UTS UNIVERSITY OF TECHNOLOGY SYDNEY

Towards Personalised Robotic Assessment and Response during Physical Human Robot Interactions

by Yujun Lai

Thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy

under the supervision of Dr. Gavin Paul and Dr. Marc Carmichael

University of Technology Sydney
Faculty of Engineering and Information Technology

March 2022

Certificate of Original Authorship

I, YUJUN LAI, declare that this thesis, is submitted in fulfilment of the requirements for the award of the degree of Doctor of Philosophy, in the School of Mechanical and Mechatronics Engineering, Faculty of Engineering and Information Technology (FEIT) at the University of Technology Sydney.

This thesis is wholly my own work unless otherwise referenced or acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis.

This document has not been submitted for qualifications at any other academic institution.

This research is supported by the Australian Government Research Training Program.

Signed: Production Note: Signature removed prior to publication.

Date: March 2022

Towards Personalised Robotic Assessment and Response during Physical Human Robot Interactions

by

Yujun Lai

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Abstract

The exponential growth of robotics in human environments have led to an explosion of human robot interactions. These interactions occur in proximity and have exposed the constraints and limitations of traditional models for robotic response which rely on task-centric measures. This has spurred on an area of research which focuses on understanding the capabilities and limitations of the human user during these interactions.

Humans are complex, autonomous agents that are difficult to model, and provide different categories of feedback that derive from biological systems. The current sensory paradigm requires an improved understanding of the limitations, the development of blended-measure models that employ human-centric measures, and a contextually connected biological human understanding into robotic frameworks.

This thesis presents a framework towards personalised robotic assessment and response with considerations on understanding the human user during physical human robot interactions. The framework approaches this by *examining* current limitations, *enabling* personalised models from human-centric measures, and *enhancing* the understanding of the human user through physiological and musculoskeletal models.

The implementation of a robotic system highlights the feasibility and limitations of using task-centric models during Physical Human Robot Interaction (pHRI). Further work

investigates inertial effects of the user during interactions in the context of a prominent predictive model, Fitts' Law. Physical Human Robot Interaction Primitives (pHRIP) extends upon Interaction Primitives (IPs) by incorporating physical interaction forces between the human user and robot, enabling the inference of user intent when generating a personalised robotic response.

Finally, the enhancement of the link between biological human understanding and robotic frameworks is explored. A validation process for a popular musculoskeletal model is conducted, comparing computational results with experimental readings. The limitations for the complex model led to the generation of an empirical model correlating forearm muscle activity and grip strength. This physiological model captured co-contractions for antagonistic muscle pairs and supplemented motion analysis for the musculoskeletal model, enhancing the computational results.

The framework combines the topics which facilitate intuitive and adaptive human-robot interactions. The advancement of such collaborative intelligence enhances complementary strengths between human and robot, and work hand in end-effector towards a safer, more interactive future.

Acknowledgements

I want to take this opportunity to thank all the people that have supported me and contributed towards this unique experience.

First and foremost, I would like to thank my supervisors, Dr. Gavin Paul and Dr. Marc Carmichael, who took a leap of faith and guided me through this journey together. The passion that you both hold for those you mentor has always been illuminated by your direction, your motivation, your optimism (even at the worst of times), and most of all, your actions. The ever-present support (directly and indirectly) that you both display has enabled me to freely explore various topics and interesting projects. This body of work would not exist without your support.

This journey also includes my unexpected involvement in the realm of Teaching and Learning. Chasing the "a-ha" moments from students through personal interactions have exposed my mind to a diverse set of lived-in experiences. Coupled with the attendance of FFYE forums, coordinated by Kathy Egea, these invaluable experiences have undoubtedly shaped part of who I am today, and inform who I will be moving forward.

I want to thank all the staff and students over the years at the Robotics Institute (formerly Centre for Autonomous Systems) for providing advice and assistance. Particularly, I want to thank those who have enriched my day-to-day life by engaging in shared celebration and commiserations, and healthy discussions on various topics: Sheila Sutjipto, Richardo Khonasty, Stefano Aldini, Julien Collart, Jonathan Woolfrey, Andrew To, Antony Tran, and Maleen Jayasuriya.

I would like to give my sincere gratitude to Prof. Takamitsu Matsubara and Dr. Yunduan Cui at the Nara Institute of Science and Technology's Robot Learning Lab for the collaboration and the opportunity to work with other students in the lab during my 2019 exchange visit. I would also like to express my thanks to Mrs. Mioko Fukuda at the International Institute for Advanced Studies for facilitating my accommodation during that time. Additionally, I would like to thank Prof. Xiong Rong at Zhejiang University, Yuquan Campus, for facilitating my short exchange visit to her lab in 2018, and Zhong Xiang Zhou for his guidance during that time.

A special thanks goes to those who stood by me, providing stability during the worst of times, and fun and exciting adventures during the best of times.

A special thanks also goes to *the boys*, who have always provided banter and support to me and each other through thick and thin: Rafid Morshedi, Anthony Zhang, Suhad Hussain, Mir Tarique, and in memory of Sibi Suresh Babu.

Finally, I would like to give the largest thanks to my parents, Sing Choon Lai and Siew Ean Ooi, for their silent but never-ending support. Being educators themselves, the importance of education has always been imprinted on me during my childhood. Without their tireless efforts to provide me with every opportunity available while growing up, including our family's move from Malaysia to Australia, I would not be the inquisitive person I am today. Many of my current achievements can be attributed to their persistence, encouragement, and care.

Contents

List of Figures

List of Tables

Acronyms & Abbreviations

AAN	Assistance as Needed
ADL	Activites of Daily Living
ANOVA	Analysis of Variance
BIPs	Bayesian Interaction Primitives
BCM	Bayesian Committee Machine
CNS	Central Nervous System
Cobot	Collaborative Robot
DBN	Dynamic Bayesian Network
DE	Differential Evolution
DLO	Deformable Linear Object
DMP	Dynamic Movement Primitives
DoF	Degrees of Freedom
DTW	Dynamic Time Warping
EEG	Electroencephalogram
EKF	Extended Kalman Filter
EM	Expectation Maximisation

EMG Electromyogram

FFT Fast Fourier Transform

FLV Force-Length-Velocity

GMM Gaussian Mixture Model

GP Gaussian Process

GP-BCM Gaussian Process-Bayesian Committee Machine

GPLVM Gaussian Process Latent Variable Model

GSR Galvanic Skin Response

HMM Hidden Markov Model

HRI Human Robot Interactions

IID Independent and Identically Distributed

IMU Inertial Measurement Unit

IPs Interaction Primitives

IVM Informative Vector Machine

LfD Learning from Demonstration

LWSS Longest Warping Subsequence

MAP Maximum A-Posteriori

MDN Mixture Density Network

ML Machine Learning

MPs Movement Primitives

MR Mixed Reality

MTU Muscle Tendon Unit

MVC Maximum Voluntary Contractions

PbD Programming by Demonstration

PCA Principal Component Analysis

PDF Probability Distribution Function

pHRI Physical Human Robot Interaction

pHRIP Physical Human Robot Interaction Primitives

POMDP Partially Observable Markov Decision Process

ProMPs Probabilistic Movement Primitives

PSD Power Spectrum Density

RBF Radial Basis Function

RCT Randomised Control Trial

RMS Root Mean Square

RMSE Root Mean Square Error

RCT Randomised Control Trial

SE Squared Exponential

SEA Series Elastic Actuator

sEMG Surface Electromyogram

SSE Sum of Squared Errors

VIA Variable Impedance Actuator

VR Virtual Reality

Glossary of Terms

Autonomous	Without human intervention.
Allocentric	Attention centered on other objects or persons. Antonym of *egocentric*.
Forward Kinematics	The process of translating system joint states into the Cartesian endpoint pose.
Gamification	The process of turning a task into a game.
Human-centric	Derived from or associated with the human.
Human-robot Dyad	A system consisting of one human and one robot.
In-silico	Testing and analysis that occurs in computational models.
In-vivo	Testing that occurs with living organisms (such as humans).
Inverse Kinematics	The process of translating a Cartesian pose into joint states of a system.
Multi-modal	Consisting of multiple *modes*, or ways.
One-shot	Using a single take or demonstration.
Osseointegration	The process of inserting a prosthetic into the residual bone of an amputee.
Proximate Cause	The immediate apparent cause of a phenomenon.
Semi-autonomous	With some human intervention.
Task-centric	Derived from or associated with the task.
Ultimate Cause	The primitive cause for a phenomenon. Usually proximate causes are the symptom of the ultimate cause.
Unimodal	Consisting of a single *mode*.

Chapter 1

Introduction

1.1 Background and Motivation

The increasing prevalence of robots in every aspect of human life is indisputable. Human behaviour has changed significantly through our interactions with these devices which are embodied in different shapes and sizes and presented through various mediums. With an industry estimated to be USD12.3 billion by 2025 [1], robots will continue to exist alongside humans and accelerate the number of instances for Human Robot Interactions (HRI).

Traditionally designed to overcome limitations of human workers in assembly lines, robots and robotic systems led the charge during the automotive industry boom in the 1970s. These industrial robots require little to no environmental feedback to achieve their predefined tasks and work in isolation to mitigate injuries and damage to their environment.

Since robots are replaceable, they are designed for applications which are predominantly dirty, dull, or dangerous. This principle (colloquially known as "The Three Ds") is still relevant in current robot deployments, including *semi-autonomous* and *autonomous* systems. The Canadarm2 [2] and Ingenuity helicopter [3] follows this principle, using contemporary technologies to enable complex tasks to be performed in a vacuum by robots.

Robots working in an isolated workspace are limited by the level of software redundancies programmed during deployment. For dynamic tasks and environments, accounting for the occurrence of all possible events creates bloated software, with ad-hoc software updates generally avoided since they may be unreliable. Thus, there has been a shift towards robotic systems that will work with humans to complete tasks since humans have superior

cognition and perception capabilities. Currently, this collaboration remains a challenging proposition for robots.

The state of affairs for HRI systems were scrutinised in the late 20^{th} century and early 21^{st} century [4, 5], identifying necessary improvements in various areas of research. Since then, the burgeoning pace of progress for robot systems and HRI has lead to its prominence in a variety of applications, including in the medical response to the recent pandemic [6], cementing its significance in our future.

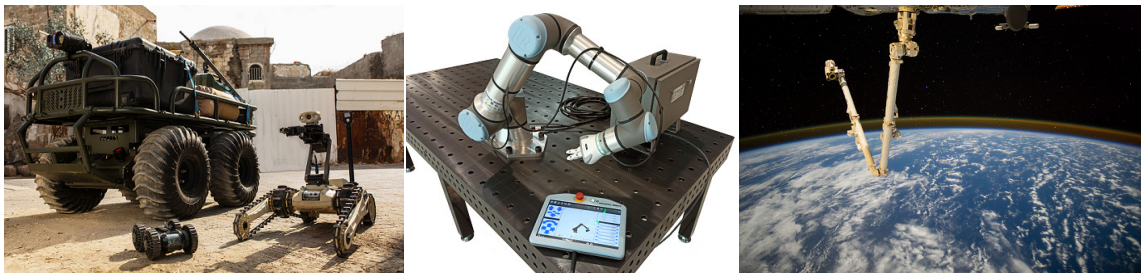

(A) (B) (C)

FIGURE 1.1: Current robots come in different shapes and sizes and operate in a variety of environments. (A) The Roboteam robots; (B) the UR16e cobot; and (C) the Canadarm2 on the International Space Station. (Wikimedia Commons for all figures)

With an increasing adoption rate for robots, there is a corresponding rise in interactions between humans and robots, leading to novel applications which highlight the limitations of HRI. In particular, tasks which rely upon human expertise are difficult for robots to replicate as humans generate novel responses through our critical thinking skills. When given a loosely defined task, humans are able to extrapolate meaning and infer the intended outcomes - a feat unmatched by traditional robots.

The performance of a robot is evaluated by metrics which are chosen based on the intended outcomes. Since early HRI applications are derived from traditional methodologies, successful collaboration and cooperation are commonly prescribed through *task-centric* measures. Some examples of these metrics are outlined in Table 1.1, and are similar or identical to those used in traditional robots, assessing the performance of the system *in the context of a pre-determined task*. While this approach can work for some contemporary applications, such as Assistive robotics, the limitations of this model become apparent as complexity increases and the objectives become more ambiguous.

Modern robotics has enabled new paradigms where human safety can be ensured while in

close proximity to a robot. The advent of Collaborative Robots (Cobots), robots with intrinsic safety features, is complemented by better control algorithms [7] and safety frameworks [8]. These improvements have enabled more collaborative ways to interact with robotic systems [9]. This has resulted in Physical Human Robot Interaction (pHRI) [10], a field of HRI incorporating physical interactions between the user and robot. pHRI considers additional aspects of robotics which are present when humans are in close proximity to robots. This includes new approaches for robotic design, actuation, motion planning, and reflexes [11], creating opportunities to transform traditional spaces, such as collaborative work cells in manufacturing assembly processes [12].

Adding information from physical interactions is challenging since a coupled system has different properties which may cause instabilities in conventional frameworks. For *human-robot dyads*, understanding the physical characteristics of the human partner, along with the task, can inform the system to improve pHRI and task outcomes [13].

The inherent variability exhibited by humans propel the need for personalised robot responses to enable productive and meaningful applications. Since traditional perspectives rely on accuracy and reliability, characteristics that are difficult for humans to achieve, *task-centric* measures are ill-equipped to facilitate pHRI. Thus, a different category of metrics is needed to support the user *without the context of any particular task*, using user-dependent characteristics and behaviour for personalisation.

Human-centric measures utilise human characteristics to shape the robot response, including demographic artefacts and individual idiosyncrasies unique to the user. This form of personalisation is currently present in the digital space, delivering personalised content on the internet using demographic and individual data. Using a similar strategy, pHRI can leverage distinct characteristics from individuals to provide a personalised response which improves user experience and potentially improves task outcomes.

Table 1.1: Examples of measures and metrics used in robotics.

	Metric Type	Examples
Task-centric	Kinematics	Position, Velocity
	Dynamics	Acceleration, Force, Torque
	Task-dependent	Trajectory Phase, Actions
Human-centric	Psychological	Emotional State, Attention, Non-verbal Communication
	Physiological	Ergonomics, Manipulability, Fatigue, Strength Capability

Human autonomy and the unique processes in our Central Nervous System (CNS) generate inherent noise, resulting in variations in human motion and outcomes. This can be seen in Activites of Daily Living (ADL) where simple objectives, such as object reaching, can be completed in various ways and exhibit traits unique to each person. This variation among the human population, and individualistic styles, can be captured using probabilistic distributions, modelling the generation of noise in our actions, a key component of user personalisation for robotics.

The simplicity, flexibility, and efficacy of probabilistic distributions makes it an attractive method to model applications where there is uncertainty and noise. During pHRI, the coupled system allows us to leverage the robot to capture human intentions and actions. These probabilistic models can then be integrated into Learning from Demonstration (LfD) frameworks, capturing variations in the possible actions of the human user.

Other than robotic information from the interactions, an alternative approach to capture the behaviour of humans is through the lens of the human physiological system. Musculoskeletal models are analogues of the human musculoskeletal system, providing functionality to compare, contrast, and investigate physiological responses through representative simulations.

This thesis presents an investigation into models which assess and assist the human user in pHRI systems with a focus on personalised interactions. Contextualising musculoskeletal models for robotics and finding methods for inferring user intent using physical interactions can provide new perspectives for future methods personalising assessment and response during pHRI.

1.2 Research Questions

Existing works for pHRI have relied on traditional metrics to arbitrate the robot response. Most rely on external sensors, such as a camera, to observe and extract information on the user. However, if the observations of the human user are unavailable, these frameworks are heavily limited in their ability to assess and assist the user.

During pHRI, additional channels of information on the human user (human-centric measures) can be made available without an external observation of the user. This warrants an investigation into models which use a mixture of task-centric and human-centric measures to assess the user and provide appropriate robot responses. Furthermore, there are

promising indications which suggest cross-disciplinary techniques and models can be integrated for pHRI. These techniques and models are heavily used in biological sciences since they provide insight into intrinsic characteristics of the user in a human-robot dyad.

To investigate the integration of task-centric and human-centric measures in coupled human-robot dyads, there are two key aspects to explore: (1) how effective are task-centric measures in personalised user assessment during pHRI, and (2) how can human-centric measures, obtained during interactions, be used to supplement task-centric models when providing personalised robotic responses? To leverage the understanding of intrinsic human characteristics in a coupled human-robot dyad, a further question is: (3) can models built upon human-centric measures supplement musculoskeletal models for human motion analysis in pHRI applications?

1.3 Scope

This research aims to explore a framework that integrates task-centric and human-centric measures for personalised robotic assessment and response during pHRI. It will investigate different perspectives for personalising robotic assessment and response, yielding insight into the differences between task-centric and human-centric models in pHRI. The framework developed in this thesis aims to: (1) *examine* both the current methods for task-centric personalisation and the available methods to overcome their limitations; (2) *enable* personalised robot responses during pHRI by integrating human-centric measures; and (3) *enhance* the understanding of the human user during pHRI, focusing on physiological and musculoskeletal models. The scope of the thesis is constrained to the investigation of each aspect in isolation. The thesis assumes the validity of existing models to be correct, relying on the inherited high-level properties with limited consideration for their corresponding importance across all perspectives. Furthermore, the feedback loop to integrate the different aspects together in a multi-model and multi-modal pHRI application is outside of the scope of this thesis.

For each aspect of the framework, there are assumptions and constraints made to enable the technical scope of this thesis to be achievable. For capturing the personalised actions of a user during pHRI, only probabilistic methods which follow Bayesian statistics are employed despite the existence of a mature field of research on manifold-based and iterative learning such as Deep Learning and Reinforcement Learning. This is mainly due to the number of opportunities to obtain data sets in pHRI applications which are repeatable and

consistent. Variations in human actions and human-centric measures can be modelled as signal noise which are assumed to be Independent and Identically Distributed (IID) and can be captured using probabilistic models.

While pHRI covers any HRI which includes physical contact (direct or indirect), this thesis focuses on dense, continuous interactions in a coupled robot-human dyad where the user interacts with the robot's endpoint using a handle. Two categories of physical interactions which are out of the scope of this thesis include indirect physical contacts and sparse contacts, commonly seen during tele-operation and corrective actions respectively. While these factors are outside the scope of this thesis, this research can provide the foundation for future work addressing these challenging issues.

Finally, for physiological and musculoskeletal models, the scope of this thesis explores qualitative validation of the results provided by these models. The vision is to utilise these models as an adjunct to quantitative models, thus tests are conducted on a single musculoskeletal model commonly used in research for human motion analysis. Furthermore, human motion analysis performed in this thesis is limited to commonly used processes which are available in open source software. In particular, the scope of the thesis does not include the active area of research on different optimisation methods to obtain muscle activity. The works presented in this thesis utilise the most commonly used optimisation process and objective function for musculoskeletal models.

1.4 Contributions

The main contributions of this thesis[1] are:

- A framework to investigate the integration of task-centric and human-centric measures for personalised robotic assessment and response.
- A derivative framework for an example of a real-world application, including an exploration of methods to overcome limitations of GP models during task-centric robotic assessment.
- A prominent information-theoretic model for user intent estimation is evaluated, and an expansion on a probabilistic method to assess user intent is conducted to contextualise pHRIs.

[1] A video presentation overview of the contributions can be found here.

- A preliminary investigation into the integration of human-centric models with musculoskeletal models for human motion analysis.

1.5 Publications

1.5.1 Directly Related Publications

- [14] **Lai, Y.**; Poon, J.; Paul, G.; Han, H.; Matsubara, T.; "Probabilistic Pose Estimation of Deformable Linear Objects", 2018 IEEE 14th International Conference on Automation Science and Engineering (CASE), 2018, 471-476, doi:10.1109/COASE.2018.8560497

- [15] **Lai, Y.**; Sutjipto, S.; Clout, M.; Carmichael, M.; Paul, G.; "GAVRe^2 : Towards Data-Driven Upper-Limb Rehabilitation with Adaptive-Feedback Gamification", 2018 IEEE International Conference on Robotics and Biomimetics (ROBIO), 2018, 164-169, doi:10.1109/ROBIO.2018.8665105

- [16] **Lai, Y.**; Sutjipto, S.; Carmichael, M.; Paul, G.; "Heuristic Detection of Recovery Progress using Robotic Data", 2019 IEEE 9th International Conference on Cybernetics and Intelligent Systems and Robotics, Automation and Mechatronics, CIS and RAM 2019, 2019, 506-511, doi:10.1109/CIS-RAM47153.2019.9095835

- [17] Sutjipto, S.; **Lai, Y.**; Carmichael, M.; Paul, G.; "Fitts' law in the presence of interface inertia", 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), 2020, 4749-4752, doi:10.1109/EMBC44109.2020.9176195

- [18] **Lai, Y.**; Paul, G.; Cui, Y.; Matsubara, T.; "User Intent Estimation during robot learning using Physical Human Robot Interaction Primitives", Autonomous Robots (AURO), 2022, Vol. 46, 421-436, doi:10.1007/s10514-021-10030-9

- [19] **Lai, Y.**; Sutjipto, S.; Carmichael, M.; Paul, G.; "Preliminary Validation of Upper Limb Musculoskeletal Model using Static Optimization", 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), 2021, 4509-4512, doi:10.1109/EMBC46164.2021.9629494

1.5.2 Partially Related Publications

- [20] Le, D.; Sutjipto, S.; **Lai, Y.**; Paul, G.; "Intuitive Virtual Reality based Control of a Real-world Mobile Manipulator", 16th IEEE International Conference on Control, Automation, Robotics and Vision, ICARCV 2020, 2020, 767-772, doi:10.1109/ICARCV50220.2020.9305517
- [21] Nguyen, D.; **Lai, Y.**; Sutjipto, S.; Paul, G.; "Hybrid Multi-Robot System for Drilling and Blasting Automation", 16th IEEE International Conference on Control, Automation, Robotics and Vision, ICARCV 2020, 2020, 79-84, doi:10.1109/ICARCV50220.2020.9305391
- [22] Aldini, S.; **Lai, Y.**; Carmichael, M.; Paul, G.; Liu, D.; "Real-time Estimation of the Strength Capacity of the Upper Limb for Physical Human-Robot Collaboration", 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), 2021, 4533-4536, doi:10.1109/EMBC46164.2021.9630230
- [23] Chotisathiantham, P.; **Lai, Y.**; Paul, G.; "Design of a Wearable Robotic Glove for Rehabilitation", Australasian Conference on Robotics and Automation (ACRA) 2021

1.6 Thesis Outline

Chapter 2 contains a survey of related work in the field of the personalisation of pHRI. This includes different metric types and their applications, probabilistic models, and conventional models for user intent recognition.

Chapter 3 introduces the framework which *examines* the limitations of current methods, *enables* the integration of human-centric measures to personalise the robot response, and *enhances* current insights into the human user during pHRI. An example of a real-world application, contextualised for robotic rehabilitation, is then presented with initial works on overcoming the limitations of current probabilistic methods outlined.

Chapter 4 evaluates the efficacy of Fitts' law, an information-theoretic model for user intent estimation, focusing on the inertial effects of the robot. Physical Human Robot Interaction Primitives, an extension on a popular HRI framework, is then presented, integrating interaction forces for personalised user intent inference and robot response.

Chapter 5 contains an exploration into musculoskeletal models and human physiology using a model regularly used in research. A supplementary model mapping between grip strength and muscular activity is presented, evaluating its efficacy and its influence as an adjunct to musculoskeletal models in pHRI tasks.

Chapter 6 presents a summary of the contributions, including a discussion on their significance in future applications. Conclusions are drawn based on the presented work and potential future works are identified.

Appendix A details the modified musculoskeletal models used for validation and application of the grip strength model. Appendix **??** presents the user survey results from the application of the framework in Chapter 3. Appendix B-C details the methodology for Bayesian Committee Machine (BCM) and the Longest Warping Subsequence (LWSS) score, along with their respective evaluative comparisons against conventional GP.

Chapter 2

Review of Related Work

Robotic assistance and assessment provide advantages which overcome limitations of subjective measures. This is apparent in traditional settings where human expertise is required, creating inconsistencies during assessments and assistance, since anecdotal experiences vary from one expert to the next [24].

Introducing robots for assessment and assistance can alleviate these challenges by utilising reliable and accurate measurements obtained from the robotic system. However, key challenges include: (1) understanding the limitations of using raw data obtained from robotic systems during HRI and pHRI; (2) exploring models which capture unique mannerisms during interactions, personalising the user's assessment and robotic response; and (3) investigating human-centric methods which provide insight into the human user when performing robotic assessments.

An overview of current applications for pHRI is presented in Section 2.1, including a closer look at the influence of human-centric measures in pHRI. The next section (Section 2.2) explores models and concepts for representing data variance in the context of task-centric measures obtained from robotic systems.

Robotic responses can be provided using both task-centric and human-centric measures. Thus, Section 2.3 introduces traditional and contemporary models which integrate both types of measures for estimating user intent and providing a robotic response. A common method for exploring human-centric factors is through the use of musculoskeletal models. Section 2.4 briefly introduces musculoskeletal models and their applications in human motion analysis. Current trends in utilising musculoskeletal models for pHRI is then

briefly outlined. Finally, related works are summarised in Section 2.5 with a discussion on their limitations.

2.1 Physical Human Robot Interactions

Physical Human Robot Interaction (pHRI) extends upon HRI by capturing applications where there is some form of physical interaction between the human and robot. These interactions can be done while coupled physically or through another interface, as shown in Figure 2.2, enabled by various algorithms and frameworks that utilise artefacts from these interactions.

While there is a focus on the utilisation of the interaction artefacts, other factors have contributed towards the adaptation of traditional HRI applications and the adoption of pHRI systems. Safer robot design and the resultant frameworks overcome prior concerns for humans and robots working in close proximity [25, 26].

2.1.1 Robot Design and Control

Early works for HRI relied on software features to incorporate safety when interacting with a robot. However, improvements in robotic hardware design have accelerated the adoption of pHRI, with robots working in proximity with humans. The introduction of Variable Impedance Actuators (VIAs) [27] and Series Elastic Actuators (SEAs) [28] provided new control methods and paradigms for safe pHRI.

These new types of actuators allow research to build upon traditional HRI control methods, enhancing the actuator's usability [29, 30]. The two most recognised frameworks are admittance control [31], which translates endpoint velocities into joint velocities, and impedance control [32], which translates endpoint forces into joint torques. Subsequent works have built upon these theoretical foundations to improve HRIs [33, 34].

Investigations into the advantages of pHRI led to frameworks which provide variable robotic assistance to the user. The assistive frameworks utilise common metrics [35] to ascertain the amount of assistance given, an approach termed Assistance as Needed (AAN). AAN frameworks estimate user strengths and deficiencies, and provide targeted assistance to improve task outcomes, with applications in manual labour [34], medical clinics [36], and robotic rehabilitation [37–41].

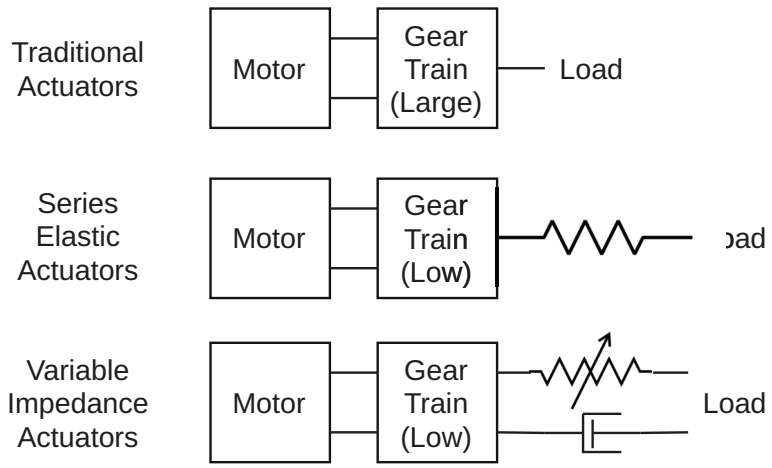

FIGURE 2.1: A block diagram of the new actuators designed for safe pHRI. Note the different gear ratios when comparing between traditional actuators, SEAs, and VIAs.

The ability for robots to work in close proximity to humans also opens up opportunities to consider anthropomorphic designs which improve the interactions. This bio-inspired perspective has produced the design of shoulder joints mirroring the dexterity of the human shoulder [42], and a full upper-limb exoskeleton design for robotic assistance [43]. This extends to control frameworks exploiting kinematic redundancy in robot arms to mirror the human arm during interactions [44]. Visual and sensorimotor feedback from interactions shape the resultant behaviour of the robot arm, using a range of measures such as estimated endpoint stiffness [45], endpoint velocities [46], visual and oral feedback [47], and kinematic configuration [48, 49].

More recently, there is an increasing body of research flipping the narrative for robotic assistance and assessment, and dropping the assumption that the human user is an expert. This perspective asserts new beliefs on the arbiter of truth during interactions, creating interesting models based on confidence and trust [50]. Confidence and trust frameworks offer a fresh perspective into human-centric assistance and assessment [51], by drawing insights from the user into the effective use of physiological responses.

2.1.2 Learning from Demonstration (LfD)

The objective of HRI and pHRI is to encourage the integration of human expertise with benefits derived from robotic systems. As such, a representative set of frameworks would primarily leverage expertise shown by humans to learn the task. Learning from Demonstration (LfD), also referred to as Programming by Demonstration (PbD) [52], builds a

model for robot trajectories based on an observed demonstration from the human user, forming the most popular approaches towards pHRI in the context of pre-determined tasks [53–56].

Notably, generalising from multiple observed demonstrations is challenging, with many works using probabilistic models to capture the variance. This body of work is expanded upon in Section 2.3.4 after a review of probabilistic models.

Traditional LfD frameworks focus on the task at hand, with contemporary works adapting the learned states based on a shift in the task outcomes. Thus, they provide personalised responses through adaptation in the task-space which rely on task-centric measures [57].

Task-centric adaptations also exist when the human is removed from direct physical interactions with the robotic system. In these circumstances, leveraging human expertise relies on an interface to communicate relevant feedback between the user and the robot.

A common method to facilitate this indirect interaction is through a phantom device providing feedback to the human user, a process termed *tele-operation* [58]. To provide the physical sense of the robot, haptic feedback is usually provided through the phantom device. This indirect feedback creates challenges for robot control to translate the user's interaction profiles during tele-operation [59, 60]. It is noted that contemporary literature using VR and Mixed Reality (MR) attempts to address this issue by augmenting communication and feedback (visual, aural, and haptic) during tele-operation between the robot, phantom device, and user [20, 61, 62]. A small sample of these types of interaction interface can be seen in Figure 2.2.

(A) (B) (C)

FIGURE 2.2: Physical interactions with a robotic system using various interfaces such as: (A) a physical hand (Mark Longair, Creative Commons); (B) tele-operation (Tecnalia, Creative Commons); and (C) Virtual Reality environment.

Control frameworks for LfD tackle this issue by modulating the apparent inertial property of the robot endpoint. This results in control algorithms which perform kinematic reconfiguration and/or torque control to modulate the endpoint stiffness [63], compliance [64, 65], and impedance [66, 67]. These control frameworks rely on data observed directly from the points of contact, using the raw task-centric measurements to adapt the endpoint behaviour. However, this reliance on task-centric measurements at the points of contact creates some drawbacks during pHRI.

Since humans are adept at creating compensatory behaviours to complete a task, the behaviour at the points of contact may remain similar despite vastly different circumstances. For example, if the interaction forces were measured when lifting a heavy object, the force profile can be identical regardless of the posture used when performing the task. Thus, there could be a degradation in the interaction experience which is not captured through task-centric measures.

2.1.3 Human-centric Measures during pHRI

Human-centric measures draw insight from the confluence of kinematic data and physiological measures to generate information on the user during interactions, providing new perspectives for pHRI. Rather than replacing task-centric measures, the consensus in the literature is to use human-centric measures to supplement conventional frameworks. This integration is inspired by investigations into our understanding of motor learning during interactions [68–71], corroborating our biological understanding of how humans learn and the implications for pHRI [72–74].

The main set of metrics which are most associated with human-centric measures are physiological responses from the human user. Since they are autonomic in nature, and extensively explored in the medical sciences, the use of physiological measures in robotic applications have been widely accepted, providing intuitive frameworks for pHRI.

Three common physiological measures used in pHRI include Electromyogram (EMG) [75], Electroencephalogram (EEG) [76], and Galvanic Skin Response (GSR) [77]. Their popularity derives from their accessibility, without the need for intrusive procedures, which facilitates their integration into robotic assistance and assessment frameworks.

Estimating the workload experienced by the user during pHRI is a difficult task since every person has a different limit. Thus, there is a significant field of research exploring the estimation of workload, both physical and cognitive, using physiological measures [78].

Cognitive workloads are generally explored in applications without a physical component, such as HRI, using GSR [79] and EEG [80].

While there are some pHRI applications which utilise EEG [81, 82], the physical interaction lends itself to using Surface Electromyogram (sEMG), indicating muscle activity, to assess user effort and provide appropriate assistance. Early works have used sEMG readings to explore the relationship between stress and limb stiffness as defined by muscle co-contractions [83]. For augmentative applications, such as in exo-skeletons, sEMG has been used to complement task-centric measures to control the exo-skeleton system and provide assistance to the user [84]. A similar framework is also applied to an EMG-controlled hand for robotic rehabilitation [85], and in the learning of a user's assistance profile using Machine Learning (ML) [86].

Supplementing interaction artefacts with physiological measures can also improve upon conventional frameworks, such as LfD, by producing stiffness and impedance profiles which match the environment and the task [67]. One other human-centric measure which affects human performance during pHRI is muscle fatigue, the acute degradation of force generation in muscle fibres, manifested in the sEMG readings. This has been leveraged to assess the user, inputting data from sEMG readings into a muscle fatigue model, to generate targeted response at the endpoint during pHRI [87] and adapt the robot response over time [88].

Shifting the focus of interactions to the human user, there are measures which have been adopted from their traditional roots in robotics such as the kinematic manipulability [89]. The endpoint accelerations of humans were used to generate the human muscular manipulability [90], while the force profiles of the human hand were used to assess user capability via their muscular manipulability [91]. Both examples, shown in Figure 2.3, assess the user using the manipulability measure, and adapt the robotic response provided to augment user capability at the endpoint.

Human-centric measures can also result from indirect measurements which extrapolate meaningful understanding of the human status during pHRI. Human ergonomic design consider multiple modes of data collection to obtain a holistic measure of the user's physical state when complete a task. These are usually based on human-focused injury risk factors which can occur when the task is being performed incorrectly. In coupled pHRI, human ergonomics are used to re-position the points of contact, mitigating injury risks for the user during physical interactions. Using this approach, researchers have explored the assessment of human ergonomics in robotics, resulting in methods that allow the user to

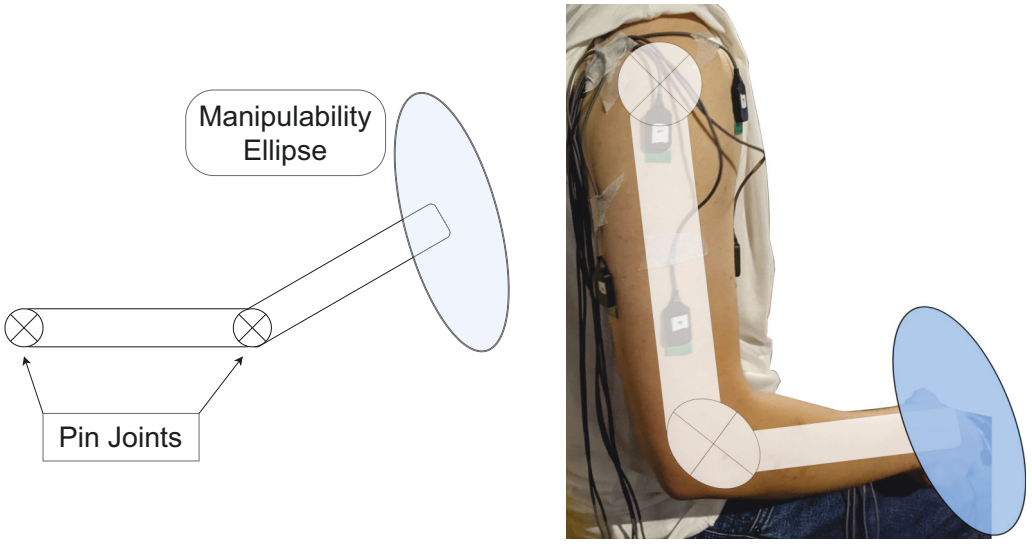

(A) (B)

FIGURE 2.3: The adoption of traditional robotic measures into pHRI, such as from: (A) robotic kinematic manipulability, into (B) human arm manipulability.

re-position their arm during coupled pHRI tasks [92]. This assessment was performed by integrating a muscle fatigue model and derivative information from a musculoskeletal model in its calculation.

Indirect measurements of the human posture, obtained from camera observations or motion capture systems, have also been performed to calculate the risk of joint overloading during a collaborative task, re-positioning the pose of the object to allow the user to return to an ergonomic pose [93]. More recently in [94], the total interaction force experienced by the user, when transporting a large object with a robot, was lowered by re-configuring the orientation of the object and the pose of the robot endpoint based on an ergonomic assessment using human pose and the interaction force profiles during the task.

While human-centric tasks have been successfully integrated into pHRI applications, examples of personalising the model for individual users are rare. To capture the noise and variance exhibited by each individual, human-centric robotics researchers have looked towards the integration of conventional methods to address these uncertainties.

2.2 Probabilistic Models

For humans, signal noise is tied into our abstract understanding of everything surrounding us. However, for robotic systems, noise is observed at all times, whether it is due to the

sensitivity of the sensor, temporary disturbances, or even signal transfer across communication cables. Signal noise and uncertainties are even more pronounced in humans where it presents in both physical and neurological systems. As you read this sentence, the signals sent to your optic nerves inherit noise from your cornea. This signal generates noise as it travels to the powerful CNS which then processes these signals, resulting in the abstract concepts of vision and comprehension.

Probabilistic models are a simple method to capture this noise, building a distribution which describes the characteristics of the noise. This is evident in recent analysis by [95] suggesting that the noise profile exhibited by humans during physical movements are bio-indicators for motor skill level.

Underlying the methods to capture the characteristics of the observed noise is the assumption that the observed system conforms to the Central Limit Theorem (CLT). This theorem suggests that when observations of IID random variables are recorded, the eventual normalised sum of the mean value approaches towards a normal distribution. This enables the characterisation of the observed noise, and can be factored into subsequent analysis.

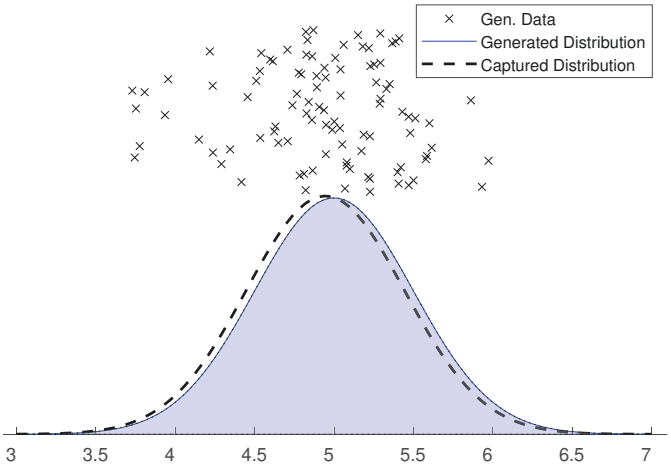

FIGURE 2.4: An example probabilistic distribution capturing inherent noise. Data points were generated using a chosen distribution (in blue), and these are captured by the other distribution in black. Note the slight distribution difference due to the noise.

With the omnipresence of noise in robotic applications, it is natural that researchers have looked to using probabilistic approaches for robotics [96]. For pHRI, a similar principle was used to characterise the noise observed for each individual user. Early works recognised the variability exhibited from one person to another, developing models based on information theory [97] to elucidate the underlying internal models for various tasks. To

overcome idiosyncrasies for each individual, early works recruited hundreds of participants to "average out" these variances.

2.2.1 Fitts' Law

One of the earliest models developed was Fitts' law [98], a predictive model characterising target-directed human movements. The model is defined as:

$$MT = a + b \cdot \log_2 \left(\frac{2D}{W} \right), \tag{2.1}$$

where MT is the movement time taken to reach a given target, D is the distance from the target, and W is the width of the target. a and b are coefficients that are obtained heuristically through linear regression for each person. Fitts' law suggests that there is a relationship in human movement that depends on D and W.

From this model, an Index of Difficulty (I_d), measured in bits, was able to be obtained, indicating the amount of information associated with a task:

$$I_d = \log_2 \left(\frac{2D}{W} \right). \tag{2.2}$$

Further research for Fitts' law have resulted in some consensus on the usability and efficacy of the model when predicting human behaviour. Empirical evidence reinforces the model's applicability, demonstrating that humans can use visual information (with Fitts' law parameters) to predict the observed behaviour [99], to exhibit kinematic patterns based on the parameters of the model [100], and to produce comparable motions when compared against simulated responses from a motor control filter [101].

While the applicability of Fitts' law seems to be generally accepted, in-depth analyses from those experiments suggest that the model is a symptom of the underlying visuomotor system. This dispute is evident in more recent works, highlighting the constraints of the model in allocentric interactions [102] and physical interactions [103].

As our understanding of the CNS increases, the knowledge on the human motor system has followed suit, deprecating the use of information theory. A similar phenomenon exists in robotics, where the increasing capability for information transfer between systems necessitated a macroscopic approach towards the environment surrounding the robot [104].

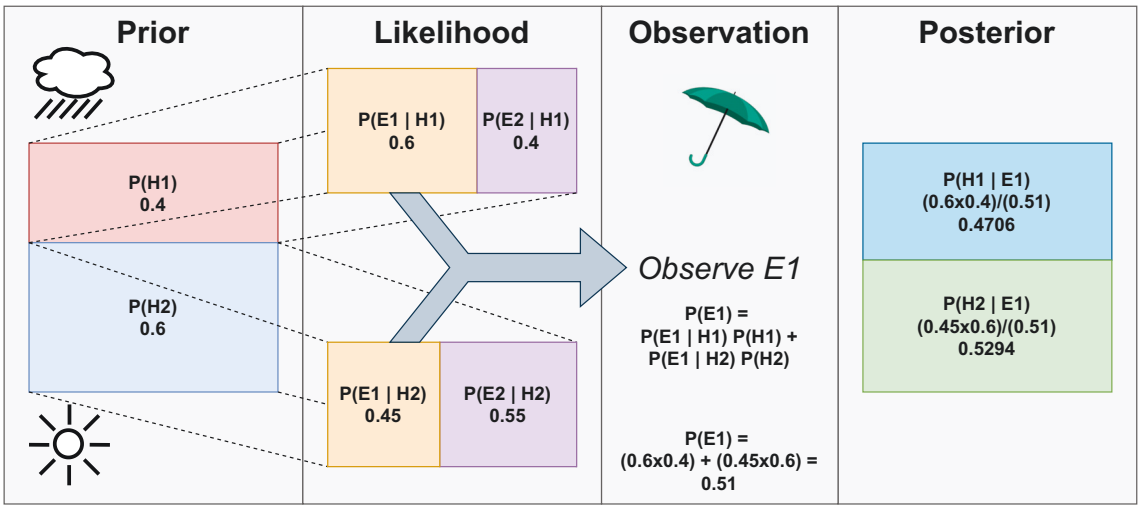

FIGURE 2.5: A visual representation of Bayes' rule using an example of determining whether it is sunny or raining given an observation that someone has an umbrella.

Rather than treating observations as information transfer, researchers shifted focus towards memory-efficient methods to capture information surrounding the agent. These approaches rely on an assumed internal model, which is then updated when new information is added. Termed *Bayesian Statistics*, these methods express a level of belief for each event, which is then updated when new observations are obtained.

2.2.2 Bayesian Statistics

Bayesian Statistics is a field of research which represents probabilities as the level of belief for an event to occur. By obtaining new information, the belief is then updated based on the confidence and usefulness of the new information, providing a quantifiable method to work with probabilistic distributions. Bayesian statistics underpins commonly used algorithms such as the Kalman Filter and Extended Kalman Filter (EKF), with complex problems leveraging these methods such as in Simultaneous Localisation and Mapping (SLAM) [105] and motion tracking [106].

The basis of Bayesian statistics is Bayes' rule which defines, for two *independent* events x and y, the probability of event y occurring, given event x occurs, is:

$$p(y \mid x) = \frac{p(x \mid y)p(y)}{p(x)}, \tag{2.3}$$

where $p(.)$ is the probability distribution of an event. A simple example of Bayes' rule is to determine whether it is raining if we *observe* someone with an umbrella. This is visually represented in Figure 2.5.

Initially, our belief is that it rains 40% of the time (our *prior*). The *likelihood* of a person bringing an umbrella will be different depending on the current weather. If it is raining, there is a 60% chance of bringing an umbrella. Conversely, if it is not raining, there is a 45% chance of bringing an umbrella. We then apply Bayes' rule to determine that, given an observation of someone with an umbrella, there is a 47.1% chance that it is currently raining, reiterating that Bayes' rule assumes the two events are independent of each other.

Rather than discrete probability values, a probability distribution displays the spread of probabilities which best represent a phenomenon. The most renowned family of probabilities is the Gaussian distribution, due to its symmetry and low number of parameters, μ and Σ. Furthermore, the symmetry provides many other properties which can be exploited in probabilistic operations. The probability of an event, x, occurring in a Gaussian (normal) distribution is defined as:

$$\mathcal{N}(x) = \frac{1}{\sqrt{2\pi\Sigma}} \exp\left(-\frac{(x-\mu)^2}{2\sqrt{\Sigma}}\right). \tag{2.4}$$

For K samples/observations, the *prior* distribution is defined by the mean and variance of the input:

$$p(x) = \mathcal{N}(x|\mu_x, \Sigma_x), \tag{2.5}$$

$$\mu_x = \frac{\sum_{i=1}^{K} x_i}{K}, \tag{2.6}$$

$$\Sigma_x = \frac{\sum_{i=1}^{K} (x_i - \mu_x)^T (x_i - \mu_x)}{K}. \tag{2.7}$$

Adding additional variables in a Gaussian distribution extends the distribution. For example, adding another variable, y, creates the distribution $\mathcal{N}(X|\mu_X, \Sigma_X)$:

$$X = \begin{bmatrix} x \\ y \end{bmatrix}, \tag{2.8}$$

$$\mu_X = \begin{bmatrix} \mu_x \\ \mu_y \end{bmatrix}, \tag{2.9}$$

$$\Sigma_X = \begin{bmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{yx} & \Sigma_{yy} \end{bmatrix}, \tag{2.10}$$

where Σ_{xy} is the covariance between x and y. We note that the symmetry in a Gaussian distribution results in the covariance matrices between x and y to be transposes of each other: $\Sigma_{xy} = \Sigma_{yx}{}^T$.

Thus, using Bayes' rule, we can update the probabilistic belief of an event, y, occurring when given an observation of x. Leveraging additional properties of the Gaussian distribution, the conditional distribution can be obtained using Gaussian conditioning:

$$\mu_{y|x} = \mu_y + \Sigma_{xy}\Sigma_{xx}^{-1}(x - \mu_x), \tag{2.11}$$

$$\Sigma_{y|x} = \Sigma_y - \Sigma_{xy}\Sigma_{xx}^{-1}\Sigma_{yx}. \tag{2.12}$$

The full derivation can be found in [107].

Apart from probabilistic operators, there are also information-theoretic measures to compare between different distributions. One prominent measure is the Kullback-Liebler Divergence [108], D_{KL}, which compares the relative entropy, or the amount of information gained, when comparing one distribution to a reference distribution. Properties from the specific distributions, such as Gaussian, Poisson, or Binomial, can then be leveraged to perform more empirical methods such as moment matching.

2.2.3 Machine Learning

Machine Learning (ML) applies probabilistic models to capture the relationship between variables and states, whether from external sensors or internal controllers. During pHRI, there is a plethora of information available to ascertain the current state of the system. However, utilising all available data might not be conducive to effective analysis and

application. Thus, ML provides an approach to systematically process and extract meaning from the data presented.

Spurred on by early human learning theories, ML leveraged early works in rough set theory to find generalised approaches to capture the relationship of inputs as a function of the outputs [109]. There are two distinct forms of output predictions in ML frameworks:

- Classification: the outputs are discrete class labels which correlate to the input sample. Classification is commonly used to label objects in a given scene making it useful for surveillance and identification tasks.
- Regression: the output is continuous in nature, predicting the expected value at the next step. Regression is commonly used to predict trends in economic fields.

Furthermore, the algorithms and methods can be sorted into four broad categories based on the nature of the training data:

- Supervised Learning: When all training data is labelled, providing explicit one-to-one data correlation for the generated model.

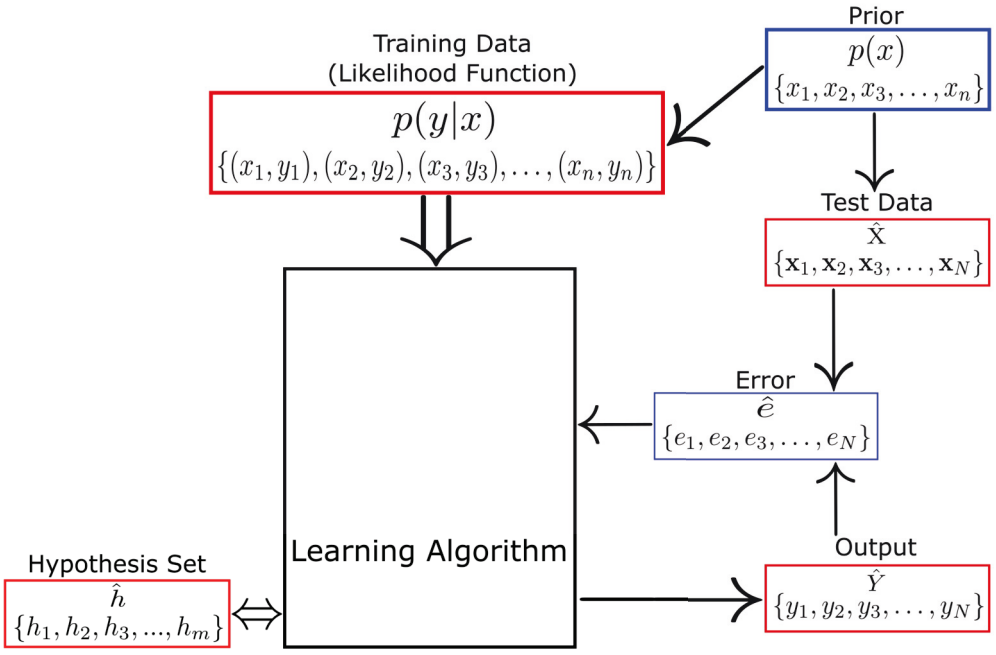

FIGURE 2.6: A general flow overview of a Machine Learning framework.

- Semi-supervised Learning: When some training data is labelled, requiring some method to correlate unlabelled data with the current model.
- Unsupervised Learning: When all training data is unlabelled, relying on topological and application constraints to find correlation between the training data for model generation and discrimination.
- Reinforcement Learning: When given no training data, an agent builds a model by exploring its environment unsupervised. The defined reward and value function provides a metric of its progress thus coupling the actions to the outcomes.

