Towards the Development of Safe, Collaborative Robotic Freehand Ultrasound

by

Nicholas A. NADEAU

THESIS PRESENTED TO ÉCOLE DE TECHNOLOGIE SUPÉRIEURE
IN PARTIAL FULFILLMENT FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
Ph.D.

MONTREAL, DECEMBER 18, 2019

ÉCOLE DE TECHNOLOGIE SUPÉRIEURE
UNIVERSITÉ DU QUÉBEC

BOARD OF EXAMINERS

THIS THESIS HAS BEEN EVALUATED

BY THE FOLLOWING BOARD OF EXAMINERS

Mr. Ilian Bonev, thesis supervisor
Department of Systems Engineering, École de technologie supérieure

Ms. Catherine Laporte, president of the board of examiners
Department of Electrical Engineering, École de technologie supérieure

Mr. Vincent Duchaine, member of the jury
Department of Systems Engineering, École de technologie supérieure

Mr. Sofiane Achiche, external independent examiner
Department of Mechanical Engineering, Polytechnique Montréal

THIS THESIS WAS PRESENTED AND DEFENDED

IN THE PRESENCE OF A BOARD OF EXAMINERS AND THE PUBLIC

ON DECEMBER 16, 2019

AT ÉCOLE DE TECHNOLOGIE SUPÉRIEURE

FOREWORD

This Ph.D. dissertation presents research carried out between 2015 and 2019 at École de technologie supérieure, under the supervision of Prof. Ilian Bonev and Prof. Pascal Bigras. This work resulted in several articles, conference presentations, and an open-source robotics framework for the Python programming language.

Originally, the goal of this work was to develop a prototype medical ultrasound application using a collaborative robot. However, following the proposal defence, it was observed that this undertaking was too broad and did not offer a unique scientific contribution. After a period of exploratory experiments, it was revealed that there was a lack of knowledge in the literature of robot fundamentals and that much of the work was focused on the design of the robotic ultrasound system for patient safety or image generation. Specifically, three critical elements were missing: contact distinction of probe-to-skin contact, calibration of a force-controlled robot for accuracy, and force controller tuning for motion performance. While the feasibility of robotic ultrasound has been the focal point of the literature, there is a lack of focus of what occurs between system design and image output. Therefore, this dissertation focuses on filling in these gaps in the literature with the intention of providing frameworks and fundamental knowledge for future ultrasound robotics research.

This dissertation focuses on the content of three distinct works, presented in Chapters 2 to 4. The Introduction chapter presents background information on ultrasound imaging, collaborative robotics, and robotic ultrasound, as well as the main problem statement, motivations, and objectives of this research. A review of relevant literature is found in Chapter 1. Finally, the Conclusion and Recommendations chapter summarizes the contributions and highlights some recommendations for future research.

ACKNOWLEDGEMENTS

First, I would like to thank my Ph.D. supervisor, Prof. Ilian Bonev, for his support throughout my graduate studies. It is thanks to his understanding and flexibility that I was able to accomplish both this PhD and develop a professional career in parallel.

Second, although no longer with us, I wish to acknowledge the support of Prof. Pascal Bigras[†], who taught my first courses in robotics control.

Further, I would like to thank the jury members, Prof. Catherine Laporte, Prof. Sofiane Achiche, and Prof. Vincent Duchaine, for accepting to review my thesis and providing meaningful feedback.

Next, I would like to thank my colleagues and peers, especially Jennifer Kwiatkowski, Ahmed Joubair, and Albert Nubiola for their contributions and support.

Finally, nobody has been more important to me in the pursuit of my goals than my family and friends, whose love and guidance are with me in whatever I pursue. They are the ultimate role models. I am especially grateful for my wife, Charlie Ohayon, for her love, support, and patience throughout this journey. You are my partner, peer, and mentor. Thank you for teaching me that my job in life was to learn, to be happy, and to know my limits; only then could I reach my true potential.

This dissertation is dedicated to all the teachers and mentors in my life. They have inspired hope, ignited my imagination, and instilled a lifelong love of learning.

[†]Died October 11, 2019

Contributions au développement de l'échographie robotique sécuritaire et collaborative

Nicholas A. NADEAU

RÉSUMÉ

L'utilisation de la robotique en médecine revêt une importance croissante pour les services de santé modernes, car les systèmes robotiques ont la capacité d'améliorer les tâches humaines, améliorant ainsi la capacité de traitement d'un prestataire de soins de santé. Dans le secteur médical, l'imagerie par ultrasons est une approche peu coûteuse sans les fortes émissions de rayonnements souvent associées à d'autres modalités, surtout lorsqu'on la compare à l'IRM et à l'imagerie par tomodensitométrie respectivement. Au cours des deux dernières décennies, des efforts considérables ont été investis dans la recherche et le développement en robotique ultrasonore à main levée.

Cependant, cette recherche s'est concentrée sur la faisabilité de l'application, et non sur les principes fondamentaux de la robotique, tels que le contrôle du mouvement, l'étalonnage et la conscience contextuelle. Jusqu'à présent, une grande partie du travail s'est concentrée sur les robots conçus sur mesure, la génération d'images ultrasonores et l'asservissement visuel, ou la commande à distance. La recherche fondée sur ces sujets souffre souvent d'importantes limitations qui entravent son utilité d'une manière adaptable, évolutive et réaliste. En particulier, alors que les robots personnalisés peuvent être conçus pour une application spécifique, les robots collaboratifs commerciaux et prêts à l'emploi constituent une solution plus robuste et économique. Par ailleurs, diverses études d'échographie robotique ont démontré la faisabilité de l'utilisation d'un contrôle de force de base, mais elles explorent rarement le réglage du contrôleur dans le contexte de la sécurité du patient et de la peau déformable dans un environnement non structuré. De plus, de nombreuses études évaluent de nouvelles approches d'asservissement visuel, mais ne tiennent pas compte de l'utilité pratique de recourir à des dispositifs de mesure externes pour contrôler le mouvement. Ces études négligent l'importance de la précision et de l'étalonnage des robots, qui permettent à un système de naviguer en toute sécurité dans son environnement tout en réduisant les erreurs d'imagerie associées au positionnement. Par conséquent, bien que la faisabilité de l'échographie robotique ait été le point central, il y a un manque de concentration sur ce qui se passe entre la conception du système et la sortie des images.

Cette thèse aborde les limites de la littérature actuelle à travers trois contributions distinctes. Étant donné la nature contrôlée par la force d'un robot à ultrasons, la première contribution présente une approche d'étalonnage en boucle fermée utilisant un contrôle d'impédance et un équipement peu coûteux. La précision est une exigence fondamentale pour la génération et le ciblage d'images ultrasonores de haute qualité. Ceci est particulièrement vrai lorsqu'un chemin spécifique est suivi le long d'un patient ou des coupes 2D sont synthétisées en une image ultrasonore 3D. Cependant, même si la plupart des robots industriels sont intrinsèquement précis, ils ne le sont pas nécessairement. Bien que l'étalonnage des robots ait fait l'objet d'études approfondies, de nombreuses approches reposent sur des équipements coûteux et très délicats.

Des essais expérimentaux ont montré que cette méthode est comparable en qualité à l'étalonnage traditionnel à l'aide d'un laser tracker. Comme démontré par une étude expérimentale et validée par un laser tracker, la précision absolue d'un robot collaboratif a été améliorée à une erreur maximale de 0.990 mm, soit une amélioration de 58.4% par rapport au modèle nominal.

La deuxième contribution explore les collisions et les événements de contact, car ils sont un sous-produit naturel des applications impliquant une interaction physique homme-robot dans des environnements non structurés. L'échographie médicale assistée par robot est un exemple de tâche où le simple fait d'arrêter le robot lors de la détection de contact peut ne pas être une stratégie de réaction appropriée. Ainsi, le robot doit être conscient de l'emplacement du contact corporel afin de planifier correctement les trajectoires contrôlées par la force le long du corps humain à l'aide de la sonde d'imagerie. C'est particulièrement vrai pour les systèmes d'échographie à distance où la sécurité et la maniabilité sont des éléments importants à prendre en compte lors de l'utilisation d'un système médical à distance via un réseau de communication. Un cadre est proposé pour la classification des contacts des robots à l'aide des données des capteurs intégrés d'un robot collaboratif. Contrairement aux études antérieures, cette classification ne fait pas de distinction entre les scénarios de contact intentionnel et non intentionnel, mais classifie plutôt le *quoi* impliqué dans l'événement de contact. Le classificateur peut discerner différentes zones corporelles spécifiques ISO/TS 15066:2016 le long d'une jambe de modèle humain avec une précision 89.37 %. Dans l'ensemble, ce cadre de distinction des contacts permet d'élaborer des stratégies de réaction plus complexes et d'adapter le comportement du robot au cours de l'interaction physique homme-robot.

Enfin, étant donné que le succès d'une tâche d'échographie dépend de la capacité du système robotique à manipuler l'interaction physique homme-robot, le contrôle de mouvement pur est insuffisant. Les techniques de contrôle de la force sont nécessaires pour obtenir un comportement efficace et adaptable d'un système robotique dans l'environnement ultrasonore non structuré tout en assurant une interaction physique sûre. Bien que le contrôle de la force ne nécessite pas une connaissance explicite de l'environnement, pour obtenir un comportement dynamique acceptable, les paramètres de contrôle doivent être réglés. La troisième contribution propose un cadre d'accord en ligne simple et efficace pour le contrôle du mouvement des ultrasons robotisés à main levée. Dans le contexte de l'échographie médicale, les différents endroits du corps humain ont une rigidité différente et nécessitent des réglages uniques. Grâce à des expériences réelles avec un robot collaboratif, le cadre a réglé le contrôle de mouvement pour des trajectoires optimales et sécuritaires le long d'un fantôme de la jambe humaine. Le processus d'optimisation a réussi à réduire l'erreur absolue moyenne de la force de contact du mouvement à 0.537 N grâce à l'évolution de huit paramètres de contrôle du mouvement. De plus, la connaissance du contexte par la classification du mouvement peut offrir un cadre pour l'optimisation de l'interaction physique homme-robot et la sécurité grâce à un comportement de mouvement prédictif avec un objectif futur de l'interaction physique homme-robot autonome. Ainsi, un pipeline de classification, formé à l'aide des données de mouvement du processus de réglage, a été en mesure de classer de façon fiable la qualité future du suivi des forces d'une session de mouvement avec une précision de 91.82 %.

Mots-clés: Systèmes adaptatifs, Étalonnage, Calcul évolutif, Contrôle des forces, Interaction homme-robot, Robotique médicale, Contrôle de mouvement, Optimisation, Contrôle robotique, Cinématique robotique, Robotique, Optimisation de trajectoire, Imagerie ultrasonique

Towards the Development of Safe, Collaborative Robotic Freehand Ultrasound

Nicholas A. NADEAU

ABSTRACT

The use of robotics in medicine is of growing importance for modern health services, as robotic systems have the capacity to improve upon human tasks, thereby enhancing the treatment ability of a healthcare provider. In the medical sector, ultrasound imaging is an inexpensive approach without the high radiation emissions often associated with other modalities, especially when compared to MRI and CT imaging respectively. Over the past two decades, considerable effort has been invested into freehand ultrasound robotics research and development.

However, this research has focused on the feasibility of the application, not the robotic fundamentals, such as motion control, calibration, and contextual awareness. Instead, much of the work is concentrated on custom designed robots, ultrasound image generation and visual servoing, or teleoperation. Research based on these topics often suffer from important limitations that impede their use in an adaptable, scalable, and real-world manner. Particularly, while custom robots may be designed for a specific application, commercial collaborative robots are a more robust and economical solution. Otherwise, various robotic ultrasound studies have shown the feasibility of using basic force control, but rarely explore controller tuning in the context of patient safety and deformable skin in an unstructured environment. Moreover, many studies evaluate novel visual servoing approaches, but do not consider the practicality of relying on external measurement devices for motion control. These studies neglect the importance of robot accuracy and calibration, which allow a system to safely navigate its environment while reducing the imaging errors associated with positioning. Hence, while the feasibility of robotic ultrasound has been the focal point in previous studies, there is a lack of attention to what occurs between system design and image output.

This thesis addresses limitations of the current literature through three distinct contributions. Given the force-controlled nature of an ultrasound robot, the first contribution presents a closed-loop calibration approach using impedance control and low-cost equipment. Accuracy is a fundamental requirement for high-quality ultrasound image generation and targeting. This is especially true when following a specified path along a patient or synthesizing 2D slices into a 3D ultrasound image. However, even though most industrial robots are inherently precise, they are not necessarily accurate. While robot calibration itself has been extensively studied, many of the approaches rely on expensive and highly delicate equipment. Experimental testing showed that this method is comparable in quality to traditional calibration using a laser tracker. As demonstrated through an experimental study and validated with a laser tracker, the absolute accuracy of a collaborative robot was improved to a maximum error of 0.990 mm, representing a 58.4% improvement when compared to the nominal model.

The second contribution explores collisions and contact events, as they are a natural by-product of applications involving physical human-robot interaction (pHRI) in unstructured environments. Robot-assisted medical ultrasound is an example of a task where simply stopping the robot upon

contact detection may not be an appropriate reaction strategy. Thus, the robot should have an awareness of body contact location to properly plan force-controlled trajectories along the human body using the imaging probe. This is especially true for remote ultrasound systems where safety and manipulability are important elements to consider when operating a remote medical system through a communication network. A framework is proposed for robot contact classification using the built-in sensor data of a collaborative robot. Unlike previous studies, this classification does not discern between intended vs. unintended contact scenarios, but rather classifies *what* was involved in the contact event. The classifier can discern different ISO/TS 15066:2016 specific body areas along a human-model leg with 89.37 % accuracy. Altogether, this contact distinction framework allows for more complex reaction strategies and tailored robot behaviour during pHRI.

Lastly, given that the success of an ultrasound task depends on the capability of the robot system to handle pHRI, pure motion control is insufficient. Force control techniques are necessary to achieve effective and adaptable behaviour of a robotic system in the unstructured ultrasound environment while also ensuring safe pHRI. While force control does not require explicit knowledge of the environment, to achieve an acceptable dynamic behaviour, the control parameters must be tuned. The third contribution proposes a simple and effective online tuning framework for force-based robotic freehand ultrasound motion control. Within the context of medical ultrasound, different human body locations have a different stiffness and will require unique tunings. Through real-world experiments with a collaborative robot, the framework tuned motion control for optimal and safe trajectories along a human leg phantom. The optimization process was able to successfully reduce the mean absolute error (MAE) of the motion contact force to 0.537 N through the evolution of eight motion control parameters. Furthermore, contextual awareness through motion classification can offer a framework for pHRI optimization and safety through predictive motion behaviour with a future goal of autonomous pHRI. As such, a classification pipeline, trained using the tuning process motion data, was able to reliably classify the future force tracking quality of a motion session with an accuracy of 91.82 %.

Keywords: Adaptive Systems, Calibration, Evolutionary Computation, Force Control, Human-robot Interaction, Medical Robotics, Motion Control, Optimization, Robot Control, Robot Kinematics, Robotics, Trajectory Optimization, Ultrasonic Imaging

TABLE OF CONTENTS

Page

LIST OF TABLES

Page

LIST OF FIGURES

Page

LIST OF ALGORITHMS

Page

LIST OF ABREVIATIONS

CMM	Coordinate Measuring Machine
DE	Differential Evolution
DOF	Degrees of Freedom
FFT	Fast Fourier Transform
FWE	Family-wise Error
IROS	IEEE/RSJ International Conference on Intelligent Robots and Systems
LBR	KUKA LBR iiwa 7 R800
MAE	Mean Absolute Error
MDH	Modified Denavit-Hartenberg Parameters
PCA	Principal Component Analysis
pHRI	Physical Human-Robot Interaction
RPC	Remote Procedure Call
SMR	Spherically-Mounted Retro-Reflector
TCP	Tool Centre Point

INTRODUCTION

Robotics are an ever-growing market in the economic world. As reported by Sander & Wolfgang (2014), the robotics sector is expected to grow to a $67 billion market in 2025, up from $15 billion in 2010. With the advent of collaborative robots in the industrial domain, there is also a paradigm shift in the way humans interact with robotic systems. A robot is no longer simply an industrial tool, but it has become a co-worker. Through a combination of strengths, human-robot collaboration can optimize a flexible process, especially when the robot takes the role of an intelligent assistant. Thus, collaborative robotics becomes a means by which a process can be further optimized when humans can no longer do so.

Human-robot collaboration can reduce the fatigue and stress of a human co-worker (Ajoudani *et al.*, 2018). Beyond optimization, human-robot collaboration can greatly reduce the occurrence of workforce injury by delegating repetitive tasks to a robot. Overall accuracy, precision, and quality of the task process would be improved through the robotic system while the human offers intuition, experience, and an understanding of correctness.

In the medical sector, ultrasound imaging is an inexpensive approach without the high radiation emissions often associated with other modalities, especially when compared to MRI and CT imaging respectively. While ultrasound imaging may be most widely associated with gynaecology, this modality has many applications, ranging from dermatology to intraoperative procedures. For example, in the case of stenosis (the narrowing of blood passage in the body), diagnosis and intervention planning requires the use of precision imaging. Ultrasound imaging is the modality of choice when assessing stenosis, due to the ability of characterizing both the blood passage anatomy and real-time blood flow (Lutz & Buscarini, 2011).

Advances in ultrasound imaging has now allowed for 3D/4D imaging. Through the real-time image synthesis of the 2D image slices, sonographers can view inner-body volumes of interest while retaining the cost-effectiveness and safety of ultrasound. However, the quality of 3D/4D

scans using freehand ultrasound are limited by the precision and speed of the transducer movement, which can be limited by human ability (Huang & Zeng, 2017).

Work-related injuries can have significant economic repercussions and impede careers. According to Hassard *et al.* (2018), worldwide work-related injuries was estimated to cost upwards of US$187 billion. Studies have shown that sonographers are prone to musculoskeletal injuries, especially carpal tunnel syndrome, from the awkward and repetitive manipulation of the ultrasound transducer. As noted in Vanderpool *et al.* (1993) and subsequently validated by Barros-Gomes *et al.* (2019), upwards of 86 % of sonographers reported one or more physical symptoms of work-related injuries. Increased workload, decreased staff, sustained posture, and equipment design have been identified as factors which could result in work-related injuries among sonographers.

0.1 Problem Statement and Motivation

The use of robotics in medicine is of growing importance for modern health services, as robotic systems have the capacity to improve upon human tasks, thereby enhancing the treatment ability of healthcare providers. While many procedures at a conceptual level follow the same steps, they still require a certain level of human intuition. However, since robots can perform difficult tasks more effectively and efficiently than humans, there are economic benefits associated with the integration of robotic systems in medicine. The more efficiently and effectively a process can be completed, the more economically advantageous it becomes. The economics of such an integration become more advantageous if we can leverage an off-the-shelf collaborative robot instead of a custom-designed system. Thus, the question becomes how a collaborative industrial robot can be used in a freehand medical ultrasound application.

While classical robot motion control is a well studied topic, combining robotics and medical ultrasound is challenging due to the complexity of the resulting application requirements

(Section 1.5). Specifically, medical ultrasound presents patient safety concerns and deformable skin in an unstructured environment. Over the past two decades, considerable effort has been invested into ultrasound robotics research and development. So far, much of the work is focused on the design of the robotic system for patient safety or image generation. Various studies have shown how both intrinsic and extrinsic safety can be achieved through design and basic force control, but then do not test the performance of such a system. Similarly, works that focus on image generation typically neglect to study and discuss the robotic control system used to obtain said images. Hence, while the feasibility of robotic ultrasound has been the focal point, there is a lack of focus of what occurs between system design and image output.

0.2 Research Objectives and Contributions

Given the previously discussed context and motivation, the objective of this research was to develop methods to assist in the development of collaborative robotic freehand ultrasound. These methods target gaps in the current literature and supply robot fundamentals that can improve the capability of existing approaches by:

1. Increase the accuracy and force tracking performance of a force-controlled robotic system
2. Develop contextual awareness for physical human-robot interaction (pHRI) in an unstructured environment

Three primary contributions are made toward this goal and a full list of publications that resulted from this research can be found in Appendix I.

0.2.1 Calibration of a Collaborative Robot

Accuracy is a fundamental requirement for high-quality ultrasound image generation and targeting. This is especially true when following a specified path along a patient or synthesizing

2D slices into a 3D ultrasound image. However, although most industrial robots are inherently precise, they are not necessarily accurate. While robot calibration itself has been extensively studied, many of the approaches rely on expensive and highly delicate equipment. Given the force-controlled nature of an ultrasound robot, we propose a closed-loop calibration approach using impedance control and low-cost equipment. As demonstrated through an experimental study and validated with a laser tracker, the absolute accuracy of a collaborative robot was improved to a maximum error of 0.990 mm, representing a 58.4% improvement when compared to the nominal model. Moreover, further testing showed that this method is comparable in quality to traditional calibration using a laser tracker.

0.2.2 Contact Distinction in an Unstructured Environment

Collisions and contact events are a natural by-product of applications involving pHRI in unstructured environments. Robot-assisted medical ultrasound is an example of a task where simply stopping the robot upon contact detection may not be an appropriate reaction strategy. Thus, the robot should have an awareness of body contact location to properly plan force-controlled trajectories along the human body using the imaging probe. This is especially true for remote ultrasound systems where safety and manipulability are important elements to consider when operating a remote medical system through a communication network. We propose a framework for robot contact classification using the built-in sensor data of a collaborative robot. Unlike previous studies, this classification does not discern between intended vs. unintended contact scenarios, but rather classifies *what* was involved in the contact event. The classifier can discern different ISO/TS 15066:2016 specific body areas along a human-model leg with 89.37 % accuracy. Altogether, this contact distinction framework allows for more complex reaction strategies and tailored robot behaviour during pHRI.

0.2.3 Motion Control Optimization in pHRI

Given that the success of an ultrasound task depends on the capability of the robot system to handle pHRI, pure motion control is insufficient. Force control techniques are necessary to achieve effective and adaptable behaviour of a robotic system in the unstructured ultrasound environment while also ensuring safe pHRI. However, while force control does not require explicit knowledge of the environment, to achieve an acceptable dynamic behaviour, the control parameters must be tuned. Within the context of medical ultrasound, different human body locations have a different stiffness and will require unique tunings. The systematic tuning of practical high degree-of-freedom (DOF) pHRI tasks is not trivial and there are many parameters to be tuned. While traditional tuning is generally performed *ad hoc* and requires knowledge of the robot and environment dynamics, we propose a simple and effective online tuning framework. Through real-world experiments with a collaborative robot, the framework tuned motion control for optimal and safe trajectories along a human leg phantom. The optimization process was able to successfully reduce the mean absolute error (MAE) of the motion contact force to 0.537 N through the evolution of eight motion control parameters. Furthermore, contextual awareness through motion classification can offer a framework for pHRI optimization and safety through predictive motion behaviour with a future goal of autonomous pHRI. As such, a classification pipeline, trained using the tuning process motion data, was able to reliably classify the future force tracking quality of a motion session with an accuracy of 91.82 %.

0.3 Thesis Outline

The work presented in this thesis is organized as follows; Chapter 1 presents relevant concepts and gives a review of relevant works on collaborative robotics, motion control, and robotic ultrasound. Chapter 2 then introduces the proposed calibration approach, based on impedance control and low-cost equipment. The work presented in this chapter corresponds to Nadeau *et al.*

(2019), which was published in *Robotics* and uses the developed open-source robotics framework for the Python programming language (Nadeau, N. A., 2019). Following this, Chapter 3 presents the contact distinction framework that classifies *what* body region was involved in a contact event. Chapter 4 introduces an online tuning framework and predictive motion classification for pHRI. This chapter corresponds to an extended version of Nadeau & Bonev (2018), published and presented at the *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Finally, Section 5.1 summarizes the main contributions of this dissertation and Section 5.2 discusses its limitations as well as possible extensions.

CHAPTER 1

LITERATURE REVIEW

1.1 Overview

Ultrasound robotics is a challenging problem that has attracted a large amount of interest from researchers worldwide. It is the synthesis of pHRI and a medical procedure, requiring many interdisciplinary components for success. Over the years, most research on ultrasound robotics has focused on the feasibility of the application, not the robotic fundamentals, such as motion control, calibration, and contextual awareness. Specifically, motion control plays an essential role in every aspect of pHRI. Calibration allows a system to accurately and safely navigate its environment while reducing the imaging errors associated with positioning. Finally, contextual awareness forms a part of the system's intelligence and gives a foundation for autonomous decision making in pHRI. The following sections present important works related to these lines of research.

1.2 Collaborative Robotics

Given the nature of pHRI in robot-assisted freehand ultrasound, the relationship to collaborative robotics is obvious. With the recent interest in collaborative robotics in the industrial and service domains, much research has been published discussing and analysing their general usage. Discussions on the domain as a whole and the concept of robot co-workers may be found in Scassellati & Tsui (2014) and Ajoudani *et al.* (2018).

The state of the art and challenges of pHRI is reviewed in Lasota *et al.* (2017). As the research notes, the context of pHRI necessitates a different set of application requirements, especially a higher level of safety. In the presentation of safe physical interaction research, Pervez & Ryu (2008) identified three distinct areas of interest: interaction safety assessment, interaction safety through design, and interaction safety through planning and control. Alami *et al.* (2006) notes that a classification for the evaluation of pHRI, safety, and dependability is missing.

More specifically, De Santis *et al.* (2008) constructed an atlas of pHRI research, identifying key elements required for future research and robot design. The authors called for a generalized safety framework for pHRI with weighted scoring and various criteria, thus totalling to an overall safety level. As the authors noted, collisions are the primary risk in human-robot interaction. Knowing that impacts will occur, the authors borrowed from the Abbreviated Injury Scale and the Head Injury Criterion of the automotive industry to classify overall human injury severity with respect to pHRI. These scales attempt to understand the physiological risks resulting from such impacts.

More recently, Lasota *et al.* (2017) considered all possible ways in which harm could come to a human collaborator. The authors explored several manners, ranging from physical contact to negative psychological effects due to disagreeable or unsafe interaction. Similar to Pervez & Ryu (2008), pHRI methods are categorized, including: safety through control, motion planning, and prediction.

Modern tangible examples of pHRI are presented in Maurtua *et al.* (2017). First, in an important step for pHRI research, Haddadin *et al.* (2008) presented a human-robot collision study. This research includes evaluations of impacts between a DLR LWR-III robot and a human arm and chest with velocities up to 2.7 $\mathrm{m/s}$. Beyond offering collision-based control schemes, the authors highlighted the importance of detecting and reacting to different contact scenarios, such as nominal conditions, where physical interaction is desired, and fault conditions, which are the result of unexpected and potentially dangerous collisions. Second, in Haddadin *et al.* (2011) the concept of the robot co-worker is developed with collaborative robotics. Using the DLR LWR-III robot, the authors developed a collaborative bin-picking application, demonstrating robust strategies for safe human-robot interaction, state-dependent behaviour, and motion control in unstructured environments. Third, Maurtua *et al.* (2017) surveyed experiments aimed to measure the trust of workers in industrial pHRI robotic applications. The acceptance of different pHRI methods was gauged and it was found that gesture-based interactions and hand-guiding interaction methods were favoured by participants. Finally, Bousquet-Jette *et al.* (2017) and Cio *et al.* (2019) present scene analysis using vision and machine learning for object grasping by an

assistive robot in a human-robot collaboration scenario. While non-industrial in nature, these systems provide a different perspective to pHRI and environment navigation where the human is the subject and target of the interaction task.

Alternatively, there are times when physical collisions are not desired. Accordingly, Corrales *et al.* (2011), Flacco *et al.* (2012), and Morato *et al.* (2014) presented collision avoidance methods, but all these methods require the use of external vision systems. First, Flacco *et al.* (2012) demonstrated a depth space approach for fast distance evaluation between a KUKA LWR IV robot and a moving object. Second, Corrales *et al.* (2011) reported methods and experimental results using a Mitsubishi PA-10 robot and geometric representations to evaluate robot-object distances and relative velocities. Finally, Morato *et al.* (2014) offered a sensing framework for human-robot collaboration and motion planning based upon a multi-Kinect camera system.

Some researchers have attempted to reach the limits of intrinsic safety with design approaches consisting of compliant manipulators and joints. As demonstrated by Quigley *et al.* (2011), with damping and mechanical flexibility being part of the design requirements of their custom robotic manipulator, energy transfer during impact is greatly reduced. Unfortunately, this compliance resulted in a low precision system, which is often a critical requirement in automation.

1.3 Ultrasound

Ultrasound imaging is an important modality in modern health services and has been found to be the standard choice for whole leg imaging (Aboyans *et al.*, 2017). Many advanced freehand ultrasound techniques are highly dependent on accurate positioning and movement of the ultrasound probe, but this required accuracy is currently limited by the human operator. As detailed by Szabo (2014), ultrasound imaging requires a high level of skill and experience to obtain quality images and diagnose. One of the most demanding aspects is keeping track of the spatial relations between the anatomy, the probe, and the resulting images.

One approach to improving the operator's spatial awareness is through 3D imaging. As reviewed by Huang & Zeng (2017), 3D ultrasound can enhance the comprehension of the anatomy

through high-quality images and more tangible spatial information. While methods for aiding 3D ultrasound have been developed (e.g., mechanical linear scanners, magnetically tracked freehand scanners), it is noted that their bulkiness and weight make them inconvenient to use. Similarly, Parraga *et al.* (2011) reviewed 3D ultrasound with respect to the carotid artery stenosis, where measurement differences and a long duration for the manual process are key limitations. Specifically, scans requiring at least 4 cm of length are not viable, due to limitations in quick and precise human motion.

Furthermore, studies have shown that sonographers are prone to musculoskeletal injuries, especially carpal tunnel syndrome, from the awkward and repetitive manipulation of the ultrasound transducer. Increased workload, decreased staff, sustained posture, and equipment design have been identified as factors which could result in work-related injuries among sonographers. As previously mentioned, upwards of 86 % of sonographers reported one or more physical symptoms of work-related injuries (Barros-Gomes *et al.*, 2019; Vanderpool *et al.*, 1993).

1.4 Medical Robotics

Medical robotics allow for improved medical processes beyond human capability and are a growing aspect of modern medical services. Systematic surveys of the domain are presented in Daneshmand *et al.* (2017) and Taylor *et al.* (2016). In particular, the surveys highlight the rapid expansion of the field over the past few decades. Current key technical barriers include motion control using real-time sensor data, accuracy, and sensitivity (i.e., tactile sense). Furthermore, Taylor *et al.* (2016) identified cost and system flexibility (i.e., designed to easily adapt to changes) as important technical barriers to wider adoption, as engineering overhead and the certification costs of medical devices is very high. While challenges and drawbacks still exist, the increase in accuracy and precision over human capability is notable. Two examples of where robotic precision is effective are presented in Mattos *et al.* (2014) and Bowthorpe & Tavakoli (2016). First, Mattos *et al.* (2014) presented robotic laser photo-microsurgery to complement the surgeon's motor skills. Compared to a standard manual manipulator, robot-assisted laser

photo-microsurgery demonstrated superior usability, accuracy, and controllability over the standard process. Second, Bowthorpe & Tavakoli (2016) demonstrated a robotic system for *beating heart* surgery where the surgeon can perform tasks on a virtually stable heart through teleoperation.

An example of where a commercial robot is applied to a medical application may be seen in Niccolini *et al.* (2016), where a KUKA robot under joint compliance control was used in a collaborative neuroendoscopy application. The motivation for using a commercial robot is twofold: most other neurosurgery robots are cumbersome and expensive. From an application standpoint, the system used closed-loop control with optical tracking to improve the accuracy of impedance control during the procedure. Overall, this application showed targeting errors well below the required threshold and positive feedback was given by the neurosurgeons, thus demonstrating a unique advantage in combining the experience and knowledge of the surgeon with the augmentation of a collaborative robotic system.

Alternatively, rather than use commercially available robots, custom robot design is a common research approach. A first example is *Dermarob*, a safe robotic system for skin harvesting in reconstructive surgery, presented in Dombre *et al.* (2003). While this novel design is surgical in nature, it presents an interesting case of requiring motion along skin (a deformable surface) and the respective control schemes. The intended application is to improve the performance of surgeons who do not regularly perform skin harvesting operations and who are not accustomed to the tools used. *In vivo* force-contact experiments with the robot on pigs showed equivalent performance to surgeons. Furthermore, Duchemin *et al.* (2005) developed biomechanical models of skin and soft tissues for the *Dermarob* robot to model their behaviour for robot contact and motion planning. In terms of non-surgical human-contact robotic applications, a unique application and design is presented by Tsumaki *et al.* (2008), where a skin care robot was developed to apply lotion on an elderly person's back for xerosis treatment. The work is notable for user safety, as the system can recognize the geometry of the human body, plan appropriate motion paths with a human subject, and execute the task without excessive forces. This research direction is continued by Papageorgiou *et al.* (2018), where a service robotic system for assisted

living is designed. As elderly people with different impairments have different needs, adaptive motion control and robot behaviour is necessary. Importantly, the authors noted that the difficulty of the task lies in operating over the subject's body, a curved, deformable, and unstructured surface.

1.5 Ultrasound Robots

Expansive reviews of robot-assisted ultrasound imaging are presented in Priester *et al.* (2013), Petrovic *et al.* (2014), Monfaredi *et al.* (2015), and Swerdlow *et al.* (2017). The following section will highlight notable studies within the ultrasound robotics domain. A summary of all relevant studies may be seen in Table II-1.

