7}$$

where $\mathbf{a}_s = (\alpha_{s,1}, \alpha_{s,2}, ..., \alpha_{s,Q})$ corresponds to the s^{th} constraint in the observations, and $\mathbf{a}_i \perp \mathbf{a}_j$ for all $i \neq j$.

Since $\tilde{\mathbf{A}}$ is full row rank, $(\tilde{\mathbf{A}}\tilde{\mathbf{A}}^\top)$ is invertible, and the pseudo-inverse has the explicit formula $\tilde{\mathbf{A}}^\dagger = \tilde{\mathbf{A}}^\top(\tilde{\mathbf{A}}\tilde{\mathbf{A}}^\top)^{-1}$, i.e.,

$$\tilde{\mathbf{A}}^\dagger = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \vdots \\ \mathbf{a}_S \end{bmatrix}^\top \begin{bmatrix} \mathbf{a}_1\mathbf{a}_1^\top & \mathbf{a}_1\mathbf{a}_2^\top & \dots & \mathbf{a}_1\mathbf{a}_S^\top \\ \mathbf{a}_2\mathbf{a}_1^\top & \mathbf{a}_2\mathbf{a}_2^\top & \dots & \mathbf{a}_2\mathbf{a}_S^\top \\ \vdots & & \ddots & \vdots \\ \mathbf{a}_S\mathbf{a}_1^\top & \mathbf{a}_S\mathbf{a}_2^\top & \dots & \mathbf{a}_S\mathbf{a}_S^\top \end{bmatrix}^{-1} \tag{5.8}$$

Noting that $\mathbf{a}_i\mathbf{a}_i^\top = \|\mathbf{a}_i\|^2$ and $\mathbf{a}_i\mathbf{a}_j^\top = 0$ for $i \neq j$, it can be shown that

$$\tilde{\mathbf{A}}^\dagger = \begin{bmatrix} \mathbf{a}_1^\top & \mathbf{a}_2^\top & \cdots & \mathbf{a}_S^\top \end{bmatrix} \begin{bmatrix} \frac{1}{\|\mathbf{a}_1\|^2} & 0 & \cdots & 0 \\ 0 & \frac{1}{\|\mathbf{a}_2\|^2} & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \cdots & 0 & \frac{1}{\|\mathbf{a}_S\|^2} \end{bmatrix} \tag{5.9}$$

$$= \begin{bmatrix} \frac{\mathbf{a}_1^\top}{\|\mathbf{a}_1\|^2} & \frac{\mathbf{a}_2^\top}{\|\mathbf{a}_2\|^2} & \cdots & \frac{\mathbf{a}_S^\top}{\|\mathbf{a}_S\|^2} \end{bmatrix}. \tag{5.10}$$

From (3.3), the projection matrix is

$$\tilde{\mathbf{N}} = \mathbf{I} - \begin{bmatrix} \frac{\mathbf{a}_1^\top}{\|\mathbf{a}_1\|^2} & \frac{\mathbf{a}_2^\top}{\|\mathbf{a}_2\|^2} & \cdots & \frac{\mathbf{a}_S^\top}{\|\mathbf{a}_S\|^2} \end{bmatrix} \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \vdots \\ \mathbf{a}_S \end{bmatrix} = \mathbf{I} - \sum_{s=1}^{S} \hat{\mathbf{a}}_s^\top \hat{\mathbf{a}}_s \tag{5.11}$$

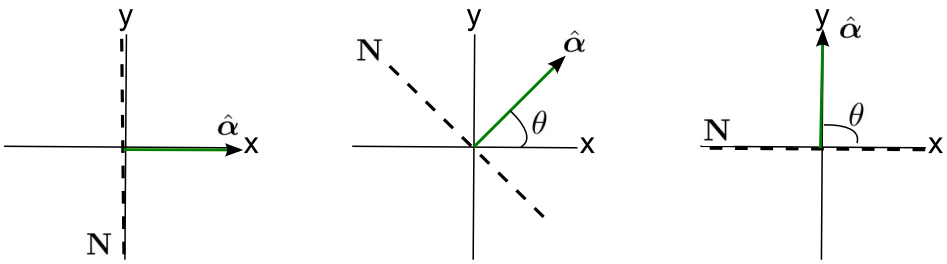

Figure 5.3: Examples of unit vectors $\hat{\mathbf{a}} \in \mathbb{R}^2$ in polar representation. Figures are (a) $\theta = 0^\circ$, (b) $\theta = 45^\circ$, and (c) $\theta = 90^\circ$ and the corresponding null-space (dashed line).

where $\mathbf{a}_i = ||\mathbf{a}_i||\hat{\mathbf{a}}_i$ is used. The last equality of (5.11), makes it clear that the projection matrix is independent of the magnitude of each of the $\mathbf{a}_s$ since only unit vectors appear in the expression. Hence, only the *direction* of these vectors need be approximated, without considering the *magnitude*.

5.2.3 *Forming the Estimate* $\tilde{\mathbf{N}}$

In the following, methods for forming the estimate $\tilde{\mathbf{N}}$ are outlined, according to the problem setting.

5.2.3.1 *Estimation for* $\mathbf{A} \in \mathbb{R}^{1\times 2}$

To begin, consider the simplest case where $\mathbf{A}$ represents a one-dimensional constraint in a two degree of freedom system, i.e., $\mathbf{A} \in \mathbb{R}^{1\times 2}$. In this case $\hat{\mathbf{a}}$ is simply the unit vector

$$\hat{\mathbf{a}} = (\cos\theta, \sin\theta) \tag{5.12}$$

where θ is the angle encoding the direction (see Fig. 5.3).

A simple method for forming the estimate in this case, is to perform a line search for the choice of θ that minimises the error criterion (5.6). Note that, the search need only be restricted to the half-space $0 \leqslant \theta < \pi$ since this covers all possible cases (the projections corresponding to the half-space $\pi \leqslant \theta \leqslant 2\pi$ are identical to those in the former).

For example, let N_θ be the number of θ's uniformly distributed between $[0, \pi)$. If $N_\theta = 180$, we calculate (5.6) for $\theta = 0^\circ, 1^\circ, ..., 179^\circ$, and the optimal θ^* minimises (5.6).

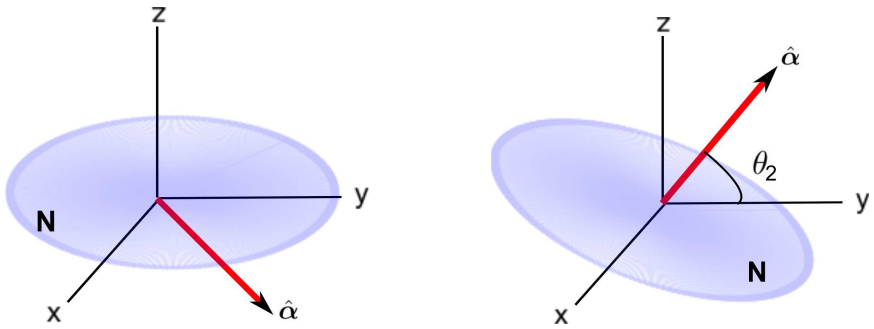

(a) (b)

Figure 5.4: Examples of unit vectors $\hat{\mathbf{a}} \in \mathbb{R}^3$ in spherical representation (red) and the corresponding null-space (the blue surface), where $\theta_1 = 45^\circ$, (a) $\theta_2 = 0^\circ$ and (b) $\theta_2 = 45^\circ$.

	$\mathbb{R}^2$	$\mathbb{R}^3$	$\mathbb{R}^4$
$\hat{a}_1$	$\cos\theta_1$	$\cos\theta_1$	$\cos\theta_1$
$\hat{a}_2$	$\sin\theta_1$	$\sin\theta_1\cos\theta_2$	$\sin\theta_1\cos\theta_2$
$\hat{a}_3$	-	$\sin\theta_1\sin\theta_2$	$\sin\theta_1\sin\theta_2\cos\theta_3$
$\hat{a}_4$	-	-	$\sin\theta_1\sin\theta_2\sin\theta_3$

Table 5.1: Representation of unit vector in $\mathbb{R}^2$, $\mathbb{R}^3$, and $\mathbb{R}^4$

5.2.3.2 *Estimation for* $\mathbf{A} \in \mathbb{R}^{1\times\Omega}$ *with* $\Omega > 2$

For problems with $\Omega > 2$ but with $\mathcal{S} = 1$ the same approach can be applied by extending the $\hat{\mathbf{a}}$ in (5.12) to handle higher degrees of freedom.

For example, to extend the formulation to $\tilde{\mathbf{A}} \in \mathbb{R}^{1\times 3}$, the search need be conducted over possible unit vectors $\hat{\mathbf{a}} \in \mathbb{R}^3$. Using spherical coordinates, the latter can be represented as

$$\hat{\mathbf{a}} = (\cos\theta_1,\ \sin\theta_1\cos\theta_2,\ \sin\theta_1\sin\theta_2) \tag{5.13}$$

where now $\theta_1 \in [0,\pi)$ and $\theta_2 \in [0,\pi)$ encode the direction of the vector in the three dimensional space. For instance, in Fig. 5.4, the additional variable θ_2 specifies the angle between $\hat{\mathbf{a}}$ and xz-plane.

Similarly, for $\mathcal{S} = 4$ the vector $\hat{\mathbf{a}}$ can be represented by three angles θ_1, θ_2 and θ_3 with elements as indicated in the third column of Table 5.1.

More generally, for $\mathcal{Q} > 2$, the unit vector $\hat{\mathbf{a}} = (\hat{a}_1, \hat{a}_2, \cdots, \hat{a}_{\mathcal{Q}})$ can be represented by $\mathcal{Q} - 1$ parameters $\boldsymbol{\theta} = (\theta_1, \theta_2, \cdots, \theta_{\mathcal{Q}-1})^\top$ where

$$\begin{aligned}
\hat{a}_1 &= \cos\theta_1 \\
\hat{a}_2 &= \sin\theta_1 \cos\theta_2 \\
\hat{a}_3 &= \sin\theta_1 \sin\theta_2 \cos\theta_3 \\
&\vdots \\
\hat{a}_{\mathcal{Q}-1} &= \prod_{\nu=1}^{\mathcal{Q}-2} \sin\theta_\nu \cos\theta_{\mathcal{Q}-1} \\
\hat{a}_{\mathcal{Q}} &= \prod_{\nu=1}^{\mathcal{Q}-1} \sin\theta_\nu
\end{aligned} \tag{5.14}$$

5.2.4 *Estimation for Multidimensional Constraints*

For systems subject to multidimensional constraints (i.e., $\mathbf{A} \in \mathbb{R}^{\mathcal{S}\times\mathcal{Q}}$ with $\mathcal{Q} > 2$ and $\mathcal{S} > 1$), $\mathbf{A}$ consisted of a set $\mathcal{S}$ orthogonal vectors as in (5.7) (see Section 5.2.2).

The same optimisation approach can be used to form the estimate $\tilde{\mathbf{N}}$ through a similar approach as in the preceding sections. Namely, by defining an orthogonal unit vector $\hat{\mathbf{a}}_s$ for each of the $\mathcal{S}$ constraints, and selecting the optimal parameter vectors $\boldsymbol{\theta}_s = (\theta_{s,1}, \theta_{s,2}, \cdots, \theta_{s,\mathcal{Q}-1})$ that minimise (5.6). However, in such a scheme the rapid increase in the size of the parameter space means this will not be feasible for very high $\mathcal{S}$ and $\mathcal{Q}$.

Instead, an iterative approach to learning may be employed, whereby a series of of constraint vectors $\hat{\mathbf{a}}_s$ are fitted to the data to form an estimate $\tilde{\mathbf{A}}$, where an $(s+1)^{\text{th}}$ vector is *only added if it does not reduce the fit under* (5.6). This exploits a second property of the projection matrix $\mathbf{N}$, namely that for multidimensional constraints the total projection can be decomposed into a set of uni-dimensional projections

$$\mathbf{N} = \mathbf{N}_1\mathbf{N}_2\cdots\mathbf{N}_{\mathcal{S}}. \tag{5.15}$$

or equivalently

$$\mathbf{u} = \mathbf{N}\boldsymbol{\pi} = \mathbf{N}_1(\mathbf{N}_2\cdots(\cdots\mathbf{N}_{\mathcal{S}}\boldsymbol{\pi})\cdots) \tag{5.16}$$

through (5.2). This suggests that a reasonable approximation $\tilde{\mathbf{N}}$ may be formed by

1. first finding the optimal $\hat{\mathbf{a}}_1^*$ (i.e., $\boldsymbol{\theta}_1^*$) under (5.6) through the fitting procedure described in Section 5.2.3.2, then
2. finding the optimal $\hat{\mathbf{a}}_2^*$ (i.e., $\boldsymbol{\theta}_2^*$), subject to $\hat{\mathbf{a}}_1^* \perp \hat{\mathbf{a}}_2^*$, and
3. repeating until the addition of a new constraint $\hat{\mathbf{a}}_{s+1}^*$ fails to reduce the error under (5.6) any further.

Note that, to find the optimal $\hat{\mathbf{a}}_s^*$ for $s > 1$, the search only need to be performed for $\hat{\mathbf{a}}_s \perp \hat{\mathbf{a}}_j^*$ for all $s > j$. To reduce the search process, we also proposed a method to define the set of $\hat{\mathbf{a}}_s$ that is orthogonal to all $\hat{\mathbf{a}}_j^*$ in Appendix B.

5.3 QUANTIFYING THE DIFFERENCE BETWEEN TWO WALKING GAITS

After exploring the solution for stationary constraints systems in the preceding section, we return to our problem of quantifying the difference between walking gaits. Namely, we consider the null-space policy is subjected to a set of constraints

$$\mathbf{A}(\mathbf{x}, t)\,\mathbf{u}(\mathbf{x}, t) = \mathbf{b}(\mathbf{x}, t)$$

where $\mathbf{b} \neq 0$ is the *task-space policy* describing the underlying task to be accomplished, and the observations can be decomposed into

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{A}^{\dagger}(\mathbf{x}, t)\,\mathbf{b}(\mathbf{x}, t) + \mathbf{N}(\mathbf{x}, t)\,\boldsymbol{\pi}(\mathbf{x})$$

As mentioned in the beginning of this chapter, in order to quantify the difference between two walking gaits, we need to know the *variation in embodiment* (i.e., this variation corresponds to the projection matrix $\mathbf{N}$ in the above functional). In the section, we discuss our method on estimating the projection matrix $\mathbf{N}$ for this system and using the learnt projection matrix for comparing the difference between walking gaits.

5.3.1 *Learning Null-space Projection for Walking Gaits*

The proposed method works on data given as $\mathcal{N}$ pairs of observed states $\mathbf{x}_n$ and observed actions $\mathbf{u}_n$. It is assumed that

1. the observations can be decomposed as $\mathbf{u} = \mathbf{A}^{\dagger}\,\mathbf{b} + \mathbf{N}\,\boldsymbol{\pi}$
2. $\mathbf{u}$ are generated using the same null-space policy $\boldsymbol{\pi}$,

3. each observation might have been constrained for some $\mathbf{A} \neq 0$ and $\mathbf{b} \neq 0$, and

4. $\mathbf{A}$ (and $\mathbf{N}$) are not explicitly known for any given observation.

In Section 4.2, we described a method such that the observations $\mathbf{u}$ can be decomposed into two orthogonal components $\mathbf{u} \equiv \mathbf{u}^{ts} + \mathbf{u}^{ns}$ such that $\mathbf{u}^{ts} \equiv \mathbf{A}^{\dagger}\,\mathbf{b}$, $\mathbf{u}^{ns} \equiv \mathbf{N}\,\boldsymbol{\pi}$, and $\mathbf{u}^{ts} \perp \mathbf{u}^{ns}$ (i.e., this is the *Step-1* of our null-space policy learning).

Similar to the same approach described in Section 5.2, $\mathbf{N}$ can be approximated by seeking an estimate such that the condition $\mathbf{N}\,\mathbf{u}^{ns} = \mathbf{u}^{ns}$ holds. Note that, by definition, $\mathbf{u}^{ts} \perp \mathbf{u}^{ns}$, so the estimate must also satisfy $\mathbf{N}\,\mathbf{u}^{ts} = 0$. Based on this insight, (5.4) can be modified by adding an additional term

$$\mathrm{E}[\tilde{\mathbf{N}}] = \sum_{n=1}^{\mathcal{N}} \|\mathbf{u}_n^{ns} - \tilde{\mathbf{N}}\mathbf{u}_n^{ns}\|^2 + \|\tilde{\mathbf{N}}\mathbf{u}_n^{ts}\|^2 \tag{5.17}$$

From (5.6), we know that the first term in (5.17) can be written as $\|\mathbf{u}_n^{ns} - \mathbf{N}\,\mathbf{u}_n^{ns}\|^2 = \mathbf{u}_n^{ns\,\top}\,\tilde{\mathbf{A}}^{\dagger}\tilde{\mathbf{A}}\mathbf{u}_n^{ns}$. Expanding the norm of the second term $\|\,\tilde{\mathbf{N}}\,\mathbf{u}_n^{ts}\|^2 = \mathbf{u}_n^{ts\,\top}\,\tilde{\mathbf{N}}^{\top}\,\tilde{\mathbf{N}}\,\mathbf{u}_n^{ts}$ and using the identities $\tilde{\mathbf{N}}^{\top} = \tilde{\mathbf{N}}$ and $\mathbf{N}^2 = \mathbf{N}$, the second term can be simplified as $\mathbf{u}_n^{ts}\,\tilde{\mathbf{N}}\,\mathbf{u}_n^{ts}$. Using the definition in (3.3),

$$\begin{aligned}
\mathbf{u}_n^{ts\,\top}\,\tilde{\mathbf{N}}\,\mathbf{u}_n^{ts} &= \mathbf{u}_n^{ts}(\mathbf{I} - \tilde{\mathbf{A}}^{\dagger}\tilde{\mathbf{A}})\mathbf{u}_n^{ts} \\
&= \mathbf{u}_n^{ts\,\top}\,\mathbf{u}_n^{ts} - \mathbf{u}_n^{ts\,\top}\,\tilde{\mathbf{A}}^{\dagger}\tilde{\mathbf{A}}\mathbf{u}_n^{ts} \\
&= \|\mathbf{u}_n^{ts}\| - \mathbf{u}_n^{ts\,\top}\,\tilde{\mathbf{A}}^{\dagger}\tilde{\mathbf{A}}\mathbf{u}_n^{ts}
\end{aligned}$$

The objective in (5.17) can be expressed in simplified form

$$\mathrm{E}[\tilde{\mathbf{N}}] = \sum_{n=1}^{\mathcal{N}} \mathbf{u}_n^{ns\,\top}\,\tilde{\mathbf{A}}^{\dagger}\tilde{\mathbf{A}}\mathbf{u}_n^{ns} + \|\mathbf{u}_n^{ts}\| - \mathbf{u}_n^{ts\,\top}\,\tilde{\mathbf{A}}^{\dagger}\tilde{\mathbf{A}}\mathbf{u}_n^{ts} \tag{5.18}$$

We define $\mathbf{v}_n^{ns} = vec(\mathbf{u}_n^{ns} \otimes \mathbf{u}_n^{ns}) \in \mathbb{R}^{1\times\mathcal{Q}^2}$, $\mathbf{v}_n^{ts} = vec(\mathbf{u}_n^{ts} \otimes \mathbf{u}_n^{ts}) \in \mathbb{R}^{1\times\mathcal{Q}^2}$, and $\mathbf{v}^{A} = vec(\mathbf{A}^{\dagger}\mathbf{A})^{\top} \in \mathbb{R}^{\mathcal{Q}^2\times 1}$ to obtain a simpler functional form

$$\begin{aligned}
\mathrm{E}[\tilde{\mathbf{N}}] &= \sum_{n=1}^{\mathcal{N}} \mathbf{v}_n^{ns}\mathbf{v}^{A} - \mathbf{v}_n^{ts}\mathbf{v}^{A} + \mathbf{u}_n^{ts\,\top}\,\mathbf{u}_n^{ts} \\
&= \sum_{n=1}^{\mathcal{N}} (\mathbf{v}_n^{ns} - \mathbf{v}_n^{ts})\mathbf{v}^{A} + \sum_{n=1}^{\mathcal{N}} \|\mathbf{u}_n^{ts}\|
\end{aligned} \tag{5.19}$$

The projection matrix can therefore be estimated by iterating $\boldsymbol{\theta}$ from 0 to π and seeking the optimal $\boldsymbol{\theta}^{*}$ that minimises (5.19).

The process is summarised in Algorithm 2.

Algorithm 2 Learning Null-space Projection

Input:

$\{\mathbf{u}^{ns}\}_{n=1}^{\mathcal{N}}$: null-space components

$\{\mathbf{u}^{ts}\}_{n=1}^{\mathcal{N}}$: task-space components

Output:

$\tilde{\mathbf{A}}$: the estimated constraint matrix

1: Estimate $\hat{\mathbf{a}}_1^*$ by minimising (5.19). Set $s \leftarrow 1$.
2: **while** $\mathbb{E}[\tilde{\mathbf{N}}]$ in (5.19) is not increasing **do**
3: $s \leftarrow s + 1$.
4: Find $\hat{\mathbf{a}}_s$ using (B.3) such that $\hat{\mathbf{a}}_s \perp \hat{\mathbf{a}}_j^* \ \forall j < s$
5: Learn $\hat{\mathbf{a}}_s^*$ by minimising (5.19)
6: Set $\tilde{\mathbf{A}} \leftarrow [\hat{\mathbf{a}}_1^*, \cdots, \hat{\mathbf{a}}_s^*]^\top$.
7: **end while**
8: Return $\tilde{\mathbf{A}}$.

5.3.2 *Approximate Policy Difference*

To estimate a projection matrix for the k^{th} phase of a new person, we seek a $\boldsymbol{\theta}_k$ such that the difference between $\tilde{\mathbf{N}}_k^{new}\mathbf{u}_k^{ns,new}$ and $\mathbf{u}_k^{ns,new}$ and the norm of $\tilde{\mathbf{N}}_k^{new}\mathbf{u}_k^{ts,new}$ is minimised. Note that the true quantities of these two components $\mathbf{u}_k^{ns,new}$ $\mathbf{u}_k^{ts,new}$ are also unknown, so the *estimated components* are used, instead.

Namely, after learning the null-space component $\tilde{\mathbf{u}}_k^{ns,new}$ for the k^{th} phase by minimising the error criterion (4.3), we obtained N_k pairs of $(\mathbf{x}_k, \tilde{\mathbf{u}}_k^{ns,new})$ and $(\mathbf{x}_k, \tilde{\mathbf{u}}_k^{ts,new})$. For each possible set of $\boldsymbol{\theta}_k \in [0, \pi)$, we can calculate

$$\mathbb{E}_{\boldsymbol{\theta}}[\tilde{\mathbf{N}}_k^{new}] = \sum_{n}^{N_k} ||\tilde{\mathbf{N}}_k^{new}\tilde{\mathbf{u}}_{k,n}^{ns,new} - \tilde{\mathbf{u}}_{k,n}^{ns,new}||^2 + ||\tilde{\mathbf{N}}_k^{new}\tilde{\mathbf{u}}_{k,n}^{ts,new}||^2. \tag{5.20}$$

The optimal $\boldsymbol{\theta}_k \in [0, \pi)$ minimises (5.20).

Given the estimated reference policy $\boldsymbol{\pi}^{ref}$ from the healthy subjects and measurements of the behaviour of the new person, we can now quantify the difference of the new person's policy through a measure that we call the Approximated Policy Difference (APD).

approximate policy difference

Specifically, given pairs of state/action observations of the new person $\mathbf{x}^{new}, \mathbf{u}^{new}$, we first divide their data into K walking phases and learn a model of the null-space

Algorithm 3 Approximate difference between two gaits

Input: $\mathbf{D} = \{\mathbf{x}^{new}, \mathbf{u}^{new}\}$: data-set of a new person

$\boldsymbol{\pi}^{ref}$: reference policy

Output: APD: approximated policy difference

1: Split $\mathbf{D}$ into $\mathbf{D}_k$ where k denotes the phase number
2: **for all** $\mathbf{D}_k$ **do**
3: Learn $\tilde{\mathbf{u}}_k^{ns,new}$ by minimising (4.3)
4: Learn $\tilde{\mathbf{N}}_k^{new}$ by minimising (5.20)
5: **end for**
6: Approximate the difference using (5.21)

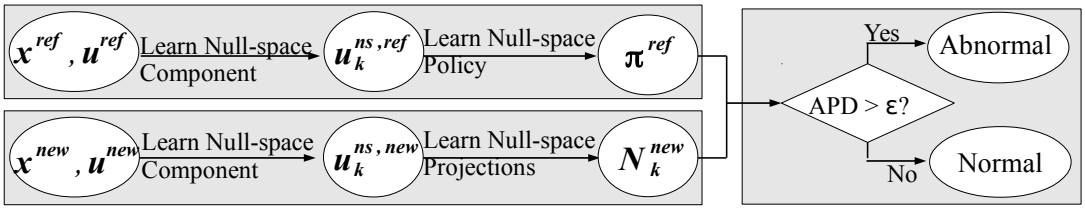

Figure 5.5: Gait abnormality detection by measure the approximated policy difference

components $\tilde{\mathbf{u}}_k^{ns,new}$ for each phase. Then, we take $\tilde{\mathbf{u}}_k^{ns,new}$ and estimate the projection matrix $\tilde{\mathbf{N}}_k^{new}$ using (5.20). The APD is computed as

$$\text{APD} = \frac{1}{N\sigma^2_{\pi^{ref}}} \sum_{k=1}^{K} \sum_{n}^{N_k} \| \tilde{\mathbf{u}}_{k,n}^{ns,new} - \tilde{\mathbf{N}}_k^{new} \boldsymbol{\pi}_{k,n}^{ref} \|^2 \tag{5.21}$$

where $\sigma^2_{\pi^{ref}}$ is the variance of the reference policy. The APD measures the difference between the reference policy and the new person in the constrained space, normalised by the variance of the reference policy. Algorithm 3 and Fig.5.5 summarises the process.

One way to interpret the quantity in (5.21) is, this is the difference between a person and the reference in an abstract policy space, where the distance resulting from various walking behaviours (speeds, step-sizes) is eliminated. Therefore, this measurement can be thought of as the quantification of how much we should correct a gait without interfering with his/her speed or step-length. Additionally, the APD can also be used for monitoring a program of rehabilitation by looking at the evolution of the APD with training.

5.4 VALIDATION

In this section, some numerical results are presented to validate our approach on learning null-space projection. In particular, we tested our approach with (i) an artificial toy 2D system with stationary constraints, (ii) a system with realistic, higher-dimensional, stationary constraints, and (iii) a simulated walking gait.

5.4.1 *Evaluation Criteria*

The goal of this work is to predict the projection matrix $\mathbf{N}$ underlying the constrained observations in order that these may be reproduced through a suitable learning scheme (e.g., Howard et al. (2009)). For testing the performance of learning, therefore, the following evaluation criteria may be defined.