The wide-ranging capability for ML has seen its application in most, if not all, areas of our everyday lives, ranging from basic models like Naïve Bayes [110] to reinforcement learning for autonomous robotic control [111]. Due to the capability for ML to capture inherent differences from one user to the next, there is a large demand in researching methods to utilise probabilistic models for personalisation. This is particularly apparent in physiotherapy, where the prescribed exercise regimen will differ based on the circumstances of the patient, as well as the diagnosis which may vary from one expert to the next.

For rehabilitation, the confluence of ML and robotics has lead to the current trend for robotic assistance during therapy, with successful Randomised Control Trials (RCTs) conducted to validate its efficacy [112]. This has lead to a surge in research focusing on rehabilitation robotics, with applications such as optimising the design of exo-skeletons [113], personalising the motor learning task [114], and adapting the targeted assistance provided as the patient improves over time [115].

The integration of ML in robotics is longstanding with applications in interaction robots [116] and assistive robots [86]. Furthermore, the need to understand the human user during HRI and pHRI has lead to the use of ML to understand human intent and detect actions which convey the state of the system to the robot [117]. These applications are detailed in Section 2.3.1.

As the amount of available data increases exponentially, new research has searched for ways to optimise the complexity of ML techniques, thus widening the scope of ML applications. Reinforcement Learning aims to mimic the learning process humans use when given an unprecedented task [118] by adapting the control policy in each "epoch" and calculating the value of that particular set of parameters. Reinforcement learning has been heavily used in recent works. The advantage is that it mitigates subconscious bias induced from assumptions made when choosing a model (in traditional ML).

Deep Learning is a subset of ML which uses hierarchical extraction methods to encapsulate higher-dimensional data in a non-tractable model [119]. The bio-mimetic network is inspired by the neural structure of human brains which allow complex visual understanding, leading to the analogy that the development of deep learning models is akin to humans building a mental model of an object. By taking snapshots from different angles, we co-ordinate the spatial information based on the different viewpoints. Deep learning aims to build the "perceptual manifold", leading to the development of frameworks that tackle high dimensional data with high levels of accuracy [120].

The contribution of deep learning for robotics cannot be underestimated with significant successes presented in computer vision. Notable frameworks include LeNet-05 [121], YOLO [122], and DeepPose [123]. The reach of deep learning is also quite broad with recent works focusing on applying these techniques for the health care industry [124], and expanding upon reinforcement learning [125] for complex robotic control [126] and pHRI [127].

2.2.4 Gaussian Process

Gaussian Process (GP) is a non-parametric family of distributions which can be used to represent priors for a target function [128]. Based on the normal Probability Distribution Function (PDF), they generalise the target function, taking into account an infinite set of hypotheses, with a stochastic process defining the properties of the function. This is similar to how Hidden Markov Model (HMM) exploits the lack of model assumptions to identify the state transition function.

The reliance on the properties of the Gaussian distribution means that it is analytically tractable and is flexible. This makes it one of the most common methods to build a probabilistic model. The flexibility of GP has lead to their use in various fields such as deformable object tracking [129], robot control [130], and human motion tracking [131].

The conversion of single GP models to ensemble models is straightforward with the Gaussian Mixture Model (GMM) [132] and other ensemble variants, known as product of experts. During inference, the output accuracy is determined by the fusion method from multiple experts. Given that each member GP provides a mean and uncertainty, the exact analytical solution for uncertainty propagation in GP-based product of experts is attainable [133]. Applying the concept of hierarchy and scalability into GP has also been successful in literature [134, 135].

One aspect of GP which is extensively explored is its computational complexity. With a computational complexity of $\mathcal{N}(\mathcal{O}^3)$, GP models are not scalable with respect to the number of training samples. Furthermore, the complexity of the training samples can decelerate the optimisation of kernel hyper-parameters, leading researchers to improve computational costs of GP-based algorithms.

Apart from the number of training samples, the computational complexity also makes GPs ill-conditioned for high-dimensional data. Thus, variants of GP were created to lower the dimensionality of the processed data. These are manifested in algorithms such as Sparse GP [136], which lowers the dimensionality in a greedy fashion using brute force; Informative Vector Machine (IVM) [137], utilising information theoretic criterion to optimise the computational complexity; and Gaussian Process Latent Variable Model (GPLVM) [138], which leverages Principal Component Analysis (PCA) to embed latent variables representing features of the high-dimensional data.

2.2.4.1 GP Regression

Overall, a non-linear target function f is assumed as:

$$\mathbf{Y} = \mathrm{f}(\mathbf{X}) + \boldsymbol{\epsilon}, \tag{2.13}$$

where $\mathbf{Y}$ is the training targets, $\mathbf{X}$ is the training inputs, and $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \Sigma_\epsilon)$ is the observation noise distributed normally with $\Sigma_\epsilon = \mathrm{diag}\{\sigma_1^2, \sigma_2^2, \ldots, \sigma_{d_y}^2\}$.

For more than 1 dimension of data, learning the inductive relationships from training data is done by fitting a GP with the prior for each dimension of y as follows:

$$\boldsymbol{y}_a \sim \mathcal{GP}(0, K_a), \tag{2.14}$$

for $a = 1, 2, 3, \ldots, d_y$. The observation noise is $\epsilon_a \sim \mathcal{N}(0, \sigma_a^2)$ and K_a is the covariance for dimension a.

2.2.4.2 Covariance Kernel

Part of the assumptions made in GP relate to how the boundaries of the target function are defined. The covariance kernel, defined by the covariance function, provides a mathematical model to calculate how close each data sample is to another. Different covariance

TABLE 2.1: A table of common covariance functions used for GP model generation. The hyper-parameters change based on the function's variables.

Covariance Function	Expression
Constant	σ_0^2
Linear	$\sum_{d=1}^{D} \sigma_d^2 \mathrm{x}_d \mathrm{x}'_d$
Polynomial	$(x \cdot x' + \sigma_0^2)^p$
Squared Exponential (SE)	$\exp\left(-\|\|\mathrm{x} - \mathrm{x}'\|\|^2 / 2l^2\right)$
Matérn [139]	$\frac{1}{2^{\nu-1}\Gamma(\nu)} \left(\frac{\sqrt{2\nu}}{l} r\right)^{\nu} K_\nu \left(\frac{\sqrt{2\nu}}{l} r\right)$
Exponential	$\exp\left(-\frac{r}{l}\right)$

functions will have different properties, lending each function well to certain applications. This choice can affect the results of the trained GP model, thus, user expertise is necessary here to choose a suitable covariance function.

The covariance kernel determines the level of similarity based on its definition of distance. Using this measure, two data samples close to each other are said to be *similar*, since intuitively a training sample right next to another would provide some information about the predictions around those two samples.

While there are infinite possible covariance functions to be used, the only stipulation for a valid covariance function is that it is positive semi-definite. The most common covariance functions are based on information theory and a list of common covariance functions is listed in Table 2.1.

Using the most commonly used kernel function, the SE covariance function, the covariance element k_a for each dimension's covariance, K_a (Equation 2.14), is as follows:

$$k_a(\boldsymbol{x}, \boldsymbol{x}') = \sigma_a^2 \exp\left(-\frac{1}{2l_a^2} ||\boldsymbol{x} - \boldsymbol{x}'||^2\right), \tag{2.15}$$

where σ_a^2 is the variance of $\boldsymbol{f}_a$, and l_a is the scaling factor to normalise the distance between x and x'.

2.2.4.3 Hyper-parameter Optimisation

While arbitrary hyper-parameter values can produce a GP model, the training data set provides examples that are used to optimise the hyper-parameter values to provide a model

which best approximates the target function. The choice of optimisation method and their relevant parameters will affect the resultant hyper-parameter values and its effect on the GP model. In particular, over-fitting or under-fitting the model to the training samples can degrade the performance of the built GP model.

Each optimisation method has unique properties and their parameters should be considered carefully when optimising the covariance function hyper-parameters. Common optimisation methods include Bayesian Optimisation [140], Conjugate Gradient Descent [141], Nelder-Mead Simplex [142], and Differential Evolution [143].

For the continuing example of the SE covariance function, the Maximum A-Posteriori (MAP) estimate of the kernel function hyper-parameter set, θ_a, occurs where $p(\mathrm{Y}_a|\mathrm{X}, \theta_a)$ is greatest. Thus, θ_a is optimised as:

$$\log p(\boldsymbol{Y}_a|\boldsymbol{X}, \theta_a) = -\frac{1}{2}\mathcal{D} - \frac{1}{2}\log|K_a| - \frac{N}{2}\log 2\pi, \tag{2.16}$$

where $\mathcal{D} = {\boldsymbol{Y}_a}^T K_a^{-1} \boldsymbol{Y}_a$, and N is the number of training data samples.

For regression inference of Gaussian Process, the joint probability $P(\boldsymbol{Y}_a, {\boldsymbol{y}_a^*}^T)$, inferring $\boldsymbol{y}_a^*$ at $\boldsymbol{x}^*$, is modelled as a multi-variate Gaussian distribution, similar to Equation 2.8:

$$P\left(\begin{bmatrix}\boldsymbol{Y}_a \\ \boldsymbol{y}_a^*\end{bmatrix}\right) = \mathcal{N}\left(0, \begin{bmatrix}K_a & K_{a*}^T \\ K_{a*} & K_{a**}\end{bmatrix}\right) \tag{2.17}$$

where $K_{a*} = [k_a(\boldsymbol{x}^*, \boldsymbol{X}_1), k_a(\boldsymbol{x}^*, \boldsymbol{X}_2), ..., k_a(\boldsymbol{x}^*, \boldsymbol{X}_N)]$ and $K_{a**} = k_a(\boldsymbol{x}^*, \boldsymbol{x}^*)$, respectively. The conditional distribution, $p(\boldsymbol{y}_a^*|\boldsymbol{Y}_a) = \mathcal{N}(K_{a*}K_a^{-1}\boldsymbol{Y}_a,\, K_{a**} - K_* K^{-1} K_{a*}^T)$ is then derived from this multivariate Gaussian.

2.2.4.4 GP Classification

The joint probability distribution from the GP model for regression provides individual posteriors for each label. However, the posterior results are analytically intractable since the joint distribution, $p(\boldsymbol{Y}, \boldsymbol{X})$, is non-Gaussian (due to the MAP estimate), making it impossible to be used for GP classification.

A way to re-instate tractability is through approximations or Monte Carlo sampling methods. While there are a plethora of approximation methods available, a popular method for binary GP classification is the Laplace Approximation [144].

Laplace Approximation takes a second-order Taylor Series expansion around the MAP, providing the following approximation:

$$q(\boldsymbol{f}|\boldsymbol{X},\boldsymbol{y}) = \mathcal{N}(\boldsymbol{f}|\hat{\boldsymbol{f}}, A^{-1}), \tag{2.18}$$

where $\hat{\boldsymbol{f}} = \text{argmax}_f\ p(\boldsymbol{f}|\boldsymbol{X},\boldsymbol{y})$, and A is the Hessian of the negative log posterior at that point $(-\nabla\nabla \text{log}p(\boldsymbol{f}|\boldsymbol{X},\boldsymbol{y})|_{\boldsymbol{f}=\hat{\boldsymbol{f}}})$.

For multi-class classification, the Laplace Approximation needs to be extended. A simple method is to obtain the softmax output $\boldsymbol{\pi}$ from training samples (Equation 2.19). The class probabilities are calculated in a one-vs.-rest fashion:

$$\boldsymbol{\pi}_i^c = \frac{\text{exp}y_c^i}{\sum_{c'} \text{exp}y_{c'}^i}, \tag{2.19}$$

for $i = 1, 2, \ldots, n$ given n training samples, and $c = 1, 2, \ldots, C$ for C classes.

Similar to the regression inference of GP, calculating the approximated joint posterior distribution will yield the predictive mean for class c as follows:

$$\mathbb{E}_q[f^c(x_*)|X,\boldsymbol{y},x_*] = k_c(x_*)^T K_c^{-1}\hat{\boldsymbol{f}}^c, \tag{2.20}$$

where $k_c(x_*)$ is the covariances between the test point and each training point of the cth covariance function, and $\hat{\boldsymbol{f}}^c$ is the subvector of $\hat{\boldsymbol{f}}$ in class c.

Unlike binary classification, a simple mean threshold does not take variability between classes into account. The most common way to overcome this is through Monte Carlo sampling of the Gaussian Laplace Approximation $q(\boldsymbol{f}_*|y)$, performing the softmax, and averaging the probabilities.

2.3 Learning from Demonstration for HRI

Originally, Learning from Demonstration (LfD) frameworks were developed to mitigate the amount of bespoke programming required when the task allocated to the robot changes. By observing an expert demonstration, through kinesthetic teaching or from external sensors, LfD frameworks enabled robots to repeat the same motion, providing a relatively quick method for robots to re-learn different motions without the financial or time costs involved. However, as humans start to encroach on the robot's workspace, there is a need to detect, understand, and respond to the user's intent and actions [145].

Implementing LfD for HRI and pHRI applications requires the system to detect user actions, contextualise those actions to the task at hand, and generate an appropriate robot response to assist the user. While each of these steps can be tackled individually, probabilistic methods have been used to couple all three steps together for HRI and pHRI.

2.3.1 Action Detection and Contextualisation

In a human-robot system, the robot needs to detect and understand the user's intent for any given task. For assembly tasks, this involves understanding the current state of the assembly as well as the possible next steps. Thus, probabilistic graph-based models are popular since they provide an awareness of both the different stages along the assembly, as well as all combinations for the next steps [146–148].

The confluence of ML for human action detection during HRI and pHRI is not unexpected since ML is able to capture individual mannerisms, with probabilistic operators which can couple multiple dimensions together. This effect can be seen in the Partially Observable Markov Decision Process (POMDP), where the contextual state of the task is linked to the observed actions from the user, facilitating HRI [149], wheelchair assistance [150, 151], and goal and trajectory adaptation during pHRI [13].

A similar framework is HMM [152] which attempts to "learn" a hidden model that explains the relationship between the observed actions and the state of the system. Its effectiveness in capturing the dynamical system for HRI applications has lead to its popularity in action detection [153] and intent estimation [154].

Efforts to improve its performance succeeded by applying techniques used in ML for the Hierarchical HMM which was used for detecting human actions from movement trajectories [155]. Other methods for detecting user actions include Dynamic Bayesian Network (DBN) [156], deep learning [157], GP Dynamical Models [131, 158], and stochastic parsing [159].

2.3.2 Movement Primitives

Generating a robotic trajectory is a complex task, especially when coordinating multiple Degrees of Freedoms (DoFs). This is further exacerbated when the robotic trajectory needs to interact with a human in a meaningful manner during HRI. The non-linear response, when a human interacts with the robot, has long been explored for LfD in both *one-shot* learning [160] and probabilistic approaches [161]. The resultant non-linearity in

the dynamical systems create complexities which complicate the interactions between the agents during HRI. Therefore, a more holistic approach is needed to simplify the trajectory generation process.

Movement Primitives (MPs) are a family of elementary operations to represent robotic trajectories for motion planning. Inspired by human locomotion [162], MPs are compact representations of complex locomotion in multi-DoF systems, forming the basis for various robotic capabilities such as learning, imitation, and trajectory generalisation. Since MPs are adapted based on their applications, they come in various forms, each alleviating certain limitations or constraints in their respective applications. Some prominent MPs and their applications include:

- Coupling Movement Primitives [163] for achieving equal interaction forces;
- Compliant Movement Primitives [164] for embedding torque profiles;
- Interaction DMP [165] for reaching equilibrium trajectories using interaction forces;
- Kernelized Movement Primitives [166] for higher compression of trajectories; and
- Periodic Movement Primitives [167] for periodic motions.

Despite the variety of available MPs, the most popular MPs used in the robotics community is Dynamic Movement Primitives (DMP) due to its versatility and track record. There is also evidence to suggest that human motions can be represented by DMP [168, 169], reinforcing its use in HRI and pHRI applications.

2.3.3 Dynamic Movement Primitives

Dynamic Movement Primitives (DMP) generate globally stable trajectories by treating the trajectory as a spring-damper system with an attractor system to encode non-linear dynamics [170]. The attractor landscape is represented by a linear system of basis functions across time or phase. This allows DMP to model the motions of expert demonstrations using only a few parameters and favourable properties such as trajectory dilation, rotational invariance, and temporal scaling [171].

While the original DMP formulation produced stable trajectories, a modified formulation overcame several undesired artefacts during edge case reproductions such as trajectory

"mirroring" and accelerations that are beyond the capabilities of the robot [172]. Each DoF in the system is modelled as:

$$\begin{aligned} \tau\dot{v} &= K(g-x) - Dv - sK(g-x_0) + sKf, \\ \tau\dot{x} &= v, \end{aligned} \tag{2.21}$$

where g is the goal, x is the position, x_0 is the starting position, v is the velocity, and K and D are the stiffness and damping of the system respectively.

The phase of the trajectory, s, is modelled as a first-order system with parameter α for temporal scaling, $\tau\dot{s} = -\alpha s$. The attractor landscape, represented by the forcing function, f, is encoded using the weighted sum of Gaussian basis functions with centres, $\boldsymbol{c}$, and is spread evenly across the phase of the trajectory:

$$\begin{aligned} f(s) &= \frac{\sum_{i=1}^{M} \psi_i(s)\omega_i s}{\sum_{i=1}^{M} \psi_i(s)} = \phi(s)^T \boldsymbol{\omega}, \\ \psi_i(s) &= \exp(-(s-c_i)^2/h), \end{aligned} \tag{2.22}$$

producing a $M \times 1$ vector of weights, $\boldsymbol{\omega}$, for the trajectory. While the Gaussian kernel has been used here, any smooth functions or mollifiers could be used as an alternative [173].

Representing the forcing function (from T observations) as a linear system allows the DMP weights, $\boldsymbol{\omega}$, to be obtained using linear least squares regression:

$$\boldsymbol{f} = \begin{bmatrix} f(s_1) \\ f(s_2) \\ \vdots \\ f(s_T) \end{bmatrix} = \begin{bmatrix} \phi_1(s_1) & \dots & \phi_M(s_1) \\ \phi_1(s_2) & \dots & \phi_M(s_2) \\ \vdots & \ddots & \vdots \\ \phi_1(s_T) & \dots & \phi_M(s_T) \end{bmatrix} \begin{bmatrix} \omega_1 \\ \omega_2 \\ \vdots \\ \omega_M \end{bmatrix}, \tag{2.23}$$

$$\boldsymbol{\omega} = (\boldsymbol{\phi}^T\boldsymbol{\phi})^{-1}\boldsymbol{\phi}^T\boldsymbol{f}. \tag{2.24}$$

Since the formulation of DMP only encodes a single trajectory, their use is limited in applications with dynamic environments such as during HRI. Thus, variations of DMP were developed to tackle various aspects of these limitations. A popular goal to overcome for DMP is avoiding obstacles in the path of the trajectory. The general approach is to append another attractor landscape as follows:

$$\tau\dot{v} = K(g-x) - Dv - sK(g-x_0) + sKf + h(s,x), \tag{2.25}$$

where $h(s, x)$ is a phase-aligned function describing the relationship between the trajectory and obstacle. This method has seen successful obstacle avoidance when planning [174, 175] and performing [176] the trajectories.

A similar approach to amend the non-linear function, $f(s)$, has been used to accomplish various aspects of the robot response in DMP. This includes weighting the Gaussian basis functions to embed the "style" of the demonstration [177, 178], appending a function, $I(s, \theta)$, to correlate the impedance profiles for an exoskeleton [179], deriving the attractor landscape function based on muscle synergies [180], and appending a variable to distinguish from one task to another [181].

Given that MPs were originally used as a method to encode a single trajectory, it is natural that hierarchical frameworks have been designed to complement this intention. Building a library using DMP provides flexibility, trading-off against complexity, for higher-level actions such as mobile robot motion planning [182]. An identical approach was employed by [183] to contextualise the sequence of motions required for object grasping.

There are also other schools of thought on the use of DMP (and MPs in general) libraries. Integrating probabilistic modelling into an observed trajectory has allowed [184] to perform both trajectory segmentation and DMP library building simultaneously. The ability to generate a library on-line enables robust generalisation for future demonstrations and promotes flexibility when generating novel trajectories. Similarly, [185] segments an interaction force-based object flipping task to create a library of MPs embedded with interaction-force profiles.

2.3.4 Movement Primitives for HRI

Initially, Movement Primitives were designed for a single expert demonstration, with the objective of imitating or repeating the same trajectory. However, in HRI, utilising a single user-directed demonstration can unintentionally introduce noise which degrades the performance of the system in its application. Thus, there was a need to work towards the analysis of multiple expert demonstrations to obtain a single set, or distribution, of parameters which best generalises the intended trajectory. A visual overview of this process is demonstrated in Figure 2.7.

Probabilistic operators are the most attractive methods to both generalise and capture noise due to their well-studied properties in the ML field. Many algorithms leveraging off

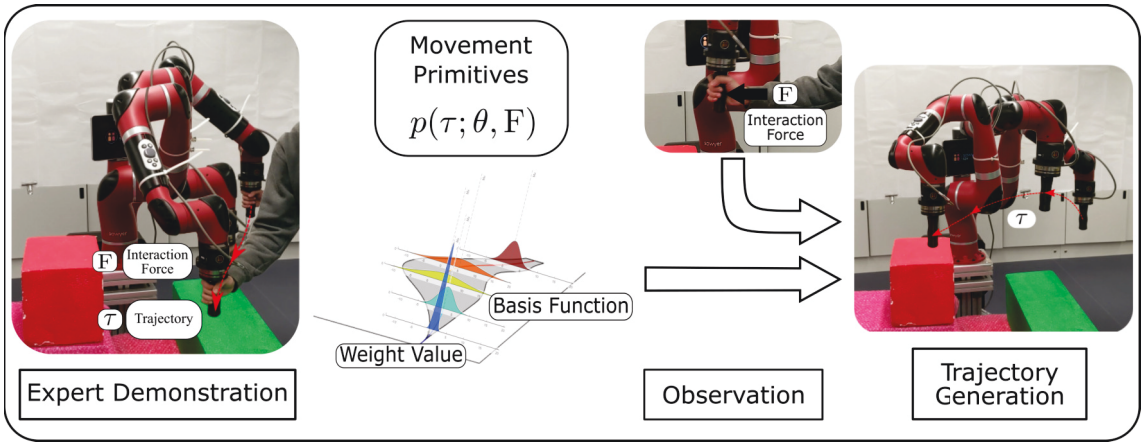

FIGURE 2.7: Overview of learning Movement Primitives for pHRI in the Bayesian context.

probabilistic operators provide tractable models, allowing simple techniques to embed the application's context into the model.

For MPs, this is most commonly used to couple multiple DoF together in a multi-DoF system, whether they are observations of the robot or the human. A prominent exemplar of this for HRI is Interaction Primitives (IPs) which has been shown in [186] to couples all dimensions of data from both the robot and human user when performing a *high-five* task. Extensions of IPs have embedded environmental context when performing a collaborative object covering task [187]. Similar methods to embed contextual information, employing probabilistic operators against parameters from MPs, include using a modified EKF to estimate the DMP endpoint (g) and time scale parameter (α) [188, 189], an Expectation Maximisation (EM)-GMM to perform an object handover task [190], and an EKF to embed contextual data from multi-modal sensors in Bayesian Interaction Primitives (BIPs) [191].

For collaborative pHRI, aligning the interaction forces with the objectives is a difficult task. Since DMP are acceleration-based dynamical systems, applying them to other data types is uncommon. Despite this, there have been some notable works that looked to imitate the impedance profile for the robot response. A naïve approach is to use a separate, amended DMP to encode the stiffness [192], although there have been other ways that varied the goal, trajectory, and speed of the robot response [193]. Alternative literature highlight ML frameworks such as HMM to couple trajectories with their respective stiffness profiles [194], and extending BIPs to integrate contact forces for phase estimate contextualisation during trajectory generation [195].

While most DMP-based frameworks build probabilistic distributions of DMP parameters, Probabilistic Movement Primitives (ProMPs) frameworks inherently encode positions and

velocities in their distribution [196]. This enables straightforward coupling of multiple DoF and the inclusion of contextual data when generating trajectories [197], e.g., inferring the user's intended instruction step during an assembly task [198].

Since DMP encodes the forcing function, $f(s)$, across the phase of the trajectory, this imposes a phase estimation requirement when attempting to infer the appropriate robot response during HRI. Traditionally, this is done using phase-alignment methods such as Dynamic Time Warping (DTW) [199], as used in IPs and its extensions.

However, probabilistic operators can also address this limitation by coupling the phase to the rest of the distribution, enabling phase estimation alongside parameter generation for robotic trajectory generation [200]. Similar approaches are used for BIPs where the phase is embedded in the EKF for parameter estimation [201]. Furthermore, probabilistic methods such as EM can also be used to phase-align observed trajectories asynchronously during tele-operation [58].

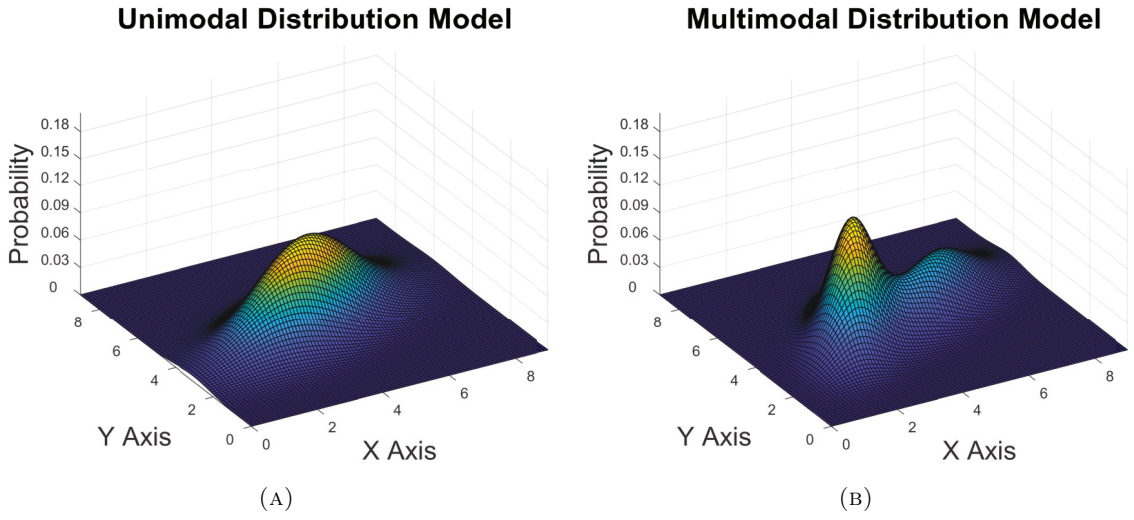

(A) (B)

FIGURE 2.8: A visual representation of a: (A) unimodal (single peak) and (B) multi-modal distribution (multiple peaks) using identical data generated from two different normal distributions.

One challenging aspect for MPs in LfD and HRI applications is the possibility of multi-modal applications. As mentioned above, MPs were originally designed for a single demonstration, creating a *unimodal* application since the primary objective is to imitate or replicate the demonstration. Multi-modal tasks, such as reaching for an object or choosing between multiple objects, can also benefit from the use of probabilistic methods, especially those which are designed for multi-modal applications. One of the most common

multi-modal models is the GMM [202], building a single distribution composed of multiple sub-populations, each corresponding to their respective *modes* of the data.

For MPs, this has been achieved using GMM to regress against the parameters of MPs, within a HMM framework, to imitate domestic feeding using robots [203]. Similar uses of GMM can be seen for task segmentation and trajectory generation for industrial object assembly during pHRI [56]. Similar techniques, such as Mixture Density Network (MDN), have been utilised to generate parameters of Via-points MPs, which are viable trajectories for multi-modal obstacle avoidance [204].

2.4 Musculoskeletal Models

Musculoskeletal models are a representative system of models which simulate the underlying mechanisms surrounding the musculoskeletal system of different organisms. Simple models can be easily designed with two bones linked with a pin joint, and controlled with a pair of antagonistic muscles.

Alternatively, sophisticated models can present the whole body structure with constrained joints, Muscle Tendon Units (MTUs) with representative intrinsic properties, and even interaction models such as the Hunt-Crossley contact model [205]. Despite the spectrum of models available, all musculoskeletal models aim to provide insight into the underlying mechanisms of the human neuro-musculoskeletal system [206, 207].

Musculoskeletal models enable physics-based simulations for various states of the model (such as position, velocity, and force/torque), using systems of differential equations. Each sub-system in a model can also utilise its own set of model equations to represent various contributing factors. By running simulations and human motion analysis, musculoskeletal models can be used to investigate human responses to various stimuli such as perturbations [208] or interaction constraints [209].

The rising popularity of musculoskeletal models in the study of human biomechanics and physiology have led to the creation of various software, facilitating the use of musculoskeletal models by other researchers in the community. Notable software includes OpenSim[1] [210] (based off the SIMM simulation engine [211]), AnyBody[2] [212], and Human Body Model[3] [213]. OpenSim is the most popular since the software is open-source.

[1] https://simtk.org/projects/opensim
[2] https://www.anybodytech.com/
[3] https://www.motekmedical.com/software/hbm/

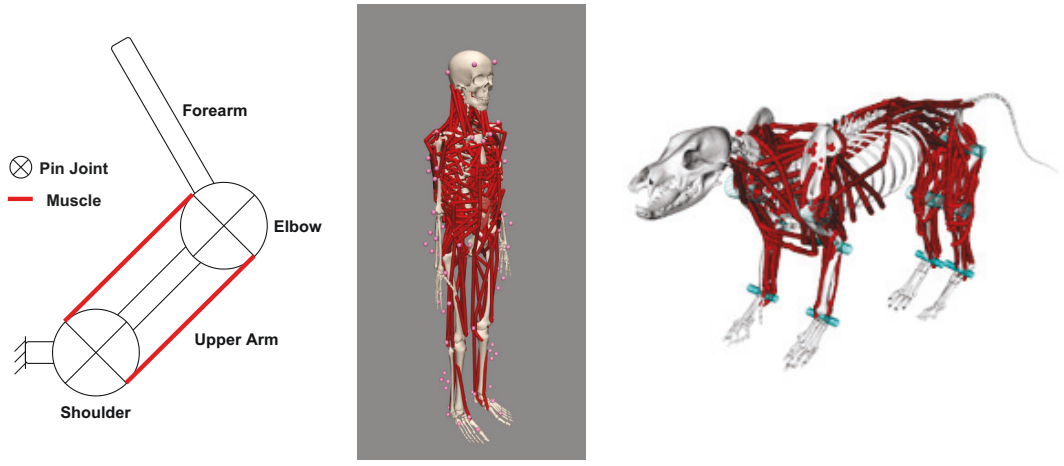

(A) (B) (C)

FIGURE 2.9: Musculoskeletal models can be: (A) simple; or (B-C) sophisticated depending on the application.

Across the body of literature, there have been many different musculoskeletal models which have been created, ranging from simplistic models to explore muscular control [214, 215] to more recent, complex whole-body models [216, 217]. Models of animals have also been created for researchers in those respective fields [218]. The application of OpenSim to musculoskeletal models was reviewed recently [219] with innovative applications for these models such as in motion analysis for osseointegrated amputees [220].

With the various software available, the comparison of results obtained from the different software is necessary. This creates a benchmark for researchers to understand the suitability of each software to their application. The benchmarking process includes the exploration of different software [221, 222], comparisons of various implementations of muscular optimisation and controllers [223–225], and verifying the obtained results using best practices [226, 227].

One significant challenge for musculoskeletal models is the personalisation of each model to individuals, using their respective model parameters. Changes in each model parameter affect the results from at least one stage along the motion analysis pipeline. For example, demographic data is regularly used to amend the bone measurements [228], which affects the results of *Inverse Kinematics*. Similarly, changing model parameters of MTUs [229], such as maximum isometric force and optimal tendon length, will affect the analysis derived from *Static Optimisation*, while inertial properties [230] will affect the resultant force and torque profiles obtained from *Inverse Dynamics*.

While models are generally validated using kinematic metrics such as generated moment arms and the range of motion, there are burdensome parameters that are difficult to obtain. Most available data for internal sub-systems, such as model parameters of MTUs, are obtained from preserved cadavers which does not reflect the "real-life" value. Other approaches build a correlation between extrinsic measurements and the internal parameters such as sEMG vs. Force [231], and sEMG onset vs. muscle activity [232].

Despite this, generic models built off demographic data [233] are still widely accepted since it is challenging to ethically obtain invasive, *in-vivo* measurements of model parameters. Similar problems are encountered when validating models using physiological signals such as sEMG that are influenced by various environmental factors when compared against invasive EMG measurements. As such, there are some reservations in the research community on the validity and consistency of simulation results for musculoskeletal models since there are large variances in the implementation and validation process [234].

Recent technological advances have enabled new, non-invasive methods to estimate some of the model parameters with promising results in estimating muscle volume using Magnetic Resonance Imaging (MRI) [235], muscle fibre composition using histological analysis [236], and muscle fibre orientation using deep learning on ultrasound images [237].

2.4.1 Human Motion Analysis

The primary intention for musculoskeletal models is to provide an *in-silico* solution to supplement and improve upon current methods for human motion analysis. In Section 2.3.1, the analysis of human motion is performed using task-centric measures without much insight from a neuro-scientific or physiological perspective.

Human-centric approaches towards human motion analysis are closely linked to a large body of literature focusing on human motor control [238]. Similar to the use of cutting edge technology for musculoskeletal models, advances in technologies to collect, aggregate, and analyse human motion data have promoted research in this field. Innovative methods include the use of marker-less motion capture systems [239, 240], calibration-less joint motion estimation algorithms [241], marker-less pose estimation algorithms [242], and the integration of Inertial Measurement Unit (IMU) data [243] for motion tracking. These techniques supplement or supersede prior methods that utilise marker-based motion capture frameworks [244, 245], as seen in Figure 2.10.

FIGURE 2.10: A representative setup for the use of a marker-based motion capture system to track the human upper limb.

Despite the lack of high-fidelity hardware available, seminal works in human motion analysis have relied on in-depth understanding of the CNS to develop an understanding of human motion. Following on from early works which developed models for *proximate causes*, the exploration of the underlying reason, or the *ultimate cause*, has led to a better understanding of human brain function. An example would be Fitts' law which employs information theory to explain and predict a human's capability. Further research has indicated that the model is a by-product of the way the human brain develops motor control skills, with [246] suggesting that the CNS utilises an intrinsic model of internal and external stimuli when planning and performing a motion.

One well regarded concept for human motor control is the role of MPs detected in the kinematic profile of human motion, suggesting that the human brain employs elementary operations to generate complex trajectories [162, 214, 247, 248]. This has inspired similar approaches for robot trajectory generation using MPs as discussed earlier in Section 2.3.2. Similar behaviour is observed using functional PCA to segment elementary non-linear time-dependent functions used for complex trajectory generation [73].

This has also led to more recent works inspecting the relationship between motion noise and motor learning [249], extending to exploring how humans regulate our limbs to support the motions [250, 251]. Further work was conducted to integrate insights from a physiological perspective into the understanding of motor learning and control, harking back to seminal works to identify the underlying principles which dictate human force and motion generation [252].

A significant portion of understanding the musculoskeletal system stems from work which explore the MTU and its influence on human motor control. These range from early

models for EMG-force relationships [253] to comprehending the core mechanics for muscle force generation [254]. Other perspectives include the elucidation on muscle force output based on optimality principles [255]. However, there is continued debate on its validity [256] with conflicting evidence for [257] and against [258, 259] its existence despite a lack of prominent alternative theories.

2.4.2 Applications

Overall, musculoskeletal models have pioneered new methods and avenues for neurological and bio-mechanical investigations into the human body (and other organic matters). The availability of open-source software, which has continued to be maintained, facilitates the adoption of these models in future research. Regardless of their limitations, the functional utility of musculoskeletal models has been proven in recent years with applications for understanding muscular dysfunction [260] and muscular impairment [261].

Musculoskeletal models have yielded insights that have developed end-user applications for strength estimation [262], impairment estimation [263], and even performing a whipping action [264]. Furthermore, they have inspired the development of more sophisticated processes which fit within the motion analysis pipeline, such as Computed Muscle Control [265] and Residual Reduction Algorithm [266], and other model-based estimations of muscle forces [267, 268].

Musculoskeletal models have also been used as a vehicle for discussion in the community for disputed topics, such as the functional existence of muscle synergies. Advocates for muscle synergies [269] suggest that the concept, where groups of muscles activate in particular pre-determined sequences, improves the efficiency of information transfer for the CNS. On the other hand, dissidents [270] argue that the available evidence is insufficient to establish a causal relationship. Recently, a synergy-based musculoskeletal model [271] has provided computational evidence for the integration of muscle synergies to generate target-reaching movements.

2.5 Summary

The ever-increasing presence of robotic systems in every facet of human life has been shown in this review of related works. The exponential growth in recent years is empowered by new hardware and frameworks to enable safe HRI. Despite significant improvements in *autonomous* frameworks, there are still many circumstances where human intervention is required. In these circumstances, *semi-autonomous* systems need to have a contextual understanding of the human user.

Current research trends treat the human user as a set of constrained systems, making assumptions for human behaviour during HRI. These assumptions are easily violated during pHRI, due to human autonomy, leading to instability or erroneous performance. Furthermore, many frameworks approach this challenge by only focusing on the robotic system while neglecting human-centric factors. This oversight adversely affects the performance of these frameworks for pHRI.

While there is a significant body of work studying human-centric factors, they are generally performed in isolation, with little to no cross-disciplinary applications. This thesis therefore presents a framework for the integration of cross-disciplinary tools investigating human-centric factors into pHRI for personalising robotic assessment and response.

To work towards the goal of human-centric semi-autonomous systems, there are three aspects to scrutinise: (1) *examine* the assessment of human performance using task-centric measures, (2) explore *enabling* methods to identify user intent and integrate human-centric measures in the robot responses, and (3) *enhance* the understanding and assessment of human physiology using cross-disciplinary tools.

A greater understanding of these three perspectives contribute towards human-centric approaches for personalised robotic assessment and response during pHRI. *Enhancing* the understanding of human influences during pHRI applications will contribute towards improving the cohesion of human-robot dyads and will present opportunities to integrate physiological understanding into traditional robotic frameworks in the future.

Chapter 3

A Framework Towards Personalised Robotic Assessment and Response

The confluence of vision, arbitration, and actuation in robotics create a complex set of parameters to navigate. This challenge is exacerbated by the integration of human-centric perspectives to enable Physical Human Robot Interaction (pHRI). The human body consists of sophisticated subsystems which interact with each other extensively, leading to causal relationships which are difficult to untangle and decode. Furthermore, there is inherent variability shown from one human to the next, further complicating this process.

For non-physical HRI applications, this level of uncertainty is traditionally mitigated by reducing or removing the external influences that do not affect the robot's operation to complete a *predetermined task*. As a result, most techniques rely on *task-centric* measures to assess human performance and adapt the robot response to complete the task. The disregard for *human-centric* factors can sow discord between human and robot since the robot response may lead to sub-optimal circumstances for the human. A common example would be taking the shortest distance rather than one which factors in human ergonomics, leading to increased risks of discomfort or injury.

The ideal paradigm combines task-centric and human-centric factors to assess the human user, and modify the robot response such that the user's capabilities can be leveraged while supporting shortfalls during interactions. The first step towards this vision is to

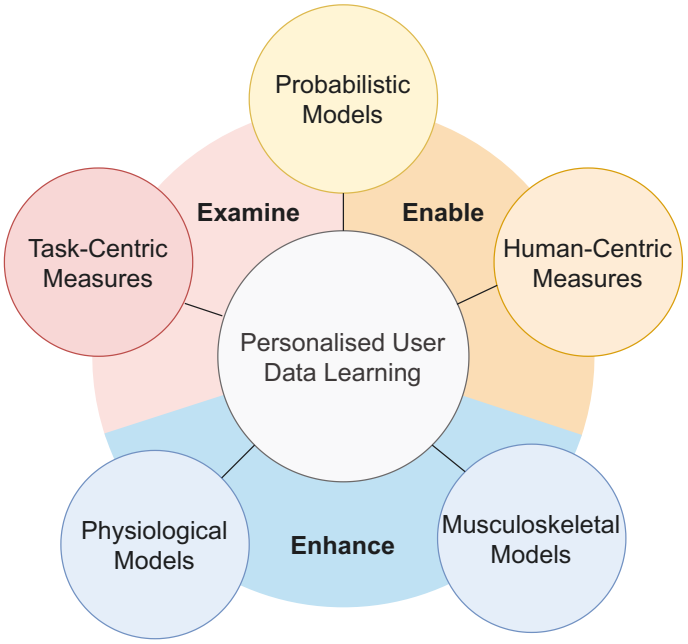

FIGURE 3.1: An overview of the framework towards personalised robotic assessment and response. The framework *examines* current task-centric measures and integrates human-centric measures in probabilistic models to *enable* learning. Musculoskeletal models and physiological models then *enhance* the robotic assessment process and potentially personalise the response provided.

identify the necessary perspectives to consider. To this end, a framework was developed to outline the various perspectives to investigate.

3.1 Description of Proposed Framework

At its core, the framework focuses on collecting human user data to generate predictive models which reflect an individual's inherent characteristics and idiosyncrasies. Consider the ideal paradigm for personalised robotic assessment and response using Physical Human Robot Interaction (pHRI). The physical interactions between the human user and the robotic system generate three key aspects that affect the performance and outcomes of the human robot system:

1. *Examining* the choice of model or method to learn from observed data,
2. *Enabling* the integration of human-centric influences and measurements in conventional assessment and response frameworks, and

3. *Enhancing* the assessment and response by looking at complex human-centric models and their capability to facilitate improved pHRI.

To unveil the effects of these three aspects for personalised robotic assessment and response, the framework presented in this thesis, shown in Figure 3.1, approaches each direction individually with a holistic viewpoint on the ideal paradigm. This facilitates a deeper understanding on current methods which are commonly used to bridge the gap between human and robot, yielding context-driven insights into pHRI and personalised robotic assessment and response.

3.2 Real-World Example: Robotic Rehabilitation

As mentioned at the start of the chapter, a disregard of human-centric factors may degrade the interaction experience. This is particularly relevant for pHRI since cognitive dissonance between the expected and actual response from the robot can negatively impact the quality of the interactions.

To explore this, a pilot real-world application of the framework was implemented, focusing on the applicability of the framework during upper limb robotic rehabilitation using only task-centric measures. The derivative implementation (Figure 3.2) forms the foundation to integrate various agnostic modules with customisation to facilitate personalised experiences. The vision for the application is to utilise the *Data Center* to collect and analyse data from each user, in order to inform the adaptive response generated to provide targeted response to the user during rehabilitation.

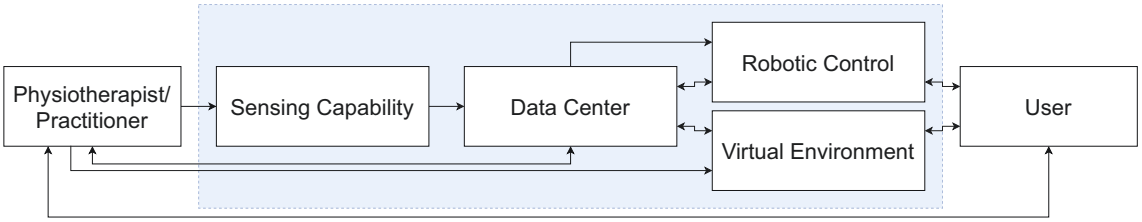

FIGURE 3.2: An overview of the implemented system for the gamification of personalised rehabilitation and assessment.

A fully-immersive game scenario is facilitated by the *Virtual Environment* during rehabilitation, while data is collected using the *Sensing Capability* of the system. Physical and haptic feedback is provided to the user through the *Robotic Control* module, and additional components can be easily integrated, providing additional modular channels for

data collection. The implementation of this pilot real-world application was realised using an integrated robotic system consisting a Cobot arm, a VR headset, and a depth camera.

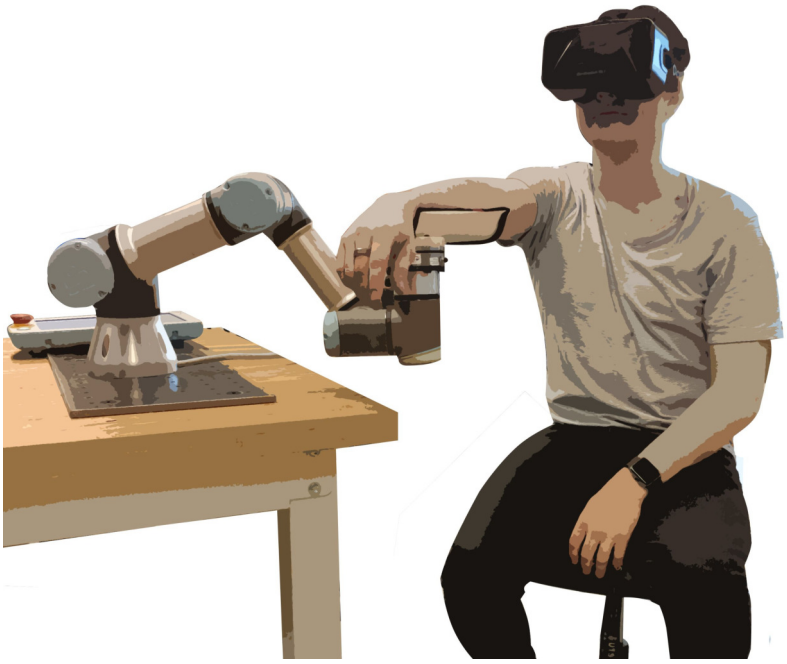

FIGURE 3.3: A visual overview of the implemented system used for the gamification of robotic rehabilitation.

3.2.1 Virtual Reality and Gamification

The implemented system was developed to explore the potential of an integrated approach towards robotic rehabilitation gamification, drawing inspiration from psychological perspectives for rehabilitation and assessment. There is evidence to suggest that gamification promotes cognitive responses [272] and user engagement [273], while the fully-immersive VR provides an "embodied recognition" for actions during rehabilitation [274].

To date, there is little work on rehabilitation gamification which utilises fully-immersive VR environments. Most works exploit semi-immersive VR environments by using external sensors, such as cameras or Inertial Measurement Units, to *insert* patients into the environment.

The VR environment was designed using the Unity game development platform, and communicates with other modules using the open-source ROS (Robot Operating System) package *ROSBridgeLib*. The package creates a web-socket that provides agnostic communication protocol connections, allowing cross-platform communications.

Development of a fully-immersive scenario with a VR headset requires considerations of portability, robustness, and application. The Oculus Rift headset was chosen due to its portability, compatibility with SteamVR, and it requires only a sole portable sensor to track head movements.

Gamification of the rehabilitation process was completed using 4 different game modes that can be selected by the therapist. The games provide targets for the patient to shoot using a virtual laser beam projected from the endpoint of the robotic arm. This requires the use of several different muscle groups in the limb to hit the targets. Hence, by carefully positioning the targets, a way of promoting the use of specific muscles can be achieved. User side effects such as nausea and disorientation [275] were mitigated through the synchronisation of the real UR3 with the virtual UR3.

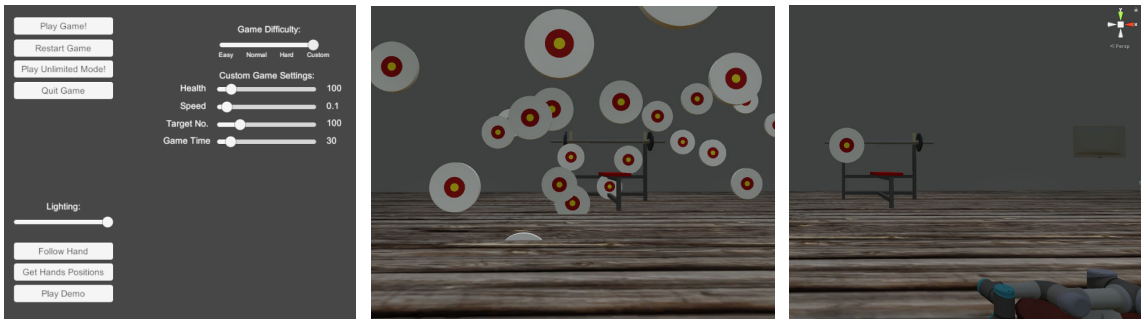

(A) (B) (C)

FIGURE 3.4: Snapshots of the VR environment: (A) the user interface for therapists to customise and choose game modes for the patient; (B) the targets spawning for Free-run Mode; (C) the target for Tracking Mode.

3.2.1.1 Free-run Mode

This game mode spawns a set number of targets where the patient attempts to destroy as many as possible within 30 seconds, as shown in Figure 3.4(B). The settings for the mode are set by the therapist through 3 default difficulties. Personalised configurations can also be set by the therapist, changing target health, speed, number of targets, and game time.

3.2.1.2 Timed Challenge Mode

For more directed muscle movements, this mode spawns single targets in series. When the target is destroyed, another target is spawned, requiring the patient to collect as many

points as possible within 30 seconds. This game mode allows therapists to design particular target placements to induce desired limb motions of the patient.

3.2.1.3 Tracking Mode

To increase patient-therapist engagement, a tracking mode was devised to allow for mirroring exercises. To facilitate patients outside of healthcare facilities, this mode allows therapists to create a target using a depth camera to track their right hand. Points are given based on which part of the target the virtual laser beam is hitting, emphasising accuracy with increasing point scores near the bullseye. This game mode is timed, allowing for consistent progress tracking for each patient. This game mode has the potential to allow real-time interactions by adding a video conferencing plug-in to enable real-time and remote feedback from therapists.

3.2.1.4 Pre-planned Path Mode

Patients and therapists have limited face-to-face consultation time during rehabilitation sessions. To maximise the benefits of robotic rehabilitation, the completion of repeated exercises is necessary. The pre-planned mode allows the therapist to record a pre-planned path for the patient to follow in their own time, changing configurations such as target speed and health. This game mode is suited to repeat exercises that patients can perform in their own time between each therapy session. In all of these game modes, each patient's scores are tracked and recorded in the data center, allowing for post-session analysis by the therapist.

3.2.2 Visual Feedback

Remote allocation of targets for the Tracking Mode and Pre-planned Path Mode are obtained through depth information from a PrimeSense Carmine RD1.08 camera. The open-source middle-ware library PrimeSense Natural Interaction Middle ware (NiTE) has been adapted to track the right hand of the user and communicate with Unity. The coordinates of the hand are transformed into the VR simulated environment frame of reference and recorded, forming the path users follow using the UR3. The path formation can be autonomous (subject to temporal constraints) or conducted manually. Both modes provide a single path for the target in the game.

3.2.3 Robotic Control

Robotic control was achieved using a UR3 robotic arm (Universal Robots, Odense, Denmark) to provide physical support and haptic feedback for the patient's upper limb. A bespoke 3D printed cast was manufactured to comfortably support and seat the patient's limb, and the controller for the robot arm was modified to enable assisted gravity compensation. To measure the interaction forces between the manipulator and the patient's upper limb, a 6-axis Axia80 force-torque sensor (ATI Industrial Automation, Apex, USA) was fitted between the limb cast and the robot endpoint.

