Autonomous Inspection of Composite Materials with Multiple Intelligent Robots through Thermal Imaging and Machine Learning

Andrew Loeliger

Department of Design, Manufacturing and Engineering Management

University of Strathclyde

Glasgow

March 2026

A thesis presented in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

Abstract

Composite materials have become key to the development of modern lightweight manufacturing, with applications in aerospace, automotive, marine, and defence industries. Their superior strength-to-weight ratio, durability, and corrosion resistance make them indispensable, yet their production remains dominated by manual processes. Manual hand lay-up techniques are time-consuming, hazardous to workers, and prone to defects that compromise the mechanical and cosmetic performance of final components. Automated Fibre Placement (AFP) machines offer an alternative by automating the layup of pre-impregnated carbon fibre tows. However, AFP remains susceptible to defects such as gaps, overlaps, wrinkles, twists, and misalignments, which can significantly reduce structural integrity. Furthermore, most existing defect detection approaches rely on post-process inspection, meaning defects are only discovered once a part has already been completed, resulting in wasted materials, higher costs, and missed opportunities for real-time repair. These challenges establish the imperative for a robust, flexible, and cost-effective in-process defect detection system that ensures higher quality, reduces waste, and increases industrial use of AFP.

To this end, this thesis investigates the integration of thermal imaging and machine learning to achieve automated, real-time AFP defect detection. A review of existing defect detection techniques was first undertaken, identifying current methods such as visual inspection, profilometry, and ultrasonic testing, alongside their limitations in detecting in-situ and hidden defects. For this work, a dual-robot experimental platform was developed in which a KUKA KR120 industrial AFP robot carries out layup operations, while a UR10e collaborative robot, equipped with an AMETEK LAND ARC thermal camera, follows in real time to capture thermal imagery of the process. The resulting thermal data was collated, cleansed, segmented, and augmented into a dedicated dataset, which was then used to train and evaluate a series of machine learning and deep learning models.

The experimental evaluation considered four machine learning models: K-Nearest Neighbours, Support Vector Machines, Decision Trees, and Naïve Bayes as well as two deep learning models, ResNet101 and ResNet152. Comparative analysis

across metrics such as accuracy, precision, recall, F1-score, and Receiver Operating Characteristic Area Under the Curve (ROC AUC) demonstrated that ResNet101 offered the most balanced and effective performance, achieving an accuracy of 78.57% and a ROC AUC score of 0.82. This outcome validates the feasibility of employing deep learning based thermographic analysis for real-time AFP defect detection. In addition to model evaluation, this thesis addresses the issue of system flexibility by designing and building a Smart Sensor Box (SSB), a modular and low-cost sensing platform with accelerated machine learning that allows fast integration of different sensors into production work cells. The SSB lowers both the technological and financial barriers for smaller manufacturers. Furthermore, multi-robot coordination was explored using RoboDK simulations, demonstrating the potential for end-to-end collaborative systems that integrate AFP robots and inspection cobots into a single defect detection framework.

The contributions of this thesis are: a comprehensive literature review identifying gaps in real-time AFP defect detection research; the creation and use of a custom new thermal imaging dataset tailored to AFP defect detection; the design, development, and validation of a low-cost Smart Sensor Box for flexible sensing integration; and the implementation of a multi robot system that combines thermal imaging and machine learning for in-process defect detection. These contributions of knowledge in automated composite manufacturing provide both theoretical insights and practical tools that enhance quality, reduce waste, and enable scalable adoption of AFP defect detection technologies across the lightweight manufacturing industry.

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.

Signed:

Date: 13/03/2026

Acknowledgements

First, I would like to express my gratitude to my supervisor Dr. Erfu Yang of the University of Strathclyde Design, Manufacturing and Engineering Management department. His knowledge and experience in academic research helped me find a way to finish this thesis. I am thankful for the faith he had in me to pursue a PhD.

I would like to also thank Mr. Mohamed Adlan Ait Ameur who provided support during the development of the machine learning and deep learning models.

My thanks also go to my second supervisor Iain Bomphray at the Lightweight Manufacturing Centre. Access to the robots and work cell at the centre were essential for my research. I would also like to thank Mr. Stig McArthur and Ms. Catherine Yokan who provided support with the Lightweight Manufacturing Centre equipment and setup.

Last but not least, I would like to express my gratitude to my family and friends. Especially my parents who have supported me through all of my research and write up.

Table of Contents

List of Figures

List of Tables

List of Acronyms

ADC	Analogue to Digital Converter
AI	Artificial Intelligence
AFP	Automated Fiber Placement
AUC	Area Under the Curve
CNN	Convolutional Neural networks
CPU	Central Processor Unit
GUI	Graphical User Interface
GPU	Graphics Processing Unit
IPC	Industrial Personal Computer
ISAAC	Integrated Structural Assembly of Advanced Composites
ISO	International Organization for Standardization
KNN	K-Nearest Neighbors
LMC	Lightweight Manufacturing Centre
MAR	Multi Autonomous Robots
ML	Machine Learning
NMIS	National Manufacturing Institute Scotland
PCB	Printed Circuit Board
PCIe	Peripheral Component Interconnect Express
PLC	Programmable Logic Controller
RB5	Robot Board 5
ROC	Receiver Operating Characteristic
ROS2	Robot Operating System 2
SSB	Smart Sensor Box
SVM	Support Vector Machine
TOPS	Tera Operation Per Second
TPU	Tensor Processor Unit
UI	User Interface
UR	Universal Robotics

Publications

Conference proceedings

A. Loeliger, E. Yang and I. Bomphray, "A smart sensor box to increase the adaptability of automated manufacturing" in Advances in Transdisciplinary Engineering. 25, p. 31-36, 8 Nov 2022.

A. Loeliger, E. Yang and I. Bomphray, "An overview of automated manufacturing for composite materials", in "2021 26th International Conference on Automation and Computing (ICAC)", Piscataway, NJ: IEEE, p. 1-6, 15 Nov 2021.

Journals

A. Loeliger, E. Yang, and M. Ameur, "In-process Defect Detection for Automated Fiber Placement Using Thermal Imaging and Machine Learning", submitted to NTD&E International Journal March 2026.

Other publications

B. Yang, E. Yang, L. Yu, and **A. Loeliger**, "High-Precision UWB-Based Localisation for UAV in Extremely Confined Environments," in IEEE Sensors Journal, vol. 22, no. 1, pp. 1020-1029, Jan.1, 2022, doi: 10.1109/JSEN.2021.3130724.

L. Yu, E. Yang, B. Yang, **A. Loeliger** and Z. Fei, "Stereo Vision-based Autonomous Navigation for Oil and Gas Pressure Vessel Inspection Using a Lowcost UAV," 2021 IEEE International Conference on Real-time Computing and Robotics (RCAR), 2021, pp. 1052-1057, doi: 10.1109/RCAR52367.2021.9517584.

Chapter 1

1 Introduction

A composite material is formed by combining two or more constituent components. These components possess different properties which together create a product with enhanced performance for a specific application [1]. Each material remains distinct from the other however the amalgamation of the separate elements allows for improved properties such as increased strength, reduced weight and the ability to develop components with complex shapes.

1.1 Research Background and Motivation

Composite materials are essential in lightweight manufacturing as they allow the development of lighter materials that are stronger and more durable. Carbon fibre composite is a composite material made up of a matrix of carbon fibres and a reinforcement of resin [2]. The carbon fibres provide lightweight strength while the resin holds the fibres together in the required shape. Carbon fibre composite materials are used in many industries such as the automotive, aerospace and marine sectors where they benefit from the lightweight strength and corrosion resistance features the materials provide.

The composite material and carbon fibre composite material markets have been growing continually for several years and are expected to maintain this trajectory for the foreseeable future. The graph in Figure 1.1 [3] shows the recent and predicted growth of composite materials globally while the graph in Figure 1.2 [4] shows data for carbon fibre material specifically.

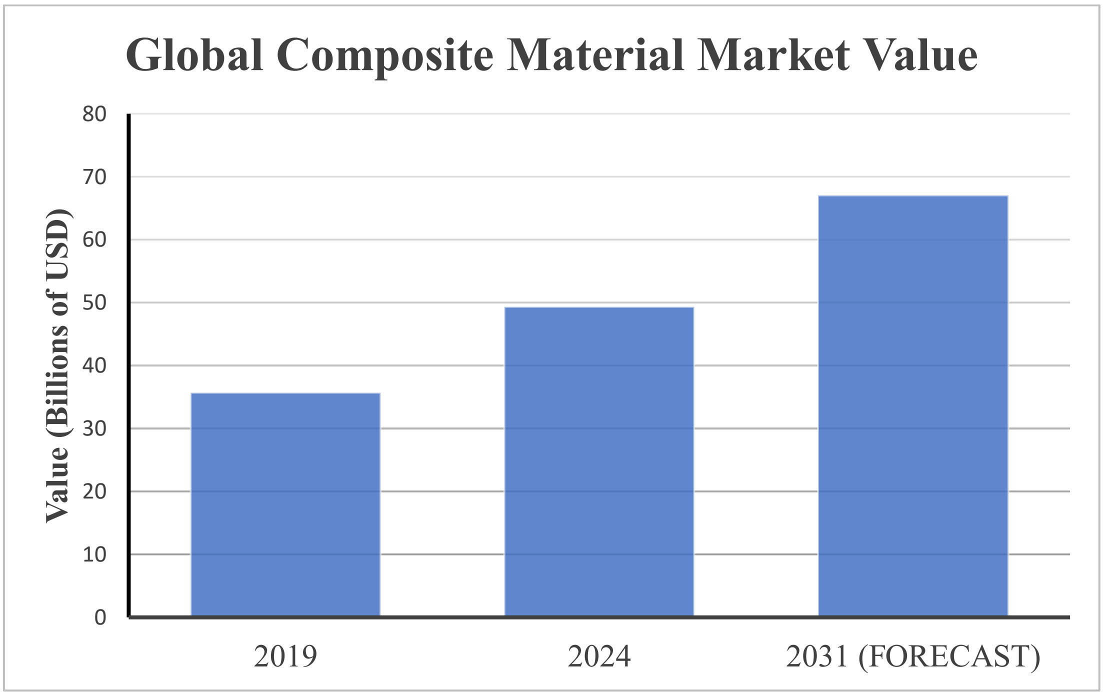

Figure 1.1 – Composite Worldwide Market

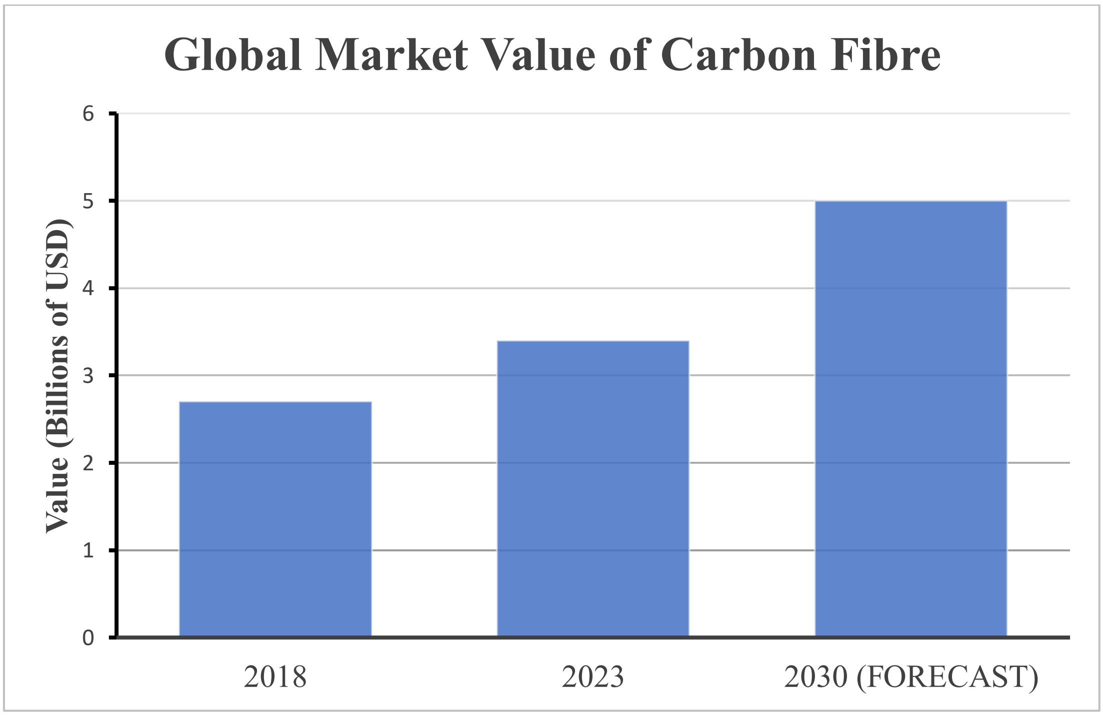

Figure 1.2 – Carbon Fibre Worldwide Market

The composite material market in general is expected to grow by 2031 to 67 billion US dollars while the carbon fibre composites market is predicted to reach 5 billion US dollars by 2030. The growth shows that this is an area of global development with opportunities for innovation and research.

Carbon fibre composite materials can be created in a manual or an automated process. Manual hand lay-up methods involve a person physically combining the required carbon fibre and resin. This is done by layering the carbon fibre segments on top of each other with layers of resin sandwiched in-between. The resin impregnates the fibres, and the layers are pushed together to ensure the resin is evenly distributed throughout the matrix. This is done by hand rollers repeatedly moved across the surface of each composite material layer. Automated methods involve using a machine to carry out the composite material creation in lieu of a person. AFP (Automated Fiber Placement) machines are an example of these systems. AFP machines work by using a robot arm to place a strip, known as a tow, on a build surface. Each tow is composed of carbon fibre material which is pre-impregnated with resin. When the tows are placed the robot arm applies the pressure required for the composite material creation.

Despite the safety issues associated with the manual hand lay-up method, it is still used for low volume, custom production requirements. Those working in the manual creation of carbon fibre composites require suitable protective equipment to limit the exposure and potential harm from interacting with the potentially toxic materials which can cause dermatitis, respiratory issues and eye problems [5]. Clearly moving to automated methods of composite material creation would be beneficial, as it removes the human from the equation, however the equipment expense can prove too high for many smaller companies. The cost to set-up AFP equipment can be too high for low volume manufacturing.

To reduce the costs involved in AFP it is essential to have a good method of defect detection. There are several methods being considered in this area and many have produced good accuracy. The majority, however, involve post-process inspection meaning composite material parts are fully laid before defects can be located and identified. Having a real-time defect detection method would allow for in-process inspection to be carried out. This would reduce waste, cut manufacturing time and eliminate the need to remake defective parts. Finding faults quickly, during the AFP process, ensures the manufacture of a defective part is stopped immediately rather than continuing to completion. This allows the opportunity to rework or repair components rather than scrapping them completely. Catching faults early not only increases quality, safety and throughput but also reduces waste, manufacturing time and cost.

This PhD research focusses on the proof of concept of, and development of, a novel, effective, low-cost approach to AFP defect detection. A multi-robot solution enables real-time defect detection using machine learning and thermography to identify and classify faults. A primary robot is used to lay the composite material tows while a secondary cobot, with thermal imaging camera, chases the robot allowing images to be taken from optimal angles. The distinction between this method and others, currently in production or proposed, is the in-situ chase feature that allows real-time detection while tows are being laid.

1.2 Research Aim, Objectives and Scope

1.2.1 Research Aim

The aim of this research is to produce adaptive control of Multi Autonomous Robots (MARs) for the Lightweight Manufacturing industry resulting in a fast, flexible and cost-effective solution for autonomous and collaborative defect detection in the manufacturing of composite material products.

1.2.2 Research Objectives

Ultimately, this research will work towards resolving challenges encountered in the lightweight manufacturing industry by creating a solution that increases flexibility, autonomy and safety. To this end, the research questions that will need to be addressed are:

Research Objective One: Carry out a comprehensive overview and review of current composite material AFP defect detection methods.

Research Objective Two: Achieve an effective, low-cost automated method of composite material AFP defect detection implementing machine learning and thermography.

Research Objective Three: Investigate how multiple robots can be automated to work collaboratively to fully automate and optimise AFP defect detection.

1.2.3 Research Scope

The scope of this thesis is limited to the development and evaluation of a proof of concept system for real-time defect detection in Automated Fibre Placement composite manufacturing. The research focuses on the integration of thermal imaging, machine learning, a Smart Sensor Box and multiple robot operation to identify AFP defects during production. The work considers experimental validation using a dedicated thermal image dataset, RoboDK simulation and a multi robot AFP inspection setup. It does not seek to provide a full industrial production system, closed loop repair process or complete inspection solution for all composite manufacturing methods. Instead, this research demonstrates the feasibility of an adaptable, low-cost, in-process, real-time defect detection framework.

1.3 Research Methodology

The research methodology adopted in this work was chosen to enable the understanding of defect detection in the area of composite materials. The methodology is further intended to identify knowledge gaps in the industry to allow the hypothesis of a unique approach that can be proven to be an effective solution through simulation and testing.

The methodology, shown in Figure 1.3, is not based on any specific pre-identified procedures but instead focussed on techniques that have been chosen to optimise the requirements of this research.

Figure 1.3 – Thesis Methodology

The key stages of the methodology can be explained as follows:

Research Proposal:

An initial investigation into the field of composite materials was carried out to identify a suitable area of research focus. This process highlighted the importance of effective and accurate defect detection in composite material manufacturing. In particular, it revealed a need to concentrate on Automated Fiber Placement (AFP) machines, which are widely used in industry and are a role in high-performance composite production.

The AFP machines could benefit from automated defect detection to increase consistency and speed of production. A new optimised computation and control system would be required to implement the defect detection.

Hypothesis Formulation:

A detailed analysis into this area of production was carried in the form of a literature review. The review identified current solutions, available, techniques used, and issues commonly encountered in the industry. The literature considered also identified specific needs and knowledge gaps that could be addressed. The analysis also enabled the hypothesis of a possible novel solution to deal with the problem of defect detection. The investigation resulted in the definition of an aim and objectives, in the form of three research questions, based on identified areas with limited or no current solution.

Research Philosophy, Strategy and Approach:

The research in this PhD is aiming to identify, and develop, a novel solution for defect detection in the process of AFP of composite materials. As the theoretical proposal would result in a solution that requires a high degree of accuracy for safe production of composite parts, it was deemed that a very practical methodology was required to tackle the process. A positivism philosophy was therefore chosen as, in comparison to the other main research philosophies of interpretivism, pragmatism and constructivism, it is the method that relies almost entirely on quantifiable and logical testing [6]. Likewise, a deductive approach using quantitative data was chosen [7]. By focussing on actual data from testing, to prove or disprove the hypotheses put forward, it was felt that a unique and flexible, yet reliable and effective, solution could be developed.

Data Collection and Analysis:

It was determined that initial testing would be carried out using simulations developed for the KUKA KR120 robot and the Universal Robotics (UR) UR10e cobot. As these pieces of equipment are extremely expensive, simulations could provide fundamental information on the collaborative working of the robot and cobot in a safe environment. A thermal camera could then be used to collect thermal images to create a data set for training machine learning models. It was decided that a working prototype would be necessary to fully test the hypotheses. This testing would give actual results in-situ with known parameters that could be replicated to not only prove proof of concept but to fully understand the proposal's limitations.

Thesis Completion:

The final section of the PhD collates all research, papers and test results to write and provide analysis on the complete research findings and proposed future work.

1.4 Contributions to Knowledge

Research identified a number of needs within carbon fibre composite material AFP defect detection that were previously unmet. In proposing and developing solutions for these needs, the following contributions have been made:

- A review of existing literature on Automated Fiber Placement defect detection has been conducted, highlighting the current methods in use while also exposing areas that remain underdeveloped and lacking in research.
- The research presents the design, development, and proof of concept of a Smart Sensing Box (SSB), which provides a flexible solution for AFP defect detection. Unlike existing approaches, the SSB is distinguished by its combination of effectiveness with machine learning and its low cost, making it not only accessible but also adaptable to different AFP manufacturing environments. The configurable standard software interfaces provide simple integration in production environments. This contribution is significant as it lowers the technological and financial barriers to implementing defect detection, thereby facilitating wider industrial adoption and enabling more consistent quality assurance across varying production scales.
- A real-time, multi-robot defect detection system has been developed, which integrates machine learning techniques with thermography to achieve accurate identification of defects. In order to provide optimised thermal images 2 robots were used: one to lay the AFP tows and one to capture the thermal images for defect detection. Machine learning used on a local compute system allowed fast accurate identification of defects. This contribution demonstrates the feasibility of advanced, collaborative robotic inspection systems that enhance manufacturing efficiency and reliability.
- First use of a new dataset of classified thermal images of AFP composite material developed for defect detection processing. This is significant because there are very few datasets of this type publicly available.

The main contributions of this thesis are the development of an automated defect detection system that can improve the quality and efficiency of AFP composite material creation and the validation of the proposed system's effectiveness through experimental evaluation.

1.5 Thesis Overview

The PhD thesis contains seven chapters, including Introduction and Conclusion. These chapters are organised as shown in the flow chart in Figure 1.4 and outlined as follows:

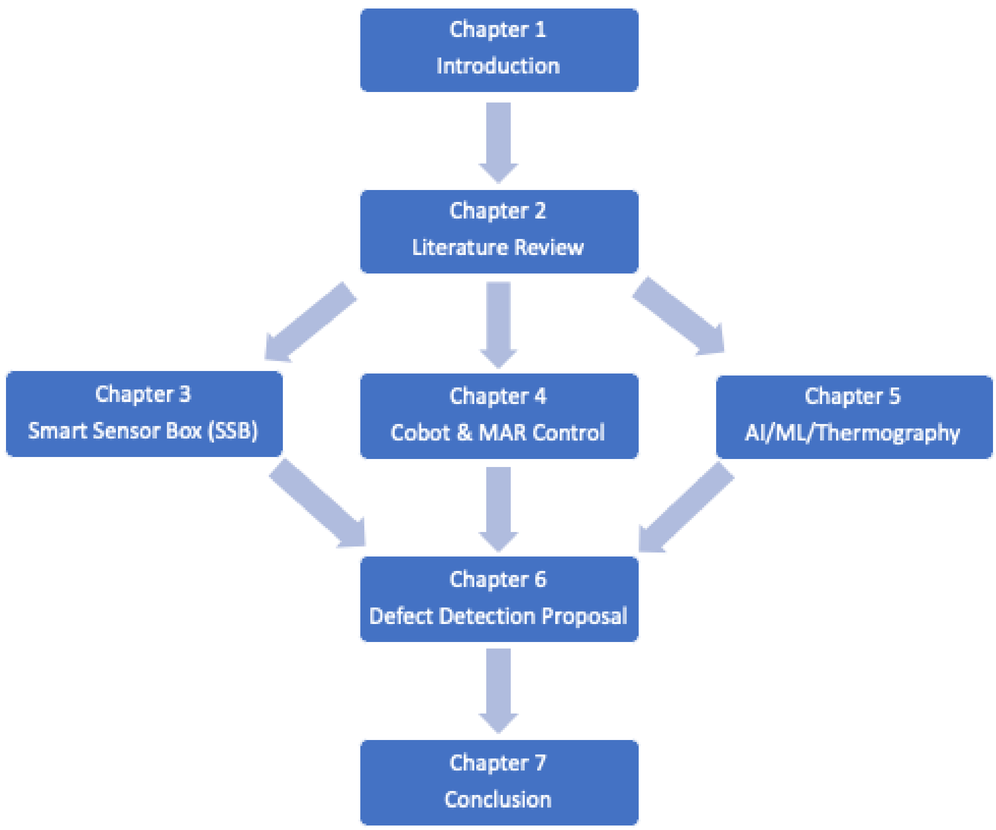

Figure 1.4 – Thesis Flow Chart

Chapter One: Introduction

This chapter focusses on establishing the concept of the research. It briefly introduces composite materials highlighting their applications and how they are manufactured. The Introduction explains the research methodology put into practice to achieve the relevant aim and objectives which are also presented in this chapter.

Chapter Two: Literature Review

This chapter provides an analysis of literature available in the field of composite material AFP defect detection. It concentrates on current techniques employed in the industry focusing specifically on the use of real-time solutions, Artificial Intelligence (AI) and machine learning models employed, collaborative interaction of multiple robots and the availability of low-cost, flexible options.

Chapter Three: SSB Development

This section introduces the design and development of a SSB which is a solution proposed to enable real-time, effective, low-cost AFP defect detection. The SSB is designed to enhance flexibility and adaptability in automated manufacturing systems, especially designed for small to medium enterprises (SMEs) aiming to remain competitive. The SSB is a modular, portable sensor hub capable of integrating multiple sensors, including cameras and advanced inputs, with support for machine learning capabilities for applications like object classification and defect detection. It presents the hardware designed, the software developed, the testing carried out and the results obtained.

Chapter Four: Cobot and MAR Control

This chapter demonstrates the integration of two robotic systems, the KUKA KR120 industrial robot and UR10e cobot as an approach to AFP and real-time defect detection. The KUKA KR120 performs high-precision fibre tow layup with a custom end effector while the UR10e cobot, equipped with a thermal camera, ensures quality by identifying defects such as overlaps or gaps through temperature variations. Both robots were simulated in RoboDK to determine path planning and defect detection capabilities before physical testing. RoboDK is a robot development environment for path planning and simulation. The real-world implementation showed the perfect coordination of these robots, where the KUKA laid composite fibres and the UR10e followed closely to provide real-time thermal imaging feedback.

Chapter Five: Machine Learning Development

This section focuses on implementing machine learning models to enhance real-time defect detection in AFP composite manufacturing using thermal imaging. Four machine learning models were implemented and evaluated: k-Nearest Neighbors (KNN), Support Vector Machine (SVM), Decision Tree, and Naïve Bayes. In addition, deep learning models were deployed. The models were trained on a dataset of segmented thermal images and Binary Classification was used to detect defects, focusing on metrics such as accuracy, precision, recall, F1-score, and ROC AUC.

Chapter Six: SSB, Cobot and MAR, ML, Thermography System Overview

This chapter integrates the SSB, Cobot and Robot, the machine learning algorithm and thermography into an overall system for optimising AFP manufacturing. The SSB acts as a central hub, coordinating the cobot's movements relative to the Robot's position, managing the collection of thermal images from the cobot-mounted thermal camera, and processing these images to identify defects in the manufacturing process through the implementation of machine learning.

Chapter Seven: Conclusion

This chapter provides an overview of the research findings and presents the conclusions drawn from each individual area of testing and analysis. It also considers possible future research that could feasibly arise from this work.

Chapters three, four and five are stand-alone chapters that explain specific areas of this PhD research. These chapters each provide the information and detail required to develop the proposal explained in chapter six.

Chapter 2

2 Literature Review

This literature review will provide an understanding of composite materials through a comprehensive overview of composites and their automation in industry. Research will then focus on investigating methods of defect detection currently employed, their benefits and any disadvantages that each solution poses. Finally, systems used to control automated production will be studied to understand how their flexibility and cost impacts manufacturing.

The purpose of the review is to identify potential areas of development and research within the field of automated defect detection in composite material production. The aim of the research carried out will be to indicate potential methods and control systems that could offer a novel, effective solution for defect detection.

2.1 Understanding Composite Materials

The concept of composites is not a new one. The use of composite materials can be traced back to the time of the Ancient Egyptians [8]. A mixture of mud and straw was used to develop bricks used for construction purposes. Another example is the Mesopotamians took thin slices of wood and glued the strips together with the grain set in alternate directions. This created the first instances of plywood, a composite still widely used in the modern age [9]. A composite still widely used today is concrete, a composite first mentioned in 25 B.C. and extensively used during the time of the Roman Empire. It is said that the concrete from ancient Rome is actually superior to the concrete used in today's day and age in terms of strength and lifespan [10].

Modern composite development did not properly occur until the early 1900s when new plastics were first created. The new plastics had the benefit of being easier and quicker to produce than the older plastics but were not sturdy enough to provide the structural strength and support required of them. This led to the requirement of a reinforcement agent to provide the structural strength and stiffness. During the 1930s, Owens Corning mass produced glass fibre used to create the product Fiberglas [11].

Fibreglas is a combination of glass fibre with a plastic polymer, produced a lightweight, sturdy and strong material, kick-starting the fibre reinforced polymer industry.

It is said war can accelerate the rate at which military technological development is carried out, and composite materials were no exception. [12] With the advent of World War 2, aircraft and ships were required to be stronger and lighter. This led to the fibre reinforced polymer industry switching from lab research to high volume mass production [13]. An additional benefit found was that fibre reinforced polymers were not visible to radio waves making it a useful material to protect and cover electronic radar equipment [2] Fibre reinforced polymers were not the only composites to be implemented during the war, another composite that was researched into for use was pykrete [14]. Pykrete is a composite consisting of wood pulp and ice, ratio 1:6. It has a slightly higher density than ice but has almost twice the crushing strength and four times the tensile strength. A project titled Habakkuk was originally conceptualised by the British military which would create an aircraft carrier out of pykrete [14]. This meant necessary metals and other materials could be used elsewhere. Ultimately, the plan was scrapped due to its very high cost and the fact that it would not be developed soon enough, although initial experimentation proved the idea was feasible, just impractical.

2.1.1 Composite Material Creation

A composite material is defined as the amalgamation of two or more differing materials, resulting in a new material. This new material can display different properties to those of the original materials used in its creation [15], Overall, the new material will have improved properties compared to the original separate materials. Although the composite material has been merged, the materials remain distinct from one another, making composite materials fundamentally different from solids, solutions and mixtures where the materials cannot be separated [15]. Composite materials are created using two key elements: the matrix and the reinforcement [15]. The matrix is implemented to bind the reinforcement together and maintain the required shape of the material. The reinforcement is implemented to provide the structural strength and stiffness required by the composite material. Additionally, the reinforcement enhances the physical properties of the matrix [16].

Composites can be classified based upon the matrix or the reinforcement used [2] [16]. The matrixes can be categorised into three sub sections: ceramic matrix composites, metal matrix composites and polymer matrix composites. The reinforcement can be separated into three main sub sections: fibre, particle and laminar as shown in Figure 2.1

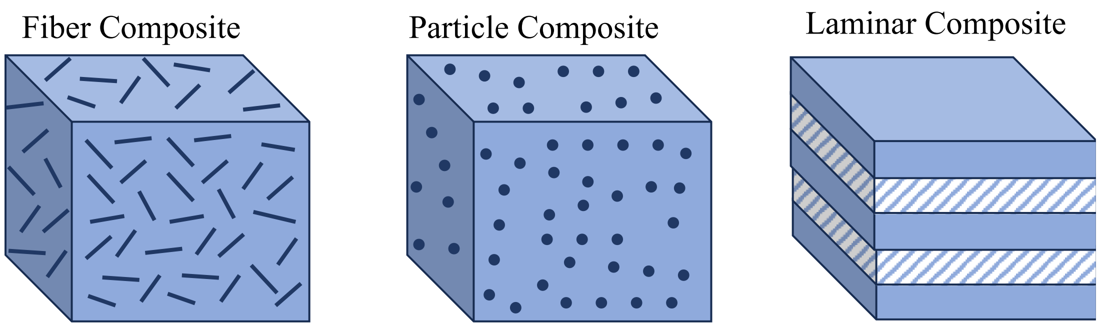

Figure 2.1 – Composite Material Reinforcement Types

Ceramic matrix composites are created by inlaying reinforcements into a ceramic matrix. Ceramic matrix composites were developed to withstand the issues of brittleness that was experienced by standard ceramics, examples include carbon and silicon carbide. Metal matrix composites are comprised of a metallic matrix with a metal, organic or ceramic reinforcement. Metal matrix composites provide a stronger material while being lightweight and experience a higher melting temperature with examples include aluminium, magnesium, titanium and copper. Polymer matrix composites contain a matrix typically made up of thermoplastics or thermosetting plastic. Continuous or short fibres are used as the reinforcement within a polymer matrix composite. Polymer matrix composites are lightweight and have a high degree of strength running through the orientation of the reinforcement. Some examples of thermosetting include epoxy, phenolic and polyurethane, while thermoplastic examples include polyethylene, polyvinyl chloride and polystyrene. Fibre reinforced composites are composed of fibres distributed throughout a binding matrix. The fibres can be woven together to create a required shape which is then put into the matrix.

Fibre examples include carbon (graphite), glass, boron, aramid (Kevlar), alumina and silicon carbide. Particle reinforced composites are composed of particle sized rather than full length strands as in fibre composites. Particles are classed as either large particle or dispersion strengthened [17]. Large particle examples include tungsten carbide in cement and sand or gravel in concrete. Dispersion strengthened examples include Thoria (thorium dioxide) in nickel for combustion engines [18]. Laminar reinforced composites are composed of thin sheets of material which are layered together and bound using the matrix material. Examples include layers of wood to make plywood and bimetals.

The most common methods for the combination of matrix and reinforcement components in a composite material are the hand lay-up and spray-up processes [2] [19]. Both are very labour-intensive procedures. The lay-up process involves the constituent parts being layered by hand and compressed while the spray-up process involves the composite being sprayed onto a surface by hand and layers being gradually built up. Both methods rely on human accuracy and can result in uneven application and areas of reinforcement with incorrect amounts of matrix.