5.4.1.1 *Normalised Projected Policy Error*

This error measure measures the difference between the policy subject to the true constraints, and that of the policy subject to the estimated constraints. Formally, the Normalised Projected Policy Error (NPPE) can be defined as

normalised projected policy error

$$E_{PPE} = \frac{1}{\mathcal{N}\sigma_{\boldsymbol{\pi}}^2} \sum_{n=1}^{\mathcal{N}} \|\mathbf{N}\boldsymbol{\pi}_n - \tilde{\mathbf{N}}\boldsymbol{\pi}_n\|^2 \tag{5.22}$$

where $\mathcal{N}$ is the number of data points, $\boldsymbol{\pi}_n$ are samples of the policy, and $\mathbf{N}$ and $\tilde{\mathbf{N}}$ are the true and the learnt projection matrices, respectively. The error is normalised by the variance of the observations under the true constraints, $\sigma_{\mathbf{u}}^2$. Note that, since $\mathbf{u} = \mathbf{N}\boldsymbol{\pi}$, (5.22) can also be written

$$E_{PPE} = \frac{1}{\mathcal{N}\sigma_{\boldsymbol{\pi}}^2} \sum_{n=1}^{\mathcal{N}} \|\mathbf{u}_n - \tilde{\mathbf{N}}\boldsymbol{\pi}_n\|^2 \tag{5.23}$$

which corresponds directly to the *normalised constrained policy error* (NCPE) discussed in Howard et al. (2009), between the constrained observations and the policy *subject to the learnt constraints*. Note that, (5.22) can only be computed given the ground truth $\boldsymbol{\pi}_n$ and $\mathbf{N}$, so is used here primarily for validation purposes.

normalised constrained policy error

5.4.1.2 *Normalised Projected Observation Error*

To evaluate the fit in absence of samples of $\boldsymbol{\pi}_n$ and $\mathbf{N}$, an alternative criterion must be used. It is proposed, therefore, to instead use the Normalised Projected Observation Error (NPOE), namely,

normalised projected observation error

$$E_{POE} = \frac{1}{\mathcal{N}\sigma_{\boldsymbol{\pi}}^2} \sum_{n=1}^{\mathcal{N}} \|\mathbf{u}_n - \tilde{\mathbf{N}}\mathbf{u}_n\|^2. \tag{5.24}$$

The NPOE indicates the quality of fit of $\tilde{\mathbf{N}}$, i.e., the extent to which the constraints in the training data are captured by the model, reaching zero only[1] when the model exactly satisfies the condition (5.3).

5.4.2 *Toy Example*

Our first experiment demonstrates our approach for learning null-space projections from a simple stationary constraint system (Section 5.2). For this, we set up a two-dimensional system with a one-dimensional constraint (i.e., $\mathbf{A} \in \mathbb{R}^{1\times 2}$). As ground truth null-space policies $\boldsymbol{\pi}$, we considered:

1. a linear policy: $\boldsymbol{\pi} = -\mathbf{L}(\mathbf{x} - \mathbf{x}^*)$ where $\mathbf{L}$ is a positive definite gain matrix.
2. a limit-cycle policy: $\dot{r} = r(\rho - r^2)$ with radius $\rho = 0.1$ m, angular velocity $\dot{\theta} = -2$ rad/s, where r and θ are the polar representation of the state (i.e., $\mathbf{x} = (r\cos\theta,\ r\sin\theta)^\top$).
3. a sinusoidal policy:
   $\boldsymbol{\pi} = 0.1(\cos x_1 \cos x_2,\ \sin x_1 \sin x_2)^\top$

The training data consists of 50 data points, drawn uniform-randomly across the space $(\mathbf{x})_i \sim U(-2, 2)$, $i \in \{1, 2\}$ and subjected to a 1-D constraint $\mathbf{A} = \hat{\mathbf{a}} \in \mathbb{R}^{1\times 2}$, in the direction of the unit vector $\hat{\mathbf{a}}$. The latter is drawn uniform-randomly, $\theta \sim U(0, \pi]$ rad, at the start of each trial of learning. A sample data set for the limit-cycle policy is presented in Fig. 5.6 (left) where projected sample points $\mathbf{u}_n$ are shown in red, with the corresponding unconstrained null-space policy predictions $\boldsymbol{\pi}$ overlaid in grey.

The projection matrix $\tilde{\mathbf{N}}$ is then learnt using this data through minimisation of the objective function (5.6) according to the scheme outlined in Section 5.2. For this, a grid

1 Assuming $\mathbf{u}_n \neq 0$ for some n.

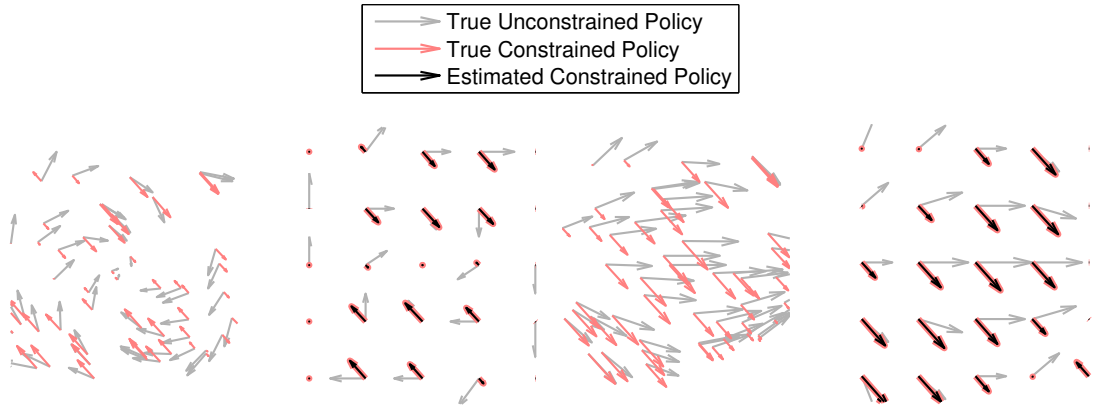

(a) (b)

Figure 5.6: A visualisation of the (a) limit-cycle and (b) sinusoid data. The left figures are the training data and the right figures are the testing data. The colours denote the true policy (grey), the true constrained policy (red), and the predicted constrained policy (black).

Policy	NPPE	NPOE
Linear	4.23 ± 0.61	2.09 ± 0.07
Limit-cycle	4.17 ± 0.38	2.09 ± 0.05
Sinusoid	20.6 ± 6.72	12.2 ± 3.54

Table 5.2: Normalised PPE and POE in predicting the projection matrix. Results are (mean$\pm$s.d.)$\times 10^{-6}$ over 50 trials.

search over 180 different values for θ, uniformly spaced between 0 and π is conducted to recover the optimal θ^*.

The experiment is repeated 50 times and the average NPPE and NPOE are evaluated on a set of 500 test data points, generated through the same procedure as described above.

Table 5.2 summarises the normalised PPE (5.22) and normalised POE (5.24) for each policy. The results are average using the hold-out testing over 50 trails. We can see that, regardless of the underlying null-space policy, our method can learn a good approximation of the projection matrix in terms of both PPE and POE.

In Fig. 5.6, the predictions of the policy under the true constraint $\mathbf{N}\boldsymbol{\pi}$ (black), and the learnt projection matrix $\tilde{\mathbf{N}}\boldsymbol{\pi}$ (red) are plotted for the limit cycle policy. As can be seen, there is good agreement between the two, verifying that in using the learnt constraint, there is little degradation in predicting constrained motion.

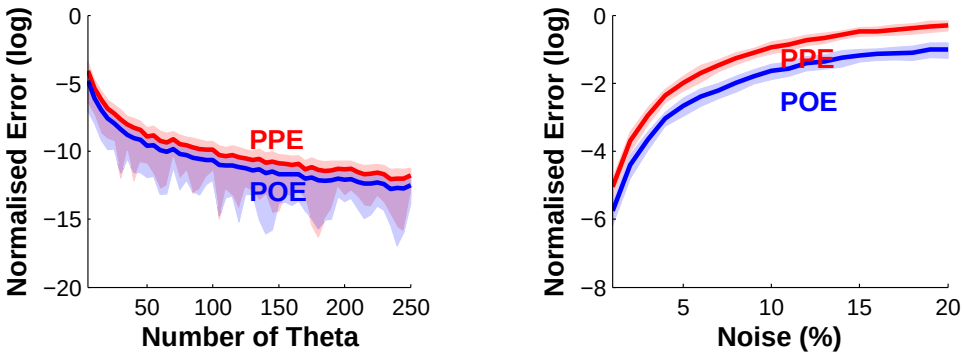

(a) (b)

Figure 5.7: Normalised PPE and POE for (a) increasing number of sampling θ and (b) increasing noise levels in the observed $\mathbf{u}$. Curves are mean±s.d. over 50 trials.

To further characterise the performance of the proposed approach, we also looked at the effect of varying the density of sample points in the grid search for θ for the limit cycle policy. We test our method for 10, 20, 30, up to 250 θ's sampled. The results over 50 trials are plotted in Fig. 5.7a. It can be seen that the NPPE and NPOE rapidly decrease as the number of θ sampled increases (please note the log scale). This is to be expected, since a higher resolution grid allows the learner to form a more accurate estimate of the constraint direction. Note that, even at relatively course sampling (number of $\theta < 50$), the error is still very low.

We also test how the levels of noise present in the training data affect the performance of our method. For this, we contaminated the limit-cycle policy $\boldsymbol{\pi}$ with Gaussian noise, the scale of which we varied to match up to 20% of the data. The resulting NPPE and NPOE follows the noise level, as plotted in Fig. 5.7b. It should be noted, however, that the error is still relatively low (NPPE$< 10^{-2}$), even when the noise is as high as 5% of the variance of the data.

5.4.3 *Kuka Lightweight Robot*

The goal of this experiment is to assess the performance of the proposed approach for learning with higher degrees of freedom, and more realistic constraints with varying dimensionality. Note that, there is no corresponding problem in walking, so a robot-reaching problem is exploited instead. For this, constrained motion data from a kinematic simulation of the 7-DOF Kuka Lightweight Robot (LWR-III) is used (Fig. 5.8).

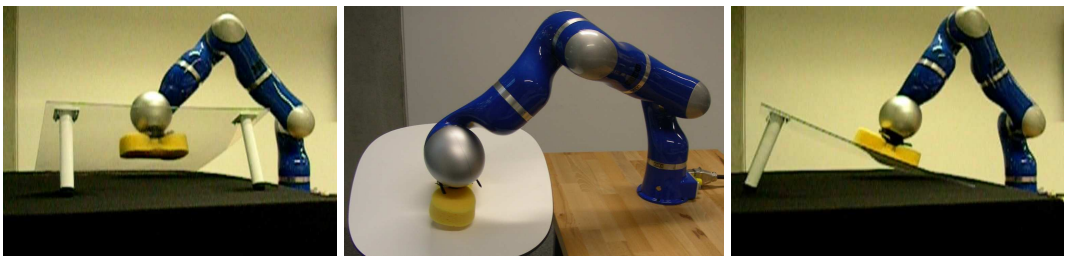

Figure 5.8: Examples of wiping on different tables Howard et al. (2008). The behaviour (wiping) is subject to various constraint imposed by the environment where the behaviour is performed (surfaces).

Here, the state and action space correspond to the joint angles and velocities, respectively, i.e., $\mathbf{x}, \mathbf{u} \in \mathbb{R}^7$, and data is gathered from the arm subject to varying constraints on its end-effector motion.

Specifically, constraints are imposed on motion in the task-space $\mathbf{r} = (x, y, z)^\top$ where x, y and z denote the translational coordinates of the end-effector. Mathematically, they are described through constraint matrices of the form

$$\mathbf{A}(\mathbf{x}) = \mathbf{\Lambda}\mathbf{J}(\mathbf{x}) \tag{5.25}$$

where $\mathbf{J}(\mathbf{x}) \in \mathbb{R}^{3\times7}$ is the manipulator Jacobian, and $\mathbf{\Lambda} \in \mathbb{R}^{8\times3}$ is a matrix specifying the coordinates to be constrained. For example, choosing $\mathbf{\Lambda} = (0, 0, 1)$, ensures that the end-effector does not move in the vertical (z) direction (a 'one-dimensional' constraint). Using $\mathbf{\Lambda} = (\lambda_1^\top, \lambda_2^\top)^\top$ with $\lambda_1 = (1, 0, 0)$ and $\lambda_2 = (0, 1, 0)$ prevents movement of the end-effector in both x and y. Note that, neither the dimensionality of the constraint, nor the subspace in which it acts, is provided a priori to the learner. The experimental procedure is as follows.

Data is gathered by recording motion of the arm, subject to different constraints, from a number of random start states. The latter are drawn randomly from the robot joint space in the half-range of the joint limits, i.e., $(\mathbf{x})_i \sim U(-0.5(\mathbf{x}^{max})_i, 0.5(\mathbf{x}^{max})_i)$ where $\mathbf{x}^{max} = (170°, 120°, 170°, 120°, 170°, 120°, 170°)^\top$. For each start state, a trajectory is recorded from a point attractor policy $\boldsymbol{\pi} = -\mathbf{L}(\mathbf{x} - \mathbf{x}^*)$ where $\mathbf{x}^* = 0$ and $\mathbf{L} = \mathbf{I}$ under the active constraint. In each data set, the latter consisted of one of the following constraints:

1. $\mathbf{\Lambda}_A = (0,\ 0,\ 1)$,
2. $\mathbf{\Lambda}_B = (0,\ \sin\frac{\pi}{3},\ \cos\frac{\pi}{3})$,
3. $\mathbf{\Lambda}_C = ((1,\ 0,\ 0), (0,\ 1,\ 0))^\top$, and

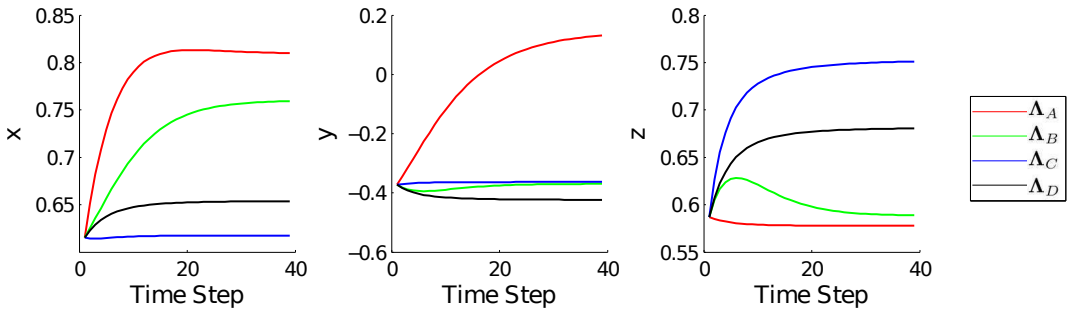

Figure 5.9: Example trajectories generated from 4 different types of constraints

4. $\boldsymbol{\Lambda}_D = (\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2)^\top$ with
   $\boldsymbol{\lambda}_1 = (0,\ \sin\frac{\pi}{3},\ \cos\frac{\pi}{3})$ and
   $\boldsymbol{\lambda}_2 = (\sin\frac{\pi}{2}\sin\frac{\pi}{3},\ \sin\frac{\pi}{2}\cos\frac{\pi}{3},\ \cos\frac{\pi}{2})$.

These correspond to various real world tasks similar to Fig. 5.8, for example, $\boldsymbol{\Lambda}_A$ applies when interacting with a flat (horizontal) surface (e.g., as when wiping or writing on a table) and $\boldsymbol{\Lambda}_B$ corresponds to the case where the surface is inclined by $\frac{\pi}{3}, \mathrm{rad}$.

In this way, data sets containing $\mathcal{K} = 100$ trajectories of length $\mathcal{T} = 40$ sample points are generated, for each of the constraints are collected, and a random $\mathcal{N} = 10$ data points are selected as training data. Examples of end-effector trajectories for each constraint are plotted in Fig. 5.9.

With this training data, Algorithm 2 is then used to form an estimate of $\boldsymbol{\Lambda}$. For each $\boldsymbol{\theta}_{s,i}$ defining the unit vectors $\hat{\mathbf{a}}_s$, 180 different values are tested, uniformly spaced in the range $\boldsymbol{\theta}_{s,i} \in [0, \pi)$. Note that, as described in Section 5.2, constraints $\hat{\mathbf{a}}_s$ are iteratively added, subject to $\alpha_s^* \perp \alpha_i \, \forall i < s$ until an increase in (5.6) is seen. In this way, no prior knowledge of the constraint dimensionality is needed for learning.

To assess performance, the errors are computed according to the evaluation criteria described in Section 5.4.1 on independent test data generated according to the same procedure as above.

The experiment was repeated for 50 trials for each data set, and each constraint. In all cases, the proposed method was able to estimate the constraint with NPPE $< 10^{-3}$ and NPOE $< 10^{-4}$.

In many scenarios, it would be useful to predict the behavioural outcomes of using a new policy $\boldsymbol{\pi}'$ to a previously seen environment where the constraints $\tilde{\mathbf{N}}$ have been learnt. This is especially the case for evaluating the use of previously untested policies (e.g., those resulting from optimisation or reinforcement learning) may be too risky to directly evaluate on the robot prior to simulation.

	NUPE	NCPE	NMSE
$\mathbf{\Lambda}_A$	7.50 ± 0.14	6.00 ± 0.09	6.00 ± 0.09
$\mathbf{\Lambda}_B$	7.26 ± 0.12	5.77 ± 0.10	5.77 ± 0.10
$\mathbf{\Lambda}_C$	6.74 ± 0.07	4.45 ± 0.05	4.45 ± 0.05
$\mathbf{\Lambda}_D$	6.72 ± 0.09	4.40 ± 0.08	4.40 ± 0.08

Table 5.3: Normalised UPE and CPE for generalising the joint-limit avoidance policy and the NMSE when applied the learnt constraints (from linear attractor policy) on the learnt policies. The results are (mean$\pm$s.d.)$\times 10^{-1}$ over 50 trails with different data sets.

To test the use of the proposed approach for this, we also evaluated the quality of the learnt constraint, in predicting the constrained actions of a new policy, not present in the training data. In the results reported here, the accuracy in predicting the constrained action of a joint-limit avoidance policy, $\boldsymbol{\pi}(\mathbf{x}) = -0.5 \bigtriangledown \Phi(\mathbf{x})$ with the potential given by $\Phi(\mathbf{x}) = \sum_{i=1}^{7} |x_i|^2$, under the learnt constraint, is evaluated. We apply the constraints learnt from the linear attractor policy to this policy, and the resulting NPPE and NPOE were $< 10^{-3}$. Note that the constraint is independent of the policies, and the learnt constraint is expected to perform well under novel constraints.

5.4.4 *Combined Constraint and Policy Learning*

The goal of this final experiment is to demonstrate the use of the proposed approach in the context of constrained motion imitation learning. In many every day behaviours, it is useful to be able to form an estimate both of the policy underlying motion, as well as the constraint itself. In this way, generalisation can be achieved both across constraints (i.e., applying the learnt policy to new constraints), as well as within constraints (i.e., applying new policies to the learnt constraint). To test this, the proposed approach is combined with that proposed in Howard et al. (2009) for learning models of both $\boldsymbol{\pi}$ and $\mathbf{N}$.

Using the data collected under constraints $\mathbf{\Lambda}_A$, $\mathbf{\Lambda}_B$, $\mathbf{\Lambda}_C$ and $\mathbf{\Lambda}_D$ from the preceding experiment (ref. Section 5.4.3), constraint consistent learning Howard et al. (2009) is used to estimate a model of the policy $\boldsymbol{\pi}$. In more detail, we used linear regression to learn a policy $\tilde{\boldsymbol{\pi}}$ for the data set generated from the linear policy. For the limit-cycle avoidance policy, we used parametric models consisting of 250 Gaussian RBFs with centres chosen according to k-means and with widths taken as the mean of the distances

between centres. To assess learning performance, the normalised unconstrained policy error (NUPE), constrained policy error (NCPE) and mean squared error (NMSE) are evaluated (see Appendix, equations (4.6), (4.7) and (4.8), respectively) over 50 trials of learning.

While we can make a nearly perfect generalisation for the linear policy (with NUPE and NCPE $< 10^{-9}$), the joint-limit-avoidance policy turned out to be a hard problem to learn. The average results over 50 trails are summarised in Table. 5.3. From the NUPE and NCPE, we can see that there are estimation error despite that fact that we tried to learn the policy with a much complex model. The NMSE measures the differences between the estimated constrained policies $\tilde{\mathbf{N}}\tilde{\boldsymbol{\pi}}$ and the observed ones $\mathbf{u}$.

For the linear attractor policy, we were able to obtain NMSE$< 10^{-9}$. The result confirms that with good enough estimation of the constraints and policies, we can accurately predict the constrained policy, even without prior knowledge of the true policy. The NMSE for the joint-limit avoidance policy is listed in Table 5.3. Note that the difference between NCPE and NMSE is that NCPE uses the true projection matrix and NMSE uses the learnt one. Since we have a nearly perfect estimation of the constraints, the resulting NCPE and NMSE is expected to be the same.

We also evaluated whether we can use the previously learnt behaviours in a new environment which was not present in the data for training the policies.

For instance, giving demonstrated wiping motion on surface with 0° and 30° inclination (e.g., $\boldsymbol{\Lambda}_A$ and $\boldsymbol{\Lambda}_B$), can we adapt the behaviour on another surface with a different slope? For this, we created a new constraint $\boldsymbol{\Lambda}_E = (\cos\frac{\pi}{4}, \sin\frac{\pi}{4}\cos\frac{\pi}{3}, \sin\frac{\pi}{4}\sin\frac{\pi}{3})$ and followed the same procedure to generate trajectories. We then approximated the constraint and applied the learnt constraint on the learnt policies.

Fig. 5.10 shows the end-effector position when using the learnt $\boldsymbol{\Lambda}_E$ on both the linear policy (Fig. 5.10a) and the joint-limit avoidance policy (Fig. 5.10b). The plots from the left to the right were the visualisation in x, y, and z position of the end-effector. The red colour denotes the true constrained movement generated from the true policy ($\mathbf{N}\boldsymbol{\pi}$), the black colour denotes the estimated constrained movement by applying the learnt constraint onto the true policy ($\tilde{\mathbf{N}}\boldsymbol{\pi}$), and the blue colour denotes the estimation by learning both constraint and the policy ($\tilde{\mathbf{N}}\tilde{\boldsymbol{\pi}}$).

For the joint-limit avoidance policy (Fig. 5.10b), there is a deviation in the end-effector position when applying the learnt constraint on the learnt policy (the blue trajectories). However, the error is probably the result of a weakly learnt policy, and this can be verified from the NCPE and NMSE in Table 5.3. On the other hand, if the con-

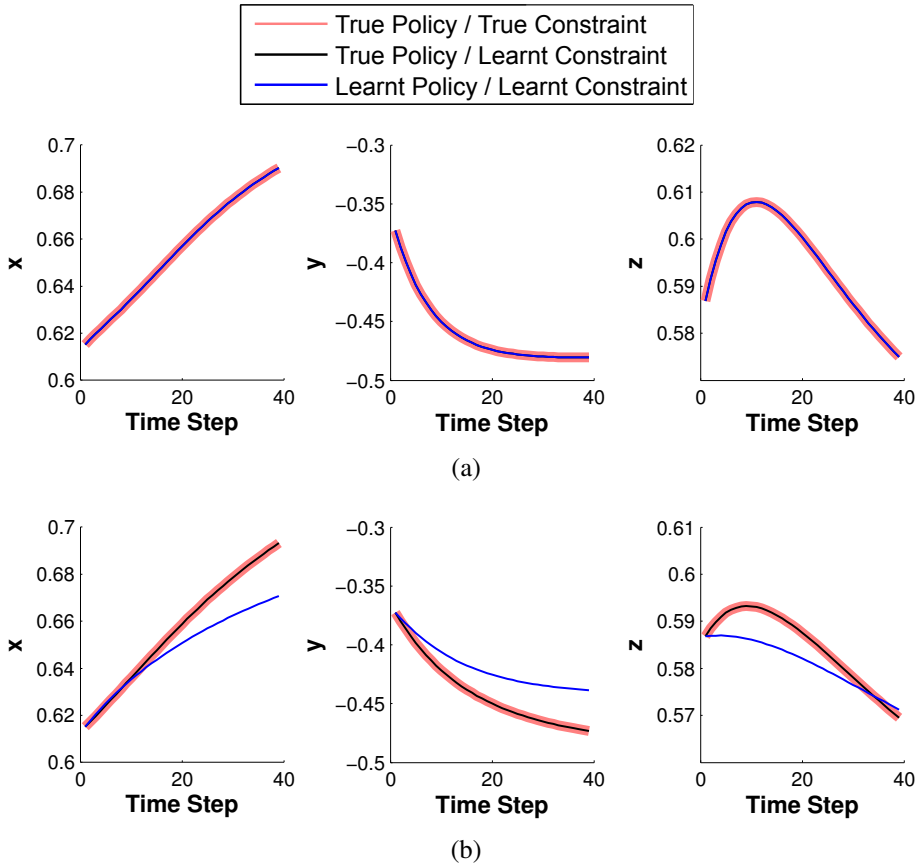

Figure 5.10: Visualisation of end-effector trajectories from (a) linear policy and (b) joint-limit avoidance policy. From the left to the right are the x, y, and z position of the end-effector. The colours denote constrained movement generated by the true constraint and the true policy (red), the constrained movement by applying the learnt constraint on the true policy (black), and by applying the learnt constraint on the learnt policy (blue).

straints and the policy were estimated accurately, such as the case for the linear policy (Fig. 5.10a), we could reproduce the same end-effector movements in all directions.

5.4.5 *Simulated Walking Data*

After validating our approach with stationary constraint systems, in this section, we demonstrate this idea using artificial walking data. Imagine a scenario in which there are four previously unseen persons, two of which have healthy gait and the other two have pathological gait. The goal is to assess the use of the APD in quantifying the

System	Null-space Policy	Task-constraints
S1	Linear	Same as reference
S2	Linear	Different from reference
S3	Sinusoid	Same as reference
S4	Sinusoid	Different from reference

Table 5.4: New systems for testing. The reference is the policy learnt from the linear dataset in Section 4.4.2

difference between the behaviour of each person and healthy walking, as captured by the learnt reference policy.

5.4.5.1 *Linear policy as reference policy*

In our first evaluation, we used the learnt linear policy from Section 4.4.2 as the reference policy (i.e., $\boldsymbol{\pi}^{\text{ref}}$ in Fig. 5.2), and compared this against data from four additional 2-link systems representing four previously unseen subjects (we will refer to them as S1, S2, S3, S4).

Among these four subjects, we used S1 and S2 to represent the normal gait, i.e., followed the same linear policy as the reference. We used S3 and S4 to represent the abnormal gait; namely, they followed a different (sinusoidal) policy $\boldsymbol{\pi}^{\sin}(\mathbf{x}) = -\sin(\mathbf{x}^* - \mathbf{x})$, where $\mathbf{x}^* = (-90^\circ, -25^\circ)$ was the null-space target.

The new subjects took the same phase-divisions and task-space policies as described in Section 3.2.2. To explore the effect of different task-constraints, S1 and S3 had the same task-constraints with one of subjects in the reference data, while S2 and S4 had a totally different one. Note that the null-space policy is independent of the task-constraints, and therefore, the results are expected to be consistent even when the task-constraints change. The setup of these four systems are listed in Table 5.4.

Algorithm 3 was applied on each new system separately. More specifically, we learnt the $\mathbf{u}^{\text{ns},new}$ and $\mathbf{N}^{new}$ for each phase of each person. The differences between each system and the reference were calculated using APD (5.21). We also tested linear regression and RBF network for comparison. We took the learnt model in Section 4.4.2, and measured the difference in joint-space between the reference policy and each new system.