To enable meaningful interactions between the robotic manipulator and the patient, an admittance control scheme is employed utilising the force-torque sensor readings to generate an appropriate motion trajectory for the robotic arm [276]. The desired velocity of the end effector ($\dot{x}$) is calculated using a wrench (F_{EE}) obtained from the force-torque sensor measuring the interaction between the patient and robot. An admittance gain matrix (K_a) is used to obtain a suitable task-space velocity command.

$$\dot{x} = K_a \cdot F_{EE} \tag{3.1}$$

Typically, the desired endpoint velocity is transformed into corresponding joint velocities using the inverse of the kinematic Jacobian of the manipulator. However in the presence of kinematic singularities the inverse of the Jacobian matrix becomes degenerate, resulting in erratic and dangerous motions of the robot. This poses a significant problem in applications such as robotic rehabilitation where human safety is of paramount importance [34]. To address kinematic singularities, damped least squares is implemented and a *damped* Jacobian inverse (J^*) is used to obtain the appropriate joint velocity for the desired task-space motion.

$$\begin{aligned} \dot{q} &= J^*(q) \cdot \dot{x} \\ J^* &= J^T(JJ^T + \lambda^2 I)^{-1} \end{aligned} \tag{3.2}$$

For cooperative robotic responses during admittance control, appropriate values for admittance gain and damping coefficients (λ) need to be ascertained. This mitigates any discord between user intent and the robotic response, leading to congruent interactions during pHRI.

3.3 Pilot Study and Results

To determine an empirical set of parameters for admittance control which best suits the user's preferences, a blind pilot study was conducted to find the relationship between the admittance gain matrix (K_a), user prior experience, user ability, and user preference.

Ten *healthy* adults participated in the study which required them to track a drawn path with a laser point projecting from the endpoint of the UR3. A candidate set of admittance gain matrices were chosen to avoid the robot triggering software safety e-stops.

Participants were not provided time to familiarise themselves with the control scheme on the robotic arm. During the study, each participant was only told, in a neutral manner, to follow the drawn path. A damping coefficient of $\lambda = 0.1$ was used for all experiments based on heuristic testing of the damping effect on exertion to manipulate the robotic arm.

Each participant completed 3 track-following tasks, one for each admittance gain matrix chosen at random, and were given a short survey shortly after to track their preference and contextualise their prior experience with Cobots. User ability was computed from the mean perpendicular error between the tracked path and the participant's actual path.

During testing, 2 out of the 10 participants' results were invalidated. This was due to human and technical errors such as discrepancies in instructions, the erroneous calibration of robotic arm prior to testing, and corrupted data sets discovered during post-hoc analysis. Table 3.1 outlines the results of the study.

TABLE 3.1: The best admittance gain matrix coefficient from eligible users in a blind test. (Number of eligible users = 8)

Admittance Gain	Best setting		Time Taken (s)		
	User chosen	Least error	Min	Median	Max
0.020	3	2	19.04	24.32	39.44
0.025	4	2	15.92	22.16	47.28
0.030	1	4	16.24	25.80	47.04

All participants took similar times to achieve the same task despite the responsiveness difference of the control scheme.

3.3.1 Discussion

The admittance gain matrix dictates the responsiveness of the endpoint motion and we postulate a higher K_a would result in a better experience of the control scheme. However, the disparity between what users experience and their objective results (Table 3.1) indicate that participants perceived the best configuration based on some subjective criteria rather than the results from the study objective.

This supports the observation that while the task was implicitly a planar task (2-Dimensional due to laser projection), all participants manipulated the robotic arm using all 3 axes to achieve the task with limited translation of their arm. This highlights the need for human-centric measures to be integrated in any pHRI frameworks since task-centric measures potentially lack enough information to capture human user responses which may originate from subjective criteria. While subjective human responses are difficult to map relative to their task-centric response, there is an opportunity for physiological measures to bridge the gap and provide insight into the human user during pHRI. For the pilot real-world application, this could potentially be realised using adaptive parameter tuning based on interaction forces and even neurological signals [22].

3.3.2 Limitations for Personalised Assessment

In the pilot application, the *Data Center* is responsible for relaying relevant information on the assessment of the user during rehabilitation. Given the abundance of robotic data available, the significant effort to interpret it may hinder the utility of these robotic systems [277]. As a result, there is a need for systematic methods to extract relevant information for the practitioner. Gaussian Process (GP) is a popular technique used in robotics due to their flexibility and simplicity. Details for GP models are outlined previously in Section 2.2.4.

One roadblock for the use of GP models are their computational complexity which is not scalable as more data is available. This constraint is exacerbated for robotic systems since most data are multi-dimensional. The Bayesian Committee Machine (BCM) [278] is an ensemble method which relieves the computational pressure by splitting the data set into several smaller data sets. The formulation of BCM and the experiment to validate its performance has been published in [14] and is briefly outlined in Appendix B.

Another challenge for relaying relevant information to the practitioner is the contextualisation of these results. For probabilistic models, such as GPs, the flexibility creates challenges when interpreting the resultant probabilities. The contextualisation of the model for robotic rehabilitation highlights the double-edged nature of GPs, as seen in a path-following task where the model is required to distinguish between an outlier path or some systematic shift in the user's capability.

Consider the outputs of a GP classifier where the recorded path is inputted into the model. A probability between 0 and 1, a representation of a "fitness" score, is produced when observing a path. This posterior probability relies on the training data used to build the model. As a result, the single probability score alone does not relay enough information to enable the practitioner to make appropriate decisions. A complementary measure, the Longest Warping Subsequence (LWSS) score, was developed to facilitate this distinction between a path that is correct, modified, or completely novel. The formulation and validation experiment for the LWSS score is outlined and presented in Appendix C.

3.4 Summary

This chapter has introduced a framework which focuses on new challenges faced when integrating human factors into personalised robotic assessment and response. The framework tackles three aspects of personalised robotic assessment and response by *examining* current task-centric models, *enabling* the integration of human-centric measures, and investigating ways to *enhance* the personalisation through physiological and musculoskeletal models. The details of each aspect are covered in the following two chapters, including a preliminary study in each chapter to accompany the main work presented.

A pilot real-world application for the framework was demonstrated by developing a robotic system for robotic rehabilitation. The application exploits sensor-modified gamification and haptic feedback using a VR environment and a robotic arm. The implementation highlights the applicability of the framework in various fields where the integration of human-centric influences may benefit the human user and improve outcomes for personalised robotic assessment and response. Additional investigations into examining roadblocks for Gaussian Process models can be found in Appendix B-C.

Chapter 4

Intent Estimation during Physical Human Robot Interaction

In the framework presented in Chapter 3, it was shown that enabling a personalised robotic assessment and response requires an understanding of the human user. The results from the pilot real-world application highlighted the need to align robotic response generation with the user's idea of the task during shared interactions. During these interactions, the integration of human-centric measures can help determine the user's intent.

One potential area where noise may be introduced, for a coupled human-robot dyad where the user is physically connected to the robot, is at the point of contact where the dynamics can be fluid. While the intent is for the robot to reflect the user's intended actions, inertial influences exerted upon the user may negatively affect the interaction. As a result, the effects from the inertial influences exerted by the robot to the user are studied on a prominent model, Fitts' law.

In Appendices B and C, the limitations of task-centric probabilistic models was highlighted when there is an absence of contextual information on the human user during pHRI. Following on the findings from the preliminary study, the focus for *enabling* personalised robotic responses aims to overcome this limitation by integrating relevant human-centric measures during coupled interactions.

This chapter presents a preliminary study on the inertial influence of robots for Fitts' law during pHRI, and then introduces an extension to a HRI framework, Interaction Primitives. Physical Human Robot Interaction Primitives (pHRIP) integrates interaction forces

between the human and robot during continuous interactions to perform probabilistic inference of user intent and to generate the appropriate robot response.

4.1 Preliminary Study: Inertial Influence on Fitts' Law for Intent Estimation

As emphasised in the results from the previous chapter, there were some observed influence of the Cobot's inertia during pHRI. To investigate this effect induced by the human-device interface, a preliminary study was conducted to determine whether coupled interactions are affected by this phenomenon. If there is indeed a significant effect induced from the human-device interface, the integration of human-centric measures during personalised robotic assessment and response would be further complicated as this confounding effect will need to be isolated from the human-centric measures.

In a human-human dyad, understanding and conveying information between the two agents are achieved through explicit and implicit interactions, such as verbal speech, body language, and haptic feedback. However, in human-robot dyads, a protocol is needed to be programmed for the robot to understand and discern human action and intent. While explicit programming of expected human actions and intent is possible, the intractable set of feasible actions that a human can perform makes this process functionally impossible. Thus, researchers have utilised elementary or commonly used movements as a benchmark when investigating functional relationships in the human body.

Fitts' law [98] is a seminal information-theoretic model which was developed to predict human movement accuracy and speed based on a free-movement target-directed reaching task. Its simplicity has lead to its use such as constraining EMG classifications [279] and developing task difficulty measures [280]. The constraints for target-directed reaching tasks, as well as the model's simplicity, makes it a prime candidate to explore the effects of the Cobot's inertial influence during physically coupled interactions.

Furthermore, the use of target-directed reaching tasks has been empirically shown to present similar characteristics from the human user [281], with spatio-temporal analysis

indicating humans "generate roughly straight hand trajectories with single-peaked, bell-shaped speed profiles" [248]. The model is defined as:

$$MT = a + b \cdot \log_2 \left(\frac{2D}{W} \right), where \tag{4.1}$$

$$I_d = \log_2 \left(\frac{2D}{W} \right). \tag{4.2}$$

The model suggests that for each human, there is a set of coefficients, a and b, which would accurately predict a user's movement time based on the difficulty of the target as defined by the Index of Difficulty (I_d). Readers are encouraged to review Section 2.2.1 for an overview of Fitts' law and its applications.

While there have been many applications of Fitts' law, most rely on mass-less interactions, such as a small stylus or the mouse pointer. A comparison of different interfaces suggest that familiarity with the interaction device influences the fit of Fitts' law [282]. Despite a similar investigation into Fitts' law for a hammering task [103], the inclusion of a metronome induces auditory stimulus which influences the results. This confluence of stimuli is outside the scope of this thesis and is not explored further.

4.1.1 Methodology

Following Fitts' law, the effect of the human-device interface is investigated by observing the fit of a human participant's movement times for a variety of targets with various index of difficulties. Statistical analysis is conducted on the fit of the coefficients (a and b in Eqn. 4.1) when correlating a target index of difficulty (I_d) to the respective movement time (MT). The setup for the study involves the participant performing upper limb target-directed movements whilst coupled to a robotic manipulator. This coupled arrangement can be seen in robotic devices for rehabilitation [283, 284] and Assistance-As-Needed systems for industrial applications [285].

Participants interacted with a 7-DoF robotic manipulator (HAHN Rethink Robotics, Rheinböllen, Germany) as shown in Figure 4.1. A 6-axis force-torque sensor (ATI Industrial Automation, Apex, NC) was affixed between a bespoke handle and the robot endpoint. The robot arm was programmed to be in a Cartesian impedance control mode native to Rethink Robotics' Intera SDK. This constrained movements to the X-Y plane parallel to the robot base at a height which was comfortable for seated participants.

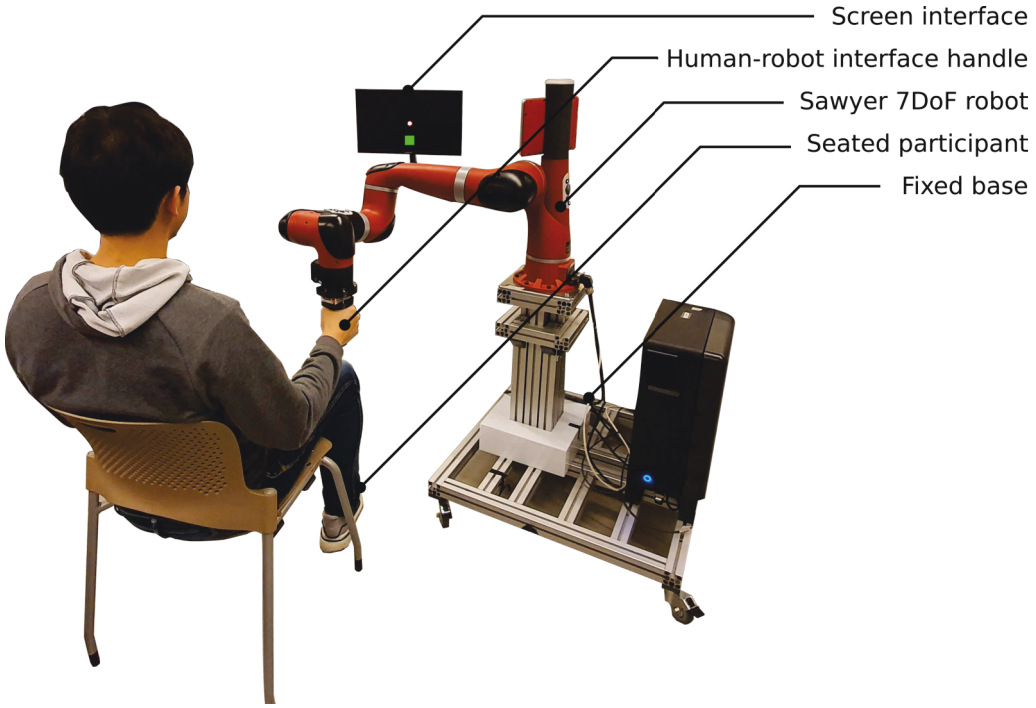

FIGURE 4.1: The setup for the experiments investigating inertial influences on Fitts' law for target-directed reaching movements.

Three different target widths (18.75mm, 37.5mm, and 56.25mm) and distances (120mm, 180mm, and 240mm) from the starting location of the endpoint were chosen for this experiment. The targets used for the experiments were arranged around the starting position of the participant's hand. Additionally, impossible-to-reach locations were excluded, thus removing targets that are occupied by the participants' body as shown in Figure 4.2. The participants did not have any prior knowledge of the possible target locations, and these locations were never visually displayed to participants until the experiment trials.

4.1.2 Human Participation Study

Ten healthy adults (9 males and 1 female) provided informed consent to participate in a target-reaching experiment approved by an ethics committee (UTS, Australia, approval number ETH18-3029). All participants were right-handed, and presented no known neuromuscular or sensory disorders.

Participants were seated and instructed to keep their back against the chair so the affixed handle in its starting location is 0.55m away from their torso. Visual feedback for the

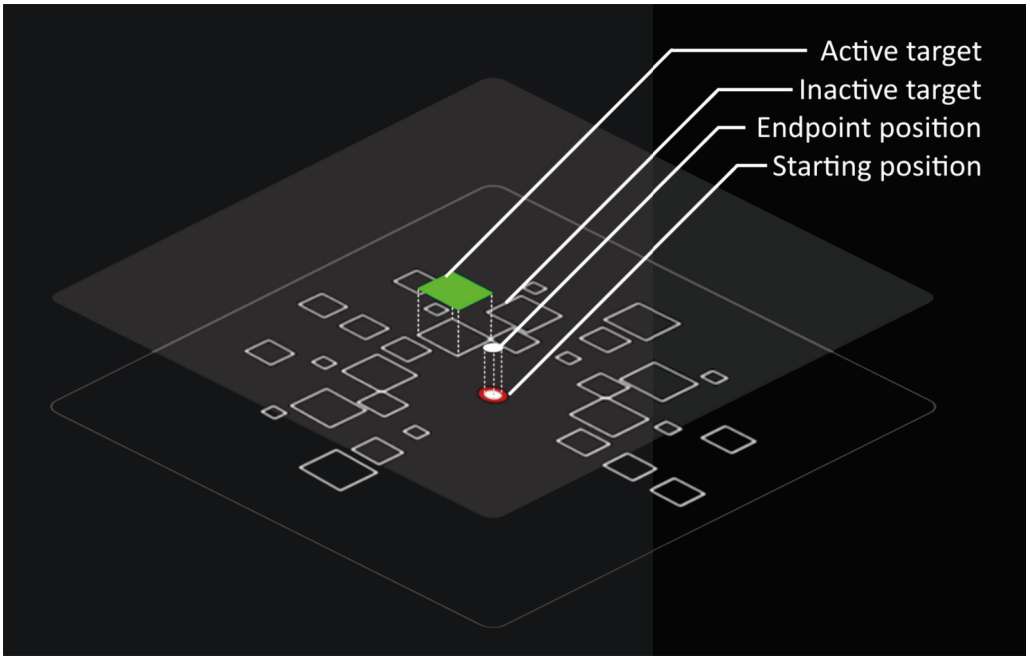

FIGURE 4.2: The array of targets that participants were required to reach using the handle affixed to the endpoint. The two planes, grey and black, represent what the participants are shown during the experiments, and the array of targets that is never revealed. The current position of the endpoint is indicated by a filled white circle, while the centre of the visual display is represented by the outline of the red circle.

experiment was communicated via a monitor screen in front of the participant, which displayed visual cues during the experiment.

Participants were instructed to move an on-screen marker, indicating the position of their hand, to the centre of a target as accurately and as quickly as possible. The participants were also instructed to move the handle towards the target while maintaining trunk posture to minimise torso movement. Specifically, they were asked to move to the target in a single motion, and to refrain from performing postural corrective movements to improve their endpoint accuracy.

Prior to the experiment, a trial run was conducted so participants could familiarise themselves with the setup and experimental procedure. This is to help remove bias due to unfamiliarity with the experiment conditions. Each participant was asked to perform 28 trials in the experiment according to a predetermined sequence, which was only displayed one-step-ahead and never displayed to the participant in its entirety.

TABLE 4.1: The participants' Fitts' law coefficients from the experiment.

FLC	Participant #									
	1	2	3	4	5	6	7	8	9	10
a	0.75	0.96	1.41	0.26	0.92	0.92	1.09	1.09	0.79	0.75
b	0.17	0.02	0.19	0.55	0.15	0.18	0.07	0.07	0.09	0.06

4.1.3 Results

Each target directed movement performed by the subjects was analysed to obtain the coefficients of Fitts' law. Table 4.1 shows the participants' Fitts' law coefficients (FLC) (Equation 2.1), calculated using linear least squares.

Although the Fitts' law model is personalised, an aggregation of the measured mean movement time and the associated I_d for all participants are shown in Figure 4.3. The 95% confidence interval for each participant's movement time, and the predicted movement times illustrate the relationship between task difficulty and movement time with a R^2 value of 0.853. Furthermore, the p-values in Table 4.2 demonstrate that the presence of inertia during interactions are statistically insignificant in this setup.

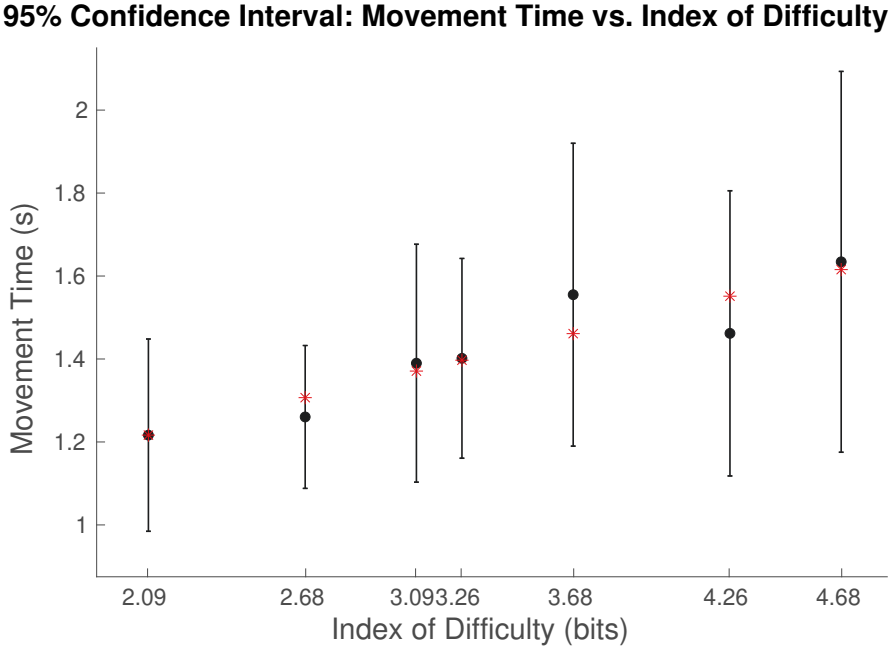

FIGURE 4.3: The 95% confidence interval produced from the aggregation of mean movement times from each participant. The circle markers, error bars, and red asterisks indicate the sample mean movement time, bounds for the 95% confidence interval, and movement times, respectively, as predicted by Fitts' law.

TABLE 4.2: The p-values for each index of difficulty calculated from the aggregation of mean movement times from each participant and the predicted movement time given by Fitts' law.

	p-values						
I_d	2.09	2.68	3.09	3.26	3.68	4.26	4.68
p-value	0.50	0.72	0.44	0.48	0.29	0.72	0.46

The targets for the trials were placed at one of three distances away from the starting position, as indicated by the 3 colours present in Figure 4.4. The shaded regions indicate the variation of velocity profiles between participants, supporting the notion of personalised parameters for Fitts' law. An aggregated endpoint velocity performed by participants was used to demonstrate the similarities that exist during the initial stage of the trajectories, as well as to highlight the increase in peak velocity when the targets were located further away from the starting point.

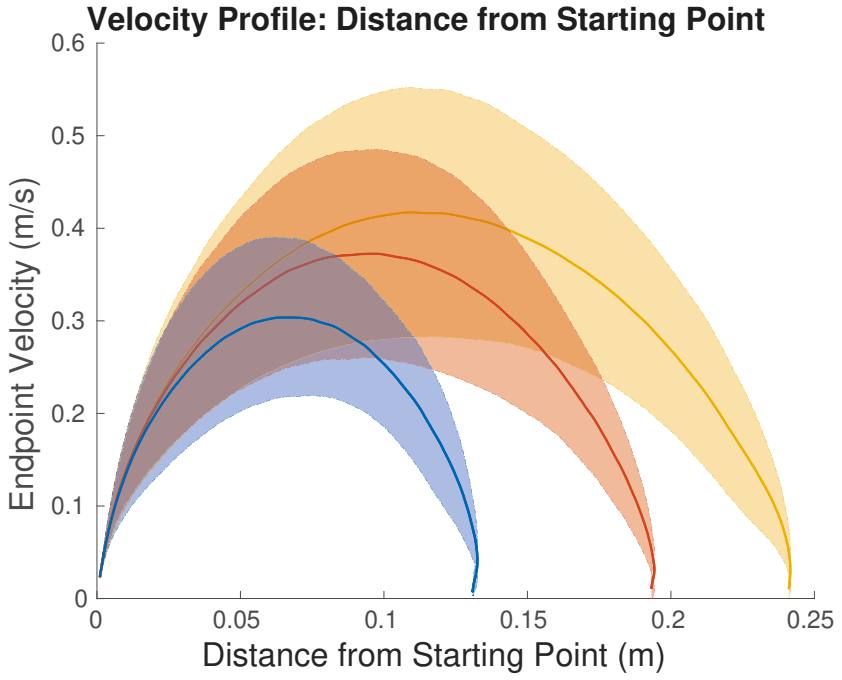

FIGURE 4.4: The average velocity profile mapping between the three target distances of the endpoint from the starting point. The standard deviation are represented by the shaded region.

The average velocity profile for the samples shown in Figure 4.5 demonstrates the single peak, bell-curved velocity profile nature of target-directed reaching movements. The two figures possess the same I_d, where target 1 is the first trial that uses that specified I_d, and target 25 is the last trial to use the I_d. For the first target, the endpoint velocities

are slower to peak and converge to the 0.1m/s threshold suggesting task unfamiliarity. Comparing against the velocity profile for the 25^{th} target, there is an indication that there is a learning component exhibited by participants although the effect is minimal.

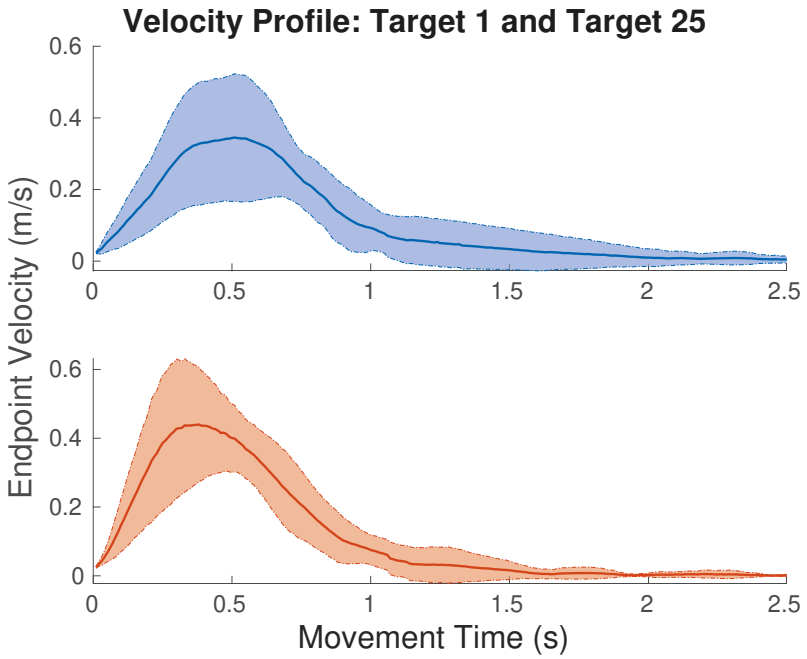

FIGURE 4.5: A comparison of the average velocity profile between an earlier trial and a later trial. The standard deviation are represented by the shaded region.

To analyse the trajectories performed during the reaching motion, a subset of the targets was chosen. The targets selected were those located directly in front of, and behind the designated starting position of the trial. Although the most direct path to the target was a straight line trajectory, all participants exhibited curved paths toward the target. Furthermore, the curved paths were not bound to either side of the x axis as shown in Figure 4.6. The curved Cartesian trajectories are potentially an effect of the mechanical configuration of the manipulator predisposing its motion to either side of the plane.

4.1.4 Discussion

One effect of the interface can be characterised by the noisy velocity profile. This feature can be attributed to the SEAs in the robotic manipulator. The elasticity from the springs within each joint can create potential parasitic dynamics when participants interact with the device. This contrasts to traditional Fitts' law tests which have extensively explored human computer interactions which rely on interaction interfaces without significant inertia and internal dynamics.

Figure 4.6 shows that the configuration of the manipulator produces varying inertia at the endpoint, thus affecting its susceptibility to move to either side of the x-axis. However, it can be seen that participants correct this motion in an attempt to finish the motion within the specified target; this correction adds another aspect that participants must consider during their movements.

Another aspect which is not explored is the visual feedback type provided to the participants. In contrast to the original Fitts' law study, the experiment conducted utilised remote visual feedback, requiring participants to rely on their hand-eye coordination to remotely estimate the location of the endpoint and target in Cartesian space. Furthermore, it is worthwhile noting that the original design of experiments for Fitts' law subjected participants to "rapid and uniform responses that have been highly overlearned" [98] which does not match the pHRI paradigm.

This preliminary study indicates that the inertial influence from the coupled interactions with a Cobot is not statistically significant when applying human-centric measures into predictive frameworks for human movement. This suggests that, for a coupled human-robot dyad, the integration of human-centric measures can be performed without any

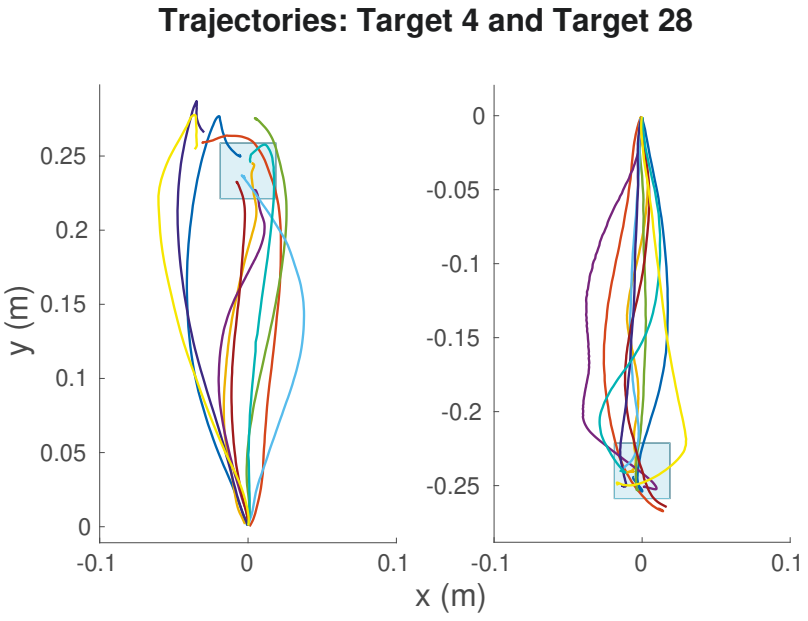

FIGURE 4.6: Variations in participant trajectories as they reach two different targets. Both targets are a pure translation along the y-axis, in front (left) of and behind (right) the starting position of [0,0]. The cyan square represents the size of the target for each respective trial.

apparent confounding influence. However, we note that although this study applies to Fitts' law, it may not directly translate to more complex models.

4.2 Leveraging Probabilistic Distributions for pHRI

Given that the preliminary study has shown the negligible effect of inertial influence on conventional predictive frameworks, the focus now shifts towards the integration of human-centric measures in conventional probabilistic frameworks.

Enabling personalised robot responses require the robot partner to learn, generalise, and adapt to different tasks while taking into account the eccentricities of the human partner. This is particularly important in pHRI applications with coupled human-robot dyads since the robot is commonly used to augment the capability of human operators with parametric frameworks. Utilising probabilistic operators in conventional frameworks can generalise each user's idiosyncrasies as well as adapt to different tasks as opposed to individual models for each task.

Intuitively, humans exert forces to indicate their intention during physical interactions in everyday situations such as pushing an object along a table or carrying an object. Thus, employing these forces can enable an estimate of the user's intent. For a coupled human-robot dyad, data from the points of contacts and interaction forces may contribute towards the understanding of the user's intent and improve the efficacy of the robot response.

Prior works for HRI, such as Interaction Primitives [186], have utilised observations of the human hand's pose to infer intent and generate an appropriate robotic trajectory in response. However, in coupled human-robot dyads, the position of the hand and the robot endpoint are identical, precluding the use of these methods which are based on task-centric measures. Thus, integrating the forces associated with the pHRI is a necessity to improve these conventional frameworks. Additionally, in applications where the robotic trajectories are very similar, the inclusion of interaction forces enables better inferences of the user's intent. It can also be hypothesised that the absence of robotic trajectories in these situations can lower the uncertainty during user intent inference.

4.3 Interaction Primitives for HRI

For HRI tasks with non-physical interactions, Interaction Primitives [186] use observations of demonstrations to: (1) build a parameter distribution, $p(\boldsymbol{\theta})$; and (2) use partial observations of phase-aligned trajectories to obtain the predictive distribution over the parameters for Dynamic Movement Primitives.

IPs model a distribution, $p(\boldsymbol{\theta})$, over the parameters of a DMP, ω, along with the goal of the trajectory, g. Given multiple demonstrations, an estimated Gaussian distribution over the parameters, θ, is generated. The predictive distribution is then obtained by using partial observations, τ_o of a trajectory to obtain the conditional distribution, $p(\boldsymbol{\theta}|\tau_o)$, by applying Bayes' rule.

The likelihood, $p(\tau_o|\boldsymbol{\theta})$, is obtained by computing the forcing function of the observed trajectory, τ_o, and the weighting matrix, Ω. The conditional distribution is then obtained by applying Gaussian conditioning on the joint distribution, $p(\tau_o, \boldsymbol{\theta})$, and the IPs parameters generate the robot response using DMP.

IPs couple multi-DoF systems by extending the parameter set, $\boldsymbol{\theta}$, to incorporate the weights and goals for each DoF, enabling multi-agent interactions. It is noted that IPs require the observations of the partial trajectories of either agent to predict the most likely set of parameters.

For de-coupled HRI systems, this spatial and temporal correlation allows for seamless interactions. However, in a coupled human-robot system, the *endpoint* of the robot and human are identical since they are located at the point of contact. A naïve approach to integrate the interaction forces in IPs is to apply DMP to the observed forces.

However, since DMPs are acceleration-based dynamical systems, the resultant parameters from interaction forces would be very noisy. Furthermore, the need for a goal in IPs makes the use of interaction forces unintuitive. Thus, this limits the application of IPs to physical interactions during pHRI.

4.4 Physical Human Robot Interaction Primitives

Physical Human Robot Interaction Primitives (pHRIP) aims to predict user intent by using the interaction forces available in coupled human-robot dyads during pHRI. pHRIP

extends upon IPs to: (1) integrate observed interaction forces during demonstrations into the parameter distribution, $p(\boldsymbol{\theta})$; and (2) use partial observations of phase-aligned interactions, χ, to obtain the predictive distribution over the pHRIP parameter set, $p(\boldsymbol{\theta}|\chi)$, and generate a robot response matching the user's intent.

4.4.1 Building the pHRIP Parameter Distribution

As reviewed in Section 2.3.3, a DMP encodes a single trajectory using a set of weights, ω, represented by M Gaussian basis functions equally spaced across the phase of the trajectory. In each trajectory, the observed interaction forces (T samples) are re-sampled into z samples, giving a $(z \times 1)$ vector, F. This phase-aligns the interaction forces to the trajectory.

Thus, for a n-DoF robotic system with d-DoF of observed interaction forces, the pHRIP parameters for a single trajectory is a $((nM + dz) \times 1)$ vector:

$$\theta = [{F_1}^T, \ldots, {F_d}^T, {\omega_1}^T, \ldots, {\omega_n}^T]^T. \tag{4.3}$$

Given K demonstrated trajectories, the distribution over the pHRIP parameters, $p(\boldsymbol{\theta})$, follows as:

$$p(\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\theta}|\mu_{\boldsymbol{\theta}}, \Sigma_{\boldsymbol{\theta}}), \tag{4.4}$$

$$\mu_{\boldsymbol{\theta}} = \frac{\sum_{j=1}^{K} \theta_j}{K}, \tag{4.5}$$

$$\Sigma_{\boldsymbol{\theta}} = \frac{\sum_{j=1}^{K} (\theta_j - \mu_{\boldsymbol{\theta}})^T (\theta_j - \mu_{\boldsymbol{\theta}})}{K}. \tag{4.6}$$

4.4.2 Phase Estimation for Partial Observations

Once the pHRIP distribution is built, partial observations are then used to predict the user's intent. The observed interaction forces, F^*, are re-sampled to the same frequency as the robotic system through cubic spline interpolation.

To understand the context of the partial observations, the current phase of the interaction is needed. This is performed using a multi-dimensional DTW algorithm [286], comparing a $(r \times n)$ reference trajectory, R, against the $(v \times n)$ observed sequence, ν.

During the DTW process, the matrix used to compare R and ν provides a $(v \times r)$ warping map, W, indicating the distance from one observation from R to another in ν. Using this warping map, a $(v \times 1)$ index vector, ρ, is obtained, which outlines the phase alignment for each observation in the observed sub-sequence, ν:

$$\rho = \left[\frac{col(min(\nu_1, R))}{r}, \ldots, \frac{col(min(\nu_v, R))}{r} \right]^T, \tag{4.7}$$

where $col(.)$ is a function to obtain the column number of the vector/matrix.

4.4.3 pHRIP Parameter Set Inference

To infer a set of pHRIP parameters which represents the user's intent, the predictive distribution is obtained using partial observations, χ, between the user and the robot during pHRI. Applying Bayes' rule, we obtain: $p(\boldsymbol{\theta}|\chi) \propto p(\chi|\boldsymbol{\theta})p(\boldsymbol{\theta})$.

The pHRIP distribution consists of the interaction forces and robotic trajectory encoded using a $(dz \times 1)$ vector representing the phase alignment of the interaction forces, and the $(nM \times 1)$ Gaussian basis function weights from DMPs. Thus, the likelihood distribution, $p(\chi|\boldsymbol{\theta})$, is modelled using a Gaussian distribution:

$$p(\chi|\boldsymbol{\theta}) \sim \mathcal{N}(\chi|\Omega\theta, I\sigma^2), \tag{4.8}$$

where σ^2 is the observation variance, and Ω is a $(v(d+n) \times (dz + nM))$ weight matrix for the observations:

$$\Omega\theta = \begin{bmatrix} \begin{bmatrix} \lambda_1 & \ldots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \ldots & \lambda_d \end{bmatrix} & \mathbf{0}_{(vd \times nM)} \\ \mathbf{0}_{(vn \times dz)} & \mathbf{0}_{(vn \times nM)} \end{bmatrix} \begin{bmatrix} F_1 \\ \vdots \\ F_d \\ \omega_1 \\ \vdots \\ \omega_n \end{bmatrix}. \tag{4.9}$$

The weight matrix for each DoF of the phase-aligned interaction forces, λ, is a $(v \times z)$ matrix, starting as a zero matrix and filled based on the index vector ρ. For each row of λ, corresponding with each sample v of ν, the value of the column numbered closest to

the integer-rounded $z\rho_v$ is filled:

$$\begin{aligned} \lambda_{x,y} &= \begin{cases} 1 & for\ y = z\rho_x \\ 0 & for\ y \neq z\rho_x \end{cases}, where \\ x &\in (1, 2, \ldots, v), \\ y &\in (1, 2, \ldots, z). \end{aligned} \tag{4.10}$$

Given the likelihood, $p(\chi|\boldsymbol{\theta})$, is modelled with the robotic trajectory, the unavailable trajectories are set to 0, giving $\chi = [F_1^{*T}, \ldots, F_d^{*T}, \mathbf{0}_{(1\times vn)}]^T$ where F^* are the interaction forces observed. The joint distribution is then defined as:

$$p(\chi, \boldsymbol{\theta}) = \mathcal{N}\left(\begin{bmatrix} \chi \\ \boldsymbol{\theta} \end{bmatrix} \middle| \begin{bmatrix} \omega\boldsymbol{\theta} \\ \mu_{\boldsymbol{\theta}} \end{bmatrix}, \begin{bmatrix} A & \Sigma_{\boldsymbol{\theta}}\Omega^T \\ \Omega\Sigma_{\boldsymbol{\theta}} & \Sigma_{\boldsymbol{\theta}} \end{bmatrix} \right), \tag{4.11}$$

where $A = \sigma^2 I + \Omega\Sigma_{\boldsymbol{\theta}}\Omega^T$, and the mean and variance of conditional distribution, $p(\boldsymbol{\theta}|\chi)$ is derived as

$$\begin{aligned} \mu_{\boldsymbol{\theta}|\chi} &= \mu_{\boldsymbol{\theta}} + \Sigma_{\boldsymbol{\theta}}\Omega^T A^{-1}(\chi - \Omega\mu_{\boldsymbol{\theta}}), \\ \Sigma_{\boldsymbol{\theta}|\chi} &= \Sigma_{\boldsymbol{\theta}} - \Sigma_{\boldsymbol{\theta}}\Omega^T A^{-1}\Omega\Sigma_{\boldsymbol{\theta}}. \end{aligned} \tag{4.12}$$

A new set of pHRIP parameters, θ_n, is then sampled from this conditional distribution. The robot then generates a new trajectory using the subset of DMP parameters and the estimated phase of the final observation, ρ_v, continuing on from its last position.

4.5 Methodology

A coupled human-robot dyad is used to validate pHRIP's ability to generate the correct, personalised robotic response. The robotic system in the coupled dyad, shown in Figure 4.7, consists of a 7-DoF robotic manipulator (HAHN Rethink Robotics, Rheinböllen, Germany) with a 6-axis force-torque sensor (ATI Industrial Automation, Apex, USA) affixed between the endpoint and a bespoke handle. The robot state is recorded at 100Hz, while wrench data from the force-torque sensor is recorded at 125Hz.

The manipulator utilises Rethink Robotics' proprietary software, Intera SDK, and an endpoint velocity threshold of $2.5cms^{-1}$ is used to determine the start and end of a trajectory.

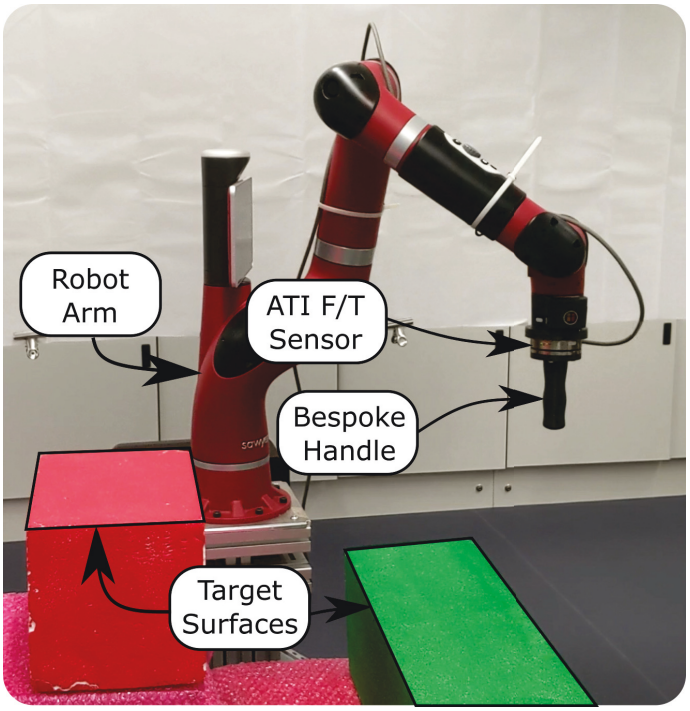

FIGURE 4.7: The experimental setup with a Sawyer robot arm, an ATI Axia80 force-torque (F/T) sensor, and a bespoke handle.

In all trials, generated trajectories were sent to the robot's native motion controller interface to be performed in an open-loop fashion. While it is possible to integrate low-level robotic feedback controllers with pHRIP, the focus is on the integration of human-centric measures for high-level trajectory generation reflecting the user's intention.

For the user-directed reaching and the Cartesian experiments, the process for collecting training, reference, and testing data are outlined in Figure 4.8.

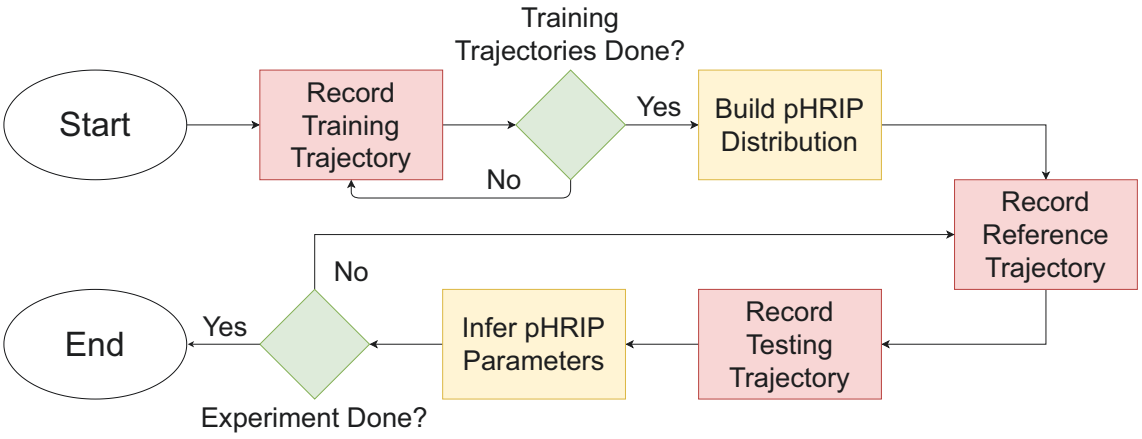

FIGURE 4.8: A flowchart of the experimental procedure to obtain the different data sets.

4.6 User-Directed Reaching Experiment

An experiment based on user-directed reaching movements was conducted to validate pHRIP across different users. Four participants (3 male and 1 female) took part in the experiment. During user interactions, the robotic arm was set up, using the native SDK, to enter an orientation-locked zero-g mode.

Participants were instructed to move the handle from a defined starting position to one of two target surfaces (Figure 4.7) in a "natural manner". Since DTW is used to phase-align the observed interaction forces, a reference trajectory is recorded for each trial to facilitate this and obtain the end pose. No other part of the reference trajectory is utilised once the phase estimate of the interaction forces is obtained. Each trial consists of a reference and a test trajectory, with participants instructed to perform both trajectories "consistently".

A total of 90 training demonstrations were recorded (45 for each target surface) using kinesthetic teaching via the bespoke handle for participants. During the test trajectory, participants were instructed to release the handle after 0.5-1.0 seconds while the new trajectory is generated and performed by the robot arm.

Partial observations for IPs consist of Cartesian trajectory, while pHRIP utilise Cartesian interaction forces only. Partial observations of both interaction forces and trajectory were used to perform further comparisons. Prior to the start of the experiment, participants were given a 5-minute window to interact with the robotic manipulator to familiarise themselves with the setup.

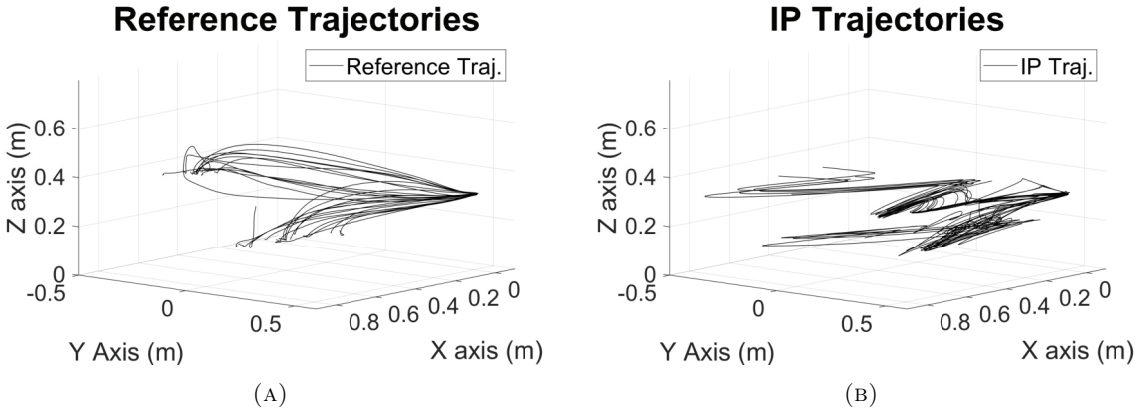

(A) (B)

FIGURE 4.9: (A) The reference trajectories from participants in the user-directed reaching experiment. (B) Trajectories generated from IPs using partial observations of trajectory.

4.6.1 Results

A total of 28 trials were conducted in the experiment. For each trial, participants demonstrated a reference trajectory as shown in Figure 4.9(A). Partial observations from the test trajectory were then used to generate new trajectories for the robot. Trajectories were generated post-hoc using IPs and pHRIP. Partial observations of trajectories were used for IPs and the resultant outputs are shown in Figure 4.9(B). For pHRIP, partial observations of interaction forces were used with the resultant outputs shown in Figure 4.10(A).

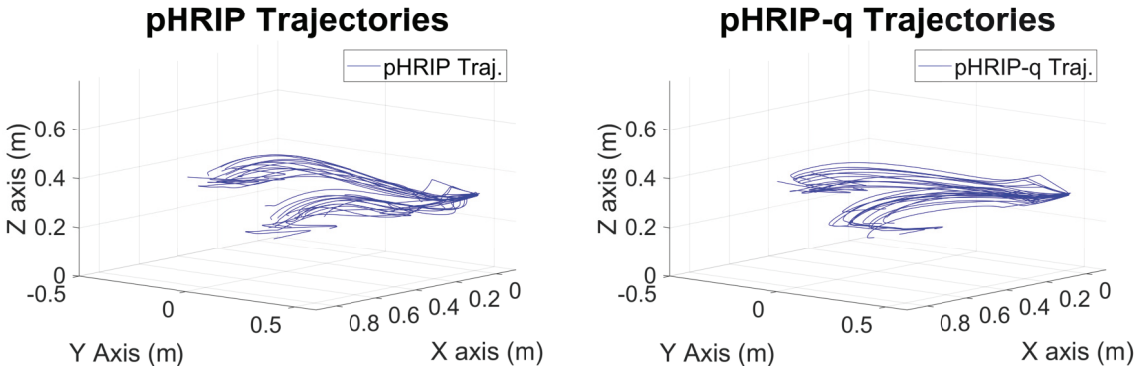

(A) (B)

FIGURE 4.10: Based on observing partial interaction forces, the resultant trajectories generated from: (A) pHRIP; and (B) pHRIP-q, a variant using joint states.

Visual inspection of the trajectories between IPs and pHRIP indicates the advantage of pHRIP over IPs despite no information on the robot trajectory. Using only interaction forces, the trajectories generated from pHRIP followed the shape of those in the reference trajectories. The performance of pHRIP is evident in the DTW scores in Table 4.3, indicating that the shape of the generated trajectories are very close to those of the reference trajectories.

The trajectories generated by IPs are quite ill-formed, with completely erroneous trajectory shapes generated, despite empirical evidence of its efficacy [186]. We posit the variance in the training demonstrations towards the end of the trajectory contributed towards the

TABLE 4.3: A comparison of the mean RMSE and DTW distance between the reference trajectory and trajectories generated from IP and pHRIP.

	IP (Trajectory Only)		pHRIP (Force Only)		pHRIP (Force & Trajectory)	
	μ	σ^2	μ	σ^2	μ	σ^2
RMSE(m)	0.2657	0.0031	0.1597	0.0030	0.1737	0.0024
DTW	4.2664	0.9732	1.6874	0.3056	1.8682	0.3800

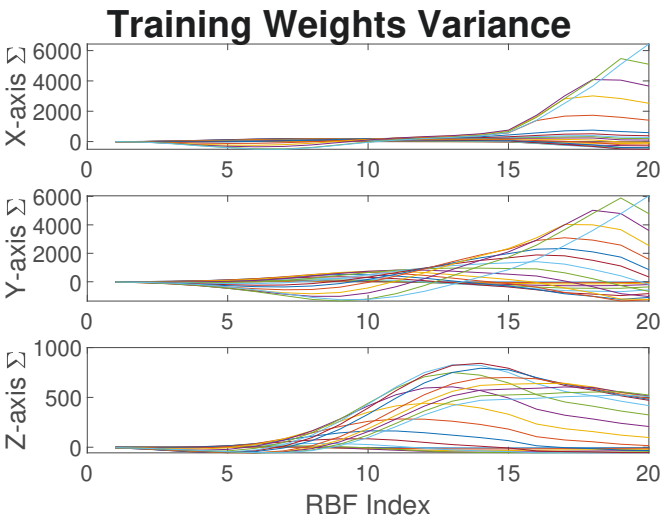

FIGURE 4.11: The variance in the training weights across the three Cartesian axis. Note the large variances towards the end of the trajectory contributing to conflicting IP conditional weights.

sensitivity of the distribution. Figure 4.11 highlights the large variances of the Gaussian basis function weights in all three axis at different sections along the trajectory.