The work presented by Boudet *et al.* (1997) and Pierrot *et al.* (1999) initiated the domain of robotic ultrasound research. The presented force-control feasibility studies and risk assessments demonstrate the minimum requirements for a successful robotic ultrasound solution. As noted in these studies, the addition of robotic automation and precision to ultrasound imaging opens new doors, further augmenting the capabilities of this modality over human-based freehand solutions. Yet, even since these studies, robotic ultrasound feasibility is still an overarching research trend and is limited to three major research topics: medical imaging and visual servoing, teleoperation, and robot design.

While these topics are important in the context of the application, they are disconnected and limit the practicality of the research. With the inclusion of implementation and application components, such as collaborative robotics (e.g., hand guiding, safety, contextual awareness), advanced control methods (e.g., beyond basic force control), and system optimization (e.g., calibration), ultrasound robotics research can grow beyond feasibility studies.

1.5.1 Robot Design

When designing an automation solution, there is always a decision to be made between focusing solely on a specific task or to generalize the solution to be able to adapt it to a range of tasks.

However, safety is often seen as the guiding principle behind ultrasound robot design. The first example of a specific ultrasound robotic solution is the *Hippocrate* ultrasound robot (Pierrot *et al.*, 1999). As described in their research, the project originally began with a Mitsubishi PA-10 industrial robot, but this was quickly identified as a hazard due to little intrinsic safety. Thus, an intrinsically safe ultrasound robot, *Hippocrate*, was designed. In the case of Abolmaesumi *et al.* (2002), their teleoperated ultrasound robot uses a unique counterbalance design to ensure an acceptable level of intrinsic safety. With a focus on the carotid artery, the authors demonstrated a dynamic controller for haptic feedback and key features required for an effective and ergonomic human-robot user interface. Next, the *MedRUE* robot may be seen in Lessard *et al.* (2006). Beyond the novel parallel design, the research detailed an in-depth analysis of the workspace requirements for ultrasound imaging along the lower limbs of the human body and investigated using static balancing methods to improve the level of system safety. As noted in Zhao, L. (2015) and Joubair *et al.* (2016), this custom robot, while novel in design and kinematics, was not accurate. Finally, Martinelli *et al.* (2007) demonstrated a robotic tele-echography system, known as *TER*, for isolated sites or remote health centres. Despite being bulky, inaccurate, and imprecise, the original master-slave cable-driven design allows the system to easily adapt to many different body types without a specific control scheme. In clinical evaluations, the robot was able to successfully perform remote abdominal aortic examinations.

Portability and compactness are key features of ultrasound imaging and several robot design approaches embrace these features when augmenting the modality with automation. Courreges *et al.* (2008) detailed the *OTELO* system, a small, teleoperated, portable ultrasound robot. Similarly, the *TERESA* project (Arbeille *et al.*, 2014) proposed a compact teleoperated system for astronauts. Alternatively, robotic gantry designs are presented in Pahl & Supriyanto (2015) and Sen *et al.* (2015), while Stewart platforms are presented in Ito *et al.* (2013) and Seo *et al.* (2015). While the studies focused primarily on design and teleoperation, these unique approaches have a limited workspace and manipulability.

While custom robots may be designed for a specific application, commercial, off-the-shelf robots are a robust and economical solution. It is for this reason that the work presented in Chapters 2

to 4 is based on the KUKA LBR iiwa 7 R800 (LBR) robot. The total cost of developing solutions is greatly reduced when a robust and off-the-shelf robotic platform is used, as the manipulator itself does not need to be debugged. Moreover, the LBR was designed specifically for pHRI, as evidenced by the integrated joint torque sensors, responsive safety controller, and KUKA's push into the medical sector with the LBR Med platform.

By the same token, Conti *et al.* (2014) presented a teleoperated, haptic feedback ultrasound robot based upon the KUKA LWR. The system offers multiple ways for the operator to control the robot, including manual control, haptic device control, and pre-programmed trajectories. The authors mentioned that informal tests by medical staff suggest that the system meets the defined goal of reducing user fatigue and is suitable for ultrasound imaging sessions. The presented force controller maintains a constant contact force between the ultrasound probe and the patient's skin while the robot moves the probe to desired anatomical locations. While the authors focused on the graphical interface, haptic control, network communication, and image processing, the study lacks insight into safety, control theory, and fundamental implementation details.

Janvier *et al.* (2008), Janvier *et al.* (2010), and Janvier *et al.* (2014) presented robotic ultrasound using a CRS F3 robot for 3D image reconstruction of lower limb arteries in the context of stenosis diagnosis. With a focus on medical imaging, feasibility of the robotic application was the primary focus of the studies. Sub-millimetre errors were achieved, and a successful *in vivo* study was performed on a human volunteer. In Janvier *et al.* (2008), a teach/replay mode was used for collaborative probe manipulation, but only in free space, not along a surface, as would be the case in an ultrasound application. This would explain why the authors obtained repeatability results like the nominal repeatability stated by the manufacturer. This limitation becomes apparent in the subsequent studies (Janvier *et al.*, 2010, 2014), as the teach/replay feature was removed due to difficulty in maintaining contact with the patient's skin. Overall, the authors focused wholly on the medical imaging aspects of robot-assisted ultrasound with simplified motion control.

Likewise, Mustafa *et al.* (2013) presented an ultrasound robot based upon a Mitsubishi MELFA RV-1 robot for liver imaging. The authors notably developed a control algorithm to compensate for patient movement due to breathing. Victorova *et al.* (2019) proposed a collaborative robot approach to manipulating an ultrasound probe over a patient's spine for scoliosis assessment. The authors noted that different degrees of spinal deformity result in significantly higher forces during scanning. While these studies included an experimental component, they were focused on imaging feasibility and lacked certain robot fundamentals, such as the methodology behind how the controller gains were determined.

1.5.2 Motion Control

Control and motion planning are key elements of robotic applications, especially for pHRI in unstructured environments, such as freehand ultrasound. A review of the fundamental motion control strategies may be seen in Whitney (1987), Zeng & Hemami (1997), Chiaverini *et al.* (1999), and Yoshikawa (2000). The origins of force-based control methods are seen in Mason (1981), Raibert & Craig (1981), and Almeida *et al.* (1999). In collaborative settings, simple motion control strategies may be inadequate, as collaborative environments are unstructured. By extending common control approaches for pHRI safety, Vick *et al.* (2013) and Zanchettin *et al.* (2016) present collaborative motion control for shared workspaces using a redundant robot, thereby prioritizing not only safety, but also robot performance.

Regarding unstructured environments, adaptive motion control strategies may be necessary to compensate for dynamic and unpredictable disturbances. A simple demonstration of this is seen in Ott *et al.* (2015), where impedance control and admittance control are implemented as an adaptive spectrum on a one-dimensional robot system. Impedance control was found to be appropriate for interaction with stiff environments, whereas admittance control was appropriate for interaction with soft environments or free space. A mixture of impedance and admittance control was dynamically applied for in-between environments.

Within the topic of robot-assisted ultrasound, many approaches choose to use external measurement and tracking devices, such as in Santos & Cortesao (2015), Mebarki *et al.* (2010), Graumann *et al.* (2016), and Huang *et al.* (2018), where visual servoing is demonstrated using 3D time-of-flight and depth cameras, but with only simple force control and no tuning analysis. A more intricate approach to vision-based trajectory planning is seen in Hennersperger *et al.* (2016), where structured light and a Kinect 3D camera are used. However, the study demonstrated poor accuracy (i.e., 2.47 ± 0.96 mm).

More accurate visual servoing was developed using the ultrasound images as the external input, such as in Nadeau & Krupa (2013), Zettinig *et al.* (2016), Chatelain *et al.* (2016), and Chatelain *et al.* (2017). However, in the case of Nadeau & Krupa (2013), a weak contact force ($\sim$1 N) makes for easier motion control along the patient's surface. While these studies all evaluated novel visual servoing approaches, the experimental methodologies did not discuss controller tuning or consider the practicality of relying on external measurement devices for motion control.

Alternatively, haptic teleoperation is a common research domain and motion control approach. The basic components for such an application are outlined in Mathiassen *et al.* (2016), where a commercial collaborative robot is used in a feasibility analysis. Notably, the authors rejected absolute accuracy as a requirement and instead focused on repeatability. In a similar manner, Abolmaesumi *et al.* (2001) presented a user interface for teleoperation for the custom robot design presented in Salcudean *et al.* (1999).

Regardless, while many of these studies present interesting approaches and tools for freehand ultrasound motion control, few studies present the controller or control algorithms. Notable exceptions are Karar (2018) and Santos & Cortesão (2018). First, Karar (2018) presented an adaptive PID force controller. While the concept is only tested in simulation, the authors correctly noted that manual tuning for pHRI is time intensive, not sufficient to obtain optimal gain values, and difficult without *a priori* knowledge of the environment dynamics. Second, Santos & Cortesão (2018) presented a combined computer vision and teleoperation approach whereby the robot dynamics are controlled by a depth camera. The method was validated in a

clinical scenario, but an external measurement system is still required for motion control, the contact event is simply modelled as a linear spring, and the controller gains were selected on an *ad hoc* basis, limiting the adaptability of the solution to other applications and environments. Overall, the authors remarked that a perfect environment model and an abundance of external sensor input are impractical in pHRI contact scenarios. Instead, rather than aiming to achieve a perfect application design, the focus should be placed on developing an estimation that improves stability and boosts force tracking performance.

1.5.3 Calibration

In robotic ultrasound, accurate probe motion allows for better motion control and image acquisition. This is especially true for haptic and teleoperation applications. By controlling an ultrasound probe through a haptic controller, teleoperation allows for remote examinations, but the probe motion must accurately replicate the sonographer's movements. An example of poor position tracking in haptic teleoperation may be seen in Sharifi *et al.* (2017), where the error between the robotic master and slave devices was measured to be 1.48 mm.

While external sensors (e.g., computer vision) can improve motion control accuracy (e.g., visual servoing), calibration is a cost-effective approach to improving the intrinsic accuracy of the robotic system. From an industrial perspective, robot calibration is a well-studied topic with several decades of robust research.

Aalamifar *et al.* (2014) presented a calibration approach using an external camera to track the robot-held probe in an ultrasound tomography application. While experiments showed more accurate results than human-held freehand ultrasound, the robot-held accuracy was only 1.79 mm. Alternatively, Joubair *et al.* (2016) used a force-based approach to ultrasound robot calibration, reducing the maximum absolute accuracy errors of a custom robot from 12 mm to 0.320 mm, but these results are from a simulation-only experiment.

1.5.4 Safety and Contact Distinction

Safety is a key issue when designing any robotic application. A review of safety approaches for pHRI and collaborative applications within an industrial setting may be found in Bicchi *et al.* (2008) and Rezazadegan *et al.* (2015). In terms of medical robots, a general overview of safety considerations may be seen in Fei *et al.* (2001). As described, the design and use of a medical robotic system should follow a safety protocol that identifies and defines hazards, limits, and user requirements to ensure application safety.

From the perspective of pHRI, robot-assisted ultrasound is an example of an application where human-contact is explicit and necessary. While the human body and clinical environment is unstructured, simply stopping the robot upon contact detection is not an appropriate reaction strategy. Alami *et al.* (2006) described safe and dependable pHRI in human environments, where it is necessary to distinguish between intentional contacts and accidental collisions. Notably, the authors noted that impedance parameter selection can be derived from contact distinction. In terms of ultrasound, the robot should have an awareness of body contact location to properly plan force-controlled trajectories, as safety and manipulability are important elements to consider, especially when operating a teleoperated system.

Collision detection and contact distinction allow the system to understand the current state of the application environment (Haddadin *et al.*, 2011). As a result, reaction strategies may be developed to improve post-contact event safety and enables the robot to navigate unexpected and human-caused events. Yet, few robot ultrasound studies offer insight into contact distinction. In Hennersperger *et al.* (2016), the authors noted that Cartesian stiffness for force-controlled trajectories can range from 125 N/m to 500 N/m depending on the patient anatomy, but the robotic system has no feedback loop to identify the *current* anatomical target or region. Furthermore, Virga *et al.* (2016) highlighted a high inter-operator variability and a low repeatability in ultrasound image acquisition due to anatomical differences in stiffness and probe force requirements. As a solution, the authors presented an online probe force estimation framework to optimize the visibility of the aorta during robotic ultrasound imaging. Though the

system is successful in a clinical scenario, it requires the use of an external vision system and co-registration with a MR-atlas for localization.

1.6 Summary

This chapter presented a vast array of techniques and research covering robotic ultrasound, pHRI, and other relevant topics. Due to complexity of robot-assisted freehand ultrasound, the research presented in the literature survey typically considers only a single subtopic to study. Most of the covered robotic ultrasound approaches tackle the problem by focusing only on the medical imaging component or the robot design, with little focus on the robotic performance. These approaches suffer from critical limitations that impede their use in real-world applications. In particular, basic robot programming and force control is used without considering adaptability and scalability. In many cases, this can result in reduced generalizability and transferability of the solution (i.e., designed to easily adapt to new applications and environments).

However, going beyond programming basics could help improve the ultrasound process while ensuring safe pHRI. The human body is a deformable surface and an unstructured environment, representing both a safety concern and a challenge for trajectory planning and control. Notably, few approaches have considered adaptable motion control for versatile probe navigation. But, the ability to classify contact events would allow for more complex robot reaction behaviour and give the system a level of contextual awareness. In various applications, if such information is readily available, the robot could autonomously respond to unforeseen events in a safe manner instead of triggering an error. Combining this adaptable and intelligent motion control with an accurate robot model could also improve ultrasound imaging performance, especially for 3D/4D image acquisition. The following chapters of this thesis present approaches and methods to address these missing elements.

CHAPTER 2

IMPEDANCE CONTROL SELF-CALIBRATION OF A COLLABORATIVE ROBOT USING KINEMATIC COUPLING

This chapter presents a closed-loop calibration approach using impedance control. The process is managed by a data communication architecture based on open-source tools and designed for adaptability. The calibration procedure uses precision spheres and a kinematic coupling, standard machine tool components, which are suitable for harsh industrial environments. As such, the required equipment is low cost, robust, and is quick to set up, especially when compared to traditional calibration devices. As demonstrated through an experimental study and validated with a laser tracker, the absolute accuracy of the KUKA LBR iiwa robot was improved to a maximum error of 0.990 mm, representing a 58.4% improvement when compared to the nominal model. Further testing showed that a traditional calibration using a laser tracker only improved the maximum error by 58 µm over the impedance control approach.

2.1 Introduction

Modern robotic applications often rely on offline programming to reduce process downtime. In a virtual environment, robot application specialists may program, visualize, and test their robotic application before uploading it to the real production environment. This offline process saves both time and costs. However, to achieve a high level of fidelity between virtual and production environments, the robot system must be accurate.

Unfortunately, even though most industrial robots are inherently precise (i.e., repeatable), they are not necessarily accurate (Nubiola & Bonev, 2013). One cost-effective approach to obtaining a more accurate robot is through calibration, where the actual kinematic and non-kinematic parameters of the robot model are identified and improved upon when compared to the nominal model. While robot calibration itself has been extensively studied, many of the approaches rely on expensive and highly delicate equipment. Such approaches include laser trackers (Nubiola & Bonev, 2013; Nubiola *et al.*, 2014; Zhao *et al.*, 2015), stereo vision systems (Svaco

et al., 2013), portable measurement arms (Joubair *et al.*, 2012), digital indicators (Gaudreault *et al.*, 2016), and coordinate measuring machines (CMM) (Borm & Meng, 1991; Nubiola *et al.*, 2014). These devices provide a guaranteed high level of accuracy, whereas they are generally very expensive.

Fuelled by the promises of more flexible automation, there is a proliferation of collaborative robots (cobots) and force-torque sensors in industrial environments. As distributed intelligence and adaptability are some of the cornerstones of modern robot applications (Lee *et al.*, 2015), impedance control techniques are becoming more common and allow for a robotic system to perform tasks as autonomously as possible. As identified by Universal Robots, through the launch of their e-Series collaborative robots, tool-centric force-torque sensing (especially when built-in) provides greater precision and sensitivity to address these applications. Specifically, impedance control regulates the dynamic interaction between a manipulator and its environment and is suitable for interaction and object manipulation. For instance, traditional jogging and manual registration of application waypoints can be replaced by hand-guided teaching.

Thus, it is through impedance control that a simple, low-cost, closed-loop calibration approach has been developed. This robot calibration strategy is automated and uses no external measurement devices, while extending upon previous work in a timely manner with application results that are relevant to today's industrial and academic needs. Through the synthesis of the magnetic coupling concept from Joubair *et al.* (2013) and the spherical constraint presented in Joubair & Bonev (2015b), the present study uses impedance control and a calibration plate with three precision spheres to calibrate a KUKA LBR iiwa 7 R800 collaborative robot. While the standard calibration process (Elatta *et al.*, 2004; Mooring *et al.*, 1991) remains the same, the use of impedance control adds an adaptable element to the approach. This adaptability comes from the fact that the absolute accuracy between the robot and the calibration tooling is not a high priority, as impedance control allows for the coupling tool to self-centre with the spheres upon contact. In such a way, workcell design can remain flexible without stringent accuracy requirements, and the robot is capable of self-navigating to its own calibration waypoints. Furthermore, as a secondary contribution, the automated calibration process is managed by a

modern communications framework that is based on open-source tools. As such, the framework is designed specifically for flexible and scalable interconnection and information transparency requirements (Lee *et al.*, 2015). Finally, the performance of this strategy is further validated through a comparison with traditional laser tracker calibration.

While the concept of using geometric constraints for robot calibration is not new, the integration of impedance control to automate and improve the process has yet to be studied. One similar approach, single endpoint contact calibration, was shown in Meggiolaro *et al.* (2000) using a ball joint fixture, but manual constraint of the robot is required, reducing automation potential. As the process uses only a single fixture point, the process calibrates the robot with respect to a virtual world frame, not a physical reference frame that would be part of the application workspace, such as that presented here. Otherwise, the plane constraint method, tested in Joubair & Bonev (2015a), constrains the tool movement to at least three planes. However, this approach requires an expensive precision touch probe and accurate plane equations that may involve mechanical characteristics (e.g., perpendicularity, flatness) of the tooling, reducing the overall simplicity and robustness of the method. Regarding the calibration of a KUKA LBR robot, the work in Besset *et al.* (2016) presented a unique calibration approach using the joint torques, but only 10 points were used to validate their calibration errors. Consequently, there were too few points for a definite measure of global accuracy. In Gaudreault *et al.* (2018), an automated measuring device was presented that also uses tooling spheres. Despite this, the device itself is relatively fragile and requires Bluetooth, which may have poor performance due to the harsh electromagnetic ambient environment of industrial settings (Neelakanta & Dighe, 2003). Given the current state of research, a driving motivation for this study was to develop a geometric constraint calibration process that leveraged modern robot control theory and capability while considering the current needs of industrial applications and environments.

The remainder of this work is organized as follows. In Section 2.2, the technical details of robot modelling and kinematics are presented, while Section 2.3 addresses aspects of kinematic parameter identification. Section 2.4 describes the methodology of the calibration processes and

Section 2.5 presents the validation results. Finally, Section 2.6 discusses the outcomes of this study and potential future work.

2.2 Calibration Model

The developed calibration procedure was tested with a seven-axis KUKA LBR iiwa 7 R800 robot (Figure 2.1). Having seven revolute joints allows this robot to have kinematic redundancy and to execute joint motions through its null space for a given Cartesian pose. The calibration model presented identifies the kinematic parameters associated with the robot and the world and tool frames. This allows for a more accurate representation of the robot application.

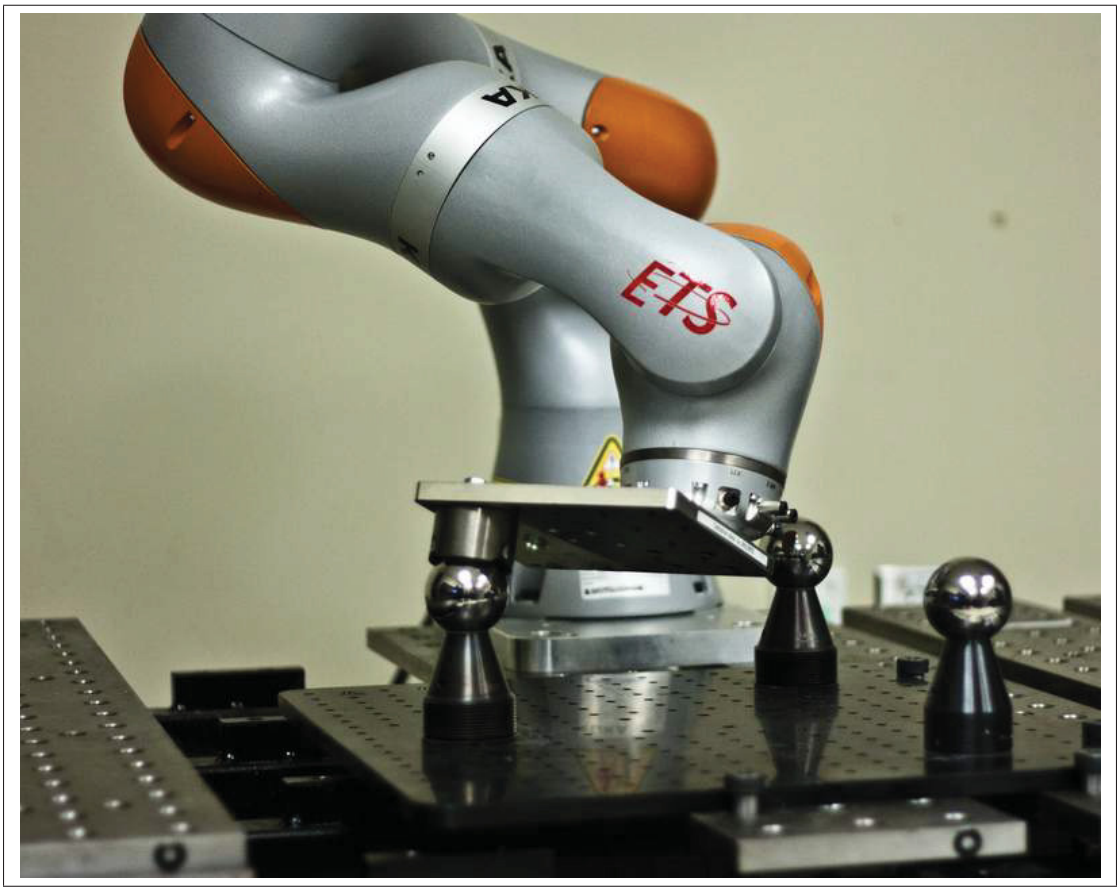

Figure 2.1 Impedance controlled calibration of a KUKA LBR iiwa 7

2.2.1 World and Tool Frames

The world frame (F_w) of the application describes the origin of the Cartesian space and is the base reference for all measurements. As shown in Figure 2.2, this frame is represented by the right-handed coordinate system defined by the three calibration spheres. Similarly, the tool frame (F_T; Figure 2.3) describes the tool centre point (TCP) of the equipped robot tool with respect to the natural flange of the robot and is located at the centre of a coupled sphere. In the present study, 6-DOF (three translation and three rotation) were considered for the transformation between the world frame and the robot base (located at the intersection of the first two joints). As position accuracy was the focus of the study, only translational DOF were considered for the transformation between the robot flange and the tool frame.

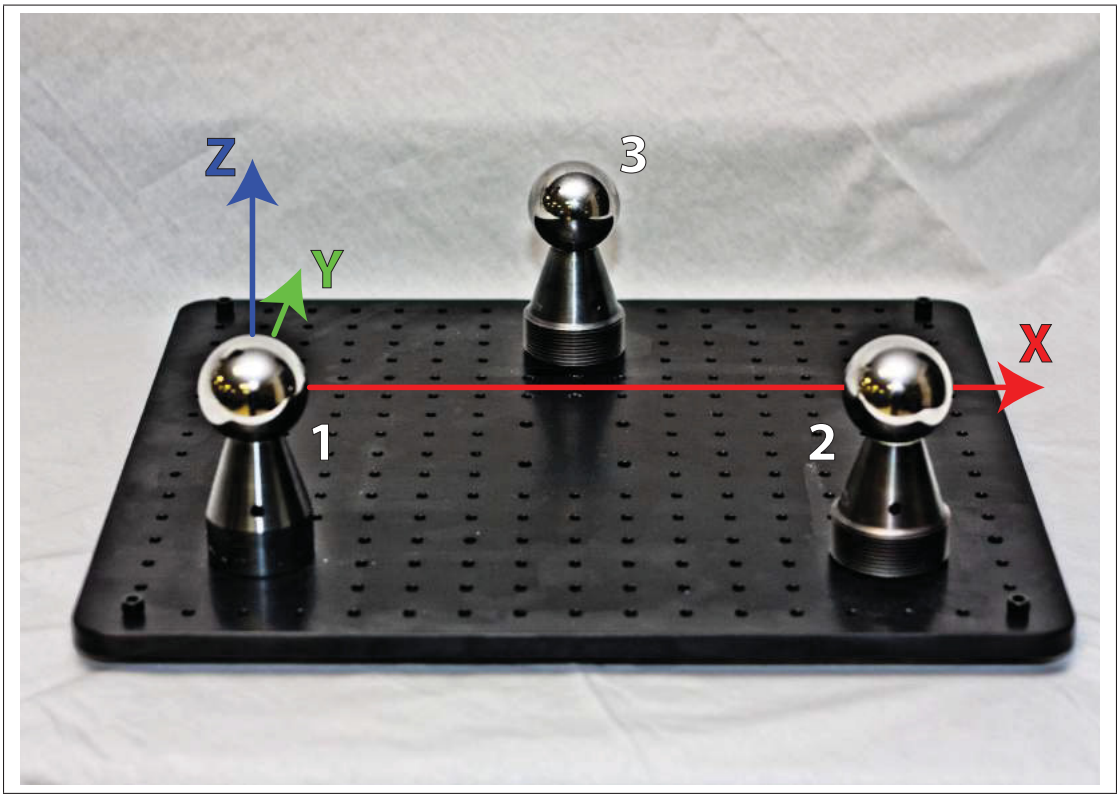

Figure 2.2 Calibration plate with three 50.8 mm precision spheres. The origin of the world frame is located at the centre of Sphere 1. The x-axis is defined by the centres of Spheres 1 and 2. The xy-plane is defined by the centres of the three spheres

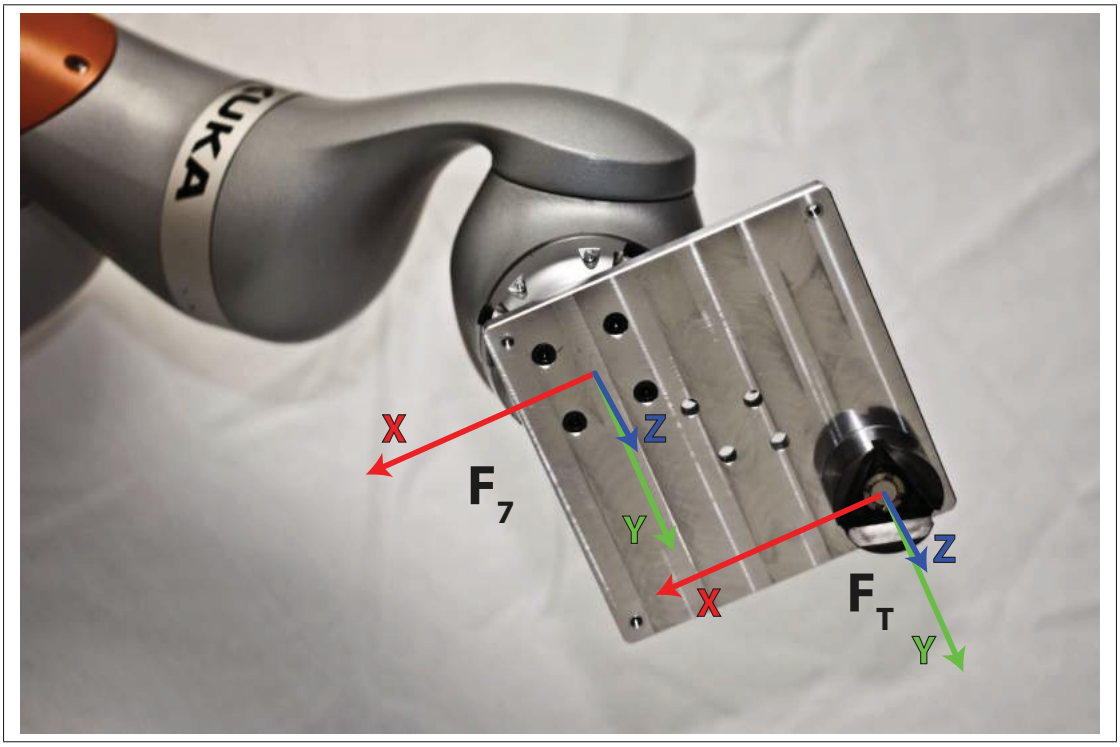

Figure 2.3 Robot calibration tool with trihedral kinematic coupling. For both the flange (F_7) and tool (F_T) frames, the xy-plane is parallel to the plate

2.2.2 Tooling

The experimental tooling, shown in Figures 2.2 and 2.3, was composed of a trihedral coupling (mounted eccentrically from the flange centre) and a calibration plate with 50.8 mm precision spheres. The total mass of this tooling was 1.5 kg. Both the coupling and spheres were sourced as stock components from Bal-Tec (https://www.precisionballs.com). Notably, the trihedral coupling has a strong embedded magnet (2.7 kg pull strength), which aids in mating the part to the steel sphere. Prior to the experiment, the absolute positions of the spheres (Table 2.1) were obtained using a Mitutoyo Bright-STRATO 7106 CMM with an uncertainty of ±2.7 µm (95% confidence level).

Table 2.1 Absolute positions of the three calibration spheres with respect to the world frame

Sphere	X (mm)	Y (mm)	Z (mm)
1	0	0	0
2	300.009	0	0
3	150.049	260.051	0

2.2.3 Kinematic Model

The KUKA LBR iiwa 7 R800 robot was modelled using modified Denavit-Hartenberg (MDH) parameters (Craig, 2005), describing the geometric parameters and the kinematic chain from the robot base (F_0) to flange (F_7). Using this convention, the forward kinematics calculation of a given set of joint angles results in a position and orientation in Cartesian space.

Given a joint angle vector:

$$q = [\theta_1 \dots \theta_7]^T \tag{2.1}$$

Any position and orientation of the tool with respect to the world frame can be described as a series of transformations along the kinematic chain of the robot:

$$^{W}_{T}T = {}^{W}_{0}T\,{}^{0}_{7}T(q)\,{}^{7}_{T}T \tag{2.2}$$

where,

$$^{0}_{7}T(q) = {}^{0}_{1}T(\theta_1)\,{}^{1}_{2}T(\theta_2) \dots {}^{5}_{6}T(\theta_6)\,{}^{6}_{7}T(\theta_7) \tag{2.3}$$

Each ith transformation along a link in the kinematic chain may be defined by a tuple of MDH parameters (i.e., α_{i-1}, a_{i-1}, θ_i, and d_i), as detailed by the following:

$$
{}^{i-1}_{i}T = \begin{bmatrix}
\cos(\theta_i) & -\sin(\theta_i) & 0 & a_{i-1} \\
\sin(\theta_i)\cos(\alpha_{i-1}) & \cos(\theta_i)\cos(\alpha_{i-1}) & -\sin(\alpha_{i-1}) & -\sin(\alpha_{i-1})d_i \\
\sin(\theta_i)\sin(\alpha_{i-1}) & \cos(\theta_i)\sin(\alpha_{i-1}) & \cos(\alpha_{i-1}) & \cos(\alpha_{i-1})d_i \\
0 & 0 & 0 & 1
\end{bmatrix} \tag{2.4}
$$

2.2.4 Non-Kinematic (Stiffness) Model

Under ideal rigid body mechanics, the geometric parameters described in Section 2.2.3 would be enough for a full calibration. However, under real-world external forces (e.g., gravity), non-kinematic parameters may have negative effects on the accuracy of the robot. As described in Nubiola & Bonev (2013) and Nubiola & Bonev (2013), a simplified model of gearbox stiffness may be used to account for these sources of error. By modelling the gearbox of each joint as a linear torsional spring with stiffness k_i, joint deflections due to the applied external joint torques (τ_i) were accounted for, as given by $\delta\theta = k\tau$. As the robot was studied under static conditions, only external torques due to gravity were observed (i.e., no dynamic forces due to motion).

Initial experiments demonstrated a negligible effect (but greatly increased computational time) by including the non-kinematic parameters due to joint stiffness during calibration. As the robot was gently coupled to a sphere, the structural closed-loop acted as a support against gravity. As such, these errors are not considered for the calibration process.

2.3 Parameter Identification

In the presented calibration, 21 kinematic error parameters were considered. While in theory, there are 28 kinematic error parameters, each associated with an MDH parameter, seven of these were found to be redundant and would degrade the efficiency of the applied optimization method. All considered kinematic errors may be seen in Table 2.2.