Fig. 5.11 summarises the results of comparing the reference policy to S1-S4 using three different methods. The y-axis shows the average differences in joint-space for

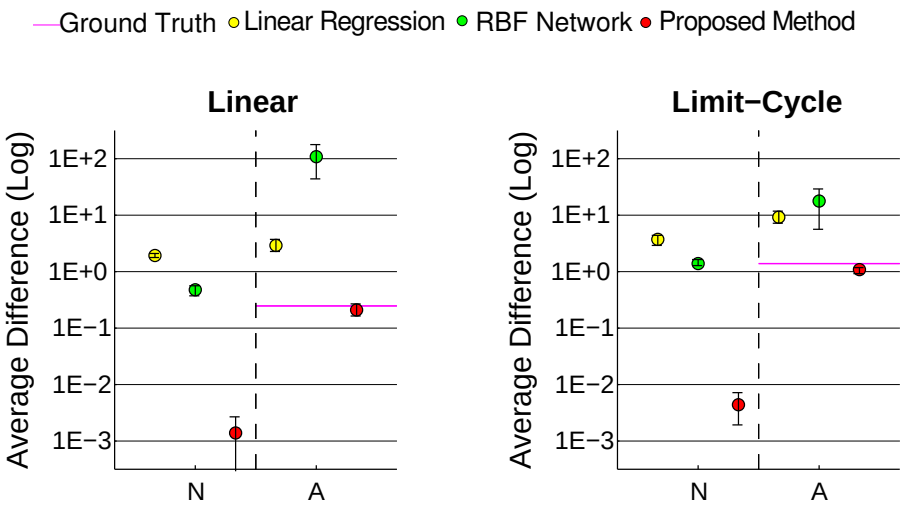

(a) (b)

Figure 5.11: Average difference between the testing systems and the reference, where the reference is (a) a linear policy and (b) a limit-cycle policy. The error-bars are mean±std.dev. in log-scale over ten experiments. The results were grouped into normal (N) and abnormal (A), where N is the average of S1 and S2, and A is the average of S3 and S4.

the regression methods and the average APD for the proposed method. (Note that, the proposed method attempts to eliminate the difference resulting from various walking behaviours; therefore, our method does not directly compare joint-angles, which is influenced by various behaviours.) The error bars are the mean±std. dev. in log scale over ten experiments.

In Fig. 5.11a, the horizontal line is the true differences between the reference (linear) and S3-S4 (sinusoidal), which is the true difference between $\mathbf{N}\boldsymbol{\pi}^{\sin}$ and $\mathbf{N}\boldsymbol{\pi}^{\mathrm{ref}}$. Note that S1 and S2 adapt the same linear policy as the reference, and the true differences between S1, S2 and the reference are 0.

In Fig. 5.11a, the yellow and green colours denote the results of standard methods. We can see that linear regression fails to differentiate linear and sinusoidal policy. Although RBF network predicts relatively higher difference for abnormal (A), the predicted difference for normal (N) is also unreasonably high. The red colour denotes the results of our approach. Our method yields lower error for normal (N), and the result confirms with the fact that S1 and S2 use the same linear policy. The error for abnormal (A) is relatively higher, which is also expected since S3 and S4 adopt the sinusoidal policy.

System	Null-space Policy	Task-constraints
S1	Limit-cycle	Same as reference
S2	Limit-cycle	Different from reference
S3	Sinusoid	Same as reference
S4	Sinusoid	Different from reference

Table 5.5: New systems for testing. The reference is the policy learnt from the limit-cycle dataset in Section 4.4.2

5.4.5.2 *Limit-cycle policy as reference policy*

After validating with a linear policy, we also tested our method on a non-linear policy. We took the limit-cycle policy learnt in Section 4.4.2 as the reference policy. S1 and S2 were generated using the same limit-cycle policy as the reference (see Table 5.5).

In Fig. 5.11b, the horizontal line shows the true difference between the reference policy (limit-cycle) and S3-S4 (sinusoidal), which is the true distance between $\mathbf{N}\,\boldsymbol{\pi}^{\sin}$ and $\mathbf{N}\,\boldsymbol{\pi}^{\text{ref}}$.

Similar to the last experiment, the standard methods (yellow and green) fail to make reasonable predictions. Our method (red), again, produces results relatively similar to the ground truth. Our experiment demonstrates good assessments on quantifying the difference between policies, even if the true policy, the constraints, and the tasks are unknown.

5.5 DISCUSSION

In this chapter, we discussed our approach for quantifying the difference between walking gaits. In particular, we begin by learning the *variations in embodiments* and examining the difference under variations.

In Section 5.2, we begin by exploring the problem of learning null-space projection for stationary constraint systems ($\mathbf{A}\,\mathbf{u} = 0$) and ways to cope with unknown and unpredictable environment. In Section 5.3, we described the method for learning the projection for our walking phase model.

For both cases, we validated our approach on simulated data. In particular, we tested our proposed method on learning the null space projection matrix of a kinematically constrained system and simulated walking data. Our experiment has demonstrated that

our approach is effective in reconstructing the constraints, without prior knowledge of the true policy and the dimensionality the constraints. We validated our method on the generalisation across constraints (i.e., applying the learnt policy to new constraints), as well as within constraints (i.e., applying new policies to the learnt constraint).

In the next chapter, we will use the proposed method to compare the difference between gaits on human walking data where the true policy and constraint are both unknown.

6

EXPERIMENTS WITH HUMAN DATA

"If a man were to walk on the ground alongside a wall with a reed dipped in ink attached to his head the line traced by the reed would not be straight but zig-zag, because it goes lower when he bends and higher when he stands upright and raises himself."

—Aristotle [1]

Previous chapters described our computational framework for generalising the characteristics of walking gait and quantifying the difference between walking gaits. We validated our approach on simulated walking data, and we demonstrated that, under ideal assumptions, our algorithms can recover an unconstrained policy that capture the characteristics of walking gait. In this chapter, we explored the utility of our approach on a more realistic setting; i.e., working with human gait data.

Our analysis will be based on kinematic and kinetic data of subjects walking with various speeds. Our goal is to see whether we can (1) extract consistency across walking behaviours and subjects and (2) use the extracted gait to quantify the difference between normal and pathological gaits.

In Section 6.1 and 6.2 we describe how kinematic and kinetic data were collected and processed. In Section 6.3, we outline our experiments based on kinematic features. In Section 6.4, we discuss the experiments on kinetic features and the results of subsequent experiments. In Section 6.5, we compare our method to a model-based approach on generalising human walking data.

1 In *parts of animals, movement of animals, progression of animals* (Translated by Peck A), Aristotle wrote his conjectures on the movement of walking gait in the vertical direction. However, in the era he lived, scientific truth could be claimed without justification by experiment (Baker, 2007).

(a) (b)

Figure 6.1: Kinematic data was collected using Xsens MVN BIOMECH motion capture system. (a) An inertial tracker compromises a 3D gyroscope, a 3D accelerometer, and a 3D magnetometers. (b) A subject wearing the Xsens MVN BIOMECH units according to Xsens configuration.

6.1 HUMAN WALKING DATA

In this section, we describe how our human walking data was collected. In particular, kinematic data was recorded using motion capture (Section 6.1.1), and kinetic features was recorded using force platforms (Section 6.1.2).

6.1.1 *Kinematic Features*

The kinematic data was collected using Xsens MVN BIOMECH system (Roetenberg et al., 2009) . The Xsens MVN consists of 17 inertial and magnetic trackers, and each tracker compromises a 3D gyroscope, a 3D accelerometer, and a 3D magnetometers (Fig.6.1a). The gyroscope measure the angular velocity of the limb, which provides the joint-angle and the joint-velocity. The accelerometer measures the acceleration due to gravity, and the magnetometer measures the earth magnetic field. These two quantities together determine the direction of the subject. *Xsens*

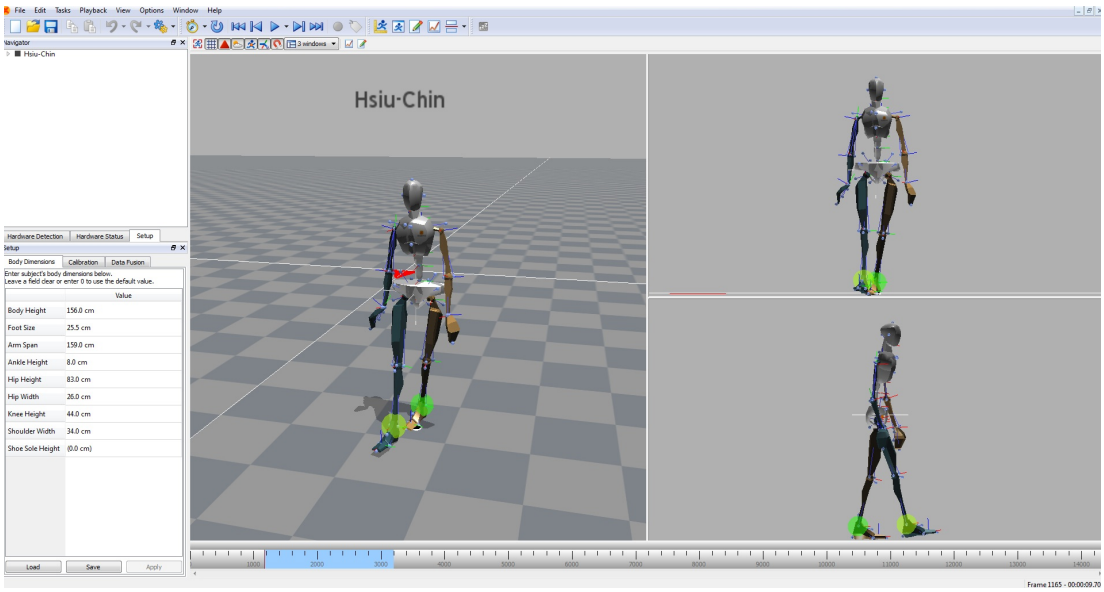

Figure 6.2: Xsens MVN Studio, a graphical interface provided by Xsens, processes and displays the captured motion in real time

The sensor units were attached to the subjects according to Xsens configuration (Fig. 6.1b). The update frequency is set to 120 FPS. The data are transmitted via wireless communication to the machine where the data can be processed and visualised in real-time.

The captured motion was processed using MVN Studio 3.0 (Fig. 6.2), a graphical interface provided by Xsens. The data was exported in MVNX format, which is an XML format with full kinematics of each segment, including position, velocity, acceleration, orientation, angular velocity and angular acceleration.

The desired parameters were calculated and preprocessed using MATLAB. Fig. 6.3a and Fig. 6.3b are examples of hip and knee joint trajectories collected using Xsens MVN BIOMECH.

6.1.2 *Kinetic Features*

V-Gait

The kinetic data was collected using V-Gait System from Motek Medical (Fig. 6.4a). The system is based on an instrumented dual-belt treadmill with two force platforms embedded. The system is controlled through the D-flow software where the user can modify the gait speeds, as well as sway, pitch, and inclination of the treadmill. Each force plate-form records the 3-dimensional *ground-reaction-force* (GRF) and computes the position of the *centre-of-pressure* (CoP).

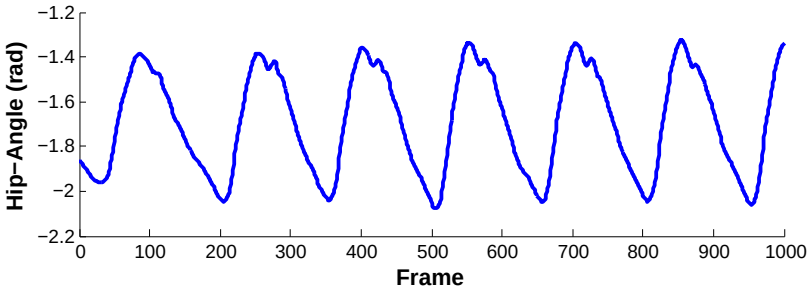

(a)

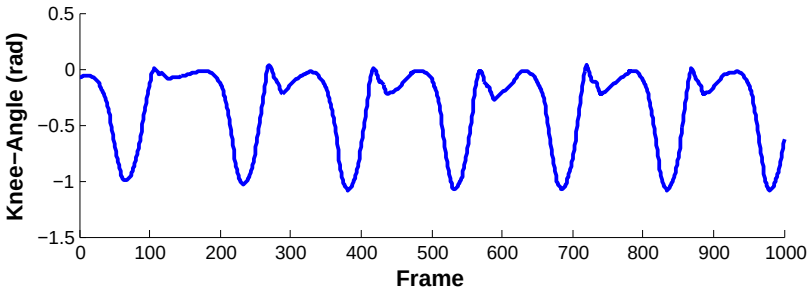

(b)

Figure 6.3: An example of (a) hip angle and (b) knee angle trajectory collected using Xsens MVN BIOMECH

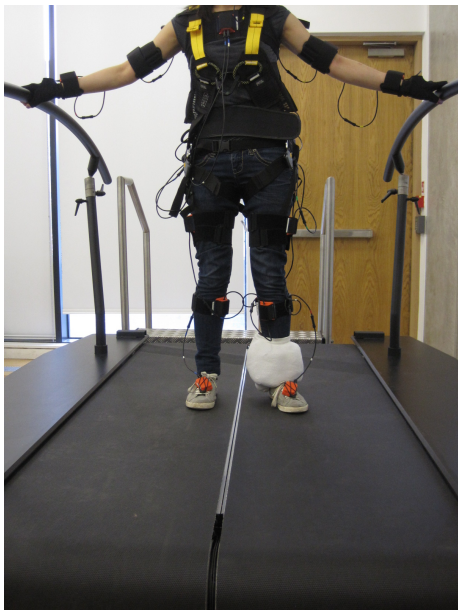

(a)

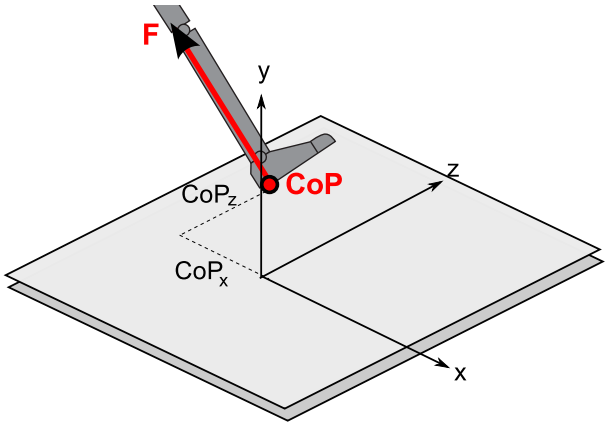

(b)

Figure 6.4: Facilities for collecting kinetic data. (a) A subject stands on the instrumented dual-belt treadmill with two force platforms embedded. (b) An illustration of a force plate-form.

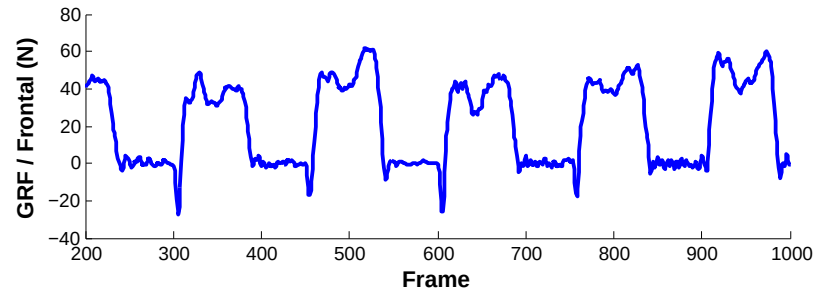

(a)

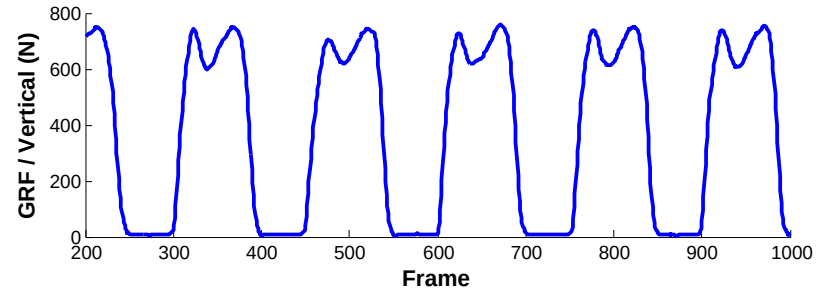

(b)

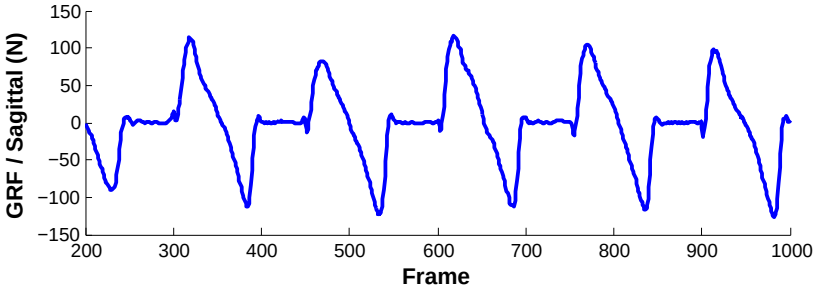

(c)

Figure 6.5: Examples of raw data collected using V-gait system. The trajectories are ground reaction force of the left leg in (a) frontal, (b) vertical, and (c) sagittal plane

The *ground-reaction-force* (GRF) is the weighted average of all forces acts on the foot during standing. We use $F = (F_x, F_y, F_z) \in \mathbb{R}^3$ to denote the direction and the magnitude of GRF where F_x, F_y, F_z are the components of GRF in the frontal, vertical, and sagittal plane, respectively.

The trajectories shown in Fig. 6.5 are GRF acquired with a subject walking in his preferred speed on V-gait, where Fig. 6.5a, Fig. 6.5b, and Fig. 6.5c are the F_x, F_y and F_z trajectories of the left leg, respectively.

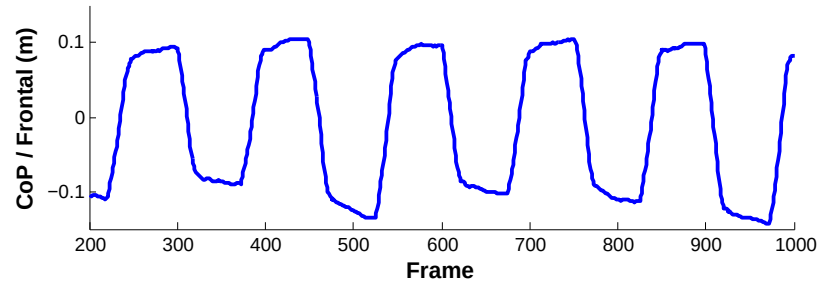

(a)

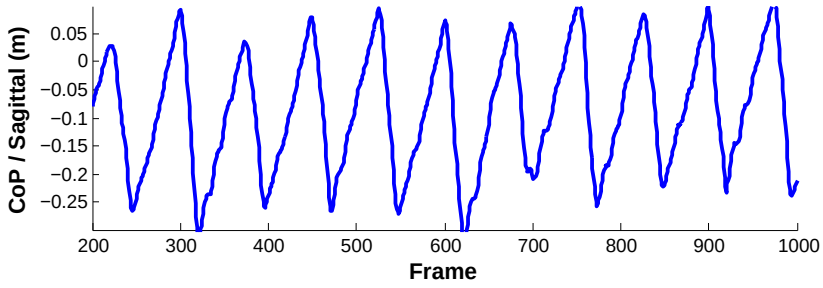

(b)

Figure 6.6: Examples of raw data collected from a subject walking with his preferred speed on V-gait. The trajectories are the position of centre-of-pressure in the (a) frontal plane and (b) Sagittal plane.

By definition, the quantity of GRF is equivalent to the sum of all products of mass and acceleration of all body segments; namely,

$$F_x = \sum_{i=1}^{N} m_i a_{x_i}, \quad F_y = \sum_{i=1}^{N} m_i a_{y_i} + g, \quad F_z = \sum_{i=1}^{N} m_i a_{z_i} \tag{6.1}$$

where m_i is mass of the i^{th} segment, a_{x_i}, a_{y_i}, a_{z_i} are the linear acceleration of CoM of the i^{th} segment, and g is the acceleration due to gravity.

Another measurement we can obtain from the force platform is the *centre-of-pressure* (CoP) position with respect to the coordinate system of the force platform (Fig. 6.4b). The CoP is the point where the GRF locates, and it is equivalent to the weighted average of all the forces over the surface in contact with the ground.

An example of CoP trajectories collected using V-gait is shown in Fig. 6.6 where Fig. 6.6a and Fig. 6.6b are the CoP position of the left leg in frontal and the sagittal plane.

During single support phase, the CoP lies within the foot of the supporting leg (Fig. 6.7a). During double support phase, there are two separate CoPs under each foot,

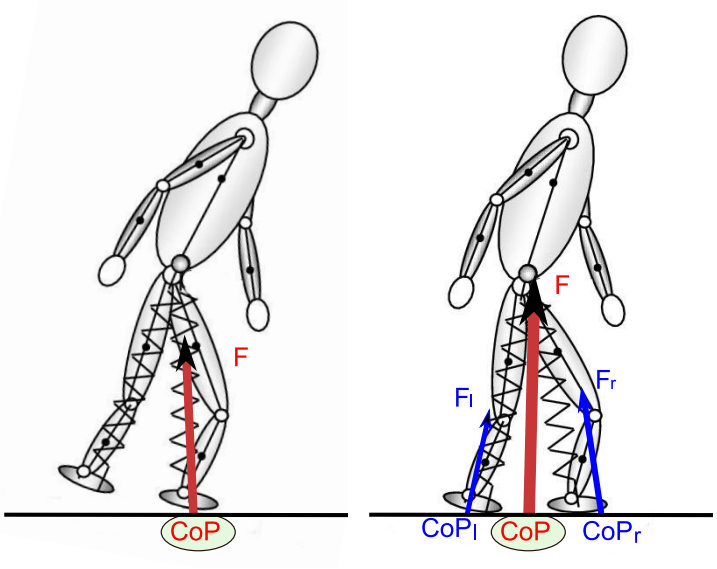

(a) (b)

Figure 6.7: An illustration of the centre-of-pressure position and ground-reaction-force. (a) During single support phase, the CoP lies within the foot of the supporting leg. During double support phase, there are two separate CoPs and GRFs

and the total CoP lies between the two feet, depending on the relative force applied by each foot (Fig. 6.7b); specifically, the whole body CoP and GRF are defined as

$$\begin{aligned} \mathrm{CoP} &= \mathrm{CoP}_l \frac{F_{y,l}}{F_{y,l}+F_{y,r}} + \mathrm{CoP}_r \frac{F_{y,r}}{F_{y,l}+F_{y,r}} \\ F &= F_l \frac{F_{y,l}}{F_{y,l}+F_{y,r}} + F_r \frac{F_{y,r}}{F_{y,l}+F_{y,r}} \end{aligned} \tag{6.2}$$

where CoP_l and CoP_r are the CoP under the left and the right foot, and $F_{y,l}$ and $F_{y,r}$ are the vertical ground reaction force under left and right foot.

6.2 PRE-PROCESSING

In this section, we describe how the data recorded in the preceding section were processed to get the input for our optimisation framework. We outline our procedures to divide the raw data into gait cycles and gait phases (Section 6.2.1 and Section 6.2.2). We demonstrate how the trajectories collected from force-platforms were transformed into centre-of-mass displacement (Section 6.2.3).

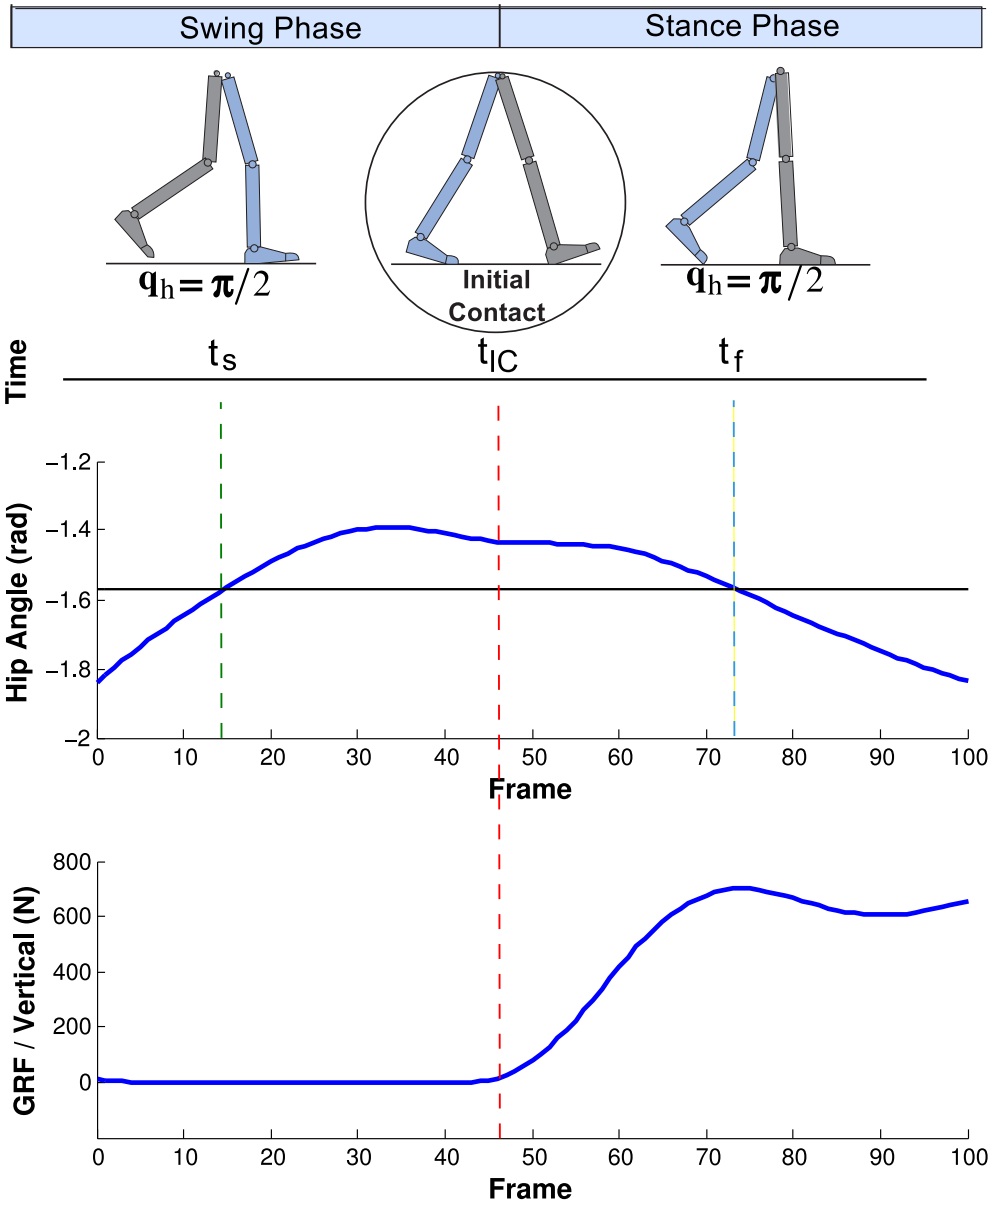

Figure 6.8: Between each pair of mid-swing and mid-stance phase, the initial-contact event can be chosen as the first sample where the vertical ground reaction force is greater than zero.