The effect of this variance can be seen in the generated conditional distribution shown in Figure 4.12, noting that weights which are completely erroneous occur during the sections of trajectories with high variance. In contrast, the interaction forces overcome this sensitivity, generating a better conditional distribution as highlighted in Figure 4.13 despite

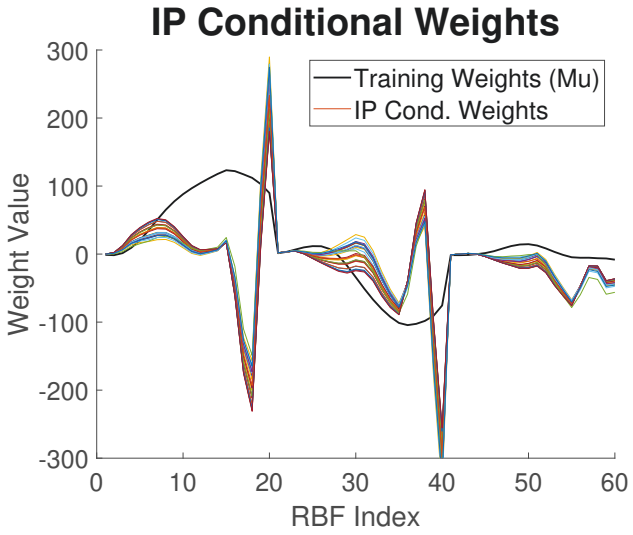

FIGURE 4.12: The conditional weights obtained given partial observations of the robot trajectory (which is identical to the hand pose).

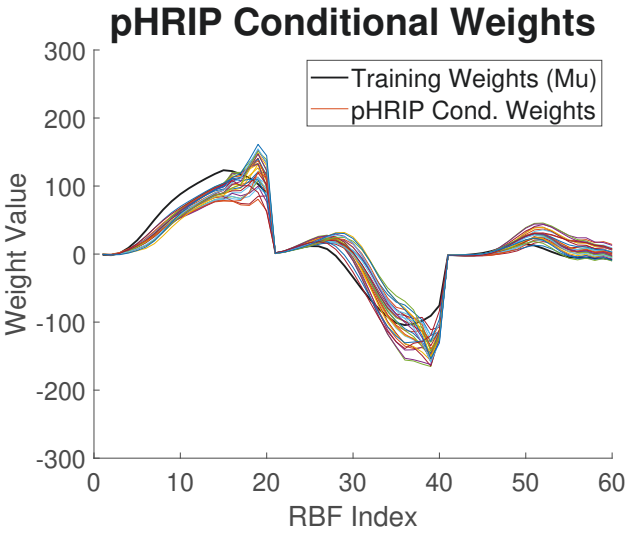

FIGURE 4.13: The conditional weights obtained given partial observations of the interaction forces at the endpoint of the robot.

the high variance. Therefore, looking at the results for pHRIP in Table 4.3, the use of force and trajectory observations account for the degradation in performance.

4.6.2 Cartesian vs. Joint Trajectories

While pHRIP has been shown to generate trajectories that reflect the user's intention in Cartesian space, for most robotic arm control systems, action policies generally operate in joint space. Further analysis was conducted using the data from the user-directed reaching experiment to investigate the application of pHRIP in joint state space (pHRIP-q). The joint states of the robot and the Cartesian interaction forces at the endpoint were used to build the weight distribution.

When assessing the performance of the pHRIP-q variant, the robot response (joint states) are re-mapped to Cartesian space (endpoint) using forward kinematics. Comparisons and analysis of the pHRIP-q are all based off the robot endpoint rather than the joint states.

TABLE 4.4: A comparison of the mean RMSE and DTW distance between the reference trajectory and trajectories generated from pHRIP-q.

	(Force Only)		(Force & Trajectory)	
	μ	σ^2	μ	σ^2
RMSE (m)	0.119123	0.004877	0.119118	0.004874
DTW	1.627105	0.572789	1.627127	0.572535

For both the original human target reaching experiment and this analysis, identical parameters were used with 20 basis functions, $K = 80N/m$, $D = 20Ns/m$, $\tau = 0.35$, $h = 0.0008$, and $\alpha = 1$. Forward kinematics was performed for the pHRIP-q trajectories to obtain their Cartesian trajectories, which showed similar trajectory shapes to the reference trajectories (see Figure 4.10(B)). This similarity is supported by the mean Root Mean Square Error (RMSE) and DTW distances as tabulated in Table 4.4.

Initial observations of the RMSE and DTW results suggest that pHRIP-q is the better variant. However, mapping the vectors showing the difference between the reference and resultant endpoints tell a different story. Figure 4.14 demonstrates the resultant ellipses using a fitting algorithm [287] which suggest that the appropriate pHRIP variant will depend on the priority of the task. For example, pick and place operations of heavy objects will prioritise the precise endpoint of the trajectory, making pHRIP more appropriate. Conversely, if the task is to conform to the shape of a trajectory performed by an expert, as is commonly seen during physical rehabilitation, it may be more suitable to use the joint variant pHRIP-q.

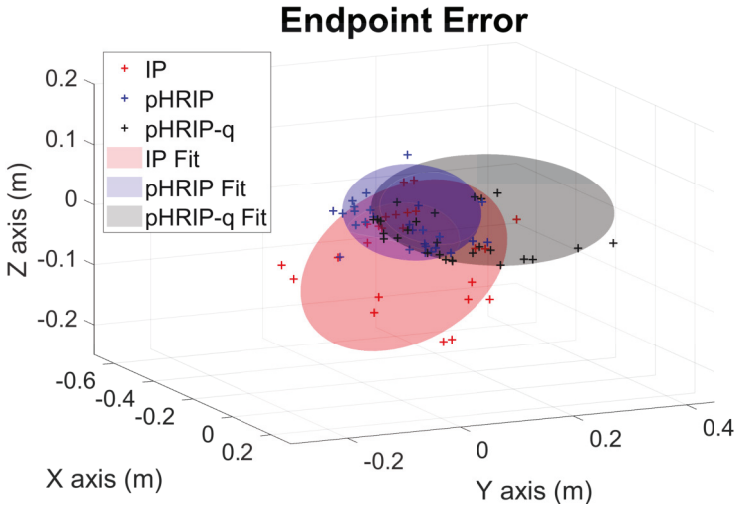

FIGURE 4.14: A visualisation of the discrepancy between the endpoint in the reference trajectories and those generated from IP, pHRIP, and pHRIP-q.

One other factor to consider when analysing the results is the range of the joint trajectories when performing the task in the experiment. In Cartesian space, the endpoint trajectory is already relatively small relative to the total workspace of the robotic arm. This is exacerbated in joint state space since the endpoint orientation was constrained, leading to an over actuated system (3-DoF endpoint movement with 7 joints). Since the range of the movement affects the pHRIP weights, the use of joint state space creates homogeneous weight parameters. This homogeneity results in a high similarity score for the RMSE and

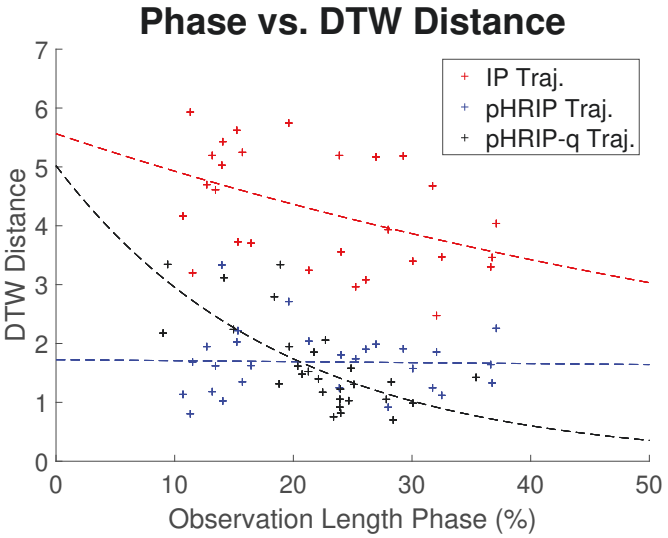

FIGURE 4.15: The relationship between observation lengths and DTW distances when generating trajectories using IPs, pHRIP, and pHRIP-q.

DTW distances when obtaining the conditional set of pHRIP parameters. If the robot was a lower DoF system, the decrease in redundant joints when moving in Cartesian space may mitigate this effect.

While the results for pHRIP show that appropriate consideration is required when choosing which variant to use, they show promising indications for the integration of haptic information during motor skill learning in pHRI applications. One potential application is in training and development systems, where expert demonstrations may be collected remotely via a haptic interface, providing intuitive motor skill learning remotely. Learning from the interaction forces on the haptic interface is transferable across various platforms provided kinesthetic teaching of the robot response is performed.

4.6.3 Influence of Observation Length

The influence of the observation length between trajectories generated using IPs and pHRIP can be seen in Figure 4.15, indicating the DTW distances for trajectories generated using pHRIP are much lower than those of IPs. As observation lengths increase, errors from the generated trajectories would approach zero, giving diminishing returns for pHRIP and IPs. While a longer observation can improve the performance of HRI applications using movement primitives, this is undesirable in pHRI applications since the goal of the robot is to contribute meaningfully as soon as possible. Trajectories from

pHRIP consistently produce better trajectories when compared against IPs, reinforcing our hypothesis that using interaction forces during pHRI can help reduce uncertainty.

It is noted that the exponential fit used to model the DTW distance and the observation length phase produces a profile for pHRIP-q which is vastly different to pHRIP. As mentioned previously, it may be possible that this particular task is well-suited for the joint state variant for pHRIP. One other possible explanation is that the model is over-fitted, with not enough samples to accurately determine the correct profile for pHRIP-q. However, the results from pHRIP-q still improve upon IPs, supporting the motivation for pHRIP.

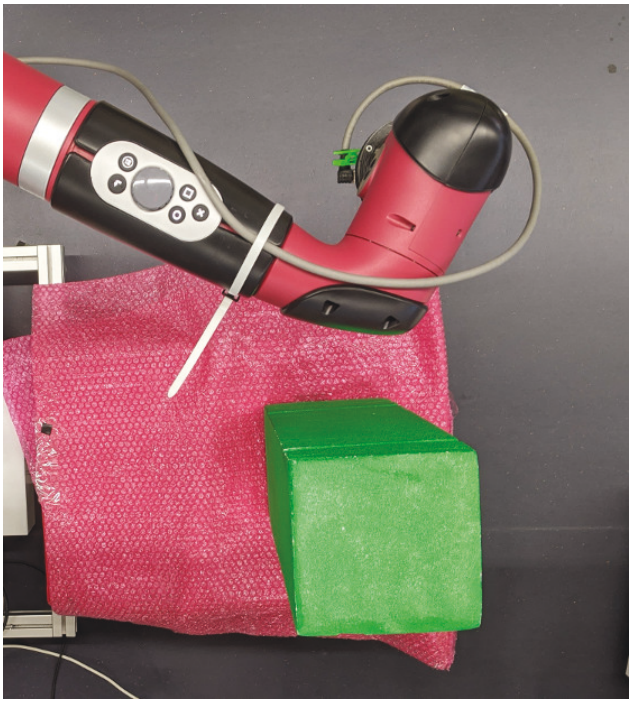

FIGURE 4.16: A top-down view of the planar validation setup.

4.7 Planar Target-Reaching Task

A planar target-reaching task was conducted to validate pHRIP, comparing against IPs, with the setup shown in Figure 4.16. An object was placed in between the start and the end of the trajectory, creating two distinct paths for the robot to reach the end position. A total of 30 training trajectories and 20 testing trajectories were recorded with an even split for each path, while the robotic arm is set up to enter an orientation-locked zero-g mode constrained to the XY plane.

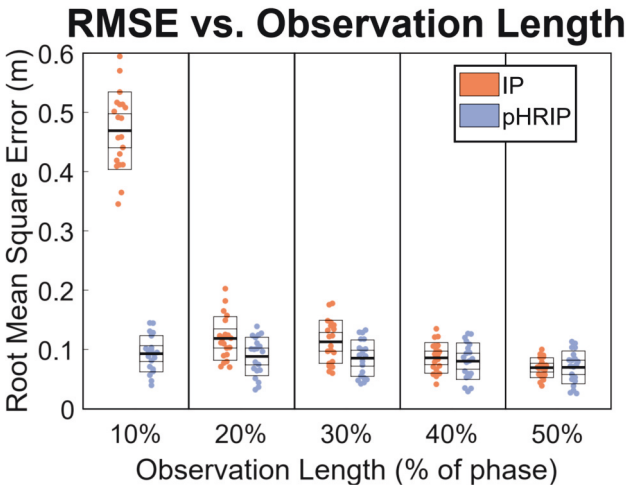

FIGURE 4.17: A modified box plot for each observation length highlights the RMSE spread between IPs and pHRIP for the 20 testing trajectories across various observation lengths. Comparing between two methods, the advantage of pHRIP over IP stands out when there are less observations since the interaction forces provide additional information on user intent.

The analysis and trajectory generation for the planar obstacle avoidance experiment were conducted post-hoc. Thus, all recorded trajectories and interaction forces were re-sampled to 400 and 500 samples respectively (matching the 100Hz and 125Hz data collection frequency). A comparison between IPs and pHRIP is performed using partial observations of: (a) trajectory only (IPs); and (b) interaction force only (pHRIP).

4.7.1 Results

For the 20 test trials conducted in the planar validation, the RMSE between the generated trajectories and their respective reference trajectories was calculated. The pHRIP and IPs tests were conducted post-hoc against observation lengths (as a % of the total trajectory) varying from 10% to 50%. Results shown in Figure 4.17 highlight the ability for pHRIP to address ambiguities in the trajectories, utilising only interaction forces to generate the intended path. The critical advantage of pHRIP over IPs is shown when there are fewer observations, such as those when only 10% of the trajectory is observed.

While the results may indicate that the advantage of pHRIP diminishes as more observations are obtained, the generated trajectories highlight an aspect of motor skill learning which is not inherited through IPs. The task to avoid the static obstacle is redundant, meaning there are multiple ways to complete the task. During the recording of training trajectories, Figure 4.18(A) shows that two distinct paths were taught kinesthetically by

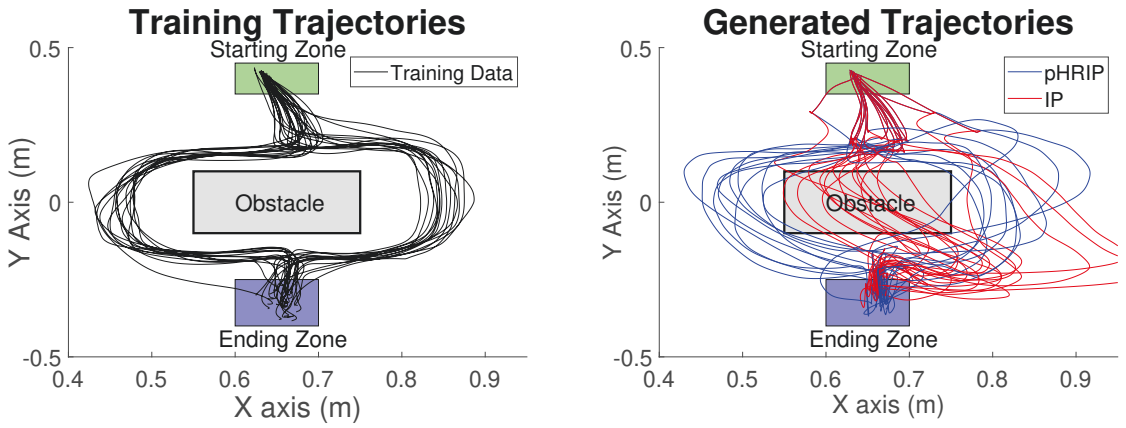

(A) (B)

FIGURE 4.18: The trajectories from the planar target-reaching task. (A) displays the training trajectories; and (B) shows the trajectories generated by pHRIP and IP when 30% of the trajectory is observed.

the demonstrator. From Figure 4.18(B), trajectories generated by IPs (in red) all collide with the obstacle severely, while the same only occur to 25% of trajectories generated by pHRIP. Observations of the pHRIP trajectories also show decreased collision severity, with only 2 trajectories out of 20 going through the core of the obstacle.

This result highlights the limitation of IPs in multi-modal tasks which is overcome by using interaction forces for pHRIP. Building the IPs distribution on only trajectories create an "averaging" effect when generating the robot response. By adding additional DoF from interaction forces, pHRIP requires less observations to achieve the correct trajectory, reinforcing our belief that integrating interaction artefacts into conventional frameworks improve motor skill learning for pHRI systems.

4.8 Cartesian Target-Reaching Application

Further experiments were conducted to reinforce the applicability of pHRIP to estimate user intention during pHRI. The user is tasked with moving the endpoint from the same starting position to various end regions. In total, 4 configurations were set up, seen in Figure 4.19, using a number of blocks which simulate changing task parameters and environments, as is common in pHRI applications. For each configuration, 10 training trajectories were recorded and are shown in Figure 4.20. The pHRIP parameters from all 40 trajectories are built into a single distribution, $p(\theta)$.

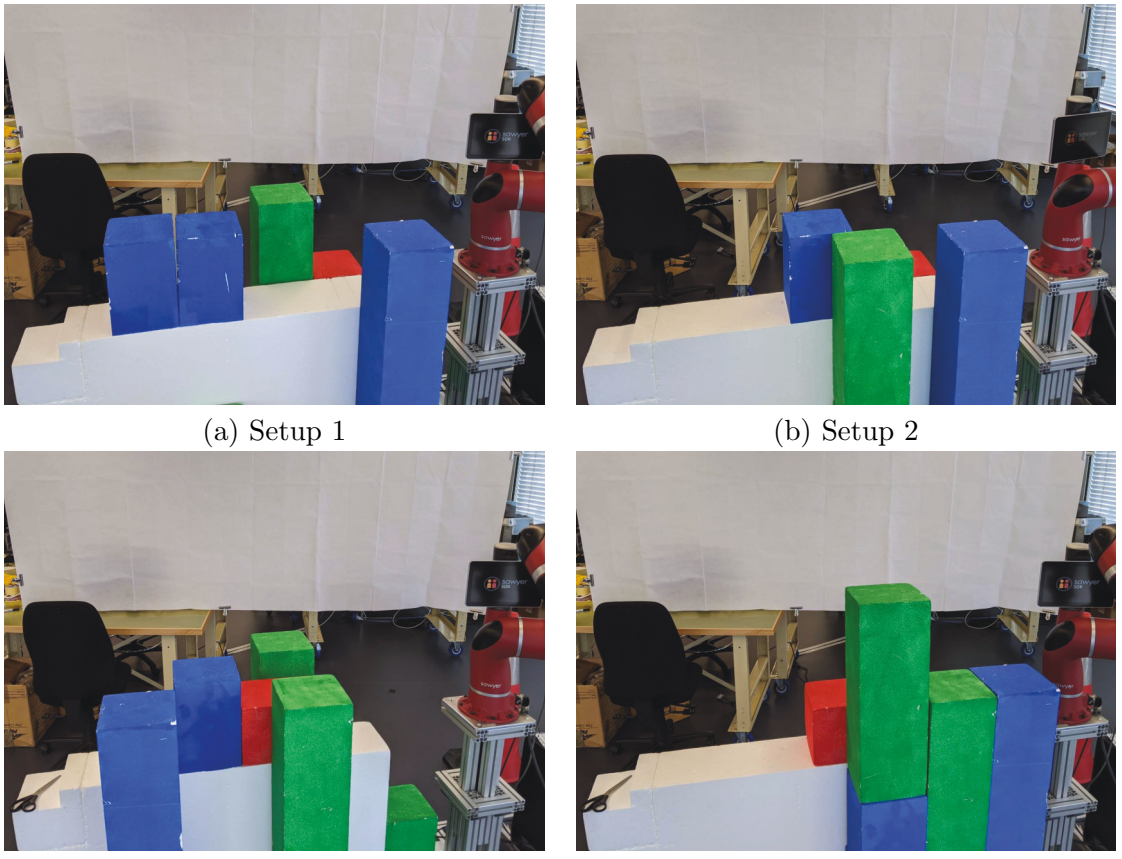

(a) Setup 1 (b) Setup 2

(c) Setup 3 (d) Setup 4

FIGURE 4.19: The four workspace configurations for the target-reaching application.

To validate pHRIP in this Cartesian target-reaching task, a total of 12 trials were conducted in Cartesian space. Similar to the user-directed target reaching task, the user releases the handle after 0.5-1.0s, allowing the robot to perform the trajectory generated by pHRIP. Successful trials are defined as generated trajectories that reach their intended end zone and follow the shape of the reference trajectory. For all trials, identical sets of parameters were used with 30 basis functions, $K = 80N/m$, $D = 20Ns/m$, $h = 0.0008$, and $\alpha = 0.8$.

One issue that arises in DMP-based frameworks, when the number of samples for each trajectory is different, is the area of effect for the Radial Basis Function (RBF) activation. A static value of τ will affect the quality of the reproduced trajectory based on the sample length. An exponential model is used to determine the relationship between the trajectory length (number of samples γ) and a τ value. Assuming that the other DMP parameters are constant, this ensures that the weighted sum of activation across all samples is above 0.5. For all trials in the Cartesian obstacle avoidance experiment, the model and its coefficients

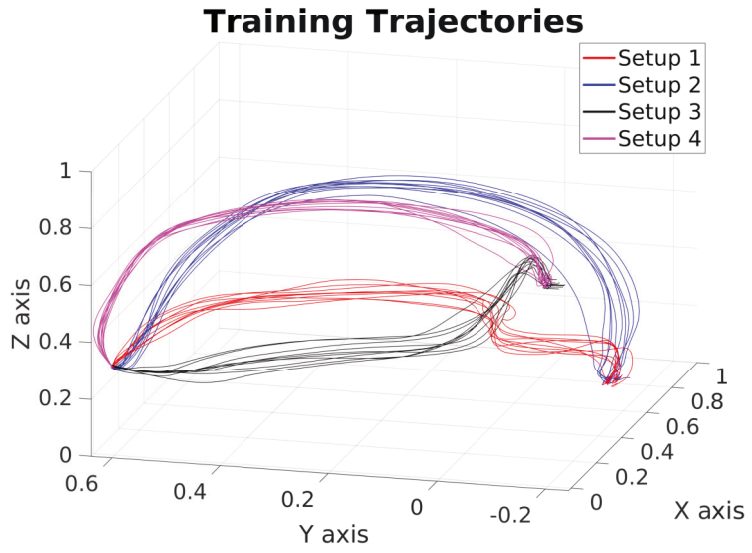

FIGURE 4.20: The training trajectories for the Cartesian target-reaching task.

used are

$$\begin{aligned} \tau &= a \cdot \exp(b * \gamma) + c \cdot \exp(d * \gamma) \\ a &= 2.22 \qquad\qquad b = -0.01292 \\ c &= 0.4684 \qquad\qquad d = -0.001595 \end{aligned} \tag{4.13}$$

For each of the 12 trials, an estimate of τ parameter is performed using the reference trajectory length.

TABLE 4.5: Measures of similarity between the reference trajectory and the trajectory generated by pHRIP. A low score indicates that the generated trajectory matches the user's intention.

	DTW Distance (unitless)		RMSE (m)
	μ	σ^2	μ
Successful	2.3482	0.2704	0.1697
Unsuccessful	3.3793	0.8941	0.2202

4.8.1 Results

Of the 12 trials, there were 5 unsuccessful trajectories, which all veered away from the intended end zone towards the end of the trajectory, as seen in Figure 4.21, despite the similar DTW scores (indicating similar trajectory shapes) seen in Table 4.5. Visual observations of the trajectories show that unsuccessful trajectories were caused by inaccurate estimates of the τ parameter. This phenomenon is evident in Figure 4.22 where the forcing function value for the unsuccessful trajectories drops to 0, causing the trajectories to

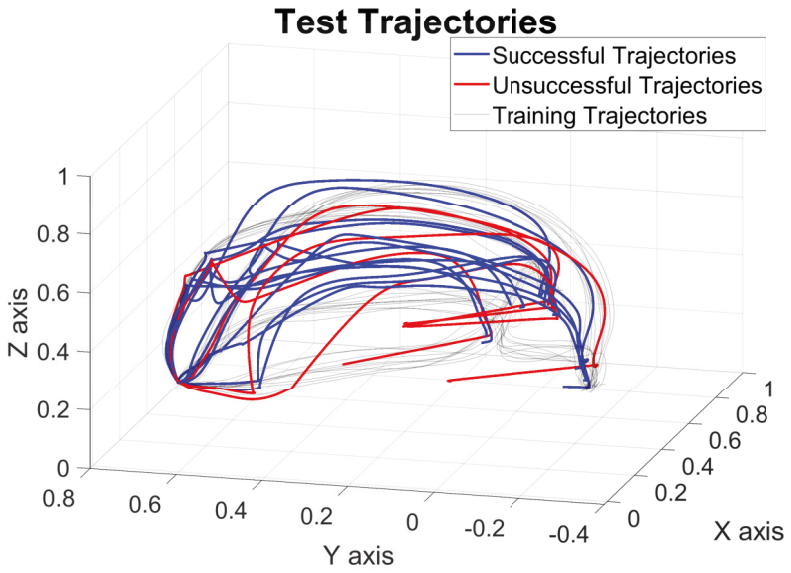

FIGURE 4.21: The test trajectories generated by pHRIP from the experiment. Successful trajectories are in blue while unsuccessful paths are shown in red.

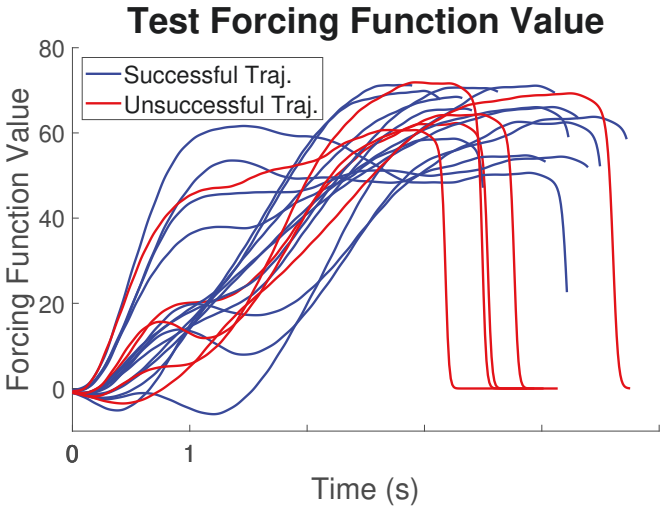

FIGURE 4.22: The forcing function values for the test trajectories generated by pHRIP during the experiment. Successful trajectories are in blue. Unsuccessful paths are shown in red.

deviate significantly. For the successful trials, pHRIP was able to correctly infer the user's intent when generating trajectories to avoid the obstacles.

4.9 Discussion and Limitations

The motivation of pHRIP is to build upon IPs by integrating physical interaction forces in a coupled system during pHRI. As IPs utilise observed trajectories, the identical location of the user endpoint and the robot endpoint precludes its use in pHRI applications. Both methods encode trajectories using DMP parameters.

A similar method for encoding trajectories is Probabilistic Movement Primitives (ProMPs) [196]. Rather than encoding the attractor landscape (also known as the "forcing function") and generating accelerations in DMP, ProMPs encodes the raw position and velocity separately using Gaussian basis functions equally spaced across the trajectory. The motivation for ProMPs is to overcome low-level robot control issues, such as acceleration profiles outside of the robot's capability, by using direct information from the position and velocities of the system.

Despite this difference, ProMPs use the same method as IPs (and pHRIP) to couple multi-DoF data together. Since the motivation for pHRIP is to extend upon IPs, the comparisons between the two (IPs and pHRIP) can be representative of the utilisation of probabilistic operators to integrate interaction forces, as performed in prior works for ProMPs [288]. Furthermore, IPs and pHRIP focus on high-level robot response rather than low-level robot control, making the comparisons more appropriate.

Other DMP-based probabilistic methods for capturing user intent include stylistic DMPs [177] and Associative Skill Memories (ASM) [289]. While both use similar probabilistic operators to capture variances across multiple demonstrations, their integration of user intent and input differ significantly from pHRIP.

Stylistic DMPs embeds a weighting factor in the Gaussian kernels when encoding the attractor landscape to provide parametric distinction of different intent. The discrete nature of the parameters precludes their use in coupled human-robot systems since there are many factors which influence the observed interaction forces. While the interaction forces could be parameterised, such as using the maximum or Root Mean Square (RMS) value, these metrics are time-dependent and are influenced by the observation window, making stylistic DMPs unsuitable.

ASM constrains user intent by attaching a Dynamic Bayesian Network to probabilistically infer the next state of the human-robot system. Integrating observed interaction forces in

this DBN can infer user intent; however, a similar issue exist for parameterising the interaction forces. Furthermore, the DBN would be more akin to hierarchical frameworks which utilise individual DMP systems where an arbitration occurs to determine the appropriate DMP to employ [182].

One limitation on pHRIP is the inheritance of DMP parameters. While they provide various functionalities, such as temporal invariance and amplitude control, the parameters (K, D, τ, and α) are still empirically obtained, with varying factors that influence which set of parameters is best suited for any particular application. As DMPs were designed for one-shot LfD, generalising multiple trajectories remains a challenge when the trajectories vary in length. This is especially pertinent when inferring intent and generating new trajectories based on partial observations, manifesting in limitations noted in Section 4.8.1.

Another constraint for pHRIP is the reliance on DTW to phase align the partial observations, requiring a reference trajectory. For a multi-modal task, it is possible to bypass this constraint by using a single reference trajectory for each trajectory mode. However, the performance of DTW degrades if there are any adaptations to the task at hand. A common approach is to use Bayesian statistics to embed the phase of the trajectory into the built distribution, such that observations can perform phase estimation and user intent simultaneously. This has been realised using Interactive ProMPs [200] and BIPs [201].

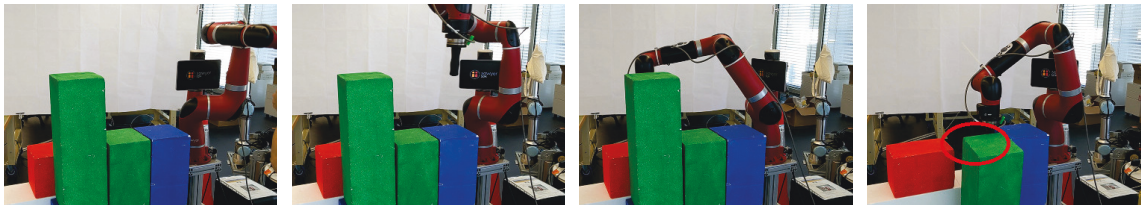

(a)*

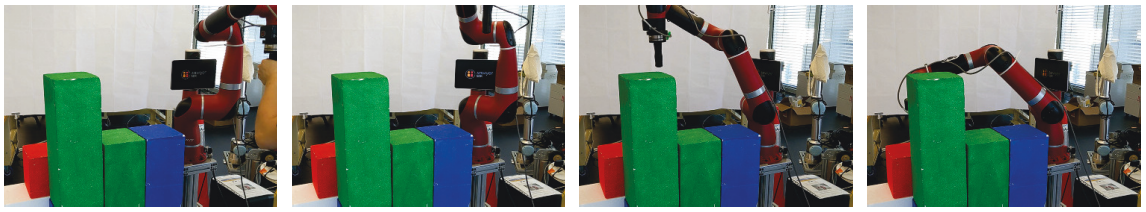

(b)

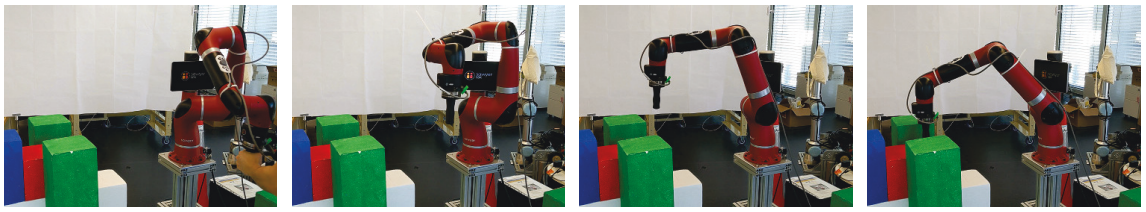

(c)

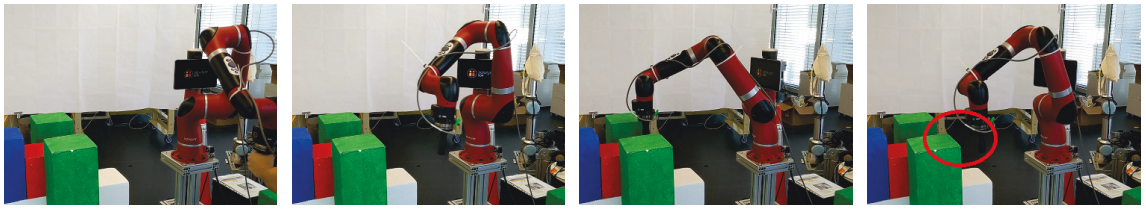

(d)*

FIGURE 4.23: The behaviour of pHRIP for known setups. * indicates an unsuccessful trial with the red ellipse showing the collisions. All collisions occurred at the end of their trajectory due to inaccurate estimates of the τ parameter as discussed in Section 4.8.1.

4.10 Summary

A preliminary study on the influence of inertia for a seminal predictive model, Fitts' law, has been presented. The results indicate that inertial influences do not affect the predictive model, suggesting that inertial influences are not significant effects which needs to be normalised during the integration of human-centric measures for personalised robotic assessment and response.

Following on this finding, an extension of a prominent HRI framework, Interaction Primitives, was presented to integrate human-centric measures into personalised robot responses. Physical Human Robot Interaction Primitives was shown to be effective in generating appropriate robot response based on the interaction forces between the human and robot. The framework overcomes limitations of Interaction Primitives for circumstances where the human hand pose is identical to the robotic endpoint. There is a caveat on pHRIP's reliance on the DMP framework; however, this is an application-dependent roadblock and is outside the scope of the thesis.

As discussed in Section 1.2, another approach towards personalised robotic assessment and response is a focus on physiological and *in-silico* models for analysing human motion and intent. This is in line with the framework which aims to connect biological human understanding to enhance the assessment pipeline. The next chapter presents this work with a study on the role of physiological models as an adjunct for musculoskeletal model-based motion analysis during pHRI.

Chapter 5

Musculoskeletal Models for Motion Analysis during pHRI

Enabling and enhancing the personalised robotic assessment and response requires an understanding of the human user during HRI. The previous chapter presented the integration of human-centric measures into conventional robotic frameworks to *enable* the appropriate personalised robot response.

While the inclusion of human-centric measures is a step forward for personalised robotic assessment and response, the information on the human user is still limited by the availability of human-centric data during physical contact. Thus, additional sources of information which inform the robotic system about the user's circumstances will enhance the human-robot system. One such tool to facilitate this enhancement is musculoskeletal models.

Musculoskeletal models provide an analogue to simulate the kinematics and dynamics borne from the interactions of a series of complex systems in the human body. Historical literature has been used to create most anatomically accurate models for humans. Model parameters are then refined using a variety of anthropomorphic studies [228] and investigations from cadaveric [290] and live human data [233].

An overview of the various software available and challenges for musculoskeletal models is presented in Section 2.4. One challenge for musculoskeletal models is the lack of validation for each model. Since kinematic measures are prioritised when developing a model, such as range of motion and moment arms, it is uncommon to validate the dynamics of the models based on the locations of the MTUs.

To include musculoskeletal models during a robotic response requires their validation to ensure the computed motion analysis results align with the human user's motion. If the results computed from musculoskeletal models do not match those of the human user during pHRI, the integration of these models during the generation of robotic responses would be counterproductive. Furthermore, any improvements to the computed results from motion analysis performed using musculoskeletal models also improves the understanding of the user, and can be employed to adapt the robotic response provided during pHRI.

Thus, this chapter first presents a study to validate the motion analysis pipeline for musculoskeletal models. Findings from the preliminary study then inform any enhancements required towards the integration of musculoskeletal models in the framework established in Chapter 3. This results in an experiment towards a grip strength model which can supplement the motion analysis pipeline, overcoming some roadblocks in the adoption of musculoskeletal models.

5.1 Preliminary Study: Upper Extremity Model

A preliminary study is conducted to validate the motion analysis pipeline for a popular musculoskeletal model as outlined in Figure 5.1. The validation compares the computed muscle activity trends against surface EMG readings that were obtained experimentally during a set of instructed movements.

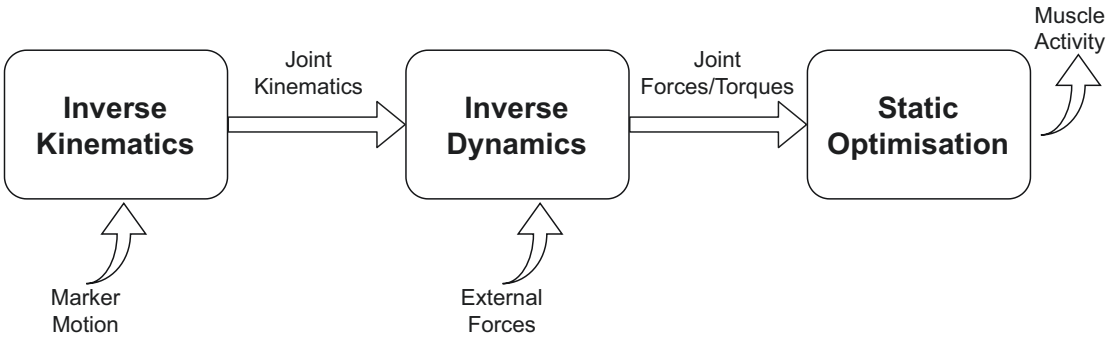

FIGURE 5.1: The motion analysis pipeline and each process' respective inputs and outputs. Inverse Kinematics and Inverse Dynamics are performed using the OpenSim API, while Static Optimisation is performed in MATLAB.

The upper extremity model used for the preliminary study is modified from an original model with 15 degrees of freedom actuated by 50 Hill-type MTUs [291]. The original model was initially validated based on the range of motion and moment arms. Most

parameters for the model were set using empirical data from magnetic resonance imaging of live healthy adults [292]. Additional modifications made to the model include:

1. Body mass and inertial properties for the humerus, radius, and capitate were changed based on literature [262], while these properties for other bodies were set as 10^{-8}.
2. The MTU models in the musculoskeletal model were updated to the more recent Millard equilibrium model [293]. Parameters for muscle activation dynamics were kept to the default values. (Activation $\tau = 0.01$, De-activation $\tau = 0.04$, fiber damping $= 0.05$)

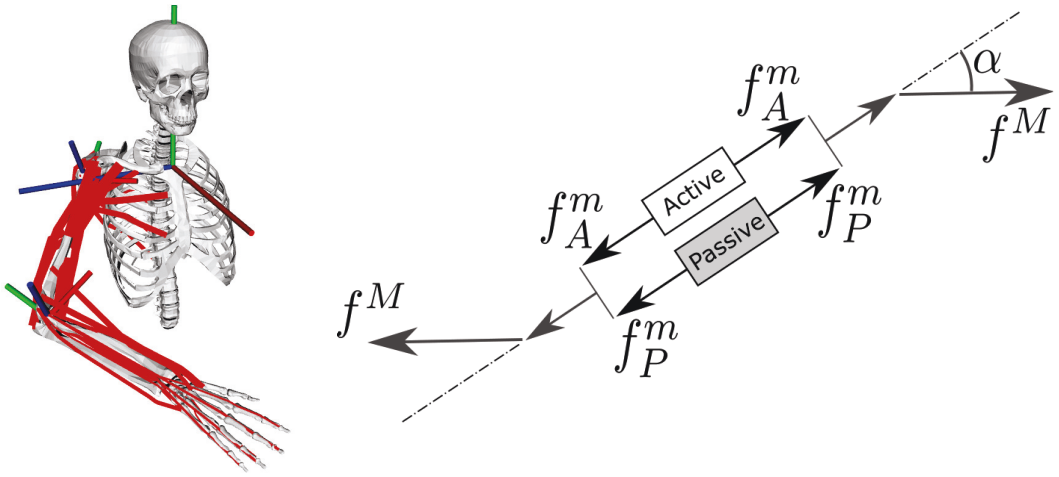

(A) (B)

FIGURE 5.2: (A) The musculoskeletal model used in OpenSim for this study. (B) A typical Hill-type MTU model with force relationships between muscles and tendons.

5.1.1 Muscular Architecture

The Hill-type MTU [294] (Figure 5.2(B)) consists of an active contractile element in parallel with a passive element. An elastic element is placed in series with the force generation mechanism representing the tendons attaching the muscle to the insertion point. The five intrinsic parameters define and normalise the muscle behaviour:

- F_0^m - Muscle maximum isometric force
- L_0^m - Muscle optimal length coinciding with F_0^m
- V_0^m - Muscle maximum contractile velocity

- α_0 - Muscle pennation at optimal length
- L_s^t - Tendon slack length.

The active contractile element force is defined as:

$$f_A^m = \tilde{\mathrm{f}}_A^m(\tilde{l}^m) \cdot \tilde{\mathrm{f}}_V^m(\dot{\tilde{l}}^m) \cdot F_0^m \cdot a, \tag{5.1}$$

where $\tilde{\mathrm{f}}_A^m(\mathbf{.})$ represents the non-linear curve between normalised muscle fibre length and active force, $\tilde{\mathrm{f}}_V^m(\mathbf{.})$ represents the curve for normalised muscle fibre velocity against normalised active force, and a is the muscle activation, ranging from 0 to 1 indicating a completely passive muscle through to a muscle with its maximum force output.

The force generated by the passive element is:

$$f_P^m = \tilde{\mathrm{f}}_P^m(\tilde{l}^m) \cdot F_0^m, \tag{5.2}$$

where $\tilde{\mathrm{f}}_P^m(\mathbf{.})$ represents the curve for normalised muscle fibre length against normalised passive force.

The total force output of the MTU is then defined by:

$$f^M = (f_A^m + f_P^m) \cos \alpha. \tag{5.3}$$

Given the m number of MTU forces in the upper limb model, $\boldsymbol{f} = [f_1^M, f_2^M, \ldots, f_m^M]^T$, the muscular torque around the k joints, defined as $\boldsymbol{\tau}^M = [\tau_1^M, \tau_2^M, \ldots, \tau_k^M]^T$, can be calculated as:

$$\boldsymbol{\tau}^M = [-\mathbf{J}_L(\boldsymbol{\theta})]^T \boldsymbol{f}, \tag{5.4}$$

where $\mathbf{J}_L(\mathbf{.})$ is the Jacobian between the MTU lengths with respect to joint position.

5.1.2 Preliminary Validation Study Setup

One healthy adult male participated in the validation trial. The participant was fitted with six 10mm electrodes (99.9% silver, Delsys, Natick, MA) on muscle groups specified in Table 5.1. sEMG data was recorded (∼1000Hz) using a data acquisition system (LabJack, Lakewood, CO). The sEMG signals were rectified before applying a zero-phase low-pass filter with 3Hz pass and 5Hz stop-band frequencies. Based on SENIAM and ISEK recommendations [295], a notch filter to remove power source signal noise was not used.

TABLE 5.1: Sensor placements of the surface EMG electrodes for participants.

Channel	**Muscle Group**	
1	Biceps Brachii	Short & Long Head
2	Triceps Brachii	Long Head
3		Lateral Head
4	Deltoideus	Medius
5		Posterior
6		Anterior

To capture the participants' kinematics, four unique motion tracking rigid body assets (each consisting of three retro-reflective markers) were affixed, each correlating to the capitate, radius, humerus, and thorax. Three-dimensional marker coordinates were recorded at 125Hz using a 12-camera motion capture system (NaturalPoint, Corvallis, USA). The dynamics of the movements were captured at 125Hz using the reaction forces from a 6-axis force-torque sensor (ATI Industrial Automation, Apex, USA). The participant held the force-torque sensor, allowing for external loads to be inserted at the lunate.

Static optimisation to obtain muscular activity was performed using the minimum sum of squared activation as the cost function and, for the sake of efficiency, only every third frame was optimised. We assume this has negligible effect on muscle activation dynamics since voluntary muscle recruitment has been found to be heterogeneous, with torque transmission delays significantly over 25ms [232]. Furthermore, from the inverse dynamic results, coordinates that experience a negligible force ($< 0.1Nm$ *or* $0.1N$) were ignored as a preliminary observation indicate that most are path points for the muscles which are under equilibrium for static optimisation. Two sets of analyses were conducted using the same data - one using ideal force generators and the other with Force-Length-Velocity (FLV) constraints [254].

For the validation trial, a series of movements were conducted with the force-torque sensor acting as a small mass. Each movement was repeated three times in series, including shoulder flexion and extension, shoulder abduction and adduction, horizontal abduction and adduction with internal and external rotations. An additional complex movement was completed, simulating a water drinking motion.

5.1.3 Results and Discussion

The resultant muscle activation trends obtained from static optimisation were compared against the six-channel sEMG readings collected during the validation trials. Figure 5.3

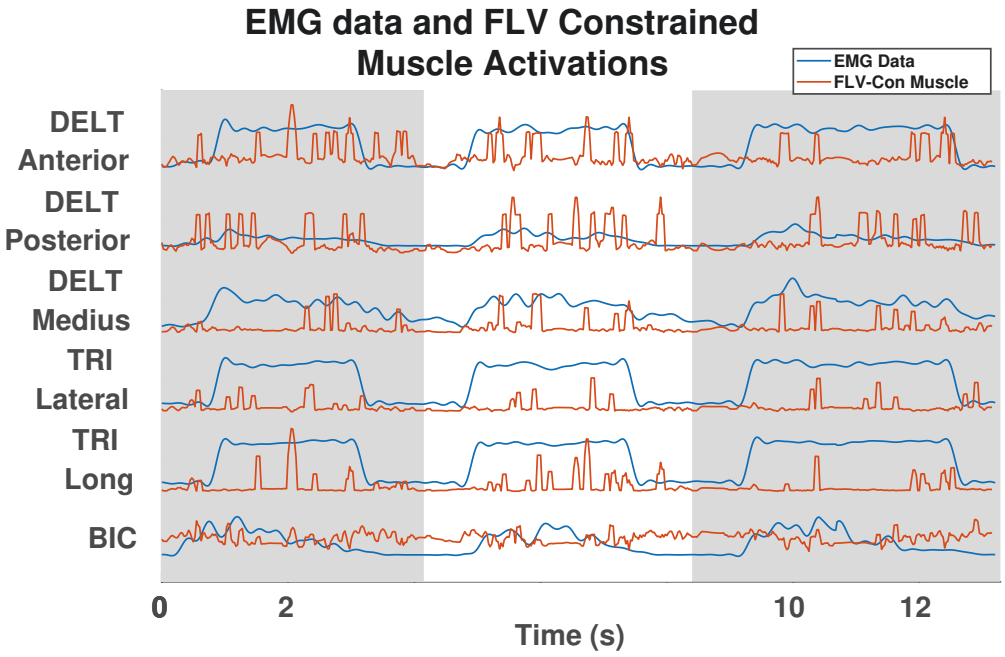

FIGURE 5.3: Results from the preliminary validation trial for shoulder flexion and extension. Static optimisation results are based on FLV constrained MTUs and EMG readings are heuristically normalised.

shows the results from the static optimisation using the FLV constrained muscle forces for the shoulder flexion and extension motion. The model treats muscle groups as homogeneous entities, despite the heterogeneity of muscle groups in real-world muscle dynamics. Combined with the equilibriation of the generalised forces, this results in the spiking nature of the muscle activation. Taking this into account, the resultant muscle activation across the six different muscle groups show that the whole musculoskeletal model matches with real-world trends. The remaining motions exhibit similar trends. Additional figures can be accessed in this repository.

Secondary results for comparisons between ideal force generators, which assume a linear force relationship, and FLV-constrained muscles can be observed in Figure 5.4. We observe that ideal muscles tend to require less activation since the dynamic properties of the MTUs typically generate forces lower than the maximum isometric force.

There are a few factors that may affect the generalisation of our findings, especially for different participants. Since we are treating real-world sEMG data as the ground truth, the evaluations might be affected by external influences such as the setup protocol, electrode placement, and inherent variability shown during the movements. One common method to overcome this is to normalise muscle activation by obtaining the Maximum Voluntary Contractions (MVC). However, this would require the participants to perform activities

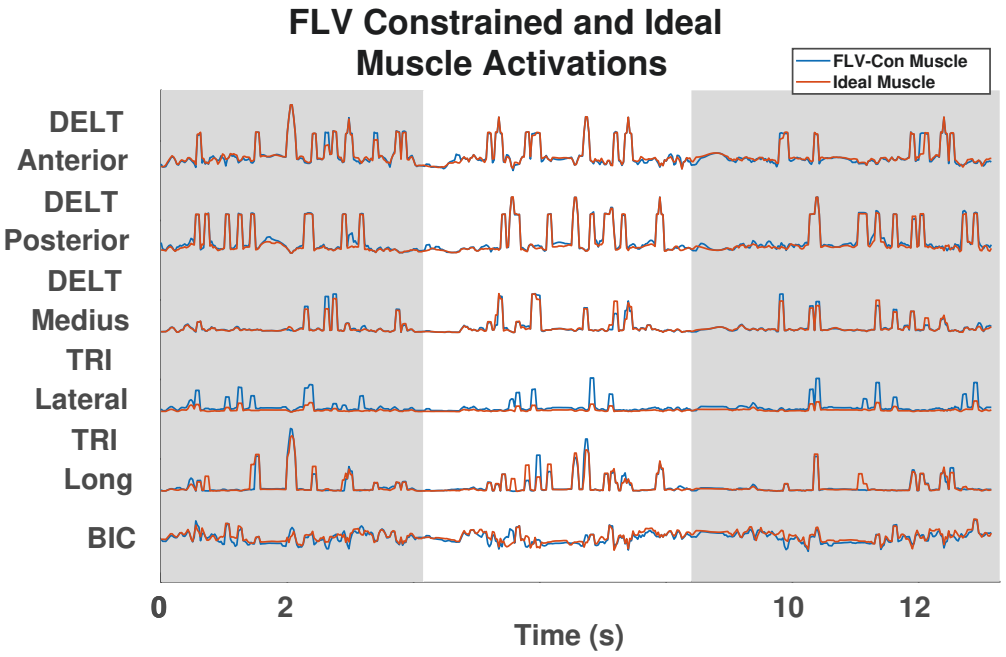

FIGURE 5.4: Results comparing between muscle activation from ideal and FLV constrained force generators.

to identify their unique MVC values. It is noted that during this process, the external influences induced by the setup protocols and electrode locations still exist.

Another factor that affects our results come from the lack of model parameters. Since accurate mass and inertial properties for particular bone segments are difficult to obtain, most models leave these parameters as zero, and the resultant muscle activation does not account for this. Furthermore, static optimisation assumes the system state is in equilibrium, meaning it does not consider other dynamic forces in continuous states such as Coriolis effects and muscle-tendon dynamics.

One limitation of performing numerical analyses with a redundant model is the inability to account for co-contracting muscle pairs, which perform other functions during human motion such as the regulation of body segment impedance [296], accuracy [297], and task variability [95]. To highlight this effect, the participant voluntarily co-contracted their arm during the same water drinking motion. The resultant difference between the optimised muscle activation and the sEMG data can be seen in Figure 5.5, highlighting the limitations of these solutions when using the musculoskeletal model.