Composite materials need to be cured to set the matrix with the reinforcements together. This curing is done in an autoclave which heats the material under pressure. Two separate methods which help to reduce the time consumption and manual labour issues while increasing quality and consistency are the use of Preform and Prepreg composites [20]. Preform is a technique of shaping the reinforcement into the shape required before the matrix is added. It can be formed into the shape by weaving or directly using a mould [21]. The matrix can then be introduced to the mould containing the shaped reinforcement under pressure. This allows the accurate control of the amount of matrix added to the composite. The advantage of Prepreg sheets is that they are composed of a predefined mix of matrix and reinforcement. The Prepreg sheets therefore do not require the user to apply the matrix which ensures the correct ratio of matrix and reinforcements across the entire material. The Prepreg is formed into the required shapes for the application and then cured. The number of layers in a composite component are dependent on the application and the industry, varying from 10s to over several hundred layers [22] [23].

2.1.2 Advantages and Disadvantages of Composite Materials

Many papers agree that there are multiple advantages in the use of composite materials [24] [25]. All papers highlight one of the key advantages as being the production of lightweight composites. The weight to strength ratio is greater than that of the standard materials that the composites are designed to replace. Another main advantage, discussed in many papers, is that composites, due to their composition, are easier to form into complex geometric shapes. Pliable, malleable composite materials are available that are versatile and can be moulded therefore requiring less labour-intensive machining methods. Papers agree that composite materials and the raw materials used to develop them are often costly. Research and development are required to create lower cost solutions [24] [26].

2.1.3 Outline of Key Composite Material Findings

Research indicates a long history of composite material use with development increasing consistently since the 1900s. [2] Development has continued due to the availability of new materials and enhanced automation techniques. The lightweight nature of composite materials makes them extremely suitable for developing key components within manufacturing.

2.2 Understanding Automation

Automated manufacturing is the implementation of robots within the production line that do not require the control of a human operator. Human workers are not required as the robot systems are designed to carry out set tasks independently by itself or part of a group of robots. With the implementation of automated manufacturing, an increase in production throughput, due to reduction in time and increase in quality, can be experienced. The consistency of robotic automation also enables higher reproducibility.

The main turning point of introducing automation into the production line occurred during the second industrial revolution. This introduced electrical systems to carry out tasks. Systems that were automated included the Spinning Jenny used to produce cloth, the water frame used to produce cotton, and the power loom used to automate weaving. The more modern method of implementing robotics within the production line was initially introduced in 1958 with the development of the first industrial robot. George Devol, dubbed the ‘Grandfather of Robotics’, and Joseph Engelberger, the ‘Father of Robotics’ started the first robotics company, Unimation [27]. The first robot, the Unimate, used hydraulics to repeat a set list of movements to carry out a task. Tasks included moving die castings within a production line and welding parts together. General Motors were the first company to implement the Unimate in their production lines in 1961 [28].

2.2.1 The Use of Robots and Cobots in Automation

The International Organization for Standardization (ISO) document ISO 8373:2021 defines a robot as a “programmed actuated mechanism with a degree of autonomy to perform locomotion, manipulation or positioning”. [29] The document further defines an industrial robot as an “automatically controlled, reprogrammable multipurpose manipulator, programmable in three or more axes, which can be either fixed in place or fixed to a mobile platform for use in automation applications in an industrial environment”. The use of robot manipulators, in the form of a robotic arm with a robotic wrist, is common in manufacturing. Different robot types have specialised capabilities making them suitable for specific functions.

Robots can be classified depending on their work envelope which describes the places in a 3D space that can be reached by the robot’s wrist. It is important to understand the classification of robots so that the optimal robot configuration can be selected for an automated manufacturing process. Two main types of joint used are prismatic and revolute. Prismatic joints consist of movement along an axis, whereas revolute consist of movement about an axis. The ISO 8373:2021 document defines six main industrial robot classifications; Cartesian, Cylindrical, Spherical, SCARA (Selective Compliance Assembly Robot Arm), Articulated and Parallel. Examples of the classifications can be seen in Figure 2.2. All of these robot types are characterised by their mechanical structure and can be created using at least three joints.

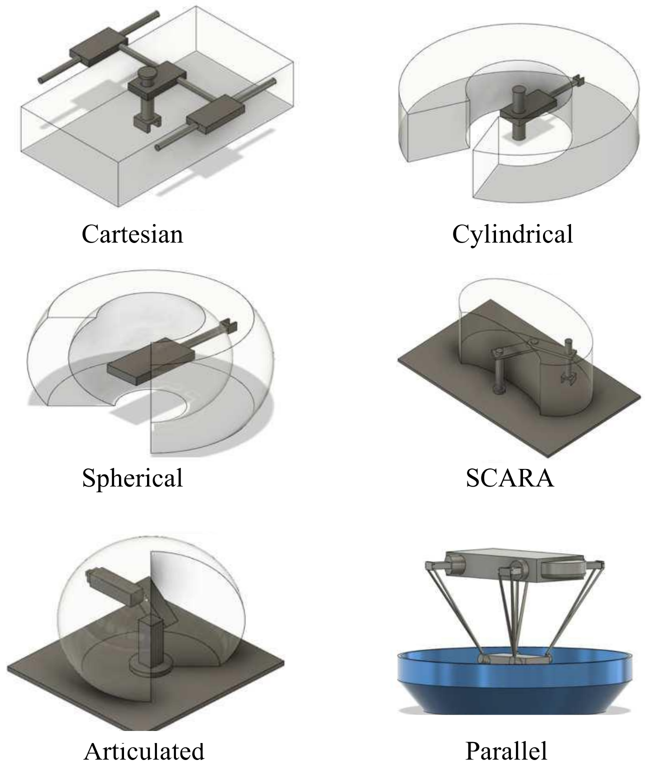

Figure 2.2 – Robot Classification

- **Cartesian**, also known as Planar, robots are constructed of three prismatic joints and work as a gantry system. The robot is able move to carry out tasks such as pick and place, arc welding and handling tools on an assembly line.
- **Cylindrical** robots implement one revolute joint and two prismatic joints to move. A robot arm can move along an axis, rotate around the same axis and extend and retract along a separate axis. Examples of use include die casting, assembly and spot welding.
- **Spherical**, also known as Polar, robots contain two revolute joints and one prismatic joint. A robot arm can rotate around an axis, rotate around a perpendicular axis and extend and retract along a separate axis. Examples of use include arc welding, spot welding and die casting.
- **SCARA** (Selective Compliance Articulated Robot Arm) robots are designed using two revolute joints and one prismatic arm. A robot arm can rotate around an axis, rotate around a second parallel axis further along the robot arm and extend and retract along a third parallel axis at the end of the robot arm. Examples of use include pick and place and handling tools on an assembly line.

- **Articulated** robots deploy three revolute joints to carry out tasks. A robot arm can rotate around an axis, rotate around a second parallel axis further along the robot arm and rotate around a third axis perpendicular to the second axis positioned further along the robot arm. Examples of use include spray painting, arc welding and part assembly.
- **Parallel** robots, also known as Delta robots, have three control joints and three additional universal joints which can be made up of revolute and/or prismatic joints. Examples of use include 3D printers, vehicle simulators and satellite trackers.

For full position and orientation utilisation, six degrees of freedom are required. Occasionally, more than six degrees of freedom are required by a robot arm. This is to work around obstacles or constraints experienced in the robot arm's work envelope.

A relatively new form of robot has slowly been making their way into the production line. These robots are called Collaborative Robots, or cobots for short. The primary use for cobots is to create human and robot collaboration in a production line. Cobots have additional safety mechanisms to ensure safety working alongside a human operator [30]. As a majority of cobots are lightweight and repositionable, this makes them a good path into the automation of a factory. Cobots are able to provide the strength and power of a robotic system with a worker who can make decisions based on situations that occur.

The concept of cobots and papers covering them have grown exponentially in the last few years as has the market for cobots within a production line. With the increase in papers being produced and production for production lines, it seems highly likely that cobots will take the forefront of human and robot interaction in the next few decades.

As cobots are still somewhat new, there is room to improve and branch out into further research. One area of research is to try and enhance the payload that cobots can handle as current systems typically cannot handle the same weights as standard robots. Another area of interest is the implementation of additional visual feedback systems allowing the cobot to incorporate extra safety features, tracking workers around it, or alternative methods to locate components it requires to interact with. The visual system could also provide a method for the cobot to predict movement of workers nearby, increasing its effectiveness. A related point of interest is the deployment of force or pressure sensors on the cobot. The sensors could be used to record and repeat the force or pressure required to be exerted to complete a task, allowing fine control and delicacy to be undertaken.

2.2.2 Advantages and Disadvantages of Automated Manufacturing

With the introduction of automation into a production line, problems and disadvantages of a new system can occur. Several problems have been highlighted including cost, limited tasks and loss of jobs [31]. As with the investment into any new system, a sizeable investment will be required to purchase, install, set up and train workers on how to use the machinery [32]. Although the robots are programmed to carry out their tasks, they are limited due to their programming. If new or unexpected problems or situations arise, the robot may not be able to react properly or in time [33]. This could cause knock-on problems throughout the production line, producing faulty products or damage to machinery. The robots also have problems easily and quickly changing to a new setup for a different task or product [32].

Throughout history, it has been seen that the introduction of some level of automation in a production line has caused a reduction in the labour required and temporary increases in unemployment. Machinery that can carry out the job of a worker has always been a worry for those working on a production line. Automation however is useful when there is a tight labour market. There are many benefits associated with automating manufacturing which is why so many companies are investing in the required technologies. The main benefits driving the next generation of automated production lines are cost, productivity, quality and jobs [34]. Initial costs are required for new robots to be added to the production line although, ultimately, the benefits experienced due to the new system will eclipse the cost factor. Due to a robot carrying out a task, the time required to complete tasks is often much quicker than when a human worker carries out the same task. As robots do not need time off, the machines can be kept running for longer periods of time, increasing the productivity and quantity of product created. A robot is more efficient at carrying out required tasks being able to reproduce the same high quality product time and time again without getting fatigued. This allows the high standard of production to not falter.

Previously stated, it was said adding automation to a production line results in the loss of jobs. While this is true that the jobs have been taken over by robots, new jobs are created replacing the jobs lost due to the automation. A similar example is the introduction of the Personal Computer (PC) [31]. The PC displaced many jobs as they were transferred to a digital form but went on to generate 15.8 million new jobs in the US alone. Workers can be retrained to work in other areas of the production line or be trained to operate and maintain the machines now working the production line.

Robotic automation has been shown to be a huge advantage in manufacturing. Many different types of robots have been developed to address different areas of automation. Many advantages of introducing robots to a production line were identified and also some disadvantages. The main advantages focussed on product throughput while the disadvantages were mostly based around human replacement.

2.3 Automation in Composite Material Manufacturing

Two significant growth areas in industry are manufacturing automation and the use of composite materials. The global automation market is forecast to grow from $76.83 billion in 2019 to $114.17 billion by 2025 [35].

Composite materials are widely used in various industries due to their high strength, low weight, and excellent resistance to corrosion and fatigue. These materials are made by combining two or more different materials to form a new material with superior properties. The most common types of composite materials are fibre-reinforced composites, which consist of fibres, such as carbon, glass, or aramid, embedded in a matrix material, such as epoxy or polymer. These composites are used in several key industries including aerospace, automotive and marine as well as in many everyday applications such as sporting equipment, hot tubs and window frames.

Automated manufacturing is the implementation of robots within the production line that do not require the control of a human operator. Human workers are not required as the robot systems are designed to carry out set tasks independently by itself or part of a group of robots. With the implementation of automated manufacturing, an increase in production throughput due to reduction in time and increase in quality and reproducibility can be experienced.

This section of the literature review will discuss the methods of automation used in specific industries when working with composite materials and will focus on current approaches and issues within this field.

2.3.1 Automated Fibre Placement

AFP is a manufacturing technology used for producing fibre reinforced composites. AFP involves the use of a robotic arm that lays down narrow tows of continuous fibres onto a mould to create a composite part. The process is automated and allows for precise fibre placement, which leads to improved strength and stiffness of the final product. AFP has several advantages over traditional manual methods. These include higher production rates, improved quality, increased consistency, and reduced labour costs. An APF end effector can be seen in Figure 2.3 and a diagram of how an AFP end effector works can be seen in Figure 2.4.

Figure 2.3 – AFP end effector

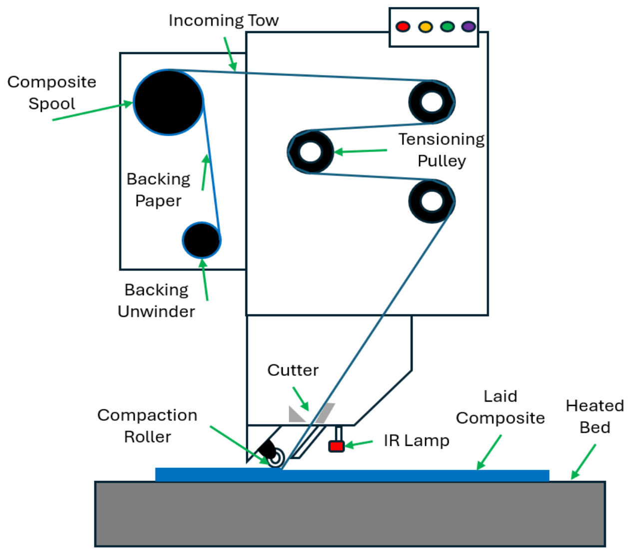

Figure 2.4 – AFP Robot End-Effector

The image shows the process of carrying out AFP, a method used to fabricate composite materials. In this process, continuous fibre tows are fed into the system and guided by a series of rollers which help control the path and tension of the tow material. Figure 2.5 shows a spool of composite tow complete with blue backing tape. The backing tape is wound onto a lower spool separating it from the composite material. The material is then pulled into the machine.

Figure 2.5 – AFP Tow Spools

A cutter is used to trim the fibre to the required length, allowing for accurate placement and dimensions. As the tow reaches the layup area, a heat source softens the resin within the fibre, increasing its ability to hold onto the tool bed surface below. The compaction roller then presses down the heated tow, ensuring proper adhesion and layering. The material rests on a bed or mould, which acts as the foundation for shaping the composite part. The entire process moves in a pre-determined layup direction, allowing for repeated automated production of complex composite structures. Figure 2.6 shows the laid tow in an octagonal design. This is the same design used in this research.

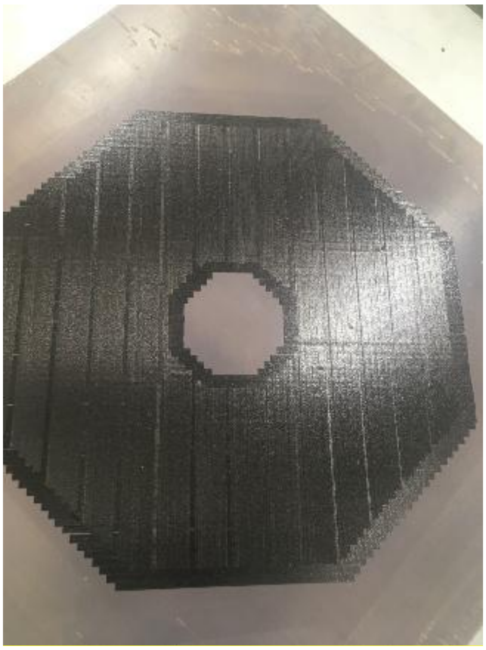

Figure 2.6 – AFP Laid Test Shape

Despite its many advantages, the AFP process can result in defects, such as voids, wrinkles, and misalignments, which can compromise the integrity of the final product.

2.3.2 Automation in Industry

Market analysts predict a growth in composite material use across many industries [36]. Industries of particular interest, due to their advances in manufacturing automation are automotive, aerospace and marine, as seen in Figure 2.7.

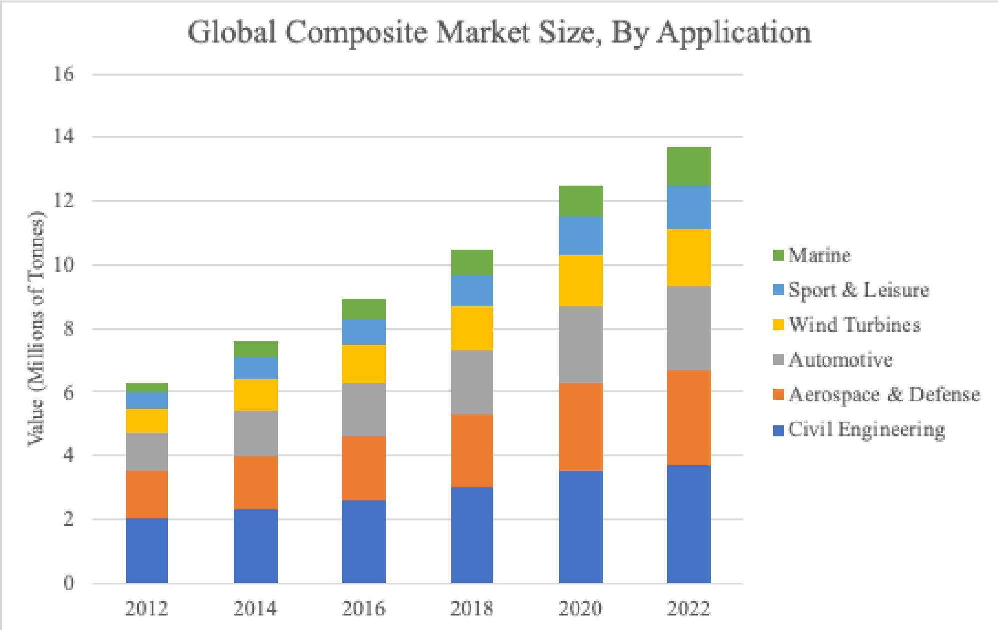

Figure 2.7 – Global Composite Market by Application

Several steps have been made in automating manufacturing with composite materials in industry. The main areas of automation, covered in many papers, are AFP and Automated Tape Laying. AFP and Automated Tape Laying are both automated methods of creating a composite material [37]. In Automated Tape Laying single large strips of material are handled on a gantry style machine and is useful for large flat shapes. In AFP, multiple smaller strips of composite material are applied in parallel and is useful for complex, detailed structures. Prepreg strips are heated and compressed in both processes which means that debulking, the compacting of composite layers to remove any air bubble impurities, is no longer required [38].

2.3.2.1 Composite Materials in the Automotive Industry

Composite materials are used in the automotive industry to produce lightweight materials that are durable and strong enough to provide protection in an impact. The weight reduction results in lower fuel consumption which is environmentally advantageous [39]. Common composite materials used in the automotive industry are carbon and glass fibre reinforced polymer composites. These materials are used to develop many components of the car including body and chassis parts, steering wheels, dashboards and racing seats. Metal matrix composites are also under investigation for use within the automotive industry as metal matrix composites exhibit increased tensile strength, thermal conductivity and working temperature [40].

Cooper Cars were the first in the 1960s to develop a chassis using a composite material [41]. The composite, made from layers of aluminium panels, aluminium honeycomb core and a glass reinforced plastic inner skin, allowed the development of a lightweight monocoque that was stiffer than previous chassis. Further developments in the field allowed this concept to be adopted by Formula 1 cars to minimise weight and enhance aerodynamics. Carbon fibre composite material was first introduced by McLaren Formula 1 in 1980 [41]. The following year, John Watson, McLaren driver, lost control of his Formula 1 car, hit a barrier but survived despite the force causing the car's engine and gearbox to be torn from the car. The monocoque protected the driver due to the ability of the carbon fibre material to absorb the energy from the crash. With John Watson walking away from the crash, the worries of carbon fibre Formula 1 cars were put to rest. John Watson is quoted saying "Had I had that accident in a conventional aluminium tub, I suspect I might have been injured because the strength of an aluminium tub is very much less than the carbon tub." [42]. High-end performance development has filtered down into standard automotive manufacturing with BMW being the first mainstream manufacturer to develop a composite material car for the mass production. The electric BMW i3 car uses carbon fibre reinforced polymer for the body and internal panels of the car to keep weight down and allow maximum efficiency [43].Figure 2.8 shows the projected rate of growth of carbon fibre composite consumption in the automotive industry growing from 3.8 thousand tons to almost 10 thousand tonnes between 2013 and 2030 [44].

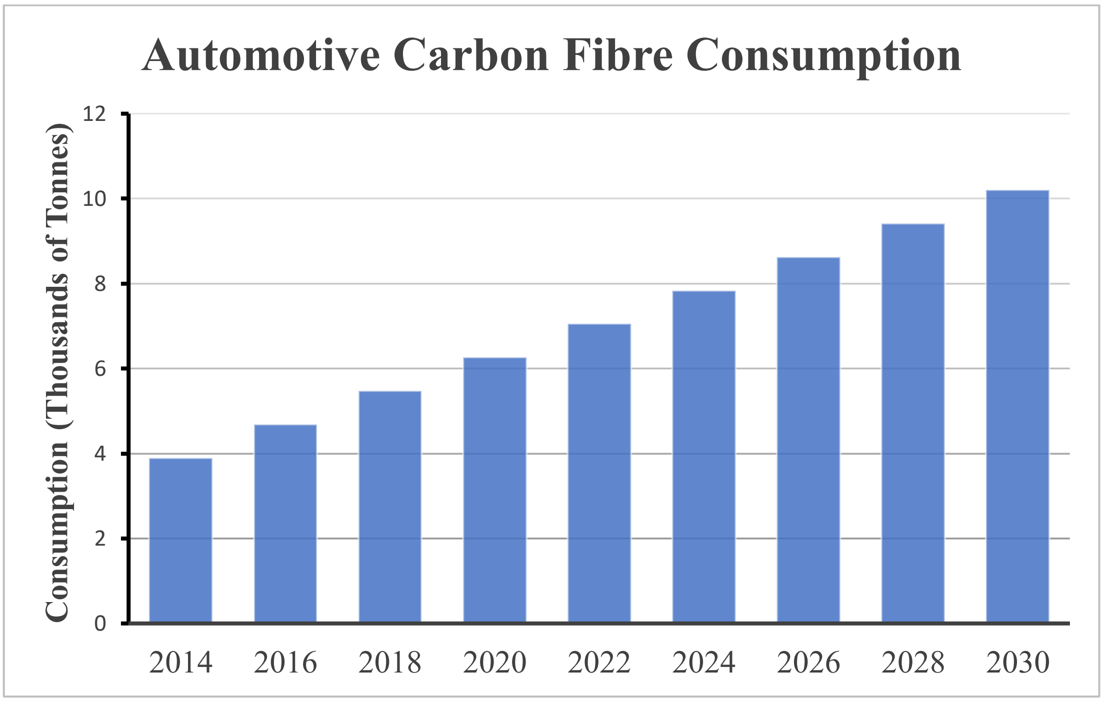

Figure 2.8 – Global Automotive Carbon Fibre Consumption

Figure 2.9 highlights the increase in vehicle applications using composite materials to make use of their lightweight characteristic [45].

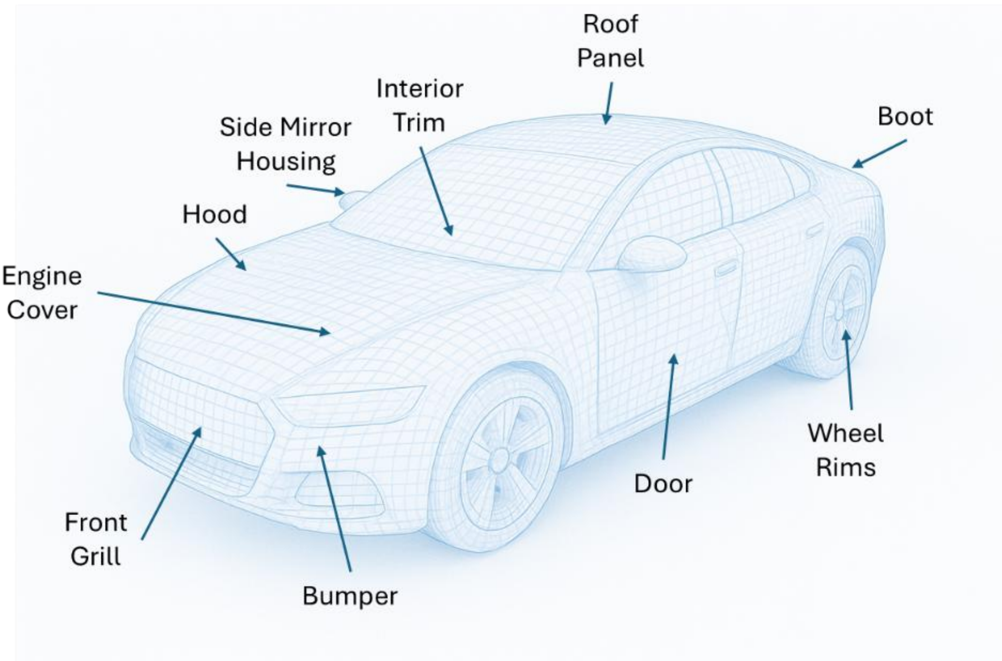

Figure 2.9 – Automotive Composite Applications

2.3.2.2 Automation in the Automotive Industry

The car manufacturer Audi is committed to investment in automated manufacturing. Their factory in Mexico is designed to use centralised production control, smart sensors, drones and 3D printing to increase productivity and efficiency [46]. Drones transport steering wheels from one area of the factory to another which relieves humans of a task that is time consuming and unskilled. Key to Audi's concept is that robots will do the bulk of the work and human operators will 'supervise and handle creative tasks'. The company Voith Composites has worked with Audi to develop a high-speed, automated process for developing carbon fibre automotive components [47]. This process is called the Voith Roving Applicator.

The Voith Roving Applicator, is an advanced AFP method which lays fibres in a directed fashion and can produce complex composite shapes with high throughput and low waste. The Voith Roving Applicator is used by Audi in the development of their A8 carbon fibre reinforced polymer rear panel. The strength produced in the carbon fibre reinforced polymer through specific placement of fibres gives the panel extra strength and rigidity.

2.3.2.3 Composite Materials in the Aerospace Industry

Composites have slowly been working their way into the aerospace industry over the past several decades. Key elements of their integration are down to the high strength and lightweight obtained from composites [48]. Modern aeroplanes are commonly constructed using aluminium, a lightweight and strong metal. Planes through the first passenger plane to the Boeing 747 were produced using aluminium. Although aluminium is lightweight, other options to make an aircraft were researched to make them lighter, allowing less fuel consumption and cost saving to occur. Composite materials can provide a similarly strong structure with almost half the weight. The most common composite materials used in the aerospace industry are carbon fibre reinforced polymers and glass fibre reinforced polymers, often Fibreglas [49]. The percentage of composite used within an aircraft has increased year on year, with Boeing going from 12% to 50% of composite used by weight in the Boeing 777 [50].

Approximately 50 years ago, the Grumman Aerospace Corporation were tasked with developing the F-14 Tomcat and Boeing the F-15 Eagle [51]. Both aircraft were high performance fighters where a high strength accompanied by low weight was paramount to their design. Both companies chose to implement boron-reinforced epoxy composites within the aircraft, with differing uses. Grumman used the boron on the horizontal tail skin whereas Boeing used the boron for the horizontal and vertical skins and for the rudder skin. Originally, composites were only implemented within additional parts of a plane and not core structures like the wings or fuselage. With the use of boron skins and further understanding of composites, the core structures slowly transitioned to composite materials, such as the fuselage, wings, doors, blades and floor panels. By 2005, the majority of aeroplanes contained over 25% of composited based on weight [50].

2.3.2.4 Automation in the Aerospace Industry

Rolls Royce is famously known for their work within the aerospace industry. Rolls Royce has developed and produced aircraft engines [52]. For new versions of engine, the quality of product and efficiency has increased. Rolls Royce run a state-of-the-art Small Engine Assembly Line (SEAL) designed to produce turboshaft engines. The production line incorporates 16 separate stations which are all produced to optimise the quality and efficiency of the station [53].

Rolls Royce recently opened a new composites technology hub in the UK. The facility is designed to develop and manufacture fan cases and blades produced from composite materials [54]. The facility will focus on using carbon-fibre composites, thus helping reduce the overall weight of the aircraft and fuel required to be used. The new engines created with the composite materials are stated to produce 25% less CO_2 compared to the first-generation engines and reduce the fuel consumed by 25% [54]. The engines will lighten an aircraft by 700kg for a twin-engine system. Rolls Royce continue to develop components with the aim of being quieter, lighter and provide more power with the assistance of composite materials.

2.3.2.5 Composite Materials in the Marine Industry

Composites are widely used within the marine industry in all aspects due to composites resistance to corrosion and degradation from the elements [55]. Boats were originally designed and crafted from wood or steel, but due to the increasing price and constant maintenance required for the boats, another material was required to be investigated for use in ship building [56]. With the use of composites thinner, lighter and easier to maintain marine craft could be developed [57]. The most common composite material used in the marine industry is fibreglass reinforced polyester, with 90% of all recreational boats being built with fibreglass reinforced polyester [56] [58]. Parts built using composites include the hull, the deck and high-performance structural sections.

In the 1930s when the Owens Corning company developed and produced Fiberglas, it led to glass fibre composite material production which drove the development of polyester resin by the American Cyamid company in 1942 [55] [56]. Around this time, Basons Industries attempted to make a fibreglass boat. Due to no releasing agent being implemented, the boat was inseparable from the mould, resulting in a failure [59]. Further development in the field led to the first proper ‘plastic’ boat to be developed in 1946 as a surfboat with the boat on display at a show in 1947 [55] [59]. The navy then went on to incorporate composites into their ships, initially with the development of mine countermeasure vessels, taking advantage of the lightweight and non-metallic nature of these materials. [58] [59].

2.3.2.6 Automation in the Marine Industry

Airborne is a technology leader in advanced composite part production [60]. Their aim is to reduce the cost of production and mass production composite parts by automating and digitalising. Current steps within the production line will be automated creating a building block approach. This allows the automated systems to be configured and placed in the relevant sections within the production line. Digitalisation allows the system to learn and optimise a production line.

Airborne is currently working with DAMEN, a shipyard group who design and build seafaring vessels. The two companies, including others, are designing and building a composite based vessel, project name RAMSSES [61]. The aim of the project is to "demonstrate the benefits of innovative materials for efficient ship designs" [62]. With the use of composites, vessels could weigh up to 40% less than steel equivalents, resulting in massive reductions in fuel consumption and emissions.

To this effect, the project partners unveiled a 6-meter-high hull section completely produced using composite materials that they have been working on for roughly three years [63]. With the hull having now been developed, the project can continue forwards and aims to prove the viability of large vessels made from composites as a sustainable and suitable solution.

2.3.3 Challenges in Automated Manufacturing with Composite Materials

Prepreg protective film removal has been discussed in the literature [64] [65]. Prepreg rolls come with a thin film backing sheet to ensure the sticky resin of the matrix is not contaminated with external materials. To use the Prepreg material, the thin backing film must be removed. This is a time-consuming process and precise work that needs to be carried out carefully to ensure the Prepreg is not damaged. Previous attempts to create a robot to automate the removal of the backing film have been undertaken [64]. The main concept was to use suction to peel the backing film off of a slightly bent corner of Prepreg material. The removal of the film using this method proved to be difficult and not possible without an initial corner being manually released from the Prepreg material. Once a section had been released it was easy to achieve an automated removal of the film. No paper found identifies a completely successful automated method of Prepreg protective film removal.

Another area of automation in the use of composite materials, described in a small number of papers, is automated draping [66] [67]. Draping is the overlaying of a Prepreg composite material onto a pre-constructed mould. This is a difficult, time-consuming process that requires great precision to avoid defects within the finished product. A research group has been investigating a possible solution for this [67].

The need for fast production turn-around and variety within product design and type leads to frequent tool set-up changes within an automated production line [68]. In the lightweight manufacturing industry this is a particular challenge where composite material advances are being sought and achieved rapidly within many industries. There is need for a solution that would allow repeated production set-up changes without causing a high cost or time penalty while maintaining expected quality standards. This is an area that would merit further investigation.

2.3.4 Outline of Key Composite Material Automation Findings

In this section, a review of composite materials and automated manufacturing has been discussed. An analysis of automation in composite product manufacturing has also been carried out. From the exponential increase of papers covering composite materials, it is clear that this is an area in the forefront of interest and one that will continue to be developed. Evidence also highlights that automated manufacturing has advanced significantly over time, and this has been incorporated into the handling of composites as well. Several areas of automation have benefitted industries in their use of composites however some areas require more attention. An area identified that would merit further research would be the requirement to increase flexibility within an automated production environment to allow rapid set-up changes without reducing quality.

2.4 Defect Detection in Automated Composite Material Manufacturing

Composite materials have gained significant popularity in various engineering fields due to their exceptional mechanical properties and lightweight nature. AFP is a widely used manufacturing process for composite materials, which involves laying down continuous fibres onto a mould using a robotic arm. However, the AFP process can result in defects, such as overlaps, twists, and misalignments, which can compromise the integrity of the final product. Therefore, it is crucial to detect and eliminate defects during the manufacturing process.

Defects in composite materials can arise from various factors, such as manufacturing errors, material inconsistencies and environmental conditions. These defects can greatly reduce the strength of the final product and can lead to catastrophic failure in some cases. Detecting and eliminating defects during the manufacturing process is key to ensuring the quality and reliability of the composite part.

Furthermore, composite materials are widely used in safety-critical applications, such as aerospace and defence, where any defect can have severe consequences. For example, a defective aircraft component made of composite materials can lead to a fatal accident. Therefore, it is essential to have an effective defect detection system that can identify defects in real-time during the manufacturing process.

