Multi-Scale Metrology for Automated Non-Destructive Testing Systems

A thesis submitted for the degree of

Doctor of Engineering

Jonathan Riise

Centre for Ultrasonic Engineering

Department of Electronic and Electrical Engineering

University of Strathclyde, Glasgow

March 16, 2020

Copyright Declaration

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.'

'The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Signed:  Date: 12/10/19

Acknowledgements

I do not believe that any human endeavour is accomplished in isolation and this work is no exception. I would first and foremost like to thank my supervisors, Prof. Gareth Pierce and Prof. Tony Gachagan at the University of Strathclyde for believing in me and providing me with the opportunity to undertake this study. Dr. Ian Nicholson at TWI Technology Centre (Wales) has been an constant pillar of support and has shown incredible patience with my work throughout this project.

This project would never have begun were it not for the diligent work of Dr. Carmelo Mineo. I am ever grateful for the guidance he has provided, conversations we have shared and endless code that has been exchanged between us. I would also like to thank Prof. Ian Cooper for involving and supporting me in the continued development of the IntACom project as well as many other challenging tasks at TWI.

I would furthermore like to acknowledge, in no particular order, Dr. Benjamin Wright, Martyn Lindop, Lance Farr, Dr. Ben Knight-Gregson, Dr. Mark Sutcliffe, Philip Wallace, Dr. Roy Brown, Rahul Summan, Dr. William Kerr and the colleagues researching robotics at the Centre for Ultrasonic Engineering. I am also grateful to the administrative staff at the University of Strathclyde, TWI and Nina Hancock at Imperial College.

Most importantly I would like to thank my family and friends who have been there for me throughout the process and helped in innumerable ways. I would never have reached this point were it not for their love and understanding. Finally, I would like to thank Kayleigh, for her patience and for providing the motivation, support and love for me to finish.

"There is no enjoying the possession of anything valuable unless one has someone to share it with"

- `Lucius Annaeus Seneca`, *Letters from a Stoic*

Abstract

The use of lightweight composite structures in the aerospace industry is now commonplace. Unlike conventional materials, these parts can be moulded into complex aerodynamic shapes, which are difficult to inspect rapidly using conventional Non-Destructive Testing (NDT) techniques. Industrial robots provide a means of automating the inspection process due to their high dexterity and improved path planning methods.

This thesis concerns using industrial robots as a method for assessing the quality of components with complex geometries. The focus of the investigations in this thesis is on improving the overall system performance through the use of concepts from the field of metrology, specifically calibration and traceability. The use of computer vision is investigated as a way to increase automation levels by identifying a component's type and approximate position through comparison with CAD models.

The challenges identified through this research include developing novel calibration techniques for optimising sensor integration, verifying system performance using laser trackers, and improving automation levels through optical sensing. The developed calibration techniques are evaluated experimentally using standard reference samples. A 70% increase in absolute accuracy was achieved in comparison to manual calibration techniques. Inspections were improved as verified by a 30% improvement in ultrasonic signal response.

A new approach to automatically identify and estimate the pose of a component was developed specifically for automated NDT applications. The method uses 2D and 3D camera measurements along with CAD models to extract and match shape information. It was found that optical large volume measurements could provide sufficiently high accuracy measurements to allow ultrasonic alignment methods to work, establishing a multi-scale metrology approach to increasing automation levels. A classification framework based on shape outlines extracted from images was shown to provide over 88% accuracy on a limited number of samples.

Contents

List of Figures

List of Figures

List of Tables

List of Abbreviations

ANN	Artificial Neural Networks
CAD	Computer Aided Design
CAM	Computer Aided Manufacturing
CFRP	Carbon Fibre Reinforced Polymers
CMM	Coordinate Measurement Machine
CUE	Centre for Ultrasonic Engineering
dB	Decibel
DLL	Dynamic Link Library
DoF	Degrees of Freedom
FMC	Full Matrix Capture
FoV	Field of View
FSH	Full Screen Height
GE	General Electric
GLOH	Gradient Location and Orientation Histogram
GPU	Graphics Processing Unit
GUI	Graphical User Interface
HOG	Histogram of Gradients
HOG	Histogram of Gradients
I/O	Input/Output
IBT	Index Based Triangulation
ICP	Iterative Closest Point
IGES	Initial Graphics Exchange Specification
iGPS	Indoor GPS
IoT	Internet of Things
ISO	International Standardization Organization
KRC	KUKA Robot Controller

KRL	KUKA Robot Language
MSER	Maximally Stable Extremal Regions
NDE	Non-Destructive Evaluation
NDT	Non-Destructive Testing
OLP	Off-Line Programming
POD	Probability of Detection
RCNDE	Research Centre for Non-Destructive Evaluation
ROV	Remotely Operated Vehicles
RSI	Rapid Sensor Interface
SAFT	Synthetic Aperture Focusing Technique
SDK	Software Development Kit
SIFT	Scale-Invariant Feature Transform
Slerp	Spherical Linear Interpolation
SMR	Spherically Mounted Reflector
STL	Stereolithography
SURF	Speeded Up Robust Features
SVD	Singular Value Decomposition
TCP	Tool Centre Point
TCP/IP	Transmission Control Protocol/Internet Protocol
TESI	Turbine Engine Sustainment Initiative
TFM	Total Focusing Method
ToF	Time of Flight
TRL	Technology Readiness Level
UAV	Unmanned Aerial Vehicles
UDP	User Datagram Protocol
USL	Ultrasonic Sciences Ltd.
UT	Ultrasonic Testing
VIEWS	Validation and Integration of Manufacturing Enablers for Future Wing Structures
XCT	X-ray Computed Tomography

XML	Extensible Markup Language

Chapter 1

Introduction

The aerospace manufacturing industry has expressed an urgency in the need for further automation to cope with the demand for new aircraft. At the same time, regulations and economic incentives also force manufacturers to use new composite materials which have high stiffness-to-weight ratios and can be moulded into complex shapes. There is a need to inspect all safety-critical parts to ensure their structural integrity, a task which is time consuming and slows down overall production [1]. As a consequence, there is a clear interest in automated Non-Destructive Testing (NDT) systems. The field of NDT covers a wide range of methods which allow the evaluation of material properties without causing damage to them. This thesis investigates methods for improving automated NDT systems and expands upon earlier work undertaken in collaboration between the University of Strathclyde and TWI Ltd.

Modern composite manufacturing is capable of producing high performing components which take advantage of the strength of reinforced fibres. The same fibres that give the component strength also makes it difficult to inspect as the anisotropic nature of the material generates scatter noise [2]. Furthermore, the flexibility of these fibres causes parts to vary during the production process. This variability is often not a problem during assembly as parts can be pushed back to their intended shape. It does, however, present a significant challenge for automated NDT, when systems rely on CAD models to plan robotic inspection paths. NDT systems must therefore be flexible enough to accommodate such variations in part geometry.

NDT inspections are often performed manually, including in the high value manufacturing chain and aerospace sector. To perform a manual inspection, a human operator must move probes over areas to be inspected, which can be time consuming for large surfaces. Humans naturally become fatigued, which leads to mistakes and missed inspection points [3]. Although the human operator is able to accomplish complex inspection tasks, industry has for many years sought automated solutions to increase reliability, traceability and lower inspection times. Semi-automated solutions have helped to some extent, but the challenges of ensuring 100% surface coverage and digitally storing large quantities of data remain. Industrial robots have become an increasingly attractive choice in recent years to overcome these challenges [4–6].

Motivated by the need for fully automated inspection solutions, a robotic inspection demonstrator system was developed and commissioned prior to the beginning of this Doctorate programme. The system was the result of the collaboration and joint vision of the University of Strathclyde and TWI Technology Centre (Wales) [7]. The system compromises a twin-robot solution for the deployment of ultrasonic NDT probes. Through the initial integration and development work, it was identified that metrology solutions were needed to obtain a more flexible and accurate system. This was mainly identified through the cumbersome, manual setup procedures needed to begin inspections and the need for a rigorous method of system verification. In particular, it was noted that better tool calibration techniques and accurate, automated part localisation were needed.

The robotic setup at TWI takes advantage of the flexibility and reach of an industrial robot while maintaining the wide-area scanning and focusing capabilities of phased array ultrasonic probes. This thesis contributes to the field of automated NDT using industrial robots by introducing a novel tool calibration method for robot-mounted phased array probes and an improved part alignment procedure based on monitoring the reflected ultrasonic signal. The thesis also investigates methods for automatically determining work-piece position and orientation using machine vision. An in-depth study of the accuracy of industrial robots and the possibility of using motion capture photogrammetry systems to improve robot dynamic accuracy is also presented.

The outcomes of this thesis have been applied to the industry-funded IntACom project led by TWI. The IntACom project is a joint research and development programme, co-funded by GKN Aerospace, Rolls-Royce, Bombardier and the Royal National Lifeboat Institution. The industry partners have all highlighted the need for improved accuracy and traceability in automated robotic inspection systems as a key step in adopting these systems into the manufacturing chain. The research presented in this thesis also aligns with core themes identified by the UK Research Centre in Non-Destructive Evaluation (RCNDE).

1.1 Motivations

Previous work [8, 9] has shown that using off-line path planning is fundamental to advancing the state of the art automated flexible inspection cells using industrial robots. In this context, flexibility means being able to inspect a variety of components with complex geometries with a minimum effort needed to change the physical layout of the cell. The current work expands on this mission by studying the methodologies needed to calibrate the off-line and real world environments. While robotic technologies overcome two main challenges in NDE, namely probe delivery and digital encoding, the systems must be tightly linked with desired spatial resolution and a solid understanding of limitations. The main motivation behind this work is to ensure that high levels of flexibility and adaptability are kept while improving and validating the overall system accuracy and performance through calibration and validation techniques.

This work constitutes applied research and is to a large extent driven by the needs of industry. The main deliverables for the aerospace industry are speed, accuracy, repeatability and ease of use. NDT can quickly become a bottle neck in production where acquiring and analysing ultrasonic data for large components can take up to a week. The required accuracy depends on the critical defect size, which in turn is a function of the type of component and potential location of the defect. For composite inspection, the positional tolerance is typically ± 1 mm. Assessing the repeatability is achieved by inspecting the same reference sample with known defects multiple times.

Improving Human-to-Machine Interfaces (HMI) is a key step to making a system easier to use.

1.2 Industrial Context

The association of engineering doctorates describes the Engineering Doctorate (EngD) degree as providing a more vocationally oriented doctorate in engineering than the traditional PhD and is better suited to the needs of industry. The EngD is a doctorate, equivalent to a PhD. However, the EngD researcher pursues a project while based within a sponsoring company. The programme merges experience of industrial involvement (and today's real-world issues) with a traditional doctorate's immersion in professionalism. The result is a unique individual with relevant industry knowledge and academic depth, grounded by up to four years of experience. During these four years, the EngD researcher may be expected to spend about 75% of their time with their sponsor company, where they are treated like a normal employee.

TWI Technology Centre (Wales)

The work presented in this thesis is the result of an EngD studentship partly sponsored by TWI. Founded in 1946, TWI Ltd. is an independent membership based research and technology organisation. TWI's headquarters are based in Cambridgeshire, with regional facilities in South Wales, Yorkshire, Middlesbrough and several other offices spread worldwide. TWI have a broad range of expertise related to all aspects of engineering, materials and joining technologies. TWI Technology Centre (Wales), located in Port Talbot, is part of TWI's integrity management group and has expertise in a wide range of techniques including ultrasound, radiography, technique validation and software development.

The IntACom Project

In response to the growing demand for high performance lightweight materials, the aerospace industry is manufacturing increasing numbers of critical components from

composite materials. The need to inspect every part coupled with the challenges of geometry, material anisotropy and the desire to drive down costs have resulted in current mainstream inspection techniques failing to keep pace with aerospace industry requirements [4].

Consequently, several aerospace partners approached TWI to investigate the possibility of using automation and advanced NDT techniques to increase inspection throughput whilst maintaining or improving existing sensitivity. The IntACom project was launched in 2010 and the resulting prototype robotic inspection cell exceeded the target of four times improvement in inspection throughput whilst providing increased capability for the inspection of complex geometries [7]. The IntACom robotic cell is shown in Figure 1.1.

Figure 1.1: The IntACom system at TWI Technology Centre Wales in 2018.

IntACom is an ongoing project at TWI, enjoying the continued support of its industry sponsors. Using off-line trajectory planning methods, the current goal of the project is to allow the operator to focus on high-level specification of the NDT task without having to worry about robotic programming. The system must be fit for the intended purpose, that is, it must be able to accurately follow the programmed paths and report back positions of any indications accurately and reliably. Achieving these goals is a major part of the work presented in this thesis.

1.3 Contributions to Knowledge

The field of automated NDT using industrial robots is relatively small and there are only a handful of research groups publishing papers on the topic at the time of writing. The majority of published research investigates either immersion scan situations or single-crystal scans using water jets. his work focuses on large complex geometry components which are scanned using water-jet coupled ultrasonic phased array probes.

Previous studies concerning robots using phased array probes have either focused on a single inspection application or given a general overview of a developed system without providing figures for obtainable accuracy. The work presented here attempts to cover both topics by investigating methods for accurate calibration, path execution and sensor integration in a complete automated NDT system. The following points summarise the main contributions to knowledge.

- Validating the obtainable positional accuracy for robotic NDT systems using laser trackers.
- A comparative study using photogrammetry systems for tracking robots with six degree-of-freedom information.
- An improved calibration procedure for phased array probes in water jet nozzles used for robotic NDT systems.
- A method for obtaining a better calibration between robot and work-piece using ultrasonic alignment.
- A study into the use of machine vision methods for automated part alignment
- The integration of motion capture technology and laser profiler scanners in automated NDT systems to increase their automation levels.
- A method for performing through-transmission inspections of components with varying thickness.

1.4 List of Publications Arising from this Thesis

- *Adapting Robot Paths for Automated NDT of Complex Structures using Ultrasonic Alignment* J. Riise, C. Mineo, S. G. Pierce, P. I. Nicholson & I. Cooper. In: 45th Annual Review Of Progress In Quantitative Nondestructive Evaluation. D. E. & Bond, L. J. (eds.). Burlington, VE.

- *Index based triangulation method for efficient generation of large three-dimensional ultrasonic C-scans* Mineo, C., Riise, J., Summan, R., MacLeod, C. N. & Pierce, S. G. 23 Jan 2018 In : Insight: The Journal of the British Institute of Non-Destructive Testing. 15 p.

- *Methods for improving the accuracy of automated NDE systems* Riise, J., Farr, L., Lindop, M., Pierce, G., Nicholson, P. I. & Cooper, I. in 12th European Conference on NDT (ECNDT) Jun 2018, Gothenburg, Sweden p. 1-9

- *Assessing the accuracy of industrial robots through metrology for the enhancement of automated non-destructive testing* Morozov, M., Riise, J., Summan, R., Pierce, S. G., Mineo, C., MacLeod, C. N. & Brown, R. H. in IEEE Multi-Sensor Fusion and Integration. Baden-baden, Germany Jul 2016 6 p. (Nominated for Best Paper Award)

- *Conformable eddy current array delivery* Summan, R., Pierce, G., MacLeod, C., Mineo, C., Riise, J., Morozov, M., Dobie, G., Bolton, G., Raude, A., Dalpé, C. & Braumann, J. 10 Feb 2016 42nd Annual Review Of Progress In Quantitative Nondestructive Evaluation: Incorporating the 6th European-American Workshop on Reliability of NDE. Chimenti, D. E. & Bond, L. J. (eds.). Melville, NY., Vol. 1706, 8 p. 170003. (AIP conference proceedings; vol. 1706)

- *Deployment of ultrasonic through-transmission inspection using twin cooperative industrial robots* Riise, J., Pierce, S. G., Nicholson, P. I., Cooper, I. & Wright, B. 8 Aug 2016 - 55th Annual British Conference of Non-Destructive Testing. Nottingham, UK p. 1-13

- *Robotic path planning for non-destructive testing through RoboNDT* Mineo, C., Pierce, S. G., Riise, J., Nicholson, P. I. & Cooper, I. 10 Sep 2015 54th Annual British Conference of Non-Destructive Testing. Telford, UK

The work presented in this thesis won the 2018 Armourers & Brasiers' Gauntlet Trust Best PhD Student of the Year Award. The Armourers & Brasiers' Company is a Livery Company of the City of London. Founded in 1322 as the Guild of St George of the Armourers, it was awarded its first Royal Charter in 1453 from King Henry VI. In 1708 the Armourers joined with the Brasiers and the Worshipful Company of Armourers & Brasiers received its current charter from Queen Anne.

1.5 Thesis Outline

This thesis is outlined as follows: Chapter 2 provides an overall literature review of automated NDT methods and why these are particularly applicable to the inspection of large composite structures. An overview of suitable NDE methods for composite material inspection is also given with a particular focus on ultrasonic testing methods. Finally an overview of metrology methods for automated NDT is included with a focus on large-volume tracking and surface scanning methods.

Chapter 3 introduces the concepts of industrial robot kinematics and programming. The various factors which can affect the positional and dynamic accuracies of industrial robots are discussed. Methods for measuring and correcting the accuracy of these robots are also introduced. A novel calibration method to calibrate the ultrasonic and robotic reference frames is presented along with experimental trials. A new method for ultrasonic alignment of a component-specific reference frame and robotic coordinate systems is also presented. Chapter 3 mainly focusses on the techniques and calibration methods needed to improve the accuracy of a robotic NDT system. Experimental data to verify these techniques are presented in Chapter 5.

Chapter 4 focusses on how automation levels can be raised through measurements, in particular how the initial unknown transformation between part and robot can be found. This is a key problem when using off-line trajectory planning as the paths

can only be executed once the robot controller gains knowledge of where the work piece is located. A study into machine vision methods for automated part recognition is presented. The challenges and foundations are presented along with experimental data. The use of motion capture technologies is also discussed as an alternative to machine vision methods.

Chapter 5 does not introduce new concepts, but is instead a collection of validation experiments which incorporate the previously discussed techniques presented in Chapters 3 and 4. These validation experiments include large volume laser and photogrammetry tracking of industrial robots and a subsequent discussion on how these results may influence inspection systems. A quantitative metric is established for validating the inspection system studied in this thesis and experiments are presented which evaluate its performance. The repeatability of the system is investigated through multiple tracking experiments and by repeat scanning of known reference samples. Improvements in absolute accuracy and signal strength are experimentally verified on complex-geometry components.

As the current thesis is undertaken in collaboration with industry, Chapter 6 is dedicated to presenting some of the industrial projects the author has been involved with throughout the EngD project. These include a metrology software toolbox to aid the IntACom project and the use of optical scanning methods for surface reconstruction. A novel method for inspecting components with varying thickness using the ultrasonic through-transmission method is also presented. Chapter 7 concludes this thesis with a summary of work done as well as recommendations for future work.

Chapter 2

Research Background

2.1 Introduction to Automated NDE

Non-Destructive Evaluation (NDE), also referred to as Non-Destructive Testing (NDT), covers the range of techniques used in industry and science to determine properties of materials or components without causing damage to them. The oldest method of NDE is undoubtedly visual inspection wherein manufactured goods, such as pottery and primitive tools, have been checked by eye for cracks or other imperfections. Blacksmiths have traditionally used an early form of sonic testing to inspect the quality of their work by listening to the ring of metals as they were hammered [10]. Following World War II, there was a rapid development of more sophisticated NDE methods such as ultrasonic testing, high-energy X-ray tubes, eddy-current testing and improved magnetic particle testing. As these methods tend to be relatively expensive, they are used in industries where the products are either very costly to manufacture or where the failure of said product has catastrophic consequences. Hence the main end users of NDE methods include the aerospace, automotive, marine, defence, petrochemical and power generation industries.

Where possible in these industries, production is steadily moving towards automation as skilled labour is becoming more costly and automation systems less expensive [11]. Besides the commercial incentive to reduce production costs, automation methods also allow for higher throughput and can reduce risk to workers. These same

principles are increasingly being used in NDE. For large and extended structures there can be an advantage in employing an automated machine to perform the inspection to ensure a 100% area coverage and to be able to pinpoint the location of defects to a higher precision compared to a human operator. Furthermore, crawling and walking robots are also being used in environments which are hazardous for humans to work within [12]. These are especially the nuclear and petrochemical industries where the length of pipes needing inspection requires automated methods, as shown in Figure 2.1. Current NDE methods including ultrasonic, eddy-current and video inspection are portable for on-site inspection, hence these lend themselves easily to automation by simply placing the NDE sensor on a robot. There is also a growing interest in using UAVs (Unmanned Aerial Vehicles) to remotely inspect large structures and ROV (Remotely Operated Vehicles) for sub-sea inspection [12–14].

Figure 2.1: The Versatrax 150 Pipe Inspection Crawler [15]. *Image credits: Inuktun*

The research motivations for this thesis only considers fixed solutions using six Degrees of Freedom (DoF) articulated robotic arms. According to the world robotics organisation the number of these machines, (broadly known as industrial robots), in operation are surpassed 2.4 million by 2018 [16]. The reason for their success since their integration into the workplace in the late 1960's is largely due to their continually lower costs along with higher flexibility and repeatability [11]. Designed to mimic the human arm in terms of flexibility and reach, these robots are ideal for repetitive operations such as assembly, welding or spraying which would previously have been done by hand.

Through automating these procedures it has been possible to produce higher quality goods at a faster rate for less money. Their widespread use and decreasing initial cost have resulted in researchers investigating using these robots for the delivery of NDE payloads in recent years [12]. The use of these industrial robots are especially relevant in the aerospace sector where the size of structures and requirement to scan 100% of load-bearing components has meant that NDE is becoming a bottleneck in production due to the time needed for inspection [17].

One of the industrial robot's main features of interest to the NDE community is their low cost, high dexterity and simple programming methods. Industrial robots have traditionally been programmed through a teach pendant, a process which can take up 10 weeks per operation [18]. As computers have become faster and off-line programming more intuitive, it has been possible to design movement patterns in a virtual environment, alleviating the need to physically manipulate the robot. This has allowed for a much wider use in the variety of tasks the robots can perform as well as a faster reprogramming time. In order to realise the potential of off-line path planning, the accuracy of the modelling environment and physical robot must be good enough to avoid collisions and ensure high quality inspections are achieved.

One main issue hindering full automation is the positional inaccuracy of a typical industrial robot. In many inspection scenarios, the NDE probe must be carefully positioned close to the surface of the component to achieve high quality inspections. Furthermore, the probe must also either be in contact or a set distance above the component. These requirements can be satisfied by performing the scans carefully by hand or by using some type of mechanical scanning device to cover the entire surface of a component. The problems with manual scanning are the large amount of time needed as well as being prone to operator errors. For very precise positioning operations a Cartesian axis gantry system can be used, but these are not able to orient the probe appropriately for components with complex surface geometries, such as those found in the aerospace industry. An industrial robot overcomes the problems of speed while retaining the ability to scan parts with non-planar geometries.

A step-change for automation has been the use of virtual planning environments.

To trust that the physical robot will move as modelled in the virtual environment, the accuracy of robotic movements relative to the component must be sufficiently high as well. This key step is often not considered in great enough detail such that many hours are spent manually adjusting off-line paths to fit the real setup, defeating the original intent of automating the process [19]. The overall goal that this EngD project wishes to contribute to is the further automation of NDE inspections. The goal can be said to be achieved when it is possible to simply place a component inside a robotic cell, tell the robot which parts should be inspected and how and then let the machine take over and do the rest.

It was identified in [20] that automated robot trajectory planning is a key step towards realising commercially viable automated inspection systems using industrial robots. This need arises from the complex geometry and size of typical aerospace components. However, this approach requires that the virtual and real environments align, which can be a challenging task. Following on from this work, the main contributions of this thesis are developing calibration methods for these types of automated non-destructive testing systems and a better understanding the limitations of robotic accuracy in the NDE context.

2.1.1 Non-Destructive Testing in Industry 4.0

In 2012, the "Industrie 4.0 Working Group", led by the German National Academy of Science and Engineering, published their final findings and coined the term *the fourth industrial revolution* [21], also referred to as Industry 4.0. In broad strokes, the first three industrial revolutions were characterised by mechanisation, electricity and digital systems. According to the authors of the report, the world is now on the cusp of the next industrial revolution which will be shaped by *cyber-physical systems* such as the Internet of Things (IoT). The idea quickly caught on in the developed world and a large amount of funding has since gone into realising this new manufacturing paradigm.

The final report on Industry 4.0 includes a number of fascinating predictions and promises, while providing a framework for current industrialised societies to arrive at these goals. Of particular interest to the field of NDE is a new approach to product

manufacturing and life-cycle management. It is anticipated that manufacturing will move away from the traditional ethos of mass producing the same component and instead let customers personalise each manufactured item such that production will have to move towards small scale, individualised batches. This requires a shift in current production facilities towards more flexible and adaptive approaches. These so-called smart factories will be fully automated, able to both take and autonomously exchange information as well as devising different ways for each product to reach their final state.

The concept of a *digital twin* has likewise been popularised. As smart factories become more responsive, each product will need to have a digital equivalent which will store all relevant information about the product, including its physical state. As an example, an aircraft turbine blade will undergo testing directly after manufacturing which will then be captured in a digital file unique to that one item. Throughout its service life, the turbine blade is able to be identified at all times and its service history, detailed geometry and expected lifespan will be captured in a single digital file.

The UK Research Centre for Non-Destructive Evaluation (RCNDE) has identified NDE in Industry 4.0 to be a strategic objective for the UK NDE community [22] and the current thesis is a contribution to future inspection systems. The RCNDE collaborates with industrial partners in both the UK and Europe to identify research areas which will benefit its members as well as running a centre for doctoral training. The work presented in this thesis is sponsored and supported by the RCNDE and the UK Engineering and Physical Sciences Research Council.

2.2 Non-Destructive Evaluation in Composites

Since the discovery of "graphite whiskers" [23] in the late 1950's along with their high stiffness and tensile strength, carbon fibres have been increasingly used for aircraft structures. Besides the high stiffness-to-weight ratio, carbon fibre composites also resist corrosion and can have high thermal conductivity [24]. These properties make Carbon Fibre Reinforced Polymers (CFRP) ideal for use in the aviation and space

industries where high stiffness combined with low weight are crucial for ensuring structural integrity and lowering fuel costs. As composite materials move away from being supplementary reinforcement structures to primary and secondary load bearing structures, NDE methods have had to follow suit [25].

Carbon fibre composites are currently used in a wide range of industries including sporting goods, marine vessels, cars and aircraft structures. As an example, the Airbus A350 XWB aircraft uses 53% CFRP in its construction, making the plane lighter and more corrosion resistant [26]. The LEAP engine (shown in Figure 2.2) from CFM International uses both CFRP in the construction of fan blades and casing, reducing the overall weight by several hundred kilos [27]. By weaving carbon fibres into preformed shapes, adding resin and curing the structure in an autoclave, it is possible to create complex, aerodynamic geometries with little or no machining.

Figure 2.2: Image showing the LEAP engine which contains both carbon fibre composite fan blades as well as fan case. The blades are lined with titanium to absorb impact and to increase durability [27]. *Image credit: CFM International*

Various types of defects are prone to occur during the composite manufacturing process, especially in the case of thick laminates (above 15 mm in thickness) [2]. A poor distribution of individual fibres can lead to some areas having larger amounts of resin and others lacking in resin. This in turn alters the local mechanical properties of the layered material, giving rise to areas which are more prone to delaminations or matrix cracking. During curing, pockets of hot air can be trapped inside the matrix,

causing voids and hence an overall weaker structure. This is what is generally known as porosity and especially weakens the interlaminar shear strength [28]. For larger laminate structures, it is difficult to obtain an evenly distributed curing temperature leading to some parts being over- or under-cured. Incorrect ply orientation can occur during CFRP layup, especially if human operators are manually preparing the material. With newer CFRP moulding techniques potential risks are resin "racetracks" and large trapped areas of air due to poor venting or the resin inlet becoming blocked [29]. Polymer matrices with metal inclusions can also suffer from residual stresses in manufacturing [2].

Finding defects in CFRP poses a different set of challenges than the traditional NDE methods used for metallic structures. The damage processes that occur from impact or fatigue include matrix cracking, debonding layers, moisture ingress, porosity from trapped gasses, delaminations, fibre wrinkles and fibre fracture [25]. As most carbon fibre composites are non-conducting, traditional magnetic particle inspection methods are not useful for inspection. Furthermore, many of the in-service defects are not visible from the surface which eliminates visual inspection. As aerospace components are inspected at the manufacturing stage, detecting defects common to this stage are mainly taken into account for the purposes of this thesis.

A further inspection challenge is encountered when other structures are bonded to CFRP in sandwich panel structures. A typical example is a honeycomb panel wherein two solid carbon fibre skins enclose the core material. This design offers increased strength at lower weights but also increase the need for further inspection to check the bonding between the skins and core materials. As the core material has different material properties from the CFRP skins, few NDE techniques can inspect both bondlines from one side and so access to both sides of a structure may be required.

A key challenge when inspecting composites in the aerospace sector arises from the size and complex geometry of the manufactured parts. The Airbus A350's wings, for example, include a 32 metre wing fixed leading edge composite structure. Figure 2.3 illustrates the size of a typical composite wing structure. Each full wing consists of a number of sub-structures which must all be inspected. As demand for these new aircraft

grows due to their increased fuel-efficiency, the inspection of these large components by human operators is economically infeasible and automated inspection systems are increasingly used [30]. At the same time as NDE methods are being developed and validated, there is also an increasing focus on designing structures for maintenance to facilitate inspections.

Figure 2.3: Image showing the manufacturing process of a 32 m long wing section of the Airbus A350 XWB. Besides the complicated damage processes found in compsosites, the size and geometry of aerospace components pose a significant inspection challenge [26]. *Image credits: Airbus S.A.S.*

2.2.1 Available NDE Methods for Composite Inspection

Some of the early NDE methods for the inspection of composites were mainly designed for glass fibre reinforced plastic gas storage tanks [31] and many of the methods are also applicable to CFRP. As the aerospace industry began to incorporate composite materials into their designs in the 1990's, a large volume of NDE [2,28,31,32] literature on inspecting CFRP was published. As with nearly any inspection scenario, there is no single technique which will manage to detect all possible defects. Each inspection scenario is unique and the chosen method must reflect the requirements of the customer along with the restrictions arising from the physical properties of the component. This section will highlight the main inspection techniques used for CFRP, along with their limitations.

Acoustic Emission

By monitoring the emission of acoustic waves arising from irreversible changes in the internal structure of a component, the growth of defects can be inferred [31]. Typically the emission of sonic or ultrasonic signals will be monitored while the component undergoes a load test to investigate crack growth. As the component deforms, stored energy (which could be strain, chemical potential or crystalline) is released and some of this manifests itself as elastic waves. Acoustic emission has been used for quality inspection of composite parts (mainly storage gas tanks) and is well suited for permanent structural health monitoring [29, 33]. Though it is possible to monitor defect growth, the technique is more qualitative than quantitative, especially in complicated structures such as composites. By applying many sensors, the origin of defect can be located and this can then be characterised with another NDT system. The main advantage of the technique is to be able to monitor a component globally with a network of sensors.

Thermal NDT Methods

By introducing heat to one side and using an infrared camera to observe how this dissipates through the component, internal flaws such as delaminations can be detected. Alternatively, this can also be done by heating and monitoring the same side or even by cooling the component and observing the temperature gradient as the component returns to room temperature [29]. The mechanism behind this relies on the different specific heat capacities of the component and inclusions such as trapped air. Thermography has proved useful in finding moisture ingress, inclusions and delaminations in composite materials, but suffers a poor signal to noise ratio for thicker components and is only able to detect indications near surfaces. The sensitivity of the technique is however very limited; typically the smallest detectable flaw has a diameter proportional to its depth below the surface.

Optical Methods

As a non-destructive method, optical (or visual) inspection is perhaps the oldest and relies on information conveyed from surface deformations to infer observations about the internal structure of a component. By projecting light patterns onto a surface, any deviations in the pattern indicate surface deformation or stress concentrations [2]. Being non-contact and easily interpreted, this method is useful for detecting near-surface delaminations, buckling, impact damage and large disbonds in composite materials. Typically, the component being tested is exposed to strains by a vacuum and the gradient of deviations in the surface indicates defects.

Vibrational Methods

Another well-established inspection method utilises the acoustic signal introduced by vibrations (tapping) to detect any discontinuities within the structure. As delaminations and disbonds cause a change in the local stiffness of the component, the reverberating frequency induced by a tap will be different for intact components and defective ones. The advantages of this method are the inexpensive setup, fast local monitoring and the ease of applicability. The sensitivity of the method is however quite limited and only very large defects are detectable [34]. Accurately pinpointing defect indications is also not possible.

Eddy-Current Testing

Although the epoxy resin in composites is non-conductive, the conductivity of the fibres is enough to allow for eddy current testing. However, as the dielectric material separates layers of conductive fibres current flows parallel to the surface and high depth resolution is difficult to achieve [32]. The technique is however good at characterising fibre volume and layout as well as near-surface delaminations [35].

Radiographic Methods

The principle of radiographic inspection methods is to let penetrating radiation (x-rays, gamma rays or neutrons) pass through the component and detect changes in intensity. Attenuation of the rays depends on the density and thickness of the material it passes through, so defects such as voids, resin and fibre rich areas, and disbonds can easily be detected [36]. As the observed intensity of the recorded radiation depends on the thickness of the component, thin defects such as delaminations which are perpendicular to the direction of the rays are not visible. However, if the defect is aligned parallel to the incoming radiation, it will show up clearly. By including metal fibres while manufacturing composite components, it is possible to obtain better defect detection as these will show up clearly against the carbon/resin background.

Other techniques such as backscattering and X-ray Computed Tomography (XCT) have better resolution and are capable of displaying the weaving of individual fibre layers. As the component (or radiation source) is rotated in XCT scans, it is also possible to detect laminations. The equipment is however very expensive and can only inspect smaller (typically up to 300 mm in diameter) components due to detector sizes [36]. Radiographic methods are in general more complicated and expensive to employ as safety concerns regarding radiation need to be addressed first.

The author of [2] notes that the usefulness of conventional radiographic methods for porosity and fibre volume measurements are limited but that XCT techniques are becoming more readily available. Conventional XCT techniques are limited by component size, inspection limitations and long scan times. Larger components could be scanned by moving the source and detector around an object, for example using robots. Such a system would be very costly and would not give the same image quality as conventional XCT due to reconstruction errors and positional inaccuracies [37]. Furthermore, few facilities have the available radiation shielding to commercialise such a solution.

Ultrasonic Methods

By applying a voltage to a piezoelectric element an ultrasonic wave can be produced which in turn can be directed into a component for testing. By either receiving the transmitted acoustic pulse or reflections from boundaries within the component, information about the internal structure of that component can be inferred [33]. Due to its portability, extended availability and ease of use, ultrasound has become a mainstream technique in composite inspection. There are however still challenges to be overcome with ultrasound such as the high attenuation of the pulse in composite materials due to their complex, layered structure. In cases with limited access and high curvature geometries the receiving of the ultrasonic pulse can become very complicated and specific transducer designs need to be manufactured for individual tasks.

As the attenuation of ultrasound is very high in air, a coupling medium is typically used. This can be achieved either by having a layer of water-based gel between the transmitter/receiver and the component or by submerging the entire setup in water. Ultrasonic inspections are used to detect voids, cracks, delaminations, inclusions, porosity and sometimes fibre wrinkling. At the same time, ultrasonic testing allows characterisation of defects in terms of size, position and shape. Laser ultrasonic delivery methods are also available for rapid large area scanning and are independent of the surface curvature, but the systems are difficult to set up and have a high up-front cost.

Summary

Each of the NDT methods reviewed has its own benefits as well as drawbacks. Often it is not sufficient to use a single NDT technique as the size and location of a defect imposes constraints on which method can be detect it. The accuracy of each method is material and situation dependent, but a few overall observations can still be made. Optical methods are only capable of detecting near-surface defects, but can do so down to hundreds of microns. Acoustic emission can potentially detect sub-mm defect events such as fibre breakages. Vibrational methods are typically limited to near-surface defects on the order of mm while thermography can detect defects with a diameter

approximately proportional to the depth of the defect. Eddy current testing is capable of detecting sub-mm flaws, but only near the surface. Radiography can resolve defects at the micron level, but is dependent on flaw orientation. The accuracy of ultrasonic testing depends on the material and frequency of sound but can measure up to 0.025 mm in thickness gauging applications.

Table 2.1 summarises the ability of the mentioned methods to detect the most common defects found in composites. The methods are ranked as either having proved (two **) or limited (one *) ability to detect a given flaw type. The limitations may be due to the defect depth, orientation or, in the case of acoustic emission, being able to detect when the defect forms. Computed tomography can detect most flaws, but as mentioned, is limited to smaller components.

Table 2.1: Defect types in composites, their detectability by different NDT methods and the automation potential of each NDT method. AE: Acoustic Emission. EC: Eddy currents. UT: Ultrasonic Testing. **: Proven ability to detect flaw. *: Limited ability to detect flaw.

Flaw Type	**AE**	**Thermal**	**Optical**	**Vibrational**	**EC**	**X-Ray**	**UT**
Porosity			*				**
Delamination	**	**	*	**	*	*	**
Disbonds	*			*	*	*	*
Fibre Alignment				*	*	*	*
Inclusions		**			*	**	**
Moisture Ingress		**				*	*
Automation Potential	Low	High	High	Low	Med	Med	High

Out of the above mentioned methods, ultrasonic testing is by far the most common technique used in the inspection of CFRP components [25] with thermographic methods becoming increasingly popular. Though each inspection will have its own specific requirements, both ultrasound and thermography are inherently safe to use and provide a method for detecting volumetric defects. These techniques have been applied to automated systems for the rapid testing at the manufacturing stage, but this thesis will solely focus on the ultrasonic testing technique as this is being used in the IntACom system.

2.3 Ultrasonic Testing

The origins of Ultrasonic Testing (UT) began during the second world war for detecting internal damage to ship hulls and has been a mainstream testing approach in industry ever since. Ultrasonic waves are elastic waves with a frequency beyond the limit of human hearing, typically around 20 kHz. Ultrasound is able to propagate through solids, liquids and gases and is used extensively in fields such as medicine, navigation and NDE. For testing metal components, frequencies between 1 MHz and 20 MHz are typically used. In composite inspection, this range is limited to around 5 MHz due to the high attenuation of sound in these anisotropic materials.

To produce sound in a material, the molecules or atoms of that medium must vibrate harmonically about their equilibrium positions. Sound waves are mechanical waves which propagate energy through a material. In the bulk of a material, sound can move either as a longitudinal (also known as a compression) wave or a transverse (shear) wave. An illustration of these wave types is shown in Figure 2.4. It is also possible for sound to move through a material as a surface wave. Shear waves are not effectively supported in gases and liquids. As ultrasonic waves move through media they lose energy through several mechanisms include scattering and absorption, known as attenuation. The loss of energy is related to the frequency of sound, with higher frequency waves experience greater attenuation [33].

2.3.1 Generation of Ultrasound

Ultrasound is conventionally generated by converting electrical energy into a mechanical displacement through a piezoelectric crystal. It was demonstrated in the late 1880s by Pierre and Jacques Curie that some crystals acquire electric charge when subjected to an external force. The reverse effect is also applicable such that an electric charge can cause these types of crystals to deform. This provides a convenient and efficient way to generate ultrasound from electric pulses. Another method of generating ultrasound in a component is to heat a small area of its surface very rapidly, using a laser. This causes the material to expand rapidly and can generate an ultrasonic wave inside the

component.

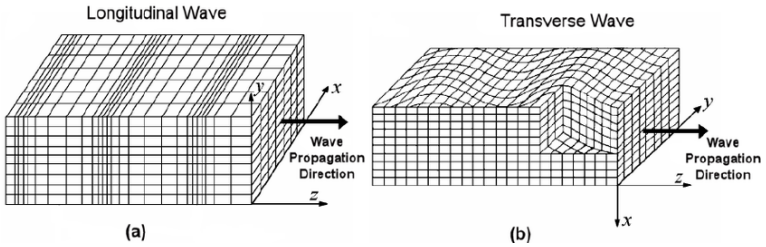

Figure 2.4: Types of waves travelling through the bulk of a material. (a) Longitudinal (compression). (b) Transverse (shear). *Image credits: [38]*

Today's piezoelectric crystals used for NDE are typically composites of either barium titanate or lead zirconate titanate. Besides the chemical composition of the crystals, their shape and size will determine their mechanical motion when subjected to an electric pulse. As an example, the frequency of the sound generated depends on the thickness of the crystals. Thin discs of crystal resonates with a wavelength twice its thickness. This causes restrictions on how high frequencies can be generated with crystals as there comes a point where the discs would be too fragile to be practical.

2.3.2 Ultrasonic NDE Techniques

A number of methods exist for inspecting components with ultrasound. The work presented in this thesis focuses on the inspection of CFRP with automated systems. The two applicable methods for the system used are Pulse-Echo (PE) and Through Transmission (TT) testing. Due to the large impedance mismatch between air and solids, water is usually used as a coupling medium. In the current research, water jets are used to transmit ultrasound from a transducer to the test-piece.

Pulse-Echo

Pulse-echo type inspections are the most commonly used ultrasonic tests and can give information about the bulk of a material as well as finding defects within the material [39]. PE inspections are carried out from a single side of a specimen and the transducer acts as both a transmitter and receiver. A pulse of ultrasound is directed into the test

specimen, reflected by boundaries or defects and then detected, as shown in Figure 2.5a. With knowledge of the speed of sound within the material it is possible to determine the depth of a reflected signal. By monitoring the attenuation of reflected signals it is also possible to infer information about changes to the material such as corrosion. PE inspections are carried out as either zero-degree or angled inspections. The zero-degree PE method is typically used for the inspection of CFRP. Angled inspections can be used to direct the ultrasonic pulse towards a certain location or to induce wave mode conversion. At non-perpendicular incident angles the operator will have to take the effects of refraction into account.

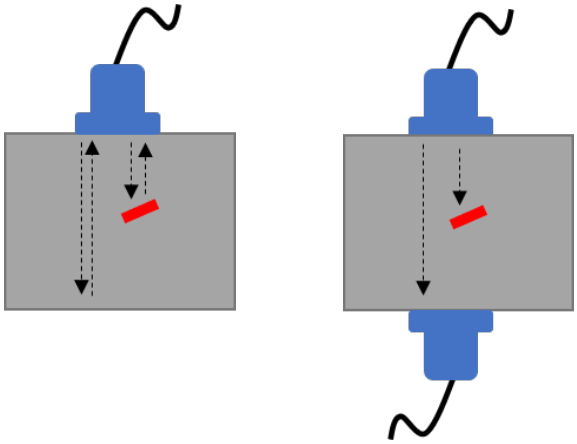

(a) Pulse-echo (b) Through Transmission

Figure 2.5: The two most common ultrasonic inspection methods

Through Transmission

The through-transmission inspection technique uses two transducers placed on opposite sides of a component, as shown in Figure 2.5b. One transducer then acts as transmitter and the other as a receiver. Information about the bulk of the material can be obtained by monitoring the attenuation of the transmitted signal. The absence of a received signal can indicate the present of a flaw. This inspection method is often used for thick or attenuating materials as the sound only has to travel through the component once. In CFRP inspection this method is often used to monitor porosity. The main advantages of the through-transmission method compared to the pulse echo method

are the full depth coverage due to the lack of a front-wall echo and less attenuation of the ultrasonic signal. On the other hand, the method cannot be used for gauging the depth of a defect and requires good alignment between the transmitting and receiving probe. To facilitate this alignment, jigs or automated devices are typically used.

Phased Array Probes

Phased array probes have been used in the medical industry for many years, but their use for industrial ultrasonic testing only began in the 1990s. Advances in micro-machining, computing power and micro-electronics all contributed to maturing phased array probe technology. Since then, the use of these probes has expanded considerably within the field of NDE [40]. Consisting of an array of smaller piezoelectric crystals, these probes allow for the simultaneous or single excitation of individual elements. Delaying the firing of single elements electronically allows the operator to sweep, steer, scan or focus the ultrasonic beam.

The need for phased array probes was driven by a number of factors. Being able to size and detect cracks at different depth from a single position, improving accuracy (for sizing, location and orientation) and covering a larger area through electronic scanning are just a few of the reasons. Computing the time delays (known as focal laws) for different elements needed to steer and scan the beam is typically done prior to inspection. To electronically scan, the same focal law is multiplexed across a sub-group of active elements at a constant angle along the length of the probe. This is the equivalent of using a conventional probe for raster scanning along one axis.

2.4 Automated Inspections Systems

Inspection of high-value components are carried out manually in a large number of scenarios by trained operators. As discussed earlier, there are some instances where it is not economic for a human to attempt to inspect a structure due to its size or complex shape. Furthermore, human operators suffer from fatigue and are prone to errors. An extensive report on human factors in NDE is presented in [3]. As such,

there is an incentive, when possible, to automate the inspection process. Additional benefits of automation include being able to keep digital records for entire components and evaluate criteria such as reliability and positional accuracy more effectively.

The use of automated inspections can generally be classified as being either in service or at the manufacturing stage. As described in Section 2.2, a number of different NDE techniques exist and some, such as visual inspection, lend themselves to automation easily by for example installing a camera over a moving conveyor belt. As automation continues to grow in many industries, the installation and use of NDE as part of an automated assembly line follows suit [8]. Automated systems for in-service inspections are typically used in situations where the environment is inaccessible or dangerous for human operators to work within. Examples of these areas include pipework, off-shore structures, storage tanks and power lines [12].

At this point it is worthwhile to draw a distinction between robots and automated systems. An automated system is typically a highly specialised tool which can consist of multiple machines such as lifts, conveyor belts, pneumatics, etc. These are built for a very specific purpose and provide a return on investment through large-volume operations, for example in a packaging operation. Robots, on the other hand, are adaptable to a number of different situations and can be reprogrammed relatively quickly and perform a number of varied tasks by selecting different tools. To do so, robots require a method of sensing, planning and executing motions to achieve a stated goal [41, 42]. Throughout the last industrial revolution, several automated solutions with and without robots have been realised to facilitate the inspection process. This section will highlight some of these systems for inspecting components at the manufacturing stage which have been developed or are in use at the time of writing.

2.4.1 Gantry Systems

Gantry type systems offer high rigidity along with high positional accuracy and have been used for ultrasonic inspections for a number of years. A typical flat-bed immersion inspection system from Ultrasonic Sciences Ltd. is shown in Figure 2.6. This type of scanner consists of a Cartesian axis motorised gantry manipulator which allows the

operator to move an ultrasonic probe over the surface of a component. Some systems allow the operator to define non-linear paths though most paths are defined on a plane. An additional manipulator will at times be attached to the end of the arm holding the probe allowing for swivel and pitch movements for a total of five programmable axes. This allows the system to inspect components with gentle curvatures, but not all types of complex geometry found in the aerospace industry.

Figure 2.6: Immersion inspection system with 5 axes on a bridge-type manipulator [43]. This is often considered a typical ultrasonic Cartesian scanner system. *Image credits: USL.*

Although these systems are widely used in industry [44] and offer high positional accuracy through their rigidity, the systems can be very expensive and are typically optimised for use with a limited set of components. Setup times can be long, especially for immersion type systems where the water typically has to sit for up to 24 hours to allow air bubbles to escape. Further to this limitation, some components are not submersible due to the possibility of water ingress or corrosion. Finally, the standard immersion system only offers one-sided access meaning that through-transmission ultrasonic inspections are not possible.

Gantry-type inspection systems using squirter nozzles are also available as commercial solutions, such as the one shown in Figure 2.7. These systems are capable of inspecting components with gentle single and double curvatures and can additionally access components from both sides, meaning that through-transmission inspections are possible. These systems, however, tend to carry a large up-front cost, require significant

floor space and only have limited flexibility in terms of reconfiguration.

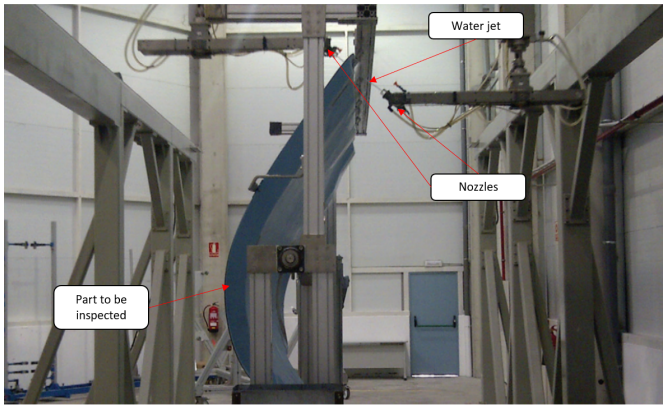

Figure 2.7: Gantry-type inspection system from ScanMaster [45]. These types of systems are able to perform dual inspections of component with single curvatures. *Image credits: ScanMaster*

2.4.2 Commercial Systems using Industrial Robots

Since their invention in the 1950's, industrial robots have become increasingly used in a large number of industries [41]. The use of industrial robots today for inspection is well-documented in a large number of publications hence this section will focus on some of the inspection systems relevant to the specific task of inspecting aerospace structures using the ultrasonic technique. Research at the University of Dayton Research Institute began work on using industrial robots equipped with phased array ultrasonic probes for inspecting turbine blades in 2001 [46]. The developed system used a combination of a single industrial robot with a turntable to provide a fully automated solution for the US Air Force [47]. Since then, a number of commercial systems have been advertised while research at several universities has looked into improving these types of inspection systems using industrial robots [6, 48–51].

The six axes of an industrial robot gives the machine greater dexterity and the ability to cope with complex 3D geometries containing large dual curvatures, as found in the aerospace industry. Robots also offer an additional advantage over gantry systems in that they are scalable such that a larger robot can replace a smaller system without having to remove the entire installation. The challenges for using industrial robots as

inspection tools have been integrating the NDE hardware, obtaining sufficiently accurate positioning, having adequate motion control for complex paths and encoding data accurately. These challenges must all be met before an inspection system can be certified to meet the relevant industry standards. The following list provides examples of some available commercial systems other than the IntACom system which was described in Section 1.2.

Tecnatom

Tecnatom (Spain) have developed a group of inspection systems, known collectively as the TAURUS inspection platform based on a collaboration with industrial robot manufacturer KUKA (Germany). The primary focus of the inspection systems is the aerospace industry. The inspection system is designed as a modular solution where one or more robots can be used, as shown in Figure 2.8. The inspection system furthermore compromises linear tracks, turntables and automated tool changers to facilitate inspection of a large range of different components. Components can either be inspected in pulse-echo or through-transmission mode using water squirter systems.

Figure 2.8: Tecnatom dual robot system with linear tracks [30]. Tecnatom have a number of solutions on the market including squirter-based ultrasonic inspection and laser ultrasonic inspection. *Image credits: Tecnatom.*

Tecnatom are undoubtedly one of the leading integrators of inspection systems using industrial robots and have developed a wide range of additional features for their TAURUS system. These include automated trajectory generation, simulation of the robot movements and a visualisation interface to display the gathered data as presented

in [30, 49]. Where is it not possible to use water jets, the TecnaLUS system can be used to inspect large components using laser ultrasound. There is unfortunately little discussion in the literature distributed by Tecnatom about the positional accuracy and obtainable repeatability of their systems.

Tecnatom have published papers discussing the challenges of inspecting components with unknown geometries and demonstrate the use of a laser surface profiler to generate point clouds for robot path trajectory planning [30]. Additionally, the company announced in 2016 that the laser profiler can also be used to identify components and perform alignment with pre-planned trajectories [52]. This alignment check is currently limited to cylindrical geometries and requires the robot to move over the surface of the component, increasing the risk of collision if an incorrect object is placed in the inspection cell.

Ultrasonic Sciences Ltd.

Ultrasonic Sciences Ltd. (USL) have been supplying ultrasonic inspection systems since 1988 and provide both gantry, flat-bed and robotic systems. One of the key features of their robotic systems is a purpose-built motion controller for a robot which can provide feedback rates of up to 2000 Hz [53]. This allows speeds of up to 700 mm/s when inspecting with single crystal transducers.

Figure 2.9: Dual robot inspection system from USL [43]. The system works with two dual-frequency single crystal transducers which couple the sound to the component through water jets. *Image credits: USL*

The majority of USL's business appears to be focussed on more conventional systems such as their gantry solutions. The approach taken by the company appears to be to adapt conventional testing to their robotic systems. It is not clear from the literature if USL use phased array transducers within the robotic solutions which would have allowed them to reach similar overall inspection times at lower robot speeds. Unlike Tecnatom, USL have not published literature concerning how the robot could adapt to unfamiliar geometries and seem to rely on rigid fixtures for part placement.

General Electric

General Electric (GE) is currently one of the largest distributors of NDE equipment in the world [54] and supply both portable ultrasonic scanners as well as automated systems. Their robotic inspection solution, Hydrastar [55] is shown in Figure 2.10. The robotic inspection solution utilises wide phased array as well as single crystal squirter nozzles and can inspect in either pulse-echo or through-transmission modes. Trajectory planning is perform on CATIA models using the FASTSUITE package from CENIT. The application software also appears to be able to teach the robots the position of a component from laser triangulation measurements, though the available literature on this function is limited.

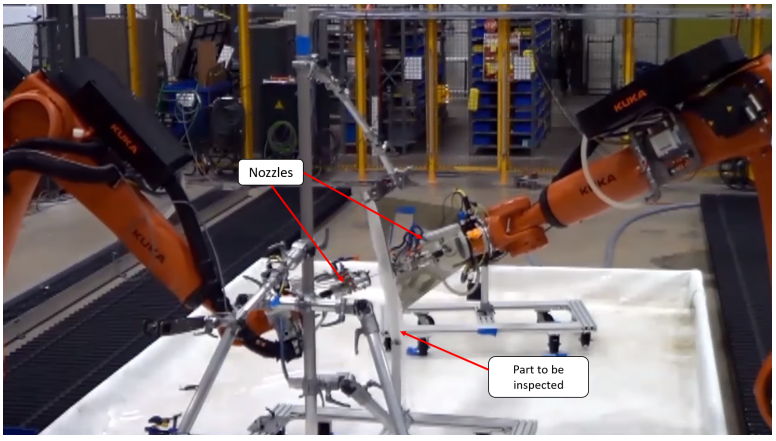

Figure 2.10: Hydrastar robotic ultrasonic inspection solution from GE. The dual robot configuration is inspecting a carbon fibre component in the through-transmission inspection mode [55]. *Image credits: GE*

Orbital ATK

Orbital ATK (now part of the Northrop Grumman Corporation) has had over 50 years of experience inspecting components for the aerospace, defence and space industries [56]. A large number of automated inspection systems were developed over this time period, notably the robotic systems introduced in the year 2000. The system consisted of dual robots to allow the inspection of diverse geometries. Furthermore a 3.7 metre turntable was added for the inspection of cylindrical structures. Throughout the years, the system has been adapted to a number of NDE methods including ultrasonic PE, TT, air-coupled UT, X-ray inspection and contact UT. The proprietary control software includes being able to define part geometry from 3D laser scanning. The authors unfortunately do not discuss the accuracy of available systems but note in concluding remarks that there is a clear need for "accurate scanning articulation" and "global-coordinate awareness" of robots [56].

FILL GmbH

FILL Gmbh is a family owned Austrian company, established in 1996, that focuses on mechanical engineering solutions for automation. The FILL Accubot, shown in Figure 2.11, consists of two STÄUBLI robots that are fully integrated in a standard control environment based on SIEMENS Sinumerik 840D [57]. To increase the accuracy of standard robot units, FILL have integrated external encoders on each robot link and calibrates the entire systems prior to installation. Combined with the enhanced control offered by the SIEMENS software, FILL are able to obtain accuracies below 0.35 mm at speeds of up to 2000 mm/s.

Figure 2.11: The Accubot dual robotic ultrasonic inspection solution from FILL GmbH [57].

Accubot is considered state of the art at the time of writing, but the need for additional controllers and calibration increases the cost of the system substantially compared to others on the market. Although FILL provide the ability to scan with phased array probes, this is only achievable in contact and through-transmission inspections are only possible with single-element probes.

IntACom

The IntACom project began as a joint-industry research project at TWI Wales in 2012 and has continued to develop since. At the core of the project is a dual KUKA robot inspection system, showing in Figure 2.12. The original target of the project was reducing the time taken for inspection of complex geometry composite components by a factor of four. To achieve this target, three areas were addressed: automation of manual inspection methods, enhancement of existing semi-automated systems through the use phased array technology and software enhancements including better off-line path planning.

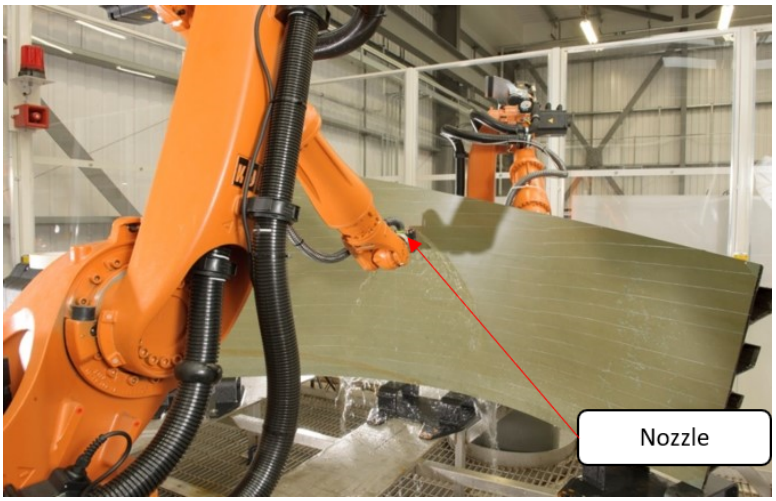

Figure 2.12: The IntACom inspection system at TWI. Twin cooperating robots are inspecting a carbon fibre winglet in through-transmission using phased array probes.

The robot system at TWI is capable of inspection a 2 m x 1.5 m x 1.5 m volume in either pulse-echo or through-transmission configurations. As IntACom progresses, a new robot system will be built which includes robots on linear tracks, capable of inspecting an 8 m x 2 m x 2 m volume. TWI have written and validated custom acquisition and visualisation software to encode ultrasonic data with positional information from the robots at up to 400 mm/s. The system is capable of using single-crystal or phased array probes and significant research has been undertaken to design and manufacture nozzle geometries which provide near-laminar flow to ensure a stable ultrasonic coupling.

2.4.3 Limitations of Current Systems

The previous examples all help illustrate that the field of robotic NDE system using ultrasonic probes is beginning to mature. Most manufacturers have previously supplied more conventional automated inspection systems which they are now tying to their robotic solutions, such that software interfaces and legacy ultrasonic hardware can be transitioned to a new positioning platform. This limits the flexibility and scalability of these commercial products. It was previously identified in [20] that there exists a need to automate the entire NDE process and develop bespoke robotic NDE software,

especially for trajectory planning.

The current EngD project is a continuation along the path of aiming to advance the Technology Readiness Levels (TRL) of automated inspection systems so they are suitable for adoption by industry. The use of dedicated path planning software has provided a valuable contribution toward this goal. However, the verification of robotic systems, especially with regards to accuracy, resolution and repeatability is still not well documented within the field of automated NDT.

Unlike most other robotic applications, a robotised NDT solution needs to accurately follow surfaces that extend several square metres. For aerospace, the typical defect sizes in composites range from 3-12 mm. To ensure adequate coverage of the smallest defect size, the positional error of a robotic system should not exceed ± 1 mm. Furthermore, the robot should stay normal to the surface within ± 2 degrees to minimise refraction of the ultrasonic beam. Many manufacturers simply state the positional resolution of each robot, but a more holistic approach is needed wherein the specific challenges and limitations of NDE hardware are taken into account as well.

The main contribution of this thesis is to further advance automation levels by using measurements from external systems as well as from the ultrasonic probes mounted on robots. Optical methods are used to determine part location prior to robotic scanning. Using advanced imaging techniques, it is shown that the geometrical relationship between robot and ultrasonic probe reference frames can be accurately found. Ultrasonic measurements are also used to provide fine-tuned adjustments to the pre-planned robotic trajectories.

2.5 Metrology and Non-Destructive Evaluation

This section discusses the ways in which dimensional measurements can be used to increase the positional accuracy and automation levels of robotic NDE systems. Metrology is the science of measurements including all the practical and theoretical aspects thereof. Furthermore, the field of metrology also encompasses the science of uncertainties in measurements. A measurement is understood as assigning a number to describe

an object or event. In order for the measurement to be meaningful, the number must be relative to a predefined standard, (in the case of distances the standard would be the meter). The establishment of standards, developing methods for the practical applications of these standards and ensuring traceability of said applications are all responsibilities of the field of metrology as well.

When discussing the role of metrology, it is worthwhile to reiterate some of the main terms that are used to classify different measurement instruments. Accuracy is a term used to describe how well measurements agree with their "true" value. Unfortunately, a true value cannot be defined, so although accuracy is used as a qualitative term, it must be associated with a spread in measurements to be meaningful. Precision describes the spread of measurements upon repetition. In practice, a nominal value is often used as a "true" value. It is possible to have a system which measures with high precision, but poor accuracy – in other words, the measurements are clustered together, but around an incorrect value. Resolution is defined as the minimum discernible difference that can be made between two adjacent measurements [58].

2.5.1 Accuracy Requirements for NDE

Defining the positional tolerance requirements for NDE is not a straightforward task as they depend on the material, component function and location within the component. Design engineers will typically want the inspection to detect flaws to a specific maximum size, but the definition of a "flaw" is often both material and geometry dependent. For composites within the aerospace the typical flaw sizes range from 3 mm to 12 mm. Furthermore, the NDE inspector must choose a method that can reliably find indications of a certain size, while being mindful of costs. These considerations led NASA in the 1970s to define Probability of Detection (POD) curves, as shown in Figure 2.13. POD curves are based on a large number of inspections of specimens with known defect sizes. They help determine the minimum indication size that will be reliably detected by the chosen NDE technique. For example, Figure 2.13 shows simulated data wherein there is a 90% probability of finding a 4.8 mm size flaw.

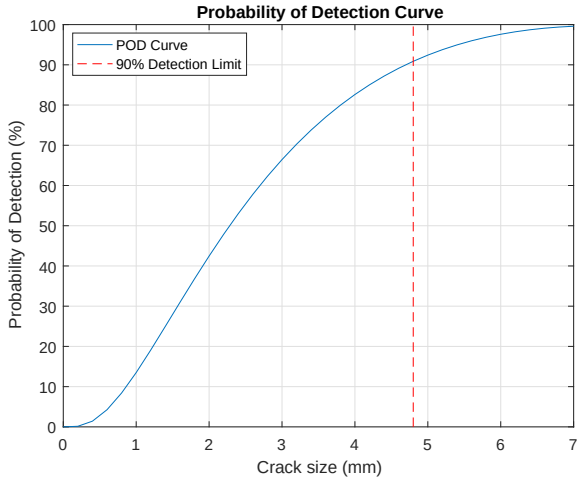

Figure 2.13: Simulated POD curve showing the threshold of 90% detection limit.

POD curves help to determine a limit on the accuracy of the inspection system. However, for dimensional sentencing on indications, a number of other factors have to be taken into account as well. Most NDE methods are indirect; ultrasonic inspections relate changes in sound intensity to flaw size for example. Besides detection, an operator typically also has to provide length, depth and position relative to some datum for each indication. Thus the critical flaw size gives an upper limit on the positional accuracy of the method as well. A rule of thumb used to be that a measuring instrument needed to have ten times better accuracy than the measurement requirement [58]. Today this is instead quantified through uncertainty budgets and measurement tolerances. For the application considered by this work, the aim is to achieve a positional tolerance of ± 1 mm.

2.5.2 Metrology for Automated NDE Systems

Figure 2.14 illustrates how traceability, calibrations and international standards are related. These concepts are translated into a number of different industries, including NDE. A series of reference blocks are typically built for each type of inspection such that signals from indications with known parameters can be compared to data obtained from a new component. The difficulty in the field of NDE, however, is that flaws can occur in a multitude of ways and only a limited number of standards can be made

when costs are considered. Often the material properties will also exhibit variations, especially in the manufacturing of composites. All together this means that proving a method is reliable is based on the POD, but this approach is highly dependent on human factors [3].

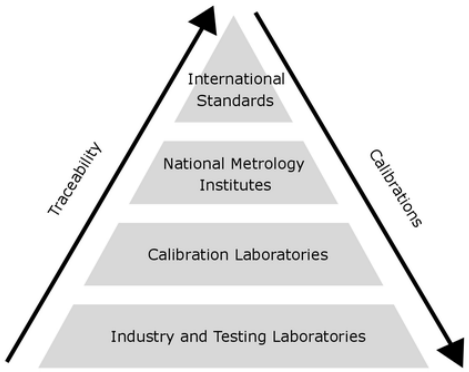

Figure 2.14: The metrology pyramid showing how measurements at one level can be traced back to a measurement made to a higher accuracy. Conversely, calibration is carried out in a top-down manner to ensure that measured values are within tolerances at each level [58]. *Image credit: NPL.*

The NDE community hopes that automated systems can help improved POD by removing human error from the inspection procedure. Metrology solutions can be used to quantify an automated systems reliability and suitability for this purpose. Broadly speaking, metrology for robotic NDE systems can be divided into two categories: 1) measurements relating to the component being inspected and 2) measurements to assess and improve the robotic manipulator.

Part Metrology

To achieve higher automation levels using robotic systems, a number of measurements must be taken before the inspection can begin. In particular, three key points must be answered about a test specimen:

- Is the object known to the inspection system and does a robot trajectory exist for this object?
- Is the component in the expected position and orientation?

- Does the geometry of the object conform to the CAD model used for trajectory planning?

If the answer to the first question is negative, then the inspection must most likely be aborted while an operator defines a trajectory. Alternatively, a fully automated path planning application could define a trajectory if the CAD model of the component exists. This is a current research theme at many companies and universities [9, 59, 60]. If the object is not known and no CAD model exists then steps will have to be taken to digitise the surface of the object, usually through optical metrology methods.

For the robot to be able to reach the object and avoid path singularities, the object must be close to its expected position, defined during the off-line path planning step. The same is true for the object's orientation in 3D space. It is possible to use jigs and fixtures to precisely position the object in the same location for every scan, but this can be costly and time consuming. A more flexible approach is to use a metrology system to measure the orientation and position of the part and then modify the trajectory (if possible) to perform an inspection.