While the use of a minimisation-based objective function creates this limitation, the redundancies between the number of MTUs and coordinates in the model makes it difficult

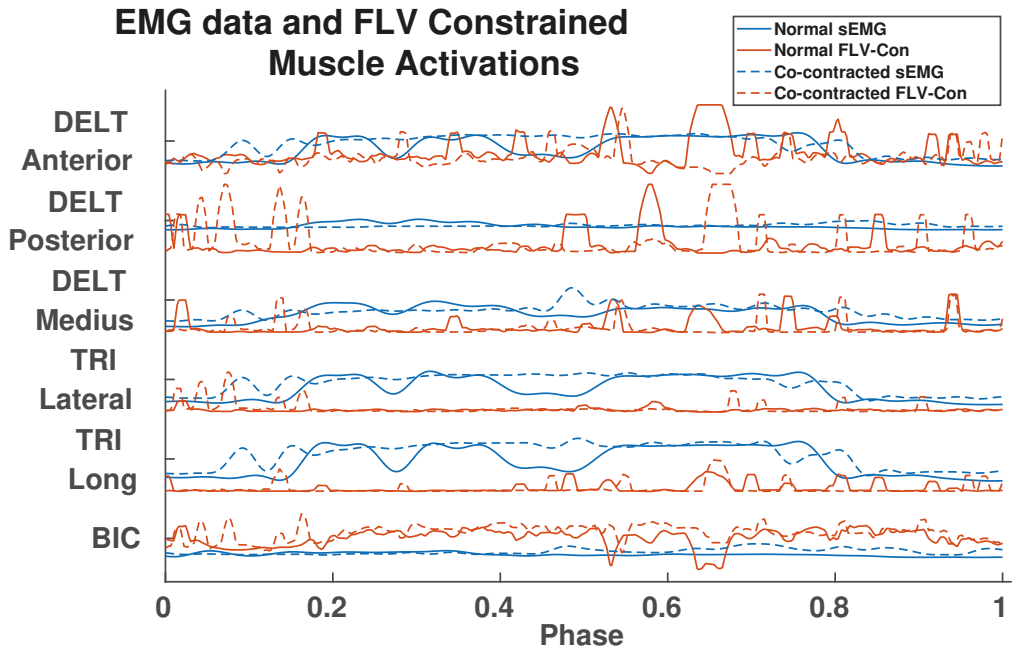

FIGURE 5.5: Results from a water drinking motion in voluntary relaxed and co-contracted states.

to develop sensible objective functions which are able to capture *all* variances in muscle activity. In particular, co-contractions in antagonistic muscle pairs has long been thought to be critical in controlling human body impedance. As a result, translating this phenomenon into the optimisation process requires additional sources of information.

5.2 Supplementary Models for Musculoskeletal Models

With multiple confounding factors affecting human motor control [238], learning [246, 249], and planning [74, 208, 298, 299], musculoskeletal models have been invaluable for researchers to isolate and investigate their influence in humans. However, the preliminary study presented above identified the need for additional sources of information to inform the optimisation process when performing human motion analysis using musculoskeletal models. With no closed form solutions available to obtain muscular activity for a particular state, the common approach utilises optimisation-based solutions. However, there is ongoing debate on the appropriate cost function to be used during this optimisation, with arguments for [255, 257] and against [256] using principles of optimal control.

The most widely accepted cost function to determine muscle activation relies on the principle of minimum energy expenditure [300], deriving the minimisation of the sum squared

muscle activation. However, there is evidence to suggest that this is not as relevant even for elite athletes whose successes rely on the optimisation of their actions [259]. An approach based on minimum energy expenditure also does not compensate for muscular deficiencies as seen in the elderly [301] and a deafferented subject [302]. Furthermore, muscular co-contraction, a critical characteristic for human motor learning [303, 304], is not considered when using minimum energy expenditure principles.

One approach to address this constraint is the use of supplemental models which bridge the gap for the muscle activation of particular muscle groups. These models are spurred on by technological advances in data collection, resulting in correlation models between muscle fibre volume and maximum isometric force generation [305, 306], muscle deformation and force output [307], and between sEMG and muscle force onset [232, 308]. However, the dynamics of the muscular architecture is still complex with many confounding factors, which affect the relationship between muscle activity and force generation [309–312].

In coupled human-robot dyads during pHRI, the physical interaction between the robot and user generates data on human grip strength, which can provide valuable information on the physiological state of the human. Human grip strength has been well studied with excellent inter-rater and test-retest reliability [313], pressure distribution dynamics [314], and significant studies to obtain demographic grip strength data in various parts of the world [315, 316]. Furthermore, there are seminal studies that have highlighted how the somatosensory contributes to grip strength control [302, 317], and other factors influencing the resultant grip strength [318].

In pHRI applications, the grip strength has been used as an indicator of human physiological response to perturbations [319]. The use of grip strength provides the opportunity for a model which maps the relationship between grip strength and muscle activation for the relevant muscle groups. Finger flexion and extension is attributed to the *flexor digitorum profundis* (FDP) and *extensor digitorum communis* (EDC) muscle groups [318], with the FDP muscle group being the primary contributor to human grip force.

As a result, a model is developed to correlate the relationship between grip strength and the muscular activity of the contributing muscle groups. While sEMG data are not true reflections on muscle activity with various factors confounding their reading, such as electrode application procedure, environmental humidity, and dermal impedance at electrode locations, the model will rely on its activation trend to correlate grip force as an analogue of grip strength. This model is then integrated into the motion analysis pipeline to partially overcome the over-actuation of the musculoskeletal model. An analysis of the

influence of the grip strength model is then presented through a popular musculoskeletal model with some modifications.

5.3 Experiment Methods

An experiment was set up to develop a grip strength model which maps the correlation between surface EMG muscle activity and grip force. Ten male adults, presenting no neuro-muscular disorders, participated in a series of double blind randomised trials to obtain their grip force at three distinct levels: 33%, 66%, and 100% of maximal voluntary grip. The experiment protocols were explained to the participants before obtaining their informed consent. This research has been carried out according to the ethical guidelines of the Human Research Ethics Committee in University of Technology Sydney (UTS HREC approval no.: ETH18-3029).

A set of standardised verbal instructions were provided to participants, with participants instructed to maintain 90° elbow flexion, no radial/ulnar deviation of the wrist, and a natural grasp on the handheld dynamometer. Across all participants, the estimated wrist extension was within ±15° with no significant changes throughout the experiment.

The experiment consisted of one reference trial, 12 isometric trials, and 6 exercise trials. The grip force exerted by each participant was measured by a digital dynamometer ($\pm 1N$). The measurements from the digital dynamometer were captured using video footage and extracted manually post-hoc. Administrators advised participants against reading their grip force results after each trial, although no physical restrictions were employed to enforce this. For all trials, the maximum grip force measurement was collected. This aligns with the sEMG data processing pipeline (Section 5.3.4) to extract the maximum RMS values.

In the reference trial, participants were instructed to "squeeze as hard as possible for 5 seconds while maintaining your trunk and arm posture". They were further instructed to "use this [*reference trial*] to relate to the 33% and 66% grip strength level for the rest of the experiment". No further feedback was provided to the participants for the rest of the experiment, and participants approximated the 33% and 66% grip strength level to the best of their ability. Since no feedback was provided to participants throughout the experiment, initial analysis of the reference trials show that the majority of participants exert supra-maximal grip force during the isometric trials. Thus, the results from the reference trial were not used in post-hoc analysis.

TABLE 5.2: The targeted muscle groups for the placement of the sEMG electrodes.

Electrode Number	Muscle Group	Muscle Acronym
1	Extensor Digitorum Communis	EDC
2	Flexor Digitorum Profundus	FDP
3	Biceps Brachii	BIC
4	Triceps Brachii	TRI
5	Deltoideus Posterior	DELT1
6	Deltoideus Medius	DELT2
7	Deltoideus Anterior	DELT3

5.3.1 Participant Measurements and Sensor Placement

The length of each participant's ulna and humerus were measured using identical protocols as those from the CDC (US) [320] and the NHS (UK) [321]. Muscular bio-potentials are collected through silver-silver electrodes using the Delsys Bagnoli 8-Channel system (Delsys, Natick, USA). Electrodes are positioned based on SENIAM guidelines[1], targeting muscle groups listed in Table 5.2.

Motion capture software was utilised to track the kinematics of each participant's upper limb. A total of 12 Optitrack Prime cameras, in conjunction with the Motive software (NaturalPoint, Corvallis, USA), tracked four user-defined rigid body assets as shown in Figure 5.6. Each asset consists of 3 passive reflective markers with known locations to provide unique poses. The rigid body assets are placed to coincide with the four major body segments of the model: chest (thorax), upper arm (humerus), forearm (ulna), and hand (lunate).

5.3.2 Isometric Trials

For isometric trials, participants were requested to maintain the arm and trunk posture as shown in Figure 5.6. In each trial, participants were provided with a 5 second visual countdown, after which participants were instructed visually to squeeze the dynamometer for 5 seconds at 3 distinct strength levels: 33%, 66%, and 100%. Between each trial, participants were given a one minute break to mitigate muscular fatigue.

In total, each participant performed 4 trials per grip strength level, with the order of the 12 trials randomised using Latin squares. The requested grip strength level for each trial

[1] http://www.seniam.org/

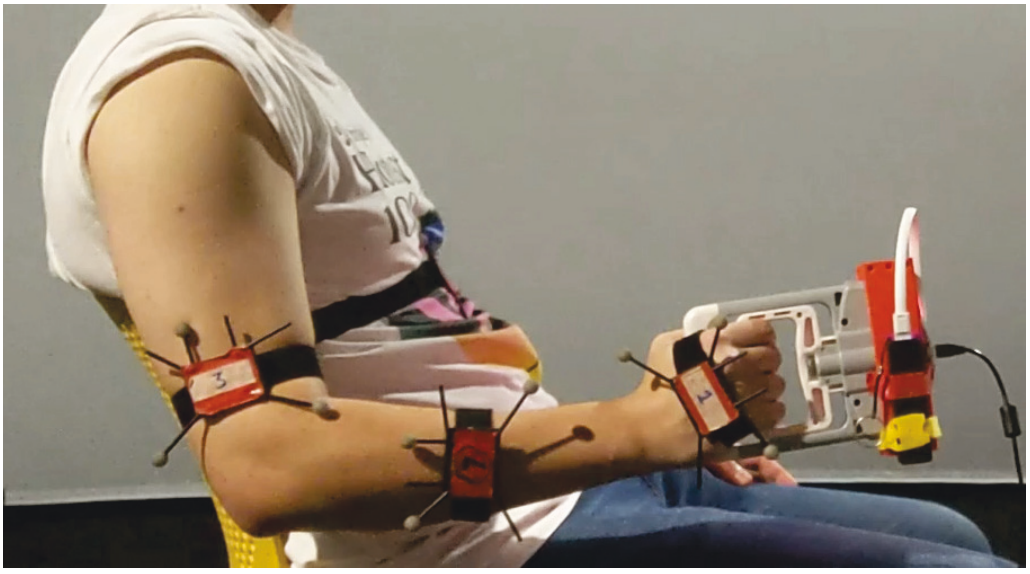

FIGURE 5.6: The posture that participants were requested to maintain during isometric trials.

was occluded from the administrator by using a participant-facing screen not visible to the administrator. For each isometric trial, sEMG data collection commenced 2 seconds into the countdown, resulting in 8 seconds worth of data each trial.

5.3.3 Exercise Trials

At the start of the exercise trials, participants were instructed to watch a video which shows a reaching exercise[2]. The exercise trials were performed in a darkened room to reduce reflections from the passive reflective markers. Thus, LED light strips were diffused and attached around the LCD display of the digital dynamometer to facilitate the reading of the grip force measurements on the video footage (see Figure 5.7(A)).

Similar to isometric trials, visual instructions for the grip strength levels were given to participants, and they were verbally instructed to "squeeze at the indicated grip strength level and start your [*sic*] movements while maintaining your [*sic*] trunk posture". No further requirements, such as task completion speed or accuracy, were imposed on participants.

Kinematic data for participants are collected from the exercise trials using the rigid body assets, starting when the countdown commences. Since there were no task completion speed requirements imposed, participants were instructed to indicate to the administrator verbally when they "feel like you [*sic*] have completed the exercise from the video". Similar

[2]The exercise video shown to participants is available here

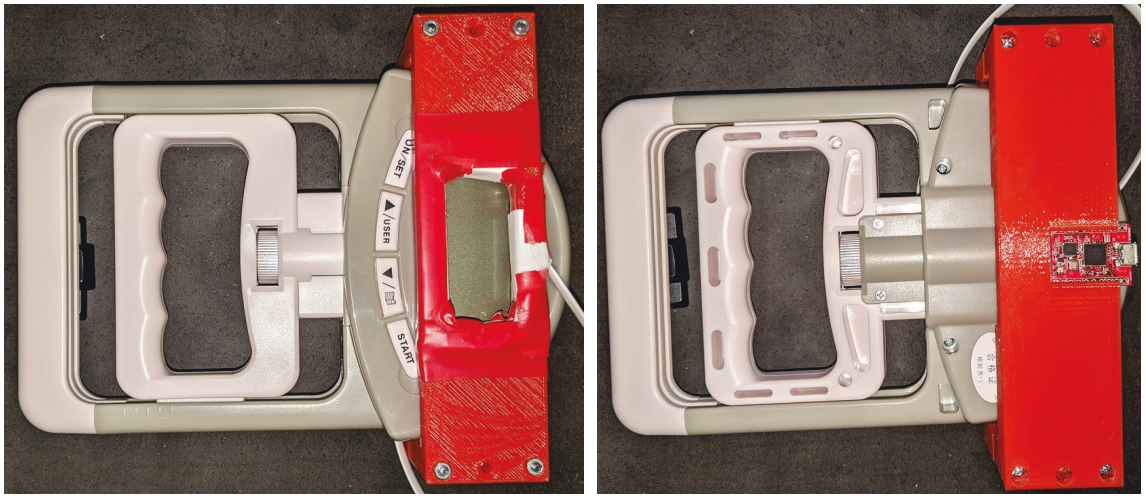

(A) (B)

FIGURE 5.7: The front and back views of the digital dynamometer fitted with a 3D printed part to house (A) the diffused LED light strip, and (B) the attitude and heading reference system (myAHRS+).

to the isometric trials, sEMG data collection commences 2 seconds into the countdown. Both sEMG and kinematic data collection is stopped manually by the administrator when they receive each participant's verbal indication.

Both kinematic and sEMG data from the exercise trials are used for post-hoc analysis with the musculoskeletal model. External forces exerted on the participant's hand are calculated from linear acceleration readings given by an attitude and heading reference system (WITHROBOT, Seoul, Korea), along with the mass of the sensor-equipped dynamometer ($m = 0.516kg$). Forces induced from gravity are also included in the dynamic analysis of the musculoskeletal model (Section 5.3.6).

5.3.4 sEMG and Kinematic Data Processing

The collected sEMG signals are rectified, before applying a 4th order Butterworth lowpass filter with a cutoff frequency of 4Hz. The first and last 0.5 seconds of the processed sEMG signals are ignored to remove artefacts arising from observations of partial phases.

The grip strength model aims to build a mapping between muscular activity and the grip force exerted by each participant. The sEMG signals are normalised to allow for quantitative analysis between participants. A common method to normalise sEMG signals is to conduct Maximum Voluntary Contractions (MVC) trials, exercises which aim to maximise the muscle activation of specific muscle groups. Different exercises are prescribed

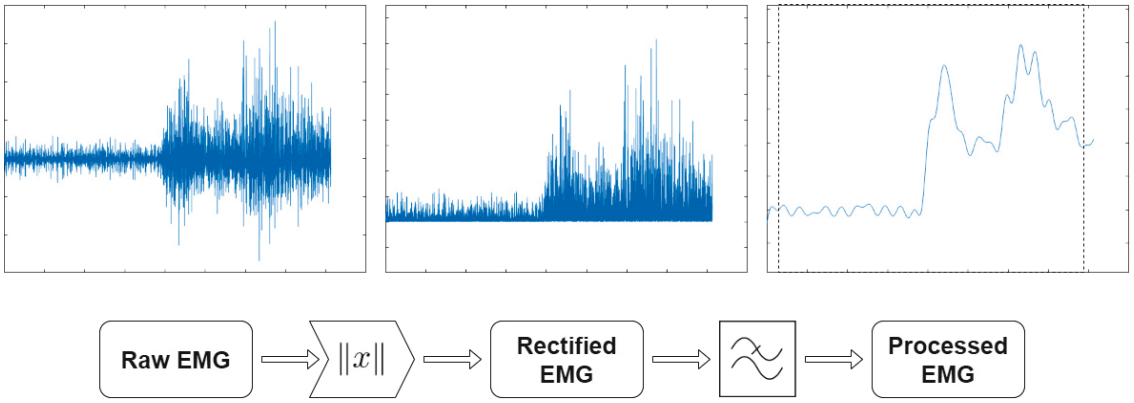

FIGURE 5.8: An illustration of the sEMG processing pipeline.

for different muscle groups despite ongoing work to determine the optimal set of exercises for inter-rater reliability and MVC values [322].

Using default MVC values to normalise the sEMG data, 6 out of the 10 participants demonstrated supra-maximal sEMG readings during the experiment. Thus, a different normalisation method based on the sEMG intensity range is used. Since participants were instructed to hold onto the handheld dynamometer (0.516kg) in a relaxed manner, the sEMG intensity for muscle activity during this time is assumed to be the minimum value.

The sEMG intensity is obtained by partitioning the signal into rolling windows which are 0.25s long (0.05s apart) and calculating the RMS value for each window. Since the experiment includes an exercise which is suitable for MVC (e.g. 100% grip strength), each participant's normalisation range is formed by the minimum RMS value between 0.5-3.0s and the maximum RMS value across all trials.

For all other targeted muscle groups, the sEMG signals are not processed since they will be used for qualitative analysis with the static optimisation results. An overview of the sEMG processing pipeline can be seen in Figure 5.8.

5.3.5 Grip Strength Model

To build the grip strength model mapping the relationship between grip strength output and muscular activity, we use isometric trial data from all participants to generate the model and validate it with the exercise trials. To build a model which is valid across all participants, the sEMG intensity and the grip force are range-normalised. While muscle force generation is constrained by their Force-Length-Velocity curves [254], which affect

muscle moment arms, participants were instructed to maintain a static posture during the isometric trials (Section 5.3.2), making the relative effect of the FLV curves negligible.

For the exercise trials, the grip strength model does not take FLV curves into account which may skew the results. An exploration of the effects from the FLV curves was conducted to ascertain the scale of this effect across different poses and velocities as extracted from participant kinematic data. Results for this analysis are shown in Section 5.5.

One popular model hypothesised for the relationship between muscular activity (% normalised EMG) and force generation (% MVC) is the two-element model by Woods and Bigland-Ritchie [253]. Figure 5.9 shows the relationship that suggests that there are intrinsically two types of muscle fibres during muscle activation. This model is supported empirically [308], while [75] identified that the parameters of the curve are unique to each muscle group. Assuming that the slow Type 1 fibres generate force asymptotically at 30% of MVC, the Woods and Bigland-Ritchie model can be parameterised using piece-wise equations to isolate the two fibre types.

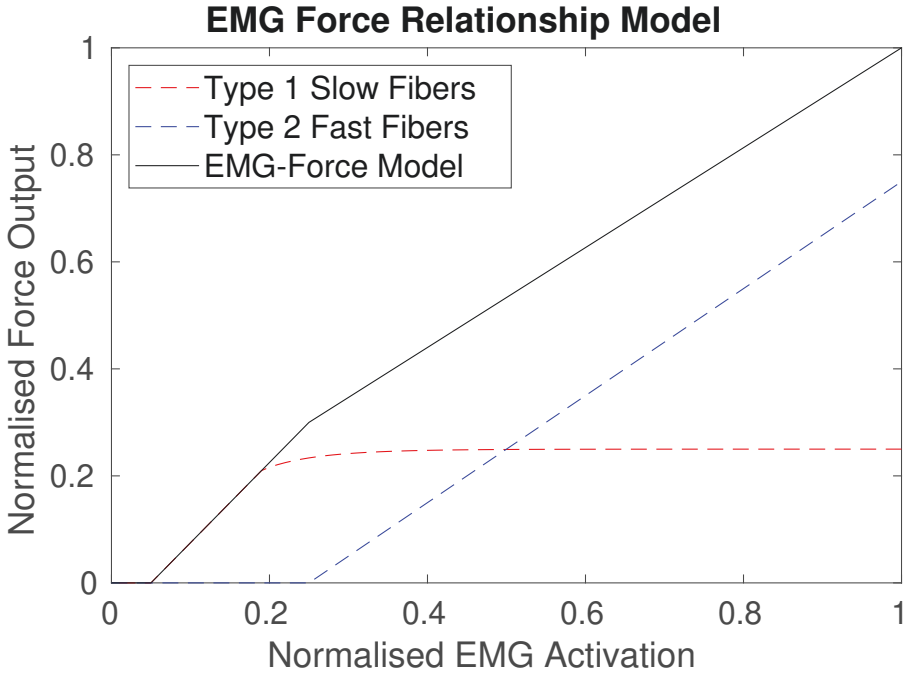

FIGURE 5.9: The two-element EMG-Force model proposed by Woods and Bigland-Ritchie used for the grip strength model (reproduced from Woods et al. (1983)).

5.3.6 Musculoskeletal Model

The upper limb musculoskeletal model used for the experiments is derived from the Stanford VA Upper Limb model [291] with updated model parameters from live human data [292]. The model consists of 15 degrees of freedom and 50 Hill-type MTUs [294]. Modifications were made to the model to simplify the model while maintaining bodies, coordinates, and actuators relevant to the experiment. Updates to the model[3] include:

- Inertial properties of the humerus, radius, and capitate were set using anthropomorphic estimation for the average male height and weight [206, 228]. Principal moments of inertia for the body segments are taken from [230].
- The original MTU model from [323] is replaced by the more recent Millard equilibrium [293] model. Available MTU model parameters are transcribed across to the new MTU model (maximum isometric force, optimal muscle tendon length, pennation angle at optimal length, tendon slack length). Other MTU parameters (Force-Length-Velocity curves) are set to the default values in OpenSim.
- Model complexity is reduced by restricting motion to 5 DoFs. These DoFs arise from the shoulder flexion/extension, abduction/adduction, internal/external rotation (3), elbow flexion/extension (1), and wrist flexion/extension (1). The remaining joints are locked to match the experiment posture (e.g., closed digits, neutral pronation/-supination, and 0° wrist abduction/adduction). MTUs which do not contribute to the 5 DoFs are disabled by ignoring them during calculations. These are determined following the process for [324]. MTUs contributing to grip strength are exempted from the modification (EDC and FDP muscle groups).

For each participant, the measurements of their humerus and ulna are used to manually scale the geometry of the model. The length of the humerus and ulna matches their respective measurements, while the length of the radius is proportionally scaled based on the ulna measurement.

For each exercise trial, upper limb motion analysis is conducted by performing inverse kinematics, inverse dynamics, and kinematic analyses. All three operations utilise the functionalities in OpenSim. The generalised forces obtained from inverse dynamics are used for static optimisation in MATLAB (using the *fmincon* function) to obtain an estimate of the muscular activity, $\tilde{A}$, during the exercise trials. A minimum sum of squared

[3]The updated model is available here.

activation cost function is used:

$$\begin{aligned} \min_{a} \quad & \sum_{i=1}^{M} a_i^2, \\ s.t. \quad & M(q)\tilde{A} = F, \end{aligned} \tag{5.5}$$

where $M(q)$ is the moment arm matrix, which depends on the current joint state, F is the net generalised forces on the model joints, and $\tilde{A} = [a_1, a_2, a_3, a_4, ...]^T$ is the activation level $(0 \rightarrow 1)$ for the M MTUs in the model.

TABLE 5.3: The replacement of the MTU activation in the musculoskeletal model with the model estimate.

Muscle Group	Model MTUs
EDC	EDCI, EDCL, EDCR, EDCM
FDP	FDPI, FDCL, FDPR, FDPM

For the purpose of demonstrating the feasibility for integrating extrinsic information channels with musculoskeletal models, the muscle activation for FDP and EDC are assumed to be identical (assuming isometric co-contraction since they do not contribute towards the movement of the upper arm). The muscular activity estimate is obtained using the normalised EMG readings for the FDP muscle group. Both values replace multiple MTUs in the musculoskeletal model as shown in Table 5.3. The generalised forces caused by the EDC and FDP activation are taken into account when performing optimisation for the remaining MTUs in the model.

5.4 Results

5.4.1 Contribution of Muscle Groups

The generation of human grip strength is attributed to the muscle groups which control finger flexion and extension: FDP and EDC. To ascertain the influence of each muscle group on the user-generated grip force, a series of one-way Analysis of Variance (ANOVA) was conducted between each muscle group and the measured grip force.

Only the subset of isometric trials, where participants were instructed to exert at 100% grip strength, were used in this particular analysis. Isometric trials for 33% and 66% grip strength are assumed to influence muscle activity as participants try to control and maintain the requested grip strength level. The need to maintain the requested grip

strength level is posited to result in a larger spread of muscle activation despite a small spread of grip force. Observations of the experimental results in Figure 5.10(C-D) confirms this finding.

The normalised EMG activation of FDP and EDC are shown to be statistically the same ($F(1, 54) = 1.0570, p = 0.3085$), with no significant difference during the 100% grip strength trials. This is reinforced by the Pearson and Spearman's Rank Order correlation score, presented in Table 5.4, which show a similar correlation between each muscle group and the measured grip force. With a similar correlation between measured grip force and FDP or EDC, a single-variable model is chosen. Since the sEMG signals are range normalised, the choice to use the FDP muscle group was based on a higher sEMG range during contractions. Using a higher range mitigates the effects of signal noise, alleviates

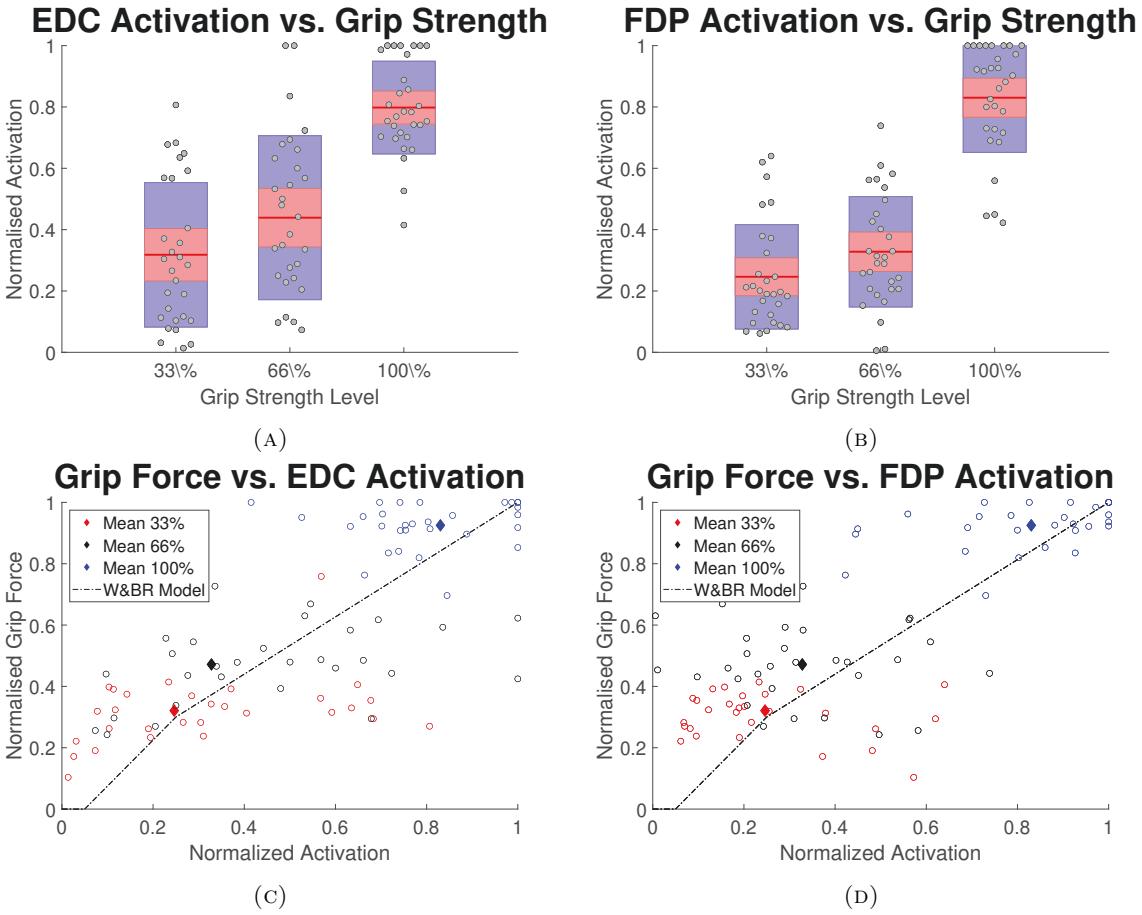

FIGURE 5.10: The relationship between the grip force measurements and muscle activation for all participants in the experiment. The EMG-force model is given as a reference. It is noted in (C-D) that there is a small normalised grip force spread (y-axis) despite a large normalised activation spread (x-axis).

the chances of supra-maximal and negative EMG activation, and can cater to different participants with a variety of muscular strength.

TABLE 5.4: Correlation scores between EDC, FDP, and grip force. Pearson's correlation scores are in the top right triangle, while the lower left triangle shows the Spearman's Rank Order correlation score.

		EDC Muscle	FDP Muscle	Grip Force
		Pearson's r		
EDC Muscle	Spearman's ρ		0.5961	0.7280
FDP Muscle		0.5644		0.7675
Grip Force		0.7203	0.6776	

5.4.2 Grip Strength Model Validation

During the experiment, a subset of data was not used due to technical issues. For some trials, the video footage of the experiment was underexposed, leading to an inability to read the measured grip force ($\sim 14\%$) from the video footage. The sEMG readings for two participants were observed to have no visible artefacts from muscle activity present across all trials and were not used in further analyses. The lack of visible artefacts derive from a flat-line response from the sEMG readings. This indicates an erroneous application of the electrodes during the setup process. Additionally, an assumption is made that the relationship between EMG activation and muscle force output is monotonic. Based on this, further visual inspection of the data from isometric trials (Figure 5.11(A)) highlighted potential outliers which may suggest deficiencies in the experiment setup. An initial analysis was conducted to determine the impact of the outlier data identified.

Visual observation of grip force vs. FDP activation indicate that an additional two participant's data are outliers (Figure 5.11(A)). Building a step-wise linear regression model (*stepwiselm* function in MATLAB) from the remaining data, the RMSE between the isometric trial data and the model was calculated. The error was more than 3 times the standard deviation of the mean RMSE. Building the grip strength model with the remaining participants resulted in a mean RMSE of $\mu = 0.1512$ with a standard deviation of $\sigma = 0.1207$.

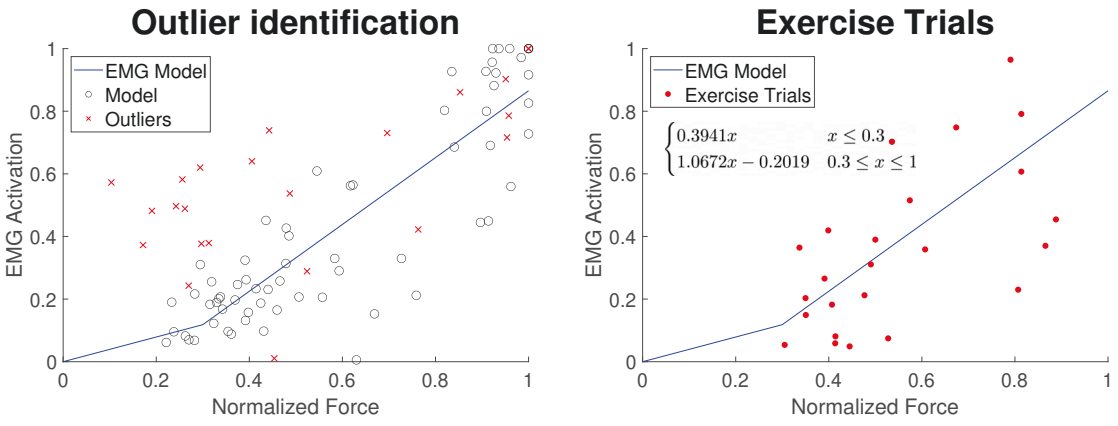

(A) (B)

FIGURE 5.11: Scatter plots of: (A) the sample points used for building the model and the outliers identified visually; and (B) the sample points from valid exercise trials to compare the fit of the grip strength model.

Visual observation in Figure 5.11(B) suggest that the data from participants' exercise trials would fit the model. The validation of the grip strength model was performed using the results from the exercise trials, showing a mean RMSE of $\mu = 0.2035$ with a standard deviation $\sigma = 0.1207$. While there is evidence to suggest a power curve would fit a similar relationship [314], this was not investigated in this work. A comparison between the instructed grip strength level and the measured grip force suggests a similar relationship. However, there are multiple confounding factors which can be attributed to the observed trend in Figure 5.12, which is outside the scope of this thesis.

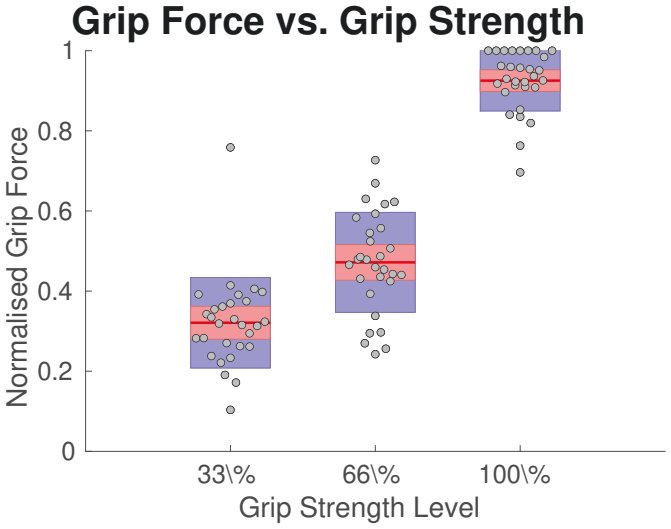

FIGURE 5.12: A modified box plot of the instructed grip strength level and the measured grip force for all participants. The mean is indicated by the red line, the standard deviation is shaded in red, and the 95% confidence interval is shaded in blue.

5.5 Discussion

From the experiment's sample size, the generated grip strength model is a feasible alternative for obtaining an estimate of the EMG activation based on a measured grip force. While the FLV constraint effects for MTUs are well understood [254], empirical evidence suggest that the Woodland and Bigland-Ritchie model can be used to correlate EMG activation and grip force.

Validation of the model is performed using the exercise trials which involve movements. To investigate FLV constraints, a comparison of the generated moment arms is conducted in three different phases of a reaching movement: the starting pose, the end pose, and a pose in-transit. Across the three poses, the mean absolute difference of moment arms (41 muscles on 94 coordinates including the FDP and EDC muscle groups) between the starting pose and the other two are $\mu_{end} = 1.6913$ (3.76% of total) and $\mu_{transit} = 1.2963$ (2.82% of total).

5.5.1 Effects of Muscle Fatigue

In the experiment, participants were provided with a minute break between each trial. However, the amount of time provided to participants in prior works have varied widely from 90 seconds [322] up to 180 seconds [253]. Thus, an investigation was conducted post-hoc into the possible effects of muscle fatigue in the results of the experiment. While there are multiple ways to extract muscle fatigue [325], the most common method is using Power Spectrum Density (PSD) plots and investigating whether there has been a shift in the peak power frequency.

The raw EMG signals from the experiment were analysed in MATLAB using the *fft* and *pwelch* functions for Fast Fourier Transform (FFT) and an estimate of the Welch PSD. Gaussian smoothing with a window size of 40 was performed on the Welch PSD to obtain clear peaks of the PSD curve. For each trial, the EMG signals were segmented into three temporal sections, with the first section starting 0.5s after the visual instruction was provided to the participants.

The PSD plots across the three temporal sections for all participants in Figure 5.13 indicate that acute effects of muscle fatigue are present within each trial. As participants are instructed to exert grip force for at least 5 seconds each trial, this phenomenon is expected.

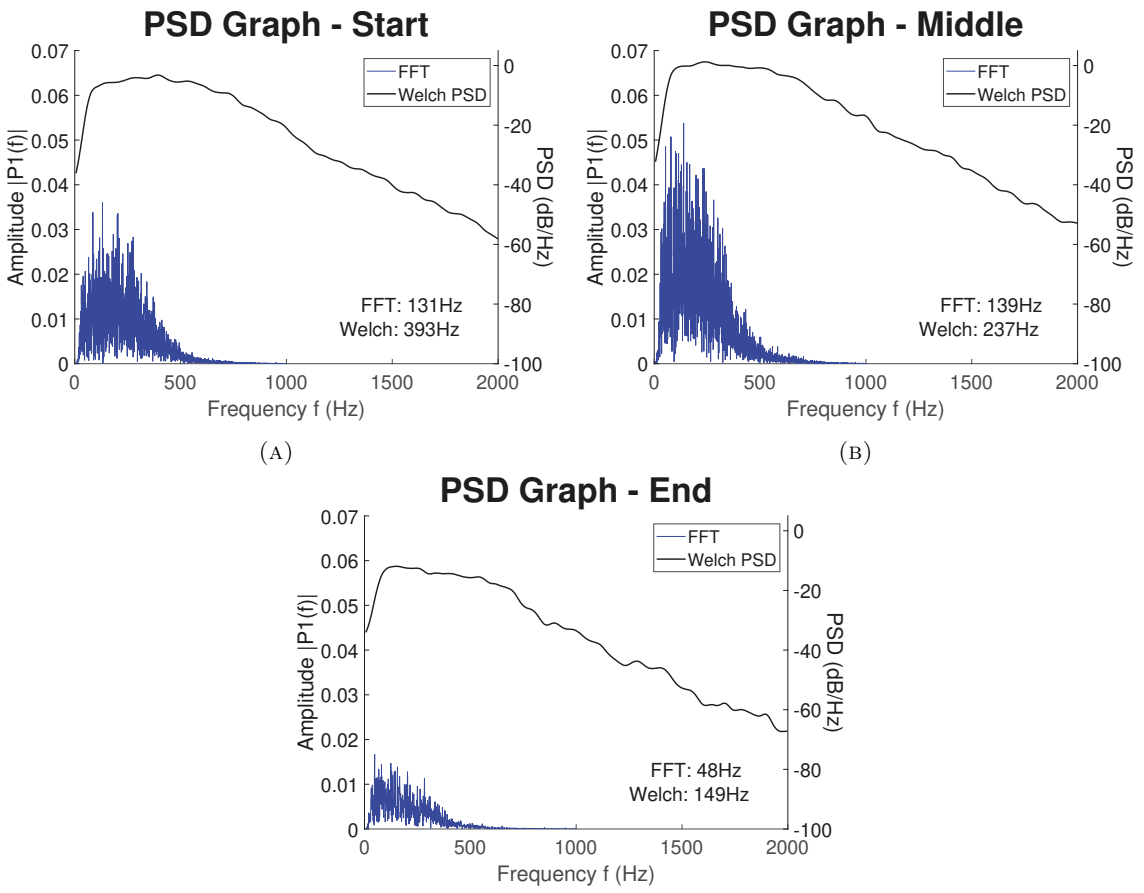

FIGURE 5.13: The Power Spectral Density plots across (A) the first third, (B) second third, and (C) final third of a trial. The values in each plot indicate the mean frequency for FFT and Welch PSD. The trend from these plots are representative across most of the isometric and exercise trials.

Thus, to assess longer-term effects of muscle fatigue, only the first temporal section of each trial was used for analysis.

For each participant, the first and last two 100% strength level trials (across both isometric and exercise trials) were used to perform a one-way ANOVA. For both FFT-derived and Welch PSD estimates, there was no significant shift in the peak frequency (FFT: $F(1, 38) = 0.9943, p = 0.3250$, Welch: $F(1, 38) < 0.001, p = 0.9823$), indicating that the effects of muscle fatigue are negligible in the experiment. It is interesting to note that the middle third of the trials represent the highest amplitude. This correlates with observations of the sEMG readings that indicate muscle activation ramps up towards the peak intensity, rather than a quasi-instantaneous response. The ramping response suggest that participants may have been regulating their muscle activity until they feel like they

have achieved the desired grip strength level. Alternatively, this artefact may be due to muscle activation dynamics since they are commonly modelled as having first order delay after excitation. However, this finding doesn't skew the objective of the grip strength model which relies on peak measurements.

5.5.2 Supplementary Information Sources for Musculoskeletal Models

The validation of the grip strength model demonstrates the feasibility of supplementary information sources during physiological-based assessments of the human in pHRI. Opportunities to personalise the model exist, which is an enticing proposition in rehabilitation and assessment [15, 40]. The grip strength model aims to contribute towards the current trend in generating supplementary information channels that are adjunct to musculoskeletal models. This will encourage their adoption in emerging fields such as rehabilitation [277], assessment [326], and pHRI [93].

While the grip strength model provides additional information sources, it may introduce bias in the muscular activity analysis. Realistic musculoskeletal models are particularly sensitive to this possibility as the upper limb complex has complicated muscle tendon paths. A series of static optimisation (using MATLAB's *fmincon* function) was conducted to investigate the influence of the grip strength model's estimates in the resultant muscle activity.

The chosen objective function for the optimisation follows the default function used in OpenSim, the minimum sum squared value of the muscle activation (Equation 5.5). Due to the use of minimisation during the optimisation process, the resultant muscle activity is generally underestimated when compared to empirical data. The discrepancies seen in Figure 5.14 for the FDP and EDC muscle groups can be attributed to the range normalisation process for the sEMG readings and implicit errors in the process of data collection. Furthermore, the estimated external loads from the dynamometer are simplified for this analysis and do not account for possible compensatory behaviour from the participant. This effect can be mitigated if the dynamometer is replaced with a lightweight force transducer.

From Figure 5.14, the results indicate that the inclusion of the grip strength model has not affected the functionality of the motion analysis pipeline for the musculoskeletal model for the muscles in the proximal sections of the upper limb. This supports the concept of employing a suite of discrete models that provide extrinsic channels of information to

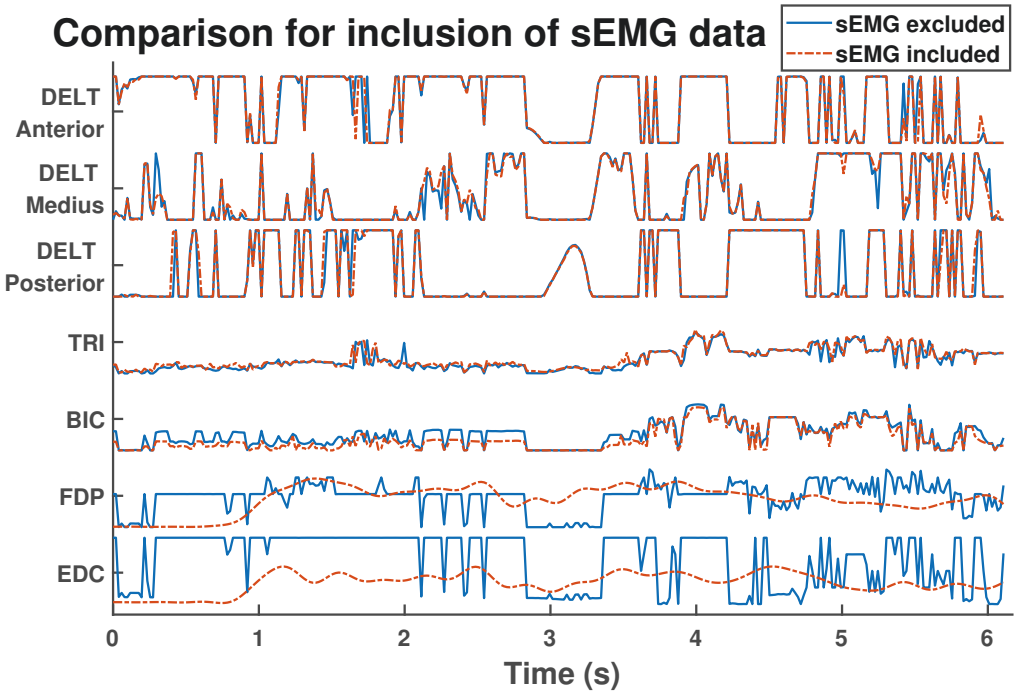

FIGURE 5.14: A comparison of the effect of the grip strength model estimate on the EMG activation of the biceps, triceps, and deltoideus muscle groups.

supplement musculoskeletal models during motion analysis. A notable phenomenon is that muscle activation from FDP and EDC may alleviate the activation necessary from biceps brachii to enable the recorded motion. This reinforces our initial observation on the bias which may be introduced into musculoskeletal models with supplementary information sources.

Overall, the grip strength model reinforces the framework presented in Chapter 3, with the feasibility of physiological models, built on task-centric and human-centric measures, to supplement musculoskeletal models during the personalisation of robotic assessment and response. Intuitively, more physiological models can be employed to complement each other and act as an adjunct to musculoskeletal models during human motion analysis and robotic assessment, alleviating bottlenecks and current limitations of using a musculoskeletal model during pHRI.

Chapter 6

Conclusion

This thesis has presented a framework towards personalised robotic assessment and assistance. The framework is informed by the various perspectives involved during Physical Human Robot Interaction. The focus is on investigations and studies to *examine* limitations in current methods, explores how to develop models to *enable* personalisation, and *enhance* the understanding of the human user during pHRI.

A real-world example application for the framework was presented, contextualised for the gamification of personalised robotic rehabilitation. As probabilistic models are the most straightforward method to capture each user's inherent variability, a popular probabilistic technique, Gaussian Process, is utilised. To investigate the limitations of GPs, explorations to alleviate computational complexity and contextualise GP outputs are conducted, with detailed formulations and experiments made available in the appendices.

An investigation into the influence of inertia during pHRI for a predictive model was presented. After this, a method to integrate interaction forces using probabilistic techniques during pHRI has been developed to provide the appropriate robot response for a coupled human-robot dyad.

A qualitative analysis of results derived from a popular musculoskeletal model was presented to investigate the motion analysis pipeline for the model. A grip strength model was then developed to complement the musculoskeletal model. The new model was experimentally validated and discussed in the context of its influence on the results derived from the musculoskeletal model.

6.1 Summary of Contributions

6.1.1 A Framework Towards Personalised Robotic Assessment and Assistance during pHRI

A framework has been developed that works towards the incorporation of the various important perspectives needed to unify personalised robot assessment and assistance during Physical Human Robot Interaction. The framework focuses on *examining* current task-centric models and how they are used for personalisation, *enabling* the integration of human-centric measures into conventional task-centric personalisation for pHRI, and *enhancing* the understanding of the human user through physiological and musculoskeletal models.

An exploration into a real-world example application for the framework was conducted. The implemented system showcases a potential pHRI application for the framework, contextualised for personalised robotic rehabilitation. of the framework was conducted, developing a derivative framework which highlights the feasibility of the framework for robotic assessment and assistance. This work was published in the IEEE International Conference on Robotics and Biomimetics (ROBIO 2018) [15].

Following on from the real-world example application, two sets of works were undertaken to study methods to overcome limitations for conventional Gaussian Process models for regression and classification. These are presented in Appendix B-C and have both been published in the IEEE International Conference on Automation Science and Engineering (CASE 2018) [14] and the IEEE International Conference on Cybernetics and Intelligent Systems and Robotics, Automation and Mechatronics (CIS-RAM 2019) [16], respectively.

6.1.2 Exploring Inertial Influence on Predictive Models

A study was performed to explore the effects of inertia when a human handles a robotic endpoint. The effects were tested based on Fitts' law where participants were instructed to move the robot endpoint to various target locations in the human upper limb planar workspace while the endpoint is compensated from gravity forces. An analysis on the fit of Fitts' law coefficients was conducted, showing the influence of inertia during coupled pHRI is statistically insignificant. This study of inertial influences has been presented and published in the International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC 2020) [17].

6.1.3 Integrating Interaction Forces during Robotic Assistance

A novel extension of a prominent HRI framework was developed, integrating interaction forces to infer user intent and generate the appropriate robot response in a pick-and-place task. The development of Physical Human Robot Interaction Primitives extend upon Interaction Primitives by embedding interaction forces in the joint distribution and inferring user intent without observations of the robot endpoint. This overcomes limitations of IPs and enables coupled human-robot dyad workspaces.

Multiple experiments were conducted to validate the efficacy of pHRIP. Results from the user-reaching, planar, and Cartesian applications demonstrate its efficacy and advantage over IPs and its ability to facilitate a human-centric approach towards personalising the robot response based on user intent. The development of pHRIP is published in Autonomous Robots (AURO) during 2022 [18].

6.1.4 Preliminary Validation of a Musculoskeletal Model

A preliminary qualitative validation was performed on a popular upper limb musculoskeletal model. An experiment, based on several exercising motions, was conducted to compare the sEMG readings against the calculated muscle activity from the human motion analysis pipeline for several muscle groups on the upper limb. Two calculation models were tested for muscle force generation in the MTU model. The results demonstrated consistency in the muscle activity trends during the observed exercises. An analysis of the results highlighted some limitations of the current pipeline indicating potential future work to improve the results obtained from the model. This preliminary validation has been accepted to be published in the International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC 2021) [19].

6.1.5 Supplementary Grip Strength Model for Musculoskeletal Models

A physiological model was developed to correlate human user grip strength and muscle activation. Participants were instructed to exert isometric grip force to a digital dynamometer that recorded their grip force at 3 distinct levels of strength. Then, were instructed to perform an exercise while exerting grip force at the same 3 levels of strength. The experiment highlighted the feasibility of the model and its fit to observed data for both grip force and sEMG readings. Additional analysis was conducted to determine the suitability of the grip

strength model to supplement the motion analysis pipeline using a musculoskeletal model. Integrating real-world sEMG readings into the analysis pipeline showed no significant effect on the calculated muscle activity, demonstrating the grip strength model's ability to supplement a portion of the calculations for musculoskeletal models.

6.2 Discussion and Limitations

The framework presents a unified approach towards robotic assessment and assistance, which is an attractive proposition for applications such as robotic rehabilitation, AAN robotics, or intuitive tele-operation. However, there are some limitations which need to be overcome before the framework can be fully realised in non-laboratory environments.

6.2.1 Reliance on Model-based Paradigm

This thesis has presented experiments, studies, and investigations that build on current conventional models for personalised pHRI. A clear limitation is the reliance on the model-based paradigm since the limitation of implementing the framework will be heavily influenced by the models chosen. Across the thesis, the level of complexity for various models have highlighted roadblocks and challenges to overcome.

For the simple models explored in this thesis, such as Fitts' law [98], the inertial influence was found to be insignificant during pHRI. The simple movements associated with the predictive model were simulated using a different interface, which has external factors influencing the participant's performance. The simplicity of the Fitts' law model means that it is more sensitive to external factors, especially if the associated task deviates from the original study. A common way to overcome this sensitivity is by increasing the number of participants to improve the confidence of the methodology as well as the observed results. In this thesis, the preliminary investigation relied on a small number of participants which may reduce the confidence on the conclusions reached. It is noted that there is evidence to suggest Fitts' law is an artefact of human motor control [101]. Thus, studies in human arm movement control [238] may be better suited for future developments for this model.