6.2.1 *Gait Cycle Extraction*

We followed the definition of Whittle (2007) to determine the gait cycle and gait events and used *initial-contact* of the left leg to demarcate the beginning of a walking cycle (see Section 2.1.1 for details).

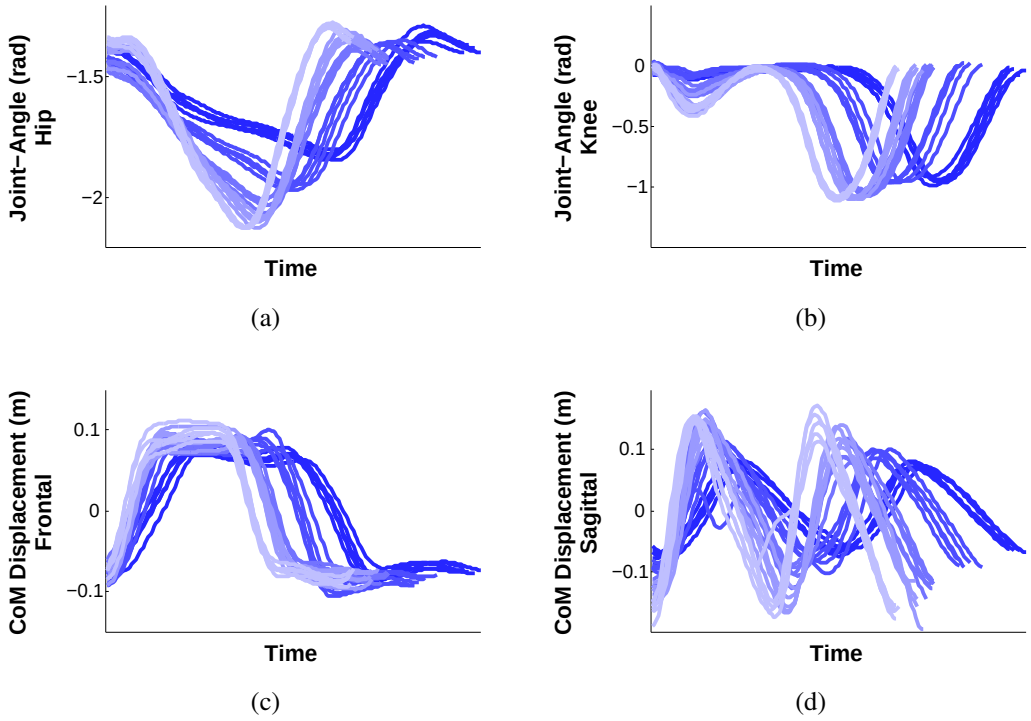

Figure 6.9: Examples of extracted gait cycles from motion-capture data of one subject walking with five different speeds. The plots are (a) hip angle, (b) knee angle, (c) CoM displacement in the frontal plane, and (d) CoM displacement in the sagittal plane. The darker the colour, the slower the speed.

First, we need to detect the *initial-contact* events from a continuous recorded signal. We proceed by noting that, there is an initial-contact event between every pair of mid-swing phase and a mid-stance phase (Fig. 6.8).

To narrow down the range of search, we can select a time point t_s from mid-swing phase and another time point t_f from the mid-stance phase. During the mid-swing phase, we know that, the femur of the left leg swings forward, and it is perpendicular to the surface of walking exactly once (the green dash line in Fig. 6.8). For this, we can find all t_s such that $\mathbf{q}_h(t_s) < -\pi/2$ and $\mathbf{q}_h(t_s + 1) > -\pi/2$. By using the same strategy, we can find t_f such that $\mathbf{q}_h(t_f) > -\pi/2$ and $\mathbf{q}_h(t_f + 1) < -\pi/2$ in the mid-stance phase (the blue dash line in Fig. 6.8).

Within this range, the initial-contact event can be detected as the first sample in which the vertical ground reaction force of the left leg F_y is greater than some threshold value ϵ (the red dash line in Fig. 6.8). Specifically, we seek the time of initial-contact t_{IC} such that $F_y(t_{IC}) > \epsilon$ and $t_s < t_{IC} < t_f$. (In our experiment, the threshold value ϵ was set to 10 N.)

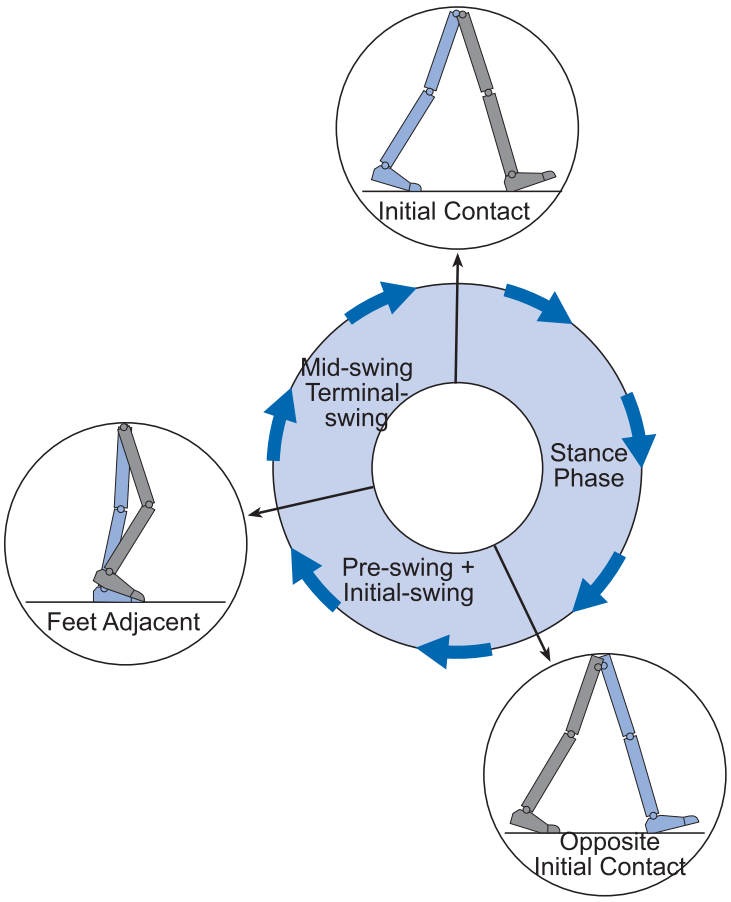

Figure 6.10: Phase division used in this experiment

From the raw data, we can detect the initial-contact events as many as possible, and the interval between two consecutive initial-contact events can be extracted as a gait cycle. Fig. 6.9 shows examples of extracted gait cycles from the trajectories, where the top row shows the kinematic features (hip and knee angles) and the bottom row shows the CoM displacement in the frontal and sagittal plane.

6.2.2 *Phase Division*

In all experiments carried out in this chapter, we divided each walking cycle into three phases: (i) stance, (ii) pre-swing and initial swing, and (iii) mid-swing and terminal swing (See Fig 6.10). Note that using three phases is not conventional or necessarily optimal. However, this seems to be the smallest sub-divisions of the walking cycle that guarantees a consistent constraint $\mathbf{A}_k$ across the entire k^{th} phase, which is a condition for our method to be effective (Section 4.3). To extract these three phases, we need to detect the *opposite initial contact* and *feet-adjacent* events

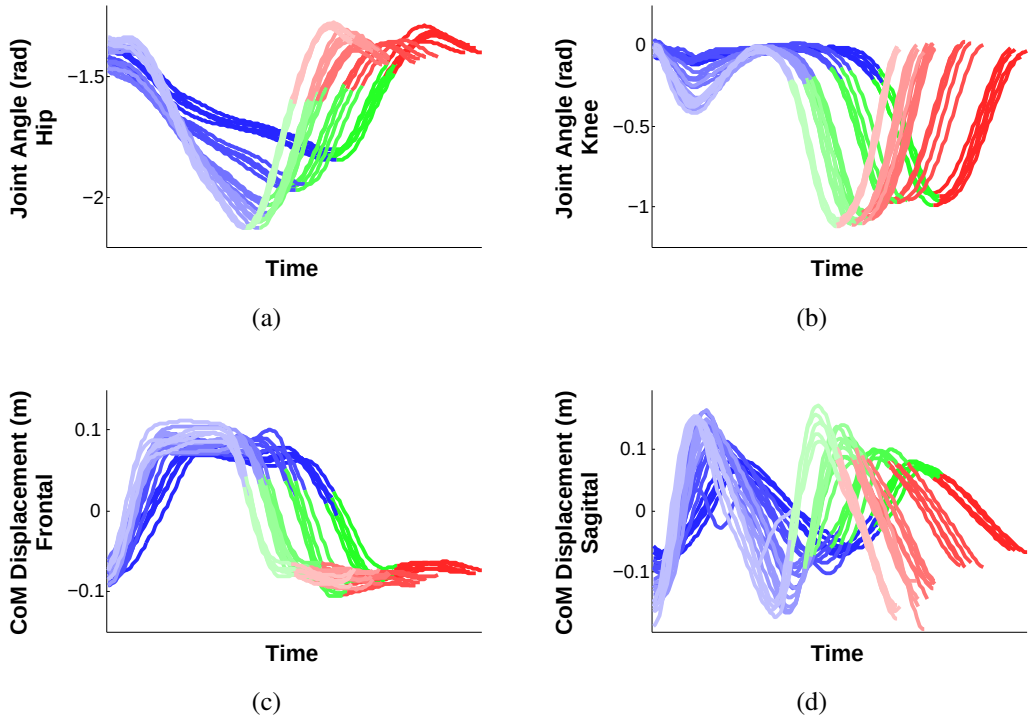

Figure 6.11: Examples of extracted gait phases from data collected with five different normal speeds. The colours denote the gait phases. The figures are (a) hip angle, (b) knee angle, (c) CoM in the frontal plane, and (d) CoM in the sagittal plane.

The *opposite initial-contact* event occurs when the right leg is in contact with the walking surface after its swing phase. This event can be determined by using a similar procedure to the one for finding the *initial-contact* event. Namely, we can seek a time t_{OI} between mid-stance phase and mid-swing phase such that the vertical component of the ground reaction force of the right leg is greater than the threshold value.

The *Feet Adjacent* is the time point when the knee angle $\mathbf{q}_k$ is at its maximum flexion in the swing phase. Hence, we seek a time where the knee angle is at the minimum between opposite initial-contact and the following initial-contact event; i.e., we seek a time t_{FA} such that $\mathbf{q}_k(t_{FA}) = \min(\mathbf{q}_k)$ for $t_{OI} < t_{FA} < t_{IC}$.

Fig. 6.11 shows examples of gait phases divided by initial-contact, opposite initial contact, and feet-adjacent. The red colour denotes the stance phase, the green colour denotes the pre-swing and the initial-swing phase, and the blue colour denotes the mid-swing and the terminal-swing phase. Note that, the gait cycles were taken from the same data in Fig. 6.9.

6.2.3 *Estimation for Centre-of-Mass Displacement*

As mentioned in Section 2.1, the dynamical stability of the gait is highly correlated to the centre-of-mass displacement, or the distance between centre-of-mass (CoM) and centre-of-pressure (CoP). Note that, the location of CoP can be measured using equipment such as force platform, but the location of CoM is not directly observable since we have no access to the mass and distribution of each body segment (see Section 3.3 for details). In order to include CoM displacement in our analysis, we use the *gravity-line-projection* (GLP) method to approximate the components of CoM in the frontal and sagittal plane (Zatsiorsky and King, 1997).

After recording walking data using force platform, we obtain trajectories of (i) the CoP position $CoP(t) = [CoP_x(t), CoP_z(t)]$ and (ii) the direction and magnitude of GRF $F(t) = [F_x(t), F_y(t), F_z(t)]$ where the subscripts x, y, z denote the frontal, vertical, and sagittal plane. The GLP approximates $x(t)$ and $z(t)$, the position of CoM in the frontal and the sagittal plane, respectively. In this section, we demonstrate this method by estimating $x(t)$ as an example.

First, the linear velocity of centre-of-mass $v(t)$ can be described as the integration of the instantaneous acceleration.

$$\begin{aligned} v(t_1) &= v(t_0) + a(t_0)dt \\ v(t_2) &= v(t_1) + a(t_1)dt = v(t_0) + a(t_0)dt + a(t_1)dt \\ &\vdots \\ v(t_n) &= v(t_0) + \sum_{j=1}^{n-1} a(t_j)dt \end{aligned} \tag{6.3}$$

The position of CoM in horizontal direction $x(t)$ can be described as the integration of the instantaneous velocity.

$$\begin{aligned} x(t_1) &= x(t_0) + v(t_0)dt \\ x(t_2) &= x(t_1) + v(t_1)dt = x(t_0) + v(t_0)dt + v(t_1)dt \\ &\vdots \\ x(t_n) &= x(t_0) + \sum_{i=1}^{n-1} v(t_i)dt \end{aligned} \tag{6.4}$$

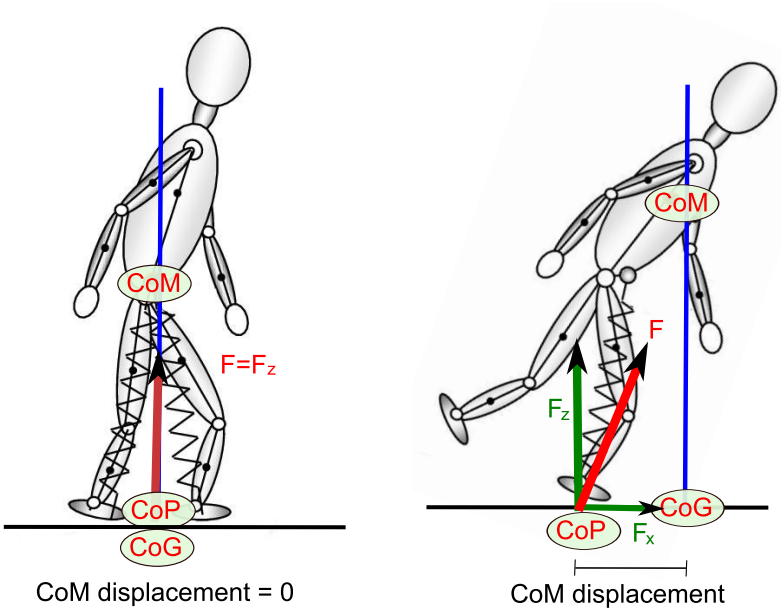

Figure 6.12: Relationship between F_x, CoP, and CoM. (a) When $F_x == 0$, CoM, CoG, and CoP coincide. (b) When $F_x \neq 0$, F_x is proportional to the linear acceleration of CoG, and the CoM displacement is the distance between CoM and CoP

By substituting $v(t_i)$ in (6.3) into (6.4), we get

$$\begin{aligned} x(t_n) &= x(t_0) + \sum_{i=1}^{n-1} \left[v(t_0) + \sum_{j=1}^{i-1} a(t_j) dt \right] dt \\ &= x(t_0) + \sum_{i=1}^{n-1} v(t_0) dt + \sum_{i=1}^{n-1} \left[\sum_{j=1}^{i-1} a(t_j) dt \right] dt \\ &= x(t_0) + n \times v(t_0) \times dt + \sum_{i=1}^{n-1} \sum_{j=1}^{i-1} a(t_j) dt^2 \\ &= x(t_0) + n \times v(t_0) \times dt + \sum_{i=1}^{n-1} \sum_{j=1}^{i-1} F_x(t_j) \frac{dt^2}{M} \end{aligned} \tag{6.5}$$

where M is the total body mass and F_x is ground reaction force in the frontal plane. Hence, the position of the CoM at any given time t_n can be described by its initial position $x(t_0)$, initial velocity $v(t_0)$, and ground reaction force F_x from t_0 to t_n.

However, to use (6.5), we need to know the values of $x(t_0)$ and $v(t_0)$, (i.e., the initial position and initial velocity of CoM), which are not observable, either. Although CoP and CoM are not directly correlated, we can see that, when the ground reaction force equals to 0, CoM, CoG, and CoP coincide on the gravitational projection line

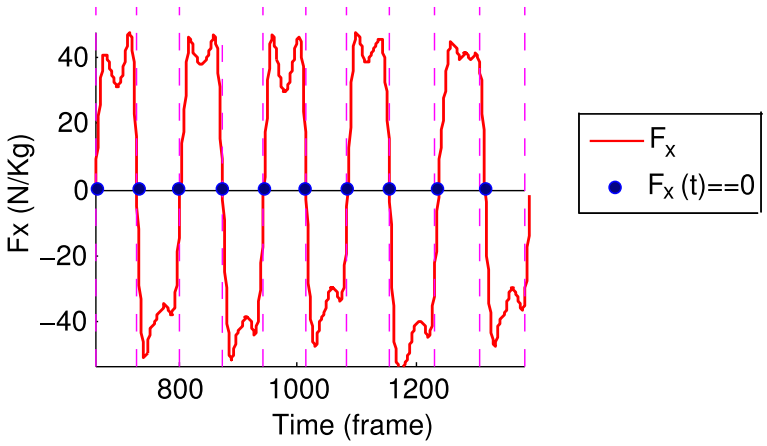

Figure 6.13: An example of ground reaction force in the frontal plane

(Fig. 6.12a). Namely, if t_s is the time such that $F_x(t_s)$ is 0, the position of CoM at t_s is

$$x(t_s) = CoP(t_s)\,|_{F_x(t_s)=0} \tag{6.6}$$

Note that, for normal subjects, this event occurs twice in a walking cycle (i.e., approximately once in the mid-swing phase and once in mid-stance phase). Fig. 6.13 is an example of the GRF trajectories F_x. When F_x crosses the zero-line (the blue points), we can use the CoP at that time as the initial position of the CoM.

After determining the value of the initial position $x(t_s)$, we also need to determine the initial velocity $v(t_s)$. Let t_f be the next time that F_x equals to 0. From (6.5), we know the position of CoG at t_f is

$$x(t_f) = x(t_s) + (t_f - t_s) \times v(t_s) \times dt + \sum_{i=s}^{f-1} \sum_{j=1}^{i-1} F_x(t_j) \frac{dt^2}{M}$$

Solving $v(t_s)$ in the above equation, we get

$$v(t_s) = \frac{x(t_f) - x(t_s) - \sum_{i=s}^{f-1} \sum_{j=s}^{i-1} F_x(t_j) \frac{dt^2}{M}}{(t_f - t_s) \times dt}$$

Since $x(t_s) = CoP(t_s)\,|_{F_x(t_s)=0}$, and $x(t_f) = CoP(t_f)\,|_{F_x(t_f)=0}$, we can substitute $x(t_s)$ and $x(t_f)$ by $CoP(t_s)$ and $CoP(t_f)$

$$v(t_s) = \frac{CoP(t_f) - CoP(t_s) - \sum_{i=s}^{f-1} \sum_{j=1}^{i-1} F_x(t_j) \frac{dt^2}{M}}{(t_f - t_s) \times dt} \tag{6.7}$$

Finally, the position of CoM at time t_n for $t_s < t_n < t_f$ is defined as

$$x(t_n) = x(t_s) + v(t_s) + \sum_{i=s}^{n-1} \sum_{j=1}^{i-1} F_x(t_j) \frac{dt^2}{M}$$

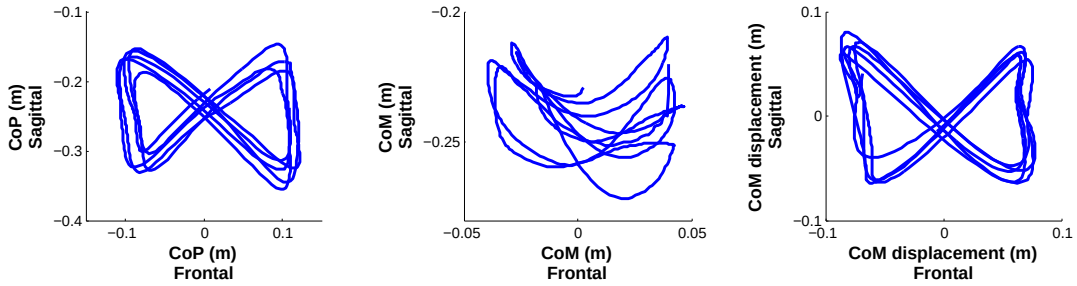

(a) (b) (c)

Figure 6.14: Preprocessed kinetic data: (a) measured centre-of-pressure position (b) approximated centre-of-mass position using gravitational-line-projection method and (c) calculated centre-of-mass displacement

where $x(t_s)$ is the position of CoM at t_s and $v(t_s)$ is the velocity of CoM at t_s defined by (6.7).

We assume that the dynamical stability of the gait is correlated to the CoM displacement, which is the distance between the CoP position and the CoM position (Fig. 6.12b). After recording the centre-of-pressure positions $CoP_x(t), CoP_z(t)$ from the force platforms and estimating the centre-of-mass position $x(t)$ and $z(t)$, we can find the CoM displacement $\rho_x(t) = x(t) - CoP_x(t)$ and $\rho_z = z(t) - CoP_z(t)$.

In Fig. 6.14, we show examples of the post-processed (a) CoP, (b) CoM, and (c) CoM displacement of data recorded from a subject walking on the treadmill at the speed of 1 meter per second for 15 seconds.

6.3 EXPERIMENTS ON KINEMATIC FEATURES

In this experiment, our goal is to test our method on kinematic features. In particular, we would like to see how well our method can generalise over subjects and walking speeds and how well the generalised policy can quantify the difference between normal and pathological gaits.

6.3.1 *Protocol*

The data were collected from two females and seven males, age between 20-29 (referred as S1-S9). These subjects have different body types, and they were chosen to

Subject	Upper Leg (cm)	Lower Leg (cm)	Gender
S1	37.7	38.9	Male
S2	41.6	41.1	Male
S3	42.6	40.5	Male
S4	44.2	43.3	Male
S5	45.1	44.2	Male
S6	42.5	41.5	Male
S7	45.8	44.8	Male
S8	41.4	36.1	Female
S9	51.0	42.1	Female

Table 6.1: Leg lengths of the subjects participated in this experiment

ensure our data contains some variations in embodiments. The leg-lengths of subjects are summarised in Table 6.1.

Data was recorded for five walking speeds: 93, 106, 119, 129, 140 steps per minute, which were taken from the speed range reported in (Öberg et al., 1993). The walking speeds were controlled through use of a metronome. The subjects were asked to walk such that heel strike coincided with the tick of the metronome. For each speed, ten walking trials were collected.

We extracted as many cycles as possible from all walking trails we collected using the pre-processing procedure described in Section 6.2.1. We obtained roughly 200 gait cycles from each subject, and 100 gait cycles were selected for analysis. In this experiment, we tested our approach with three walking phases according to the descriptions in Section 6.2.2.

6.3.2 *Pathological Gait*

To create 'pathological' gait, A 3.5 kg bag of sand was strapped to the subjects' left leg (see Fig. 6.15). We used the same setup (speed, number of trails, etc) to collect kinematic data of abnormal gait from each subject.

For both "normal" and "pathological gait", we used the method discussed in Section 6.2.2 to divide the collected data into walking phases. We obtained roughly 200 gait cycles from each subject and selected the first 100 cycles for analysis. Fig. 6.16 shows the hip and knee angle from one of the subjects. One trajectory from each walk-

Figure 6.15: A 3.5 kg weight was attached to the subject's leg to create 'abnormal' gaits

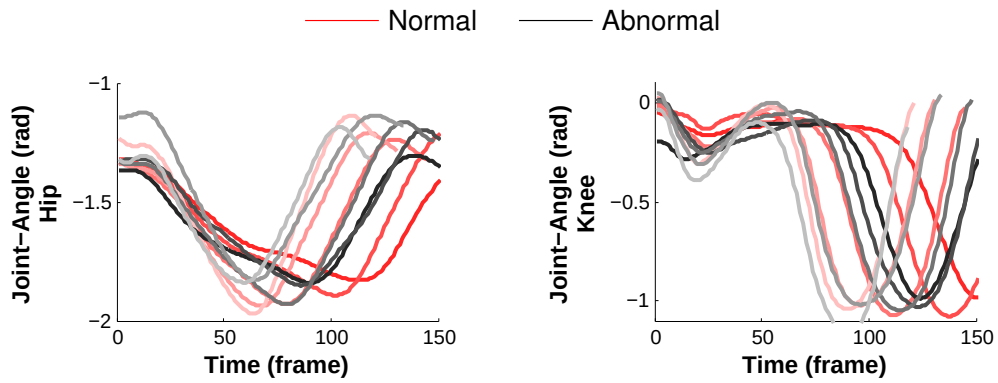

(a) (b)

Figure 6.16: (a) Hip and (b) Knee angle from one subject walking at five different speeds (93, 106, 119, 129, 140 steps per minute). The colours denote the normal gaits (red) and abnormal gait (black)

ing speed was selected. Note that normal gaits (red) and abnormal gaits (black) look very similar from direct observation.

6.3.3 *Baseline*

For comparison, we also trained models using (i) linear regression and (ii) RBF network on raw observations $(\mathbf{x}, \mathbf{u})$ from the normal gaits, and tested if we can see a difference between normal and abnormal gaits in joint-space.

6.3.4 *Learning Reference Policy*

We used seven subjects, including five males (S1-S5) and two females (S8-S9), for learning the reference policy (Algorithm 1). For each walking phase of each subject, we learnt a model for null-space component $\tilde{\mathbf{u}}_{i,k}^{ns,ref}$, which yielded 21 models. Each $\tilde{\mathbf{u}}_{i,k}^{ns,ref}$ consisted of M Gaussian radial basis functions (RBFs) where M varied from 16 to 100. The null-space policy $\boldsymbol{\pi}^{ref}$ was also modelled as parametric model with Gaussian RBFs.

In order to effectively learn a reference policy, we need to have a consistent policy within the dataset, but we are not clear whether the walking behaviour across all subjects can be described by a single null-space policy. Many research have shown that males and females are different in both kinematics and dynamics (Kerrigan et al., 1998; Chumanov et al., 2008). For instance, by examining kinematic features, females normally have larger hip flexion and less knee extension before initial contact. From kinetic features, females have greater knee moment in pre-swing and greater joint power absorption.

From the above observations, the gender difference might affect the performance of our method, so we also train another model using only the male subjects (S1-S5) and see if we can learn a better model. (Since we only have two female subjects, we could not run another separated experiment for females.)

We do not have access to the true null-space policy for human data, so the only way to evaluate the model is by inspecting the optimisation error E_2 (Equ. 4.5). It can be proven that the optimisation error E_2 is the lower bound of the policy error (Howard et al., 2009), if E_2 cannot be minimised, then the resulting model cannot make reasonable predictions.

We ran a 10-folds cross-validations, and the average E_2 for the all subjects (include both male and female) is 0.1493 ± 0.0065 while the average E_2 for the male only data is 0.0433 ± 0.0089. Since gender might be a negative factor, female subjects were eliminated from all the experiments presented later in this thesis. We used the learnt policy from male subjects as the reference policy for gait abnormality detection.