A fundamental assumption of off-line path planning is that the geometry of test specimens match their CAD model. If this is not true, the trajectory must be corrected for an optimal inspection result. However, methods are needed to check for conformity, especially within the composite manufacturing sector where deviations from CAD are well-documented [61, 62].

Manipulator Metrology

Depending on their size, robots can cover large areas and the operator must be confident that data has been gathered with sufficient accuracy. Paths planned off-line are assumed to be followed as simulated, but variations can occur due to deviations between the simulated robot and actual robot controller. The metrology requirements for automated robotic NDE systems can be summarised as follows:

- What is the accuracy, resolution and repeatability of a robotic NDE system?
- How accurately does the industrial robot follow the expected trajectory?

- How accurately is the robot reporting the position of the ultrasonic transducer?
- When using cooperating robots, how confident can the operator be that the robots maintained alignment?
- How accurately can the location of an indication be encoded taking into consideration the limitations of the robots and ultrasonic probes?

Unlike part metrology, these questions help the NDE users quantify and understand the limitations of a robotic NDE system. A key contribution of this thesis is the validation of these systems. A further discussion on data presentation and reliability is found in [63].

As evident from the above, two separate systems are needed wherein one provides measurements about a component that is to be inspected while the other provides positional data about the robots themselves. An optimal solution would be one wherein the two systems can complement each other. For example, a system used to track the robots could simultaneous be used to provide positional data about the part if the robot was equipped with a tool such as a laser scanner. As with any metrology system, there is always a trade-off between cost, the complexity of the system, measurement speeds and accuracy, reliability as well as how the data is handled. Suitable solutions for part and manipulator metrology are reviewed in the following sections.

2.5.3 Metrology Methods for Large-Volume Tracking

Due to their main areas of application, industrial robots are optimised for repeatability, but not for absolute accuracy. To improve the accuracy a robot can either be calibrated or a positional feedback system can be used to correct its motion. Either approach requires some sort of metrology instrument and optical measurement systems are the most popular currently. Due to the size of robots, the ability to measure over large scales is denoted large volume metrology.

Laser Trackers

Laser trackers work by measuring one distance and two angles. The tracker emits a laser beam that is reflected off a target held against a feature to be measured. The reflected laser light is directed to a distance meter, which is either an interferometer or Absolute Distance Meter (ADM). Interferometry relies on the overlap of monochromatic light, to calculate distances. By considering constructive and destructive interference, the difference between a reference wave and reflected wave can be found and the distance inferred [64, 65].

Besides the laser light itself, a laser tracker also contains two angle encoders measuring the azimuth and elevation axes. An illustration of these angles and distance are shown in Figure 2.15. Using knowledge of the two angles and a distance, the 3D position of the reflector can be calculated. Although the accuracy of the interferometer is around 0.5 microns per metre, the uncertainty from the angles mean that the overall system has an accuracy between 25 to 50 microns.

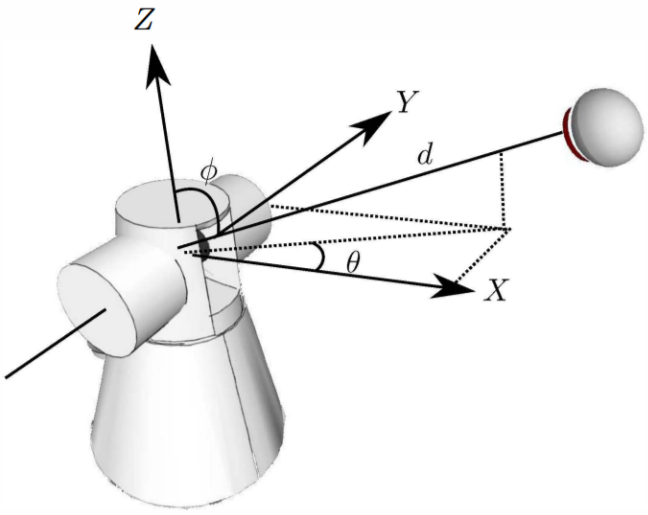

Figure 2.15: Illustration of a laser tracker. Accurate distance and angle measurements are combined to give the position of a point in 3D space. *Image credits: [65]*

The main distributors of commercial laser trackers are: Leica, FARO and API. Laser trackers are used for both dynamic and static applications and typically feature positional feedback rates of several thousand points per second. Though expensive,

laser trackers provide the most accurate commercially available measurements at large distances (ranging up to 50 m). As with any optical measurement, the accuracy deteriorates the further away the target is from the device. Most laser trackers only provide three DoF positions, but some manufacturers include extra options to give full six DoF information. This extra functionality, however, tends to also come with a large increase in cost.

Photogrammetry

Another method for providing accurate measurements in large volumes is through photogrammetry. By identifying the same point in different camera views, it is possible to reconstruct the 3D position of that point through triangulation, as shown in Figure 2.16. Photogrammetry can be used both for large scale applications like aerial mapping and to reconstruct surfaces of small items, like tooling or fixtures [66]. To obtain data from different viewpoints the camera is usually moved, but it is also possible to use two or more cameras with a known spatial separation to reconstruct a 3D scene. This is the principle behind stereovision cameras.

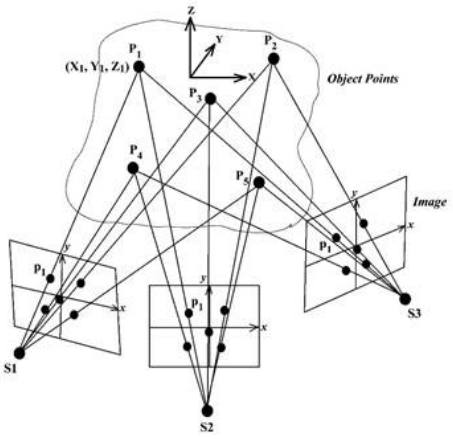

Figure 2.16: Reconstruction of 3D points through photogrammetry. Camera positions are denoted by **S1-3**. Each point is identified across multiple views to form a 3D location. *Image credits: [67]*

One of the key challenges in using photogrammetry is the identification of the same point across several images. Solving this problem was an intense focus of the machine

vision community for many years [68]. These feature points are typically defined by rapid changes in intensity, but natural features can change appearance due to variations in illumination or viewing angle [69]. For featureless surfaces (like composite components), one may choose to fix markers to an object and design a computer program to track this marker instead. This is the basis of motion capture systems. Markers are typically chosen to be flat reflective targets or spheres coated in a reflective material.

To be able to measure across a large volume, several cameras can be fixed in the environment and as long as markers can be seen by at least two cameras, the position of said marker can be found. This creates a more flexible system than a laser tracker as no single line of sight to the object is needed. Several objects can also be tracked simultaneously at high feedback rates. Defining a set of markers as a rigid body allows one to infer orientation as well. The disadvantage is that these systems are not as accurate due to noise in the measurements and approximations in reconstruction algorithms.

Indoor GPS

Indoor GPS (iGPS) system works in a similar way to photogrammetry in that several transmitters are placed around an object to be tracked [70]. Each transmitter consists of a rotating laser head theodolite which emits strobes of light at a particularly frequency. A sensor can then pick up signals from a number of transmitters and distinguish between them based on the frequency of strobes. With a distributed network of transmitters, a sensor can be tracked over large volumes without the need for a single line of sight. The distributor of this system, Nikon, advertise that a 200μ m accuracy is possible.

2.5.4 Metrology Methods for Surface Mapping

At the most basic level, a physical measurement can only take place if there is some sort of information conveyed between the object needing to be measured and the tool measuring. This can be done either via physical contact or by light being reflected off the surface and captured by a system. Metrology devices can therefore broadly be categorised as either being contact (e.g. touch probes) or non-contact (e.g. laser scanners). Many contact based metrology devices are familiar hand tools such as callipers,

gauges, and micrometers.

The industry standard for obtaining high-accuracy surface measurements is the Coordinate Measurement Machine (CMM). This is a contact-based system which uses a rigid gantry to move a touch probe to a number of points on a component. CMMs are calibrated to very high accuracies (below 10 microns) but are limited to smaller samples (typically around 1.5-2 m), are expensive and take up significant floor space.

A useful diagram, shown in Figure 2.17, compares the relative resolution of various optical measurement methods. One of the issues that an end-user of a metrology system faces in their decision on what supplier to buy from is the difference between the stated (under optimal conditions) accuracy and the accuracy achieved in practice under normal working conditions. Other studies [71, 72] have compared the performance of different measurement systems to traceable calibrated artefacts, but in general there is not yet an internationally agreed upon standard for evaluating optical coordinate measurement systems.

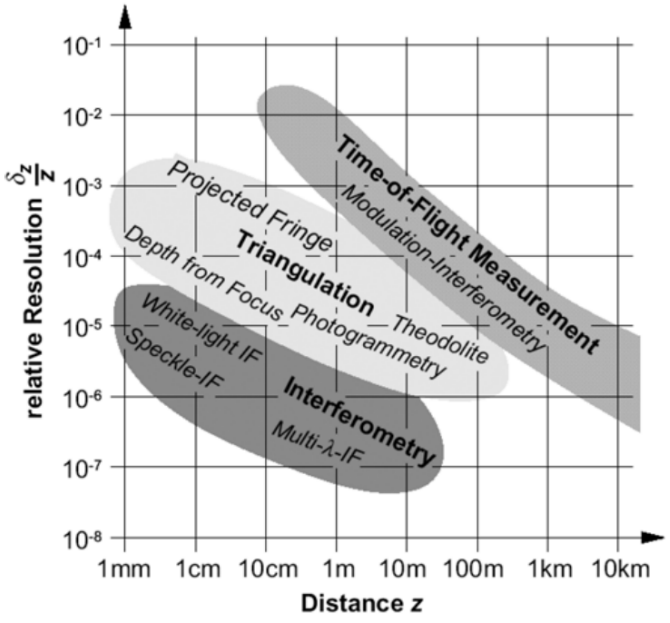

Figure 2.17: Resolution comparison of different optical metrology systems for obtaining 3D information [73].

Laser Triangulation

This technique uses a laser line emitted on a surface and a camera to record the position of the reflected light. The laser, object and camera positions form a triangle, hence the name laser triangulation. The basic mechanism of a laser triangulation device is shown in Figure 2.18. If the surface is non-specular (so some part of the light will reflect at an angle) the detecting lens will focus the reflected laser light onto a detector. Knowing the location of the highest charged pixel will then indicate the angle of reflection which in turn can be interpreted as a distance through trigonometry. As the charge builds up for longer exposure times, sensitivity to weaker reflections increases over time [74]. For specular (mirror-like) surfaces where the light from the laser would bounce back to the laser itself, the laser light is emitted at an angle as shown in Figure 2.18.

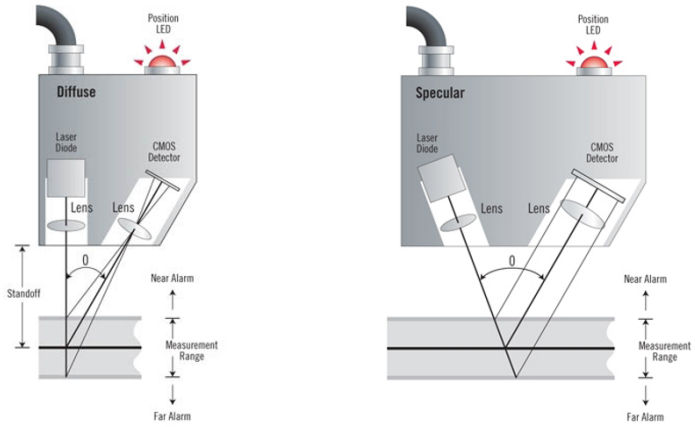

Figure 2.18: Principle of a laser triangulation distance sensor [74]. Two different camera and laser orientations are employed for either diffuse or specular surfaces.

An issue with these types of devices is their limited measurement range which is dependent on the angle between the laser emitter and camera. A typical laser triangulation device will have a measurement range of a 20-100 mm but very high resolution (better than 50 μm). The resolution here is defined as the smallest distance which can be measured reliably. The spatial resolution of the system depends on the size of the laser spot as features smaller than the spot itself cannot be accurately measured. The advantages of using laser triangulation devices is their lower cost (compared to laser

interferometers), non-contact measurements as well as high frequency sampling rates (several kHz). Disadvantages include that accuracy falls off rapidly with increasing range as well as the environment in which they are used has to be kept very clean. Another issue that needs to be considered is the scanning pattern. Sudden changes in the height of an object can lead to hidden areas where the laser scanner is unable to capture data if the camera is unable to observe the laser line.

Structured Light

By projecting a pattern of light onto a surface, any irregularities or deformations of that surface will appear as distortions in the geometry of the projection. The projected patterns are typically in the form of lines, dots or fringes and light frequencies outside the visible spectrum may also be used to avoid interference with other sources of light in the area. One of the most common techniques is to project a series of stripes upon an object as illustrated in Figure 2.19. An image processing algorithm counts the stripes and measures any change in width of a stripe to provide information about the depth of the surface. Thinner stripes correspond to steeper surfaces and a combination of pattern recognition and Fourier transform are implemented to convert the observed pattern into 3D coordinates [75]. The clear advantage of these types of system is the large area that can be covered from a single viewpoint.

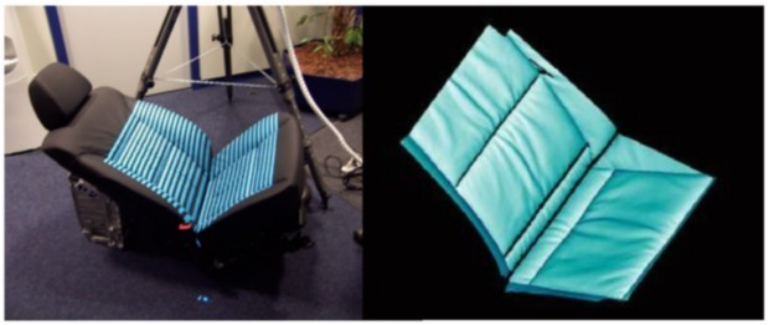

Figure 2.19: Example of structured light patterns used for surface reconstruction. *Image source: Wikimedia Commons*

Typical resolutions are on the sub-millimetre scale and the range of these scanners is up to several meters depending on the power of the projector as well as the en-

vironmental influences. As this technique is becoming increasingly used, several more inexpensive systems are now available. A popular device from the entertainment industry is the Microsoft Kinect camera which uses an infrared laser as well as a combined infrared and RGB camera [76]. However, the accuracy of this system ranges from a few mm to several cm at the maximum range. Other disadvantages are the depth of field (the camera has to be able to focus on the scene) as well as the lack of standardised calibration methods.

2.5.5 Examples of Metrology for Automated Inspection Systems

As mentioned in Section 2.4.3, most commercial systems rarely mention how the discussed metrology challenges are addressed in their systems. This section reviews two systems that have been developed by academic groups, namely the University of Dayton Research Institute (USA) and the University of Strathclyde (UK). Other academic groups who are investigating these types of systems include the London South Bank University, University West in Trollhättan (Sweden), Beijing Institute of Technology and Ecole Supérieure des Technologies Industrielles Avancées in France.

TESI Robotic System

Kramb et al. [46, 47], however, provides a thorough analysis of the obtainable accuracy of their engine turbine blade inspection system. The system compromises a phased array probe attached to an industrial robot which scans engine disks rotated on a precision turntable. A calibration of the industrial robot was performed beforehand to increase the positional accuracy to be better than 0.25 mm per metre moved [46]. It is assumed this accuracy is measured without the ultrasonic probe mounted to the robot. The inspection system, known as Turbine Engine Sustainment Initiative (TESI) is shown in Figure 2.20.

Besides the obvious improvement in robot positional accuracy through calibration, the TESI system furthermore includes an automated probe alignment check. This is performed in immersion by scanning an alignment check block with machined targets. By knowing the actual position of machined targets, accuracy of each reported indi-

Figure 2.20: The Turbine Engine Sustainment Initiative (TESI) inspection system, developed by researchers at the University of Dayton Research Institute (UDRI) for the US Air Force [46].

cation by the phased array probe can be verified. If the system does not meet the tolerances set by the scan plan, testing is aborted. This type of check allows the operator to be confident that the system is reporting back the correct position of the piezoelectric elements and their orientation which can be checked by monitoring the amplitude of the reflected signal. After 50 runs, it was reported that the average positional deviation was around 0.5 mm and the average orientation error was well below 0.5° [47].

The TESI system is one of the few automated NDE systems in the literature where a traceability and calibration procedure is laid out in detail. Besides this important aspect of the system, another metrology application is also incorporated. To ensure the system is fully automated, the robot must be able to ascertain where the engine disk is with respect to the robot's reference frame. This is achieved using an industrial vision system consisting of a 4 megapixel (MP) camera with an attached illumination ring. The robot positions the camera one metre above the component on the turntable and a commercial machine vision package performs a series of checks. These include if the correct component is present, if that component is in the right orientation and whether

or not it is off-centre. These checks are performed with an accuracy of approximately 0.075 mm [47]. It should be noted that this procedure is only demonstrated on a single type of geometry and it is not clear if the method can quickly be modified for other components.

VIEWS at the University of Strathclyde

The Validation and Integration of Manufacturing Enablers for Future Wing Structures (VIEWS) project at the Advanced Forming Research Centre in Glasgow has developed a robotic inspection cell in collaboration with Spirit AeroSystems [50, 77]. The cell compromises three KUKA industrial robots on linear tracks wherein one gathers ultrasonic data using a roller probe while the others perform surface metrology. An annotated image of the cell is shown in Figure 2.21. The VIEWS cell overcomes the challenge of part position and variation by using a force-torque sensor to correct ensure the ultrasonic probe follows the surface at all times. This approach imposes a limit to the curvature that can be inspected due to the size of the probe, but it does helps solve the question of how accurately the robot follows a trajectory. As the path is constantly adapted to the part geometry, the need for an exact match between off-line and real world trajectories is removed.

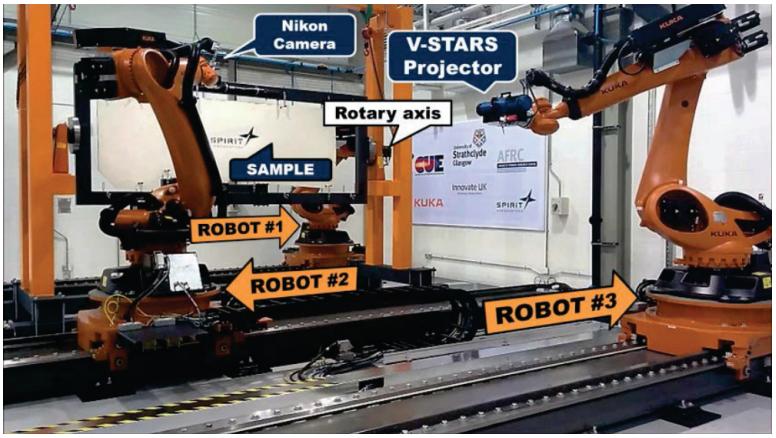

Figure 2.21: VIEWS robotic inspection cell at the Advanced Forming Research Centre. Image credits: [77].

While the ultrasonic probe is capable of dealing with variations in geometry, the

VIEWS cell also uses non-contact photogrammetry to determine conformance of dimensions with respect to a CAD model. This is achieved by attaching a projector to one robot while the other robot moves a high-resolution camera around to photograph the sample from different viewpoints. By projecting up to 22,000 points onto the sample, the camera is provided with key points for 3D reconstruction. Additionally, a custom-built scale bar must be present to determine real-world distances from pixel values. This process, however, takes considerable time to set up and does not provide consistently good surface data. The novelty of this system is the integration of NDE and dimensional metrology, two fields of testing which are closely related, but often used at different stages of the manufacturing process.

2.6 Motivations for Chosen Approach

It is clear from the literature that automated NDT systems using industrial robots have become prevalent on the market within the last four years, coinciding with the work undertaken for this thesis. Traditional teach programming is rapidly being phased out and replaced by off-line path planning platforms, of which there are now a large number of distributors. This reflects a broader change in manufacturing which allows for more complex parts to be made such that systems must have the ability to generate paths from CAD models. There is however a clear lack of rigour between moving from the digital world where paths can be simulated perfectly to the real situation where the component may not be in the right place, or conform to the expected geometry. The following gaps were identified in the literature on automated NDT for large composite structures using industrial robots:

- The calibration between digital and the real-world environments needs to be addressed
- A higher automation level is needed to determine the position and location of parts to be inspected
- Validation techniques are needed to determine the true accuracy and reliability

of robotic automated inspection systems

- Methods are needed to address variations in component geometry for new carbon-fibre reinforced polymer parts used in the aerospace industry.

The limitations of current off-line-path planning platforms are clearly in the inaccuracies between workpiece location and modelled work environment as well as manufacturing tolerances between CAD models and actual parts. Furthermore, tool calibration methods specifically for NDT are not addressed adequately in the literature. The current work investigates methods for overcoming these limitations through a novel tool calibration method, the use of computer vision and motion capture systems for part localisation. This thesis also investigates a method for surface scanning, error estimation through tracking and ways of adapting robotic paths to account for misalignments using front-wall ultrasonic echoes.

Chapter 3

Positional Accuracy for Automated Inspections

3.1 Introduction

This chapter introduces Off-Line Path (OLP) planning, the method by which most modern tasks are designed for industrial robots and how the accuracy of the robot is fundamental to successfully carrying out these tasks. The kinematic model of an industrial robot is presented and causes of inaccuracies are identified. Methods for calibrating robots and integrating sensors with an industrial robot is presented in the context of NDE inspections. The calibration of sensors is vital to ensure that any positional errors kept to a minimum. This chapter presents research conducted to develop and optimise calibration procedures between the relevant coordinate reference frames for NDE inspection processes. In particular, a tool calibration method based on the reflected ultrasonic signal is presented. To increase the accuracy and automation level of the inspection system, a new method for accurate part position calibration will also be presented in this chapter.

3.2 Describing Industrial Robots

According to the International Standardization Organization (ISO), an industrial robot is an "*automatically controlled, reprogrammable, multipurpose manipulator, programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications*" [78]. Though the field of robotic and automated NDE covers a large number of robots, this thesis only concerns the use of fixed six-axis industrial robots for NDE, as shown in Figure 3.1. These types of robots are used in a large number of industrial situations such as welding, palletising, pick-and-place or spraying applications where the flexibility and working envelope of the robot is critical to the successful execution of the task.

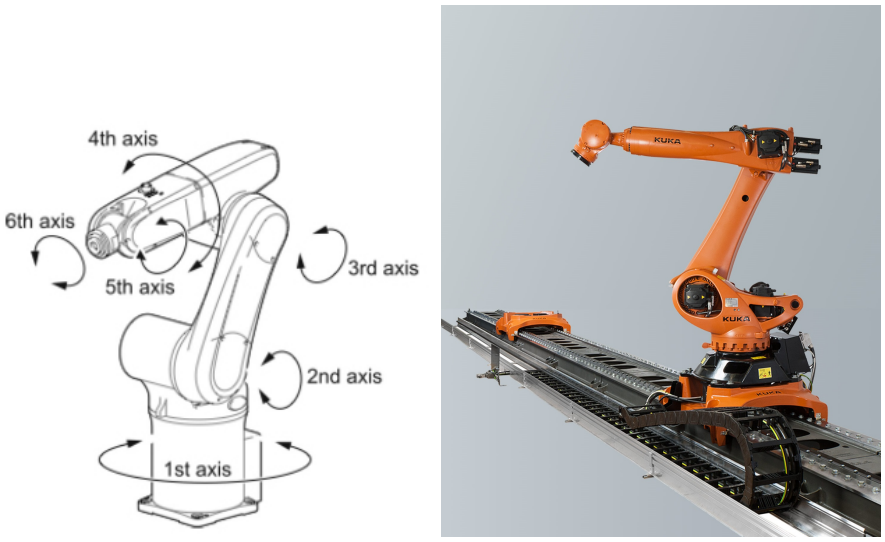

Figure 3.1: Outline of a six-axis industrial robot with labels denoting each axis [79].Six-axis robotic arm with a linear positioner rail. [80].

A six-axis robot allows a larger reach and is able to reach into confined spaces by virtue of having several links connected serially. Of particular interest to the NDE sector, is the ability to manipulate an inspection tool across surfaces or follow complex features such as edges. Having six separate axes, these robots can position the final joint (often denoted the "wrist") with six degrees of freedom anywhere in their working envelope. This arrangement of serial links allows the robot to both mimic the human

arm, but also overcome some of the biological constraints, for example by being able to spin the 4^{th} and 6^{th} axes continuously. This is an important feature as NDE probes designed to be operated by a human can easily be attached to a robotic arm with little modification and allows the operator to program the robot more intuitively. The reach of a robot is limited by the length of each of the links, but larger components can be encompassed by introducing additional axes such as turntables or linear rails as shown in Figure 3.1.

3.2.1 Mechanics

As stated in the ISO definition, an industrial robot fundamentally consists of a physical manipulator (a series of rigid links with attached motors) and a control system which dictates how each motor should move to position the robot in 3D space. Each link in an industrial robot will typically have its own actuator which is optimised for high-torque and low rotational speeds to ensure smooth and accurate motions. Pre-tensioned timing belts are used to minimise backlash and to transmit torque more effectively, while high gear ratios enable robots to have good repeatability, even with large payloads [81]. The industrial robot market has grown significantly in recent years [82] and today it is possible to find a wide range of robots for specific applications from a world-wide selection of companies. Throughout this thesis, KUKA robots were used as their models have a combination of payload, accuracy, repeatability and data transfer rates which are suitable for NDE applications.

The movements of an industrial robot must be executed within the range of possible movements. The robot must not, for example, rotate the third axis to such an extent as to cause a collision with the second axis. To avoid such situations, most industrial robots have both a software limit on how far each joint can rotate and also hard-wired switches in each link for safety. The range of all possible movements, restrained by these limits and the length of each link is known as the *working envelope* of the robot.

To execute a pre-programmed trajectory, the robot controller (which typically located in a cabinet next to the robot), applies a voltage to each motor to move each joint by a certain amount, as calculated by the kinematic model. This control system

is by no means straightforward and most of the time, each motor will have to move at different speeds. Throughout this movement, the robot controller will apply acceleration and deceleration curves to the trajectory. Furthermore, the gravitational and inertial forces of each link as well as the tool at the end of the robot has to be taken into account.

3.2.2 Kinematic Models

An industrial robot is driven as an open-loop control system, insofar there is no independent feedback about the position of the tool. Instead, motors are driven until each joint encoder reaches a specified value, determined by the kinematic model of the robot. The encoder values are then also used to provide the position of the robot end-effector. A forward kinematic model is used to compute to the pose of the tool given joint encoder values. Conversely, the inverse kinematic model computes which actuator positions give a desired tool position. Hence, any error in the kinematic model can give rise to an incorrect tool pose and without a method of external feedback, the robot will have no way to compensate for this inaccuracy.

There are a number of different kinematic models which can be used to describe six-axis robots, but by far most are based on the Denavit-Hartenberg (DH) convention [11, 83] due to its geometrical intuitiveness and simplicity. The pose of any joint can be found by considering the pose of any previous joints, knowing the rotation and translation between each link. This can be described as the product of a series of matrices, as shown in Equation 3.1. Here T represents the position and orientation of link j with respect to link *j-1*. Further information on representing orientation and position in 3D is given in Appendix A. The twist of each link is described by α while the rotation angle about each joint is θ. The length of a link, a, is the distance between the centre of rotation of each joint and d is the offset between joints along the axis of rotation [41]. This relationship between rotational and translational parameters is illustrated in Figure 3.2.

$$T_j^{j-1} = f(\theta_j, \alpha_j, d_j, a_j) = R_z(\theta_j)T_z(d_j)T_x(a_j)R_x(\alpha_j) \tag{3.1}$$

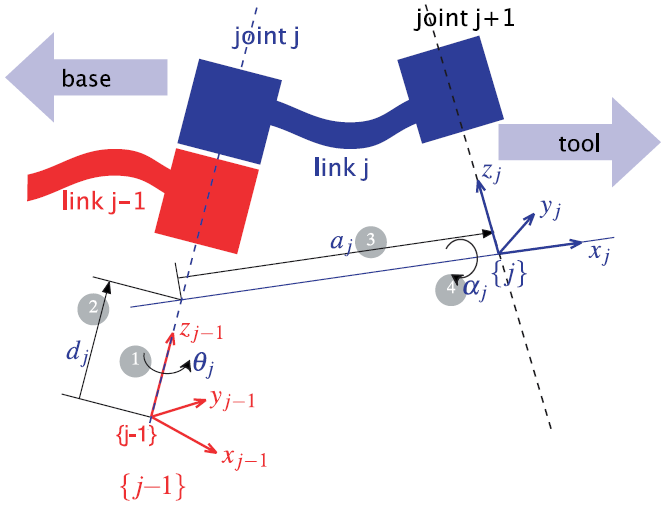

Figure 3.2: Illustration of the Denavit-Hartenberg kinematic model convention showing the connection between two links [41]. This relationship is expressed mathematically through Equation 3.1.

Equation 3.1 can be expressed as a single matrix incorporating rotations and translations by expanding a 3x3 rotation matrix to include a translational part, forming a 4x4 *homogeneous matrix*. Homogeneous matrices are discussed in further detail in Appendix A. The full expansion of Equation 3.1 is written as

$$T_j^{j-1} = \begin{bmatrix} cos(\theta_j) & -sin(\theta_j)cos(\alpha_j) & sin(\theta_j)sin(\alpha_j) & a_j cos(\theta_j) \\ sin(\theta_j) & cos(\theta_j)cos(\alpha_j) & -cos(\theta_j)sin(\alpha_j) & a_j sin(\theta_j) \\ 0 & sin(\alpha_j) & cos(\alpha_j) & d_j \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{3.2}$$

The forward kinematics of the robot consisting of N links is then described by Equation 3.3. This equation illustrates the usefulness of expressing rotation and translation simultaneously as a homogeneous matrices since the final pose is simply the product of the individual link transformation matrices.

$$T_{End}^0 = T_1^0 T_2^1 \dots T_N^{N-1} \tag{3.3}$$

The forward kinematic model uses the information captured in the DH model along with a dynamic model of the robot to drive each actuator to the desired angle to

obtain a commanded pose in 3D space. The inverse problem is how to determine the combination of angles of each joint to obtain a desired 3D pose. For a six axis robot, there are up to eight different configurations for a given pose. This can lead to potential conflicts, especially if some part of the robot's payload requires the robot to stay in a particular configuration. This can for example happen when wires or hoses travel from behind the robot to the wrist. These can easily be entangled or get caught in the moving joints of a robot. Furthermore, some poses are physically impossible as the robot would have to bend in on itself, but these scenarios are avoided as mentioned earlier.

Another problem which can arise in robot programming is singularities. These occur when two axes are aligned such that one becomes redundant or when the robot would need to rotate a joint infinitely fast to achieve continuous end-effector motion. This situation is also known as *gimball lock*. In terms of the kinematic model, these singularities occur when the axes of joints become aligned such that no unique solution exists for the end-effector pose. To find the relationship between joint angle velocity and end-effector velocity, the Jacobian matrix of the forward kinematics in joint space is used. If the determinant of the Jacobian is zero, the joints of the robot would have to rotate infinitely fast to obtain a constant end-effector velocity [41].

The requirements of robots today are to accurately perform a task, preferably at very high speeds with minimum wear on the robot itself while avoiding collisions. These high-performance requirements are difficult to meet if a complete model of the robot needs to be solved at every program iteration. Hence, a model-reduction approach is employed to find the balance between capturing the most important aspects of the kinematic model of a robot while meeting "real-time" interpolation targets [84]. The identification of key parameters is a difficult task, to say the least, as each robot represents a multi-variable, non-linear system that is constantly affected by the task at hand. These effects can for example be the inertial forces of the tool used or speed- and acceleration-dependent ripple.

3.3 Robot Programming

Since the invention of the first industrial robot by Devol and Engelberger in 1956 [85], the method of programming robots has evolved dramatically. The first ever commercial robot, *the Unimate*, was driven by encoding a series of instructions for joint angles on magnetic drums. Though cumbersome, the process nonetheless achieved the goal of creating an automatically controlled, reprogrammable manipulator that could be used for a variety of purposes. As technology developed and processing power became more portable, robots began to come equipped with their own micro-controller units. This allowed an operator to type in a series of Cartesian coordinates using a text-editor or Graphical User Interface (GUI) . Before computers became as ubiquitous as they are today, industrial robots were mainly programmed using a teach-pendant, as shown in Figure 3.3(a).

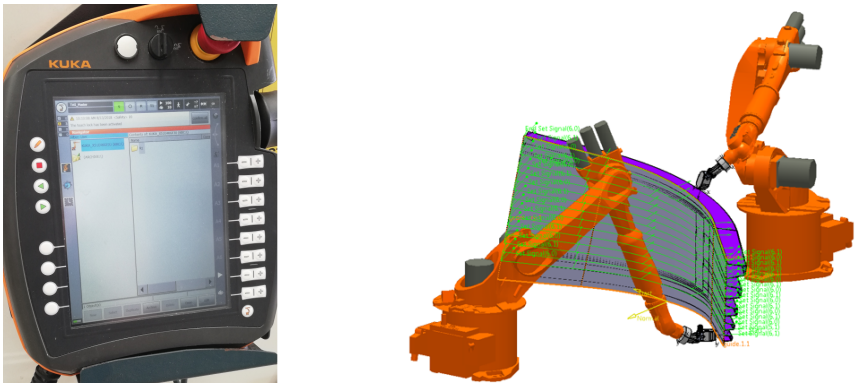

(a) (b)

Figure 3.3: (a) An example of a touch-screen teach pendant. (b) Off-line path planning application.

Using a teach-pendant allows an operator to easily drive the robot to a series of positions and record the encoder's value at each location. The robot controller then interpolates between the saved poses to create a series of basic movements such as linear or circular motions. The manual approach to robot programming is sufficient for many of the traditional robot operations such as welding, moving objects, painting, etc. For safety reasons, the robot is driven in "teach mode" which limits the maximum velocity

the robot can move at to avoid injury. The benefit of this method is that the accuracy of the robot becomes irrelevant as only repeatability is important. Although conceptually easy, manual teaching requires skilled operators and programming can be very time consuming. Some robot manufacturers also offer a "lead-through" programming method, where the robot motors and brakes are uncoupled, allowing the operator to manually move the robot in 3D space and record encoder values [84].

KUKA robots are programmed using a touch-screen teach pendant (Figure 3.3a) using modules written in the KUKA Robot Language (KRL). This is a domain specific language which is similar to Pascal [86]. KRL scripts are loaded into the KUKA Robot Controller (KRC) environment, of which several versions exist. During the work presented in this thesis, only KRC2 and KRC4 environments were used. The transfer of files is possible using USB or Ethernet connections. Depending on the additionally installed packages, it is also possible to obtain robot coordinates and pass signals to the KRC through an Ethernet connection.

The manual, on-line, teaching method is only adequate in situations where the number of programmed positions is relatively small and where the operator can intuitively instruct the robot how to accomplish a set of movements. In many cases, however, the required path is very large and complicated or the robot may have to fit into a larger automation scenario. As society moves towards Industry 4.0, the need for short turnaround jobs will mostly increase as well, requiring frequent reprogramming of robots. In these cases, it becomes impractical to attempt to program industrial robots by hand and since the 1980's [87], simulated robotic work cells have been used for off-line trajectory planning. By creating a virtual model of the robot and its environment, the needed joint-level trajectories can be calculated and simulated without having to move the physical robot.

3.3.1 Off-Line Path Planning

OLP brings a variety of benefits to robot programming. Paths can be planned without the need to stop production work and can be checked for collisions if the physical environment around the robot is modelled accurately. CAD models of components can

be imported into the simulation as well and paths can be generated on their surfaces as shown in Figure 3.3b - something a human operator working manually would struggle to do accurately. For large scale operations, it is possible to construct the necessary control architecture through digital or analogue signals to allow cooperation between several robots or other equipment such as conveyor belts. Cycle times can also be calculated and optimised, resulting in time and cost savings for the operator.

Today there are a large number of robot simulation packages which allow the user to both define robot movements, check for collisions, interact with components and other machines. Commercial Computer Aided Manufacturing (CAM) solutions include Delmia/Cenit FASTSURF, Siemens Kineo, KUKA Sim, ABB Robot Studio, PowerMill from Delcam and RoboDK, to mention a few. Non-CAM OLP solutions also exist, such as ROS-Industrial, KUKA|PRC, and V-REP. Besides software solutions which are brand-specific, most packages have the ability to choose between a large number of robots from a wide selection of manufacturers.

A drawback of the off-line approach to path planning is that it relies completely on the accuracy of the calibration between the virtual and real environments. To be able to detect and avoid collisions, the user must be confident that the surroundings of the real robot are in the same place as in the software. This is certainly a valid assumption in highly structured environments such as a production line in a factory. However, in line with the ambitions of the next industrial revolution, robots will most likely begin to work in more unstructured environments and even alongside humans. So not only must the accuracy be sufficient to ensure the path is correctly executed, but the robot must also be able to adapt to its environment. These changing requirements mean that OLP must be complemented by further sensor integration which can adapt the ideal paths [88].

In the context of the current thesis, the use of OLP for the inspection of complex, composite parts is favoured compared to the manual teaching method. The reason for this is that an inspection path typically requires the full surface of the component to be scanned, which is prohibitively complex for a human to accomplish using the teach pendant programming approach. Further to this, inspection is already a bottleneck in

the aerospace composite sector, so the ability to quickly generate robotic tool-paths is desired. As an example, the inspection of large engine casings can take up to two working days to inspect manually, but the same component can be inspected in a few hours using industrial robots. The author of [20] demonstrated the challenges of developing an OLP software solution in the MATLAB environment specifically aimed at the challenges in automated NDT.

The present thesis focuses on the further integration of sensors into the automated inspection system based on industrial robots. These sensors, (cameras and laser profilers), not only help improve the accuracy of inspections but further the automation levels by adding extra inputs to the inspection setup. Based on the work presented in [20], it is possible to expand on the schema outlining the robotic inspection procedure using OLP and external sensors. This updated schematic is shown in Figure 3.4 where the contributions of the current thesis are shown in green with red arrows indicating their relationship to the existing setup.

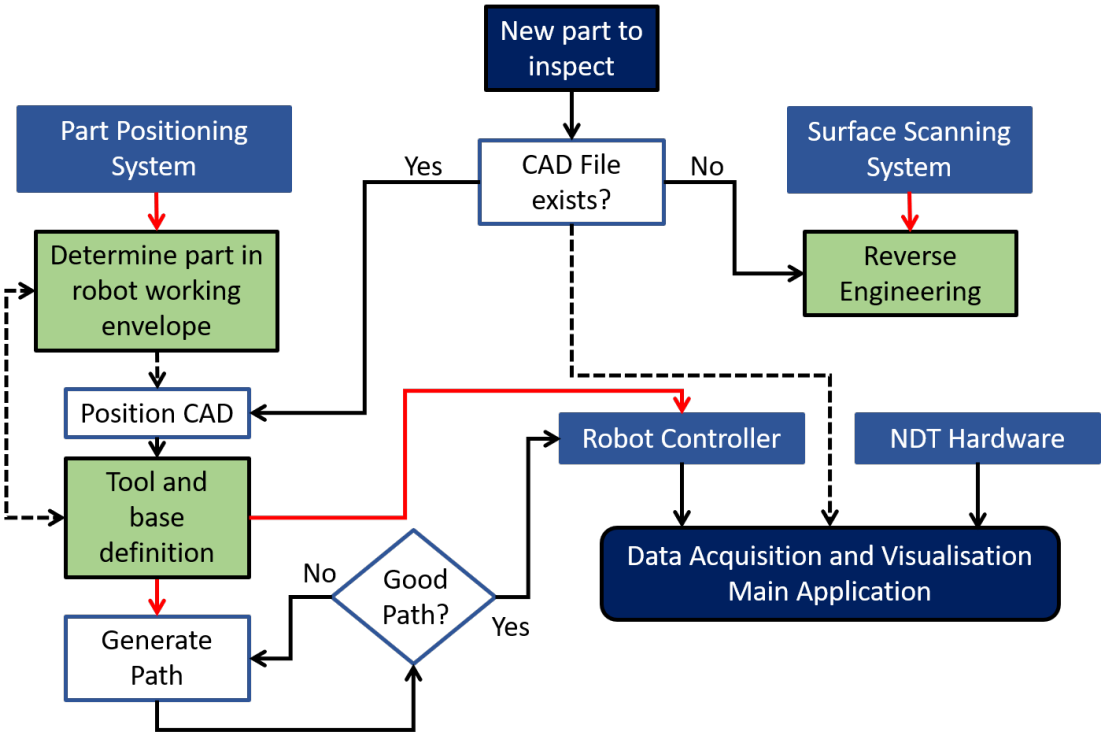

Figure 3.4: Schematic showing the contributions of the thesis to the inspection procedure based on off-line trajectory planning (green boxes and red arrows).

3.3.2 External Control

There are many ways an inspection system can be realised. Points on the surface of a part can be discrete and the robot's kinematic controller can interpolate between each point. Alternatively, the robot can be instructed to move a certain distance in each kinematic interpolation cycle. This approach is known as external control and is typically implemented when the robot's path is determined by sensor input, like a camera. Particular attention must be paid to planning trajectories in a way which uses reasonable acceleration and deceleration curves to ensure a smooth motion. Previous work by [20] showed how trajectory planning can be realised in the context of NDE. The author provides an in-depth discussion on the merits of external control versus conventional off-line path planning. Discrete points are used and the robot controller is relied upon to interpolate between them in most situations throughout this work.

3.4 Interfacing with NDT Equipment

It is possible to interface with KUKA robots through digital input/output (I/O) signals, analogue I/O or through an Ethernet port. If the Rapid Sensor Interface (RSI) [89] additional technology package is installed on the robot control unit then positional information can be exchanged with external applications at either 12 ms or 4 ms intervals. Alternatively another technology package, FastSend [90], can be used to obtain positions at feedback rates of up to 1 kHz. The RSI package is used throughout this work as it allows for expanded user customisation as well as the ability to transmit coordinates and signals back to the robot controller.

There are two available data communication protocols to send and receive signals via an Ethernet cable: User Datagram Protocol (UDP) and Transmission Control Protocol/Internet Protocol (TCP/IP). To be able to send data via a TCP/IP connection, the systems participating in the exchange must first acknowledge their presence on the network before data can be transmitted. This is typically known as a "hand-shake" between server and client. This aspect of the communication protocol makes the connection more reliable as each system must acknowledge when data is sent and received.

Furthermore, data is ensured to travel along the same network path and is thus guaranteed to arrive in the sequence it was sent.

Alternatively, UDP transmits data without first establishing a connection between server and client. UDP is thus a faster connection protocol but there is no guarantee that the data will arrive in the order that it is sent, especially for large network hubs with many connections. This can lead to congestion on the network which is why TCP/IP is often the preferred choice. For real-time feed-back and to ensure accurate encoding, a UDP connection is established between a server and the robots and extra code is written to determine if packets arrive out of order by monitoring the KUKA-specific interpolation tag known as the "IPOC" number.

The current work builds upon the previous system architecture discussed in [20] and uses a dedicated server application to acquire NDT data and robotic positional feedback simultaneously. The robotic data is sent to the server through a custom written RSI framework and is received roughly every 4ms. NDT data in the form of digitised ultrasonic signals are also received by the server application when triggered by a signal sent from the robot. To be able to encode ultrasonic data, the two data streams are timestamped with the same clock which starts when the server application is launched. The next section discusses how to encode received ultrasonic data accurately using this approach.

3.4.1 Accurately Encoding Ultrasonic Data

To fuse robot positions and ultrasonic data, each probe firing is associated with a single 6 DoF robot pose. It is assumed that the robot speed is low enough such that the probe does not move significantly during the sound beam's travel time. Both the robot positions and ultrasonic data are encoded with the same clock, but may not have arrived synchronously. It is therefore necessary to interpolate between the robot's positions to infer the position and orientation of the probe when data was gathered. This process is commonly referred to as pose estimation.

Initially, linear interpolation was used between manipulator positions, but it became clear that this can lead to errors if the robot moves on a curved path. As an example,

imagine the robot is following a sinusoidal path in space, as shown in Figure 3.5. The robot will report its position at regular intervals (up to every 4 ms for a KRC4 controller), shown as blue dots. The ultrasonic hardware may send data at the same frequency, but offset in time due to, for example, latency in the system. If a linear interpolation method is used, it is clear that the interpolated values will be further away from the true position of the robot than if a cubic interpolation method is used, with the exception of start and end points where the cubic method is less constrained. Figure 3.5 shows the difference in relative error from the actual curved path. The use of an appropriate interpolation method is a key consideration when trying to accurately fuse data from two different sensors. Cubic interpolation is favourable in the current context due to the low computational cost compared to the increase in accuracy.

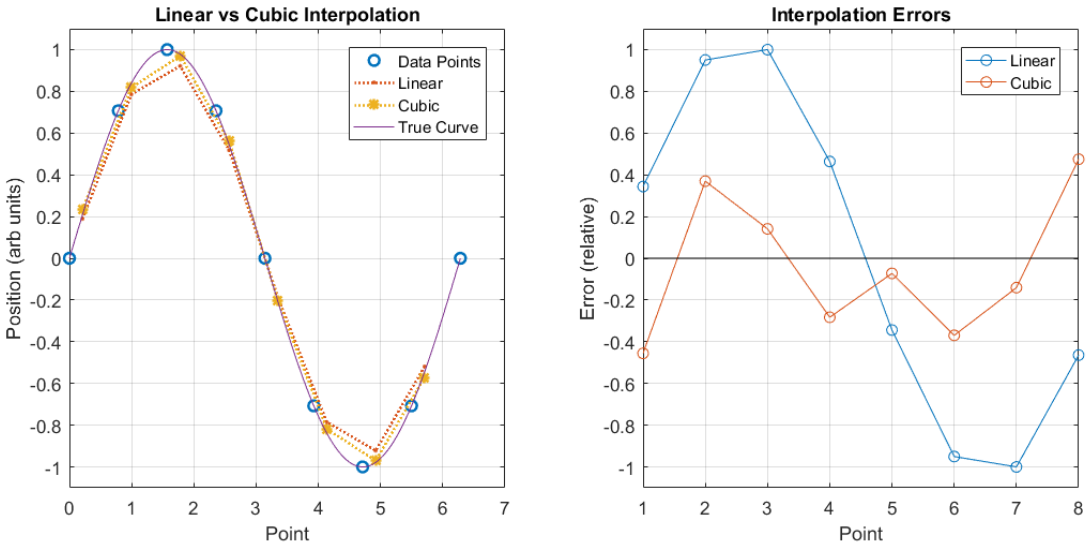

Figure 3.5: Interpolation methods: using a cubic interpolation algorithm for curved motions yields a more accurate positional encoding compared to linear interpolation.

Cubic interpolation can be applied to each of the three spatial coordinates using the associated time stamps to determine the 3D position of the probe. Along with Cartesian coordinates, Euler angles representing the orientation of the probe are also supplied. To estimate the orientation, the reported Euler angles were initially also linearly interpolated as separate variables. During the inspection of complex geometry components, the ultrasonic probe is likely to undergo rotations which a naïve interpo-

lation would fail to interpret correctly. Euler angles are usually limited to the range of [180°,-180°[, such that a rotation above 180 degrees reverts to the equivalent rotation anticlockwise from the origin.

However, if the tool moves across this threshold, an interpolated value would fall far from the real rotation. As an example, robot positional data might be captured while the tool is rotating about one axis from 177° to 178° and then -176° during a movement. Say the ultrasonic data was captured halfway between each of these robot positions. Treating each angle as an independent variable, a cubic interpolation would subsequently encode the inspection points with angles 177.88° and 67.44° which is obviously incorrect.

To avoid such situations, a different notation for representing rotations in 3D space is needed. Quaternions offer a simple framework for describing rotations as an angle about a rotation axis [91]. A quaternion is a four dimensional vector described by a scalar part s and a vector part, $\mathbf{v}$. Expressed together, the quaternion becomes $q = [s, \mathbf{v}], s \in \mathbb{R}, \mathbf{v} \in \mathbb{R}^3$. Appendix A describes quaternions in further detail. The advantage of using quaternions over the matrix representation is their compact notation and smooth interpolation methods. The disadvantage is that quaternions only capture rotation and not translation and that the mathematics is more complicated.

A standard interpolation method for quaternions is known as the "Spherical Linear Interpolation" (Slerp) method [91]. The interpolation method interpolates between two points on a unit-radius circle assuming a constant rotational velocity between the two poses. The algorithm needs the two endpoints and an interpolation parameter, t, which is obtained using the time difference in the current context. Using the same example as previously, the interpolated angular values using Slerp would be 177.52° and -178.98°, as expected.

Using cubic interpolation to determine the 3D position and quaternion interpolation for orientation, the full 6 DoF pose of the probe can be found. Ultrasonic echoes can consequently be mapped to a location in 3D space by projecting along the normal of the transducer face. For a single crystal probe, this will result in a single point. For phased array probes, the spatial separation between pulse generation sites must be

taken into account when using individual elements (or sub-apertures). To enable this type of encoding, one of the axes of the probe coordinate system is chosen to coincide with the active axis of the probe. Thus the accuracy of the tool calibration is crucial for encoding ultrasonic data correctly, which is the topic of Section 3.6.

When electronically rastering with a phased array probe, a number of ultrasonic signals are generated for each probe position. To project the origin of each signal from the ultrasonic probe reference frame, P_{tool}, the orientation of the probe is converted into a rotational matrix $R = R_z R_y R_x$ using the ZYX Euler convention for consistency with the KUKA programming environment. The tool coordinate system is such that $\hat{x}$ points along the active axis, so a vector, $V_{ut} = [\pm 1\, 0\, 0]$, can be formed. The position of each ultrasonic signal (P_i) is then encoded through Equation 3.4 where N is the total number of signals per scan, i is the current signal and p is the pitch of the probe.

$$P_i = P_{tool} + R \cdot V_{ut}^T \cdot p \left(\frac{N - i}{2} \right) \tag{3.4}$$

3.5 Accuracy of Industrial Robots

There are two important aspects defined by the international standard ISO 9283 [92] regarding the performance of industrial robots: accuracy and repeatability. Repeatability is determined by a robot's ability to return to the same point in 3D space after having moved to at least four other poses. As industrial robots have traditionally been used for repetitive tasks such as pick-and-place operations, their motions were only programmed once by a human operator and left to run continuously for extended periods of time. Thus, the ability of the robot to return to the same point again and again became the focus of manufacturers and is quoted in the specification data sheets for each robot model. The typical repeatability of a modern industrial robot is in the range of 10-100μ m [93].

The accuracy of an industrial robot is defined as the deviation between a commanded position in 3D position and the actual obtained position by the manipulator. Both the repeatability and accuracy of a robot are determined through measurements

by an external measurement device such as a laser tracker, theodolite or similar apparatus. There are a number of aspects regarding accuracy; first of all there is a difference between positional accuracy and dynamic path accuracy. Positional accuracy describes the ability of a robot to accurately reach a point while dynamic path accuracy describes how well a path can be followed by the robot. These two types of accuracy are illustrated in Figure 3.6.

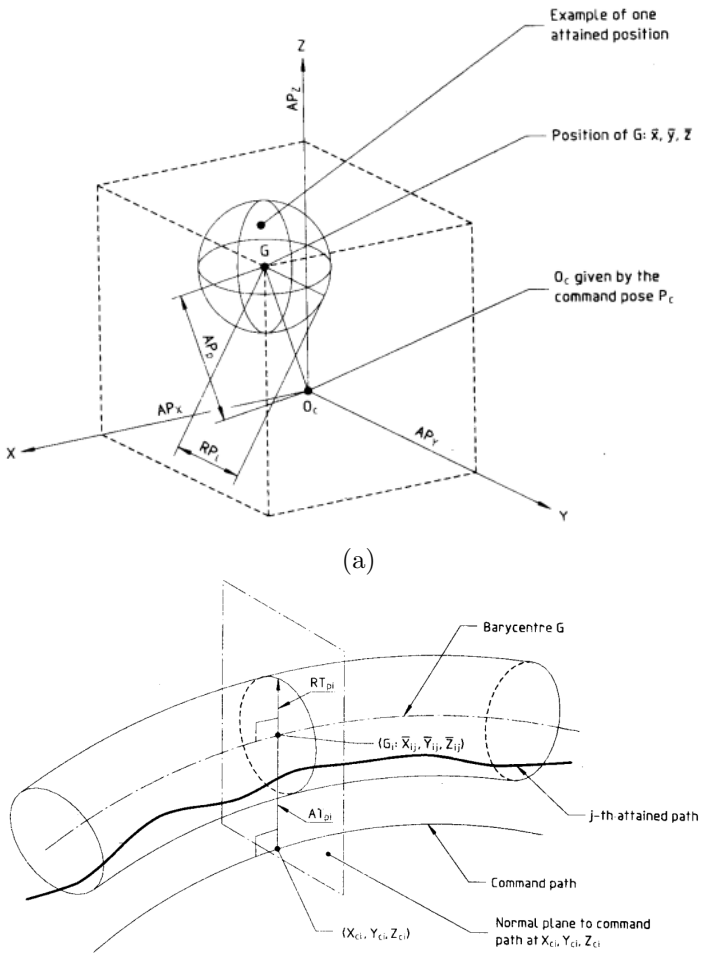

Figure 3.6: (a) Measurement of robot pose accuracy where AP denotes pose accuracy and RP is a measure of pose repeatability. (b) Measurement of path accuracy (AT_p) [92]

Besides positional and dynamic accuracy, one can also distinguish between absolute and relative accuracy. Absolute accuracy would then be the ability of the robot to

move to a point in 3D space in a predefined coordinate system (such as the internal reference system of the robot) while relative accuracy would be measured relative to a point the robot has already moved to. It is rare to see accuracy figures mentioned in robotic sales brochures, as the obtainable accuracy is highly dependent on the end-effector and coordinate reference system calibrations. In general, deviations between the commanded pose and the obtained pose are caused by a number of factors, including [94]:

- Static effects due to low manufacturing tolerances, joint off-sets, gear backlash, wear on joints and bearings, expansion of links due to temperature or low stiffness of links and joints.
- Errors arising from deviations between the kinematic model of the robot and the physical manipulator.
- Dynamic effects due to friction and inertial forces or other dynamic influences not incorporated into the kinematic model.
- Interpolation errors due to low resolution encoders or choice of algorithm in the robot controller.
- Gravitational effects due to the centre of mass moving as the robot arm is extended. These can also be caused if the weight of the tool is not taken into account.

One way of overcoming the above listed limitations is through calibration which uses measurements to update the kinematic model of the robot so that it more closely resembles the physical robot. A secondary calibration can in turn be carried out to counter dynamic effects, though this is substantially more difficult [95]. The following sections discuss the different types of calibrations relevant to a robotic automated inspection system.

3.5.1 Reference Frames

Due to changing manufacturing ethos that is currently sweeping industries throughout the developed world (see Section 2.1.1), demands are shifting from large scale production to small-batch, customised products. This demand requires a flexible approach to manufacturing and OLP is playing an increasingly larger role as robots take over traditional manufacturing processes. To facilitate this approach to flexible manufacturing, the digital model of the robot and its environment must match the physical setup. If the digital setup is not a so-called "twin" of the actual setup, the owner risks extensive material damage due to collisions or potentially danger to human operators in the area. It is of course possible to install sensors enabling robots to sense their environments, but this can be costly and cumbersome process which in turn relies on the strength of adaptive algorithms to compensate for changes in the environment. A cheaper and more straightforward way to overcome discrepancies is to carry out calibrations to align the digital and real world setups.

In general, three key calibrations are required in a robotic inspection system. These three calibrations (labelled T_y^x) and their relations to each other are shown in Figure 3.7. An industrial robot is supplied from the manufacturer's side with a mathematical model establishing the relationship between a global origin (Robot Frame, {O}) between a point inside the primary axis and the centre of the final link as discussed in Section 3.2.2. The second calibration establishes the geometric relationship between the final axis of the Wrist Frame, {W} and the reference frame of the tool being used, the Tool Frame, {U}. Finally, a relationship between the robot's internal origin and some coordinate frame fixed to the component being worked on is required as well, shown as the Part Frame, {B} in Figure 3.7.

Although the ISO 9283 standard [92] provides a mechanism and vocabulary for determining the accuracy of an industrial robot, the standard is only aimed at determining the accuracy of the T_W^O transformation, which in turn is determined by how well the chosen kinematic model describes the physical robot. This is an important parameter as it establishes a baseline value for the accuracy that an integrator or operator

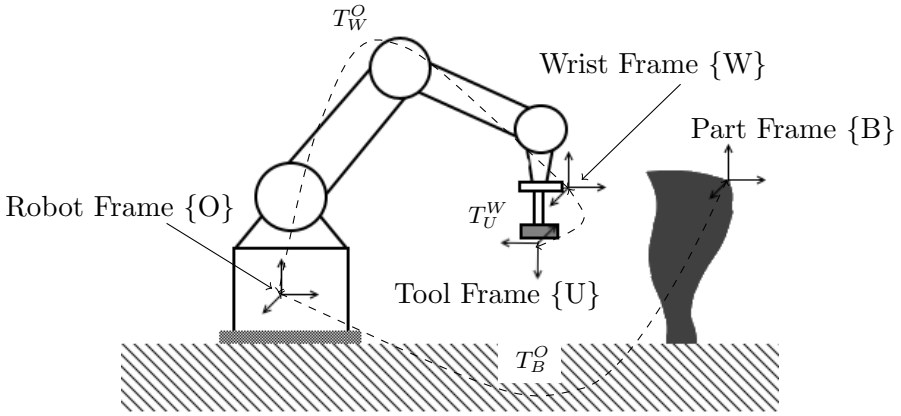

Figure 3.7: Schematic of the different reference systems used in a robotic inspection system.

can expect from the robotic system. In an inspection setting, however, the objective is to ensure an NDE probe is moved in an optimal path over the surface of a component while accurately encoding signals which could indicate potential defects. Thus the parameter of interest is not only the accuracy of the robot itself, but the overall accuracy of the system as a whole, taking all three key calibrations into account.

The obtainable accuracy of an inspection system naturally depends on the type of application. Thermography inspections, for example, generally have poor depth resolution due to the diffusivity of heat in most materials. X-ray Computed Tomography (CT) inspections can achieve extremely high resolution images but depend on knowing the exact positions of source and detector. For ultrasonic inspections, the choice of probe and the material properties of the component to be inspected both put limits on the minimum size of indications which can reliably be observed, which in turn limits the spatial resolution of the inspection. No matter what the choice of NDE technique is, calibration methods are needed to minimise the error in each of the three key geometrical relations shown in Figure 3.7.

3.5.2 Industrial Robot Calibration

The definition of the Denavit-Hartenberg parameters becomes essential in ensuring that the robot is able to move accurately and precisely to a specified point in space. However,

as the control programme for the robots is typically designed for a specific *type* of robot, and not for the individual robot, small deviations in manufacturing can lead to the robot moving inaccurately. Thus robots are typically calibrated upon delivery to ensure that the accuracy of the robot matches that stated by the manufacturer. Calibration can be either static or dynamic. Static calibration identifies static parameters such as link lengths, coupling factors and deviations in axis orientations. Dynamic calibration identifies parameters such as friction and loading forces and requires considerable effort [94].

A standard calibration procedure involves having the robot visit a number of poses and using an external metrology system capable of 6 DoF measurements to record these movements. Coordinates of the robot controller are compared to those obtained using an external metrology system for a large number of points (typically over 100). The difference in position between the kinematic model prediction and the observed points are then minimised by updating the kinematic model parameters through a non-linear least squared method [96, 97].

The deviations in segment length are one of the main factors influencing robot accuracy. According to the authors of [98], a length deviation of ± 0.1 mm and angle uncertainty of 0.000375° can lead to a final positional inaccuracy of approximately ± 0.3 mm. This is only due to machining tolerances and does not include weight loading or temperature effects. As previously mentioned, further additions to the overall error include friction, inertia, gear backlash, lack of rigidity and the fundamental limit of how accurately each encoder can describe a movement. It is also worth noting that different poses (orientations) put slightly different strains on each joint motor, so the inaccuracy varies within the working envelope. Typical inaccuracies of modern six axis industrial robots depend on the size of the robot (due to the larger loads it can carry and manufacturing tolerances) but with better compensation models and calibration procedures, even large robots can attain a positional accuracy of approximately ± 0.5 mm according to [99].

3.5.3 Improving Accuracy through Measurements

A key problem with robot calibration is that fundamental assumptions about the model or environment are easily violated. Non-linearities in the kinematics are notoriously difficult to model and environmental changes such as temperature variations can have significant effects. Often the inertial moments and off-set centre of gravity of a tool are not taken into account in the dynamic model of the robot either. As an alternative to extensive, time-consuming calibration procedures, many research groups have focussed on using external feedback through measurements to correct the pose and movements of industrial robots.

Using an external sensor allows direct comparison between robot feedback and actual position, thereby alleviating the need to update the kinematic model. Laser trackers appear to be the main choice due to their high accuracy, long range and fast measurement rates, as discussed in Chapter 2. Examples of laser guided industrial robots are mainly from the aerospace industry [100–102]. Some systems like Nikon's adaptive robot control [103] correct paths by performing an initial run wherein each position can be measured by an external metrology system and corrected. The positions are then saved as new encoder positions instead of coordinates. Thus, the systems utilises the repeatability of the robots to achieve an accurate path. The downside is the long time needed to perform a primary run to be corrected before any operations can be performed. The system also needs an extensive fixture system which must be accurately measured as well for each run. Most applications in the literature are restricted to point-type movements for positioning components or high-accuracy drilling. For NDT purposes, the entire path needs to be accurate to ensure correct alignment between probe and part.

Droll [104] presents a fast feedback algorithm which corrects joint positions in real time. The method is able to decrease the deviations of a robot moving at 100 mm/s to about 0.1 mm. The corrections are only in Cartesian positioning though as the laser only provides 3 DoF and hence orientation is not corrected. Summers [105] uses an optical CMM from Krypton Metrology to provide full 6 DoF for adaptive robot control

in an attempt to reach accuracies of ±0.2 mm. The focus was once again on pose accuracy and not path accuracy. Mosqueira et al. [106] used the Nikon iGPS to improve robotic positional accuracy and found that 0.38 mm absolute accuracy was possible. Once again, these were static positional measurements and not path-corrected.

It is evident from the literature that improving the dynamic and positional accuracy of an industrial robot is by no means a simple task. Laser trackers provide high feedback and precise measurements, but they require a single line of sight, do not work well in wet environments and if 6 DoF information is needed the price increases dramatically. Optical CMMs overcome some of these restrictions, but have relatively small working volumes and require line of sight from a single side. Motion capture technology may be able to overcome these challenges, but their accuracy performance is still not well understood in industrial environments.

To establish a baseline of obtainable accuracy, an early part of the EngD project investigated how accurate a standard KUKA industrial robot could move its end-effector throughout the working volume. Two different robots were tested, a KUKA KR5 HW Arc robot available at the University of Strathclyde and a KUKA KR16-L6 robot at TWI Technology Centre's facilities. Both robots were measured with laser trackers to evaluate performance. Another set of tests to establish the accuracy of a Vicon motion capture system was also conducted.

3.6 Tool Calibration Methods

As described in previous sections, the mathematical framework which makes up the kinematic model of a robot allows for the calculation of the pose of the final joint of a robot. In nearly any application, however, the final joint simply provides a platform to attach an end-effector to the robot such as a gripper or, in the case of NDT, an ultrasonic probe. After attaching a tool, the operator is no longer interested in the pose of the final joint of the robot, but the pose of the tool is instead. The centre and orientation of the tool is commonly referred to as the Tool Centre Point (TCP). The TCP is defined at a position which is sensible for the operator to visualise and

work with, such as the tip of an arc-welder or the centre of an ultrasonic probe. A representation of a TCP is shown in Figure 3.8.

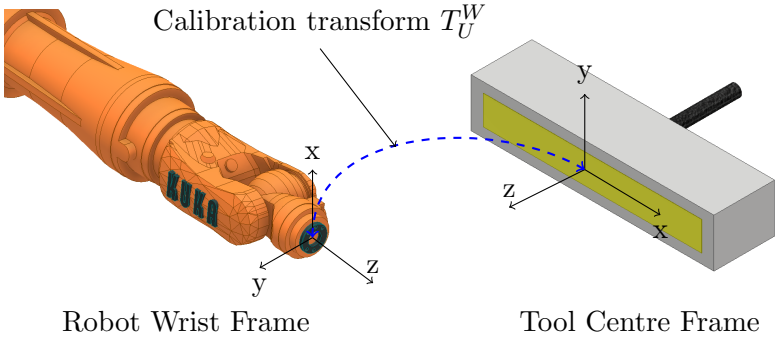

Figure 3.8: Illustration showing the relationship between robot wrist and Tool Centre Point (TCP) for a phased array probe. The unknown transform is highlighted in blue.

The kinematic model of a robot described in Equation 3.1 can easily be extended under the assumption that the tool attached to the last axis of the robot does not move with respect to the robot. This is generally the case in many industrial situations where the tool is a rigid construction. As such, the pose of the TCP can be found by multiplying the position of the flange of the robot by a transformation matrix which captures the geometrical relationship (rotation and orientation) between the reference frames of the TCP and the flange. Finding this transform is known as *tool calibration*. The problem of tool calibration has been studied extensively in the literature and a number of numerical frameworks exist [107].

Determining the reference frame calibration between an ultrasonic probe and robots for NDT has been the topic of a limited number of research groups. Kramb [47] presents a method used in the automated TESI system which uses a phased array probe mounted on a rod extending from the wrist of a robot into a water tank. The probe pitch and yaw angles are found by monitoring the ultrasonic signal reflected from a flat surface. The roll angle is found by mechanically moving the probe over a wire target and the positional accuracy is verified by scanning a series of flat bottomed holes in a reference sample.

A larger body of research has been conducted in the medical field for encoding ultrasonic data using tracking systems. The authors of [108] propose a closed-form

solution to the problem of relating ultrasonic images and an external reference frame. Their solution requires the scanning of N-shaped wire targets (phantoms) in immersion using a hand-held phased array probe. In general, cross-wire phantoms tend to be used to define points in 3D space, but a large amount of image processing needed to identify the wire itself in a standard B-scan. These methods are however tricky to use as the phantom must lie in the mid-plane of the emitted ultrasonic beam [109]. To improve on the accuracy, researchers at John Hopkins University [110] used active ultrasonic elements to calibrate the probe-to-robot transform of a phased array probe mounted on a co-bot from Universal Robotics. The active element emits an ultrasonic pulse which allows the researchers to align the probe to the same point in space by monitoring the intensity of the received signal.