Similarly, in Chapter 4, pHRIP is an extension on a DMP-based framework. While the employment of DMP has resulted in compact representations of trajectories, the reliance on the framework, originally designed for one-shot imitation learning, has created complications when performing online implementations. As discussed in Section 4.8, the generation

of the robot response with different trajectory lengths is difficult to attain. More work needs to be done to improve its reliability to generate robust trajectories that fit the task. While there are other MPs that increase flexibility by removing the reliance on DMP, notably ProMPs [197], this flexibility creates a need to place additional constraints to fit the application. Furthermore, MPs operate under a pick-and-place paradigm where the robot response is defined based on discrete demonstrations.

In this thesis, all experiments and validation studies with a robotic arm were performed using its native controller, Rethink Robotics Intera SDK. The ability to use the gravity compensation mode allows free manipulation of the endpoint, providing human user manipulation with ease. However, the robot responses generated were limited by the reliance on Intera SDK's control system. While the native controller achieves the trajectories generated by pHRIP, the pick-and-place paradigm restricts the implementation of continuous assistance in coupled human-robot dyads. Approaching personalised robotic assessment and assistance during pHRI from the control perspective may improve the performance. Responsive control frameworks such as impedance control or hybrid control frameworks can potentially incorporate optimisation to vary the support given to the human user [91].

6.2.2 Limitations of Musculoskeletal Models

The use of a musculoskeletal model in the thesis presents some limitations on the accuracy and reliability of the results derived in the motion analysis pipeline. The primary limitation for a musculoskeletal model is the complex anatomical, physiological, and neurological systems in humans which are difficult to replicate. While the musculoskeletal model used in this thesis was shown to be adequately accurate in its joint strength [291], moment arm generation [327], and limb stiffness [328], there are still other influential factors which have not been taken into account. For example, sophisticated models require a large number of parameters to represent the intrinsic differences between the different muscle groups, joints, and bones. The musculoskeletal model used in this thesis has 50 MTUs, which consist of 5 parameters and 3 non-linear functions (for the FLV constraints) to describe each MTU.

Personalising the musculoskeletal models require knowledge on each of these parameters to best represent the human user, accounting for inherent variation in the population. Currently the process to obtain these parameters *in-vivo* is unclear or challenging. However, recent technological advances and methods may help overcome these limitations. Examples include employing ultrasound imaging to validate fibre orientation for MTUs [237]

and muscle fibre histology to clarify the composition of various muscle fibre types [236]. Due to these limitations, in this thesis, the employment of a musculoskeletal model is to perform qualitative trend analysis rather than directly using the raw results since they will not be a reflection of the real world.

In this thesis, the human motion analysis pipeline includes a static optimisation process. It assumes that the forces and torques generated by the MTUs are in equilibrium given the kinematics and dynamics of the model as calculated by Inverse Kinematics and Inverse Dynamics. This assumption of equilibrium ignores the non-linear dynamics imposed by changing muscle fibre lengths when the arm is not stationary [329]. While this effect is mitigated in the methodology for Section 5.3.2, where participants exerted isometric grip force, the non-linear dynamics still exist and influence the accuracy of the results. For non-isometric movements analysed in this thesis, the Millard equilibrium MTU model is used; however, there is another variant of the MTU model, associated with acceleration dynamics [293], which may have mitigated this assumption during motion analysis. Additionally, the validation of the results obtained from the musculoskeletal models are achieved by comparing against real-world sEMG readings. Despite best practices with sEMG data collection, there is variability during the procedure. These introduce bias into the readings such as non-linear noise during contraction, with a lower signal to noise ratio during low contraction intensity [330].

As previously discussed in Section 1.3, the scope of this thesis excludes an investigation into the active field of research focusing on different objective functions during muscle activity optimisation. The limitations on using the minimum sum of squared activation objective function is evident in the "spiking" muscle activation since each frame of the motion is optimised in isolation while under equilibrium. This calculation ignores the dynamic forces induced during motions. However, this method continues to be the most commonly used objective function during analysis. One straightforward constraint to incorporate would be to limit the change in muscle activation from one time frame to the next since muscle groups are heterogeneous with individual motor action potential for each muscle fibre. An alternative method to validate the muscle activation may be the utilisation of indirect calorimetry [331] to investigate inconsistencies between the equivalent energy expenditure and the measured real-world energy consumption.

6.3 Future Research

6.3.1 A Validation Task for the Multi-Model Implementation

The framework presented in this thesis has explored various perspectives for robotic assessment and assistance during pHRI, with experiments and studies conducted in isolation. However, the complexities involved when integrating the multiple models and various perspectives needs to be considered in future work. This is particularly pertinent for computational complexity and the communication framework between the various models during pHRI.

For near real-time implementations, which include human motion analysis using musculoskeletal models, the computational complexity of the analysis pipeline needs to be examined to enable near real-time calculations. Alternatively, the application should be reconfigured such that the slower computation time does not adversely affect the performance. The reduction of computational time can be seen in [324] by simplifying optimisation constraints for estimating endpoint strength during bimanual operations, while [88] built an offline GP model that correlates the endpoint impedance to the joint kinematics.

With such a variety of models, each with its individual motivations, the integration of multiple models can be challenging. As a result, a communication framework outlining how collected data would interact and communicate with each internal model will facilitate the validation task based on a multi-model implementation. This is particularly pertinent when integrating physiological results and human-centric measures into robotic frameworks. Examples of this pipeline include DMPSynergies [180], which leverage muscle synergies to build the joint probabilistic distribution for the DMP parameters, and the use of reinforcement learning to learn the appropriate assistance based on sEMG readings [86].

As pointed out in Section 6.2.1, the lack of low-level robotic control limits the ability for shared control towards personalised robotic assistance. Thus, the multi-model implementation will also validate the integration of the musculoskeletal model, including the supplemental grip strength model, during robotic arbitration for pHRI. A low-level control framework can enable adaptive robotic responses, which would improve the quality of the pHRI and enable robot control. In turn, this control can more effectively utilise the computed results, such as for strength estimation [22] and in AAN applications [285].

6.3.2 Trials to Validate Personalisation

The motivation for the framework presented in this thesis relies on probabilistic models, under the foundation of Bayesian statistics, to facilitate personalisation. The employment of probabilistic models in Chapter 4 for pHRIP is appropriate since they provide a flexible method to directly couple observed interaction forces with the robotic trajectory. The implicit personalisation of the pHRIP parameter sets for each user is promising for personalised robotic assessment and assistance during pHRI. However, explicit personalisation of probabilistic distributions should be explored in future work. This explicit personalisation can be applied to the physiological models supplementing musculoskeletal models during human motion analysis.

Using the grip strength model presented in Chapter 5 as an example, a parametric prior distribution can be built using demographic data, which is then personalised by obtaining individual data during a "calibration" process to update the posterior predictive distribution. The main challenge to overcome for this approach is the right set of parameters to best represent the model, and the number of participants to: (1) build the demographic prior distribution, and (2) validate the personalisation process.

One factor to consider when using this approach is the assumption of IID observations when building and updating the parametric distributions. While there might be a strong correlation that can be observed, the performance of the model can falter if it does not inherently capture the ultimate causal relationship. Thus, the choice of model complexity and its causal relationship with the observations will influence the predictive power and reliability of the posterior distribution.

6.3.3 Sensitivity Analysis into Model Complexity

This thesis has presented multiple perspectives for personalisation during pHRI for robotic assessment and assistance. A common challenge borne from the reliance on model-based paradigms (Section 6.2.1) is the contextual functionality of the model. While simple human-centric models, such as Fitts' law, do not yield enough insight into human motor learning or locomotion, the other spectrum of representative models, such as sophisticated musculoskeletal models, present challenges in obtaining the correct parameter set and validating their accuracy for real-world applications. As a result, future work in this topic will require a systematic method to investigate and perform a sensitivity analysis on the balance between model complexity and its functional outcomes for pHRI.

For complex musculoskeletal models, simplifications are usually performed by either discretising the continuous workspace [22, 88] or ignoring a portion of the model [242, 262]. However, the difficulty remains in finding the correct set of intrinsic parameters to ensure the accuracy of the results. Thus, an examination of models with varying complexities is needed to find the balance between (1) representative modelling of human anatomy and physiology, and (2) functional outcomes for multi-model implementations for personalised robotic assessment and assistance during pHRI.

A similar investigation into model complexity is also needed from the perspective of probabilistic techniques. While probabilistic distributions offer a variety of options for personalisation, the difficulty in contextualising the posterior predictive distribution indicates that a hybrid approach, like that of pHRIP, is appropriate. The sensitivity analysis would be on the underlying framework to simplify robotic trajectories/control and provide robust robotic assistance during pHRI. Furthermore, there is a need to examine how to best utilise the data collected during interactions. This includes a look at assessing the user from multiple models with the same data, such as using interaction forces and limb kinematics to assess endpoint impedance, grip strength, and user intent.

Future work exploring this balance will be difficult since the level of flexibility for personalisation required will be dependent on the application. Giving too little flexibility during personalisation will lead to less ability to capture unique variations for the human user, while too much flexibility will require some form of contextualisation for the captured distribution. One other aspect to consider for personalised robotic assessment and assistance during pHRI is the user experience and their opinions during interactions.

Appendix A

Musculoskeletal Model Parameters and Modifications

The original upper limb musculoskeletal model for the studies and experiments presented in Chapter 5 consists of 15 DoF with 50 MTU models [291], with the model representing the upper limb including the extremities.

A.1 Modifications for Preliminary Validation

For the preliminary validation, a reduction in complexity of the model is performed by removing or ignoring MTUs and joints which controlled and represented the extremities. MTUs which did not produce torque around the shoulder and elbow joints were ignored in the calculations. Further reductions of complexities were conducted and are detailed in Section 5.1.2. These simplifications were performed to reduce computation time for the static optimisation considering the number of frames to be calculated during the validation study.

TABLE A.1: The list of generalised coordinates in the upper limb musculoskeletal model used and ignored for the preliminary validation (Section 5.1).

Coordinate Name	**Utilised in Model?**
elv_angle	Yes
shoulder_elv	Yes
shoulder_rot	Yes
elbow_flexion	Yes
pro_sup	Yes
deviation	Yes
flexion	Yes
cmc_flexion	No
cmc_abduction	No
mp_flexion	No
ip_flexion	No
2mcp_flexion	No
2mcp_abduction	No
2pm_flexion	No
2md_flexion	No

TABLE A.2: The list of muscle tendon units in the upper limb musculoskeletal model ignored for the preliminary validation (Section 5.1).

MTU Acronym	**Name**
PQ	Pronator Quadratus
FDSM	Flexor Digitorum Superficialis (Median)
FDSL	Flexor Digitorum Superficialis (Left)
FDPL	Flexor Digitorum Profundus (Left)
FDPM	Flexor Digitorum Profundus (Median)
FDPR	Flexor Digitorum Profundus (Right)
FDPI	Flexor Digitorum Profundus (Interphalangeal)
EIP	Extensor Indicis Proprius
EPL	Extensor Pollicis Longus
EPB	Extensor Pollicis Brevis
FPL	Flexor Pollicis Longus
APL	Abductor Pollicis Longus

TABLE A.3: The list of muscle tendon units in the upper limb musculoskeletal model utilised for the preliminary validation (Section 5.1).

MTU Acronym	**Name**
DELT1	Anterior Deltoid
DELT2	Acromial Deltoid
DELT3	Posterior Deltoid
SUPSP	Supraspinatus
INFSP	Infraspinatus
SUBSC	Subscapularis
TMIN	Teres Minor
TMAJ	Teres Major
PECM1	Pectoralis Major (Clavicle)
PECM2	Pectoralis Major (Sternum)
PECM3	Pectoralis Major (Ribs)
LAT1	Latissimus Dorsi Tvert
LAT2	Latissimus Dorsi Lvert
LAT3	Latissimus Dorsi Iliac
CORB	Coracobrachialis
TRIlong	Triceps Brachii Long
TRIlat	Triceps Brachii Lateral
TRImed	Triceps Brachii Medial
ANC	Anconeus
BIClong	Biceps Brachii Long
BICshort	Biceps Brachii Short
BRA	Brachialis
BRD	Brachioradialis
ECRL	Extensor Carpi Radialis Longus
ECRB	Extensor Carpi Radialis Brevis
ECU	Extensor Carpi Ulnaris
FCR	Flexor Carpi Radialis
FCU	Flexor Carpi Ulnaris
PL	Palmaris Longus
PT	Pronator Teres
FDSL	Flexor Digitorum Superficialis (Left)
FDSR	Flexor Digitorum Superficialis (Right)
EDCL	Extensor Digitorum Communis (Left)
EDCR	Extensor Digitorum Communis (Right)
EDCM	Extensor Digitorum Communis (Medial)
EDCI	Extensor Digitorum Communis (Interphalangeal)
EDM	Extensor Digiti Minimi

A.2 Modifications for Grip Strength Model

For the comparisons using the grip strength model, the complexity of the model is reduced by ignoring the MTUs and joints which controlled and represent the extremities. Torques produced in the wrist joint are taken into account during calculations and further reductions in the computations during static optimisation are detailed in Section 5.3.6. In particular, compared to the modifications performed for the preliminary validation, the MTUs that contribute to grip force generation are included.

TABLE A.4: The list of generalised coordinates in the upper limb musculoskeletal model used and ignored for the grip strength model validation (Section 5.2).

Coordinate Name	Used in Model?
elv_angle	Yes
shoulder_elv	Yes
shoulder_rot	Yes
elbow_flexion	Yes
pro_sup	No
deviation	No
flexion	No
cmc_flexion	No
cmc_abduction	No
mp_flexion	No
ip_flexion	No
2mcp_flexion	No
2mcp_abduction	No
2pm_flexion	No
2md_flexion	No

TABLE A.5: The list of muscle tendon units in the upper limb musculoskeletal model ignored for the grip strength model validation (Section 5.2).

MTU Acronym	Name
PL	Palmaris Longus
PT	Pronator Teres
PQ	Pronator Quadratus
EPL	Extensor Pollicis Longus
EPB	Extensor Pollicis Brevis
FPL	Flexor Pollicis Longus
APL	Abductor Pollicis Longus

TABLE A.6: The list of muscle tendon units in the upper limb musculoskeletal model utilised for the grip strength model validation (Section 5.2).

MTU Acronym	**Name**
DELT1	Anterior Deltoid
DELT2	Acromial Deltoid
DELT3	Posterior Deltoid
SUPSP	Supraspinatus
INFSP	Infraspinatus
SUBSC	Subscapularis
TMIN	Teres Minor
TMAJ	Teres Major
PECM1	Pectoralis Major (Clavicle)
PECM2	Pectoralis Major (Sternum)
PECM3	Pectoralis Major (Ribs)
LAT1	Latissimus Dorsi Tvert
LAT2	Latissimus Dorsi Lvert
LAT3	Latissimus Dorsi Iliac
CORB	Coracobrachialis
TRIlong	Triceps Brachii Long
TRIlat	Triceps Brachii Lateral
TRImed	Triceps Brachii Medial
ANC	Anconeus
BIClong	Biceps Brachii Long
BICshort	Biceps Brachii Short
BRA	Brachialis
BRD	Brachioradialis
ECRL	Extensor Carpi Radialis Longus
ECRB	Extensor Carpi Radialis Brevis
ECU	Extensor Carpi Ulnaris
FCR	Flexor Carpi Radialis
FCU	Flexor Carpi Ulnaris
FDSL	Flexor Digitorum Superficialis (Left)
FDSM	Flexor Digitorum Superficialis (Medial)
FDSR	Flexor Digitorum Superficialis (Right)
FDPL	Flexor Digitorum Profundus (Left)
FDPM	Flexor Digitorum Profundus (Median)
FDPR	Flexor Digitorum Profundus (Right)
FDPI	Flexor Digitorum Profundus (Interphalangeal)
EIP	Extensor Indicis Proprius
EDCL	Extensor Digitorum Communis (Left)
EDCR	Extensor Digitorum Communis (Right)
EDCM	Extensor Digitorum Communis (Medial)
EDCI	Extensor Digitorum Communis (Interphalangeal)
EDM	Extensor Digiti Minimi

Appendix B

Bayesian Committee Machine (BCM)

Addressing the roadblocks for implementing GP will improve its ability to provide concise and contextualised results to personalise the assessment of the end user. One such problem for GPs is the computational cost of inference, $\mathcal{O}(N^3)$ for N training samples, which scales poorly as the number of training samples increase.

The Bayesian Committee Machine (BCM) [278] is an ensemble extension for Gaussian Process to alleviate computational bottlenecks when performing inference. For N training samples, the computational complexity for GP is $\mathcal{O}(N^3)$, which scales poorly as the number of training samples increase.

BCM tackles this issue by breaking down the training data set into M equally sized data sets, $D^{1,..,M}$, building models for each D^i, where the number of data, $n \ll N$, and then combining their estimates. Note that a BCM is built for each dimension of $\boldsymbol{y}$. Given

$$P(y^*|D^{i-1}, D^i) \propto P(y^*)P(D^{i-1}|y^*)P(D^i|D^{i-1}, y^*), \tag{B.1}$$

and the approximation,

$$P(D^i|D^{i-1}, y^*) \approx P(D^i, y^*), \tag{B.2}$$

Bayes' rule then yields,

$$P(y^*|D^{i-1}, D^i) = Q \times \frac{P(y^*|D^{i-1})P(y^*|D^i)}{P(y^*)}, \tag{B.3}$$

where Q is a constant. The resultant predictive distribution is then approximated by:

$$\hat{P}(y^*|D^{1,\ldots,M}) = Q \times \frac{\prod_{i=1}^{M} P(y^*|D^i)}{P(y^*)^{M-1}}. \tag{B.4}$$

From M GP inferences, $\mathcal{N}(\mu_i, \sigma_i^2)$ at x^*, a prediction for the overall $\hat{y}^* = \mathcal{N}(\hat{\mu}^*, \hat{\sigma}^2)$ can be obtained:

$$\hat{\mu}^* = C^{-1} \times \sum_{i=1}^{M} \sigma_i^{-2} \mu_i, \tag{B.5}$$

where $C = \hat{\sigma}^{-2} = -(M-1)K_{**}^{-1} + \sum_{i=1}^{M} \sigma_i^{-2}$.

B.1 Experiment

Experiments were conducted to evaluate the performance of Gaussian Process-Bayesian Committee Machine (GP-BCM) through two variants, offline-trained and online learning, and are compared against standalone GP models. For both variants, the data subset size is explored to evaluate its effects on both computation time and estimation error. This experiment was contextualised for inferring the node locations of a Deformable Linear Object (DLO) during occlusions.

The complexity of describing a DLO in Cartesian space makes it prohibitive for conventional GP modelling. Rather than mapping the incoming points in a latent space [332] to reduce the dimensionality, the geometric properties of DLOs are utilised instead.

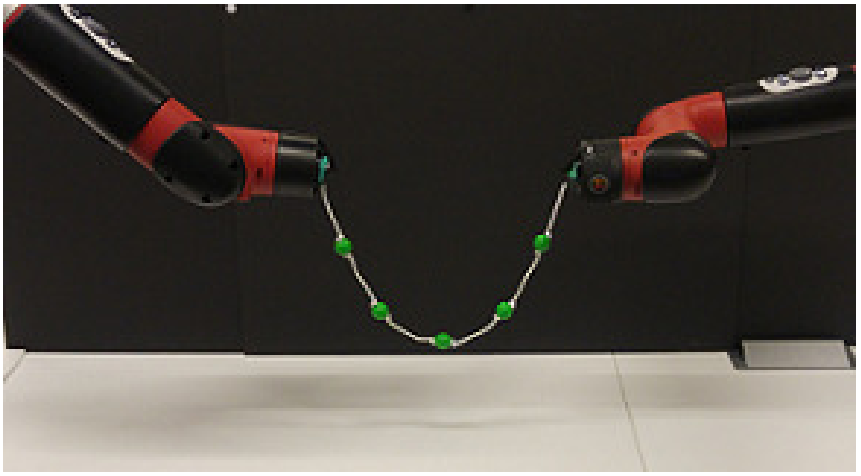

FIGURE B.1: A Deformable Linear Object (DLO) handled by two robotic manipulators.

Here the DLO is assumed to be hanging freely and subject to no other external forces besides gravity. Additionally, the DLO is taken as a series of nodes equally spaced along its length. As the DLO is inherently planar in this scenario, Figure B.2 shows how the structure can be defined as a series of angular offsets, α, between adjacent nodes.

Since the input for the GP models are based on angular representation, the possibility of multiple training targets around a specific input is high. A major effect is the under-fitting of inferences resulting in both erroneous and irrelevant estimates. Furthermore, training times would be expected to increase due to the complexity of the hyper-parameter optimisation, undermining the benefits of lower computation complexity. To address this, training data was binned to yield a target mean and uncertainty for each discretised query. For all models, hyper-parameter optimisation was performed using Differential Evolution (DE) [143].

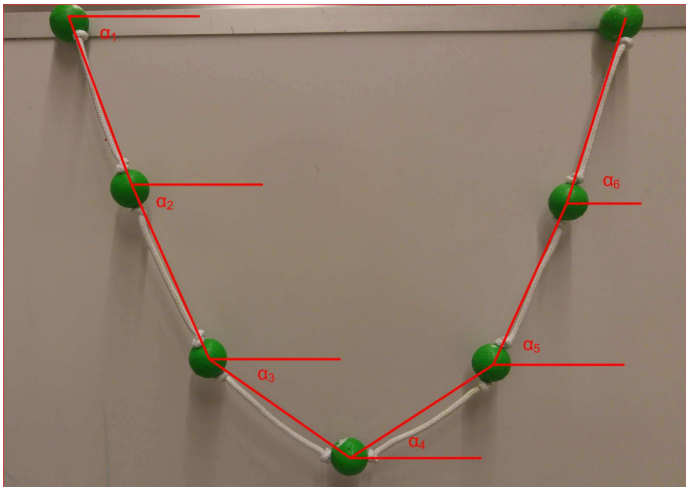

FIGURE B.2: Representation of the angular offsets between adjacent nodes along the DLO.

The inference objective is defined as $\boldsymbol{\alpha} = \{\alpha_1, \alpha_2, ..., \alpha_{n-1}\}^T$ for n physical nodes. Two GP-BCMs are built, accumulating data over time to model α as a function of its neighbours, incorporating the ascending and descending uncertainty from both GP-BCMs; i.e. $\alpha_i = g(\alpha_{i-1}, \alpha_{i+1})$.

For clarity, here the index, i is dropped. From either of the BCM's estimates (Eqn. B.5) of the first α in a sequence of several consecutive missing α values, the mean μ_f and its associated standard deviation, σ_f is obtained. This procedure is repeated on the other BCM to obtain μ_g and σ_g. An estimate of the next α's mean is simply $\hat{\mu_f} = f(\mu_f)$. Then,

from

$$(\mu_{lb}, \sigma_{lb}^2) = f(\mu_f - \sigma_f), \tag{B.6}$$

$$(\mu_{ub}, \sigma_{ub}^2) = f(\mu_f + \sigma_f), \tag{B.7}$$

we approximate a bound for its ascending and descending uncertainties, σ_{asc} and σ_{desc} as,

$$\sigma_{asc} = (\mu_{ub} + \sigma_{ub}) - (\mu_{lb} - \sigma_{lb}), \tag{B.8}$$

$$\sigma_{desc} = (\mu_{lb} + \sigma_{lb}) - (\mu_{ub} - \sigma_{ub}), \tag{B.9}$$

$$\hat{\sigma}_f \approx \frac{1}{2} \times \max(|\sigma_{asc}|, |\sigma_{desc}|). \tag{B.10}$$

$\hat{\mu}_f$ and $\hat{\sigma}_f$ are then propagated in the same manner for the remaining α in the missing sequence. Although faster than the exact analytical solution [133], this approximation assumes that $\hat{\sigma}_{f/g}$ remains reasonably constrained and that the latent BCM functions are not saddle points at $\hat{\mu}_{f/g}$. As the two GP-BCMs infer inwards towards the center of the hanging DLO from opposite ends, propagated uncertainties are softened through weights based on each GP-BCM's variance. Low variance estimates are prioritised and favoured, providing improved confidence in the weighted estimates. These estimates are not included in the BCM's GP model generation, satisfying the assumption that variances are reasonably constrained.

From this, the estimated angle, $\hat{\alpha}$, can be obtained from predictions, $(\hat{\mu}_f, \hat{\sigma}_f^2)$, $(\hat{\mu}_g, \hat{\sigma}_g^2)$ and weights, $w_f = \hat{\sigma}_f^{-2}$, $w_g = \hat{\sigma}_g^{-2}$ as,

$$\hat{\alpha} = \mathcal{N}\left(\frac{w_f \hat{\mu}_f + w_g \hat{\mu}_g}{w_f + w_g}, \frac{1}{w_f + w_g}\right). \tag{B.11}$$

B.1.1 Setup

One end of the DLO was manipulated, while the other end was affixed to a robot endpoint. A representative image of this scenario is shown in Figure B.1. When collecting data for experiments, all nodes are visible to serve as an occlusion-free ground truth for error calculations. Occlusions are simulated by intentional omission of α measurements along a middle portion of the DLO, varying from no occlusions to 3 node occlusions. Such occlusions account for 25% of all data collected and simulate the natural occurrence of occlusions as illustrated by the sequence in Figure B.3. The number of training samples for each GP was varied to evaluate the estimate accuracy.

The experiments were conducted based on 4 collected data sets. The data sets vary in length from 401 to 743 unique frames, each with 5 α measurements. The size of the data subset was explored ranging from 200 to 600 samples per GP model. Experiments were conducted on a machine with an Intel Core i7-7700@3.60GHz×8 CPU, 16GB RAM, and a Nvidia GeForce GTX 970 GPU.

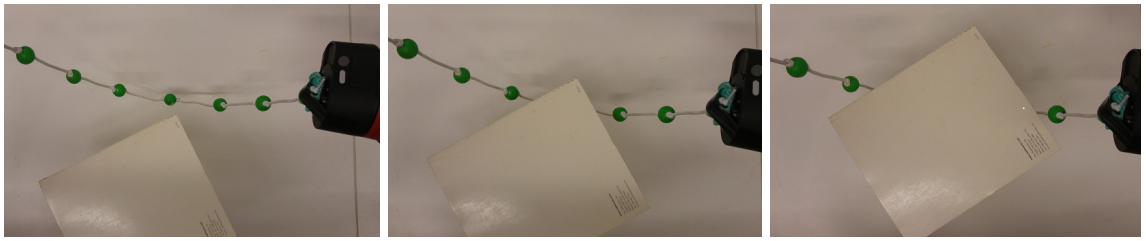

(A) (B) (C)

FIGURE B.3: The omission of α measurements simulate this sequence of occurrence where the data set has: (A) no occlusions, (B) one node occlusion, and (C) three node occlusions.

B.1.2 Offline-trained GP-BCM

The offline-trained GP-BCM approach yields a BCM with M GPs of n training samples:

$$\alpha_{i,m} = g_m(\alpha_{i-1}, \alpha_{i+1}), \tag{B.12}$$

$$\alpha_{i,j} = g_j(\alpha_{i-1}, \alpha_{i+1}), \tag{B.13}$$

where $m, j < M$. The multiple underlying GPs in the BCM can mitigate the effects of noisy training data due to the weight-based estimations. However, we expect the noise filtering effect to eliminate any features present in the test data. The expected total training cost of this variant of GP-BCM reduces to $M \times n^3 < N^3$ where $n \ll N$.

B.1.3 Online GP-BCM

The online GP-BCM starts with no data samples and accumulates each subset over time. Once the accumulated data samples exceeds a maximum subset size, a GP model is trained utilising the latest set of data. After more than one model has been trained, the framework will commence estimation of occluded α as they occur. Visible data is then used to build the subsequent GP models within the BCM, providing continuous estimation and training.

Noisy input data can be expected to be filtered out as more underlying GP models are built, lowering the sensitivity of the estimates to noise. Similar to the offline GP-BCM

variant, as more underlying models are built over time, local features will be dissipated unless explicitly recovered.

B.2 Results

The efficacy of BCM is assessed based on the predictive accuracy of the DLOs, using the RMSE between the inferred nodes and the ground truth, and computation complexity, using training and running time. Further analysis looks at the training and runtime efficiency of the BCM using the s/GP measure. As the data is binned prior to GP model creation, the training time is capped to the number of bins. This processing step is required as demonstrated by Table B.1 where the number of samples is non-linearly proportional to the training time required for a conventional GP using unbinned data.

TABLE B.1: Total time (in seconds) for conventional offline and unbinned GP.

Dataset	1	2	3	4
No. samples	2005	2685	3715	3655
Training Time	300.49	530.52	1298.8	1223.1
Run Time	0.0394	0.0423	0.0474	0.0504
Efficiency (s/sample)	0.1499	0.1976	0.3496	0.3347

The comparison for total time taken (Training Time + Run Time) is evident when compared against the results presented in Table B.2 where the longest time taken is $36.07s$. Furthermore, we observe that the two GP-BCM variants have similar efficiency. We do note that online GP-BCM has a marginally lower efficiency when compared to offline GP-BCM. However, the training time required for online GP-BCM is significantly lower. The online GP-BCM variant also allows continuous inference of α while training new models, providing scalability for applications.

TABLE B.2: Total time taken, in seconds (s), for the offline and online variants of GP-BCM to perform inference in data set 1 (2005 samples).

GP-BCM Variant	No. samples per GP	No. GPs	Training Time	Run Time	Efficiency (s/GP)
Offline	200	10	35.9875	0.0828	3.60703
	400	5	17.9483	0.0569	3.60104
	600	3	10.9066	0.0577	3.65477
Online	200	4	3.7616	11.0546	3.70405
	400	2	3.7548	3.7545	3.75465
	600	1	3.7451	0.0591	3.8042

TABLE B.3: Root Mean Square Error from experiment results conducted with GP-BCM variants and offline-trained GP, and Sum of Squared Errors of each variant for all datasets.

RMSE	Online GP-BCM					Offline GP-BCM					GP
	No. Samples					No. Samples					
(rads)	200	300	400	500	600	200	300	400	500	600	
Dataset 1 (401 sample frames)											
α_2	0.19	0.15	0.17	0.19	0.17	0.66	0.49	0.35	0.92	0.19	0.08
α_3	0.32	0.30	0.32	0.33	0.32	0.75	1.21	0.67	1.39	0.89	0.19
α_4	0.32	0.25	0.26	0.56	0.31	0.60	1.20	0.73	0.31	1.05	0.22
α_5	0.11	0.13	0.11	0.117	0.10	0.57	0.41	0.08	0.09	0.55	0.11
SSE	0.25	**0.19**	0.21	0.47	0.24	1.68	3.31	**1.11**	2.88	2.22	0.10
Dataset 2 (537 sample frames)											
α_2	0.01	0.02	0.01	0.01	0.01	2.74	2.98	3.55	1.68	0.87	0.02
α_3	0.06	0.09	0.05	0.04	0.04	1.92	2.73	2.89	1.50	0.70	0.05
α_4	0.30	0.36	0.03	0.04	0.02	0.38	0.43	0.38	0.33	0.53	0.06
α_5	3.01	0.99	0.24	0.26	0.17	0.21	0.15	0.13	0.15	0.25	0.03
SSE	9.16	1.13	0.06	0.07	**0.03**	11.4	16.5	21.1	5.20	**1.59**	0.01
Dataset 3 (743 sample frames)											
α_2	0.09	0.05	0.05	0.07	0.11	1.61	0.45	0.10	0.17	0.61	0.05
α_3	0.26	0.25	0.28	0.25	0.23	1.97	2.87	0.52	0.26	0.77	0.15
α_4	4.66	0.66	0.47	0.95	0.45	2.00	1.70	0.79	1.19	1.85	0.19
α_5	0.24	0.22	0.16	0.05	0.38	0.65	0.77	0.98	1.03	1.35	0.08
SSE	21.8	0.55	**0.32**	0.98	0.41	10.89	11.92	**1.87**	2.56	6.22	0.07
Dataset 4 (731 sample frames)											
α_2	0.20	0.20	0.20	0.18	0.20	0.49	0.65	0.21	0.22	0.27	0.15
α_3	0.18	0.28	0.20	0.18	0.22	0.57	0.58	0.37	0.29	0.70	0.14
α_4	0.65	1.28	0.45	0.18	0.15	0.39	0.95	0.35	0.53	1.04	0.17
α_5	0.18	0.27	0.18	0.19	0.19	0.56	0.63	0.41	0.46	0.60	0.15
SSE	0.52	1.83	0.31	**0.13**	0.15	1.04	2.06	**0.47**	0.62	1.99	0.09

Given that there are 3 occluded nodes in the data set (requiring 4 α inferences), an overall metric is required to evaluate the different variants within GP-BCM. The Sum of Squared Errors (SSE) is chosen to evaluate each variant's sensitivity to outliers due to the quadratic nature of the error metric, similar to RMSE. SSE's sensitivity to outlier data forms a biased metric against outliers, making it optimal for the application since large inference errors can adversely affect subsequent analysis when interacting with a DLO.

From the results shown in Table B.3, we can see that the online GP-BCM variant is comparable in performance to conventional GP. The sample sizes which produced the best result for online and offline GP-BCM were bolded, identifying a key characteristic of GP-BCM. We can see that the GP sample size affects the ability for the underlying GPs to provide accurate estimates as small sample sizes can potentially have a smaller

range of training inputs, causing the model to be overconfident on incorrect estimates. A representative example of this phenomenon can be seen in data set 2 for offline GP-BCM where α_2 and α_3 were consistently inaccurate when compared against α_4 and α_5.

The overall results of the framework were worse off in data sets 3 and 4 due to the feature-rich nature of the data set itself. A noticeable dissipation of features within the data set can be seen when the number of samples per sub-model decreases. However, the preservation of local features can be seen when the number of models reaches a moderate value. In contrast to the feature-rich data sets 3 & 4, the GP-BCM performed well when compared to conventional GP with the feature-sparse data sets 1 & 2 as seen in Figure B.4.

Further investigation into highly erroneous outputs observed for the offline variants indicate that the underlying models for those data sets had the least variance due to insufficient training data. This skewed the resultant output since that overconfident GP ensemble member, with inaccurate regression inferences, was strongly weighted. On the other hand, the ability for GP-BCM to capture and integrate local features and trends with prior online data provides a distinct advantage over standalone GPs. Furthermore, the uncertainties attached to the resultant output distribution provides data-rich information for subsequent frameworks without additional overheads.

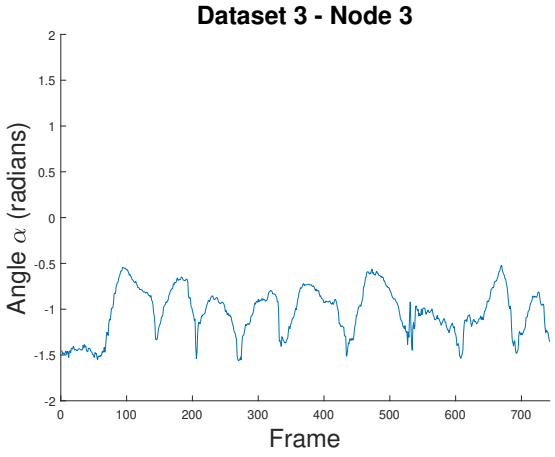

(A)

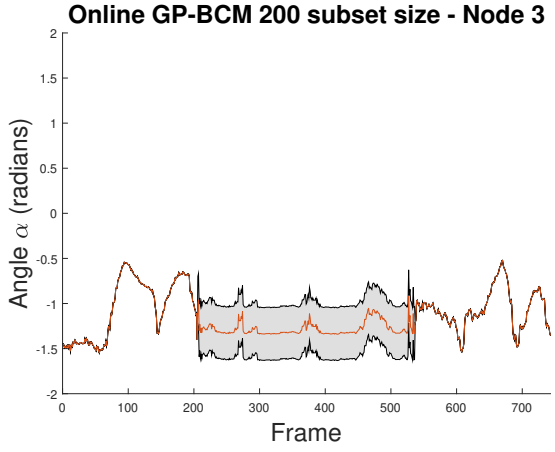

(B)

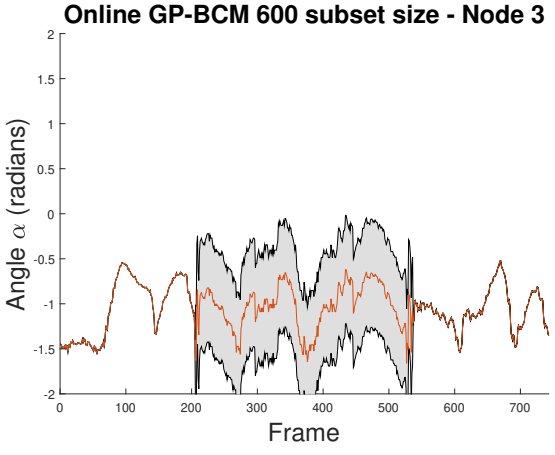

(C)

FIGURE B.4: The results of α_3 node from: (A) raw dataset; (B) the 200 sample size online GP-BCM; (C) the 600 sample size online GP-BCM. In (A), local features are filtered through the 4 (at the time of inference) underlying GP models built. The 95% confidence interval is indicated by the shaded region. In contrast, (C) highlights the preservation of local features leading to better results for α_3 node for the data set.

Appendix C

Longest Warping Subsequence (LWSS)

The Longest Warping Subsequence (LWSS) score was developed to work in tandem with the RMSE to distinguish between outliers and systematic changes. Given two multi-dimensional data sets, conventional comparison methods compute a single scalar metric of similarity between the two data sets. Common methods include the paired t-test, RMSE, and the Auto-regressive Moving Average. For a path following task, commonly seen in rehabilitation, these methods do not indicate *when* and *where* the divergence occurs, making it difficult to extract relevant information.

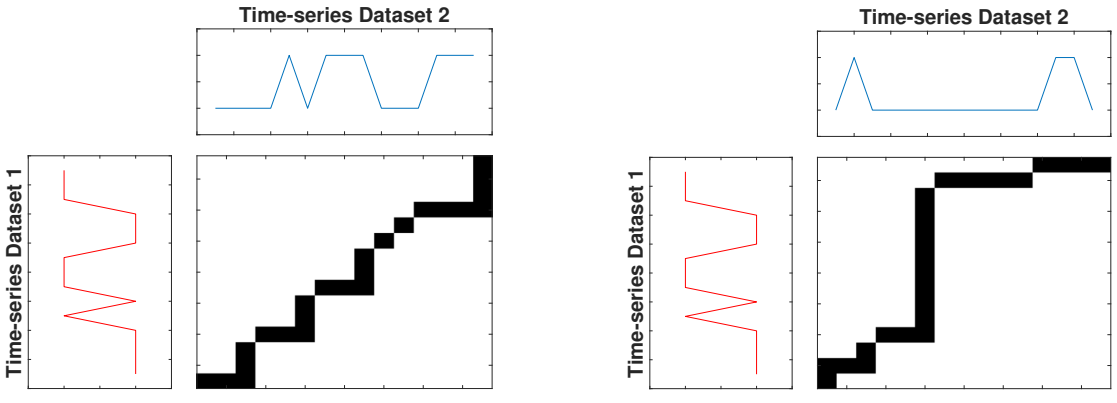

(A) (B)

FIGURE C.1: A comparison of the warping path between two time-series data sets: (A) two near-identical data sets produce a near-diagonal warping path representing a 1-to-1 match between the data samples; and (B) two distinctly different data sets will cause long sequences in the warping path indicating its inability to find data sample matches.

Dynamic Time Warping [199] provides a distance metric and a warping path, describing how data samples are matched, using one data set as a reference. Thus, change detection is possible using the longest sequence in the warping matrix. A long sequence, as shown in Figure C.1(B), indicates a region of change.

The LWSS is the ratio of the longest sequence normalised to the data set sizes:

$$LWSS(w,z) = \frac{n \times L_{max}(w,z)}{m \times \operatorname{argmax} size(w,z)}, \tag{C.1}$$

where w is a $(n \times D)$ data set, z is a $(m \times D)$ data set, and L_{max} is the longest sequence within the warping matrix.

In scenarios where two data sets are discordant, ascertaining the magnitude of the difference is challenging with the warping path. Thus, RMSE can be used as a complementary measure of the difference between the two data sets. When a changed path is detected, a buffer, W is initialised to accumulate subsequent test paths. The buffer looks for changes in a window of test paths, using the RMSE and the LWSS ratio as follows:

$$\boldsymbol{g}(W,c) = \sum_{a=1}^{A} \Big\{ RMSE(w_a, z_c) \times LWSS(w_a, z_c) \Big\}, \tag{C.2}$$

where A is the number of accumulated test paths and z is the model of the trained path for class c.

When the joint RMSE-LWSS score, $\boldsymbol{g}(W,c)$, is consistently less than a specified threshold, the presented input is treated as an outlier. Conversely, when the RMSE-LWSS score is over the previous threshold, the presented input is assumed to be novel, providing an opportunity to create a separate model or ignoring the specific input.

C.1 Experiment

An experiment was conducted to investigate the performance of the Longest Warping Subsequence (LWSS) score and to compare against GP classification. Twenty-five *healthy* participants were asked to interact with a 7 Degree-of-Freedom (DoF) robotic manipulator (HAHN Rethink Robotics, Rheinböllen, Germany) with a bespoke handle affixed to the robot endpoint. The robot arm was programmed to enter a XY-plane constrained zero-g mode native to the Rethink Robotics' Intera SDK. Participants were asked to follow the path displayed on a screen. The relative position of the robot endpoint is presented in the

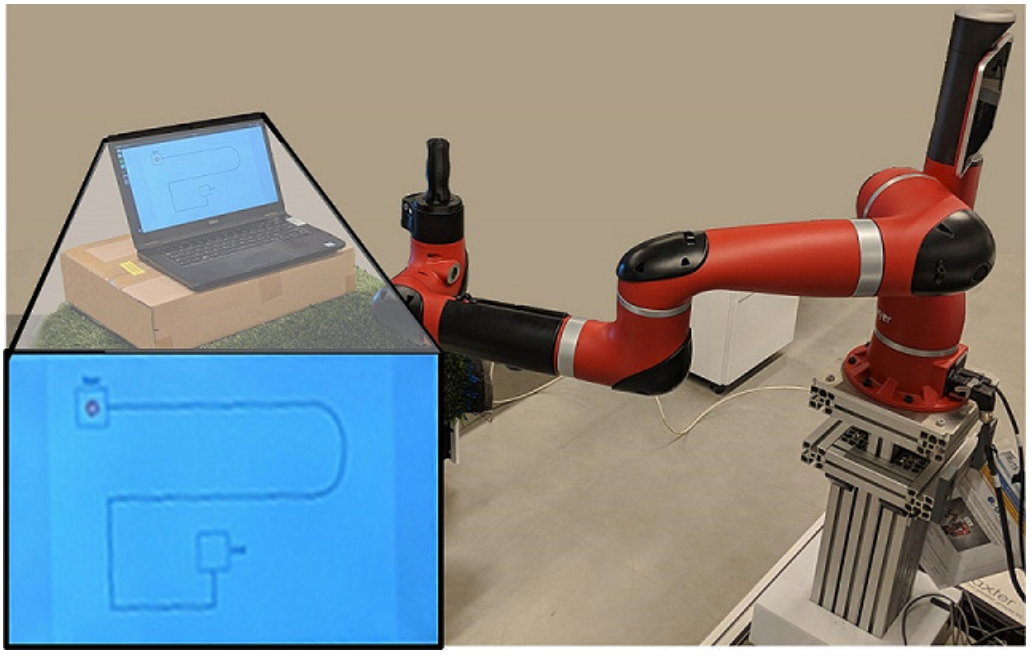

FIGURE C.2: The setup for the Sawyer manipulator, constrained to the X-Y plane with gravity compensation, and the target path shown on the screen.

screen as seen in Figure. C.2. Notably, in this experiment, only the task-centric endpoint position data from the robotic arm was recorded.

Participants were asked to perform the path following task on a total of 6 unique paths. These included 4 original paths for training a Gaussian Process model, and 2 different paths which represent outliers and systematic changes (This compensatory behaviour results in paths that deviates from the original tracks, and completely different track profile indicating dysfunction). In total, 100 trajectories were collected, of which 60 were used for training the GP model, 20 were used to validate the generated model, and 20 were used for testing (all trajectories with the 2 different paths are in this data subset).

Since participants completed the task with different velocity profiles, the data was temporally aligned and normalised to facilitate the generation of the time-series GP model. The data is embedded in a multivariate environment. Thus, a multi-dimensional DTW algorithm [286] was used to warp the data sets, with the data sets z-normalised to follow the algorithm's recommendation.

A covariance matrix of the DTW distances between the training data sets was created and averaged to determine the reference data set for each path during DTW. The warped paths are then temporally standardised and used to train a Gaussian Process model to generalise the optimal path and classify future input paths.

C.2 Results

The paths performed by the participants are indicated as Paths 1-4. The 15-participant training data set produced a GP model for each of the original 4 paths (with one seen in Figure C.3). The 5-participant holdout set used the original paths, while the 5-participant test set was obtained from 2 original paths, 1 modified path, and 1 novel path. All data sets were temporally normalised to 200 samples through linear interpolation. Each of the 4 GP models in the training data set was then trained on a GP Laplace Approximation classifier in a one-vs.-rest fashion.

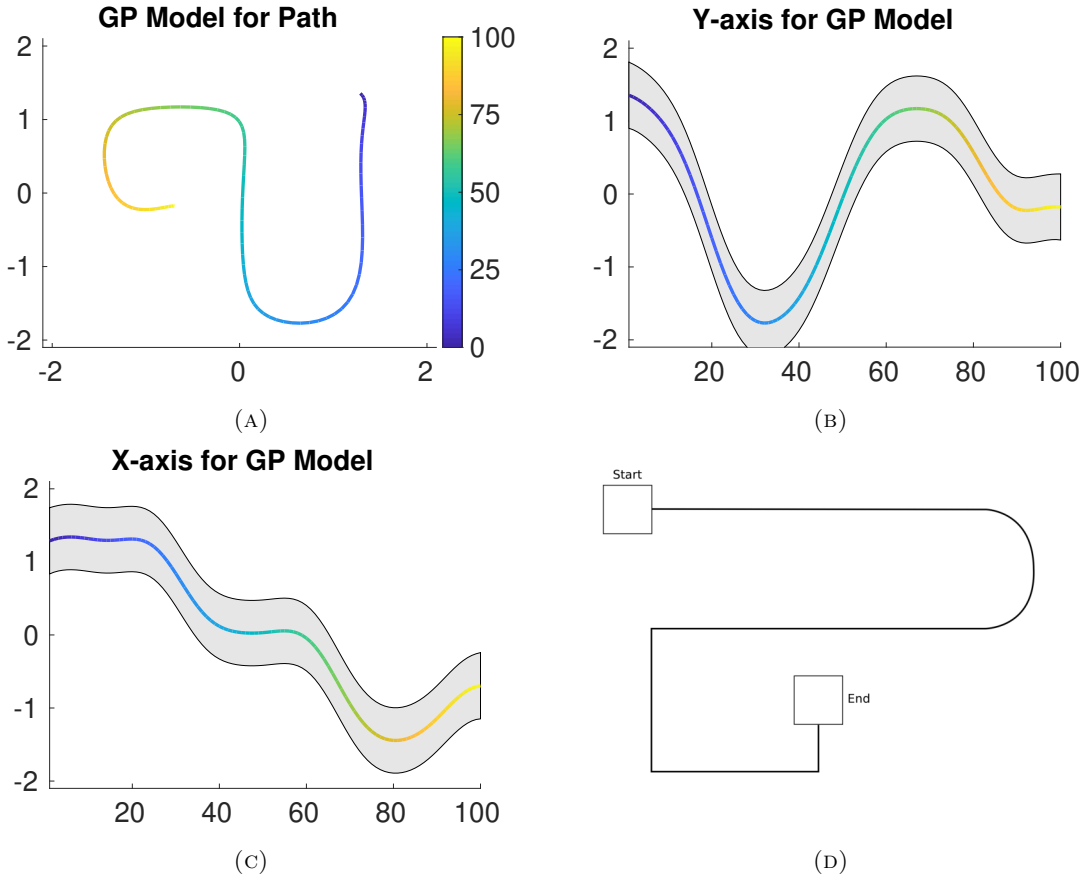

FIGURE C.3: The time-series GP model built for classification based on the data sets for the path shown in the bottom right.

To avoid class imbalance, training samples were reduced to keep a 1:1 ratio of positive and negative training samples. The holdout set was tested against the classification models built to confirm the validity of the trained classifier. During data processing, one of the data sets in the holdout set was found to be erroneous, since it only consisted of half a

TABLE C.1: The posterior probabilities from the GP classification and RMSE-LWSS score from the holdout set.

Path No.	**Mean Class Posterior Probabilities**			
	Class 1	Class 2	Class 3	Class 4
1	**0.69**	0.16	0.26	0.27
2	0.27	**0.79**	0.18	0.22
3	0.28	0.18	**0.73**	0.31
4	0.37	0.30	0.34	**0.72**
	Mean RMSE-LWSS Score			
1	**0.002**	0.238	0.096	0.068
2	0.534	**0.004**	0.145	0.037
3	0.116	0.121	**0.001**	0.016
4	0.143	0.053	0.016	**0.002**

second's worth of data, and was thus discarded. The results from the holdout set are shown in Table C.1 along with the calculated RMSE-LWSS score for comparison.

In the 5-participant test data set, Paths 1 & 2 are identical to the learned models (Class

TABLE C.2: The mean posterior probabilities, RMSE-LWSS score, RMSE, and LWSS score of each class/path in the testing data set.

Path No.	**Mean Class Posterior Probabilities**			
	Class 1	Class 2	Class 3	Class 4
1	**0.797**	0.229	0.352	0.357
2	0.202	**0.812**	0.259	0.297
3	0.240	0.185	**0.509**	0.405
4	0.256	0.268	0.456	**0.493**
	Mean RMSE-LWSS Score			
1	**0.004**	0.271	0.021	0.108
2	0.114	**0.003**	0.133	0.212
3	0.042	0.153	**0.023**	0.034
4	**0.018**	0.048	0.051	0.068
	Mean RMSE			
1	0.164	1.002	**0.137**	0.428
2	0.856	**0.156**	0.974	1.265
3	0.214	1.123	**0.178**	0.278
4	**0.114**	0.897	0.280	0.413
	Mean LWSS			
1	**0.025**	0.272	0.180	0.246
2	0.133	**0.018**	0.136	0.167
3	0.193	0.136	0.130	**0.115**
4	0.155	**0.053**	0.184	0.162

TABLE C.3: The GP Laplace Approximation classification posterior probabilities for z-normalised and raw data, along with position only data and force-inclusive position data.

Path No.	**Mean Class Posterior Probabilities**			
	Class 1	Class 2	Class 3	Class 4
	Z-normalised Position Data			
1	**0.80**	0.23	0.35	0.36
2	0.20	**0.81**	0.26	0.30
3	0.24	0.18	**0.51**	0.41
4	0.26	0.27	0.46	**0.49**
	Z-normalised Position & Force Data			
1	**0.87**	0.21	0.40	0.32
2	0.21	**0.88**	0.38	0.31
3	0.28	0.14	0.44	**0.51**
4	0.31	0.28	**0.42**	0.41
	Raw Position Data			
1	**0.77**	0.20	0.36	0.41
2	0.23	**0.82**	0.25	0.35
3	0.28	0.27	**0.56**	0.44
4	0.32	0.38	**0.44**	0.43
	Raw Position & Force Data			
1	**0.85**	0.50	0.33	0.28
2	0.41	**0.86**	0.30	0.21
3	**0.59**	0.44	0.58	0.49
4	0.55	0.56	0.55	0.56

1 & 2), Path 3 is a modified path of Class 3, and Path 4 is a novel path which has not been seen by any of the GP classifiers. The posterior probabilities from the Laplace Approximation classification and the RMSE-LWSS scores is presented in Table C.2 with the best results highlighted in bold. The component measures of mean RMSE and mean LWSS score are also presented in the table.