Currently, defect detection in AFP composite material creation is mainly performed using visual inspection, which can be time-consuming, costly, and subjective. In recent years, researchers have explored various techniques to detect defects in AFP composite materials. Ultrasonic testing and Eddy-current testing are methods used for inspection of composite parts post-curing cycle and has been shown to be effective at detecting defects embedded within the laminates. It does however lack the ability to detect defects during the manufacturing process. Research has been ongoing in utilising profilometry for detecting surface level defects during the manufacturing process through analysis of the surface topology. The visual and profilometry methods have limitations in detecting certain types of defects, such as voids and delamination, which can be hidden from the surface. Another promising technique is thermal imaging, which involves analysing the difference in heat distribution across the surface of the composite material as a result of the heat transfer through the laid material. Differences of recorded heat are indicative of differences in layup structure and hence can be used to detect the presence of defects. Table 2.1 shows AFP defect detection methods that are currently being researched. From the systems identified, only the profilometry research achieved true real-time defect detection. Very little information was identified for the use of eddy currents in AFP defect detection. The research found is still at an experimental stage, so its real-time status is unknown. Research involving both thermography and ultrasonic methods had solutions that achieved near real-time detection where data is collected in real-time, but defects are examined post-layup.

Table 2.1 – Defect Detection Methodologies

Method	Research Group	Technique	Defects Detected	Real Time
Thermography	NASA Langley – ISTIS [69]	IR thermal imaging mounted on AFP head	Gaps, overlaps, folds, twists, foreign objects	Near real-time
Thermography	NASA Langley (ISAAC AFP system) [70]	IR camera mounted on AFP robotic cell	Gaps, overlaps, wrinkles	Near real-time
Thermography	Harbin Institute of Technology [71]	IR vision + AFP process monitoring research system	Position error, overlaps	Near real-time
Ultrasonic	NASA LAR-TOPS-327 system [72]	Automated ultrasonic scanning during cure	Evolving defects	Near real-time (not AFP layup)
Ultrasonic	University of Bristol (UNDT group) [73]	Ultrasonic imaging of CFRP laminates	Wrinkles, gaps, overlaps	No (offline)
Ultrasonic	DLR (Germany Aerospace Center) [74]	Automated phased array	Pucker, wrinkles	No
Profilometry	Electroimpact RIPITx [75]	Laser profilometry + machine vision	Gaps, overlaps, twists, foreign objects	Yes
Profilometry	University of Bristol closed-loop AFP control [76]	Profilometry linked with adaptive AFP process control	Wandering tow, position error	Yes
Eddy Current	University of Bristol directional EC wrinkle probe [77]	Directional asymmetric eddy-current sensing	Wrinkles	Unknown

Investigation of each method presented thermography as a suitable option for this research with the potential to develop a real-time defect detection solution. An example thermal image showing gap defects can be seen in Figure 2.10.

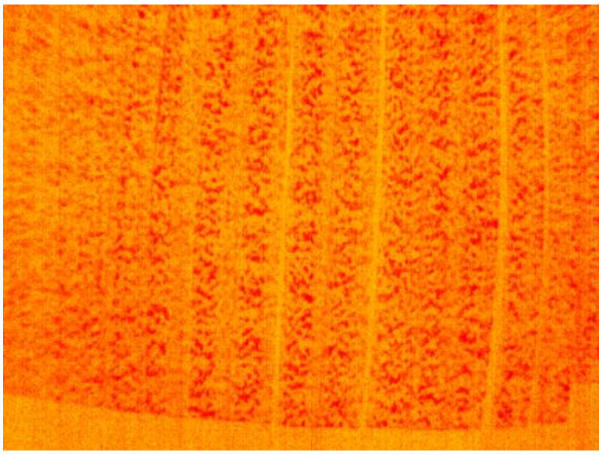

Figure 2.10 – Thermal Image with Defects

Defects can occur for many reasons in the AFP process. In a paper by Pantoji et al. [78], several causes of gap/overlap defect are described. These include inaccuracy in the robot arm positioning, movement of the tow on the roller, variance in the width of the tow, and the amount on compaction under the roller.

Several studies have investigated defect detection in AFP composite materials. A paper by R. Harik et al [79] is important for initial coverage as they list different classifications identified for AFP defects. In total, fourteen different defect types are described throughout the paper. A description and example illustration of each defect type is shown in Table 2.2.

Table 2.2 – Defect Types

Defect Type	What it Means	Why it Matters	Sample Defects
Gap / Overlap	Adjacent tows not placed edge to edge. Gap leaves exposed space and overlap stacks one tow over another.	Creates local resin-rich or fibre-rich regions, thickness variation, voids and possible strength reduction.	
Pucker	Local raised area or bunching of tow material.	Causes local waviness, poor compaction and possible loss of strength.	
Wrinkle	Tow creates visible waviness or buckling instead of lying flat.	Produces fibre waviness, poor load transfer and possible delamination.	
Bridging	Tow covers a concave area without fully contacting the tool or previous ply.	Leaves voids or resin-rich regions under the tow increasing the chance of delamination.	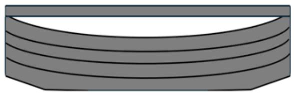
Boundary Coverage	Tow does not properly cover the required region.	Affects final part shape and can create weak edge sections if not trimmed or repaired.	
Angle Deviation	Tow placed at an angle different from the programmed fibre orientation	Changes local laminate stiffness. May also create overlaps or resin-rich areas.	

Fold	Tow is folded over itself, usually across its width.	Creates local thickness increase, fibre misorientation and resin-rich areas.	
Twist	Tow rotates before being compacted, causing tow to be laid with a twisted cross-section.	Causes local thickness change and fibre misalignment.	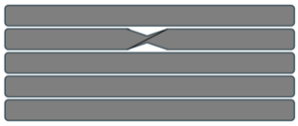
Wandering Tow	Tow drifts away from its intended path, especially after cutting or before compaction.	Can create unexpected gaps, overlaps, or boundary errors.	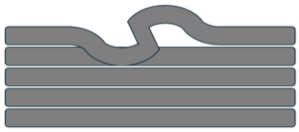
Loose Tow	Tow is placed but not firmly bonded or compacted to the surface.	Can lift, shift, wrinkle, bridge or lead to a missing tow later.	
Missing Tow	A full tow absent from where it should have been placed.	Similar to a large gap. Creates local thickness loss and resin-rich area.	
Splice	Two ends are joined in the material, creating a thicker section.	Causes local thickness increase may become a failure site.	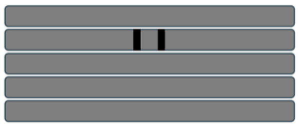
Position Error	Tow start/end point is placed in the wrong position.	Produces misalignment, boundary defects, gaps, or overlaps.	
Foreign Object	Unwanted debris is trapped on the layup surface.	Prevents proper adhesion, can create voids, bumps, delamination, or local weakness.	

Each defect presents differently but all can lead to failure in the composite material production and waste of partially or completely laid material. The vast number of defect types make detection difficult as different methods are better at recognising certain errors than others. In a study carried out by Bell Helicopter Textron Canada Limited in 2015, it was found that gaps/overlap defects accounted for 57% of all defects [80]. Based on these findings, only gap/overlap defects were investigated. It is difficult to quantify the defect rate of hand layup as it is highly dependent on the experience of the person carrying out the layup.

While there is no standard definition of what constitutes a defect, anecdotally, within the industry, a defect is classed as a 2mm gap/overlap. This is not an absolute rule and becomes more relevant when issues are compounded. The definition of a defect is also dependent on the application and type of part being created.

2.4.1 Robots in Defect Detection Automation

In the work authored by Wu et al. [70], the authors present their innovative solution for AFP through the development of a system named Integrated Structural Assembly of Advanced Composites (ISAAC). ISAAC is a sophisticated and precise robotic platform designed to facilitate comprehensive research in the domains of advanced composite materials and structures, encompassing aspects such as design, analysis, manufacturing, and evaluation.

This contribution underscores the importance of incorporating advanced robotic platforms, such as ISAAC, within the realm of AFP research. Such systems enable researchers to explore new avenues in composite material engineering and establish novel methodologies for enhanced design, analysis, manufacturing, and evaluation processes. The development of ISAAC marks a notable advancement in the field, offering substantial opportunities for further advancements and breakthroughs in the domain of advanced composites. However, the inspection processes for the system still relies largely on manual defect detection, highlighting the need for an automated approaches to enable efficient and scalable defect detection in AFP manufacturing.

2.4.1.1 Thermal Imaging in Defect Detection Automation

Thermal imaging measurements are a non-destructive testing technique that has been used to detect defects in composite materials by analysing the heat distribution on the surface. The technique involves capturing the surface temperature of the composite material during the manufacturing process and using it to detect the presence of defects. In recent years, machine learning has been integrated with thermal imaging to automate the defect detection process and improve its accuracy.

Using a thermal imaging technique for AFP defect detection has been shown to be a practical solution for composite materials. The general equation for heat transfer in a material is based on Fourier's Law [81] which is shown in Eq. 2.1

$$\frac{Q}{t} = \frac{kA\Delta T}{L} \tag{2.1}$$

where Q/t is rate of heat transfer, k is thermal conductivity, A is cross-sectional area, ΔT is temperature difference between two points and L is material thickness. The equation shows that heat transfer is indirectly proportional to thickness. As the thickness of the material being examined increases, the rate of heat transfer decreases. Fourier's Law relates primarily to homogeneous materials with a single material type which means that composite materials, which are non-homogeneous, require more complex equations that take into account the different heat transfer properties of each material [82] [83]. Despite this extra complexity, the relationship still exists that rate of heat transfer depends on the thickness of the material being considered making thermography a suitable method for detecting defects in AFP tows such as gaps and overlaps.

Thermal imaging has several advantages over traditional defect detection methods. It is non-contact, non-destructive, and can be performed in real-time during the manufacturing process. This allows for defects to be detected and corrected immediately, reducing the risk of defects going unnoticed and resulting in a defective final product. Additionally, thermal imaging can detect defects that may be hidden from the surface and therefore not visible to visual inspection or profilometry.

Thermographic inspection has been shown to provide comprehensive defect detection capabilities, with [84] writing about its effectiveness in identifying and categorising all different types of defects. The paper highlights machine learning as a developing research area for defect detection in this field, which is addressed in this paper. In a related study, [85] carried out deployment of an edge-detection algorithm along with thermography. This allowed both the localisation of individual tows and the identification of thermal anomalies. Another paper, Schmidt et al. [86] integrated a Convolutional Neural Network (CNN) with thermographic imaging, demonstrating that machine learning can significantly enhance the accuracy and reliability of defect detection based on thermography. However, their study was limited in scope as it only evaluated two deep learning models and considered a small set of defect types.

In a study conducted by Peter D. Juarez et al. [87], a thermographic inspection technique was proposed as a solution for defect detection in an AFP systems. The authors suggested the use of a quartz heating lamp placed before the tow placement process. By attaching an infrared camera to the AFP head and capturing images of the region immediately behind the compaction roller, the heat conduction from the substrate to the outer surface or from the outer surface into the composite material could be evaluated to assess its state.

The study demonstrated that in situ thermal imaging successfully detected relevant defects in AFP without requiring visual inspection of each ply layup. Furthermore, the authors highlighted the benefits of performing a post-layup line scan to gain insights into the final material state prior to curing. The application of both in situ and ex situ thermal imaging techniques holds the potential to create more efficient production environments and contribute to the advancement of process development for new material systems and structural designs. These findings underscore the value of thermal imaging in defect detection and its wider implications for improving manufacturing processes in the field of composite materials. However, defect identification in this work relies largely on manual interpretation of the thermal data, highlighting the need for automated analysis methods to enable scalable and reliable defect detection in industrial AFP applications.

2.4.2 Machine Learning in Defect Detection Automation

Machine learning algorithms can be integrated with thermal imaging to automate the defect detection process and improve its accuracy. These algorithms can learn from large amounts of data and identify complex patterns that may be difficult for human operators to detect. By training a machine learning algorithm on a dataset of known defects, along with known defect-free samples, the algorithm can learn to detect similar defects in real-time during the manufacturing process. This allows for a more objective and consistent approach to defect detection, reducing the reliance on subjective human judgment.

In the work by Sacco et al. [88], an inspection system leveraging machine learning algorithms is presented, aiming to detect and characterise defects in profilometry scans. The system incorporates a data storage system and a User Interface (UI) to enable informed manufacturing processes. The machine learning algorithm employed is based on fully convolutional neural networks, enabling pixel-level classification. This enables the software to accurately identify the size, shape, and location of defects on the inspected part.

The inclusion of a user interface empowers the system operator with precise control over the inspection process and provides an opportunity to rectify any inaccuracies in the predictions made by the machine learning algorithms. The continuous training of the fully convolutional neural network using these corrections facilitates ongoing improvements in the accuracy of the inspection process.

A primary objective of the ongoing development of this AFP inspection platform is the identification of multiple defects beyond the scope of conventional automated inspection systems, which primarily focus on gap and overlap defects. Although the algorithms employed in the system can often detect the presence of defects, they may occasionally misclassify the specific defect type, particularly in cases involving rare defect types.

The paper "Review of machine learning applications for defect detection in composite materials" by Daghigh et al. [89] gives a comprehensive analysis of the deployment of machine learning methodologies in the context of defect detection in composite materials. It highlights the potential of machine learning techniques, including supervised and unsupervised learning approaches, for predicting composite behaviour under several conditions and environments. The review shows the transition from traditional physics-based modelling, like finite element methods, to data-driven and hybrid models, such as physics-informed neural networks. Despite the advantages of machine learning, challenges remain, particularly in incorporating physical laws into machine learning frameworks for interpretable and reliable results. The paper explores a range of machine learning methods such as regression trees, support vector machines, and deep learning for applications ranging from defect identification to material optimisation. It also underscores the need for explainable AI and benchmark datasets to enhance the utility and generalisation of these models. Overall, the review serves as a view to the future for advancing defect detection and material design through machine learning, showing the way for improved efficiency and reliability in composite material applications.

The reviewed paper [90] explores the integration of machine learning with structural health monitoring and non-destructive evaluation methods to enhance damage state evaluation in composite materials, with a key focus on fibre-reinforced polymers. It highlights machine learning's potential to address complex challenges such as detecting, classifying, locating, and predicting damage mechanisms like overlaps, matrix cracks, and fibre breaks, which conventional models struggle to achieve without explicit material microstructure details. The paper delves into supervised learning models like ANNs, SVMs, and ensemble methods, and unsupervised clustering techniques like k-means, detailing their application in structural health monitoring or non-destructive evaluation scenarios for tasks such as damage detection and prognosis. Key contributions include the innovative use of machine learning to pair structural health monitoring data with finite element models for real-time analysis and the exploration of hybrid machine learning approaches to enhance prediction accuracy and reduce computational loads. By demonstrating significant economic and safety advantages in applications like aerospace, the study shows how machine learning reduces inspection costs and downtime, improves maintenance planning, and extends the lifespan of composite materials. Additionally, it identifies limitations like small dataset availability, noise sensitivity, and the need for scalable models, offering future directions for advancing real-time and multi-task machine learning systems for dynamic industrial environments. This comprehensive review underscores machine learning's transformative role in advancing reliability, safety, and cost-effectiveness in composite material damage evaluation.

The paper [91] presents an in-depth study on the integration of machine learning and advanced composite manufacturing, particularly in the AFP process. The authors propose a novel inspection method deploying machine learning and computer vision to detect, classify, and characterise defects in AFP-manufactured composite parts. Using tools like convolutional neural networks fully convolutional networks, and ResNet architectures, the study demonstrated a cutting-edge defect detection system integrated with the Automated Composite Structure Inspection System called ACSIS, equipped with laser profilometers for high-precision defect profiling. The system leverages semantic segmentation to identify and quantify defects, offering improved defect classification and real-time feedback capabilities. A user interface allows operators to interactively reclassify and train the machine learning model, enhancing its adaptability and accuracy. Validation was conducted on various AFP-produced components, including laminates and cylindrical parts, showcasing the system's ability to identify defects like wrinkles and gaps, reduce inspection time, and inform real-time corrective actions. Furthermore, the system's integration with third-party finite element modelling software underscores its utility in assessing structural impacts of defects. By addressing challenges like manual inspection's time-intensity and lack of precision, this paper highlights machine learning's transformative role in streamlining quality assurance and advancing composite manufacturing, with future prospects in online inspection and process optimisation to increase throughput and efficiency.

2.4.3 Machine Learning and Thermal Imaging in Defect Detection

The combination of thermal imaging and machine learning has shown promising results in defect detection for AFP composite material creation. By detecting defects in real-time during the manufacturing process, it is possible to reduce the risk of defective final products and ensure the quality and reliability of composite parts. Additionally, the use of thermal imaging and machine learning can improve production efficiency by reducing the time and cost associated with traditional defect detection methods. Therefore, the development of a thermal imaging and machine learning-based defect detection system is of significant importance for AFP composite material creation. Building on this, a key question to address is which machine learning models can be optimised to provide the highest accuracy and reliability for defect detection in AFP processes.

The research conducted by Caizhi Li et al. [92] explores the application of an infrared image and signal fusion algorithm (1D-YOLOv4) for detecting damaged areas using images and identifying damage types using signals.

The study [92] begins by preparing damaged specimens of composite materials and employing infrared non-destructive testing methods to detect damage within the materials. Subsequently, the 1D-YOLOv4 network is developed to enable intelligent fusion detection of infrared images and signals. A comparative analysis is then conducted, comparing the proposed model with various other models. The results demonstrate the superior detection capabilities of the developed model.

Through iterative enhancements, the 1D-YOLOv4 network achieves an accuracy rate of 98.3%, an average precision (AP) of 91.9%, and a kappa coefficient of 0.997. The algorithm effectively detects both infrared images and signals simultaneously. It accurately identifies the damaged areas of composite materials and provides insights into the specific type of damage, thereby introducing a novel approach to composite material detection.

Another notable study by Arrabiyeh et al. [93] explores the integration of artificial intelligence and low-cost machine vision systems to enhance defect detection and quality control in composite manufacturing processes. The study focuses on implementing a budget-friendly setup using a webcam paired with an NVIDIA Jetson Nano and AI algorithms to address fiber misalignment during the wet fiber placement (WFP) process. By employing YOLO algorithms for real-time monitoring, the system achieves 91.3% detection accuracy after 50 training epochs, demonstrating the feasibility of low-budget AI-driven quality control in advancing Industry 4.0 applications. This research highlights the importance of fiber alignment for ensuring mechanical integrity, while showcasing the potential of cost-effective, AI-enhanced systems for improving manufacturing efficiency. The findings align with this study's focus on leveraging advanced technologies like machine learning and vision systems to improve the AFP defect detection process. However, none of the studies looked at using a vision system on a secondary robot arm to provide better flexibility and optimisation of the imaging system for capturing images to be processed by the machine learning system.

Other studies have also investigated the use of thermal imaging and machine - learning for defect detection in composite materials, including the use of deep learning algorithms such as recurrent neural networks and long short-term memory networks. The results of these studies have shown that thermal imaging combined with machine learning can be an effective method for defect detection in composite materials.

2.4.4 Outline of Key Defect Detection Findings

Several types of faults were identified which was essential as a means of listing, sorting and classifying results in AFP defect detection systems. Thermal imaging has been used successfully for defect detection systems in AFP. However, as a standalone method it is not sufficient as it is a manual system requiring humans to analyse and review the images that are output. Machine Learning for defect detection has had a large amount of research dedicated to it. It can be an effective system however current systems are based on visual systems and laser profilometry which have difficulty in detecting faults in tows below the top layer. Some investigations have been done looking at the combination of machine learning and thermal imaging in the field of AFP defect detection. This is an area that is still undergoing more research to enable the development of effective systems with the challenge to provide low-cost solutions in real-time.

2.5 Automated Control Systems

In automated manufacturing, once a production line has been set up, it is difficult to change or add sensors to processes. Change is however essential to allow flexible manufacturing that enables a focus on quality and cost efficiency. There is need for a powerful, simple and adaptable solution which would allow the addition of multiple external sensors using a standard software interface across multiple robot types.

Robot arms come with, as standard, a control box that manages the robot functionality and some allow the addition of sensors and connection to other systems. Two of the leading robotics companies are KUKA and Universal Robotics. KUKA offers several variations of control box whereas UR have the one control box which covers all robots. Four control box options for KUKA and one for UR are shown in Table 2.3.

Table 2.3 – Control Box Comparisons

Attribute\Device	**D3076-K**	**D3236-K**	**KR C5 Micro**	**KR C5**	**UR Control Box**
USB2.0	8	6	2	0	1
USB3.0	0	2	0	0	1
Ethernet	1	1	1	1	1
Digital Input	0	0	16	0	16
Digital Output	0	0	16	0	16
Analog Input	0	0	0	0	2
Analog Output	0	0	0	0	2

These control boxes are limited in the type and number of sensor inputs offered. Current solutions to add sensor capability fall into two categories: low cost with limited functionality and high cost with limited flexibility. Low-cost options are dominated by Programmable Logic Controller (PLC) whereas Industrial PCs (IPC) are predominately the solution for higher cost options. A PLC normally contains a limited number of inputs, usually 8-12, and even fewer outputs, 4-6. Depending upon the inputs received, commonly digital with one or two inputs capable of analogue compatibility, the PLC will carry out basic commands based on custom programming. More expensive PLCs are capable of networking, allowing readings and values from the inputs to be forwarded onto connected systems. The Siemens S7 series of networked PLCs is used as a viable option for robotic arm manipulation [94]. The Siemens S7-300 PLC has 10 digital inputs, 6 digital outputs and costs around £800 [95]. PLCs have a good level of flexibility for a reasonably low cost but lack the processing power to offer features such as machine learning and do not have the capability to handle complex inputs such as a camera.

Industrial rail or rack mount PCs offer more powerful solutions for automated manufacturing. A rack mount PC is a specialised design of PC which provides the full capabilities of a full PC system but can be mounted into a standard instrumentation rack. The Siemens Simatic Industrial Computer is an example of a rack mounted PC. The system is housed in a large chassis of 117x430x446 mm, provides powerful processing capabilities and contains a large memory capacity. One major drawback of this as a flexible manufacturing option is the cumbersome size and weight of the PC making it difficult to manoeuvre within a working environment. Pricing for an industrial PC can also be considered a drawback as systems often cost several thousand pounds. The Siemens Simatic Industrial Computer retails for around £2,500 [96]. A solution that attempts to solve the size and cost issue is the Beckhoff C60xx Industrial PC range [97]. The systems are designed as low cost, space saving, high computing power industrial PCs. The series ranges from the C6015 Ultra Compact which is 82x82x40 mm with an estimated cost of around £500 [98] to the C6032 which is 129x133x104 mm with an estimated cost of around £2150 [98]. While the C6015 solves the size and cost issue of other industrial PCs, it has limited processing and I/O capability with low AI processing performance. The C6032 solves the size issue and has greater processing capability but at a higher cost.

2.5.1 Outline of Key Control Systems Findings

Research highlighted that PLCs, and industrial PCs are commonly used as control systems in production environments. While PLCs and industrial PCs provide viable options for many applications, there is an opportunity for a system that sits between the two to provide more flexibility and performance for production systems.

2.6 Literature Review Findings and Proposals

Research has demonstrated that the composite materials are becoming an increasingly essential factor in the manufacture of components for multiple applications. Their strength and durability, while minimising weight, make them ideal for the lightweight manufacturing industry.

Analysis of the carbon fibre composite material industry clearly shows that successfully automating the process of defect detection is essential for three key reasons: reducing the interaction of human operatives with hazardous materials, increasing material quality through consistency, cost management through waste reduction.

The research carried out demonstrates that there is no effective, low-cost solution to add machine learning capability easily to current production systems. Research showed that machine learning can be used for carrying out defect detection on composite materials. There are many different types of defect detection that machine learning is used in conjunction with thermography being one such method. Thermography has been used before, but the majority of systems have been found to be expensive. In addition, almost no machine learning thermal data sets have been publicly released making it difficult for this kind of defect detection method to be widely utilised.

Literature analysis revealed that although there are many automated defect detection methods there were only a few which offered real-time solutions. Real-time detection is essential as it allows for defects to be dealt with in-situ rather than after complete designs have been created. This allows the option to repair or restart a component which saves cost and production time. Analysis also highlighted that no current systems provided a complete solution covering AFP defect detection implementing machine learning and thermography in a real-time capacity.

Literature reviewed indicated that some headway has been made in the automation of composite material manufacturing with single robots but not with multiple robots. Multiple robots would be essential to allow a fully automated system capable of handling the complex task of aligning and inspecting composite material components itself. This is important as employing a second robot allows images to be captured at optimal angles for defect detection.

The final part of this research examined the difficulties encountered in automated production, especially for cost sensitive applications and smaller companies. Analysis highlighted the lack of flexibility in current production methods. A need was identified for a system that could allow the optimisation of processing capabilities through the simple and fast addition of sensors.

There is an evident gap between the current PLCs and industrial PCs that are utilised to control production systems and robots. There is need for a small, low cost, flexible industrial control solution with accelerated AI, complex input capabilities and wireless data capture. The proposed solution developed from the literature analysis, was the SSB. The SSB is a solution for highly interconnected smart factories requiring high speed set up changes. A comparison of control options, including the SSB, can be seen in Table 2.4.

Table 2.4 – Option comparison

Control Option	**PLC**	**SSB**	**Industrial PC**
Moveability	Flexible	Flexible	Static
Cost	Low	Low	High
Processing Capabilities	Minimal	AI/ML Capabilities	AI/ML Capabilities

Having a control system with flexible moveability makes it quick and easy to reposition for different or new tasks quickly. Cost is always a factor and needs to be kept low when possible. Processing capabilities can be a differentiating factor and having a good machine learning performance provides support for future advanced development.

In conclusion, analysis of available literature has established that there is a need for a real-time, effective, low cost, fully automated AFP defect detection system that would protect a workforce while maintaining a high level of quality.

The proposal of this research, therefore, is the development and of an SSB using machine learning and thermography in real-time whereby machine learning is required to perform defect detection, an additional cobot is needed to act as an in-situ, versatile thermal camera and the SSB is essential to enable the defect detection to be carried out in real-time.

Chapter 3

3 Smart Sensor Box Development and Deployment

To support the additional machine learning requirements of the AFP system, new hardware was required to provide the necessary real-time computational performance. An increased number of input and output channels was also necessary to support the integration with the networked production system. Additionally, there was a need to control the coordination of multiple robots in the production area. A solution was required that was small and flexible so it could be easily added to existing robots and machines.

3.1 System Introduction

The developed solution was a sensor hub called the SSB, capable of connecting multiple sensors together and accelerating machine learning computation in any area of integrated manufacturing. The SSB was designed so it could be added at any point in a production line to increase the number of sensor inputs and control outputs available. Adding multiple SSBs to any given process would allow a greater number of sensors and controls to be implemented. By standardising the software interface, portable smart sensor boxes are able to be moved around and used along the production line as required. The primary function of the SSB was to connect sensors to the production process and provide localised additional processing capabilities, including machine learning. The goal of the SSB was adaptability and ease of use of both hardware and software interfaces.

3.2 SSB Outline

A high-level overview of the SSB is shown in Figure 3.1. At the centre of the SSB is a single board computer connected to two versatile I/O bays surrounded by a larger number of interfaces.

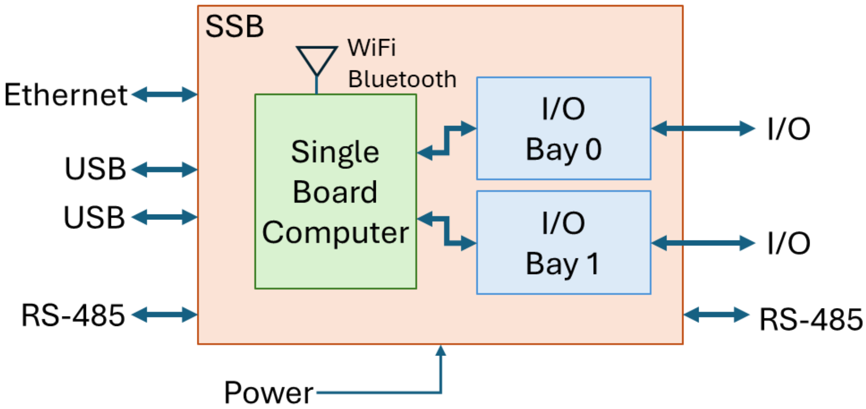

Figure 3.1 – SSB Overview

To accommodate different requirements, four different versions of the SSB were developed. The main difference between the different versions of the SSB is the single board computer it is based on and the machine learning accelerator performance. A list of the SSB versions can be seen in Table 3.1.

Table 3.1 – SSB versions

SSB version	SBC	CPU	AI TOPS	AI Accelerator	Linux Kernel
Version 1	Raspberry Pi 4	4x ARM A72 @ 1.8 GHz	-	-	6.6
Version 2	Raspberry Pi 4	4x ARM A72 @ 1.8 GHz	4	Coral	6.6
Version 3	Qualcomm RB5	4x ARM A77 @ 2.8 GHz	15	Hexagon DSP	4.15
Version 4	Raspberry Pi 5	4x ARM A76 @ 2.4 GHz	26	Hailo-8	6.6

All of the SSB's single board computer options use ARM Cortex-A class Central Processor Unit (CPU) cores. These are high-performance high-speed cores designed to run rich operating systems like Linux. The Cortex-A (Application) CPUs were designed for mobile and embedded applications. All of the cores have hardware floating point, ARM's NEON single instruction multiple data (SIMD) accelerator, multiple levels of cache and a large 128-bit wide system bus. The deep CPU pipelines allow the CPU cores to run at high frequency. The Raspberry Pi uses a newer version of the Linux kernel. This is important to provide long term support for the software in a production environment.

Adding machine learning acceleration was a key feature of the SSB. The industry standard of maximum possible AI performance is measured in Tera Operation Per Second (TOPS). An operation is a simple mathematical operation of add or multiply. To determine the TOPS for a system all additions and multiplications that can be done at the same time for each arithmetic unit (AU) need to be combined as shown in Eq. 3.1.

$$TOPS_{total} = \sum AU(addition_{ops} + multiplication_{ops})/second \quad (3.1)$$

The standard data size is 8-bit integer numbers. AI accelerator hardware is designed to use 8-bit integers and machine learning tools quantise the models to use 8-bit data size. The total performance on the Raspberry Pi 5 is equal to the CPU and Hailo-8 accelerator combined as shown in Eq. 3.2.

$$TOPS_{Raspberry\ Pi\ 5} = CPU\big(SIMD\ AU_{ops}\big) + Hailo8\big(AU_{ops}\big) \quad (3.2)$$

Using 8-bit integer data, the ARM NEON SIMD accelerator can do 16 parallel operations per clock cycle. There are 4 CPU cores with NEON running at 2.4GHz. The MAC (multiply accumulate) counts as 2 operations because it does a multiply and an accumulate. The Hailo-8 accelerate provides 26 TOPS. The complete Raspberry Pi with the Hailo-8 accelerator is shown in Eq. 3.3 which is then solved in Eq. 3.4.

$$TOPS_{Raspberry\ Pi\ 5} = \left(2_{MAC\ ops} * \ 4_{cores} * 2.4G_{clock} * 16\frac{ops}{clock}\right) + 26TOPS \quad (3.3)$$

$$TOPS_{Raspberry\ Pi\ 5} = 0.307\ TOPS + 26\ TOPS = 26.307\ TOPS \quad (3.4)$$

The Raspberry Pi 5 with Hailo-8 AI+ HAT (Hardware Attached on Top) accelerator provides the highest performance option with 26.3 AI TOPS which compares to the development laptop systems with a NVIDA GTX 4060 Graphics Processing Unit (GPU) providing 233 AI TOPS [99].

The SSB is housed in a custom designed enclosure that is 3D printed. This is a practical solution because of low-cost 3D printers being readily available. It enables a small compact design that is optimised for easy relocation and mounting in the production area or on a robot. The flexibility of 3D printing and design means that it would be easy and quick to create new versions of the enclosure for custom applications, for example a custom enclosure for a new type of robot used. The design can be 3D printed with different material if different characteristics are required. Most 3D print filament is non-conductive, and some types are self-extinguishing fire retardant, meaning it will slow down and eventual stop a flame. Another advantage is that if the case is damaged a replacement can be quickly and easily 3D printed.

3.2.1 SSB Overview

The SSB was designed as a hub to take multiple inputs for several different types of sensors and provide access to the acquired data via multiple different output methods. With simple sensors connected, the smart sensor box can take non-interconnected sensors and make results gained more accessible for integrated production processes. The SSB can make the results available in different formats and over different interfaces (serial or network) to meet the requirements of the production system.

The multiple inputs allow real-time feedback from the SSB to the connected production process. Multiple smart sensor boxes can be placed in parallel to allow an even greater number of sensor inputs to be connected. A detailed internal view of the SSB is shown in Figure 3.2.

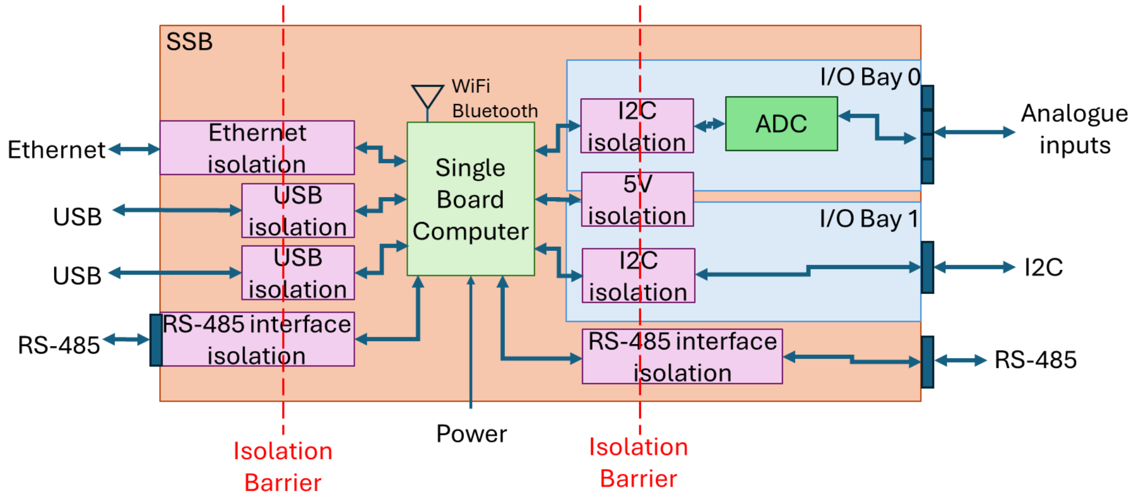

Figure 3.2 – SSB Internal Detailed View

The SSB is composed of several separate components organised in a single housing. This allows components, mainly isolation interfaces, to be replaced separately. It also allows flexibility in the interfaces that are used because it is possible to replace them with different interfaces.