The methods presented in the literature all rely on finding the tool calibration in immersion and not with water jets. Tecnatom [49] discuss ultrasonic calibration, but only in the context of setting adequate inspection parameters such as receiver gain. The authors must use a form of tool calibration to enable robot paths to be used, but it is unclear how and if the method is automated. A gap has therefore been identified in the literature regarding the calibration of tools used for automated NDT using robots. The method presented in this section establishes a procedure for tool calibration when the tool is a phased array probe embedded in a water-jet nozzle.

The method presented here is novel in a number of ways. The method presented in [47] is only possible when the probe orientation is closely aligned with the robot wrist reference frame in the first instance. Although the current method images a phantom (a sphere), it does not require the probe to be aligned with the centre of the sphere. Finally, using synthetic aperture imaging techniques improves the resolution of the ultrasonic data allowing simple image processing techniques to be used. This section discusses both manual techniques as well as presenting a new method for automated tool calibration.

3.6.1 Manual methods

If the exact geometric transform between a TCP and robot wrist is known, a homogeneous matrix containing rotation and translation can be constructed. This may be the case if the tool has a straightforward geometric shape or is designed to high tolerances such that a CAD software environment can be used to determine the TCP. In many cases, however, the TCP may not be known *a priori* and the operator will have to perform a calibration procedure. If the TCP is chosen to coincide with a physical point in space, such as a spike or the end of a welding gun, then a manual calibration can be performed. A manual calibration is performed by driving the chosen TCP to the same point in space from at least four different orientations.

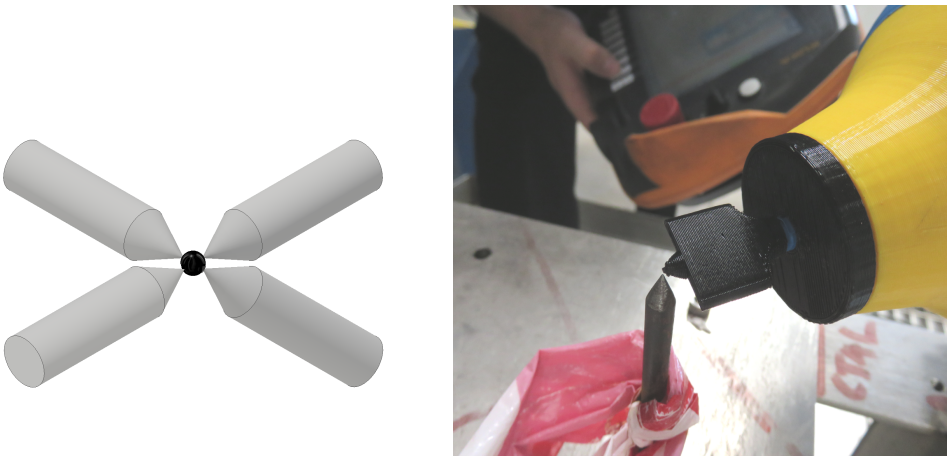

(a) (b)

Figure 3.9: (a) Manual calibration is performed by driving a robot to the same position from four different orientations. (b) Picture of a phased-array squirter nozzle on a robot being manually calibrated using a specially designed calibration artefact.

If the TCP of the attached tool reaches the same position in 3D space from different orientations, the wrist must move across the surface of a sphere centred on the chosen position. This is visually represented in Figure 3.9a. As the robot is capable of identifying the position of the wrist from the kinematic model at each pose, the following system of equations can be constructed to find the unknown $x_{tool}, y_{tool}, z_{tool}$, shown in Equation 3.5.

$$R^2 = (x_i - x_{tool})^2 + (y_i - y_{tool})^2 + (z_i - z_{tool})^2 \text{ for } i \geq 4 \tag{3.5}$$

The set of equations can be rearranged by substitution such that the following matrices can be constructed (Equation 3.6), when using four different poses. To solve for the unknown TCP, a set of linear equations can be expressed as $Ah = B$, the solution of which can be found using linear least squares: $h = A^{-1}B$.

$$2\begin{bmatrix}(x_2 - x_1) & (y_2 - y_1) & (z_2 - z_1)\\ (x_3 - x_2) & (y_3 - y_2) & (z_3 - z_2)\\ (x_4 - x_3) & (y_4 - y_3) & (z_4 - x_3)\end{bmatrix}\begin{bmatrix}x_{tool}\\ y_{tool}\\ z_{tool}\end{bmatrix} = \begin{bmatrix}{x_2}^2 - {x_1}^2 + {y_2}^2 - {y_1}^2 + {z_2}^2 - {z_1}^2\\ {x_3}^2 - {x_2}^2 + {y_3}^2 - {y_2}^2 + {z_3}^2 - {z_2}^2\\ {x_4}^2 - {x_3}^2 + {y_4}^2 - {y_3}^2 + {z_4}^2 - {z_3}^2\end{bmatrix} \tag{3.6}$$

Further points can of course be added to increase the accuracy of the calibration, but due to time constraints and the, at times, cumbersome nature of this approach, four points are usually used. A measure of how well the calibration was performed can be expressed as the variation in R, calculated using Equation 3.5.

3.6.2 Tool Calibration Procedure for Phased Array Probes

For automated NDT using water jets and phased array probes, a manual calibration procedure is slightly more involved. Unlike manual inspections where the probe sits on the surface of a component, the automated system should not physically touch the component. Furthermore, the probe sits inside a cavity as part of the water delivery system and so access to the probe itself is not possible. It is, however, important to know exactly where the face of the probe is and what orientation it is facing such that indications can be correctly encoded in 3D space as described in Section 3.4.1. Additionally, calibrating the probe as the TCP accurately is vital to ensuring that the robot will follow an intended path across the surface of a part.

The first step in calibration is to establish the nature of the TCP, as shown in Figure 3.10a. The z-axis is chosen to be normal to the face of the phased array probe and the x-axis runs along the active axis of the probe. The centre of the probe face is chosen to be the TCP origin. As mentioned previously, it is not possible to access the probe

face directly as the probe sits within a composite structure, as shown in Figure 3.10b. As an early part of the EngD project, a number of artefacts (such as the 3D printed tip shown in Figure 3.9b) were designed. This allows the robotic operator to perform a manual TCP calibration without damaging the probe. The use of artefacts assumes that the probe is located perpendicularly to the end of the nozzle enclosure, which is not always the case.

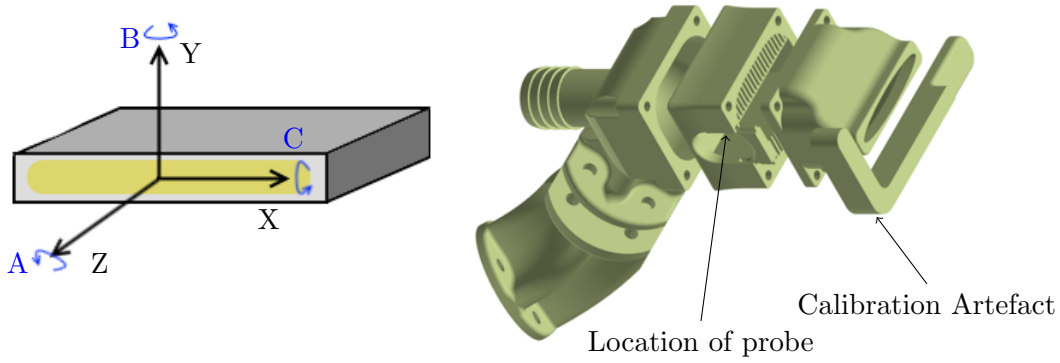

Figure 3.10: (a) Phased array probe coordinate system. (b) 3D printed nozzle and custom designed calibration artefact, illustrating the complexity of the tool being calibrated.

Besides being both time and skill-demanding, there are also a number of reasons why the manual method does not provide an accurate calibration. First of all, the probe may not sit completely aligned with the end of the nozzle where the artefact is attached and as the distance between probe and nozzle exit may be up to 50 mm, this can lead to significant misalignments. Secondly, the individual elements may not actually sit flush with the backing material that protects them. The operator assumes that this is the case due to the geometry of the probe, but it is possible they are slightly misaligned, the effect of which would again be magnified over the distance between probe and part.

To counter these problems and to develop a method which lends itself to automation, a calibration procedure for determining the robot-to-sensor relationship of a laser scanner [111] was adapted to ultrasonic calibrations. The principle of calibration is similar to the manual method in that the same point in 3D space is identified from a number of orientations and used to determine the unknown TCP transformation.

However, instead of manually driving the robot to the same point, a sphere is imaged ultrasonically and the centre of the sphere is used as the consistent 3D point. To be able to identify a 3D position in space, the surface of a sphere is imaged, as shown in Figure 3.11.

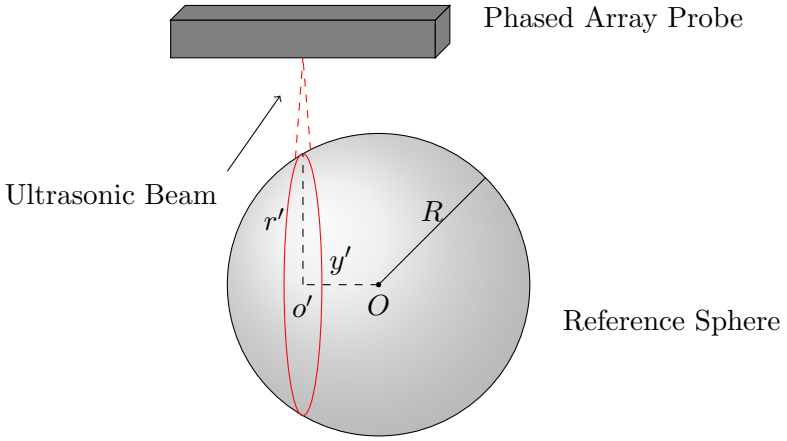

Figure 3.11: Illustration showing principle of behind the developed method for phased array probe calibration. The difference between O and o' provides the third dimension needed for calibration.

The problem of finding the robot-to-sensor calibration is commonly referred to as *hand-eye* calibration, due to large amount of work done on calibrating camera sensors (which act as an "eye" for the robot). A phased array probe is able to capture 2-dimensional data if the separation between elements and speed of sound in the coupling medium is known. This gives depth and axial data, or in terms of Figure 3.10a, z and x coordinates. The y coordinate is assumed to always be zero, ignoring the passive width of the probe. It is of interest to be able to obtain a full 6 DoF calibration in a single step, which is why the sphere method is used. By determining the radius at an arbitrary pose above the sphere, it is possible to determine the distance to the centre of the sphere, providing the full 3D position in the reference frame of the probe.

3.6.3 Automated Ultrasonic Tool Calibration

Ultrasonic industrial NDT is mainly concerned with the volumetric imaging of components to find sub-surface defects. Thus, research does not generally focus on accurately

imaging surface profiles. The calibration technique described in the previous section, however, needs to extract the circular profile of the ultrasonic beam reflecting off the surface of a sphere. It should be noted that calibration procedures do exist for single crystal probes in gantry-type systems, but as the full 6 DoF information is required for robotic NDT, these are not applicable in the current context.

Surface profiles of an object can be extracted from ultrasonic data accurately if the speed of sound in the coupling medium is known. In an immersion setup, the first echo that is recorded arises from the water-solid interface. If the sampling frequency of the receiver is known, the Time-of-Flight (ToF) can be measured and distance to the object calculated. If using a phased array probe and electronic scanning, echoes from different sub-apertures are then plotted directly below each transmission site. This approach is unsuitable for curved surfaces as illustrated in Figure 3.12a.

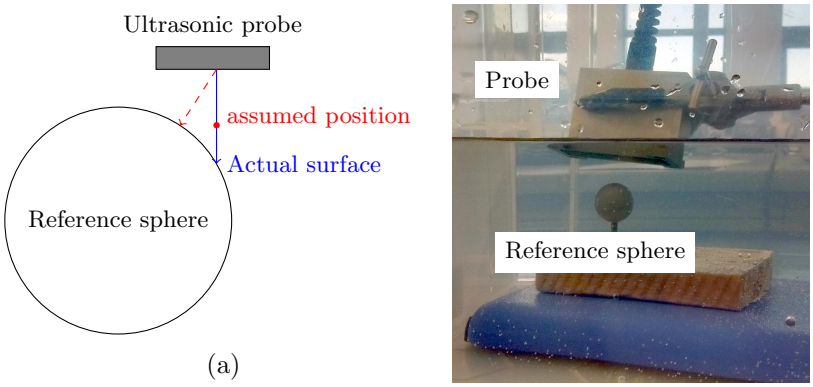

Figure 3.12: (a) Illustration showing how direct pulse-echo methods would incorrectly plot the received pulse. (b) Setup showing phased array probe held above a reference sphere.

Figure 3.12a illustrates how, as sound spreads out from the transmitting elements, the first echo will arrive from the closest point on the sphere. If the sound is plotted to originate straight from underneath the probe, the ToF will be too short, leading to an incorrect positioning of the echo's origin.

It is therefore clear that a simple encoding strategy is not suitable to image a sphere. In recent years, synthetic aperture techniques have been researched in the NDT community which allow higher resolution and focussing capabilities [112]. The authors of [113] discuss a similar setup wherein three simple shapes are imaged using

synthetic aperture techniques. The Full Matrix Capture (FMC) is used to gather a full dataset and two synthetic aperture techniques are used to generate images: the Total Focusing Method (TFM) and Synthetic Aperture Focusing Technique (SAFT). The authors argue that TFM provides superior resolution, but that SAFT shows fewer imaging artefacts. The TFM imaging method covers a larger area, however, due to the use of the full array compared to the narrow sub-aperture used in SAFT.

It was decided to use FMC with the TFM imaging technique to find a surface profile where the ultrasonic beam intersects a sphere. Each FMC dataset is gathered from a pulser-receiver and processed in the Crystal software package developed at TWI [114]. An example TFM image of a foam sphere is shown in Figure 3.13a, normalised to the maximum amplitude. As clearly seen in the image, the maximum intensity occurs where the probe face is tangential to the sphere and rapidly drops off to around -20dB afterwards. There are also imaging artefacts at the edges of the image, which was also noted in [113].

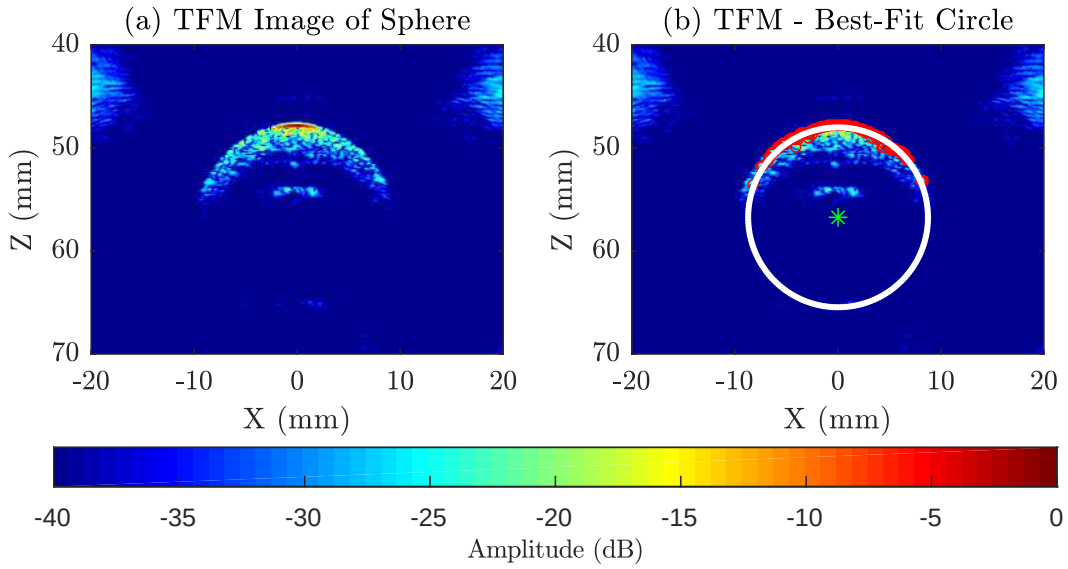

Figure 3.13: TFM images of a sphere generated from FMC data. Z denotes distance below the probe while X is along the active axis of the probe. (a) shows the original image and (b) shows the result of fitting a circle to surface points extracted from a threshold operation. Red points indicate extracted points, white points are best-fit circle and the calculated centre is shown with a green star.

To determine the 3D position of the centre of the sphere, a threshold is applied to

each column in the TFM image at -23dB, which was found to give an adequate number of surface points. Figure 3.13b shows the result of extracting the surface of a sphere and fitting a circle to the points. The red dots in Figure 3.13b correspond to identified surface points while the centre of the circle is shown as a green star. The best fit circle, shown in white, is found by least squares minimisation in MATLAB.

Referring to Figure 3.11, the centre of the sphere in 3D can be calculated after fitting a circle to the surface of the sphere below the probe. Knowing the measured (r') and nominal (R) radii, the missing distance y' is simply given by the Euclidean distance shown in Equation 3.7.

$$|y'| = \sqrt{(R^2 - r'^2)}. \tag{3.7}$$

To give the actual 3D coordinate, it must be known whether the centre lies in front or behind the probe. The 3D centre of the sphere is then given by the $(o'_x, \pm y', o'_z)$ where o' is the calculated centre of the best-fit circle. Using this method, a 3D position for the sphere can be found for a number of probe positions, allowing the TCP to be calculated using the method presented in the following section.

Calculating the TCP

Through knowing the position and orientation of the robot's wrist at a number of positions, the unknown tool transformation, expressed as the homogeneous matrix X, can be found through the relationship $AX = XB$ for robot poses A and B. Traditionally, each A_i and B_i are collected as 6 DoF information, for which a number of numerical solutions exist [107]. A phased array probe is, as stated, only able to collect 2 DoF information with a third, inferred, coordinate from the sphere setup shown in Figure 3.11. Thus, a different method is needed wherein one set of data (robot pose) is collected in 6 DoF and the other (ultrasonic data) is collected in 3 DoF.

Other work by [109] also uses a spherical imaging phantom to find the unknown transformation between an ultrasonic probe and an external positional reference system. The authors use a robot to move the sphere, but the probe is not attached as a robotic

tool and the full 6 DoF calibration of the probe is not considered. Furthermore, the method presented relies on approximating when the probe is directly over the sphere unlike the method presented here. As mentioned previously, the method used here is adapted from [111], with the addition of using the FMC/TFM method to enable it to be used for ultrasonic data.

If the sphere centre, can be located using the approaches described in Section 3.6.3, the relationship between the 3D point in the probe's frame, $P_u = (x_u, y_u, z_u)^T$, and the robot frame, $P_o = (x_o, y_o, z_o)^T$ is given by Equation 3.8.

$$P_o = T_w^o \cdot T_u^w \cdot P_u \tag{3.8}$$

Where T_w^o denotes the transform between robot origin and wrist frames while T_u^w is the unknown tool calibration (see Figure 3.7 for reference). For a number of robot poses, the following set of equations can be generated:

$$\begin{cases} P_o = T_{w1}^o \cdot T_u^w \cdot P_u^1 \\ P_o = T_{w2}^o \cdot T_u^w \cdot P_u^2 \\ \vdots \\ P_o = T_{wi}^o \cdot T_u^w \cdot P_u^i \end{cases} \tag{3.9}$$

If the physical location of the sphere (P_o) remains constant in the robot frame, then Equation set 3.9 can be rearranged to give (for each $j \neq i$):

$$T_{wj}^o \cdot T_u^w \cdot P_u^j = T_{wi}^o \cdot T_u^w \cdot P_u^i \tag{3.10}$$

The system of linear equations can be solved by expanding Equation 3.10 using the following conventions: $T_{wi}^o = \begin{bmatrix} R_{(3x1)}^i & T_{(3x1)}^i \\ 0 & 1 \end{bmatrix}$, $T_u^w = \begin{bmatrix} r_{(3x1)}^1 & r_{(3x1)}^2 & r_{(3x1)}^3 & t_{(3x1)} \\ 0 & 0 & 0 & 1 \end{bmatrix}$ and $P_u = [x_u^i, y_u^i, z_u^i, 1]^T$ to give

$$(R^i x_u^i - R^j x_u^j)r^1 + (R^i y_u^i - R^j y_u^j)r^2 + (R^i z_u^i - R^j z_u^j)r^3 + (R^i - R^j)t = T^j - T^i \tag{3.11}$$

Equation 3.11 can be rewritten in matrix format by using the different combinations of subtracting the i^{th} and j^{th} equations as shown in Equation 3.12.

$$\begin{bmatrix} R^i x_u^i - R^j x_u^j & R^i y_u^i - R^j y_u^j & R^i z_u^i - R^j z_u^j & R^i - R^j \\ R^i x_u^i R^{j+1} x_u^{j+1} & R^i y_u^i - R^{j+1} y_u^{j+1} & R^i z_u^i - R^{j+1} z_u^{j+1} & R^i - R^{j+1} \\ & & \vdots & \\ R^i x_u^i R^N x_u^N & R^i y_u^i - R^N y_u^N & R^i z_u^i - R^N z_u^N & R^i - R^N \end{bmatrix} \cdot \begin{bmatrix} r^1 \\ r^2 \\ r^3 \\ t \end{bmatrix} = \begin{bmatrix} T^j - T^i \\ T^{j+1} - T^i \\ \vdots \\ T^N - T^i \end{bmatrix} \quad (3.12)$$

If the first matrix containing the linear coefficients is not singular then the system of equations can be solved as they are now in the form $Ah = b$. To find the unknown tool transformation, h, a least-squares method is used using the MATLAB least squares solver `lscov` function. The 12 resulting values from the second vector in Equation 3.12 are then used to construct a homogeneous matrix for the unknown TCP transformation. The rotational part of this resultant matrix is not guaranteed to be orthonormal, so singular value decomposition is used to ensure orthogonality. The benefit of using this method over other available numerical methods is that no rounding errors are propagated between the orientation and translation parts as these are solved simultaneously.

Iterative Optimisation

A drawback using the proposed method is that it relies on the accurate determination of the sphere centre in 3D, but is inferred from circle fitting as described in Section 3.6.3. An iterative optimisation procedure was written in MATLAB to attempt to increase the accuracy. To optimise, an error function was chosen which expressed the spread of points around the assumed 3D point of the sphere centre. The optimisation algorithm then searches for the values of y in the reference frame of the probe which minimise the spread. It is assumed that because the measurements are taken at a number of different poses, the TCP calibration that minimises the spread will be closer to the true sphere centre.

For each measurement pose, the sphere centre's position was transformed into the robot's reference frame using the initial TCP calibration provided by the steps outlined

in Section 3.6.3. The MATLAB `fminimax` function from the Optimisation Toolbox finds a point which minimises the maximum of a set of objective functions. These objective functions were set to minimise the distances between projected sphere centres by varying the values of the actual, measured points. Constraints were imposed such that only y values allowed within the known radius of the sphere were used.

In each iteration, the algorithm tests new values of y for the sphere centre and recalculates the TCP using these. The error is then compared to the previous iterations to see if the chosen values improved the TCP estimate. If the error function falls below a threshold of 1 mm or reaches a maximum of 20 iterations, the optimisation stops and the new best estimate of the TCP calibration is supplied to the user.

3.6.4 Experimental Results

A reference sphere was scanned from 8 different orientations using a 64 element, 5 MHz linear phased array probe, sitting inside the 3D printed nozzle shown in Figure 3.14. The nozzle was mounted as a tool on a KUKA KR16-L6 robot at TWI's facilities. The probe had a pitch of 0.6 mm and data was acquired using a PeakNDT LPTA pulser-receiver unit. The reference sphere is constructed from brass and was kindly lent by Dr. William Kerr at the Advanced Forming Research Centre in Glasgow. The sphere has a nominal diameter of 25 mm and measurements using callipers showed that the diameter is accurate to 0.05 mm. A plane is cut on the sphere allowing it to be drilled and mounted on a brass rod. Mounting the sphere on a rod allows the robot to move around it without colliding with the mounting table, as shown in Figure 3.14.

Figure 3.14: Image showing the robot positioning the phased array probe nozzle over the reference sphere.

To ensure that the FMC and TFM algorithm imaged the sphere accurately, the robot was manually driven across the top of the sphere and the position of the maximum signal response and minimum distance to the surface was recorded. The assumption was made that this point would correspond to the probe being aligned with the centre of the sphere and thus the method presented in Section 3.6.3 should report the radius as being half of the nominal diameter: 12.5 mm. As the wavelength of sound in water at 5 MHz is 0.29 mm, the reported radius could at best be expected to fall within half of this value.

Figure 3.15 shows both a C-Scan image of the top of the sphere as well as the intensity of the reflected signal at different positions across the sphere. As seen in Figure 3.15a, the maximum intensity is observed around a central region, which is assumed to be the top of the sphere. Figure 3.15b shows a large increase in intensity at x: 23.4 mm and a normally distributed intensity profile across the sphere at each location. This is anticipated due to the curvature of the sphere.

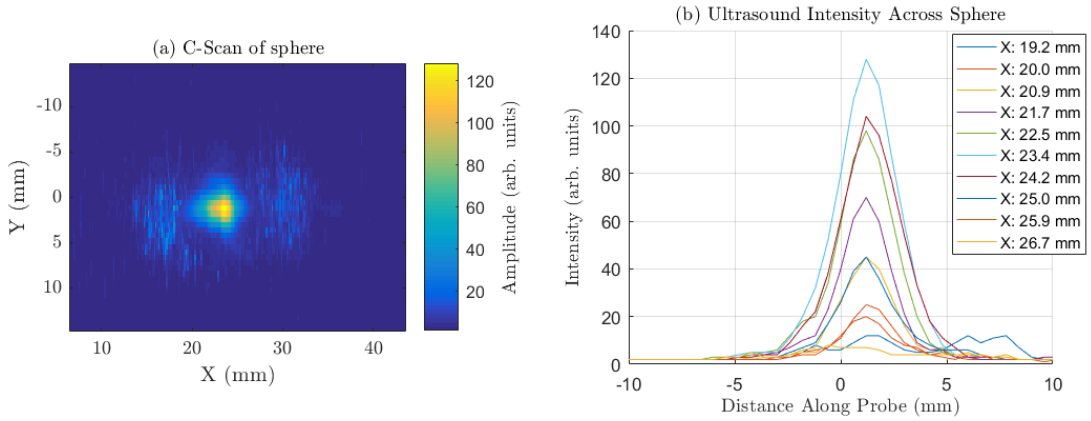

Figure 3.15: Variation in ultrasonic intensity across a sphere. (a) shows a C-Scam image of the top of the sphere as the robot moves a probe along the x axis. (b) The variation in amplitude at different points along the x-axis, showing a Gaussian curve.

An FMC dataset was acquired at the position with the highest intensity and a TFM image was formed using the minimum pixel resolution available in the Crystal software, 0.0625 mm. The obtained TFM image is shown in Figure 3.16 with the fitted circle and calculated radius. In the image, the best fit circle is shown as white dots while the extracted surface points are shown as red marks. The overall arc extracted from the sphere is around 10 mm wide, which covers approximately 25% of the hemisphere.

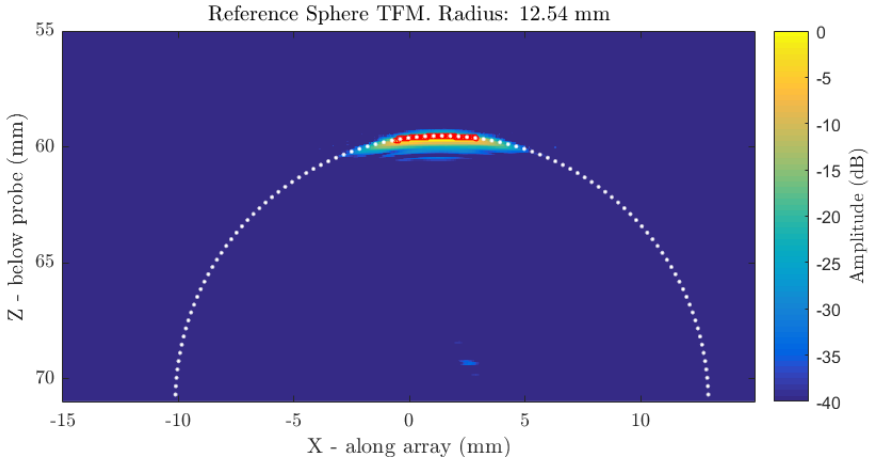

Figure 3.16: TFM image of brass reference sphere with a nominal diameter of 25 mm. The threshold was set to -23 dB for best results. The red points are extracted surface points while the while indicate the best fit circle.

It was found that the most accurate results were obtained when using a threshold

of -23 dB to extract the surface of the brass sphere. Using this threshold, the radius was calculated to be 12.54 mm, indicating the method is accurate. To get an estimate of the tolerances of this method, a range of threshold values from -25 dB to -18 dB was used. The average value of the calculated radius was 12.34 mm with a spread of 0.80 mm (2σ). As seen in Figure 3.15b, the sound intensity varies in a Gaussian manner as the probe is moved along the sphere, indicating that a standard deviation provides a true representation of the error. Thus, it is likely that the proposed method using FMC and TFM will be accurate to within a millimetre, which is acceptable for most NDT applications.

Eight robot positions were manually chosen at different points on the sphere. For each pose, a FMC dataset was gathered and the coordinates of the robot's wrist (denoted as `TOOL 0` in the KUKA controller) were obtained by running a simple RSI server and connecting via MATLAB to the robot through UDP. As mentioned, the accuracy of a calibration can be verified by using the estimated TCP calibration to project the static 3D position of the sphere into the robot's reference frame. Prior to optimisation, the average spread was 2.93 mm from the mean of the positions. After running the optimisation algorithm described in the previous section, the average spread was reduced to 0.98 mm.

The results of optimisation indicate that the overall calibration should be true to within 1 mm of the actual centre of the probe. To investigate the difference between automated and manual methods, the phased array nozzle was calibrated manually using the artefact shown in Figure 3.10b. The manual calibration can only calibrate the centre of the nozzle exit (which may not be aligned with the probe face) due to the geometry of the probe holder. To find the position of the TCP, the artefact centre point is projected in its negative z direction by 50 mm, corresponding to the distance between probe face and nozzle as measured on the CAD model.

As shown in Table 3.1, the manual and automated method results differ slightly, especially in the tool's x and y directions. This is most likely due to probe not being aligned perfectly with the end of the nozzle which violates the assumption made when using the manual method. The manual method is both time-consuming and prone to

operator error while the calibration method presented in this section is easily automated. It could, for example, be possible to permanently install a reference sphere somewhere in the robot's working envelope to allow for daily calibration checks. To truly test which calibration mechanism provides a more accurate result a reference sample with indications in known locations must be ultrasonically imaged using each of the calibrations above.

Table 3.1: Table showing differences in values for the TCP when using manual and automated calibration methods.

Manual Method				Automated Method			
X:	65.264 mm	A:	103.071°	X:	71.767 mm	A:	101.291°
Y:	-34.113 mm	B:	8.893°	Y:	-41.116 mm	B:	8.208°
Z:	181.094 mm	C:	30.379°	Z:	179.883 mm	C:	27.494°

Sensitivity Analysis

The method for calibrating the TCP of a phased array probe presented in this section relies on accurately inferring the centre of a sphere by analysing a single arc measurement. To better understand the robustness of the method, a sensitivity analysis was performed. For each of the eight robot positions, a 6 dB range of threshold values was used to provide an estimate in the centre position of extracted circle (as shown in Figure 3.13) of each FMC image. It was found that the x-coordinate varies with approximately 0.79 mm while the variation in the y-coordinate is 1.31 mm. It is assumed that the larger variation in y is due to a smaller arc being observed. Fitting a circle to a larger arc will give a more accurate estimate of the centre and radius, but as shown in Figure 3.12, it can be difficult to observe a large arc using ultrasound.

To understand how sensitive the technique is, a Monte Carlo simulation was used. A number of arcs, spanning 4 to 150 degrees from the top of a sphere were generated using a circle with a nominal radius of 12.5 mm. Random noise with a standard deviation equal to the variation seen in the estimate of circle centres were added to the arc points. A circle was fitted to this noisy data and the steps were repeated 1000 times. The calculated radius plotted against the actual radius is shown in Figure 3.17.

As seen in the figure, if more than 50 degrees of arc is observed, the method converges to the correct solution, even in the presence of noise. Analysis of the 8 robot positions revealed that on average, about 20% of a full arc (roughly 40 degrees) was seen in the FMC images. This is shown as a red line in Figure 3.17.

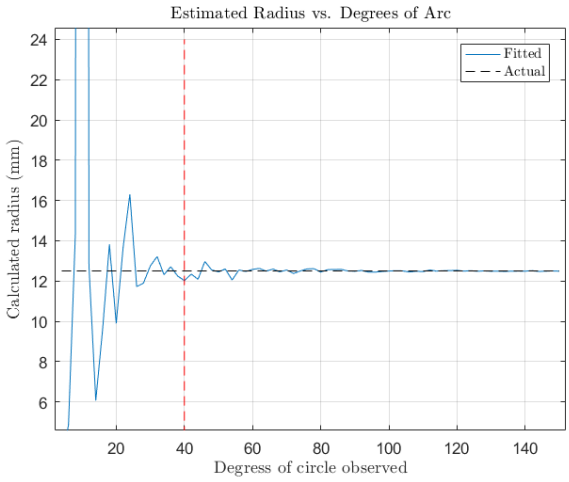

Figure 3.17: Calculated radius versus actual radius for the size of arc observed when imaging a sphere. Noise was added to each arc based on variations seen during ultrasonic imaging. As seen, the estimate converges when approximately 50 degrees can be seen. The red line indicates the average arc length seen in the experiments.

Another source of uncertainty is the error propagation in the values shown in Equation 3.12. The solution for the unknown TCP parameters is found using MATLAB's least squares solver, `lscov`. This function also provides an estimate of the error as well as the standard deviation of each term. As a reminder, Equation 3.12 sets up a system of equations in the form $Ah = b$. The model then becomes $b = Ah + \epsilon$, so that the error is described by $\epsilon^2 = (b - Ah)^2$. The mean square error, MSE is then described by Equation 3.13 for the best estimate of the vector h.

$$MSE = \mathbb{E}||b - Ah||^2 \tag{3.13}$$

Using this definition of the MSE, the error in the least squares solution is 0.318, though this does does not represent a physical distance as Equation 3.12 solves for rotation and translation simultaneously. Looking at the standard deviation of each

parameter in the vector h, the expected error in the x,y and z coordinates of the TCP is around 1.5 mm. The expected variation in angles is 2 degrees around z, 1.8 degrees around y and 0.3 degrees around x. It was expected that the variation around z and y would be largest as the y-coordinate cannot be measured (as it is the passive axis of the probe). The uncertainty in y is identified as being the largest source of error.

To better understand how much the uncertainty in y affects the results, another Monte Carlo simulation was run using an idealised situation wherein the TCP transform is already known. Using the 8 known robot poses, the extracted circle centres were calculated using the idea TCP transform. Gaussian noise was then added to the y-coordinate at each pose and the TCP solution was calculated again. This was repeated 3000 times using noise with a 0.5 mm standard deviation. Figure 3.18 shows the results of these simulations.

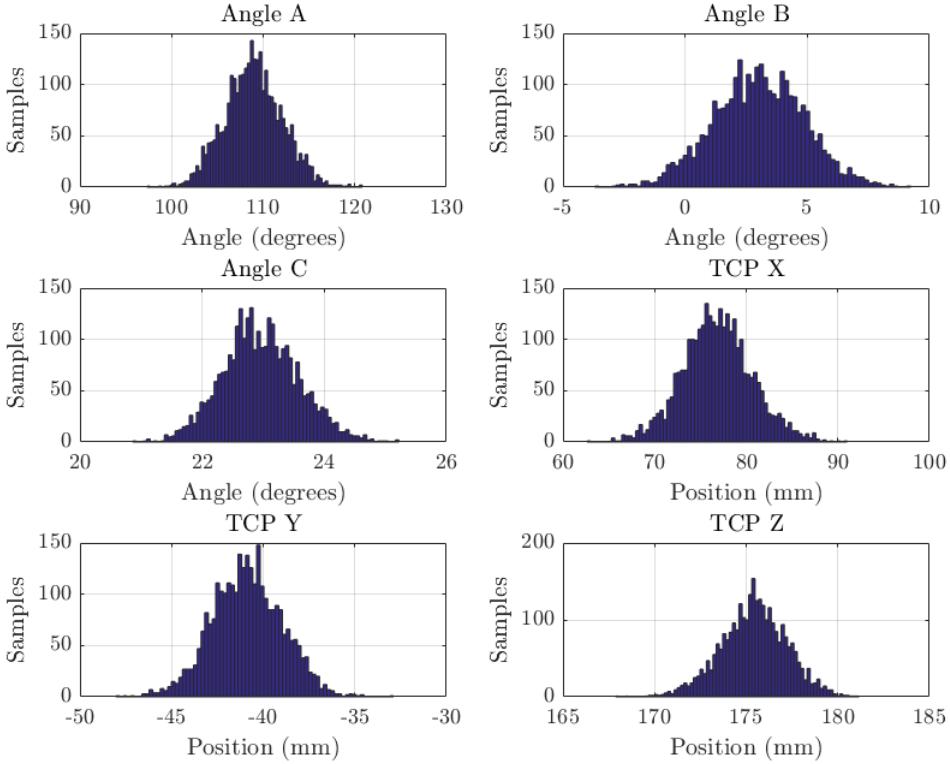

Figure 3.18: Monte Carlo simulations showing how noise with 0.5 mm standard deviation added to the y coordinate of the sphere centre influences the final TCP calibration.

As seen in Figure 3.18, even a relatively low level of noise (0.5 mm standard deviation) has a large impact on the final TCP parameters. The standard deviations in x, y and z are 3.8 mm, 1.9 mm and 1.8 mm, respectively while the deviations in angles A, B, and C are 3.1, 1.9 and 0.6 degrees. As seen before, the deviation in angles A and B is larger, which could mean that the robot poses used are not sufficiently far apart to solve for all three angles accurately. Interestingly, the variation in these 6 parameters scales almost linearly with the standard deviation of the noise, such that the variation in the TCP translation is roughly 4 times the noise level in the y-coordinate.

3.6.5 Suggestions for Future Work

Accurate tool calibration mechanisms are crucial for any robotic application, but especially so in automated NDT systems. A camera inspection system can, for example, still provide good results even if the camera is not completely normal to the surface of the part. An ultrasonic probe, on the other hand, needs to be perpendicular to the sample at all times when using longitudinal waves to minimise refraction at the surface. The current system being discussed uses water jets as a coupling medium for the sound, which also provides a measure of compliance between the probe and part to avoid damage. The distance that a quasi-laminar flow can cover after leaving the nozzle is limited to around 10-20 mm so an accurate tool calibration is also needed to avoid loss of coupling or collisions.

The method presented in this section provides a novel method for tool calibration, but can still be improved further. The sensitivity analysis identified that a large arc covering more than 1/3 of the full sphere should ideally be imaged to accurately fit a circle to the surface points. Furthermore, even small variations in the estimated centre of the sphere can lead to large deviations in the calculated TCP. Future work should investigate ways of choosing optimum robot poses to minimise variation in the calculated TCP values.

To ensure that the centre of the sphere is estimated correctly, another scanning strategy could be implemented, keeping the probe orientation fixed and scanning across the top of the sphere. This would reduce the reliance on estimating the sphere centre

from a single point. Another idea could be to use a smaller sphere to ensure more of the surface could be imaged at each point. As mentioned previously, the standard techniques from the medical industry use wire-based phantoms. It may be possible to obtain better results using synthetic aperture imaging methods with these types of phantoms compared to standard imaging techniques.

3.7 Ultrasonic Base Alignment

Moving on from the challenge of accurately calibrating a phased-array probe as a robotic tool for NDT purposes, this section addresses how the calibration between robot and component can be found. As a reminder, this calibration is labelled T_B^O in Figure 3.7 or *base calibration* in KUKA robot terminology. Being able to accurately locate a component's position and orientation are of course crucial to any robotic application to avoid collisions and ensure a path is followed accurately. During traditional robot programming this was less of an issue as paths were designed in a static environment. Now that trajectory planning is done off-line, the resultant NDT paths are defined relative to a datum on a CAD model which then subsequently needs to be identified during the actual setup.

Identifying a component and its exact pose (position and orientation) with respect to a robot is a significant challenge to obtaining a more flexible system in line with the Industry 4.0 ethos. The step is traditionally performed by manually driving the robot to the part, a process which is time-consuming and prone to operator error. The work presented in this section shows that it is possible to perform sub-mm alignment between the component and robot coordinate frames using ultrasonic ToF information.

A visual representation of the developed method for alignment is shown in Figure 3.19. During the path planning process, several key points are selected which cover 10-30% of the part surface. To avoid collisions while moving between these key points, the robot follows the surface between each target location. The ultrasonic pulser-receiver is set up such that the front-wall echo is recorded at each location. After the scan finishes, the ToF information, encoded with robot positional data is used to create a

point cloud representing the surface of the component (P_s). The key points identified during the path planning stage form another (sparser) point cloud, representing the CAD models expected position (P_{CAD}).

It should be noted that the development of the method presented here pre-dates the similar methodology presented by the authors of [115] and was developed for squirter nozzle inspections and not immersion testing. Unlike the method in [115], the method shown here uses key points which are determined more robustly, reducing the number of calculations needed. At each P_{CAD} key point, a search algorithm finds a set of closest points in P_s and fits a plane to the them. The normal vector, average time-of-flight distance and amplitude are then used to create a matched key point. A list of matching key points is created and the deviation between P_s and P_{CAD} can be calculated. This information is then transferred back to the robot controller to update the coordinate reference system for that inspection.

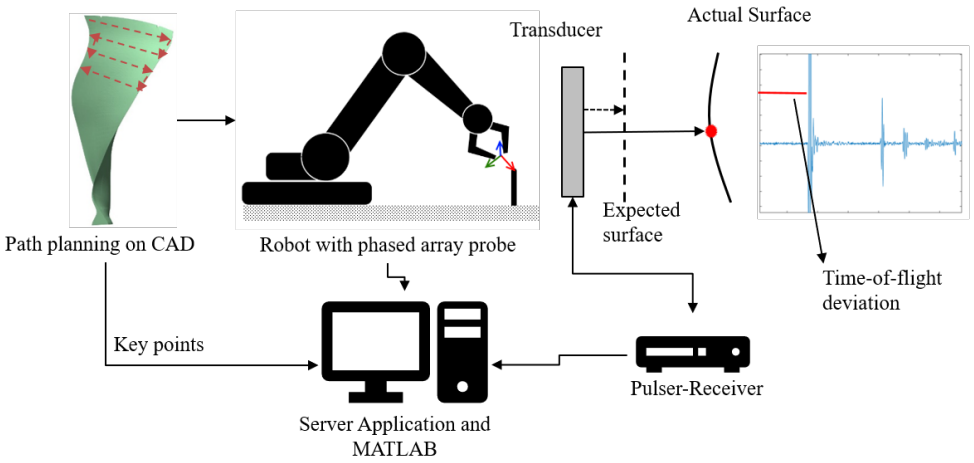

Figure 3.19: Illustration of the ultrasonic alignment technique. The difference in ToF between expected and actual surface helps to calibrate the part-to-robot base more accurately.

3.7.1 Matching Algorithm

Point cloud registration is a topic which has been around for many years [116] and the approach is generally to search for corresponding points in different models and find a rigid transformation which estimates pose deviation. One of the key challenges

is to define which points correspond to one another. The current method avoids this problem by only considering points in a defined region of space around each key point. To ensure the data is adequate, two checks are performed on the data before a key point is generated from data in P_s. First, the average amplitude must be above a set threshold and secondly, the normal vector must be within a set tolerance.

For each key point in P_{CAD}, a specified volume is defined in 3D space and points, $P_i = (P_{ix}, P_{iy}, P_{iz})^T$ from P_s which fall inside this volume in are used to calculate a single, corresponding, key point. The average amplitude of $P_1 \dots P_N$, (where N is the number of points inside the volume) should be above a set threshold, dependent on the reference amplitude of the scan. If the average amplitude falls below this value, the key point is discarded from the list of key points in P_{CAD}.

The normal vector is calculated by finding the normal vector of a plane fitted to the data using a least-squares optimization. This is done by first removing the centroid of the subset of P_s and computing the Singular Value Decomposition (SVD) of the resulting coordinates through $Y = U\Sigma V^T$ where Σ is a 3x3 matrix wherein the diagonals correspond to the eigenvalues. The normal vector of the plane is then defined as the singular vector from V corresponding to the zero value eigenvalue. The normal vector for each plane is compared to the expected normal vector of the key point by taking the dot product of the two vectors. If the direction is too dissimilar, the key point is again discarded from the list in P_{CAD}.

If a set of points in P_{CAD} meet the above conditions, then a corresponding single point is created. The 3D coordinates of this matched key point are the average of the set of points in the chosen volume. In this way, the entire point cloud generated from the ToF data does not need to be searched which improves the performance of the algorithm. Once all key points have been matched or discarded, the rigid transformation between the expected and observed points cloud is found through the algorithm described in [117]. First the centroids of each point set (K_{CAD} and K_S) is calculated and removed from their respective point sets, as shown in Equation 3.14.

$$K_{CAD} = P_{CAD} - \frac{1}{N}\sum_{i=1}^{N} P_{CAD}^{i} \quad K_s = P_s - \frac{1}{N}\sum_{i=1}^{N} P_s^i \tag{3.14}$$

The covariance matrix of the two sets is then given by Equation 3.15.

$$X_{(3x3)} = \sum_{i=1}^{N} K_{CAD}^{i} \cdot K_{s}^{i} \tag{3.15}$$

SVD is then performed on X to obtain the factorization $U\Sigma V^T$. The rotation matrix aligning the orientation of the point sets is then given by $R = UV^T$. The 3x3 rotation matrix can then be represented as three Euler angles as expected by the KUKA robot control software (see Appendix A for details). The translation vector (displacements in x, y and z) aligning the two sets is given by $T = K_S - R \cdot K_{CAD}$. The matching algorithm, implemented in MATLAB 2016b, is represented as a flow chart shown in Figure 3.20.

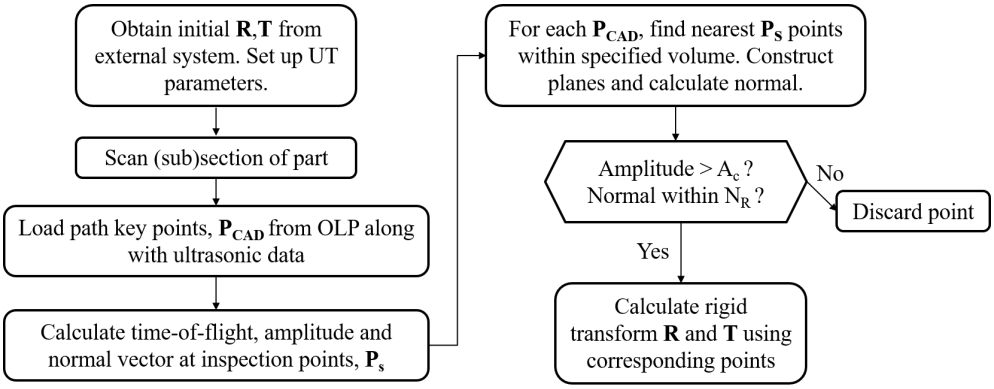

Figure 3.20: Flowchart illustrating the different matching steps. A_c and N_R are thresholds set at the beginning of the process.

The described method allows for the accurate alignment of robot trajectory and part based on a few assumptions. First, it is anticipated that the part is within a few millimetres of the expected position such that the robot can move without collision and that the front-wall ultrasonic echo can be observed at all points. Secondly, the alignment scan may not be able to scan the entire surface, so if the part is invariant under translation and rotation the result for this degree of freedom may have to be ignored. Key features such as corners or fixturing holes would be ideal to use, but these can be difficult to accurately identify using a water-jet coupled ultrasonic probe. Finally, the curvature of the part must be sufficiently low so that a plane can be

approximated at each P_{CAD} inspection point. In practice this means that P_{CAD} key points cannot include areas close to edges and corners.

3.7.2 Experimental results

An experiment was carried out in a controlled setting to test the accuracy of the developed approach. A flat plastic sample, measuring 150 mm x 100 mm x 10 mm, machined to a tolerance of 0.1 mm, was fixed in place in the robot's working envelope. The sample was ultrasonically scanned in a raster pattern with a 5 MHz linear phased array probe, with a pitch of 0.6 mm using an electronically scanned sub-aperture of 15 elements to match the elevation of the probe. Ultrasonic data was gathered using a Peak NDT Micropulse 5PA unit and encoded with the robot's positional information using the IntACom acquisition software. The probe TCP was calibrated using the method presented in Section 3.6.

The resulting scan is shown in Figure 3.21a where it can clearly be seen that there is little variation in the front wall echo. The robot's reference frame was then changed such that a misalignment of 1.1 mm in x, -0.7 mm in y, 1 mm in z, 0.2° about $\hat{z}$, -0.8° about $\hat{y}$ and 0.5° about $\hat{x}$ was applied. These values were randomly chosen in a range that represent typical misalignments seen in practice and are also shown in Table 3.2. The sample was then scanned again, as shown in Figure 3.21b which clearly shows the variation in time-of-flight across the surface.

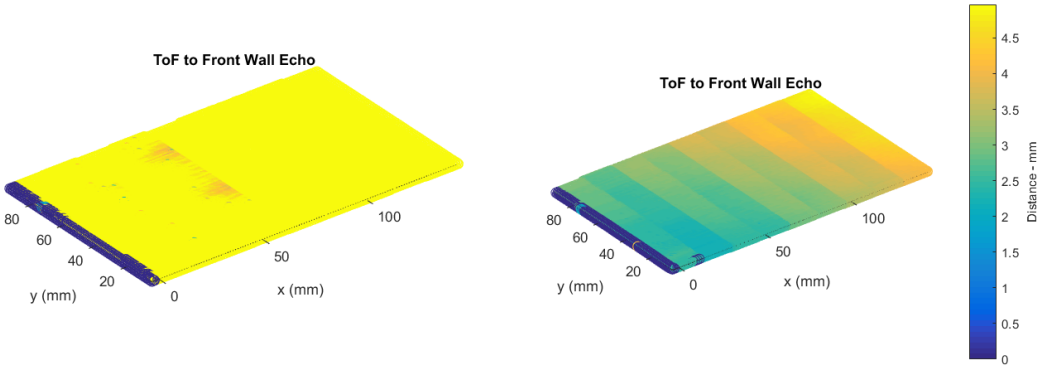

(a) (b)

Figure 3.21: Reference scan of a flat sample. When the sample is well aligned (a), there is almost no variation in the ToF. When purposely misaligned (b), the ToF varies across the sample. The goal of the method is to calculate this misalignment.

After this second scan had been collected, the ultrasonic alignment algorithm was applied to the data to test how well the original calibration could be recovered. The calculated corrections are shown in Table 3.2. As evidenced in Table 3.2, the sample is invariant under transformation in x, y and around $\hat{z}$ hence the corrected values for these dimensions are not usable. However, the results of the alignment algorithm for the remaining three dimensions are within 0.03 mm and 0.02°.

Table 3.2: Applied and recovered misalignment data for a flat sample. Angles A, B and C correspond to rotations about $\hat{z}$, $\hat{y}$ and $\hat{x}$.

Dimension	**Applied Offset**	**Recovered Offset**	**Error**
X	1.10 mm	0.25 mm	0.85 mm
Y	-0.70 mm	0.13 mm	-0.83 mm
Z	1.00 mm	0.97 mm	0.03 mm
A	0.2°	0.01°	0.19°
B	-0.8°	-0.82°	0.02°
C	0.5°	0.49°	0.01°

As the results from scanning a flat plate show, the alignment algorithm does not work on components that exhibit translational symmetries. Future research will investigate ways to improve ultrasonically imaging reference points such as corners, edges or holes of a component to provide stronger matching features. Using edges is, however, tricky as the turbulence of the water jet causes a large amount of noise. This could

potentially help provide a better absolute alignment for all six degrees of freedom. Furthermore, it has been noted that components manufactured from Carbon Fibre Reinforced Polymers (CFRP) have a tendency to "spring-back" from their mould after curing. This means that a simple rigid-body transformation is no longer adequate as the trajectory designed on a CAD model will never match the geometry perfectly. Future work will look at methods of using non-rigid point cloud alignment methods to not only update the robot base reference data but the path itself as well.

3.8 Conclusions

Understanding and improving the accuracy of automated NDT system using industrial robots is the main focus of this thesis. The current chapter has presented the fundamentals of industrial robot mechanics, programming, kinematic models and how ultrasonic data can be encoded using these types of machines. Various aspects related to the positional accuracy of industrial robots have also been discussed. In particular, the three types of calibrations needed to ensure the overall accuracy of the system is high were studied in detail.

The merits of calibrating robots to provide a more accurate kinematic model versus using external measurement systems were discussed. Although calibration can significantly lower pose deviations, the assumptions that are made in the kinematic model are often violated. Further complications occur when attempting to improve the dynamic path accuracy, which is a key parameter in automated ultrasonic NDT. It was therefore decided that a feedback control mechanism in terms of an optical measurement system would be a better choice to compensate for errors.

Improving the positional accuracy of the robot in isolation is however not sufficient to achieve high quality inspection results. This chapter also discussed two other calibrations, namely the tool and base calibrations needed to carry out NDT paths. A novel method for calibrating the geometric relationship between a phased array and robot coordinate reference system was presented. Experimental data illustrated the accuracy and applicability of using advanced ultrasonic imaging methods.

A new method for accurately calibrating the position and orientation of a work-piece with respect to an industrial robot was also presented. This work is only possible as a consequence of being able to accurately calibrate the ultrasonic probe as measurements are taken by considering the time-of-flight of the ultrasonic pulses. Experimental data was presented to validate the technique and demonstrated that sub-millimetre accuracy was possible on a test sample.

Chapter 4

Increasing Automation Levels through Optical Measurements

4.1 Introduction

The focus of this chapter is on how automation levels in robotic NDT systems can be improved through the use of optical measurements. In particular, the challenge of automatically determining the position and orientation of a specimen located in a robot's working envelope is addressed through a custom developed machine vision framework centred on the use of CAD models. A discussion of suitable methods for this task was previously presented in Chapter 2. This chapter expands upon this topic and reviews a number of implementations through experimental trials. Since standard, off-the shelf robots do not typically contain any sensors for determining the environment around them, external sensors are needed. Optical systems are particularly well suited for sensing the environment as they are capable of accurate, non-contact measurements often covering a large field of view. Giving a robot "eyes" is often the aim of a machine vision application and stems from a desire to mimic the visual capabilities of humans. Machine vision is a field of computer engineering which has experienced incredible growth in recent years, spurred by the availability of cheap cameras, increased computing power and large, open-source online image and software repositories. The

omnipresence of cameras in personal devices such as smart phones is also a major contributor to this growth.

Chapter 3 presented a method for performing precise alignment between a robot and part coordinate reference frames using the ultrasonic echo from a surface reflection. While the method proved to be successful, it still relies on knowing the approximate position of a component to avoid collision with the robot and obtain an adequate ultrasonic response. Furthermore, the method presented only works if the component's shape closely resembles the CAD model used for path planning. The objective of this chapter is to introduce large-volume measurement methods to determine the pose of a component to a precision and accuracy suitable for ultrasonic alignment.

4.1.1 Automation Challenges for Robotic NDE Systems

This chapter identifies three main challenges for realising a flexible and adaptable robotic inspection system. These are: (1) automated path planning, (2) automated part recognition and (3) adapting to changes in the environment. As discussed in previous chapters, industrial robots in their current configuration are typically programmed once and repeat the same procedure for extended periods of time. The operations are usually individual point motions such as moving an object, spot welding or painting. In automated NDT the robot must accurately follow the surface of a component to within a millimetre with velocities determined by the material properties, geometry and type of inspection method. These requirements make manual programming impractical and hence the first challenge facing robotic NDT systems lies in automated path planning for this specific purpose, as discussed in detail in Mineo's thesis [20].

If automated off-line path planning can be achieved, the path will consist of a series of command poses in a suitably chosen coordinate system. Efforts are made to choose a sensible origin and unit axes relative to a datum on a part. An industrial robot does not, however, have *a priori* knowledge of where in 3D space a component is located or if the component has an associated inspection path. Thus, the second challenge consists of automatically recognising parts and accurately determining their pose. This problem is not unique to robotic NDT operations and is faced by automation engineers

in a large number of industries and situations such as recognising parts on conveyor belts or picking objects from bins [18].

A robotic system which can automatically identify and locate large parts for inspection can provide a flexible and easily reconfigurable solution for the manufacturing environment of the future. To be fully adaptable, however, the system must also be able to identify changes in the inspection process. In the field of composite manufacturing these changes are typically in the form of deviations between an actual part and its nominal CAD model. The third identified automation challenge is being able to automatically capture new surface data to update a previously defined inspection path.

The rest of this chapter will be devoted to addressing the second automation challenge as this thesis builds upon previous work on automated path planning for robotic NDT systems [20]. The third automation challenge is tackled separately in Chapter 6. To address the difficulty of part localisation, a number of measurement methods for locating and identifying parts relevant to aerospace composite inspection were considered. A standard approach when using industrial robots is to construct a rig or fixture which can present the part to the robot in the same pose for each process iteration [18]. This strategy is inflexible and typically only economically sensible for high-volume production processes. Due to variances in composite production, even well-designed fixtures may need to be manually adjusted to account for part shape variation. Another standard, but time-consuming approach, is to manually register a part using for example a bar code label and a manual measurement method. A manual method is naturally undesirable if the objective is to increase automation levels.

The approach suggested in this chapter is to allow the part to be approximately correctly positioned in front of the robot and use measurements to adapt the off-line paths to fit the real situation. Some type of fixture will undoubtedly be needed to present the part in a suitable orientation, but these can be manufactured to lower tolerances, for example using plastic polymer additive manufacturing techniques. An operator can present the part to the robot in an suitable location within the robot's reach, alleviating the need to spend time taking accurate measurements. A remote

sensing system will then measure the actual pose of the part and update the robot's base reference values such that the desired scan can be carried out. Figure 4.1 illustrates the considered metrology solutions for remote sensing.

Each method in Figure 4.1 has its own advantages and disadvantages. Laser trackers provide high accuracy but require specialised targets and do not work well with water, which is used to provide ultrasonic coupling. Motion capture systems use several cameras and photogrammetry to determine the 3D positions of markers. Although markers need to be physically mounted on parts to be tracked, these can work in wet environments and systems are able to track hundreds of markers simultaneously, providing a large volume tracking solution. Machine vision provides a flexible measurement solution, but can require significant software development and calibrated cameras. Using visual information without the need to attach markers to neither fixtures nor parts is appealing as it reduces setup times and the need for manual labour. Machine vision offers such a solution by being able to extract meaningful data about a scene from pixel intensity values.

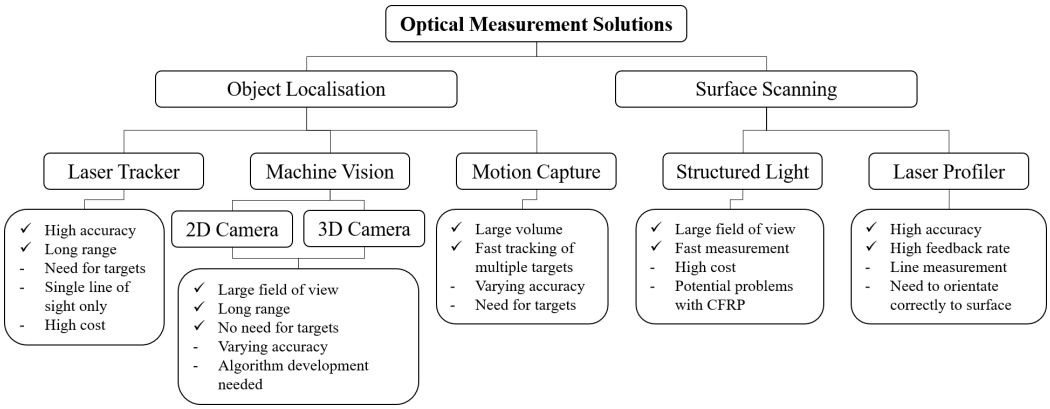

Figure 4.1: Overview of various remote sensing methods for object detection and surface scanning.

Figure 4.1 also shows methods which are suitable for meeting the third challenge of how to deal with variations in surface shape. Obtaining surface geometry accurately was previously discussed in Chapter 2. However, the use of traditional metrology methods, like CMMs, is not feasible in the current context as the part would need to be scanned in a separate facility. Furthermore, CMM scanning would be far too slow

to meet inspection time targets. Instead, a local metrology method which is accurate enough to provide updates for NDT paths and compact enough to be mounted on a robot is needed. Remotely measuring the surface of a component is also desirable to avoid any potential damage to the part. Structured light projection and laser scanning both offer this type of solution.

Both a laser scanner and projected light system utilise information about how light is reflected off a surface to infer information about the shape of the surface. Both methods were discussed in detail in Section 2.5.4. It was decided that a laser scanner was most suitable for surface mapping after discussions with manufacturers (Stemmer Imaging employees, private communication, 22 March, 2016) who indicated that laser line scanners typically perform better on CFRP materials. Additionally, the cost of a laser profiler is up to ten times lower and the compact size was more suitable to attach to the available robots.

4.1.2 Contributions to Knowledge

It is outside the scope of the current EngD project to arrive at a completely novel method for pose estimation and shape recognition. Instead, the author's contribution to knowledge lies in the application and combination of several computer vision methods in an industrial setting. The key differences between the presented work and previously published papers are the distance between the camera and object in an industrial setting, the size of objects to be recognised and the accuracy required for executing robotic scan paths. Furthermore, the consequences of incorrectly matching components is more severe than previously presented applications which focus purely on image recognition in datasets where there are no severe consequences for equipment if an image is incorrectly labelled.

To tackle this issue, this thesis advocates the use of a flexible, multi-scale measurement approach to metrology which is applicable to automated inspections. This ranges from measurements taken at a distance using a calibrated camera, to fine adjustments based on the reflected ultrasonic signal. These measurements are not accurate enough for geometric tolerance measurements, but sufficient for NDT inspections, where the

identification and approximate location of defects is more important. This approach minimises the need to make large adjustments to the existing robotic cell such as ensuring uniform lighting and the proposed solution can be implemented with inexpensive and readily available cameras. Combining measurements in this way is known as data fusion whereby more useful and accurate information can be provided than that of a single sensor.

This chapter presents a method which centres on the use of CAD models for object recognition and pose estimation using both 2D and 3D image data. The proposed method has been shown to work as a proof of concept that can reliably classify the correct component based on a shape outline, with a high degree of accuracy. Unfortunately, the demands from industry were focussed on improving calibration procedures and less on trialling an automated part recognition system. Furthermore, although CAD models of parts were available, the actual parts were no longer available on site which also hindered real-world experimental trials. Thus most of the results presented in this chapter are based on simulated data and future work should investigate how the developed framework can be integrated into a fully automated solution.

4.2 Automated Part Registration

The following sections investigate the use of machine vision techniques for automated part recognition and pose measurement. An overview of machine vision techniques and research for the purpose of part registration and localisation is presented. Methods for simultaneous recognition and pose measurements are discussed. The use of CAD files for this purpose is also considered as an alternative to direct image comparison. The use of motion capture systems for part localisation and recognition is also discussed.

4.2.1 Introduction to Machine Vision

Machine vision is the field of computer science and engineering which seeks to extract meaningful data from digital images. The aim of machine vision is to convert, through analysis and processing, pixel intensity values into information that can be used to make

decisions in an automated manner. The field began with research at universities in the 1960s and it was initially thought that the task of automatically describing scenes from images could be solved as a student summer project [118]. Since then, the field has expanded massively and now covers a number of domains such as scene reconstruction, video tracking, image restoration, object recognition, pose estimation, image learning and image content retrieval, to mention a few.

In an industrial context the main application areas for machine vision are position measurement, inspection, sorting, counting and object detection [119] as illustrated in Figure 4.2. Position measurements and sorting are the most common applications wherein a camera placed over a conveyor belt measures the position of objects to be picked up by a robot. Visual inspection is used extensively within manufacturing to check labels, circuit boards or for the presence of unexpected objects in the food industry. To facilitate the implementation of these applications, the environment is often designed to create uniform lighting conditions, remove background clutter and highlight the object in front of the camera. Depending on the complexity of the task, other sensors may be integrated alongside a vision system to ensure reliable detection.

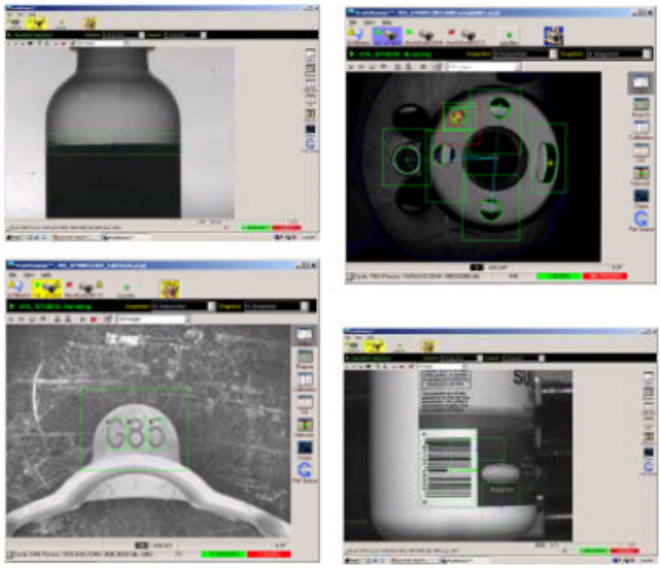

Figure 4.2: Examples of a machine vision application for inspecting manufactured goods [120].