C.3 Discussion

The results indicate that an experimental threshold of 0.01 is sufficient to identify input paths which are outliers or novel paths. Looking into the individual component scores, the RMSE and the LWSS, can provide the distinction between the two. The outlier (modified) path is distinguishable by the combination of low RMSE and high RMSE-LWSS score, while the novel path is identified by a high score for both RMSE-LWSS and LWSS.

To explore the efficiency of data utilisation, additional post-hoc analysis was conducted to determine if the use of *all* available data would improve the functional use of the LWSS score. An additional set of binary GP classifiers was built using both position and endpoint force data derived from the robot joint torques. The results are tabulated in Table C.3. While the inclusion of force data mildly improved the classifier's confidence for known paths, its inclusion is highly detrimental to the results for the outlier (modified) and novel paths. It also provides no insight when compared against results obtained from the position only GP classifiers.

One aspect which did affect the performance of the LWSS score is the z-normalisation of all data, as recommended by authors for the multi-dimensional [286]. Z-normalising the data provides properties of size-invariance and robustness which outweighs the benefits of using raw data. This difference in the representations is clearly seen in Figure C.4. Furthermore, the normalisation of data will provide a more agnostic approach towards the integration of RMSE-LWSS score in robotic systems for assessments and assistance.

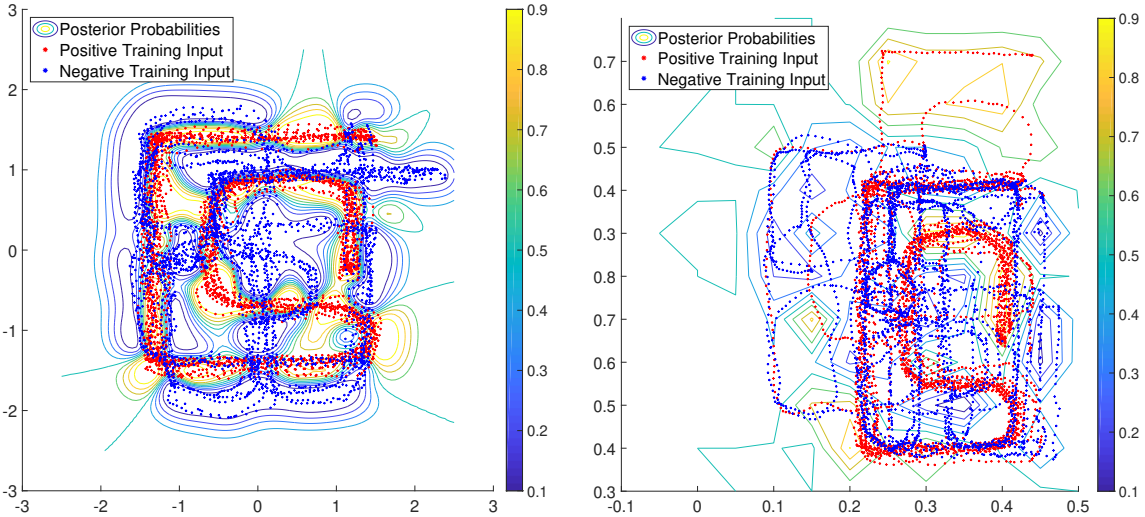

(A) (B)

FIGURE C.4: Two contour plots comparing how data representation can affect the posterior probabilities from the learning model: (A) using z-normalised data; and (B) using raw data.

Bibliography

[1] Bram Vanderborght. Unlocking the potential of industrial human–robot collaboration. Technical report, Publications Office of the EU, Brussels, Belgium, 2 2020.

[2] Graham Gibbs and Savi Sachdev. Canada and the international space station program: Overview and status. *Acta Astronautica*, 51(1-9):591–600, 2002. ISSN 00945765. doi: 10.1016/S0094-5765(02)00077-2.

[3] J. "Bob" Balaram, MiMi Aung, and Matthew P. Golombek. The Ingenuity Helicopter on the Perseverance Rover. *Space Science Reviews*, 217(4):56, 6 2021. ISSN 0038-6308. doi: 10.1007/s11214-021-00815-w.

[4] Sebastian Thrun. When robots meet people. *IEEE Intelligent Systems and Their Applications*, 13(3):27–29, 5 1998. ISSN 10947167. doi: 10.1109/5254.683178.

[5] Jennifer Casper and Robin Roberson Murphy. Human-robot interactions during the robot-assisted urban search and rescue response at the World Trade Center. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 33(3): 367–385, 6 2003. doi: 10.1109/TSMCB.2003.811794.

[6] Timothy D. Barfoot, Jessica Burgner-Kahrs, Emily E. Diller, Animesh Garg, Anna Goldenberg, Jonathan Kelly, Xiaoyun Liu, Hani E. Naguib, Goldie Nejat, Angela P. Schoellig, Florian Shkurti, Hallie Siegel, Yu Sun, Steven L. Waslander, and . Making Sense of the Robotized Pandemic Response: A Comparison of Global and Canadian Robot Deployments and Success Factors. Technical report, University of Toronto Robotics Institute, 9 2020.

[7] Sami Haddadin, Alessandro De Luca, and Alin Albu-Schäffer. Robot collisions: A survey on detection, isolation, and identification. *IEEE Transactions on Robotics*, 33(6):1292–1312, 12 2017. ISSN 15523098. doi: 10.1109/TRO.2017.2723903.

[8] Valeria Villani, Fabio Pini, Francesco Leali, and Cristian Secchi. Survey on human–robot collaboration in industrial settings: Safety, intuitive interfaces and applications. *Mechatronics*, 55:248–266, 11 2018. ISSN 09574158. doi: 10.1016/j.mechatronics.2018.02.009.

[9] Richard Bloss. Collaborative robots are rapidly providing major improvements in productivity, safety, programing ease, portability and cost while addressing many new applications. *Industrial Robot*, 43(5):463–468, 8 2016. ISSN 0143991X. doi: 10.1108/IR-05-2016-0148.

[10] Arash Ajoudani, Andrea Maria Zanchettin, Serena Ivaldi, Alin Albu-Schäffer, Kazuhiro Kosuge, and Oussama Khatib. Progress and prospects of the human–robot collaboration. *Autonomous Robots*, 42(5):957–975, 6 2018. ISSN 15737527. doi: 10.1007/s10514-017-9677-2.

[11] Sami Haddadin and Elizabeth Croft. Physical Human-Robot Interaction. In B Siciliano and O. Khatib, editors, *Springer Handbook of Robotics*, chapter 69, pages 1835–1874. Springer, 2016. ISBN 978-3-319-32552-1. doi: https://doi.org/10.1007/978-3-319-32552-1_69.

[12] Emanuele Magrini, Federica Ferraguti, Andrea Jacopo Ronga, Fabio Pini, Alessandro De Luca, and Francesco Leali. Human-robot coexistence and interaction in open industrial cells. *Robotics and Computer-Integrated Manufacturing*, 61:101846, 2 2020. ISSN 07365845. doi: 10.1016/j.rcim.2019.101846.

[13] Andrea Bajcsy, Dylan P Losey, Marcia K O'malley, and Anca D Dragan. Learning Robot Objectives from Physical Human Interaction. *Proceedings of Machine Learning Research*, 78:217–226, 2017.

[14] Yujun Lai, James Poon, Gavin Paul, Haifeng Han, and Takamitsu Matsubara. Probabilistic Pose Estimation of Deformable Linear Objects. In *2018 IEEE 14th International Conference on Automation Science and Engineering (CASE)*, pages 471–476, Munich, Germany, 2018. IEEE. doi: 10.1109/COASE.2018.8560497.

[15] Yujun Lai, Sheila Sutjipto, Matthew D. Clout, Marc G. Carmichael, and Gavin Paul. GAVRe2 : Towards Data-Driven Upper-Limb Rehabilitation with Adaptive-Feedback Gamification. In *2018 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 164–169. IEEE, 2018. ISBN 978-1-7281-0377-8. doi: 10.1109/ROBIO.2018.8665105.

[16] Yujun Lai, Sheila Sutjipto, Marc Carmichael, and Gavin Paul. Heuristic Detection of Recovery Progress using Robotic Data. In *Proceedings of the IEEE 2019 9th International Conference on Cybernetics and Intelligent Systems and Robotics, Automation and Mechatronics, CIS and RAM 2019*, pages 506–511. IEEE, 2019. ISBN 9781728134581. doi: 10.1109/CIS-RAM47153.2019.9095835.

[17] Sheila Sutjipto, Yujun Lai, Marc G. Carmichael, and Gavin Paul. Fitts' law in the presence of interface inertia. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 4749–4752. IEEE, 2020. ISBN 978-1-7281-1990-8. doi: 10.1109/EMBC44109.2020.9176195.

[18] Yujun Lai, Gavin Paul, Yunduan Cui, and Takamitsu Matsubara. User Intent Estimation during robot learning using Physical Human Robot Interaction Primitives. Under Review in Autonomous Robots (AURO), 2021.

[19] Yujun Lai, Sheila Sutjipto, Marc Carmichael, and Gavin Paul. Preliminary Validation of Upper Limb Musculoskeletal Model using Static Optimization. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2021.

[20] Dinh Tung Le, Sheila Sutjipto, Yujun Lai, and Gavin Paul. Intuitive Virtual Reality based Control of a Real-world Mobile Manipulator. In *16th IEEE International Conference on Control, Automation, Robotics and Vision, ICARCV 2020*, pages 767–772. IEEE, 12 2020. ISBN 9781728177090. doi: 10.1109/ICARCV50220.2020.9305517.

[21] Dac Dang Khoa Nguyen, Yujun Lai, Sheila Sutjipto, and Gavin Paul. Hybrid Multi-Robot System for Drilling and Blasting Automation. In *16th IEEE International Conference on Control, Automation, Robotics and Vision, ICARCV 2020*, pages 79–84. IEEE, 12 2020. ISBN 9781728177090. doi: 10.1109/ICARCV50220.2020.9305391.

[22] Stefano Aldini, Yujun Lai, Marc Carmichael, Gavin Paul, and Dikai Liu. Real-time Estimation of the Strength Capacity of the Upper Limb for Physical Human-Robot Collaboration. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2021.

[23] Petchpansa Chotisathiantham, Yujun Lai, and Gavin Paul. Design of a Wearable Robotic Glove for Rehabilitation. Under Review in Australasian Conference on Robotics and Automation (ACRA), 2021.

[24] Thomas Platz, Cosima Pinkowski, Frederike van Wijck, In-Ha Kim, Paolo di Bella, and Garth Johnson. Reliability and validity of arm function assessment with standardized guidelines for the Fugl-Meyer Test, Action Research Arm Test and Box and Block Test: a multicentre study. *Clinical Rehabilitation*, 19(4):404–411, 6 2005. doi: 10.1191/0269215505cr832oa.

[25] Christian Vogel, Maik Poggendorf, Christoph Walter, and Norbert Elkmann. Towards safe physical human-robot collaboration: A projection-based safety system. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3355–3360. IEEE, 12 2011. ISBN 978-1-61284-456-5. doi: 10.1109/iros.2011.6094550.

[26] Rachid Alami, Alin Albu-Schaeffer, Antonio Bicchi, Rainer Bischoff, Raja Chatila, Alessandro De Luca, Agostino De Santis, Georges Giralt, Jeremie Guiochet, Gerd Hirzinger, Felix Ingrand, Vincenzo Lippiello, Raffaela Mattone, David Powell, Soumen Sen, Bruno Siciliano, Giovanni Tonietti, and Luigi Villani. Safe and dependable physical human-robot interaction in anthropic domains: State of the art and challenges. In *IEEE International Conference on Intelligent Robots and Systems*. IEEE, 2006. doi: 10.1109/IROS.2006.6936985.

[27] Bram Vanderborght, Alin Albu-Schaeffer, Antonio Bicchi, Etienne Burdet, Darwin G. Caldwell, Raffaella Carloni, Manuel Catalano, Oliver Eiberger, Werner Friedl, Ganesh Ganesh, Manolo Garabini, Markus Grebenstein, Giorgio Grioli, Sami Haddadin, Hannes Hoppner, Amir Jafari, Matteo Laffranchi, Dirk Lefeber, Florian Petit, Stefano Stramigioli, Nikolas Tsagarakis, Michael Van Damme, Ronald Van Ham, Ludo C. Visser, and Sebastian Wolf. Variable impedance actuators: A review. *Robotics and Autonomous Systems*, 61(12):1601–1614, 12 2013. ISSN 09218890. doi: 10.1016/j.robot.2013.06.009.

[28] Gill A. Pratt and Matthew M. Williamson. Series elastic actuators. In *IEEE International Conference on Intelligent Robots and Systems*, volume 1, pages 399–406. IEEE, 1995. doi: 10.1109/iros.1995.525827.

[29] Gill A. Pratt, Pace Willisson, Clive Bolton, and Andreas Hoftnan. Late motor processing in low-impedance robots: Impedance control of series-elastic actuators. In *Proceedings of the American Control Conference*, volume 4, pages 3245–3251. IEEE, 2004. ISBN 0780383354. doi: 10.23919/acc.2004.1384410.

[30] Haoyong Yu, Sunan Huang, Gong Chen, Yongping Pan, and Zhao Guo. Human–Robot Interaction Control of Rehabilitation Robots With Series Elastic

Actuators. *IEEE Transactions on Robotics*, 31(5):1089–1100, 10 2015. doi: 10.1109/TRO.2015.2457314.

[31] Kazuhiro Kosuge, H. Yoshida, and T. Fukuda. Dynamic control for robot-human collaboration. In *Proceedings of 1993 2nd IEEE International Workshop on Robot and Human Communication*, pages 398–401. IEEE, 1993. ISBN 0-7803-1407-7. doi: 10.1109/ROMAN.1993.367685.

[32] Neville Hogan. Impedance control: An approach to manipulation: Part I-theory. *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME*, 107(1):1–7, 3 1985. ISSN 15289028. doi: 10.1115/1.3140702.

[33] Giuseppe Averta and Neville Hogan. Enhancing Robot-Environment Physical Interaction via Optimal Impedance Profiles. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, volume 2020-November, pages 973–980. IEEE, 11 2020. ISBN 9781728159072. doi: 10.1109/BioRob49111.2020.9224382.

[34] Marc G Carmichael, Dikai Liu, and Kenneth J Waldron. A framework for singularity-robust manipulator control during physical human-robot interaction. *The International Journal of Robotics Research*, 36(5-7):861–876, 6 2017. doi: 10.1177/0278364917698748.

[35] Aaron Steinfeld, Terrence Fong, David Kaber, Michael Lewis, Jean Scholtz, Alan Schultz, and Michael Goodrich. Common metrics for human-robot interaction. In *Proceeding of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction - HRI '06*, pages 33–40. ACM Press, 2006. ISBN 1595932941. doi: 10.1145/1121241.1121249.

[36] Ravindra Ranasinghe, Lakshitha Dantanarayana, Antony Tran, Stefan Lie, Michael Behrens, and LiYang Liu. Smart hoist: An assistive robot to aid carers. In *2014 13th International Conference on Control Automation Robotics & Vision (ICARCV)*, pages 1285–1291. IEEE, 12 2014. ISBN 978-1-4799-5199-4. doi: 10.1109/ICARCV.2014.7064501.

[37] David J. Reinkensmeyer. How to Retrain Movement after Neurologic Injury: A Computational Rationale for Incorporating Robot (or Therapist) Assistance. In *Annual International Conference of the IEEE Engineering in Medicine and Biology - Proceedings*, volume 2, pages 1479–1482. IEEE, 2003. doi: 10.1109/iembs.2003.1279616.

[38] Rodrigo Pérez-Rodríguez, Carlos Rodríguez, Úrsula Costa, César Cáceres, Josep M. Tormos, Josep Medina, and Enrique J. Gómez. Anticipatory assistance-as-needed control algorithm for a multijoint upper limb robotic orthosis in physical neurorehabilitation. *Expert Systems with Applications*, 41(8):3922–3934, 6 2014. doi: 10.1016/J.ESWA.2013.11.047.

[39] James L. Patton, Mary Ellen Stoykov, Mark Kovic, and Ferdinando A. Mussa-Ivaldi. Evaluation of robotic training forces that either enhance or reduce error in chronic hemiparetic stroke survivors. *Experimental Brain Research*, 168(3):368–383, 1 2006. doi: 10.1007/s00221-005-0097-8.

[40] Hermano I. Krebs, Jerome J. Palazzolo, Laura Dipietro, Mark Ferraro, Jennifer Krol, Karen Rannekleiv, Bruce T. Volpe, and Neville Hogan. Rehabilitation Robotics: Performance-Based Progressive Robot-Assisted Therapy. *Autonomous Robots*, 15 (1):7–20, 2003. doi: 10.1023/A:1024494031121.

[41] Vincent Crocher, Justin Fong, Marlena Klaic, Ying Tan, and Denny Oetomo. Direct versus Indirect Visual Feedback: The Effect of Technology in Neurorehabilitation. In *International IEEE/EMBS Conference on Neural Engineering, NER*, volume 2019-March, pages 105–109. IEEE, 5 2019. ISBN 978-1-5386-7921-0. doi: 10.1109/NER.2019.8716925.

[42] Ruwan A. R. C. Gopura and Kazuo Kiguchi. Mechanical designs of active upper-limb exoskeleton robots: State-of-the-art and design difficulties. In *2009 IEEE International Conference on Rehabilitation Robotics*, pages 178–187. IEEE, 6 2009. ISBN 978-1-4244-3788-7. doi: 10.1109/ICORR.2009.5209630.

[43] Joel C. Perry, Jacob Rosen, and Stephen Burns. Upper-limb powered exoskeleton design. *IEEE/ASME Transactions on Mechatronics*, 12(4):408–417, 8 2007. ISSN 10834435. doi: 10.1109/TMECH.2007.901934.

[44] Arash Ajoudani, Marco Gabiccini, Nikos G. Tsagarakis, and Antonio Bicchi. Human-like impedance and minimum effort control for natural and efficient manipulation. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 4499–4505, 2013. ISBN 9781467356411. doi: 10.1109/ICRA.2013.6631216.

[45] Yuqiang Wu, Fei Zhao, Wansoo Kim, and Arash Ajoudani. An Intuitive Formulation of the Human Arm Active Endpoint Stiffness. *Sensors*, 20(18):5357, 9 2020. ISSN 1424-8220. doi: 10.3390/s20185357.

[46] Jeng Feng Yang, John P. Scholz, and Mark L. Latash. The role of kinematic redundancy in adaptation of reaching. *Experimental Brain Research*, 176(1):54–69, 2007. ISSN 00144819. doi: 10.1007/s00221-006-0602-8.

[47] Jennifer Goetz, Sara Kiesler, and Aaron Powers. Matching robot appearance and behavior to tasks to improve human-robot cooperation. In *The 12th IEEE International Workshop on Robot and Human Interactive Communication*, pages 55–60. IEEE, 2003. ISBN 0-7803-8136-X. doi: 10.1109/ROMAN.2003.1251796.

[48] Andrea Maria Zanchettin, Luca Bascetta, and Paolo Rocco. Achieving Humanlike Motion: Resolving Redundancy for Anthropomorphic Industrial Manipulators. *IEEE Robotics & Automation Magazine*, 20(4):131–138, 12 2013. doi: 10.1109/MRA.2013.2283650.

[49] Shiqiu Gong, Jing Zhao, Ziqiang Zhang, and Biyun Xie. Task motion planning for anthropomorphic arms based on human arm movement primitives. *Industrial Robot*, 47(5):669–681, 6 2020. ISSN 0143991X. doi: 10.1108/IR-12-2019-0261.

[50] Antony Tran. *Robot confidence modeling and role change in physical human-robot collaboration*. PhD thesis, University of Technology Sydney (UTS), Sydney, Australia, 2019.

[51] Amos Freedy, Ewart DeVisser, Gershon Weltman, and Nicole Coeyman. Measurement of trust in human-robot collaboration. In *2007 International Symposium on Collaborative Technologies and Systems*, pages 106–114. IEEE, 5 2007. doi: 10.1109/CTS.2007.4621745.

[52] Aude Billard, Sylvain Calinon, Rüdiger Dillmann, and Stefan Schaal. Robot Programming by Demonstration. In Bruno Siciliano and Oussama Khatib, editors, *Springer Handbook of Robotics*, pages 1371–1394. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. doi: 10.1007/978-3-540-30301-5_60.

[53] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5): 469–483, 5 2009. ISSN 09218890. doi: 10.1016/j.robot.2008.10.024.

[54] Leonel Rozo, Sylvain Calinon, Darwin G. Caldwell, Pablo Jiménez, and Carme Torras. Learning Physical Collaborative Robot Behaviors From Human Demonstrations. *IEEE Transactions on Robotics*, 32(3):513–527, 6 2016. ISSN 15523098. doi: 10.1109/TRO.2016.2540623.

[55] Mattia Racca, Joni Pajarinen, Alberto Montebelli, and Ville Kyrki. Learning in-contact control strategies from demonstration. In *IEEE International Conference on Intelligent Robots and Systems*, pages 688–695. IEEE, 2016. ISBN 9781509037629. doi: 10.1109/IROS.2016.7759127.

[56] Maria Kyrarini, Muhammad Abdul Haseeb, Danijela Ristić-Durrant, and Axel Gräser. Robot learning of industrial assembly task via human demonstrations. *Autonomous Robots*, 43(1):239–257, 1 2019. doi: 10.1007/s10514-018-9725-6.

[57] Mahdi Khoramshahi and Aude Billard. A dynamical system approach to task-adaptation in physical human–robot interaction. *Autonomous Robots*, 43(4):927–946, 4 2019. ISSN 15737527. doi: 10.1007/s10514-018-9764-z.

[58] Affan Pervez, Hiba Latifee, Jee Hwan Ryu, and Dongheui Lee. Motion encoding with asynchronous trajectories of repetitive teleoperation tasks and its extension to human-agent shared teleoperation. *Autonomous Robots*, 43(8):2055–2069, 12 2019. doi: 10.1007/s10514-019-09853-4.

[59] Ed Colgate and Neville Hogan. Analysis of contact instability in terms of passive physical equivalents. In *IEEE International Conference on Robotics and Automation*, pages 404–409. IEEE, 1989. doi: 10.1109/robot.1989.100021.

[60] Luuk M. Doornebosch, David A. Abbink, and Luka Peternel. The Force-Feedback Coupling Effect in Bilateral Tele-Impedance. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, volume 2020-November, pages 152–157. IEEE, 2020. ISBN 9781728159072. doi: 10.1109/BioRob49111.2020.9224296.

[61] Eric Rosen, David Whitney, Elizabeth Phillips, Gary Chien, James Tompkin, George Konidaris, and Stefanie Tellex. Communicating and controlling robot arm motion intent through mixed-reality head-mounted displays. *International Journal of Robotics Research*, 38(12-13):1513–1526, 10 2019. ISSN 17413176. doi: 10.1177/0278364919842925.

[62] David Whitney, Eric Rosen, Elizabeth Phillips, George Konidaris, and Stefanie Tellex. Comparing Robot Grasping Teleoperation Across Desktop and Virtual Reality with ROS Reality. In Nancy M. Amato, Greg Hager, Shawna Thomas, and Miguel Torres-Torriti, editors, *Robotics Research*, pages 335–350. Springer International Publishing, 2020. doi: 10.1007/978-3-030-28619-4_28.

[63] Anna Shafer and Ashish D. Deshpande. Human-like Endtip Stiffness Modulation towards Stable, Dexterous Manipulation. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, pages 1122–1127. IEEE, 2020. ISBN 9781728159072. doi: 10.1109/BioRob49111.2020.9224331.

[64] Andrea Calanca, Riccardo Muradore, and Paolo Fiorini. A review of algorithms for compliant control of stiff and fixed-compliance robots. *IEEE/ASME Transactions on Mechatronics*, 21(2):613–624, 4 2016. ISSN 10834435. doi: 10.1109/TMECH.2015.2465849.

[65] Federica Ferraguti, Chiara Talignani Landi, Lorenzo Sabattini, Marcello Bonfè, Cesare Fantuzzi, and Cristian Secchi. A variable admittance control strategy for stable physical human–robot interaction. *The International Journal of Robotics Research*, 38(6):747–765, 5 2019. ISSN 0278-3649. doi: 10.1177/0278364919840415.

[66] Daniel S. Walker, Robert P. Wilson, and Günter Niemeyer. User-controlled variable impedance teleoperation. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 5352–5357, 2010. ISBN 9781424450381. doi: 10.1109/ROBOT.2010.5509811.

[67] Arash Ajoudani, Nikos Tsagarakis, and Antonio Bicchi. Tele-impedance: Teleoperation with impedance regulation using a body–machine interface. *The International Journal of Robotics Research*, 31(13):1642–1656, 11 2012. ISSN 0278-3649. doi: 10.1177/0278364912464668.

[68] Takahiro Nozaki and Hermano Igo Krebs. Characteristics of Human Behavior in Force Modulation while Performing Force Tracking Tasks. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, pages 240–245. IEEE, 2020. ISBN 9781728159072. doi: 10.1109/BioRob49111.2020.9224348.

[69] Robrecht P.R.D. Van der Wel, Guenther Knoblich, and Natalie Sebanz. Let the Force Be With Us: Dyads Exploit Haptic Coupling for Coordination. *Journal of Experimental Psychology: Human Perception and Performance*, 37(5):1420–1431, 10 2011. ISSN 00961523. doi: 10.1037/a0022337.

[70] Andrew Sawers, Tapomayukh Bhattacharjee, J. Lucas McKay, Madeleine E. Hackney, Charles C. Kemp, and Lena H. Ting. Small forces that differ with prior motor

experience can communicate movement goals during human-human physical interaction. *Journal of NeuroEngineering and Rehabilitation*, 14(1):1–13, 1 2017. ISSN 17430003. doi: 10.1186/s12984-017-0217-2.

[71] Shou-Han Zhou, Denny Oetomo, Ying Tan, Iven Mareels, and Etienne Burdet. Effect of sensory experience on motor learning strategy. *Journal of Neurophysiology*, 113 (4):1077–1084, 2015. doi: 10.1152/JN.00470.2014.

[72] Hermano I. Krebs, Michael Krams, Dimitris K. Agrafiotis, Allitia DiBernardo, Juan C. Chavez, Gary S. Littman, Eric Yang, Geert Byttebier, Laura Dipietro, Avrielle Rykman, Kate McArthur, Karim Hajjar, Kennedy R. Lees, and Bruce T. Volpe. Robotic Measurement of Arm Movements After Stroke Establishes Biomarkers of Motor Recovery. *Stroke*, 45(1):200–204, 1 2014. doi: 10.1161/STROKEAHA.113.002296.

[73] Giuseppe Averta, Cosimo Della Santina, Edoardo Battaglia, Federica Felici, Matteo Bianchi, and Antonio Bicchi. Unvealing the principal modes of human upper limb movements through functional analysis. *Frontiers Robotics AI*, 4, 8 2017. ISSN 22969144. doi: 10.3389/frobt.2017.00037.

[74] Gergely Buza, John Milton, Laszlo Bencsik, and Tamas Insperger. Establishing metrics and control laws for the learning process: ball and beam balancing. *Biological Cybernetics*, 114(1):83–93, 2020. ISSN 14320770. doi: 10.1007/s00422-020-00815-z.

[75] Carlo J. De Luca. The Use of Surface Electromyography in Biomechanics. *Journal of Applied Biomechanics*, 13(2):135–163, 1997. doi: 10.1123/jab.13.2.135.

[76] Paul L. Nunez and Ramesh Srinivasan. *Electric Fields of the Brain: The neurophysics of EEG*. Oxford University Press, 2009. ISBN 9780199865673. doi: 10.1093/acprof:oso/9780195050387.001.0001.

[77] Robert A. McCleary. The nature of the galvanic skin response. *Psychological Bulletin*, 47(2):97–117, 1950. ISSN 00332909. doi: 10.1037/h0059810.

[78] Domen Novak, Benjamin Beyeler, Ximena Omlin, and Robert Riener. Workload Estimation in Physical Human–Robot Interaction Using Physiological Measurements. *Interacting with Computers*, 27(6):616–629, 11 2015. doi: 10.1093/iwc/iwu021.

[79] Yu Shi, Natalie Ruiz, Ronnie Taib, Eric Choi, and Fang Chen. Galvanic skin response (GSR) as an index of cognitive load. In *Conference on Human Factors in*

Computing Systems - Proceedings, pages 2651–2656. ACM, 2007. ISBN 1595936424. doi: 10.1145/1240866.1241057.

[80] Alan Gevins and Michael E. Smith. Neurophysiological measures of cognitive workload during human–computer interaction. *Theoretical Issues in Ergonomics Science*, 4(1-2):113–131, 2003. ISSN 1464536X. doi: 10.1080/14639220210159717.

[81] Stefano Aldini, Ashlesha Akella, Avinash K. Singh, Yu Kai Wang, Marc Carmichael, DIkai Liu, and Chin Teng Lin. Effect of mechanical resistance on cognitive conflict in physical human-robot collaboration. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 6137–6143. IEEE, 2019. ISBN 9781538660263. doi: 10.1109/ICRA.2019.8793748.

[82] Andres F. Salazar-Gomez, Joseph Delpreto, Stephanie Gil, Frank H. Guenther, and Daniela Rus. Correcting robot mistakes in real time using EEG signals. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 6570–6577. IEEE, 2017. ISBN 9781509046331. doi: 10.1109/ICRA.2017.7989777.

[83] Editha M. van Loon, Richard S. W. Masters, Christopher Ring, and David B. McIntyre. Changes in limb stiffness under conditions of mental stress. *Journal of Motor Behavior*, 33(2):153–164, 2001. ISSN 00222895. doi: 10.1080/00222890109603147.

[84] Panagiotis K. Artemiadis and Kostas J. Kyriakopoulos. EMG-Based Position and Force Estimates in Coupled Human-Robot Systems: Towards EMG-Controlled Exoskeletons. In O. Khatib, V. Kumar, and G. J. Pappas, editors, *Experimental Robotics*, pages 241–250. Springer, Berlin, Heidelberg, 2009. ISBN 978-3-642-00195-6. doi: 10.1007/978-3-642-00196-3_29.

[85] Daniele Leonardis, Michele Barsotti, Claudio Loconsole, Massimiliano Solazzi, Marco Troncossi, Claudio Mazzotti, Vincenzo Parenti Castelli, Caterina Procopio, Giuseppe Lamola, Carmelo Chisari, Massimo Bergamasco, and Antonio Frisoli. An EMG-controlled robotic hand exoskeleton for bilateral rehabilitation. *IEEE Transactions on Haptics*, 8(2):140–151, 2015. ISSN 19391412. doi: 10.1109/TOH.2015.2417570.

[86] Masashi Hamaya, Takamitsu Matsubara, Tomoyuki Noda, Tatsuya Teramae, and Jun Morimoto. Learning assistive strategies for exoskeleton robots from user-robot physical interaction. *Pattern Recognition Letters*, 99:67–76, 2017. ISSN 01678655. doi: 10.1016/j.patrec.2017.04.007.

[87] Luka Peternel, Nikos Tsagarakis, Darwin Caldwell, and Arash Ajoudani. Robot adaptation to human physical fatigue in human–robot co-manipulation. *Autonomous Robots*, 42(5):1011–1021, 6 2018. doi: 10.1007/s10514-017-9678-1.

[88] Luka Peternel, Cheng Fang, Nikos Tsagarakis, and Arash Ajoudani. A selective muscle fatigue management approach to ergonomic human-robot co-manipulation. *Robotics and Computer-Integrated Manufacturing*, 58:69–79, 8 2019. doi: 10.1016/J.RCIM.2019.01.013.

[89] Tsuneo Yoshikawa. Manipulability of Robotic Mechanisms. *International Journal of Robotics Research*, 4(2):3–9, 7 1985. ISSN 02783649. doi: 10.1177/027836498500400201.

[90] Kohji Ohta, Yoshiyuki Tanaka, Isao Kawate, and Toshio Tsuji. Human muscular mobility ellipsoid: End-point acceleration manipulability measure in fast motion of human upper arm. *Journal of Biomechanical Science and Engineering*, 9(3):14–00207, 2014. ISSN 1880-9863. doi: 10.1299/jbse.14-00207.

[91] Tadej Petrič, Luka Peternel, Jun Morimoto, and Jan Babič. Assistive Arm-Exoskeleton Control Based on Human Muscular Manipulability. *Frontiers in Neurorobotics*, 13:30, 5 2019. doi: 10.3389/fnbot.2019.00030.

[92] Luka Peternel, Cheng Fang, Marco Laghi, Antonio Bicchi, Nikos Tsagarakis, and Arash Ajoudani. Human Arm Posture Optimisation in Bilateral Teleoperation through Interface Reconfiguration. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, pages 1102–1108. IEEE, 2020. ISBN 9781728159072. doi: 10.1109/BioRob49111.2020.9224339.

[93] Wansoo Kim, Jinoh Lee, Luka Peternel, Nikos Tsagarakis, and Arash Ajoudani. Anticipatory Robot Assistance for the Prevention of Human Static Joint Overloading in Human-Robot Collaboration. *IEEE Robotics and Automation Letters*, 3(1):68–75, 1 2018. ISSN 23773766. doi: 10.1109/LRA.2017.2729666.

[94] Linda van der Spaa, Michael Gienger, Tamas Bates, and Jens Kober. Predicting and Optimizing Ergonomics in Physical Human-Robot Cooperation Tasks. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 1799–1805. IEEE, 2020. ISBN 9781728173955. doi: 10.1109/ICRA40945.2020.9197296.

[95] Dagmar Sternad. It's not (only) the mean that matters: variability, noise and exploration in skill learning. *Current Opinion in Behavioral Sciences*, 20:183–195, 4 2018. doi: 10.1016/J.COBEHA.2018.01.004.

[96] Sebastian Thrun, Wolfram. Burgard, and Dieter. Fox. *Probabilistic robotics*. MIT Press, 2005. ISBN 0262201623.

[97] Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 7 1948. ISSN 0005-8580. doi: 10.1002/J.1538-7305.1948.TB01338.X.

[98] Paul M. Fitts. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47(6):381–391, 1954. ISSN 00221015. doi: 10.1037/h0055392.

[99] Marc Grosjean, Maggie Shiffrar, and Günther Knoblich. Fitts's Law Holds for Action Perception. *Psychological Science*, 18(2):95–99, 2 2007. doi: 10.1111/j.1467-9280.2007.01854.x.

[100] Reinoud J. Bootsma, Laure Fernandez, and Denis Mottet. Behind Fitts' law: kinematic patterns in goal-directed movements. *International Journal of Human-Computer Studies*, 61(6):811–821, 12 2004. doi: 10.1016/J.IJHCS.2004.09.004.

[101] Gerard P. van Galen and Willem P. de Jong. Fitts' law as the outcome of a dynamic noise filtering model of motor control. *Human Movement Science*, 14(4-5):539–571, 11 1995. doi: 10.1016/0167-9457(95)00027-3.

[102] Jos J. Adam, Robin Mol, Jay Pratt, and Martin H. Fischer. Moving Farther but Faster. *Psychological Science*, 17(9):794–798, 9 2006. ISSN 0956-7976. doi: 10.1111/j.1467-9280.2006.01784.x.

[103] Tadej Petrič, Cole S. Simpson, Aleš Ude, and Auke J. Ijspeert. Hammering Does Not Fit Fitts' Law. *Frontiers in Computational Neuroscience*, 11, 5 2017. ISSN 1662-5188. doi: 10.3389/fncom.2017.00045.

[104] Bruno Siciliano and Oussama Khatib. *Springer Handbook of Robotics*. Springer International Publishing, Cham, 2016. ISBN 978-3-319-32550-7. doi: 10.1007/978-3-319-32552-1.

[105] Gamini Dissanayake, Paul Newman, Steven Clark, Hugh F. Durrant-Whyte, and M. Csorba. A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Transactions on Robotics and Automation*, 17(3):229–241, 6 2001. ISSN 1042296X. doi: 10.1109/70.938381.

[106] Yiu-Tong Chan, A. G.C. Hu, and John B. Plant. A Kalman Filter Based Tracking Scheme with Input Estimation. *IEEE Transactions on Aerospace and Electronic Systems*, AES-15(2):237–244, 1979. ISSN 0018-9251. doi: 10.1109/TAES.1979.308710.

[107] Christopher M. Bishop. *Pattern recognition and machine learning*. Springer, 2006. ISBN 9781493938438.

[108] Solomon Kullback. *Information Theory and Statistics*. Dover Publications Inc., Gloucester, Massachusetts, 1978. ISBN 0-8446-5625-9.

[109] Jeffrey C. Schlimmer and Richard H. Granger. Incremental learning from noisy data. *Machine Learning*, 1(3):317–354, 1986. doi: 10.1007/BF00116895.

[110] David D. Lewis. Naive (Bayes) at forty: The independence assumption in information retrieval. In *Machine Learning: ECML-98*, pages 4–15. Springer, Berlin, Heidelberg, 1998. doi: 10.1007/BFb0026666.

[111] Haifeng Han, Gavin Paul, and Takamitsu Matsubara. Model-based reinforcement learning approach for deformable linear object manipulation. In *2017 13th IEEE Conference on Automation Science and Engineering (CASE)*, pages 750–755. IEEE, 8 2017. ISBN 978-1-5090-6781-7. doi: 10.1109/COASE.2017.8256194.

[112] Albert C. Lo, Peter D. Guarino, Lorie G. Richards, Jodie K. Haselkorn, George F. Wittenberg, Daniel G. Federman, Robert J. Ringer, Todd H. Wagner, Hermano I. Krebs, Bruce T. Volpe, Christopher T. Bever, Dawn M. Bravata, Pamela W. Duncan, Barbara H. Corn, Alysia D. Maffucci, Stephen E. Nadeau, Susan S. Conroy, Janet M. Powell, Grant D. Huang, and Peter Peduzzi. Robot-Assisted Therapy for Long-Term Upper-Limb Impairment after Stroke. *New England Journal of Medicine*, 362(19): 1772–1783, 5 2010. doi: 10.1056/NEJMoa0911341.

[113] Daniel F.N. Gordon, Takamitsu Matsubara, Tomoyuki Noda, Tatsuya Teramae, Jun Morimoto, and Sethu Vijayakumar. Bayesian Optimisation of Exoskeleton Design Parameters. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, volume 2018-August, pages 653–658. IEEE Computer Society, 10 2018. ISBN 9781538681831. doi: 10.1109/BIOROB.2018.8487720.

[114] Camilla Pierella, Christian Giang, Elvira Pirondini, Nawal Kinany, Martina Coscia, Jenifer Miehlbradt, Cecile Magnin, Pierre Nicolo, Adrian G. Guggisberg, and Silvestro Micera. Personalizing Exoskeleton-Based Upper Limb Rehabilitation Using

a Statistical Model: A Pilot Study. In *Biosystems and Biorobotics*, pages 117–121. Springer, 21 edition, 2019. doi: 10.1007/978-3-030-01845-0_23.

[115] Francisco Javier Badesa, Ricardo Morales, Nicolas M. Garcia-Aracil, Jose M. Sabater, Loredana Zollo, Eugenia Papaleo, and Eugenio Guglielmelli. Dynamic Adaptive System for Robot-Assisted Motion Rehabilitation. *IEEE Systems Journal*, 10(3):984–991, 9 2016. doi: 10.1109/JSYST.2014.2318594.

[116] Heriberto Cuayáhuitl, Martijn van Otterlo, Nina Dethlefs, and Lutz Frommberger. Machine learning for interactive systems and robots. In *Proceedings of the 2nd Workshop on Machine Learning for Interactive Systems Bridging the Gap Between Perception, Action and Communication - MLIS '13*, pages 19–28, New York, New York, USA, 2013. ACM Press. ISBN 9781450320191. doi: 10.1145/2493525.2493530.

[117] Ruikun Luo, Rafi Hayne, and Dmitry Berenson. Unsupervised early prediction of human reaching for human–robot collaboration in shared workspaces. *Autonomous Robots*, 42(3):631–648, 3 2018. ISSN 15737527. doi: 10.1007/s10514-017-9655-8.

[118] Sergey Levine. Reinforcement Learning and Control as Probabilistic Inference: Tutorial and Review. Technical report, UC Berkeley, Berkeley, CA, USA, 5 2018.

[119] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521 (7553):436–444, 5 2015. doi: 10.1038/nature14539.

[120] H. Sebastian Seung and Daniel D. Lee. The Manifold Ways of Perception. *Science*, 290(5500):2268–2269, 12 2000. doi: 10.1126/SCIENCE.290.5500.2268.

[121] Yann Lecun, Leon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.

[122] Joseph Redmon and Ali Farhadi. YOLO9000: Better, Faster, Stronger. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6517–6525. IEEE, 7 2017. ISBN 978-1-5386-0457-1. doi: 10.1109/CVPR.2017.690.

[123] Alexander Toshev and Christian Szegedy. DeepPose: Human pose estimation via deep neural networks. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1653–1660. IEEE Computer Society, 2014. ISBN 9781479951178. doi: 10.1109/CVPR.2014.214.

[124] Riccardo Miotto, Fei Wang, Shuang Wang, Xiaoqian Jiang, and Joel T Dudley. Deep learning for healthcare: review, opportunities and challenges. *Briefings in Bioinformatics*, 19(6):1236–1246, 11 2018. doi: 10.1093/bib/bbx044.

[125] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. Deep Reinforcement Learning: A Brief Survey. *IEEE Signal Processing Magazine*, 34(6):26–38, 11 2017. doi: 10.1109/MSP.2017.2743240.

[126] Yoshihisa Tsurumine, Yunduan Cui, Eiji Uchibe, and Takamitsu Matsubara. Deep dynamic policy programming for robot control with raw images. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1545–1550. IEEE, 9 2017. ISBN 978-1-5386-2682-5. doi: 10.1109/IROS.2017.8205960.

[127] Ali Ghadirzadeh, Judith Bütepage, Atsuto Maki, Danica Kragic, and Mårten Björkman. A sensorimotor reinforcement learning framework for physical human-robot interaction. In *IEEE International Conference on Intelligent Robots and Systems*, pages 2682–2688. IEEE, 2016. ISBN 9781509037629. doi: 10.1109/IROS.2016.7759417.

[128] Carl Edward. Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006. ISBN 026218253X.

[129] John Schulman, Alex Lee, Jonathan Ho, and Pieter Abbeel. Tracking deformable objects with point clouds. In *2013 IEEE International Conference on Robotics and Automation*, pages 1130–1137. IEEE, 5 2013. ISBN 978-1-4673-5643-5. doi: 10.1109/ICRA.2013.6630714.

[130] Marc Peter Deisenroth, Dieter Fox, and Carl Edward Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):408–423, 2 2015. ISSN 01628828. doi: 10.1109/TPAMI.2013.218.

[131] Jack M. Wang, David J. Fleet, and Aaron Hertzmann. Gaussian Process Dynamical Models for Human Motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):283–298, 2 2008. doi: 10.1109/TPAMI.2007.1167.

[132] Richard O. Duda and Peter Elliott Hart. *Pattern classification and scene analysis*. Wiley, New York, 1973. ISBN 0471223611.

[133] Joaquin Quiñonero-Candela, Agathe Girard, and Carl Edward Rasmussen. Prediction at an Uncertain Input for Gaussian Processes and Relevance Vector Machines

Application to Multiple-Step Ahead Time-Series Forecasting. Technical report, Technical University of Denmark, Kongens Lyngby, 2003.

[134] Sunho Park and Seungjin Choi. Hierarchical Gaussian Process Regression. In Masashi Sugiyama and Qiang Yang, editors, *Proceedings of 2nd Asian Conference on Machine Learning*, pages 95–110. MLR, 2010.

[135] Marc Peter Deisenroth and Jun Wei Ng. Distributed Gaussian Processes. In Francis Bach and David Blei, editors, *Proceedings of Machine Learning Research*, volume 37, pages 1481–1490. MLR, 2015.

[136] Alex J. Smola and Peter L. Bartlett. Sparse Greedy Gaussian Process Regression. In *Advances in Neural Information Processing Systems 13 (NIPS)*, pages 619–625, 2001.

[137] Neil Lawrence, Matthias Seeger, and Ralf Herbrich. Fast sparse Gaussian process methods: the informative vector machine. In *Proceedings of the 15th International Conference on Neural Information Processing Systems*, pages 625–632. MIT Press, 2002.

[138] Neil D. Lawrence. Gaussian Process Latent Variable Models for Visualisation of High Dimensional Data. In *Advances in Neural Information Processing Systems 16 (NIPS 2003)*, pages 329–336, 2003.

[139] Bertil Matérn. *Spatial Variation*, volume 36 of *Lecture Notes in Statistics*. Springer, 1986. ISBN 978-0-387-96365-5. doi: 10.1007/978-1-4615-7892-5.

[140] Martin Pelikan, David E. Goldberg, and Erick Cantu-Paz. BOA: the Bayesian optimization algorithm. In *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation - Volume 1*, pages 525–532. Morgan Kaufmann Publishers, 1999. ISBN 1558606114. doi: https://dl.acm.org/doi/pdf/10.5555/2933923.2933973.

[141] Katta G. Murty. Quadratic Programming. In Saul I. Gass and Michael C. Fu, editors, *Encyclopedia of Operations Research and Management Science*, pages 1207–1215. Springer US, Boston, MA, 2013. ISBN 978-1-4419-1137-7. doi: https://doi.org/10.1007/978-1-4419-1153-7_838.

[142] John A. Nelder and Roger Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313, 1 1965. doi: 10.1093/comjnl/7.4.308.

[143] Rainer Storn and Kenneth Price. Differential Evolution – A Simple and Efficient Heuristic for global Optimization over Continuous Spaces. *Journal of Global Optimization*, 11(4):341–359, 1997. doi: 10.1023/A:1008202821328.

[144] Christopher K. I. Williams and David Barber. Bayesian classification with Gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12): 1342–1351, 1998. doi: 10.1109/34.735807.

[145] Dylan P. Losey, Craig G. McDonald, Edoardo Battaglia, and Marcia K. O'Malley. A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction. *Applied Mechanics Reviews*, 70(1), 2018. ISSN 00036900. doi: 10.1115/1.4039145.

[146] Kelsey P. Hawkins, Shray Bansal, Nam N. Vo, and Aaron F. Bobick. Anticipating human actions for collaboration in the presence of task and sensor uncertainty. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2215–2222. IEEE, 5 2014. ISBN 978-1-4799-3685-4. doi: 10.1109/ICRA.2014.6907165.

[147] Kelsey P. Hawkins, Nam Vo, Shray Bansal, and Aaron F. Bobick. Probabilistic human action prediction and wait-sensitive planning for responsive human-robot collaboration. In *IEEE-RAS International Conference on Humanoid Robots*, pages 499–506. IEEE, 2015. ISBN 9781479926183. doi: 10.1109/HUMANOIDS.2013.7030020.

[148] Yue Wang, Jie Cai, Yabiao Wang, Youzhong Hu, Rong Xiong, Yong Liu, Jiafan Zhang, and Liwei Qi. Probabilistic graph based spatial assembly relation inference for programming of assembly task by demonstration. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4402–4407. IEEE, 9 2015. ISBN 978-1-4799-9994-1. doi: 10.1109/IROS.2015.7354002.

[149] Tarek Taha, Jaime Valls Miro, and Gamini Dissanayake. A POMDP framework for modelling human interaction with assistive robots. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 544–549, 2011. ISBN 9781612843865. doi: 10.1109/ICRA.2011.5980323.

[150] Tarek Taha, Jaime Valls Miró, and Gamini Dissanayake. POMDP-based long-term user intention prediction for wheelchair navigation. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 3920–3925, 2008. ISBN 978-1-4244-1646-2. doi: 10.1109/ROBOT.2008.4543813.

[151] James Poon, Yunduan Cui, Jaime Valls Miro, and Takamitsu Matsubara. Learning Mobility Aid Assistance via Decoupled Observation Models. In *2018 15th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pages 1903–1910. IEEE, 11 2018. ISBN 978-1-5386-9582-1. doi: 10.1109/ICARCV.2018.8581375.

[152] Andrew M. Fraser. *Hidden Markov models and dynamical systems*. Society for Industrial and Applied Mathematics, 2008. ISBN 9780898716658. doi: https://doi.org/10.1137/1.9780898717747.

[153] Junji Yamato, Jun Ohya, and Kenichiro Ishii. Recognizing human action in time-sequential images using hidden Markov model. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 379–385. IEEE, 1992. ISBN 0818628553. doi: 10.1109/CVPR.1992.223161.

[154] Takamitsu Matsubara, Jaime Valls Miro, Daisuke Tanaka, James Poon, and Kenji Sugimoto. Sequential intention estimation of a mobility aid user for intelligent navigational assistance. In *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, pages 444–449. IEEE, 2015. ISBN 9781467367042. doi: 10.1109/ROMAN.2015.7333580.

[155] Nam T. Nguyen, Dinh Q. Phung, Svetha Venkatesh, and Hung Bui. Learning and detecting activities from movement trajectories using the hierarchical hidden markov model. In *Proceedings - 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2005*, volume II, pages 955–960. IEEE Computer Society, 2005. ISBN 0769523722. doi: 10.1109/CVPR.2005.203.

[156] Jimmy Baraglia, Maya Cakmak, Yukie Nagai, Rajesh PN Rao, and Minoru Asada. Efficient human-robot collaboration: When should a robot take initiative? *The International Journal of Robotics Research*, 36(5-7):563–579, 6 2017. doi: 10.1177/0278364916688253.

[157] Yifan Shi, Yan Huang, D. Minnen, A. Bobick, and I. Essa. Propagation networks for recognition of partially ordered sequential action. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004*, volume 2, pages 862–869. IEEE, 2004. ISBN 0-7695-2158-4. doi: 10.1109/CVPR.2004.1315255.

[158] Raquel Urtasun, David J. Fleet, and Pascal Fua. 3D people tracking with Gaussian process dynamical models. In *Proceedings of the IEEE Computer Society Conference*

on Computer Vision and Pattern Recognition, volume 1, pages 238–245. IEEE, 2006. ISBN 0-7695-2597-0. doi: 10.1109/CVPR.2006.15.

[159] Yuri A. Ivanov and Aaron F. Bobick. Recognition of visual activities and interactions by stochastic parsing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):852–872, 2000. doi: 10.1109/34.868686.

[160] Auke Jan Ijspeert, Jun Nakanishi, and Stefan Schaal. Trajectory formation for imitation with nonlinear dynamical systems. In *IEEE International Conference on Intelligent Robots and Systems*, volume 2, pages 752–757, 2001. doi: 10.1109/iros.2001.976259.

[161] S. Mohammad Khansari-Zadeh and Aude Billard. Learning stable nonlinear dynamical systems with Gaussian mixture models. *IEEE Transactions on Robotics*, 27(5): 943–957, 10 2011. ISSN 15523098. doi: 10.1109/TRO.2011.2159412.

[162] Neville Hogan and Dagmar Sternad. Dynamic primitives in the control of locomotion. *Frontiers in Computational Neuroscience*, 7:71, 6 2013. doi: 10.3389/fncom.2013.00071.