Table 2.2 Considered kinematic parameters for the calibration and validation process. The redundant parameters are denoted with an asterisk (*)

i	α_{i-1} (rad)	a_{i-1} (mm)	θ_i (rad)	d_i (mm)
1	0*	0*	θ_1*	0*
2	$-\frac{\pi}{2} + \delta\alpha_1$	δa_1	$\theta_2 + \delta\theta_2$	δd_2
3	$\frac{\pi}{2} + \delta\alpha_2$	δa_2	$\theta_3 + \delta\theta_3$	$400 + \delta d_3$
4	$\frac{\pi}{2} + \delta\alpha_3$	δa_3	$\theta_4 + \delta\theta_4$	δd_4
5	$-\frac{\pi}{2} + \delta\alpha_4$	δa_4	$\theta_5 + \delta\theta_5$	$400 + \delta d_5$
6	$-\frac{\pi}{2} + \delta\alpha_5$	δa_5	$\theta_6 + \delta\theta_6$	0*
7	$\frac{\pi}{2} + \delta\alpha_6$	0*	θ_7*	$126 + \delta d_7$

2.3.1 Non-Identifiable and Redundant Parameters

Several parameters in the robot model were not identifiable due to Cartesian and kinematic redundancies. A direct effect of these redundant parameters is two-fold. First, they increase the number of dimensions in the calibration solution space, increasing the computational requirements and complexity. Second, they reduce the ability of the optimizer (Section 2.3.2) to converge to an optimal solution using the change in the optimization vector as a stopping criterion.

Given that the position and orientation of the robot base were tied directly to the definition of the world frame, the first four kinematic error parameters associated with the transformation between the robot base and the first joint (i.e., α_0, a_0, θ_1, and d_1) were excluded, as the world frame was identified separately. Additional kinematic parameters were excluded through an analysis of the rank and condition number of the robot position Jacobian (Joubair & Bonev, 2015b), as shown below.

Given a robot TCP position (P) as a function of the MDH parameters,

$$P = \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} f_x(\alpha_{i-1}, a_{i-1}, \theta_i, d_i) \\ f_y(\alpha_{i-1}, a_{i-1}, \theta_i, d_i) \\ f_z(\alpha_{i-1}, a_{i-1}, \theta_i, d_i) \end{bmatrix} \tag{2.5}$$

the position Jacobian (J_P) may be described as:

$$J_P = \begin{bmatrix} \frac{\partial x}{\partial \alpha_{i-1}} & \frac{\partial x}{\partial a_{i-1}} & \frac{\partial x}{\partial \theta_i} & \frac{\partial x}{\partial d_i} \\ \frac{\partial y}{\partial \alpha_{i-1}} & \frac{\partial y}{\partial a_{i-1}} & \frac{\partial y}{\partial \theta_i} & \frac{\partial y}{\partial d_i} \\ \frac{\partial z}{\partial \alpha_{i-1}} & \frac{\partial z}{\partial a_{i-1}} & \frac{\partial z}{\partial \theta_i} & \frac{\partial z}{\partial d_i} \end{bmatrix} \tag{2.6}$$

where the number of columns equals the number of MDH parameters currently being considered. Following the derivation of the Jacobian, the nominal MDH values were substituted into the J_P. Next, the corresponding position (x, y, z) of 29 uniformly-distributed joint configurations across the whole joint space was computed and substituted into J_P, giving an augmented Jacobian ($J_{P'}$) with 87 rows. A minimum of 10 configurations were required to generate an overdetermined system, as we had 28 unknowns with three constraints per joint configuration. By defining the Jacobian around these joint configurations, we represented the relationship between an infinitesimal MDH parameter change (i.e., $\partial\alpha$, ∂a, $\partial\theta$, ∂d) and the respective infinitesimal displacement (i.e., ∂x, ∂y, ∂z) of the TCP across the joint configuration space covered by the 29 joint positions. The rank ($r_{J_{P'}}$) of the augmented Jacobian quantifies the number of identifiable kinematic parameters, while the condition number ($c_{J_{P'}}$) represents the sensitivity of this relationship. The condition number is defined as the L2-norm of the Jacobian times the L2-norm of the inverse of the Jacobian (computed directly using SVD), and the calculation was implemented by the Python scientific computing package NumPy in the function *numpy.linalg.cond*. The non-identifiable parameters (columns of the Jacobian) were removed using Algorithm 2.1 to generate a better conditioned problem. In the present study, δd_6, δa_6,

and $\delta\theta_7$ were found to be redundant. A more detailed analysis of this selection algorithm may be seen in Joubair *et al.* (2013) and Joubair & Bonev (2015b).

Algorithm 2.1 Redundant parameter identification

calculate condition number of $J_{P'}$, $c_{J_{P'}}$
calculate rank of $J_{P'}$, $r_{J_{P'}}$
$minConditionNumber \leftarrow c_{J_{P'}}$
$columnToRemove \leftarrow null$
while $r_{J_{P'}}$ < number of columns in $J_{P'}$ **do**
  for column i in $J_{P'}$ **do**
    $J^* \leftarrow J_{P'}$
    remove column i from J^*
    calculate condition number of J^*, c_{J^*}
    calculate rank of J^*, r_{J^*}
    if $r_{J^*} == r_{J_{P'}}$ **then**
      if $c_{J^*} < minConditionNumber$ **then**
        $minConditionNumber \leftarrow c_{J^*}$
        $columnToRemove \leftarrow i$
      end if
    end if
  end for
  remove column $columnToRemove$ from $J_{P'}$
end while

2.3.2 Optimization and Numerical Modelling

With the identifiable parameters, an optimal set may be found that improves the accuracy of the robot. Like Joubair & Bonev (2015b) and Gaudreault *et al.* (2016), distance errors between spheres were used for robot model optimization, removing the dependency (and associated errors) of the world frame from the calibration step. The relative distance fitness function (f_{fit})

for a current set of robot parameters is defined below through Equations (2.1) to (2.3) and the real sphere positions (P_{sphere_i} and P_{sphere_j}).

Given the absolute position of two spheres and two associated measurement configurations,

$$P_{sphere} = \begin{bmatrix} x_{sphere} \\ y_{sphere} \\ z_{sphere} \end{bmatrix} \tag{2.7}$$

$$P_{TCP}(q) = \begin{bmatrix} {}^{w}_{t}T_{x} \\ {}^{w}_{t}T_{y} \\ {}^{w}_{t}T_{z} \end{bmatrix} \tag{2.8}$$

The real distance ($D_{sphere_{i,j}}$) between spheres i and j were found,

$$D_{sphere_{i,j}} = P_{sphere_i} - P_{sphere_j} \tag{2.9}$$

Next, the forward kinematic TCP distance between measured spheres i and j were found,

$$D_{TCP_{i,j}}(m, q) = P_{TCP}(m, q_i) - P_{TCP}(m, q_j) \tag{2.10}$$

where m represents the current set of MDH parameters. These parameters were the array of parameters to be optimized during calibration.

Finally, the relative fitness was calculated for all given combinations,

$$f_{fit}(m, q) = \sum (D_{sphere_{i,j}} - D_{TCP_{i,j}}(m, q))^2 \tag{2.11}$$

where i and j correspond to the ith and jth measured configuration coupled with a sphere.

The Nelder-Mead method (equivalent to *fminsearch* in MATLAB) (Nelder & Mead, 1965) was used for optimization with the above fitness function and initialized with the nominal robot model. This method is robust to noise and works well with non-locally smooth data, such as discrete and noisy measurement data. At each solver iteration, the forward kinematic relative distance errors were evaluated for a given set of robot parameters. The numerical methods and kinematic modelling were implemented using Pybotics (Nadeau, 2019), allowing multicore processing and a direct interface to the optimization methods implemented by the SciPy Python package.

2.4 Methodology

Three primary components formed the basis of this study: communication, measurement, and calibration. First, a modern communications architecture was developed to transmit data and command packets between an external PC controller, the KUKA LBR iiwa robot, and a FARO ION laser tracker. Next, with a common communications and data processing framework, impedance control and traditional laser measurement and calibration processes were developed.

2.4.1 Communication

A calibration framework viable for adaptability and scalability requires the use of modern data processing protocols. As such, the calibration communications architecture was designed to satisfy Levels 1–3 of the cyber-physical system requirements, as defined in Lee *et al.* (2014). Plug and play decentralized services allow for messaging patterns where data can be integrated with various systems in a smart factory network. From this interconnection, meaningful information can be inferred from the data, such as robot calibration quality and workcell health. Over the lifecycle of the robotic system, machine performance and historical data can be used to predict future behaviour.

Given these criteria, protocol buffers (https://developers.google.com/protocol-buffers/) were selected as the mechanism for serialized structured data (i.e., packets), and gRPC (https://grpc.io/)

was selected as the transmission mechanism. Figure 2.4 illustrates the control and data flow of this architecture.

Specifically, the external controller that manages the automation of the entire calibration process uses protocol buffer packets and gRPC to communicate across the network with the robotic devices and calibration database. This distributed computing architecture is based on remote procedure call (RPC), a type of a client-server model. With the RPC protocol, a program can request a service or function from another program without having to understand the network layout, such as the relative location of the two programs (i.e., on the same computer or on different devices). The requesting program is a client (e.g., external controller), and the service providing program is the server (e.g., FARO ION controller, KUKA LBR controller, calibration database). In the presented architecture (Figure 2.4), the command packets tell the robot and laser tracker where, when, and how to move. Following the movement command, these devices reply with a response packet that contains the necessary data for calibration (e.g., joint values, Cartesian position). These responses are collected by the external controller and are stored in a database.

The primary advantage of this model is the ability to develop a microservices-style architecture where device control (e.g., robot control, laser tracker control) and data processing become abstract services from the perspective of the external PC controller. A second advantage is the programming language and platform neutral nature of the chosen mechanisms, allowing for straightforward integration with the robot and laser tracker firmware and APIs.

2.4.2 Sphere Measurement and Calibration

The impedance control measurement phase focuses on the collection of data using controlled coupling events. First, whether through manual jogging, offline programming, or hand-guiding, the general location of the three spheres is recorded. In the case of this study, as the KUKA LBR robot is a collaborative robot, these waypoints were taught by quickly hand-guiding the robot to each sphere.

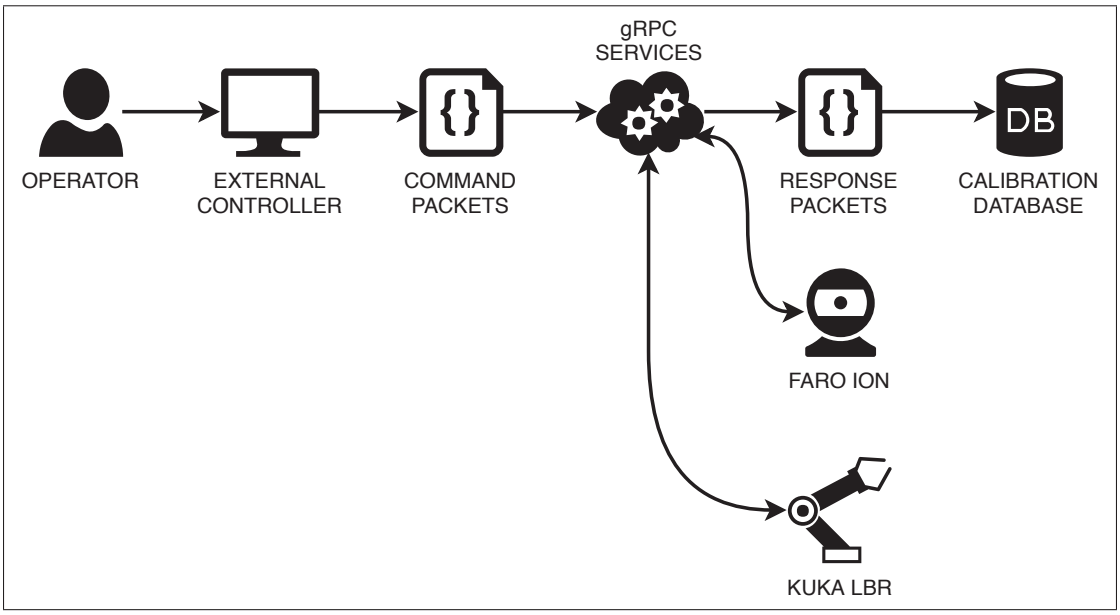

Figure 2.4 Calibration communications framework. Protocol buffers were selected as the mechanism for serialized structured data (i.e., packets). gRPC was selected as the transmission mechanism. During the measurement phase, collected data are uploaded to a database for future processing

Next, the automated calibration process begins and is outlined in Algorithm 2.2. For each sphere, the automated process followed a simple series of steps: contact, record, and decouple. As the robot was under impedance control with a kinematic coupling and magnet, the accuracy of the saved sphere waypoints was not a high priority, and the robot will self-centre upon contact (Figure 2.5), as illustrated in Figure 2.6.

The control architecture exploited the built-in real-time impedance control of the LBR robot as the base control strategy (*KUKA Sunrise.OS* v1.10). The basic model was a virtual spring-damper system with configurable values for stiffness and damping, allowing the KUKA LBR to be highly sensitive and compliant (Figure 2.7). As outlined in Algorithm 2.2, the axis stiffness values, K, were independently updated throughout different phases of the measurement process. The axis damping coefficients (Lehr's damping ratio) remained constant and were set to $\zeta = 0.7$.

Subsequently, the robot will iterate through a series of coupled poses commanded by the external controller. At each pose, the robot held a position under impedance control, allowing the strong

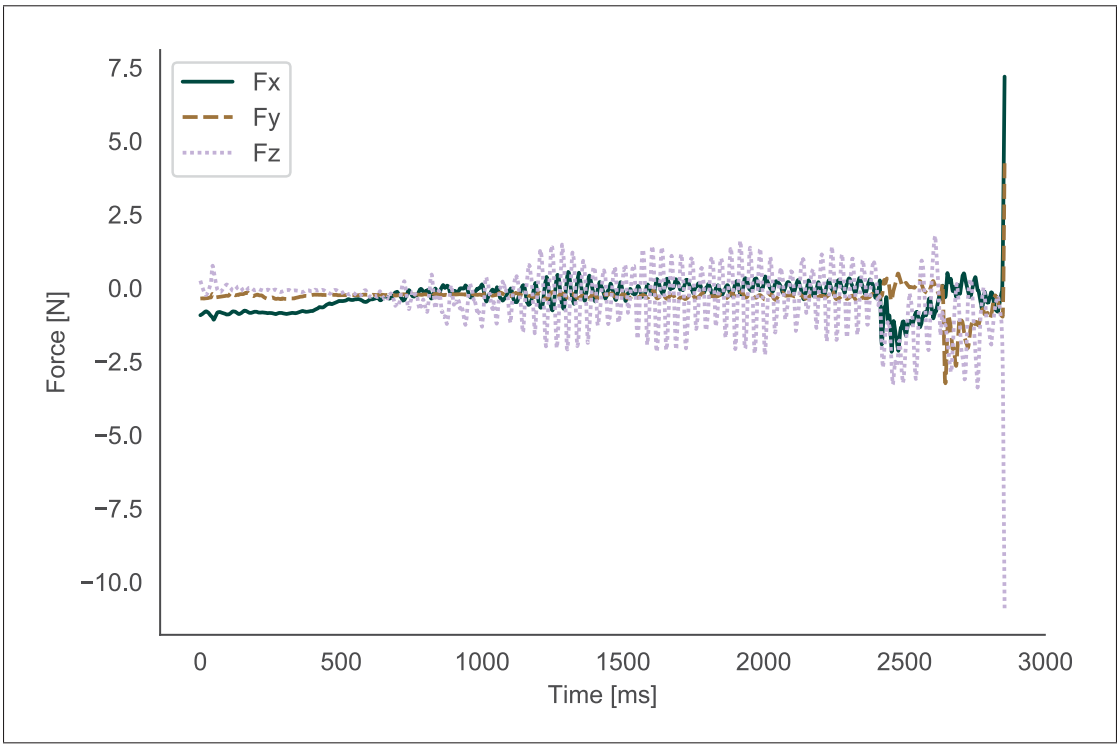

Figure 2.5 Recorded force data during the sphere contact event. The robot controller logs the Cartesian forces applied to the tool frame at a rate of 1000 Hz. As described in Algorithm 2.2, the kinematic coupling event is triggered by a 7.5 N force condition

magnet to maintain proper coupling. The robot controller then logged the robot joint angles and torque values at a rate of 1000 Hz for 1 s, allowing for mechanical stabilization. These data were transmitted back to the external controller and uploaded to a database for future processing. During the calibration process, the mean joint values from this logging process were used. Further mechanical stability analysis demonstrated that even over a period of 30 s, the maximum deviations from the mean values were negligible for the joint angles. Following the completion of the recording phase, the robot decoupled itself from the sphere and repeated the process on the next sphere.

All calibration poses were defined as relative frames with respect the robot base orientation and centred at the sphere. Thus, for each frame, the position of the TCP remained constant, but the tool coupling rotated about the sphere. The frames were generated from a distribution of $-20°$

Algorithm 2.2 Sphere calibration process. All motions are under impedance control. The *approachFrame* may be programmed offline or taught using hand-guiding. The *homeFrame* represents a safe starting point for robot motions

```
robot.move(homeFrame)
for sphere i in [1,2,3] do
  // move above sphere i
  robot.move(approachFrame)

  // prepare for sphere contact
  robot.setXStiffness(50)    // N/m
  robot.setYStiffness(50)    // N/m
  robot.setZStiffness(100)   // N/m
  robot.setForceTrigger(7.5)   // N

  // contact sphere
  while !isContact do
    robot.moveWithVelocity(25)   // mm/s
  end while

  // allow the strong magnet to couple fully
  robot.setXYZStiffness(5)   // N/m
  robot.holdPosition(5)   // s

  // receive poses from external controller and record joints
  while nextPose ← getPose() ! = null do
    // move to pose
    robot.setXYZStiffness(20)   // N/m
    robot.move(nextPose)

    // record joint data over a period of time
    robot.setXYZStiffness(5)   // N/m
    robot.startRecording()
    robot.holdPosition(1)   // s
    robot.stopRecording()

    // transmit joint data to database
    jointData ← robot.getRecording()
    uploadData(jointData)
  end while

  // decouple from sphere
  robot.decouple()
  robot.move(homeFrame)
end for
```

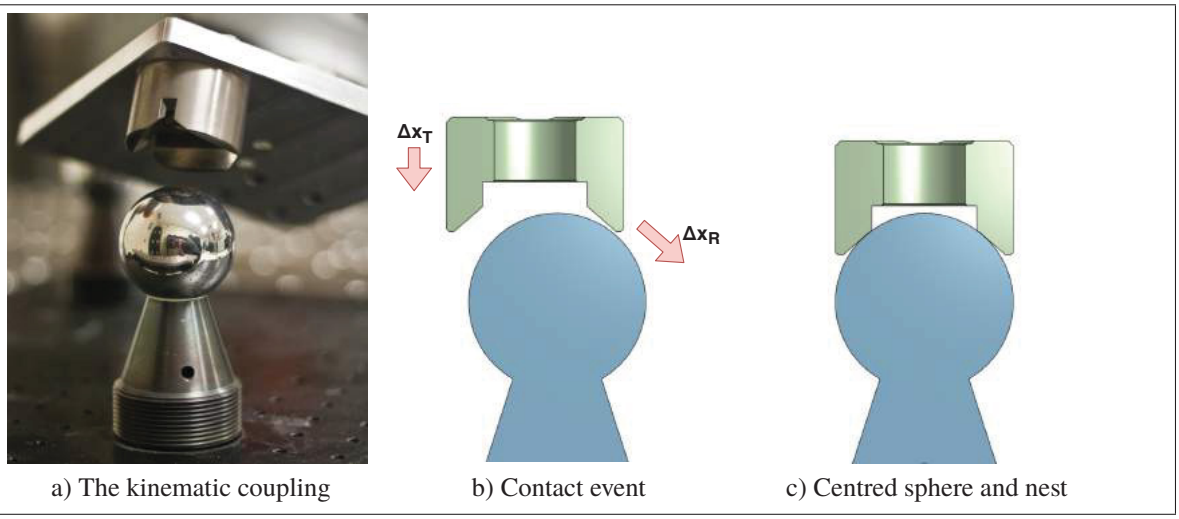

a) The kinematic coupling b) Contact event c) Centred sphere and nest

Figure 2.6 Self-centring of the kinematic coupling under impedance control. (**a**) The combination of the kinematic coupling, a strong magnet, and impedance control allows for a lower accuracy of the saved sphere waypoints. (**b**) Upon contact, a commanded tool motion, Δx_T, will produce a resulting motion, Δx_R, that centres the nest onto the sphere (**c**)

to 20° in 10° steps about all three rotational axes. As the KUKA LBR is a redundant robot with seven axes, each pose was further expanded by altering the redundant α angle (i.e., the angle formed between the elbow of the robot and the vertical base axis) through the same angular distribution (Figure 2.8). A total of 625 poses were generated and recorded on each sphere.

Using the data collected and the optimization algorithms and models presented in Section 2.3.2, the robotic system was calibrated using relative distance errors. As 625 poses per sphere would result in 1,756,875 combination pairs, the dataset was reduced to 100 pairs per sphere (44,850 combinations) through random selection to prevent selection bias. Moreover, mini-batch optimization with a batch size of 32 was used for computational efficiency (Bengio, 2012).

2.4.3 Laser Measurement and Calibration

To validate the quality of the impedance control calibration method properly, traditional calibration with a FARO ION laser tracker (uncertainty of ±49 µm) and spherically-mounted retro-reflectors (SMR) was performed (Figure 2.9). Following the procedure outlined in

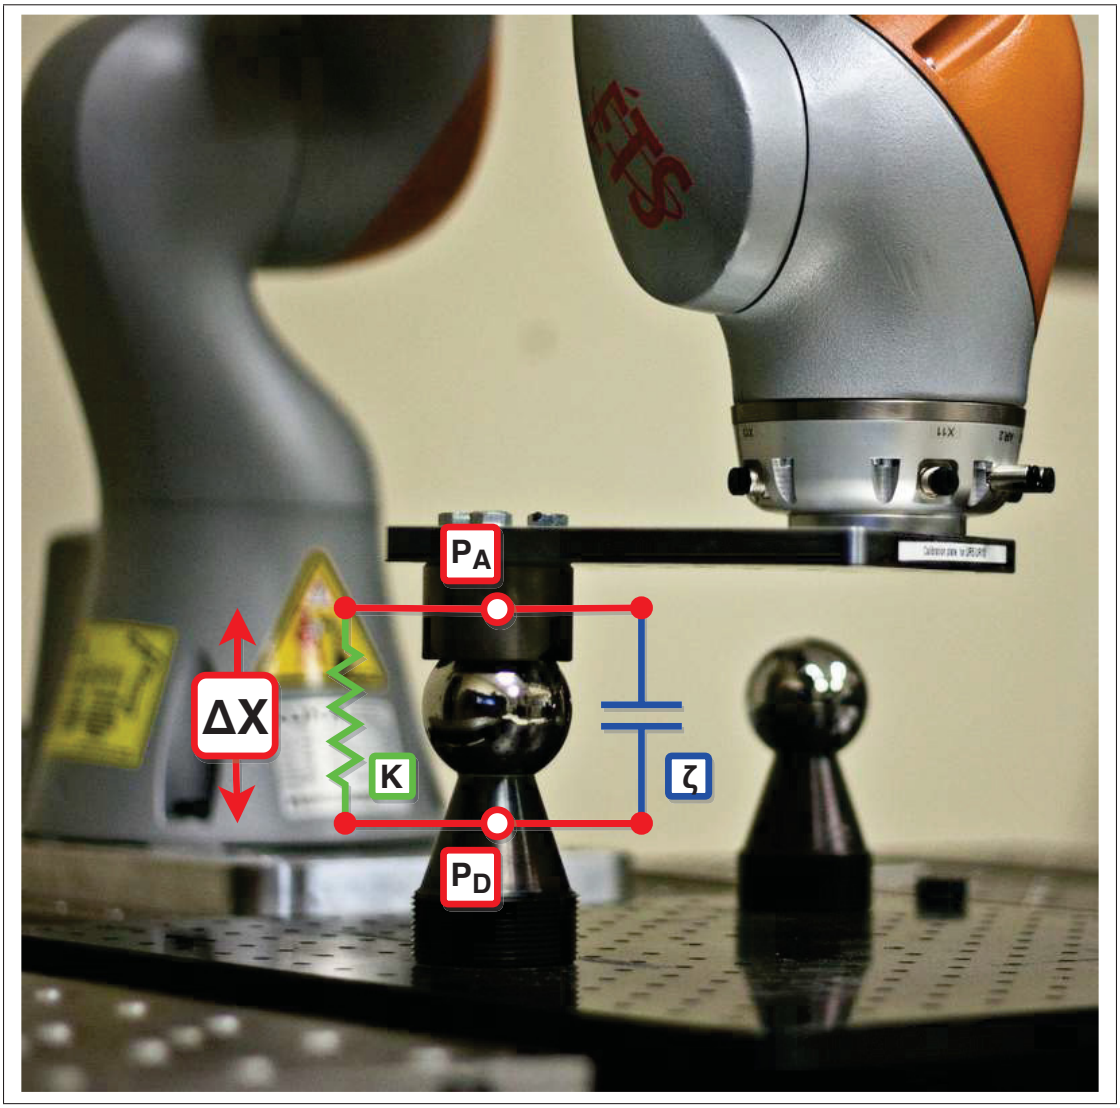

Figure 2.7 The built-in real-time impedance control of the KUKA LBR iiwa robot. The model is a virtual spring-damper system with configurable values for stiffness (K) and damping (ζ), allowing the KUKA LBR to be highly sensitive and compliant. The resulting interaction forces are calculated from the displacement (ΔX) between the desired position (P_D) and the actual position (P_A) in each Cartesian axis

Nubiola & Bonev (2013), 284 uniformly-distributed joint configurations were recorded within the workspace of the robot for calibration (Table 2.3). An additional dataset of 125 configurations was recorded for validation (Table 2.4). The 3D distribution of these points may be seen in Figure 2.10.

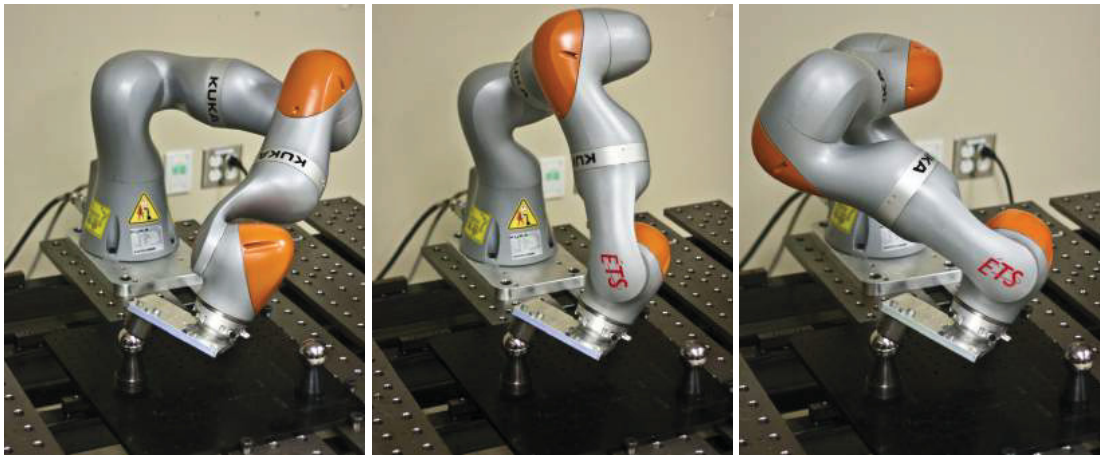

a) Robot pose with the redundant α angle parameter set to $-20°$

b) Robot pose with the redundant α angle parameter set to $0°$

c) Robot pose with the redundant α angle parameter set to $20°$

Figure 2.8 Expansion of coupled poses through redundant joint configurations. All three pictured configurations represent the same Cartesian pose. The KUKA LBR is a redundant robot with seven axes, and a non-singular pose has infinite joint configurations. The redundant α angle parameter is defined as the angle formed between the elbow of the robot and the vertical base axis

Table 2.3 Statistics of the 284 configurations used for laser calibration

	mean (mm)	**std (mm)**	**min (mm)**	**max (mm)**
x	525.872	160.031	−7.304	947.042
y	−69.712	220.397	−560.339	452.969
z	650.021	186.840	198.322	1197.046

Table 2.4 Statistics of the 125 configurations used for laser validation

	mean (mm)	**std (mm)**	**min (mm)**	**max (mm)**
x	501.986	211.902	−24.214	995.460
y	−7.856	252.262	−578.336	523.221
z	750.930	227.457	137.178	1135.052

a) A FARO ION laser tracker (uncertainty of ±49 µm)

b) Spherically-mounted retro-reflectors (SMR) mounted to the same tooling described in Section 2.2.2

Figure 2.9 Traditional laser tracker calibration equipment

To calibrate the robot base frame and the tool frame, 150 and 153 additional measurements were recorded, respectively. The base frame measurements were generated by independently rotating only Joints 1 and 2 to produce two 3D arc trajectories whose intersection gave the location of the base frame. Similarly, the tool frame measurements were generated by independently rotating only Joints 5, 6, and 7. The independently-measured base and tool frames were used not only during the laser calibration, but also during validation, acting as control variables for the model comparison. As a result, the calibrated robot models were isolated for comparison.

As in Section 2.3.2, the identified kinematic parameters were independently optimized using absolute accuracy errors and the Nelder-Mead method. For absolute position optimization, using Equations (2.1) to (2.3) and (2.8), the absolute position error ($E_{abs_i}(q)$) per measurement configuration is described as follows:

$$P_{measure} = \begin{bmatrix} x_{measure} \\ y_{measure} \\ z_{measure} \end{bmatrix} \tag{2.12}$$

$$E_{abs_i}(q) = |(P_{measure_i} - P_{TCP}(q))| \tag{2.13}$$

where i is the ith measured configuration measured by the laser tracker.

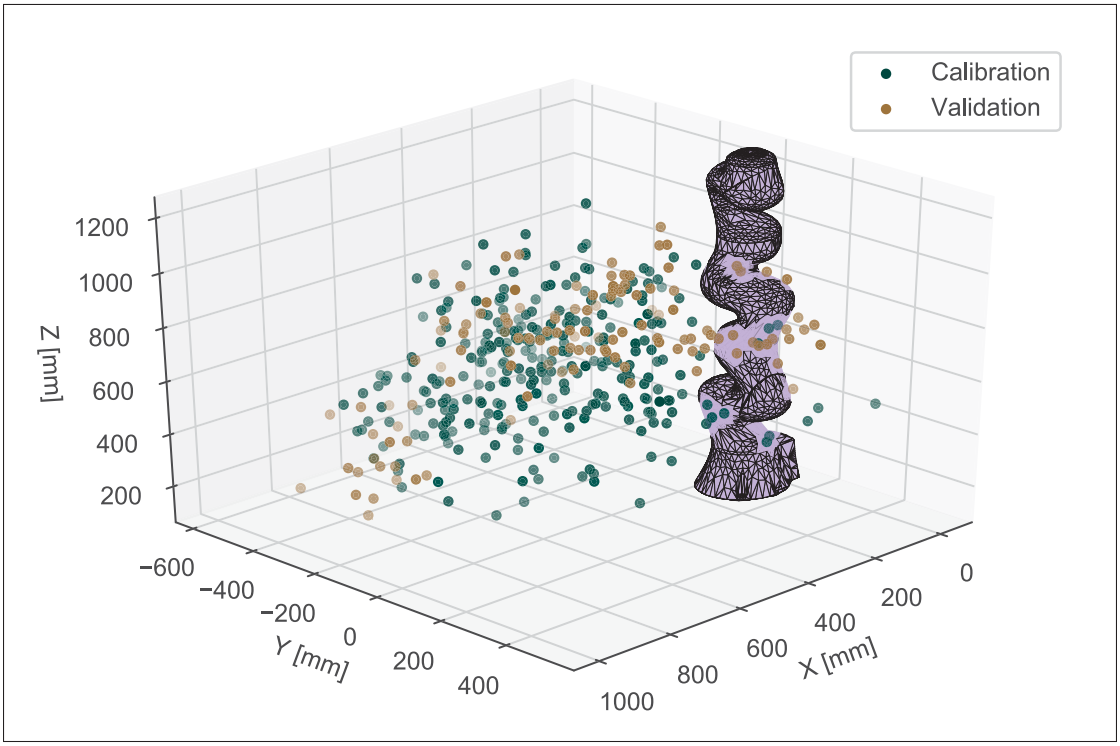

Figure 2.10 3D distribution of calibration and validation points. Two hundred eighty four configurations were used for laser calibration. One hundred twenty five configurations were used for laser validation

2.5 Results

Following impedance control and laser calibrations, the absolute accuracy of the calibrated models was compared against the nominal model using the 125 measurements mentioned in Section 2.4.3 and Equation (2.13). Error distributions between the nominal model and the calibrated models are visualized in Figures 2.11 and 2.12. As shown in Figure 2.11, both calibration methods offered significant accuracy improvements over the nominal model. Figure 2.12 offers a more detailed comparison between traditional laser tracker calibration and

the impedance control method. Error statistics of all the models are summarized in Table 2.5. The calibrated kinematic model of the KUKA LBR iiwa 7 using impedance control may be seen in Table 2.6.