6.3.5 *Identifying Pathological Gaits*

Five subjects (S1-S5) were used to collect five normal and five pathological gaits (using the leg loading). To investigate how well the learnt policy can generalise across sub-

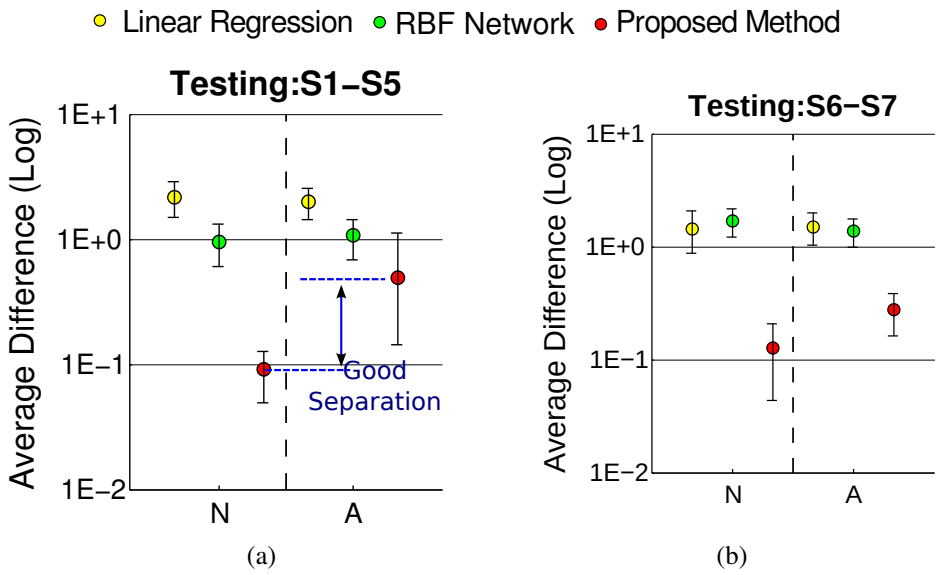

Figure 6.17: Average difference between the testing subjects and the reference policy, where the testing subjects are (a) S1-S5 and (b) S6-S7. The error-bars are mean±std.dev. in log scale over ten experiments and over subjects. The results are grouped into normal (N) and abnormal (A).

jects, we also performed the same experiment on the subjects whose data had not been used for training the reference policy (S6-S7) – we collected normal and pathological gaits for each of these.

We used Algorithm 2 to form an estimate of null-space projection $\mathbf{N}$ for each gait phase of each subject separately. Algorithm 3 was applied to quantify the difference between each subject and the reference gait.

Fig. 6.17 shows the the average results over all subjects. The yellow and green colours denote the results using standard methods, and the red colour denotes the results using our method. The error bars are the mean±std. dev. in log scale over ten experiments.

Fig. 6.17a shows the average results over S1-S5. We can see that the standard methods cannot differentiate normal and abnormal gait. Our approach achieved relatively lower difference (<0.01) when comparing with normal gaits and higher difference when comparing with abnormal gaits. Even if we have no access to the true policy, constraints, nor tasks, our reference policy is more effective in differentiating between normal and abnormal gaits.

Fig. 6.17b shows the the average results over S6 and S7. In this case, RBF based approximation predicts that the abnormal gait is more similar to the reference gait.

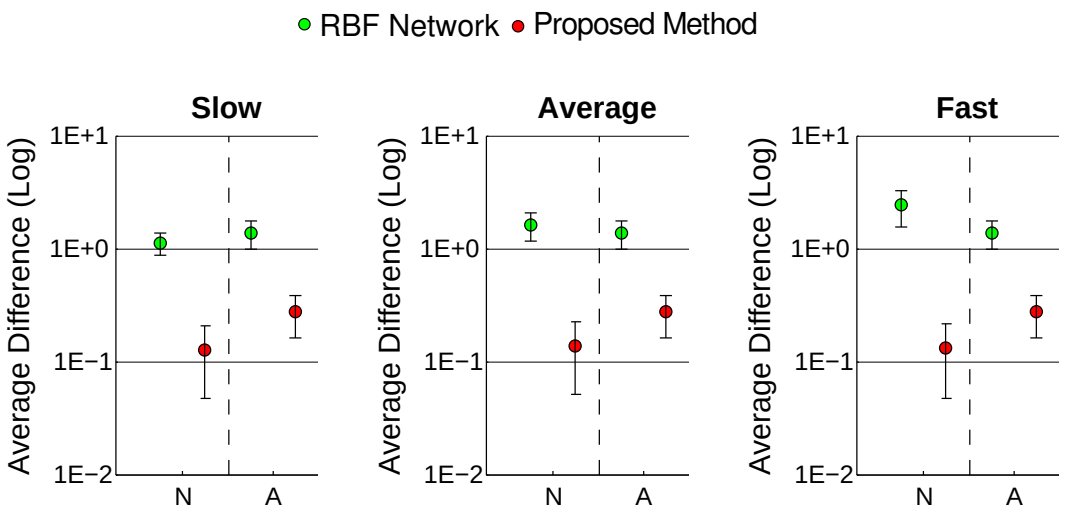

Figure 6.18: Average difference between the new subjects (S6 and S7) and the reference gait (learnt from S1-S5). The testing data was divided by three different speeds: (a) slow, (b) average, and (c) fast.

This outcome reflects the problem of using average template, where the reference gait fails to adapt to the new subjects. Our proposed method, on the other hand, can still show some difference between normal and abnormal, even if S6 and S7 are different from the subjects used to train the reference policy.

6.3.6 *Applying the Learnt Policies on Various Behaviours*

We also evaluated how well the learnt policies can generalise across various walking behaviours. In this experiment, we tested slow, average, and fast walks separately, and see how the results might be affected by various speeds.

Similar to the last experiment, we used the normal and abnormal gait from S6 and S7 to represents two normal and two abnormal persons, and compare them with the reference policy learnt from S1, S2, S3, S4, and S5. Fig. 6.18 shows the results of predicted difference between the new subjects and the reference policy, where the data was divided by three different speeds: (a) slow, (b) average, and (c) fast.

From Fig. 6.18, we can see that, by using RBF network (green), the predicted difference between the normal gait and the reference increases as the walking speed increases. This is equivalent to considering those faster walks are deviations from the normal gait.

In contrast, our proposed method (red) yields consistent results regardless of walking speeds, and this outcome confirms the fact that our method attempts to eliminate the difference coming from various walking speeds. Implication in real world application are that, our quantification method deals with different walking behaviours consistently, and the patients can choose to walk faster or slower.

6.4 EXPERIMENTS ON KINETIC FEATURES

In health care, patients falling down often result in injuries that have significant impact on their morbidity and mortality. Therefore, it will be useful if we can effectively measure the ”stability” of a subject. As mentioned in Section 2.1, maintaining the CoM displacement is a hypothetical factor that influences the dynamical stability of the person, yet how the CoM displacement relates to the stability remains unknown.

In the preceding section, we examined the kinematics data across various walking speeds to see if any consistency can be found, and we have shown some promising results in gait abnormality detection. In this experiment, we would like to carry out similar analysis with CoM displacement as an additional feature and see if we can improve the performance of the model for generalising gaits and quantifying the difference between gaits.

6.4.1 *Protocol*

In this experiment, nine subjects with slightly different body types were selected to ensure some variations in embodiments (We refer them as S1-S9). The leg-lengths and weights of subjects are summarised in Table 6.2.

To collect both kinematic and kinetic data, the subjects were asked to wear Xsens motion capture system and walk on the the instrumented treadmill of the V-gait system (see Section 6.1). A snapshot of the experiment in progress is shown in Fig. 6.19. For both system, the update frequency was set to 120 FPS.

To produce some variations in behaviours, the subjects were asked to walk with five different walking speeds, and the speed were chosen as 0.5, 0.75, 1, 1.25, 1.5 meter-per-second. For each walking speed, two minutes of walking is recorded. The walking speed was enforced by changing the speeds of the treadmill, while the cadence was uncontrolled, and the subjects were asked to walk with their comfortable step-length.

Subject	Upper Leg (cm)	Lower Leg (cm)	Weight (kg)
S1	50.5	44.1	65.1
S2	51.0	42.2	72.7
S3	52.1	43.2	68.1
S4	48.1	44.3	70.4
S5	57.0	45.4	73.5
S6	48.7	47.0	65.3
S7	55.1	46.2	84.7
S8	51.4	46.1	73.7
S9	47.7	45.0	74.6

Table 6.2: Leg lengths and weights of the subjects participated in this experiments. All subjects were male age between 20-29 (referred as S1-S9).

The inclination of the treadmill remained flat throughout the experiment to ensure no variation in environment.

After five walking speeds were collected, we followed the same procedure to collect induced pathological gait by attaching 3.5 kg weight to the subjects left leg (see described in Section 6.3.2).

For all data we collected, we used the method discussed in Section 6.2 and Section 6.2.2 to extract the gait cycles and gait phases. For each subject, we could extracted approximately 150 to 200 gait cycles, and we sampled 20,000 data points for analysis.

After recording data using the force platforms, we obtained trajectories of the centre-of-pressure (CoP) positions. We used the gravitational-line-projection method to approximate the centre-of-mass (CoM) positions in the frontal and the sagittal plane, and the CoM displacements were calculated as the distance between CoP and CoM (see Section 6.2.3 for details).

Fig. 6.20 shows examples of the CoM displacement of normal (red) and pathological (black) gait in the frontal and the sagittal plane. One trajectory from each walking speed was plotted. From direct observation, we can see a large variations in the sagittal plane (Fig. 6.20b)

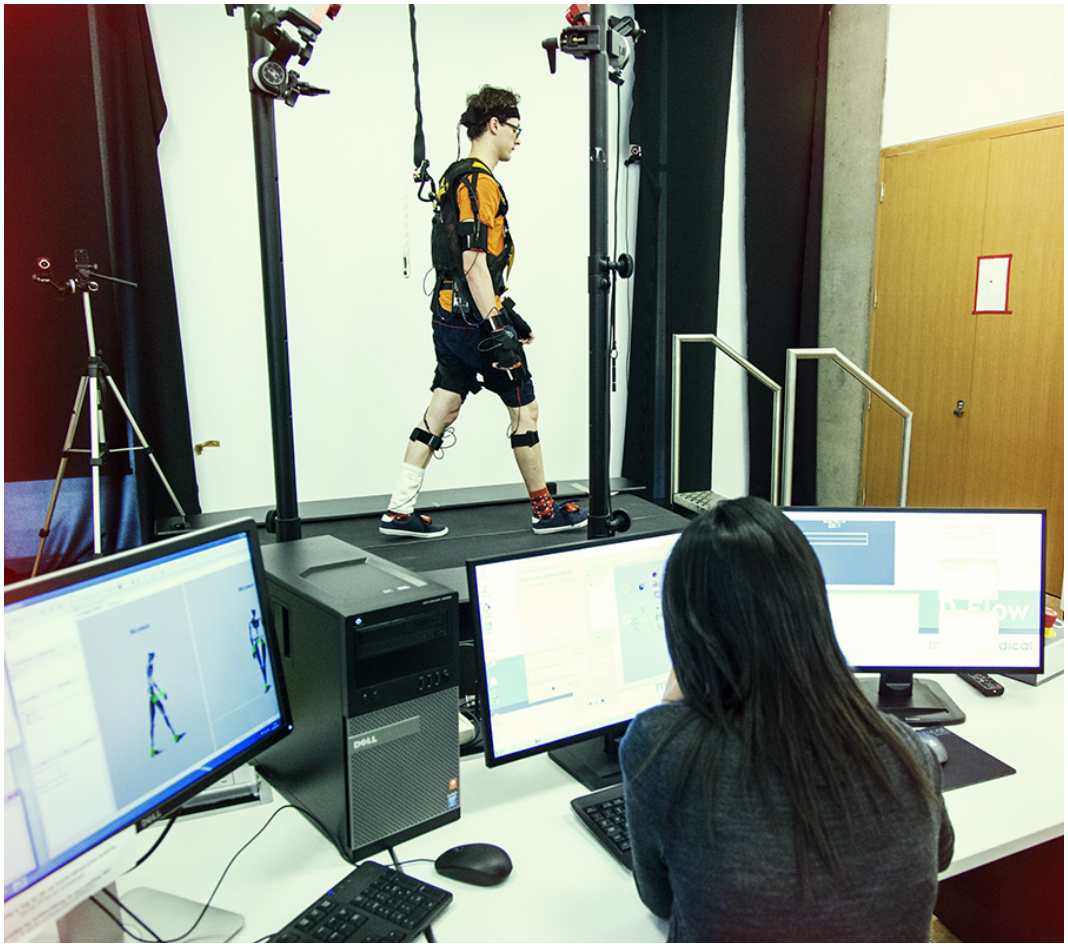

Figure 6.19: A snapshot of experiment in progress. A subject is wearing Xsens motion capture system and walking on the V-gait system.

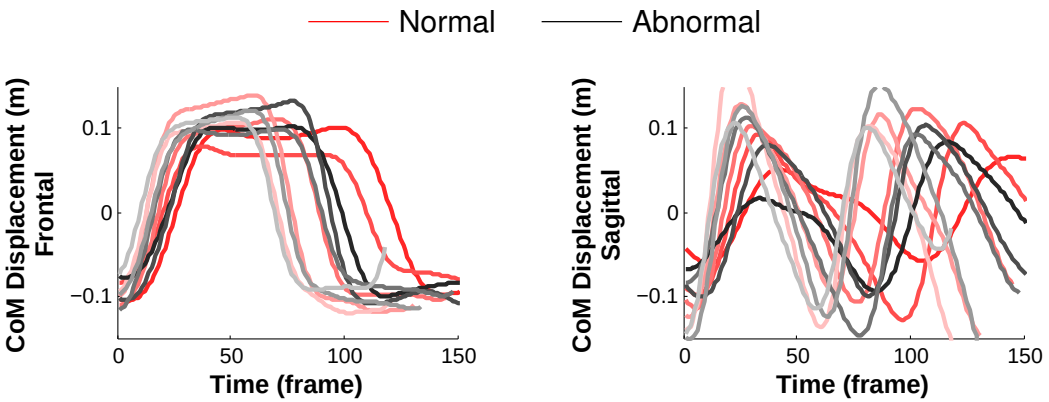

(a) (b)

Figure 6.20: CoM displacement in (a) the frontal and (b) the sagittal plane from one subject walking with different speeds. The colours denote the normal gaits (red) and abnormal gait (black)

6.4.2 *Learning and Quantifying Centre-of-Mass Displacement*

In the first set of experiment, we briefly explore how well we can differentiate normal and abnormal gait by using CoM displacement features instead of kinematic features; namely, by using the positions and linear velocities of CoM displacement as the state $\mathbf{x}$ and action $\mathbf{u}$ of the system.

$$\mathbf{x} \in \mathbb{R}^2 = \begin{bmatrix} \rho_x \\ \rho_z \end{bmatrix}, \quad \mathbf{u} \in \mathbb{R}^2 = \begin{bmatrix} \dot{\rho_x} \\ \dot{\rho_z} \end{bmatrix}$$

where the subscript x and z denote the frontal and sagittal plane, $\dot{\rho_x}$ and $\dot{\rho_z}$ are the linear velocity of CoM displacement.

For learning the reference trajectories, we selected five subjects, S1, S2, S3, S4, S5 (Algorithm 1). Specifically, for each walking phase of each subject, we learnt a model for null-space component $\tilde{\mathbf{u}}_{i,k}^{ns,ref}$, and each $\tilde{\mathbf{u}}_{i,k}^{ns,ref}$ was consisted of M Gaussian radius basis functions (RBFs) where M varied from 5 to 20. The null-space policy $\boldsymbol{\pi}^{ref}$ was modelled as 50 parametric model with Gaussian RBFs, and this learnt policy $\boldsymbol{\pi}^{ref}$ was used as the reference policy for gait abnormality detection in this section.

Algorithm 2 was applied on each gait phase of each subject separately to form an estimate of null-space projection $\mathbf{N}$, and Algorithm 3 was used to quantify the difference between a subject and the reference policy.

For comparison, we also trained models using (i) linear regression, (ii) RBF regression, and (iii) PCA regression on raw observations $(\mathbf{x}, \mathbf{u})$ from the normal gaits, and tested if we can see a difference between normal and abnormal gaits.

On average, our proposed method estimate the null-space component and null-space policy with the optimisation error $E_1 = 0.0499 \pm 0.0437$ (4.3) and $E_2 = 0.8347 \pm 1.1649$ (4.5). In terms of E_2, it turns out that, it is hard to fit a good model by examining the positions of CoM displacement alone.

Fig. 6.21 shows the average results over all subjects. The yellow, green, and blue colours denote the results using standard methods, and the red colour denotes the results using our method. The error bars are the mean±std. dev. in log scale over ten experiments.

In contrast with the assumption that CoM displacement is a good criteria for gait abnormality detection, none of the method we tested was able to discriminate normal and abnormal gait. From Fig. 6.21a, the normal gait of S1-S5 were considered to be very different from the reference policy, even if they were used to train the reference policy.

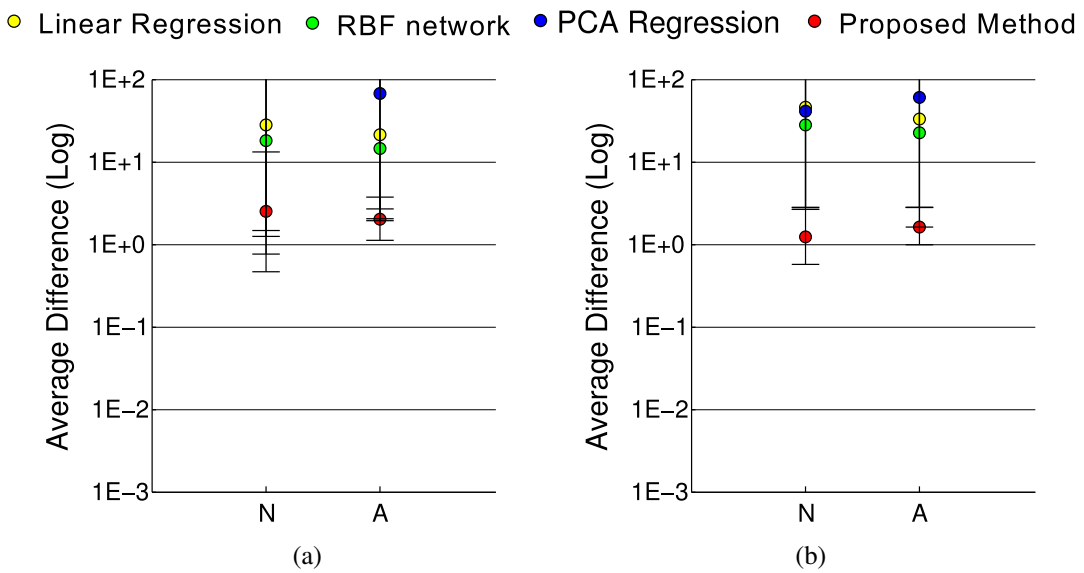

Figure 6.21: Average difference between the testing subjects and the reference policy, where the testing subjects are (a) S1-S5 and (b) S6-S9. The error-bars are mean±std.dev. in log scale over ten experiments and over subjects. The results are grouped into normal (N) and abnormal (A).

In Orendurff et al. (2004), the author discuss that the CoM displacement changes substantially with walking speed, and even healthy subjects show significant difference. Based on this, we assume that the consistency in CoM displacement is too small to recovered.

6.4.3 *Combining Kinematic and Kinetic Features*

In this experiment, the goal is to (i) investigate whether there is a consistency by examining kinematic and kinetic features together and (ii) test the scalability of our approach for higher dimensional data. For this, we state and action space were

$$\mathbf{x} \in \mathbb{R}^4 = \begin{bmatrix} \mathbf{q}_h \\ \mathbf{q}_k \\ \rho_x \\ \rho_z \end{bmatrix}, \mathbf{u} \in \mathbb{R}^4 = \begin{bmatrix} \dot{\mathbf{q}}_h \\ \dot{\mathbf{q}}_k \\ \dot{\rho_x} \\ \dot{\rho_z} \end{bmatrix}$$

We follow the same process as described in the preceding section to learn the reference gait and projection matrix. (Note that, the only difference between the experiment here and the preceding section is the choice of state and action space.)

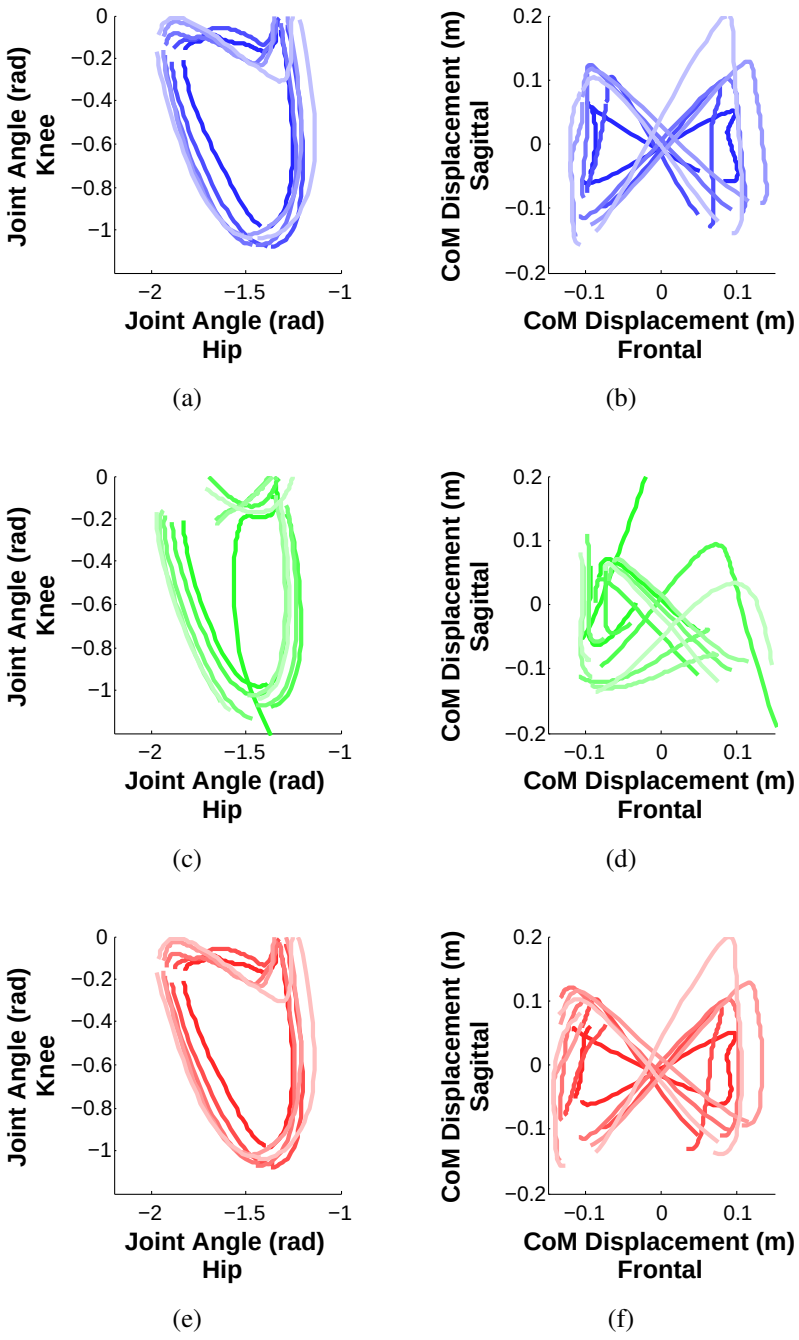

Figure 6.22: Examples of reconstructed trajectories. (a) joint-angles recorded from motion capture (b) CoM displacement data estimated from force-plate data (c) joint-angles learnt using RBF regression (d) CoM displacement learnt using RBF regression (e) joint-angles learnt using the proposed method (f) CoM displacement learnt using the proposed method

For learning the reference policy, our algorithm achieved a very low optimisation errors of $E_1 = 0.0007 \pm 0.0001$ (4.3) and $E_2 = 0.0028 \pm 0.0013$ (4.5). Note that, com-

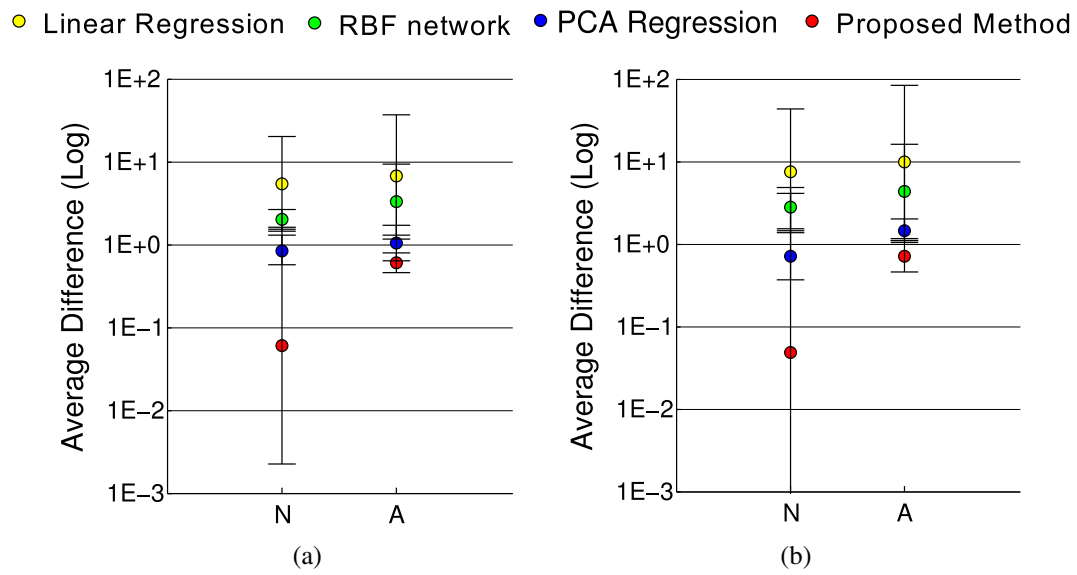

Figure 6.23: Average difference between the testing subjects and the reference policy, where the testing subjects are (a) S1-S5 and (b) S6-S9. The error-bars are mean±std.dev. in log scale over ten experiments and over subjects. The results are grouped into normal (N) and abnormal (A).

bining kinematic and kinetic features resulted in lower optimisation error as compared to the preceding experiment which used kinetic features alone.

Fig. 6.22 shows a visualisation of the true and the learnt trajectories over five gait cycles, and each gait cycle takes a different walking speed. The top row and the bottom row are the trajectories of joint angles and CoM displacement, respectively.

Fig. 6.22a and Fig. 6.22a are the true data recorded from human walking. Fig. 6.22c and Fig. 6.22d are the recovered trajectories using RBF regression, and we can see that this approach is unable to handle centre-of-mass trajectories. Our proposed method (Fig. 6.22e and Fig. 6.22f), on the other hand, reproduced the movement in both joint space and CoM displacement space

For quantifying the difference between gaits, Fig. 6.23 shows the average results over all subjects where Fig. 6.23a were the average results for S1-S5 and Fig. 6.23b were the average results of S6-S9. The yellow, green, and blue colours denote the results using standard methods, and the red colour denotes the results using our method. The error bars are the mean±std. dev. in log scale over ten experiments.

We can see that, for both cases (Fig. 6.23a and Fig. 6.23b), the linear regression (yellow), RBF regression (green), and PCA regression (blue) cannot differentiate normal and abnormal gait. By using our proposed approach (red), however, we could obtain a

much wider separation between normal and abnormal gait as comparing to the baseline method.