3.2.2 SSB Inputs and Outputs

The interfaces provided on the SSB are: 2 USB, 2 RS-485, 2 I/O bays, Ethernet and wireless. A key feature of these interfaces is that they all provide Galvanic isolation. This is essential in an industrial production environment to protect from damaging the SSB with high voltages.

The I/O bays are a way to customise the interfaces available on the SSB. Each bay has an isolated 5-volt supply and an isolated I2C interface. In addition, there is an opening in the enclosure where screw terminal blocks or other connectors can be exposed. Research indicates that the main use of the I/O bays would be to add analogue inputs to the SSB, however other interfaces, like relays or digital I/O, are also possible.

3.2.2.1 Analogue Inputs

The first board for the I/O bays is an analogue input board with four pairs of screw terminals that are situated on the edge of the box to allow a range of sensors to be connected. This board contains an analogue to digital converter which converts input voltages from 0 to 5 volts into a digital number that can be used by the control system. As separate circuits are connected to each pair of screw terminals, multiple types of sensors can be implemented at the same time on the SSB. The four input channels can be configured as four single-ended ground referenced inputs or 2 differential inputs. Differential inputs are useful when measuring signals that do not reference or go down to ground. Each input channel has a low pass filter with current limiting resistor and diodes to clamp the inputs for protection.

To meet the requirement of the Analogue to Digital Converter (ADC) interfaces and physical layout in the enclosure for the SSB a custom Printed Circuit Board (PCB) was designed for the analogue inputs. The PCB was designed using the advanced cross platform open source KiCAD tools [100]. KiCAD allowed the generation of the schematics, PCB layout and manufacturing files. It also provided design rule checking (DRC) to help detect errors in the PCB. The PCB traces were routed manually to meet part placement requirements. In order to provide low noise on the ADC inputs a 4-layer PCB was designed. The internal layers are 5-volt power and ground with the top and bottom layer used for signal routing. This standard arrangement provides low noise with the internal layers and flexibility in routing the signals on the top and bottom layers. The PCBs included standard plated through holes and surface mount devices which were all hand soldered to the boards.

The PCB was designed to take plug-in modules from the company Adafruit. The use of plug-in modules allowed them to be easily replaced if they got damaged. This was a reasonable precaution because the inputs connect to external circuits that can have a wide range of voltages on them. The PCB has an I2C isolator and an ADC module on it. The schematic for the analogue input PCB is shown in Figure 3.3. Two footprints for the ADC module were provided to support two versions of the module. The footprints for the two modules are on top of each other to save space since only one will be used at a time. The layout of the PCB is shown in Figure 3.4.

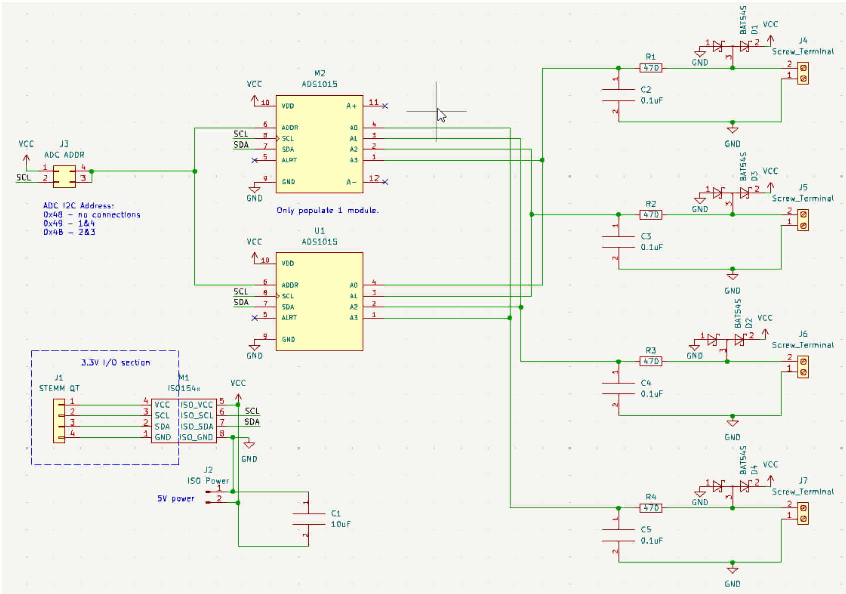

Figure 3.3 – Analogue Input Schematic

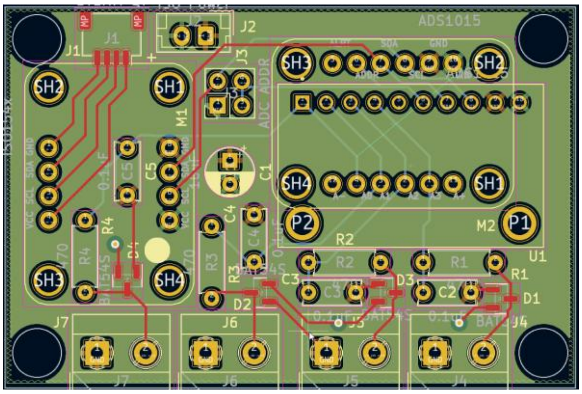

Figure 3.4 – Analogue Input PCB

Currently there are 2 ADC modules that can be used. The modules are pin and size compatible so the ADCs can be changed to meeting the specific requirements of the production system. The two modules are based on the devices listed in Table 3.2.

Table 3.2 – ADC devices

ADC Device	Conversion Speed	Resolution
ADS1115	Up to 860 samples/sec.	16-bit
ADS1015	Up to 3.3 k samples/sec.	12-bit

The Texas Instruments made ADCs are sigma delta based converters with programable gain amplifiers to provide accurate results over a wide range of input voltages. A sigma delta ADC converter provides low noise high resolution results, however at a low conversion speed compared to other ADC architecture. If the magnitude of the input signal is smaller than the reference supply (5-volts) then the programmable gain amplifier can be used to magnify the signal effectively creating a higher resolution result, however if the input range is exceeded the result value will be clipped.

3.2.2.2 Testing and Results

Initial testing of the SSB analogue inputs was carried out to prove the repeatability and accuracy of data collection using simple sensors. Various resistive loads were connected to the SSB and repeated values taken. The results, including standard deviation, are shown in Figure 3.5. The small variance between readings highlights the consistency of the SSB.

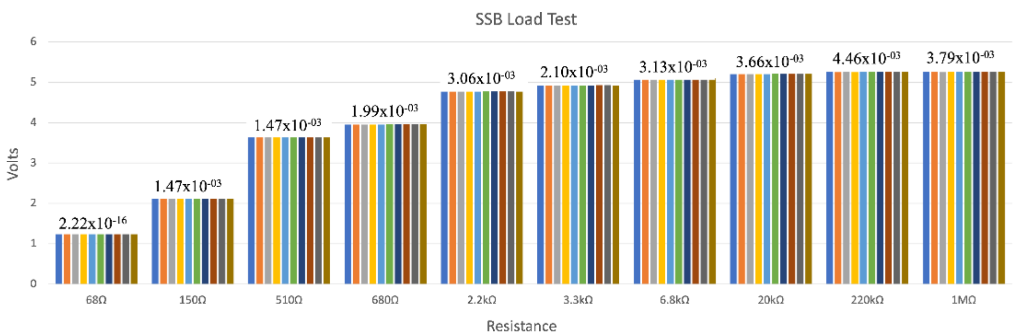

Figure 3.5 – Resistance and Load Testing Results with Standard Deviation Values

3.2.2.3 Smart Inputs

The second board for the I/O bays is a smart sensor interface. It provides isolated power and an I2C interface. I2C is a widely used synchronous serial half duplex interface that supports multiple devices on the bus. It uses a separate data and clock signal that are single ended. The default configuration of the I2C bus on the SSB supports 100kbit/s but can be configured to support up to 400kbit/s. Multiple devices can be connected to the I2C with a 7-bit addresses part of the protocol to select a specific device. Some devices used multiple addresses on the bus, for example the ADC converters on the analogue input PCB use 4 addresses for configuration and results.

The I2C interface is available on a 4-pin connector with 2 pins for power (5 volts and ground) and 2 pins for I2C (SDA – serial data and SCL – serial clock). I2C is a common interface used on many sensors that can measure conditions like temperature, pressure, light intensity, colour, humidity, air quality, gas concentrations and particulates. This provides a flexible, simple low-cost, short-range connection for local sensors.

3.2.2.4 RS-485 Interface

RS-485 also known as TIA-485(-A) or EIA-485, is a half-duplex communication standard. Networks using it can communicate over long distances (up to 1200m) and in electrically noisy conditions. The interface can transmit data at up to 10Mb/s, however it depends on the length of the network. As an estimate the speed in bit/s multiplied by the length in metres should be less than 10^8. It uses a differential signal and can be configured as a multidrop linear bus. The communication is low cost and simple requiring only a single twist pair of wires. These characteristics make it useful for industrial control applications.

The galvanically isolated RS-485 interfaces on the SSBs use the M5Stack isolated RS-485 module which converts a standard serial UART interface to RS-485. The interface uses the ChipAnalog CA-IS3082W device to isolate and convert the control signals and a transformer to isolate the power. The M5Stack interface supports up to 500kb/s communications.

The RS-485 standard only defines the physical communication interface not the protocol for data transfer. In industrial automation environments Modbus and Profibus are commonly used for automation of communications over the RS-485 interface.

The SSB has 2 RS-385 interfaces. The first interface is designed to connect to the automation control system of the production environment. The SSB can provide data and status information to the production control system when operating in client mode. The second RS-485 is used by the SSB to communicate with sensors and devices that it is controlling and using, acting as the local master in server mode controlling communications. The RS-485 ports are not dedicated to the proposed functions, either port can be used in server or client mode. The only differences in the ports are the software and configuration.

3.2.2.5 Ethernet

All versions of the SSB have an RJ-45 Ethernet connector. This connection is galvanically isolated with a transformer, which is part of the Ethernet standard. The Ethernet port supports a 1Gb/s connection. It can be used when a wired network connection is needed. The Ethernet interface is the highest speed port on the SSB and useful when large amounts of data or fast response times are needed.

3.2.2.6 USB

There are two USB2 ports that support USB2 full speed transfers up to 12Mb/s. The ports are galvanically isolated using the Analog Devices' ADUM3160 device. The power is galvanically isolated using a 5-volt to 5-volt converter. The USB ports are important for connecting advanced peripherals like cameras. They are also useful for adding new network interfaces like a 5G modem. USB is a widely used interface standard and there are a large number of devices and sensors that can be attached using it. It is possible to use an external USB hub to increase the number of USB ports, however any protection between the ports on the hub would need to be provided externally.

3.2.2.7 Wireless

All versions of the SSB support Wi-Fi and Bluetooth communications. A minimum of dual band IEEE 802.11ac Wi-Fi and Bluetooth 5.0 with Bluetooth Low Energy are supported. The Wi-Fi is useful when the SSB needs to be added to an existing system and there is no wired network access available. It is also useful as a backup network interface if the wired connection is disrupted.

The Bluetooth interface can be used to add sensors located close to the SSB when they cannot be connected with wires. Bluetooth sensors can use the Bluetooth Low Energy protocol. This is useful for sensors because they run on batteries and a low energy protocol will extend battery life and because the Bluetooth Low Energy protocol does not require a pre-configured connection to the SSB. The Bluetooth Low Energy sensors broadcast the data they can provide to the SSB. The SSB can then ask for the data updates as needed.

3.2.3 Version 1

The first version of the SSB is based on the Raspberry Pi 4 single board computer. This is a good option because it is small, low power, low cost, good performance and has plenty of I/O expansion capabilities. The Raspberry Pi has significant hardware and software support provided by the Raspberry Pi Holdings Plc company and a large enthusiast community with very active discussion forums. There are special programs and support for using the Raspberry Pi boards in commercial applications for example the Raspberry Pi compute module 3 is used in the Wallbox home EV charger [101]. The Raspberry Pi computers are used in commercial PLCs and industrial controllers from over ten different companies [102]. Raspberry Pi Holdings Plc have stated that the Raspberry Pi 5 will stay in productions until at least 2036 [103] making it practical for the long lifetime of an industrial product.

A key advantage of the Raspberry Pi is its strong software support and updates. Software releases and updates are maintained for all the versions of the Raspberry Pi hardware. The current release of Debian Linux 12 (Bookworm) operating system using version 6.6 of the Linux Kernel is available for the Raspberry Pi. There is also a wide and large selection of additional software packages available that are pre-compiled and easy to install. This includes many machine learning programs like PyTorch, TensorFlow and TensorFlow lite making it easy to develop and train machine learning models on a powerful desktop computer and then translate them to the Raspberry Pi. There is good software available with examples to facilitate the development and translation of the machine learning models.

Other operating systems are available for the Raspberry Pi. The popular Robot Operating System 2 (ROS2) is available as an installable image with Ubuntu Linux [104]. Ubuntu on the Raspberry Pi has full support available from the ROS organisation.

The Raspberry Pi has a 40 pin I/O connector and a standard specification for mechanical and electrical connections to it. Boards that interface to it are called HATs (Hardware Attached on Top). This interface is standard across all Raspberry Pi's (except the original Raspberry Pi 1) so hardware add-on boards can be designed once and used on all versions of the Raspberry Pi. The Raspberry Pi based SSBs use this feature with a simple HAT breakout board to connect power and the serial interfaces (RS-485 and I2C) to the isolated SSB interfaces.

Robust tools are provided to program the non-volatile memory (NVM). This includes the embedded Multi-Media Card (eMMC) memory on the Raspberry Pi compute modules and the Secure Digital (SD) card for the regular Raspberry Pis and the compute module lite versions. These tools follow current legal regulations, for example no default password, a custom password must be specified. The Raspberry Pi versions, 4 and later, support secure boot which is where the boot image code must be verified before it is executed to prove it is authenticated to be executed. In an industrial production application secure boot is required by the European Union Cyber Resilience Act [105] and expected to be required by the United Kingdom (UK) Cyber Security and Resilience Bill which is being created.

3.2.4 Version 2

The second version of the SBB, the SSB2, uses the same hardware as the SSB1 but with the addition of a Google Coral Edge Tensor Processor Unit (TPU) connected to a Universal Serial Bus 3 (USB3) port. The Coral Edge TPU is a machine learning accelerator for neural network applications developed by Google. The Coral is supported with Google's TensorFlow software. The USB3 port is used to allow high speed access (up to 5Gbps) to the Coral, a USB2 port can be used however the performance would be lower due to its lower transfer speed (up to 480Mbps). With the addition of a Coral Accelerator, machine learning processes can be carried out twenty-five times faster than a Raspberry Pi 4 on its own (using USB3.0 running Mobilenet v2) [12]. An internal photograph of the SSB2 is shown in Figure 3.6. The cable connecting the Coral to the Raspberry Pi 4 has been omitted for clarity.

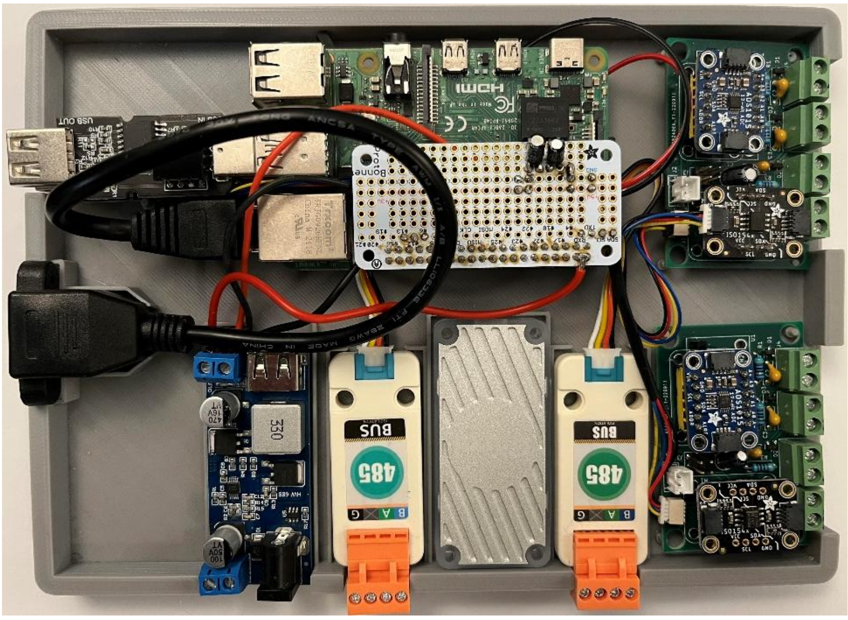

Figure 3.6 – SSB2 Internal View

Google released the Coral Edge TPU in 2019 with support for TensorFlow Lite models. Google appears to have stopped supporting the Coral and have not updated Coral Edge TPU runtime in a long time. The Coral runtime no longer works with the current version of the TensorFlow Lite tool chain. The company Ultralytics has continued support for the Coral hardware and provides software updates to get the Coral Edge TPU running with the current TensorFlow Lite tools on the Raspberry Pi. The tool chain flow for deployment of a model on the Coral hardware is shown in Figure 3.7.

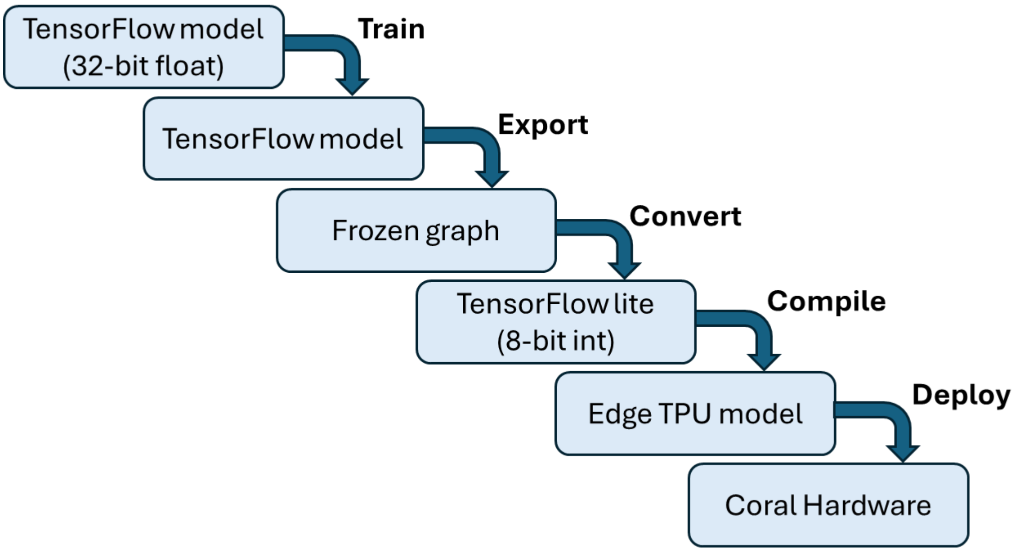

Figure 3.7 – Machine Learning Model to Coral Hardware

The process of converting a model starts with the development of the model using TensorFlow with 32-bit floating point numbers. The model is then exported and converted to convert it to a quantised model in TensorFlow lite using 8-bit integers. Finally, the model is mapped to the hardware architecture of the Coral.

3.2.5 Version 3

Qualcomm released a robotics development platform called the Robot Board 5 (RB5) which looked like a good upgrade from the Raspberry Pi 4. It had faster CPUs, an integrated AI accelerator, high speed camera support and an option for 5G communication which would be good for the SSB. On the downside it was larger, used high density connectors, was higher power and much more expensive. The QRB5165 processor in the RB5 is much more powerful but also much more complex to use.

The main challenge with using the RB5 was software and support. The RB5 runs Ubuntu Linux 18.04. Ubuntu 18.04 is a LTS (Long Term Support) version of Ubuntu with support for 5 years after release and a future 5 years of security updates. It is currently only being updated for security issues. It is based on the Linux kernel version 4.15.

There are no pre-compiled boot images so a separate computer running Ubuntu 18.04 is required to create the boot image file. This effectively required a 2 computer development setup since the machine learning development workflow was done on Windows. The Windows computer also ran the Android Debugger Bridge to initialise and flash the RB5's memory. The flash tool over the Android Debugger Bridge interface did not work reliability. It was not always able to program the on-board FLASH memory. This caused issues and slowed development.

The RB5 could be used in a headless mode (without a display) or with a display and keyboard. In headless mode a remote network connection using SSH (Secure Shell) with the Putty tool was used. Putty is a windows application that can communicate using the SSH protocol.

Several of the demonstration programs from Qualcomm, providing information about the device, Wi-Fi and ROS2, worked on the RB5. Upgrading the Ubuntu operating system and software components system however was inconsistent, and software packages were not always available.

The main issue encountered was that the TensorFlow software would not install, compile and run correctly. None of the software was available from the RB5 repository and the generic versions of the code would not compile because there were missing support modules. The TensorFlow software could not be compiled, and it could not be fixed. Attempts were made to get support directly from Qualcomm with very little response and no solutions. Queries were also made on the RB5 support forums. Again, there was very little response and no solutions. The RB5 seemed like a good solution for a higher performance SSB, but without good hardware and software support it was not a viable single board computer option for the SSB3. Qualcomm is not providing full support for the RB5 anymore. These led to the development of the SSB3 being stopped.

3.2.6 Version 4

The fourth and final version of the SSB, the SSB4, is based on the current highest performance Raspberry Pi, the Raspberry Pi 5. The Raspberry Pi 5 provides 2-3 times the performance of the Raspberry Pi 4. The Raspberry Pi 5 board is the same size so can fit into the existing SSB housing. The only differences in the Raspberry Pi boards are that the Ethernet and USB2 ports have swapped locations, and the Raspberry Pi 5 needs a higher current power supply. The same HAT board in the SSB1 and SSB2 is used with the Raspberry Pi 5 in the SSB4.

An important new feature on the Raspberry Pi 5 is a Peripheral Component Interconnect Express (PCIe) port. PCIe is a high-speed serial computer expansion standard from the personal computer industry. It was possible to have PCIe on the Raspberry Pi 4 compute module with its I/O board, however it had limited support and compatibility. A PCIe port can support up to sixteen data lanes which consist of a pair of transmit and receive differential signals. The PCIe port on the Raspberry Pi 5 supports a single data lane. Different versions, called generations, of the PCIe specification support different data transfer speeds. The Raspberry Pi 5 PCIe is certified as a Gen. 2 interface which provides a transfer speed of 4Gb/s.

With the addition of the PCIe interface it is possible to attach new high-performance accelerators to the Raspberry Pi 5. There is an official AI HAT+ from Raspberry Pi. The AI HAT+ features the Hailo-8 AL processor that provides 26 TOPS performance. The hardware PCIe interface on the Raspberry Pi 5 supports the higher speed Gen. 3 standard however it has not officially been approved for this speed for all applications. The AI HAT+ uses the PCIe and 40-pin I/O connector on the Raspberry Pi 5. The PCIe interface to the AI HAT+ has been tested to Gen. 3 operation with 8Gb/s transfers. This provided a much higher interface bandwidth compared to the Coral on a USB3 port. It is possible to use multiple Hailo-8 devices in a system to provide over 200 TOPS performance, however PCIe adapters and expanders would be required.

Complete software support is provided for the AI HAT+. Standard development frameworks like TensorFlow, TensorFlow Lite and PyTorch are supported. The software seamlessly integrates with the Raspberry Pi AI development tools.

The addition of an AI HAT+ makes the stack of PCBs in the SSB slightly higher so the top half to the SSB case was modified to support the SSB3 hardware.

3.3 Software

There are two parts to the software on the SSB: the machine learning defect detection and the general PLC control. The two parts run in parallel simultaneously. The machine learning defect detection software includes coordinating the movement of the cobot arm with the thermal camera attached and analysing the thermal images collected. The custom developed machine learning AFP defect detection software (detailed in section 5) analyses the thermal images for defects and provides feedback to the production system.

3.3.1 General PLC software

Rather than developing custom PLC software the open source secure OpenPLC package was used. It is a powerful PLC system that supports the industry standard IEC 61131-3. This is the standard for programming PLC controllers. The IEC 61131-3 defines 3 graphical and 2 text based programming languages. The graphical programming languages are Ladder diagram (LD), Function block diagram (FBD) and Sequential function chart (SFC). The Ladder diagrams provide a representation that is similar to mechanical systems to provide an easy interface for production engineers not familiar with programming. The text-based languages are Structured text (ST) and Instruction list (IL). Open PLC provides a Graphical User Interface (GUI) for the graphical languages, as shown in Figure 3.8 which gives an example of the GUI for a Ladder diagram. These languages are in common use in the industrial control industry.

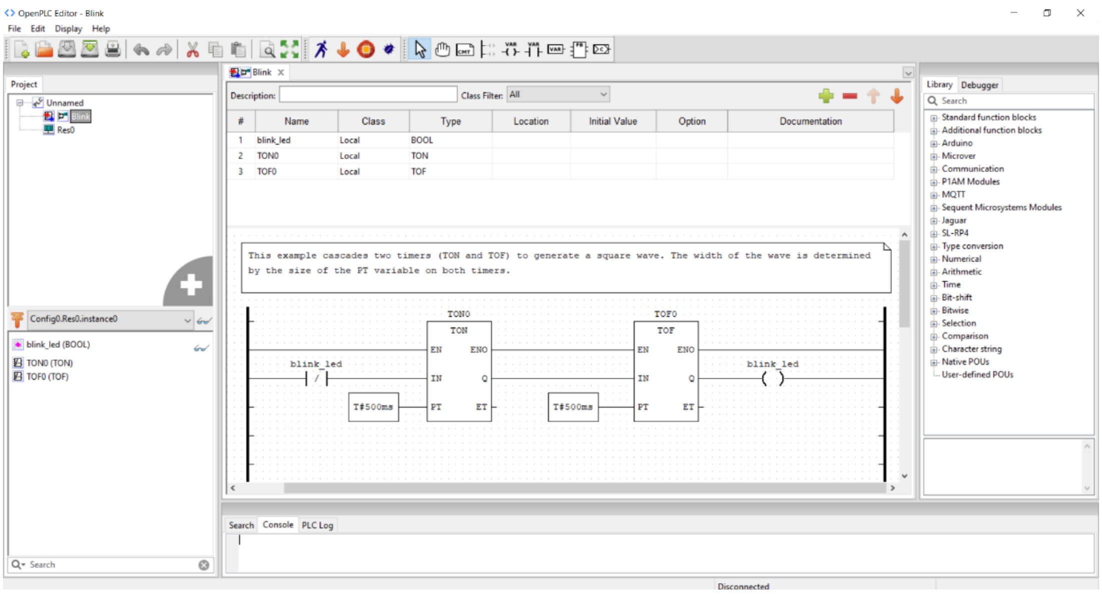

Figure 3.8 – OpenPLC GUI

In addition to programming and I/O control, communication is key to a PLC. OpenPLC supports controlling I/O directly connected to the SSB or over a communications interface. Modbus is the de facto standard communication protocol used by PLC controllers. [106] OpenPLC and the SSB support Modbus over either network or serial connections. Modbus TCP is the network version of Modbus and it can be used over Ethernet or Wi-Fi. Modbus RTU is the serial version of Modbus and it can be used over RS-485 or USB interfaces. OpenPLC supports acting as a server controller for I/O or as a client responding to requests.

3.3.2 Defect Detection Software

The custom developed machine learning AFP detect detection software has two tasks. First, it controls the cobot sending instructions to drive the cobot's position and movement. Second, it provides a simple pass / fail to the production system based on the machine learning processing of the thermal images. It uses the OpenPLC Modbus (either TCP or RTU) interface to send results to the production control system.

3.4 Example SSB Applications

The SSB can be used in many ways throughout an automated production line to allow the easy addition of multiple sensors and data management. The following scenarios highlight uses for a smart sensor box.

3.4.1 Scenario One

Sensors need to be added to production line systems to implement new production features or to enhance quality systems. The SSB can be used to add additional sensors and cameras along the production line or within a production cell as shown in Figure 3.9. A camera can be set up at the start of a production cell to monitor incoming material. Images from the camera are then processed by the SSB. This processed data is then sent via fifth generation (5G) cellular data or other network interface to a process control unit which uses the input to manage the next steps for the material. A second camera can be set up with another SSB at the end of the production cell. Here images from the camera can be used for quality control to monitor the outgoing product from the cell.

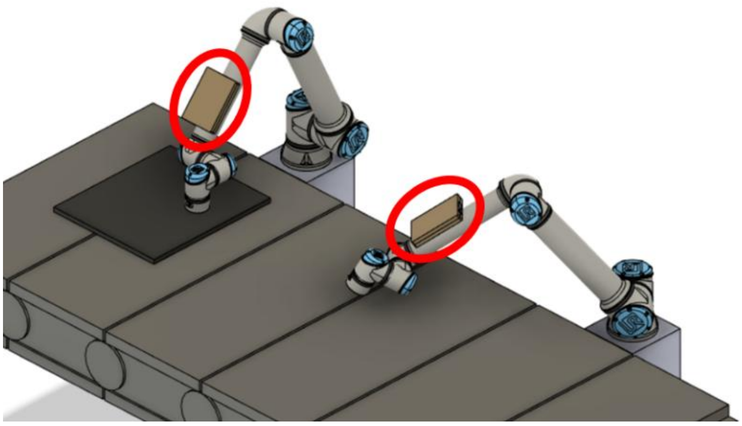

Figure 3.9 – SSB Additional Sensors

3.4.2 Scenario Two

It would be advantageous to have real time automated defect detection within the area of composite material manufacturing. Cameras and additional sensors can be deployed to analyse AFP systems allowing quality and part inspection to be carried out. In an agile production system, where many different designs are produced, machine learning acceleration on the SSB allows defect detection to be carried out. Collaborative robots can be utilised to create composite materials and capture images suitable for machine learning defect detection. An example cell can be seen in Figure 3.10.

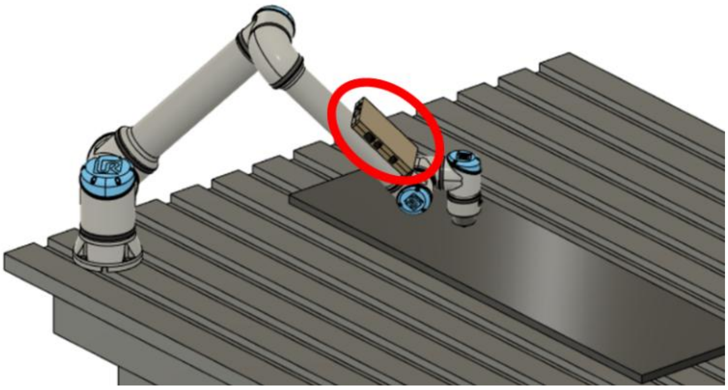

Figure 3.10 – SSB Defect Detection

3.5 SSB Discussion and Summary

This section presented a new solution to increase the flexibility and functionality of automated production systems. Testing displayed accurate and repeatable measurements taken via the sensor box system. The SSB acts as a highly intelligent smart sensor hub allowing the addition of multiple sensors anywhere along a manufacturing line. The key advantages of this solution are that it is flexible, accurate, easily repositioned, low cost and has the capability to provide smart sensor functionality to enhance existing simple sensors. This section introduced key application areas for the SSB highlighting its ability to manage cobot movements and machine learning object classification processes. This novel proposal, with its ability to add multiple sensors within a production process, offers adaptability, increased quality control and greater cost effectiveness to any production system.

There was interest from the Scottish Lightweight Manufacturing Centre (LMC) in the SSB as part of this research project and interest has been expressed by the National Manufacturing Institute Scotland (NMIS) in possibly using it in a future project or some related industrial automation research.

Chapter 4

4 Robot and Cobot Control

The purpose of this chapter is to show how control and synchronisation of multiple robots was achieved, highlighting the hardware and software research and implementation. It will cover the roles and capabilities of these robotic systems, the mechanisms employed to control their movements, and the simulation tools used to ensure their effective integration. Together, these systems demonstrated the importance of advanced robotics in composite material manufacturing by improving quality and efficiency.

4.1 Robot Analysis

The integration of robotic systems in the automated manufacturing processes plays a pivotal role in enhancing efficiency, precision, and safety. This is particularly evident in the production of advanced composite materials through AFP, where composite fibre tows must be laid with exceptional accuracy and repeatability to ensure a defect-free layup. The complexity and demands of AFP processes necessitate the use of high-precision robots and collaborative systems to achieve the desired quality, repeatability, and adaptability. Cobots, are particularly beneficial in these applications as they can safely work alongside human operators and other robotic systems, enabling greater flexibility and real-time adjustments in response to variations in the manufacturing process. By incorporating collaborative and adaptable robotic systems, manufacturers can improve production speed, minimise defects, and enhance overall process reliability, making AFP more versatile for a wide range of applications.