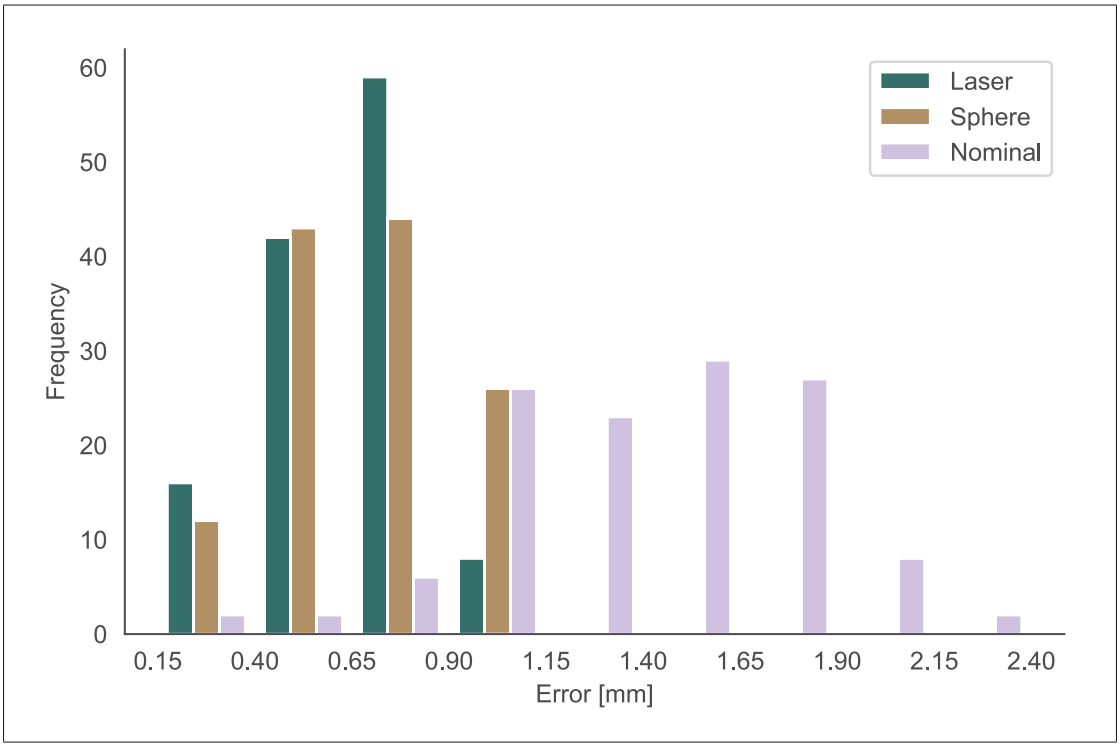

Figure 2.11 Histogram of validation errors comparing the nominal and calibrated models

Table 2.5 Summary of validation errors between the nominal KUKA LBR kinematic model, sphere calibration with impedance control, and traditional laser tracker calibration

	mean (mm)	**std (mm)**	**min (mm)**	**max (mm)**
Laser	0.645	0.192	0.168	0.932
Sphere	0.677	0.206	0.225	0.990
Nominal	1.403	0.392	0.325	2.381

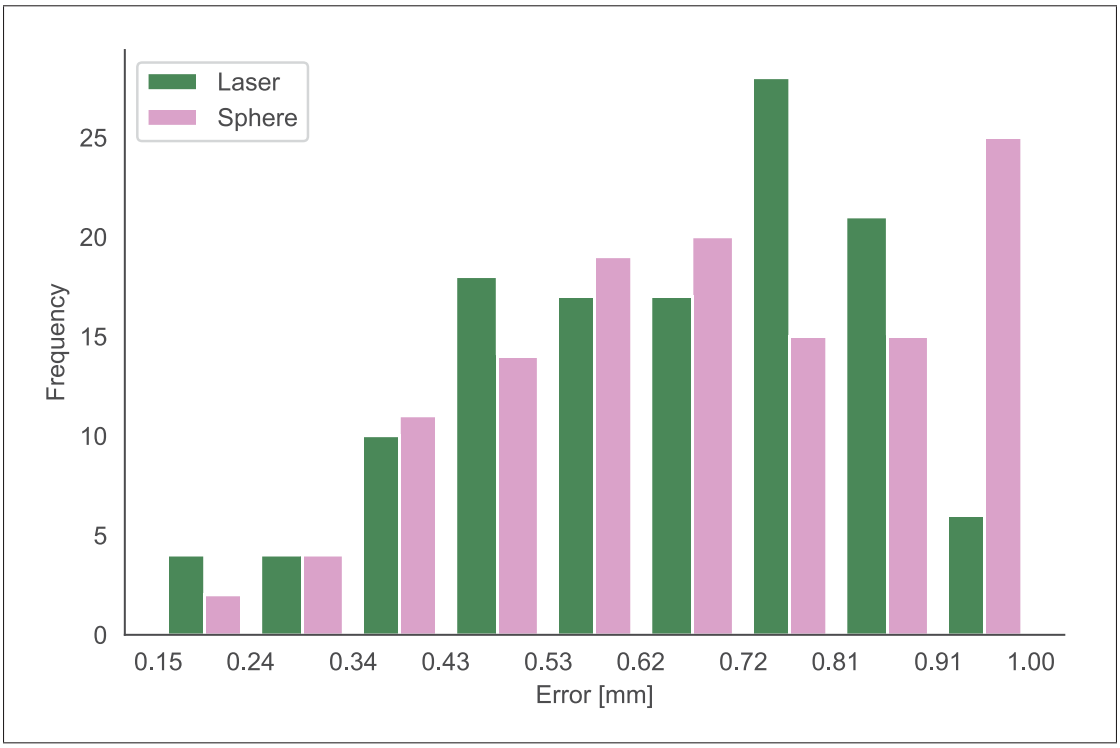

Figure 2.12 Histogram of validation errors comparing the calibrated models

Table 2.6 Calibrated kinematic model of the KUKA LBR iiwa 7 using the presented impedance control approach

i	α (rad)	a (mm)	θ (rad)	d (mm)
0	0.000000	0.000000	0.000000	340.000000
1	−1.570759	0.000061	0.000062	0.000062
2	1.571575	0.000056	−0.000490	400.574635
3	1.569842	0.000061	0.000061	0.000063
4	−1.570623	0.000062	0.000059	400.629881
5	−1.571451	0.000062	0.000061	0.000000
6	1.571136	0.000000	0.000000	125.505736

2.6 Discussion

This work presented an impedance control approach to robot calibration and compared the novel method to traditional calibration with a laser tracker. The calibration tooling has the advantage of being low-cost, robust, and simple to use, with little start-up time required. As shown through absolute accuracy validation, this calibration approach can improve the accuracy of a KUKA LBR iiwa 7 R800 collaborative robot upwards of 58.4% to a maximum error of 0.990 mm when compared to the nominal model.

It should be noted that the sheer simplicity of the tooling used in this study makes for an interesting comparison to typical calibration equipment, especially when considering the cost difference. While the traditional laser tracker calibration could reduce the maximum error to 0.932 mm, the tooling presented in this study is affordable and easily available as off-the-shelf components from mechanical tooling vendors, while a laser tracker is generally quite expensive. Furthermore, the robot modelling software (Nadeau, 2019) and the communication packages are open-source and available for both industrial and academic use. It thus becomes an engineering economics issue to justify the cost difference for an additional percentage point of improvement.

With respect to previous studies, the present study advances the literature in several distinct ways. First, while the nominal errors of the KUKA LBR agreed with Besset *et al.* (2016), the presented calibration approach has been validated with 12.5-times the number of validation points, offering a more robust and confident examination of calibration quality. Compared to the self-calibration approach presented in Gaudreault *et al.* (2018), while the resulting accuracies due to the methods cannot be directly compared (two different robots), the difference between their method and the (same) laser tracker was relatively equivalent. Yet, the method presented here is even less expensive and more durable for industrial environments. While often undervalued, durability and robustness are key considerations, as repairs and workcell downtime can be expensive.

Although this study used a redundant robot, this method can easily be applied to less sophisticated robots with no major differences in procedure. While redundant motion allows many joint configurations for a single Cartesian pose, the results of this process with a non-redundant robot

will produce a dataset without redundant configurations. Simply put, redundant motion produces more joint configurations with less tool movement.

From a manufacturing and operations perspective, the authors of Zhang & Van Luttervelt (2011) and Zhang *et al.* (2017) described a set of guidelines for a resilient manufacturing system, whereby the system is fault tolerant and capable of recovering from failure. In the case of the presented calibration methodology, failure may occur during any of the steps outlined in Section 2.4.2. During the contact, coupling, and motion steps, force limits are set to prevent damage to the tooling, and the recovery process is to simply decouple, move away, and restart the coupling procedure. Returning to a known stable position is not required and slows the overall process. During the recording phase, the system follows redundancy guidelines and records more poses than necessary for calibration optimization. Further reliability can be achieved through the robot communication and control framework described in Section 2.4.1, where data collection can be used to develop continuous improvement initiatives. This communication framework integration follows Guidelines II–IV presented in Zhang & Van Luttervelt (2011). In an ideal integration, the presented calibration tooling would become part of the robot workcell, allowing the robot to self-validate and calibrate. As noted in Lee *et al.* (2014), with a connected fleet of machines and an awareness of machine condition, predictive maintenance and workload balancing become possible and maximize machine performance and organization transparency.

CHAPTER 3

EVOLUTIONARILY OPTIMIZED PHYSICAL HUMAN-ROBOT CONTACT DISTINCTION IN AN UNSTRUCTURED ENVIRONMENT

Collisions and contact events are a natural by-product of applications involving pHRI in unstructured environments. Accordingly, the ability to classify contact events would allow for more complex robot reaction behaviour. This chapter presents an experimentally tested, machine learning framework for robot contact classification using the built-in sensor data of a KUKA LBR iiwa robot. Unlike previous studies, this classification does not discern between intended vs. unintended contact scenarios, but rather classifies *what* was involved in the contact event. The framework can perfectly differentiate between high-level *hard*, *free space*, and *human* motion events at low speeds with an average prediction time of 2.71 ms. Through evolutionary optimization of the classification pipeline, the classifier can discern different ISO/TS 15066:2016 specific body areas along a human-model leg with 89.37 % accuracy. Altogether, this contact distinction framework allows for more complex reaction strategies and tailored robot behaviour during pHRI.

3.1 Introduction

With the recent trend of collaborative robots in industrial environments, human operators no longer need to be physically protected and separated from their robot co-worker. Nevertheless, collisions and contact events are a natural by-product of applications involving pHRI in unstructured environments. Whether the human-contact is part of the application (e.g., robotic medical ultrasound) or accidental, the initial human-robot contact event holds important information for application and safety control.

It is important to note that safe pHRI is composed of both intrinsic (i.e., inherent) and extrinsic (i.e., application design) safety. A successful collaborative application utilizes the strengths of both the robotic system and the human operator, forming a flexible team environment. Existing

safety systems often use collision avoidance (Morato *et al.*, 2014) or intrinsic safety through robot design (Kock *et al.*, 2011), but these strategies break down when a collision event occurs.

Alternatively, collision detection and contact distinction allow the system to understand the current state of the collaborative application environment. As a result, reaction strategies may be developed to improve post-contact event safety. Examples of collision detection may be seen in Haddadin *et al.* (2008) and Geravand *et al.* (2013), where the authors were successful in detecting and reacting to collisions using external torque measurements and joint motor current, respectively. While several reaction strategies were proposed, they were mainly designed to simply abandon the current robot task, "flee" from the disturbance, or switch to a zero-gravity mode.

As for contact distinction, Cho *et al.* (2012) found that the rate of change of the joint torques can be used to discern intended (e.g., hand guiding) vs. unintended (e.g., collision) pHRI, but the experiments were limited to a 2-DOF manipulator. More recently in Kouris *et al.* (2016) and Kouris *et al.* (2018), the authors used frequency domain segmentation and monitoring for contact distinction. While the study successfully distinguished between intended vs. unintended pHRI, the experiments were performed at relatively high speeds (200 mm/s to 1000 mm/s), where unintended collisions produce a significantly different signal than intended interaction. Noticeably, the authors remarked that low-speed collisions cannot be noticed by their distinction method. As a comparison, KUKA defines speeds lower than 250 mm/s as safe speeds for robot *teach mode*.

From a machine learning perspective, Golz *et al.* (2015) developed a set of handcrafted features from joint torque sensor data for contact distinction supervised learning with a support vector machine. However, the study displays several limitations, such as a limited description of the classifier architecture and a small sample size that could lead to overfitted test results. Lu *et al.* (2005) presented a neural network approach to detect the collisions, but the framework is only tested with virtual collisions.

These studies focused solely on intended vs. unintended pHRI. When a contact event occurs, it is difficult for the system to understand *what* it encountered, such as differentiating between hard obstacles (e.g., steel table) or a human, or even specific human body regions. Such knowledge regarding the current environment state would benefit robot decision making and finite-state machine designs to allow for more advanced robot behaviour and reactions. As noted in Lu *et al.* (2005), stopping the robot due to a contact event is not necessarily an optimal method for human safety and risk avoidance. If significant local information can be extracted from the event, it becomes possible to control the robot such that the task does not need to be abandoned.

This work presents an experimentally tested, machine learning framework for robot contact classification. Unlike previous studies, this classification does not discern between intended vs. unintended contact scenarios, but rather classifies *what* was involved in the contact event. Through evolutionary optimization of the classification pipeline, the framework can differentiate between high-level motion events and even discern different ISO/TS 15066:2016 specific body areas along a human-model leg. Moreover, in contrast to Kouris *et al.* (2018), this framework is capable of contact distinction at low speeds with a smaller time window of data. Like Geravand *et al.* (2013), Kouris *et al.* (2016), and Kouris *et al.* (2018), no *a priori* information about the robot or environment is used in the classifier. Finally, only a single force signal is used, allowing this technique to be used simply with a 1D external force sensor on non-collaborative robots.

The study is structured as follows: Section 3.2 defines the application context and problem, and theoretically and experimentally resolves the transient energies involved in a clamping contact scenario. Section 3.3, describes the process of building a classification pipeline for contact distinction and presents the results. Finally, the implications of the solution are discussed in Section 3.4.

3.2 Problem Formulation

Typical contact detection with a collaborative robot involves setting force and energy limits for a motion. Any force vectors (magnitude and direction) detected that surpass the given limits will

trigger a secondary process, such as stopping the robot or logging the event. While for many robotic tasks stopping the robot is a valid reaction strategy for safety, discarding the current task upon contact detection is inefficient. If reaction strategies are built around understanding the current environment and contact state, then adaptable reaction strategies may be developed continue tasks without loss of progress.

Robot-assisted medical ultrasound is an example of such a task where simply stopping the robot upon contact detection may not be an appropriate reaction strategy. As discussed in Gilbertson & Anthony (2015), because ultrasound is used to image soft, deformable tissues, the image quality and repeatability are degraded by human performance due to force control. Improvements to freehand ultrasound force control could lead to improvements in medical care. Thus, the robot should have an awareness of body contact location to properly plan force-controlled trajectories along the human body using the imaging probe. This is especially true for remote ultrasound systems, such as the one presented in Koizumi *et al.* (2009), where safety and manipulability are important elements to consider when operating a remote medical system through a communication network. As an example ultrasound scenario for this work, the healthcare professional would teach lower limb imaging targets to the robot system through hand guiding. In this situation, the human and robot are co-workers, with the human offering intuition and application knowledge, and the robot offering precision and accuracy. The robot would then go to a selected target and contact the patient's leg at a desired force (Fig. 3.1). This results in a clamping contact situation as defined by ISO/TS 15066:2016.

Beyond healthcare, equivalent scenarios may also be seen in manufacturing and assembly environments (Rosenstrauch & Krüger, 2017). The following subsections will decompose and evaluate this contact scenario and experimental setup using ISO/TS 15066:2016 protocols to develop a better understanding of the forces involved.

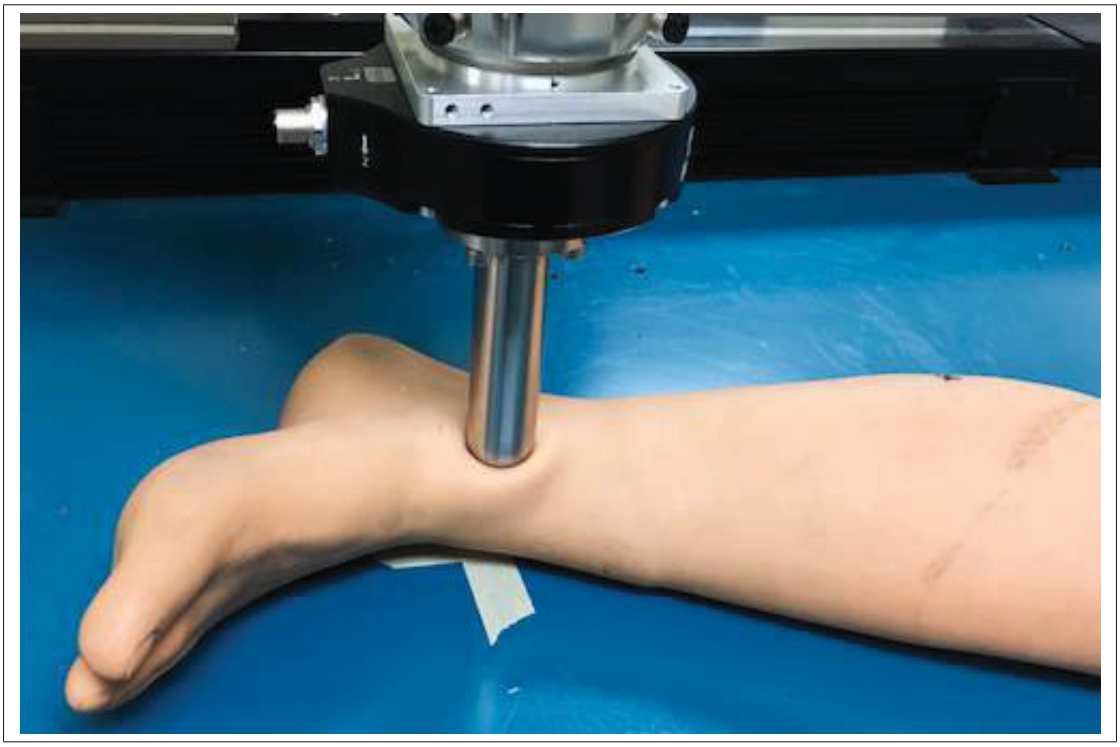

Figure 3.1 Ankle contact event

3.2.1 Experimental Setup

The experimental setup may be seen in Fig. 3.2 where a KUKA LBR iiwa 7 R800 robot was equipped with a rounded probe tool. The probe tool was designed to approximate the type of tool referenced in Annex A of ISO/TS 15066:2016. The LBR collaborative robot has built-in torque sensors at every joint, allowing for calculations of Cartesian forces acting on the flange and tool. A Robotiq FT150 force-torque sensor was used to externally validate the forces measured by the robot prior to the experiment and was disconnected for the actual scenarios to avoid additional forces from the hanging wire of the sensor. The mass data of the tool assembly was calibrated using the built-in calibration procedure of the robot to remove the tool load from the force measurements. A mannequin leg with an internal skeletal structure was used to simulate the patient.

The robot was programmed to perform a linear downwards motion until the force condition was triggered (Fig. 3.1) at which point the event was logged and the robot retracted along the same path. For *human* contact events, four lower limb targets were selected: *thigh*, *knee*, *calf*, and *ankle*. These targets were selected to correspond with the body model outlined in ISO/TS 15066:2016. The mannequin leg was then removed for *hard* contact events where the force condition was triggered when the probe collided with the steel table. Finally, *free space* motion events without contact was performed as a control scenario where the downwards motion was programmed to finish just above the base to avoid the contact event.

Each scenario was performed with random Cartesian speeds (50 mm/s to 250 mm/s) and force conditions (5 N to 30 N). The speed range was chosen to be within the robot manufacturer's collaborative guidelines and the force range was chosen to be representative of medical ultrasound (Mathiassen *et al.*, 2016). 1580 sessions were performed under Cartesian impedance control with random stiffness in the contact axis (0 N/m to 5000 N/m; $damping = 0.7$). The remaining sessions were performed under standard position control. In total, 3806 sessions were performed. For all sessions, the internal data recorder of the LBR was used to log the robot data at a rate of 1000 Hz.

3.2.2 Theoretical Contact

During collaborative application design, the energy of motions needs to be calculated to ensure that safe limits are respected. According to ISO/TS 15066:2016, the maximum allowed quasi-static force is 220 N for the *thighs and knees* region and 130 N for the *lower legs* region. The transient energy for a motion may be calculated from Annex A of ISO/TS 15066:2016, as shown below.

Given the robot moving mass, M, and the mass of the effective payload (e.g., tooling), m_L, the effective robot mass, m_R, may be calculated by:

$$m_R = \frac{M}{2} + m_L \tag{3.1}$$

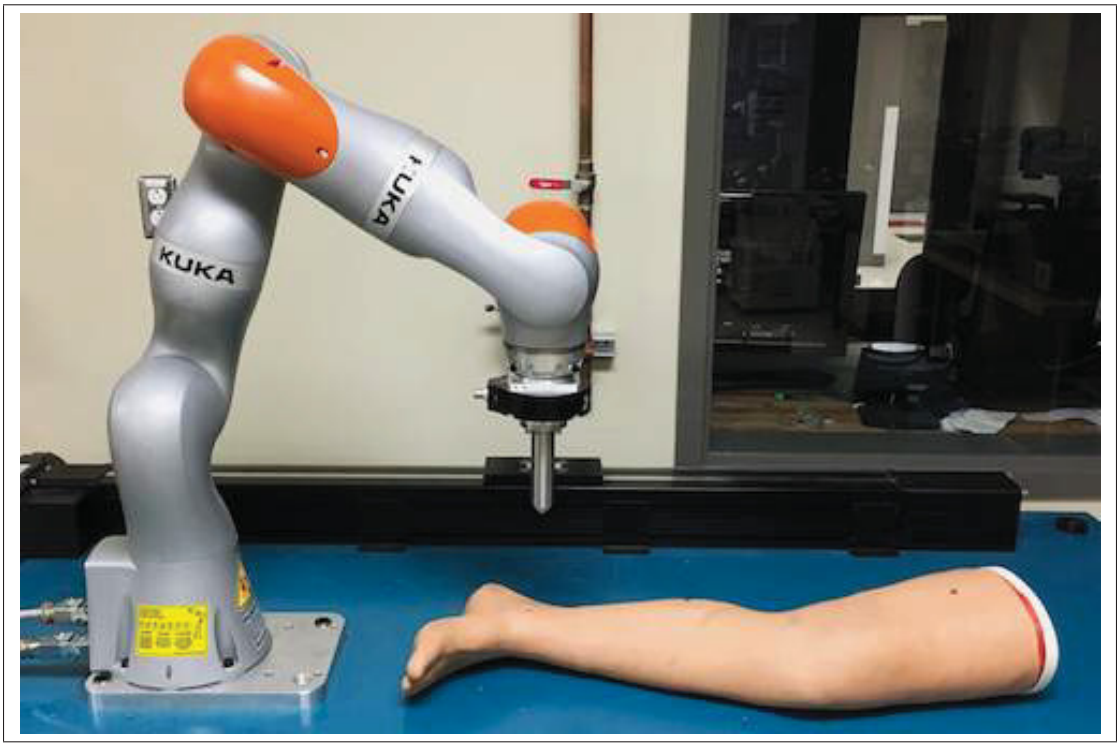

Figure 3.2 The KUKA LBR iiwa robot was equipped with a probe tool and a Robotiq FT150 sensor for external force measurement validation. A mannequin leg with internal skeletal structure was used to simulate a human subject. The probe tip has a radius of 5 mm

As the exact mass distribution of a robot may be proprietary information, the robot moving mass may be estimated as a percentage of the total mass of the robot. Since the contact event is simplified as a two-body system, the effective mass of a given human body region, m_H, is retrieved from Table A.3 of ISO/TS 15066:2016. For both the *thighs and knees* region and the *lower legs* region, the effective human mass is 75 kg. The reduced mass of the two-body system, μ, may now be calculated as:

$$\mu = \left(\frac{1}{m_H} + \frac{1}{m_R} \right)^{-1} \tag{3.2}$$

Finally, the transferred energy due to an inelastic contact, E, event is given by:

$$E = \frac{\mu v^2}{2} \tag{3.3}$$

where v is the relative velocity between the robot and human contact bodies.

Given the above equations and the calibrated tooling mass data ($m_L = 1.96\,\text{kg}$), the transient energy transfer for this study may be calculated. Assuming an inelastic contact scenario, the maximum allowed transferred energy according to ISO/TS 15066:2016 is 1.9 J for the *thighs and knees* region and 0.52 J for the *lower legs* region. Even the most conservative calculation, with a robot moving mass equal to the total mass of the LBR robot (23.9 kg) and a velocity of 250 mm/s, gives energy values (0.37 J) below the lowest of the two thresholds, implying that the proposed contact scenario is theoretically safe and collaborative.

3.2.3 Experimental Contact

Reaction and stopping time are well known to robotic integrators and are generally considered in the risk assessments of a specific application. Figure 3.3 demonstrates a sample of the results of the experimental contact scenarios. Even with a simple scenario and while the robot force condition was triggered during each session, all motions at all speeds surpassed the desired force limit, due to the reaction time of the system. Interestingly, the reaction time of the LBR for software level motion force conditions is not characterized. Instead, safety monitoring functions that result in an emergency stop (e.g., velocity monitoring, axis range monitoring) are characterized as having reaction times between 22 ms to 154 ms (according to the manufacturer's specifications). Unfortunately, an emergency stop reaction abandons the current task.

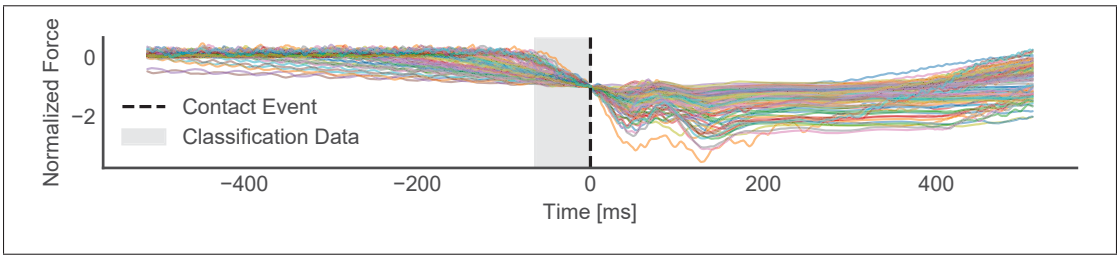

Figure 3.3 Measured F_z forces normalized with respect to the force limit of the session during *calf* contact. The time domain is centred about the force condition trigger. The shaded region highlights the data used for contact classification (64 ms)

3.3 Contact Classification

Given the results of Fig. 3.3 and that contact events may occur in unstructured collaborative environments, the contact event itself holds important information for application and safety control, highlighting the importance of contact classification. Instead of abandoning the current task or performing an emergency stop, complex robot behaviour and decision making can react in a way that will maximize the likelihood of successful task completion. Such a decision-making process would have different reaction strategies for different types of contact, such as the *human*-contact and *hard*-contact events. As part of this post-contact recovery strategy, classification of the contact event allows the system to understand the current state of the application environment.

This section details the process of constructing a machine learning pipeline for human-contact classification. Using the experimental setup and data labels described in Section 3.2.1, the collected data was used in a supervised learning pipeline for contact classification. This pipeline was built using *Python* v3.6.1, *scikit-learn* v0.19.1 (Pedregosa *et al.*, 2011), and *SciPy* v1.1.0 (Jones *et al.*, 2001). All classifier training was performed with a training set of 3044 sessions, representing 80 % of the recorded data.

3.3.1 Dimensionality Reduction

All data analysis and supervised learning was performed on a time window that preceded the force trigger. This self-imposed limitation replicates a real-time robotic classification system, as this data would be accessible at the moment of a force limit trigger (Fig. 3.3).

As the force data was recorded at a high frequency, variance analysis and dimensionality reduction were the first steps in the classification pipeline. High data dimensionality leads to high computational cost and potential overfitting. Dimensionality reduction preserves the most relevant information to train accurate models. Noticeably, as seen in Fig. 3.4, most of the variance found in the raw force data is located during the timespan immediately preceding the contact event. Accordingly, the first dimensionality reduction step isolated a data window of 64 ms. For *free space* motion, the time window preceded the instant of lowest Z position in the up-down motion.

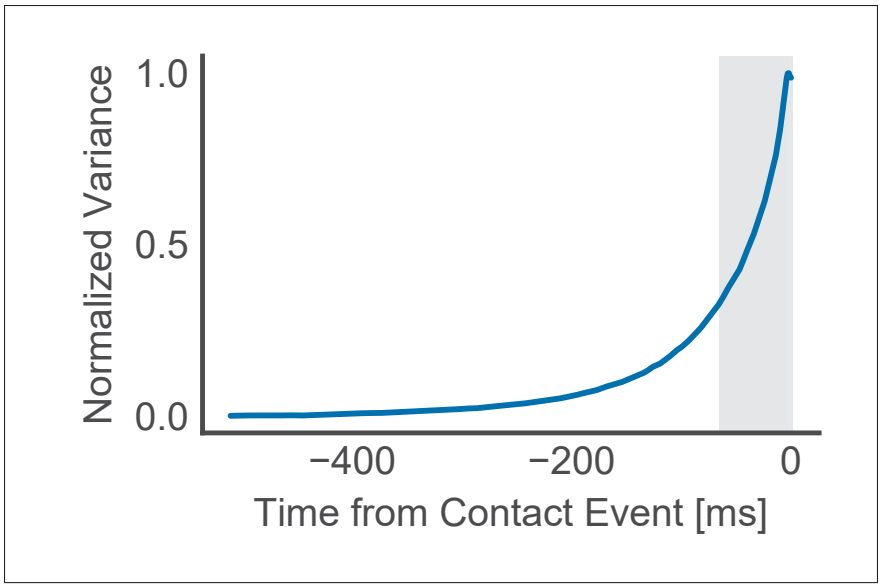

Figure 3.4 Time series variance across all force data samples. The shaded region represents the 64 ms used in the supervised learning

As noted in Kouris *et al.* (2016) and Kouris *et al.* (2018), the frequency domain also contains relevant information to contact classification. When the previously extracted 64 ms of data is processed by a fast Fourier transform, the lowest frequency components appear to contain the most variance (Fig. 3.5).

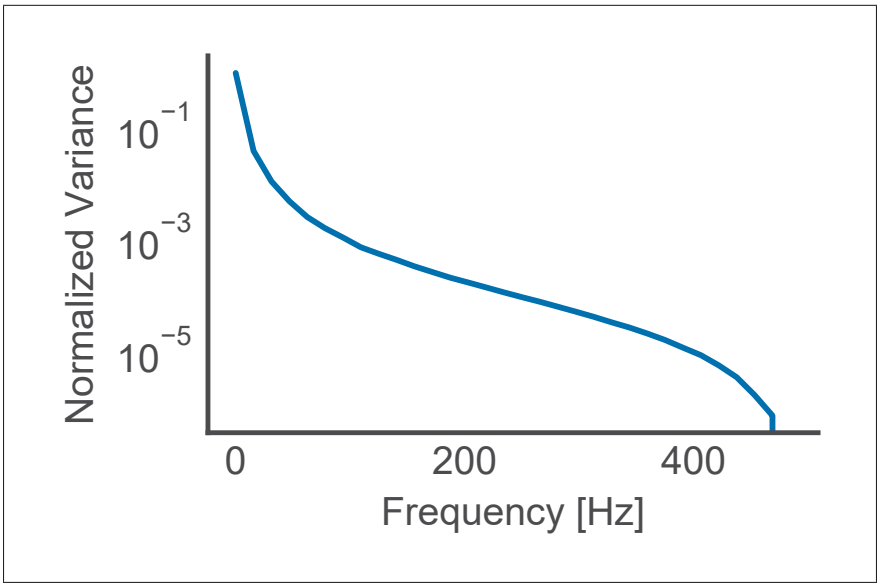

Figure 3.5 Frequency domain variance of the 64 ms preceding the contact event

Further processing with principal component analysis (PCA) on both the normalized force and frequency data found that 99.88 % of force data variance may be found with six components. Similarly, just four frequency domain components represent 99.97 % of variance.

3.3.2 Handcrafted Pipeline

Supervised learning develops a solution model from a given set of inputs (i.e., features) and outputs (i.e., labels). Given the extracted features presented in Section 3.3.1 and the labels presented in Section 3.2.1, a handcrafted pipeline was developed to produce a simple but effective classifier.

This simple pipeline, illustrated in Fig. 3.6, combines the dimensionality reduction and feature extraction steps previously outlined with a k-nearest neighbours (k-NN) classification algorithm. The pipeline was designed for simplicity and allows for the parallel processing of force and frequency data.

The pipeline steps are summarized as follows: First, the raw force data is truncated to the 64 ms preceding the contact event. Next, the pipeline splits between force (blue) and frequency (purple) data processing. PCA decomposition is then used to further reduce the dimensionality of the data. Prior to consolidation, the data is scaled such that the value of each feature is between 0 and 1. Finally, classification is provided by the k-NN algorithm. A sample is classified by a majority vote of its neighbours and is assigned to the label most common among its k nearest neighbours ($k = 3$). This technique is considered one of the simplest algorithms in machine learning. All classifier hyperparameters were experimentally tuned using the training data.