A digitised image consists of individual blocks of captured light, called pixels. The

basic principle of a machine vision application is a bottom-up approach where information from each pixel are grouped together in stages to reach a point where decisions can be made about an image. Pixels are considered to be relevant or irrelevant based on global or local thresholds. Relevant pixels can then be grouped together to form structures representing edges, lines or even feature points based on nearby changes in intensity. Once the data has been collected into a suitable format it can then be compared to a model, depending on the nature of the application. Finally, inference about the image can be made based on statistics or by training machine learning methods such as neural networks using previous data. To provide a robust and reliable solution the model must either be well defined or a large amount of training data must be available if machine learning methods are required.

Computer vision and machine learning has grown substantially in the last decade, driven by the availability of processing power and the desire to install cameras in everyday objects such as phones, watches and cars. Humans (and other mammals) rely on visual information for a tremendous amount of seemingly trivial tasks such as recognition, estimating distances, and guiding body movements [121]. Hence, it is no wonder that we also try to tap into this massive resource from a machine learning perspective in an effort to simulate or imitate the human cognitive abilities. In an effort to simulate the brain, Artificial Neural Networks (ANN) have become popular in the last decade and have shown an impressive accuracy in finding patterns in large datasets, including recognising humans, animals and objects in images [122, 123]

A number of vision-based inventions are prominently featured in the media at the time of writing, including driver-less cars and facial recognition through social media. The speed at which innovation moves forward in these areas can easily be correlated with the potential earnings that are at stake. For critical applications, the use of several sensors can increase reliability and safety. Self-driving cars, for example, use at least three other sensors besides cameras to avoid collisions and navigate, relying on many years of accurate mapping with GPS [124]. A combination of LIDAR and visual cues give enhanced adaptive capabilities through machine learning algorithms which are enabled by an increase in available computer processing power. Even so, training

a neural network with the amount of data still requires tens of thousands of driving hours with operators behind the wheel.

Similarly, the Amazon GO supermarket has also unveiled a new concept of a till-less supermarkets where items chosen by shoppers are automatically registered by a central system. The patent filing [125] shows pressure sensors, infrared sensors, a scale, load cells, a volume displacement sensor, a light curtain, microphones, RFID and haptic input sensors are all used in conjunction to achieve this feat. Once again, the availability of cheap processing power and the use of machine learning techniques to handle the enormous amounts of data has led to these automation improvements. In comparison, automated part placement in a robotic cell encompasses a very small automation task and consequently using the same amount of sensors and data is economically and logistically infeasible. This scope of work instead focusses on the challenges faced by robotic NDT systems and how visual information can be used to solve these.

4.2.2 Object Recognition and Pose Estimation

Although the field of machine vision encompasses a large number of domains and applications, the work presented here only focuses on the task of recognising objects and determining their 3D pose. As machine vision is used in a large number of applications with different requirements and constraints, there are a multitude of approaches for object recognition. Each machine vision application will be unique based on local requirements and constraints, but some high-level criteria can still be used to classify common approaches to object recognition as discussed in [68, 119, 121]. The main criteria can be summarised as follows:

- *Representation*: To be able to locate and recognise an object, the chosen machine vision approach must have a way to represent the object. The two main methods are either shape or appearance based representations. A shape representation captures information about the geometry of the object, typically through the outline or surfaces of the object. This method requires a geometric model of the object which can be obtained through training images or by using a shape model such as CAD models. On the other hand, the appearance representation

relies on the image characteristics of the object, as a human would perceive it. This representation usually requires extensive image data of the object to build a model and relies on the object being presented to the camera in a consistent manner. There are advantages and disadvantages to both approaches. A shape representation cannot tell the difference between a lemon and a lime, for example, while an appearance based method would struggle to identify blank, texture-less objects.

- *Data Model*: After choosing how to represent the object (shape or appearance), the next choice is how to construct a data model. The data captured about the object can either be local (such as corners or key features) or global (overall area, perimeter length, average texture). Local data can be gathered from a number of sites in the image resulting in a higher dimensional feature matrix whereas a global data model will typically have a single feature vector for each object. Global features are often more suitable for simple structures or in applications where each object is visually unique. Local features are often better suited for complex objects as features can be searched individually and still result in recognition. The data used to construct features can either be low-level such as raw pixel intensities or high-level such as a number of identified edge pixels being grouped together to form a line segment.

- *Expected Variations*: Once the choice of representation and data model have been made, the algorithm must be designed with the expected object and environmental variations in mind. If a number of similar objects have to be recognised, the algorithm must be capable of dealing with intra-class variation (such a different types of circuit boards for example). The image quality variation must also be taken into account if, for example, the amount of light in the scene changes throughout the day.

- *Matching*: To recognise an object, the algorithm must at some point compare the observed data to a model. This step is known as matching and is a rich and active research field within computer science. Common strategies include

matching observed features to a database of trained features or match a geometric model to an image through projection. In nearly all instances, the matching step seeks to optimise a similarity parameter which in turn is defined by the choice of data model. Matching using high-level features is usually more reliable as these contain more information, but they can be difficult to detect even when scene variations are taken into account.

4.2.3 Related Work

Early work on object recognition focused on using geometric models [126] and the main objective of many research groups was establishing correspondences between 3D models and their 2D projections in images. Unfortunately, these methods were limited to mainly planar objects and had little success with images in real world environments. In response, research has since been targeted towards appearance based descriptors (for either points or small regions) [127–130]. Another driver for these types of descriptors is the increased amount of real-world images where geometric models of objects are not available.

These techniques have had tremendous success in a large number of fields, particularly in instance recognition [131–133] but also for pose recognition [134–136]. In recent years, research groups are once again focussing on the use of 3D shapes for recognition, partly due to the pose estimation limitations of point feature detectors but also due to the abundance on CAD models available on the internet [135]. Another reason for the re-emergence of the shape based recognition approach is the introduction of low-cost 3D cameras such as the Kinect camera from Microsoft [137].

One of the largest challenges for automated image recognition is that an object will appear slightly differently depending on the viewing angle, giving rise to a continuum of possible 2D views. A recognition approach based on point or region descriptors will often struggle with unfamiliar perspectives as the data model is built on a limited number of viewpoints. To partially overcome this limitation, modern descriptors are normalised such that they are insensitive to changes in scale, illumination and position. Combined with a scheme for pose similarity comparison, features can be found on

objects even if they are presented from a new viewpoint. Point and region detectors tend to work best on textured objects and their reliability depends on the type and amount of training data available.

Besides needing a large dataset to obtain reliable recognition, point based descriptors also struggle with accurate pose estimation. For many purposes, this may not be relevant. This is because the task at hand perhaps only requires a true/false indication of the presence of an object. If more accurate pose estimation is needed, a larger framework must be built which captures pose information at the training stage, by taking pictures of an object from all angles which can be time consuming [138]. In general, pose estimation and recognition are treated as two sides of the same coin, as an object must first be recognised before its pose can be estimated.

Even when no prior image data of objects is available, CAD models with texture information can be employed to train appearance based matching algorithms. The authors of [132] exemplify this by training a Histogram of Gradients (HOG) based classifier using computer rendered CAD data to detect furniture from IKEA. Though the authors report better results than standard appearance based methods, their approach still suffers from a number of false positives when objects appear in cluttered environments. Pose estimation accuracy is not addressed specifically by the authors, which is a general trend within the literature as recognition is often sufficient. The pose in [132] is defined as an average precision parameter compared to a bounding box threshold. Few papers provide a reference pose measured by an external system and often cite either an error measurement or use a relative pose compared to a checkerboard pattern as a baseline reference.

Tulsiani and Malik [139] describe an algorithm which uses a Convolutional Neural Network (CNN) to provide a rough pose estimate for an object described by a number of keypoints. The estimated viewpoint is then used as input to predict local features in a feedback process. This is inspired by the observation that humans will often perceive the global appearance of an object before focusing on the finer details. A pose estimate is deemed correct if the predicted keypoints are within a set distance of the observed keypoints. The authors demonstrate the performance of their method on

the challenging PASCAL-VOC 2012 detection set where bounding boxes are also given along with viewpoint information. The training of neural networks for both the initial pose estimate and keypoint determination requires a considerable prior data set and long computational time.

Rosenhahn's thesis on the pose estimation problem [140] approaches the situation from a different perspective. The overall aim of his thesis is to develop a visual landmark based navigation system for robots but the main contribution lies in exploring the geometric modelling of the pose of an object. Instead of projecting an object to the image plane and comparing image data directly, he modelled the image as a 3D entity such that the comparison could be done in 3D. This gave the advantage of providing a true Euclidean distance error measure for pose refinement as well as increasing the speed of the estimation process. A method similar to the Iterative Closest Point (ICP) algorithm is then used to compare the two 3D shapes for object recognition and pose estimation.

In the field of automated NDT, few papers have been published which address the need for automated part recognition and pose estimation. Stubbs *et al.* demonstrate how a camera mounted on a robot can be used to determine the location of a part on a turntable prior to ultrasonic scanning [46]. The camera uses commercial software to determine alignment to within 0.075 mm as well as making sure the right part is present. Lighting is mentioned as being very important and a ring light capable of illuminating a 500 mm diameter circle is used. The alignment takes about 50 s and is limited to the two-dimensional case of a small object sitting on a turntable. Tecnatom have also presented a method called GeoID for part recognition and alignment [52]. This method uses a laser scanner to perform surface scans of a part which are then compared to CAD models. The method requires a robot to move the laser scanner which can be time consuming.

4.2.4 Requirements and Limitations

As the previous sections have illustrated, there is not a single "one-size-fits-all" solution for object recognition using machine vision methods. The particular requirements and

limitations of any application can vary widely, which has resulted in a large number of approaches and solutions both in industry and the academic literature [119]. Thus, it is pertinent to consider what particular requirements and constraints are present in the context of automated NDE systems using industrial robots.

As ultrasonic scan trajectories are programmed in an off-line environment, there is an expectation that the path for a particular component already exists when the component is at a specific pose with respect to the robot. To raise automation levels, a vision system should be able to confirm the presence of an object and calculate the pose of the component to within a tolerance where a calibration scan can be executed using the water jet nozzles. As the stand-off tolerances for the water jets currently used are 5-10 mm, this effectively determines the accuracy requirement of the vision system. If objects can be recognised to this degree of accuracy, an ultrasonic alignment procedure (described in Section 3.7) can be executed to provide final adjustments for the actual scan path. Conversely, if the component pose is too far from the expected location, the pre-defined path may no longer be valid due to the robot's reach or poses which result in singularities.

There are several limitations to be considered in the current context. Firstly, parts are not typically available at the robot cell before they are due to be inspected. Hence, it would require significant organisational efforts to establish an image database of each component prior to inspection. Consequently, matching strategies which rely on training recognition models using previously seen images are not possible. The CAD models, however, are supplied in advance as they are used to design the scan trajectory at the path planning stage. A camera can either be mounted on a robot or at a point overlooking the working envelope with the latter being a safer choice and alleviating the need for changing robotic tools. Placing a camera at a point overlooking the cell also provides a larger measurement volume but reduces the obtainable accuracy.

Further assumptions can be made about the scene the camera will observe. The object can be expected to be found on the jigging table, establishing boundaries for any search algorithm. The object can be rotated, but most likely only about its vertical axis if it is to remain upright. Background and foreground clutter can be removed from

the scene, making it easier to use segmentation algorithms. The TWI IntACom robot cell sits in a large open space, hence illumination is not guaranteed to be constant and it would be difficult to minimise the variation in light levels. Invariance with regards to illumination is therefore desirable when considering which data model to use. An example image of the IntACom cell from a camera viewpoint which captures the jigging table and robots is shown in Figure 4.3.

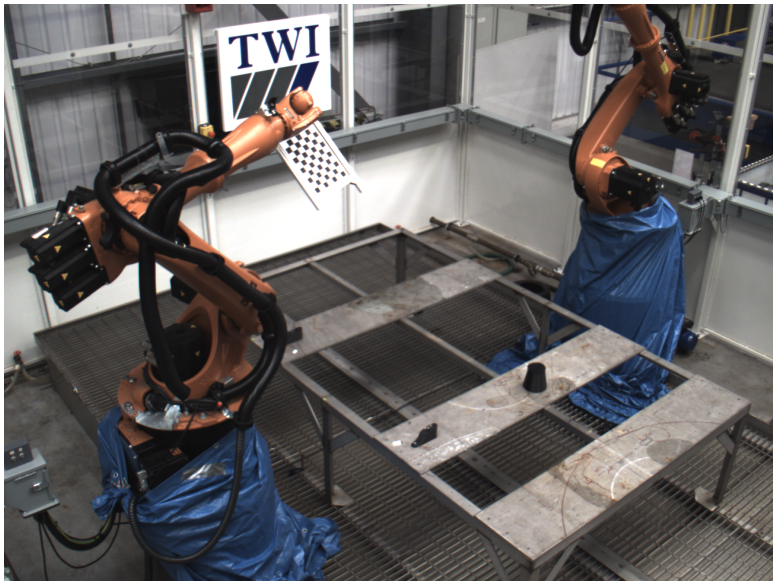

Figure 4.3: Image showing robot cell layout from a fixed camera vantage point. The left-hand side robot is moving a checkerboard calibration pattern throughout the volume to find the transform between camera and robot reference systems.

4.2.5 Camera Hardware

Another main consideration that needs to be taken into account when designing a machine vision application is the camera hardware, especially the choice of camera resolution and lens focal length. To determine these parameters the physical placement of the camera needs to be considered. The camera must be able to have a field of view (FoV) that can encompass the entire scene where objects are present. This FoV must be covered by enough pixels to provide adequate spatial resolution. The smallest feature to be detected needs to be established - in the current scenario, this requirement is determined by the stand-off tolerance of the water jet nozzles. The distance between

camera and object must also be estimated. If the camera is placed on the robot cell enclosure, the distance is roughly five metres.

The required sensor resolution (r_{Camera}) is given by the ratio between the field of view and spatial resolution ($r_{Spatial}$) multiplied by 2 to ensure that at least 2 pixels cover the smallest feature, as shown in Equation 4.2.

$$r_{Camera} = 2\left(\frac{FoV}{r_{Spatial}}\right) \tag{4.1}$$

Once the required camera resolution is known, the sensor size needs to be known to establish the focal length. The sensor size is manufacturer dependent, but ranges from 0.4 to 1 inch for most industrial machine vision cameras. The required focal length is given by the ratio of the physical sensor size (S) multiplied by the distance between camera and object (D) to the field of view, as shown in Equation 4.2. Lenses that match the desired focal length can then be acquired if they match the sensor size of the camera.

$$f = \frac{S * D}{FoV} \tag{4.2}$$

In the current application, the FoV is approximately 1500 mm x 2000 mm and the smallest detectable feature should be around 2 mm. This gives a camera resolution of about 2000 px in the horizontal axis. Commercially available sensor sizes are around 6-8 mm wide for this resolution. Assuming a 6mm sensor width and a distance of 5m to the object, a focal length of around 15mm is required. Further hardware requirements include ensuring that data transfer rates from the camera are suitable, especially in high-speed applications. To meet these requirements, a machine vision camera from FLIR (formerly Point Grey) was acquired. The camera has a 5 MP colour sensor (2592 x 1944 pixels) measuring 0.4 inches (5.76 mm by 4.29 mm) along its diagonal. A Fujinon machine vision lens with a focal length of 16 mm was also acquired along with a 5 m Ethernet cable.

4.3 Image Processing for Recognition and Pose Estimation

This section outlines some of the main image processing steps used in machine vision applications for object recognition and pose estimation. As mentioned previously, two standard approaches within this field are appearance or shape based descriptors. Methods for both types of descriptors are covered in this section. Additionally, this section also covers methods for image segmentation whereby an image is spilt into regions corresponding to either object or background pixels. This is a particularly important image processing step needed for accurate shape representation as the outline of an object needs to be extracted. The Bayesian classification framework is then explained in the context of object recognition based on the discussed representations. Finally, the use of several cameras to generate 3D point cloud data based on visual clues is discussed as well.

4.3.1 Feature Extraction

As mentioned in Section 4.2.2, a method of representation is needed to describe an object to a machine. There are various approaches to representation [121], but broadly these can be condensed into two categories: appearance or shape based descriptors. Each descriptor is usually represented as a numerical vector where each value codifies a feature of an object, such as an edge or texture patch. There exists a large number of feature extraction methods in the literature which again reflects the complexity of tasks handled by computer vision. No single definition of a feature exists, but in essence a feature has to be able to detect something "unusual" in an image. This section describes some of the main descriptors which have been identified as being potentially useful for recognition and pose estimation.

Local Features

Descriptors can either relate to a single pixel or a region of pixels, depending on the feature being described. These features can be categorised as either local or global.

Local features need to be sufficiently unique to be identifiable across images but local enough to not be affected by clutter in images. Global features describe an entire object and often require the image to be segmented to be extracted. As mentioned in Section 4.2.3, local feature descriptors have become increasingly popular in the last 15 years as the desire to process natural images grew. Standard local features include edges and corners [119, 121]. Edges are identified as sets of points which have a high magnitude of intensity gradient while corners are points where edges change direction quickly. Common examples of these descriptors include the Canny edge detector [141], the Harris corner detector [142], the Laplacian of Gaussian and Difference of Gaussian detectors.

Around the turn of the century, the limitations of simple features was clear and a new approach of using local region descriptors was pioneered by researchers like David Lowe [131]. The idea was to use a two step approach of detecting key points such as corners and then construct a region descriptor based on neighbouring pixels around each key point. The purpose of this approach is to ensure enough information about a key point is captured such that the descriptor gets invariant to changes like viewpoint or illumination. Furthermore, points can also be ranked based on how "strong" each region is, typically based on local variations in pixel intensity.

These key points can then be used for object recognition by comparing detected features to a database containing models. The model is usually established by collecting features from a number of images during a training phase. Recognition is then performed by both comparing the numerical content of each descriptor but also by checking the positional relations between key points. Examples of local region based descriptors include [69, 119, 127]: SIFT (Scale-Invariant Feature Transform), SURF (Speeded Up Robust Features), GLOH (Gradient Location and Orientation Histogram), HOG (Histogram of Gradients) and MSER (Maximally Stable Extremal Regions). Table 4.1 also lists these feature types. A main advantage of using these types of descriptors is that no assumptions are made about object appearance. This can however also be a drawback as key points still need to be detected which requires variations in pixel intensity.

Table 4.1: Table showing common examples of features used for object recognition in machine vision.

Feature Types		
Point Descriptors	Local Region Features	Global Features
Canny Edge	SIFT	Correlation
Harris Corner	SURF	Image Moments
Difference of Gaussian	GLOH	Fourier Descriptor
Laplacian of Gaussian	HOG	Shape Descriptors
	MSER	

Global Features

Contrary to focusing on a smaller region around a key point, one can instead focus on describing the set of pixels that define an object in an image. This is the aim of a global feature description approach. Each model now represents the entire object by describing characteristics such as area, perimeter, major axes, etc. The simplest form of global feature recognition is to define a template of the object to be found and looking for direct correlations between this object and a test image. This method is however very sensitive to changes in illumination or scale, but can work well if objects are likely to always appear the same way to the camera, i.e. conveyor belt applications.

Instead of comparing pixel values directly, features relating to the statistics of the region defining an object can be extracted instead. The area can be calculated by counting all pixels inside the region and edges can be found to determine the contour of an object. An image moment is a simple but powerful global feature calculated from weighing the pixels belonging to an object in different ways. Moments can be a useful way of describing a region's size and shape as well its location and orientation. The region moment of order $p + q$ (labelled m_{pq}) is given by Equation 4.3 where u, v are pixel coordinates.

$$m_{pq} = \sum_{u,v} u^p v^q \tag{4.3}$$

Some moments have a physical interpretation, akin to classical mechanics. The moment m_{00}, for example, describes the area of the region and can be used to scale

other moments. These normalised moments are thereby made invariant to changes in scale. By relating moments to the centre of gravity of an object, they can also be made tolerant to changes in position. A class of invariant moments, known as the Hu moments, [143] are particularly useful for object recognition.

If the contour of a region can be extracted, then this information can also be used to classify an object. By considering the distance from each contour pixel to the centre of the object, a distance function can be built. This can be considered to form a unique signature for each shape and in order to compress this information further, the Fourier transform of the distance function can be taken. The result is a compact feature vector which can be used for classifying object based on their contour signature [144]. Using the contour or shape of an object is useful if the geometry of the component is known beforehand. Using geometry for object recognition is covered in more detail in Section 4.3.3.

A limitation of global feature descriptors is that the object must be separated from the rest of the image. This is a crucial image processing step known as segmentation and is covered in the following section. Global features are preferred in the current context of localising objects for robotic NDT as the CAD model of the object is available prior to inspection and the environment can be structured to avoid occlusions which can otherwise limit the usefulness of global feature approaches. Table 4.1 provides an overview of all descriptor types covered in this section.

4.3.2 Segmentation

Image segmentation is one of the most important steps in extracting meaningful information from images. The goal of segmentation is to separate an image into areas which show strong correlation. These areas can correlate with respect to a number of properties such as brightness, colour or texture. For object recognition purposes the aim is typically total segmentation such that an image is divided into regions of pixels corresponding to either “object” or “not object”.

Segmentation is widely used for machine vision tasks such as counting products on an assembly line, identifying blood cells or reading printed characters [121]. The

task is usually made easier by ensuring that images are taken when objects are placed against a uniform background. In other situations, such as surveillance, it may not be possible to achieve such complete segmentation and instead, partial segmentation is used to provide an input to higher level processes which utilise prior information about the scene to distinguish objects from background.

Under controlled conditions, global segmentation can be easily achieved by setting a simple threshold on a grey-scale image to separate the image into foreground and background. This is particularly suitable for applications such as text character recognition. The optimum threshold can in most cases be readily achieved using a histogram of pixel intensities to split the image into bright and dark areas. Segmentation can be difficult in natural scenes, mainly because the real world often violates the assumptions of a model, such as having a non-static background [145]. Model violations can occur in a number of ways: presence of noise, occlusions, and variations in lighting or texture [146]. Exploiting prior information about the scene or using motion can aid segmentation in these situations.

To identify the background, a number of images are usually taken of a scene with no objects present. Lighting is varied between each image to represent typical changes that would occur in a working environment. Once the data is gathered, the average intensity for each pixel in each colour channel is calculated as well as its variance. A pixel threshold map for each pixel intensity is then set as a predefined number of standard deviations from the mean. In other words, if a new image pixel intensity is found to be outside the bounds established by the mean and variance it is classified as belonging to an object, if the intensity is within the bounds it is instead labelled as a background pixel. This method can be useful if the only expected variations are changes in lighting, but a new background set will have to be constructed if the scene changes.

Segmentation is a very active research field within computer vision and a large number of more advanced methods exist. K-means clustering tries to cluster all pixels into K groups based on similarities in pixel intensities (or colour). There is no guarantee that the algorithm segments objects from a background, it will only define similar

regions. Region-growing methods rely on assuming that pixels within distinct regions have similar values. An image is seeded with input points and from each point a region is iteratively grown by comparing neighbouring pixels. The performance of this algorithm is heavily dependent on the choice of seed locations and is susceptible to noise. More advanced methods use intensity gradients in images as analogous to forces which propagate a curve towards edges defining regions. A well-established framework for this type of approach is the level set method [119, 147].

The Level Set Method

The level set method was developed in the 1980's and has extensively been used as an analysis tool for surfaces and more recently as an image segmentation approach. Instead of performing segmentation based on pixel values alone, the level set considers the zero crossing of a characteristic function of an image, based on the gradients of intensity of pixels. This crossing then defines an edge which can be used to segment the image into two distinct regions. The method has certain advantages over similar frameworks in that a curve does not need to be explicitly specified, instead only the evolution of an embedding function is considered [148].

A contour, C, is represented as the intersection of a 3D surface ($\phi(x)$) with the zero level on the image domain, ($\Omega \subset \mathbb{R}^2$) as shown in Equation 4.4.

$$C = \{x \in \Omega | \phi(x) = 0\} \tag{4.4}$$

We want to know how this curve evolves over time, so we will let the surface change with time, t. The curve is then presented by $\phi(x(t), t) = 0$. In order for the curve to evolve, some sort of "force" must act on it, and we assume the direction of the curve's movement is normal to the curve itself, i.e. $\frac{\nabla\phi}{||\nabla\phi||}$. If we consider the evolution of the surface with respect to time, the curve evolution will be given by Equation 4.5.

$$\frac{\partial\phi(x(t), t)}{\partial t} = 0 \tag{4.5}$$

Given a force, F, this partial differential expression can be expressed as shown in

Equation 4.6.

$$\frac{d\phi}{dt} = \phi + \Delta t F ||\nabla \phi|| \tag{4.6}$$

This relatively simple formulation provides a powerful framework for evolving a contour [148]. All that is needed is a initial value for the surface ϕ and a method for describing the force, F. In image segmentation, the force at each pixel is usually derived from the gradient of pixel intensities at that point. Using gradients in the image creates a vector field which naturally drives the curve towards strong edges in an image. A popular implementation is the geodesic active contour formulation developed by Caselles et al. [149]. The formulation is presented in Equation 4.7. In this equation, the first term drives the speed of curve evolution based on pixel intensities while the second term $\nabla g(I)$ attracts the curve to strong edges in the image.

$$\frac{d\phi}{dt} = g(I)||\nabla \phi|| div \left(\frac{\nabla \phi}{||\nabla \phi||} \right) + \nabla g(I) \cdot \nabla \phi \tag{4.7}$$

Figure 4.4 shows an illustration of how the level set formulation can be used for image segmentation. The initial contour is chosen so it partly covers all three shapes. As the surface evolves based on the "forces" in the image (generated by darker pixels), the contour evolves and splits until all three shapes are segmented from the background. The 3D visualisation shows how the surface evolves over time.

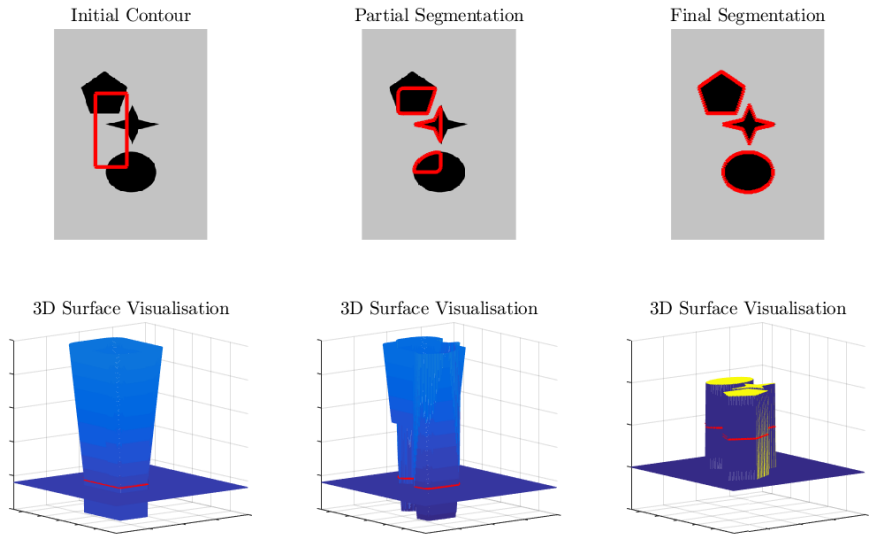

Figure 4.4: Illustration of the the level set method and how a curve evolves. The 3D surface intersecting the image plane generates the contour outline.

The level set method has a number of benefits compared to other segmentation approaches. No assumptions are made about the background and there is no need to find an optimal threshold level for pixel intensities. The approach can expand a contour based on an initial guess, ϕ_0, which then evolves until the edges of an object (or several objects) are found. This is particularly useful in the context of a robotic inspection cell where only a single object will be present and an initial estimate of where the object needs to be can be obtained from off-line path planning.

The limitations of the level set method is that the initial guess has a large influence on the time it takes for the method to converge and that strong edges need to be present in the image for the curve to evolve as expected. The method converges quickly if the initial guess is close to the intended shape. This can be achieved by resizing the image (using image pyramids, for example) to achieve a coarse segmentation. This segmentation can then be used as an initial guess for a finer resolution stage, until the original image is recovered [150].

4.3.3 Using Geometry

As seen in the previous section, the level set method provides a powerful framework for image segmentation but needs an initial guess to be efficient. As CAD data is available, previous knowledge about the shape of an object exists. However, the real world (and CAD models) are three dimensional, while the projection in an image is two dimensional. To be able to integrate 3D shape information into an image, a model of the object has to be projected onto the image plane. To be able to achieve this, the pose of the object in the scene must be known. In other words, we have to find the translation, rotation and scale of the component and then extract the outline of the model. This concept is illustrated in Figure 4.5 where a CAD model is rotated, **R** and translated, **t**, to create a desired outline on the image plane.

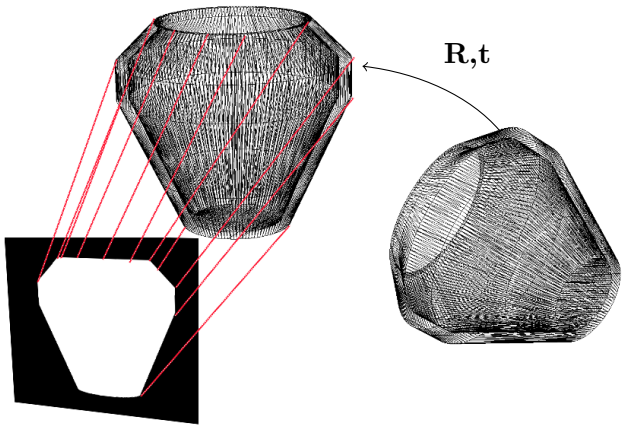

Figure 4.5: Illustration showing how knowledge of the pose (rotation, **R** and translation, **t**) of a CAD model can be used to generate a 2D image.

As mentioned previously, the reason for taking this approach is that components are typically not present prior to inspection which means that a library of images cannot be built up to achieve feature extraction. Instead, the CAD model can be used to provide recognition by comparing extracted contours against a library of projected 3D shapes. To build such a library, images can be generated at different locations on the viewing sphere, as shown in Figure 4.6. Additionally to providing recognition through the methods described in Section 4.3.1, matching can also provide a pose estimate, using pose data captured when building the library.

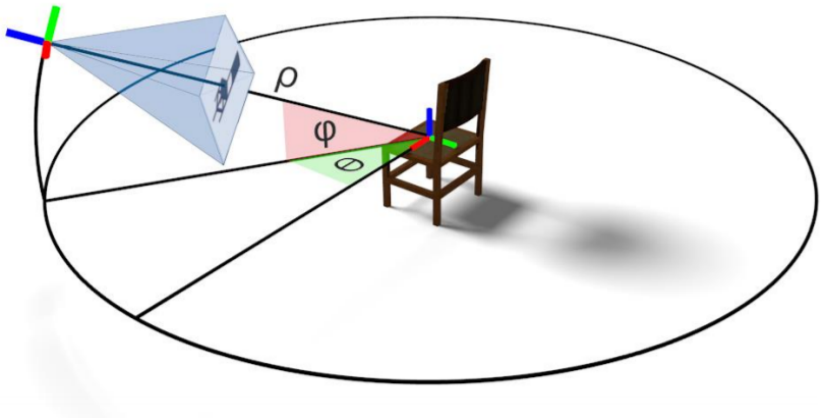

Figure 4.6: Illustration of how images of an object can be generated by defining a "viewing sphere" [135].

Projecting 3D shapes onto the image plane requires the use of perspective projection. Perspective projection uses a model of a camera along with a 3D point cloud to construct outlines of objects given the pose, as shown conceptually in Figure 4.5.

4.3.4 Perspective Projection

Projections of object models can be formed using perspective projection in which a central pin-hole camera model is defined which will correspond to the actual camera used during inspection (as shown in Figure 4.7). The pin-hole model projects 3D points onto the plane defined at $z = f$, where the centre of the projection is taken to be the origin. This central projection maps from Euclidean 3-space onto Euclidean 2-space as $(X, Y, Z)^T \mapsto (fX/Z, fY/Z)$ as shown in Equation 4.8 where P_C are 2D image points, A is the intrinsic matrix, K the extrinsic matrix and P_W the world 3D coordinates [151].

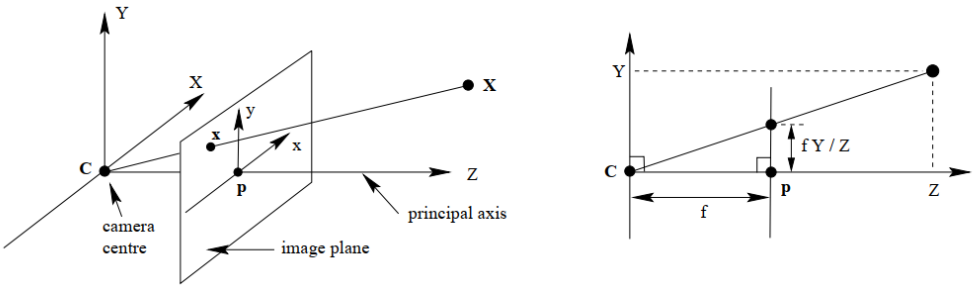

Figure 4.7: Pinhole model of a camera showing the global origin and camera origin, O. A world point in 3D space, P_W, is projected onto the image plane at point P_C. Image taken from Hartley and Zisserman [151].

The parameters which define the intrinsic characteristics of the camera are the focal length (f), the position of the central focal point on the imaging sensor (u_h and u_v), any lens distortions and skew (γ) arising from non-square pixels. The intrinsic parameters are captured in the matrix A, in Equation 4.8. The position of the camera's origin is captured by the extrinsic parameters which is a matrix consisting of the rotational and translational parts, denoted as K in Equation 4.8. These parameters are all found during a calibration step, which is a feature of many image processing software packages such as MATLAB.

$$P_C = A \times K \times P_W \tag{4.8}$$

$$A = \begin{bmatrix} f & \gamma & u_h & 0 \\ 0 & f & u_v & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, K = \begin{bmatrix} R & t \\ 0_{1\times 3} & 1 \end{bmatrix}, P_W = \begin{bmatrix} X_n \\ Y_n \\ Z_n \\ 1 \end{bmatrix}$$

Using a pinhole model of the camera, a set of image points P_C can be generated using the 3D point locations used in a CAD model. Once the points are projected onto the image plane (which consists of the same number of pixels as the actual camera used), the edges of the object are identified as described in Section 4.3.1. The overall process is illustrated in Figure 4.8 for a CAD model representing an aerospace engine

fan blade.

As the images are generated from points contained in a CAD file, the resolution of the CAD file has an effect on subsequent recognition and pose estimation functions. Throughout this work, Stereolithography (STL) files are used which consist of a number of small triangles which generate a dense point cloud as shown in step 2 of Figure 4.8. Other file formats such as Initial Graphics Exchange Specification (IGES) define geometries in terms of arcs, points and lines which would require additional post-processing to be projected accurately onto the image plane (such that lines and arcs would be drawn correctly). The advantages of using IGES files would be that contours would be drawn more accurately, but as long as STL files contain enough points the accuracy is sufficient for the intended application discussed in this chapter.

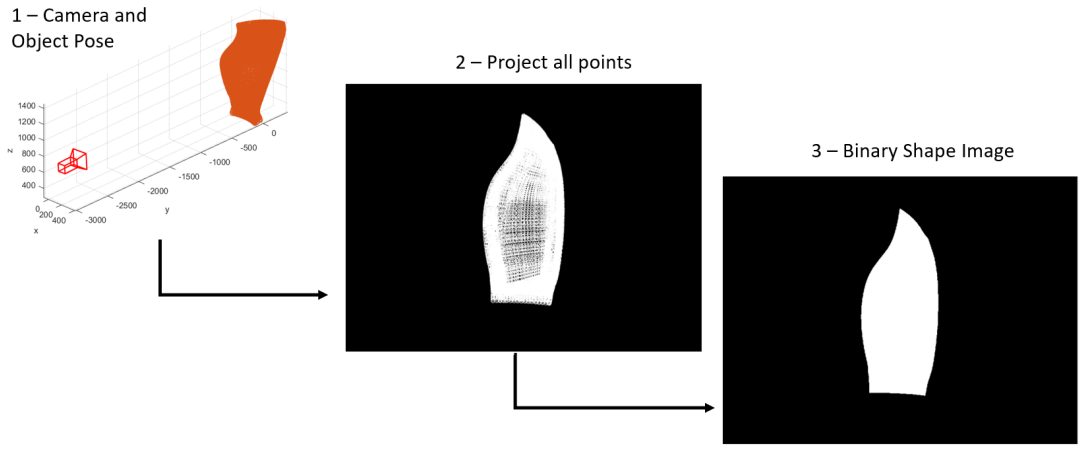

Figure 4.8: The three steps for generating binary images from CAD models. First the pose of the object and camera are defined. Then each point in the CAD model is projected using the pinhole camera model. Finally, points are connected to form a single binary shape.

This section shows how it is possible to generate projected object shapes from knowing the position of the camera and a geometric model of a component. This is a useful approach for quickly generating large repositories of shape images for the purpose of recognition which does not rely on the part being available. Another approach would be to attempt to generate 3D points directly using several cameras which is discussed in the following section.

4.3.5 Using Several Cameras

As light is projected through a camera lens onto the image plane, depth information is lost, as shown in Equation 4.8. All we know is that a point in 3D space lies along a ray projected from the centre of the camera through the pixel of interest, as visualised in Figure 4.5. To recover the lost depth information, the same point in space can be viewed from two different angles such that the full 3D information can be found via triangulation. This is illustrated in Figure 4.9a where P_1 and P_2 denote the principal points of two cameras. If the distance between the two cameras is known, the distance AC from point A to the epipolar line connecting the two cameras can be found. To be able to perform triangulation, however, corresponding points within each image have to be identified. Finding the same point in two images is the key challenge when using more than one camera. The main approach is to use key points, as described earlier in Section 4.3.1.

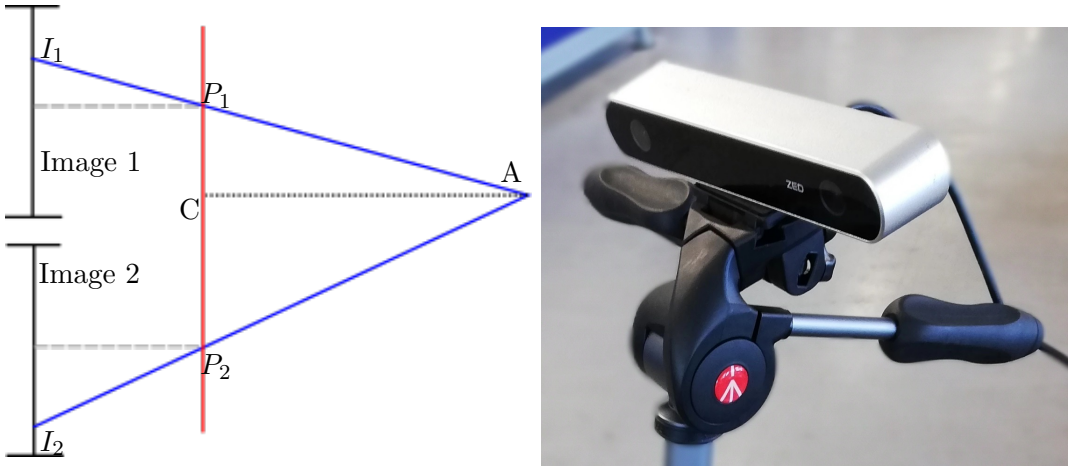

(a) (b)

Figure 4.9: (a) Principle of triangulation for two camera locations imaging the same point. (b) Example of a stereovision camera from the company Stereolabs.

A number of instruments using multiple cameras have since 2010 become available to the mass market, in part driven by the entertainment industry, but also because of the increasing popularity of robotics [152]. These systems, often known as 3D cameras, depth cameras or RGB-D cameras are all systems which provide 3D information about a scene. Stereo cameras use two cameras with a known separation, such as the ZED camera shown in Figure 4.9b. Other systems like the Microsoft Kinect camera use two cameras and an infrared projector to create a structured light pattern to identify key points. Other sensors use the Time-of-Flight (ToF) principle to determine distance. These systems have become increasingly popular for robotic applications, especially for scene reconstruction and navigation.

The ZED camera is a recent addition to the market which relies solely on image features (instead of structured light) and is aimed at long-distance 3D mapping with a range of 0.5 m to 20 m, combining both colour data and 3D points [153]. The company claims a positional accuracy of ± 1 mm during pose tracking, but this is most likely under ideal conditions. As absolute resolution is a function of distance to an object, suppliers instead state the image sensor resolution. The ZED camera uses two 4 MP cameras. In 2013, Gonzales-Jorge et al. [154] found that the first generation of Kinect cameras had an absolute accuracy of 5-10 mm (depending on angle and distance). The

following iteration of the camera, the Kinect V2 (which uses ToF), was found to have an accuracy between 1.5 and 10 mm [155].

Reconstructing 3D geometry is not limited to two cameras, but can either be performed by having several cameras or a moving a single camera to a number of positions. The science of taking measurements from images is known as photogrammetry. Another application considered in this thesis is the uses of motion capture systems which uses several cameras to track reflective markers.

The benefit of using camera systems which provide a 3D point cloud of a scene is that full 3D comparisons can now take place. This removes the need to map from 3D space to 2D space as discussed in Section 4.3.3. As CAD models of objects are available, this means that poses can be found directly, with no need for image projections. However, 3D cameras provide a lot of additional data which must be filtered to segment the objects of interest. Additionally, as points are reconstructed from image features, the measurements can be both noisy and inaccurate, as seen in the literature. Recognition and object classification is still possible using 3D data, however, the descriptors discussed earlier would have to be adapted from their 2D versions.

4.3.6 Classification

Once feature vectors have been extracted from images through image processing, the next step is trying to decide what object these features most likely belong to. This task is known as classification. Identifying which label to assign to a feature with reference to a data model is a key subject in machine learning and pattern recognition which has grown extensively in recent years (as discussed in Section 4.2.2). This section will discuss some basic classification principles. A more thorough treatment of classification and machine learning can be found in [156].

To be able to predict which label a new unseen feature vector needs to have, a set of training data must be obtained first. This will create a series of points in a space with the same dimensionality as the feature vectors. As an example, this could be the width and height, in pixels, of a segmented object as shown in Figure 4.10. The most straightforward way of comparing feature vectors is using the nearest neighbour

method. This simply calculates the Euclidean distances in feature space from a new point to every previously seen point and the minimum distance then determines the label given. This works well if features are clustered for each object, but separated from other objects.

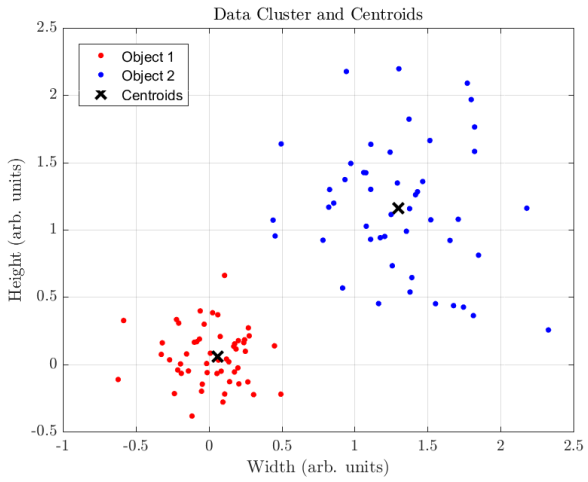

Figure 4.10: Example of data points from two objects plotted in feature space. This data is illustrative only. Using centroids can be a useful way of labelling new images.

If feature vectors for each object are not tightly clustered, the distance to n nearest neighbours can be considered instead. Alternatively, the mean of each cluster can be used, which can be obtained by K-means clustering. A linear classification scheme can also be used where simple rules divide the feature space into regions belonging to different classes. If a new point falls within a region, it is automatically assigned the label of that region. If a large number of training images are used, the variation of each feature point can be calculated. This allows for more robust classification schemes such as naïve Bayesian classification.

If very large datasets exist, neural network techniques can be used which attempt to determine a suitable way of linking inputs to outputs by varying weighted nodes in the network. As discussed in Section 4.2.1, these frameworks have had success in a number of fields due to availability of large data. There is, however, a question of traceability as the process of arriving at an output is hidden to the user. These are therefore not suitable for the current application where a wrong classification can lead to structural

damage. Thus, the way outputs are arrived at must be clearly defined and based on known statistical models. Naïve Bayesian classification, although conceptually simple, provides such a solution [148].

Naïve Bayesian Classification

Considering the example shown in Figure 4.10, a new measured feature point must either be classified as belonging to "Object 1" or "Object 2". Bayesian analysis allows us to incorporate prior knowledge and likelihood estimates to calculate the probability of a new point belonging to either category. Each of the previous objects has been measured 50 times, so we have no prior indication that a new point should be more likely to belong to any of the two groups. The prior probability is therefore 50% for each object. We can also see that the two groups of feature points are not evenly clustered. Thus, we have a 2D probability distribution for each cluster, as shown in Figure 4.11. This forms likelihood estimates for each object in feature space.

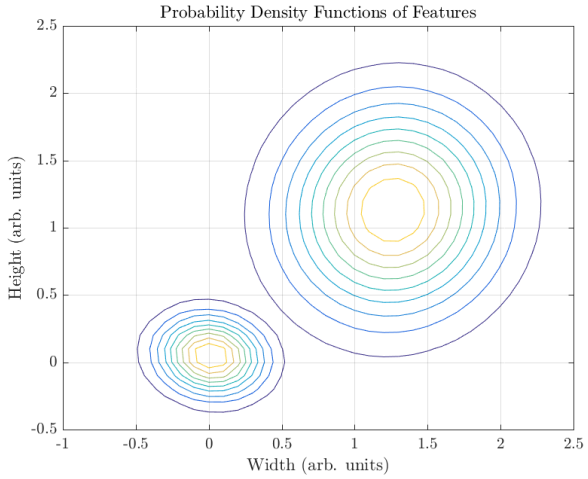

Figure 4.11: Probability density functions of points shown in Figure 4.10

Naïve Bayes classification combines the prior probability with the likelihood function to find the posterior probability. This is described by Equation 4.9 where a number of objects $k \in [1, 2...N]$ are labelled and described by feature vectors $\mathbf{x}$. The likelihood probabilities, $p(\mathbf{x}|k)$ have to be estimated during training phase and $p(\mathbf{x})$ can be estimated by summing $\sum_{k=1}^{N} p(\mathbf{x}|k) \cdot p(k)$.

$$p(k|\mathbf{x}) = \frac{p(\mathbf{x}|k) \cdot p(k)}{p(\mathbf{x})} \tag{4.9}$$

Through this relatively simple relationship between probabilities, it is possible to classify images. The method is however called "naïve" for a reason as it assumes all features are independent and relies heavily on the right choice of model for the likelihood estimation [119]. This can work favourably in situations where the datasets are small as long as the assumptions are clearly understood. Bayesian classification can also be used as a machine learning tool where the posterior probability of the last time step can be used as a prior distribution for the next [157].

4.4 Experimental Results

As discussed in the introduction to this chapter, there were a number of logistical challenges to obtaining the level of experimental evidence which was originally intended. This section describes the experiments which were done and which assess the suitability of the methods described in the previous section. In particular, the use of point-like descriptors is evaluated on model images while segmentation and recognition are performed on CAD model projections. Due to the robotic inspection cell being used for industrial work, a small-scale demonstrator was set up in a laboratory at TWI Wales and used for experimental trials. The small scale setup is shown in Figure 4.12.

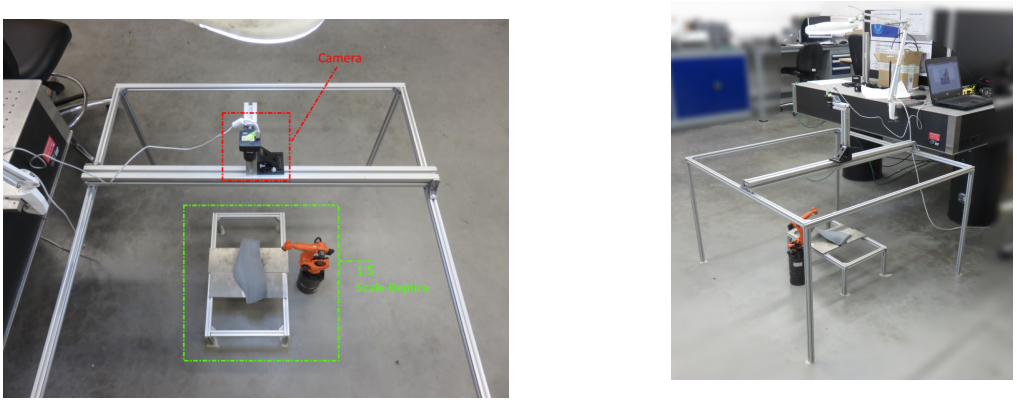

Figure 4.12: Small scale setup to test machine vision algorithms on suitable test samples.

The methods and algorithms described in the previous section were tested on CFRP samples to determine the optical appearance and how well features could be detected on this type of material. A scaled-down version of a aerospace fan blade was also 3D printed to provide a suitable geometry for shape-based recognition methods.

4.4.1 Using Local Feature Descriptors

As mentioned in Section 4.3.1, local features have been used extensively in recent years for object detection in images. To test how these would perform on composite materials, a small number of CFRP samples were photographed in the setup shown in Figure 4.12. It was found that SIFT, SURF and MSER features could all be detected using the MATLAB Computer Vision Toolbox [158]. The points that were detected, however, were all on the borders of the object. This was to be expected, given that the object itself appears dark and feature-less and the largest gradients are found on the border between the object and background, as shown in Figure 4.13. The dull, black appearance of CFRP is a fundamental optical property of the material.

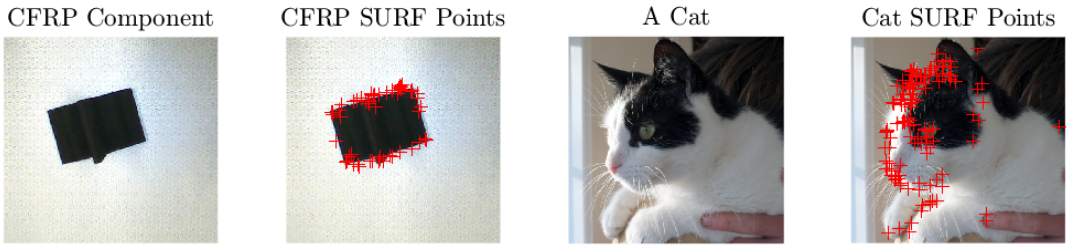

Figure 4.13: Figure showing how SURF point features are not well detected on dark, featureless components such as CFRP compared to natural images, like that of a cat.

Figure 4.13 shows the resulting difference in detected key points between a CFRP component and a natural image of a cat. The images clearly shows that most points are detected on the outside of the CFRP component while the many of the points on the cat image relate to its fur. Furthermore, the number of points detected in the CFRP image was 3 times less than those found in the cat image. Thus there are two main reasons for not using local feature descriptors when considering an automated detection system for robotic NDT. Firstly, it is time consuming to obtain images of

components before they are inspected and unlikely to be economical as components are typically not available prior to inspection. Secondly, the features that are likely to be detected will not be a property of the object itself, but be found at the edge of the component. Therefore it would make more sense to use the edge itself which can be obtained through image segmentation.

Another approach which could be investigated in this context is to artificially create feature points using structured light. As the surface itself appears feature-less, the use of directed light patterns, such as stripes, can be used to optically map the surface. This technique is the foundation of structured light 3D scanning (see Section 2.5.4). Instead of attempting to calculate the 3D surface, it may be possible to use a different pattern, such as a series of dots or speckles, to create local changes in light intensity. It would have to be investigated if these are effected by the surface properties of CFRP such that a unique pattern of local descriptors could be captured for each part.

4.4.2 Segmentation

Two experimental approaches were taken to image segmentation. Firstly 30 images were taken of the small scale setup shown in Figure 4.12 without an object present while varying the light intensity in each image. This was done to represent typical changes that would occur in a working environment. The average intensity for each pixel in each colour channel (RGB) was then calculated as well as its variance. A pixel threshold map for each pixel in each colour domain was then set as a predefined number of standard deviations from the mean.

An object was then placed in the scene and a new image was captured. The difference in intensity between each new pixel and the background map was found. Segmentation was achieved by assigning each pixel to zero if the intensity was within the tolerance zone and a one if the variation was outside. This is illustrated in Figure 4.14.

As seen in Figure 4.14, this approach to segmentation does not separate the object and background perfectly as the edges of the object remain fuzzy. Furthermore, a line exists through the middle of the component where the background and part are both dark. Morphological image operations such as opening and closing the region

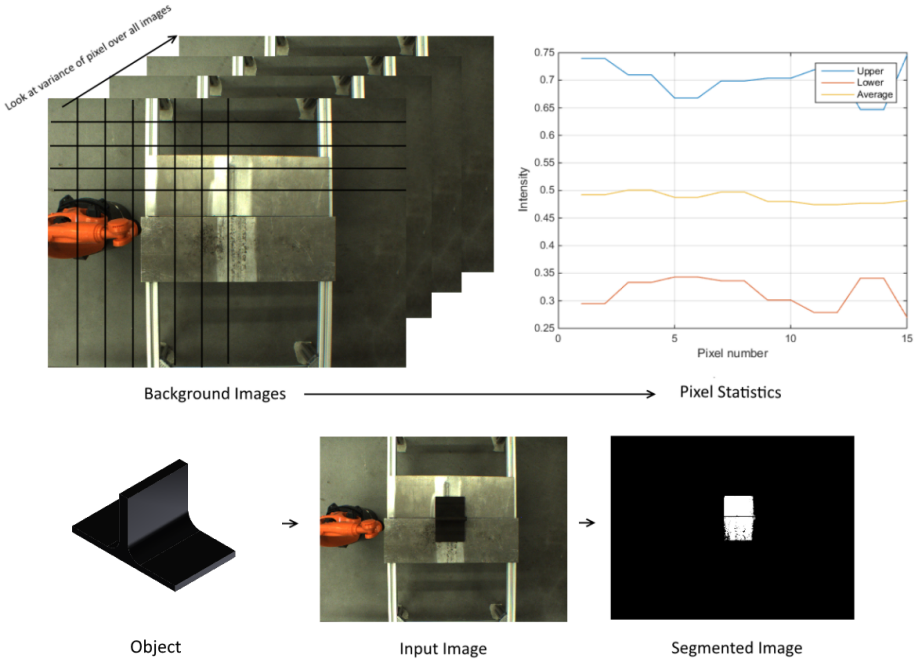

Figure 4.14: Image background removal (image segmentation). By calculating the variation of each pixel value for a range of different environmental conditions, thresholds can be established to distinguish background from object.

can be used to overcome this problem, but the technique is still limited by how well the variation in lighting can be captured. The advantage of this technique is that no prior knowledge of the object itself is needed, only information about the background scene needs to be captured. This approach is unlikely to work well within the IntACom robotic cell unless modifications were made to ensure light levels did not vary significantly. The need for further, potentially manual morphological operations also makes the approach less appealing.

A level set formulation was developed in MATLAB based on the theory described in Section 4.3.2 to segment objects from the background in an image. The algorithm was tested on a number of test images, as shown in Figure 4.15. A rectangular region covering about 50% of the area of each component was manually drawn to provide an initial contour, ϕ_0. This contour was then allowed to expand until it reached the edges of the component, needing about 500-800 time steps for each sample. Rosenhahn et al. demonstrate the use of the level set method for real-time tracking, showing it is

possible for the method to work much faster [150].

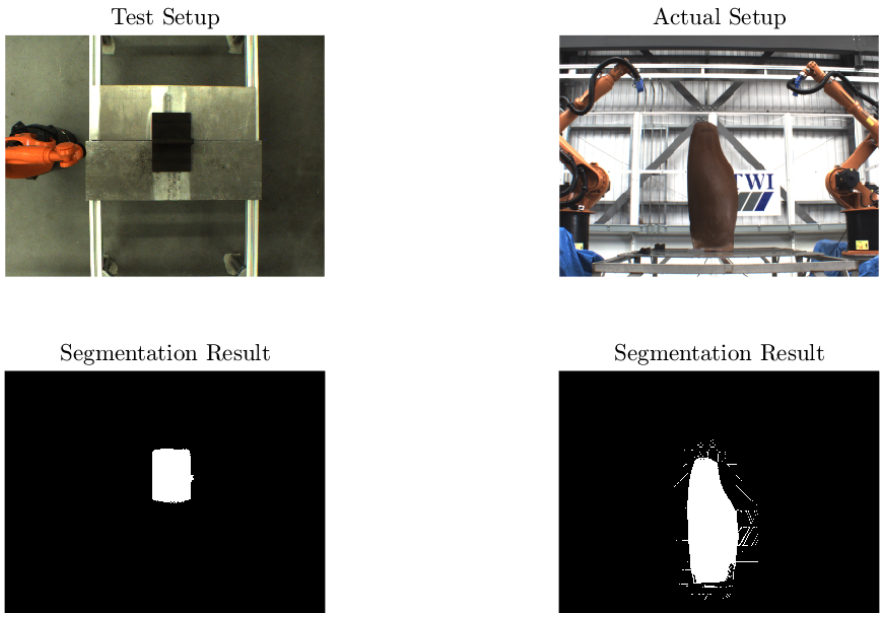

Figure 4.15: Segmentation results using the level set method

As seen in Figure 4.15, the level set method segments the object shown in the small-scale setup (shown on the left) with clearer edges and less noise. The edges are clearly defined and unlike the segmentation shown in Figure 4.14, there is no gap in the middle of the region. The level set method tries to minimise sharp edges in the contour meaning the method is more likely to segment an entire object and not leave gaps in the middle. Another advantage of the level set method is that no information about the background is needed and the method is more tolerant to changes in lighting or background clutter.

To test the applicability of the level set method in a real world environment, an image was taken of a CFRP aerospace component in the IntACom robotic cell. This is shown on the right hand side in Figure 4.15. Once again, a rectangular region covering about half of the component was manually chosen as a starting contour for the algorithm. As seen in Figure 4.15, the level set method is able to extract just the object while leaving most of the background out. Segmenting the singular object could not be achieved with a simple threshold method and does not require the user to first establish a background variance model.

It should be noted, however, that some of the background is still captured in the real-world segmentation example. This could be improved by either trying to make the background look more uniform or by using a bright light source behind the object which would highlight the edge of the component. It was however in the interest of the industrial partner to avoid repainting the background as well as investigating what could be achieved using the available light. Still, the example shown in Figure 4.15 illustrates how the level set can be used in a real world scenario. To further minimise any operator environment, a number of initial contours could generated from the CAD model and the contour which converges the fastest would then most likely be the actual component.

4.4.3 Obtaining Shape Models from CAD

One of the main challenges which this section attempts to overcome is how to recognise objects automatically when the appearance of the object is unknown. The only prior information is the digital 3D model of the component, from which we can deduce the shape. The idea is to try to recognise the shape of the object using the segmentation approach described in the previous section. To obtain the shape, a perspective projection of the model onto the image plane is needed, as discussed in Section 4.3.3. This section will go into more details about how these shape models are obtained in practice.

Three things are needed to construct a shape outline from a CAD file. First of all, the CAD file is needed. For this EngD work, only STL files were used as these simply contain a list of 3D points and an index for which triangles to draw by connecting the points. Next, a model of a camera is needed, as described in Section 4.3.4. To be able to relate this back to the actual camera used for experimental trials, the camera used was first calibrated. A camera calibration procedure is incorporated into the MATLAB Computer Vision Toolbox [158] and calibration is achieved by taking several images of a checker-board pattern from a number of different poses. By knowing the spacing between each square in the pattern, real world values such as the focal length and pixel size can be estimated.

The final piece of information that is needed is the pose of the CAD model with

respect to the camera. For the camera to be able to see the object, the object must naturally be in front of the camera. As the number of pixels is known, an image sensor can be modelled. Each 3D point can then be projected onto the screen as shown in Equation 4.8. To generate an image, the sensor is split into pixels and each projected point is assigned to one pixel value only. In reality the optics are going to cause a slight blurring but the method is considered accurate enough to extract a shape. The edge of the shape is extracted by finding the non-zero pixels closest to the edges of the image. An example is shown in Figure 4.16 where a model of an aerofoil shape is positioned 1500mm in front of the camera to ensure the entire object is visible.

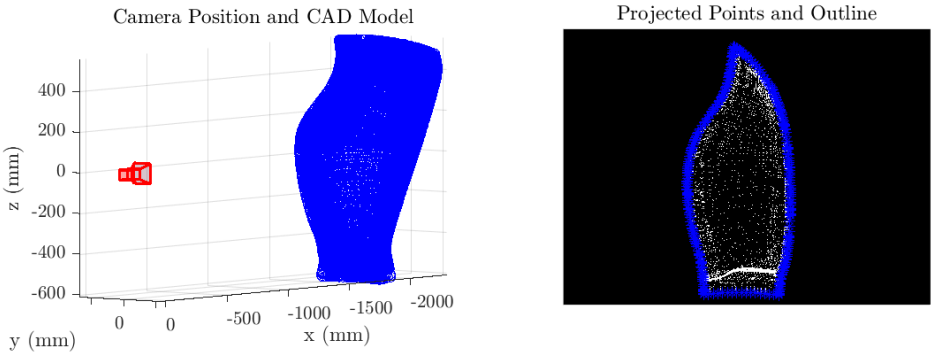

Figure 4.16: Example of extracting shapes from CAD models. Once the camera model and pose of the CAD are known, each 3D location in the model can be projected onto the image plane.

4.4.4 Camera and Geometric Calibrations

In the example shown in Figure 4.16, the pose of the object was manually chosen to fit within an image. For this approach to be practical, the pose of the object with respect to the camera needs to be known beforehand such that accurate shape models can be generated in their expected locations. Fortunately, this knowledge can be relatively

easily obtained as the pose of the component with respect to the robots is chosen during the path-planning phase (see Section 3.3.1). Therefore, if the pose of the camera with respect to the robot is known, then the pose of the component relative to the camera can also be found. To find the camera to robot pose, another calibration is needed, as shown in Figure 4.17.

As the camera is intended to remain fixed within the robotic cell, it is only necessary to carry out this calibration once. To define where the camera is, a global origin first needs to be defined. The choice of this origin is somewhat arbitrary, as long as it remains consistent. CAD models each have a separate origin, but these are component dependent. Therefore it makes sense to use the robot's internal world coordinate as a global origin. There are two reasons for this choice. Firstly, it is a fixed and well-defined point from where all other robot positional measurements are taken. Secondly, the inspection paths are all defined relative to a component origin which in turn is calibrated (in the first instance) using the robot. Thus, if the camera-to-robot transformation is known, so is every subsequent camera-to-component transformation.

The calibration procedure consists of attaching a rigid checker-board pattern to the end of the robot arm and moving the board to a number of poses throughout the robot volume. At each pose, the position of the wrist was recorded in the robot world frame and an image was taken. The position of the wrist was given by the internal kinematic model of the robot controller (transformation **A** shown in Figure 4.17(a)). For each position the MATLAB Computer Vision Toolbox [158] was first used to calibrate the camera and remove any lens distortion from the images. The position and orientation of the checker-board was then determined in the camera coordinate system (transformation **B** in Figure 4.17(a)).

A total of 15 robot poses were used to create an over-determined system of equations given by Equation 4.10. In Equation 4.10, **A** and **B** are both constructed as 15 4-by-4 matrices. The system of matrix equations given in Equation 4.10 can then be solved through a number of methods, as investigated by Shah et al. [107] who also provide open access to MATLAB implementations of their code.

$$\mathbf{A}\mathrm{X} = \mathrm{Y}\mathbf{B} \tag{4.10}$$

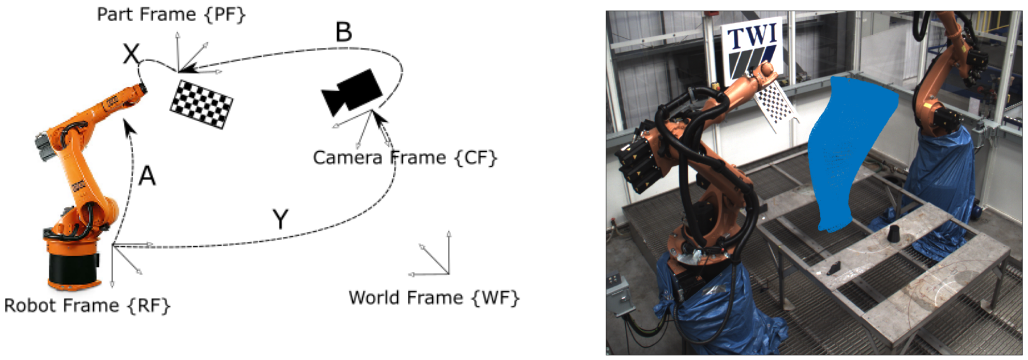

(a) (b)

Figure 4.17: (a) illustration of the transformations need for robot-to-camera calibration and (b) CAD model of an engine turbine blade projected into image of robot cell using robot to camera calibration.

Using the mentioned function, the two calibrations were determined simultaneously, unfortunately with a high re-projection error of 12.4 px. This error is expressed as the number of pixels between the projected solution of where the origin of the checker board should be and where it was actually observed to be. This error is most likely due to modelling the camera and lens as a single pin-hole camera. Although lens distortion is accounted for by using the in-built MATLAB camera calibration toolbox, the distance between camera and board meant that the black and white edges were not clearly discernible. This could be improved by using a higher quality lens and calibrating the camera on a bench-top, such that only the position and orientation of the board needs to be calculated.

During the experiment, robot poses were saved as a KRL programme which means that the procedure can quickly be repeated in the future. The origin of the camera is defined at one focal length away from the image sensor. The edges of the image sensor define a plane with x and y pointing along the horizontal and vertical axes of the sensor, respectively. The positive z-axis points from the centre of the projection out through

the lens of the camera, normal to the image plane. Using this coordinate system, the estimated position of the camera origin was at $x = 0.77$ m, $y = 1.29$ m, $z = 3.59$ m from the robot origin. This position was manually verified using a measuring tape and found to be plausible.

This transform can then be used with the chosen component pose from the path planning phase to project an object outline when components are in their expected positions. As an example, the aerofoil shape shown in Figure 4.16 was projected into one of the calibration images, as shown in Figure 4.17(b). Using the prior information from the path planning step allows the system to both construct images which can be used as initial contours for segmentation, but also to perform recognition based on the extracted outline.