This research focused on the development and implementation of a robotic system for AFP that incorporated real-time defect detection using thermal imaging and machine learning. The experimental setup included two key robotic systems: the KUKA KR120 industrial robot and the UR10e collaborative robot (cobot). These robots were used because the KUKA KR120 was equipped with the AFP end effector at the Lightweight Manufacturing Centre (LMC) and the UR10e because it could be added to the AFP work cell. The UR10e was available at both the LMC and University of Strathclyde for consistent development. The KUKA KR120 performed the AFP process by laying composite tows with a bespoke end effector, while the UR10e, equipped with a thermal imaging camera, followed the KUKA robot to detect defects during the layup process.

An AFP tow refers to a narrow strip of composite material that is used in the AFP processes. These tows are laid down sequentially by robotic systems onto a mould or surface to create composite structures layer by layer. AFP allows for precise placement of individual tows, enabling manufacturers to create highly optimised composite structures. Tows can be placed at varying angles and orientations to enhance strength and structural performance in specific directions. Unlike larger composite sheets, tows minimise material waste by only placing material where needed. AFP systems use these tows in a continuous and automated manner, reducing manual labour and increasing production speed. AFP tows are typically thin and narrow strips rather than wide sheets. The small width allows for intricate and variable layup designs, making it easier to create complex 3D designs. The thin design enables precise control over fibre orientation, optimising load paths and reinforcing specific areas as needed. If a defect does occur in a tow, it can be removed and replaced without affecting the entire composite structure. Multiple tows can be placed in parallel or at different angles to create complex and strong composite laminates.

Using the two robots, a KUKA KR120 for AFP and a UR10e cobot for defect detection, struck a balance between efficiency, precision, and system complexity. A single robot would have been inefficient, as it would have had to switch between tow placement and defect detection, significantly slowing down the process. Having three or four robots could have introduced unnecessary complexity, requiring additional coordination, workspace, and synchronisation, increasing costs and potential interference. Two robots ensured continuous operation, with one focusing on AFP and the other on real-time defect detection, optimising workflow without excessive redundancy.

4.1.1 KUKA Robot

The KUKA KR120 is a six-axis industrial robot specifically designed for applications that require high precision and reliability with a positional accuracy of 0.05mm [107]. AFP is a critical manufacturing process utilised in industries such as aerospace and automotive to produce lightweight and high-strength composite materials. In this research, the KUKA KR120 was responsible for accurately laying down continuous tows of composite fibres following a predefined pattern. To fulfil this role, the robot was equipped with a customised end effector capable of handling the composite tows and ensuring their proper application.

The AFP process involved the use of a tow delivery mechanism integrated into the robot's end effector. This mechanism continuously fed the composite fibres while a heated lamination tool ensured proper adhesion of the material to the mould.

An AFP end effector is a specialised robotic tool designed to precisely place, compact, and cut composite fibre tows during the AFP process. Mounted on the KUKA KR120 robot, it featured a tow feed system that independently controlled a single narrow composite strip, allowing for customised fibre placement. A compaction roller applied pressure to ensure proper adhesion, while an integrated infrared heating system activated the resin for strong bonding. The tow cutting mechanism enabled precise starts and stops, allowing smooth fibre orientation changes without overlaps.

In AFP, a heated bed is a temperature-controlled surface onto which composite tows are laid. The heated bed helped maintain dimensional stability, preventing warping or fibre misalignment due to thermal stresses. By keeping the tows slightly tacky, the heated bed improved layup accuracy, preventing shifting or lifting after placement. The heating system used embedded electrical elements, one in the centre of each quadrant, to provide uniform temperature control to the entire bed.

The KUKA KR120's ability to follow complex, curved paths with high precision was crucial to achieving the desired structural integrity and performance of the composite parts. The robot's movements were controlled by pre-programmed trajectories that utilise inverse kinematics to calculate the joint angles necessary for the end effector to reach specific points. Path planning algorithms further ensured that the robot's movements were smooth and accurate, avoiding deviations that could compromise the quality of the composite layup.

In this study, the KUKA KR120 industrial robot arm, featuring six degrees of freedom, was equipped with a bespoke single tow Accudyne AccuFP 1000+ end effector specifically designed for composite material layup. The robot was programmed to follow a predefined trajectory using GTech Vericut Composites Programming (VCP) software. GTech VCP software was used to optimise AFP layup paths, simulate defects, and integrate with robotic systems for precise fibre placement. Its key features included layup path optimisation, defect simulation, real-time robot integration, customisable programs, and process parameter control, ensuring accuracy and efficiency in composite manufacturing. This software facilitated precise control over various aspects of the layup process, including the speed, direction, and orientation of the robot arm. The composite material was applied onto a heat-controlled tooling plate, with the AFP end effector additionally equipped with an infrared heat lamp to preheat the tooling area directly ahead of the oncoming tow. This preheating mechanism ensured optimal adhesion of the composite material to the tooling plate.

4.1.2 UR10e Cobot

The UR10e, a collaborative robot developed by Universal Robots, is designed to operate safely alongside humans in industrial settings. Unlike traditional industrial robots, the UR10e incorporates integrated force-limiting technology, allowing it to halt operation immediately upon sensing an obstruction. In this research, the UR10e served as a secondary system tasked with capturing thermal images of the composite material laid down by the KUKA KR120. These images were analysed in real-time to identify potential defects in the layup process.

The UR10e is a cobot from Universal Robots, featuring six degrees of freedom and a 1300 mm reach, making it ideal for applications requiring extended movement. It has a payload capacity of 12.5 kg, a speed per joint of 1 m/s and weighs 33.5 kg. It includes advanced safety features such as force and torque sensors, collision detection, and emergency stop functions, allowing it to operate alongside humans safely.

A camera was required to take thermal images of the tows, as they were laid, in order to provide data to be evaluated for defect detection with the machine learning models. An initial thermal imaging development system was built and tested for its efficacy as a means of providing the required data.

The initial development system was based on a Raspberry Pi 4 with two image sensors. The portable unit was housed in a 3D printed case with an internal battery, LCD touch screen on the back and a shutter button on the top as shown in Figure 4.1. A program was developed to enable images, from both sensors, to be captured and displayed on the LCD when the shutter was pressed.

Figure 4.1 – Initial Thermal Imaging Development System Front View

The two image sensors were the Melexis MLX90640[108] and the OmniVision OV5647[109]. The MLX90640 had a far-infrared sensor with a range from -40 to 300°C, however it only had a resolution of 32 x 24 pixels. This sensor could potentially have been suitable for this application, but the resolution was too low. The OmniVision OV5647 had a 5 MP (2592 x 1944 pixels) camera. This was a standard optical sensor with the infrared filter removed. This sensor was not suitable for the application because it was unable to detect temperature variations which can be seen from the image of a hot cup of liquid in

Figure 4.2.

Figure 4.2 – Initial Thermal Imaging Development System Back View

As these sensors proved unsuitable for the application requirements, an alternative thermal imaging camera was sourced. Table 4.1 shows a selection of general-purpose industrial thermal cameras for comparison.

Table 4.1 – Thermal Camera Comparison

Camera	**Detector Resolution**	**Temp Range**	**Price (USD)**	**FOV (°)**	**Interface**
Fluke RSE300	320 x 240	-10 °C to +1200 °C	$15,000 – $25,000	34	Ethernet / Camera-link
Optris PI450i	382 x 288	-20 to 900°C	$10,000 – $20,000	18-80	USB 2.0 / Ethernet
FLIR T560-42	640x480	-20°C to 120°C	$30,000 – $45,000	14-42	USB / Wi-Fi / SD Card
AMETEK LAND ARC-8-60-500-HF	384 x 288	0°C to 500°C	$25,000 – $40,000	11-60	Ethernet / Camera-link

The AMETEK Land ARC thermal imaging camera was selected. It was mounted onto the UR10e cobot, enabling it to capture thermal images of the composite material during the AFP process. The robot was programmed to follow a predefined trajectory that allowed the thermal camera to capture images at regular intervals. These images were subsequently transmitted to a computer for analysis using a machine learning algorithm developed for defect detection. This integration of robotic movement and thermal imaging facilitated real-time monitoring and quality control of the AFP process.

The AMETEK Land ARC thermal imaging camera was a general-purpose thermal imager designed for industrial applications. It offered a temperature measurement range from 0°C to 500°C and was equipped with Viewer Software for user-friendly monitoring and visualisation of thermal data. Additionally, the ARC camera supported direct connection through standard industrial Ethernet providing simple, stand-alone operation. The camera offered a resolution of 384 x 288 pixels. The AMETEK Land ARC thermal imaging camera was interfaced with a laptop via Ethernet, allowing for seamless communication and control. Using the Viewer Software installed on the laptop, the camera was successfully integrated with the software, enabling access to various functional settings. The software provided several key functionalities, including the ability to set the temperature range, adjust the colour scale for different temperature levels, and record video for continuous thermal monitoring. Additionally, it allowed the capture of individual thermal images, as well as the configuration of a delayed timer function for automatic image capture at set intervals. This setup ensured efficient data collection and analysis, making it easy to monitor and document thermal variations in real-time.

The end effector for the UR10e was specifically designed to securely attach the thermal camera to the robot while ensuring sufficient structural support. The design process began with obtaining a 3D model of the thermal camera from the AMETEK website, which was used as the basis for creating the customised end effector.

A custom-designed and 3D-printed adapter was created to securely attach the AMETEK Land ARC thermal imaging camera to the UR10e cobot. The adapter was designed using Fusion 360, a computer aided design (CAD) software developed by

Autodesk, which allowed for precise 3D modelling, simulation, and design optimisation. To ensure a perfect fit, an exact 3D model of the camera was used, and the adapter was designed around it within Fusion 360, incorporating a square frame with holes on three sides to secure the camera using two screws per side. The top of the adapter featured a round connector designed to interface with the cobot's end-effector, including four screw holes for a secure attachment. The final custom designed 3D model of the adapter with camera is shown in Figure 4.3.

Figure 4.3 – 3D Model of Thermal Camera and Custom Adapter

The adapter was 3D printed using a Prusa MK4S, a high-precision, user-friendly 3D printer known for its reliability and advanced features like automatic bed levelling and high-speed printing. PLA (Polylactic Acid) filament was chosen for the print instead of materials like ABS or PETG due to its ease of printing, low warping, and sufficient strength for the application while avoiding issues like shrinkage. A key design consideration was that the adapter needed to be attached to the cobot before securing the camera, as the cobot's connector screw holes would otherwise be inaccessible once the thermal camera was in place. This custom adapter ensured a secure and stable mount, allowing for precise positioning of the thermal camera during AFP defect detection. The adapter and thermal camera can be seen mounted on the UR10e in Figure 4.4.

Figure 4.4 – Thermal Camera Mounted on Cobot Using Custom Adapter

The cobot's ability to move dynamically allowed it to follow the recently laid composite material and capture thermal images from multiple angles, significantly enhancing defect detection accuracy.

4.1.3 UR10e Cobot Teach Pendant

A teach pendant on a cobot is a handheld control device that allows users to program, control, and monitor the robot in an intuitive way. The teach pendant, used in this research, is shown in Figure 4.5. It served as the primary interface to manually guide the cobot through movements, set waypoints, and define tasks without needing complex coding. The teach pendant featured a touchscreen interface with simple drag-and-drop programming, making it extremely user-friendly. It was essential for initial setup, calibration, and troubleshooting, allowing safe movement of the cobot in manual mode before running automated programs. The teach pendant was required to provide direct, real-time control, making it easy to adjust movements, test new paths, and ensure accuracy before executing fully automated sequences.

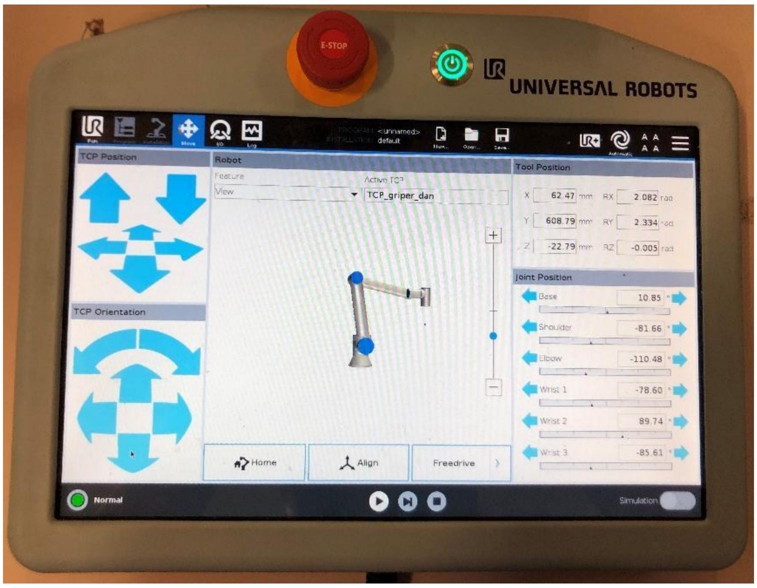

Figure 4.5 – UR Teach Pendent

4.2 Work Cell

A work cell was identified for use at the LMC. The cell was set up to enable the manufacture of composite material and consisted of a KUKA KR120 robot and a worktable. As this set up only had the capability to carry out AFP it had to be enhanced and reconfigured to allow defect detection. The work cell table featured a slotted T-design worktop, allowing various components to be securely mounted to the tabletop. This design enabled the cobot to be firmly attached, ensuring stability and consistency in positioning when repeating experiments. The cobot was placed in a specific region of the table, strategically positioned to avoid interference with the main workspace of the KUKA robot, while still maintaining access to the required area for defect detection. Additionally, a thermal camera was integrated into the system, providing thermal images essential for detecting defects in the AFP process.

The UR10e cobot, equipped with a thermal camera, detected temperature variations on the surface of the composite material. Such variations often indicate defects like delamination, voids, or gaps/overlaps. The thermal camera operated by capturing infrared radiation emitted by the surface, creating a visual representation of the temperature distribution across the material. The UR10e end effector maintained a distance of approximately 20 cm from the KUKA KR120 AFP end effector and moved in a coordinated pattern to capture high-quality images. By dynamically adjusting its position, the cobot ensured that areas of interest were imaged with higher resolution when necessary. The 20 cm gap between the end effectors of the KUKA KR120 and the UR10e cobot was essential for ensuring smooth operation and effective defect detection. This distance prevented physical interference between the robots, allowing the AFP process to continue without obstruction while the cobot captured thermal images. The gap provided enough room for adjustments if necessary, preventing unexpected collisions while still keeping the cobot close enough to detect defects immediately after layup. This balance between proximity and clearance optimised workflow efficiency and ensured a seamless integration between tow placement and real-time quality control.

The 2-robot system allows a more flexible solution. The AFP robot solely focusses on laying the tow while the cobot is entirely responsible for thermal data collection. Having a second robot that is dedicated to the camera enables the optimisation of photo angles. This in turn provides a method to collect clearer defect images. It takes 5 seconds for the cobot to move into position, take the picture, and process it.

The physical arrangement of the robot and cobot was designed to maximise efficiency while preventing interference. The KUKA KR120 was mounted on one side of the work cell, with its arm extending over the workspace to carry out the AFP process. On the opposite side, the UR10e cobot was securely positioned with its end effector facing downward, equipped with a thermal camera for defect detection. The KUKA robot had a wide range of movement, allowing it to place composite tows across the worktable, while the UR10e operated with a more precise motion, scanning the surface in a structured pattern. Their movements were carefully coordinated to maintain a 20 cm gap, ensuring the cobot could inspect the tows immediately after layup without disrupting the AFP process. The overlap in their operational areas was limited to the fibre placement region, where defect detection occurred, enabling seamless integration of tow placement and quality inspection while minimising the risk of mechanical interference.

The seamless coordination between the KUKA KR120 and UR10e was critical to the system's effectiveness. While the KUKA robot focused on laying down the tows, the UR10e monitored the process in real-time, identifying any defects as they occurred. The cobot's position was continuously adjusted based on the KUKA robot's movements, maintaining optimal alignment for thermal imaging. This dynamic feedback loop ensured that the AFP process remained uninterrupted and allowed for immediate corrective actions to address any detected defects.

To maintain the reliability of the experimental setup, several control measures were implemented. These included the temperature range setting of the thermal camera to ensure accurate temperature measurements, synchronisation of the robots' movements, and regulation of environmental conditions such as ambient temperature and lighting. The work area was in a limited access environment with a constant 21ºC temperature and air filtration. The composite material tows used were also standardised to ensure consistency. Figure 4.6 and Figure 4.7 depict the setup and positioning of the KUKA and UR10e robots, respectively.

Figure 4.6 image shows the work cell containing both the KUKA robot and the cobot, with the KUKA being the orange industrial robot and the cobot identifiable by its silver and blue design. The KUKA robot was equipped with an AFP end effector head and was actively performing AFP, laying tows onto a heated tool bed that was securely mounted on the work cell table. Meanwhile, the cobot was positioned off to the side, ensuring it did not interfere with the AFP process while the KUKA completed the tow layup.

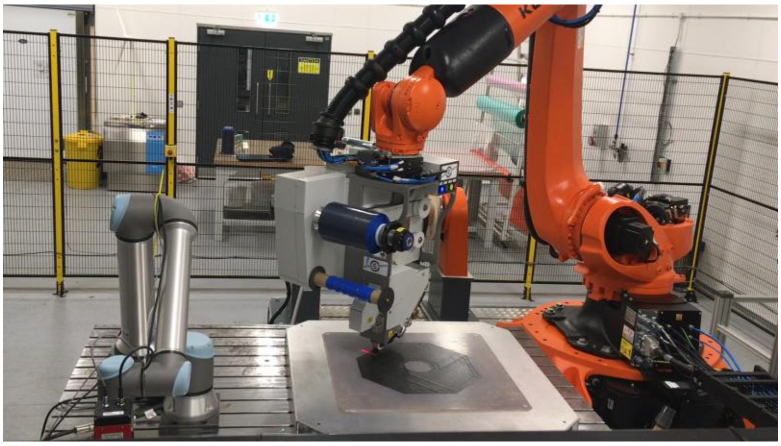

Figure 4.6 – KUKA AFP Tow Laying

In Figure 4.7, both the robot and cobot are visible again, but here the KUKA has moved off to the side after completing a full AFP layer. With the layup process finished, the cobot moved into the workspace to begin collecting thermal images for defect detection. It followed a raster scan path across the heated tool bed, at height of 60mm, systematically capturing images to analyse potential defects in the placed tows.

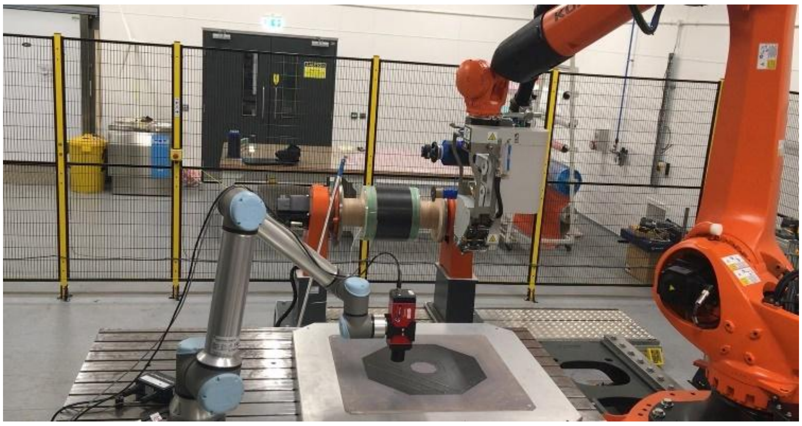

Figure 4.7 – Cobot Thermal Image Raster Scan

This setup ensured efficient collaboration between the KUKA KR120 and UR10e, with minimal interference between tasks. Their movements are carefully managed to maintain workflow continuity, allowing precise defect detection without disrupting the AFP process.

4.3 AFP Creation

A custom AFP end effector, heated tooling bed, and path design were implemented to facilitate data collection. The AFP layup followed a predefined octagonal pattern with a central hole, using Hexcel M91/34%/UD194/T700GC continuous fibre prepreg tows as the material.

The design included a four-ply laminate with a [-45/0/45/90] degree structure, which took approximately 10 minutes per layer to complete. The chosen pattern was a test pattern designed specifically for evaluating the Automated Fibre Placement system. The test design followed an octagonal shape, allowing for a controlled and repeatable layup process. To assess the system's effectiveness, four different layup orientations were selected, with layers applied at -45 degrees, 0 degrees, 45 degrees, and 90 degrees. This multi-layered approach simulated real-world composite manufacturing, particularly in regard to the way heat dissipates through the tow material during processing. Only flat designs were evaluated for this study as this was the configuration supported by the equipment at the LMC. Prior to this research the LMC has only used the APF machine for specific test designs.

The ability of thermal imaging to detect defects is highly dependent on the heat transfer available from the tool bed through the composite layers. As the number of tows increases, the thermal measurement decreases. When the temperature difference between the previous layer and current layer is no longer great enough, thermal imaging is unable to detect defects accurately.

The tow material used was Hexcel M91/34%/UD194/T700GC, a high-performance carbon fibre composite, with each tow strip measuring 6.35mm wide. This specific material and tow size are commonly used in aerospace and high-strength composite applications, making it ideal for testing thermal effects, adhesion, and layer consistency in the AFP process. Figure 4.8 illustrates the final AFP laminate design.

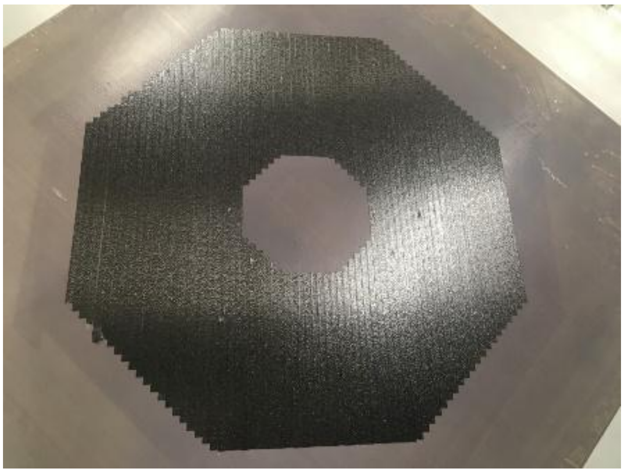

Figure 4.8 – AFP Test Pattern

The tool bed contained four heating elements, one in each quadrant, allowing for controlled heating up to 100°C. Since each carbon fibre tow reel had slight variations, testing showed that 80°C was the optimal bed temperature to ensure that the tows remained securely in place during the layup process. The heated tooling bed was initially set to 80°C. The first two plies were laid with minimal issues, allowing for the collection of several thermal videos. However, for the subsequent layers, the composite material's adhesion deteriorated due to the material's extended exposure to ambient conditions, which reduced its tackiness. Carbon fibre tow tackiness is the adhesive property of the pre-impregnated resin, which increases when heated, helping the tows to stay in place during the AFP process. This tackiness is essential for preventing movement, ensuring strong interlayer bonding, and improving automation efficiency, ultimately leading to precise fibre placement and defect-free composite structures.

A second tow reel was required after the first became out of date and lost its tackiness, making it unsuitable for the AFP process. A second spool of composite material was obtained, and the tooling bed temperature was adjusted to 55°C to prevent the material from adhering to the compaction roller. Since the new reel had different material properties, additional testing was necessary to determine the optimal heated bed temperature. After evaluation, the heated bed was set to 55°C, ensuring proper tow adhesion and placement during layup.

This adjustment proved effective, and further experimentation was conducted without adhesion issues. Gap and overlap defects were the primary focus of this study, with examples shown in Figure 4.9.

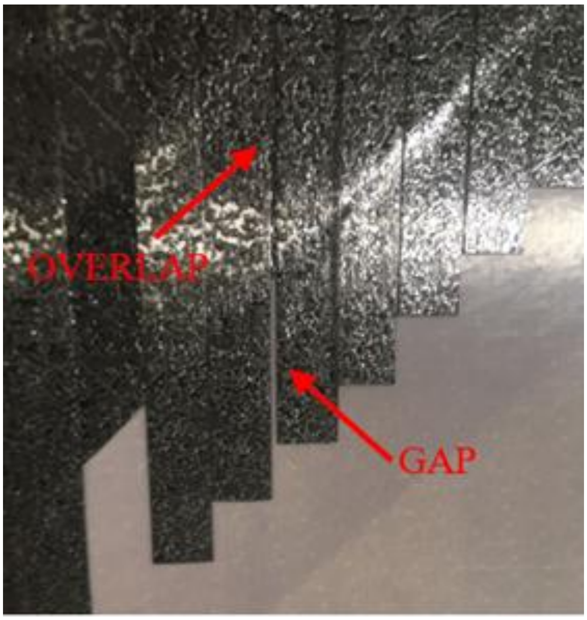

Figure 4.9 – Layup with Defects

This image shows AFP tows laid onto a surface, highlighting two common defects: overlap and gap. The overlap occurred when adjacent tows were placed too closely together, causing one tow to partially cover another. This defect can lead to thickness variations, uneven curing, and stress concentrations that weaken the composite structure. Conversely, the gap formed when tows were placed too far apart, leaving exposed areas between them. This defect can result in resin-rich zones, reduced structural integrity, and increased risk of delamination or void formation. These defects could have arisen due to incorrect path programming, material handling inconsistencies, or variations in tow tackiness, all of which need to be minimised to ensure high-quality, consistent composite manufacturing. The defects shown in the image, such as overlaps and gaps, naturally occur during the AFP process due to slight variations in tow placement, material properties, or machine calibration, making precise control essential to minimise errors and maintain composite quality.

To generate defect images for the dataset, artificial defects were intentionally introduced during the AFP process. The octagonal layup program was modified to lay four tows as normal, while the fifth tow was placed with a half-tow offset. This pattern was repeated, with the next four tows laid correctly, followed by another offset fifth tow. This systematic alteration created controlled gaps and overlaps, allowing for the collection of thermal images that accurately represented common AFP defects for dataset development.

4.4 Robot Control Software

The management of each robot, and coordination of the two robots together, was managed by dedicated software specifically designed for robotic control.

4.4.1 RoboDK Development Software

RoboDK is a robot simulation and programming tool made by RoboDK Inc. a Canadian company focused on industrial robotics software. It is used to simulate, program, and control a wide range of robots and cobots from brands like KUKA, ABB, Fanuc, and Universal Robots.

One big advantage of RoboDK, over other manufacturer-specific software, was that it was simple and had a user-friendly interface that enabled robot programming without needing to write complicated code. It also offered great simulation tools, so movements could be digitally tested before running them on a real robot, helping to avoid mistakes and save time. Other programs, like ABB, RobotStudio or MATLAB’s Robotics Toolbox, were more specialised and could be harder to learn, whereas RoboDK worked with many different robots and was much more intuitive.

At the University of Strathclyde lab, RoboDK was provided, making it an obvious choice for research and testing. The GUI made it simple to connect with real cobots and robots, allowing for direct programming and testing. Plus, its simulations were highly accurate making it an excellent tool for running digital tests before committing to real-world execution. Overall, RoboDK was a versatile, easy to use, and reliable option for both learning and real automation work.

The integration of the AFP robot and cobot was initially validated using RoboDK simulation software. RoboDK provided a virtual environment for offline programming, simulation, and optimisation of robotic systems. In this development research, RoboDK was used to simulate the movements of the KUKA KR120 and UR10e, ensuring their coordination before transitioning to a physical implementation.

To evaluate the robots and cobots, four different tests were conducted to assess their movement, controllability, and integration with both simulation and real-world systems. The first test focused solely on the cobot, where its movement and controllability were tested in simulation, followed by verifying its connectivity with the real cobot and confirming that it could be controlled through the software. The second test involved simulating the KUKA robot in a work cell, ensuring its motion and workspace compatibility before physical deployment.

The third test simulated one cobot and one robot working within the same environment, assessing their ability to operate together without collisions and verifying coordinated movements. Finally, the fourth test involved simulating two cobots in a single work cell, where one cobot was programmed to follow the position of the other. This was followed by connecting to both real cobots and running a live test to validate synchronisation and control. After successfully simulating the two-cobot system, further testing was conducted on the SSB to ensure it could control both cobots, demonstrating its ability to manage real-time communication and coordination between them.

In the simulation, a 3D model of the worktable was created to represent the AFP manufacturing process. The KUKA KR120's path planning algorithms were tested within this environment, verifying that the robot could execute the composite layup accurately and efficiently. Similarly, the UR10e's movements were simulated to ensure that the thermal camera captured images at the appropriate intervals and angles. The thermal camera captured images directly downward, ensuring consistency and a standardised angle for accurate thermal information. Images were taken at five-second intervals, allowing the cobot to move into position without disrupting the process. This timing also prevented motion blur, ensuring clear and precise defect detection.

By simulating the cobot's flexibility, the system's ability to view potential defects from various perspectives was confirmed. The cobot was able to move into various angles and orient the connected thermal camera to optimal positions for capturing thermal images of the tows. With its six degrees of freedom system, the cobot was capable of adjusting its angle to achieve better views of defects, improving analysis and defect detection.

Once the individual robots were tested, their integration was simulated with the cobot maintaining a constant offset from the KUKA AFP robot. This simulation allowed fine-tuning of the control algorithms and ensured that the system functioned seamlessly.

When testing RoboDK, the first step was to launch the software and observe the graphical user interface (GUI) as shown in

Figure 4.10. The main window consisted of a menu bar at the top, a toolbar with commonly used icons, a station tree on the left listing robots, tools, and programs, and a central 3D simulation workspace where the robotic arm and objects were visualised. The status bar at the bottom provided system messages and logs. To verify functionality, sample projects could be loaded, a robot could be added from the library, and a simple simulation was able to be run. Checking if the robot moved as expected and interacted with objects in the workspace confirmed that the software was working correctly. Additionally, testing connectivity with external hardware and exporting generated code to a real robot further ensured proper functionality.

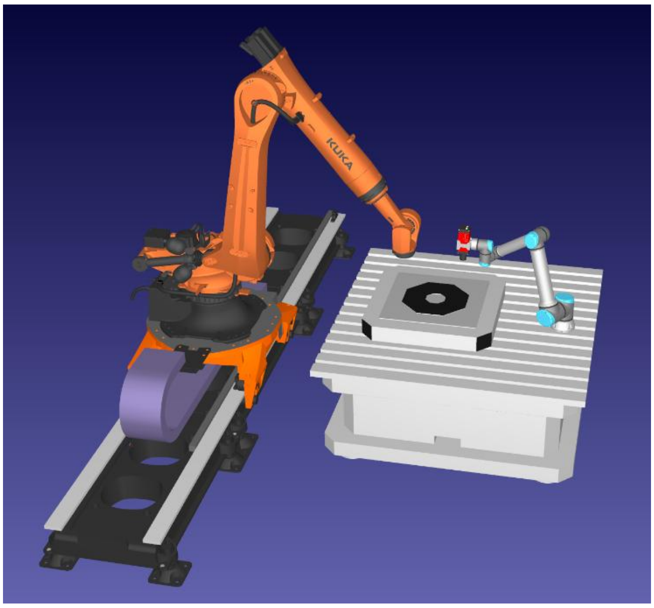

Figure 4.10 – RoboDK

4.4.2 Single Cobot Model

The UR10e cobot was controllable using RoboDK[110] and a Python[111] script, allowing for precise programming and execution of robotic tasks. To achieve this, several steps were carried out. First, the UR10e was connected to RoboDK via its Ethernet connection, enabling seamless communication and program transfer. Next, a robot program was created within RoboDK's user-friendly interface, specifying the end-effector path and other relevant parameters. Once the program was set up, a Python script was generated using RoboDK's Python API, which contained all the necessary commands to execute the programmed movements on the cobot. The script was then transferred to the UR10e, allowing it to execute the planned sequence of operations. Finally, RoboDK's monitoring and control tools were used to oversee the execution, ensuring that the cobot performed as expected and allowing for real-time adjustments if needed.

4.4.3 Multiple Cobot Models

To fully research the simultaneous use and control of multiple robots, analysis was carried out in simulation before moving to investigations with physical robot systems.

To enable simulation with multiple robots, the RoboDK simulation environment, previously set up for single robot use, was expanded to include a second UR10e collaborative robot (cobot). A second UR10e model was imported from the RoboDK library and positioned alongside the initial robot.

The workspace of the second cobot was configured to prevent any interference with the first robot's movements. RoboDK's simulation tools were employed to synchronise the operations of both cobots. The cobots were programmed to execute raster scan movements on distinct sections of a larger predefined area.

A raster scan is a systematic method of capturing thermal images by moving a thermal camera or sensor in a line-by-line pattern across a surface. This structured scanning approach ensured full coverage and consistency, preventing defects from being missed due to gaps in data collection. In thermal defect detection, raster scanning is particularly useful because it ensures that temperature variations, indicating voids, delaminations, or material inconsistencies, are accurately mapped and analysed. This method provided uniform temperature readings, making defect detection more reliable by avoiding inconsistent measurements. Additionally, using a cobot to perform a raster scan automated the process, making it repeatable and efficient for detecting defects in composite materials during AFP. The number of thermal images required for each row of the raster scan varies with the position on the octagonal shape. At the narrowest part of the octagon 4 images were required while at the widest part 7 were required. Each octagonal shape required 10 or 11 rows of scans. This approach enhances the accuracy, precision, and reliability of thermal defect detection in manufacturing. The simulation was rigorously tested to ensure precise coordination and seamless operation between the two robots.

To move from a simulation environment to a practical operation, two physical cobots were connected to the SSB through an Ethernet switch. Unique IP addresses were assigned to each cobot to prevent network conflicts. The Python program controlling the cobots was updated to communicate with both robots independently using their respective IP addresses. Multithreading or distinct control loops were implemented to manage simultaneous communication with both robots. The connection was validated by sending independent and synchronised commands to the cobots to ensure proper functionality.