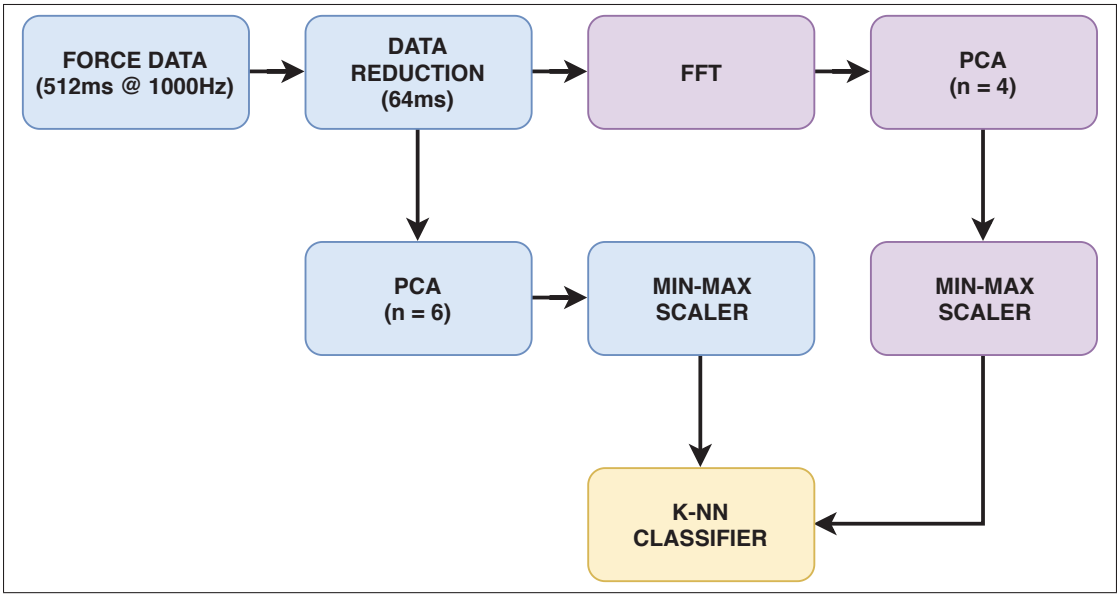

Figure 3.6 Handcrafted classification pipeline using fast Fourier transform, principal component analysis (PCA), and k-nearest neighbours (k-NN). The pipeline was designed for simplicity, with parallel processing of force data (blue) and frequency data (purple) for contact classification

3.3.2.1 Results

All classification results were generated from a reserved test set of 762 sessions, representing 20 % of the recorded data. The classification results for the handcrafted pipeline are visualized using a confusion matrix (Fig. 3.7). A confusion matrix C is such that $C_{i,j}$ is equal to the percent of observations known to be in group i and predicted to be in group j. The pipeline achieved an overall accuracy of 77.82 % across all labels. Interestingly, *thigh* contact events have the poorest outcome and are often mistakenly classified as *calf* contact. The inverse relationship also exists. This is most likely due to the similar tissue softness of the body regions. Moreover, the pipeline has perfect accuracy with respect to distinction between *hard*, *free space*, and *human* contact events.

A more explicit classification report may be seen in Table 3.1. *Precision* (P) is defined as the number of true positives (T_p) with respect to the sum of true positives and false positives (F_p):

$$P = \frac{T_p}{T_p + F_p} \tag{3.4}$$

Recall (R) is defined as the number of true positives with respect to the sum of true positives and false negatives (F_n):

$$R = \frac{T_p}{T_p + F_n} \tag{3.5}$$

Simply put, *precision* is the ability to correctly apply a label while *recall* is the ability to find all samples belonging to a label. *F1 Score* is the mean of *precision* and *recall*. *Support* is the number of occurrences of each label.

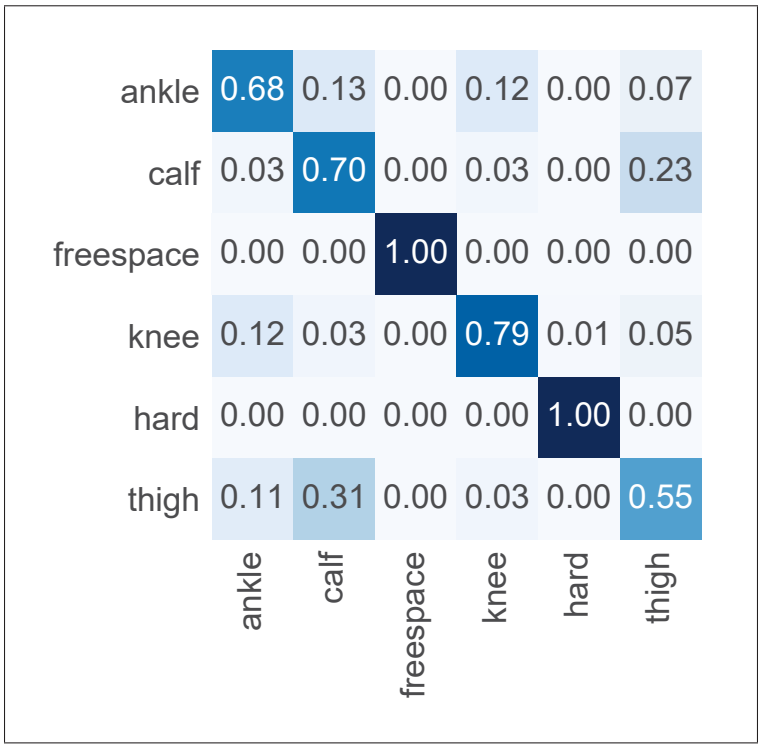

Figure 3.7 Handcrafted pipeline confusion matrix

Table 3.1 Classification summary for the handcrafted pipeline

Label	Precision (%)	Recall (%)	F1 Score (%)	Support (#)
knee	80	79	79	122
free space	100	100	100	86
ankle	71	68	69	122
table	99	100	100	146
thigh	57	55	56	124
calf	66	70	68	162

3.3.3 Evolutionary Pipeline

While the handcrafted classification pipeline presented in Section 3.3.2 provided perfect accuracy for high-level distinction of *hard* vs. *free space* vs. *human* contact events, the intra-*human* accuracy was lacking. As an additional step, the classification pipeline was optimized using

evolutionary genetic algorithms. Specifically, the *TPOT* (v0.9.3) (Olson *et al.*, 2016) and *Distributed Evolutionary Algorithms in Python* (*DEAP*; v1.2.2) (Fortin *et al.*, 2012) frameworks were used to develop a more effective classifier at the cost of complexity.

Given the pipeline nodes provided by *scikit-learn*, the *TPOT* framework allowed for the automation of exploring thousands of possible pipelines to optimize the classifier architecture. To automatically generate these pipelines, evolutionary genetic algorithms provided by *DEAP* built a population of 100 candidate pipelines to optimize for classification accuracy. As desscribed by Olson *et al.* (2016), the evolution of machine learning pipelines can provide a significant improvement over a basic machine learning analysis while requiring little to no input nor prior knowledge from the user. Moreover, overly complex pipelines can be prevented by integrating Pareto optimization, which produces compact pipelines without reduced classification accuracy. The sequence of pipeline operators and hyperparameters were evolved over 100 generations.

The resulting optimized pipeline is visualized in Fig. 3.8. Architecturally, the pipeline is composed of the original handcrafted design (without the k-NN classifier) and an evolutionarily optimized segment. Experimental tests with the training dataset found that this combination provided the best classification results.

The evolutionarily optimized segment extends the original pipeline with an additional parallel process of force data. Instead of directly decomposing the force data into a fixed number of components, a PCA process is followed by a *variance threshold* ($threshold = 0.6$). Next, a *polynomial features* node ($degree = 2$) generates a new feature matrix consisting of all polynomial combinations of the features with degree less than or equal to the specified degree. Subsequently, an *FWE selector* node provides feature selection by selecting the p-values ($alpha = 0.016$) corresponding to family-wise error (FWE) rate as scored by an ANOVA. Afterwards, the feature values are scaled in the same manner as the handcrafted design. Finally, classification is now provided by an extra-trees classifier. While more complex than k-NN, extra-trees fits a number of randomized decision trees ($n_{estimators} = 100$) on various data subsamples and uses averaging to improve the accuracy and control overfitting.

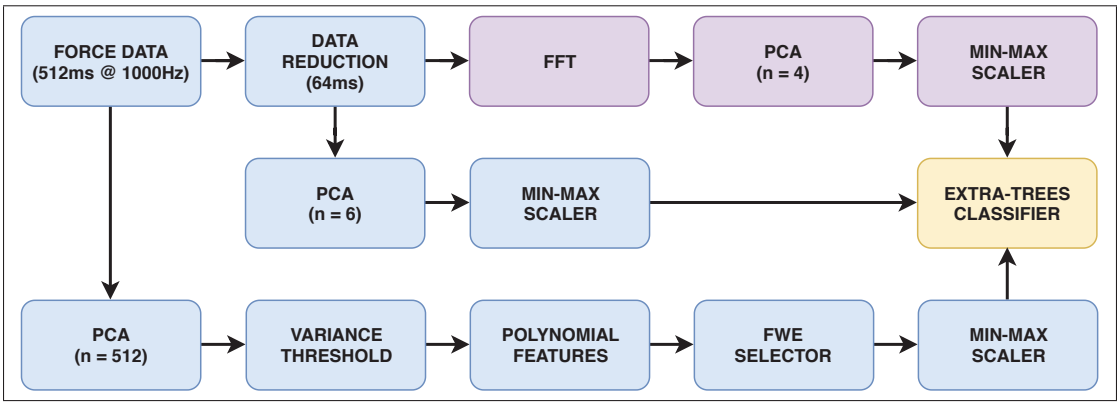

Figure 3.8 Evolutionarily optimized classification pipeline using *TPOT* and *DEAP*. The pipeline was optimized for prediction accuracy, with parallel processing of force data (blue) and frequency data (purple) for contact classification

3.3.3.1 Results

The classification results for the evolutionarily optimized pipeline are visualized in Fig. 3.9. The pipeline achieved an overall accuracy of 89.37 % across all labels, representing an 11.55 percentage point increase over the original pipeline. Like the handcrafted pipeline, *thigh* contact events are still often mistakenly classified as *calf* contact. However, all other confusion matrix errors have been greatly reduced. Furthermore, the new pipeline retains the accuracy with respect to distinction between *hard*, *free space*, and *human* contact events.

The full classification report may be seen in Table 3.2. Nearly all metrics experienced double-digit improvement. Most remarkably, *ankle* and *thigh* classification *precision* increased by 26 and 32 percentage points, respectively.

3.3.4 Pipeline Characterization

As noted in Kouris *et al.* (2018), classifiers require time to detect and process a given sample. In the context of robotics, this constraint has implications on the feasibility of real-time implementation. Furthermore, collecting training samples often requires process downtime and loss of productivity, potentially reducing the practicability of the approach.

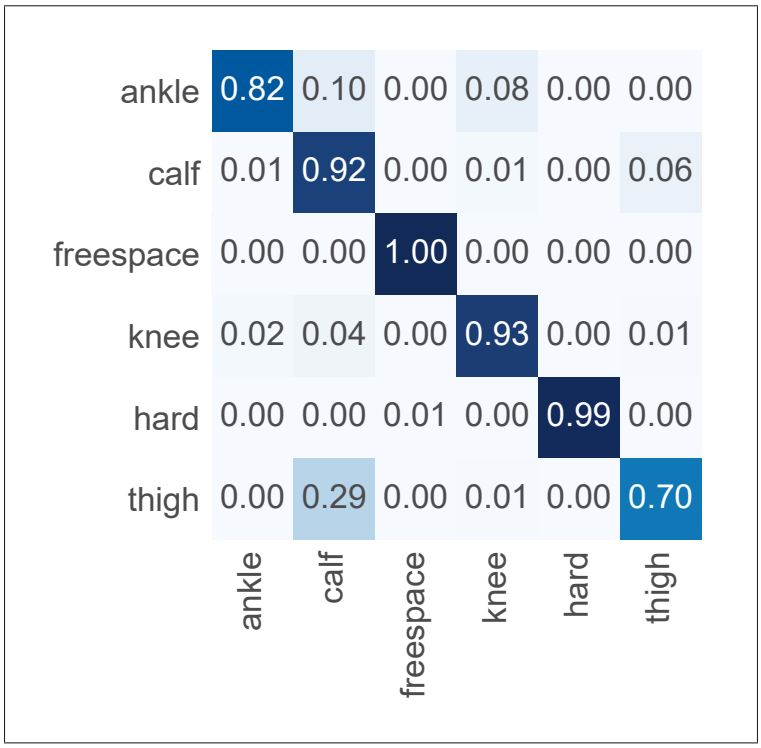

Figure 3.9 Evolutionarily optimized pipeline confusion matrix

Table 3.2 Classification summary for the evolutionarily optimized pipeline

Label	**Precision (%)**	**Recall (%)**	**F1 Score (%)**	**Support (#)**
knee	90	93	92	122
free space	99	100	99	86
ankle	97	82	89	122
table	100	99	100	146
thigh	89	70	78	124
calf	74	92	82	162

As such, the pipelines were characterized with respect to distinction speed and required experience. The pipeline learning curves are shown in Fig. 3.10 and the results are the average of the validations from a stratified k-fold cross-validation ($k = 10$). Of the k subsamples, a single subsample is retained as the validation data for testing the model, and the remaining $k - 1$

subsamples are used as training data. The cross-validation process is then repeated k times, with each of the k subsamples used exactly once as the validation data. Each process uses a new model. The k results are then averaged to produce a single score. The folds are made by preserving the percentage of samples for each label class (Kohavi *et al.*, 1995). It may be noted that at least 1500 training samples (3.9 s per training session, on average) are required before achieving relatively high accuracy with respect to contact classification prediction.

In terms of computation speed, the handcrafted pipeline required an average of 2.71 ms ±192 µs ($mean \pm std.\ deviation$) per sample. In contrast, the optimized pipeline required an average of 27.7 ms ±483 µs per sample. Speed is an important metric to consider, as classification should occur within a similar window as the safety monitoring of the robot (Section 3.2.3). It should be noted that the current implementation of the classifier uses *Python*, an interpreted language. A native code implementation can provide additional performance improvements. Furthermore, computation speed tests were performed with a 2.80 GHz Intel Core i5 processor.

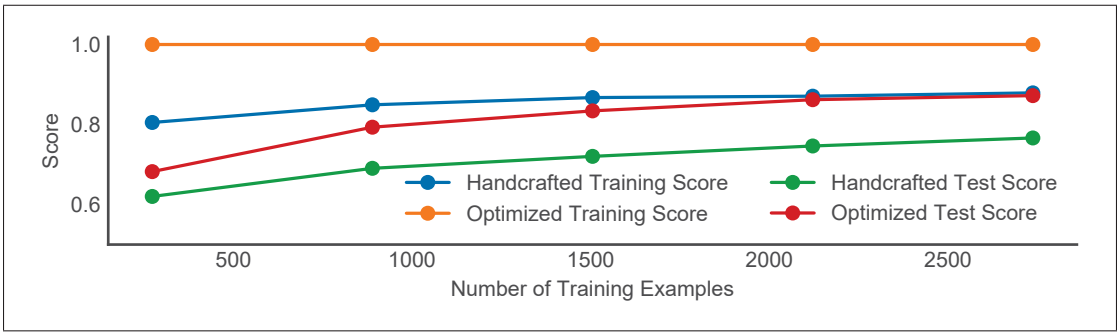

Figure 3.10 Handcrafted pipeline learning curves. All test scores are the average of the validations from a stratified k-fold cross-validation ($k = 10$)

3.4 Discussion

Collisions and contact events are a natural by-product of applications involving pHRI in unstructured environments. Accordingly, the ability to classify contact events allows for more complex robot reaction behaviour. This work presented a high-accuracy supervised learning

approach to robot contact classification using the built-in sensor data of a KUKA LBR iiwa robot.

As demonstrated in Section 3.3, a simple classification pipeline can discern between *hard*, free space, and *human* contact events with perfect accuracy. Through evolutionary optimization, the pipeline has an accuracy of 89.37 % across all labels, including the ISO/TS 15066:2016 lower limb body regions: *thigh*, *knee*, *calf*, and *ankle*. Moreover, the mean computation times of 2.71 ms and 27.70 ms per sample for the handcrafted and optimized pipelines, respectively, are comparable to the emergency stop reaction time of the LBR robot (22 ms to 154 ms). This indicates that the classification pipeline can be used in conjunction with existing reaction strategies, allowing for more information to be extracted from the contact event and a more appropriate reaction strategy to take place, instead of simply stopping.

With respect to previous studies, the presented classification pipeline has four distinct advantages. First, as compared to Kouris *et al.* (2018), the pipeline uses significantly less time series data to make a prediction. Second, Kouris *et al.* (2018) are only capable of contact distinction at high speeds ($\geq$200 mm/s), whereas our lowest speed was 50 mm/s. Third, in contrast to Haddadin *et al.* (2008), Cho *et al.* (2012), Geravand *et al.* (2013), and Golz *et al.* (2015), the presented concept is model-free and only requires an end-effector force sensor, allowing this concept to be applied to robots without joint torque sensors. Finally, this work unlocks a new contact classification paradigm beyond intended vs. unintended contact distinction, allowing the system to understand *what* was involved in the contact event.

Stopping the robot due to a contact event is not necessarily an optimal method for human safety and risk avoidance. Rather, if local information can be extracted from the event, it becomes possible to control the robot such that the task does not need to be abandoned. In a medical robotics context, it is interesting to consider the classification of different specific body areas along the human-model leg. This classification knowledge can be directly applied to ultrasound applications for probe placement and body area targeting, complimenting the work of Koizumi *et al.* (2009) and Gilbertson & Anthony (2015). Alternatively, as specific body areas have

different safety criteria, robot reaction strategies can be further tailored in real-time to respect and compensate for these varying safety levels. Altogether, contact classification allows for complex reaction strategies and robot behaviour, such as behaving in a more conservative manner when human contact occurs as opposed to *hard* contact.

When reviewing the deformations seen in Fig. 3.1 and the results of Fig. 3.3, the authors are in complete agreement with the observations of Rosenstrauch & Krüger (2017). Simply put, a programmed force condition trigger for a motion does not directly translate into a safe-contact application. Even with safety limits below those given in ISO/TS 15066:2016, there is still a significant lack of directly applicable guidelines for safe pHRI. It should be noted that the force limits described in ISO/TS 15066:2016 are for the human pain threshold and are the result of a single study. Though the human pain threshold is useful knowledge, explicit pHRI applications (i.e., medical ultrasound) will want to be well below this threshold. Therefore, future pHRI research should ask "What is a comfortable force threshold?", particularly if it can replicate existing human interaction. True application safety is not trivial, especially for unstructured environments.

Beyond the contact distinction application, this work has implications in process data decomposition and feature extraction for machine learning in the context of robotics. Similar in concept to the dictionary learning presented in Roberge *et al.* (2016), force sensor signals can be compressed into smaller representations for classification applications. In this case, it was found that splitting the signal into raw force and frequency components (FFT) followed by a dimensionality reduction step (PCA) provided a robust yet compact representation of the process data. As the larger the number of variables allowed, the greater is the chance of an overfitted model, this dimensionality reduction approach allows a better generalization to other datasets.

Future work should be devoted to the integration of contact event classification into robot reaction strategies for a system to automatically navigate human contact events or fail gracefully without requiring an emergency stop. This study should also be extended to investigate other real-time control approaches, such as recurrent neural networks, and explore the possibility of transfer

learning from simulation and phantom models to reduce deployment overhead. Altogether, while collaborative robot safety has not yet reached an ideal stage, a robot can understand and react to its environment in an appropriate manner.

CHAPTER 4

EVOLUTIONARY MOTION CONTROL OPTIMIZATION IN PHYSICAL HUMAN-ROBOT INTERACTION

This chapter presents a simple and effective online tuning framework using differential evolution (DE) to optimize the motion parameters for parallel force/impedance control in a pHRI and medical ultrasound motion application. Furthermore, a classification pipeline, trained using the tuning process motion data was developed to extend and combine the work presented in Chapter 3. Given that the success of an interaction task depends on the capability of the robot system to handle physical contact with its environment, pure motion control is often insufficient. This is especially true in the context of medical freehand ultrasound where the human body is a deformable surface and an unstructured environment, representing both a safety concern and a challenge for trajectory planning and control. While traditional tuning is generally performed *ad hoc* and requires knowledge of the robot and environment dynamics, the systematic tuning of practical high DOF pHRI tasks is not trivial and there are many parameters to be tuned. Through real-world experiments with a KUKA LBR iiwa 7 R800 collaborative robot, the DE framework tuned motion control for optimal and safe trajectories along a human leg phantom. The optimization process was able to successfully reduce the mean absolute error (MAE) of the motion contact force to 0.537 N through the evolution of eight motion control parameters. Moreover, the predictive motion classifier was able to reliably predict the future quality of a motion session (with respect to the force tracking error) with an accuracy of 91.82 %.

4.1 Introduction

Collaborative industrial robots are increasingly being used outside of traditional manufacturing where pHRI plays an important role in the application. Given that the success of an interaction task depends on the capability of the robot system to handle physical contact with its environment, pure motion control is often insufficient. This is especially true in the context of medical freehand ultrasound where the human body is a deformable surface and an unstructured environment, representing both a safety concern and a challenge for trajectory planning and control.

Within the research domain of robot-assisted freehand ultrasound, previous studies have focused on ergonomic benefits (Salcudean *et al.*, 1999), image quality optimization (Chatelain *et al.*, 2017) and reconstruction (Merouche *et al.*, 2015b), pHRI safety (Pierrot *et al.*, 1999), and application design (Mathiassen *et al.*, 2016). Regardless of the objective, force control techniques are necessary to achieve effective and adaptable behaviour of a robotic system in the unstructured ultrasound environment while also ensuring safe pHRI. However, while force control does not require explicit knowledge of the environment, to achieve an acceptable dynamic behaviour, the control parameters (e.g., gain, stiffness, damping) must be tuned.

The performance of a force control strategy is greatly affected by the tuned parameters. Past studies have explored effective pHRI control and tuning including, iterative tuning (Yu *et al.*, 2013), adaptive control (Alqaudi *et al.*, 2016; Gribovskaya *et al.*, 2011; Jung *et al.*, 2004), fuzzy learning (Du *et al.*, 2017), impact collision modelling (Roveda *et al.*, 2016), and passivity-based control (Albu-Schäffer *et al.*, 2007). Regardless of the approach, the tuning is generally performed on an *ad hoc* basis (Sage *et al.*, 1999) or requires knowledge of the robot and environment dynamics (Kelly, 1995).

As noted by Yu *et al.* (2013), the systematic tuning of practical high DOF robotic tasks is not trivial; there are many parameters to be tuned simultaneously and heuristically, in addition to the coupling dynamics that must be considered. Moreover, Villani & De Schutter (2008) point out that the choice of control parameters depends on the application and involves a compromise between interaction forces and position accuracy. Within the context of medical ultrasound, different human body locations have different stiffnesses (ISO/TS 15066:2016) and thus will require different tunings.

In a similar manner, pHRI safety and process optimization through planning and control is an often-discussed topic, but lacks practical solutions for direct contact tasks, such as medical ultrasound. For example, Haddadin *et al.* (2011) considered contextual awareness and reaction strategies for developing post-contact event safety and improved navigation but did not offer concrete approaches. Interestingly, Virga *et al.* (2016) noted variability and repeatability issues

in ultrasound image acquisition due to probe force effects, thus hinting at the possibility of process optimization through contextual pHRI motion planning.

Given these challenges, we first present an online optimization framework for parallel force/impedance control using differential evolution (Storn & Price, 1997) in the context of pHRI and medical ultrasound. While previous research has explored the integration of DE in robot control, such as simulated control tuning of a PUMA-560 (Zheng *et al.*, 2014) or PID tuning in a five-bar mechanism (Villarreal-Cervantes & Alvarez-Gallegos, 2016), there is a lack of knowledge with regards to the potential use of evolutionary algorithms in pHRI motion control. The current study takes inspiration from Jeon & Tomizuka (1993), Buchli *et al.* (2011), and Inoue *et al.* (2017) and leverages task repetition to tune motion control for optimal and safe trajectories along a human leg phantom. In our approach, DE can enable the robot system to explore the large parameter domain for candidate tuning solutions using metaheuristics without any assumptions about the problem space or environment. Further, a primary advantage of the framework is the preservation of a simple control structure, as no additional complexities have been added (i.e., fuzzy logic, neural networks).

Second, following DE tuning optimization, a future motion quality classifier was trained using the motion process data and the techniques presented in Chapter 3. This extension of the adaptive control framework allows for safety and pHRI process optimization through contextual pHRI motion planning.

The rest of the paper is organized as follows. First, Section 4.2 explains the problem of controlled motions along a deformable surface. Second, Section 4.3 describes the proposed optimization method and implementation of DE. Third, Section 4.4 outlines the experiment and the validation of the framework. Fourth, Section 4.5 presents the future motion quality classifier and test results. Finally, Section 4.6 concludes the paper and discusses potential future work.

4.2 Problem Description

The context of medical freehand ultrasound on a human leg is the basis of this study with the primary goal of maintaining a constant normal force with the probe, as compression effects would distort the resulting images (Gilbertson & Anthony, 2015). In the ideal application, robot-assisted ultrasound would allow for a sonographer to collaborate with the robot through hand-guidance while the robot performs the repetitive force-based motion tasks. Fundamentally, this application is a surface following operation on an unstructured and deformable body, a classical exercise in robotics and force control. In-depth surveys may be seen in Whitney (1987) and Siciliano & Villani (2012). Furthermore, given the pHRI context of this study, impedance control is an essential component to ensure safe motion and to compensate for the unstructured environment (Ficuciello *et al.*, 2015).

In this work we combine force and impedance control in a parallel control model. As shown in Chiaverini & Sciavicco (1993), parallel control strategies allow for control in the full-dimensional space without selection matrices. Parallel control loops override one another given a prioritization policy. For this task, the impedance control loop allows for safety and environment compensation, giving compliance to the tool frame in all Cartesian axes. At a higher priority level, task-orientated force control allows for constant force-tracking along the surface.

Thus, given an initial contact with the surface, the task can be achieved by assigning the following constraints (Villani & De Schutter, 2008):

- a non-zero linear velocity along the x axis
- zero linear velocity along the y axis
- zero angular velocity about the x, y, and z axes
- a non-zero force along the z axis

where the above motion task constraints and coordinate system are illustrated in Fig. 4.1.

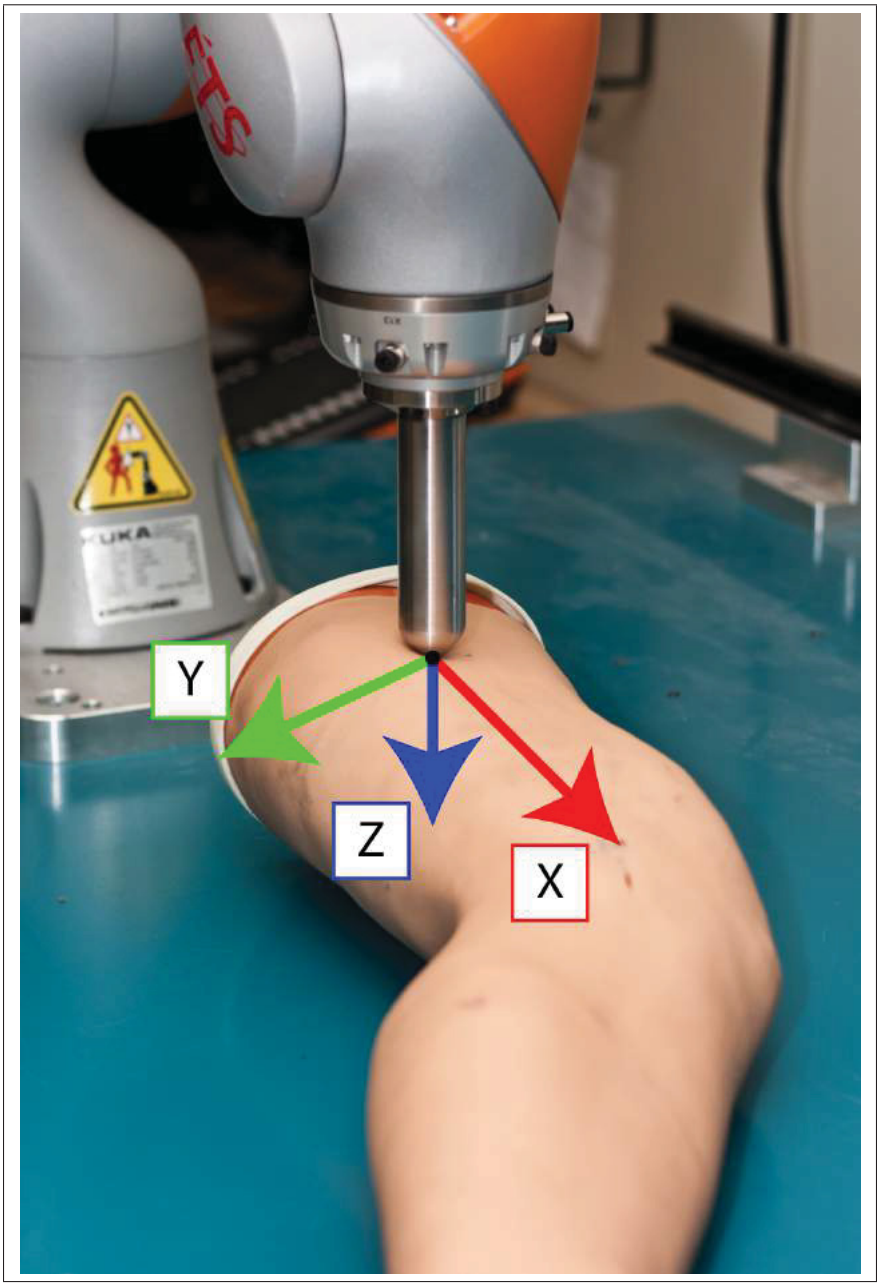

Figure 4.1 The motion task coordinate system. The primary goal was to maintain a constant normal force (f_z) with the probe while performing a linear motion along a human leg phantom (x axis direction)

4.3 Differential Evolution Framework

DE, first presented in Storn & Price (1997) and reviewed in Das & Suganthan (2011), is a simple yet powerful evolutionary algorithm for global optimization, being able to minimize nonlinear and non-differentiable continuous space functions. Typical direct search optimization methods (e.g., Nelder–Mead) use a greedy criterion to select newly generated parameters. Under this criterion, the new parameters are only kept if it reduces the fitness function value, risking convergence to local minima. In contrast, DE maintains a population of candidate solutions and creates new solutions through the strategic and stochastic combination of existing ones, allowing the full (but bounded) solution space to be continuously explored. This is especially important with regards to robotics, as the solution space fitness of a high DOF motion task is assumed to be noisy and is not guaranteed to be differentiable. An illustrated overview of DE may be seen in Fig. 4.2 with the specific implementation shown in Algorithm 4.1.

The DE solver for the present study was implemented using *Python* v3.6.1 and *SciPy* v1.0.0 on an external PC as a client to the robot controller over a local network. As optimization of the solver itself was outside the scope of the current study, the solver hyperparameters were set using the recommendations of both Storn & Price (1997) and the *SciPy* package. Notably, the population size was set to $5 \times D$, where D is the dimension of the input vector. The mutation constant was set in the range of $[0.5, 1]$, whereby dithering randomly changes this constant every generation, helping to speed convergence. The recombination constant (i.e., crossover probability) was set to 0.7. The population was initialized using Latin hypercube sampling to maximize coverage of the available parameter space. Finally, the *best1bin* evolution strategy was selected, where for each mutation, two candidates of the population are randomly chosen, and their difference is used to mutate the best candidate into a mutant vector. The crossover process then combines the given candidate with the mutant, based on a binomial distribution, where a random number in $[0, 1)$ is generated. If this number is less than the recombination constant, then the parameter is loaded from the mutant, otherwise it is loaded from the original candidate. The resulting trial vector is to be evaluated by a robot motion session, generating a

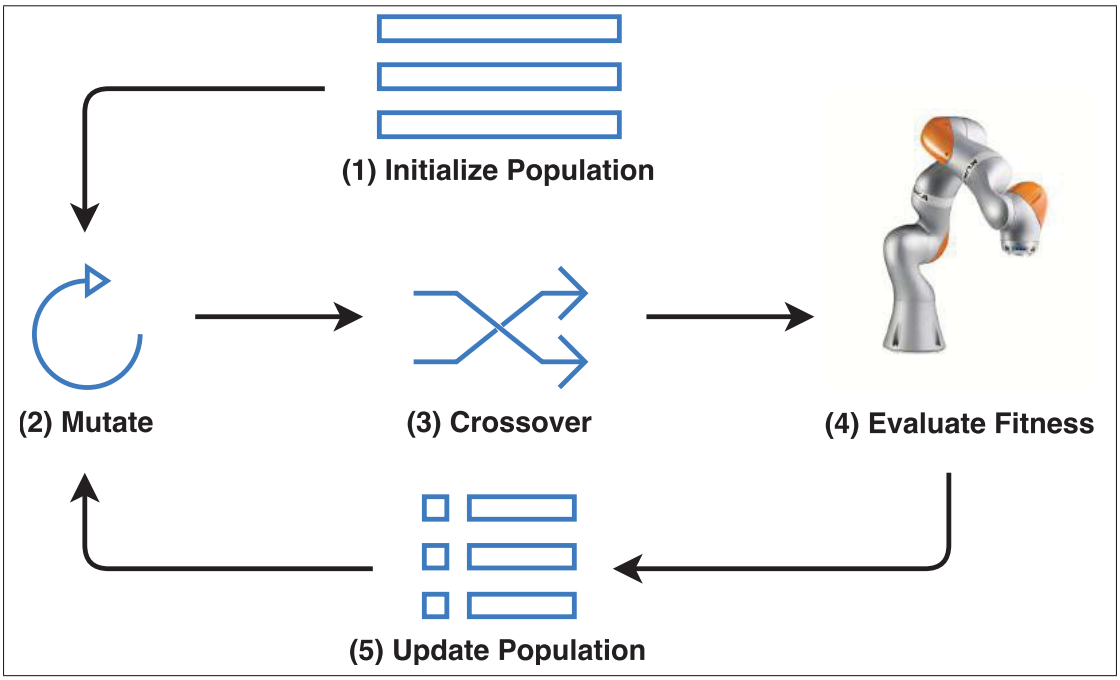

Figure 4.2 The high-level steps of differential evolution (DE): (1) A population of candidate solution vectors is initialized; (2) Mutations are generated for each candidate; (3) The mutations are crossed with the original candidates; (4) The fitness of the crossed candidates is scored; (5) The population is updated based on the scored fitness. Steps 2–5 are repeated until a convergence criterion is met

fitness score. All classifier hyperparameters were experimentally tuned using the training data. A detailed overview of this process is seen in Algorithm 4.1.