4.4.5 Shape Based Recognition

To demonstrate the recognition principles discussed in Section 4.3.6, six available CAD models were chosen which represent typical aerospace parts. For each CAD model, a suitable camera-to-part pose was chosen manually to allow a shape image to be formed based on the formulation presented in the previous section. The six different models are shown in Figure 4.18. To establish a training set, a suitable range of pose variations was chosen which could represent real world variations. These variations were all implemented as a change in the orientation of the component with rotation angles ranging from -10 to 10 degrees randomly generated about the y-axis and -2 to 2 degrees about the z-axis. These values were chosen to represent typical viewpoint variations. Unlike standard image rotations, these rotations change the projection in a non-affine manner as each rotation leads to a new viewpoint of the object.

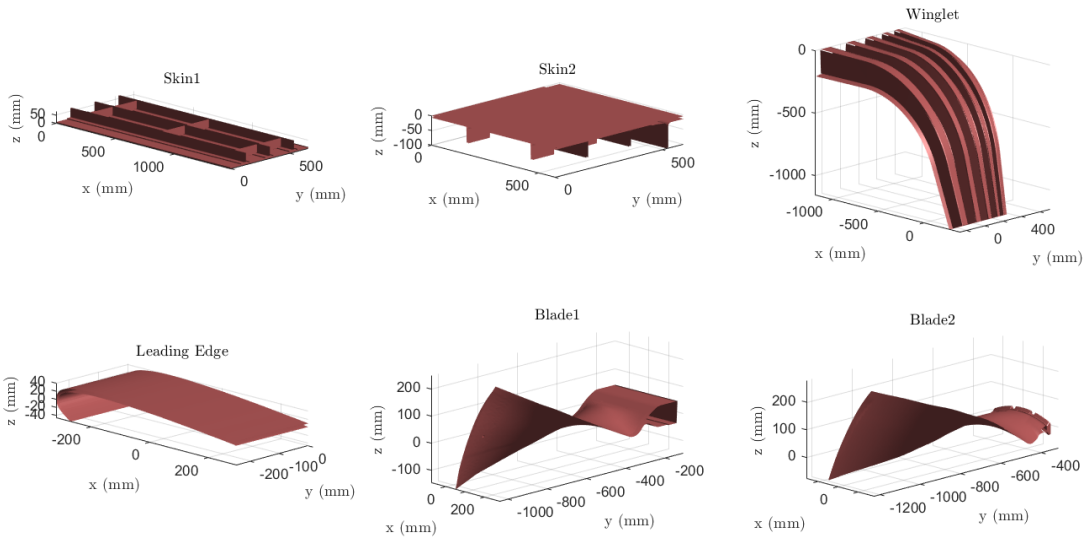

Figure 4.18: CAD models used for recognition.

A training set consisting of 25 different poses per CAD model was generated in MATLAB, leading to an overall training set of 150 images. A small subset, showing three different poses per CAD, of this training data is shown in Figure 4.19. Feature extraction was then performed on this dataset such that a 28-element Fourier descriptor vector and 7 Hu moments were generated for each image as described in Section 4.3.1. The pose and CAD model label were also saved for each descriptor. A smaller test set consisting of 3 images per model was then generated with random variations in the same range as the training set about the y- and z-axis. The same feature extraction process was applied to each of these images.

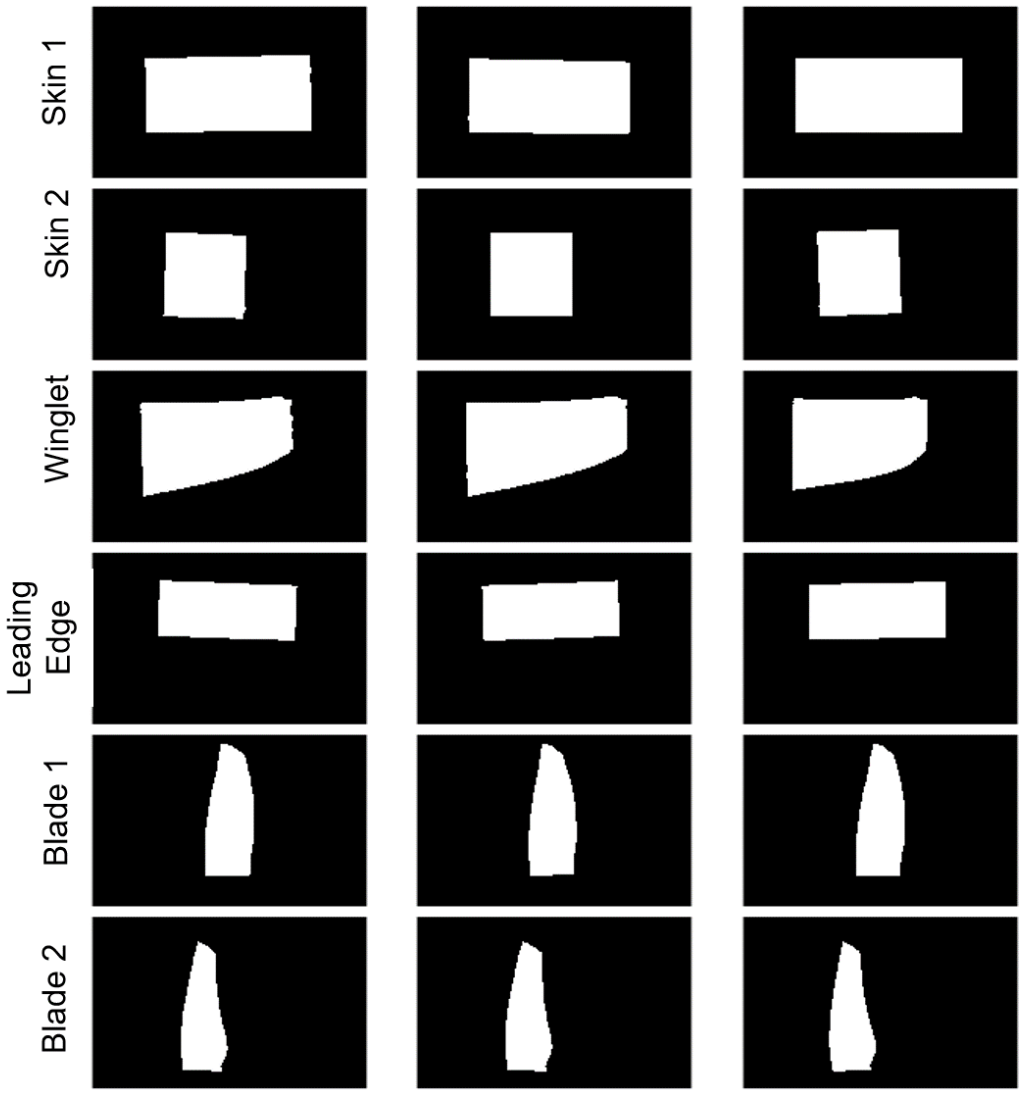

Figure 4.19: Generated shape images for six different CAD components, showing three different poses for each model.

Once the training image set was generated, the data was analysed to determine if the chosen descriptors (7 Hu moments and 28-element Fourier vector) were suitable to discriminate between different shapes. It was found that only the first two Hu moments served as a suitable descriptor as the rest of the moments did not yield any additional information. The average value and variation of the moments for shape images generated for each CAD model were calculated. A probability density function was then fitted to each of these datasets to construct a likelihood function which could later be used for Bayesian inference. Figure 4.20a shows the probability density function for the first moment. As the data was not normally distributed, a kernel distribution was fitted to the data using the MATLAB function `pfit`.

Fourier descriptors were generated using the generic Fourier descriptor method presented in the paper by Zhang and Lu [144] which generate feature vectors whose length is determined by the radial and angular frequency resolution. A vector containing 28 points was found to be a good compromise between being descriptive but not overfitting and quick to calculate. The distance between feature vectors can be used to match Fourier descriptors. To measure the distance between vectors, the function `pdist` was used in MATLAB using the standardised Euclidean distance measure. A feature vector from each component was compared against the others to form the matrix shown in Figure 4.20b.

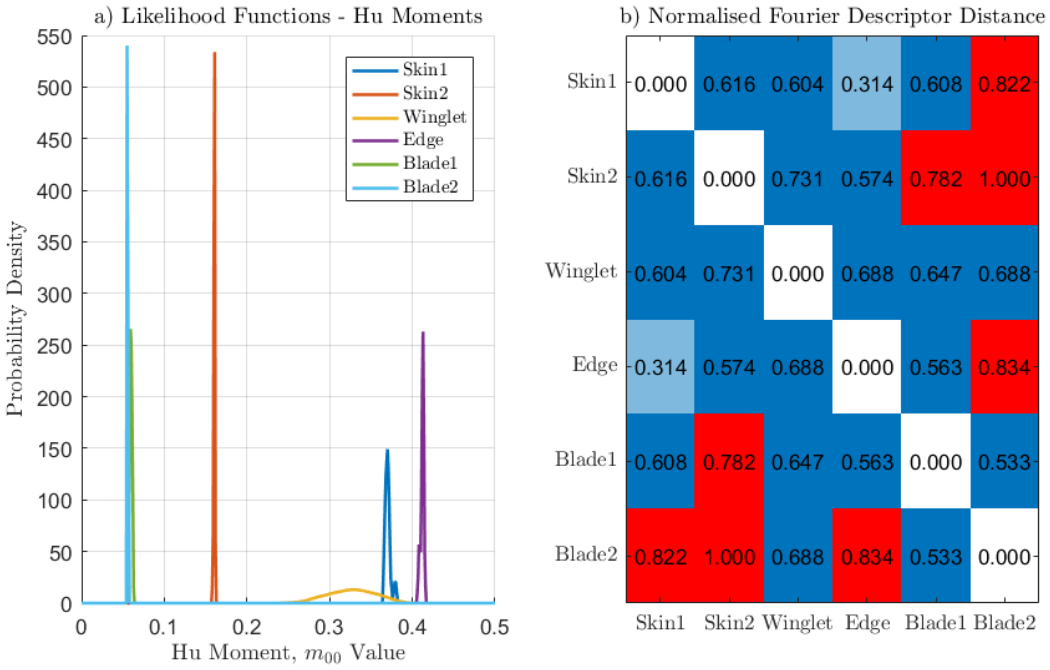

Figure 4.20: Descriptor suitability. a) shows the fitted probability density functions for the first Hu moment for each CAD model. b) shows the pairwise distance between each Fourier descriptor feature vector.

As seen in Figure 4.20, the two types of chosen descriptors each have their strengths and limitations. The Hu moments clearly distinguish between most models except for *Blade1* and *Blade2*, which are the two fan blade models, shown at the bottom right in Figure 4.18. The variation of each Hu moment is also very narrow indicating that for the chosen CAD viewpoints, the overall shape area and aspect ratio do not change drastically, except for the *Winglet* model which shows more variation.

Similarly, Figure 4.20b shows the normalised pairwise distance between the Fourier descriptor vectors between the different CAD models. The largest distance indicates the most dissimilar objects. Using this type of descriptor, the *Edge* and *Skin1* models appear most similar, but these are clearly separated when using Hu moments. Although the two blade models appear to have a similar shape, they can be relatively well distinguished using Fourier descriptors.

The fact that neither of the two descriptors is able to distinguish all six CAD models reliably shows the importance of fusing information from different sources. The

Bayesian Classification scheme, described in Section 4.3.6 provides a straightforward way of combining these different data sources. This framework was used to classify each of the 18 test images, giving the results shown in Table 4.2.

Table 4.2: Confusion matrix showing average classification probabilities for each of the six CAD models.

CAD Model Classification Probabilities

	Skin 1	Skin 2	Winglet	Edge	Blade1	Blade2
Skin 1	99.996	0	0.004	0	0	0
Skin 2	0	100	0	0	0	0
Winglet	0	0	100	0	0	0
Edge	0	0	0	100	0	0
Blade1	0	0	0	0	100	0
Blade2	0	0	0	0	11.182	88.818

The aim is to determine the probability of an object belonging to a particular model, M, given feature data $\mathbf{x}$ ($\mathrm{P}(M|\mathbf{x})$). A probability for each model is given for a new Hu moment measurement using the generated likelihood functions. For the Fourier descriptors, the average distance between a new feature vector and all vectors for each model is calculated. The probability is then defined as the ratio of each distance to the sum of all distances. As an equal number of images for each CAD model were generated, the prior probability was set to 1/6 for each model. Equation 4.9 was then used to calculate the probability for each object and the maximum probability was used to determine the classification result, as shown in Table 4.2.

As seen in Table 4.2, the classification framework nearly gave a 100% confidence for every model, except the second blade model. It is unusual to obtain zero probability for other models, but this is due to the very narrow Hu moment likelihood functions shown in Figure 4.20a. This illustrates that Hu moments are an adequate descriptor, at least for these types of shapes, as there is little ambiguity.

As expected, the two blade models with similar blade shapes are the hardest to distinguish between, but the combination of the two feature types meant that each blade was classified correctly in all test cases. Another test set was generated with a larger variation of angles around the z-axis ranging from -10 to 10 degrees. In this case, *Blade2* was incorrectly identified 75% of the time. This does not necessarily mean

that the approach is not valid, but helps illustrate that the training set must be set up correctly to reflect what will be seen during the actual setup.

As the training data was set up with CAD model poses similar to those seen in the actual robotic cell, the image of a real-life blade, shown in Figure 4.21a was used as input to the classification algorithm. Two additional image processing steps were performed to the result of the level set segmentation. First the MATLAB `bwconncomp` function was used to find the largest pixel region and then a morphological opening step was used to remove noise from the edges, as shown in Figure 4.21b. The resulting probability was 100% confidence that the shape seen was a composite fan blade, which was indeed the case.

a) Real Image

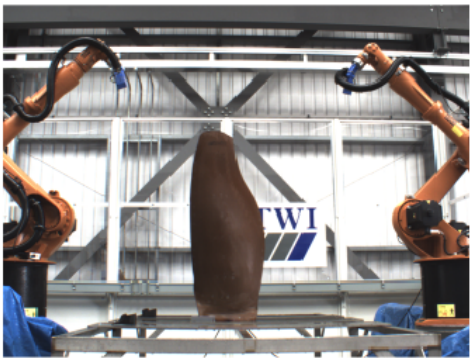

b) Segmented Image

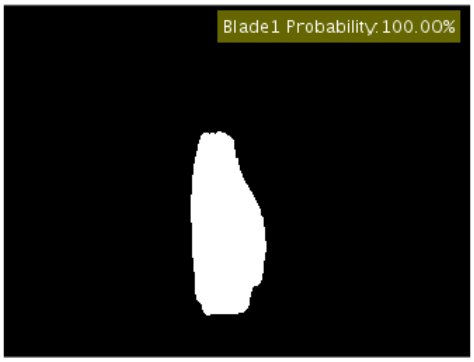

Figure 4.21: a) Real world image of a carbon fibre blade in the IntACom robotic cell. b) Segmentation and classification result.

The Bayesian classification framework using two different global image descriptors has been shown to provide good classification results on simulated shape images. There are naturally limitations to the chosen approach, for example when two CAD models have very similar shapes. The use of several features helps to overcome this problem. It should be noted that these results are only generated from simulated data although using a real-world image produced the correct classification verdict. Further work would be needed to test the performance in a real-world application.

One limitation of the way the recognition is implemented is that only a limited number of poses are generated for each model and that the shape image of each model has to be resized and moved to the centre of the image to efficiently generate the Fourier

descriptors. Hence, even though the descriptors are useful for recognition, they are not as suitable to determine the real-world pose of the component accurately. For this purpose, a system which can generate 3D data directly is more appropriate.

4.4.6 Using a 3D Camera

Although using standard global image features like Hu moments can lead to good recognition results, the fact that these descriptors are designed to be invariant to rotation and scale changes in images means they are less suitable for determining the actual pose of the component. To be able to determine the real world pose of an object, given its CAD model, a 3D camera can be used instead. As described in Section 4.3.5, a stereo-vision camera uses a known baseline distance between two calibrated cameras to infer the missing depth to a point identified in both images.

The ZED camera (Figure 4.9b) from Stereolabs was set up on a tripod and calibrated using a checker-board pattern and the MATLAB image processing toolbox. An image was then acquired of a static background and the same scene with a grey box measuring 60cm x 40cm x 40cm in it. This box was used as a CAD model of it could easily be generated in MATLAB. Pictures of the scene with and without a box in it are shown in Figure 4.22.

As mentioned previously, the ZED camera generates a 3D point cloud using optimised local image features to find corresponding points in both images. The camera can be instructed to discard points that are measured as being further away than a set distance, but any observed point closer than this threshold is displayed. This creates a very noisy 3D point cloud of the entire scene. Several approaches can be taken to identifying the 3D points of interest in a scene such as this. For example, if it is known that the object is on a flat surface, planes can found in the point cloud and anything not directly on the plane of interest can be discarded [159].

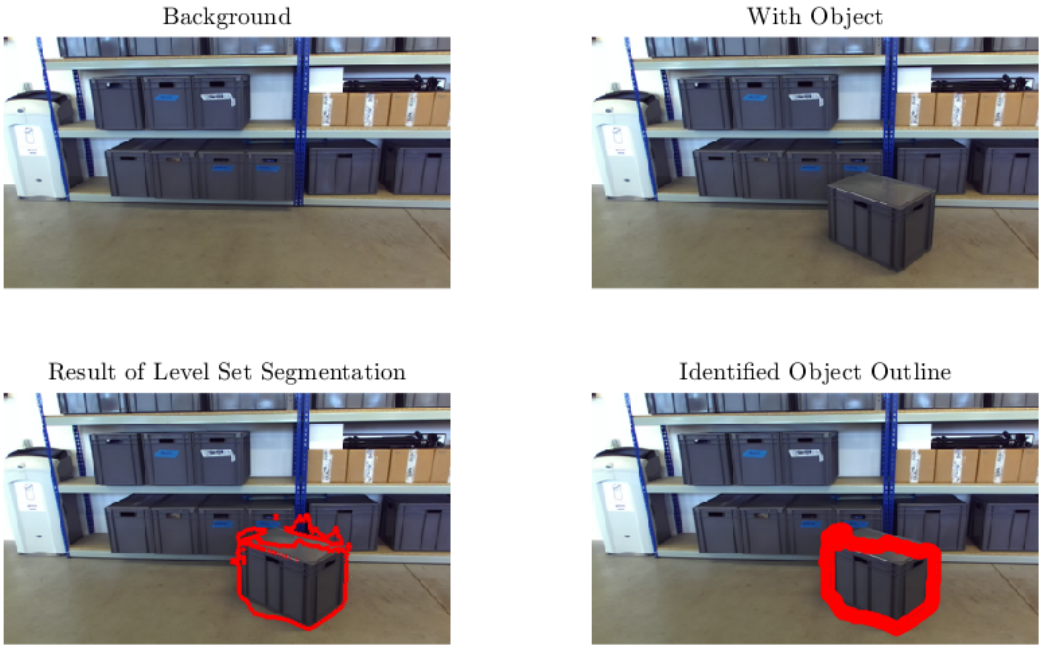

Figure 4.22: Image segmentation of an object in 2D images from the ZED camera.

Instead of attempting to segment the 3D point cloud directly, the approach taken in this work is to segment the 2D image and use this information to only obtain the 3D coordinates of interest. This is because segmenting based on 2D image features is more straightforward than trying to segment 3D data. By using a simple background subtraction, a starting mask for the level set method was obtained. The level set method was then used to segment the image (lower left image in Figure 4.22). The result of this segmentation was with limited success due to the similar boxes in the background. After another image processing step isolating the largest region of the segmented image, the object was identified as shown in the lower right corner of Figure 4.22.

The ZED camera allows the user to access either the 2D images of the scene, a 3D point cloud of the scene or a depth-map. The depth-map contains the same number of pixels as each image, but instead of intensity values, each pixel value is a distance along the viewing axis of the camera. This is important as there is no one-to-one relationship between pixels and the point cloud which means that identified object pixels cannot be directly related to a 3D position. By calibrating the camera first, the x and y positions

of a pixel of interest can be found by inverting the intrinsic matrix in Equation 4.8. The missing depth coordinate can then be obtained from the depth-map. This creates a mapping between the 2D image and the 3D point cloud such that segmentation can be performed, as shown in Figure 4.23.

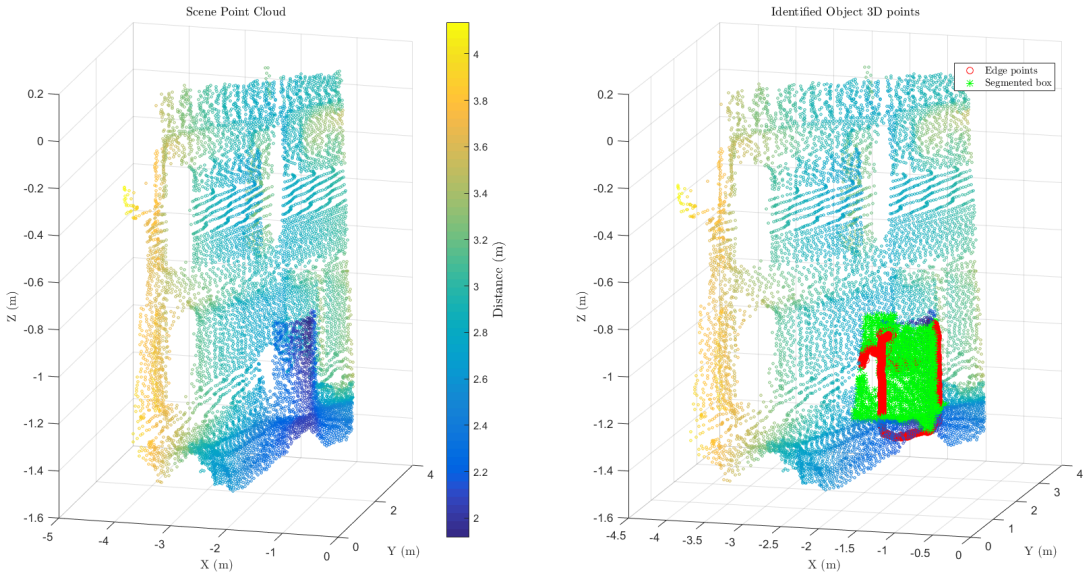

Figure 4.23: Left: Point cloud of the scene shown in Figure 4.22 with the box present. Right: Point cloud with segmented points corresponding to the box.

As seen in Figure 4.23, the point cloud contains areas with missing data where no features where detected by the ZED camera. The detected points create a noisy point cloud, but using the segmentation approach on a 2D image allows one to segment the point cloud such that points belonging to the box can be identified (shown on the right). To segment the entire box, the edge points where used to define a region in 3D space and all points within this region were extracted.

As a 3D point cloud of the object has now been extracted, this can be compared against a CAD model. To align the extracted data and a CAD model, the ICP method can be used [116]. An important aspect to take into account at this stage is that the extracted point cloud only represents a partial model as the rear of the object cannot be seen by the camera. Thus, the standard ICP method will never yield accurate results.

To overcome this, one could define key points in each point cloud and attempt to

match correspondences. The approach taken here is to only include surfaces from the CAD model which can be seen by the camera, assuming its viewpoint is known. As the viewpoint is established during calibration, this is a valid approach. A CAD model can relatively easily be segmented by considering the normal vector of each surface. By excluding any surfaces with normal vectors whose direction falls outside a set threshold (for example, 70 degrees), a partial CAD model can be constructed, such as the one shown in Figure 4.24.

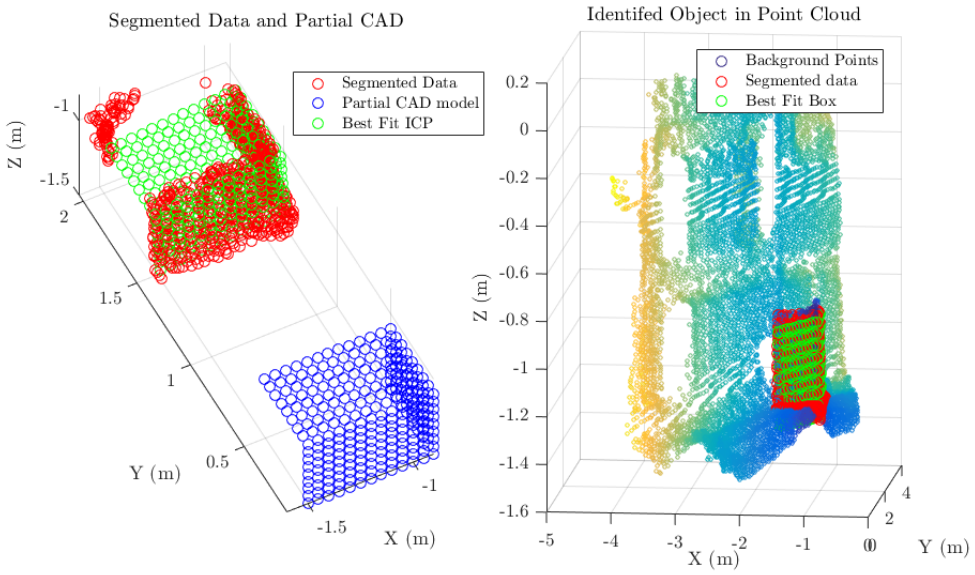

Figure 4.24: Left: extracted point cloud data aligned with partial CAD model. Right: Full point cloud with identified box.

As seen in Figure 4.24, using a partial CAD model leads to the extracted data and model being nearly perfectly aligned. The obtained transformation can then be used to align the full model with the background point cloud, as shown on the right. As seen in the figure, the position of the box is correctly identified in the 3D point cloud. There are several limitations to using this technique. First of all, the segmentation results must be quite accurate as noisy points affect how well the model is aligned to the data, as seen on the left. Secondly, the CAD model must be close to the expected pose for the ICP algorithm to converge correctly.

Even with the limitations stated previously, the use of a 3D cameras illustrates

how the actual pose of a component can be obtained by utilising the same concepts developed for the classification stage. Although the results here are only illustrative, the work flow for segmenting an image, classifying the segmented shape and obtaining the pose using a 3D camera and the segmentation result has been realised. Assuming the accuracy of the 3D camera is as stated by the manufacturer, this approach could be sufficiently accurate to determine the pose of an object to allow the ultrasonic alignment method to be successful. To implement this, however, the limitations mentioned above must be kept in mind. Future work should focus on investigating how accurate the proposed method is in a real-world environment and can be validated using an external measurement system such as the robot itself or a laser tracker.

4.5 Discussion and Future Work

As discussed in Section 4.1.2, the aim of the current chapter was to investigate possible methods for automated part registration in the specific context of robotic NDE applications. The type of part must be identified such that correct scan path can be selected. The parts pose must be determined to a tolerance that will allow an initial ultrasonic alignment path to be executed. A number of methods were considered for part recognition and a machine vision application was chosen due to its versatility and low cost.

One of the main limitations of the current setup is that only CAD data is available prior to inspection, which means standard computer vision techniques using direct image comparison is not feasible. This chapter has presented a classification and pose estimation strategy based on the shape outline of a part obtained from a CAD model. Using a pin-hole model of a camera, it is possible to generate shape images from these models and extract features from these. Actual part shapes can be segmented efficiently from real world images using the level-set method. Classification is then possible by extracting features from the resulting shape images to training a Bayesian framework. Finally, the same segmented image can be used to partition a point cloud obtained by a 3D camera to estimate the actual pose of the object. This process is visually

represented in Figure 4.25 which clearly highlights the central role the CAD model plays in the process.

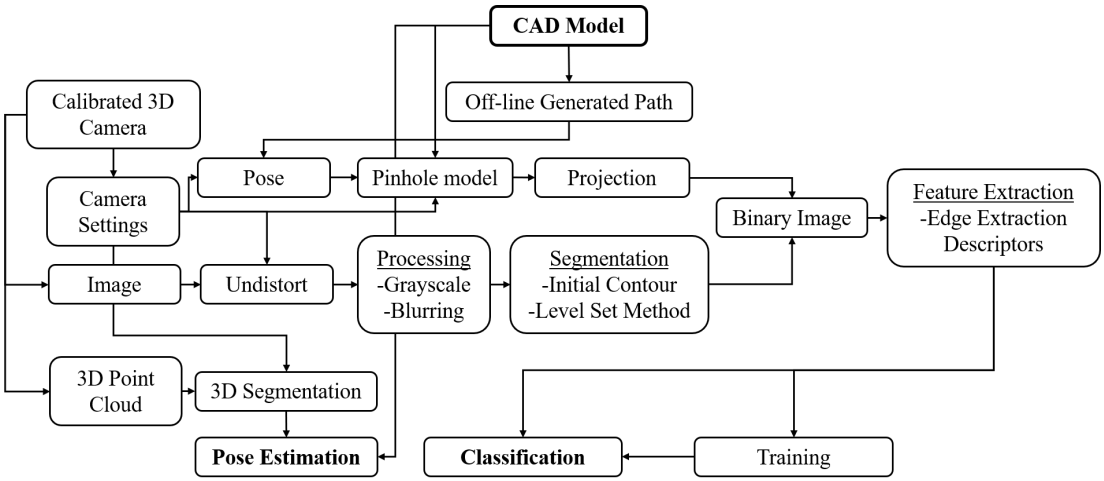

Figure 4.25: Flowchart showing how the CAD model and 3D camera interact to perform both classification and pose estimation.

A successful approach to object recognition which has widely been publicised in the literature is the use of local features such as SURF or SIFT. The use of these were considered in the current context but was not pursued for two main reasons. Firstly, a classification algorithm must be trained on a large number of images of the object to be recognised, which is not feasible as the parts typically aren't available prior to inspection. Secondly, carbon fibre components have a fairly uniform, dark appearance which means that strong correlation between feature points is hard to obtain.

Realising that the standard approach of using local features was unreliable, the approach was changed to use the shape of a component as a unique feature. This was motivated by the fact that CAD models of components are nearly always available prior to inspection as they are used for path planning. To extract the shape of a component seen in an image, it is necessary to segment the object from the background. Once the part is segmented, the edge of the component can easily be obtained and global features such as Hu Moments and Fourier descriptors can be extracted.

Segmentation is an important and active area of research within the computer vision community and a large number of segmentation methods are available. After reviewing

these methods, the level set method was favoured as it has a number of advantages. By not making assumptions about the background in an image, the technique is relatively insensitive to changes in lighting. Although an initial contour is needed to initialise the method, knowledge of the approximate part position can be used to define an initial contour. Dark CFRP components tend to produce clearly defined edges in images, so the initial contour will evolve until it encompasses the entire shape. The framework also allows the contour to wrap around holes in components which is another useful property. Determining exactly when to stop the contour and the best initialisation are both difficult tasks which are left as future work.

Experimental results showed the usefulness of the level set method in both laboratory trials and in the actual IntACom cell. Once the shape of a component has been extracted, it can be compared against a number of CAD shapes. Using a pin-hole camera model to project CAD points proved to be a fast and efficient way of generating shape images compared to using other rendering tools such as standard CAD software packages. The method can only work if the pose of the camera and CAD model are known. This is not an issue in the current context where both parameters can be found from a camera-to-robot calibration (presented in Section 4.4.4) and during the path planning setup. Assuming the camera remains in a fixed position, this calibration can be done once during the setup phase.

By choosing a suitable range of poses, a large dataset of shape images can be generated in the presented manner. It was shown how global image features can be extracted from these shapes and later used for recognition. In experimental trials, the variation in pose was limited to orientation changes around two axes. This limitation arose due to the time needed to generate datasets. Future work should focus on generating larger datasets and inferring descriptor information directly from 3D point clouds without the need to generate images first.

Object recognition was demonstrated by generating a dataset of 150 training images and using the Bayesian classification network to label 18 test images with nearly 100% classification accuracy. This high accuracy demonstrates that the chosen approach is capable of determining what part is in the cell. Even with such a small training set,

it was still possible to correctly identify the shape of a real-world component whose shape was obtained using the level-set method. Further work with larger, real-world datasets and a larger range of variations is needed to get a more realistic figure for the classification accuracy. The classifier must also be adapted such that it can identify unknown objects and accordingly classify this as "unknown".

It was initially hoped that the CAD data combined with a camera model would be sufficient to perform both part recognition and pose estimation. The approach of finding the optimal rotation and translation of a CAD model to fit a partial shape extracted from an image has been shown to work in the literature and is a feature in some modern metrology software packages [145, 160, 161]. This method was found to be require significantly more software development and was therefore not explored in this work. Instead, this work explored how a 3D camera, which indiscriminately generates a point cloud of an entire scene, can be used in conjunction with the developed segmentation approach to perform pose estimation. Future work should compare the developed technique against local feature descriptors such as 3D SIFT and assess the accuracy of the 3D camera in an industrial setting.

The industrial partners emphasised the need for improved calibration methods which meant that less focus was placed on automating the initial part alignment procedure. Due to these constraints, it was not possible to integrate the developed part classification and pose recognition framework into the existing robotic cell. This chapter has however presented a clear proof of concept (illustrated in Figure 4.25) using real and simulated data. Future work should focus on how this system performs under actual industrial conditions and how it can be used to automate inspections further by interfacing with industrial robots. For KUKA robots, this will mainly consist of sending updated base coordinates via the RSI connection, along with a command to run a specific sub-program for the recognised part.

4.6 Conclusions

The philosophy behind the existing IntACom robotic inspection cell is to obtain a flexible and adaptable automated NDT tool. Automated path planning is utilised such that robotic paths can be designed prior to the part actually arriving on site. To ensure flexibility, these paths are defined relative to a part coordinate reference frame instead of as an absolute position in the robot's work space. Low cost, 3D printed fixtures and standard aluminium extrusion rails are used to fix parts into approximate positions. Together this creates a more attractive solution to industry as setup times are reduced and the need for expensive jigs eliminated. Current practice, however, needs a human operator to manually determine the position of the part by manually moving the robot. An objective of this thesis was to investigate how to automate this step of the process.

This chapter has discussed different ways of automatically identifying parts using computer vision and estimating their pose relative to a CAD model. Using a camera and computer vision techniques was seen as advantageous as it provides a large volume measurement tool which can work with minimal alterations to the existing setup. It furthermore provides a low cost and versatile approach that can be optimised to the specific application. An approach designed around the use of CAD models and shape features was specifically designed to meet the challenges and limitations faced by robotic NDT systems.

Table 4.3 summarises the approaches and solutions adopted in this chapter to solve the challenge of object recognition and pose estimation for robotic NDT. The listed criteria were discussed in Section 4.2.2 and must be addressed whenever one considers a machine vision application. It was shown that objects could be recognised based on their shape outline with a high degree of confidence using a Bayesian classification approach. The use of a 3D camera alongside 2D segmentation approaches illustrated how it is possible to perform pose estimation if the CAD model is already known.

A limitation encountered throughout this chapter was being able to obtain a unique image reference for CFRP components due to their optical appearance. The use of segmentation methods overcomes this problem to some extent, but still relies on back-

Table 4.3: Approaches and solutions presented in this chapter to solve the problem of automatically recognising objects and estimating their pose.

Summary of Implementation		
Criteria	**Approach**	**Solution**
Representation	Shape-based model	CAD models, shape images
Data model	Global features	Hu moments, Fourier descriptors
Expected variance	Lighting and pose variation	Level set method, generating set of shapes
Matching	Bayesian classification	Likelihood functions, distance between features
Pose estimation	3D camera	Segmentation, calibrated cameras, ICP algorithm

ground lighting to provide clear edge definitions. Future work should investigate if a projected light source, such as a structured light system, could aid in this respect. Perhaps by project stripes onto a component, the reflection from the CFRP surface will more clearly indicate where the edges are which in turn can be used with the developed edge matching strategy.

From an overall system integration perspective, the Bayesian formulation could be extended to not only encompass different image descriptors, but also the knowledge of part placement. As an initial pose estimate is provided at the path planning stage, this can be used as the prior information. Further measurements taken using optical images or ultrasonic surface reflections can then update this prior to achieve a better understanding of uncertainties in where the object actually is. Future work should investigate how such a formulation could be realised and whether it would provide a better pose estimate.

Future work will also focus on implementing the developed approach into the IntA-Com inspection cell. This will involve accurately calibrating the relationship between cameras, robots and off-line path planning and establish whether the system can deliver accurate enough results to allow the ultrasonic alignment method (presented in Chapter 3) to be used to perform fine-scale adjustments. If this is possible, it will further support the over-arching aim of this thesis to create a flexible, multi-scale approach to metrology for robotic NDT.

Chapter 5

Validation Experiments

5.1 Introduction

The aim of non-destructive testing is to assess the structural integrity of a part. To ensure the best possible data is available for an operator, the quality and reliability of the inspection procedure itself must first be assessed. The overarching topic of this thesis is how to ensure that robotised inspection systems provide accurate and reliable results each time they are run. The commissioning and integration of a robotic NDT system is a complex task, so the user must still be confident that any observed indications stem from a physical phenomena pertaining to the part's structure and not an error in the system setup. As an example, the reduction of surface echo intensity could be due to probe misalignment, but may be interpreted by the operator as the surface being more absorptive at this point. Situations such as this must be avoided as they can lead to unnecessary and expensive repair operations.

The purpose of this chapter is to introduce and discuss validation experiments for the topics discussed in Chapter 3. In particular, this chapter will seek to quantify, through experimental trials, the accuracy and applicability of the developed calibration techniques. As discussed in Section 3.5.1, three calibrations are needed to obtain accurate results from a robotic ultrasonic inspection system. Firstly, the internal model of the robot must be calibrated such that the robot wrist is accurately moved in 3D space. Next, the geometrical relationships between wrist and tool must be defined to

let the robot know how to move its end-effector. Finally, the geometrical relationship between robot and work-piece must be established. Only when all three calibrations have been performed can the user rely on external path and trajectory planning to ensure that data gathered is accurate. Each of these calibrations is discussed in this chapter along with an assessment of the overall positional accuracy of the IntACom system.

The positional accuracy of a large-volume commercial motion tracking system is also assessed. The purpose of this is to ascertain if such a system could be used to provide real-time feedback to improve robot accuracy. The chapter is structured as follows: firstly, a baseline dynamic and static positional accuracy of an industrial robot and photogrammetry are assessed simultaneously. Next, experiments which empirically determine the maximum allowable misalignment between probe and part are presented. This is followed by experimental trials in the IntACom system implementing the developed calibration techniques. Finally a discussion of the experimental results is presented along with suggestions for further work.

5.2 Hardware used

As this thesis is a results of a collaboration between the University of Strathclyde and TWI Ltd, experimental work was carried out at both institutions. This section describes the robotic, optical and ultrasonic hardware used throughout this chapter.

5.2.1 Robot Systems

The experiments presented in this chapter were conducted at two different facilities as discussed in Chapter 1. Initial experiments were carried out using a KUKA KR5 arcHW robot at the University of Strathclyde. This robot was part of a dual robot setup, established within the Centre for Ultrasonic Engineering (CUE) in 2010 to investigate the uses of robots for automated NDT. The outcome of this early research led to the establishment of a twin robot cell at TWI's facilities in Wales in 2013. The robot cell at TWI consists of two KUKA KR16 L6-2 robots [162]. The research presented here

was firstly carried out at the University of Strathclyde in 2014-15 and at TWI from 2015-18. Although both robot systems used are from KUKA, the robots in CUE use the second generation KRC system while the robots at TWI use a fourth generation robot controller. The two different types of robots used are shown in Figure 5.1.

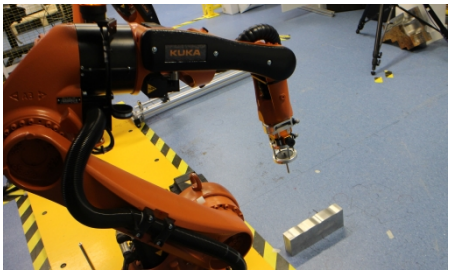

(a)

(b)

Figure 5.1: (a) A KUKA KR5 robot at the University of Strathclyde. (b) The IntACom inspection cell with twin KUKA KR16-L6 robots at TWI.

Besides the difference in robot control system, the reach and size of each robot is also different. The KR5 robot has a maximum reach of 1423 mm while the the KR16 can reach 1911 mm. The maximum payload of each robot is 5 kg and 6 kg, respectively. The robots have a positional repeatability (discussed in Section 3.5) of 0.04 mm and 0.05 mm, respectively. The KR16 robots at TWI were configured using the collaborative technology package RoboTeam from KUKA [163] which allows a master-slave relationship to be defined between the two robots. This package allows up to 15 robots to perform coordinate motions in five different ways. The robots can either work independently, synchronise wait points, synchronise motion times, as well as being directly and in-directly coupled.

The main advantage of using the RoboTeam package is that ultrasonic through-transmission paths can be realised by having one robot's motions mirrored by the other. Besides the RoboTeam package needed for cooperative motion, the RSI [89] and EthernetKRL packages are also installed. These two additional packages were also installed on the robots in CUE and allow for positional information to be streamed back to a server as well as exchanging digital signals between an external computer and

robot controller.

5.2.2 Ultrasonic Systems

A MicroPulse 5PA pulser and receiving unit from PeakNDT was used to produce and record ultrasonic data. The MicroPulse unit features both conventional and phased array channels along with four positional encoder inputs [164]. The unit is particularly well suited for the research and advanced NDT community as a Dynamic Link Library (DLL) is openly available with an extensive command reference manual. The MicroPulse unit is shown in Figure 5.2.

Figure 5.2: MicroPulse 5PA ultrasonic hardware from PeakNDT [164].

To gather ultrasonic data the IntACom software was used which communicates with the Micropulse through a high speed Ethernet connection. The Micropulse is capable of sampling data points at 100 MHz with a 12 bit signal amplitude resolution. In the current setup the equipment is set to pulse continuously and received ultrasonic data is then time-stamped to enable positional interpolation as described in Section 3.4.

5.2.3 Optical Systems

As discussed previously, high accuracy measurements are possible using optical systems, especially laser trackers (see Section 2.5.3). The limitations of using laser trackers are the need for a single line of sight and their high cost. Motion capture systems provide an alternative optical measurement system which can cover a larger volume and track several objects simultaneously. The accuracy with which these can measure reflective markers is however not well defined. Measurements provided by laser trackers were

used as reference values due to their high accuracy. The experiments presented in this chapter concern both the positional accuracy of industrial robots as well as the obtainable positional accuracy of a commercially available motion capture system.

Experiments which took place at the University of Strathclyde used a Leica Absolute Tracker AT901 [165]. This tracker has a measurement range of 80 m and can measure at up to 1000 points per second. The accuracy of the tracker is $\pm 15\mu m + 6\mu m/m$ and was used at a distance of about 2 m. At TWI's facilities, a FARO Vantage^{E} [166] was used to track the movements of both robots individually. This tracker has a measurement range of 60 m and can also measure up to 1000 points per second. The accuracy of the tracker is $\pm 20\mu m + 5\mu m/m$ and was used at a distance of about 3 m. The two laser trackers are shown in Figure 5.3

(a) (b)

Figure 5.3: (a) A Leica Absolute Tracker AT901 [165]. (b) A FARO Vantage laser tracker [166].

At the University of Strathclyde a Vicon T160 motion capture system was used in a small volume covering the robotic cell available in CUE. The volume covered was approximately 1.5 m^3, which also roughly corresponds to the working envelope of the robots at TWI. The Vicon system works by emitting pulses of infrared light and detecting how this light is reflected from specially designed targets placed on an object. A minimum of three targets are needed to provide an identification of an object and at least two cameras must have a line of sight to the objects. The system used in the case studies consisted of five Vicon cameras.

5.3 Motion Capture Systems

Chapter 4 illustrated some of the difficulties in using machine vision as a tool for improving the automation levels of a flexible robot cell. Although a single camera system is a more elegant and robust solution from a hardware perspective, the requirements for advanced algorithms, large datasets and controlling the lighting conditions are key challenges identified in this thesis. The use of two cameras can provide full 3D coordinates as shown in Section 4.4.6, but the segmentation and identification of objects still hinder the ease of integration and use.

As discussed in the research background chapter (Section 2.5.3), motion capture systems provide a method for large volume object tracking. The use of several calibrated cameras allows each point in the volume to be found with a higher degree of positional accuracy compared to a stereovision system and the use of specific reflective targets alleviate the need for complicated segmentation algorithms for finding key points. As more cameras are present throughout the volume, objects moving along complex paths can also be tracked easily as long as a line of sight to two cameras exists.

Motion capture systems are widely used in a number of industries today, but their main use is either for animation films or in medical applications to study body movement. The development of motion capture has mainly been driven by the animation industry, starting from tracing live actor's movements onto paper, frame by frame in 1915, to the modern day use of reflective markers in the 1980s. The technology has since evolved with the increased availability of computer processing power such that markers can automatically be identified in real time [167]. The ability to identify markers automatically and stream their 3D position thousands of times per second has also led motion capture to be used in robotics, especially for tracking drones and validating kinematic models.

The advantage of using motion capture systems for robotics, besides the automated identification of markers, is that a external measurements can be taken of a robot's position and orientation simultaneously. Most robots, as discussed earlier, are open-loop feedback systems where a forward model is used to estimate the pose of the

robot. Unlike a laser system, a motion capture system does not need a single line of sight to a reflective target, but can still track objects as long as at least two cameras can see a marker. This gives the system the advantage of being more flexible, but with the disadvantage that the obtainable positional accuracy is generally lower than that of laser trackers. The loss of accuracy stems from several factors such as noise in measurements, the choice of reconstruction algorithm, lens distortion, and limited depth of field (depending on the aperture of the cameras used).

Bearing these limitations in mind, it was investigated if a motion capture system could provide a solution to several challenges identified in this thesis. Firstly, if the accuracy of the system is high enough to track objects to within a few millimetres then it can be used for automated part positioning, as discussed in Chapter 4. Although it is unlikely that reflective markers can be attached directly to an object which is to be inspected, the object itself has to be mounted in a rigid fixture prior to inspection. By attaching markers to the fixture instead of the object it would still be possible to infer the pose of the object with respect to the robot.

Secondly, the ability to track several objects simultaneously with 6 DoF means that motion capture systems could potentially be used to either calibrate multiple industrial robots or be used to update the path of a robot in real time if the feedback rate of the cameras are high enough. A limitation of most laser trackers is that data is only obtained in 3D while a motion capture system will also be able to monitor the orientation of a probe, which is another crucial parameter in NDT.

5.3.1 Principle of Operation

This section discusses the principle of operation behind two of the world's leading commercial manufacturers of motion capture systems: Vicon (Oxford,UK) and Qualisys (Gothenburg, Sweden). Both systems work by using infrared illumination and calculate position of a marker using photogrammetry. This means that any point in 3D space imaged by different cameras can be located if the angles and distances between cameras are known. The problem of reconstructing m points from N views can be cast as a minimisation problem, commonly known as bundle adjustment. A projection matrix

(P^i) can be constructed for each camera, containing both the internal (focal lengths and pixel sizes) and external (position and orientation in world frame) parameters. Bundle adjustment seeks to estimate the matrices P^i given a set of 3D points $\mathbf{X}_j$ and their corresponding image coordinates, $\mathbf{x}_j^i$ for the j^{th} point in the frame of the i^{th} camera such that $P^i\mathbf{X}_j = \mathbf{x}_j^i$ [151].

As real world measurements are imperfect, the above equation is never able to be solved exactly so instead a maximum likelihood solution is sought assuming that the noise in measurements is Gaussian. The projection matrices and 3D points are then estimated in such a way that the distances between the re-projected image points and measured image points are minimised. This can be expressed as Equation 5.1, where $\hat{P}^i$ and $\hat{\mathbf{X}}_j$ denote the best estimates of projection matrices and 3D positions, respectively.

$$\min_{P^i, \mathbf{X}_j} \sum_{i,j} d(\hat{P}^i \hat{\mathbf{X}}_j, \mathbf{x}_j^i) \tag{5.1}$$

Equation 5.1 is a refinement step, which needs a good initial estimate to converge. Furthermore, to establish a metric reconstruction, a real-world distance must be provided. Motion capture systems make this process less complicated in a number of ways. First of all, pre-calibrated cameras are used such that internal camera parameters do not need to be estimated in the projection matrix. Secondly, the number of points that need to be reconstructed is minimised by using a ring of infrared light projectors and special coated targets which are particularly reflective in the infrared spectrum. Finally, a calibration step is performed prior to any measurements which estimates the position and orientation of each camera with respect to a fixed coordinate system as well as providing a precisely known distance. Both Vicon and Qualisys provide carbon fibre calibration "wands", shown in Figure 5.4. During the calibration phase, a wand is moved around the volume which is to be tracked and the systems can subsequently use the estimated camera poses to reduce the number of parameters which need to be minimised in Equation 5.1.

Figure 5.4: Calibration wand kit used to determine camera poses and provide a real-world distance to enable the system to accurately reconstruct 3D points.

In theory there is no upper limit to the number of cameras which can be used for motion capture, but each reflective marker must be seen by at least two cameras simultaneously to provide a 3D position. Using reflective tape and infrared light means that the marker should show up clearly in the image, leaving the background dark. This allows the obtained image to be binarised based on a set intensity threshold. By using parallel processing architecture, such as Graphics Processing Units (GPUs), calculations of registered image points can be done faster, which leads to output rates of several hundred Hertz for both systems. Further increases in output speed can also be achieved by sub-sampling the detector (effectively lowering the resolution), resulting in 2000 Hz output rates at a resolution of 1.3 megapixels (Vicon) or 1100 Hz at 3 megapixels (Qualisys). Motion capture manufacturers typically supply packages consisting of at least four cameras, depending on the type of application.

Several previous studies [168, 169] have demonstrated that photogrammetry systems, such as the Vicon MX, can be very accurate (20-200 μm) within small volumes, as found in the biomedical sciences. In the field of engineering, Vicon motion capture systems have been used to track and manoeuvre drones and validate wheeled robot motion control systems. Validating the accuracy of a photogrammetry system has been an active area of research in recent years. Studies at the University of Strathclyde [170] have shown the type of calibration artefact used has a significant impact on the performance of the system. This was achieved by measuring the accuracy of the Vicon T160 system using 8 cameras over a large volume measuring approximately 100 m^3 [170].

An initial part of this thesis work sought to determine if the accuracy of a standard motion capture system was sufficiently high to be useful for real-time robot path corrections. To be able to quantify the usefulness of the system for this purpose, an experiment was designed to compare the accuracy of a KUKA KR5 robot and a Vicon motion capture system.

5.4 Robotic and Vicon Positional Accuracy

An experiment was designed to measure both the positional accuracy of a KUKA KR5 Arc HW robot as well as the positional accuracy of a Vicon T160 motion capture system with five cameras covering approximately 1.5 m^3. A Leica AT901 laser tracker was used as ground truth comparison for both the robot and Vicon positions. The two optical systems were set up around one of the robots available at the University of Strathclyde as shown in Figure 5.5. Tripods were used to position the cameras at suitable heights providing a good overview of the robot's working envelope. The results of the experiment described here were presented at the IEEE 2016 Multisensor Fusion and Intelligence Conference in Baden-Baden, Germany and published in the conference proceedings [171].

Figure 5.5: Vicon, robot and laser tracker setup. A Leica laser tracker (located behind the column in the left of the image) was used as ground truth.

Both the Qualisys and Vicon systems allow the user to define objects through the identification of a pattern of markers. This allows the system to define the orientation

of objects by fitting a plane to at least three markers. By assuming the distribution of markers does not change, the object can still be tracked even if a single marker is obscured. This ability helps the user keep track of several objects simultaneously by giving each a unique label. It should be noted that patterns of markers must be distinct for the system to track more than one object.

To enable comparison of laser tracker data and Vicon positional data, a custom made target plate was designed in collaboration with Dr. Rahul Summan, as shown in Figure 5.6. The Leica laser tracker is capable of tracking a range of stainless steel Spherically Mounted Reflectors (SMR) which can be mounted to surfaces using magnetic holders. For the experiment, a 1.5" SMR was used and angled towards the tracker to give the best readings. Four tape-coated metal reflective targets with a nominal diameter of 1.5" ($\pm 100\mu m$) were ordered from Vicon and placed around the SMR, as shown in Figure 5.6. The plate, constructed from aluminium to not exceed the weight limit of the robot, was also fitted with a custom clamp allowing it to be rigidly mounted to the robot. As the coefficient of thermal expansion for aluminium is 21-24 $\mu m/m$ per degree, the temperature of the laboratory in which the experiment took place was monitored using the in-built thermometer of the laser tracker.

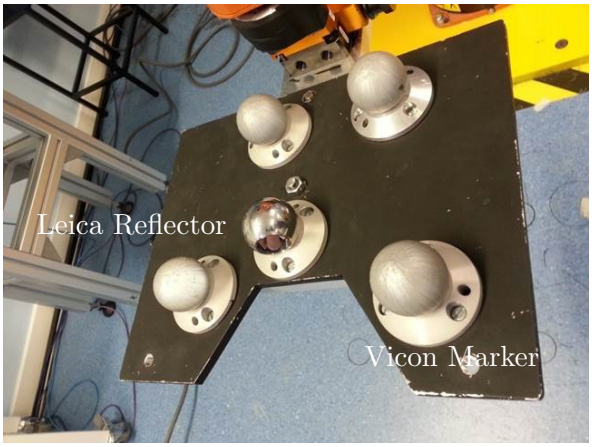

Figure 5.6: Custom designed plate attached to robot with one Leica reflector and four Vicon tracking markers.

While three markers are sufficient to define an object, four were chosen to increase accuracy and provide redundancy in case one marker was not seen by some of the

cameras. The position of each magnetic holder was measured using the laser tracker. As all five magnetic nests were measured in the same reference frame, their relative position could be calculated. This allowed the centre of the defined object in the Vicon Tracker software to coincide with the centre of the SMR. This calibration step allowed subsequent comparison of Vicon and Leica positions as both measured the same position in 3D space.

To assess the accuracy of the robot coordinates, the TCP of the robot was also defined at the centre of the SMR. This was done by creating a virtual "spike" by measuring the position of the SMR using the laser tracker. By driving the robot to the same coordinates from 4 different orientations (having moved the joints of the robot), the robot controller provides an estimate for TCP location. This implements the same procedure as described in Section 3.6.1 which, as discussed, is sensitive to operator error.

A path was designed to follow the guidelines set out by the IS0 9283:1998 standard (discussed in Section 3.5). This standard states that a diagonal plane of a cube in the most used work space should be used to determine the accuracy and repeatability of an industrial robot [92]. A robot path for the execution of this experiment was designed using the KUKA—PRC module in the Grasshopper package of the Rhinoceros 5 CAD software, shown in Figure 5.7a. The path visited five points along the diagonals of a plane and was repeated 30 times as specified by the ISO standard. To test the accuracy of the robot and Vicon system throughout a larger volume, a raster path covering approximately a 1000 mm x 600 mm x 500 mm volume was also programmed, as shown in Figure 5.7b.

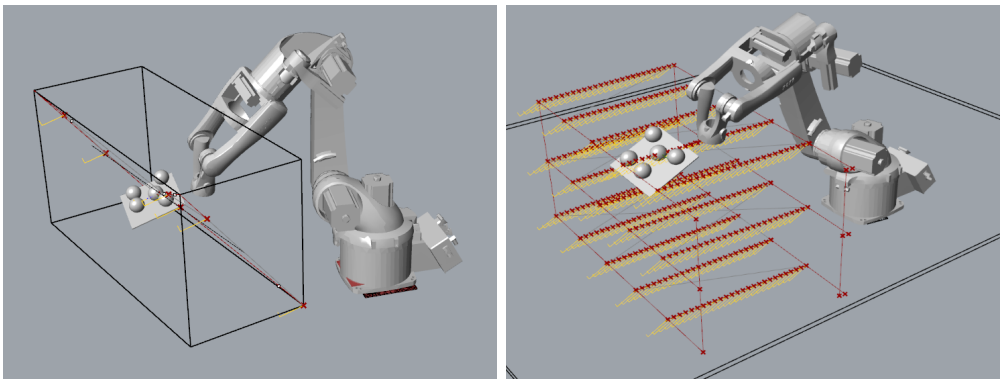

(a) (b)

Figure 5.7: Visualisation of a path designed to conform with the ISO 9283:1998 standard. (b) Raster-type path to test the accuracy over a larger volume.

To collect data from the each measurement system, existing C# code written by a colleague, Carmelo Mineo, at the University of Strathclyde was modified to include functionality allowing the robot to repeat a path a set number of times. This code provides an interface to gather data from the KUKA encoders through RSI, as well as connecting to the Vicon and Leica laser tracker software interfaces. At each measurement location, a triplet of files (one from each system) was saved as a text file.

For each point, each of the three measuring methods provided 100 readings and an average of these readings was taken as the centre. The standard deviation of the KUKA measurements was negligible, while the Leica and Vicon measurements had an average deviation of 0.03 mm to 0.05 mm in each Cartesian direction. The maximum standard deviation in Leica measurements was 1.14 mm. This large deviation was only seen once and discarded as it was most likely a spurious reflection. The maximum standard deviation seen in the Vicon measurements was 0.08 mm. Only static measurements were taken for this experiment as the robot came to a halt at each pose. Dynamic measurements were not taken at this time due to software integration issues and time constraints.

5.4.1 Results

Although the calibration plate allowed for the simultaneous measurement of the same point in 3D space by the robot, Leica and Vicon systems, the three do not share a common coordinate system. Hence, it was necessary to transform the coordinates of the Vicon and laser data into the global KUKA reference frame. Through this transformation (rotation, R and translation t) it is possible to compare the actual pose of the robot to the commanded pose, as instructed by the ISO standard. To obtain the rotation and translation matrices, the method described by Arun et al. [117] was used. This method solves the optimisation problem of aligning N points in two point sets (a and b), shown in Equation 5.2, using a least squares solution.

$$\min_{\mathbf{R},\boldsymbol{t}} \sum_{i=1}^{N} \|\boldsymbol{R}a_i + \mathbf{t} - b_i\| \tag{5.2}$$

To ensure the rigid transformation (rotation R and translation t) between coordinate systems is accurate, the error $\sigma^2 = var(\epsilon)$ can be defined as shown in Equation 5.3 where KUKA coordinates are defined by $\mathbf{P}_{kuka}$ and observed Leica positions by $\mathbf{P}_{Leica}$. The average error for the transformation was found to be 0.043 mm.

$$\begin{aligned} \epsilon &= ||\mathbf{P}_{kuka} - R \cdot \mathbf{P}_{Leica} + t||^2 \\ \sigma^2 &= \frac{1}{N}\sum_{i=1}^{N}(\epsilon_i - \hat{\epsilon})^2 \end{aligned} \tag{5.3}$$

Table 5.1 shows the result of the experiment which followed the IS0 9283:1998 path. As mentioned, each position was measured 30 times. Robot positional accuracy is defined as the average Euclidean distance between the position of the SMR as reported by the laser and the commanded pose. The Vicon positional accuracy is defined as the average Euclidean distance between the reported position of the SMR by the laser and the origin of the rigid body defined in the Vicon software. Repeatability measures the spread of reported positions about the average pose. This is defined by Equation 5.4 where RP_l is the repeatability of the $\mathrm{l}^t h$ pose, $\bar{l}$ denotes the average distance between

each measured point for $j = 1...N$ repeated movements to the same pose.

$$\begin{aligned} RP_l &= \bar{l} + 3S_l \\ S_l &= \sqrt{\frac{\sum_{j=1}^{N}(l_j - \bar{l})^2}{n-1}} \end{aligned} \tag{5.4}$$

Table 5.1: Accuracy and repeatability results for KUKA pose repeatability study. All measurements in mm.

	KUKA		**Vicon**	
Pose	Accuracy	Repeatability	Accuracy	Repeatability
1	1.346	0.047	1.510	0.145
2	0.511	0.052	0.350	0.024
3	0.665	0.045	1.720	0.016
4	0.303	0.085	2.820	0.018
5	1.152	0.119	1.960	0.040
Average	0.796	0.070	1.672	0.049

As seen in Table 5.1, following the ISO 9283 standard is not particularly informative as the five points used potentially have nothing to do with the application the robot is needed for. A slightly more useful way of describing the positional error is to look at the vector between the observed and commanded positions the robot reached at each pose. Figure 5.8 illustrates this as a vector plot for each of the 5 positions for both the KUKA and Vicon systems. It was observed that each of the commanded positions varied in the same direction for each of the 30 trials. This again demonstrates the repeatability of the robot, but could also indicate that the TCP calibration was off, causing a systematic offset.

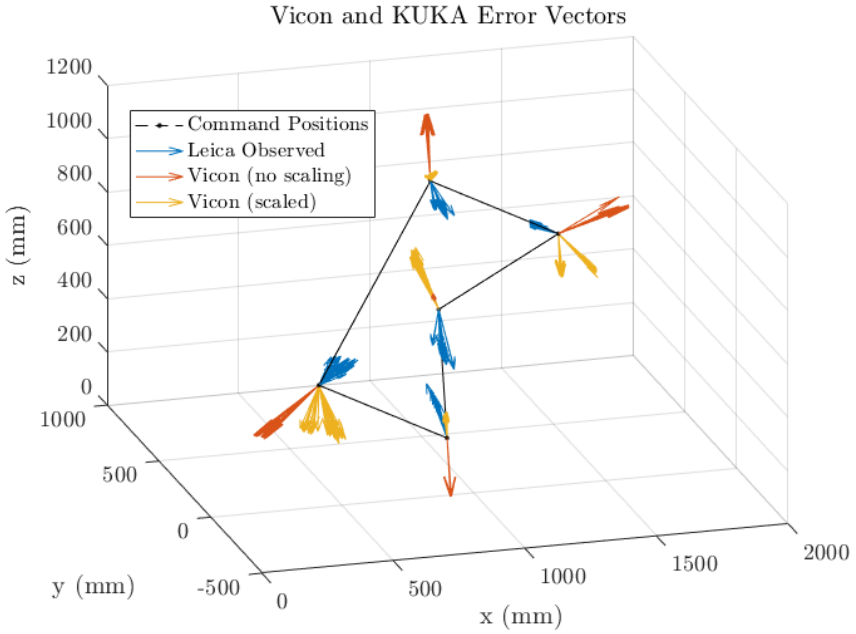

Figure 5.8: Vector plot showing deviations between the commanded KUKA positions, the position actually reached by the robot (blue). The observed positions by the Vicon system are also shown in red and yellow. It was found a scaling factor improved the Vicon data fit.

Looking at Figure 5.8, an interesting observation can be made about the Vicon deviations. For each pose, the deviations are always away from the centre. This indicates that the coordinates are scaled, such that a rigid transformation (as per Equation 5.2) is not sufficient to align the coordinate sets. To estimate the scale, s, Equation 5.5 can be used, as shown in [172].

$$\min_{\mathbf{R},\boldsymbol{t},s} \sum_{i=1}^{N} \|s \cdot \boldsymbol{R}a_i + \mathbf{t} - b_i\| \tag{5.5}$$

Including the scale factor aligns the Vicon data set better with the Leica positions, such that the average variation across the five poses and 30 trials reduces from 1.67 mm to 0.33 mm. This indicates that the Vicon system contains a scaling error, which is not uncommon when using camera systems as a scale is needed to transform pixel coordinates to real-world distances. This observations was also made by Summan *et al.* in [170]. Figure 5.8 shows (in yellow) how the Vicon errors become more randomly

oriented after the scaling factor is applied. This result indicates that the Vicon system is capable of higher accuracies, but requires a better initial calibration (or a software modification) to output more accurate results.

The initial path, designed to measure the accuracy and repeatability of the robot, only gives a limited view of the positional accuracy at 5 poses. To test if a well-defined error distribution could be found, the larger path (shown in Figure 5.7b) was also tracked by both the Vicon and laser tracker. For this experiment, the cameras were raised slightly higher on the tripods to provide coverage over a larger volume. The path was only repeated once and accuracy was defined as the Euclidean distance between the SMR's position reported by the tracker and the two other systems. The KUKA and Vicon coordinates were transformed to coincide with the laser reference frame. Figure 5.9 shows how this error is distributed across the covered volume for the KUKA and Vicon systems respectively. Table 5.2 shows the numerical results for this larger volume consisting of 387 points.

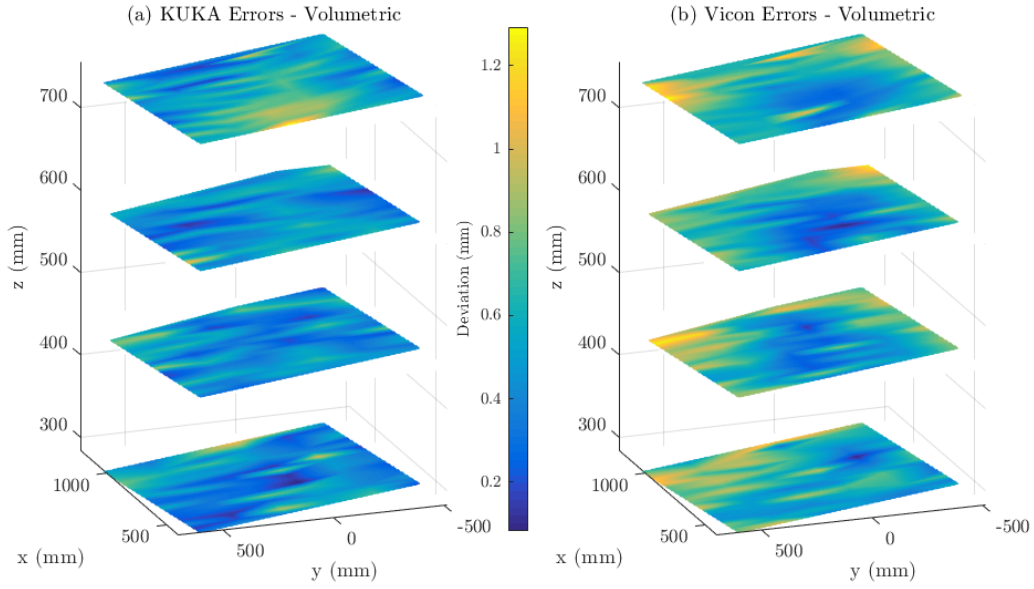

Figure 5.9: Accuracy results visualised over a larger volume for (a) KUKA encoders and (b) Vicon system. The Vicon accuracy appears to be highest in the centre where more camera viewing angles intersect.

Table 5.2: Accuracy and variation for KUKA encoders and the Vicon system for the path shown in Figure 5.9 containing 387 points. All measurements are in mm.