The Python program for raster scan movement was further modified to integrate commands for both cobots. The scan area was divided into two distinct regions, with each cobot assigned to a specific section. The program was designed to operate both cobots concurrently while ensuring collision-free movements. The program was executed, and the cobots' performance in following the raster scan path in unison was monitored closely. Parameters related to synchronisation and movement were fine-tuned to optimise overall performance. This configuration effectively demonstrated the capability to coordinate multiple cobots for complex operations, utilising the SSB as a centralised control system.

This Python program utilised the RoboDK API to simulate two UR10e collaborative robots, Cobot1 and Cobot2, that were positioned 2.5 meters apart and facing each other. The primary objective of the program was to have Cobot2 execute a raster scan motion in a 1-meter by 1-meter square in front of it while Cobot1 followed the same path but with an offset of 20 cm closer to its base. The script started by initialising the RoboDK environment and loading two UR10e cobots from the RoboDK library. Cobot1 was positioned at the origin, while Cobot2 was placed 2.5 meters along the x-axis and rotated 180 degrees so that it faced Cobot1. The program then defined a raster scan pattern for Cobot2, covering a 1-meter by 1-meter square area with a series of parallel horizontal lines spaced evenly apart. These waypoints were generated to create a structured path, ensuring that the end effector of Cobot2 systematically moved back and forth across the scan area. Once the scan path was defined, Cobot2 began executing the raster scan by moving linearly between waypoints. At every half-second interval, the program queried the current position of Cobot2 and applied a transformation to create an offset of 20 cm towards Cobot1's base. This transformed position was then used to move Cobot1, ensuring that it followed the trajectory of Cobot2 but with the specified offset. This approach maintained a synchronised motion where Cobot1 effectively shadowed Cobot2 with a controlled spatial displacement. The use of linear motion commands, MoveL, ensured smooth transitions between waypoints, minimising abrupt movements. The program efficiently utilised RoboDK's built-in functions for setting base frames, defining movement paths, and executing coordinated robotic motion. By leveraging linear interpolation for smooth transitions between waypoints, the simulation ensured accurate path tracking. This implementation is particularly useful in applications requiring synchronised robotic operations, such as industrial automation, material handling, and coordinated inspection tasks. The program can serve as a foundation for further development, including dynamic speed adjustments, real-time obstacle detection, and sensor-based feedback to enhance robotic collaboration.

4.4.4 Follower Robot Matrix Calculation

The UR10e robot was required to follow the KUKA robot's end-effector at a fixed horizontal distance of 200 mm (20 cm). This was achieved using homogeneous transformation matrices, which defined the position and orientation of the robot's end-effector within the RoboDK station's global coordinate system as shown in Eq. 4.1. Each robot position was represented as a 3×3 rotation matrix (R) in the top-left and a 3×1 position vector (*t*) in the top-right.

$$T = \begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix} \tag{4.1}$$

This expanded to a 4×4 matrix that combines rotation and translation components. The layout of the homogeneous transformation matrix is shown in Eq. 4.2.

$$T = \begin{bmatrix} r11 & r12 & r13 & x \\ r21 & r22 & r23 & y \\ r31 & r32 & r33 & z \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.2}$$

The KUKA robot's base frame was defined in RoboDK station with a position of 250 mm in the X direction, -2000 mm in the Y direction, and 750 mm in the Z direction, and a rotation of 90 degrees around the Z-axis. These values were determined by the placement of the KUKA robot in the RoboDK environment. The robots were placed to make their relative position the same as in the real work cell.

This results in the following transformation matrix for the KUKA base with respect to the station's global reference frame are shown in Eq. 4.3.

$$T_{KUKA\_Base} = \begin{bmatrix} 0 & -1 & 0 & 250 \\ 1 & 0 & 0 & -2000 \\ 0 & 0 & 1 & 750 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.3}$$

The tool frame attached to the KUKA end-effector was located 1550 mm on the X-axis and 530 mm on the Z-axis. It was rotated 180 degrees around the X-axis, swapping the Y and Z directions. The transformation from the KUKA base to its tool frame is shown in Eq. 4.4.

$$T_{KUKA\_Tool} = \begin{bmatrix} 1 & 0 & 0 & 1550 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 530 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.4}$$

To get the global position of the KUKA's end effector, the tool matrix and the base matrix were multiplied together as shown in Eq. 4.5.

$$T_{KUKA\_EE} = T_{KUKA\_Base} * T_{KUKA\_Tool} \tag{4.5}$$

To calculate the position of the KUKA robot's end-effector in the RoboDK station's global frame, the base frame T_{KUKA_Base} and the tool frame of the end-effector T_{KUKA_Tool} transformation matrices were multiplied together as shown in Eq. 4.6.

$$T_{KUKA\_EE} = \begin{bmatrix} 0 & -1 & 0 & 250 \\ 1 & 0 & 0 & -2000 \\ 0 & 0 & 1 & 7500 \\ 0 & 0 & 0 & 1 \end{bmatrix} * \begin{bmatrix} 1 & 0 & 0 & 1550 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 530 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.6}$$

The matrix multiplication can be done by row performing a dot product on the rows of T_{KUKA_Base} and the columns of T_{KUKA_Tool}. The first row is shown in Eq. 4.7.

$$\begin{aligned}
T_{KUKA_EE}[0][0] &= 0 * 1 + (-1) * 0 + 0 * 0 + 250 * 0 = 0 \\
T_{KUKA_EE}[0][1] &= 0 * 0 + (-1) * (-1) + 0 * 0 + 250 * 0 = 1 \\
T_{KUKA_EE}[0][2] &= 0 * 0 + (-1) * 0 + 0 * (-1) + 250 * 0 = 0 \\
T_{KUKA_EE}[0][3] &= 0 * 1550 + (-1) * 0 + 0 * 530 + 250 * 1 = 250
\end{aligned} \tag{4.7}$$

The next 2 rows are shown in Eq. 4.8 and Eq. 4.9.

$$\begin{aligned}
T_{KUKA_EE}[1][0] &= 1 * 1 + 0 * 0 + 0 * 0 + (-2000) * 0 = 1 \\
T_{KUKA_EE}[1][1] &= 1 * 0 + 0 * (-1) + 0 * 0 + (-2000) * 0 = 0 \\
T_{KUKA_EE}[1][2] &= 1 * 0 + 0 * 0 + 0 * (-1) + (-2000) * 0 = 0 \\
T_{KUKA_EE}[1][3] &= 1 * 1550 + 0 * 0 + 0 * 530 + (-2000) * 1 = -450
\end{aligned} \tag{4.8}$$

$$\begin{aligned}
T_{KUKA_EE}[2][0] &= 0 * 1 + 0 * 0 + 1 * 0 + 750 * 0 = 0 \\
T_{KUKA_EE}[2][1] &= 0 * 0 + 0 * (-1) + 1 * 0 + 750 * 0 = 0 \\
T_{KUKA_EE}[2][2] &= 0 * 0 + 0 * 0 + 1 * (-1) + 750 * 0 = -1 \\
T_{KUKA_EE}[2][3] &= 0 * 1550 + 0 * 0 + 1 * 530 + 750 * 1 = 1280
\end{aligned} \tag{4.9}$$

Since this was a homogeneous matrix the third row remained unchanged as shown in Eq. 4.10.

$$T_{KUKA\_EE}[3] = [0 \quad 0 \quad 0 \quad 1] \tag{4.10}$$

The complete translation matrix was obtained by combining the results from all of the row calculations and is shown in Eq. 4.11.

$$T_{KUKA\_EE} = \begin{bmatrix} 0 & 1 & 0 & 250 \\ 1 & 0 & 0 & -450 \\ 0 & 0 & -1 & 1280 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.11}$$

To keep the UR10e at a distance of 200 mm behind the KUKA, an offset transformation matrix was required. The offset is a purely translational movement along the negative X direction of the tool frame. The matrix is shown in Eq. 4.12.

$$T_{Offset} = \begin{bmatrix} 1 & 0 & 0 & -200 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.12}$$

The target position for the UR10e's end effector in terms of the global frame coordinates can be calculated by applying the offset to the KUKA's tool position as shown in Eq 4.13.

$$T_{UR10e\_Target} = T_{KUKA\_EE} * T_{Offset} \tag{4.13}$$

The target matrix was calculated as shown in Eq. 4.14.

$$\begin{aligned} T_{UR10e\_Target} &= \begin{bmatrix} 0 & 1 & 0 & 250 \\ 1 & 0 & 0 & -450 \\ 0 & 0 & -1 & 1280 \\ 0 & 0 & 0 & 1 \end{bmatrix} * \begin{bmatrix} 1 & 0 & 0 & -200 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\ T_{UR10e\_Target} &= \begin{bmatrix} 0 & 1 & 0 & 250 \\ 1 & 0 & 0 & -650 \\ 0 & 0 & -1 & 1280 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{aligned} \tag{4.14}$$

As the UR10e was not meant to track vertical movement, the Z-coordinate from the calculated target position needed to be replaced with the actual required height of the UR10e end effector. The UR10e's original Z-height was 657.8 mm, therefore the third element of the translation vector in the position needed to be updated as shown in Eq. 4.15.

$$t_{UR10e\_Target} = \begin{bmatrix} 250 \\ -650 \\ 1280 \end{bmatrix} => t_{UR10e\_Target} = \begin{bmatrix} 250 \\ -650 \\ 657.8 \end{bmatrix} \tag{4.15}$$

This change ensured the UR10e robot kept a constant height regardless of the KUKA's movement in the Z-axis while still following its X/Y-axis movements.

4.4.5 Cobot and Robot Testing

After successful simulation, the next step was to connect the robots to RoboDK running on a laptop and test them in a real-world environment. The KUKA KR120 and UR10e were linked to the laptop, and their programs were uploaded to the robots.

With both robots connected, the AFP layup process began. The KUKA KR120 robot laid the fibre tows and the UR10e followed behind, capturing thermal images of the layup using the Python program developed previously to follow the AFP layup. A delay was introduced in the KUKA AFP laying program at the end of each tow line, allowing the Cobot to complete thermal image collection and move out of the KUKA's path before it began laying the next tow.

The thermal camera scanned the composite surface, identifying defects in real-time. Any defects detected by the thermal camera were logged and could be used to adjust the AFP process, improving the final quality of the part.

In the work cell the UR10e cobot did not do path planning. The cobot was instead connected to the SSB which did the path planning and sent the coordinates to the cobot.

4.5 SSB Integration

Setting up RoboDK on a Raspberry Pi involved a structured approach to ensure proper installation and functionality. The Raspberry Pi first had to be equipped with a compatible operating system, such as Raspberry Pi OS. Essential dependencies, including Python, necessary libraries, and OpenGL support, had to be installed on the system. The Raspberry Pi compatible version of RoboDK was then downloaded from the official RoboDK website [112]. Once installed, the software had to be tested to confirm its operational integrity by launching it and verifying the graphical interface.

The integration of a UR10e collaborative robot (cobot) into the RoboDK simulation began with accessing the RoboDK online library, which offered a diverse range of robotic models. The UR10e model was able to be imported into the simulation environment and positioned appropriately within a workspace to allow for unrestricted movement. Using RoboDK's calibration tools, the UR10e's dimensions and operational parameters were configured accurately to ensure that the virtual representation mirrored the physical robot. This alignment was essential to avoid discrepancies during physical implementation.

Once the UR10e model was integrated, its movements had to be thoroughly tested in the simulation environment. The RoboDK interface was used to design and execute simple motion sequences, such as translating the end effector along the X, Y, and Z axes and rotating its joints. The simulated movements were analysed to ensure they complied with the cobot's operational specifications, including speed and range limits. This testing phase was critical to validate the robot's performance in the virtual setting before transitioning to real-world operations.

To enable the cobot to perform a raster scan movement, a Python program was developed using RoboDK's Python API. The script was designed to direct the UR10e's end effector in a systematic back-and-forth (zigzag) pattern over a predefined area. The raster pattern was defined by specifying parameters such as start and end points, step sizes, and movement speed. Nested loops were used to control the motion along the X and Y axes while maintaining a fixed Z-coordinate. The program was tested within the simulation environment to ensure the motion met the specified parameters and performed as intended. Iterative adjustments were made to refine the script for optimal performance.

Connecting the SSB to the physical UR10e cobot required the establishment of a reliable network connection via an Ethernet cable. Both devices had to be configured to operate on the same network, with the cobot assigned an IP address that corresponded to the Raspberry Pi's communication settings. Necessary drivers or communication libraries, such as URScript, were installed on the Raspberry Pi to facilitate seamless interaction with the cobot.

A URScript is a text-based programming language used to control Universal Robots cobots, such as the UR10e, by executing custom commands directly on the robot's controller. It allows for precise movement control, sensor integration, and task automation without relying solely on graphical programming interfaces. Using URScript was beneficial because it enabled greater flexibility, allowing the fine-tuning of robot motion, interaction with external systems, and automation of complex processes beyond standard teach pendant programming. It was particularly useful in research and automation when integrating custom scripts, external sensors, or third-party software like RoboDK for seamless robot control.

The connection was validated by sending basic commands, such as moving the cobot's joints, to predetermined positions. Once successful communication was confirmed, the Python raster scan program was executed on the physical cobot, allowing it to perform the defined movements in a real-world scenario. This final step bridged the gap between virtual simulation and practical application, demonstrating the effective deployment of the robotic system.

4.6 Robot and Cobot Control Discussion and Summary

The integration of the KUKA120 industrial robot and the UR10e collaborative robot for AFP and real-time defect detection has proven highly effective in improving composite manufacturing. Using RoboDK for simulation enabled the system to be designed, tested, and optimised in a virtual environment before physical implementation, ensuring seamless coordination and reducing risks in real system control.

The KUKA120's precise fibre placement and the UR10e's flexibility in defect detection demonstrated the benefits of combining high-precision industrial robotics with real-time inspection capabilities. The UR10e's thermal imaging system enabled immediate identification of defects, improving process reliability and final product quality. Additionally, incorporating machine learning into the defect detection process enhanced adaptability, allowing the system to evolve and improve over time.

This research highlighted the potential of robotic systems to transform advanced manufacturing by increasing efficiency, accuracy, and defect-free production. The success of synchronised multi-robot systems, as seen with the dual robot configuration, further emphasises the scalability of such approaches. This work provides a foundation for future advancements in automated composite manufacturing and intelligent robotic integration.

Chapter 5

5 Machine Learning Development for Defect Detection

This research explored real-time defect detection in AFP composite material creation using thermal imaging data processed by various machine learning models. A binary classification approach was employed to identify defects. This is useful in a production environment where material must be processed based on a binary pass or fail.

5.1 Introduction

Four machine learning models were evaluated: k-Nearest Neighbors, Support Vector Machine, Decision Tree, Naive Bayes, and a deep learning model, ResNet. KNN was evaluated with 7 different values of k and ResNet was evaluated with 2 different sizes. The models selected were chosen for evaluation because they are common, well supported models that are capable of using supervised learning with binary outputs. In general machine learning models work better on small data sets while deep learning models typically rely on larger data sets. Although a small data set was used in this project, two ResNet deep learning models were included in the analysis to ensure the correct model was found for the application. The models were evaluated by comparing confusion matrices, accuracy, recall, F1-score, and the Receiver Operating Characteristic (ROC) curve with Area Under the Curve (AUC) scores.

5.2 Industrial Thermal Imaging

To carry out thermal image capturing, a rugged, industrial, thermal camera was procured as low-cost thermal cameras proved unsuitable due to their low resolution and low temperature range. The AMETEK LAND ARC Imager was chosen for implementation, as it is capable of capturing both video and images at 30Hz and has a pixel resolution of 384 x 288. This thermal camera has a temperature range of 0°C to

500°C, making it suitable for capturing thermal data in a wide range of scenarios, including AFP.

To operate the thermal camera, AMETEK's custom software was utilised to interface with it. The software allows for the capture of thermal data videos, as well as the setting of thermal limits. These limits can be used to establish minimum and maximum temperature values for the camera feed, which results in a visual colour temperature range scaling with the set temperature range. This was critical to ensure a constant colour pallet was kept for the processing of images for machine learning. Any temperature values above or below the limits were displayed as a black region. During experimentation, it was found that a temperature range of 40°C to 60°C worked best due to the temperature of the heated tool bed. When the tows are laid down the temperature fluctuations stabilise after 1-2 seconds. As long as the tool bed is on, this temperature stability is maintained. An example thermal image with the thermal colour palette is shown in Figure 5.1.

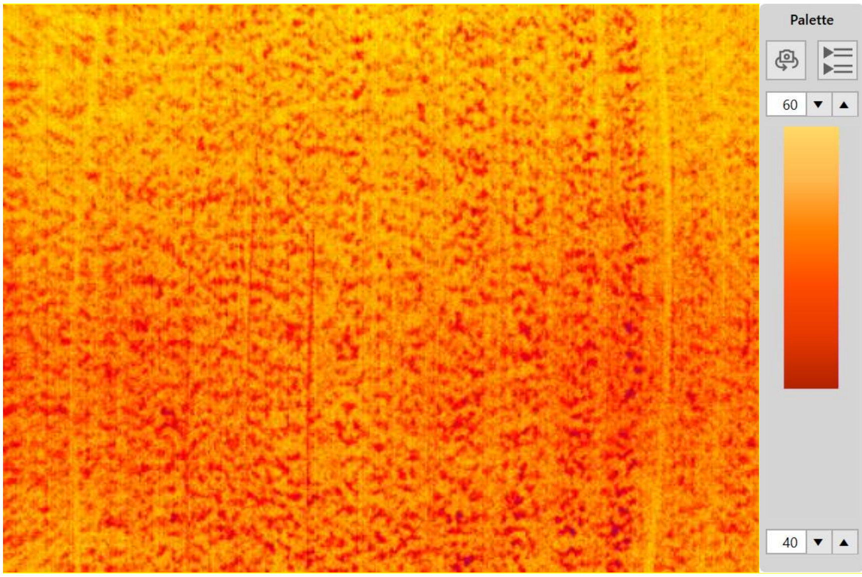

Figure 5.1 – Thermal Image with Colour Palette

5.3 Dataset Processing

This section covers the processing of the raw thermal images to create a suitable data set for the machine learning models. This was required because there were no publicly available, usable data sets.

5.3.1 Data Cleansing

The integrity and quality of the dataset is key to the success of any analytical undertaking. Upon collecting the thermal image data from the AFP layup process, a meticulous review and assessment process is required. Each image undergoes individual scrutiny to ensure its suitability for inclusion in the final dataset. This assessment involves a careful examination of various factors, including clarity, focus, and integrity. Images that exhibit blurriness, corruption, or lack clarity in depicting the main focus of interest, such as images taken of the underlying testbed not covered with a composite fibre tow, are promptly identified and flagged for exclusion from the final dataset.

In this study, the dataset underwent diligent scrutiny to ensure its integrity and suitability for subsequent analyses. Each image was subject to individual assessment, revealing various flaws that necessitated their removal from the final dataset.

One such instance involved images captured during the transit of the thermal imaging camera, rather than waiting for the Cobot to come to a complete stop before collection. Consequently, the resulting images exhibited significant blurring, obscuring the delineation between AFP tows and impeding accurate analysis due to the lack of clear delimiters. Another identified issue arose from an image exhibiting light refractions resulting from the interaction between overhead lighting, material reflectivity, and glare affecting the camera lens. These refractions introduced distortions that undermined image fidelity, rendering them unsuitable for incorporation into the dataset. Furthermore, a third image flaw primarily depicting the experimental bed upon which AFP tows were laid, failed to capture the main focus of interest in the experiment. This lack of focus detracted from the relevance of the image, leading to their exclusion from the collated dataset.

5.3.2 Image Segmentation and Naming Standardisation

The next step in dataset processing involved segmentation. Each image was broken up into a 3x3 grid, resulting in nine separate images. This segmentation was not only carried out to increase the dataset's size by a factor of nine but also to enable a more detailed analysis of the thermal fluctuations captured within each segmented image. This step increased the dataset size from 246 thermal images to 2214 thermal images, while maintaining usable image sizes for manual classification.

An example of a thermal image taken from the collected dataset can be seen in Figure 5.2 showing multiple tows. It represents an area of 90mm x 67.5mm. The thermal images were each stored utilising the naming convention ‘imageX.png’.

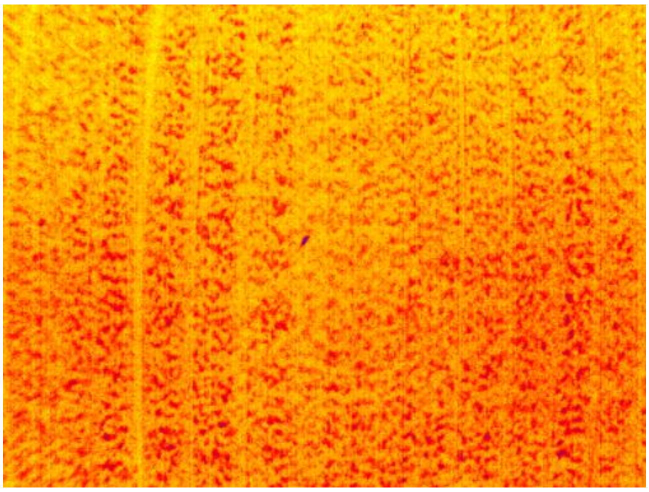

Figure 5.2 – Example Thermal Dataset Image

After undergoing segmentation and renaming via a Python script, the resulting image was split into the images displayed in Figure 5.3. The separate images were each stored with the naming conventions ‘imageX_a.png’ through ‘imageX_i.png’, thus creating the 2214 image dataset.

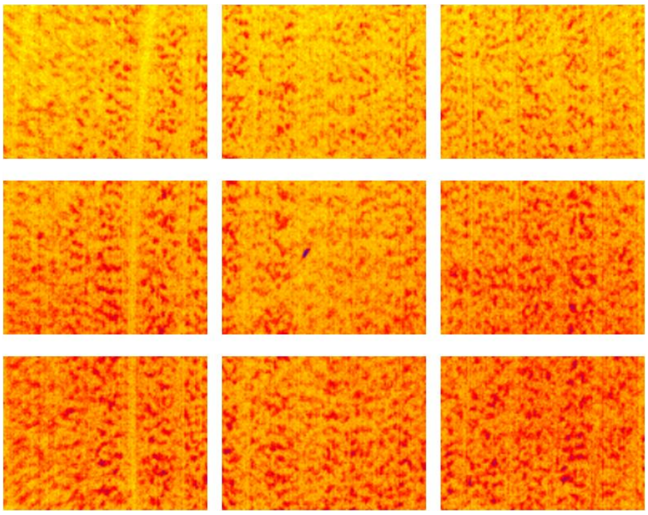

Figure 5.3 – Example Segmented Thermal Images

Following segmentation, time was taken to ensure all images employed the same naming convention. This ensured consistency throughout the dataset and minimisation of analysis errors. Standardised names contribute to the automation process of image classification tagging which allows for streamlined processing. In order to establish consistency in image naming, all images employed the convention 'pyimageX.png'.

5.3.3 UI Image Automated Tagging

A Python program was developed to streamline the inspection of all 2214 thermal images to determine the presence of defects and relocate them to the appropriate folder. The program was developed to perform the following tasks:

- Choose the full image dataset folder
- Create 'good' and 'bad' folders if absent
- Verify the correct extension for all images in the dataset folder
- Provide instructions for the user
- Define the range of images to review
- Display enlarged images
- Await user input
- Execute operations based on user input

The Python program was created to allow a user to enter four possible actions:

- 1: Image marked as 'good' and copied to relevant folder
- 2: Image marked as 'bad' and copied to relevant folder
- S: Skip to next image
- Q: Quit program

After running the program to process all thermal images in the dataset, only two folders remained: one for 'good' images and another for 'bad' images. Out of the complete dataset, 737 images were identified as containing defects, while 1477 images were labelled as defect-free. This categorisation ensured that all images were properly labelled and organised in a standardised format in preparation for deployment with machine learning models.

5.3.4 Machine Learning Data Augmentation

The next four data processing steps required were Greyscale Conversion, Resizing, Histogram Equalisation and Normalisation. These stages were implemented to standardise and enhance the quality of image data for use in machine learning models. The steps ensured uniformity, improved contrast, and optimised the images for computational efficiency and model performance.

- **Greyscale Conversion** - the first step is greyscale conversion, where each image was transformed into greyscale. This removed colour variations and ensured consistency in representation while retaining critical thermal information. By simplifying the image structure, this step enhanced the effectiveness of subsequent processing tasks. Any transparency information, known as alpha information, in the pixel was removed and the Red, Green, Blue colour data was converted using Eq. 5.1. The equation uses the international standard coefficient values for converting colour to greyscale.

$$pixel = 0.2989 * R + 0.587 * G + 0.114 * B \tag{5.1}$$

- **Resizing** - image resizing was performed to standardise all images to a predetermined size of 48 by 48 pixels from a size of 208 by 156. This ensured uniformity across the dataset, facilitating seamless integration into later analyses while also reducing computational complexity and minimising variability among images. The resizing was accomplished using bilinear interpolation as shown in Figure 5.4.

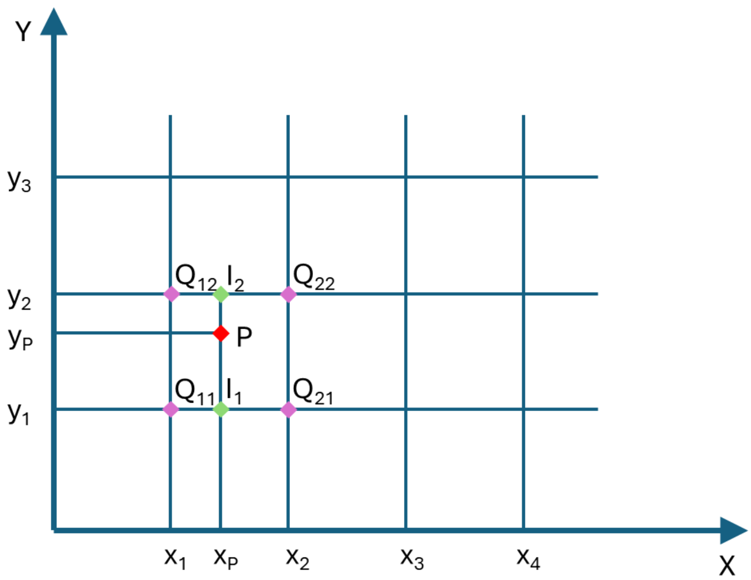

Figure 5.4 – Bilinear interpolation

The greyscale value of point P was determined using a weighted average of the 4 nearest neighbours (Q_{11}, Q_{21}, Q_{12} and Q_{22}). First greyscale values of points I_1 and I_2 awee calculated using Eq. 5.2 and Eq. 5.3 along the x-axis.

$$I_1(x, y) = Q_{11} * \frac{x_2 - x_p}{x_2 - x_2} + Q_{21} * \frac{x_p - x_1}{x_2 - x_2} \tag{5.2}$$

$$I_2(x, y) = Q_{12} * \frac{x_2 - x_p}{x_2 - x_2} + Q_{22} * \frac{x_p - x_1}{x_2 - x_2} \tag{5.3}$$

Then the I_1 and I_2 points were interpolated along the y-axis to get the final value at point P using Eq. 5.4.

$$P(x,y) = I_1 * \frac{y_2 - y_p}{y_2 - y_2} + I_2 * \frac{y_p - y_1}{y_2 - y_2} \tag{5.4}$$

The interpolation was calculated for every pixel in the new size image.

- **Histogram Equalisation** - to further refine the image quality, histogram equalisation was applied to enhance contrast and improve the distribution of pixel intensities. This technique redistributed pixel values more evenly across the intensity spectrum, resulting in better perceptual quality and increased interpretability. This step was useful for highlighting slight surface features and defects that may be difficult to otherwise distinguish.
- **Normalisation** - finally, normalisation was performed, where pixel values were scaled to a range between 0 and 1 by dividing by 255. This step ensured consistency in numerical representation, mitigating potential scaling issues and enhancing the stability of the machine learning models.

Following data augmentation, the dataset was split into training, validation and testing subsets. A validation dataset comprising 20% of the total data was separated, with 10% for testing and the remaining 70% being allocated for training purposes. This split allowed for model evaluation and optimisation, ensuring better classification of unseen data and performance in real-world applications.

5.4 Machine Learning and Deep Learning Overview

Different models of machine learning and deep learning, as well as methods of implementation, were considered for defection detection in this research.

5.4.1 Different Model Types (Machine Learning and Deep Learning)

Machine learning is a branch of artificial intelligence that enables computers to learn patterns from data and make predictions or decisions without being explicitly programmed. Traditional machine learning algorithms rely on structured data and use statistical techniques to improve performance over time. These algorithms are broadly categorised into supervised learning, unsupervised learning, and reinforcement learning. Supervised learning involves training a model on labelled data, where each input is associated with the correct output. Reinforcement learning allows models to learn through interactions with an environment.

Deep learning is a specialised subset of machine learning that utilises artificial neural networks to model complex relationships in data. These networks are inspired by the human brain and contain multiple layers of interconnected neurons, allowing them to learn complex patterns. In deep learning, the term 'deep' refers to neural networks with multiple hidden layers, enabling them to learn hierarchical and complex features. Deep learning excels at handling structured data such as images, text, and audio. Convolutional neural networks are widely used in image processing, facial recognition, and medical imaging, while recurrent neural networks and long short-term memory networks specialise in sequential data tasks such as speech recognition and natural language processing. Deep learning models typically require large datasets and high computational power, often relying on specialised hardware like Graphics Processing Units (GPUs) or Tensor Processing Units (TPUs) to achieve efficient performance.

While both machine learning and deep learning are used for predictive modelling, they differ significantly in their approach, complexity, and requirements. Machine learning models generally perform well with small to medium-sized datasets and require manual feature engineering. These models are easier to interpret, making them ideal for applications where transparency and explainability are crucial. Deep learning automatically extracts features from raw data, eliminating the need for manual entry. This makes it particularly effective for processing images and videos, where feature extraction is challenging. However, deep learning models are computationally intensive and require extensive training time, often taking days or weeks to reach optimal performance. Therefore, their complexity makes them harder to interpret, leading to challenges in understanding how a model arrives at its decisions.

Choosing between machine learning and deep learning depends on several factors, including the availability of data, computational resources, and the need for interpretability. Machine learning is a better choice when dealing with structured data and when computational efficiency is a priority. It is also preferable when the dataset is relatively small, as traditional machine learning models can still achieve high accuracy with well-designed features. Deep learning, on the other hand, is the preferred approach for applications requiring advanced pattern recognition. When large amounts of data and high-performance hardware are available, deep learning can significantly outperform traditional machine learning methods, as it can learn more intricate representations of data.

Ultimately, the decision to use machine learning or deep learning depends on the specific problem at hand. While deep learning has achieved amazing performance across various fields, it is not always the most practical solution due to its high data and computational requirements.

5.4.2 Transfer Learning

Transfer learning is a method in deep learning where a model developed for one task is repurposed for another, often related, task. Instead of training and building a model from scratch, which can be time-consuming and data-intensive, transfer learning employs the knowledge of a model that has already been developed from a pre-existing dataset. In computer vision, models trained on large datasets like ImageNet are trained to detect common visual features such as edges, shapes, and textures. These learned behaviours are useful and can serve as a strong beginning for a wide range of other image classification problems.

The main idea behind transfer learning is to keep the useful information learned by a pre-trained model and then deploy it for a new, more specific task. This is done by keeping most of the model's architecture unchanged and updating only the final layers to suit the new model and output requirements, such as changing from multi-class to binary classification.

For a deep learning project involving image classification, transfer learning allowed for faster training, improved accuracy, and better predictability, especially when working with a smaller, more specialised dataset.

5.4.3 PyTorch

In Python, the statement '*import torch*' was used to access the PyTorch library which is an open-source framework for machine learning and deep learning applications. PyTorch provided a set of tools for working with tensors, which are efficient, multi-dimensional arrays that are key to numerical computation. The library supported GPU acceleration for improved performance and was capable of handling modular components such as *torch.nn* for creating neural network architectures and *torch.optim* for implementing optimisation algorithms. Additionally, it included functionality for managing datasets and streamlining the training and evaluation processes.

5.5 Model Development

A custom dataset was prepared for use in machine learning to train models which were each then analysed.

5.5.1 Dataset Separation

The Python program was designed to split an image dataset into three separate sets: training (70%), validation (20%), and testing (10%). This process was essential for the machine learning and deep learning workflows to ensure that each model was trained on one portion of the data, validated on another for hyperparameter tuning, and tested on a completely unseen set to evaluate performance.

The Python script began by importing necessary libraries. The OS module handled file paths and directory creation, allowing the script to navigate and manipulate files within the dataset. The random module was used to shuffle images before splitting them, ensuring that the dataset was distributed randomly and not ordered in a way that could introduce bias. The *shutil* Python module was used to copy image files from the original dataset into new directories created for training, validation, and testing.

At the core of the script was the *split_dataset()* function, which takes four parameters: *input_dir, output_dir, train_ratio, val_ratio*, and *test_ratio*. The *input_dir* specified the location of the original dataset, which contained images organised into class-labelled subfolders such as "good" and "bad." The *output_dir* was where the new structured dataset was stored after the split. The function accepted optional arguments for dataset proportions, with default values set at 70% for training, 20% for validation, and 10% for testing. Before proceeding, the function verified that these three proportions summed to 1.0. If the total deviated from one, an error was raised to prevent an incorrect dataset split. Separation occurred before any data augmentation was carried out.