Algorithm 4.1 The Differential Evolution (DE) solver process

```
// initialize solver
CR ← 0.7 // crossover probability
D ← getNumParameters() // number of parameters
populationSize ← 5 × D
population ← initializePopulation(populationSize)

// each loop represents a single generation
loop
   // get a random mutation within the given range
   F ← random(0.5, 1)
   for all candidate ∈ population do
      // best1bin mutation
      r0 ← population.getRandomCandidate()
      r1 ← population.getRandomCandidate()
      Δr0,1 ← r0 − r1
      bc ← population.getBestCandidate()
      mutantCandidate ← bc + F × Δr0,1

      // best1bin crossover
      trialCandidate ← candidate
      for i to D do
         cr ← random(0, 1)
         if cr < CR then
            trialCandidate[i] ← mutantCandidate[i]
         end if
      end for
      // one random parameter is always from the mutant
      i ← randomInt(0, D)
      trialCandidate[i] ← mutantCandidate[i]

      // perform a session with the trial candidate
      // this blocks until the motion session is complete
      result ← performSession(trialCandidate)

      // evalute the fitness and update the population
      fitness_trial ← evaluateFitness(result)
      if fitness_trial < fitness_candidate then
         candidate ← trialCandidate
      end if
   end for
end loop
```

4.4 Experiment

The proposed optimization framework was evaluated with a 7-DOF KUKA LBR iiwa 7 R800 collaborative robot and the experiment setup may be seen in Fig. 4.3. The LBR robot was equipped with a rounded probe tool that was designed to approximate the tool referenced in Annex A of ISO/TS 15066:2016. In addition, the robot is equipped with torque sensors at every joint, allowing for calculations of Cartesian forces acting on the tool. The mass data of the tool assembly (0.948 kg) was calibrated using the built-in calibration procedure of the robot to remove the tool load from the force measurements.

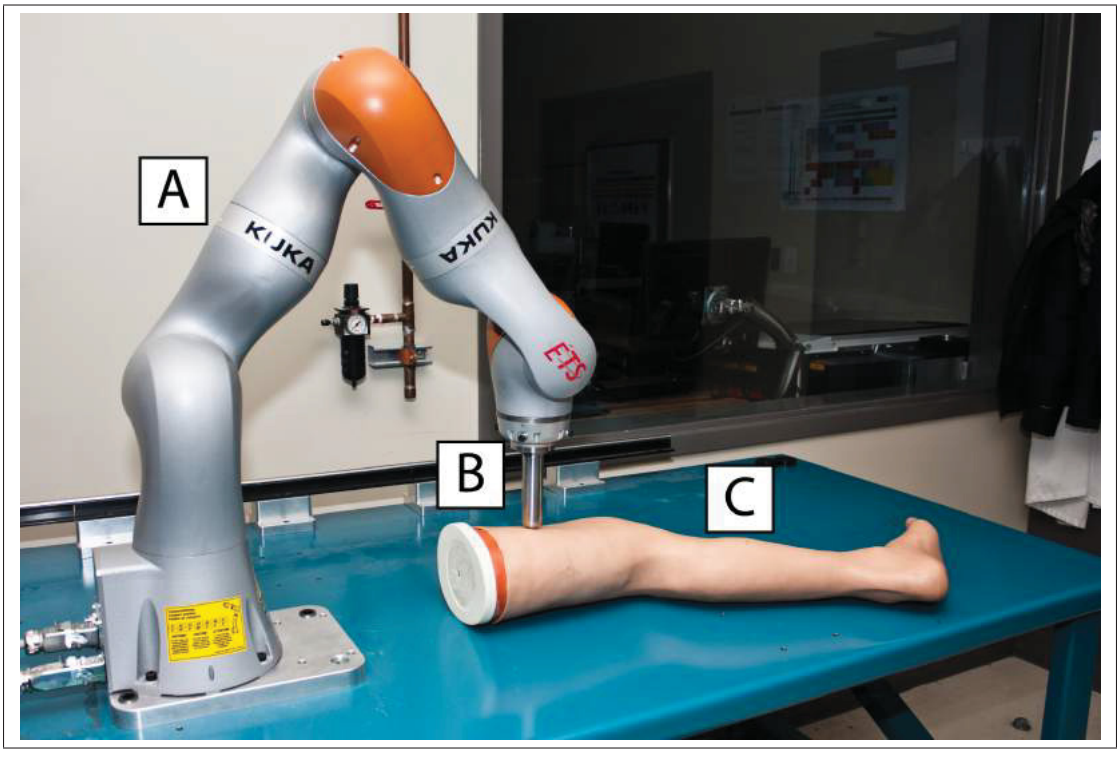

Figure 4.3 The motion tuning experiment setup. A KUKA LBR iiwa 7 R800 collaborative robot (A) was equipped with a rounded probe tool (B). A mannequin leg with an internal skeletal structure (C) was used to simulate the ultrasound subject and water-based lubricant was used on the leg to fully simulate real ultrasound examination and reduce friction

A mannequin leg with an internal skeletal structure was used to simulate the ultrasound subject and water-based lubricant was used on the leg to fully simulate real ultrasound examination

and reduce friction. The desired contact force for the motion was derived from a review of previous robot-assisted ultrasound studies. Forces between 1 N to 5 N, 3.70 N to 4.60 N, and 5 N to 20 N were reported in Pierrot *et al.* (1999), Chatelain *et al.* (2017), and Courreges *et al.* (2008), respectively. As such, a setpoint force of 5 N was selected.

As discussed in Section 4.2, parallel force/impedance control was selected as the control strategy and was implemented as two layers. First, the control architecture exploited the built-in real-time impedance control of the LBR robot as the base control strategy (*KUKA Sunrise.OS* v1.10). The basic model is a virtual spring-damper system with configurable values for stiffness and damping, allowing the LBR to be highly sensitive and compliant (Fig. 4.4). Second, the high-priority task-oriented force control was implemented through the *Sunrise.Connectivity DirectServo* control loop interface. This allows for non-deterministic soft real-time control and fast corrections to the robot path.

Each motion session begins with the DE solver process providing a candidate *Session State* (*Protocol Buffers* v3.5.1) to the robot controller through an RPC interface (*gRPC* v1.9.0). The *Session State* input vector defines the parameters that form the optimization problem and is given by:

$$SessionState = \begin{bmatrix} k_x & k_y & k_z & \zeta_x & \zeta_y & \zeta_z & d_x & g_z \end{bmatrix} \tag{4.1}$$

where $k_{x,y,z}$ and $\zeta_{x,y,z}$ are the Cartesian stiffnesses (0 N/m to 5000 N/m) and Lehr's damping ratios (0.10 to 1), respectively, for the LBR impedance control. d_x defines the x axis position displacement per servo iteration (0.10 mm to 10 mm) and g_z defines the gain of the force-tracking (0.10 to 10). The bounds of the impedance control parameters are hard limits set by the robot while the remaining bounds were selected experimentally. Furthermore, these parameters remain constant for the duration of a single session. Finally, the stiffness and damping of the rotation axes were set to a constant 300 N m/rad and 0.70, respectively, as they were not part of the optimization process to reduce the number of independent variables.

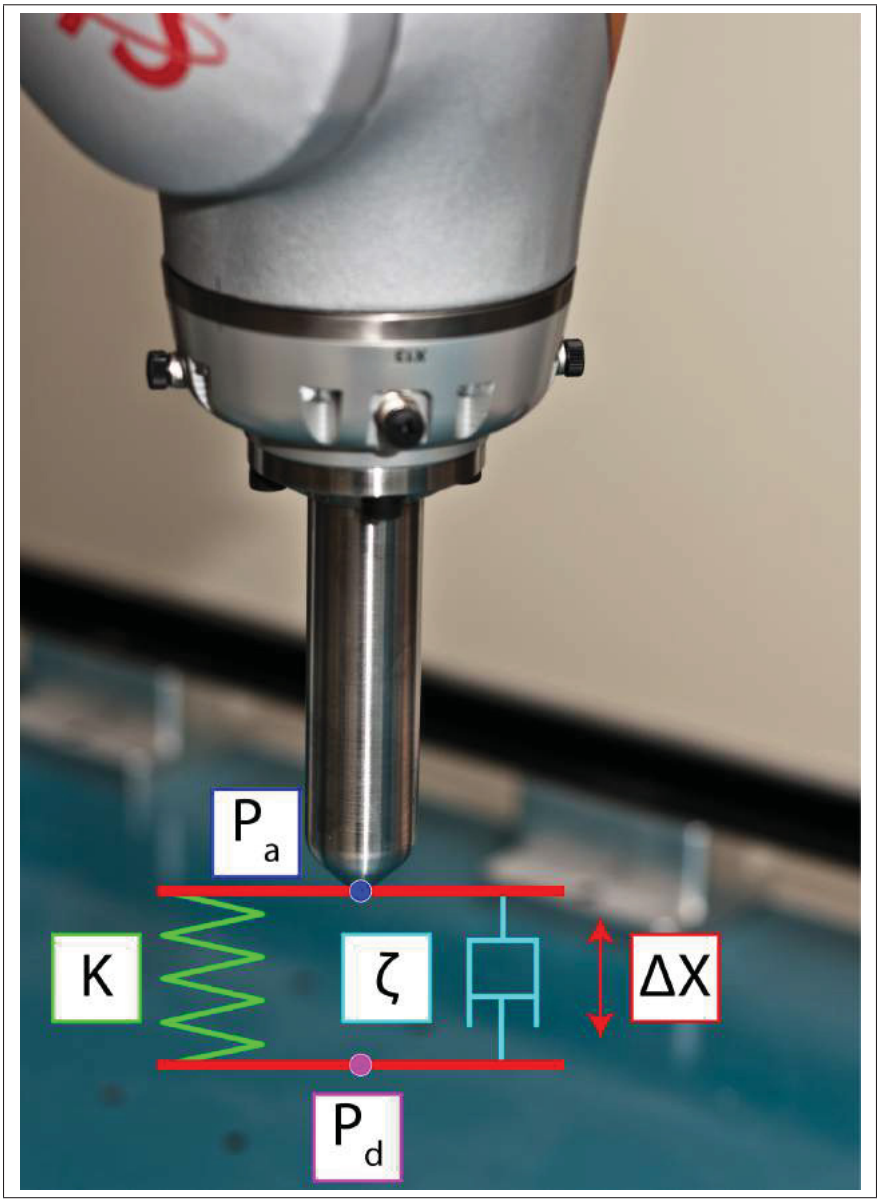

Figure 4.4 The built-in real-time impedance control of the KUKA LBR iiwa robot. The model is a virtual spring-damper system with configurable values for stiffness (K) and damping (ζ), allowing the LBR to be highly sensitive and compliant. The resulting interaction forces are calculated from the displacement (ΔX) between the desired position (P_d) and the actual position (P_a) in each Cartesian axis

Following session initialization, the robot performed a linear downwards motion until the force condition was triggered at which point the servo loop was initiated for the ultrasound motion.

A motion session is judged as a success after 100 mm of travel in the x axis direction from the initial contact point (visualized in Fig. 4.1). A timeout of 3000 servo iterations was also implemented to avoid endless loops (successful sessions took an average of 1129 iterations). Throughout the session, the z axis force (f_z) was logged at 1000 Hz.

Following the completion of a motion session, the robot controller updated the external client with a *Session Result*, consisting of the recorded f_z data and the session success status. The DE solver used the MAE of the f_z data with respect to the setpoint force as a fitness score. The first second of data was discarded to avoid edge effects with respect to the initial contact event (highlighted in Figs. 4.6 and 4.7). However, if a timeout occurred, the session was penalized with a score of infinity. Iteratively, motion sessions were performed, evaluating candidate control parameter solutions as generated by the DE solver with the parameters evolving over time. A detailed description of the motion optimization process and framework is specified in Algorithm 4.2.

4.4.1 Results

A total of 795 sessions were performed, representing nearly 20 generations of candidate solutions, with 709 successful sessions (89.2 %). The optimization process was terminated due to the fitness score leveling off, as clearly seen in Fig. 4.5. The mean servo loop period was 3.12 ms (320.49 Hz), on par with the responsive ultrasound force control reported by Chatelain *et al.* (2017).

Comparisons of motion behaviour throughout the process are shown in Figs. 4.6 to 4.8. The worst performing sessions noticeably suffer from strong oscillations or completely miss the desired contact force, as compared to stable and accurate optimized motion. It should be noted that the optimization of the initial contact event was outside the scope of this study, as this transient event would require its own set of motion control parameters, thus resulting in the force overshoots seen outside the shaded area in Fig. 4.6. In an actual robot-assisted application, the initial contact would most likely be achieved through the sonographer hand-guiding the robot.

Algorithm 4.2 The motion control session

```
// receive candidate vector from DE solver
LBR ← SessionState

// initialize session
numServoIterations ← 0
maxServoIterations ← 3000
f_{z,d} ← −5 // Newtons, desired force
Δx_d ← 100 // millimeters, desired travel distance
P̄_0 ← getCurrentCartesianPosition()

// initialize internal impedance control layer
LBR_{k_x,k_y,k_z} ← SessionState_{k_x,k_y,k_z}
LBR_{k_rx,ry,rz} ← 300 // Nm/rad
LBR_{ζ_x,ζ_y,ζ_z} ← SessionState_{ζ_x,ζ_y,ζ_z}
LBR_{ζ_rx,ry,rz} ← 0.7 // Lehr's damping ratio

isServoRunning ← True
while isServoRunning do
  numServoIterations ← numServoIterations + 1
  if numServoIterations > maxServoIterations or x_i − x_0 > Δx_d then
    isServoRunning ← False
  else
    // force control layer
    F̄_i ← getForceVector()
    Δf_z ← f_{z,d} − F̄_{z,i}
    Δz ← Δf_z / k_z
    Δz ← Δz × SessionState_{g_z}

    // calculate next servo position
    P̄_i ← getCurrentCartesianPosition()
    P̄_{x,i+1} ← P̄_{x,i} + SessionState_{d_x}
    P̄_{y,i+1} ← P̄_{y,0}
    P̄_{z,i+1} ← P̄_{z,i} + Δz
    P̄_{rx,ry,rz,i+1} ← P̄_{rx,ry,rz,0}
    setDestination(P̄_{i+1})
  end if
end while
```

Over the evolution of the candidate parameters, correlations (with absolute Pearson correlation coefficients greater than 0.4) were found between three key parameters (ζ_z, g_z, and k_z), the

fitness score, and the session index. These correlations give insight into the evolutionary process of the DE solver as it generates candidate solutions. The raw parameter evolution may be seen in Fig. 4.9 and Fig. 4.10. It is also interesting to note how even at later indices the DE solver still explores the full parameter domain and does not simply converge to local minima in parameter space.

Finally, the evolution of the best candidates is visualized in Fig. 4.11 and the optimal parameters are summarized in Table 4.1. While hundreds of motion sessions were performed, the MAE fitness reached 0.607 N after only 28 sessions and the optimal value of 0.537 N after 689 sessions. Close examination of the evolution in Fig. 4.11 shows relatively wide parameter exploration in sessions prior to session 28 and relative stabilization afterwards. The trend towards high z-direction rigidity (Figs. 4.9 and 4.10) suggests the application favours stronger position control, but the results in Table 4.1 demonstrate that the optimal parameters were not simply at maximum. Moreover, Fig. 4.8 clearly demonstrates that a simple fixed z-depth motion would not be a valid approach to this application.

Table 4.1 Parameters and fitness of the best differential evolution (DE) candidate. The fitness is calculated as the mean absolute error (MAE) between the desired contact force and the recorded force through a motion session

Parameter	**Value**
d_x	8.865 mm
g_z	9.736
k_x	3703.233 N/m
k_y	2777.930 N/m
k_z	4885.149 N/m
ζ_x	0.132
ζ_y	0.407
ζ_z	0.327
MAE	0.537 N

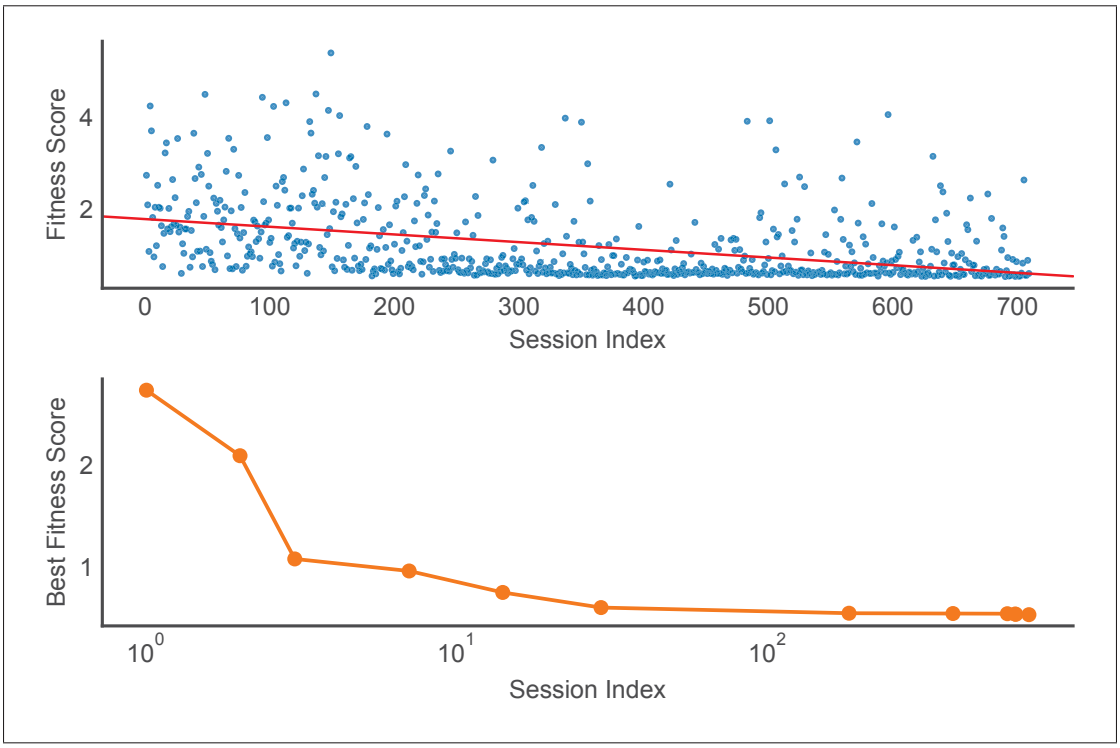

Figure 4.5 Raw fitness scores and linear regression (top) and best fitness scores (bottom) vs. session index

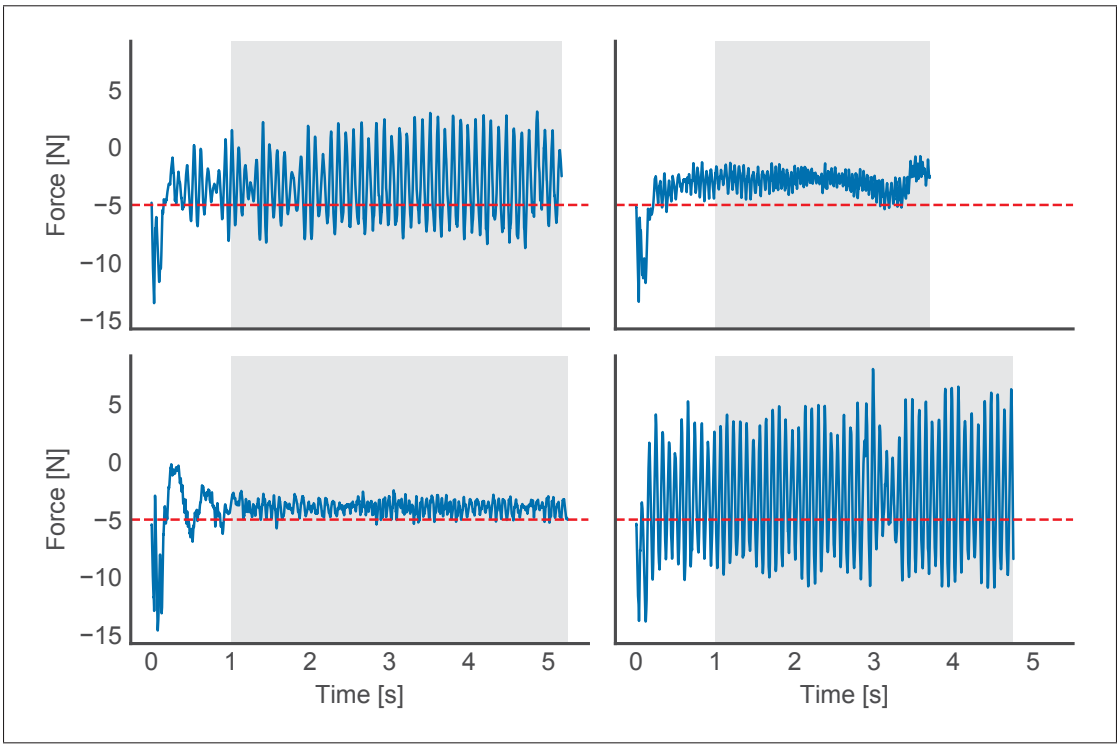

Figure 4.6 First few sessions of the probe contact force, f_z, throughout the differential evolution optimization process. The shaded region highlights the motion data used to score the candidate fitness

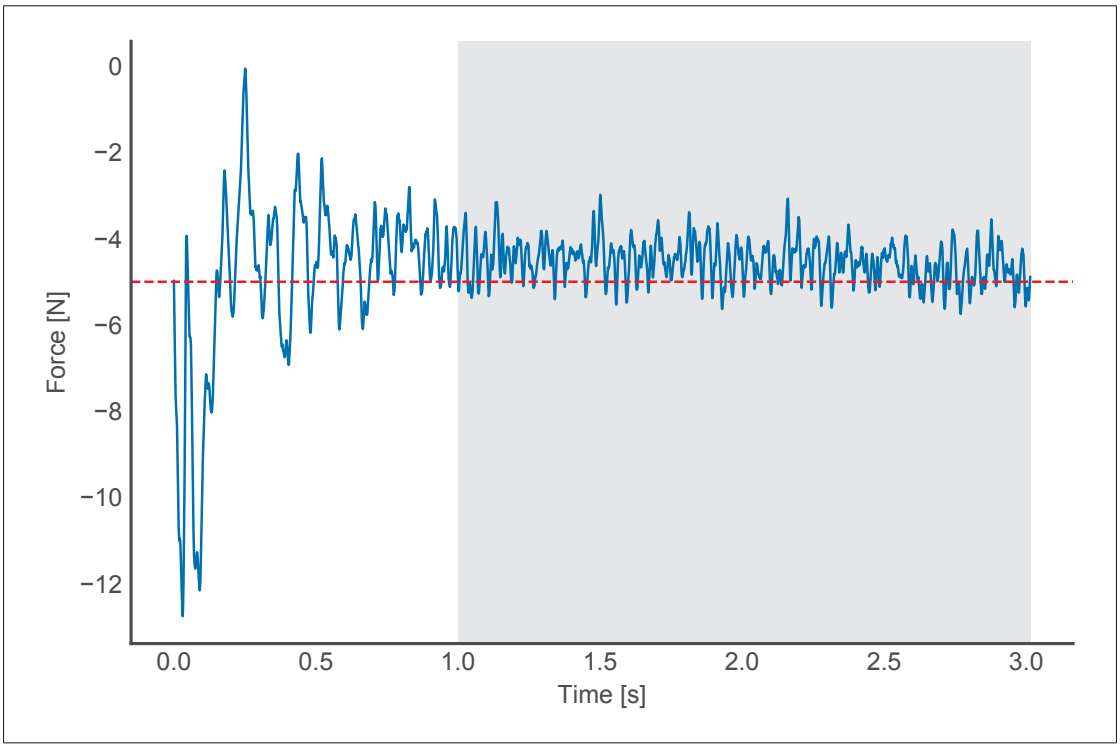

Figure 4.7 Best session probe contact force, f_z, from the differential evolution optimization process. The shaded region highlights the motion data used to score the candidate fitness

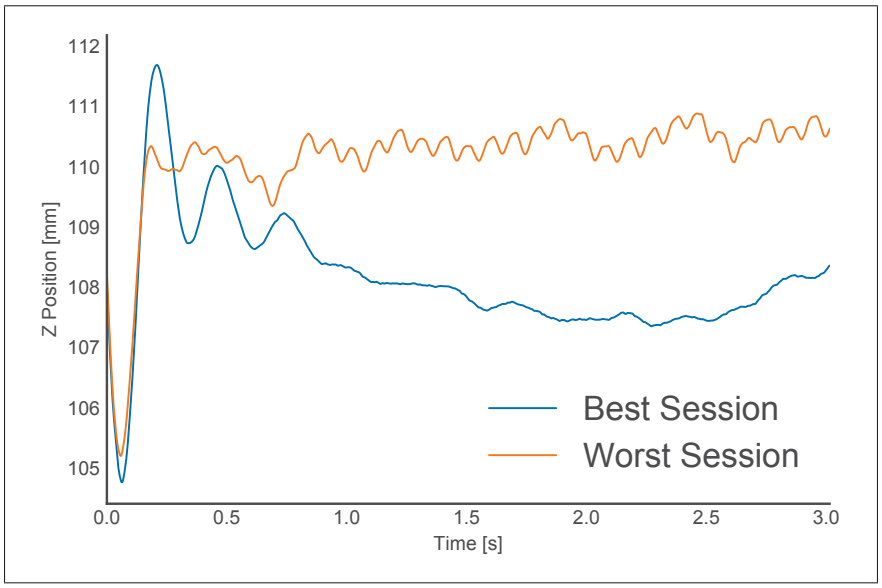

Figure 4.8 Probe z-position of the best and worst sessions of the differential evolution optimization process. The plot demonstrates that a simple fixed z-depth motion would not be a valid approach to this application

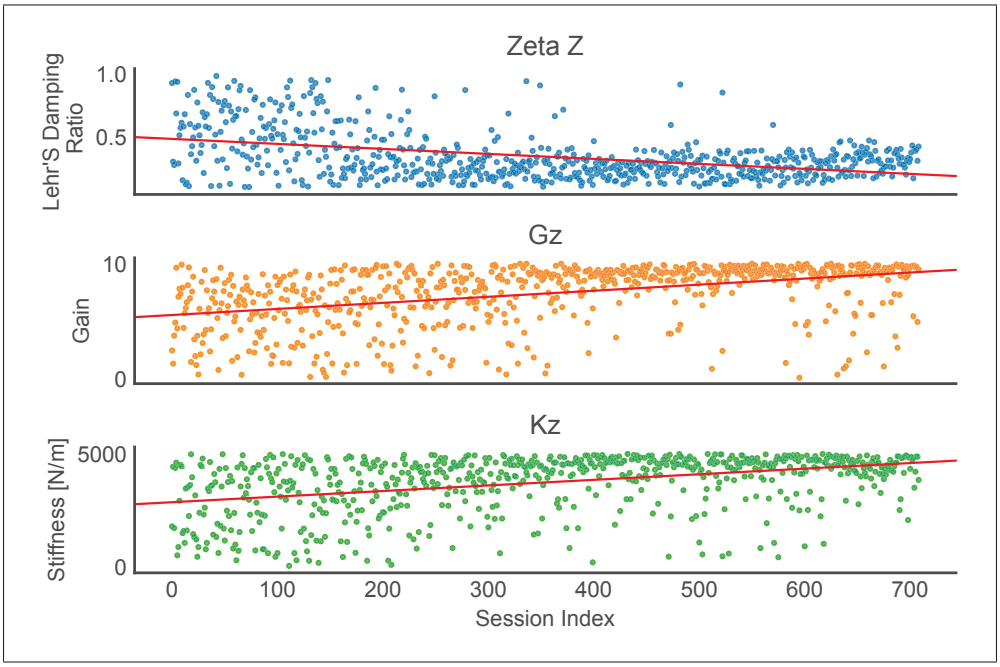

Figure 4.9 Evolution of key parameters with respect to the optimization process. Linear regressions are represented by the red lines

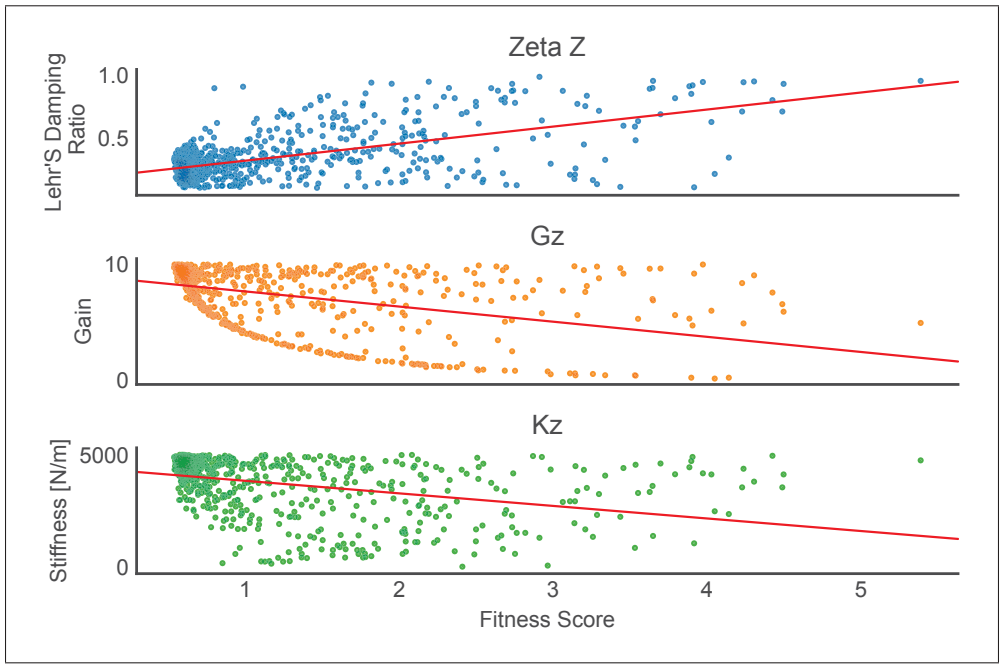

Figure 4.10 Evolution of key parameters with respect to the fitness score. Linear regressions are represented by the red lines

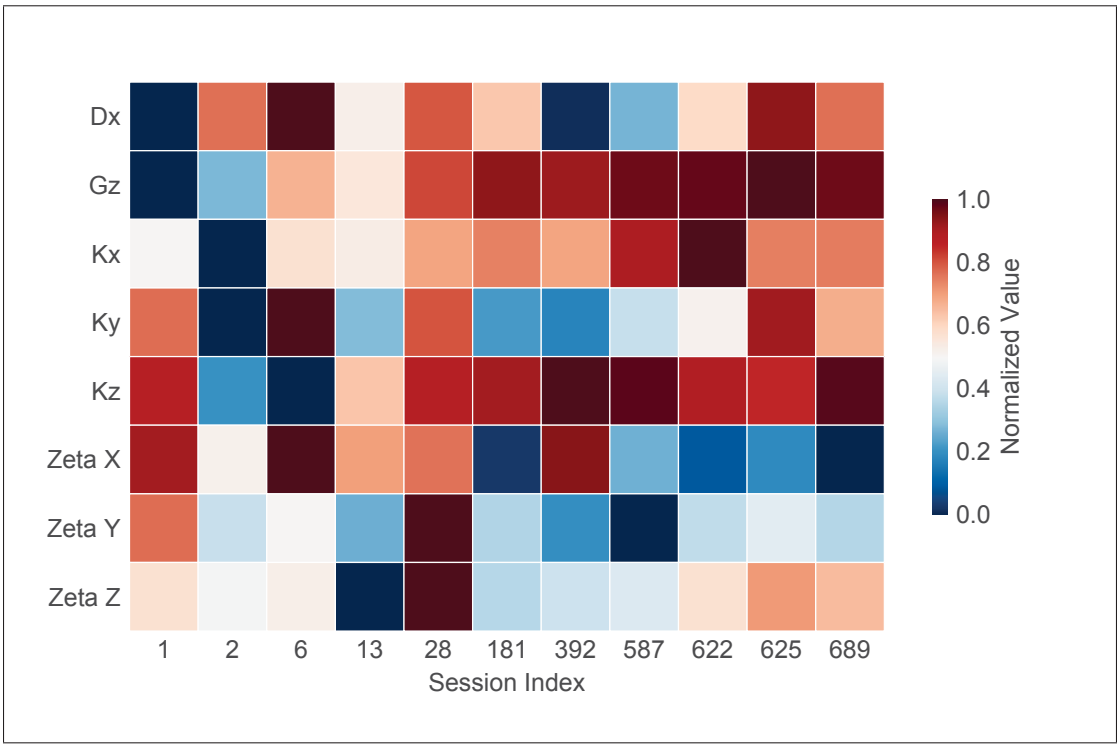

Figure 4.11 Evolution of the motion control parameters in the best candidates. The columns represent the parameters of best candidate up to a given session. Each parameter row is normalized over its evolution

4.5 Predictive Motion

Given the process data and the fitness scores of the motion sessions (Section 4.4), it was hypothesized that the future quality of a motion session with respect to the force tracking error could be predicted following the initial contact event between the subject and the robot. This predictive motion would allow for safety through intelligent reaction of the robot system and process optimization by not wasting time on a low-quality motion session. Contextual awareness through predictive motion would give a basis for autonomous decision making in pHRI.