	Average Error	Max. Error	Min. Error	Standard Deviation
KUKA	0.480	1.315	0.077	0.206
Vicon	0.829	1.721	0.093	0.322

5.4.2 Discussion

As evident from the data presented in Table 5.1, the internal positioning system of the KUKA robot is in fact more accurate than the Vicon T160 system without a scaling factor applied to the Vicon data output. As mentioned earlier, and as demonstrated in the literature, the Vicon system's accuracy is very dependent on the camera placement as well as the distance between the tracking objects and camera. The presence of external sources of infrared radiation such as direct sunlight or reflective metal surfaces can also influence the performance of the system. The Vicon software contains features to alert the user to counter these sources of error (by highlighting areas of large reflectivity), but optimising the camera system is still a cumbersome, iterative task that is done manually.

Table 5.1 shows that the robot repeatability is slightly worse than the manufacturer's specification, which is stated as 0.04 mm [80]. As mentioned in Chapter 3, certain poses can impose higher loads on the joints due to the weight of the robot which can lead to deviations in positional accuracy. The Vicon and KUKA measurements both indicate the repeatability of each system is much better than the accuracy. The Vicon system is capable of achieving sub-mm accuracies, as seen in Pose 2 (Table 5.1), but as the variation between poses is large, it is hard to conclusively state a figure for the overall system accuracy. Potential sources of error is definition of the TCP which was calibrated manually as well as a bias in the transform obtained using Equation 5.2.

Over a larger volume the average accuracy of each system improves. The accuracy of the robot may be better as the raster-pattern used was covering a smaller volume than the ISO 9283:1998 path. It can clearly be seen in Figure 5.9a that the robot is least accurate at the top of the path, which may be because of the joint configuration needed

to reach those points. The Vicon system, on the other hand, indicates that the volume of highest accuracy is centrally located, where the lines of sight from more cameras intersect. This highlights the importance of camera placement to ensure adequate coverage to obtain better accuracy. The Vicon camera software allows the user to visualise the viewing cone of each camera such that overlap between different viewpoints can be ensured. This method is, however, time consuming as the cameras need to be recalibrated after being moved.

It is worth noting that although the average accuracy of the Vicon system is lower than that of the KUKA output, it is still possible to obtain sub-mm accuracy over a larger volume (as seen by the average value over a larger volume). So while the Vicon system might not be suitable for providing on-the-fly corrections to the robot path, it could still be used during the setup phase. If the transformation between the camera and robot coordinate systems is found, the motion capture system can provide the robot a base-to-part transformation if markers are located on the part itself. Although a motion capture system is more expensive than a single camera, removing the need from a single line of sight and being able to accurately measure several objects simultaneously are clear benefits.

The results from these experiments only indicate what accuracy is obtainable for static robot movements. In most NDE applications, the probe will have to move across the surface of a part and thus the static measurements are mainly relevant for point-type inspections. In some thermographic or X-ray inspection scenarios, the NDE payload needs to be stationary while measuring so these results may be of value in these scenarios. The use of a larger volume of points helps illustrate which parts of the robot's working envelope would be more appropriate to use if high-accuracy positioning is needed. This highlights a shortcoming of the ISO 9283:1998 standard which only considers pose accuracy for a limited number of points, which might not reflect the actual pose accuracy of the intended robot application.

5.5 Robot Dynamic Accuracy

The previous section quantified the obtainable static positional accuracy and repeatability of a typical industrial robot. For the purposes of automated NDE, however, full area measurements are needed. It is therefore important to know what dynamic positional accuracy can be achieved using off-the-shelf industrial robots. This section presents and discusses experiments conducted in the IntACom robot cell using a FARO Vantage laser tracker. Previous experiments at the University of Strathclyde (discussed in [171]) have shown that the dynamic path accuracy, according to ISO 9283:1998, can reach deviations of up to 4.464 mm when travelling at full robot speed. It was, however, found that the TCP calibration method used may have been a contribution to this high error. An improved calibration procedure was adapted to eliminate this potential source of error.

5.5.1 Improved Calibration

To be able to track the robots in the IntACom cell, a spherical reflector was mounted on a custom designed holder attached to the robot wrist, as shown in Figure 5.10a. The actual reflector sat in a magnetic nest containing a small pin which could fit into the hole on the platform, ensuring the safety of the reflector during robotic movements. To compare the outputs of the laser tracker and robot encoders, the two must both be reporting the same position in 3D space. The robot reports the position of the wrist while the laser records the position of the reflector which was located approximately 60 mm below the wrist. Hence there are two unknowns in this scenario which must be determined to be able to compare the two measurement sets.

As mentioned in Section 3.6.3, an unknown TCP can be found by solving $AX = XB$ for a number of robot poses moving about a fixed point. The relationship between robot, laser and reflector is shown in Figure 5.10b; In this case, however, two different measurement systems are used such that the equation becomes $AX = YB$ instead. This well known hand-eye calibration has been solved numerous times when both measurement systems record 6 DoF information (see Section 4.4.4 and [107]), but has not

been as widely researched when one system measures with 3 Dof and the other with 6 DoF.

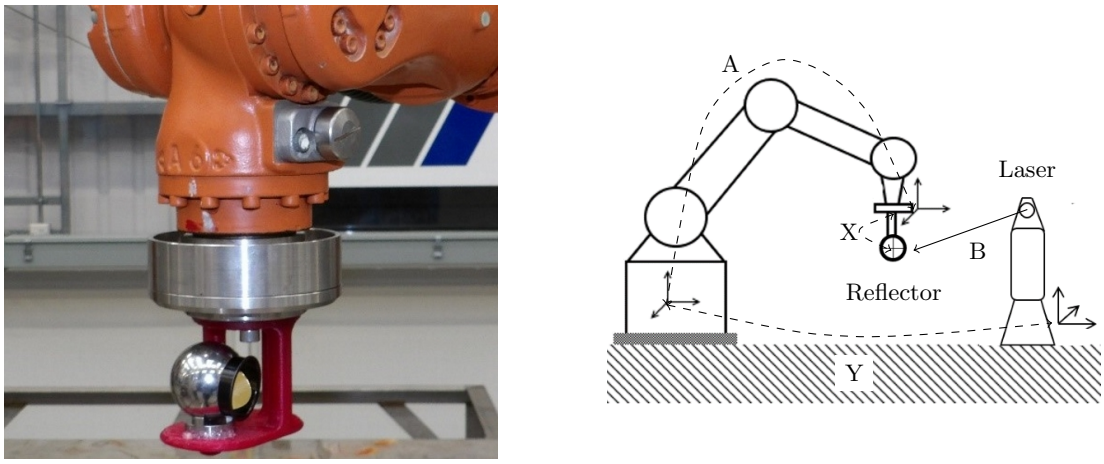

(a) (b)

Figure 5.10: (a) 3D printed nest holding a spherical reflector attached to a robotic arm. (b) Geometric relationships between robot, laser and reflector.

To overcome this limitation, a MATLAB script was written implementing the approach presented in the work published by Liao Wu and Hongliang Ren [173] which solves both unknowns simultaneously using a least-squares and iterative approach. The approach is, unlike many other calibration methods, aimed at these exact situations where one measurement tool only measures in 3 DoF. By measuring at least 8 points at different positions, the method finds the best-fit solution for both the TCP calibration as well as the relationship between the robot and laser coordinate reference frames (transformation Y in Figure 5.10b). This approach has the added advantage that the geometric transform between two or more cooperating robots can be found in an external reference frame.

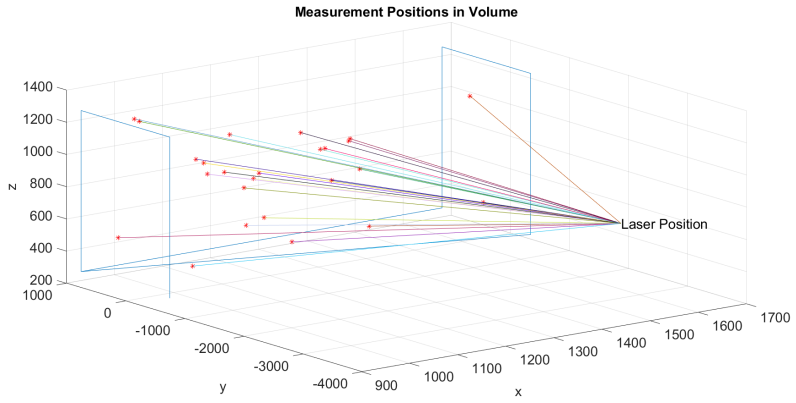

Figure 5.11: A script was written to randomly generate 20 robot poses within a specified volume. These points were then used to calibrate the TCP and laser-to-robot transform.

To facilitate data gathering, another MATLAB script was written to randomly pick twenty points inside each of the robots' working envelopes as well as a direction in which the TCP should be facing. An example of the distribution of points is shown in Figure 5.11, along with the direction vector of the TCP. The robot was then driven to each of these points and stopped for 5 seconds while the laser recorded 100 measurements at each pose and reported the average. The tool calibration method was used for both the master and slave robot with the results presented in Table 5.3.

Table 5.3: Measured tool calibrations using new method.

	Master Tool (mm)	Slave Tool (mm)
X	-28.04	-27.91
Y	0.29	0.15
Z	65.97	65.22

As seen in Table 5.3, the tool calibrations for the master and slave robots are very similar, which is to be expected as the same tool was mounted on both robots in turn. Small differences are to be expected as the tool-coupling plates on each robot are slightly different. The method does not provide any values for the direction of the tool (angles A, B and C) as the laser tracker only measures in three degrees of freedom. This means that the laser tracker is only partly suitable for tool calibrations. The average re-projection error for each of the tool calibrations was 0.599 mm. This is a

better accuracy value than what was achieved in previous experiments using a "laser-spike" manual method and is furthermore based on more than four points, increasing the confidence the user can have in the result.

5.5.2 Measurement Approach

After the tool and base calibrations were complete, a raster scan was carried out as series of linear motions at different heights. Positional data was collected using the robot's positional encoders (which have a feedback rate of 250 Hz) and the laser tracker. The FARO tracker SDK allows for continuous output with either a time or distance based frequency. Laser data was acquired based on travelled distance, such that positions were only saved when the SMR had moved 1 mm in 3D space.

The laser positional data was transformed into the KUKA reference frame such that easy comparisons were possible. This was necessary as the raster paths had been specifically designed to check for variation along a straight path in the robot's base reference system. This allows investigation of variations which only arise as a result of robot inaccuracy and are not linked to tool or base calibrations. In other words, the robot was instructed to move 800 mm along in its world x-direction with no variation in y for a series of fixed z-values. A total of six linear paths separated by 150 mm in height were tracked. This allows the analysis of each of the Cartesian coordinates in turn and such that the variation can be quantified.

Two systematic offsets were identified due to time constraints and inexperience using the FARO tracker. The reported tool definition was not set to the wrist coordinate system which resulted in a 10 mm z-offset in the KUKA feedback positions. Furthermore, all points measured by the tracker were shifted by one radius of the spherical reflector in the y-direction. This may have been caused by a switch in the FARO software settings occurring when measuring a continuous path instead of individual points. Being systematic, these two offsets were corrected for which allowed the paths to be aligned.

5.5.3 Results

When viewing the data in Figure 5.12, it is clear that the new calibration method has aligned the data sets as expected and immediate differences between the laser and robot values during the transition between raster steps is noticeable. However, the variation is visually exaggerated in the Figure as the y-axis has a 0.2 mm grid spacing while the x and z axes have 500 mm and 200 mm spacings, respectively. Figure 5.13 presents a top-down view which clearly shows the variation in the y-direction. It is worth noting that the robot encoders are very consistent throughout each raster step while the laser shows large variations. This would indicate that the robot encoders are precise, but do not accurately capture the motion of the robot. The resolution of the robot encoders is 0.01 mm.

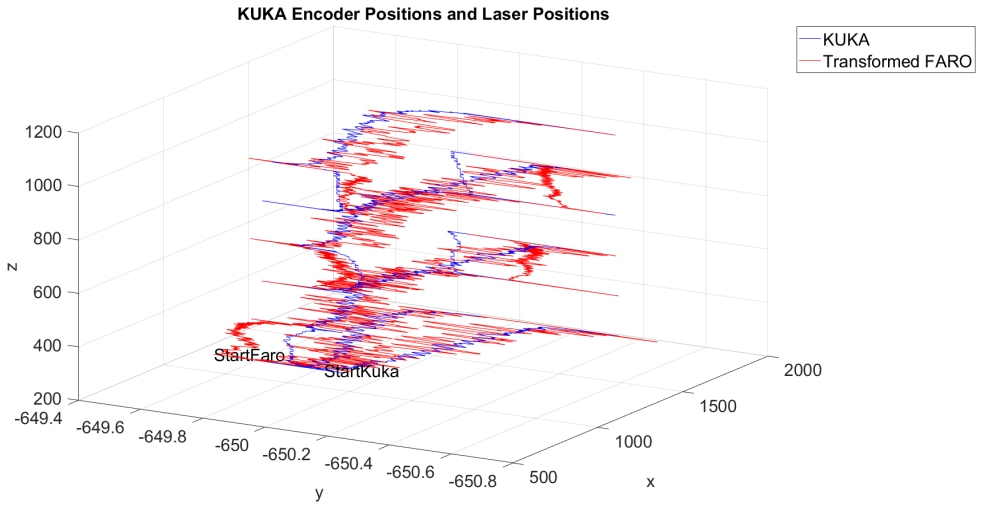

Figure 5.12: 3D view of paths showing the KUKA encoder values and the transformed laser data.

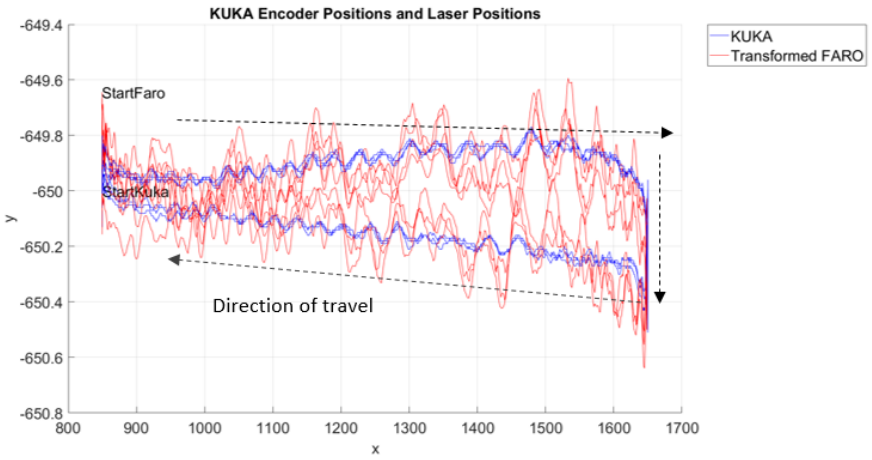

Figure 5.13: Top-down view of the data showing variation in y as raster paths are executed.

Another clear conclusion to draw from the plot in Figure 5.13 is that the robot encoders report that the robot *does not* follow a straight path. The path was planned to have no variation in the y-direction, but it is clear that the robot moves along two different paths depending on the direction of motion. This has also been documented in the work by Kubela et al. [174] and has also previously been observed during ultrasonic scanning at TWI. It is unknown why the robot is not able to correct its motion if the encoders have accurate information, but it could be because the path is executed with expected inertial forces of the tool's weight taken into consideration. Another reason could be related to the fundamental limitation of approximating a linear motion using rotational joints. As a linear path, defined in Cartesian coordinates is translated into a series of individual joints motions (which are all rotational), the kinematic model may have be designed to sacrifice precise linear motion in favour of a smoother, less precise motion which requires less strain on the robot motors. The experiment indicates that there is not a problem with the fundamental kinematic model but instead with the trajectory execution.

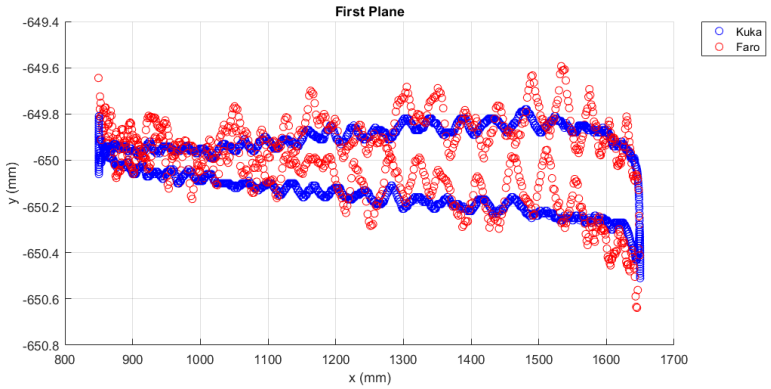

Figure 5.14: Top-down view of the first raster plane only. Error can be quantified by measuring the distance between points at set intervals along the x-axis.

To be able to quantify the dynamic accuracy of the KUKA robot, a way of finding correspondences in the two datasets was needed. As there was no temporal or spatial synchronisation during data capture, a simple one-to-one correspondence was not possible. It is however known that the two systems recorded the same motion, so the separation between points can be used to find approximate locations to compare points from the two systems. Analysing the data from each source, it is obvious that the laser and robot did not capture data in an evenly distributed manner. In fact, 75% of all robot points were recorded while the robot was not in motion. The laser, on the other hand, only recorded while the robot was in motion – with a resolution of one point per millimetre.

It was attempted to match the two datasets by comparing the total distance travelled during one plane, as shown in Figure 5.14. The ratio of robot positions to laser positions shown in Figure 5.14 is nearly unity, so naively one would assume that corresponding points in the ordered dataset would correspond to the same recorded point in physical space. However, this is not the case as the frequency of data capture was different between each dataset. By looking at the nearest robot position to each laser position we can estimate the path accuracy, though this method may introduce it's own error. To achieve a better result, the two ways of recording positional data must be synchronised in some way.

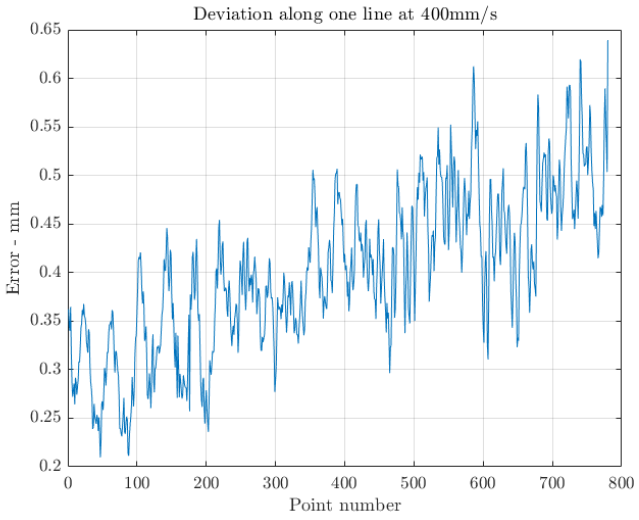

Figure 5.15: Error in positional reporting for one raster step. Note this error is underestimating the true error as no one-to-one correspondence of points is available.

An estimate of the path error is shown in Figure 5.15, but this is most likely smaller than the true error. As seen in the plot, the physical distance between the robot's actual and intended position increases along the direction of travel. This trend is reversed when the robot travels in the other direction, as shown in Figure 5.13. The average deviation was found to be 0.358 mm with a standard deviation of 0.176 mm and a maximum error of 0.792 mm. The speed of the robot while executing the motion was set to 400 mm/s with an acceleration of 100 mm/s^2. These speed and acceleration values are typical of automated NDE paths for phased array probes.

It should be noted that the above discussion is aimed at quantifying how accurately the robot can report the position of the tool being used. This is important as ultrasonic data is encoded using the robot positional feedback. Another interesting question to ask, however, is how accurately the robot relative to the commanded path. This can be readily calculated using the knowledge of the intended path and the laser data which has been transformed to the KUKA reference frame. The distance between a point, $\mathbf{x}_0$ and a line defined by two points, $\mathbf{x}_1, \mathbf{x}_2$ in 3D space [175] can be calculated using Equation 5.6. Knowing the transformation between laser and robot along with the intended path means the error for each point recorded by the laser can be found. Using

this method shows that the average robot path variation is 0.297 mm with a standard deviation of 0.160 mm.

$$d = \frac{|(\mathbf{x}_0 - \mathbf{x}_1) \times (\mathbf{x}_0 - \mathbf{x}_2)|}{|\mathbf{x}_2 - \mathbf{x}_1|} \tag{5.6}$$

Parameter Impact Analysis

Further to tracking a raster path with standard settings used during NDE inspections, a series of motion parameters in the robot's control unit were changed for the same path and then tracked. The purpose of this was to test if the speed, acceleration or weight and moment of inertia of the tool can have detrimental effects on the path accuracy. These parameters should ideally be set correctly during the forward kinematic motion planning [41]. However, in the current application, it can be difficult to accurately measure the weight of a nozzle full of water. Furthermore, as water flows through the nozzle, the mass varies throughout the inspection. It is therefore of interest to determine if a mismatch between expected and actual weight has a significant influence on how robot paths are executed.

To investigate this effect, a series of experiments were carried out. The speed, acceleration and weight parameters were changed between each experiment. The description of each trial is listed below. A figure of the path deviation, calculated using Equation 5.6 is also shown for each path.

Trial 1: Raster Path 50 mm/s

The path was executed with the tool weight set to 10 kg. The observed deviation was clearly larger towards the ends of the raster lines where the acceleration and deceleration occur. A large spike in error value is potentially due to a spurious laser reflection.

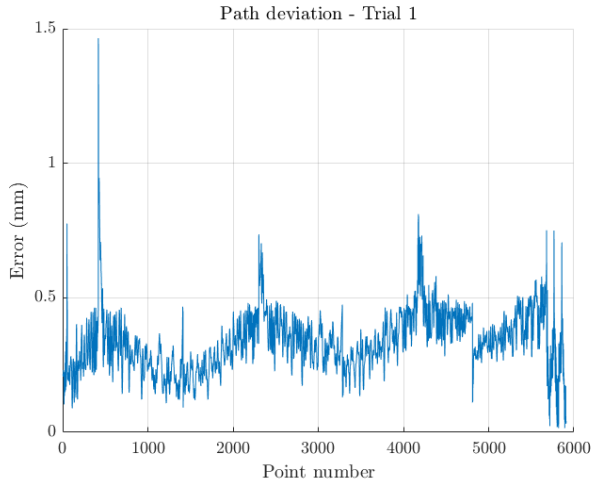

Figure 5.16: Trial 1: Path speed was set to 50 mm/s and the tool weight to 10kg.

Trial 2: Raster Path 400 mm/s

The path was executed with the tool weight set to 10 kg.

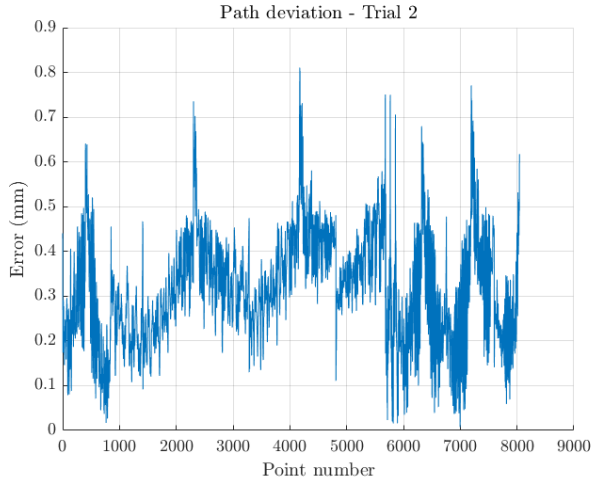

Figure 5.17: Trial 2: Path speed was set to 400 mm/s and the tool weight to 10kg.

Trial 3: Raster Path 400 mm/s with 1kg Load Settings

The path was executed with the tool weight set to 1 kg. The change in weight did not seem to have a large effect on path deviations.

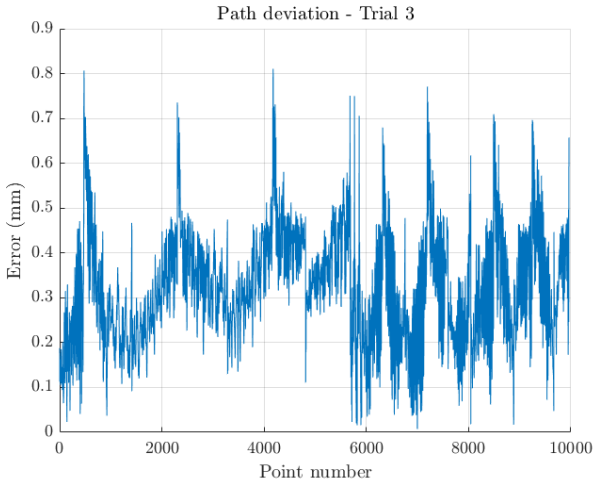

Figure 5.18: Trial 3: Path speed was set to 400 mm/s and the tool weight to 1kg.

Trial 4: Raster Path 400 mm/s with Lower Acceleration

The path was executed with the tool weight set to 10 kg and the acceleration changed to 10 mm/s^2. The change in acceleration caused the path to be slowed significantly which also affected the number of points recorded by the laser for some reason.

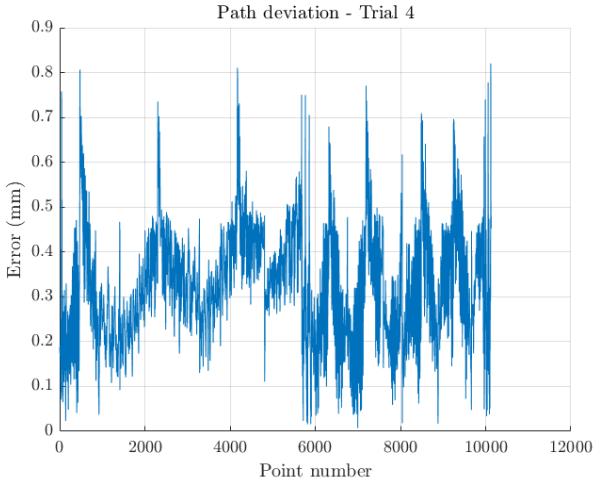

Figure 5.19: Trial 4: Path speed was set to 400 mm/s, acceleration was changed to 10 mm/s^2. The tool weight was set to 10kg.

Trial 5: External Control Raster Path at 100 mm/s

A path was designed in MATLAB which would perform a raster scan with

the same step and size as the previous paths. A script was used to convert the path into a suitable format for external control. The path was executed with the tool weight set to 10kg and the acceleration changed to 50 mm/s^2 to avoid damaging the robot. There appears to be a large variation towards the end of each raster step which could be due to an issue with the speed pattern applied to the motion between raster segments (which accelerated faster than intended).

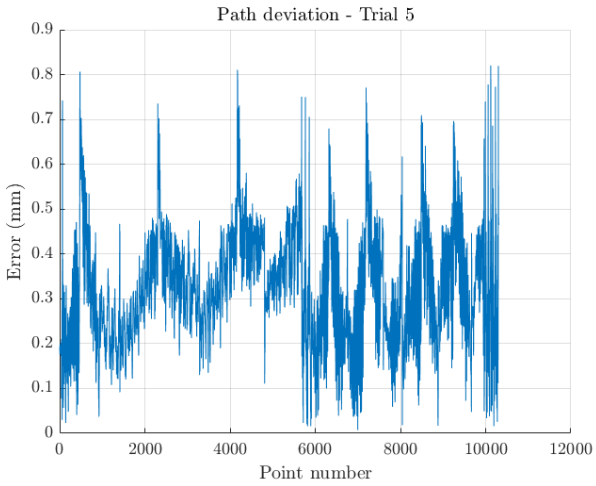

Figure 5.20: Trial 5: External control. Path speed was set to 100 mm/s, acceleration was set to 50 mm/s^2. The tool weight was set to 10kg.

Trial 6: External Control Raster Path at 250 mm/s

The path was again executed using external control with the tool weight set to 10 kg and the acceleration again set to 50 mm/s^2.

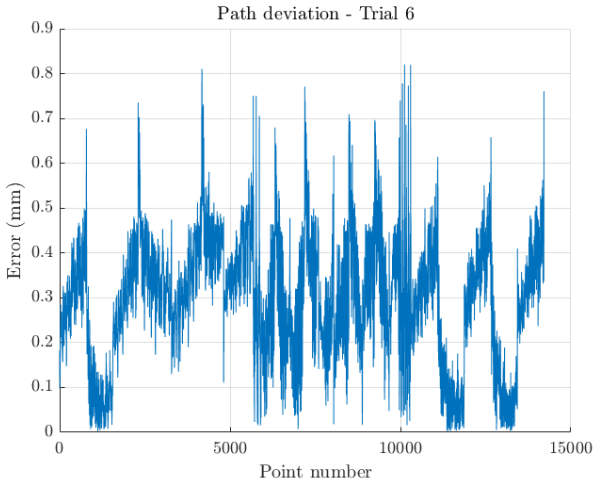

Figure 5.21: Trial 6: External contorl. Path speed was set to 250 mm/s, acceleration was set to 10 mm/s^2. The tool weight was set to 10kg.

Table 5.4 shows the variation in y and z coordinates from the intended path as measured using the laser tracker.

Table 5.4: Effects of changing acceleration and weight parameters on raster path deviations in the y and z directions. Experiments 5 and 6 were executed via external control while the rest of the experiments used the internal KUKA path interpolation mode. A list with a full description of each experiment is included below.

	Speed (mm/s)	Tool Mass	Acc. (mm/s^2)	Avg. ΔY(mm)	Max ΔY(mm)	Avg. ΔZ(mm)	Max ΔZ(mm)	σY(mm)
1	50	10kg	100	0.665	0.983	0.229	1.312	0.151
2	400	10kg	100	0.703	1.112	0.127	0.683	0.242
3	400	1kg	100	0.747	1.225	0.442	0.702	0.214
4	400	10kg	10	0.641	1.273	0.045	0.691	0.219
5	100	10kg	50	0.451	0.745	0.185	0.558	0.133
6	250	10kg	50	0.449	0.798	0.186	0.555	0.134

5.5.4 Discussion

The experiments in this section have presented a quantitative measure of how accurately a standard industrial robot (KUKA KR16-L6) can follow the type of path needed for automated NDE. It was identified in the previous section that the manual TCP calibration procedure was prone to errors and thus a new algorithm presented was

implemented for simultaneous tool and laser-to-robot reference frame calibration. This provided a much more robust way of comparing results from the two systems. The results of tracking a large raster pattern covering an area of approximately 0.6 m^2 showed the average path variation was less than 0.4 mm. This is lower than values found in previous experiments at the University of Strathclyde and may be due to both the robot type being different but also in part due to the improved calibration method.

The parameter analysis shows that a lower speed generally improves the path accuracy and that changing the tool mass in the robot controller has little significance on the path accuracy. Lowering the acceleration appears to improve positional accuracy, which is to be expected as inertial effects are minimised. Interestingly, it appears that moving the robot in external control mode increases the path accuracy by up to 60% from 0.75 mm to 0.45 mm average deviation. This suggests that the trajectory planning approach investigated earlier at TWI and the University of Strathclyde [176] could lead to an overall better performance.

As discussed in this section, it was not possible to synchronise robot encoder and laser tracker measurements due to the limited amount of time the laser tracker was available. Due to the improved calibration procedure and design of raster path, however, it is possible to compare the robot's actual position to the intended path. Using this method, the average path accuracy was found to be below 0.3 mm. As the robot positional accuracy is used to encode ultrasonic data, the accuracy which with the robot reports its position is mainly of interest.

Taking the mean of the average variations of each experiment can provide a better idea of the standard dynamic accuracy obtainable. The expected variation in path accuracy, calculated this way, is then 0.609 $\pm$ 0.182 mm. This value is similar to the results presented by Santiago Droll in [104], who used a similar KUKA robot model (the KR30-HA) and found the average path variation to be 0.87mm when travelling at 100 mm/s.

The values obtained for dynamic path accuracy can be used to set a baseline figure for the minimum deviations an operator can expect from the IntACom system without further refinements. These refinements, as shown in [104] include performing additional

robot calibrations or correcting the path using an external metrology system. The purpose of the work presented in this thesis is to demonstrate that the ultrasonic signal obtained from surface reflections can be used to improve the accuracy of automated NDE systems through novel tool and base calibration procedures.

5.6 Alignment and Signal Strength

Prior to discussing the results of the validation experiments for the tool and base calibration methods, presented in Chapter 3, the acceptable range of probe-to-part distances must be considered. As discussed throughout this thesis, the IntACom system uses water jet nozzles to provide ultrasonic coupling between probe and part. Understanding how the signal is affected by a longer water travel path is relevant for two main reasons. Firstly, when performing a tool calibration, the actual TCP that is calculated sits on the face of the probe, as discussed in Section 3.6.2. The TCP must be moved along the z-axis of the probe such that it is located outside the physical nozzle to avoid collisions. Knowing how far to place the TCP outside the nozzle has to be measured empirically unless accurate simulations of the ultrasonic beam's interaction with the nozzle can be modelled.

Secondly, the ultrasonic base alignment procedure presented in Section 3.7 can only take place if the part is already in an approximate position such that echoes from the surface can be measured. Hence it must first be determined how closely the part must be to its intended position. Methods for automatically determining the position and orientation of parts in a robot cell prior to inspection are addressed in Chapter 4. Limits on the needed accuracy for the external part positioning system need to be established.

To quantify these two parameters, the current section presents two different experiments. The first experiment measures the quality of an ultrasonic signal as the probe is moved away from a surface. The second experiment investigates how the reflected surface signal deteriorates as the probe is misaligned with a surface.

5.6.1 Signal and Distance to Part

An experiment was carried out to investigate how the ultrasonic signal is affected when the distance between a probe and part is increased. This experiment consisted of having one robot move a receiving probe along the back surface of an acrylic component while a second robot placed the transmitting probe on the opposite side of the acrylic at various distances from the surface. As the receiving probe moves further away from the part, the water jet distance increases which causes more attenuation of the signal. Furthermore, turbulence and gravitational pull on the water column can also lead to a lower signal being observed. Figure 5.22 illustrates the experimental setup.

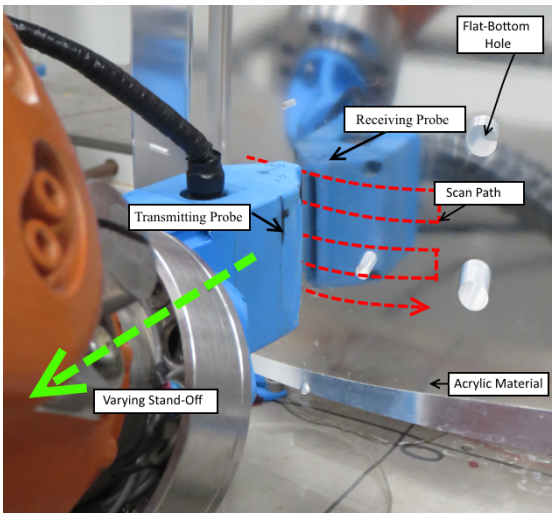

Figure 5.22: Raster path carried out by twin cooperating robots. The first robot rasters along the surface and the second robot moves back in increments of 2-5 mm.

Ultrasonic data was gathered using two water jet nozzles, shown in Figure 5.22, each containing a 64 element phased array probe with a central frequency of 5 MHz. The pitch of each array probe was 0.6 mm and a sub-aperture of 15 elements was used such that the overall active area was an approximately 9 mm x 9 mm square, the idea being to produce a nearly symmetrical Gaussian beam profile. The beam was electronically incremented with a step of 1.2 mm and data was recorded at every 0.6 mm in the scanning direction.

The receiving nozzle was kept at 5 mm $\pm$0.6 mm offset from the surface of the

component which was verified by monitoring the front wall signal in a separate pulse-echo scan. The second robot was moved backwards in steps of 5 mm while staying perpendicular to the surface of the component. A KRL file was written to produce this path behaviour while recording data through the IntACom software. The results of the experiment are show in Figure 5.23 which illustrates the loss of signal as the transmitting probe is moved farther from the surface. The effects of turbulence as the water path increases are clearly visible from offsets larger than 25 mm.

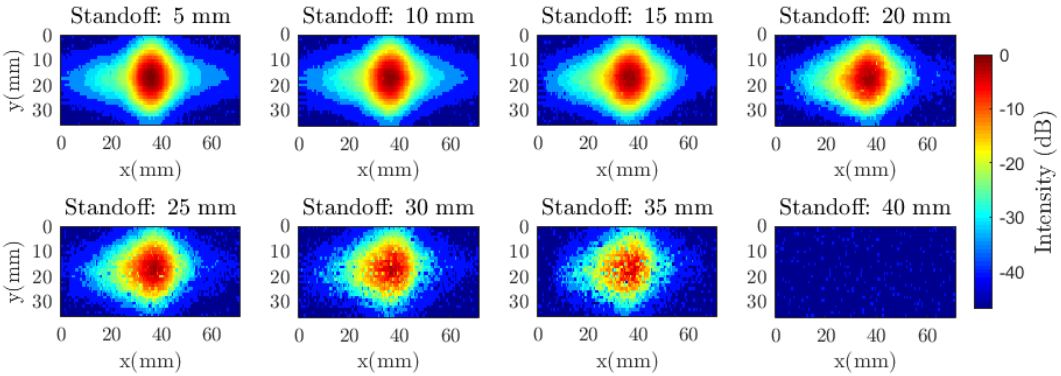

Figure 5.23: Ultrasonic C-Scan data for nozzle stand-off distances varying from 5 mm to 40 mm. It can clearly be seen how the signal quality diminishes with increasing stand-off.

Generating a contour plot for each stand-off position allows for the analysis of a region of less than 2 dB of intensity loss as a function of stand-off distance. A 2dB signal loss from an 80% of screen height front-wall echo was chosen as an appropriate measure for areas of good alignment, as this is common practice in industry. The raster experiment was also replicated using a 2mm stand-off increment instead of 5mm. Figure 5.24 shows how this area drops as the transmitting probe is moved away from the surface. The peak at close to 7mm is caused by the natural focal point of the beam hitting the surface which was confirmed through CIVA simulations.

The resulting graph shows that an area of nearly 30 mm^2 of good signal intensity is received through the acrylic component for stand-off distances between 5 mm and 10 mm. Assuming a circular focal spot, this area provides roughly 3 mm of misalignment tolerances in the x and y directions wherein the signal will fluctuate less than 2dB.

This value is of course not tolerable when measuring the position of defects, but it does indicate that small misalignments arising from robot positional inaccuracies are not likely to have a large influence on signal intensity. This relatively large area is most likely due to the size of the ultrasonic beam spreading as the pulse travels through the material.

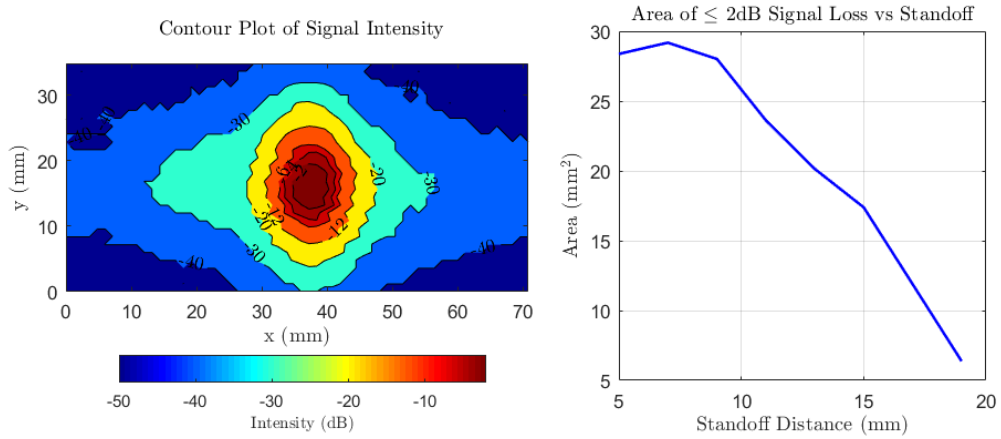

Figure 5.24: Left: Contour plot of signal intensity at 5mm stand-off. Right: Area of signal intensity being greater than 2dB as a function of stand-off distance.

5.6.2 Signal and Probe Orientation

To minimize influences of refraction at the component interface and effects of anisotropy, the ultrasonic probe must be kept perpendicular to the surface throughout the inspection. An experiment was conducted to determine the tolerances for probe alignment. A linear 5 MHz phased array ultrasonic probe containing 64 elements, with a pitch of 0.6 mm was rotated above the flat surface of a calibration block. A central sub aperture of 15 elements was used to match the probe elevation. To avoid collision, the water jet nozzle containing the probe is kept approximately 10 mm above the surface of the part. The probe was rotated about the x and y directions (angles B and C shown in Figure 5.25a) and the amplitude of the front wall reflection was recorded at each angle increment. The rotation around the z axis (angle A) does not lead to a change in signal amplitude as the probe stays normal to the surface.

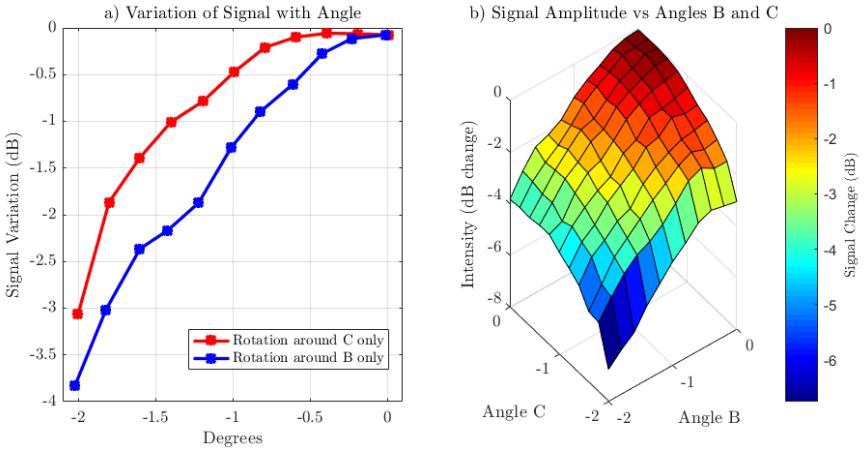

Figure 5.25: (a) Variation of signal amplitude with angles B and C. (b) Surface representation of the amplitude drop as a function of probe misalignment along two axes

The result of incrementing the probe from perpendicular (zero degrees) to -2 degrees in 0.2 degree steps is shown in Figure 5.25a. Positive rotations are not included due to the symmetry of the setup. It is evident that the rotation around B has a larger detrimental effect on the signal amplitude, which is expected from the probe geometry. The decrease in amplitude is even greater when the probe is misaligned in both axes, as shown in Figure 5.25b. A limit of 2 dB variation is tolerable in most inspection scenarios. As evident from Figure 5.25b, this means the probe must be kept normal to the surface within ± 1 degree. To obtain this level of accuracy it is crucial that the TCP of the ultrasonic probe and the part position and orientation are correctly calibrated. Although it is important that the probe follows the surface at a specified distance to avoid collisions, this off-set distance has less effect on signal amplitude due to the relatively low attenuation of ultrasound in water.

5.6.3 Discussion

The two experiments presented in this section have empirically measured two key values which must be taken into account when trying to optimise the ultrasonic data gathered using the IntACom system. Firstly, it was shown in Figure 5.24 that the TCP should ideally be placed 7 mm in front of the water jet nozzle to obtain the largest amount of ultrasonic energy entering the part. The second experiment clearly indicated that the

probe should be kept within 1 degree of the normal direction to the surface to maintain an adequate amount of ultrasonic energy entering the part.

It should be noted that these two experiments were conducted for a single type of probe and nozzle configuration. The methodology, however, lends itself to quickly be used for other probes such that limits on nozzle stand-off and orientation tolerances can be measured. Another results from the current experiments is that the automated part positioning system should aim at locating the part within 7 mm and within 1 degree of the intended pose to allow the ultrasonic alignment method to work as intended.

5.7 Obtainable Accuracy using the IntACom System

As the inherent accuracy of a standard industrial robot has been established, the task of optimising the calibrations between wrist and tool as well as robot and part needs to be addressed. Presenting new methods for this task was the main topic of Chapter 3 and this section provides experimental validation of those methods. In particular, it will be shown how the automated tool calibration methods provides more accurate defect indications and how the ultrasonic alignment method gives a better signal-to-noise ratio, even on complex geometries.

This section presents two sets of validation trials. In the first experiment, an accurately machined calibration block containing four flat-bottomed holes was scanned using the IntACom system after calibrating the TCP of the robot using the automated method presented in Section 3.6. The reported positions of the holes were then compared to their nominal positions. The same block was then scanned after calibrating the tool manually. Comparing the two scans shows that the automated calibration procedure improves absolute accuracy and signal intensity.

In the second set of trials, a carbon-fibre aerospace component with complex surface geometry was mounted in brackets within the IntACom cell. Although a CAD model of the component was available, there were no well-defined points to establish the robot-to-part base calibration. A manual calibration therefore resulted in a path with a large variation in signal intensity. The ultrasonic alignment procedure presented in Section

3.7 was used to refine the manual base calibration. Using the results of the method, the base reference was updated and its signal variation was minimised.

5.7.1 Automated Tool Calibration Procedure

A series of ultrasonic scans were carried out in the IntACom cell to determine the accuracy of the developed calibration method presented in Section 3.6. As discussed previously, the calibration method uses a metal sphere to define a unique point in 3D space. An indication of accuracy is given by calculating the variation in this calculated 3D point for each robot pose. To determine the accuracy with respect to an external reference frame, an aluminium calibration block measuring 40.0 mm x 130.0 mm x 60.0 mm containing four flat bottomed holes of nominal diameter 5 mm at depths of 4.98 mm, 9.87 mm, 24.88 mm and 49.90 mm was ultrasonically imaged. A cross-sectional view of the block is shown in Figure 5.26a. By defining the robot base coordinate system at the corner of the reference block, it is possible to compare nominal and observed hole positions.

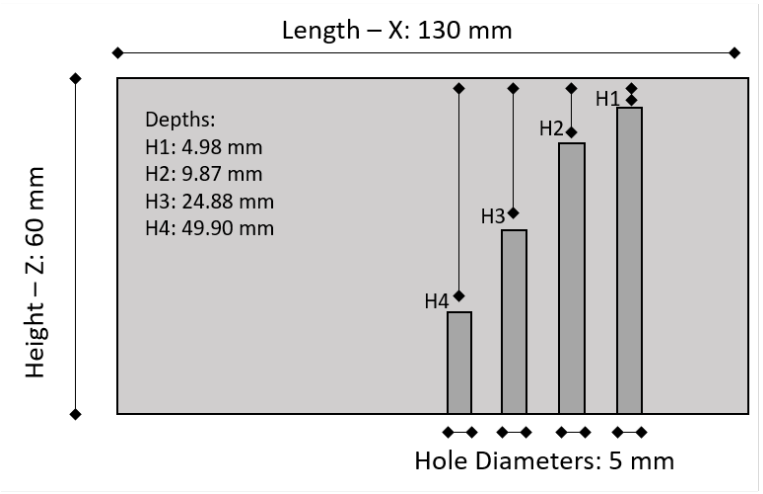

(a)

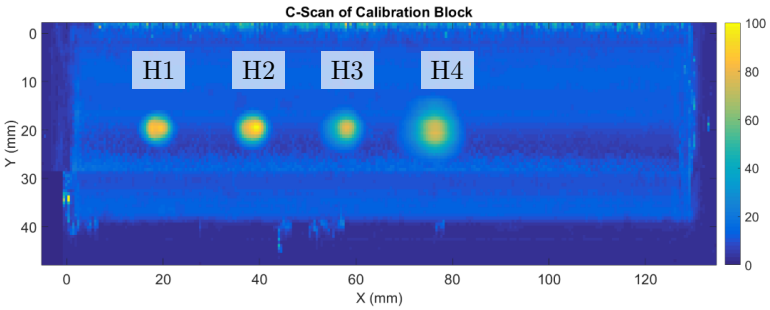

(b)

Figure 5.26: (a) C-scan of calibration block clearly showing the four flat bottomed holes with overlaid labels. Colour represents signal amplitude. (b) CAD model of the reference block

The observed centre of each hole will be a function of both the robot's accuracy as well as the TCP and base calibrations. The position of the reference block with respect to the robot base was accurately measured using a laser profiler attached to the robot (as discussed in Section 6.2). To test if the developed TCP calibration method improved the accuracy of the observed indications, a scan of the block was conducted using both a manual TCP calibration and the new, automated TCP calibration. The scans were performed five times with the tool's y direction travelling in the direction of the longest axis of the block as well as five times perpendicular to the longest axis. The ultrasonic settings were defined in the IntACom software to obtain an 80% full-screen height reflection from the shallowest hole. These settings were kept unchanged for each

scan.

Data was gathered using the IntACom acquisition software using a 5 MHz, 64 element phased array probe from Vermon. The pitch of the probe was 0.6 mm and a sub-aperture of 16 elements was used to create a uniform beam profile. As shown in Figure 5.26a, the deepest defect is not resolved as clearly as the shallower defect. To identify the centre of the hole position, code was written in MATLAB to analyse in the binary IntACom data, measure the time of flight using nominal speeds of sound and measure the hole centre positions automatically.

To calculate the centre positions, the ultrasonic data was segmented into four 20 mm x 20 mm regions around each hole. As the signal from the hole reflection is larger than the back-wall signal, three thresholds of -3 dB, -6 dB and -9 dB from the largest amplitude in each region were chosen. This segmented the data into A-scans coming from reflections from each indication and back-wall reflections. The average x and y coordinates of these segmented regions were used to define the centre. The average time of flight value inside the material was used to define the depth of the hole. It was observed that the -6 dB threshold gave the most accurate indication of hole position, so this method was used for all subsequent analyses.

The presented MATLAB script processed each of the 10 scans for the probe moving longitudinally and perpendicularly along the calibration block. Figure 5.26 shows a C-scan image from a single scan where hole indications can clearly be seen. Table 5.5 summarises the average distance between the actual and observed locations of the 4 holes for each tool calibration method. As seen in both the Figure and Table, the last indication (H4) is the least accurate due to its large depth (nearly 50 mm) which causes a larger beam spread. Looking at the three components of each error, a slightly higher error in x was observed which is most likely due to the pitch of the ultrasonic elements. The average deviation for the new method is 0.88 mm and 3.29 mm for the manual method, giving an absolute improvement in accuracy of 2.42 mm or 73.43%.

Table 5.5: Results from using old and new tool calibration method

	New Method				**Manual Method**			
	H1	H2	H3	H4	H1	H2	H3	H4
Mean Error (mm)	0.632	0.435	0.749	1.683	3.420	2.824	2.881	2.927
Standard Dev. (mm)	0.249	0.050	0.226	0.2644	0.855	0.610	0.607	0.461
Mean Signal (%)	83	94	76	75	62	66	49	51

Figure 5.27 shows a graphical presentation of the errors for the two TCP calibration methods. The graph clearly shows the overall trend of deviations for the 4 holes being similar which indicates that the manual method introduces a systematic off-set to the calibration. This could potentially be because of inaccuracies in the 3D printed artefact (see Figure 3.10a), meaning that a more accurate manual calibration is possible. However, as discussed earlier, the problem with manual calibrations is that the tip of the nozzle is used as a reference and not the probe face itself. By using the ultrasonic signal to calibrate the TCP position, this source of uncertainty is removed, leading to a more accurately encoded scan as seen in Figure 5.27.

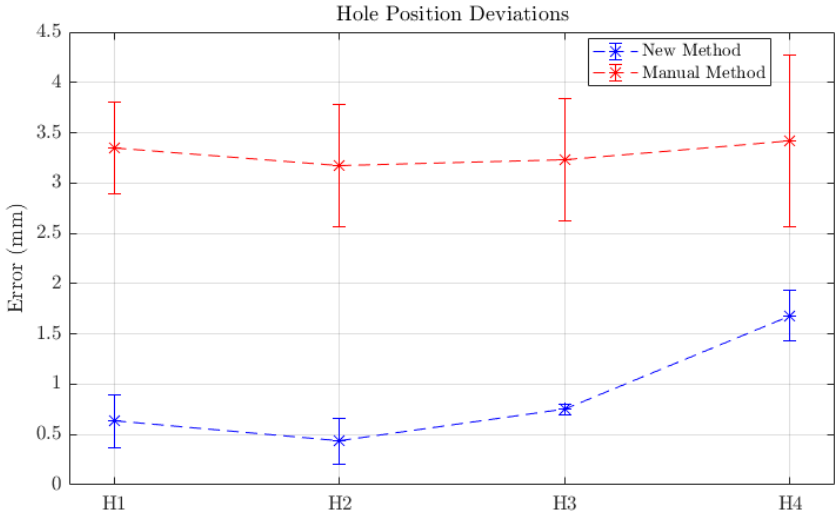

Figure 5.27: Graph showing positional error of each reported hole position using the new and manual TCP calibration methods.

Besides improving the overall absolute accuracy of inspections, Figure 5.27 also illustrates that the deviation between each scan is smaller. Looking at the data in

more detail, this was observed to be due to the manual calibration suffering larger misalignments when rotated 90 degrees for the perpendicular scans. This means that the new method not only improves the absolute accuracy with which the centre of the probe is defined as the TCP, but also the orientation of the probe with respect to the robot wrist. The robot is able to rotate the tool more accurately about the centre point when using the new TCP method.

The improved definition of orientation is also clearly expressed in the improved signal amplitude, as seen in Table 5.5. On average, the received signal was 30.6% lower when using the manual method, equivalent to 3.19 dB. As shown in Section 5.6.2, maintaining good alignment between probe and surface is key to obtaining a good signal response from defects. The lower amplitude seen when using the manual method thus clearly indicates that this TCP calibration method does not provide an accurate probe orientation.

5.7.2 Complex Geometry Ultrasonic Alignment

The ultrasonic alignment method presented in Section 3.7 demonstrated that the reflected surface echo can be used to improve the robot-to-part calibration for a simple, flat geometry. Although the initial experiment provided good results, the simple geometry of the sample is not representative of most parts which are inspected using the IntACom system. To test the algorithm further the developed method was used on a complex aero-foil shape provided by an industrial partner of TWI. A CAD model of the blade was supplied by the partner and an inspection path covering the entirety of the actual blade section (not the root), was generated using the commercial off-line path planning software FASTSURF. When seeing the actual blade it became clear that accurately calibrating the part pose (the base frame in the KUKA controller) would be a challenging task. This was due to a mismatch between the CAD and actual geometry as well as the lack of any clear features such as sharp corners on the physical blade. An attempt to calibrate the component pose was achieved by manually driving a metal spike mounted as the robot's TCP to three chosen reference points and recording their positions.

The component, measuring roughly 1.5 m in height and shown as a CAD model in Figure 5.28, was scanned using a 1 inch single-crystal 1 MHz probe with a raster step of 2 mm. The single element probe was mounted inside a custom-designed water jet nozzle which was manually calibrated such that the TCP of the tool was approximately 10 mm from the surface of the part. If aligned properly, an ideal scan should keep the probe within ±1.5 mm of the reference offset. The quality of the ultrasonic scan was assessed based on the stand-off distance and surface echo signal amplitude. The IntACom software was used to define ultrasonic settings such that a reference front-wall signal height was obtained on a flat section of the component. Inspections covering the whole sample were then determined to be adequate if the variation of the front-wall echo remained within 2 dB throughout the scan. The figure of 2 dB and 1.5 mm were agreed upon with the industrial partner.

An alignment scan path was designed using off-line path planning to cover roughly 25% of the total surface of the component. This area was chosen as it would provide enough information to use the new alignment method while keeping the overall inspection time short. A complete scan of the part was completed to provide a reference data set of data quality when using manual base alignment methods. The alignment scan was then carried out and corrections for the robot-to-part calibration were calculated using the method described in Section 3.7. After updating the robot base reference system, the entire component was scanned again.

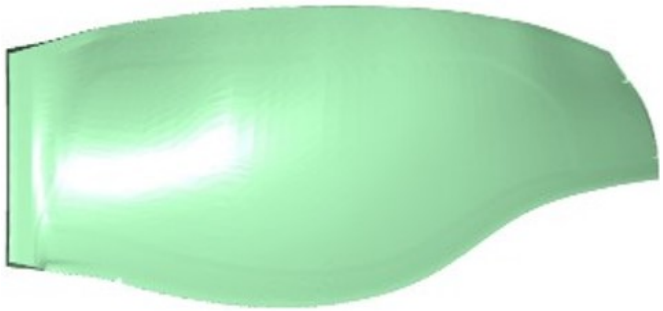

Figure 5.28: A representative CAD model of a composite blade showing the complex geometry of the component. The actual blade was found to deviate slightly from its CAD model.

Results

In order to evaluate how well the robot followed the surface of the component, the variation in time of flight (ToF) and in amplitude of the front wall signal was calculated for 95% of all the points (excluding obvious outliers where the robot moved off the surface to ensure full coverage). To visualise the effects of misalignment, each inspection point in 3D space was coloured blue if the amplitude was within the 2 dB range and red if outside. As the component geometry was not entirely consistent with the CAD geometry, the complete scan path included over-spray points at the edges. These points made up about 5% of the total number of inspection points and were removed for the variation analysis.

Figure 5.29 shows the variations when the part was scanned using manual alignment. It can clearly be seen in Figure 5.29a that there are large patches of high variation in amplitude, indicating the tool is not well-aligned to the surface. The variation from the reference front-wall signal is between -7 dB and 3.3 dB. As seen in Figure 5.29b, a range of 8 mm was seen in the ToF data, with a standard deviation of 2.06 mm. Again, this shows that the tool is not following the front wall particularly well which can also lead to inaccurate positional encoding of any defects.

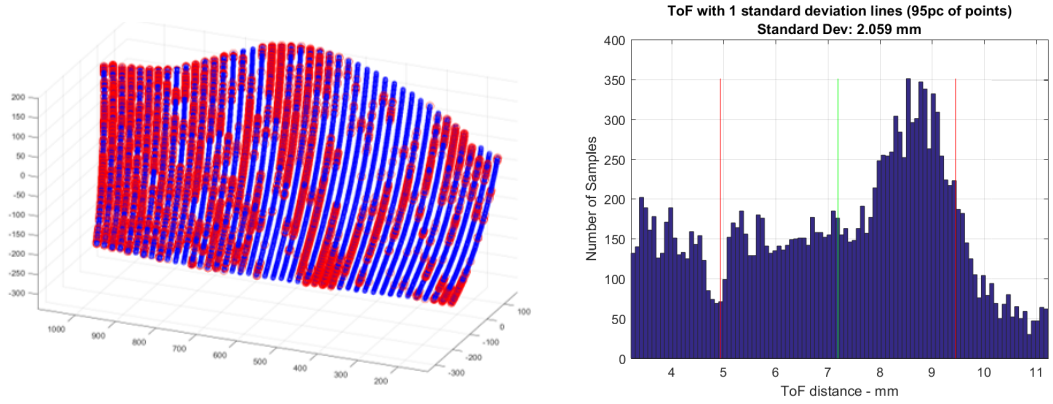

(a) (b)

Figure 5.29: Manual base calibration: (a) Variation in front-wall echo amplitude. Blue points are within the 2dB threshold while red points fall outside. (b) Variation in ToF distance between probe and component surface.

After the developed alignment method was used, the variation in signal amplitude

was between 2.8 dB and -2.7 dB, though as evident from Figure 5.30, most of the inspection points are within the desired 2 dB range. Interestingly, a pattern forms when the outliers are plotted in 3D space as shown in Figure 5.30a. This could mean that the drop in signal amplitude is not caused by misalignment but by surface properties which absorbs more of the sound. As seen in the lower right hand corner, a small patch of outlier points are clustered together. It was understood that the corner of the blade often shows the largest deviation from CAD, which means that the off-line generated path would not be valid in this area, even if the pose of the component was calibrated perfectly.

The variation range in ToF values was 3.7 mm with a standard deviation of 0.76 mm as shown in Figure 5.30b. From the bar graph, it can be seen that most points are clustered centrally about the reference stand-off value, as expected from a well aligned scan. Ideally the distance between the probe and the surface should be constant throughout the entire inspection but factors such as the inherent inaccuracy of the robot and gravitational effects on the water jet will influence the ToF. The data, however, does fall within the limit of ±2 mm about the reference stand-off distance.

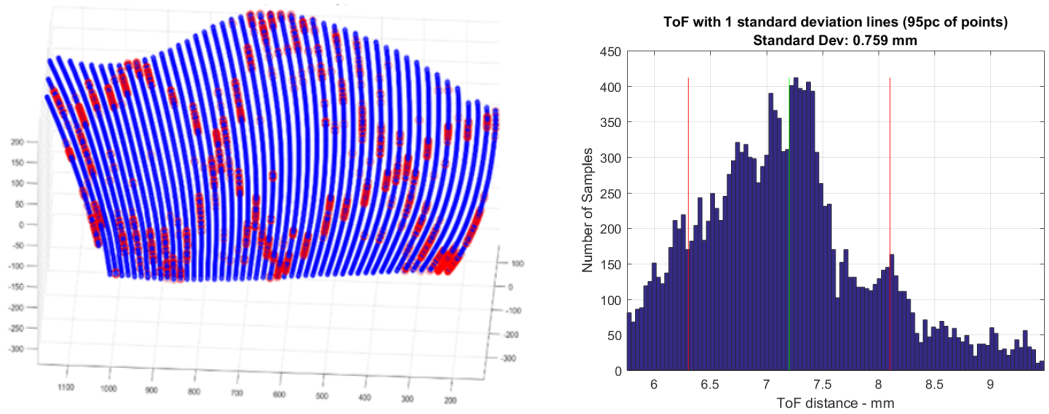

(a) (b)

Figure 5.30: Base calibrated using ultrasonic alignment: (a) Variation in front-wall echo amplitude. Blue points are within the 2 dB threshold while red points fall outside. (b) Variation in ToF distance between probe and component surface.

The results of using the manual method and the new ultrasonic method for calibrating the position of the base reference system are summarised in Table 5.6. The

initial manual calibration resulted in large areas where the variation in amplitude of the front wall echo exceeded the 2 dB threshold, as shown in Figure 5.29a. As a measure of improvement, the average change in signal amplitude was 32.2% better after the alignment had taken place. The improvement in the range of ToF distances is 4.3 mm with a 61% reduction in variation from 2.059 mm to 0.759 mm. These results underline the usefulness of the developed technique and its ability to provide high-accuracy alignment between the robot and part reference frames, even when scanning component which deviate from their CAD models.

Table 5.6: Difference in front-wall amplitude and time of flight variation about the reference value using the standard manual method and the new ultrasonic base alignment method.

Calibration Method	Ref. Signal (% FSH)	Max Amp (dB)	Min Amp (dB)	Max ToF (mm)	Min ToF (mm)
Manual	80	3.085	-7.002	4.042	-3.954
Ultrasonic	55	2.573	-2.673	2.300	-1.419

5.8 Discussion and Further Work

This chapter has presented a number of validation experiments which help to quantify the obtainable positional accuracy of an automated NDT system using standard industrial robots. Experiments to quantify both the static and dynamic positional accuracies which can be achieved using standard KUKA industrial robots showed that sub-millimetre accuracy is achievable at speeds which are typically used in NDE scenarios. There are, however, several caveats to note about positional accuracy.

Firstly, there are no clear patterns to where one can expect a higher positional accuracy except that it appears the robot is less accurate at the extremes of the working envelope. One would intuitively expect such a behaviour as the mass of the robot is further away from the central axis which is bolted to the floor. Further work should investigate if the variation of positional accuracy can be explained by, for example, a simple mass distribution model. If a model could describe where the robot is likely to be less accurate then perhaps the robot paths could be designed to avoid those parts

of the working envelope.

Secondly, the dynamic accuracy experiments demonstrated an interesting periodic behaviour whereby the robot deviated from the path in a quasi-sinusoidal manner. The direction of travel appeared to have an influence on how well the robot could follow a path, which may mean that a mechanical phenomena related to the rotary joints and electric motors could be responsible. This could potentially be gear backlash or torque ripple effects. Another explanation might be the way the robot approximately translates a straight path in Cartesian space to a set of rotary joint angle movements. Further work should try to investigate the causes of the seen behaviour and whether or not it can be modelled and corrected for.

Even with the observed behaviours in both static and dynamic cases, the positional accuracy of a standard industrial robot has been shown to be sufficient for most NDE applications. This is a key contribution of this thesis as few research groups have published literature establishing a fundamental baseline for the obtainable accuracy of such systems in the context of NDE. This could be because most integrators believe that a highly repeatable robot will be sufficient for most applications, but this thesis argues that a solid understanding of robot accuracy is needed if off-line path planning is used to design inspection paths. By conducting tracking experiments with highly accurate optical measurement systems, this work shows that an operator can expect the robot to follow a path to within half a millimetre with variations of up to roughly one mm. Attempting to size defects smaller than one millimetre should therefore not be attempted as the operator would not have the needed confidence in the system's positional feedback.

An approach to improving the positional accuracy of robotic systems relies on using external measurement systems to provide feedback to the robot. This approach has been demonstrated in the work by Nikon [103] and New River Kinematics [104]. The Accubot system from Fill [57] instead relies on an extremely accurate calibration of the robot at installation along with a completely independent computed numerical control system, thereby bypassing the internal kinematic control of the robot. While both approaches give higher accuracies, the expense and integration challenges associated

with these make them less attractive to industry. This thesis argues that better sensor integration with standard industrial robots creates a cheaper and easier to maintain automated NDE system.

Once the baseline accuracy of the robots themselves was established, this chapter presented validation experiments for the methods presented in Chapter 3. The aim of the developed calibration procedures is to try and reach the same accuracy as the robot such that the limiting factor in an automated inspection system is the hardware itself and not an inaccurate calibration. It was shown that the tool and base calibration methods were able to substantially improve the absolute accuracy of the IntACom system when compared to manual methods. This is another key result of this thesis as it illustrates how ultrasonic sensor integration can be improved, both in terms of automation but also in terms of system performance.

The automated TCP calibration procedure based on the ultrasonic signal was shown to bring the absolute accuracy of the robot equipped with a phased array probe and squirter nozzle to the level of the robot itself, at just below one millimetre. The base alignment procedure demonstrated an improvement of over 30% in the quality of the received signal and a 61% reduction in stand-off variation. The experiments presented in Section 5.6 showed that as long as an automated part placement system can locate a part to within 5mm and 1 degree of accuracy, the developed base alignment method should be able to refine the measurements to optimise the inspection. It was shown that motion capture system would be able to provide this level of accuracy while being more versatile and robust than a single camera solution.