The data separation function first created a structured folder hierarchy in the output directory. It iterated through three dataset splits: training, validation, and testing. Within each split, it created subdirectories for the two classes, ensuring that the images remained organised according to their labels. This structured organisation made it easier to load images for training later.

Once the necessary folders were created, the function processed each class separately. It retrieved a list of all image filenames from the respective class directory while ensuring that only files were selected, preventing any errors caused by directories within the dataset. The images were then shuffled using *random.shuffle()* to ensure randomness in their distribution. After shuffling, the dataset was divided into three parts: 70% of the images were allocated to the training set, 20% to the validation set, and the remaining 10% to the test set. The number of images assigned to each set was printed for both the "good" and "bad" classes, providing a summary of the dataset distribution.

To transfer images into their designated directories, the function defined a helper function called *copy_images()*. This function took a list of image filenames and a dataset split name as arguments. It iterated through each image in the list, constructed the source and destination file paths, and used *shutil.copy2()* to copy the image while preserving its metadata, such as file timestamps. This copying process was carried out separately for the training, validation, and test sets.

The script is designed to run as a standalone program when executed directly from the command line in the operating system. Within the *if __name__ == "__main__":* block, the script sets the *input_directory* to "all", which is assumed to be the location of the original dataset, and the *output_directory* to "split", where the processed dataset will be stored. Before calling *split_dataset()*, the script sets a random seed using *random.seed(42)*. This ensures that the dataset split remains consistent across different runs, making results reproducible.

By structuring the dataset into three well-defined splits, the script ensured that the training set (70%) was used for model learning, the validation set (20%) was used to fine-tune model parameters and prevent overfitting, and the test set (10%) provided a final, unbiased evaluation of the model's performance on unseen data. This systematic approach was essential in developing reliable machine learning models capable of generalisation in predicting classifications on unseen data.

5.5.2 Model Testing Machine Learning

The Python code was designed to implement a ML classifier for binary image classification. The program followed a structured design that included loading and preprocessing grayscale images, performing data augmentation, encoding class labels, standardising feature values, training a ML model, evaluating its performance, plotting an ROC curve for binary classification, and saving the trained model and preprocessing objects for future use.

The Python script started by importing the necessary libraries. The *os* module was used for file handling, allowing the script to navigate and manipulate directories and file paths. cv2 (OpenCV) was used for image processing, particularly for loading, resizing, normalising, and augmenting images. numpy facilitated numerical operations, while *matplotlib.pyplot* was used for visualisation, such as plotting the ROC curve. *sklearn* was used to define the classifier. Other Scikit-learn modules, such as *accuracy_score, classification_report, confusion_matrix*, and *roc_curve*, were used for model evaluation. *StandardScaler* was used to standardise image features, and *LabelEncoder* converted class labels from string values to numerical representations. The *joblib* module was used to save and load trained models efficiently.

The Python script defined the dataset's base directory as "split", which was expected to contain three subdirectories: train, val, and test, each with class labelled subfolders (good and bad). The Python script included an option to apply data augmentation during training, controlled by the *data_augmentation* flag, which was set to True. The *augmentation_factor* was set to 2, meaning each original image would have two additional augmented versions, effectively increasing the size of the training set.

To load and preprocess images, the Python script iterated over the dataset splits (train, val, and test). For each split, it processed the good and bad class directories separately. It first checked whether the directory existed to avoid errors. The images were loaded in grayscale using *cv2.imread()* with the *cv2.IMREAD_GRAYSCALE* flag, ensuring that all images were processed in a consistent format. If an image failed to load, it was skipped. The images were then resized to 48×48 pixels using *cv2.resize()*, ensuring uniformity in input dimensions. Histogram equalization was applied using *cv2.equalizeHist()*, enhancing contrast by spreading out pixel intensity values. Finally, the images were normalised by dividing pixel values by 255.0, scaling them to a range of [0,1] to improve model performance.

For the training set, the Python script applied data augmentation if *data_augmentation* was enabled. Each original image was duplicated, and an augmented version was generated using horizontal flipping (*cv2.flip(image, 1)*). This increased the diversity of training samples, helping the model generalise better. Each original and augmented image was flattened into a one-dimensional feature vector and stored along with its class label.

After collecting the processed images, the Python script converted the training, validation, and test datasets into NumPy arrays. These arrays were then printed to verify the dataset's dimensions. Each dataset consisted of an X array (image features) and a y array (class labels).

To prepare the data for machine learning, the Python script encoded class labels using *LabelEncoder()*, converting the good and bad labels into numerical values (0 and 1). The encoder was applied to the training labels and then applied to the validation and test labels to ensure consistent encoding across all datasets. The image feature vectors were then standardised using *StandardScaler(),* which scaled features to have zero mean and unit variance. The scaler was applied to the training data, the validation and test sets, ensuring that all data followed the same scaling distribution.

With the data pre-processed, the model was trained using the *fit()* method, where it learned to classify images based on the closest training samples in the feature space.

After training, the Python script evaluated the model's performance on both the validation and test sets. Predictions were made using *predict()*, and accuracy was computed using *accuracy_score()*. A classification report was printed, providing precision, recall, and F1-score for each class. A confusion matrix was also displayed, showing the number of correct and incorrect predictions for each category.

An ROC curve was plotted for the validation set. The classifier’s *predict_proba()* method was used to obtain probability scores for the positive class. These scores were used to compute False Positive Rate and True Positive Rate at various threshold levels using *roc_curve()*. The AUC score was calculated, providing a single metric that quantified the model’s ability to distinguish between classes. A ROC curve plot was generated, showing how well the classifier performed across different decision thresholds.

Finally, the script saved the trained model and preprocessing objects using *joblib.dump()*. These files allowed the trained model to be reused without retraining, making future predictions efficient.

In summary, this Python script provided a complete workflow for training and evaluating a KNN classifier for binary image classification by:

- Load images
- Apply data augmentation
- Encode labels
- Standardise feature vectors
- Train ML model
- Evaluate performance using classification metrics
- Save trained model

This approach ensured an organised and efficient pipeline for building an image classification system using ML in Python and OpenCV.

5.5.3 Model Testing Deep Learning

A Python script was created to implement a deep learning based binary image classifier using *PyTorch* and *torchvision*. The script followed a structured pipeline that included dataset preparation, model selection, training with class balancing, evaluation using multiple performance metrics, and model saving for future use.

The Python script began by importing essential libraries. PyTorch (*torch*) provided tools for defining and training deep learning models, while *torch.nn* contained necessary components such as layers and loss functions. *torch.optim* was used to optimise the model during training. Additional libraries included NumPy for numerical operations, Matplotlib and Seaborn for visualising model performance, and Scikit-learn for computing classification metrics like confusion matrices, ROC curves, and classification reports. *torchvision* was used for handling image datasets and performing transformations, while *tqdm* provided progress bars to track training progress. The *collections.Counter* module was used to analyse dataset class distributions.

To ensure efficient computation, the Python script set the device to GPU mode if available, otherwise defaulting to CPU. This allowed *PyTorch* to leverage GPU acceleration, when possible, significantly improving training speed.

Several helper functions were defined to aid in model evaluation. The *plot_confusion_matrix()* function took ground truth labels and predicted labels to generate a heatmap using Seaborn, providing insights into the model's classification accuracy for each class. The *plot_roc_curve()* function computed and visualised the ROC curve, along with the AUC score, which measured the model's ability to distinguish between classes. The *compute_class_weights()* function calculated class weights based on dataset distribution to address class imbalance. This function counted occurrences of each class in the dataset using *Counter()*, computed inverse frequencies, and returned a tensor of class weights that would later be used to balance the loss function.

The Python script then prepared the dataset by loading images using *torchvision.datasets.ImageFolder*, assuming they were organised into train, val (for validate), and test directories. The images were transformed using *torchvision.transforms.Compose()* to standardise input data. Each image was converted to grayscale but expanded to three channels to match the input requirements of the pre-trained *ResNet* model. The images were resized to 48×48 pixels, ensuring uniform input dimensions across the dataset. Data augmentation was applied to the training dataset, including random horizontal flipping, to improve generalisation. Finally, all images were normalised using ImageNet statistics, which stabilised the training process. The transformed datasets were then loaded into *torch.utils.data.DataLoader*, which created batches of images for efficient training and evaluation.

Next, the Python script set up a deep learning model using ResNet, a pre-trained convolutional neural network designed for image classification. The *models.resnet(weights=True)* function loaded the model with pre-trained weights from ImageNet. The final fully connected layer was replaced with a single neuron to accommodate binary classification. The *BCEWithLogitsLoss()* function was used as the loss function, as it was well-suited for binary classification tasks, combining a sigmoid activation function with binary cross-entropy loss. The optimiser chosen was *AdamW*, which was an improved version of Adam that included weight decay to reduce overfitting.

The training loop was structured to run for up to 50 epochs, but an early stopping mechanism was implemented to prevent overfitting. If validation loss did not improve for five consecutive epochs, the training was stopped early. For each epoch, the model iterated over both training and validation datasets. During the training phase, the model learned by making predictions, computing loss, and updating its parameters using backpropagation. In the validation phase, the model was evaluated without updating its weights. Metrics such as loss and accuracy were tracked throughout the process. If the validation loss improves, the best model weights were saved.

After training, the model was tested on unseen test data. Predictions were made without gradient computation *(torch.no_grad())*, ensuring efficient evaluation. The test loss and accuracy were computed, and predictions were stored for further analysis. The script then generated a classification report displaying precision, recall, and F1-score for each class. A confusion matrix was plotted to visualise how well the model differentiated between classes, and a ROC curve was plotted to assess the model's sensitivity and specificity.

In summary, this Python script provided a complete workflow for training and evaluating a ResNet for binary image classification by:

- Load images
- Apply data augmentation
- Setup ResNet model
- Train ResNet model
- Evaluate performance using classification metrics
- Save trained model

Finally, the trained model's weights were saved using *torch.save()*, allowing it to be reused without retraining. The model was stored as "*pretrained_model.pt*", which could later be loaded for inference or fine-tuning. The script provided a complete pipeline for binary image classification, handling everything from data preprocessing to model evaluation and saving, making it a robust deep learning solution using PyTorch.

5.6 Machine Learning Data Analysis

To conduct the initial testing of the overall system, a binary classification approach was employed to differentiate between the presence and absence of defects. This approach was chosen as the ultimate outcome of the experimentation was not affected by the specific type of defect detected, but rather by the detection of any defect.

To fully analyse the dataset, the processed images underwent evaluation using four distinct machine learning models: KNN, SVM, Decision Tree, and Naïve Bayes. Each model was thoroughly evaluated based on several key performance metrics to measure its effectiveness in detecting defects in AFP composite material creation.

The analysis primarily focused on comparing the models across six specific metrics: Confusion Matrix, Accuracy, ROC curve and AUC score, Recall, F1-Score and Support.

Confusion Matrix – this metric provided a detailed breakdown of a model's performance by categorising predicted outcomes into true positive, true negative, false positive, and false negative instances. This matrix, shown in

Table 5.1, offered a comprehensive view of a model's predictive capabilities, highlighting potential areas of strength and weakness.

Table 5.1 – Confusion Matrix

Predicted\Actual	**Positive**	**Negative**
Positive	True Positive (TP)	False Positive (FP)
Negative	False Negative (FN)	True Negative (TN)

- **Accuracy** – served as a fundamental measure of a model's overall correctness in predicting defect occurrences. It quantified the proportion of correctly classified instances among the total instances evaluated and provided insight into the model's overall performance [113]. Accuracy is calculated as shown in Eq. 5.5.

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN} \tag{5.5}$$

- **ROC Curve and AUC score** – the ROC curve is a graphical representation of a model's ability to correctly classify data. AUC is the area under a ROC curve and is used to calculate the predicted performance of a model. ROC curves, an example of which is shown in Figure 5.5, were achieved by plotting the True Positive Rate (TPR) of a model against the False Positive Rate (FPR).

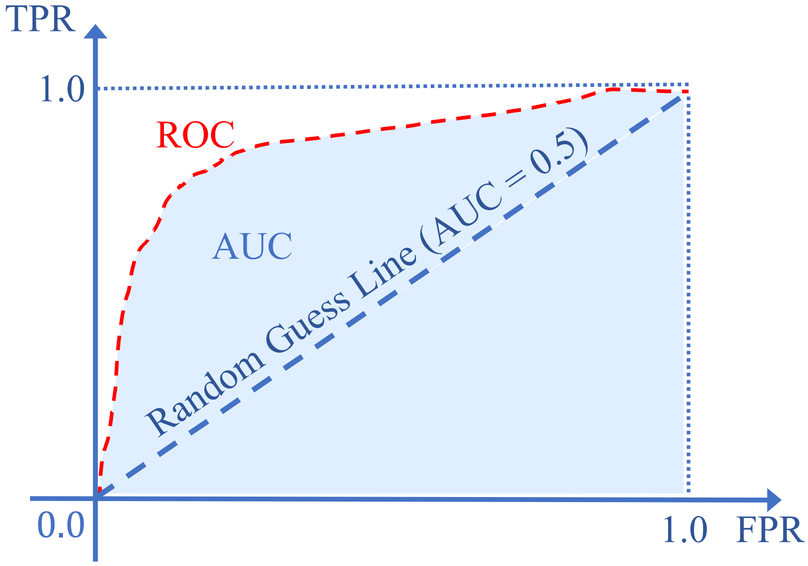

Figure 5.5 – ROC-AUC Graph

The True Positive Rate, or sensitivity of a model, refers to the proportion of true positives correctly identified while the False Positive Rate refers to the proportion of true negatives incorrectly identified as positive.

The ROC AUC score, calculated using Eq. 5.6, represents the area under the ROC curve with a value in the range $0 \leq \text{AUC} \leq 1$.

$$\text{ROC AUC} = \int \text{TPR}(\text{FPR})\text{dFPR} \qquad (5.6)$$

A ROC AUC score of 0.5 lies on the line TPR = FPR and indicates a 50% rate for correct classification which is the same success rate as random guessing. Higher AUC scores indicate superior predictive performance. The closer the ROC curve is to the line TPR = 1 the greater the model's ability to discriminate between defect and non-defect instances [113].

- **Recall** – the Recall metric, also known as sensitivity or true positive rate, measured the proportion of actual positive instances that were correctly identified by the model. It signified the model's ability to capture all relevant instances of defects in AFP composite material creation, thus minimising the risk of false negatives [113]. Recall is calculated as shown in Eq. 5.7.

$$Recall = \frac{TP}{TP + FN} \tag{5.7}$$

- **F1-Score** – the F1-score provided a balanced measure of a model's precision and recall, offering insight into its overall performance. It is calculated as the harmonic mean of precision and recall, thereby accounting for both false positives and false negatives. A higher F1-score indicated better model performance, striking a balance between precision and recall [113]. F1-Score is calculated as shown in Eq. 5.8.

$$F1\,Score = \ 2 * \left[\frac{Precision * Recall}{Precision + Recall}\right] \tag{5.8}$$

- **Support** – Support refers to the number of actual occurrences of each class in the dataset. It provided context for interpreting the performance of the model, particularly in scenarios where class imbalances may exist. Understanding the support for each class helped in assessing the reliability and generalisability of the model's predictions across different classes [113].

By comparing all of the metrics across the four machine learning models, the analysis aimed to find the strengths and weaknesses of each approach within the area of defect detection in AFP composite material creation. The metrics were compared using averages:

- **Macro Average** – this was the standard mathematical mean where the sum of the data values was divided by the number of data values

- **Weighted Average** – this was similar to the macro average however the data values were weighted, and the sum of the data values was divided by the sum of the weighting values

By employing a varied set of evaluation criteria, an insight into the comparative performance and applicability of various machine learning algorithms in tackling real-world manufacturing obstacles was able to be obtained. Following data processing, the outcomes for each machine learning model were organised and presented in the tables in the following sections.

5.6.1 k-Nearest Neighbours Model

A KNN model is a machine learning algorithm used for classification and regression tasks. In KNN, the classification of a data point is determined by the majority class among its k nearest neighbours in the feature space.

The parameter k indicates the number of nearest neighbors to evaluate when classifying data. If the k value is too high, it may lead to correlations being missed and if it is too low it can cause the model to be overly influenced by random data points that can skew the outcome. To determine which neighbors are nearest a new data point, various methods of distance calculation can be used. The most commonly used equations are the Euclidean, Manhattan and Minkowski distance equations.

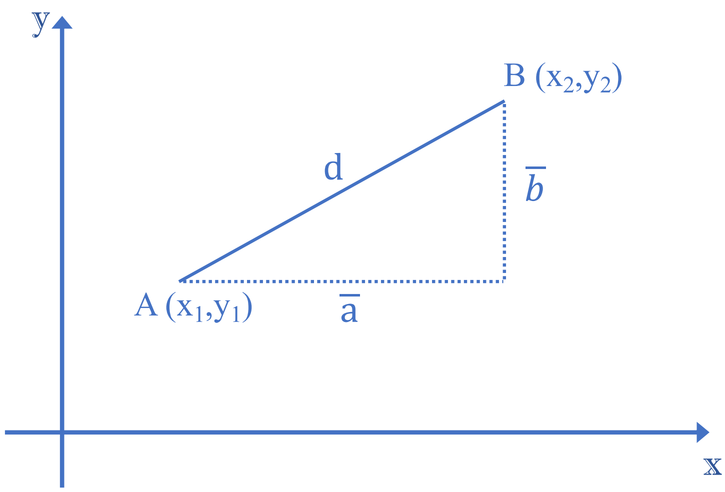

Figure 5.6 – Distance Equation

For two points in space, A and B, as shown in Figure 5.6, Euclidean distance is calculated as the direct distance, d, between the two points and can be generalised as shown in Eq. 5.9 where n represents the number of dimensions. The Manhattan distance equation, shown in Eq. 5.10, is calculated using the horizontal and vertical paths required to get from initial to final data point and the Minkowski distance, calculated using Eq. 5.11, is a combination of the Euclidean and Manhattan formulae. In this equation, if the parameter p=1 it acts as the Manhattan equation and if p=2 it provides the Euclidean equation.

$$d = |AB| = (\sum_{i=1}^{n} (x_i - y_i)^2)^{1/2} \tag{5.9}$$

$$d = |\bar{a}| + |\bar{b}| = \sum_{i=1}^{n} |x_i - y_i| \tag{5.10}$$

$$d = (\sum_{i=1}^{n} (|x_i - y_i|)^p)^{1/p} \tag{5.11}$$

The KNN algorithm has the advantage of being simple to implement, easy to understand and good for classification applications however it utilises a large amount of memory, requires the careful setting of the k parameter and is sensitive to spurious outlying data points.

For binary classification tasks in thermal imaging, KNN can be beneficial due to its simplicity and effectiveness in capturing relationships in the data. Thermal imaging data often exhibit patterns that can be crucial for detecting anomalies or classifying objects. KNN employs these characteristics by considering the similarity of nearby data points.

K values of 2, 3, 5, 6, 7, 8 and 9 were evaluated with the detailed findings of 2, 5, 7 and 9 presented. The results of 3, 6 and 8 were not presented because they were not significantly different to the other results.

5.6.1.1 kNN k=2

The results of the KNN k = 2 machine learning model can be found in Table 5.2.

Table 5.2 – KNN k=2 Model Results

Model: KNN=2	Precision	Recall	F1-Score	Support
0 (no defect)	0.44	0.23	0.3	75
1 (defect)	0.69	0.85	0.76	149
Macro Average	0.56	0.54	0.53	224
Weighted Average	0.6	0.64	0.61	224
Confusion Matrix	17 (TP)	58 (FP)		
	22 (FN)	127 (TN)		
ROC AUC	0.501867			
Accuracy	0.642857			

The KNN model with k=2 showed moderate performance in detecting defects during AFP composite material creation, achieving an overall accuracy of 64.29% on the test set.

For class 0 (non-defective), the model recorded a precision of 44% and a recall of 23%, resulting in an F1-score of 30%. This indicates that out of all instances predicted as non-defective, only 44% were correctly identified, and the model successfully detected just 23% of the actual non-defective instances. In contrast, class 1 (defective) demonstrated stronger performance, with a precision of 69% and a high recall of 85%, leading to an F1-score of 76%. This suggests the model was more effective at identifying defective parts, correctly classifying 127 out of 149 actual defective instances.

The weighted average F1-score of 61% reflects a better overall performance for the majority class (defective), while the macro average F1-score of 53% highlights the imbalance in class performance. The confusion matrix further illustrates this discrepancy: out of 75 actual non-defective instances, only 17 were correctly classified, while 58 were mistakenly labelled as defective. Meanwhile, out of 149 defective instances, 127 were correctly identified, and 22 were misclassified as non-defective. The ROC AUC score of 0.5019 suggests the model performed only slightly better than random guessing when distinguishing between defective and non-defective parts.

Overall, while the KNN with k=2 model exhibited some capability in detecting defects, particularly for the defective class, its performance on non-defective instances was limited. Based on the results obtained and shown in Table 5.2, the following ROC curve in Figure 5.7 was generated for the KNN k=2 model.

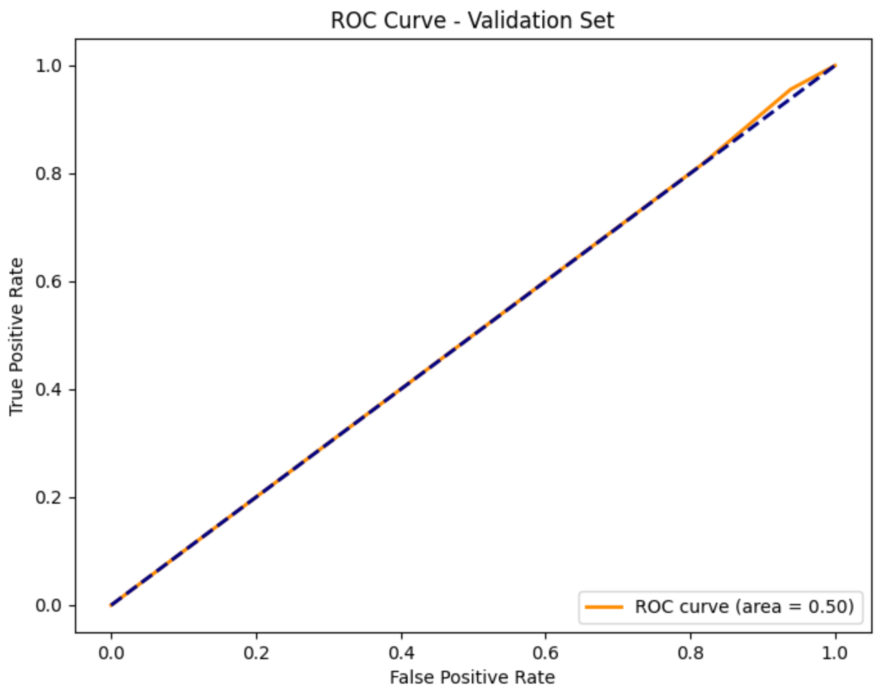

Figure 5.7 – KNN k=2 ROC Curve

5.6.1.2 kNN k=5

The results of the KNN k = 5 machine learning model can be found in Table 5.3.

Table 5.3 – KNN k=5 Model Results

Model: KNN=5	Precision	Recall	F1-Score	Support
0 (no defect)	0.59	0.13	0.22	75
1 (defect)	0.69	0.95	0.8	149
Macro Average	0.64	0.54	0.51	224
Weighted Average	0.65	0.68	0.6	224
Confusion Matrix	10 (TP)	65 (FP)		
	7 (FN)	142 (TN)		
ROC AUC	0.516199			
Accuracy	0.678571			

The KNN model with k=5 demonstrated modest performance in identifying defects, achieving an overall test accuracy of 67.86%.

For class 0 (non-defective), the model yielded a precision of 59% and a recall of only 13%, resulting in an F1-score of 22%. This indicates that while over half of the predicted non-defective instances were correct, the model struggled to identify actual non-defective samples, correctly detecting just 10. On the other hand, class 1 (defective) showed stronger results, with a precision of 69% and a high recall of 95%, producing an F1-score of 80%. This means the model correctly identified 142 out of 149 actual defective parts, showing a clear strength in detecting defective material.

The weighted average F1-score of 60% and macro average F1-score of 51% highlight the imbalance in performance between the two classes, with a strong bias towards classifying defective items correctly. The confusion matrix supports this observation: of the 75 actual non-defective instances, only 10 were correctly classified, while 65 were misclassified as defective. In contrast, of the 149 defective instances, 142 were correctly labelled, and only 7 were missed. Lastly, the ROC AUC score of 0.5162 suggests the model performs only marginally better than chance in distinguishing between defective and non-defective cases.

In summary, the KNN model with k=5 was effective at detecting defective material but performed poorly when identifying non-defective components. While its high recall for defects may be advantageous in avoiding missed defects, its low precision for non-defective parts may lead to redundant rework or inspection. Based on results in Table 5.3, the ROC curve in Figure 5.8 was generated for the KNN k=5 model.

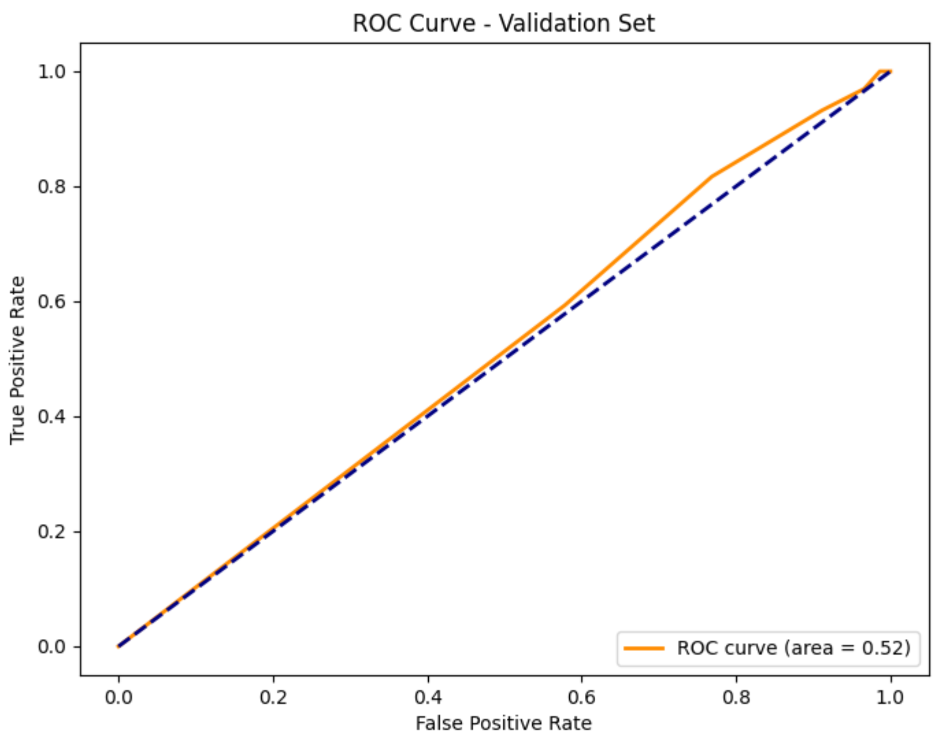

Figure 5.8 – KNN k=5 ROC Curve

5.6.1.3 kNN k=7

The results of the KNN k = 7 machine learning model can be found in Table 5.4.

Table 5.4 – KNN k=7 Model Results

Model: KNN=7	Precision	Recall	F1-Score	Support
0 (no defect)	0.69	0.155	0.24	75
1 (defect)	0.69	0.97	0.81	149
Macro Average	0.69	0.56	0.52	224
Weighted Average	0.69	0.69	0.62	224
Confusion Matrix	11 (TP)	64 (FP)		
	5 (FN)	144 (TN)		
ROC AUC	0.505004			
Accuracy	0.691964			

The KNN model (k=7) showed slightly improved performance in defect detection, achieving an overall accuracy of 69.20% on the test set.

For class 0 (non-defective), the model recorded a precision of 69% but a low recall of 15.5%, leading to an F1-score of 24%. This suggests that while most predictions labelled as non-defective were correct, the model struggled to correctly identify actual non-defective samples, correctly classifying only 11 out of 75. In contrast, class 1 (defective) continued to be well handled by the model, with a precision of 69%, a high recall of 97%, and a strong F1-score of 81%. The model accurately detected 144 out of 149 defective instances, making it highly reliable for catching defects.

The weighted average F1-score of 62% reflects decent overall performance, with a tendency to favour the defective class. The macro average F1-score of 52% indicates continued imbalance between class performances. The confusion matrix further illustrates this disparity: of the 75 actual non-defective parts, only 11 were correctly identified, while 64 were misclassified as defective. In comparison, the model correctly labelled 144 out of 149 defective parts, misclassifying just 5. However, the ROC AUC score of 0.5050 reveals that the model's ability to distinguish between defective and non-defective instances remains close to random.

In summary, while the KNN model with k=7 effectively detected defective parts with high recall and solid precision, it performed poorly in identifying non-defective items. Its skewed performance may be beneficial for catching critical defects but poses risks of over-flagging good images. The results from Table 5.4 for the KNN k=7 model produced the ROC curve shown in Figure 5.9.

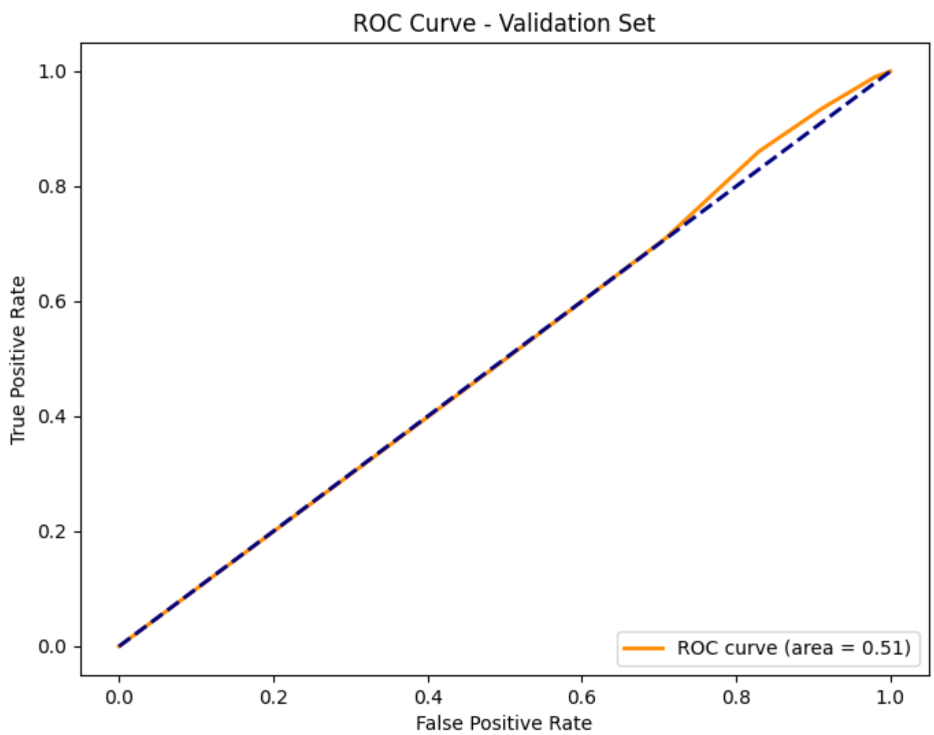

Figure 5.9 – KNN k=7 ROC Curve

5.6.1.4 kNN k=9

The results of the KNN k = 9 machine learning model can be found in Table 5.5.

Table 5.5 – KNN k=9 Model Results

Model: KNN=9	Precision	Recall	F1-Score	Support
0 (no defect)	0.78	0.09	0.17	75
1 (defect)	0.68	0.99	0.81	149
Macro Average	0.73	0.54	0.49	224
Weighted Average	0.72	0.69	0.59	224
Confusion Matrix	7 (TP)	68 (FP)		
	2 (FN)	147 (TN)		
ROC AUC	0.534832			
Accuracy	0.6875			

The KNN model with k=9 achieved a test accuracy of 68.75% in detecting defects during AFP composite material production.

Performance for class 0 (non-defective) remained limited, despite a high precision of 78%. The model's recall was just 9%, resulting in a low F1-score of 17%. This indicates that while the model is cautious in predicting non-defective instances (making few such predictions), it failed to identify most actual non-defective cases, correctly classifying only 7 out of 75. Conversely, class 1 (defective) showed excellent detection, with a precision of 68%, a very high recall of 99%, and a strong F1-score of

81%. The model correctly identified 147 out of 149 defective items, misclassifying only 2.

The weighted average F1-score of 59% reflects the skewed nature of the model's predictions, which favoured detecting defects over preserving accuracy in identifying non-defective parts. The macro average F1-score of 49% further emphasises the imbalance in class-level performance. The confusion matrix highlights this disparity clearly: out of 75 actual non-defective instances, only 7 were correctly identified, while 68 were misclassified as defective. For the 149 defective instances, 147 were accurately classified, and only 2 were missed. Although the ROC AUC score of 0.5348 indicates a slight improvement over random guessing, the model's overall ability to differentiate between classes was still quite limited.

In summary, the KNN model with k=9 was highly effective at identifying defective parts but continued to struggle with correctly classifying non-defective ones. Based on the results in Table 5.5, the ROC curve in Figure 5.10 was generated for the KNN k=9 model.