Following the data acquisition and motion tuning presented in Sections 4.3 and 4.4, the resulting motion process data was evaluated using the techniques presented in Chapter 3. As discussed in Section 4.6, process data decomposition and feature extraction can compress force sensor signals into smaller representations for classification applications. As in Section 3.3, it was found that splitting the signal into raw force and frequency components followed by a dimensionality reduction step provided a robust yet compact representation of the process data for a classification pipeline.

Previously described in Section 4.4, the fitness was calculated as the MAE between the desired contact force and the recorded force through a motion session. Further analysis of the process data found that the mean successful fitness score was 1.19 N. A histogram of fitness scores may be seen in Figure 4.12. As such, for future motion quality classification purposes, the dataset was split into *good* and *bad* motion sessions, where a *good* session was defined as $MAE < 1\,\text{N}$. This resulted in 432 *good* sessions and 363 *bad* sessions. Moreover, it was also found that the minimum and mean successful motion durations, from initial contact to final position, were 2.86 s and 3.50 s, respectively.

Using this transformed dataset and additional insights, a classification pipeline was constructed to classify future motion session quality. The architecture of the pipeline is shown in Figure 4.13. This pipeline is nearly identical to the one presented in Section 3.3.2 but uses post-contact force sensor data to qualify the motion session.

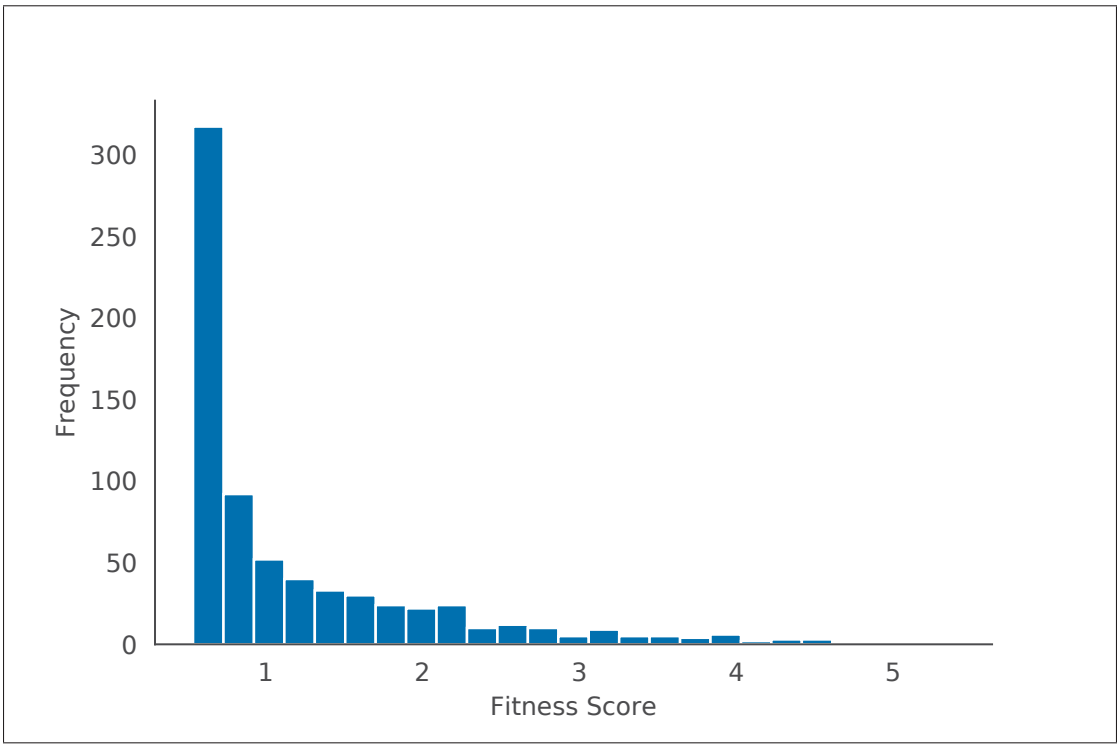

Figure 4.12 Histogram of fitness scores from the 709 successful sessions

The pipeline steps are summarized as follows: First, the raw force data is truncated to the 1024 ms following the contact event (i.e., the initiation of the desired motion). This time period is well below the minimum motion duration and would allow for an early reaction event if necessary. Next, the pipeline splits between force (blue) and frequency (purple) data processing. PCA decomposition is then used to further reduce the dimensionality of the data. Prior to consolidation, the data is scaled such that the value of each feature is between 0 to 1. Finally, classification is provided by an extra-trees classifier. Extra-trees fits a number of randomized decision trees ($n_{estimators} = 10$) on various data subsamples and uses averaging to improve the accuracy and control overfitting.

4.5.1 Results

This classification pipeline was trained on a subset of the overall dataset, leaving 20 % of the data (159 sessions) for testing. Detailed classification results for may be seen in Table 4.2. The

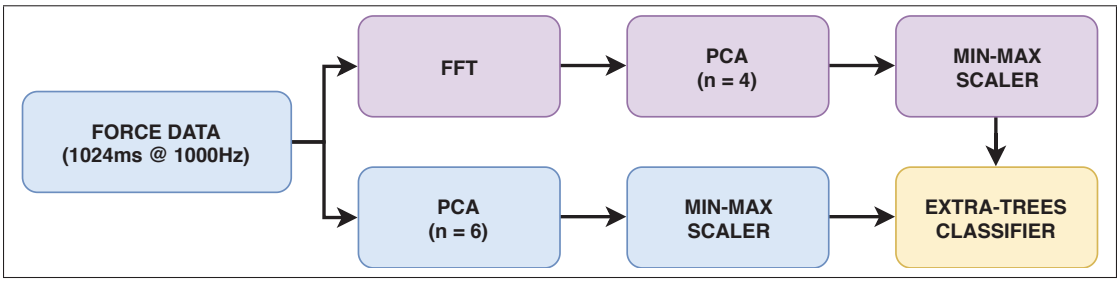

Figure 4.13 Classification pipeline using fast Fourier transform (FFT), principal component analysis (PCA), and an extra-trees classifier. The pipeline was designed for data decomposition and feature extraction, with parallel processing of force data (blue) and frequency data (purple) for future motion quality classification

pipeline achieved an overall accuracy of 91.82 %. As previously described in Section 3.3.2.1: *Precision* is the ability to correctly apply a label. *Recall* is the ability to find all samples belonging to a label. *F1 Score* is the mean of *precision* and *recall*. *Support* is the number of occurrences of each label.

Table 4.2 Classification summary of the test data for the predictive motion pipeline

Performance Label	**Precision (%)**	**Recall (%)**	**F1 Score (%)**	**Support (#)**
Good	90	95	93	88
Bad	94	87	91	71

4.6 Discussion

The systematic tuning of practical high DOF robotic tasks is not trivial and there are many parameters to be tuned, especially in the context of pHRI force control. While traditional tuning is generally performed *ad hoc* and requires knowledge of the robot and environment dynamics, we propose a simple and effective online framework using DE to optimize the motion parameters for parallel force/impedance control in a medical ultrasound motion application.

Through real-world experiments with a KUKA LBR iiwa 7 R800 collaborative robot, probe tool, and a mannequin leg with an internal skeletal structure, the DE framework was able

to successfully reduce the MAE of the motion task to 0.537 N through the optimization and evolution of eight motion control parameters. Although real-world fitness function evaluations are a relatively expensive process, the total optimization time was just 55 min for 20 generations, well under the 50-generation benchmark set in Zheng *et al.* (2014).

The simplicity of the framework has three key benefits. First, inferring the motion parameters from the application is not trivial, especially since the human leg has varying stiffness, according to ISO/TS 15066:2016. Second, it can easily be used to extend and improve existing control schemes through a simple communication network between the DE solver and the robot. Third, it allows for the optimization process to be well understood by the user, unlike black-box methods. As this approach is model-based, the motion parameters may be easily updated with respect to changing environments (e.g., change in leg orientation), whereas a black-box model may have difficulty adapting.

Furthermore, following the tuning process, a classification pipeline was developed to predict the quality of future motion sessions. The pipeline was able to reliably classify *good* and *bad* motion sessions with an accuracy of 91.82 %. This contextual awareness through motion classification can offer a framework for pHRI optimization and safety through predictive motion behaviour with a future goal of autonomous pHRI.

Consequently, while the mannequin leg may not perfectly represent a human subject, this framework can be adopted as a starting point for more advanced control paradigms, such as reinforcement learning. Future motion control architectures for pHRI and intelligent systems will likely rely on forms of continuous learning, highlighting the importance of an optimization framework that allows the robot system to explore, evolve, and evaluate the solution space for more effective motion behaviour.

As stability analysis was outside the scope of the current study, it provides a foundation for future work. Other potential extensions of this study include the optimization of the DE solver's hyperparameters, the exploration of new motion parameters and control architectures, and

the development of smooth transitions between initial contact and surface following motion control.

CONCLUSION AND RECOMMENDATIONS

The following chapter provides a summary of the thesis' contributions and recommendations for addressing its limitations.

5.1 Summary of Contributions

In Chapter 2, we proposed a novel closed-loop calibration approach using impedance control. Additionally, a data communication architecture, based on open-source tools, was introduced that manages the calibration process in an adaptable and scalable manner. The calibration procedure uses standard machine tool components, such as precision spheres and a kinematic coupling, which are suitable for harsh environments. As such, the required equipment is low cost, robust, and is quick to set up, especially when compared to traditional calibration devices. As no external measuring devices are required and the procedure uses force control, this approach fills a practical gap in robotic ultrasound presented in Section 1.5.3. Experimental studies and validations with a laser tracker demonstrated that the absolute accuracy of the KUKA LBR iiwa robot was improved to a maximum error of 0.99 mm, representing a 58.4% improvement when compared to the nominal model. As a comparison, further testing showed that a traditional calibration using a laser tracker only improved the maximum error by 58 µm over the impedance control approach.

With respect to the current literature, Chapter 2 directly addresses the current lack of practical solutions for accurate probe motion for medical ultrasound image acquisition. As discussed in Szabo (2014), many advanced medical ultrasound techniques are highly dependent on accurate positioning and movement of the ultrasound probe. It is through this work that past robot-assisted ultrasound research, such as Aalamifar *et al.* (2014), Joubair *et al.* (2016), Hennersperger *et al.* (2016), and Sharifi *et al.* (2017), can be improved beyond the low-accuracy results previously presented.

Chapter 3 presented an experimentally tested, machine learning framework for robot contact classification using the built-in sensor data of a KUKA LBR iiwa robot. As collisions and contact events are a natural by-product of applications involving pHRI in unstructured environments the ability to classify contact events would allow for more complex robot reaction behaviour. Safety is a key issue when designing any robotic application, especially medical procedures. Unlike previous studies presented in Section 1.5.4, this classification does not discern between intended vs. unintended contact scenarios, but rather classifies *what* was involved in the contact event. The framework can perfectly differentiate between high-level *hard*, *free space*, and *human* motion events at low speeds with an average prediction time of 2.71 ms. Through evolutionary optimization of the classification pipeline, the classifier can discern different ISO/TS 15066:2016 specific body areas along a human-model leg with 89.37 % accuracy. Overall, this contact distinction framework allows for more complex reaction strategies and tailored robot behaviour during robotic freehand ultrasound.

As discussed in Fei *et al.* (2001), a medical robotic system should be able to identify and react to events to ensure application safety. Chapter 3 not only offers a practical example for pHRI contact distinction but offers a safety framework beyond what was presented in Alami *et al.* (2006) and Haddadin *et al.* (2011). In addition, the work presented in Hennersperger *et al.* (2016) and Virga *et al.* (2016) can now be enhanced with anatomical distinction for improved force-controlled trajectories.

In Chapter 4, a simple and effective online tuning framework using differential evolution was presented. While traditional tuning, as seen in Section 1.5.2, is generally performed *ad hoc* and requires knowledge of the robot and environment dynamics, this framework can optimize the motion parameters for parallel force/impedance control in a medical ultrasound motion application. The human body represents both a safety concern and a challenge for trajectory planning and control and the success of an interaction task depends on the capability of the

robot system to handle physical contact with its environment. In the context of medical freehand ultrasound where the human body is a deformable surface and an unstructured environment, pure motion control is insufficient. Furthermore, the systematic tuning of practical high DOF pHRI tasks is not trivial and there are many parameters to be tuned. Through real-world experiments with a KUKA LBR iiwa 7 R800 collaborative robot, the DE framework tuned motion control for optimal and safe trajectories along a human leg phantom. The optimization process was able to successfully reduce the MAE of the motion contact force to 0.537 N through the evolution of eight motion control parameters.

As seen in Section 4.5, the work presented in Chapter 3 and Chapter 4 can be combined for a full end-to-end motion control solution for robot-assisted freehand ultrasound. The presented classification pipeline was able to reliably predict the quality of future motion sessions with an accuracy of 91.82 %. This contextual awareness through motion classification can offer a framework for pHRI optimization and safety through predictive motion behaviour with a future goal of autonomous pHRI. Specifically, pHRI distinction and adaptable in-process motion tuning would greatly augment the previously cited literature.

As noted in Section 1.5.2, manual pHRI tuning is time intensive, not sufficient to obtain optimal gain values, and difficult without *a priori* knowledge of the environment dynamics. While the current literature often neglects motion tuning, Chapter 4 offers a comprehensive solution for both tuning and perception in the context of robot-assisted ultrasound pHRI. The integration of this work with external measurement and tracking devices, such as presented in Santos & Cortesao (2015), Mebarki *et al.* (2010), Graumann *et al.* (2016), and Huang *et al.* (2018), would offer all the safety fundamentals for a well-designed pHRI application.

5.2 Limitations and Recommendations

While the contributions of this thesis were developed around gaps in the current robot-assisted ultrasound literature, the limitations of the presented frameworks and innovations must be addressed. Most notably, the methods proposed in this thesis were experimentally tested and proven in laboratory settings. While a human leg phantom provides a more realistic model than simulation or non-anthropomorphic items, it cannot completely replace the fidelity of a clinical scenario. A possible next step for this work would be to integrate the methods into a clinical scenario. This could be done by replicating existing work, such as Mathiassen *et al.* (2016) or Conti *et al.* (2014), and measure performance improvements.

Likewise, the contact distinction framework presented in Chapter 3 has only been trained and tested with human leg regions from a single leg phantom. This limitation could be addressed by procuring different phantoms of varying body regions, training and testing on different body regions, and training with human subjects.

The calibration method proposed in Chapter 2 relies on the integration of additional equipment into the clinical environment. While machine tool components come in stainless steel varieties and can be sterilized, there may be an impedance to adding more items into the already meticulous clinical workspace.

Another limitation of the methods proposed in Chapter 3 and Chapter 4 is the need for training data. For safety and efficiency reasons, a human subject cannot act as the training model for the robotic system, nor can a medical professional wait for a system to be trained prior to starting a task. Future work should evaluate the transferability of training data from phantom models (such as that presented here) and simulation to human models.

Given these limitations, the foundation of future research can be defined. First, there is a significant amount of high-quality data available in real-world pHRI applications that is unused

or not shared across domains. A promising extension of this research would be to investigate the cross-training of machine learning models from different pHRI scenarios to develop pre-trained modules (i.e., reaction behaviours, trajectory control procedures) that can be integrated into existing applications in a plug and play manner. The overall goal, much like this thesis, would be to build a reusable, adaptable, and scalable framework for all pHRI applications.

Second, much of the literature has yet to address other humans involved in the application. Aside from the pHRI between the ultrasound subject and the robot, there is still the concern of how such intelligent systems will respond and adapt to a collaborative co-worker (e.g., a sonographer) and a moving patient. While collaborative hand-guiding was presented in Section 2.4.2 and Section 4.2, there are many more types of collaborative events that may occur. Particularly, in a medical environment the ultrasound process may be just one of several parallel procedures taking place on the subject at once, not all of which are necessarily robotic.

Third and in a like manner, failing gracefully in a medical environment is a topic of interest in the context of automated pHRI. Depending on the context, a medical procedure can quickly turn from benign to critical in a matter of moments. Beyond the intrinsic and extrinsic design requirements for safety, all medical systems can be considered safety-critical. While the contact and motion distinction methods presented in Chapter 3 and Chapter 4 offer a foundation for intelligent perception, ultrasound application events are not the only concern, but external issues or necessary human intervention may also be the cause of reaction events. While the problem space may appear to be infinite, future work exploring potential adaptive solutions would be very interesting.

Finally, given a fully functional, intelligent, and adaptive robot-assisted ultrasound platform, how will patients and medical professionals trust it? Patients may fear harm. Medical professionals may not trust that the data and automated motion trajectories are valid versus their own abilities. While robot application developers can claim performance metrics, there is still much work

to be done to integrate robotic systems into the anthropomorphic domain and human social psychology.

APPENDIX I

PUBLICATIONS DURING PHD STUDY

Kwiatkowski, J., Roberge, J., Nadeau, N. A., L'Écuyer-Lapierre, L. & Duchaine, V. (2018, Oct). An Extrinsic Dexterity Approach to the IROS 2018 Fan Robotic Challenge. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4139-4144. doi: 10.1109/IROS.2018.8594224

Nadeau, N. A. & Bonev, I. A. (2018, Oct). Evolutionary Motion Control Optimization in Physical Human-Robot Interaction. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1347-1353. doi: 10.1109/IROS.2018.8593598

Nadeau, N. A., Bonev, I. A. & Joubair, A. (2019). Impedance Control Self-Calibration of a Collaborative Robot Using Kinematic Coupling. *Robotics*, 8(2). doi: 10.3390/robotics8020033

Nadeau, N. A. (2019). Pybotics: Python Toolbox for Robotics. *Journal of Open Source Software*, 4(41), 1738. doi: 10.21105/joss.01738

APPENDIX II

ROBOTIC ULTRASOUND LITERATURE

Table-A II-1 Summary of current robotic ultrasound literature

Reference	Primary Topics	Comments
Boudet *et al.* (1997)	Force Control; Robot design	Mitsubishi PA10 used to quantify atheromatous plaques.
Pierrot *et al.* (1999)	Force Control; Robot design	Introduction of custom *Hippocrate* robot.
Salcudean *et al.* (1999)	Robot design	Custom counterbalanced design that focuses on intrinsic safety.
Salcudean *et al.* (2000)	Robot design; Visual servoing	Follow up to Salcudean *et al.* (1999) using teleoperation.
Zhu *et al.* (2000)	Robot design; Visual servoing	Same custom robot as Salcudean *et al.* (1999) but focused on controller design.
Abolmaesumi *et al.* (2001)	Robot design; Visual servoing	Same custom robot as Salcudean *et al.* (1999) but focused on user interface design.
Abolmaesumi *et al.* (2002)	Robot design; Visual servoing	Same custom robot as Salcudean *et al.* (1999) but focused on visual servoing controller.
Vilchis *et al.* (2003)	Robot design; Teleoperation	Introduction of custom *TER* robot for isolated sites or remote health centres.
Delgorge *et al.* (2005)	Robot design; Teleoperation	Introduction of custom teleoperated and mobile robot, *OTELO*.
Lessard *et al.* (2006)	Robot design	Introduction of custom *MedRUE* robot.
Lessard *et al.* (2007a)	Robot design	Follow up to Lessard *et al.* (2006), focusing on design optimization.
Lessard *et al.* (2007b)	Robot design	Follow up to Lessard *et al.* (2006), focusing on design optimization.
Martinelli *et al.* (2007)	Robot design; Teleoperation	Clinical evaluation of custom robot presented in Vilchis *et al.* (2003).
Courreges *et al.* (2008)	Robot design; Teleoperation	Survey of teleoperated robotic ultrasound with strong focus on the *OTELO* robot.
Janvier *et al.* (2008)	Medical imaging	Feasibility study of robotic ultrasound using a CRS F3 robot with little focus on control.

Table-A II-1 Summary of current robotic ultrasound literature (continued)

Reference	Primary Topics	Comments
Janvier *et al.* (2010)	Medical imaging	Follow up to Janvier *et al.* (2008) focusing on image reconstruction.
Mebarki *et al.* (2010)	Visual servoing	Ultrasound-based visual servoing using a robot like the *Hippocrate*.
Nakadate *et al.* (2011)	Visual servoing	A custom robotic device for ultrasound-based visual servoing.
Ito *et al.* (2013)	Robot design; Teleoperation	A custom, patient-mounted robot for emergency medical care during transport to a hospital.
Mustafa *et al.* (2013)	Force Control; Medical imaging	Used a Mitsubishi MELFA RV-1 robot for liver imaging with a control algorithm for compensation of patient movement due to breathing.
Nadeau & Krupa (2013)	Medical imaging; Visual servoing	Ultrasound-based visual servoing using an Omron Adept Viper.
Vieyres *et al.* (2003)	Robot design; Teleoperation	Introduction of custom *TERESA* robot.
Aalamifar *et al.* (2014)	Calibration; Medical imaging	Camera-based calibration of robotic ultrasound system (using a UR5) with low-accuracy results.
Arbeille *et al.* (2014)	Robot design; Teleoperation	Follow up to Vieyres *et al.* (2003) where teleoperation is performed over an Internet connection.
Conti *et al.* (2014)	Teleoperation	A teleoperated, haptic feedback ultrasound robot based upon the KUKA LWR but the study focuses on the user interface.
Janvier *et al.* (2014)	Medical imaging	Follow up to Janvier *et al.* (2008) focusing on image quality.
Gilbertson & Anthony (2015)	Force control; Robot design	Hand-held custom device for freehand ultrasound but the device tends to chatter when placed in contact with non-soft tissues.
Merouche *et al.* (2015a)	Medical imaging; Visual servoing	Ultrasound-based visual servoing with a CRS F3 robot but no controller discussions.

Table-A II-1 Summary of current robotic ultrasound literature (continued)

Reference	**Primary Topics**	**Comments**
Pahl & Supriyanto (2015)	Robot design; Teleoperation	Custom and bulky gantry robot for transabdominal ultrasound imaging.
Santos & Cortesao (2015)	Teleoperation; Visual servoing	Servoing with a 3D time-of-flight camera with a Barrett Technology WAM robot.
Sen *et al.* (2015)	Robot design; Teleoperation	A bulky and passive custom design for collaborative probe manipulation. It has limited mobility, cannot self navigate, and requires optical tracking.
Seo *et al.* (2015)	Robot design; Teleoperation	A custom teleoperated Stewart platform for probe manipulation in paramedic and remote doctor scenarios.
Avgousti *et al.* (2016)	Teleoperation	Cardiac ultrasonography over 4G wireless networks using an *OTELO*-based robot.
Chatelain *et al.* (2016)	Medical imaging; Visual servoing	Ultrasound image optimization using an Omron Adept Viper s650.
Graumann *et al.* (2016)	Visual servoing	A KUKA LBR iiwa robot under impedance control and with a depth camera for vision-based trajectory planning.
Hennersperger *et al.* (2016)	Visual servoing	Vision-based trajectory planning with structured light and a Kinect 3D camera but poor accuracy.
Joubair *et al.* (2016)	Calibration	Simulated force-based calibration of *MedRUE* robot.
Mathiassen *et al.* (2016)	Teleoperation	Feasibility robotic ultrasound study using a UR5 robot with little discussion of methodology.
Virga *et al.* (2016)	Medical imaging; Visual servoing	Online probe force estimation framework to optimize the visibility of the aorta but requires the use of an external vision system and co-registration with a MR-atlas.

Table-A II-1 Summary of current robotic ultrasound literature (continued)

Reference	Primary Topics	Comments
Wang *et al.* (2016)	Medical imaging; Robot design; Visual servoing	Custom robot for ultrasound-based probe tracking but safety and force sensing are not discussed.
Wu & Huang (2016)	Robot design; Visual servoing	Custom gantry design that uses a Kinect 3D camera for visual servoing but no controller or accuracy discussions.
Zettinig *et al.* (2016)	Medical imaging; Visual servoing	Ultrasound-based visual servoing feasibility study using a KUKA LBR iiwa 7 R800 robot.
Aalamifar *et al.* (2017)	Visual servoing	Follow up to Aalamifar *et al.* (2014) with the same limitations.
Chatelain *et al.* (2017)	Medical imaging; Visual servoing	Follow up to Chatelain *et al.* (2016) using a KUKA LWR iiwa R800 to optimize the ultrasound acoustic window.
Fang *et al.* (2017)	Force control; Robot design	Force-controlled UR5 but no discussion of how controller gains are selected or tuned.
Guan *et al.* (2017)	Robot design; Teleoperation	Custom robot for teleoperation through the Internet.
Seo *et al.* (2017)	Robot design; Teleoperation	Follow up to Seo *et al.* (2015) focusing on experimental evaluations but with low-accuracy results.
Sharifi *et al.* (2017)	Robot design; Teleoperation	Custom benchtop robot to study impedance control and haptics during teleoperation.
Du *et al.* (2018)	Medical imaging	Feasibility study using a small, non-industrial, 3-axis robot to hold ultrasound probe during stenosis evaluation.
Huang *et al.* (2018)	Visual servoing	A visual servoing feasibility study using a Kinect 3D camera and checkerboard calibration but limited to slow speeds and no gains tuning.
Karar (2018)	Force control	One of the only studies to identify the difficulty in controller gain tuning but the study is simulation only.

Table-A II-1 Summary of current robotic ultrasound literature (continued)

Reference	**Primary Topics**	**Comments**
Lan & Huang (2018)	Robot design; Visual servoing	Follow up to Wu & Huang (2016) with the same limitations.
Santos & Cortesão (2018)	Teleoperation; Visual servoing	Follow up to Santos & Cortesao (2015) where the contact event is modelled as a linear spring.
Victorova *et al.* (2019)	Medical imaging	Feasibility study with UR3 robot and basic force control without controller tuning methodology.

REFERENCES

Aalamifar, F., Khurana, R., Cheng, A., Taylor, R. H., Iordachita, I. & Boctor, E. M. (2014). Enabling technologies for robot assisted ultrasound tomography: system setup and calibration. *Medical Imaging 2014: Ultrasonic Imaging and Tomography*, 9040, 90401X.

Aalamifar, F., Khurana, R., Cheng, A., Guo, X., Iordachita, I. & Boctor, E. M. (2017). Enabling technologies for robot assisted ultrasound tomography. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 13(1), e1746.

Abolmaesumi, P., Salcudean, S. S., Zhu, W., Sirouspour, M., DiMaio, S., Wen-Hong Zhu, W.-H., Sirouspour, M. & DiMaio, S. (2002). Image-guided control of a robot for medical ultrasound. *IEEE Transactions on Robotics and Automation*, 18(1), 11–23. doi: 10.1109/70.988970.

Abolmaesumi, P., Salcudean, S., Zhu, W.-H., DiMaio, S. P. & Sirouspour, M. R. (2001). A user interface for robot-assisted diagnostic ultrasound. *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No. 01CH37164)*, 2, 1549–1554.

Aboyans, V., Ricco, J.-B., Bartelink, M.-L. E., Björck, M., Brodmann, M., Cohnert, T., Collet, J.-P., Czerny, M., De Carlo, M., Debus, S. et al. (2017). 2017 ESC guidelines on the diagnosis and treatment of peripheral arterial diseases, in collaboration with the European Society for Vascular Surgery (ESVS) document covering atherosclerotic disease of extracranial carotid and vertebral, mesenteric, renal, upper and lower extremity arteries endorsed by: the European stroke organization (ESO) the task force for the diagnosis and treatment of peripheral arterial diseases of the European Society of Cardiology (ESC) and of the European Society for Vascular Surgery (ESVS). *European heart journal*, 39(9), 763–816.

Ajoudani, A., Zanchettin, A. M., Ivaldi, S., Albu-Schäffer, A., Kosuge, K. & Khatib, O. (2018). Progress and prospects of the human–robot collaboration. *Autonomous Robots*, 42(5), 957–975.

Alami, R., Albu-Schäffer, A., Bicchi, A., Bischoff, R., Chatila, R., De Luca, A., De Santis, A., Giralt, G., Guiochet, J., Hirzinger, G. et al. (2006). Safe and dependable physical human-robot interaction in anthropic domains: State of the art and challenges. *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1–16.

Albu-Schäffer, A., Ott, C. & Hirzinger, G. (2007). A Unified Passivity-based Control Framework for Position, Torque and Impedance Control of Flexible Joint Robots. *The International Journal of Robotics Research*, 26(1), 23–39. doi: 10.1177/0278364907073776.

Almeida, F., Lopes, A. & Abreu, P. (1999). Force-impedance control: a new control strategy of robotic manipulators. *Recent advances in Mechatronics*, 1, 126–137.

Alqaudi, B., Modares, H., Ranatunga, I., Tousif, S. M., Lewis, F. L. & Popa, D. O. (2016). Model reference adaptive impedance control for physical human-robot interaction. *Control Theory and Technology*, 14(1), 68–82. doi: 10.1007/s11768-016-5138-2.

Arbeille, P., Provost, R., Zuj, K., Dimouro, D. & Georgescu, M. (2014). Teles-operated Echocardiography Using a Robotic Arm and an Internet Connection. *Ultrasound in Medicine & Biology*, 40(10), 2521–2529. doi: 10.1016/j.ultrasmedbio.2014.05.015.

Avgousti, S., Panayides, A. S., Jossif, A. P., Christoforou, E. G., Vieyres, P., Novales, C., Voskarides, S. & Pattichis, C. S. (2016). Cardiac ultrasonography over 4G wireless networks using a tele-operated robot. *Healthcare technology letters*, 3(3), 212–217.

Barros-Gomes, S., Orme, N., Nhola, L. F., Scott, C., Helfinstine, K., Pislaru, S. V., Kane, G. C., Singh, M. & Pellikka, P. A. (2019). Characteristics and Consequences of Work-Related Musculoskeletal Pain among Cardiac Sonographers Compared with Peer Employees: A Multisite Cross-Sectional Study. *Journal of the American Society of Echocardiography*.

Bengio, Y. (2012). Practical recommendations for gradient-based training of deep architectures. 437–478.

Besset, P., Olabi, A. & Gibaru, O. (2016). Advanced calibration applied to a collaborative robot. *Power Electronics and Motion Control Conference (PEMC), 2016 IEEE International*, pp. 662–667.

Bicchi, A., Peshkin, M. A. & Colgate, J. E. (2008). Safety for physical human–robot interaction. *Springer handbook of robotics*, 1335–1348.

Borm, J.-H. & Meng, C.-H. (1991). Determination of Optimal Measurement Configurations for Robot Calibration Based on Observability Measure. *The International Journal of Robotics Research*, 10(1), 51–63. doi: 10.1177/027836499101000106.

Boudet, S., Gariepy, J. & Mansour, S. (1997). An integrated robotics and medical control device to quantify atheromatous plaques: experiments on the arteries of a patient. *Proceedings of the 1997 IEEE/RSJ International Conference on Intelligent Robot and Systems. Innovative Robotics for Real-World Applications. IROS '97*, 3, 1533–1538. doi: 10.1109/IROS.1997.656562.

Bousquet-Jette, C., Achiche, S., Beaini, D., Cio, Y. L.-K., Leblond-Ménard, C. & Raison, M. (2017). Fast scene analysis using vision and artificial intelligence for object prehension by an assistive robot. *Engineering Applications of Artificial Intelligence*, 63, 33–44.

Bowthorpe, M. & Tavakoli, M. (2016). Ultrasound-Based Image Guidance and Motion Compensating Control for Robot-Assisted Beating-Heart Surgery. *Journal of Medical Robotics Research*, 01(01), 1640002. doi: 10.1142/S2424905X1640002X.

Buchli, J., Stulp, F., Theodorou, E. & Schaal, S. (2011). Learning variable impedance control. *The International Journal of Robotics Research*, 30(7), 820–833.

Chatelain, P., Krupa, A. & Navab, N. (2016). Confidence-driven control of an ultrasound probe: Target-specific acoustic window optimization. *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3441–3446.

Chatelain, P., Krupa, A. & Navab, N. (2017). Confidence-Driven Control of an Ultrasound Probe. *IEEE Transactions on Robotics*, 33(6), 1410–1424. doi: 10.1109/tro.2017.2723618.

Chiaverini, S. & Sciavicco, L. (1993). The parallel approach to force/position control of robotic manipulators. *IEEE Transactions on Robotics and Automation*, 9(4), 361–373. doi: 10.1109/70.246048.

Chiaverini, S., Siciliano, B. & Villani, L. (1999). A survey of robot interaction control schemes with experimental comparison. *IEEE/ASME Transactions On Mechatronics*, 4(3), 273–285.

Cho, C.-N., Kim, J.-H., Kim, Y.-L., Song, J.-B. & Kyung, J.-H. (2012). Collision detection algorithm to distinguish between intended contact and unexpected collision. *Advanced*

Robotics, 26(16), 1825–1840.

Cio, Y.-S. L.-K., Raison, M., Ménard, C. L. & Achiche, S. (2019). Proof of concept of an assistive robotic arm control using artificial stereovision and eye-tracking. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*.