[163] Andrej Gams, Bojan Nemec, Auke Jan Ijspeert, and Aleš Ude. Coupling movement primitives: Interaction with the environment and bimanual tasks. *IEEE Transactions on Robotics*, 30(4):816–830, 2014. ISSN 15523098. doi: 10.1109/TRO.2014.2304775.

[164] Miha Denisa, Andrej Gams, Ales Ude, and Tadej Petric. Learning Compliant Movement Primitives Through Demonstration and Statistical Generalization. *IEEE/ASME Transactions on Mechatronics*, 21(5):2581–2594, 10 2016. ISSN 10834435. doi: 10.1109/TMECH.2015.2510165.

[165] Tomas Kulvicius, Martin Biehl, Mohamad Javad Aein, Minija Tamosiunaite, and Florentin Wörgötter. Interaction learning for dynamic movement primitives used in cooperative robotic tasks. *Robotics and Autonomous Systems*, 61(12):1450–1459, 12 2013. ISSN 09218890. doi: 10.1016/j.robot.2013.07.009.

[166] Yanlong Huang, Leonel Rozo, João Silvério, and Darwin G Caldwell. Kernelized movement primitives. *The International Journal of Robotics Research*, 38(7):833–852, 6 2019. ISSN 0278-3649. doi: 10.1177/0278364919846363.

[167] Andrej Gams, Tadej Petrič, Martin Do, Bojan Nemec, Jun Morimoto, Tamim Asfour, and Aleš Ude. Adaptation and coaching of periodic motion primitives through

physical and visual interaction. *Robotics and Autonomous Systems*, 75:340–351, 1 2016. ISSN 09218890. doi: 10.1016/j.robot.2015.09.011.

[168] Miguel Prada, Anthony Remazeilles, Ansgar Koene, and Satoshi Endo. Dynamic Movement Primitives for Human-Robot interaction: Comparison with human behavioral observation. In *IEEE International Conference on Intelligent Robots and Systems*, pages 1168–1175, 2013. ISBN 9781467363587. doi: 10.1109/IROS.2013.6696498.

[169] James Hermus, Dagmar Sternad, and Neville Hogan. Evidence for Dynamic Primitives in a Constrained Motion Task. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, volume 2020-November, pages 551–556. IEEE Computer Society, 11 2020. ISBN 9781728159072. doi: 10.1109/BioRob49111.2020.9224352.

[170] Stefan Schaal, Jan Peters, Jun Nakanishi, and Auke Ijspeert. Learning movement primitives. *Springer Tracts in Advanced Robotics*, 15:561–572, 2005. ISSN 1610742X. doi: 10.1007/11008941_60.

[171] Auke Jan Ijspeert. Central pattern generators for locomotion control in animals and robots: A review. *Neural Networks*, 21(4):642–653, 5 2008. ISSN 08936080. doi: 10.1016/j.neunet.2008.03.014.

[172] Peter Pastor, Heiko Hoffmann, Tamim Asfour, and Stefan Schaal. Learning and generalization of motor skills by learning from demonstration. In *IEEE International Conference on Robotics and Automation*, pages 763–768. IEEE, 2009. doi: 10.1109/robot.2009.5152385.

[173] Michele Ginesi, Nicola Sansonetto, and Paolo Fiorini. Overcoming Some Drawbacks of Dynamic Movement Primitives. *arXiv preprint arXiv:1908.10608*, 2021.

[174] Heiko Hoffmann, Peter Pastor, Dae-Hyung Park, and Stefan Schaal. Biologically-inspired dynamical systems for movement generation: Automatic real-time goal adaptation and obstacle avoidance. In *2009 IEEE International Conference on Robotics and Automation*, pages 2587–2592. IEEE, 2009. ISBN 978-1-4244-2788-8. doi: 10.1109/robot.2009.5152423.

[175] Michele Ginesi, Daniele Meli, Andrea Calanca, Diego Dall'Alba, Nicola Sansonetto, and Paolo Fiorini. Dynamic Movement Primitives: Volumetric Obstacle Avoidance. In *2019 19th International Conference on Advanced Robotics (ICAR)*, pages

234–239. Institute of Electrical and Electronics Engineers (IEEE), 12 2019. ISBN 9781728124674. doi: 10.1109/icar46387.2019.8981552.

[176] Martin Karlsson, Fredrik Bagge Carlson, Anders Robertsson, and Rolf Johansson. Two-Degree-of-Freedom Control for Trajectory Tracking and Perturbation Recovery during Execution of Dynamical Movement Primitives. *IFAC-PapersOnLine*, 50(1): 1923–1930, 7 2017. ISSN 24058963. doi: 10.1016/j.ifacol.2017.08.383.

[177] Takamitsu Matsubara, Sang Ho Hyon, and Jun Morimoto. Learning stylistic dynamic movement primitives from multiple demonstrations. In *IEEE/RSJ 2010 International Conference on Intelligent Robots and Systems, IROS 2010*, pages 1277–1283. IEEE, 2010. ISBN 9781424466757. doi: 10.1109/IROS.2010.5651049.

[178] Takamitsu Matsubara, Sang Ho Hyon, and Jun Morimoto. Learning parametric dynamic movement primitives from multiple demonstrations. *Neural Networks*, 24 (5):493–500, 6 2011. ISSN 08936080. doi: 10.1016/j.neunet.2011.02.004.

[179] Rui Huang, Hong Cheng, Jing Qiu, and Jianwei Zhang. Learning Physical Human-Robot Interaction with Coupled Cooperative Primitives for a Lower Exoskeleton. *IEEE Transactions on Automation Science and Engineering*, 16(4):1566–1574, 10 2019. ISSN 15583783. doi: 10.1109/TASE.2018.2886376.

[180] Elmar Rückert and Andrea d'Avella. Learned parametrized dynamic movement primitives with shared synergies for controlling robotic and musculoskeletal systems. *Frontiers in Computational Neuroscience*, 7(OCT):138, 10 2013. ISSN 16625188. doi: 10.3389/fncom.2013.00138.

[181] Affan Pervez and Dongheui Lee. Learning task-parameterized dynamic movement primitives using mixture of GMMs. *Intelligent Service Robotics*, 11(1):61–78, 1 2018. ISSN 18612784. doi: 10.1007/s11370-017-0235-8.

[182] Zhuang Mei, Yang Chen, Minghao Jiang, Huaiyu Wu, and Lei Cheng. Mobile robots path planning based on dynamic movement primitives library. In *Chinese Control Conference, CCC*, pages 6906–6911. IEEE, 2017. ISBN 9789881563934. doi: 10.23919/ChiCC.2017.8028446.

[183] Peter Pastor, Mrinal Kalakrishnan, Franziska Meier, Freek Stulp, Jonas Buchli, Evangelos Theodorou, and Stefan Schaal. From dynamic movement primitives to associative skill memories. *Robotics and Autonomous Systems*, 61(4):351–361, 4 2013. ISSN 09218890. doi: 10.1016/j.robot.2012.09.017.

[184] Rudolf Lioutikov, Gerhard Neumann, Guilherme Maeda, and Jan Peters. Learning movement primitive libraries through probabilistic segmentation. *International Journal of Robotics Research*, 2017. ISSN 17413176. doi: 10.1177/0278364917713116.

[185] Jens Kober, Michael Gienger, and Jochen J. Steil. Learning movement primitives for force interaction tasks. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 3192–3199. IEEE, 2015. ISBN 9781479969234. doi: 10.1109/ICRA.2015.7139639.

[186] Heni Ben Amor, Gerhard Neumann, Sanket Kamthe, Oliver Kroemer, and Jan Peters. Interaction primitives for human-robot cooperation tasks. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 2831–2837. IEEE, 2014. doi: 10.1109/ICRA.2014.6907265.

[187] Yunduan Cui, James Poon, Jaime Valls Miro, Kimitoshi Yamazaki, Kenji Sugimoto, and Takamitsu Matsubara. Environment-adaptive interaction primitives through visual context for human–robot motor skill learning. *Autonomous Robots*, 43(5): 1225–1240, 6 2019. doi: 10.1007/s10514-018-9798-2.

[188] Dominik Widmann and Yiannis Karayiannidis. Human Motion Prediction in Human-Robot Handovers based on Dynamic Movement Primitives. In *2018 European Control Conference, ECC 2018*, pages 2781–2787. IEEE, 2018. ISBN 9783952426982. doi: 10.23919/ECC.2018.8550170.

[189] Antonis Sidiropoulos, Yiannis Karayiannidis, and Zoe Doulgeri. Human-robot collaborative object transfer using human motion prediction based on dynamic movement primitives. In *2019 18th European Control Conference, ECC 2019*, pages 2583–2588. IEEE, 2019. ISBN 9783907144008. doi: 10.23919/ECC.2019.8796249.

[190] Guilherme J. Maeda, Gerhard Neumann, Marco Ewerton, Rudolf Lioutikov, Oliver Kroemer, and Jan Peters. Probabilistic movement primitives for coordination of multiple human–robot collaborative tasks. *Autonomous Robots*, 41(3):593–612, 3 2017. doi: 10.1007/s10514-016-9556-2.

[191] Joseph Campbell, Simon Stepputtis, and Heni Ben Amor. Probabilistic Multimodal Modeling for Human-Robot Interaction Tasks. In *Proceedings of Robotics: Science and Systems*, 2019. doi: 10.15607/RSS.2019.XV.047.

[192] Chenguang Yang, Chao Zeng, Cheng Fang, Wei He, and Zhijun Li. A DMPs-Based Framework for Robot Learning and Generalization of Humanlike Variable Impedance

Skills. *IEEE/ASME Transactions on Mechatronics*, 23(3):1193–1203, 6 2018. ISSN 10834435. doi: 10.1109/TMECH.2018.2817589.

[193] Bojan Nemec, Nejc Likar, Andrej Gams, and Aleš Ude. Human robot cooperation with compliance adaptation along the motion trajectory. *Autonomous Robots*, 42 (5):1023–1035, 6 2018. ISSN 15737527. doi: 10.1007/s10514-017-9676-3.

[194] Emre Ugur and Hakan Girgin. Compliant parametric dynamic movement primitives. *Robotica*, 38(3):457–474, 3 2020. doi: 10.1017/S026357471900078X.

[195] Joseph Campbell and Katsu Yamane. Learning Whole-Body Human-Robot Haptic Interaction in Social Contexts. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 10177–10183. IEEE, 5 2020. ISBN 9781728173955. doi: 10.1109/ICRA40945.2020.9196933.

[196] Alexandros Paraschos, Christian Daniel, Jan R Peters, and Gerhard Neumann. Probabilistic Movement Primitives. In C J C Burges, L Bottou, M Welling, Z Ghahramani, and K Q Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 2616–2624. Curran Associates, Inc., 2013.

[197] Alexandros Paraschos, Christian Daniel, Jan Peters, and Gerhard Neumann. Using probabilistic movement primitives in robotics. *Autonomous Robots*, 42(3):529–551, 3 2018. ISSN 15737527. doi: 10.1007/s10514-017-9648-7.

[198] Dorothea Koert, Joni Pajarinen, Albert Schotschneider, Susanne Trick, Constantin Rothkopf, and Jan Peters. Learning Intention Aware Online Adaptation of Movement Primitives. *IEEE Robotics and Automation Letters*, 4(4):3719–3726, 10 2019. ISSN 23773766. doi: 10.1109/LRA.2019.2928760.

[199] Hiroaki Sakoe and S. Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):43–49, 2 1978. doi: 10.1109/TASSP.1978.1163055.

[200] Guilherme Maeda, Marco Ewerton, Gerhard Neumann, Rudolf Lioutikov, and Jan Peters. Phase estimation for fast action recognition and trajectory generation in human–robot collaboration. *The International Journal of Robotics Research*, 36(13-14):1579–1594, 12 2017. doi: 10.1177/0278364917693927.

[201] Joseph Campbell and Heni Ben Amor. Bayesian Interaction Primitives: A SLAM Approach to Human-Robot Interaction. In Sergey Levine, Vincent Vanhoucke, and Ken Goldberg, editors, *Proceedings of the 1st Annual Conference on Robot Learning*,

volume 78 of *Proceedings of Machine Learning Research*, pages 379–387. MLR, 11 2017.

[202] Carl Edward Rasmussen. The Infinite Gaussian Mixture Model. In S. A. Solla, T. K. Leen, and K.-R. Muller, editors, *Advances in information processing systems 12*, pages 554–560. MIT Press, 2000. ISBN 9780262194501.

[203] Sylvain Calinon, Florent D'Halluin, Eric L. Sauser, Darwin G. Caldwell, and Aude G. Billard. Learning and reproduction of gestures by imitation. *IEEE Robotics and Automation Magazine*, 17(2):44–54, 6 2010. ISSN 10709932. doi: 10.1109/MRA.2010.936947.

[204] You Zhou, Jianfeng Gao, and Tamim Asfour. Movement primitive learning and generalization: Using mixture density networks. *IEEE Robotics and Automation Magazine*, 27(2):22–32, 6 2020. ISSN 1558223X. doi: 10.1109/MRA.2020.2980591.

[205] Kenneth H. Hunt and Erskine Crossley. Coefficient of restitution interpreted as damping in vibroimpact. *Journal of Applied Mechanics, Transactions ASME*, 42(2): 440–445, 1975. ISSN 15289036. doi: 10.1115/1.3423596.

[206] David A. Winter. *Biomechanics and Motor Control of Human Movement*. John Wiley & Sons, Inc., Hoboken, NJ, USA, 4 edition, 9 2009. ISBN 9780470549148. doi: 10.1002/9780470549148.

[207] Vladimir M. Zatsiorsky and Boris I. Prilutsky. *Biomechanics of Skeletal Muscles*. Human Kinetics Publisher, 2012. ISBN 9780736080200.

[208] Katrin Stollenmaier, Winfried Ilg, and Daniel F. B. Haeufle. Predicting Perturbed Human Arm Movements in a Neuro-Musculoskeletal Model to Investigate the Muscular Force Response. *Frontiers in Bioengineering and Biotechnology*, 8:308, 2020. ISSN 2296-4185. doi: 10.3389/fbioe.2020.00308.

[209] Richardo Khonasty, Marc G Carmichael, Dikai Liu, and Stefano Aldini. Effect of External Force and Bimanual Operation on Upper Limb Pose during Human-Robot Collaboration. In *Australasian Conference on Robotics and Automation (ACRA)*, pages 81–89. Australian Robotics and Automation Association, 2017.

[210] Ajay Seth and Marcus G. Pandy. A neuromusculoskeletal tracking method for estimating individual muscle forces in human movement. *Journal of Biomechanics*, 40 (2):356–366, 2007. ISSN 00219290. doi: 10.1016/j.jbiomech.2005.12.017.

[211] Scott L. Delp and J. Peter Loan. A graphics-based software system to develop and analyze models of musculoskeletal structures. *Computers in Biology and Medicine*, 25(1):21–34, 1995. doi: 10.1016/0010-4825(95)98882-E.

[212] John Rasmussen, Michael Damsgaard, Egidijus Surma, Søren T Christensen, Mark De Zee, and Vit Vondrak. AnyBody-a software system for ergonomic optimization. In *Fifth World Congress on Structural and Multidisciplinary Optimization*, 2003.

[213] Antonie J. van den Bogert, Thomas Geijtenbeek, Oshri Even-Zohar, Frans Steenbrink, and Elizabeth C. Hardin. A real-time system for biomechanical analysis of human movement and muscle function. *Medical & Biological Engineering & Computing*, 51(10):1069–1077, 2013. doi: 10.1007/s11517-013-1076-z.

[214] Neville Hogan. The mechanics of multi-joint posture and movement control. *Biological Cybernetics*, 52(5):315–331, 9 1985. doi: 10.1007/BF00355754.

[215] Domenico Buongiorno, Michele Barsotti, Edoardo Sotgiu, Claudio Loconsole, Massimiliano Solazzi, Vitoantonio Bevilacqua, and Antonio Frisoli. A neuromusculoskeletal model of the human upper limb for a myoelectric exoskeleton control using a reduced number of muscles. In *2015 IEEE World Haptics Conference (WHC)*, pages 273–279. IEEE, 6 2015. ISBN 978-1-4799-6624-0. doi: 10.1109/WHC.2015.7177725.

[216] Apoorva Rajagopal, Christopher L. Dembia, Matthew S. DeMers, Denny D. Delp, Jennifer L. Hicks, and Scott L. Delp. Full-Body Musculoskeletal Model for Muscle-Driven Simulation of Human Gait. *IEEE Transactions on Biomedical Engineering*, 63(10):2068–2079, 10 2016. ISSN 0018-9294. doi: 10.1109/TBME.2016.2586891.

[217] Alexander G. Bruno, Mary L. Bouxsein, and Dennis E. Anderson. Development and Validation of a Musculoskeletal Model of the Fully Articulated Thoracolumbar Spine and Rib Cage. *Journal of Biomechanical Engineering*, 137(8), 8 2015. ISSN 1528-8951. doi: 10.1115/1.4030408.

[218] Heiko Stark, Martin S. Fischer, Alexander Hunt, Fletcher Young, Roger Quinn, and Emanuel Andrada. A three-dimensional musculoskeletal model of the dog. *Scientific Reports 2021 11:1*, 11(1):1–13, 5 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-90058-0.

[219] Ajay Seth, Jennifer L. Hicks, Thomas K. Uchida, Ayman Habib, Christopher L. Dembia, James J. Dunne, Carmichael F. Ong, Matthew S. DeMers, Apoorva Rajagopal, Matthew Millard, Samuel R. Hamner, Edith M. Arnold, Jennifer R. Yong,

Shrinidhi K. Lakshmikanth, Michael A. Sherman, Joy P. Ku, and Scott L. Delp. OpenSim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement. *PLoS Computational Biology*, 14(7), 2018. doi: 10.1371/journal.pcbi.1006223.

[220] Vishal Raveendranathan and Raffaella Carloni. Musculoskeletal Model of an Osseointegrated Transfemoral Amputee in OpenSim. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, pages 1196–1201. IEEE, 2020. ISBN 9781728159072. doi: 10.1109/BioRob49111.2020.9224422.

[221] Antoine Falisse, Sam Van Rossom, Johannes Gijsbers, Frans Steenbrink, Ben J.H. van Basten, Ilse Jonkers, Antonie J. van den Bogert, and Friedl De Groote. OpenSim Versus Human Body Model: A Comparison Study for the Lower Limbs During Gait. *Journal of Applied Biomechanics*, 34(6):496–502, 2018. doi: 10.1123/jab.2017-0156.

[222] Katherine R. Saul, Xiao Hu, Craig M. Goehler, Meghan E. Vidt, Melissa Daly, Anca Velisar, and Wendy M. Murray. Benchmarking of dynamic simulation predictions in two software platforms using an upper limb musculoskeletal model. *Computer Methods in Biomechanics and Biomedical Engineering*, 18(13):1445–1458, 10 2015. doi: 10.1080/10255842.2014.916698.

[223] Yi Chung Lin, Tim W. Dorn, Anthony G. Schache, and Marcus G. Pandy. Comparison of different methods for estimating muscle forces in human movement. *Journal of Engineering in Medicine*, 226(2):103–112, 2012. ISSN 09544119. doi: 10.1177/0954411911429401.

[224] Thomas S Buchanan, David G Lloyd, Kurt Manal, and Thor F Besier. Neuromusculoskeletal modeling: estimation of muscle forces and joint moments and movements from measurements of neural command. *Journal of applied biomechanics*, 20(4): 367–95, 2004.

[225] Darryl G. Thelen and Frank C. Anderson. Using computed muscle control to generate forward dynamic simulations of human walking from experimental data. *Journal of Biomechanics*, 39(6):1107–1115, 2006. doi: 10.1016/j.jbiomech.2005.02.010.

[226] Jennifer L. Hicks, Thomas K. Uchida, Ajay Seth, Apoorva Rajagopal, and Scott L. Delp. Is My Model Good Enough? Best Practices for Verification and Validation of Musculoskeletal Models and Simulations of Movement. *Journal of Biomechanical Engineering*, 137(2):020905, 2015. doi: 10.1115/1.4029304.

[227] Tito Bassani, Elena Stucovitz, Zhihui Qian, Matteo Briguglio, and Fabio Galbusera. Validation of the AnyBody full body musculoskeletal model in computing lumbar spine loads at L4L5 level. *Journal of Biomechanics*, 58:89–96, 2017. ISSN 18732380. doi: 10.1016/j.jbiomech.2017.04.025.

[228] Claire C. Gordon, Cynthia L. Blackwell, Bruce Bradtmiller, Joseph L. Parham, Patricia Barrientos, Stephen P. Paquette, Brian D. Corner, Jeremy M. Carson, Joseph C. Venezia, Belva M. Rockwell, Michael Mucher, and Shirley Kristensen. 2012 Anthropometric Survey of U.S. Army Personnel: Methods and Summary Statistics. Technical report, U.S. Army Natick Soldier Research, Development and Engineering Center, Natick, MA, 2012.

[229] Darryl G. Thelen. Adjustment of muscle mechanics model parameters to simulate dynamic contractions in older adults. *Journal of Biomechanical Engineering*, 125 (1):70–77, 2 2003. ISSN 01480731. doi: 10.1115/1.1531112.

[230] R. F. Chandler, C. E. Clauser, J. T. McConville, H. M. Reynolds, and J. W. Young. Investigation of Inertial Properties of the Human Body. Technical report, Aerospace Medical Research Laboratory, 1975.

[231] Bjarne Laursen, Bente Rona Jensen, Gunnar Németh, and Gisela Sjøgaard. A model predicting individual shoulder muscle forces based on relationship between electromyographic and 3D external forces in static position. *Journal of Biomechanics*, 31(8):731–739, 8 1998. doi: 10.1016/S0021-9290(98)00091-8.

[232] Angela V. Dieterich, Alberto Botter, Taian Martins Vieira, Anneli Peolsson, Frank Petzke, Paul Davey, and Deborah Falla. Spatial variation and inconsistency between estimates of onset of muscle activation from EMG and ultrasound. *Scientific Reports*, 7(1):1–11, 2017. ISSN 20452322. doi: 10.1038/srep42011.

[233] Brian A. Garner and Marcus G. Pandy. Musculoskeletal Model of the Upper Limb Based on the Visible Human Male Dataset. *Computer Methods in Biomechanics and Biomedical Engineering*, 4(2):93–126, 2001. doi: 10.1080/10255840008908000.

[234] David W. Wagner, Vahagn Stepanyan, James M. Shippen, Matthew S. DeMers, Robin S. Gibbons, Brian J. Andrews, Graham H. Creasey, and Gary S. Beaupre. Consistency Among Musculoskeletal Models: Caveat Utilitor. *Annals of Biomedical Engineering 2013 41:8*, 41(8):1787–1799, 6 2013. ISSN 1573-9686. doi: 10.1007/S10439-013-0843-1.

[235] Katherine R. Saul, Meghan E. Vidt, Garry E. Gold, and Wendy M. Murray. Upper Limb Strength and Muscle Volume in Healthy Middle-Aged Adults. *Journal of Applied Biomechanics*, 31(6):484–491, 12 2015. doi: 10.1123/jab.2014-0177.

[236] Tobias Winkler, Falk Mersmann, Philipp von Roth, Ralf Dietrich, Stefanie Bierbaum, and Adamantios Arampatzis. Development of a Non-invasive Methodology for the Assessment of Muscle Fibre Composition. *Frontiers in Physiology*, 10:174, 2019. ISSN 1664-042X. doi: 10.3389/FPHYS.2019.00174.

[237] Ryan Cunningham, María B. Sánchez, Gregory May, and Ian Loram. Estimating Full Regional Skeletal Muscle Fibre Orientation from B-Mode Ultrasound Images Using Convolutional, Residual, and Deconvolutional Neural Networks. *Journal of Imaging*, 4(2):29, 1 2018. doi: 10.3390/JIMAGING4020029.

[238] Francisco M.M.O. Campos and Joao M.F. Calado. Approaches to human arm movement control-A review. *Annual Reviews in Control*, 33(1):69–77, 2009. ISSN 13675788. doi: 10.1016/j.arcontrol.2009.03.001.

[239] Nils Hasler, Bodo Rosenhahn, Thorsten Thormahlen, Michael Wand, Juergen Gall, and Hans-Peter Seidel. Markerless Motion Capture with unsynchronized moving cameras. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 224–231. IEEE, 2010. ISBN 978-1-4244-3992-8. doi: 10.1109/cvpr.2009.5206859.

[240] Stefano Corazza, Lars Mündermann, Ajit M. Chaudhari, T. Demattio, Claudio Cobelli, and Thomas P. Andriacchi. A markerless motion capture system to study musculoskeletal biomechanics: Visual hull and simulated annealing approach. *Annals of Biomedical Engineering*, 34(6):1019–1029, 2006. ISSN 00906964. doi: 10.1007/s10439-006-9122-8.

[241] Max Van Lith, Justin Fong, Vincent Crocher, Ying Tan, Iven Mareels, and Denny Oetomo. Calibration free upper limb joint motion estimation algorithm with wearable sensors. In *2016 14th International Conference on Control, Automation, Robotics and Vision, ICARCV 2016*, page Su26.5. IEEE, 1 2017. ISBN 978-1-5090-3549-6. doi: 10.1109/ICARCV.2016.7838634.

[242] Richardo Khonasty, Marc G Carmichael, Dikai Liu, and Kenneth J Waldron. Upper Body Pose Estimation Utilizing Kinematic Constraints from Physical Human-Robot Interaction. In *Proceedings of Australasian Conference on Robotics and Automation*, 2017.

[243] Dimitar Stanev, Konstantinos Filip, Dimitrios Bitzas, Sokratis Zouras, Georgios Giarmatzis, Dimitrios Tsaopoulos, and Konstantinos Moustakas. Real-Time Musculoskeletal Kinematics and Dynamics Analysis Using Marker- and IMU-Based Solutions in Rehabilitation. *Sensors*, 21(5):1804, 2021. ISSN 1424-8220. doi: 10.3390/s21051804.

[244] Aimee Cloutier, Robyn Boothby, and Jingzhou Yang. Motion Capture Experiments for Validating Optimization-Based Human Models. In *International Conference on Digital Human Modeling*, pages 59–68. Springer, 2011. doi: 10.1007/978-3-642-21799-9_7.

[245] Juan García-López and Pedro Abal del Blanco. Kinematic Analysis of Bicycle Pedalling Using 2D and 3D Motion Capture Systems. *ISBS Proceedings Archive*, 35(1), 2017.

[246] Francesca Gandolfo, Ferdinando A Mussa-Ivaldi, and Emilio Bizzi. Motor learning by field approximation. *Proceedings of the National Academy of Sciences of the United States of America*, 93(9):3843–6, 4 1996. doi: 10.1073/pnas.93.9.3843.

[247] Hermano I Krebs, Mindy L Aisen, Bruce T Volpe, and Neville Hogan. Quantization of continuous arm movements in humans with brain injury. *Proceedings of the National Academy of Sciences of the United States of America*, 96(8):4645–9, 4 1999. doi: 10.1073/pnas.96.8.4645.

[248] Tamar Flash and Neville Hogan. The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience*, 5(7):1688–1703, 1985. ISSN 02706474. doi: 10.1523/jneurosci.05-07-01688.1985.

[249] Elias B. Thorp, Konrad P. Kording, and Ferdinando A. Mussa-Ivaldi. Using noise to shape motor learning. *Journal of Neurophysiology*, 117(2):728–737, 2017. ISSN 0022-3077. doi: 10.1152/jn.00493.2016.

[250] David W. Franklin and Theodore E. Milner. Adaptive control of stiffness to stabilize hand position with large loads. *Experimental Brain Research*, 152(2):211–220, 9 2003. ISSN 00144819. doi: 10.1007/s00221-003-1540-3.

[251] Randy D. Trumbower, Matthew A. Krutky, Bing Shiang Yang, and Eric J. Perreault. Use of self-selected postures to regulate multi-joint stiffness during unconstrained tasks. *PLoS ONE*, 4(5), 5 2009. ISSN 19326203. doi: 10.1371/journal.pone.0005411.

[252] Roy D. Crowninshield and Richard A. Brand. A physiologically based criterion of muscle force prediction in locomotion. *Journal of Biomechanics*, 14(11):793–801, 1 1981. doi: 10.1016/0021-9290(81)90035-X.

[253] J J Woods and Brenda Bigland-Ritchie. Linear and non-linear surface EMG/force relationships in human muscles. An anatomical/functional argument for the existence of both. *American journal of physical medicine*, 62(6):287–99, 12 1983.

[254] Felix E. Zajac. Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control. *Critical reviews in biomedical engineering*, 17(4): 359–411, 1989.

[255] Emanuel Todorov. Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9):907–915, 2004. ISSN 10976256. doi: 10.1038/nn1309.

[256] Gerald E. Loeb. Optimal isn't good enough. *Biological Cybernetics*, 106(11-12): 757–765, 2012. ISSN 14320770. doi: 10.1007/s00422-012-0514-6.

[257] Bastien Berret, Enrico Chiovetto, Francesco Nori, and Thierry Pozzo. Evidence for composite cost functions in arm movement planning: An inverse optimal control approach. *PLoS Computational Biology*, 7(10):1–18, 2011. ISSN 1553734X. doi: 10.1371/journal.pcbi.1002183.

[258] Ryan Koeppen, Meghan E. Huber, Dagmar Sternad, and Neville Hogan. Controlling physical interactions: Humans do not minimize muscle effort. In *ASME 2017 Dynamic Systems and Control Conference, DSCC 2017*, volume 1. American Society of Mechanical Engineers, 2017. ISBN 9780791858271. doi: 10.1115/DSCC2017-5202.

[259] Paul S. Glazier and Sina Mehdizadeh. In search of sports biomechanics' holy grail: Can athlete-specific optimum sports techniques be identified? *Journal of Biomechanics*, 94:1–4, 2019. ISSN 18732380. doi: 10.1016/j.jbiomech.2019.07.044.

[260] M. Hongchul Sohn, Daniel M. Smith, and Lena H. Ting. Effects of kinematic complexity and number of muscles on musculoskeletal model robustness to muscle dysfunction. *PLOS ONE*, 14(7):e0219779, 2019. ISSN 1932-6203. doi: 10.1371/JOURNAL.PONE.0219779.

[261] Katrin Stollenmaier, Ilka S. Rist, Fabio Izzi, and Daniel F.B. Haeufle. Simulating the response of a neuro-musculoskeletal model to assistive forces: Implications for the design of wearables compensating for motor control deficits. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics*

and Biomechatronics, pages 779–784. IEEE, 2020. ISBN 9781728159072. doi: 10.1109/BioRob49111.2020.9224411.

[262] Marc Gary Carmichael and Dikai Liu. Estimating Physical Assistance Need Using a Musculoskeletal Model. *IEEE Transactions on Biomedical Engineering*, 60(7): 1912–1919, 2013. doi: 10.1109/TBME.2013.2244889.

[263] Marc G. Carmichael and Dikai Liu. Upper limb strength estimation of physically impaired persons using a musculoskeletal model: A sensitivity analysis. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 2438–2441. IEEE, 8 2015. ISBN 978-1-4244-9271-8. doi: 10.1109/EMBC.2015.7318886.

[264] Moses C. Nah, Aleksei Krotov, Marta Russo, Dagmar Sternad, and Neville Hogan. Dynamic Primitives Facilitate Manipulating a Whip. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, pages 685–691. IEEE, 2020. ISBN 9781728159072. doi: 10.1109/BioRob49111.2020.9224399.

[265] Darryl G. Thelen, Frank C. Anderson, and Scott L. Delp. Generating dynamic simulations of movement using computed muscle control. *Journal of Biomechanics*, 36(3):321–328, 2003. doi: 10.1016/S0021-9290(02)00432-3.

[266] Chand T. John, Frank C. Anderson, Jill S. Higginson, and Scott L. Delp. Stabilisation of walking by intrinsic muscle properties revealed in a three-dimensional muscle-driven simulation. *Computer Methods in Biomechanics and Biomedical Engineering*, 16(4):451–462, 4 2013. doi: 10.1080/10255842.2011.627560.

[267] Ahmet Erdemir, Scott McLean, Walter Herzog, and Antonie J. van den Bogert. Model-based estimation of muscle forces exerted during movements. *Clinical Biomechanics*, 22(2):131–154, 2 2007. doi: 10.1016/J.CLINBIOMECH.2006.09.005.

[268] Andrea Zonnino and Fabrizio Sergi. Model-Based Estimation of Individual Muscle Force Given an Incomplete Set of Muscle Activity Measurements. In Lorenzo Masia, Silvestro Micera, Metin Akay, and José L. Pons, editors, *Proceedings of the 4th International Conference on NeuroRehabilitation (ICNR2018)*, pages 800–804. Springer, 2018. doi: 10.1007/978-3-030-01845-0_160.

[269] Andrea D'Avella, Philippe Saltiel, and Emilio Bizzi. Combinations of muscle synergies in the construction of a natural motor behavior. *Nature Neuroscience*, 6(3): 300–308, 3 2003. ISSN 10976256. doi: 10.1038/nn1010.

[270] Matthew C. Tresch and Anthony Jarc. The case for and against muscle synergies. *Current Opinion in Neurobiology*, 19(6):601–607, 12 2009. ISSN 0959-4388. doi: 10.1016/J.CONB.2009.09.002.

[271] Jiahao Chen and Hong Qiao. Muscle-Synergies-Based Neuromuscular Control for Motion Learning and Generalization of a Musculoskeletal System. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, pages 1–14, 1 2020. ISSN 2168-2216. doi: 10.1109/tsmc.2020.2966818.

[272] Darryl K Charles and Suzanne McDonough. A Participatory Design Framework for the Gamification of Rehabilitation Systems. In *Recent advances on using virtual reality technologies for rehabilitation*, pages 2–9. Nova Science Publishers, 2015. ISBN 9781634840286.

[273] Rachel Proffitt. Gamification in Rehabilitation: Finding the "Just-Right-Challenge". In Daniel Novák, Bengisu Tulu, and Håvar Brendryen, editors, *Handbook of Research on Holistic Perspectives in Gamification for Clinical Practice*, chapter 7, pages 132–157. Idea Group,U.S., Harrisburg, PA, 2016. ISBN 9781466695221. doi: 10.4018/978-1-4666-9522-1.ch007.

[274] Michael Yates, Arpad Kelemen, and Cecilia Sik Lanyi. Virtual reality gaming in the rehabilitation of the upper extremities post-stroke. *Brain Injury*, 30(7):855–863, 6 2016. doi: 10.3109/02699052.2016.1144146.

[275] Sarah Sharples, Sue Cobb, Amanda Moody, and John R. Wilson. Virtual reality induced symptoms and effects (VRISE): Comparison of head mounted display (HMD), desktop and projection display systems. *Displays*, 29(2):58–69, 3 2008. doi: 10.1016/J.DISPLA.2007.09.005.

[276] Marc G. Carmichael and Dikai Liu. Admittance control scheme for implementing model-based assistance-as-needed on a robot. In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 870–873. IEEE, 2013. ISBN 978-1-4577-0216-7. doi: 10.1109/EMBC.2013.6609639.

[277] Jacob Brackenridge, Lynley V. Bradnam, Sheila Lennon, John J. Costi, and David A. Hobbs. A Review of Rehabilitation Devices to Promote Upper Limb Function Following Stroke. *Neuroscience and Biomedical Engineering*, 4(1):25–42, 5 2016. doi: 10.2174/2213385204666160303220102.

[278] Volker Tresp. A Bayesian Committee Machine. *Neural Computation*, 12(11):2719–2741, 11 2000. doi: 10.1162/089976600300014908.

[279] Erik J. Scheme and Kevin B. Englehart. Validation of a selective ensemble-based classification scheme for myoelectric control using a three-dimensional fitts' law test. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 21(4):616–623, 2013. ISSN 1534-4320. doi: 10.1109/TNSRE.2012.2226189.

[280] Ziheng Wang and Ann Majewicz Fey. Human-centric predictive model of task difficulty for human-in-the-loop control tasks. *PLOS ONE*, 13(4):e0195053, 4 2018. ISSN 1932-6203. doi: 10.1371/JOURNAL.PONE.0195053.

[281] J. M. Hollerbach and C. G. Atkeson. Deducing planning variables from experimental arm trajectories: Pitfalls and possibilities. *Biological Cybernetics 1987 56:5*, 56(5): 279–292, 7 1987. ISSN 1432-0770. doi: 10.1007/BF00319509.

[282] Lawrence Sambrooks and Brett Wilkinson. Comparison of gestural, touch, and mouse interaction with Fitts' law. *Proceedings of the 25th Australian Computer-Human Interaction Conference: Augmentation, Application, Innovation, Collaboration, OzCHI 2013*, pages 119–122, 2013. doi: 10.1145/2541016.2541066.

[283] David J. Reinkensmeyer, Eric T. Wolbrecht, Vicky Chan, Cathy Chou, Steven C. Cramer, and James E. Bobrow. Comparison of Three-Dimensional, Assist-as-Needed Robotic Arm/Hand Movement Training Provided with Pneu-WREX to Conventional Tabletop Therapy After Chronic Stroke. *American Journal of Physical Medicine & Rehabilitation*, 91(11 Suppl 3):S232–S241, 11 2012. doi: 10.1097/PHM.0b013e31826bce79.

[284] Shraddha Srivastava, Pei Chun Kao, Darcy S Reisman, John P Scholz, Sunil K Agrawal, and Jill S Higginson. Robotic Assist-As-Needed as an Alternative to Therapist-Assisted Gait Rehabilitation. *International journal of physical medicine & rehabilitation*, 4(5):370, 2016. doi: 10.4172/2329-9096.1000370.

[285] Marc G. Carmichael, Stefano Aldini, Richardo Khonasty, Antony Tran, Christian Reeks, Dikai Liu, Kenneth J. Waldron, and Gamini Dissanayake. The ANBOT: An Intelligent Robotic Co-worker for Industrial Abrasive Blasting. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8026–8033. IEEE, 2019. ISBN 978-1-7281-4004-9. doi: 10.1109/iros40897.2019.8967993.

[286] Mohammad Shokoohi-Yekta, Bing Hu, Hongxia Jin, Jun Wang, and Eamonn Keogh. Generalizing DTW to the multi-dimensional case requires an adaptive approach. *Data Mining and Knowledge Discovery*, 31(1):1–31, 1 2017. doi: 10.1007/s10618-016-0455-0.

[287] Qingde Li and John G. Griffiths. Least squares ellipsoid specific fitting. In *Proceedings - Geometric Modeling and Processing 2004*, pages 335–340, 2004. ISBN 0769520782. doi: 10.1109/gmap.2004.1290055.

[288] Oriane Dermy, Alexandros Paraschos, Marco Ewerton, Jan Peters, François Charpillet, and Serena Ivaldi. Prediction of intention during interaction with iCub with probabilistic movement primitives. *Frontiers Robotics AI*, 4:45, 10 2017. ISSN 22969144. doi: 10.3389/frobt.2017.00045.

[289] Peter Pastor, Mrinal Kalakrishnan, Ludovic Righetti, and Stefan Schaal. Towards associative skill memories. In *IEEE-RAS International Conference on Humanoid Robots*, pages 309–315. IEEE, 2012. ISBN 9781467313698. doi: 10.1109/HUMANOIDS.2012.6651537.

[290] Wendy M Murray, Thomas S Buchanan, and Scott L Delp. Scaling of peak moment arms of elbow muscles with upper extremity bone dimensions. *Journal of biomechanics*, 35(1):19–26, 2002. doi: https://doi.org/10.1016/S0021-9290(01)00173-7.

[291] Katherine R. S. Holzbaur, Wendy M. Murray, and Scott L. Delp. A Model of the Upper Extremity for Simulating Musculoskeletal Surgery and Analyzing Neuromuscular Control. *Annals of Biomedical Engineering*, 33(6):829–840, 2005. doi: 10.1007/s10439-005-3320-7.

[292] Katherine R.S. Holzbaur, Wendy M. Murray, Garry E. Gold, and Scott L. Delp. Upper limb muscle volumes in adult subjects. *Journal of Biomechanics*, 40(4):742–749, 2007. doi: 10.1016/J.JBIOMECH.2006.11.011.

[293] Matthew Millard, Thomas Uchida, Ajay Seth, and Scott L. Delp. Flexing computational muscle: Modeling and simulation of musculotendon dynamics. *Journal of Biomechanical Engineering*, 135(2), 2013. ISSN 01480731. doi: 10.1115/1.4023390.

[294] Archibald V. Hill. The heat of shortening and the dynamic constants of muscle. *Proceedings of the Royal Society of London. Series B - Biological Sciences*, 126(843): 136–195, 1938. doi: 10.1098/rspb.1938.0050.

[295] Peter Konrad. The abc of emg. Technical report, Noraxon U.S.A., Inc., Arizona, USA, 2005.

[296] Paula Silva, Sergio Fonseca, Juliana Ocarino, Gabriela Gonçalves, and Marisa Mancini. Contributions of cocontraction and eccentric activity to stiffness regulation. *Journal of Motor Behavior*, 41(3):207–218, 2009. ISSN 00222895. doi: 10.3200/JMBR.41.3.207-218.

[297] Paul L. Gribble, Lucy I. Mullin, Nicholas Cothros, and Andrew Mattar. Role of cocontraction in arm movement accuracy. *Journal of Neurophysiology*, 89(5):2396–2405, 5 2003. ISSN 00223077. doi: 10.1152/jn.01020.2002.

[298] Aaron L. Wong and Adrian M. Haith. Motor planning flexibly optimizes performance under uncertainty about task goals. *Nature Communications*, 8(1):1–10, 2017. ISSN 20411723. doi: 10.1038/ncomms14624.

[299] Timothy Verstynen and Philip N. Sabes. How each movement changes the next: An experimental and theoretical study of fast adaptive priors in reaching. *Journal of Neuroscience*, 31(27):10050–10059, 2011. ISSN 02706474. doi: 10.1523/JNEUROSCI.6525-10.2011.

[300] Boris I. Prilutsky and Vladimir M. Zatsiorsky. Optimization-based models of muscle coordination. *Exercise and Sport Sciences Reviews*, 30(1):32–38, 2002. ISSN 00916331. doi: 10.1097/00003677-200201000-00007.

[301] Justus D. Ortega and Claire T. Farley. Individual limb work does not explain the greater metabolic cost of walking in elderly adults. *Journal of Applied Physiology*, 102(6):2266–2273, 2007. ISSN 8750-7587. doi: 10.1152/japplphysiol.00583.2006.

[302] D. A. Nowak, Stefan Glasauer, and Joachim Hermsdörfer. How predictive is grip force control in the complete absence of somatosensory feedback? *Brain*, 127(1):182–192, 2004. ISSN 1460-2156. doi: 10.1093/brain/awh016.

[303] James B. Heald, David W. Franklin, and Daniel M. Wolpert. Increasing muscle co-contraction speeds up internal model acquisition during dynamic motor learning. *Scientific Reports*, 8(1), 12 2018. ISSN 20452322. doi: 10.1038/s41598-018-34737-5.

[304] Kurt A. Thoroughman and Reza Shadmehr. Electromyographic correlates of learning an internal model of reaching movements. *Journal of Neuroscience*, 19(19):8573–8588, 1999. ISSN 02706474. doi: 10.1523/jneurosci.19-19-08573.1999.

[305] Liliam Fernandes de Oliveira and Luciano Luporini Menegaldo. Individual-specific muscle maximum force estimation using ultrasound for ankle joint torque prediction

using an EMG-driven Hill-type model. *Journal of Biomechanics*, 43(14):2816–2821, 2010. ISSN 00219290. doi: 10.1016/j.jbiomech.2010.05.035.

[306] Fabien Dal Maso, Mickaël Begon, and Maxime Raison. Methodology to customize maximal isometric forces for hill-type muscle models. *Journal of Applied Biomechanics*, 33(1):80–86, 2017. ISSN 15432688. doi: 10.1123/jab.2016-0062.

[307] Laura A. Hallock, Akash Velu, Amanda Schwartz, and Ruzena Bajcsy. Muscle deformation correlates with output force during isometric contraction. In *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, pages 1188–1195. IEEE, 2020. ISBN 9781728159072. doi: 10.1109/BioRob49111.2020.9224391.

[308] Thomas J. Roberts and Annette M. Gabaldón. Interpreting muscle function from EMG: Lessons learned from direct measurements of muscle force. *Integrative and Comparative Biology*, 48(2):312–320, 2008. ISSN 15407063. doi: 10.1093/icb/icn056.

[309] George P. Pappas, Deanna S. Asakawa, Scott L. Delp, Felix E. Zajac, and John E. Drace. Nonuniform shortening in the biceps brachii during elbow flexion. *Journal of Applied Physiology*, 92(6):2381–2389, 2002. ISSN 87507587. doi: 10.1152/japplphysiol.00843.2001.

[310] Thorsten Rudroff, Didier Staudenmann, and Roger M. Enoka. Electromyographic measures of muscle activation and changes in muscle architecture of human elbow flexors during fatiguing contractions. *Journal of Applied Physiology*, 104(6):1720–1726, 6 2008. ISSN 8750-7587. doi: 10.1152/japplphysiol.01058.2007.

[311] Robert D. Herbert and Simon C. Gandevia. Changes in pennation with joint angle and muscle torque: in vivo measurements in human brachialis muscle. *The Journal of Physiology*, 484(2):523–532, 4 1995. ISSN 14697793. doi: 10.1113/jphysiol.1995.sp020683.

[312] Simon Bouisset. EMG and Muscle Force in Normal Motor Activities. In *New Concepts of the Motor Unit, Neuromuscular Disorders, Electromyographic Kinesiology*, volume 1, pages 547–583. S. Karger AG, 1973. doi: 10.1159/000394059.

[313] Virgil Mathiowetz, Karen Weber, Gloria Volland, and Nancy Kashman. Reliability and validity of grip and pinch strength evaluations. *Journal of Hand Surgery*, 9(2): 222–226, 3 1984. ISSN 03635023. doi: 10.1016/S0363-5023(84)80146-X.

[314] Raghu Gurram, Subhash Rakheja, and Gerard J. Gouw. A study of hand grip pressure distribution and EMG of finger flexor muscles under dynamic loads. *Ergonomics*, 38(4):684–699, 1995. ISSN 13665847. doi: 10.1080/00140139508925140.

[315] Virgil Mathiowetz, Nancy Kashman, Gloria Volland, Karen Weber, Mary Dowe, and Sandra Rogers. Grip and pinch strength: normative data for adults. *Archives of physical medicine and rehabilitation*, 66(2):69–74, 1985.

[316] Suzy L. Wong. Grip strength reference values for Canadians aged 6 to 79: Canadian Health Measures Survey, 2007 to 2013. *Health Reports*, 27(10):3–10, 10 2016.

[317] Roland S. Johansson, Ronald Riso, Charlotte Häger, and Lars Bäckström. Somatosensory control of precision grip during unpredictable pulling loads - I. Changes in load force amplitude. *Experimental Brain Research*, 89(1):181–191, 4 1992. ISSN 00144819. doi: 10.1007/BF00229015.

[318] Satyajit Ambike, Florent Paclet, Vladimir M. Zatsiorsky, and Mark L. Latash. Factors affecting grip force: Anatomy, mechanics, and referent configurations. *Experimental Brain Research*, 232(4):1219–1231, 2014. ISSN 14321106. doi: 10.1007/s00221-014-3838-8.

[319] Antony Tran, Dikai Liu, Ravindra Ranasinghe, Marc Carmichael, and Chuanbo Liu. Analysis of Human Grip Strength in Physical Human Robot Interaction. *Procedia Manufacturing*, 3:1442–1449, 2015. ISSN 23519789. doi: 10.1016/j.promfg.2015.07.320.

[320] National Center for Health Statistics. NHANES Anthropometry Procedures Manual, 2020. URL `https://wwwn.cdc.gov/nchs/data/nhanes/2019-2020/manuals/2020-Anthropometry-Procedures-Manual-508.pdf`.

[321] NIHR Southampton Biomedical Research Centre. Procedure for Measuring ADULT ULNA LENGTH, 2014. URL `https://www.uhs.nhs.uk/Media/Southampton-Clinical-Research/Procedures/BRCProcedures/Procedure-for-adult-ulna-length.pdf`.

[322] Samuel Rota, Isabelle Rogowski, Stéphane Champely, and Christophe Hautier. Reliability of EMG normalisation methods for upper-limb muscles. *Journal of Sports Sciences*, 31(15):1696–1704, 2013. ISSN 02640414. doi: 10.1080/02640414.2013.796063.

[323] Lisa M. Schutte, Mary M. Rodgers, Felix E. Zajac, and Roger M. Glaser. Improving the Efficacy of Electrical Stimulation-Induced Leg Cycle Ergometry: An Analysis

Based on a Dynamic Musculoskeletal Model. *IEEE Transactions on Rehabilitation Engineering*, 1(2):109–125, 6 1993. ISSN 1063-6528. doi: 10.1109/86.242425.

[324] Marc G. Carmichael. *A musculoskeletal model-based Assistance-As-Needed paradigm for assistive robotics*. PhD thesis, University of Technology Sydney (UTS), 2013.

[325] Mario Cifrek, Vladimir Medved, Stanko Tonković, and Saša Ostojić. Surface EMG based muscle fatigue evaluation in biomechanics. *Clinical Biomechanics*, 24(4):327–340, 2009. ISSN 02680033. doi: 10.1016/j.clinbiomech.2009.01.010.

[326] Anne Schwarz, Christoph M. Kanzler, Olivier Lambercy, Andreas R. Luft, and Janne M. Veerbeek. Systematic Review on Kinematic Assessments of Upper Limb Movements After Stroke. *Stroke*, 50(3):718–727, 3 2019. doi: 10.1161/STROKEAHA.118.023531.

[327] Christopher J. Gatti, Clark R. Dickerson, Edward K. Chadwick, Amy G. Mell, and Richard E. Hughes. Comparison of model-predicted and measured moment arms for the rotator cuff muscles. *Clinical Biomechanics*, 22(6):639–644, 7 2007. ISSN 0268-0033. doi: 10.1016/J.CLINBIOMECH.2007.02.001.

[328] Xiao Hu, Wendy M. Murray, and Eric J. Perreault. Muscle short-range stiffness can be used to estimate the endpoint stiffness of the human arm. *Journal of Neurophysiology*, 105(4):1633–1641, 4 2011. doi: 10.1152/JN.00537.2010.

[329] Edith M. Arnold, Samuel R. Hamner, Ajay Seth, Matthew Millard, and Scott L. Delp. How muscle fiber lengths and velocities affect muscle force generation as humans walk and run at different speeds. *Journal of Experimental Biology*, 216(11):2150–2160, 6 2013. ISSN 0022-0949. doi: 10.1242/JEB.075697.

[330] Laura Frey Law, Chandramouli Krishnan, and Keith Avin. Modeling nonlinear errors in surface electromyography due to baseline noise: A new methodology. *Journal of Biomechanics*, 44(1):202–205, 1 2011. ISSN 0021-9290. doi: 10.1016/J.JBIOMECH.2010.09.008.

[331] David C. Frankenfield. On heat, respiration, and calorimetry. *Nutrition*, 26(10):939–950, 2010. ISSN 0899-9007. doi: 10.1016/J.NUT.2010.01.002.

[332] Neil D. Lawrence and Joaquin Quiñonero-Candela. Local distance preservation in the GP-LVM through back constraints. In *Proceedings of the 23rd international conference on Machine learning*, pages 513–520. Association for Computing Machinery, 2006. doi: 10.1145/1143844.1143909.