We consider this result is notably good, given that the dimensionality of the state/action space is higher (and normally implies a more complex problem) but the learnt model can discriminate normal and abnormal gait better than using only kinematic or kinetic features alone. Additionally, Fig. 6.23b were the average results of subjects who were not in the dataset for training the reference policy. This outcome demonstrated that our method is robust to generalise gaits across subjects.

6.5 COMPARISON TO MODEL-BASED APPROACH

In this experiment, we aim to compare our method and model-based approaches. Passive dynamic principle is a prominent approach for bipedal walking robot (Section 2.2.2. In Section 4.5, we showed that our method can effectively model simulated data generated from a passive dynamic walker. In this experiment, we used the data collected in the last experiment to compare the performance of our method and a passive dynamic walker on learning human data.

6.5.1 *Actuated Compass-Gait model*

Similar to Section 4.5, we approximate the mass distribution m_h, m_t, m_s and centre-of-mass location b_1, b_2 using the anthromopherical table Winter (2009). For each subject, we collected 100 gait cycles, and the initial condition of each cycle is also taken from the experimental data.

If we used the same compass-gait model from Section 4.5, the walking gait is stable if the initial conditions were within the base of attraction. However, even if the resulting gait is stable, it cannot reproduce the human walking data. Therefore, external forces are needed to track the trajectories and to stablise the walker. One way to increase the basin of attraction and match the input data is by adding hip actuation. For this, we need a method predict the amount of joint-torque required to track the input trajectories

Assuming that human walking can be represented as a passive dynamic walker, the desired joint-torque of the input data can be calculated using the Lagrangeâs equations of motion. $\tau^* = \mathbf{M}(\mathbf{q}, \dot{\mathbf{q}})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{G}(\mathbf{q})$ where τ^* is the desired joint torque, $\mathbf{M}$

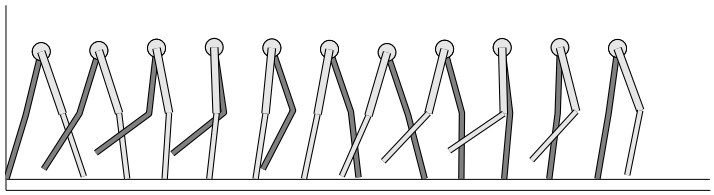

(a)

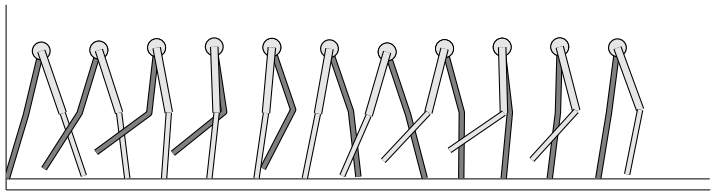

(b)

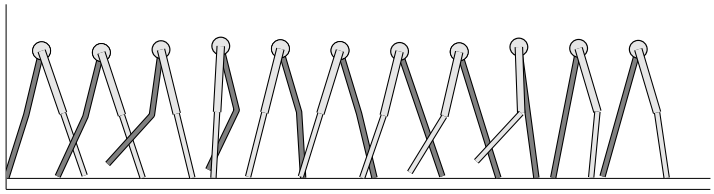

(c)

Figure 6.24: An example of walking cycle: (a) motion capture data (b) motion generated from the proposed method (c) motion generated from actuated compass-gait walker

is the inertia matrix, $\mathbf{C}$ is the centrifugal and Coriolis matrix, and $\mathbf{G}$ is the gravitational vector. Following the idea in Wisse (2004), the joint-torque can be formed as

$$\tau = -k(\mathbf{q} - \mathbf{q}^{sp}) - c\dot{\mathbf{q}}$$

where $\mathbf{q}^{sp}$ is the fixed point of the compass-gait walker, k is the stiffness factor, and c is the damping factor. We use regression method learn the parameters k, c, and $\mathbf{q}^{sp}$.

6.5.2 *Results*

For our method, we used Algorithm 1 to learn the null-space policy. Fig. 6.24 shows an example of resulting trajectories. Fig. 6.24a is the motion capture data, Fig. 6.24b is the motion generated from our method, which can closely reproduce the human walking data. Fig. 6.24c is the motion generated from the actuated compass-gait walker, where we can see a small difference between human walking and the compass-gait walker.

We repeated this experiment 10 times and evaluated the normalised root-mean-squared-error (NMSE) between true $\mathbf{u}$ and predicted $\mathbf{u}$. On average, the NMSE for compass-gait walker is 2.31 ± 1.83, while the NMSE of the proposed method is 0.0858 ± 0.0689. Our method has outperformed the compass-gait walker in both cases. We made a few assumptions on what might cause this difference.

First, for the compass-gait walker, the mass-distribution and centre-of-mass location are approximated. If the subject's body type is dissimilar to the majorities, the approximation might not reflect the right inertia, centrifugal and Coriolis, and gravitational matrix. Second, knee-strike and heel-strike are modeled as elastic collisions, and the post-collision velocities are estimated using conservation of angular momentum around the point of impact. If this assumption is wrong, these estimation would not be correct. In contrast to the model-based approach, our method, which is model-free, do not have the above issue.

Also, the compass-gait walker constraints the knee movement between knee-strike and the next heel-strike. In Fig. 6.24, we can see that the stance knee of the subject flexes between initial-contact and opposite toe-off. Although it is not visible in this example, the stance knee normally flexes slightly after mid-stance (Winter, 2009).

6.6 DISCUSSION

In this chapter, we tested our proposed method on human walking data. In particular, we analysed kinematic and kinetic features of subjects walking with various speeds, and we aim to see whether we can (i) find any consistency across walking behaviours and/or subjects and (ii) use the generalised walking to measure the difference between normal and pathological subjects.

First, we described our motion capture system for collecting kinematic data (joint-angles) and our instrumented treadmill for collecting kinetic features (centre-of-pressure and ground reaction force). We then went on to outline our procedures to automatically divide the recorded trajectories into gait cycles and gait phases. Also, we described how to transform the collected kinetic data into centre-of-mass and centre-of-mass displacement.

In our preliminary experiment, our analysis was based on kinematic features; namely, by learning a mapping between joint-angles and joint-velocities. Our initial results demonstrated that the proposed method was more effective than standard regression methods for generalising walking gait and measuring gait abnormality.

We also tried to improve our model by incorporating kinetic features (i.e., the CoM displacement in the frontal and the sagittal plane). In our first attempt, we examined the positions and velocities of CoM displacement instead of joint-space, and none of the methods we tested was able to model CoM displacement well. In the subsequent experiment, we combined the kinematic and kinetic features, and the results demonstrated that, by using these features, our method was more robust in quantifying the difference between normal and abnormal.

In our last experiment, we compared our method with a model-based approach in robotic community. In particular, we were interested in the performance of our method and a model based on passive dynamic principles. Experimental results have shown that a model-free approach is superior since it does not rely on assumptions of the bio-mechanical model of human.

We consider these results as an excellent proof of concept under challenging conditions since (i) our data was collected from different walking speeds, and our method was able to eliminate the effect of various walking speeds, (ii) our proposed method was able to generalise to new subjects. (This is particularly important in gait rehabilitation since the patients are often the subjects that we have never encountered.), and (iii) For human walking, we have no prior knowledge about the underlying characteristics (null-space policy) and the constraints of the subjects, but we were able to approximate these two quantities and use them to measure the distance between normal and abnormal gait.

7

CONCLUSION

In this thesis, we have explored the problem of representing, generalising, and comparing periodic gaits; particularly, we have focused on walking gaits that are subject to intra- and inter-personal variations. In the final chapter, we summarise the results of our analysis and point to directions for future work.

Our work was motivated by the current issue in clinical gait analysis and gait-assisted rehabilitation (see Section 2.3). A more robust algorithm for analysing movements and generating the *reference gait* is fundamental to successful and effective robot rehabilitation; on the other land, even the simplest gait (e.g., normal walking) produces various behaviours, and consequently makes robot-assisted walking a difficult problem.

Our aim was to arrive at a novel methodology for learning this reference gait as well as a method for measuring the difference between gaits. From the results in quantitative gait analysis we mentioned in Section 2.1, we hypothesised that locomotion can be described as a combination of consistent *characteristics* of the gait and *variations* from environment, embodiments, and behaviours. Our approach is to extract the underlying consistent components (as the reference gait) and separating those from variations.

In Chapter 3, we introduced a *walking phase model* for representing walking gait such that the characteristics and variations can be separated. In the light of this model, we analysed the way in which variations affect the walking gait. We then went on to describe several examples of representations and control schema for this model.

In Chapter 4, we proposed a supervised learning algorithm for generalising the consistent characteristics of walking from various embodiments and behaviours. Utilising previous work in operational space control (Section 2.4), the proposed algorithm aims at capturing the consistency in the input data by learning an *unconstrained policy*. We have validated this method with simulated walking data. We have shown that, under ideal assumptions, our algorithm is effective in reconstructing the unconstrained policy without explicit knowledge of the variations. Even if the preconditions of the method are not met, our method can still make reasonable prediction.

If Chapter 5, we described a novel approach for measuring the difference between walking gaits subject to unknown variations. One of the potential applications in gait analysis and gait rehabilitation is to quantify to degree of "pathology" of a mobility-impaired patient. For this, we proposed a framework to estimate the *variations in embodiments* in the form of *constraints* and calculate the distance under these constraints. We have validated this method with several examples that can be framed in terms of performing some task subject to variations, and we have shown that our approach is effective in reconstructing the constraints, without prior knowledge of the true policy and the dimensionality the constraints.

In Chapter 6, we explored the utility of our approach on human walking data; specifically, our analyses were based on kinematic (e.g., joint-angles) and kinetic features (e.g., centre-of-mass displacement) of subjects walking with various speeds. We aimed to see whether we can (i) uncover the consistency across embodiments and behaviours and (ii) measure the difference between normal and pathological subjects. In our preliminary experiment (based on kinematic features), the results have demonstrated that our framework has distinct benefits over the standard approach for generalising walking gait and measuring gait abnormality. In the subsequent experiment, we discovered that, by combining the kinematic and kinetic features, the novel approach showed a significant improvement in performance over standard regression techniques.

Note that, for human walking, we have no explicit knowledge about the underlying characteristics and the variations in embodiments and behaviours. For this, we consider that we have demonstrated an excellent proof of concept under challenging conditions since (i) the experimental data was collected from various speeds, and the method proposed in Section 4.2.1 was able to eliminate the effect of these variations, (ii) the proposed null-space learning algorithm (Algorithm 1) reconstructed an unconstrained policy that captures the characteristics of walking, which can generalise to new subjects, and (iii) we were able to approximate the variations in embodiments (Algorithm 2), and use them to measure the distance between normal and abnormal gaits.

7.1 FUTURE WORK

There are a few future extensions that might potentially improve the work presented in this thesis.

Alternative Representations of Walking Gaits

Experiments carried out in this thesis were based on the kinematic and kinetic features of movement; i.e., the state of the system was based on joint angles or centre-of-mass displacement. Some interesting future work is to examine other kinetic features (e.g., joint-moments, joint-works) and muscle electrical activity (e.g., EMG).

Force Control Scheme

The method we proposed in this thesis can be extended to various control space with some modification. As mentioned in Section 2.4, the constraint model can be applied to higher order control kinematic policies (i.e., mapping from positions and velocities to accelerations) or dynamic control (i.e., mappings from positions and velocities to control torques). Since many developing rehabilitation devices from the past five years have shifted the control paradigm from position control to force control (Section 2.3), The later could be an interesting future topic of work.

Alternative Optimisation Strategies for Learning Null-space Policy

In Chapter 4, we introduced a null-space policy learning algorithm to recover the underlying characteristics of walking, and the learning process was based on least-squares optimisation (i.e., by modeling the policy into radial basis function and minimise some objective functional). An interesting direction of future work could be to reformulate the current optimisation structure with other regression strategies, such as support vector regression or Bayesian network.

Learning Null-space Projection for Learning Null-space Policy

The null-space policy learning method proposed in Chapter 4 generalises an unconstrained policy by minimising the inconsistency error when the policy is projected onto the learnt null-space components (i.e., this refers to the 1-D projection of $\tilde{\mathbf{u}}^{\text{ns}}$ in (4.5)). The performance of our generalisation algorithm is highly dependent on the accuracy of this projection. In Chapter 5, we proposed a method to estimate the projection matrix $\tilde{\mathbf{N}}$ by minimising the error functional (5.18). An interesting direction of future work is to replace the 1-D projection in (4.5) by $\tilde{\mathbf{N}}$ and see if we can improve the performance of our algorithm.

Learning Null-space Projection from Demonstration for Robot Manipulator

In Section 5.4, we validated our approach of learning null-space projection on the

Kuka lightweight robot, and simulation results demonstrated that our approach is effective in reconstructing the null-space projection across constraints (i.e., applying the learnt policy to new constraints), as well as within constraints (i.e., applying new policies to learnt constraint). This is particularly useful to interact with an unknown environment. In the future, we will use the same experimental protocol for human demonstration.

Part I

APPENDIX

A

SUPPLEMENTARY MATERIALS FOR LEARNING NULL-SPACE POLICIES

In this appendix, a proof of convergence for generalising the null-space policy (Algorithm 1 in Section 4.2) is given. As discussed in the main text, the true null-space policy $\boldsymbol{\pi}$ is not known in human data, so we cannot directly evaluate the performance of our learnt model by the error between the true policy $\boldsymbol{\pi}$ and learnt policy $\tilde{\boldsymbol{\pi}}$.

In the following, we will show that, if the requirements for Algorithm 1 listed in Section 4.3 are satisfied, and the optimisation functions of Step-1 (4.3) and Step-2 (4.5) are minimised, then the null-space policy $\boldsymbol{\pi}$ can be exactly reconstructed.

We begin by showing that, if the observations contain *enough variations* in the task-space, then there is a unique decomposition of the null-space and the task-space (Appendix A.1). Then, we present the convergence analysis of learning null-space components (i.e., the *Step-1*) in Appendix A.2 and null-space policies (i.e., the *Step-2*) in Appendix A.3.

A.1 VARIATIONS IN TASK-SPACE AND THE DECOMPOSITION OF TASK- AND NULL-SPACE

In Section 4.3, we mentioned that one of the requirements for Algorithm 1 to perform well is to have *enough variations* in the task-space. In the following, we will show that, if there is *enough variations*, then there is an unique decomposition of null-space and task-space.

To begin, we formally define the notion of *enough variations* in the task-space. Let $(\mathbf{v}_1, \mathbf{v}_2, ..., \mathbf{v}_d)$ be an orthonormal basis in $\mathbb{R}^d$ such that

$$\begin{aligned} \text{task-space } \mathbb{W}^{ts} &= \text{Span}\{\mathbf{v}_1, \mathbf{v}_2, ..., \mathbf{v}_k\} \\ \text{null-space } \mathbb{W}^{ns} &= \text{Span}\{\mathbf{v}_{k+1}, \mathbf{v}_{k+2}, ..., \mathbf{v}_d\} \\ \mathbb{W}^{ts} &\perp \mathbb{W}^{ns} \end{aligned}$$

for $1 \leqslant k < d$. Let $\mathbf{u}_i^{ts}$ be a vector in the task-space; i.e., $\mathbf{u}_i^{ts}$ can be expressed as the sum of the scaled orthonormal basis in $\mathbb{W}^{ts}$

$$\mathbf{u}_i^{ts} \in \mathbb{W}^{ts} = s_1\mathbf{v}_1 + s_2\mathbf{v}_2 + ... + s_k\mathbf{v}_k$$

Let $\mathbf{u}_j^{ts} \in \mathbb{W}^{ts}$ be another vector in the task-space, and $\triangle\ \hat{\mathbf{u}}_{ij}$ denotes the normalised difference between $\mathbf{u}_i^{ts}$ and $\mathbf{u}_j^{ts}$; i.e.,

$$\triangle\ \hat{\mathbf{u}}_{ij} \in \mathbb{W}^{ts} = \frac{\mathbf{u}_i^{ts} - \mathbf{u}_j^{ts}}{\| \mathbf{u}_i^{ts} - \mathbf{u}_j^{ts} \|}$$

Note that $\triangle\ \hat{\mathbf{u}}_{ij}$ is also in the task-space, and $\triangle\ \hat{\mathbf{u}}_{ij}$ can be expressed as the sum of the scaled orthonormal basis.

Definition A set of observations is considered having *enough variations* if (i) there exists at least one $\triangle\ \hat{\mathbf{u}}_{ij}$ such that $s_m \neq 0$ for $1 \leqslant m \leqslant k$, and (ii) for each pair of basis vectors $\mathbf{v}_m, \mathbf{v}_n$ for $1 \leqslant m, n \leqslant k, m \neq n$, there exists at least two $\triangle\ \hat{\mathbf{u}}_{ij}$ such that their s_m or s_n are different.

Note that, if $s_m = 0$ for all $\triangle\ \hat{\mathbf{u}}_{ij}$, then there is no variation along the direction of $\mathbf{v}_m$, and we cannot determine whether $\mathbf{v}_m \in \mathbb{W}^{ts}$ or $\mathbf{v}_m \in \mathbb{W}^{ns}$.

If all $\triangle\ \hat{\mathbf{u}}_{ij}$ are identical, the difference between all pairs of $\mathbf{u}_i^{ts}$ and $\mathbf{u}_j^{ts}$ are either parallel or anti-parallel. In this case, there exists a vector $\triangle\ \mathbf{u}_\perp^{ts} \in \mathbb{W}^{ts}$ such that $\triangle\ \mathbf{u}_\perp^{ts}$ is orthogonal to all $\triangle\ \hat{\mathbf{u}}_{ij}$. Note that, when $\mathbf{u}_i^{ts}$ are projected onto this $\triangle\ \mathbf{u}_\perp^{ts}$, the inconsistency is minimised. (i.e., $\tilde{\mathbf{u}}^{ns} \equiv \triangle\ \mathbf{u}_\perp^{ts}$ is a solution such that the error in (4.3) is minimised) This $\triangle\ \mathbf{u}_\perp^{ts}$ and the set of vectors orthogonal to $\triangle\ \mathbf{u}_\perp^{ts}$ can yield another decomposition of task-space and null-space in $\mathbb{R}^d$; however, it is different from the true decomposition. Since we have no prior knowledge of the decomposition of task- and null-space, having more than one possible solution would be problematic.

In Fig. A.1 and Fig. A.2, we demonstrate two examples in $\mathbb{R}^3$. In both figures, the xy-plane forms the task-space, and the z-dimension is the null-space; i.e., $\mathbf{v}_x,\ \mathbf{v}_y \in \mathbb{W}^{ts}$ and $\mathbf{v}_z \in \mathbb{W}^{ns}$.

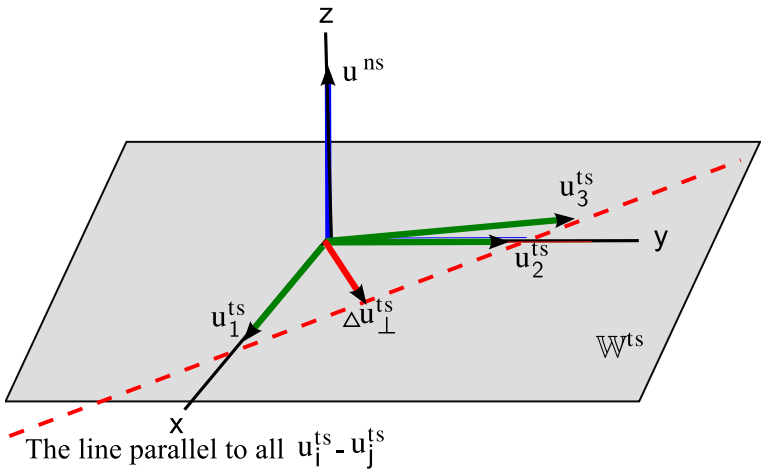

Figure A.1: An example of insufficient variations in task-space. The difference between all pairs of $\mathbf{u}_i^{ts}$ and $\mathbf{u}_j^{ts}$ are parallel to the red dash-line. (i.e., $\triangle\ \hat{\mathbf{u}}_{12}$ and $\triangle\ \hat{\mathbf{u}}_{23}$ lie on the same line). In this case, there exists a vector $\triangle\ \mathbf{u}_\perp^{ts}$ (red) such that $\triangle\ \mathbf{u}_\perp^{ts}$ is orthogonal to this red dash-line. Note that, $\triangle\ \mathbf{u}_\perp^{ts}$ and the set of vectors orthogonal to $\triangle\ \mathbf{u}_\perp^{ts}$ can form another decomposition of task-space and null-space that can minimise the objective function of Step-1 (4.3), but the solution is different from the true $\mathbf{u}^{ns}$

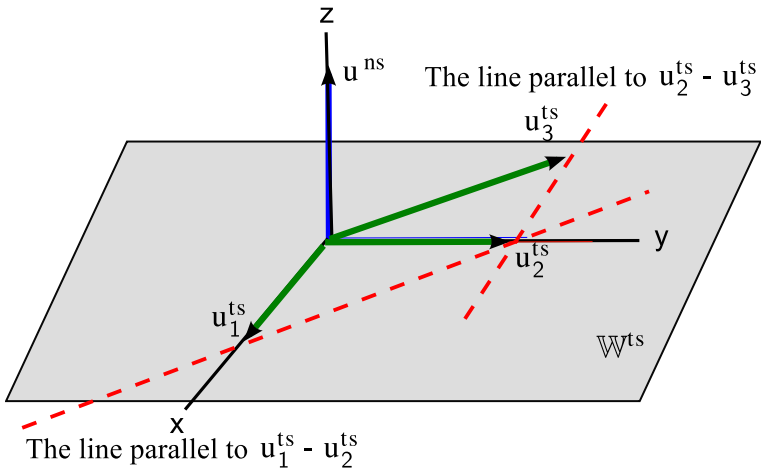

Figure A.2: An example of enough variations in task-space. If $\triangle\ \hat{\mathbf{u}}_{12}$ is different from $\triangle\ \hat{\mathbf{u}}_{23}$, there is no solution in xy-plane that is orthogonal to both $\triangle\ \hat{\mathbf{u}}_{12}$ and $\triangle\ \hat{\mathbf{u}}_{23}$, so there exists a unique solution for $\mathbf{u}^{ns}$.

Fig. A.1 illustrates an example of insufficient variations in task-space. In this figure, there are three different task-space component, $\mathbf{u}_1^{ts}$, $\mathbf{u}_2^{ts}$, $\mathbf{u}_3^{ts}$ (green). The difference between all pairs of $\mathbf{u}_i^{ts}$ and $\mathbf{u}_j^{ts}$ are parallel or anti-parallel to the red dash-line. (i.e., $\triangle\ \hat{\mathbf{u}}_{12}$ and $\triangle\ \hat{\mathbf{u}}_{23}$ lie on the same line). In this case, there exists a vector $\triangle\ \mathbf{u}_\perp^{ts}$ (red) such that $\triangle\ \mathbf{u}_\perp^{ts}$ is orthogonal to this red dash-line. Note that, $\triangle\ \mathbf{u}_\perp^{ts}$ and the set of vectors

orthogonal to $\triangle\ \mathbf{u}^{ts}_{\perp}$ can form another decomposition of task-space and null-space that can minimise the objective function of Step-1 (4.3), but the solution is different from the true $\mathbf{u}^{ns}$, but the solution is different from the true $\mathbf{u}^{ns}$ (blue).

In Fig. A.2, we show an example of *enough variations*. In this figure, $\triangle\ \hat{\mathbf{u}}_{12}$ and $\triangle\ \hat{\mathbf{u}}_{23}$ are not parallel, so we have enough information to recover the true $\mathbf{u}^{ns}$ (blue).

A.2 CONVERGENCE ANALYSIS OF LEARNING NULL-SPACE COMPONENTS

Given a set of observations $(\mathbf{x}, \mathbf{u})$, we assume that the observed action $\mathbf{u}$ can be decomposed into two orthogonal components $\mathbf{u}^{ts} \perp \mathbf{u}^{ns}$ such that $\mathbf{u} = \mathbf{u}^{ts} + \mathbf{u}^{ns}$ for some $\mathbf{u} \neq 0$ and $\mathbf{u}^{ns} \neq 0$. The goal of Step-1 (Section 4.2.1) is to learn a model that approximates the null-space component $\mathbf{u}^{ns}$ by minimising the error function (4.3)

$$E_1 = \sum_n \| \mathbf{P}\,\mathbf{u} - \tilde{\mathbf{u}}^{ns} \|^2 \text{ where } \mathbf{P} = \frac{\tilde{\mathbf{u}}^{ns}\,\tilde{\mathbf{u}}^{ns\top}}{\tilde{\mathbf{u}}^{ns\top}\,\tilde{\mathbf{u}}^{ns}}$$

In Section 4.3, we mentioned that, for Step-1 to perform well, one requirement is to have consistent null-space components and enough variations in the task-space components. In the following, the convergence analysis is presented.

Proposition A.2.1. *If the observations contain consistent null-space components* $\mathbf{u}^{ns}$ *and enough variations in the task-space components* $\mathbf{u}^{ts}$*, and the objective function* E_1 *(4.3) is minimised, then the learnt null-space component* $\tilde{\mathbf{u}}^{ns}$ *is a good approximation to the true null-space component* $\mathbf{u}^{ns}$*.*

Proof. Consider that we have k distinct observations $\mathbf{u}$ at a point $\mathbf{x}$ in the state space, there are k distinct $\mathbf{u}^{ts}$ and a consistent $\mathbf{u}^{ns} \neq 0$. If $E_1 = \sum_n^{\mathcal{N}} \| \mathbf{P}_n\,\mathbf{u}_n - \tilde{\mathbf{u}}^{ns}_n \|^2 = 0$, then

$$\begin{aligned} \mathbf{P}\,\mathbf{u}_1 &= \tilde{\mathbf{u}}^{ns} \\ \mathbf{P}\,\mathbf{u}_2 &= \tilde{\mathbf{u}}^{ns} \\ &\vdots \\ \mathbf{P}\,\mathbf{u}_k &= \tilde{\mathbf{u}}^{ns} \end{aligned} \tag{A.1}$$

By definition, $\mathbf{u}_i = \mathbf{u}^{ts}_i + \mathbf{u}^{ns}$ for $1 \leqslant i \leqslant k$, (A.1) can be written as

$$\begin{aligned} \mathbf{P}\,\mathbf{u}^{ts}_1 + \mathbf{P}\,\mathbf{u}^{ns} &= \tilde{\mathbf{u}}^{ns} \\ \mathbf{P}\,\mathbf{u}^{ts}_2 + \mathbf{P}\,\mathbf{u}^{ns} &= \tilde{\mathbf{u}}^{ns} \\ &\vdots \\ \mathbf{P}\,\mathbf{u}^{ts}_k + \mathbf{P}\,\mathbf{u}^{ns} &= \tilde{\mathbf{u}}^{ns} \end{aligned} \tag{A.2}$$

Since $\mathbf{P}\,\mathbf{u}^{ns}$ and $\tilde{\mathbf{u}}^{ns}$ are consistent for all $\mathbf{u}$, we must have

$$\mathbf{P}\,\mathbf{u}_1^{ts} = \mathbf{P}\,\mathbf{u}_2^{ts} = ... = \mathbf{P}\,\mathbf{u}_k^{ts}$$

For each pair of $\mathbf{u}_i^{ts}$ and $\mathbf{u}_j^{ts}$ such that $1 \leqslant i, j \leqslant k$ and $i \neq j$,

$$\mathbf{P}\,\mathbf{u}_i^{ts} - \mathbf{P}\,\mathbf{u}_j^{ts} = \mathbf{P}(\mathbf{u}_i^{ts} - \mathbf{u}_j^{ts}) = 0$$

Since all $\mathbf{u}_i^{ts}$ are distinct, we must have $\mathbf{u}_i^{ts} - \mathbf{u}_j^{ts} \neq 0$. If $\mathbf{P}(\mathbf{u}_i^{ts} - \mathbf{u}_j^{ts}) = 0$, then $\tilde{\mathbf{u}}^{ns}$ must be orthogonal to all $\mathbf{u}_i^{ts} - \mathbf{u}_j^{ts}$. If the requirement for Step-1 are met, the set of $\triangle\,\hat{\mathbf{u}}_{ij}$ spans the task-space $\mathbb{W}^{ts}$. This implies that $\tilde{\mathbf{u}}^{ns}$ is orthogonal to the task-space and $\mathbf{P}\,\mathbf{u}_i^{ts} = 0$. For each observation, $\mathbf{P}\,\mathbf{u}_i^{ts} + \mathbf{P}\,\mathbf{u}^{ns} = 0 + \mathbf{P}\,\mathbf{u}^{ns} = \tilde{\mathbf{u}}^{ns}$. Therefore, if $E_1 = 0$, we must have $\mathbf{P}\,\mathbf{u}^{ns} = \tilde{\mathbf{u}}^{ns}$ and $\mathbf{P}\,\mathbf{u}^{ts} = 0$.