Further work should look at quantifying the dynamic accuracy of a motion capture system to investigate if these could be used to provide external positional feedback for a moving robot. It would also be worthwhile to investigate the developed TCP calibration method further to see if this can be automated further, perhaps using a permanently installed calibration artefact (like the brass sphere) in an inspection cell. The positional accuracy of the TCP method, shown in Section 5.7.1, should also be quantified throughout a larger volume, perhaps using different ultrasonic targets.

Finally, further work should aim at improving the ultrasonic part alignment pro-

cedure, potentially by using a full CAD model instead of the control points as well as being able to adapt the path in case of deviations from CAD. As shown in the following chapter, a laser profiler integrated with a robot provides a method for overcoming situations where the part and CAD model do not match.

Chapter 6

Industrial Projects

The RCNDE Engineering Doctorate programme is specifically targeted at research projects which are intended to transfer lower level TRL research ideas from the academic environment to industry. The current thesis is no exception. Previous collaboration between TWI and the University of Strathclyde [1] had proven that industrial robots can be used for automated NDT through a series of demonstration experiments. This earlier work focussed on path planning approaches, ultrasonic data encoding and preliminary system integration. This thesis expands upon the earlier work in a number of ways, in particular with regards to assessing robot accuracy and increasing automation levels. This chapter presents three industrial use-cases which are direct outcomes of the work presented in this thesis. These outcomes have helped develop and mature the IntACom programme to a point where, at the time of writing, it is regularly being used for single client and collaborative industrial projects.

The chapter is structured in the following manner: first, an overview of twin-robot through transmission inspections is presented along with the contributions of this thesis to their improvement in the IntACom system. Next, the integration of a laser profile sensor with the IntACom robot cell for surface measurements is presented. Finally, a Graphical User Interface (GUI), developed in MATLAB, for the calibration methods developed throughout this thesis is presented. Although this thesis has contributed to a number of other industrial projects, these three case studies are chosen as they highlight and exemplify how measurement solutions have improved overall system performance.

6.1 Adaptive Stand-off Through-Transmission Technique

As described in Section 2.3.2, the Through Transmission (TT) technique uses two probes, requires access to both sides of a component and is useful for inspecting thick components. Obtaining good alignment between transmitting and receiving probes is the main obstacle to using this technique. In robotised inspection systems, this obstacle can be overcome by either using two cooperating robots or by having a single robot manipulate a yoke which can hold both probes. Using a yoke is less demanding from a programming and interfacing perspective, but may not be feasible for large components and assumes a uniform component thickness. The IntACom robot cell uses two cooperating robots as this provides more flexibility, especially when inspecting parts with varying thickness as will be shown in this section. This section also highlights the importance of calibrating the position of each robot with respect to the other.

6.1.1 Cooperative Robot Motion for NDT

A number of applications require the coordinated motion of several robots to complete a complicated task. A typical example of such instances is when the weight of a component is too large for a single robot to lift but can be achieved by using two robots cooperatively. Other examples include procedures like fitting, gluing and painting where two robots need to work together within a limited time to achieve a task. In robotised inspection scenarios, robot cooperation is mainly used for through transmission inspections but could also be used for complex computed tomography paths using an X-ray source and detector.

Many modern composite components require ultrasonic inspection in both PE and TT mode, due to the high attenuation of CFRP and to detect porosity more effectively. A small number of research papers have been published which discuss through transmission inspections with industrial robots. Lu et al. [177] highlight the importance of probe alignment during robotic through transmission inspections and demonstrate a method for obtaining synchronised motion with two industrial robots using a custom built server application. The authors highlight the importance of maintaining control

of the orientation of the receiving probe so that Snell's law can be used to obtain the optimal signal.

Schwabe et al. [178] discuss the use of water jet nozzles for ultrasonic through transmission inspections using both gantry-type inspection systems and industrial robots. The authors mention that achieving less than 2 dB of signal fluctuation is challenging with standard industrial robots. Nieto et al. [179] discuss the challenges of designing a water jet nozzle for phased array ultrasonic through transmission inspections. The authors demonstrate a design which enables robotic through-transmission inspections with a 20 mm diameter water jet at 750 mm/s.

Following the work of [180], cooperative motion between industrial robots can be categorised accordingly: concurrent cooperation (synchronised start points), coupled synchronous cooperation (master-slave relationship) and combined synchronous operation, illustrated in Figure 6.1. The three types of motion have been explored for through-transmission inspections. As seen in Figure 6.1, the three types of motion can be used to inspect geometries of increasing complexity.

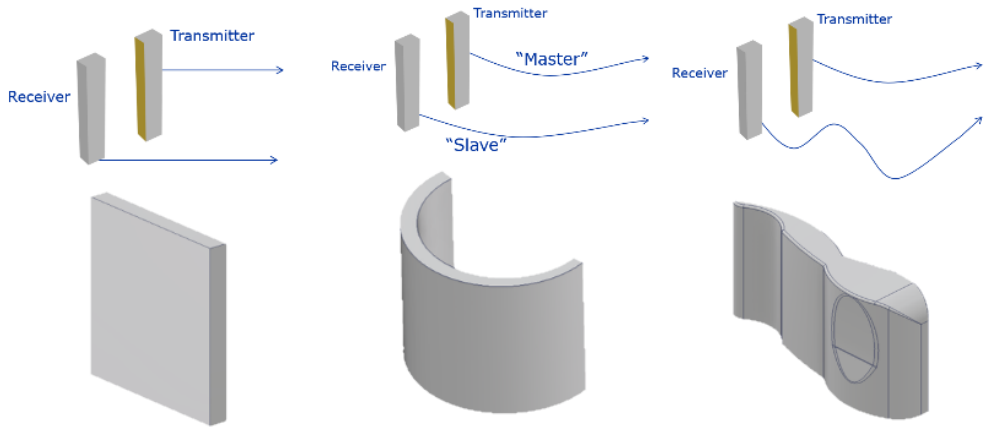

(a) (b) (c)

Figure 6.1: Cooperative motion applied to three geometry type inspections. (a) Flat surfaces can be inspected with concurrent cooperation. (b) Simple curves require coupled synchronous cooperation. (c) Complex curves require combine synchronous operation.

If flat surfaces need to be inspected it may be sufficient to simply synchronise the start and end points of a motion and ensure the velocities of the two robots are equal.

This type of motion is referred to as concurrent cooperation. Coupled synchronous cooperation is useful if the curvature of the part is regular and simple as the two robots keep the distance and orientation of the probes with respect to each other constant throughout the inspection. Combined synchronous operation is needed when the geometry is complex, for example when the component has sections of varying thickness as seen in Figure 6.1c. If water jets are used, small changes in thickness can be overcome using coupled synchronous cooperation as long as the change is smaller than the water column such that collisions are avoided.

The RoboTeam technology package from KUKA [163] allows all three types of motions to be executed. Synchronised `Wait` functions allow concurrent cooperation. Coupled synchronous cooperation is referred to as geometric coupling, or "GeoLinking", and is used extensively in the IntACom cell. Combined synchronous operation can be achieved using motion synchronisation as shown in Figure 6.2. This type of cooperation forces the robots to complete paths in the same motion time. A key industrial contribution of this thesis was to develop a method to inspect parts using combined synchronous operation. The motivation for doing so is to enable the rapid inspection of components with varying thickness using twin cooperating robots.

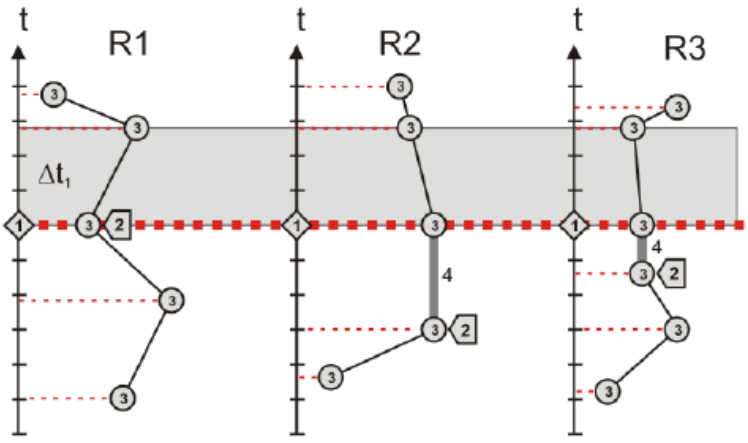

Figure 6.2: Combined synchronous operation behaviour. Irrespective of their previous motions, each of the three robots (R1-R3) are forced to complete the next motion within the same time interval [163].

6.1.2 Path Planning for Cooperative Paths

Most commercially available path planning software packages only contain limited functionality for robot cooperation and synchronisation. The available software at TWI, (FASTSURF from CENIT), allows some program synchronisation, notably synchronising the start, end-points and velocities through the use of wait signals. A TT method using this approach was covered in the PhD work of Carmelo Mineo [20], but was found to provide insufficient alignment across larger components. The main reason for this was the difficulty in matching the velocities of both robots as they covered large distances.

Instead, the GeoLink synchronisation method was used for TT inspections. To achieve a GeoLinked cooperative motion, a single path is planned for the "master" robot which is then mimicked by the "slave" robot. For this approach to work, the two robots must first be instructed to establish an initial alignment position which will define the off-set and orientation of the two probes throughout the rest of the path. Finding the alignment point and optimising the ultrasonic signal is done manually which is both time consuming and prone to operator error. The need for a single path is the main advantage of this technique as it requires less set-up time. The drawbacks of the method are the need for accurate initial positioning and being unable to inspect components with non-parallel surfaces.

As mentioned, the GeoLink approach assumes that the change in thickness of a component is relatively small, which is not always the case. Several components were encountered which exhibited such complex geometries that combined synchronous operation was needed, which was not available in any commercial path planning software solution. To achieve this type of motion, two paths are defined, (one for each robot), which are then completed simultaneously. Developing a path planning approach to overcome this problem is the main industrial contribution of this section.

6.1.3 Twin Robot Calibration

Before addressing the challenges of combined synchronous operation, it is worthwhile discussing how GeoLinked motions can be made more accurate. As discussed in Section 3.2.2, a robot moves with respect to its own internal reference frame. To perform coordinated motions, two or more robots need to establish a common reference frame and a way of exchanging positional information. This is necessary to be able to calculate the position of each end effector with respect to the others as well as synchronise trajectory speeds. A single base frame is typically defined on the part to be inspected or the reference frame of a "master" robot is defined as an origin. To share a single robotic base coordinate system, the transformation between robot bases must be found.

Section 5.5.1 illustrated how the reference frame of several robots can be measured with an external optical system. The external system can either be a laser tracker, or a motion capture system as discussed in Section 5.8. The position and orientation of each robot reference frame can be calculated through the calibration methods discussed in Section 5.5.1. This information then allows for the calculation of the relative transformation between two robots (referred to as master and slave), using Equation 6.1, wherein T is a 4x4 homogeneous matrix. The term external in Equation 6.1 refers to any external measurement system capable of determining position and orientation.

$$T_{Slave}^{Master} = T_{External}^{Master} T_{Slave}^{External} \tag{6.1}$$

Knowing the inter-robot calibration is an important quantity as it is the basis on which GeoLinked motion is carried out in through-transmission inspections. Within the KUKA control system, a structured variable saves this information which is consequently used to calculate slave robot poses during synchronised motion. Section 5.6 demonstrated how even a small misalignment between two probes causes a reduction in received signal amplitude. Misalignment between the two probes affects ultrasonic results both as a consequence of the probes not being opposite each other, but also due to their orientation being misaligned.

The measurements taken with the FARO laser tracker (described in Section 5.5)

were used to update the slave-to-master calibration on the slave robot control unit. The maximum translational difference from the existing calibration was found to be 0.64 mm and maximum angular deviation 0.035 degrees. These small deviations can arise as a result of environmental changes such as thermal expansion of the robot pillars. Figure 6.3 shows a 3D colour coded view of how positions in the robot envelope are affected by a 2 mm and 0.5 degrees deviation in the robot-to-robot calibration. Each coloured dot shows the distance between where that point would have been projected to using the old calibration and where it is using the new calibration. As can be seen in the figure, even sub-mm deviations can cause errors of up to several millimetres and the error is not constant throughout the volume. Hence, it is important to try to calibrate within the space most often used by both robots.

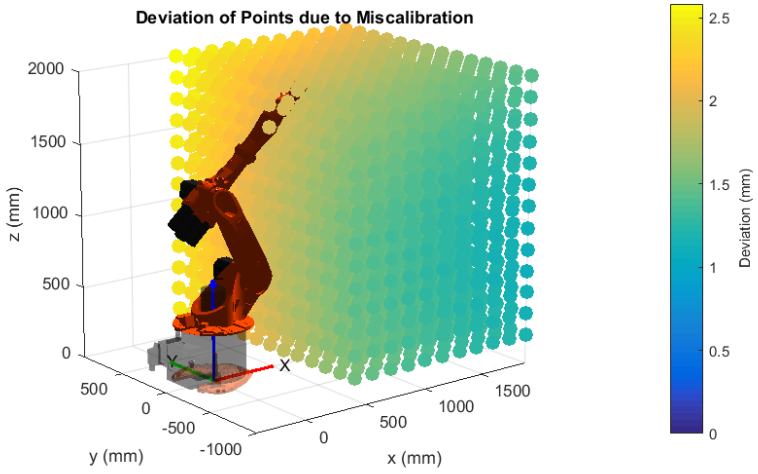

Figure 6.3: Colourised 3D view of the effect of sub-mm twin robot misalignment in the robot working envelope.

Previously the robot-to-robot calibration has been carried out manually using a metal spike mounted on each robot which is both time consuming and prone to operator variability and error. Using an optical system along with the calibration techniques described in Section 5.5.1 is both quick (the procedure took around 15 minutes to complete), more accurate and could be performed as a weekly check or at the beginning of a through-transmission experiment. TWI have invested in a motion capture system

and a potential use of the system is to be able to provide this calibration on a regular basis. Future work should investigate to what extent the ultrasonic signal is affected over a large volume by any errors in the robot-to-robot calibration.

Knowing the geometrical relationship between the reference frames of cooperating robots is useful, not only for improving path accuracy, but also for automating the GeoLink path planning. As mentioned earlier, although the actual inspection path can be planned offline, the point at which the two robots initiate cooperative motion needs to be manually defined. By knowing the transform between two robots and accurately calibrating the TCP of each phased array probe (using the method presented in Section 3.6), the position of one probe can be found in the reference frame of the other robot using Equation 6.1. A suitable point can thus be found in one robot's reference frame and the other robot can be instructed to move to the same point. As the two tools need to face one another, one tool orientation will need to be rotated around the active axis of the probe, which is a simple geometric transformation.

6.1.4 Concurrent Synchronised Operation

To achieve synchronised motion with two robots scanning each side of a component, each robot must follow different surfaces in the same time interval. The KUKA add-on package, RoboTeam, allows this type of motion to be programmed using the `SyncCmd()` command. An illustration of how this motion synchronisation is achieved was illustrated in Figure 6.2. An approach was developed during this EngD programme to generate two synchronised paths by projecting the position of a point on one side a CAD model through the model to find a corresponding pose for the other robot.

It must be stressed that it was not within the scope of the EngD project to develop a complete path planning technique, such as the one developed during the preceding PhD work [20]. The approach taken in this work relies on other software to generate an initial scan path for a single robot. This can be achieved using commercial solutions as discussed in Section 3.3.1. The initial path only needs to be generated on one side of the component over a suitable inspection area. The approach is to then use these points to create a corresponding path for the second robot on the opposite surface.

This both saves time and fits with the existing industrial work flow for generating TT paths using the GeoLink method. The following sections describe the approach in more detail while Appendix B provides further details on the specific programming needed to construct KRL scripts using this type of synchronised motion.

Projecting Through a CAD Model

As the principle of through-transmission inspections relies on one transducer receiving the sound emitted from another, the position of one robot will be constrained by the location of the other. A path was designed in FASTSURF on an acrylic component with a complex geometry, shown in Figure 6.4. The path was designed with a point spacing of 150 mm, between which the robot would interpolate a linear motion. Each point was projected directly through the CAD model to generate a second path on the opposing surface.

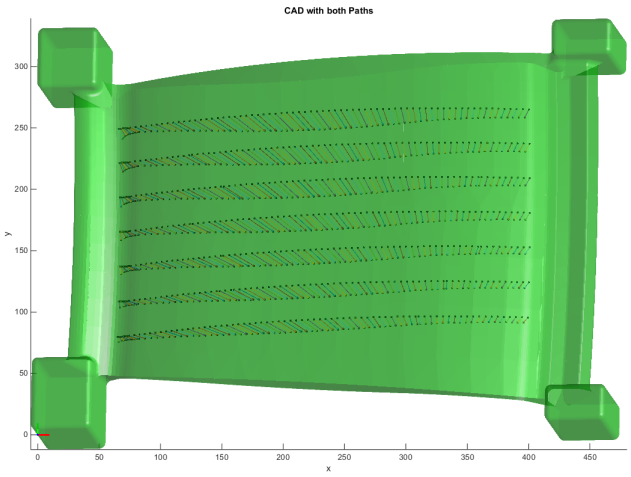

Figure 6.4: Paths projected through a component. A path is defined on one side of the component using off-line path planning software and traced through the CAD model.

The developed algorithm uses the *xyz* position and direction of the probe at each point to define a vector. Each point is assumed to be on or near the surface of the CAD model, so only intersections being further than 1 mm away (unless the component is very thin) are considered. As the CAD model consists of a list of triangles, the objective

becomes finding which triangle is opposite each point. The surface normal of that triangle can then determine the orientation of the opposing probe. The intersection of a ray with a triangle in three dimensions is a well known problem from computer graphics and an efficient method was described by Tomas Möller and Ben Trumbore [181]. A ray is described by an origin O, direction, D and distance along the line, t as shown in Equation 6.2.

$$R(t) = O + tD \tag{6.2}$$

A point at (u, v) on a triangle is given by Equation 6.3, where (u, v) are barycentric coordinates which must satisfy $u \geq 0, v \geq 0$ and $u + v \leq 1$. V_0, V_1 and V_2 are the three triangle vertices. To find the intersection between the ray and triangle the two equations must be equal, $R(t) = T(u, v)$, as shown in Equation 6.4.

$$T(u, v) = (1 - u - v)V_0 + uV_1 + vV_2 \tag{6.3}$$

$$O + tD = (1 - u - v)V_0 + uV_1 + vV_2 \tag{6.4}$$

Rearranging Equation 6.4 leads to Equation 6.5 which is a linear system of equations that can be solved using MATLAB's `linsolve` function.

$$\begin{bmatrix} -D & V_1 - V_0 & V_2 - V_0 \end{bmatrix} \begin{bmatrix} t \\ u \\ v \end{bmatrix} = O - V_0 \tag{6.5}$$

Once an intersection is found which is within a set limit (more than 1 mm away but less than 20 mm, for example), the distance to that triangle is then recorded (t in Equation 6.2). This allows the algorithm to work with both thin components as well as components which curve back unto themselves (which would result in two intersections). The second robot position is then given by moving the recorded distance, t, along the initial vector. Through this method, it is possible to create a series of points on the opposite surface, as shown in Figure 6.4. Another script was then written which outputs

the two paths in the KUKA Robot Language (KRL) using unique identifiers to achieve synchronisation using the `SyncCmd()` command. An example of the output is shown in Appendix B, Figure B.1.

CAD Partitioning

To project a vector through a component, the CAD data (in form of an STL file) and the path must be in the same coordinate system. This is achieved by identifying a coordinate system which is fixed to the part when designing the path and then using the relative transform between the origin of the CAD model and the new origin to transform the coordinates of the vertices of each triangle. The same origin can then be used in the actual inspection to define the local robot base coordinate system. As an STL file typically consists of tens of thousands of triangles, an iterative search of all triangles to find the point of intersection is very slow, even on a modern laptop (Intel i5 processor, 4GB RAM).

To alleviate this problem an algorithm was written which partitions CAD data into cubic volumes until each box contains less than a specified number of triangles. This is done through what is commonly known as the "Octree" method [182], which is used to recursively divide a 3D space into eight smaller spaces, as shown in Figure 6.5. The partitioning of the CAD model only takes a few seconds for a model with more than 100,000 vertices. When defining a new partition, the algorithm also finds the Cartesian coordinates of the centres of each box, which are later used in a search algorithm. By requiring that three vertices (forming a triangle) are kept together as the CAD is partitioned, the algorithm avoids splitting up triangles between adjacent boxes.

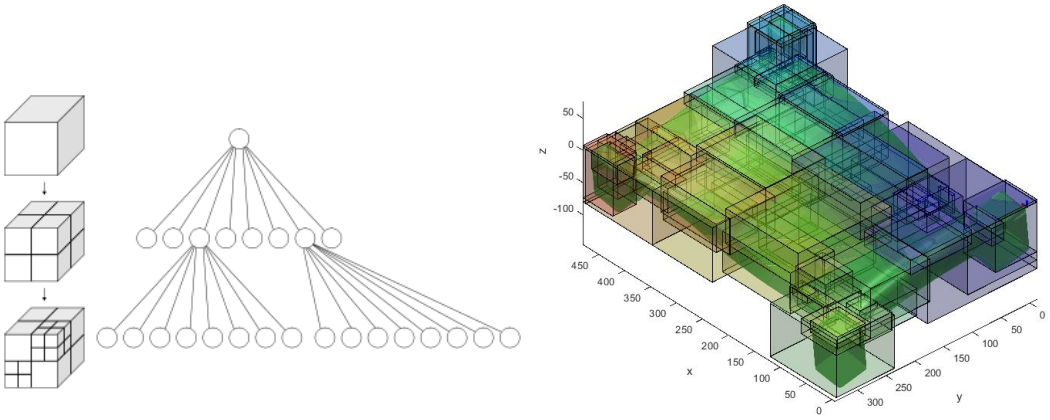

(a) (b)

Figure 6.5: (a) The Octree partitioning principle (from [182]). (b) Split CAD model of an acrylic component showing the result of partitioning. Each large box contains less than 10,000 vertices and get progressively smaller. This partitioning scheme allows for faster searching.

To find the point of intersection, a search algorithm was written. The search algorithm identifies which box the current coordinate is in and then uses a nearest neighbour algorithm to rank the other boxes based on the Cartesian distance from the origin box. The next step is to then check for any intersections within the same box that the origin is in and, failing this, move along to the next box in the list and look through all triangles in this box and so on. This is then iterated until a intersection is found which is further away than a defined minimum distance.

By partitioning the CAD data into a relatively low number of boxes containing fewer triangles, the projection algorithm becomes much more effective. Compared to the iterative method, the current implementation showed approximately a ten-fold increase in performance (from 2 hours to 10 minutes) for a representative CAD model. The search algorithm also has some error handling functionality, such as in the case where a vector does not intersect the back wall. In this case the search algorithm alerts the user to the missing intersection and uses the last known "good" distance as an interpolation. Before the user can export any paths they are given a full 3D view of all points on both sides of the CAD model to ensure that the correct paths have been constructed for both robots.

Synchronised Robot Orientation

The Cartesian coordinates of the slave robot are given once the intersection point has been found using the described method. The coordinate is given in the local coordinate frame of the component which means that the slave robot will need to be taught exactly the same base reference frame as the master robot. Defining the orientation of the slave robot at each point is however not as straightforward. One solution is to simply mimic the orientation of the master, but mirrored, such that the active area of the probe points towards the master robot (assuming both tools are defined similarly with respect to the robot flange). Another solution is to orientate the robot perpendicular to the other surface. These two concepts are illustrated in Figure 6.6.

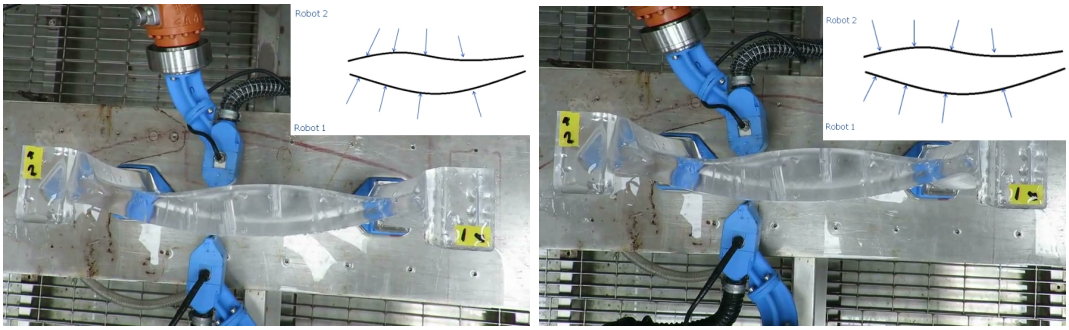

(a) (b)

Figure 6.6: (a) Robot 2 mirrors the position of Robot 1 at each point on a complicated surface profile. (b) Robot 2 is perpendicular to the surface at each point along the path.

Simply mirroring the orientation is not optimised in terms of the sound path. Ideally, knowledge of the material would be used to calculate the refraction of the sound beam as it travels through the component (using Snell's law or modelling software solutions such as PZFlex [183] or CIVA [184]). In the developed method, the sound is modelled as a ray moving through the component which is sufficient for most industrial applications. As the results of Section 5.6.1 show, if the probes are quite close (approximately 7mm) from the surface of the component, the water will act as a wave-guide, countering the effects of refraction. This is most likely only the case for surfaces which are not highly curved and may not be applicable in all cases.

To obtain the best possible ultrasonic scan, the slave location should be at the point where the highest intensity of the refracted beam will exit the component. Snell's law in three dimensions could be used for isotropic materials but a more complex modelling solution would be needed for anisotropic, composite materials. As mentioned previously, the surface normal at each intersection triangle can be used to determine the orientation of the second robot's TCP such that it is perpendicular to the surface at each point on the path. It was however found that this can lead to irregular robot movements if the CAD model mesh is made up of large triangles. As mentioned, it is also possible to simply mirror the first robot's orientation as shown in Figure 6.6.

Using this approach for finding the normal of the surface also allows for the calculation of Snell's law at each point. However, it was found that the calculated normals were not optimal during the execution of the robot path as the sound at times refracted backwards, causing the second robot to quickly change direction of travel. It may be possible to overcome these irregular behaviours by considering the average of a number of surface normals near each inspection point and filtering the angles to achieve a smoother motion.

6.1.5 Experimental Trials

Two samples were scanned using the developed approach to cooperative robot motion for through transmission inspections. The first sample was a custom designed acrylic component with non-parallel surfaces, shown in Figure 6.4. The sample contained a number of flat-bottomed holes simulating defects. The second sample was a CFRP turbine blade, previously discussed in Section 5.7.2. The blade contained large changes in thickness due to its aerodynamic shape, changing from around 1mm thick at the edges to over 20 mm thick in the middle section. The blade also contained internal reference reflectors on each side of the component.

The acrylic sample was scanned using a 5 MHz, 64 element ultrasonic phased array probe. An initial path was generated on one side of the component and two paths were generated, as illustrated in Figure 6.6. One path had the second (slave) robot mirroring the orientation of the master robot while the other path used the surface normals to

determine the orientation of the slave robot. The purpose of generating these two paths was to investigate which approach would be better for TT inspections. As the surfaces exhibit significant curvature, the sound beam is expected to refract significantly and knowing how to orientate the two robots becomes critical for obtaining a good signal.

Figure 6.7 shows the difference in signal intensity for the two paths. The plots clearly show that mirroring the orientation of the master robot provides a much improved ultrasonic scan with higher signal intensity and clearly defined reflectors. Although not intuitively obvious, it was found that the KUKA synchronisation technique does not offer good control over the angular velocities of each robot. This leads to misalignments when the curvature of the component changes rapidly. By mirroring the orientation, the angular velocities of each robot were equal which led to better overall alignment. However, as one knows the sound beam will refract, this constraint does not allow for the optimal placement of the receiving transducer. This is to some extent overcome by the water jet acts like a wave-guide for the ultrasound.

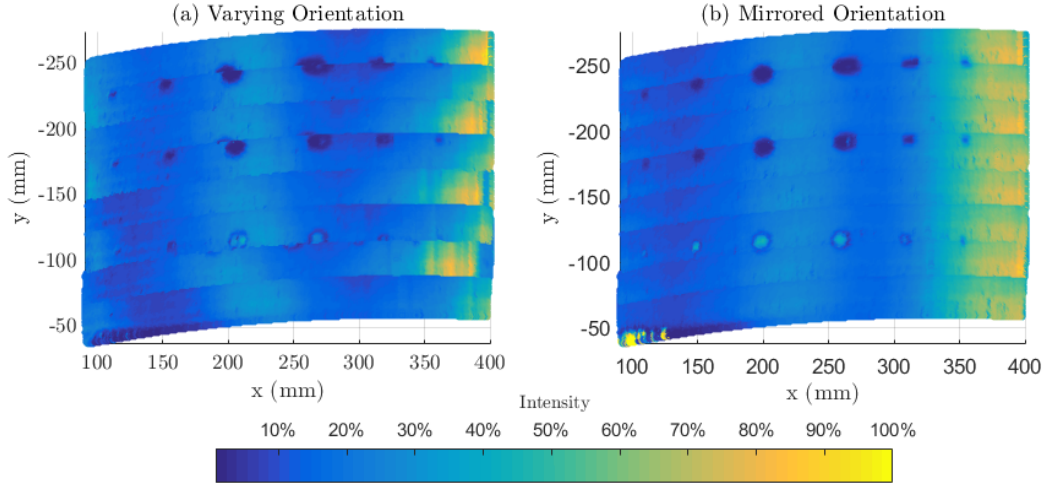

Figure 6.7: Difference in signal intensity from (a) having each robot's orientation be determined by separate surfaces and (b) constraining the two robots to mirror orientations.

Another path was also generated using Snell's law to attempt to calculate the exit point of the sound beam as the component was made of an isotropic material. Though

it was possible to calculate the angles, the refraction caused points to be projected backwards which led to a very irregular robot motion. Furthermore some points would have caused a collision between the end-effector and component due to the angle of refraction being very steep. An optimal approach to finding the best orientation of the second robot may be to simply measure the signal intensity at a number of key points and manually move the robot until the signal is maximised.

Following the results of the trials using an acrylic component, a path was designed using FASTSURF to inspect one side of the CFRP turbine blade, as shown in Figure 6.8a. During pulse-echo inspections, it was found that the thin sections near the edges of the component needed significantly different ultrasonic settings to be inspected, so the path was reduced to only inspect the middle area of the component. Two different KRL scripts were written to compare the new method to the existing GeoLink approach. For the GeoLink approach, a suitable alignment position was found in the robot cell working envelope and the two robots were manually aligned to maximise the observed ultrasonic signal. Figure 6.8b shows the actual inspection and position of the blade with respect to both robots.

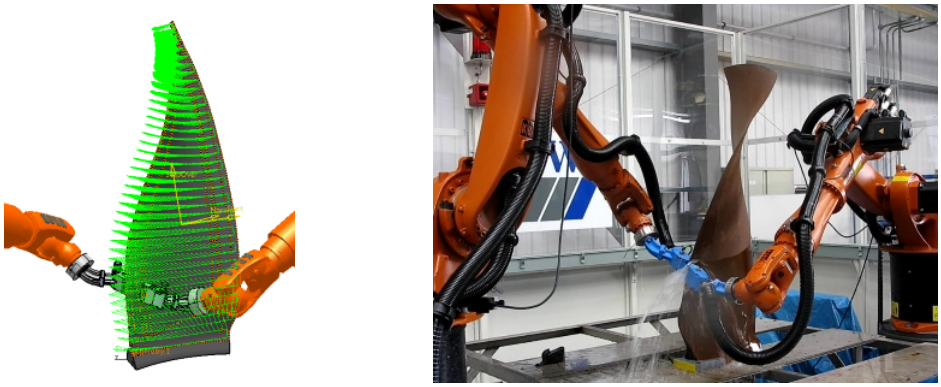

(a) (b)

Figure 6.8: (a) Off-line path planning and simulation for through transmission scans of a turbine blade. (b) Image from actual inspection using the IntACom system.

Figure 6.9 shows the results of the two types of TT scans. As mentioned previously, the scans were confined to a central area of the blade as the edges would otherwise appear over-saturated. As seen in the plots, the GeoLinked approach does not provide

any useful information about the part except for near the edge where the two surfaces were more parallel. This is because the two probes had to be separated far enough apart to avoid colliding with the thickest part of the blade. As a result, the probe was too far from the surface of the component and the signal deteriorated in the water path. As seen in Figure 6.6, these particular probes have to be around 5 to 10 mm from the surface due to the large water jet aperture and probe size.

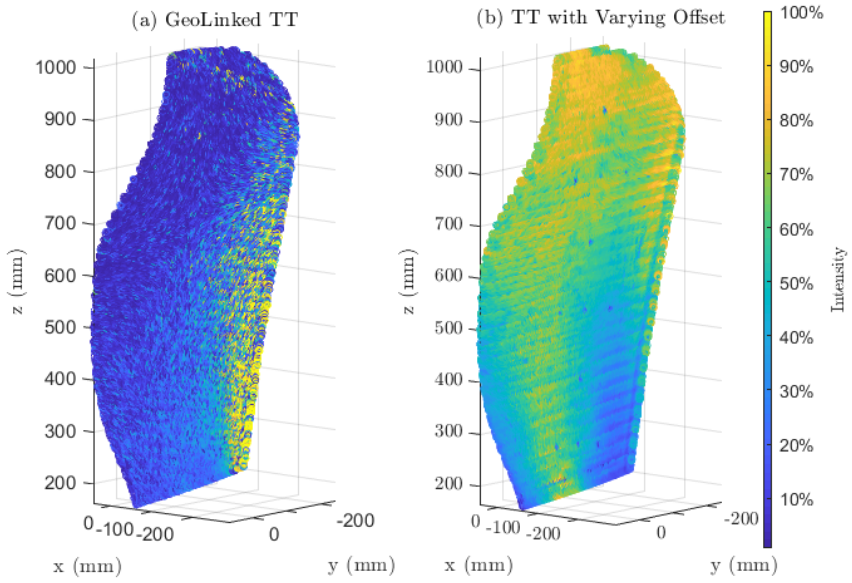

Figure 6.9: Amplitude of received signal during ultrasonic through-transmission scans using (a) the GeoLink method with a fixed offset and (b) the new method allowing fhe two robots to follow separate paths.

On the other hand, as seen in Figure 6.9b, the new method provides a much better signal across most of the inspected area. There are still some regions near the bottom of the blade where the orientation of the two surfaces causes the sound to refract too much. As the orientation of the slave probe mirrored that of the master, these areas could perhaps be improved by having different probe orientations. Another possible improvement could be to scan those regions along the length of the blade as the curvature is smaller in this direction. Using the technique developed in this thesis, it is possible to see the reference defects in the ultrasonic scan, thus achieving a result which

was not previously possible using the existing GeoLink method.

6.1.6 Discussion

Using twin robots for through-transmission inspections is one of the most useful features of the IntACom system. Other automated systems, such as immersion scanners, do not have this ability as access to both sides of a part is needed. Gantry-type systems have the ability to inspect from both sides but often lack the dexterity to cope with complex curvatures. To be able to deal with the complexity of aerospace geometries, it is necessary to be able to control both the transmitting and receiving probe positions for optimal results. This thesis presents an improvement on the previous method, whereby two robots are linked in a master-slave relationship by using combined synchronous operation. Besides being able to to inspect parts with large changes in thickness, there are also a number of other advantages to the new technique.

As discussed in Section 6.1.3, the GeoLink method relies on knowing the robot-to-robot calibration accurately, which can be achieved using external measurement devices. The new method only relies on accurately knowing where the part is in the reference frame of each robot. This means the part position can be calibrated for each robot, using the methods presented in Section 3.7 and Chapter 4. A further advantage is that the method can use path planning for either robot unlike the GeoLink method where the master-slave relationship is hard-wired to only work one-way, such that all paths have to be generated for the master robot. Allowing the orientations of each robot to vary (although the implementation of this needs to be improved) can potentially lead to much better results than what is achievable with the GeoLink method.

The drawbacks of the new method are the time needed to determine the position of the second robot by projecting each point through a CAD model. Though Section 6.1.4 demonstrated how this can be improved, the technique is still in a development phase and slow for large CAD components. Using the GeoLink method is in general faster both in terms of path planning and setup as the part position only needs to be calibrated once, but the GeoLink method is also much more sensitive to initial misalignment.

Future work should investigate how the projection algorithm can be improved in terms of processing time and how trajectory planning to match the orientation velocities can be achieved. To ensure the best possible signal at each point, one robot could possibly be instructed to automatically search for the pose which maximises the ultrasonic signal, which would alleviate the need for calculating the beam refraction. Another avenue to explore would be to use the external control method (through the RSI add-on) to stream coordinates to both robot simultaneously. This would involve significantly more work in terms of path and trajectory planning but would ultimately give the operator full control over each robot. This is the approach described in [180] and could potentially work well with KUKA robots due to the high update rate of the RSI module.

6.2 Optical Surface Scans

Chapter 4 discussed the challenges associated with automatically determining the position and orientation of an object in a robot's working envelope. As discussed in Chapter 4, the CAD model is an essential part of this process, both for path planning, but also for the proposed camera-based localisation method. This section discusses how to approach a situation wherein the component's geometry does not match its CAD model. This is a particularly pertinent problem in the composite manufacturing industry where deviations from the intended shape occur due to residual stress in the carbon fibres. The phenomenon is commonly known as "spring-back" and typically occurs at the post-curing phase of manufacture [2]. Due to the light weight and flexibility of CFRP, this deviation is often ignored as parts can manually be forced back to their intended shape during assembly.

During a robotic inspection, however, such deviations can lead to suboptimal scans. As paths are designed to follow a nominal shape, any deviations cause the probe to be misaligned with the component surface, or in extreme cases even collide with the surface. Forcing the component back into its intended shape would require bespoke fittings and extensive setup time, which defeats the purpose of a flexible inspection cell.

A solution to overcome this problem is to measure the actual surface geometry using a non-contact optical source (as mentioned in Section 2.5.4). It was identified that a laser-triangulation sensor provided the best compromise between price, ease of integration, scanning speed, resolution and accuracy. A MicroEpsilon 2950-50 laser triangulation sensor was acquired as part of the IntACom programme. The laser scanner, shown in Figure 6.10a, has a vertical measurement range of 65-125 mm, a depth resolution of 4 microns and can output profiles at up to 2000 Hz.

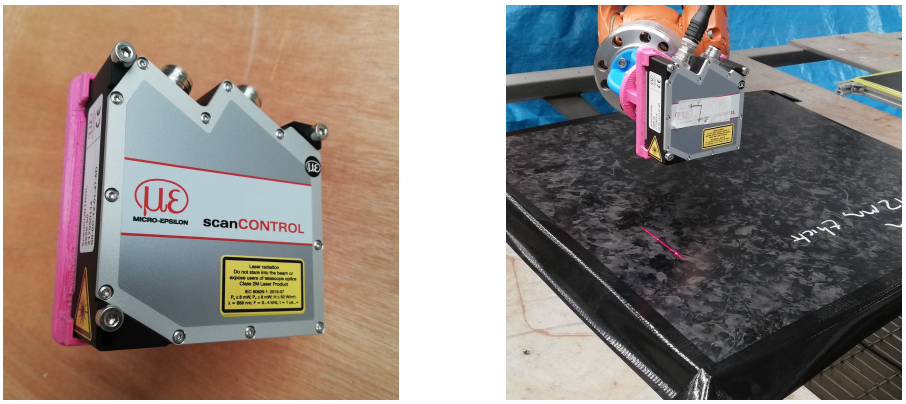

(a) (b)

Figure 6.10: (a) Laser triangulation sensor from MicroEpsilon. (b) Laser scan of CFRP plate which showed significant deviations from the expected geometry.

The MicroEpsilon triangulation sensor only supports four width resolutions, 1280, 640, 320 and 160 points per line. The spatial separation between these points depends on the distance to the surface that is being imaged, such that the spacing increases with depth. At the midrange, where the depth equals 95 mm, the width is 50 mm giving a spatial resolution of $\approx 39\mu$m. Using the maximum 1280 points per profile at the maximum depth results in a minimum horizontal resolution of $\approx 47\mu$m. Conversely, a maximum resolution of $\approx 31\mu$m is given using 1280 points at the minimum distance.

As discussed in Chapter 4, optical measurements can be used to ensure that the NDE trajectories can be executed correctly. A laser triangulation sensor can be used to both obtain new surface data and to accurately calibrate the robot to part coordinate transformation. The resolution of a laser profiler provides higher accuracy calibrations than what is achievable with machine vision along with simpler data analysis. Due

to the short measurement range of the laser profiler, however, the robot will have to move close to the part to obtain measurements. Thus, the use of a machine vision for obtaining an overview of the cell is still desirable to ensure that the correct component is in the cell and roughly in the correct position, prior to obtaining high accuracy measurements with a laser profiler. The two solutions can in this way complement one another, particularly in situations where new surface geometry data is required, which is not feasible to acquire with the machine vision solutions discussed previously.

6.2.1 Laser Profiler Integration

MicroEpsilon provide a Software Development Kit (SDK) which allows the user to send commands to the triangulation unit through an Ethernet cable. A custom C++ software script, based on the previous IntACom software architecture, was written as part of this thesis. A server application sends commands to the laser sensor to set parameters for the scan while also communicating with the robot controller through a separate UDP thread. All communication packets are sent via an Ethernet cable and network switch. This allows a single program to monitor the incoming data streams independently, which is necessary as the feedback rates of the two systems differ.

One thread listens to the UDP stream of coordinates and signals from the robot controller while the other thread handles communications with the laser sensor. As in ultrasonic inspections, the robot path contains both positional instructions and digital signals which are sent through the RSI framework. Throughout the robot path, if a designated digital signal is set to true, the two threads will begin to save robot coordinates and laser data to separate text files, both timestamped using the C++ function, `QueryPerformanceFrequency`, which establishes a counter with a set frequency which is consistent across all processors. The previous work by Dr. Mineo describes this approach to sensor integration in much greater detail [20].

To fuse the robot and laser data, each laser line is associated with a single robot position, where it is assumed that the robot speed is low enough such that the line does not move significantly during the exposure time. Both the robot positions and laser data are encoded with the same clock, but may not have arrived synchronously. It is

therefore necessary to interpolate between the robot's positions to infer the position of the laser when data was gathered. As long as both data streams are encoded with the same clock, it is the resolution and not the absolute accuracy which determines how well data points can be subsequently matched through temporal interpolation. The resolution can be increased arbitrarily by dividing the number of "ticks" from the C++ timer function into smaller intervals. A resolution of roughly 1 microsecond was used for this application.

The temporal resolution relates to the spatial accuracy as a function of the speed of the end-effector. If, for example, the robot moves the TCP at a constant speed of 500 mm/s the uncertainty in time encoding only relates to approximately $500\ \text{mm/s} \cdot 10^{-6}\ \text{s} = 0.5\,\mu\text{m}$, taking into account that each tick might not be exactly one microsecond. This provides a measure of the positional uncertainty due to the time it takes the server application to time-stamp a packet. Another main source of error is the latency between the laser, server and robots. While the laser moves at a constant speed, one can interpolate between different positions to obtain the required spacing between points (typically 1 mm). If linear interpolation is used, the overall uncertainty is then just the sum of the uncertainty in both points ($2 \cdot 0.5\,\mu\text{m} = 1\,\mu\text{m}$).

When operating normally, the laser triangulation sensor utilises the entire field of view of the sensor to measure the position of the laser beam. During full-frame operation, the profile measurement frequency is limited to 144 Hz. This limits the overall robot speed if high spatial resolution surface data is needed. The measuring field of the laser can be modified to only use a subsection of the camera's frame which allows for higher profile measurement rates, extending to 2000 Hz for the smallest regions. As mentioned in Section 3.4.1, cubic interpolation helps improve the accuracy of encoding, but to avoid errors, the best choice is to increase the frequency of measurements or to lower the speed of the robot to ensure higher sampling rates.

If the object to be scanned is well known then it may be acceptable to only use a reduced field of view for higher measurement speeds. In situations where uncertainty about the object's geometry exists, then it may be more advisable to use the full sensor to ensure all features of the geometry are captured in the field of view. These

considerations need to be taken into account when planning the robot path as the speed at which the robot moves at will determine the actual spacing of acquired laser data. For NDT applications, a surface spatial resolution of around one millimetre is sufficient to generate a new mesh which can be used for path planning.

6.2.2 Calibration

A fixture was designed using the CAD model of the MicroEpsilon unit to mount it as the end-effector of the robot. The fixture was 3D printed, as shown in Figure 6.10b. Unlike a normal robot tool, it is not possible to physically access the origin of the profiler's coordinate system as this is chosen by the manufacturer to be the point where the laser beam is emitted. Any attempt to use a traditional tool calibration method (as described in Section 3.6.1) would likely cause damage to the sensor or put the operator in danger as their eyes would be close to the laser beam.

The problem of finding the transformation between the reference frames of a laser profiler and robot has been previously addressed in the literature. The method presented by [185] was implemented in MATLAB to find a least squares solution to solving the unknown TCP transform. The method has already been presented in Section 3.6.3, as it is the same that was adapted to calibrate the tool transformation for a phased array probe. In the current instance, however, no additional imaging steps are required as the laser line provides a high resolution profile of the sphere directly in Cartesian coordinates.

To calibrate the TCP of the laser, the robot was manually driven such that the laser pointed towards a brass reference sphere with a precisely known diameter. A measurement was then taken of the sphere to create a surface profile. The robot was then driven to at least five other orientations and profiles were taken from these poses as well. An example calibration data profile is shown in 6.11a. Although some noise is present in the laser profile, this can be eliminated by using a digital moving average filter. A circle is fitted to the points using a best-fit, least squares solution. This further reduces the influence of any spurious reflections.

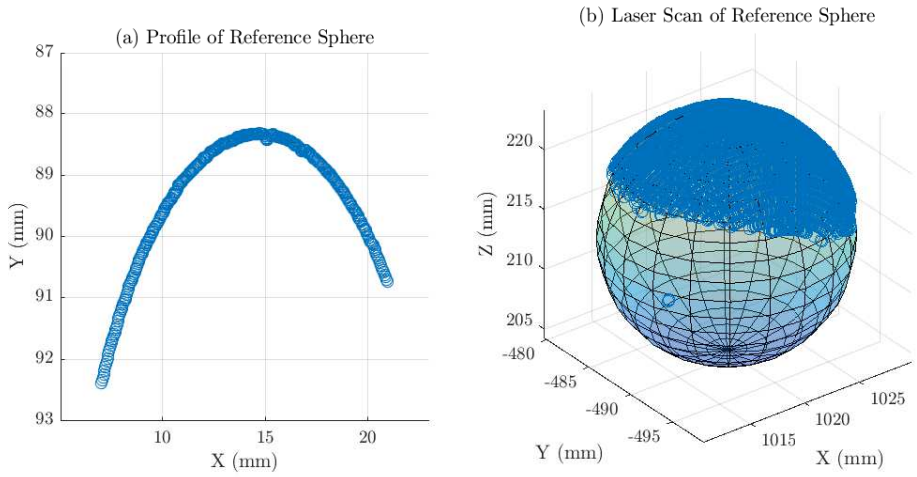

Figure 6.11: (a) Profile of a reference sphere acquired with a laser profiler. (b) Scan of top surface of reference sphere (blue points) along with best-fit sphere.

As described in Section 3.6.3, the centre and radius of each circular profile, acquired from different orientations, allows one to calculate the centre of a sphere in 3D space. Six orientations were chosen to calibrate the TCP of the laser and the error was found to be 0.274 mm. The error was calculated as described previously, by considering the spread of the projected centres of a sphere in 3D space. Furthermore, the Mean Square Error, described by Equation 3.13 was calculated as being 0.04 which further increases the confidence in the calculated TCP calibration.

To further test the accuracy of the obtained calibration, the entire top of a reference sphere with nominal diameter 19 mm ($\pm 0.01\,mm$) was scanned by moving the robot across the top in a fixed orientation. A best-fit sphere was calculated from the scanned points in MATLAB. This best fit solution seeks to minimise Equation 6.6 wherein the centre of the sphere is denoted by x_c, y_c, z_c and the radius of the sphere is r.

$$\min_{centre,r} \sum_{i=1}^{N} ((x_i - x_c)^2 + (y_i - y_c)^2 + (z_i - z_c)^2 - r^2)^2 \tag{6.6}$$

The result of fitting the observed data is shown in 6.11b. A least-squares fit of a

sphere to the data was performed and the calculated diameter was found to be 18.986 mm. The fitting error, Hence, the error is within 0.07% of the actual value, validating the TCP calibration of the laser on the robot.

The obtained results demonstrate the validity of the chosen calibration approach and its implementation in MATLAB. It is possible that the presented method can be automated, for example by installing a permanently mounted reference sphere within reach of the robot. This reference sphere could be used to perform accuracy checks on a regular basis to confirm the validity of the calibration. Signal processing steps such as filtering and excluding data outside a set range could be used to avoid manually selecting surface points. As discussed in Section 3.6, the same reference sphere could also be used to calibrate the TCP of a phased array probe.

From the calibration scans it may appear that the robot-laser system has an accuracy on par with that of purpose-built metrology systems. On small scales this may be the case, but as laser data is encoded with robot positions, the overall accuracy is limited by the robot's positional accuracy. It is therefore not envisioned that the system can be used for geometrical tolerance checks, which would furthermore require a better control of environmental factors such as temperature and lighting.

It should be noted that the objective of using the laser scanner is not to perform geometrical conformance checks, but rather to generate new CAD models of objects which are either legacy components or do not conform to their nominal CAD, as will be shown in the following section. Further to this, the laser can also be used to determine the position and orientation of objects to update the base coordinate reference system in the KUKA control software. This is done by measuring key features which are subsequently used to define a base reference system. Currently this process is done manually by driving a metal spike attached to a robot to key points.

6.2.3 Triangulation of Points

The direct output from a laser profiler is a 2D list of points with x and z values. By moving the laser across a surface using an encoder mechanism (a robot arm in the current implementation), a 3D point cloud can be generated for a number of profiles,

as shown in Figure 6.11b. Measurements can be taken on this point cloud and the overall shape can be visualised by plotting each point in 3D space. For robotic path planning purposes, a surface mesh must be generated such that the orientation with respect to surface normals can be calculated. The problem of generating a mesh from a point cloud is a common challenge for reverse engineering. There are several software solutions available for surface reconstruction and these tools are used within many fields ranging from archaeology, medicine and many areas of engineering. Common examples include: GeoMagic, Verisurf and Meshlab as well as a number of plug-ins for commercial CAD packages.

To reconstruct a surface, the observed points need to be connected to form a mesh. The simplest surface that can be generated is a triangle which only requires three points. The difficulty lies in determining which 3D points must link together to form the vertices of each triangle. Real world data often contains noise, spurious reflections, and areas of incomplete coverage. Hence the first step in the meshing process is often data cleaning, filling holes and re-sampling to ensure an even mesh coverage. To completely reverse engineer a part as it was designed, a laborious manual process is needed whereby 2D slices are taken from the point cloud data, re-sketched and extruded to form closed surfaces. For NDT purposes, however, just knowing the surface mesh is sufficient to generate an inspection path.

Different methods exist for creating meshes from just points and a popular choice is the Delaunay triangulation algorithm [186]. The Delaunay algorithm attempts to triangulate all points while ensuring that the circumsphere associated with each triangle contains no other points within it. Another widely adopted approach is to perform Poisson surface reconstruction which computes new points instead of using the original points in the point cloud. This allows the algorithm to vary the point density based on the local 3D curvature and create smoother surfaces. Poisson surface reconstruction is often used to mesh laser generated point cloud data, but tends to over-smooth the data [187]. Furthermore, as additional points are generated, the mesh may deviate from the actual surface geometry.

In collaboration with colleagues at the University of Strathclyde, a new approach

to visualising ultrasonic C-scan data was developed. This new method, known as the Index Based Triangulation (IBT), makes use of the fact that ultrasonic A-scan data is obtained in a sequential and structured manner. Hence, it is possible to efficiently determine which points must form triangles by attaching an index to each point as it is captured. This allows each data point to form the vertex of a triangle and the signal intensity at each vertex can then be used to shade the triangle to illustrate signal amplitude. The full details of the technique were presented in *Insight: The Journal of the British Institute of Non-Destructive Testing* [188].

Although the IBT method was intended to visualise ultrasonic data, the method can be extended to work with laser data as well. As the output of the laser line scanner is also sequential, each recorded point can have an index ranging from 1 to N associated with it. The IBT method can then be used to construct triangles from these vertices in an ordered manner. To illustrate this effect, a metal calibration step wedge (Figure 6.12) was scanned with the laser profiler. The part was placed on a metal plate within the IntACom cell and scanned using a single linear motion in one direction.

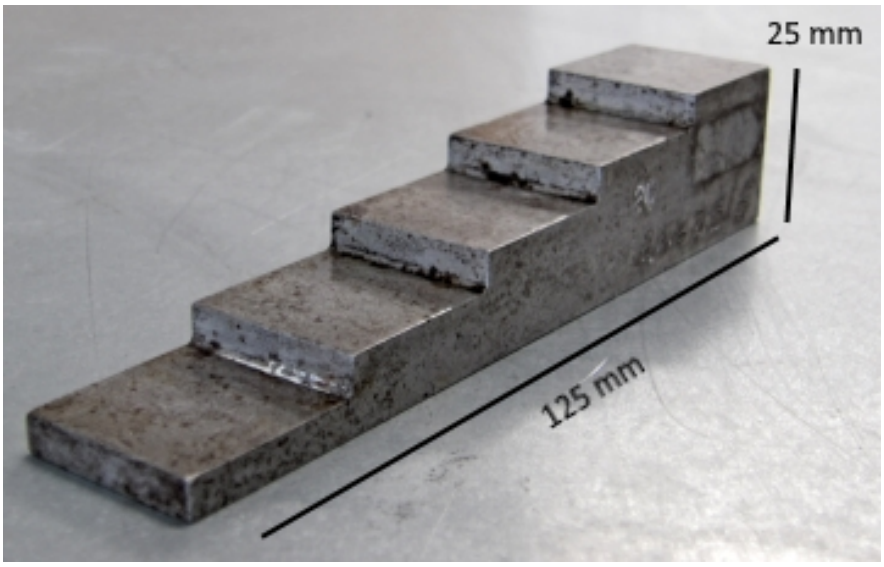

Figure 6.12: Metal reference block which was scanned with the laser profiler to test triangulation algorithm.

The IBT method was used to triangulate the points obtained from the laser scanner. To provide a reference mesh, the open-source software package MeshLab [189] was used

as well. The results of the two methods are shown in Figure 6.13. As seen in the figure, the IBT method provides a better representation of actual object geometry. As the IBT algorithm makes use of the knowledge of the order in which points were gathered, it is possible to accurately construct the triangles between each wedge step. The standard triangulation algorithm attempts to connect all triangles with the shortest distance which results in sloped lines as seen in Figure 6.13a. Measurements of the length of the piece taken from the mesh showed that the reconstruction method was accurate to within 0.2mm.

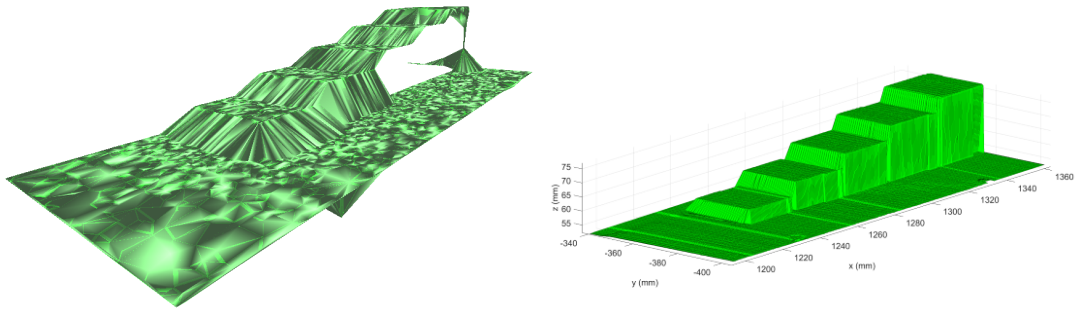

(a) (b)

Figure 6.13: Triangulation of the point cloud obtained from laser scanning the part shown in Figure 6.12 (a) Using the open-source software MeshLab (b) Using the developed IBT algorithm.

As seen in Figure 6.13b, the IBT algorithm provides both a truer representation of actual geometry with less manual input but also contains a more evenly sampled mesh. It is possible to re-mesh the surfaces obtained in MeshLab, but this is a time-consuming process and requires extensive training. The drawbacks of the IBT method are that laser data tends to contain sporadic reflections which will either lead to a noisy surface or unwanted triangles. Holes in objects also present a problem as these have a default distance of zero, causing unwanted surfaces to be generated. It is possible to remove these outliers, but this breaks the criteria for sequentiality needed by the IBT algorithm. Future work should investigate ways of filtering out these points or modifying the IBT algorithm to allow it to leave gaps in surfaces.

6.2.4 Carbon Fibre Plate Scan

This section demonstrates how a laser line profiler can be used in practice to improve ultrasonic scanning. A CFRP sample from an industrial partner was identified as significantly varying from the nominal geometry. The sample was a rectangular plate measuring approximately 500 x 500mm with a nominal thickness of 3.12 mm . The supplied CAD model indicated that the sample was supposed to be flat, but after visually inspecting the part, it was clear that the surface buckled in the middle and towards the edges. A preliminary ultrasonic scan was performed using a linear phased array (5 MHz, 0.6 mm pitch, 64 element) probe in a water jet nozzle moving approximately 10 mm above the part. Ultrasonic data was gathered using the Micropulse 5PA unit and a sub-aperture of 15 elements was linearly swept across the array such to match the elevation of the probe.

The resulting amplitude and time-of-flight scans are shown in Figure 6.14. As seen in the scans, there is significant variation in both the signal amplitude response as well as the front wall time-of-flight data. The probe had initially been calibrated using the technique presented in Chapter 3 of this thesis to ensure the TCP calibration did not affect the obtained results. Although some correlation appears to exist between the maximum amplitude and time-of-flight data, it is clear that the plate was warped such that an orientation mismatch was present. If the two datasets were correlated, the scan could probably be improved by simply changing the probe stand-off based on the time-of-flight.

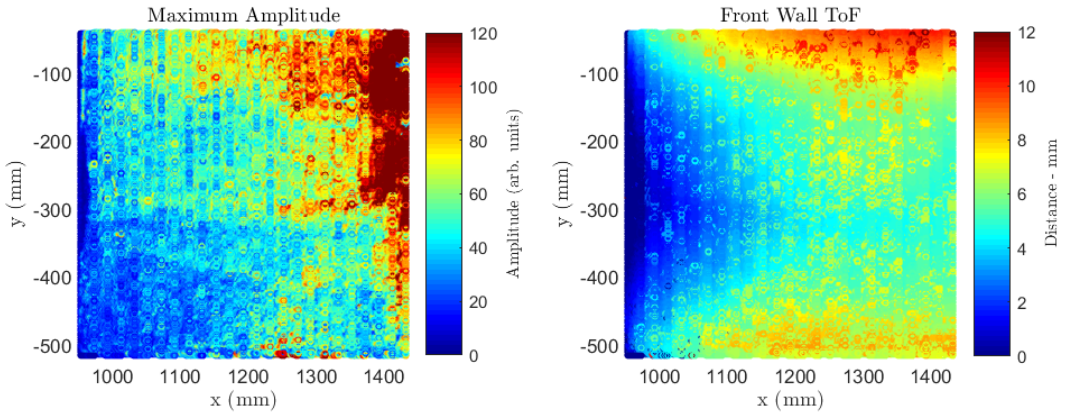

(a) (b)

Figure 6.14: Ultrasonic scans of a warped CFRP plate which does not conform to the expected geometry. (a) Maximum amplitude representation. As the probe becomes misaligned with surface, variation occurs. (b) Time of flight map showing variation in distance across the sample.

To generate a new scan path, the CFRP sample was scanned using the calibrated laser triangulation sensor, shown in Figure 6.10b. The plate was allowed to lay flat on a metal plate in the IntACom cell such that the surface could be mapped when the piece was not subjected to any load. A laser scan was performed in the midrange of the scanner's field of view, at approximately 100 mm from the component surface. The sample was scanned in a raster-type fashion with a step size of 50 mm which ensured sufficient overlap between each pass. As discussed, a MATLAB script was written to interpolate the laser and robot data to construct a 3D point cloud.

The IBT algorithm was used to create a meshed surface of the part, shown in Figure 6.15. As seen in the figure, a few spurious reflections caused unwanted triangles to be formed along the edges of the part. As expected, the part buckled significantly, showing deviations of -2 to 4 mm from the height at the centre of the plate. This resulted in a saddle-type geometry, as seen in Figure 6.15. The meshed data was afterwards exported as an STL file and turned into a CATIA solid to be used with the path planning application FASTSURF. A new robot scan path was generated based on

the scanned geometry and the component was ultrasonically scanned again.

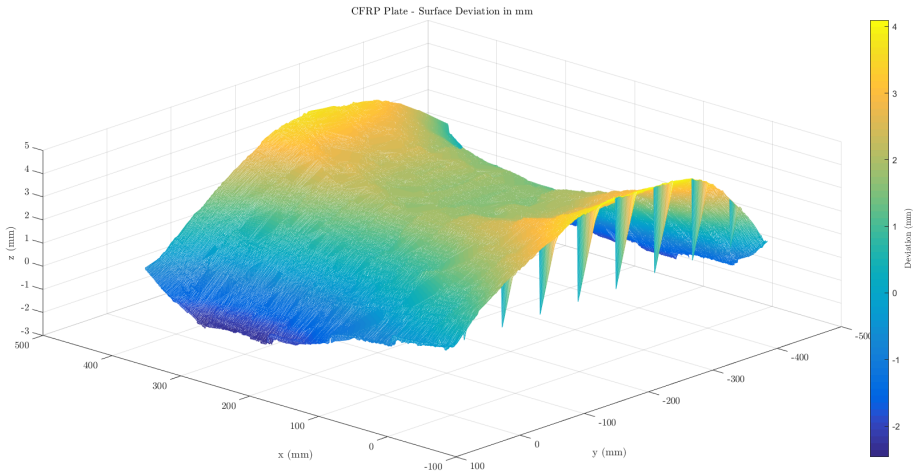

Figure 6.15: Meshed surface of the CFRP plate after triangulating the obtained point cloud from laser scanning using the IBT algorithm.

The results of the ultrasonic scans obtained using the new CAD model are shown in Figure 6.16 using the same colour representation as in Figure 6.14. From the time-of-flight map it is clear to see that the probe is now well aligned with the surface, except for small patches which show interference either due to pooling water or surface finish. The maximum amplitude map also shows that the signal response from the front wall is very consistent across the entire sample. The small areas of variation may be due to inconsistencies in pressure from the water pumps.

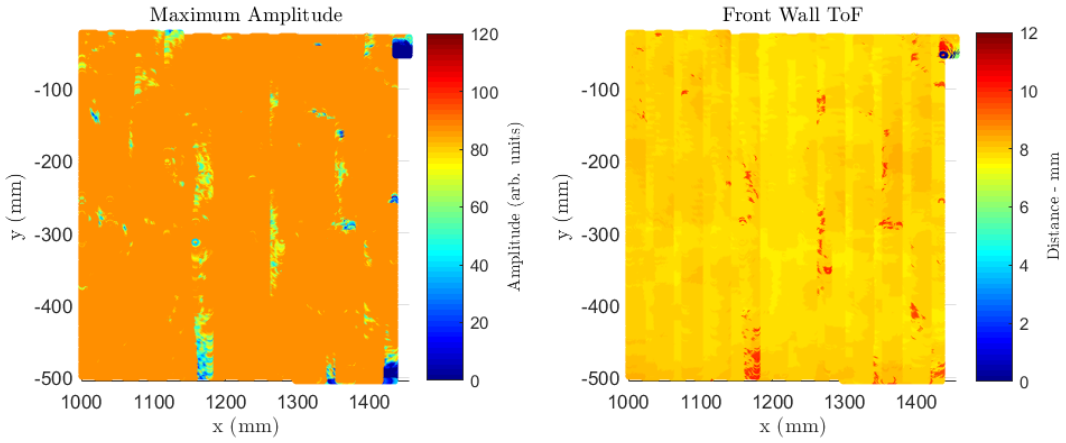

(a) (b)

Figure 6.16: Ultrasonic scans of a warped CFRP plate after a new CAD was generated from laser scans. (a) Maximum amplitude representation. (b) Time of flight map.

Figures 6.14 and 6.16 highlight the need and importance of being able to obtain a new CAD model when the component deviates from its nominal geometry. Looking at the variation in the two scans, the standard deviation in time of flight was 1.72 mm when the sample was scanned assuming a flat geometry. After updating the scan paths with the new CAD model, this variation was reduced to 0.12 mm, showing an improvement of over 1.5 mm on average. This work demonstrates the contribution of this thesis to industrial projects as it is unlikely this scan quality could have been obtained using the existing tools available in the IntACom project.

It should be noted that optically scanning components is a time consuming task as new paths have to be planned and the data processed. Using the IBT algorithm considerably lowers the time needed to obtain a surface mesh, but overall the procedure still takes several hours. Reverse engineering components and obtaining laser scanning data has become a large industry in recent years, driven to a large extent by the additive manufacturing industry. Several hand-held laser scanners are able to reconstruct large surface geometries in minutes using a combination of several laser line profilers and machine vision algorithms. Future work should look at optimising the speed and work-

flow of the integrated laser profiler to streamline the process for obtaining new CAD data.

6.3 Metrology Toolbox

To implement the calibration solutions developed throughout this thesis and create a tool which would increase the flexibility and adaptability of the IntACom robot cell, a GUI application was designed in MATLAB. By compiling the application, it is possible for it to run using only the MATLAB runtime environment, which allows it to be distributed within TWI. The GUI interface is shown in Figure 6.17. The GUI consists of four main parts which will be explained in this section.

The first part allows the operator to rotate and translate the TCP (or any pose in 3D) by known amounts. The second part determines the centre and orientation of a circle from 3 points which is useful for calibrating symmetrical components. The third part allows the user to find the pose of one robot (master or slave) in the reference frame of the other, as discussed in Section 6.1.3. Finally, the fourth section of the GUI allows the operator to calibrate a base frame using the ultrasonic alignment method presented in Section 3.7 and 5.7.2. The GUI also contains a graphical environment which visualises the results of each section.