Conti, F., Park, J. & Khatib, O. (2014). Interface Design and Control Strategies for a Robot Assisted Ultrasonic Examination System. In Khatib, O., Kumar, V. & Sukhatme, G. (Eds.), *Experimental Robotics SE - 7* (vol. 79, pp. 97–113). Springer Berlin Heidelberg. doi: 10.1007/978-3-642-28572-1_7.

Corrales, J., Candelas, F. & Torres, F. (2011). Safe human–robot interaction based on dynamic sphere-swept line bounding volumes. *Robotics and Computer-Integrated Manufacturing*, 27(1), 177–185.

Courreges, F., Vieyres, P. & Poisson, G. (2008). Robotized Tele-echography. In *Teleradiology* (pp. 139–154). Springer Berlin Heidelberg. doi: 10.1007/978-3-540-78871-3_13.

Craig, J. J. (2005). *Introduction to Robotics Mechanics and Control* (ed. 3). Pearson Education International.

Daneshmand, M., Bilici, O., Bolotnikova, A. & Anbarjafari, G. (2017). Medical robots with potential applications in participatory and opportunistic remote sensing: A review. *Robotics and Autonomous Systems*, 95, 160–180.

Das, S. & Suganthan, P. N. (2011). Differential Evolution: A Survey of the State-of-the-Art. *IEEE Transactions on Evolutionary Computation*, 15(1), 4–31. doi: 10.1109/tevc.2010.2059031.

De Santis, A., Siciliano, B., De Luca, A. & Bicchi, A. (2008). An atlas of physical human–robot interaction. *Mechanism and Machine Theory*, 43(3), 253–270. doi: 10.1016/j.mechmachtheory.2007.03.003.

Delgorge, C., Courreges, F., AlBassit, L., Novales, C., Rosenberger, C., Smith-Guerin, N., Bru, C., Gilabert, R., Vannoni, M., Poisson, G., Vieyres, P., Courrèges, F., Al Bassit, L., Novales, C., Rosenberger, C., Smith-Guerin, N., Brù, C., Gilabert, R., Vannoni, M., Poisson, G., Vieyres, P., Courreges, F., AlBassit, L., Novales, C., Rosenberger, C., Smith-Guerin, N., Bru, C., Gilabert, R., Vannoni, M., Poisson, G. & Vieyres, P. (2005). A Tele-Operated Mobile Ultrasound Scanner Using a Light-Weight Robot. *IEEE Transactions on Information Technology in Biomedicine*, 9(1), 50–58. doi: 10.1109/TITB.2004.840062.

Dombre, E., Duchemin, G., Poignet, P. & Pierrot, F. (2003). Dermarob: A safe robot for reconstructive surgery. *IEEE Transactions on Robotics and Automation*, 19(5), 876–884.

Du, Y.-C., Shih, J.-B., Wu, M.-J. & Chiou, C.-Y. (2018). Development of an AVF Stenosis Assessment Tool for Hemodialysis Patients Using Robotic Ultrasound System. *Micromachines*, 9(2), 51.

Du, Z., Wang, W., Yan, Z., Dong, W. & Wang, W. (2017). Variable Admittance Control Based on Fuzzy Reinforcement Learning for Minimally Invasive Surgery Manipulator. *Sensors*, 17(12), 844. doi: 10.3390/s17040844.

Duchemin, G., Maillet, P., Poignet, P., Dombre, E. & Pierrot, F. (2005). A hybrid position/force control approach for identification of deformation models of skin and underlying tissues. *IEEE Transactions on Biomedical Engineering*, 52(2), 160–170.

Elatta, A. Y., Gen, L. P., Zhi, F. L., Daoyuan, Y. & Fei, L. (2004). An Overview of Robot Calibration_2004_china. *Information Technology*, 3(1), 74–78.

Fang, T.-Y., Zhang, H. K., Finocchi, R., Taylor, R. H. & Boctor, E. M. (2017). Force-assisted ultrasound imaging system through dual force sensing and admittance robot control. *International journal of computer assisted radiology and surgery*, 12(6), 983–991.

Fei, B., Ng, W. S., Chauhan, S. & Kwoh, C. K. (2001). The safety issues of medical robotics. *Reliability Engineering & System Safety*, 73(2), 183–192. doi: 10.1016/S0951-8320(01)00037-0.

Ficuciello, F., Villani, L. & Siciliano, B. (2015). Variable impedance control of redundant manipulators for intuitive human–robot physical interaction. *IEEE Transactions on Robotics*, 31(4), 850–863.

Flacco, F., Kröger, T., De Luca, A. & Khatib, O. (2012). A depth space approach to human-robot collision avoidance. *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pp. 338–345.

Fortin, F.-A., De Rainville, F.-M., Gardner, M.-A., Parizeau, M. & Gagné, C. (2012). DEAP: Evolutionary Algorithms Made Easy. *Journal of Machine Learning Research*, 13, 2171–2175.

Gaudreault, M., Joubair, A. & Bonev, I. A. (2016). Local and closed-loop calibration of an industrial serial robot using a new low-cost 3D measuring device. *Proceedings - IEEE International Conference on Robotics and Automation*, 2016-June, 4312–4319. doi: 10.1109/ICRA.2016.7487629.

Gaudreault, M., Joubair, A. & Bonev, I. (2018). Self-calibration of an industrial robot using a novel affordable 3D measuring device. *Sensors*, 18(10), 3380.

Geravand, M., Flacco, F. & De Luca, A. (2013). Human-robot physical interaction and collaboration using an industrial robot with a closed control architecture. *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pp. 4000–4007.

Gilbertson, M. W. & Anthony, B. W. (2015). Force and position control system for freehand ultrasound. *IEEE Transactions on Robotics*, 31(4), 835–849.

Golz, S., Osendorfer, C. & Haddadin, S. (2015). Using tactile sensation for learning contact knowledge: Discriminate collision from physical interaction. *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pp. 3788–3794.

Graumann, C., Fuerst, B., Hennersperger, C., Bork, F. & Navab, N. (2016, may). Robotic ultrasound trajectory planning for volume of interest coverage. *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 736–741. doi: 10.1109/ICRA.2016.7487201.

Gribovskaya, E., Kheddar, A. & Billard, A. (2011, may). Motion learning and adaptive impedance for robot control during physical interaction with humans. *2011 IEEE International Conference on Robotics and Automation*. doi: 10.1109/icra.2011.5980070.

Guan, X., Wu, H., Hou, X., Teng, Q., Wei, S., Jiang, T., Zhang, J., Wang, B., Yang, J. & Xiong, L. (2017). Study of a 6DOF robot assisted ultrasound scanning system and its simulated control handle. *2017 IEEE International Conference on Cybernetics and Intelligent Systems (CIS) and IEEE Conference on Robotics, Automation and Mechatronics (RAM)*, pp. 469–474.

Haddadin, S., Albu-Schaffer, A., De Luca, A. & Hirzinger, G. (2008). Collision detection and reaction: A contribution to safe physical human-robot interaction. *Intelligent Robots and*

Systems, 2008. IROS 2008. IEEE/RSJ International Conference on, pp. 3356–3363.

Haddadin, S., Suppa, M., Fuchs, S., Bodenmüller, T., Albu-Schäffer, A. & Hirzinger, G. (2011). Towards the robotic co-worker. In *Robotics Research* (pp. 261–282). Springer.

Hassard, J., Teoh, K. R., Visockaite, G., Dewe, P. & Cox, T. (2018). The cost of work-related stress to society: A systematic review. *Journal of occupational health psychology*, 23(1), 1.

Hennersperger, C., Fuerst, B., Virga, S., Zettinig, O., Frisch, B., Neff, T. & Navab, N. (2016). Towards MRI-based autonomous robotic US acquisitions: a first feasibility study. *IEEE transactions on medical imaging*, 36(2), 538–548.

Huang, Q. & Zeng, Z. (2017). A review on real-time 3D ultrasound imaging technology. *BioMed research international*, 2017.

Huang, Q., Lan, J. & Li, X. (2018). Robotic arm based automatic ultrasound scanning for three-dimensional imaging. *IEEE Transactions on Industrial Informatics*, 15(2), 1173–1182.

Inoue, T., Magistris, G. D., Munawar, A., Yokoya, T. & Tachibana, R. (2017, Sept). Deep reinforcement learning for high precision assembly tasks. *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 819-825. doi: 10.1109/IROS.2017.8202244.

Ito, K., Sugano, S., Takeuchi, R., Nakamura, K. & Iwata, H. (2013). Usability and performance of a wearable tele-echography robot for focused assessment of trauma using sonography. *Medical Engineering & Physics*, 35(2), 165–171. doi: 10.1016/j.medengphy.2012.04.011.

Janvier, M.-A., Durand, L.-G., Cardinal, M.-H. R., Renaud, I., Chayer, B., Bigras, P., de Guise, J., Soulez, G. & Cloutier, G. (2008). Performance evaluation of a medical robotic 3D-ultrasound imaging system. *Medical image analysis*, 12(3), 275–90. doi: 10.1016/j.media.2007.10.006.

Janvier, M.-A., Soulez, G., Allard, L. & Cloutier, G. (2010). Validation of 3D reconstructions of a mimicked femoral artery with an ultrasound imaging robotic system. *Medical physics*, 37(7), 3868–79. doi: 10.1118/1.3447721.

Janvier, M.-A., Merouche, S., Allard, L., Soulez, G. & Cloutier, G. (2014). A 3-D Ultrasound Imaging Robotic System to Detect and Quantify Lower Limb Arterial Stenoses: In Vivo Feasibility. *Ultrasound in Medicine & Biology*, 40(1), 232–243. doi: 10.1016/j.ultrasmedbio.2013.08.010.

Jeon, D. & Tomizuka, M. (1993). Learning hybrid force and position control of robot manipulators. *IEEE Transactions on Robotics and Automation*, 9(4), 423–431. doi: 10.1109/70.246053.

Jones, E., Oliphant, T. & Peterson, P. (2001). SciPy: Open source scientific tools for Python.

Joubair, A. & Bonev, I. A. (2015a). Non-kinematic calibration of a six-axis serial robot using planar constraints. *Precision Engineering*, 40, 325–333. doi: 10.1016/j.precisioneng.2014.12.002.

Joubair, A. & Bonev, I. A. (2015b). Kinematic calibration of a six-axis serial robot using distance and sphere constraints. *International Journal of Advanced Manufacturing Technology*, 77(1-4), 515–523. doi: 10.1007/s00170-014-6448-5.

Joubair, A., Slamani, M. & Bonev, I. A. (2012). A novel XY-Theta precision table and a geometric procedure for its kinematic calibration. *Robotics and Computer-Integrated Manufacturing*, 28(1), 57–65. doi: 10.1016/j.rcim.2011.06.006.

Joubair, A., Slamani, M. & Bonev, I. A. (2013). Kinematic calibration of a five-bar planar parallel robot using all working modes. *Robotics and Computer-Integrated Manufacturing*, 29(4), 15–25.

Joubair, A., Zhao, L., Bigras, P. & Bonev, I. (2016). Use of a Force-Torque Sensor for Self-Calibration of a 6-DOF Medical Robot. *Sensors*, 16(6), 798. doi: 10.3390/s16060798.

Jung, S., Hsia, T. C. & Bonitz, R. G. (2004). Force tracking impedance control of robot manipulators under unknown environment. *IEEE Transactions on Control Systems Technology*, 12(3), 474–483.

Karar, M. E. (2018). A simulation study of adaptive force controller for medical robotic liver ultrasound guidance. *Arabian Journal for Science and Engineering*, 43(8), 4229–4238.

Kelly, R. (1995). A tuning procedure for stable PID control of robot manipulators. *Robotica*, 13(2), 141–148. doi: 10.1017/S0263574700017641.

Kock, S., Vittor, T., Matthias, B., Jerregard, H., Källman, M., Lundberg, I., Mellander, R. & Hedelind, M. (2011). Robot concept for scalable, flexible assembly automation: A technology study on a harmless dual-armed robot. *Assembly and Manufacturing (ISAM), 2011 IEEE International Symposium on*, pp. 1–5.

Kohavi, R. et al. (1995). A study of cross-validation and bootstrap for accuracy estimation and model selection. *Ijcai*, 14(2), 1137–1145.

Koizumi, N., Warisawa, S., Nagoshi, M., Hashizume, H. & Mitsuishi, M. (2009). Construction Methodology for a Remote Ultrasound Diagnostic System. *IEEE Transactions on Robotics*, 25(3), 522-538. doi: 10.1109/TRO.2009.2019785.

Kouris, A., Dimeas, F. & Aspragathos, N. (2018). A Frequency Domain Approach for Contact Type Distinction in Human-Robot Collaboration. *IEEE Robotics and Automation Letters*, 3(2), 720-727. doi: 10.1109/LRA.2017.2789249.

Kouris, A., Dimeas, F. & Aspragathos, N. (2016). Contact distinction in human-robot cooperation with admittance control. *Systems, Man, and Cybernetics (SMC), 2016 IEEE International Conference on*, pp. 001951–001956.

Kwiatkowski, J., Roberge, J., Nadeau, N. A., L'Écuyer-Lapierre, L. & Duchaine, V. (2018, Oct). An Extrinsic Dexterity Approach to the IROS 2018 Fan Robotic Challenge. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4139-4144. doi: 10.1109/IROS.2018.8594224.

Lan, J. & Huang, Q. (2018). Automatic Three-Dimensional Ultrasound Scanning System Based on RGB-D Camera. *2018 2nd International Conference on Robotics and Automation Sciences (ICRAS)*, pp. 1–5.

Lasota, P. A., Fong, T., Shah, J. A. et al. (2017). A survey of methods for safe human-robot interaction. *Foundations and Trends® in Robotics*, 5(4), 261–349.

Lee, J., Kao, H.-A. & Yang, S. (2014). Service innovation and smart analytics for industry 4.0 and big data environment. *Procedia Cirp*, 16, 3–8.

Lee, J., Bagheri, B. & Kao, H.-A. (2015). A cyber-physical systems architecture for industry 4.0-based manufacturing systems. *Manufacturing Letters*, 3, 18–23.

Lessard, S., Bonev, I., Bigras, P., Durand, L.-G., Soulez, G., Cloutier, G. & De Guise, J. A. (2006). Parallel robot for medical 3D-ultrasound imaging. *2006 IEEE International Symposium on Industrial Electronics*, 4, 3102–3107.

Lessard, S., Bigras, P. & Bonev, I. A. (2007a). A new medical parallel robot and its static balancing optimization. *Journal of medical Devices*, 1(4), 272–278.

Lessard, S., Bonev, I., Bigras, P., Briot, S. & Arakelyan, V. (2007b). Optimum static balancing of the parallel robot for medical 3D-ultrasound imaging. *IFTOMM 2007: 12th World Congress in Mechanism and Machine Science*.

Lu, S., Chung, J. H. & Velinsky, S. A. (2005). Human-robot collision detection and identification based on wrist and base force/torque sensors. *Proceedings - IEEE International Conference on Robotics and Automation*, 2005(April), 3796–3801. doi: 10.1109/ROBOT.2005.1570699.

Lutz, H. & Buscarini, E. (2011). *Manual of diagnostic ultrasound*. World Health Organization.

Martinelli, T., Bosson, J.-L., Bressollette, L., Pelissier, F., Boidard, E., Troccaz, J. & Cinquin, P. (2007). Robot-based tele-echography: clinical evaluation of the TER system in abdominal aortic exploration. *Journal of ultrasound in medicine : official journal of the American Institute of Ultrasound in Medicine*, 26(11), 1611–6.

Mason, M. T. (1981). Compliance and Force Control for Computer Controlled Manipulators. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(6), 418–432. doi: 10.1109/TSMC.1981.4308708.

Mathiassen, K., Fjellin, J. E., Glette, K., Hol, P. K. & Elle, O. J. (2016). An ultrasound robotic system using the commercial robot UR5. *Frontiers in Robotics and AI*, 3, 1.

Mattos, L. S., Deshpande, N., Barresi, G., Guastini, L. & Peretti, G. (2014). A novel computerized surgeon–machine interface for robot-assisted laser phonomicrosurgery. *The Laryngoscope*, 124(8), 1887–1894.

Maurtua, I., Ibarguren, A., Kildal, J., Susperregi, L. & Sierra, B. (2017). Human–robot collaboration in industrial applications: Safety, interaction and trust. *International Journal of Advanced Robotic Systems*, 14(4), 1729881417716010.

Mebarki, R., Krupa, A. & Chaumette, F. (2010). 2-D Ultrasound Probe Complete Guidance by Visual Servoing Using Image Moments. *IEEE Transactions on Robotics*, 26(2), 296–306. doi: 10.1109/TRO.2010.2042533.

Meggiolaro, M., Scriffignano, G. & Dubowsky, S. (2000). Manipulator calibration using a single endpoint contact constraint. *Proceedings of the 26th Biennial Mechanisms and Robotics Conference of the 2000 ASME Design Engineering Technical Conferences*, 1–9.

Merouche, S., Allard, L., Montagnon, E., Soulez, G., Bigras, P. & Cloutier, G. (2015a). A robotic ultrasound scanner for automatic vessel tracking and three-dimensional reconstruction of B-mode images. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 63(1), 35–46.

Merouche, S., Allard, L., Montagnon, E., Soulez, G., Bigras, P. & Cloutier, G. (2015b). A robotic ultrasound scanner for automatic vessel tracking and three-dimensional reconstruction of B-mode images. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 63(1), 35–46.

Monfaredi, R., Wilson, E., Azizi koutenaei, B., Labrecque, B., Leroy, k., Goldie, J., Louis, E., Swerdlow, D. & Cleary, K. (2015). Robot-assisted ultrasound imaging: Overview and development of a parallel telerobotic system. *Minimally Invasive Therapy & Allied Technologies*, 24(1), 54–62.

Mooring, B. W., Roth, Z. S. & Driels, M. R. (1991). Fundamentals of manipulator calibration. Wiley-interscience.

Morato, C., Kaipa, K. N., Zhao, B. & Gupta, S. K. (2014). Toward safe human robot collaboration by using multiple kinects based real-time human tracking. *Journal of Computing and Information Science in Engineering*, 14(1), 011006.

Mustafa, A. S. B., Ishii, T., Matsunaga, Y., Nakadate, R., Ishii, H., Ogawa, K., Saito, A., Sugawara, M., Niki, K. & Takanishi, A. (2013, dec). Development of robotic system for autonomous liver screening using ultrasound scanning device. *2013 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pp. 804–809. doi: 10.1109/ROBIO.2013.6739561.

Nadeau, C. & Krupa, A. (2013). Intensity-based ultrasound visual servoing: Modeling and validation with 2-d and 3-d probes. *IEEE Transactions on Robotics*, 29(4), 1003–1015.

Nadeau, N. A. & Bonev, I. A. (2018, Oct). Evolutionary Motion Control Optimization in Physical Human-Robot Interaction. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1347-1353. doi: 10.1109/IROS.2018.8593598.

Nadeau, N. A. (2019). Pybotics: Python Toolbox for Robotics. *Journal of Open Source Software*, 4(41), 1738. doi: 10.21105/joss.01738.

Nadeau, N. A., Bonev, I. A. & Joubair, A. (2019). Impedance Control Self-Calibration of a Collaborative Robot Using Kinematic Coupling. *Robotics*, 8(2). doi: 10.3390/robotics8020033.

Nakadate, R., Solis, J., Takanishi, A., Minagawa, E., Sugawara, M. & Niki, K. (2011, may). Out-of-plane visual servoing method for tracking the carotid artery with a robot-assisted ultrasound diagnostic system. *2011 IEEE International Conference on Robotics and Automation*, pp. 5267–5272. doi: 10.1109/ICRA.2011.5979594.

Neelakanta, P. S. & Dighe, H. (2003). Robust factory wireless communications: a performance appraisal of the Bluetooth/spl trade/and the ZigBee/spl trade/colocated on an industrial floor. *Industrial Electronics Society, 2003. IECON'03. The 29th Annual Conference of the IEEE*, 3, 2381–2386.

Nelder, J. A. & Mead, R. (1965). A simplex method for function minimization. *The computer journal*, 7(4), 308–313.

Niccolini, M., Castelli, V., Diversi, C., Kang, B., Mussa, F. & Sinibaldi, E. (2016). Development and preliminary assessment of a robotic platform for neuroendoscopy based on a lightweight robot. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 12(1), 4–17.

Nubiola, A. & Bonev, I. A. (2013). Absolute calibration of an ABB IRB 1600 robot using a laser tracker. *Robotics and Computer-Integrated Manufacturing*, 29(1), 236–245.

Nubiola, A., Slamani, M., Joubair, A. & Bonev, I. A. (2014). Comparison of two calibration methods for a small industrial robot based on an optical CMM and a laser tracker. *Robotica*, 32(3), 447–466.

Olson, R. S., Bartley, N., Urbanowicz, R. J. & Moore, J. H. (2016). Evaluation of a Tree-based Pipeline Optimization Tool for Automating Data Science. *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, (GECCO '16), 485–492. doi: 10.1145/2908812.2908918.

Ott, C., Mukherjee, R. & Nakamura, Y. (2015). A Hybrid System Framework for Unified Impedance and Admittance Control. *Journal of Intelligent & Robotic Systems*, 78(3-4),

359–375. doi: 10.1007/s10846-014-0082-1.

Pahl, C. & Supriyanto, E. (2015). Design of automatic transabdominal ultrasound imaging system. *2015 20th International Conference on Methods and Models in Automation and Robotics (MMAR)*, pp. 435–440.

Papageorgiou, X. S., Chalvatzaki, G., Dometios, A. C. & Tzafestas, C. S. (2018). Human-centered service robotic systems for assisted living. *International Conference on Robotics in Alpe-Adria Danube Region*, pp. 132–140.

Parraga, G., Fenster, A., Krasinski, A., Chiu, B., Egger, M. & Spence, J. (2011). 3D Carotid Ultrasound Imaging. In Suri, J. S., Kathuria, C. & Molinari, F. (Eds.), *Atherosclerosis Disease Management SE - 12* (pp. 325–350). Springer New York. doi: 10.1007/978-1-4419-7222-4_12.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V. et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12(Oct), 2825–2830.

Pervez, A. & Ryu, J. (2008). Safe physical human robot interaction-past, present and future. *Journal of Mechanical Science and Technology*, 22(3), 469.

Petrovic, P., Lukic, N. & Danilov, I. (2014). Robot-Assisted 3D Medical Sonography. In *New Trends in Medical and Service Robots* (pp. 45–61). Springer.

Pierrot, F., Dombre, E., Dégoulange, E., Urbain, L., Caron, P., Boudet, S., Gariépy, J. & Mégnien, J.-L. (1999). Hippocrate: A safe robot arm for medical applications with force feedback. *Medical Image Analysis*, 3(3), 285–300.

Priester, A. M., Natarajan, S. & Culjat, M. O. (2013). Robotic ultrasound systems in medicine. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 60(3), 507–523.

Quigley, M., Asbeck, A. & Ng, A. (2011). A low-cost compliant 7-DOF robotic manipulator. *2011 IEEE International Conference on Robotics and Automation*, 6051–6058. doi: 10.1109/ICRA.2011.5980332.

Raibert, M. H. & Craig, J. J. J. (1981). Hybrid Position/Force Control of Manipulators. *Journal of Dynamic Systems, Measurement, and Control*, 103(2), 126. doi: 10.1115/1.3139652.

Rezazadegan, F., Gengb, J., Ghirardi, M., Menga, G., Murèb, S., Camuncolib, G. & Demichelac, M. (2015). Risked-based design for the physical human-robot interaction (pHRI): An overview. *Chemical Engineering Transactions*, 43, 1249–1254.

Roberge, J.-P., Rispal, S., Wong, T. & Duchaine, V. (2016). Unsupervised feature learning for classifying dynamic tactile events using sparse coding. *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2675–2681.

Rosenstrauch, M. J. & Krüger, J. (2017). Safe human-robot-collaboration-introduction and experiment using ISO/TS 15066. *Control, Automation and Robotics (ICCAR), 2017 3rd International Conference on*, pp. 740–744.

Roveda, L., Iannacci, N., Vicentini, F., Pedrocchi, N., Braghin, F. & Tosatti, L. M. (2016). Optimal Impedance Force-Tracking Control Design With Impact Formulation for Interaction Tasks. *IEEE Robotics and Automation Letters*, 1(1), 130–136. doi: 10.1109/lra.2015.2508061.

Sage, H. G., Mathelin, M. F. D. & Ostertag, E. (1999). Robust control of robot manipulators: A survey. *International Journal of Control*, 72(16), 1498–1522. doi: 10.1080/002071799220137.

Salcudean, S. E., Bell, G., Bachmann, S., Zhu, W. H., Abolmaesumi, P. & Lawrence, P. D. (1999). Robot-Assisted Diagnostic Ultrasound – Design and Feasibility Experiments. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI'99* (pp. 1062–1071). Springer Berlin Heidelberg. doi: 10.1007/10704282_115.

Salcudean, S. E., Zhu, W. H., Abolmaesumi, P., Bachmann, S. & Lawrence, P. D. (2000). A robot system for medical ultrasound. In *Robotics Research* (pp. 195–202). Springer.

Sander, A. & Wolfgang, M. (2014). The Rise of Robotics.

Santos, L. & Cortesao, R. (2015, may). A dynamically consistent hierarchical control architecture for robotic-assisted tele-echography with motion and contact dynamics driven by a 3D time-of-flight camera and a force sensor. *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2931–2937. doi: 10.1109/ICRA.2015.7139600.

Santos, L. & Cortesão, R. (2018). Computed-Torque Control for Robotic-Assisted Tele-Echography Based on Perceived Stiffness Estimation. *IEEE Transactions on Automation Science and Engineering*, 15(3), 1337–1354.

Scassellati, B. & Tsui, K. M. (2014). Co-Robots: Humans and Robots Operating as Partners: The Confluence of Engineering-Based Robotics and Human-Centered Application Domains. *Handbook of Science and Technology Convergence*, 1–10.

Sen, H. T., Bell, M. A. L., Zhang, Y., Ding, K., Wong, J., Iordachita, I. & Kazanzides, P. (2015, jul). System integration and preliminary in-vivo experiments of a robot for ultrasound guidance and monitoring during radiotherapy. *2015 International Conference on Advanced Robotics (ICAR)*, pp. 53–59. doi: 10.1109/ICAR.2015.7251433.

Seo, J., Cho, J., Woo, H. & Lee, Y. (2015, oct). Development of prototype system for robot-assisted ultrasound diagnosis. *2015 15th International Conference on Control, Automation and Systems (ICCAS)*, pp. 1285–1288. doi: 10.1109/ICCAS.2015.7364834.

Seo, J., Cho, J. H., Cha, J., Kim, C. & Kwon, O. (2017). Design and experimental evaluations of robot-assisted tele-echography system for remote ultrasound imaging. *2017 14th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI)*, pp. 592–594.

Sharifi, M., Salarieh, H., Behzadipour, S. & Tavakoli, M. (2017). Tele-echography of moving organs using an Impedance-controlled telerobotic system. *Mechatronics*, 45, 60–70.

Siciliano, B. & Villani, L. (2012). *Robot force control*. Springer Science & Business Media.

Storn, R. & Price, K. (1997). Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization*, 11(4), 341–359.

Svaco, M., Sekoranja, B., Suligoj, F. & Jerbic, B. (2013). Calibration of an Industrial Robot Using a Stereo Vision System. *DAAAM International Symposium on Intelligent Manufacturing and Automation*, 69, 459–463. doi: 10.1016/j.proeng.2014.03.012.

Swerdlow, D. R., Cleary, K., Wilson, E., Azizi-Koutenaei, B. & Monfaredi, R. (2017). Robotic arm–assisted sonography: Review of technical developments and potential clinical applications. *American Journal of Roentgenology*, 208(4), 733–738.

Szabo, T. L. (2014). *Diagnostic Ultrasound Imaging: Inside Out (Second Edition)* (ed. Second Edition). Boston: Academic Press. doi: https://doi.org/10.1016/C2011-0-07261-7.

Taylor, R. H., Menciassi, A., Fichtinger, G., Fiorini, P. & Dario, P. (2016). Medical robotics and computer-integrated surgery. In *Springer handbook of robotics* (pp. 1657–1684). Springer.

Tsumaki, Y., Kon, T., Suginuma, A., Imada, K., Sekiguchi, A., Nenchev, D. N., Nakano, H. & Hanada, K. (2008). Development of a skincare robot. *2008 IEEE International Conference on Robotics and Automation*, pp. 2963–2968.

Vanderpool, H. E., Friis, E. A., Smith, B. S. & Harms, K. L. (1993). Prevalence of Carpal Tunnel Syndrome and Other Work-Related Musculoskeletal Problems in Cardiac Sonographers. *Journal of Occupational and Environmental Medicine*, 35(6), 604–610. doi: 10.1097/00043764-199306000-00018.

Vick, A., Surdilovic, D. & Kruger, J. (2013, jul). Safe physical human-robot interaction with industrial dual-arm robots. *9th International Workshop on Robot Motion and Control*, pp. 264–269. doi: 10.1109/RoMoCo.2013.6614619.

Victorova, M., Navarro-Alarcon, D. & Zheng, Y.-P. (2019). 3D ultrasound imaging of scoliosis with force-sensitive robotic scanning. *2019 Third IEEE International Conference on Robotic Computing (IRC)*, pp. 262–265.

Vieyres, P., Poisson, G., Courrèges, F., Mérigeaux, O. & Arbeille, P. (2003). The TERESA project: from space research to ground tele-echography. *Industrial Robot: An International Journal*, 30(1), 77–82. doi: 10.1108/01439910310457742.

Vilchis, A., Troccaz, J., Cinquin, P., Masuda, K. & Pellissier, F. (2003). A new robot architecture for tele-echography. *IEEE Transactions on Robotics and Automation*, 19(5), 922–926. doi: 10.1109/TRA.2003.817509.

Villani, L. & De Schutter, J. (2008). Force Control. In Siciliano, B. & Khatib, O. (Eds.), *Springer Handbook of Robotics* (pp. 161–185). Berlin, Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/978-3-540-30301-5_8.

Villarreal-Cervantes, M. G. & Alvarez-Gallegos, J. (2016). Off-line PID control tuning for a planar parallel robot using DE variants. *Expert Systems with Applications*, 64, 444 - 454. doi: https://doi.org/10.1016/j.eswa.2016.08.013.

Virga, S., Zettinig, O., Esposito, M., Pfister, K., Frisch, B., Neff, T., Navab, N. & Hennersperger, C. (2016). Automatic force-compliant robotic ultrasound screening of abdominal aortic aneurysms. *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 508–513.

Wang, S., Singh, D., Johnson, D., Althoefer, K., Rhode, K. & Housden, R. J. (2016). Robotic ultrasound: View planning, tracking, and automatic acquisition of transesophageal echocardiography. *IEEE Robotics & Automation Magazine*, 23(4), 118–127.

Whitney, D. E. D. (1987). Historical Perspective and State of the Art in Robot Force Control. *The International Journal of Robotics Research*, 6(1), 3–14. doi: 10.1177/027836498700600101.

Wu, B. & Huang, Q. (2016). A Kinect-based automatic ultrasound scanning system. *2016 International Conference on Advanced Robotics and Mechatronics (ICARM)*, pp. 585–590.

Yoshikawa, T. (2000). Force control of robot manipulators. *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, 1, 220–226. doi: 10.1109/ROBOT.2000.844062.

Yu, W., Li, X. & Carmona, R. (2013). A novel PID tuning method for robot control. *Industrial Robot: An International Journal*, 40(6), 574–582. doi: 10.1108/ir-09-2012-406.

Zanchettin, A. M., Ceriani, N. M., Rocco, P., Ding, H. & Matthias, B. (2016). Safety in human-robot collaborative manufacturing environments: Metrics and control. *IEEE Transactions on Automation Science and Engineering*, 13(2), 882–893. doi: 10.1109/TASE.2015.2412256.

Zeng, G. & Hemami, A. (1997). An overview of robot force control. *Robotica*, 15(5), 473–482. doi: 10.1017/S026357479700057X.

Zettinig, O., Fuerst, B., Kojcev, R., Esposito, M., Salehi, M., Wein, W., Rackerseder, J., Sinibaldi, E., Frisch, B. & Navab, N. (2016). Toward real-time 3D ultrasound registration-based visual servoing for interventional navigation. *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 945–950.

Zhang, T., Zhang, W. & Gupta, M. (2017). Resilient Robots: Concept, Review, and Future Directions. *Robotics*, 6(4), 22.

Zhang, W. & Van Luttervelt, C. (2011). Toward a resilient manufacturing system. *CIRP annals*, 60(1), 469–472.

Zhao, L. (2015). *Design and validation of a system for controlling a robot for 3D ultrasound scanning of the lower limbs*. (Ph.D. thesis, École de technologie supérieure).

Zhao, L., Joubair, A., Bigras, P. & Bonev, I. A. (2015). Metrological evaluation of a novel medical robot and its kinematic calibration. *International Journal of Advanced Robotic Systems*, 12(9), 126.

Zheng, Q., Simon, D., Richter, H. & Gao, Z. (2014, June). Differential particle swarm evolution for robot control tuning. *2014 American Control Conference*, pp. 5276-5281. doi: 10.1109/ACC.2014.6858721.

Zhu, W.-H., Salcudean, S., Bachmann, S. & Abolmaesumi, P. (2000). Motion/force/image control of a diagnostic ultrasound robot. *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, 2, 1580–1585. doi: 10.1109/ROBOT.2000.844822.