By definition, $\mathbf{u}^{ts} \perp \mathbf{u}^{ns}$. We have shown that $\tilde{\mathbf{u}}^{ns}$ is orthogonal to $\mathbb{W}^{ts}$ and $\tilde{\mathbf{u}}^{ns} \perp \mathbf{u}^{ts}$. If $\mathbf{u}^{ts}$ is also orthogonal to $\mathbf{u}^{ns}$, the relationship between $\tilde{\mathbf{u}}^{ns}$ and $\mathbf{u}^{ns}$ is either parallel or orthogonal. Since we already show that $\mathbf{P}\,\mathbf{u}^{ns} = \tilde{\mathbf{u}}^{ns}$ and we only consider that cases where $\tilde{\mathbf{u}}^{ns} \neq 0$, we know that $\tilde{\mathbf{u}}^{ns}$ and $\mathbf{u}^{ns}$ are not orthogonal. If $\mathbf{u}^{ns}$ is parallel to $\tilde{\mathbf{u}}^{ns}$ and $\mathbf{P}\,\mathbf{u}^{ns} = \tilde{\mathbf{u}}^{ns}$, we must have $\tilde{\mathbf{u}}^{ns} = \mathbf{u}^{ns}$. □

A.3 CONVERGENCE ANALYSIS OF LEARNING NULL-SPACE POLICIES

Given a set of observed state and null-space components $(\mathbf{x}, \mathbf{u}^{ns})$ where $\mathbf{u}^{ns} \neq 0$. We consider $\mathbf{u}^{ns} \equiv \mathbf{N}\,\boldsymbol{\pi}$ where $\mathbf{N}$ is a projection matrix, and $\boldsymbol{\pi}$ is the null-space policy. The goal of Step-2 is to learn a model that approximates $\boldsymbol{\pi}$ by minimising the error function (4.5)

$$E_2 = \sum_n || \mathbf{N}_n\,\tilde{\boldsymbol{\pi}}(\mathbf{x}_n) - \mathbf{u}_n^{ns} ||^2 \text{ where } \mathbf{N}_n = \frac{\mathbf{u}_n^{ns}\,\mathbf{u}_n^{ns\top}}{\mathbf{u}_n^{ns\top}\,\mathbf{u}_n^{ns}}$$

As discussed in Section 4.3, one requirement for Step-2 is that $\mathbf{u}^{ns}$ spans the action space. (This implies that we need to have various $\mathbf{N}$) The convergence analysis is provided as the following.

Proposition A.3.1. *If $\mathbf{u}^{ns}$ spans the action space, and the objective function E_2 (4.5) is minimised, then the learnt null-space policy $\tilde{\boldsymbol{\pi}}$ is a good approximation to the true null-space policy $\boldsymbol{\pi}$.*

Proof. Given observations d distinct $\mathbf{u}_i^{ns} = \mathbf{N}_i \boldsymbol{\pi}$ for $1 \leqslant i \leqslant d$ at a point $\mathbf{x}$ in the state space such that the set of $\mathbf{u}_i^{ns}$ spans $\mathbb{R}^d$, by definition,

$$\begin{aligned} \mathbf{N}_1 \boldsymbol{\pi} &= \mathbf{u}_1^{ns} \\ \mathbf{N}_2 \boldsymbol{\pi} &= \mathbf{u}_2^{ns} \\ &\vdots \\ \mathbf{N}_d \boldsymbol{\pi} &= \mathbf{u}_d^{ns} \end{aligned} \tag{A.3}$$

If $E_2 = 0$, we have learnt a policy $\tilde{\boldsymbol{\pi}}$ such that

$$\begin{aligned} \mathbf{N}_1 \tilde{\boldsymbol{\pi}} &= \mathbf{u}_1^{ns} \\ \mathbf{N}_2 \tilde{\boldsymbol{\pi}} &= \mathbf{u}_2^{ns} \\ &\vdots \\ \mathbf{N}_d \tilde{\boldsymbol{\pi}} &= \mathbf{u}_d^{ns} \end{aligned} \tag{A.4}$$

From (A.3) and (A.3), we can observe the difference between the learnt policy $\tilde{\boldsymbol{\pi}}$ and the true policy $\boldsymbol{\pi}$ as

$$\begin{aligned} \mathbf{N}_1(\tilde{\boldsymbol{\pi}} - \boldsymbol{\pi}) &= 0 \\ \mathbf{N}_2(\tilde{\boldsymbol{\pi}} - \boldsymbol{\pi}) &= 0 \\ &\vdots \\ \mathbf{N}_d(\tilde{\boldsymbol{\pi}} - \boldsymbol{\pi}) &= 0 \end{aligned}$$

Since $\mathbf{u}_i^{ns} \neq 0$ (which implies $\mathbf{N}_i \neq 0$) for $1 \leqslant i \leqslant d$, we must have either $(\tilde{\boldsymbol{\pi}} - \boldsymbol{\pi}) = 0$ or $(\tilde{\boldsymbol{\pi}} - \boldsymbol{\pi})$ is orthogonal to all $\mathbf{u}_i^{ns}$. If the set of $\mathbf{u}_i^{ns}$ spans $\mathbb{R}^d$, there exists no $(\tilde{\boldsymbol{\pi}} - \boldsymbol{\pi})$ in $\mathbb{R}^d$ that is orthogonal to all $\mathbf{u}_i^{ns}$. The only possiblility is that $(\tilde{\boldsymbol{\pi}} - \boldsymbol{\pi}) = 0$. Therefore, $\tilde{\boldsymbol{\pi}} = \boldsymbol{\pi}$ □

B

SUPPLEMENTARY MATERIALS FOR LEARNING NULL-SPACE PROJECTIONS

In Section 5.2, we proposed a method to estimate the null-space projection matrix with multi-dimensional constraints; i.e., the constraint matrix $\mathbf{A}$ is formed from a set of $S > 1$ linearly independent vectors

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \vdots \\ \mathbf{a}_S \end{bmatrix}$$

To find the optimal $\hat{\mathbf{a}}_s^*$ for $s > 1$, the search only need to be performed for $\hat{\mathbf{a}}_s \perp \hat{\mathbf{a}}_j^*$ for all $s > j$. In this appendix, we proposed a method to define the set of $\hat{\mathbf{a}}_s$ that is orthogonal to all $\hat{\mathbf{a}}_j^*$ to reduce the search process.

Our approach is to begin with a set of unit vectors $\hat{\mathbf{a}}_1^0, \hat{\mathbf{a}}_2^0, \ldots, \hat{\mathbf{a}}_{s-1}^0$ that is known to be mutually orthogonal and a set of $\hat{\mathbf{a}}_s^c$ that is orthogonal to $\hat{\mathbf{a}}_j^0$ for all $s > j$ such that $\hat{\mathbf{a}}_j^* = \hat{\mathbf{a}}_j^0 \mathbf{R}_j$ and $\hat{\mathbf{a}}_s = \hat{\mathbf{a}}_s^c \mathbf{R}_s$ for some tranformation matrix $\mathbf{R}_j$ and $\mathbf{R}_s$. We will discuss our choice of $\hat{\mathbf{a}}_j^0$, $\hat{\mathbf{a}}_s^c$, $\mathbf{R}_j$, and $\mathbf{R}_s$ in the next two sections.

B.1 ORTHONORMAL BASIS

One simple and well defined orthonormal basis is the *standard basis*, which is the set of unit vectors pointing in the direction of the axes. Let $\hat{\mathbf{a}}_1^0, \hat{\mathbf{a}}_2^0, \ldots, \hat{\mathbf{a}}_S^0$ be the set of vectors that forms the standard basis in $\mathbb{R}^S$ where $\hat{\mathbf{a}}_j^0 \in \mathbb{R}^S$ and the subscript j denotes the axis; i.e., in $\mathbb{R}^3$, $\hat{\mathbf{a}}_1^0 = [1,0,0]$, $\hat{\mathbf{a}}_2^0 = [0,1,0]$, and $\hat{\mathbf{a}}_3^0 = [0,0,1]$ (Fig.B.1a).

With the representation in (5.14), each $\hat{\mathbf{a}}_j^0$ is characterised by $(S-1)$ parameters $\theta_{j,1}, \theta_{j,2}, \ldots, \theta_{j,S-1}$. Note that, $\hat{\mathbf{a}}_1^0$ is represented by parameters $\theta_{1,k} = 0$ for $k =$

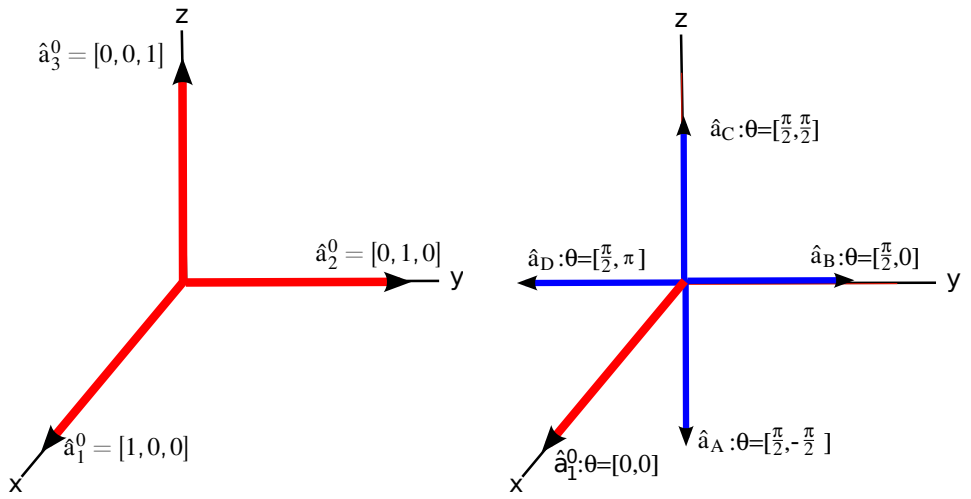

Figure B.1: Examples of $\hat{\mathbf{a}}^0$ and $\hat{\mathbf{a}}^c$ in $\mathbb{R}^3$: (a) The standard basis is formed by $\hat{\mathbf{a}}_1^0 = [1,0,0]$, $\hat{\mathbf{a}}_2^0 = [0,1,0]$, and $\hat{\mathbf{a}}_3^0 = [0,0,1]$. (b) $\hat{\mathbf{a}}_1^0$ is characterised by $\theta_{1,1} = \theta_{1,2} = 0$, and a vector with its first parameter equal to $\frac{\pi}{2}$ (e.g., $\hat{\mathbf{a}}_A^c, \hat{\mathbf{a}}_B^c, \hat{\mathbf{a}}_C^c, \hat{\mathbf{a}}_D^c$) lies on the yz plane and is orthogonal to $\hat{\mathbf{a}}_1^0$.

$1, 2, ..., \mathcal{S}-1$, and any vectors with its first parameter equal to $\frac{\pi}{2}$ is orthogonal to $\hat{\mathbf{a}}_1^0$. An example in $\mathbb{R}^3$ is illustrated in Fig. B.1b. $\hat{\mathbf{a}}_1^0 = [1,0,0]$ is the first vector in the standard basis and $\theta_{1,1} = \theta_{1,2} = 0$. $\hat{\mathbf{a}}_A$, $\hat{\mathbf{a}}_B$, $\hat{\mathbf{a}}_C$, $\hat{\mathbf{a}}_D$ lie on the yz-plane and are orthogonal to $\hat{\mathbf{a}}_1^0$.

Therefore, if $\hat{\mathbf{a}}_2^c$ be the set of vectors orthogonal to $\hat{\mathbf{a}}_1^0$, characterised by parameters $\boldsymbol{\theta}_2$, then one way to define $\hat{\mathbf{a}}_2^c$ is by setting the first parameters to $\frac{\pi}{2}$ and $\theta_{2,k} \sim U(0,\pi)$ for $1 < k < \mathcal{S}-1$.

More generally, let $\hat{\mathbf{a}}_{s+1}^c$ be the set of vectors orthogonal to the first s basis $\hat{\mathbf{a}}_1^0$, $\hat{\mathbf{a}}_2^0$, $\ldots$, $\hat{\mathbf{a}}_s^0$ for $s < \mathcal{S}$. $\hat{\mathbf{a}}_{s+1}^c$ can be represented by parameters $\boldsymbol{\theta}_{s+1} = \theta_{s+1,1}, \theta_{s+1,2}, \cdots, \theta_{s+1,\mathcal{S}-1}$ where $\theta_1 = \theta_2 = \cdots = \theta_s = \frac{\pi}{2}$ and $\theta_{s+1,k} \sim U(0,\pi)$ for $s < k < \mathcal{S}-1$.

B.2 ROTATION IN N-DIMENSIONAL SPACE

We can consider that an arbitrary unit vector $\hat{\mathbf{a}}_1 \in \mathbb{R}^{\mathcal{S}}$ is the result of performing rotations from $\hat{\mathbf{a}}_1^0 \in \mathbb{R}^{\mathcal{S}}$. For instance, in $\mathbb{R}^2$, $\hat{\mathbf{a}}$ is obtained by rotating $\hat{\mathbf{a}}_1^0$ counter-clockwise in the xy-plane by θ_1 radians.

An example of $\hat{\mathbf{a}}_1^0$ and $\hat{\mathbf{a}}_1^c$ is illustrated in Fig. B.2. In Fig. B.2a, $\hat{\mathbf{a}}_1^0 = [1,0,0]$ is the first vector in the standard basis (red) and $\hat{\mathbf{a}}_1^c$ are the vectors orthogonal to $\hat{\mathbf{a}}_1^0$ (blue). An

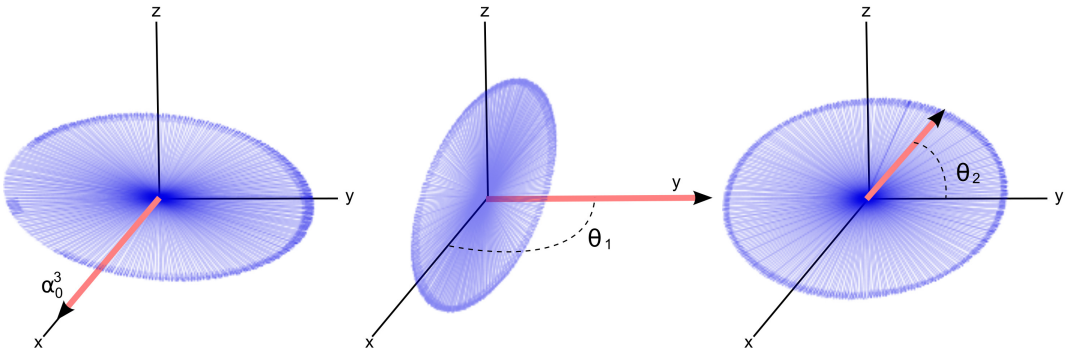

(a) (b) (c)

Figure B.2: Examples of rotation in $\mathbb{R}^3$: (a) $\hat{\mathbf{a}}_1^0 = [1,0,0]$ (red) is the first vector in the standard basis and $\hat{\mathbf{a}}_2^c$ (blue) are vectors orthogonal to $\hat{\mathbf{a}}_1^0$ (b) rotate counter-clockwise in the xy-plane of θ_1 and (c) rotate counter-clockwise in the yz-plane of θ_2

unit vector $\hat{\mathbf{a}}$ with θ_1, θ_2 can be calculated by (1) rotating $\hat{\mathbf{a}}_1^0$ counter-clockwise in the xy-plane by θ_1 radians (Fig. B.2b) and (2) rotating the result of (1) counter-clockwise in the yz-plane by θ_2 radians (Fig. B.2c).

A plane rotation in an arbitrary plane and dimension can be represented by a Givens Rotation (Givens, 1958):

$$\mathbf{G}(i,j,\theta) \in \mathbb{R}^{\mathcal{S}\times\mathcal{S}} = \begin{bmatrix} 1 & \dots & 0 & \dots & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & & \vdots & & \vdots \\ 0 & \dots & \cos(\theta) & \dots & -\sin(\theta) & \dots & 0 \\ \vdots & & \vdots & \ddots & \vdots & & \vdots \\ 0 & \dots & \sin(\theta) & \dots & \cos(\theta) & \dots & 0 \\ \vdots & & \vdots & & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & \dots & 0 & \dots & 1 \end{bmatrix} \tag{B.1}$$

where i,j specifies the plane of rotation, and $1 \leqslant i < j \leqslant \mathcal{S}$, θ is the degree of counter-clockwise rotation, and the non-zero elements are defined by

$$\begin{aligned} g_{k,k} &= 1 \text{ for } 1 \leqslant k \leqslant \mathcal{S}, k \neq i, k \neq j \\ g_{i,i} &= \cos(\theta) \\ g_{j,j} &= \cos(\theta) \\ g_{i,j} &= -\sin(\theta) \\ g_{j,i} &= \sin(\theta) \end{aligned}$$

As an example in $\mathbb{R}^3$, the counter-clockwise rotation in xy-plane of θ_1 radians and in yz-plane of θ_2 radians can be described as

$$\mathbf{G}(x, y, \theta_1) = \begin{bmatrix} \cos(\theta_1) & -\sin(\theta_1) & 0 \\ \sin(\theta_1) & \cos(\theta_1) & 0 \\ 0 & 0 & 1 \end{bmatrix} \text{ and } \mathbf{G}(y, z, \theta_2) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta_2) & -\sin(\theta_2) \\ 0 & \sin(\theta_2) & \cos(\theta_2) \end{bmatrix}$$

The transformation matrix $\mathbf{R}_1(\boldsymbol{\theta}) = \mathbf{G}(y, z, \theta_2)\mathbf{G}(x, y, \theta_1)$ is the product of these two plane rotations. For higher dimensional space $\mathbb{R}^S$, we term the axis as $d_1, d_2, \ldots, d_S$, the rotation matrix which transform from $\hat{\mathbf{a}}_i^0$ is

$$\mathbf{R}_i(\boldsymbol{\theta}) = \mathbf{G}(d_{S-1}, d_S, \theta_{S-1})\mathbf{G}(d_{S-2}, d_{S-1}, \theta_{S-2}) \cdots \mathbf{G}(d_i, d_{i+1}, \theta_i) \tag{B.2}$$

B.3 TRANSFORMATION FROM STANDARD BASIS

As described in the preceeding sections, a more efficient way to find the $(s+1)^{\text{th}}$ optimal $\hat{\mathbf{a}}_{s+1}^*$ is to iterate through the set of candidates $\hat{\mathbf{a}}_{s+1}$ subject to $\hat{\mathbf{a}}_{s+1} \perp \hat{\mathbf{a}}_s^*$.

After finding the optimal $\hat{\mathbf{a}}_1^*$ and $\boldsymbol{\theta}_1^*$ under (5.6), we begin with the set of vectors $\hat{\mathbf{a}}_2^c$, represented by parameters $\theta_{2,1} = \frac{\pi}{2}$ and $\theta_{2,k} \sim U(0, \pi)$ for $1 < k < S-1$, so that $\hat{\mathbf{a}}_2^c \perp \hat{\mathbf{a}}_1^0$. Note that, if $\hat{\mathbf{a}}_1^* = \hat{\mathbf{a}}_1^0$, an efficient way to find the second optimal $\hat{\mathbf{a}}_2^*$ is by iterating over the set of $\hat{\mathbf{a}}_2^c$ and performing a line search for the choice of $\theta_{2,k}$ for $1 < k < S-1$.

However, it is unlikely that $\hat{\mathbf{a}}_1^* = \hat{\mathbf{a}}_1^0$. Nevertheless, we can consider that $\hat{\mathbf{a}}_1^*$ and $\hat{\mathbf{a}}_2$ are the results of some linear transformation such that $\hat{\mathbf{a}}_1^* = \left(\mathbf{R}_1(\boldsymbol{\theta}_1^*)(\hat{\mathbf{a}}_1^0)^\top\right)^\top = \hat{\mathbf{a}}_1^0\mathbf{R}_1(\boldsymbol{\theta}_1^*)^\top$ and $\hat{\mathbf{a}}_2 = \hat{\mathbf{a}}_2^c\mathbf{R}_1(\boldsymbol{\theta}_1^*)^\top$ where $\mathbf{R}_1(\boldsymbol{\theta}_1^*) \in \mathbb{R}^{S\times S}$ is the rotation matrix consisting of a sequence of $S-1$ plane rotations on $\hat{\mathbf{a}}_1^0$ (B.2).

Note that, orthogonal transformation preserves lengths of the vectors and angles between vectors. Namely, if $\hat{\mathbf{a}}_2^c \perp \hat{\mathbf{a}}_1^0$, we must have $\hat{\mathbf{a}}_2^c\mathbf{R}_1(\boldsymbol{\theta}_1^*)^\top \perp \hat{\mathbf{a}}_1^0\mathbf{R}_1(\boldsymbol{\theta}_1^*)^\top$ and $\hat{\mathbf{a}}_2 \perp \hat{\mathbf{a}}_1^*$. Therefore, we can test the set of vectors $\hat{\mathbf{a}}_2$ to find the optimal $\hat{\mathbf{a}}_2^*$ that minimises (5.18).

The same approach can be extended to find $\hat{\mathbf{a}}_{s+1}^*$. Namely, let $\hat{\mathbf{a}}_{s+1}^c$ be the set of vectors orthogonal to the first s vectors $\hat{\mathbf{a}}_1^0, \cdots, \hat{\mathbf{a}}_s^0$ in the standard basis. $\hat{\mathbf{a}}_{s+1}^c$ can be represented by parameters $\theta_{s+1,1} = \theta_{s+1,2} = \cdots = \theta_{s+1,s} = \frac{\pi}{2}$ and $\theta_{s+1,k} \sim U(0, \pi)$ for $s < k < S-1$. The set of candidates for $\hat{\mathbf{a}}_{s+1}^*$ can be defined by

$$\hat{\mathbf{a}}_{s+1} = \hat{\mathbf{a}}_{s+1}^c \prod_{i=1}^{s} \mathbf{R}_i(\boldsymbol{\theta}_i) \tag{B.3}$$

Then, we can search over the set of $\hat{\mathbf{a}}_{s+1}$ for the choice of $\theta_{s+1,k} \sim U(0,\pi)$ for $s < k < \mathcal{S}-1$. Note that, by using the proposed method, we can reduce the search procedure from $\Theta(\mathbf{N}_\theta^{\mathcal{S}-1})$ to $\Theta(\mathbf{N}_\theta^{\mathcal{S}-1-s})$.

BIBLIOGRAPHY

Agrawal, S., Banala, S., Fattah, A., Sangwan, V., Krishnamoorthy, V., Scholz, J., and Wei-Li, H. (2007). Assessment of motion of a swing leg and gait rehabilitation with a gravity balancing exoskeleton. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 15(3):410 –420. (Cited on page 30.)

Aoi, S. and Tsuchiya, K. (2005). Locomotion control of a biped robot using nonlinear oscillators. *Autonomous Robots*, 19(3):219–232. (Cited on page 16.)

Aoyagi, D., Ichinose, W., Harkema, S., Reinkensmeyer, D., and Bobrow, J. (2007). A robot and control algorithm that can synchronously assist in naturalistic motion during body-weight-supported gait training following neurologic injury. *IEEE Transation on Neural System Rehabilitation Engneering*, 17(3):387–400. (Cited on pages 18 and 30.)

Argall, B., Chernova, S., Veloso, M., and Browning, B. (2009). A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483. (Cited on page 22.)

Arikan, O. and Forsyth, D. (2002). Interactive motion generation from examples. *ACM Transactions on Graphics*, 21(3):483–490. (Cited on page 16.)

Arsenault, A., Winter, D., and Marteniuk, R. (1986). Is there a 'normal' profile in emg activity in gait? *Medical and Biological Engineering and Computing*, 24:337–343. (Cited on page 13.)

Artemiadis, P. and Krebs, H. (2011). On the potential field-based control of the mit-skywalker. In *Proceedings of International Conference on Robotics and Automation*, pages 1427–1432. (Cited on pages 2, 18, 20, and 21.)

Atkeson, C. and Schaal, S. (1997). Robot learning from demonstration. In *Proceedings of International Conference on Machine Learning*, volume 97, pages 12–20. (Cited on page 24.)

Baillieul, J. (1986). Avoiding obstacles and resolving kinematic redundancy. In *Proceedings of International Conference on Robotics and Automation*, volume 3, pages 1698–1704. (Cited on page 26.)

Baker, R. (2007). The history of gait analysis before the advent of modern computers. *Gait and posture*, 26(3). (Cited on page 88.)

Banala, S., Kim, S., Agrawal, S., and Scholz, J. (2009). Robot assisted gait training with active leg exoskeleton (alex). *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 17(1):2–8. (Cited on pages 18 and 21.)

Billard, A., Calinon, S., Dillmann, R., and Schaal, S. (2007). Robot programming by demonstration. In *Springer Handbook of Robotics*, chapter 59, pages 1371–1394. MIT Press. (Cited on page 22.)

Boulic, R., Thalmann, N., and Thalmann, D. (1990). A global human walking model with real-time kinematic personification. *The Visual Computer*, 6:344–358. (Cited on page 1.)

Breniere, Y., Do, M., and Bouisset, S. (1987). Are dynamic phenomena prior to stepping essential to walking? *Journal of Motor Behavior*, 19:62–76. (Cited on page 41.)

Butz, M., Herbort, O., and Hoffmann, J. (2007). Exploiting redundancy for flexible behavior: Unsupervised learning in a modular sensorimotor control architecture. *Psychological Review*, 114:1015–1046. (Cited on page 23.)

Calinon, S., D'halluin, F., Sauser, E., Caldwell, D., and Billard, A. (2010). Learning and reproduction of gestures by imitation. *IEEE Robotics Automation Magazine*, 17(2):44 –54. (Cited on page 23.)

Chen, H. (2007). *Passive dynamic walking with knees: A point foot model*. PhD thesis, Massachusetts Institute of Technology. (Cited on pages 16 and 58.)