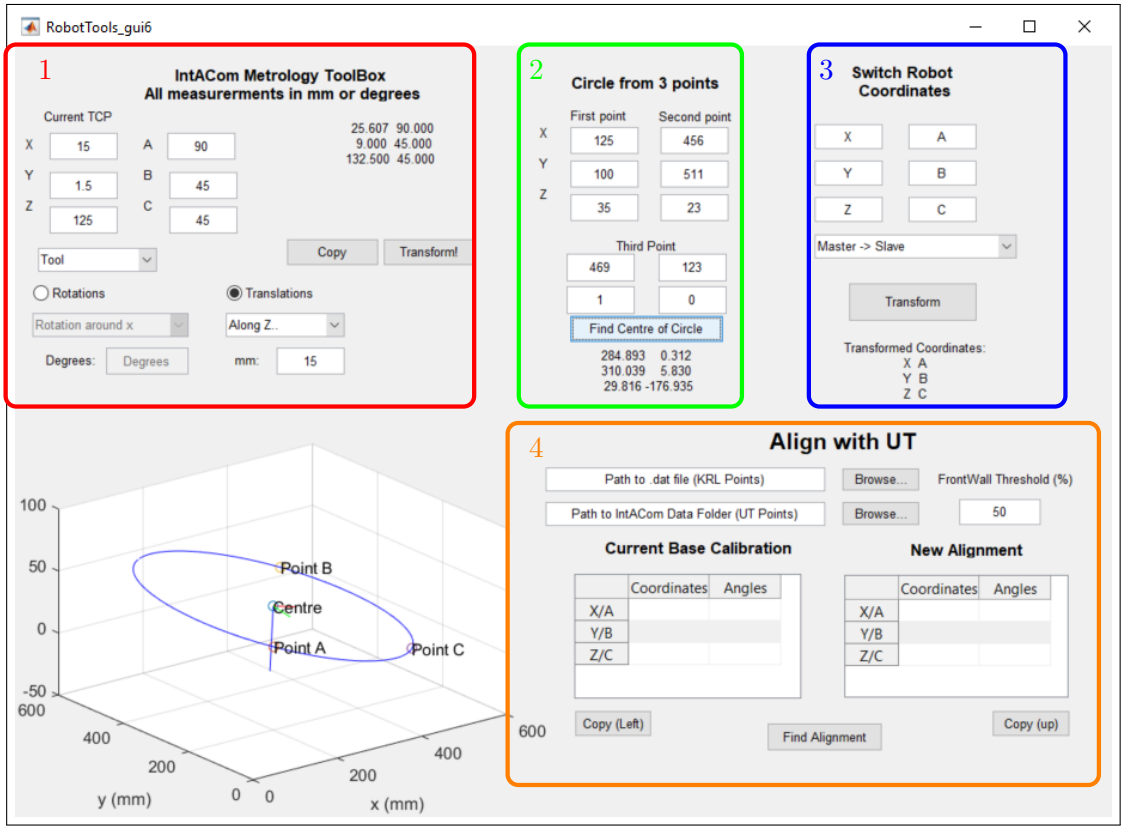

Figure 6.17: Metrology toolbox GUI for IntACom. 1) TCP 3D manipulation tool. 2) Defining a base coordinate system for symmetric objects. 3) Switching coordinates between master and slave robots. 4) Refining base coordinates using ultrasonic data.

To demonstrate the usefulness of the developed toolbox, the inspection of a cylindrical fan casing for jet engines will be used as an example. These engine casings, provided by an industrial partner of TWI, are regularly scanned using the IntACom robotic cell and present a number of challenges which will be addressed in the following sections.

6.3.1 TCP Adjustments

Although the overall shape of the fan casing is cylindrical, there is a significant variation in geometry including changes in thickness. To obtain a response from the back wall in pulse-echo inspections, the TCP of the robot is instructed to follow the opposing surface during path planning. Physically, however, the TCP of each water jet nozzle

is located just 10 mm from the opening of the nozzle. Thus, the nozzle would collide with the part at the thickest sections as the probe follows the internal (or external) surface. To avoid this, a new calibration artefact (as shown in Figure 3.9b) could be manufactured.

3D printing new calibration pieces for each inspection would be time consuming and infeasible. Instead, the TCP manipulation tool in the toolbox can be used to project the current TCP forward by a distance chosen by the operator. By entering the current TCP calibration, the operator can move or rotate the TCP along any of its three axes. By clicking "transform", the new TCP coordinates are displayed for the operator to manually input into the KUKA controller. If several transformations are needed in sequence, the "copy" button can be used to save the latest results, sparing the operator from having to type in new numbers.

6.3.2 Base Calibration of Cylinders

Another challenge encountered during the inspection of a fan casing was determining a way to calibrate the component base coordinate system. As discussed in Section 3.5.1, the part reference system is usually defined through three key points on an object, identified during the path planning stage. For symmetrical components such as a fan casing, there are no distinct features to unambiguously define a reference frame from. As an alternative, the component reference frame can be chosen to lie at the centre of a circle, defined by three points in 3D space. The orientation of the reference system is defined using the normal of the plane defined by the three points as well as the projection of the robot x-axis onto this plane.

To obtain 3D coordinates, one of the robots can used as a measuring device by recording the position of the TCP in the global reference frame of the robot. The tip of the designed calibration artefacts can be used as a measuring pointer or alternatively an ultrasonic reflection from the front wall interface can be used. The reason for only using three points is the time needed to drive the robot to different positions. Another possibility is to use the laser profiler, as discussed in Section 6.2. The ability to quickly and precisely define a reference frame by considering symmetries of the component

geometry has been used extensively for industrial inspections. An example of the graphical output from this feature is shown on the graph in Figure 6.17.

6.3.3 Cooperative Robot Base Frames

During through-transmission inspections, the alignment position of two probes must be defined at a synchronisation point, as discussed in Section 6.1.3. This point can be found manually, or it can be calculated using the third part of the metrology toolbox. By inputting the pose of a TCP in the world reference frame of one robot, the corresponding coordinates can be found in the other robot's frame of reference. The TCP adjustment tool can then be used to rotate that pose 180 degrees such that the two tools face each other. This saves an estimated 30-45 minutes of setup time during GeoLinked through-transmission inspections of large engine casings, as it is not always possible to manually align both probes due to the size of the part. Furthermore, the toolbox can also be used to obtain the pose of a component by measuring key points with one robot and translating these into the reference frame of the other. This alleviates the need to attach a pointer to the first robot.

6.3.4 Ultrasonic Base Alignment

The finally feature of the GUI implements the developed method described in Section 3.7, which uses the ultrasonic signal to fine adjust the base reference calibration. To use this feature, a reference scan of the component, covering 10-20% of the total inspection area is carried out. The ultrasonic settings need to be adjusted such that the front wall echo is clearly visible in the data. The data is gathered using the IntACom software and saved as a binary data file. The GUI needs to know the location of both the ultrasonic data file and the KRL text file which contains key points the robot moved to when data was obtained. As these points are assumed to lie on the surface of the component, deviations between the observed and intended signal arrival times are used to calculate a global offset. The GUI also reads the current base reference frame from the IntACom setup file and calculates the new base reference data which the operator can manually enter into the KUKA controller.

The developed method should only be used while keeping a few caveats in mind. As discussed in Section 3.7, the method is only able to update coordinates where unique features are present. For the fan casing inspections, for example, the cylinder is symmetric about the z-axis so any translation in z cannot be identified unless the edge of the cylinder is also scanned. The method is also only able to provide small corrections as finding deviations is cast as an optimisation problem. To converge on the correct solution, the method requires that a good guess is supplied as input. Even with these caveats, the ultrasonic alignment is used as a refinement step in a number of industrial inspections as it improves overall performance, as demonstrated in Section 5.7.2.

6.4 Conclusion

This chapter has demonstrated three key industrial contributions of the work presented in this thesis. Improving through-transmission inspections using dual robots, providing new surface data for warped components and developing a calibration toolbox are all results of the fundamental research presented in Chapters 2 to 5. An overarching thread within this thesis is how measurements can be used to improved existing industrial robot inspection systems. The three applications described in this chapter further highlight how this goal is achieved.

By measuring the robot-to-robot geometrical relationship, deviations during through-transmission inspections can be minimised. Using measurements taken from a CAD model of a component, it was shown that a new method for inspecting components with varying geometry in through-transmission is possible. Measurements of a component's surface geometry using a laser profiler allows parts without an associated CAD model to be inspected, yielding high quality ultrasonic results. Finally, four different calibration methods developed in the preceding chapters come together in a single GUI application.

There is still ample room for improvements to further automate and improve actual industrial inspections. Future work should consider expanding the metrology toolbox

and perhaps merge it with the RoboNDT toolbox developed by Dr. Mineo during his PhD [20]. The toolbox can also be further improved by incorporating graphical interfaces for tool calibrations; both ultrasonic and laser based. As discussed in Section 3.6.2 and 6.2, both TCP calibration procedures require a sphere to be scanned. Visualising this data would help the operator ensure the calibration is adequate. Next, the GUI toolbox could also automatically save existing tool and base frames from the robot KRL by requesting these via the RSI interface. Having the ability to send new tool and base coordinates back to the robot would also save considerable time and reduce the risk of an operator copying numbers into the KRL incorrectly.

Chapter 7

Conclusions

As the number of air passengers continues to increase every year, the aviation industry is preparing itself by modernising fleets and investing in new, fuel-efficient aircraft. The improved fuel efficiency is largely driven by the use of light-weight composite materials as primary load-bearing structures in planes. This has, in turn, meant that automated, faster and more reliable inspection methods are needed to ensure the structural integrity of these new components. Unlike metals, modern composites can be directly moulded into near-ready shapes, minimising the need for subtractive processes. The result is that parts arrive for inspection with complex, aerodynamic shapes that are difficult and time-consuming to inspect using manual or semi-automated NDT methods.

The use of industrial robot-based inspection systems has been shown to provide a step change in the inspection process of these parts. Not only can six-axis robots follow complex surface curvatures reliably, but the use of off-line path planning and phased array ultrasonic probes has reduced inspection times significantly. Furthermore, these robotic cells can be integrated at the fraction of the cost of a purpose-built automated scanning solutions such as large gantry systems.

At the time of writing there are a number of commercially available solutions which utilise industrial robots for ultrasonic inspections. Despite this, there are key areas relating to the obtainable accuracy and verification of these systems which has remained unexplored. A survey of the literature revealed faster and more accurate calibration methods to align the digital and real-world environments are needed. These calibra-

tions ensure that trajectories planned off-line are completed as expected during the actual inspection. Further challenges include methods for automatically determining the pose of components to be inspected, thorough validation techniques based on external measurement systems and integration of surface metrology measurements to cope with deviations from CAD.

This thesis is a continuation of previous collaborative efforts between the University of Strathclyde and TWI which established a prototype version of a robot-based automated NDT inspection system, known as the IntACom project. This thesis further expands upon the previous work by using a variety of measurements and calibration techniques to improve and verify the IntACom system's performance to allow it to be used in industry. To achieve this, the work in this thesis focuses on how dimensional measurements, along with the CAD model of a component can be used to improve automation levels, verify system performance and improve overall inspection quality.

Investigating what can be achieved in terms of absolute positional accuracy of six-axis robots through measurements and sensor integration is a central part of this thesis. Relevant aspects of six-axis manipulator system mechanism, programming and sensor interfacing were examined. It was determined that off-line path planning is a crucial part of a modern, flexible and adaptable robotic inspection system. However, it was also found that accurately executing these paths in real life is difficult due to the number of reference frames involved and how these relate to each other. Relevant aspects from the field of metrology were examined, where calibrations and traceability were identified as lacking in the field of robotised NDE.

It was investigated how different optical tracking systems could be used to monitor robot movements and measure positional accuracy. A study comparing the positions between a laser tracker, a KUKA KR5 industrial robot and a Vicon motion capture system was performed. It was found that the robot is accurate to around 0.5 mm on average when moving from point to point. The motion capture system was found to be less accurate (on average 0.8 mm), but the accuracy was highly dependent on camera placement. A laser tracker was also used to track a robot moving along a raster-pattern path at 400 mm/s where the path deviation was found to be 0.609 ± 0.182 mm.

Extensive research was committed to developing NDE-specific calibration procedures which would accurately define the tool-point and base reference parameters of the robot controller. Concepts from the medical industry and laser-sensor integration were adopted to be used with phased array transducers mounted as a tool on the end of a robot. A novel calibration procedure for determining the tool centre point, based on full-matrix capture ultrasonic imaging, was developed and tested. It was shown that the new calibration procedure improved the absolute accuracy of ultrasonic inspections by 73%, compared to manual calibration methods. The developed method also increased the signal amplitude by 30% as the probe was better aligned to the component surface.

Another calibration method to optimise the robot-to-part calibration was also developed by using the ultrasonic reflection from the part surface. By considering the difference in signal arrival time between the expected and actual signals, it was possible to calculate a global pose correction. Optimising calibrations in this way led to an overall reduced ultrasonic signal variation during actual inspections. For aerospace components with complex surface geometries, the new method resulted in an improvement of over 30% in the quality of the received signal and a 61% reduction in stand-off variation. This method works by only scanning a fraction of the component surface and has since been used extensively during industrial inspections.

Part of this thesis also investigated methods of using computer vision and CAD models to increase automation levels. Specifically, it was examined how optical measurements taken at a distance could be used to automatically determine the position, orientation and type of component in the robot cell. This would allow robots to begin automatically scanning with a minimum of operator input. A number of object recognition approaches were investigated and it was found that using geometrical shape models for identification and pose estimation was the most appropriate choice for NDE applications. This was due to the optical appearance and size of CFRP components as well as the lack of available image data, which meant standard computer vision methods were not suitable.

A work-flow based on the CAD model and calibrated cameras was developed and tested in a laboratory setting. It was shown that objects could be recognised based

on their shape outline with a high degree of confidence using a Bayesian classification approach. The use of a 3D camera alongside 2D segmentation methods illustrated how it is possible to perform pose estimation using the CAD model of a part. It was found that using cameras for large volume measurements could provide an accuracy which was sufficiently high to allow the ultrasonic part alignment method to work, establishing a multi-scale metrology approach to increasing automation levels.

The research presented in this thesis is a result of an RCNDE Engineering Doctorate project (grant number EP/L015587/1) between the University of Strathclyde and TWI. The Engineering Doctorate programme is specifically aimed at building bridges between industry and academia such that fundamental research can be transferred and implemented in real-world applications. Part of the work presented in this thesis investigated ways in which measurements and CAD models could be used to solve inspection challenges encountered when inspecting actual industrial samples. A method for improving through-transmission ultrasonic inspections using twin robots was realised. This method uses a new synchronisation technique and enables the IntACom system to inspect components with large changes in thickness, something which was not previously possible.

To further increase the usability of the IntACom system, a laser profiler was integrated as a tool the robots can manipulate and a novel method for reconstructing 3D surfaces from structured data points was developed. Using a laser profiler to reconstruct the geometry of parts which deviate significantly from the CAD models was seen as a key enabler by the industrial partners in the IntACom project. A simple to use graphical interface was also developed to combine the various calibration methods into a single toolbox.

Unlike other robot-based inspection systems on the market, the IntACom system remains a research tool which serves the need of the industrial members of TWI. To deal with a wide variety of components, the system is intentionally designed to be adaptable and flexible. This means that the use of expensive, rigid fixtures and machined tool holders is avoided in favour of modifiable, 3D printed jigs. Taking this approach means that proof-on-concept trials can quickly be realised and optimised for prototype parts,

allowing industry to test and determine the best inspection solutions before investing in new equipment.

Off-line programming is a key step in realising this system approach. Simulation software vastly reduces set-up times and enables complex paths to be realised, but relies on a number of assumptions, including that robots move accurately and know where parts are located. These assumptions are rarely true and must be compensated for by a combination of calibration, compliance, and sensing. The work presented in this thesis provides solutions to all three aspects by using novel calibration techniques, exploiting the compliance offered by water jets and integrating large volume optical sensing.

7.1 Future Work

The results of this research work inspired many ideas which the author hopes will be investigated in the future. It was shown that better calibration techniques can accurately determine the relationships between a robot, end-effector and external measurement systems. Knowing these relationships should enable the external system to monitor and correct the robot's motion. To do this, however, the motion of the robot would have to be externally controlled, which can lead to serious safety concerns if not implemented correctly.

Being able to control robot motion based on external sensor feedback could also lead to an improved calibration technique for through-transmission inspections. As discussed in this work, finding the optimal alignment between transmitting and receiving probes is influenced by how sound is refracted within a component. By moving the receiver until a signal is maximised, the optimal dual-robot alignment could be determined.

The work in this thesis highlighted an unexpected, weaving pattern when a robot was instructed to move along a straight line. Future work should investigate the source of this error, especially whether it arises as a result of mechanical interactions within the robot motors or is a consequence of the translation between Cartesian coordinates and joint-angle values.

Future work should continue to investigate how large-volume measurements can be taken using machine vision techniques based on the CAD model of an object. To overcome some of the limitations identified in this work, large area structured light systems could be used to clearly identify key features on otherwise optically uniform CFRP structures. Information from path planning, optical measurements and ultrasonic data could be combined using the Bayesian formalism to create a more robust pose estimation procedure which takes uncertainties into account to update a probabilistic model. Integrating measurements taken this way with the robot controller would further increase automation levels.

The IntACom project enjoys continued support from its industrial partners and will likely enter its fourth iteration in the near future. The newest addition to the project are two larger robots on linear tracks and a 4 metre diameter turntable. Calibration techniques for these external axes are already an active area of research at TWI. Areas of further development include developing collaborative robot capabilities for on-site inspection, introducing advanced ultrasonic imaging such synthetic aperture focusing techniques and introducing machine learning methods for assisted defect detection.

Bibliography

[1] I. Cooper, I. Nicholson, D. Liaptsis, B. Wright, and C. Mineo, "Development of a fast inspection system for complex aerospace structure," in *6th International Symposium on NDT in Aerospace*, Madrid, 2014.

[2] M. E. Ibrahim, "Nondestructive evaluation of thick-section composites and sandwich structures: A review," *Composites Part A: Applied Science and Manufacturing*, vol. 64, pp. 36–48, 2014.

[3] A. Hobbs, "An Overview of Human Factors in Aviation Maintenance," Australian Transport Safety Bureau, Tech. Rep. AR-2008-055, 2008.

[4] R. Bogue, "The role of robotics in non-destructive testing," *Industrial Robot: An International Journal*, vol. 37, no. 5, pp. 421–426, aug 2010.

[5] T. P. Sattar and A. Brenner, "Robotic system for inspection of test objects with unknown geometry using NDT methods," *Industrial Robot: An International Journal*, vol. 36, no. 4, pp. 340–343, jun 2009.

[6] A. Maurer, W. Deodorico, R. Huber, and T. Laffont, "Aerospace Composite Testing Solutions using Industrial Robots," *18th World Conference on Nondestructive Testing*, no. April, p. 7, 2012.

[7] C. Mineo, S. G. Pierce, B. Wright, I. Cooper, and P. I. Nicholson, "PAUT inspection of complex-shaped composite materials through six DOFs robotic manipulators," *Insight - Non-Destructive Testing and Condition Monitoring*, vol. 57, no. 3, pp. 161–166, mar 2015.

[8] C. Mineo, D. Herbert, M. Morozov, and S. G. Pierce, "Robotic Non-Destructive Inspection," in *51st Annual Conference of the British Institute of Non-Destructive Testing*, vol. 44, 2012, pp. 345–352.

[9] C. Mineo, S. G. Pierce, P. I. Nicholson, and I. Cooper, "Robotic path planning for non-destructive testing – A custom MATLAB toolbox approach," *Robotics and Computer-Integrated Manufacturing*, vol. 37, pp. 1–12, 2016.

[10] C. J. Hellier, *The Handbook of Nondestructive Evaluation.* New York, N.Y.: McGraw-Hill, 2001.

[11] J. J. Craig, *Introduction to Robotics: Mechanics and Control.* Pearson/Prentice Hall, 2005.

[12] R. Bogue, "The role of robotics in non-destructive testing," *Industrial Robot: An International Journal*, vol. 37, no. 5, pp. 421–426, 2010.

[13] I. Golightly and D. Jones, "Visual control of an unmanned aerial vehicle for power line inspection," in *ICAR '05. Proceedings., 12th International Conference on Advanced Robotics, 2005.* IEEE, 2005, pp. 288–295.

[14] D. Zhang, K. Burnham, L. Mcdonald, C. Macleod, G. Dobie, R. Summan, and G. Pierce, "Remote Inspection of wind turbine blades using UAV with photogrammetry payload," in *56th Annual British Conference of Non-Destructive Testing-NDT*, no. 5, 2017, pp. 1–11.

[15] S. Roschi, "These Robots Crawl Into Every Nook and Cranny," 2017, Date accessed: 2017-09-21. [Online]. Available: https://drive.tech/en/stream-content/these-robots-crawl-into-every-nook-and-cranny

[16] International Federation of Robotics – IFR, "Executive Summary World Robotics 2019 - Industrial Robots," IFR International Federation of Robotics, Frankfurt, Tech. Rep., 2019.

[17] I. Cooper, I. Nicholson, D. Yan, B. Wright, D. Liaptsis, and C. Mineo, "Development of a Fast Inspection System for Complex Composite Structure - The Intacom Project," in *5th International Symposium on NDT in Aerospace*, Singapore, 2013, pp. 13–15.

[18] H. Colestock, *Industrial Robotics: Selection, Design, and Maintenance.* McGraw-Hill, 2005.

[19] M. Kleinkes and R. Loser, "Laser Tracker and 6DoF measurement strategies in industrial robot applications," in *CMSC 2011: Coordinate Metrology System Conference*. Phoenix, Arizona: CMSC, 2011, pp. 25–28.

[20] C. Mineo, "Automated NDT inspection for large and complex geometries of composite materials," Ph.D. dissertation, University of Strathclyde, 2015.

[21] K. Henning, W. Wolfgang, and H. Johannes, "Recommendations for implementing the strategic initiative INDUSTRIE 4.0," ACATECH – National Academy of Science and Engineering, München, Final Report of the Industrie 4.0 Working Group April, 2013.

[22] P. Thayer, "Enabling the Fourth Industrial Revolution (4IR) and the role of NDE and monitoring," *Insight - Non-Destructive Testing and Condition Monitoring*, vol. 59, no. 9, pp. 469–472, 2017.

[23] R. Bacon, "Growth, structure, and properties of graphite whiskers," *Journal of Applied Physics*, vol. 31, no. 2, pp. 283–290, 1960.

[24] American Chemical Society National Historic Chemical Landmarks., "High Performance Carbon," Date accessed: 2018-08-02. [Online]. Available: http://www.acs.org/content/acs/en/education/whatischemistry/landmarks/carbonfibers.html

[25] R. Bogue, "New NDT techniques for new materials and applications," *Assembly Automation*, vol. 32, no. 3, pp. 211–215, 2012.

[26] G. Marsh, "Airbus A350 XWB update," *Reinforced Plastics*, vol. 54, no. 6, pp. 20–24, 2010.

[27] T. Kellner, "GE Reports: These Materials Scientists Are Teaching Robots Awesome New Tricks," GE Aviation, Tech. Rep., 2014, Date accessed: 2018-01-23. [Online]. Available: https://www.ge.com/reports/post/101419232330/these-materials-scientists-are-teaching-robots/

[28] D. K. Hsu and A. Minachi, "Defect Characterisation in Thick Composites by Ultrasound," *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 9, pp. 1481–1488, 1990.

[29] P. Cawley, "Inspection of Composites – Current Status and Challenges," in *9th European Confernce on NDT (ECNDT) 2006*, 2006, pp. 1–10.

[30] E. Cuevas and S. Hernández, "Robot-based solution To Obtain an Automated , Integrated And Industrial Non-Destructive Inspection Process," *6th International Symposium on NDT in Aerospace*, no. November, pp. 12–14, 2014.

[31] J. Summerscales, *Non-destructive testing of fibre-reinforced plastics composites.* Springer Netherlands, 1990.

[32] I. Scott and C. Scala, "A review of non-destructive testing of composite materials," *NDT International*, vol. 15, no. April, pp. 75–86, 1982.

[33] R. Halmshaw, R. Honeycombe, and P. Hancock, *Non-destructive testing.* Edward Arnold, 1991.

[34] P. Cawley, R. A. Smith, and R. J. Freemantle, "Composites in NDE Course Material," 2014, Imperial College London.

[35] G. Mook, R. Lange, and O. Koeser, "Non-destructive characterisation of carbon-fibre-reinforced plastics by means of eddy-currents," *Composites Science and Technology*, vol. 61, pp. 865–873, 2001.

[36] A. E. Birt, "The Applicability of X-Radiography to the Inspection of Composites," 2014.

[37] M. Krumm, C. Sauerwein, V. Hämmerle, S. Heile, T. Schön, A. Jung, and M. Sindel, "Rapid Robotic X - ray Computed Tomography of Large Assemblies in Automotive Production," in *8th Conference on Industrial Computed Tomography, Wels, Austria (iCT 2018).* NDT.net, 2018, pp. 1–9.

[38] R. S. C. Cobbold, *Foundations of biomedical ultrasound.* Oxford University Press, 2006.

[39] J. Krautkrämer and H. Krautkrämer, *Ultrasonic testing of materials*, 4th ed. Springer-Verlag Berlin Heidelberg, 1990.

[40] Noël Dubé, *Introduction to Phased Array Ultrasonic Technology Applications: R/D Tech Guideline*. Waltham,MA: Olympus NDT, 2007.

[41] P. Corke, *Robotics, Vision and Control: Fundamental Algorithms In MATLAB®*, 2nd ed. Springer-Verlag Berlin Heidelberg, 2013.

[42] J. R. Parks and D. A. Bell, "Industrial robots and machine intelligence," *Physics Bulletin*, vol. 21, pp. 549–553, 1970.

[43] USL Ultrasonic Sciences, "Automated Ultrasonic Inspection Systems," 2018, Date accessed: 2018-08-07. [Online]. Available: http://www.ultrasonic-sciences.co.uk/products/

[44] M. Chahbaz, "The Use of Industrial Robots for NDT Applications," *Quality Magazine (online)*, sep 2016, Date accessed: 2017-12-14. [Online]. Available: https://www.qualitymag.com/articles/93488-the-use-of-industrial-robots-for-ndt-applications

[45] ScanMaster, "ScanMaster Ultrasonic Inspection Solutions," Date accessed: 2018-08-07. [Online]. Available: http://scanmaster-irt.com/

[46] D. Stubbs, R. Cook, D. Erdahl, I. Fiscus, D. Gasper, J. Hoeffel, W. Hoppe, V. Kramb, S. Kulhman, R. Martin, R. Olding, D. Petricola, N. Powar, and J. Sebastian, "An automated ultrasonic system for inspection of aircraft turbine engine components," *Insight: Non-Destructive Testing and Condition Monitoring*, vol. 47, no. 3, pp. 157–162, 2005.

[47] V. A. Kramb, "Use of phased array ultrasonics for automated aerospace testing applications," *Materials Evaluation*, vol. 65, no. 1, pp. 67–73, 2007.

[48] P. Louviot, A. Tachattahte, and D. Gardener, "Robotised UT Transmission NDT of Composite Complex Shaped Parts," in *4th International Symposium on NDT in Aerospace, Berlin (Germany)*, 2012, pp. 1–8.

[49] E. Cuevas, M. López, and M. García, "Ultrasonic Techniques and Industrial Robots : Natural Evolution of Inspection Systems," in *4th International Symposium on NDT in Aerospace*. Berlin, Germany: NDT.net, 2012, pp. 1–12.

[50] C. Mineo, C. MacLeod, M. Morozov, S. G. Pierce, T. Lardner, R. Summan, J. Powell, P. McCubbin, C. McCubbin, G. Munro, S. Paton, D. Watson, and D. Lines, "Fast ultrasonic phased array inspection of complex geometries delivered through robotic manipulators and high speed data acquisition instrumentation," in *IEEE International Ultrasonics Symposium, IUS*, vol. 2016-Novem. Tours, France: IEEE, 2016, pp. 1–4.

[51] F. Khodayar, F. Lopez, C. Ibarra-Castanedo, and X. Maldague, "Optimization of the Inspection of Large Composite Materials Using Robotized Line Scan Thermography," *Journal of Nondestructive Evaluation*, vol. 36, no. 2, 2017.

[52] Tecnatom-NDT, "October 2016 Newsletter," *Tecnatom Aircraft & Space Newsletter*, no. October (11), pp. 1–4, 2016.

[53] Ultrasonic Sciences Ltd., "A new NDT development using high speed robots," *The e-journal of nondestructive testing & ultrasonics*, no. 2014-03, pp. 1–3, 2014.

[54] B. I. of Non-destructive Testing (BINDT), "Global non-destructive testing equipment market forecast to 2024," apr 2018, Date accessed: 2018-02-15. [Online]. Available: http://www.bindt.org/News/April-2018/global-non-destructive-testing-equipment-market-forecast-to-2024/

[55] GE Inspections, "Hydrastar: Robotic Ultrasonic Solution for Composite Inspection," 2015. [Online]. Available: https://www.gemeasurement.com/inspection-ndt/automated-ut-systems/hydrastar-robotic-ultrasonic-solution-composite-inspection

[56] B. Vanderheiden, C. Thomson, I. Ivakhnenko, and C. Garner, "Transition to high rate aerospace NDI processes," in *44th Annual Review of Progress in Quantitative Nondestructive Evaluation*, vol. 37, 2018, pp. 020 003–1–020 003–13.

[57] FILL Gmbh, "Ultrasonic Inspection Dual Robot - ," Date accessed: 2018-11-14. [Online]. Available: https://www.fill.co.at/en/aerospace/composites/

non-destructive-inspection/ultrasonic/ultrasonic-inspection-dual-robot/highlight/1412p604.html

[58] D. Flack and H. John, "Fundamental Good Practice in Dimensional Metrology - No. 80," NPL, Tech. Rep. 80, 2005.

[59] C. Mineo, S. G. Pierce, P. I. Nicholson, and I. Cooper, "Introducing a novel mesh following technique for approximation-free robotic tool path trajectories," *Journal of Computational Design and Engineering*, vol. 4, no. 3, pp. 192–202, 2017.

[60] C. Eitzinger, A. Walch, W. Palfinger, and S. Beyer, "Reactive coverage planning for robotic NDT of complex parts," in *12th European Conference on Non-Destructive Testing (ECNDT 2018)*. Gothenburg: NDT.net, 2018, pp. 1–8.

[61] B. Simpson and P. J. Dicken, "Integration of Machining and Inspection in Aerospace Manufacturing," in *IOP Conference Series: Materials Science and Engineering*, vol. 26. IOP Publishing, dec 2011, pp. 12–14.

[62] N. Ersoy, T. Garstka, K. Potter, M. R. Wisnom, D. Porter, and G. Stringer, "Modelling of the spring-in phenomenon in curved parts made of a thermosetting composite," *Composites Part A: Applied Science and Manufacturing*, vol. 41, no. 3, pp. 410–418, 2010.

[63] D. S. Forsyth, "Experiences in practicing the assessment of nondestructive testing performance," in *44th Annual Review of Progress In Quantitative Nondestructive Evaluation*, vol. 37, 2018.

[64] B. Bridges, "How Laser Trackers Work," *Quality Digest (online)*, p. 2, jun 2009. [Online]. Available: https://www.qualitydigest.com/inside/twitter-ed/how-laser-trackers-work.html

[65] T. Ulrich, "Uncertainty Estimation for Kinematic Laser Tracker Measurements," in *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. Sydney: IEEE, 2012, pp. 1–10.

[66] T. Luhmann, "ISPRS Journal of Photogrammetry and Remote Sensing Close range photogrammetry for industrial applications," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 65, no. 6, pp. 558–569, 2010. [Online]. Available: http://dx.doi.org/10.1016/j.isprsjprs.2010.06.003

[67] C. Fraser, "What Is Photogrammetry?" *Quality Digest (online)*, jul 2015, Date accessed: 2018-03-23. [Online]. Available: https://www.qualitydigest.com/inside/metrology-article/070115-what-photogrammetry.html

[68] D. Zhang and G. Lu, "Review of shape representation and description techniques," *Pattern Recognition*, vol. 37, no. 1, pp. 1–19, 2004.

[69] D. G. Lowe, "Distinctive image features from scale invariant keypoints," *Int'l Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.

[70] J. E. Muelaner and P. G. Maropoulos, "Large volume metrology technologies for the light controlled factory," *Procedia CIRP*, vol. 25, no. C, pp. 169–176, 2014.

[71] B. R. Barbero and E. S. Ureta, "Computer-Aided Design Comparative study of different digitization techniques and their accuracy," *Computer-Aided Design*, vol. 43, no. 2, pp. 188–206, 2011.

[72] R. Kurtz and R. Nesbitt, "Improving the Accuracy of Surface Metrology," *Optical Engineering*, vol. 50, 2011.

[73] R. Schwarte, "Principles of 3-D Imaging Techniques," *Handbook of Computer Vision and Applications*, vol. 3, 1999.

[74] MTI Instruments, "Laser Triangulation Systems: Non-contact Measurement Sensors, Solutions and Systems - MTI Instruments," Date accessed: 2015-01-09. [Online]. Available: https://www.mtiinstruments.com/technology-principles/laser-triangulation-sensors/

[75] C. Chen and A. Kak, "Modeling and calibration of a structured light scanner for 3-D robot vision," *Proceedings. 1987 IEEE International Conference on Robotics and Automation*, 1987.

[76] P. Garbat, W. Skarbek, and M. Tomaszewski, "Structured light camera calibration," *Opto-Electronics Review*, vol. 21, no. 1, pp. 23–38, dec 2012.

[77] C. Mineo, C. MacLeod, M. Morozov, S. G. Pierce, R. Summan, T. Rodden, D. Kahani, J. Powell, P. McCubbin, C. McCubbin, G. Munro, S. Paton, and D. Watson, "Flexible integration of robotics, ultrasonics and metrology for the inspection of aerospace components," in *43rd Annual Review of Progress in Quantitative Nondestructive Evaluation*, vol. 1806, no. 1. AIP Conference Proceedings, 2017.

[78] International Standardization Organization (ISO), "ISO 8373: 2012. Robots and robotic devices–Vocabulary," 2012.

[79] Mathieu Bélanger-Barrette, "How to Choose the Right Industrial Robot?" 2014, Date accessed: 2018-05-28. [Online]. Available: https://blog.robotiq.com/bid/70408/How-to-Choose-the-Right-Industrial-Robot

[80] KUKA Roboter GmbH.

[81] A. Lauletta, "The Basics of Harmonic Drive Gearing," *Gear Product News*, pp. 32–36, apr 2006, Date accessed: 2018-05-21. [Online]. Available: http://www.gearproductnews.com/issues/0406/gpn.pdf

[82] International Federation of Robotics, "World Robotic Report - Executive Summary 2013," International Federation of Robotics, Tech. Rep., 2013.

[83] R. S. Hartenberg and J. Denavit, "A Kinematic Notation for Lower-Pair Mechanisms Based on Matrices," *Transaction ASME J Appl Mech*, vol. 22, pp. 215–221, 1955.

[84] T. Brogårdh, "Robot control overview: An industrial perspective," *Modeling, Identification and Control*, vol. 30, no. 3, pp. 167–180, 2009.

[85] International Federation of Robotics – IFR, "History of Industrial Robots - From the first installation until today," IFR, Tech. Rep., 2012.

[86] L. S. Ginani and J. M. S. T. Motta, "Theoretical and practical aspects of robot calibration with experimental verification," *Journal of the Brazilian Society of Mechanical Sciences and Engineering*, vol. 33, no. 1, pp. 15–21, 2011.

[87] M. L. Hornick and B. Ravani, "Computer-Aided Off-Line Planning and Programming of Robot Motion," *The International Journal of Robotics Research*, vol. 4, no. 4, pp. 18–31, 1986.

[88] F. Proctor, M. Franaszek, and J. Michaloski, "Tolerances and Uncertainty in Robotic Systems," in *Proceedings of ASME 2017 International Mechanical Engineering Congress & Exposition*. Tampa, FL: ASME, 2017.

[89] KUKA Roboter GmbH, "KUKA Robot Sensor Interface (RSI) 3.1 Documentation," 2010.

[90] Kuka Roboter GmbH, "KUKA FastSendDriver 4.0 Documentation," 2014.

[91] E. B. Dam, M. Koch, and M. Lillholm, "Quaternions , Interpolation and Animation," DTU, Lyngby, Denmark, Tech. Rep., 1998.

[92] International Standardization Organization (ISO), "BS EN ISO 9283:1998.: Manipulating industrial robots - Performance criteria and related test methods," 1998.

[93] M. Slamani, A. Nubiola, and I. Bonev, "Assessment of the positioning performance of an industrial robot," *Industrial Robot: An International Journal*, vol. 39, no. 1, pp. 57–68, 2012.

[94] R. Bernhardt and S. Albright, *Robot Calibration*. Springer Science & Business Media, 1993.

[95] M. Abderrahim, A. Khamis, S. Garrido, and L. Moreno, "Accuracy and Calibration Issues of Industrial Manipulators," *Industrial Robotics - Programming, Simulation and Applications*, no. December, pp. 131–147, 2006.

[96] S. Aoyagi, "Selection of Optimal Measuring Points Using Genetic Algorithm in the Process to Calibrate Robot Kinematic Parameters," in *Genetic Algorithms in Applications*, R. Popa, Ed. InTech, 2012, pp. 3–19.

[97] A. Nubiola and I. A. Bonev, "Absolute calibration of an ABB IRB 1600 robot using a laser tracker," *Robotics and Computer-Integrated Manufacturing*, vol. 29, no. 1, pp. 236–245, 2013.

[98] K. L. Conrad and T. C. Yih, "Robotic Calibration Issues: Accuracy, Repeatability and Calibration," in *8th Mediterranean Conference on Control & Automation*, no. July, 2000, pp. 17–19.

[99] T. Brogårdh, "Present and future robot control development—An industrial perspective," *Annual Reviews in Control*, vol. 31, no. 1, pp. 69–79, jan 2007.

[100] B. Rooks, "Automatic wing box assembly developments," *Industrial Robot: An International Journal*, vol. 28, no. 4, pp. 297–301, 2001.

[101] Albin Sunnanbo, "Laser feedback control for robotics in aircraft assembly," Master Thesis, Linkoping Universitet, 2003.

[102] J. E. Muelaner, Z. Wang, and P. G. Maropoulos, "Concepts for and Analysis of a High accuracy and High Capacity (HAHC) Aerospace Robot," in *The 21st International Computer-Aided Production Engineering Conference (CAPE)*, 2010.

[103] N. Metrology, "Adaptive Robot Control Brochure," 2018.

[104] S. Droll, "Real Time Path Correction of Industrial Robots with Direct End-Effector Feedback from a Laser Tracker," *SAE International Journal of Aerospace*, vol. 7, pp. 222–228, 2014.

[105] M. Summers, "Robot capability test and development of industrial robot positioning system for the aerospace industry," *SAE transactions*, vol. 114, no. 1, pp. 1108–1118, 2005.

[106] G. Mosqueira, J. Apetz, K. Santos, E. Villani, R. Suterio, and L. Trabasso, "Analysis of the indoor GPS system as feedback for the robotic alignment of fuselages using laser radar measurements as comparison," *Robotics and Computer-Integrated Manufacturing*, vol. 28, no. 6, pp. 700–709, dec 2012.

[107] M. Shah, R. D. Eastman, and T. Hong, "An Overview of Robot-Sensor Calibration Methods for Evaluation of Perception Systems," *Proceedings of the Workshop on Performance Metrics for Intelligent Systems - PerMIS '12*, p. 15, 2012.

[108] E. Boctor, A. Viswanathan, M. Choti, R. Taylor, G. Fichtinger, and G. Hager, "A novel closed form solution for ultrasound calibration," *2004 2nd IEEE Inter-*

national Symposium on Biomedical Imaging: Macro to Nano, vol. 2, pp. 527–530, 2004.

[109] L. E. Bø, E. F. Hofstad, F. Lindseth, and T. A. Hernes, "Versatile robotic probe calibration for position tracking in ultrasound imaging," *Physics in Medicine and Biology*, vol. 60, no. 9, pp. 3499–3513, 2015.

[110] F. Aalamifar, A. Cheng, Y. Kim, X. Hu, H. K. Zhang, X. Guo, and E. M. Boctor, "Robot-assisted automatic ultrasound calibration," *International Journal of Computer Assisted Radiology and Surgery*, vol. 11, no. 10, pp. 1821–1829, 2016.

[111] S. Yin, Y. Ren, Y. Guo, J. Zhu, S. Yang, and S. Ye, "Development and calibration of an integrated 3D scanning system for high-accuracy large-scale metrology," *Measurement: Journal of the International Measurement Confederation*, vol. 54, pp. 65–76, 2014.

[112] C. Holmes, B. W. Drinkwater, and P. D. Wilcox, "Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation," *NDT&E International*, vol. 38, pp. 701–711, 2005.

[113] W. Kerr, S. G. Pierce, and P. Rowe, "Investigation of synthetic aperture methods in ultrasound surface imaging using elementary surface types," *Ultrasonics*, vol. 72, pp. 165–176, 2016.

[114] TWI Ltd., "Advanced ultrasonic full matrix capture inspection software launched," 2016, Date accessed: 2017-09-06. [Online]. Available: https://www.twi-global.com/news-events/news/2016-09-advanced-ultrasonic-full-matrix-capture-inspection-software-launched/

[115] H. Zhang, C. Xu, and D. Xiao, "Offline correction of tool path deviations for robot-assisted ultrasonic nondestructive testing," *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, vol. 233, no. 8, pp. 2879–2893, 2019.

[116] P. J. Besl and N. D. McKay, "A Method for Registration of 3D-Shapes," *IEEE Transactions on Pattern Analysis and Machine Inteligence*, vol. 14, no. 2, pp. 239–256, 1992.

[117] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-Squares Fitting of Two 3-D Point Sets." *IEEE Trans. Pattern Analysis Mach. Intell. ()*, vol. 9, no. 5, pp. 698–700, 1987.

[118] M. A. Boden, *Mind as machine: A history of cognitive science.* Oxford University Press, 2008.

[119] M. A. Treiber, *An introduction to object recognition: selected algorithms for a wide variety of applications.* Springer Science & Business Media, 2010.

[120] Omron Microscan Systems Inc, "What Is Machine Vision?" 2015, Date accessed: 2015-06-01. [Online]. Available: https://www.microscan.com/en-us/resources/know-your-tech/what-is-machine-vision

[121] M. Sonka, V. Hlavac, and R. Boyle, *Image processing, analysis, and machine vision*, 3rd ed. Thompson Learning, 2008.

[122] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," *arXiv preprint arXiv:1409.1556*, pp. 1–14, 2014.

[123] A. Krizhevsky and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in neural information processing systems*, 2012, pp. 1097—-1105.

[124] R. K. Narla, T. Technology, and U. States, "The Evolution of Connected Vehicle Technology : From Smart Drivers to Smart Cars to ... Self-Driving Cars," *ITE Journal*, vol. 83, no. July, pp. 22–26, 2013.

[125] A. A. Taylor, S. Kessel, D. T. Mellott, B. R. Mills, and L. M. Walsh, "Electronic commerce functionality in video overlays," Oct. 11 2018, uS Patent App. 16/004,693.

[126] M. Stark, M. Goesele, and B. Schiele, "Back to the Future: Learning Shape Models from 3D CAD Data," in *Proceedings of the British Machine Vision Conference*. BMVA Press, 2010, pp. 106.1–106.11.

[127] D. Lowe, "Local feature view clustering for 3D object recognition," *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, vol. 1, pp. 682–688, 2001.

[128] D. Schmitt and N. McCoy, "Object Classification and Localization Using SURF Descriptors," pp. 1–5, 2011.

[129] S. Leutenegger, M. Chli, and R. Y. Siegwart, "BRISK: Binary robust invariant scalable keypoints," in *2011 IEEE International Conference on Computer Vision (ICCV)*. IEEE, 2011, pp. 2548–2555.

[130] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 1. IEEE, 2005, pp. 886–893.

[131] D. G. Lowe, "Object Recognition from Local Scale-Invariant Features," *Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2, no. 8, pp. 1150–1157, 1999.

[132] J. J. Lim, H. Pirsiavash, and A. Torralba, "Parsing IKEA Objects : Fine Pose Estimation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2013, pp. 2992–2999.

[133] S. Belongie, J. Malik, and J. Puzicha, "Shape Matching and Object Recognition Using Shape Contexts," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 24, pp. 509–522, 2002.

[134] I. Gordon and D. G. Lowe, *What and Where : 3D Object Recognition with Accurate Pose*. Springer Berlin Heidelberg, 2004.

[135] M. Aubry, D. Maturana, A. A. Efros, B. C. Russell, and J. Sivic, "Seeing 3D chairs : exemplar part-based 2D-3D alignment using a large dataset of CAD models," in *Proceedings of the IEEE conference on computer vision and pattern recognition*. IEEE, 2014, pp. 3762–3769.

[136] D. Collet, Alvaro; Berenson, Dmitry; Srinivasa, Siddhartha S.; Ferguson, "Object Recognition and Full Pose Registration from a Single Image for Robotic Manip-

ulation," in *2009 IEEE International Conference on Robotics and Automation*. Kobe, Japan: IEEE, 2009, pp. 48–55.

[137] B.-K. Seo and H. Wuest, "A Direct Method for Robust Model-Based 3D Object Tracking from a Monocular RGB Image," in *Computer Vision–ECCV 2016 Workshops*. Springer, 2016, pp. 551–562.

[138] C. Wöhler, *3D Computer Vision*. Springer-Verlag, 2013.

[139] S. Tulsiani and J. Malik, "Viewpoints and keypoints," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 07-12-June, pp. 1510–1519, 2015.

[140] B. Rosenhahn, "Pose Estimation Revisited," Ph.D. dissertation, Christian-Albrechts-Universitat, Kiel, 2003.

[141] J. Canny, "A computational approach to edge detection," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, no. 6, pp. 679–698, 1986.

[142] C. Harris and M. Stephens, "A combined corner and edge detector," in *Alvey Vision Conference*, 1988, pp. 147–151.

[143] M.-K. Hu, "Visual pattern recognition by moment invariants," *IRE Transactions on information Theory*, vol. 8, no. 2, pp. 179–187, 1962.

[144] D. Zhang and G. Lu, "Shape-based image retrieval using generic Fourier descriptor," *Signal Processing: Image Communication*, vol. 17, no. 10, pp. 825–848, 2002.

[145] B. Rosenhahn, T. Brox, and J. Weickert, "Three-Dimensional Shape Knowledge for Joint Image Segmentation," *International Journal of Computer Vision*, vol. 73, no. 3, pp. 243–262, 2007.

[146] G. Bradski and A. Kaehler, *Learning OpenCV: Computer vision with the OpenCV library*. O'Reilly Media Inc., 2008.

[147] R. C. Gonzalez and R. E. Woods, "Image processing," *Digital image processing*, vol. 2, 2007.

[148] R. Szeliski, *Computer Vision: Algorithms and Applications*, 1st ed. New York, N.Y.: Springer-Verlag New York, 2010.

[149] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *International journal of computer vision*, vol. 22, no. 1, pp. 61–79, 1997.

[150] T. Brox, B. Rosenhahn, and J. Weickert, "Three-Dimensional Shape Knowledge for Joint Image Segmentation and Pose Tracking," *International Journal of Computer Vision*, vol. 73, no. 3, pp. 243–262, 2007.

[151] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.

[152] M. Zollhöfer, P. Stotko, A. Görlitz, C. Theobalt, M. Nießner, R. Klein, and A. Kolb, "State of the Art on 3D Reconstruction with RGB-D Cameras," in *Computer Graphics Forum*, vol. 37, no. 2, 2018, pp. 625–652.

[153] Stereolabs, "ZED Stereo Camera," Date accessed: 2018-11-28. [Online]. Available: https://www.stereolabs.com/zed/

[154] H. Gonzalez-jorge, B. Riveiro, E. Vazquez-fernandez, J. Martínez-sánchez, and P. Arias, "Metrological evaluation of Microsoft Kinect and Asus Xtion sensors," *Measurement*, vol. 46, no. 6, pp. 1800–1806, 2013.

[155] A. Corti, S. Giancola, G. Mainetti, and R. Sala, "A metrological characterization of the Kinect V2 time-of-flight camera," *Robotics and Autonomous Systems*, vol. 75, pp. 584–594, 2016.

[156] C. Bishop, *Pattern Recognition and Machine Learning*. Springer-Verlag, 2006.

[157] K. M. Hanson, "Introduction to Bayesian image analysis," in *Medical Imaging: Image Processing*, vol. 731. International Society for Optics and Photonics, 1993, pp. 716–731.

[158] MATLAB, "Computer Vision Toolbox (R2016b)," 2016.

[159] A. Aldoma, M. Vincze, N. Blodow, D. Gossow, S. Gedikli, R. B. Rusu, and G. Bradski, "CAD-model recognition and 6DOF pose estimation using 3D cues,"

in *Proceedings of the IEEE International Conference on Computer Vision*, 2011, pp. 585–592.

[160] B. Rosenhahn, T. Brox, D. Cremers, and H. Seidel, "A Comparison of Shape Matching Methods for Contour Based Pose Estimation," in *International Workshop on Combinatorial Image Analysis*, no. June. Springer, 2006, pp. 263–276.

[161] M. Zhu, K. G. Derpanis, Y. Yang, S. Brahmbhatt, M. Zhang, C. Phillips, M. Lecce, and K. Daniilidis, "Single image 3D object detection and pose estimation for grasping," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. Hong Kong, China: IEEE, 2014, pp. 3936–3943.

[162] KUKA Roboter GmbH, "KR 16 L6-2 Technical Specification," 2003.

[163] ——, "KST RoboTeam 2.0 V4," 2014.

[164] Peak NDT, "MicroPulse 5PA," 2018. [Online]. Available: https://www.peakndt.com/products/micropulse-5pa/

[165] Leica Geosystems, "Leica Absolute Tracker AT901. Product Brochure," 2010.

[166] FARO Technologies Inc., "FARO Vantage," 2013.

[167] J. Conditt, "100 years of motion-capture technology," *Engadget UK*, feb 2018, Date accessed: 2018-08-20. [Online]. Available: https://www.engadget.com/2018/05/25/motion-capture-history-video-vicon-siren/

[168] H. Liu, C. Holt, and S. Evans, "Accuracy and repeatability of an optical motion analysis system for measuring small deformations of biological tissues." *Journal of biomechanics*, vol. 40, no. 1, pp. 210–214, jan 2007.

[169] P.-F. Yang, M. Sanno, G.-P. Bruggemann, and J. Rittweger, "Evaluation of the performance of a motion capture system for small displacement recording and a discussion for its application potential in bone deformation in vivo measurements," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, vol. 226, no. 11, pp. 838–847, jul 2012.

[170] R. Summan, S. G. Pierce, C. N. Macleod, G. Dobie, T. Gears, W. Lester, P. Pritchett, and P. Smyth, "Spatial calibration of large volume photogrammetry based metrology systems," *Measurement*, vol. 68, pp. 189–200, 2015.

[171] M. Morozov, J. Riise, R. Summan, S. Pierce, C. Mineo, C. MacLeod, and Brown R.H., "Assessing the Accuracy of Industrial Robots through Metrology for the enhancement of Automated Non-Destructive Testing," in *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, vol. 44. Baden-Baden: IEEE, 2016, pp. 335–340.

[172] S. Umeyama, "Least-Squares Estimation of Transformation Parameters Between Two Point Patterns," *Transactions on Pattern Analysis & Machine Intelligence*, vol. 13, no. 4, pp. 376–380, 1991.

[173] L. Wu and H. Ren, "Finding the Kinematic Base Frame of a Robot by Hand-Eye Calibration Using 3D Position Data," *IEEE Transactions on Automation Science and Engineering*, vol. 14, no. 1, pp. 314–324, 2016.

[174] T. Kubela, A. Pochyly, and V. Singule, "Investigation of Position Accuracy of Industrial Robots and Online Methods for Accuracy Improvement in Machining Processes," in *2015 International Conference on Electrical Drives and Power Electronics (EDPE)*. IEEE, 2015, pp. 385–388.

[175] E. W. Weisstein, "Point-Line Distance–3-Dimensional," Date accessed: 2018-06-15. [Online]. Available: http://mathworld.wolfram.com/Point-LineDistance3-Dimensional.html

[176] C. Mineo, M. Morozov, G. Pierce, I. Nicholson, and I. Cooper, "Computer-Aided Tool Path Generation for Robotic Non-Destructive Inspection," in *52nd Annual Conference of the British Institute of Non-Destructive Testing 2013: NDT 2013*, 2013.

[177] Z. Lu, C. Xu, X. Zhao, L. Zhang, H. Wang, and Q. Pan, "Ultrasonic transmission testing of twin-robots coordinated control," *2013 IEEE International Conference on Mechatronics and Automation, IEEE ICMA 2013*, pp. 1256–1260, 2013.

[178] M. Schwabe, A. Maurer, and R. Koch, "Ultrasonic Testing Machines with Robot Mechanics – A New Approach to CFRP Component Testing," in *2nd International Symposium on NDT in Aerospace*, 2010, pp. 1–5.

[179] M. Á. Nieto, Z. Re, and E. Cuevas, "TTU Phased Array : Quality and Productivity," in *6th International Symposium on NDT in Aerospace*, no. November, Madrid, 2014, pp. 12–14.

[180] Y. Gan, X. Dai, and J. Li, "Cooperative path planning and constraints analysis for master-slave industrial robots," *International Journal of Advanced Robotic Systems*, vol. 9, 2012.

[181] T. Möller and B. Trumbore, "Fast, minimum storage ray/triangle intersection," in *ACM SIGGRAPH 2005 Courses*. ACM, 2005, p. 7.

[182] Wikipedia, "Octree," Date accessed: 2016-05-02. [Online]. Available: https://en.wikipedia.org/wiki/Octree

[183] Weidlinger Associates Inc, "PZFlex," Date accessed: 2016-07-02. [Online]. Available: http://www.pzflex.com/

[184] EXTENDE, "EXTENDE - CIVA Simulation Software," 2013, Date accessed: 2016-06-11. [Online]. Available: http://www.extende.com/

[185] J. Li, M. Chen, X. Jin, Y. Chen, Z. Dai, Z. Ou, and Q. Tang, "Calibration of a multiple axes 3-D laser scanning system consisting of robot, portable laser scanner and turntable," *Optik*, vol. 122, no. 4, pp. 324–329, 2011.

[186] B. Delaunay and Others, "Sur la sphere vide," *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, vol. 7, no. 793-800, pp. 1–2, 1934.

[187] M. Kazhdan and H. Hoppe, "Screened poisson surface reconstruction," *ACM Transactions on Graphics (ToG)*, vol. 32, no. 3, p. 29, 2013.

[188] C. Mineo, J. Riise, R. Summan, C. N. MacLeod, and S. G. Pierce, "Index-based triangulation method for efficient generation of large three-dimensional ultrasonic C-scans," *Insight - Non-Destructive Testing and Condition Monitoring*, vol. 60, no. 4, 2018.

[189] P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, and G. Ranzuglia, "MeshLab: an Open-Source Mesh Processing Tool," in *Eurographics Italian Chapter Conference*, V. Scarano, R. D. Chiara, and U. Erra, Eds. The Eurographics Association, 2008.

Appendix A

Representing Pose in 3D

Any rotation in Cartesian space can be expressed as a combination of three unit rotations, i.e. rotations about each axes of a coordinate system. Rotations in 3D are generally not commutative, so the order in which the rotations are applied is important. The KUKA control systems uses the ZYX Euler convention, such that rotations are applied first around the z-axis, then about the (now rotated) y-axis and finally about the x-axis. The angles, expressed in degrees, used by KUKA are denoted as angles A,B and C, corresponding to the ZYX rotations.

A rotation about the z-axis, by α degrees, followed by a rotation of β degrees about the y-axis, and terminated by a rotation about the x-axis of γ degrees, can be described by the combined rotation matrices shown in Equation A.1:

$$R(\alpha, \beta, \gamma) = R(\alpha)R(\beta)R(\gamma) \tag{A.1}$$

Expanding each of the rotation matrices leads to Equation A.2.

$$
\begin{aligned}
&= \begin{bmatrix} \cos\gamma & -\sin\gamma & 0 & 0 \\ \sin\gamma & \cos\gamma & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos\beta & 0 & \sin\beta & 0 \\ 0 & 1 & 0 & 0 \\ -\sin\beta & 0 & \cos\beta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\alpha & -\sin\alpha & 0 \\ 0 & \sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\
&= \begin{bmatrix} \cos\beta\cos\gamma & \cos\gamma\sin\alpha\sin\beta - \cos\alpha\sin\gamma & \cos\alpha\cos\gamma\sin\beta + \sin\alpha\sin\gamma & 0 \\ \cos\beta\cos\gamma & \cos\alpha\cos\gamma + \sin\alpha\sin\beta\sin\gamma & -\cos\gamma\sin\alpha + \cos\alpha\sin\beta\sin\gamma & 0 \\ -\sin\beta & \cos\beta\sin\alpha & \cos\alpha\cos\beta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}
\end{aligned} \tag{A.2}
$$

Alternatively, it is also possible to construct a rotation matrix from three orthogonal vectors, assuming the normal points in the z-direction. The rotation matrix is thus constructed as shown in Equation A.3. The rotation matrix in Equation A.3 consists of three columns wherein the third column is the normal vector, the second column is the tangential vector, and the first is the cross-product of the two others.

$$
R = \begin{bmatrix} R_{11} & R_{12} & R_{13} \\ R_{21} & R_{22} & R_{23} \\ R_{31} & R_{32} & R_{33} \end{bmatrix} = - \begin{bmatrix} V_X & T_X & N_X \\ V_Y & T_Y & N_Y \\ V_Z & T_Z & N_Z \end{bmatrix} \tag{A.3}
$$

To reverse this operation and move from a rotation matrix to the Euler-angle representation, it is necessary to use the inverse trigonometric functions. However, both *arccos* and *arcsin* are only defined in the interval from 0 to π. To cover the entire range from $]-\pi, \pi]$, both sin and cos are used by defining a function: $\mathrm{sgn}(x) = \frac{x}{|x|}$ and sng(0) = 0. Thus $v = \mathrm{sng}(\sin(v))\arccos(\mathrm{v})$. From Equation A.2 it can be readily identified that $\sin\beta = -R_{31}$. The following relations can then be generated to calculate angles A,B and C:

$$
B = atan2\left(\frac{-R_{31}}{\sqrt{R_{11}^2 + R_{21}^2}}\right) \tag{A.4}
$$

$$A = \begin{cases} 0, & \text{for } |B| = \frac{\pi}{2}; \\ atan2\left(\frac{R_{21}}{R_{11}}\right), & \text{for } |B| \neq \frac{\pi}{2}; \end{cases} \tag{A.5}$$

$$C = \begin{cases} \frac{B}{|B|} \cdot atan2(R_{12}), & \text{for } |B| = \frac{\pi}{2}; \\ atan2\left(\frac{R_{32}}{R_{33}}\right), & \text{for } |B| \neq \frac{\pi}{2}; \end{cases} \tag{A.6}$$

A.1 Combining Orientation and Translation

To be able to express the pose of a rigid body in a 3D coordinate system, a compact notation representing both orientation and translation is needed. Two popular representations are homogeneous transformation matrices and quaternion vector pairs. The homogeneous transformation matrix is used extensively throughout this thesis due to its intuitive appeal. A homogeneous transformation matrix expressing the coordinates of point B, in the reference frame of A, is given by Equation A.7.

$$\begin{bmatrix} {}^{A}x \\ {}^{A}y \\ {}^{A}z \\ 1 \end{bmatrix} = \begin{pmatrix} {}^{A}\mathbf{R}_{B} & \mathbf{t} \\ 0_{1X3} & 1 \end{pmatrix} \begin{bmatrix} {}^{B}x \\ {}^{B}y \\ {}^{B}z \\ 1 \end{bmatrix} \tag{A.7}$$

The translation vector between the origin of the reference frames A and B is denoted $\mathbf{t}$ and orientation is denoted by the 3x3 orthonormal matrix $\mathbf{R}$, as shown in Equation A.1. The combination of rotation and translation into a single matrix creates the 4x4 homogeneous transformation matrix, commonly denoted as ${}^{A}\mathbf{T}_{B}$.

A.2 Quaternion Algebra

One of the drawbacks of using Euler angles is that the same representation of an angle can be arrived at in a number of different ways. It is also possible to arrive at singularities when two axes are aligned (also known as gimbal-lock) such that there is no way to express rotation about the "lost" axis. To avoid this situation, the quaternion

formulation can be used instead wherein two different orientations can be described by a *single* rotation about some axis in 3D space. By convention, the angle is denoted θ and the vector, $\mathbf{v}$. Determining exactly which way this vector points is not intuitive, but it helps to think of a the direction the sound is emitted from an ultrasonic probe, for example.

A quaternion is an extension of a complex number and written as a scalar plus a vector. Equation A.8 expresses using the fundamental quaternion units i,j and k. These units are defined such that $i^2 = j^2 = k^2 = ijk = -1$, making them orthogonal complex numbers. Rotations are typically expressed with unit quaternions who are of unit magnitude such that $s^2 + v_1^2 + v_2^2 + v_3^2 = 1$.

$$\dot{q} = s + \mathbf{v}$$
$$\dot{q} = s + v_1 i + v_2 j + v_3 k \tag{A.8}$$

To move between representation orientation using quaternion and Euler angles, Rodrigues' rotation formula (Equation A.9) can be used. In Equation A.9, $\mathbf{S}$ denotes a skew-symmetric matrix.

$$\mathbf{R} = \mathbf{I}_{3X3} + \sin\theta\mathbf{S}(\mathbf{v}) + (1 - \cos\theta)(\mathbf{v}\mathbf{v}^T - \mathbf{I}_{3X3}) \tag{A.9}$$

Orientation and translation can be combined using a translation vector and a quaternion, such that $\xi = (\mathbf{t}, \dot{q})$, where $\mathbf{t}$ is a regular translation vector in 3D and $\dot{q}$ is a unit quaternion.

Appendix B

Synchronised Robot Motion

Unlike GeoLinked motion, which leaves the slave unable to move independently, the `SyncCmd` tag allows both robots to move along different paths, but in the same time interval. To accomplish this, the two robots must have RoboTeam 2.0 (or later) installed. Using the expert level syntax, it is possible to include a different synchronisation tag for each motion step to form a sequence of synchronised motions. The tags are necessary to tell each robot which motions must be performed in the same time interval. The tag must be identical in each program (for each robot) for the interpreter to execute the motion.

The use of the `SyncCmd` allows each interpreter to change the velocity of either robot, dependent on the path length. A MATLAB script was written which outputs points with the correct KRL (KUKA Robot Language) syntax into two files. The MATLAB programs also output the expected signals needed for communication with the RSI (Robot Sensor Interface) server, as defined by the IntACom software. These programs can be downloaded to the KRC (KUKA Robot Controller) and executed as regular KRL programs, alleviating the need for external control.

As mentioned in Chapter 6, path planning must still be performed using dedicated software, such as FASTSURF. To obtain the points from FASTSURF, the CENIT-generated KRL files need to be parsed to obtain xyz coordinates in MATLAB. Two new files then have to be written to achieve synchronised motion; one for the master and one for the slave robot. The reading of KRL programs is achieved through a text

scan function in MATLAB which is a standard feature of most programming languages. Besides reading the desired coordinates, the parser also identifies the type of motion (linear -LIN, or circular -CIRC), such that a combination of both motion types can be used. A lookup-table is created in MATLAB to store the motion type for each point along with its coordinates and orientation.

Two separate paths with equal number of points are created once the original path has been projected through the CAD model and the new orientations have been defined for the slave robot. To write the programs to a KRL file, the MATLAB program uses both sets of coordinates. A script to write KRL files was designed which uses the KUKA expert-level user syntax LIN and CIRC commands, along with the required robot pose. Using the lookup-table, the script can output a combination of LIN and CIRC motions, providing a smoother motion for complex geometries. The script also monitors the direction of motion and can automatically add signal tags before and after the raster steps at the end of each scan line. Each pose contains a unique `SyncCmd()` identifier to activate the RoboTeam's synchronisation method.

Figure B.1 shows an example output of the MATLAB-generated KRL code. The C_DIS command is an important parameter which is used to obtain a smooth motion. The default motion in the KUKA operating system is a motion that reaches a commanded pose and stops momentarily before proceeding to the next motion. This is undesirable during an ultrasonic scan as it increases the scan speed and causes an uneven sampling of the surface. To avoid stopping at every point on the path, the robots can be told to approximate points along the path. The approximation algorithm is triggered either by a velocity, position or angular criteria. The current implementation uses a distance criterion where the approximation begins once the robot is within 5mm of the intended pose. This results in a smooth motion along the surface of the component while staying close to the intended path.

Besides the list of coordinates and synchronisation identifiers, the KRC script also adds additional lines of code for initialisation and setting up the data server interface to communicate with the external PC. As each pose is defined without tool, base or robot configuration information, a few points have to be added to the program manually

```
SyncCmd Example KRL Module

PTP HOME

$VEL.CP = 0.4 ;TCP velocity
$ACC.CP=1 ;TCP acceleration
$VEL.ORI1=10 ;Orientational velocity
$ACC.ORI1=5 ; Orientational acceleration
$APO.CDIS=5 ; Start approximation at 5mm
...
LIN {X 396.937910,Y 258.326834,Z -30.746092,A -1.695254,B
7.693063,C -15.941791} C_DIS
SyncCmd(#MotionSync,"M1",3)
LIN {X 392.117802,Y 258.415469,Z -30.285934,A -1.632506,B
7.425511,C -15.925662} C_DIS
SyncCmd(#MotionSync,"M2",3)
LIN {X 387.402120,Y 258.473124,Z -29.945494,A -1.500565,B
6.861853,C -15.887828} C_DIS
SyncCmd(#MotionSync,"M3",3)

PTP HOME

END
```

Figure B.1: Example output from MATLAB. Each commanded position has a unique identifier (i.e. "M1") to synchronise the motion between the two robots. The C_DIS parameter determines the path smoothness, velocity and acceleration.

during setup. This allows the operator to set approaching and retracting motions depending on the physical set-up in the robotic cell. It is of crucial importance that the operator chooses the correct tool and base calibrations for each robot during this set-up phase. As each pose is only relative to the previous points, the incorrect definition of tool and base parameters could have undesired and expensive consequences.