Chevallereau, C., Westervelt, E., and Grizzle, J. (2005). Asymptotically stable running for a five-link, four-actuator, planar bipedal robot. *The International Journal of Robotics Research*, 24(6):431–464. (Cited on page 16.)

Chumanov, E., Wall-Scheffler, C., and Heiderscheit, B. (2008). Gender differences in walking and running on level and inclined surfaces. *Clinical biomechanics*, 10:1260–1268. (Cited on page 105.)

Cleveland, W. and Devlin, S. (1988). Locally weighted regression: an approach to regression analysis by local fitting. *Journal of the American Statistical Association*, 83(403):596–610. (Cited on page 23.)

Craig, J., Hsu, P., and Sastry, S. (1987). Adaptive control of mechanical manipulators. *The International Journal of Robotics Research*, 6(2):16–28. (Cited on page 23.)

Department for Work and Pensions (2012). Family resources survey, 2010-2011. *UK Data Archive*. (Cited on page 17.)

Duschau-Wicke, A., von Zitzewitz, J., Caprez, A., Lunenburger, L., and Riener, R. (2010). Path control: A method for patient-cooperative robot-aided gait rehabilitationsky. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 18(1):38–48. (Cited on pages 2, 18, 21, and 30.)

Ekkelenkamp, R., Veneman, J., and van der Kooij, H. (2005). Lopes: selective control of gait functions during the gait rehabilitation of cva patients. In *Proceedings of International Conference on Rehabilitation Robotics*, pages 361–364. (Cited on page 21.)

Emken, J., Benitez, R., and Reinkensmeyer, D. (2007). Human-robot cooperative movement training: Learning a novel sensory motor transformation during walking with robotic assistance-as-needed. *Journal of NeuroEngineering and Rehabilitation*, 4:8. (Cited on page 20.)

Fasoli, S., Krebs, H., Stein, J., Frontera, W., Hughes, R., and Hogan, N. (2004). Robotic therapy for chronic motor impairments after stroke: Follow-up results. *Archives of Physical Medicine and Rehabilitaion*, 85:1106–1111. (Cited on page 17.)

Ferris, D., Czerniecki, J., and Hannaford, B. (2005). An ankle-foot orthosis powered by artificial pneumatic muscles. *Journal of Applied Biomechanics*, 21(2):189. (Cited on page 20.)

Fukuoka, Y., Kimura, H., and Cohen, A. (2003). Adaptive dynamic walking of a quadruped robot on irregular terrain based on biological concepts. *The International Journal of Robotics Research*, 22(3-4):187–202. (Cited on page 1.)

Garcia, M., Chatterjee, A., Ruina, A., and Coleman, M. (1998). The simplest walking model: stability, complexity, and scaling. *Journal of biomechanical engineering*, 2:281–288. (Cited on page 15.)

Geng, T., Porr, B., and Wörgötter, F. (2006). Fast biped walking with a sensor-driven neuronal controller and real-time online learning. *The International Journal of Robotics Research*, 25(3):243–259. (Cited on page 16.)

Ghahramani, Z. and Jordan, M. (1994). Supervised learning from incomplete data via an em approach. In *Advances in Neural Information Processing Systems*. (Cited on page 23.)

Gienger, M., Janssen, H., and Goerick, C. (2005). Task-oriented whole body motion for humanoid robots. In *Proceedings of International Conference on Humanoid Robots*, pages 238–244. (Cited on page 27.)

Givens, W. (1958). Computation of plain unitary rotations transforming a general matrix to triangular form. *Journal of the Society for Industrial and Applied Mathematics*, 1:26–50. (Cited on page 132.)

Gomes, M., Silveira, G., and Siqueira, A. (2011). Gait pattern adaptation for an active lower-limb orthosis based on neural networks. *Advanced Robotics*, 25(15):1903–1925. (Cited on page 23.)

Goswami, A. (1999). Postural stability of biped robots and the foot-rotation indicator (fri) point. *The International Journal of Robotics Research*, 18(6):523–533. (Cited on page 15.)

Goswami, A., Thuilot, B., and Espiau, B. (1998). A study of the passive gait of a compass-like biped robot: symmetry and chaos. *International Journal of Robotics Research*, 17(12):1282–1301. (Cited on pages 16 and 58.)

Grizzle, J., Chevallereau, C., Choi, J., and Morris, B. (2007). *Feedback control of dynamic bipedal robot locomotion*. Boca Raton: CRC press. (Cited on page 14.)

Grollman, D. and Jenkins, O. (2007). Dogged learning for robots. In *Proceedings of International Conference on Robotics and Automation*, pages 2483–2488. (Cited on page 23.)

Guan, Y., Neo, E., Yokoi, K., and Tanie, K. (2006). Stepping over obstacles with humanoid robots. *IEEE Transactions on Robotics*, 22(5):958 –973. (Cited on page 15.)

Hayashi, T., Kawamoto, H., and Sankai, Y. (2005). Control method of robot suit hal working as operator's muscle using biological and dynamical information. In

Proceedings of International Conference on Intelligent Robots and Systems, pages 3063–3068. (Cited on pages x, 19, and 20.)

Herman, R. (1976). *Neural control of locomotion*, volume 18. Plenum Publishing Corporation. (Cited on page 11.)

Herr, H. (2009). Exoskeletons and orthoses: classification, design challenges and future directions. *Journal of NeuroEngineering and Rehabilitation*, 6:21. (Cited on page 20.)

Hesse, S. and Uhlenbrock, D. (2000). A mechanized gait trainer for restoratin of gait. *Journal of Rehabilitation Research Development*, 37(6):701–708. (Cited on pages x, 18, and 30.)

Hidler, J., Nichols, D., Pelliccio, M., and Brady, K. (2005). Advances in the understanding and treatment of stroke impairment using robotic devices. *Topics Stroke Rehabilitation*, 12:22–35. (Cited on pages 2 and 20.)

Hirai, K., Hirose, M., Haikawa, Y., and Takenaka, T. (1998). The development of honda humanoid robot. In *Proceedings of International Conference on Robotics and Automation*, volume 2, pages 1321–1326. (Cited on pages 13 and 15.)

Hofmann, A., Popovic, M., and Herr, H. (2009). Exploiting angular momentum to enhance bipedal center-of-mass control. In *Proceedings of International Conference on Robotics and Automation*, pages 4423–4429. (Cited on page 27.)

Hollerbach, J. and Suh, K. (1987). Redundancy resolution of manipulators through torque optimization. *IEEE Journal of Robotics and Automation*, 3(4):308–316. (Cited on page 26.)

Honda (2009). Walk assist and mobility devices. Retrieved 2013 from http://corporate.honda.com/innovation/walk-assist/. (Cited on pages x and 19.)

Howard, M., Klanke, S., Gienger, M., Goerick, C., and Vijayakumar, S. (2008). Learning potential-based policies from constrained motion. In *Proceedings of International Conference on Humanoid Robots*, pages 714–735. (Cited on pages xiv and 79.)

Howard, M., Klanke, S., Gienger, M., Goerick, C., and Vijayakumar, S. (2009). A novel method for learning policies from variable constraint data. *Autonomous Robots*, 27(2):105–121. (Cited on pages xii, 3, 49, 52, 75, 81, and 105.)

Howard, M. and Vijayakumar, S. (2007). Reconstructing null-space policies subject to dynamic task constraints in redundant manipulators. In *Workshop on Robotics and Mathematics*. (Cited on page 52.)

Hutter, M., Hoepflinger, M., Gehring, C., Bloesch, M., Remy, C. D., and Siegwart, R. (2012). Hybrid operational space control for compliant legged systems. In *Robotics: Science and Systems*, Sydney, Australia. (Cited on page 27.)

Ijspeert, A. J. (2008). Central pattern generators for locomotion control in animals and robots: a review. *Neural Networks*, 21(4):642–653. (Cited on page 16.)

Ijspeert, A. J., Nakanishi, J., and Schaal, S. (2003). Learning attractor landscapes for learning motor primitives. In *Advances in Neural Information Processing Systems*, pages 1523–1530. (Cited on pages 16 and 24.)

Inamura, T., Toshima, I., Tanie, H., and Nakamura, Y. (2004). Embodied symbol emergence based on mimesis theory. *The International Journal of Robotics Research*, 23(4):363–377. (Cited on page 23.)

Ivanenko, Y., Poppele, R., and Lacquaniti, F. (2004). Five basic muscle activation patterns account for muscle activity during human locomotion. *The Journal of Physiology*, 556:267–282. (Cited on page 2.)

Jezernik, S., Colombo, G., and Morari, M. (2004). Automatic gait-pattern adaptation algorithms for rehabilitation with a 4-dof robotic orthosis. *IEEE Transactions on Robotics and Automation*, 20(3):574–582. (Cited on pages 2 and 21.)

Jolliffe, I. (2005). *Principal components analysis*. John Wiley and Sons, Ltd. (Cited on page 3.)

Kadaba, M., Ramakrishnan, H., Wootten, M., Gainey, J., Gorton, G., and Cochran, G. (1989). Repeatability of kinematic, kinetic, and electromyographic data in normal adult gait. *Journal of Orthopaedic Research*, 7:849–860. (Cited on page 12.)

Kajita, S. and Espia, B. (2008). Legged robots. *Springer handbook of robotics*, pages 361–389. (Cited on pages x and 14.)

Kajita, S., Kanehiro, F., Kaneko, K., Fujiwara, K., Harada, K., Yokoi, K., and Hirukawa, H. (2003). Biped walking pattern generation by using preview control of zero-moment point. In *Proceedings of International Conference on Robotics and Automation*, volume 2, pages 1620–1626. (Cited on page 15.)

Kaneko, K., Harada, K., Kanehiro, F., Miyamori, G., and Akachi, K. (2008). Humanoid robot HRP-3. In *International Conference on Intelligent Robots and Systems*, pages 2471–2478. (Cited on page 13.)

Kanoun, O., Laumond, J., and Yoshida, E. (2011). Planning foot placements for a humanoid robot: A problem of inverse kinematics. *The International Journal of Robotics Research*, 30(4):476–485. (Cited on page 27.)

Kato, I. (1973). Development of WABOT 1. *Biomechanism*, 2:173–214. (Cited on page 13.)

Kawato, M. (1990). Feedback-error-learning neural network for supervised motor learning. *Advanced neural computers*, 6(3):365–372. (Cited on page 23.)

Kazerooni, H., Racine, J., Huang, L., and Steger, R. (2005). On the control of the berkeley lower extremity exoskeleton (BLEEX). In *Proceedings of International Conference on Robotics and Automation*, pages 4353–4360. (Cited on page 19.)

Kerrigan, D., Casey, M. T., and Croce, U. D. (1998). Gender differences in joint biomechanics during walking. *Normative Study in Young Adults*, 1:2–7. (Cited on page 105.)

Khatib, O. (1986). Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1):90–98. (Cited on page 26.)

Khatib, O. (1987). A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal of Robotics and Automation*, 3(1):43–53. (Cited on page 26.)

Kim, S., Banala, S., Brackbill, E., Agrawal, S., Krishnamoorthy, V., and Scholz, J. (2010). Robot-assisted modifications of gait in healthy individuals. *Experimental brain research*, 202(4):809–824. (Cited on page 21.)

Kingma, I., Toussaint, H., Commissaris, D., Hoozemans, M., and Ober, M. (1995). Optimizing the determination of the body center of mass. *Journal of Biomechanics*, 28:1137–1142. (Cited on page 41.)

Kolter, J., Abbeel, P., and Ng, A. (2008). Hierarchical apprenticeship learning with application to quadruped locomotion. In *Advances in Neural Information Processing Systems*, pages 769–776. (Cited on page 24.)

Krebs, H., Hogan, N., Aisen, M., and Volpe, B. (1998). Robot-aided neurorehabilitation. *IEEE Transactions on Rehabilitation Engineering*, 6(1):75–87. (Cited on pages 2 and 19.)

Krebs, H., Palazzolo, J., Dipietro, L., Ferraro, M., Krol, J., Rannekleiv, K., Volpe, B., and Hogan, N. (2003). Rehabilitation robotics: Performance-based progressive robot-assisted therapy. *Autonomous Robots*, 15:7–20. (Cited on page 21.)

Kuo, A. (1999). Stabilization of lateral motion in passive dynamic walking. *The International journal of robotics research*, 18(9):917–930. (Cited on page 16.)

Lee, S. and Sankai, Y. (2002). Power assist control for walking aid with HAL-3 based on emg and impedance adjustment around knee joint. In *Proceedings of International Conference on Intelligent Robots and Systems*, volume 2, pages 1499–1504. (Cited on page 19.)

Lee, S. and Sankai, Y. (2005). Virtual impedance adjustment in unconstrained motion for an exoskeletal robot assisting the lower limb. *Advanced Robotics*, 19:773–795(23). (Cited on page 21.)

Liégeois, A. (1977). Automatic supervisory control of the configuration and behavior of multibody mechanisms. *IEEE Transation on Systems, Man and Cybernetics*, 7:868–871. (Cited on page 26.)

Lin, H., Howard, M., and Vijayakumar, S. (2014a). A novel approach for generalising walking gaits across subjects and walking speeds. In *Proceedings of International Conference on Biomedical Robotics and Biomechatronics*, pages 1009–1015. (Cited on page 6.)

Lin, H., Howard, M., and Vijayakumar, S. (2014b). A novel approach for representing and generalising periodic gaits. *Robotica*, pages 1–20. (Cited on page 6.)

Lin, H., Howard, M., and Vijayakumar, S. (2015). Learning null-space projection. In *International Conference on Robotics and Automation*. submitted. (Cited on page 6.)

Lum, S., Reinkensmeyer, D., and Lehman, S. (1993). Robotic assist devices for bimanual physical therapy: preliminary experiments. *IEEE Transactions on Rehabilitation Engineering*, 1(3):185–191. (Cited on pages 2 and 19.)

Lünenburger, L., Colombo, G., and Riener, R. (2007). Biofeedback for robotic gait rehabilitation. *Journal of neuroengineering and rehabilitation*, 4(1):1–11. (Cited on pages x and 18.)

Marchal-Crespo, L. and Reinkensmeyer, D. (2009). Review of control strategies for robotic movement training after neurologic injury. *Journal of NeuroEngeering and Rehabilitation*, 6. (Cited on page 20.)

McGeer, T. (1990a). Passive dynamic walking. *The International Journal of Robotics Research*, 9(2):62–82. (Cited on page 15.)

McGeer, T. (1990b). Passive walking with knees. *Proceedings of International Conference on Robotics and Automation*. (Cited on page 16.)

Miura, H. and Shimoyama, I. (1984). Dynamic walk of a biped. *The International Journal of Robotics Research*, 3(2):60–74. (Cited on page 15.)

Miura, K., Morisawa, M., Nakaoka, S., Kanehiro, F., Harada, K., Kaneko, K., and Kajita, S. (2009). Robot motion remix based on motion capture data towards human-like locomotion of humanoid robots. In *Proceedings of International Conference on Humanoid Robots*, pages 596–603. IEEE. (Cited on page 27.)

Mochon, S. and McMahon, T. (1980). Ballistic walking. *Journal of biomechanics*, 13(1):49–57. (Cited on page 15.)

Multon, F., France, L., Cani-Gascuel, M., and Debunne, G. (1999). Computer animation of human walking: a survey. *The Journal of Visualization and Computer Animation*, 10(1):39–54. (Cited on page 1.)

Nakamura, Y. and Hanafusa, H. (1986). Inverse kinematic solutions with singularity robustness for robot manipulator control. *Journal of dynamic systems, measurement, and control*, 108(3):163–171. (Cited on page 26.)

Nakanishi, J., Morimoto, J., Endo, G., Cheng, G., Schaal, S., and Kawato, M. (2004). Learning from demonstration and adaptation of biped locomotion. *Robotics and Autonomous Systems*, 47(2-3):79–91. (Cited on pages 16 and 23.)

Nakaoka, S., Nakazawa, A., Yokoi, K., Hirukawa, H., and Ikeuchi, K. (2003). Generating whole body motions for a biped humanoid robot from captured human dances. In *Proceedings of International Conference on Robotics and Automation*, volume 3, pages 3905–3910. (Cited on page 16.)

Neu, G. and Szepesvári, C. (2007). Apprenticeship learning using inverse reinforcement learning and gradient methods. In *Uncertainty in Artificial Intelligence*, pages 295–302. (Cited on page 24.)

Nguyen-Tuong, D., Seeger, M., and Peters, J. (2009). Model learning with local gaussian process regression. *Advanced Robotics*, 23(15):2015–2034. (Cited on page 23.)

Öberg, T., Karsznia, A., and Öberg, K. (1993). Basic gait parameters: reference data for normal subjects, 10-79 years of age. *Journal of rehabilitation research and development*, 30:210–210. (Cited on page 103.)

Ogura, Y., Aikawa, H., Shimomura, K., Kondo, H., Morishima, A., Lim, H., and Takanishi, A. (2006). Development of a new humanoid robot wabian-2. In *Proceedings of International Conference on Robotics and Automation*, pages 76–81. (Cited on page 15.)

Okada, M., Tatani, K., and Nakamura, Y. (2002). Polynomial design of the nonlinear dynamics for the brain-like information processing of whole body motion. In *Proceedings of International Conference on Robotics and Automation*, volume 2, pages 1410–1415. (Cited on page 16.)

Orendurff, M., Segal, A., Klute, G., Berge, J., Rohr, E., and Kadel, N. (2004). The effect of walking speed on center of mass displacement. *Journal of Rehabilitation Research & Development*, 41(6):829–834. (Cited on pages 13, 41, and 112.)

Pedotti, A. (1977). A study of motor coordination and neuromuscular activities in human locomotion. *Biological Cybernetics*, 26:53–62. (Cited on page 13.)

Pelossof, R., Miller, A., Allen, P., and Jebara, T. (2004). An svm learning approach to robotic grasping. In *Proceedings of International Conference on Robotics and Automation*, volume 4, pages 3512–3518. (Cited on page 23.)

Peters, J. and Schaal, S. (2008). Learning to control in operational space. *The International Journal of Robotics Research*, 27(2):197–212. (Cited on page 22.)

Plagemann, C., Mischke, S., Prentice, S., Kersting, K., Roy, N., and Burgard, W. (2008). Learning predictive terrain models for legged robot locomotion. In *Proceedings of International Conference on Intelligent Robots and Systems*, pages 3545–3552. (Cited on page 23.)

Pollard, N., Hodgins, J., Riley, M., and Atkeson, C. (2002). Adapting human motion for the control of a humanoid robot. In *Proceedings of International Conference on Robotics and Automation*, volume 2, pages 1390–1397. (Cited on page 16.)

Pongas, D., Mistry, M., and Schaal, S. (2007). A robust quadruped walking gait for traversing rough terrain. In *Proceedings of International Conference on Robotics and Automation*, pages 1474–1479. (Cited on page 1.)

Popovic, M., Goswami, A., and Herr, H. (2005). Ground reference points in legged locomotion: Definitions, biological trajectories and control implications. *The International Journal of Robotics Research*, 24(12):1013–1032. (Cited on page 15.)

Popovic, M., Hofmann, A., and Herr, H. (2004). Angular momentum regulation during human walking: biomechanics and control. In *Proceedings of International Conference on Robotics and Automation*, volume 3, pages 2405–2411. (Cited on page 15.)

Pratt, J., Chew, C., Torres, A., Dilworth, P., and Pratt, G. (2001). Virtual model control: an intuitive approach for bipedal locomotion. *The International Journal of Robotics Research*, 20(2):129–143. (Cited on page 27.)

Pratt, J., Krupp, B., Morse, C., and Collins, S. (2004). The RoboKnee: an exoskeleton for enhancing strength and endurance during walking. In *Proceedings of International Conference on Robotics and Automation*, volume 3, pages 2430–2435. (Cited on page 18.)

Raibert, M. (1986). *Legged robots that balance*. Cambridge, MA: MIT press. (Cited on page 15.)

Ramachandran, D. and Amir, E. (2006). Bayesian inverse reinforcement learning. In *Proceedings of International Joint Conference on Artificial Intelligence*, pages 2586–2591. (Cited on page 24.)

Ren, L., Patrick, A., Efros, A., Hodgins, J., and Rehg, J. (2005). A data-driven approach to quantifying natural human motion. *ACM Transactions on Graphics*, 24(3):1090–1097. (Cited on page 16.)

Riener, R., Lunenburger, L., Jezernik, S., Anderschitz, M., Colombo, G., and Dietz, V. (2005). Patient-cooperative strategies for robot-aided treadmill training: first experimental results. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(3):380–394. (Cited on pages 2 and 21.)

Roetenberg, D., Luinge, H., and Slycke, P. (2009). Xsens mvn: full 6 dof human motion tracking using miniature inertial sensors. Xsens Technologies BV, Enschede, The Netherlands. (Cited on page 89.)

Russell, S. (1998). learning agents for uncertain environments. In *Annual conference on Computational learning theory*, pages 101–103. (Cited on page 24.)

Sentis, L. and Khatib, O. (2005). Synthesis of whole-body behaviors through hierarchical control of behavioral primitives. *International Journal of Humanoid Robotics*, 2(4):505–518. (Cited on page 27.)

Sentis, L. and Khatib, O. (2006). A whole-body control framework for humanoids operating in human environments. In *Proceedings of International Conference on Robotics and Automation*, pages 2641–2648. (Cited on page 26.)

Shumway-Cook, A. and Woollacott, M. (1995). *Motor control: theory and practical applications*. Philadelphia, United States. (Cited on pages 12 and 27.)

Siciliano, B. and Khatib, O. (2007). *Springer Handbook of Robotics*. Springer-Verlag New York, Inc., Secaucus, NJ, USA. (Cited on page 14.)

Spong, M. W. (1998). Underactuated mechanical systems. In *Control Problems in Robotics and Automation*, pages 135–150. Springer Berlin Heidelberg. (Cited on page 16.)

Steger, R., Kim, S., and Kazerooni, H. (2006). Control scheme and networked control architecture for the berkeley lower extremity exoskeleton (bleex). In *Proceedings of International Conference on Robotics and Automation*, pages 3469–3476. (Cited on page 20.)

Stephens, B. (2007). Integral control of humanoid balance. In *Proceedings of International Conference on Intelligent Robots and Systems*, pages 4020–4027. (Cited on page 27.)

Stilman, M. and Kuffner, J. (2008). Planning Among Movable Obstacles with Artificial Constraints. *The International Journal of Robotics Research*, 27(11-12):1295–1307. (Cited on page 27.)

Stokes, V., Andersson, C., and Forssberg, H. (1989). Rotational and translational movement features of the pelvis and thorax during adult human locomotion. *Journal of Biomechanics*, 22(1):43–50. (Cited on page 2.)

Suleiman, W., Yoshida, E., Kanehiro, F., Laumond, J., and Monin, A. (2008). On human motion imitation by humanoid robot. In *Proceedings of Internaitonal Conference on Robotics and Automation*, pages 2697–2704. (Cited on page 16.)

Surdilovic, D., Zhang, J., and Bernhardt, R. (2007). String-man: Wire-robot technology for safe, flexible and human-friendly gait rehabilitation. In *Proceedings of the International Conference on Rehabilitation Robotics*, pages 446–453. (Cited on page 20.)

Sutton, R. and Barto, A. (1998). *Reinforcement learning: an introduction*. MIT Press, Cambridge, MA. (Cited on page 24.)

Taga, G., Yamaguchi, Y., and Shimizu, H. (1991). Self-organized control of bipedal locomotion by neural oscillators in unpredictable environment. *Biological cybernetics*, 65(3):147–159. (Cited on page 16.)

Tedrake, R., Zhang, T. W., Fong, M.-F., and Seung, H. S. (2004). Actuating a simple 3d passive dynamic walker. *IEEE International Conference on Robotics and Automation*, pages 4656–4661. (Cited on page 16.)

Towell, C., Howard, M., and Vijayakumar, S. (2010). Learning nullspace policies. In *Proceedings of International Conference on Intelligent Robots and Systems*, pages 18–22. (Cited on pages 3 and 47.)

Udwadia, F. and Kalaba, R. (2007). *Analytical Dynamics: A New Approach*. Cambridge University Press. (Cited on page 25.)

Vapnik, V. (2000). *The nature of statistical learning theory*. Springer. (Cited on page 23.)

Veneman, J., Ekkelenkamp, R., Kruidhof, R., van der Helm, F., and van der Kooij, H. (2006). A series elastic- and bowden-cable-based actuation system for use as torque actuator in exoskeleton-type robots. *International Journal Robotic Research*, 25:261–281. (Cited on pages 2 and 21.)

Veneman, J., Kruidhof, R., Hekman, E., Ekkelenkamp, R., van Asseldonk, E., and van der Kooij, H. (2007). Design and evaluation of the LOPES exoskeleton robot for interactive gait rehabilitation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 15(3):379–386. (Cited on pages 2, 18, and 30.)

Vijayakumar, S. and Schaal, S. (2000). Locally weighted projection regression: An O(n) algorithm for incremental real time learning in high dimensional space. In *Proceedings of International Conference on Machine Learning*, pages 1079–1086. (Cited on page 23.)

Vukobratović, M. and Stepanenko, J. (1972). On the stability of anthropomorphic systems. *Mathematical Biosciences*, 15(1):1–37. (Cited on page 14.)

Wang, J., Fleet, D., and Hertzmann, A. (2008). Gaussian process dynamical models for human motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):283–298. (Cited on page 23.)

Whittle, M. (2003). *Gait analysis: an introduction*. Butterworth-Heinemann, 3 edition. (Cited on page 7.)

Whittle, M. (2007). *Gait Analysis: an introduction*. Butterworth-Heinemann, 4 edition. (Cited on pages x, 8, 10, and 95.)

Winter, D. (1984). Kinematic and kinetic patterns in human gait: Variability and compensating effects. *Human Movement Science*, 3(1-2):51–76. (Cited on pages 13 and 35.)

Winter, D. (1995). Human balance and posture control during standing and walking. *Gait and Posture*, 3(4):193–214. (Cited on pages 11 and 40.)

Winter, D. (2009). *Biomechanics and motor control of human movement*. John Wiley and Sons. (Cited on pages 59, 115, and 117.)

Wisse, M. (2004). Three additions to passive dynamic walking; actuation, an upper body, and 3d stability. *International Conference on Humanoid Robots*, pages 113–132. (Cited on pages 16 and 116.)

Wolpert, D. and Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, 11:1317–1329. (Cited on page 24.)

World Health Organization (2011). World report on disability. *WHO press*. (Cited on page 17.)

Zatsiorsky, V. and King, D. (1997). An algorithm for determining gravity line location from posturographic recordings. *Journal of Biomechanics*, 31:161–164. (Cited on pages 41 and 99.)

Zeilig, G., Weingarden, H., Zwecker, M., Dudkiewicz, I., Bloch, A., and Esquenazi, A. (2012). Safety and tolerance of the rewalk exoskeleton suit for ambulation by people with complete spinal cord injury: A pilot study. *The Journal of Spinal Cord Medicine*, 35:101–196. (Cited on page 18.)

Zhao, Y., Kim, D., Fernandez, B., and Sentis, L. (2013). Phase space planning and robust control for data-driven locomotion behaviors. In *International Conference of Humanoid Robots*. (Cited on page 23.)

Zoss, A., Kazerooni, H., and Chu, A. (2005). On the mechanical design of the berkeley lower extremity exoskeleton (BLEEX). In *Proceedings of International Conference on Intelligent Robots and Systems*, pages 3465–3472. (Cited on pages x and 19.)