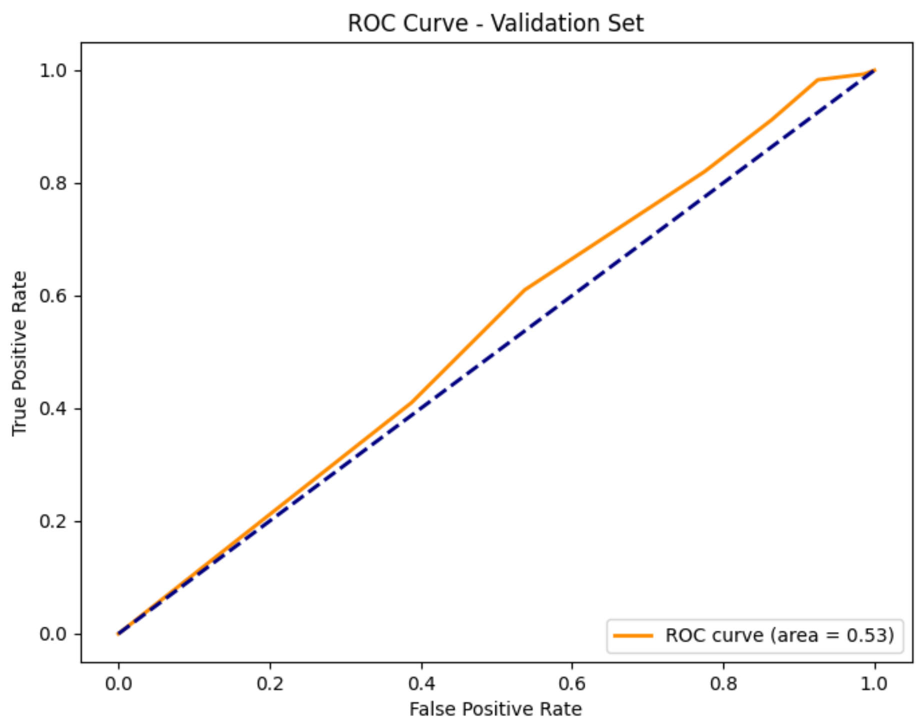

Figure 5.10 – KNN k=9 ROC Curve

5.6.2 Decision Tree Model

A decision tree model is a supervised machine learning algorithm used for both classification and regression tasks. It segments the feature space into regions, assigning labels to instances based on the majority class or mean value within each region. Decision trees are constructed, as shown in Figure 5.11, by splitting the data based on the feature that best separates the classes. Starting with the root node, which represents the complete data set, branches are created at the feature decision nodes ultimately leading to terminal leaf nodes. Each branch creates a node with a set of data that is purer than the last. This process is repeated until a predetermined end point is reached.

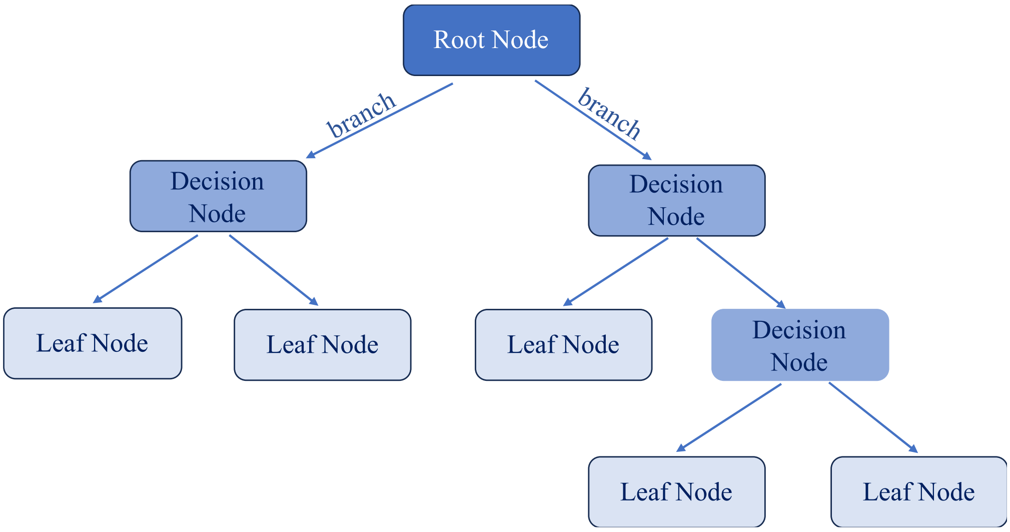

Figure 5.11 – Decision Tree

Decision trees are well-suited for binary classification in thermal imaging due to their interpretability, ability to handle non-linear relationships, robustness to noise, and processing efficiency. However, they may suffer from overfitting when the tree becomes too complex. The results of the decision tree machine learning model can be found in Table 5.6.

Table 5.6 – Decision Tree Model Results

Model: Decision Tree	Precision	Recall	F1-Score	Support
0 (no defect)	0.36	0.33	0.35	75
1 (defect)	0.68	0.7	0.69	149
Macro Average	0.52	0.52	0.52	224
Weighted Average	0.57	0.58	0.58	224
Confusion Matrix	25 (TP)	50 (FP)		
	44 (FN)	105 (TN)		
ROC AUC	0.515876			
Accuracy	0.580357			

The decision tree model demonstrated limited effectiveness in detecting defects, achieving a test set accuracy of 58.04%.

For class 0 (non-defective), the model yielded a precision of 36% and a recall of 33%, resulting in an F1-score of 35%. This reflects modest performance in identifying non-defective items, with the model correctly classifying just 25 out of 75 such instances. For class 1 (defective), the model performed notably better, with a precision of 68%, a recall of 70%, and an F1-score of 69%. It correctly identified 105 out of 149 defective parts, though 44 were still misclassified.

The macro average F1-score of 52% indicates relatively even, albeit low, performance across both classes. The weighted average F1-score of 58% reflects the dominance of the defective class in the dataset and the model's stronger performance in detecting it. The confusion matrix further reveals this imbalance: among 75 non-defective instances, 25 were correctly identified, while 50 were misclassified as defective. Out of 149 defective cases, 105 were accurately classified, with 44 missed. The ROC AUC score of 0.5159 suggests that the model had only a slight edge over random guessing in distinguishing between defective and non-defective instances.

Overall, while the decision tree model showed moderate success in identifying defective parts, its accuracy and recall for non-defective instances were weak. The ROC curve in

Figure 5.12 was generated for the Decision Tree model results shown in Table 5.6.

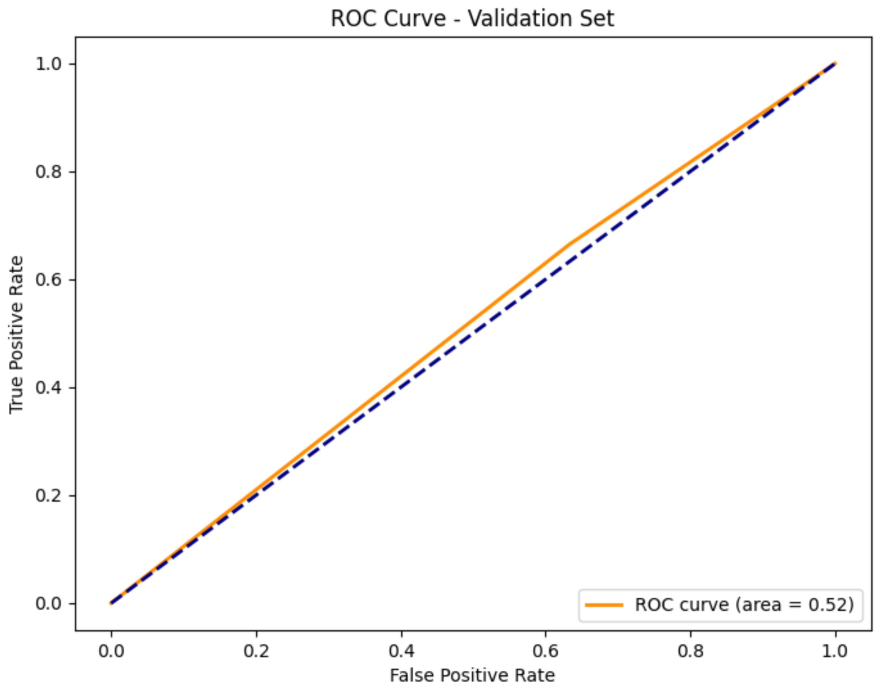

Figure 5.12 – Decision Tree ROC Curve

5.6.3 Naïve Bayes Model

A Naïve Bayes model is a probabilistic machine learning algorithm based on the Bayes' theorem with an assumption of independence between features. It calculates the probability of a class label given a set of features. The Bayes theorem is shown in Eq. 5.12.

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \tag{5.12}$$

where P(A|B) is the posterior probability of A given that B has already been observed, and P(B|A) is the likelihood probability of B given that A is true. P(A) and P(B) are the probabilities of A and B respectively. In Naïve Bayes classification algorithms, A refers to the classes that can be assigned to the available data or images and B refers to the set of features identified as belonging to the required classes.

Naïve Bayes is well-suited for binary classification due to its simplicity, speed, robustness to high-dimensional data and interpretability. However, its performance may be limited by strong assumptions of feature independence, which might not be relevant in all cases. The results of the Naïve Bayes model can be found in Table 5.7.

Table 5.7 – Naïve Bayes Model Results

Model: Naïve Bayes	**Precision**	**Recall**	**F1-Score**	**Support**
0 (no defect)	0.36	0.48	0.41	75
1 (defect)	0.68	0.56	0.62	149
Macro Average	0.52	0.52	0.51	224
Weighted Average	0.57	0.54	0.55	224
Confusion Matrix	36 (TP)	39 (FP)		
	65 (FN)	84 (TN)		
ROC AUC	0.572754			
Accuracy	0.535714			

The Naïve Bayes model demonstrated limited, but somewhat balanced, performance in detecting defects, with an overall accuracy of 53.57% on the test set.

For class 0 (non-defective), the model achieved a precision of 36% and a recall of 48%, resulting in an F1-score of 41%. This indicates the model identified just under half of the actual non-defective instances correctly (36 out of 75), though its relatively low precision suggests frequent misclassification of defective parts as non-defective. For class 1 (defective), performance was stronger, with a precision of 68%, a recall of 56%, and an F1-score of 62%. The model correctly classified 84 out of 149 defective instances but misclassified 65 as non-defective.

The macro average F1-score of 51% and weighted average F1-score of 55% highlight moderate, though still sub-optimal, performance across classes. The confusion matrix confirms this balance: among the 75 non-defective instances, 36 were correctly identified and 39 misclassified as defective; for the 149 defective cases, 84 were correctly predicted and 65 misclassified. The ROC AUC score of 0.5728 suggests the model had a modest ability to discriminate between defective and non-defective instances—better than chance, but still far from ideal.

In summary, while the Naïve Bayes model offered a more even performance across both classes compared to other models tested, its overall effectiveness remained limited. Based on the results obtained and shown in Table 5.7, the following ROC curve in Figure 5.13 was generated for the Naïve Bayes model.

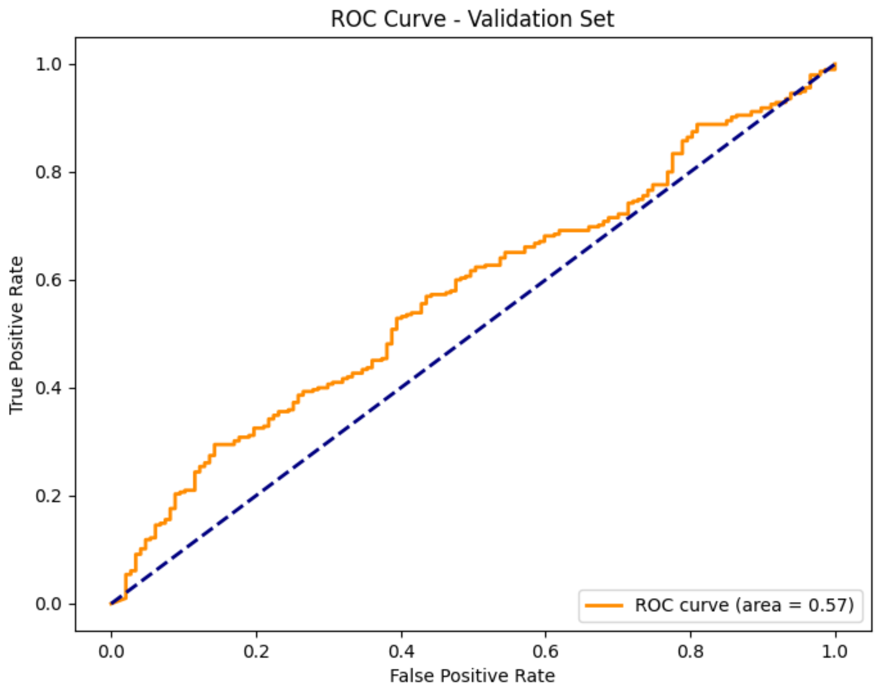

Figure 5.13 – Naïve Bayes ROC Curve

5.6.4 Support Vector Machine Model

A Support Vector Machine (SVM) model is a supervised learning algorithm used for both classification and regression tasks. It works by finding the optimal hyperplane that separates data points of different classes in the feature space, maximising the margin between the classes. Figure 5.14 shows the hyperplane decision boundary separating two classes. The support vectors are the data points nearest the central boundary. These support vectors determine the width and position of the margin.

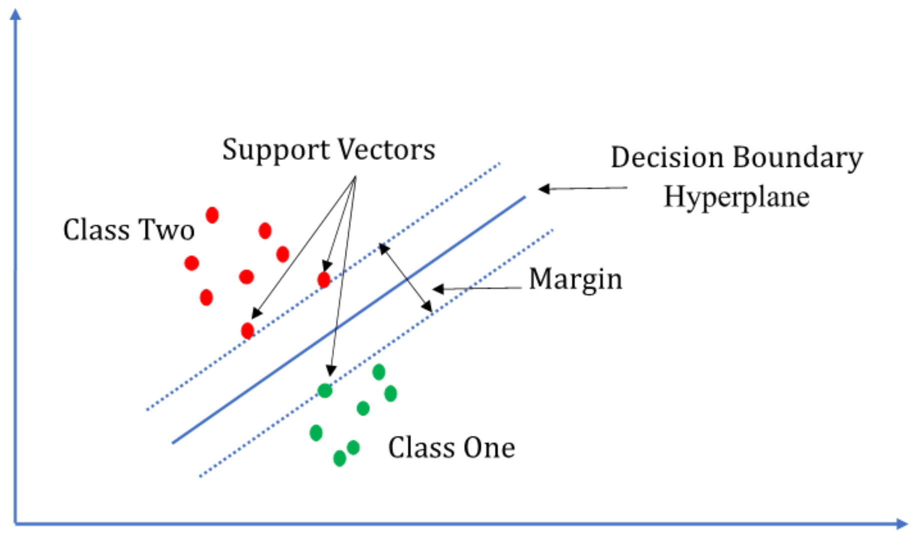

Figure 5.14 – Support Vector Machine

SVMs are well-suited for binary classification in thermal imaging due to their effectiveness in high-dimensional spaces, robustness to overfitting, ability to capture non-linear relationships, and margin maximisation properties. However, SVMs may require careful selection of hyperparameters and can have significant computation requirements for large datasets. The results of the SVM machine learning model can be found in Table 5.8.

Table 5.8 – SVM Model Results

Model: SVM	Precision	Recall	F1-Score	Support
0 (no defect)	0.69	0.15	0.24	75
1 (defect)	0.69	0.97	0.81	149
Macro Average	0.69	0.56	0.52	224
Weighted Average	0.69	0.69	0.62	224
Confusion Matrix	11 (TP)	64 (FP)		
	5 (FN)	144 (TN)		
ROC AUC	0.679257			
Accuracy	0.691964			

The SVM model achieved a test accuracy of 69.20% in detecting defects in AFP composite material production, showing strong performance for the defective class but less success with non-defective identification.

For class 0 (non-defective), the model attained a precision of 69% but only a recall of 15%, resulting in a modest F1-score of 24%. This means that while most predicted non-defective instances were correct, the model failed to detect the majority of actual non-defective cases, correctly identifying just 11 out of 75. In contrast, the model exceled in identifying class 1 (defective), achieving a precision of 69%, an impressive recall of 97%, and a high F1-score of 81%. It correctly classified 144 out of 149 defective parts, making it highly effective in defect detection.

The macro average F1-score of 52% and weighted average F1-score of 62% reflect the performance skew toward the defective class. The confusion matrix reinforces this: among 75 non-defective samples, 64 were misclassified, while 144 out of 149 defective samples were correctly identified. The ROC AUC score of 0.6793 suggests a reasonably good ability to distinguish between defective and non-defective parts, higher than what was observed in other models like Naïve Bayes and Decision Tree.

In summary, the SVM model offered a solid performance in identifying defects, with high recall and F1-score for the defective class. However, its low recall for non-defective items indicates a significant trade-off. The results obtained and shown in Table 5.8 for the SVM model generated the ROC curve shown in Figure 5.15.

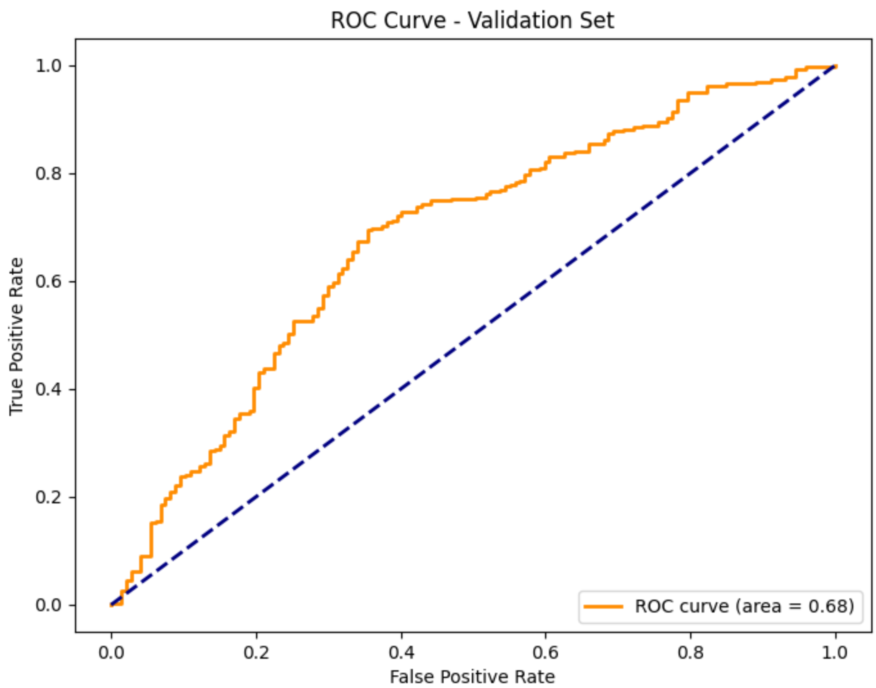

Figure 5.15 – SVM ROC Curve

5.7 Machine Learning Metrics Comparisons

A comprehensive evaluation across several machine learning models revealed varied performances in detecting AFP defects. The results, shown in Figure 5.16, identify that SVM and KNN models with k=7 and k=9 stand out with the highest overall accuracy of approximately 69%. While the accuracy metric alone suggests solid performance, a deeper dive into precision, recall, and F1-scores reveals important nuances.

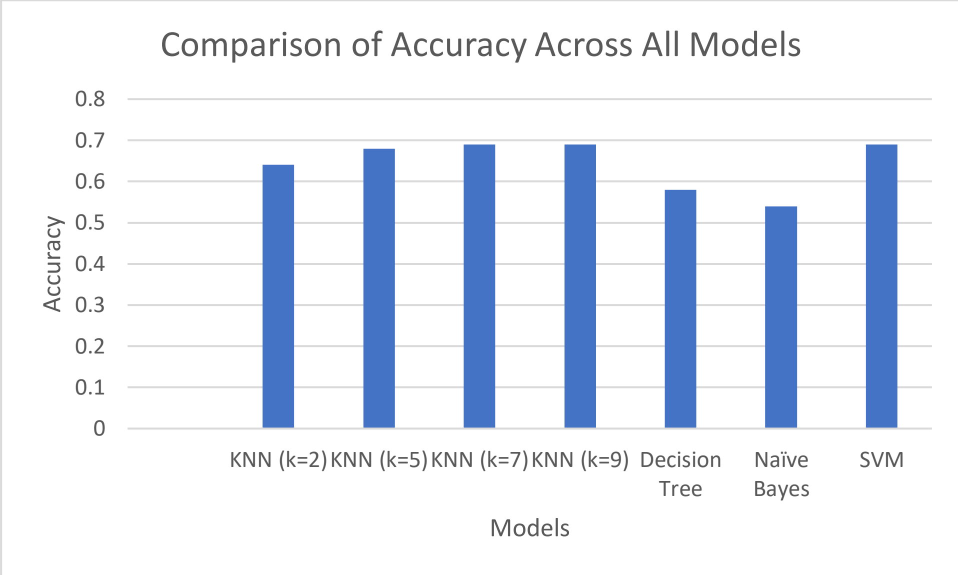

Figure 5.16 – ML Model Accuracy Evaluation

When it came to identifying defective parts (class 1), all models generally performed well, with F1-scores ranging from 0.62 (Naïve Bayes) to 0.81 (SVM, KNN k=7 and k=9). The recall for class 1 was particularly high in these top models, reaching 97–99%, indicating their reliability in catching most defects. This is critical in manufacturing, where missing a defect can be far more costly than falsely flagging a non-defective part.

However, performance on non-defective parts (class 0) was consistently weak across all models. The best F1-score for this class was just 0.41 from the Naïve Bayes model, and most other models performed below 0.30, with KNN (k=9) and SVM struggling to recall more than 15% of non-defective instances. This highlights a strong model bias toward predicting defects, likely due to class imbalance.

The weighted F1-scores, which account for the imbalance in class distribution, give more evenly distributed results. SVM and KNN (k=7) provided the best scores of around 0.62, while Naïve Bayes and Decision Tree achieved scores of 0.55–0.58.

The ROC AUC metric, which evaluates the model's ability to distinguish between the two classes, places SVM as the top performer with a score of 0.679, suggesting it has the best trade-off between sensitivity and specificity. The Naïve Bayes and KNN with k=9 models follow with scores of 0.5728 and 0.5348 while the lowest scoring models, sitting just above chance level, were the KNN models with k=2 and k=7, scoring 0.5019 and 0.5050. A graph showing a comparison of different models' ROC AUC metric results is shown in Figure 5.17.

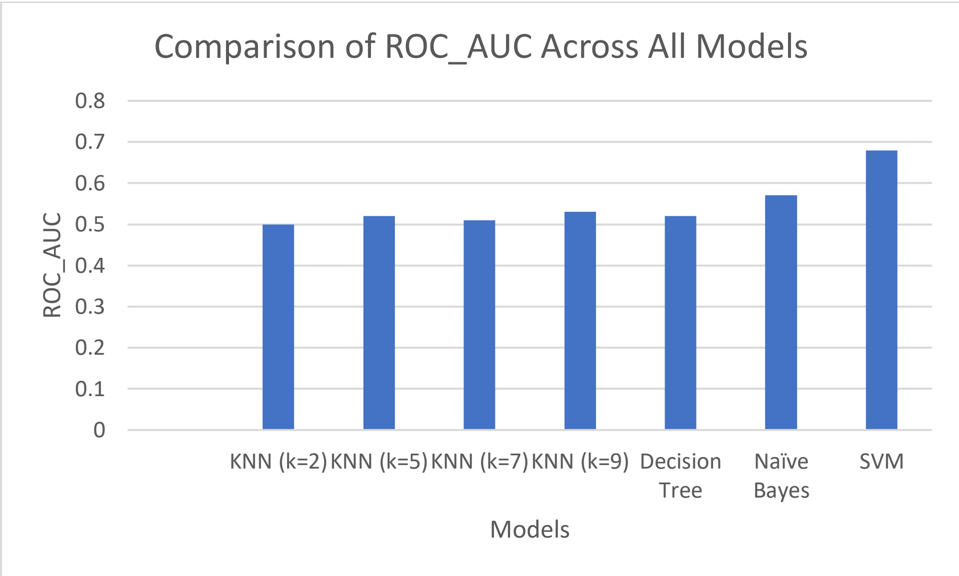

Figure 5.17 – ML Model ROC AUC Evaluation

In conclusion, while several models exhibited strong recall for detecting defects, their ability to correctly identify non-defective instances was limited. SVM emerged as the most balanced and discriminative model overall, but further modelling may be required to improve performance on non-defective predictions and ensure a more practical deployment of the model.

5.8 Deep Learning Data Analysis

In addition to the standard machine learning models, two versions of the ResNet deep learning model were also evaluated. ResNet, which stands for Residual Networks, is a convolutional network with skip residual connections. The residual path can be seen in Figure 5.18 bypassing the convolutional layers.

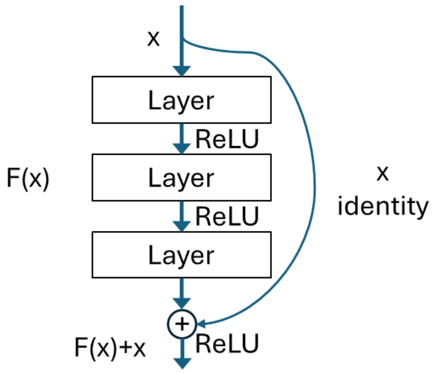

Figure 5.18 – ResNet Residual Path

The layers and residual paths are combined using the ReLU (Rectified Linear Unit) function.

5.8.1 ResNet101

ResNet101 is a version of ResNet with 101 layers. The architecture of ResNet101 is shown in Figure 5.19.

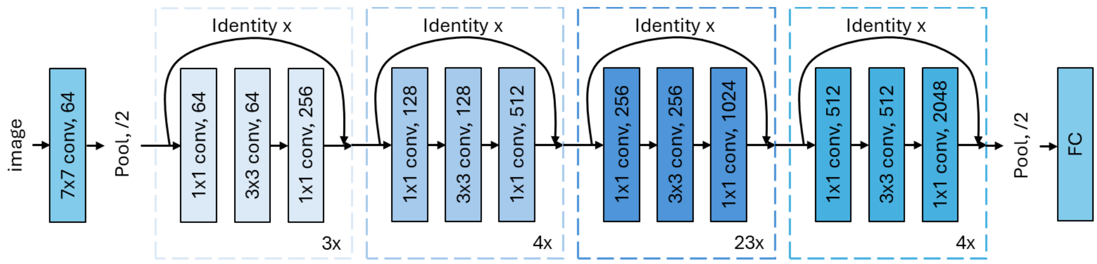

Figure 5.19 – ResNet101

There is an initial 7x7 convolutional layer, followed by 4 types of repeated blocks with a residual path and finally a fully connected layer. The ResNet101 model showed strong performance in identifying both defective and non-defective images, achieving a high accuracy of 78.57% and a test loss of 0.4892 after 37 training epochs. The results of the ResNet101 deep learning model can be found in Table 5.9.

Table 5.9 – ResNet101 Model Results

Model: ResNet101	Precision	Recall	F1-Score	Support
0 (no defect)	0.83	0.85	0.84	149
1 (defect)	0.69	0.65	0.67	75
Macro Average	0.76	0.75	0.76	224
Weighted Average	0.78	0.79	0.78	224
Confusion Matrix	49 (TP)	26 (FP)		
	22 (FN)	127 (TN)		
End Epoch	37/50			
Test Loss	0.4892			
ROC AUC	0.82			
Accuracy	0.7857			

For class 0 (non-defective), the model showed great performance with a precision of 83%, recall of 85%, and F1-score of 84%. This shows ResNet101 identified non-defective samples with both high precision and sensitivity. Performance on class 1 (defective) was also strong, with a precision of 69%, recall of 65%, and an F1-score of 67%. While lower than class 0, this still displays an ability to detect defective parts.

The macro average F1-score of 76% and weighted average F1-score of 78% reflect balanced, model behaviour across both classes, even with class imbalance. The confusion matrix shows that out of 149 non-defective samples, 127 were correctly identified, while 22 were misclassified. Among 75 defective instances, the model correctly identified 49 and misclassified 26.

Notably, the ROC AUC score of 0.82 indicates that ResNet101 had strong discriminative ability in distinguishing between defective and non-defective images.

5.8.2 ResNet 152

The ResNet152 model, trained for 36 epochs, demonstrated solid performance in detecting defects, achieving a test accuracy of 75.00% with a test loss of 0.5081. The results of the ResNet152 deep learning model can be found in Table 5.10.

Table 5.10 – ResNet152 Model Results

Model: ResNet152	**Precision**	**Recall**	**F1-Score**	**Support**
0 (no defect)	0.80	0.83	0.82	149
1 (defect)	0.64	0.59	0.61	75
Macro Average	0.72	0.71	0.71	224
Weighted Average	0.75	0.75	0.75	224
Confusion Matrix	44 (TP)	31 (FP)		
	25 (FN)	124 (TN)		
End Epoch	36/50			
Test Loss	0.5081			
ROC AUC	0.80			
Accuracy	0.75			

For class 0 (non-defective), the model achieved a precision of 80%, recall of 83%, and an F1-score of 82%. This suggests strong capability in correctly identifying non-defective items, with only 31 out of 149 misclassified. For class 1 (defective), performance was not as good with a precision of 64%, recall of 59%, and F1-score of 61%. Of the 75 defective parts, 124 were correctly detected, while 25 were misclassified as non-defective.

The macro average F1-score was 71%, and the weighted average was 75%, indicating a decent overall balance considering the class imbalance. Notably, the model achieved a ROC AUC score of 0.80, which reflects strong discriminative ability, slightly below ResNet101's 0.82, but still a clear improvement over traditional machine learning models.

5.9 Model Accuracy

The analysis of the machine learning and deep learning models highlights significant variations in performance across different approaches. The KNN models, with varying values of k, showed fluctuating validation and test accuracy scores, generally within the range of 0.60 to 0.69. The highest validation accuracy among KNN models was observed at k=9 with 0.681, while the lowest was at k=2 with 0.606. Similarly, test accuracy remained relatively stable, peaking at 0.692 for k=7 and k=9. Despite decent overall accuracy, the KNN models exhibited a strong bias towards the majority class, as indicated by their recall values. Lower values of k generally resulted in higher variance, which explains some of the inconsistencies in performance.

Other traditional machine learning models, such as Naïve Bayes and Decision Trees, underperformed relative to KNN and SVM. Naïve Bayes achieved a validation accuracy of 0.572 and a test accuracy of 0.536, showing weak generalisation capability. Decision tree performed slightly better with a validation accuracy of 0.566 and a test accuracy of 0.580, indicating that it was learning some useful patterns but still struggling with generalisation. These models exhibited relatively balanced precision-recall scores but did not outperform more sophisticated methods, suggesting they might not be the best choices for this classification task.

SVM emerged as one of the top-performing traditional machine learning models, achieving a validation accuracy of 0.699 and a test accuracy of 0.692. This suggests that SVM is capable of capturing the underlying data patterns better than KNN, Naïve Bayes, and Decision Trees. The use of the radial basis function kernel likely aids in separating the classes effectively, which is reflected in the relatively high recall for the majority class.

Deep learning models, specifically ResNet101 and ResNet152, outperformed all traditional models. ResNet101 achieved a validation accuracy of 0.778 and a test accuracy of 0.786, while ResNet152 recorded a validation accuracy of 0.742 and a test accuracy of 0.750. These results indicate the superior capability of deep learning models in handling complex data structures. The early stopping mechanism suggests that both models were trained optimally to avoid overfitting. Furthermore, the classification reports show that deep learning models maintained a more balanced performance across different classes compared to traditional approaches.

5.10 Machine Learning Discussion and Summary

A review of both machine learning and deep learning models shows that ResNet101 stands out as the top performing model for detecting defects. It achieved the highest accuracy at 78.6%, along with a great F1-score of 0.84 for non-defective parts and a strong F1-score of 0.67 for defective parts. Also, its ROC AUC score of 0.82 indicates a high ability to distinguish between defective and non-defective samples, making it the most well-rounded and suitable model among all those tested.

ResNet152 was the next best preforming model with an accuracy of 75% and good metrics for the non-defective class (F1-score of 0.82). However, its performance in detecting defective parts was slightly lower (F1-score of 0.61) compared to ResNet101, and it fell short on overall discrimination with a ROC AUC score of 0.80. Despite this, ResNet152 remains a highly viable option.

Among the classical models, SVM demonstrated the most competitive results, achieving an accuracy of 69.2% and a high F1-score of 0.81 for the defective class. While it had problems with non-defective predictions (F1-score of 0.24), its ROC

AUC value of 0.679 was the best among non-deep learning models, displaying a strong bias toward identifying defects, a useful trait in situations where false negatives are more problematic than false positives.

In contrast, other models such as Naïve Bayes and Decision Tree showed weaker overall performance, with F1-scores and ROC AUC values that showed limited predictive reliability. The KNN models, particularly with k=7 and k=9, showed high recall and F1-scores for the defective class (around 0.81) but suffered from extremely low recall and F1-score for non-defective instances, causing them to be imbalanced and less useful for practical use.

In conclusion, deep learning models, particularly ResNet101, clearly surpassed traditional approaches in both accuracy and balance. For environments where catching defects is critical while minimising false alarms, ResNet101 provides the best choice going forward.

Chapter 6

6 Integration Framework for a Defect Detection System

This chapter presents the integration framework that was proposed and developed to bringing each aspect of the PhD research together into one coherent system. It looks at the real-time performance of the complete AFP defect detection system by combining dual robot control with machine learning and thermography to identify defects using the coordination, processing and communication capabilities of the SSB.

6.1 Introduction

Defect detection is a critical aspect of automated manufacturing, particularly in high-precision industries, where composite materials are extensively used. The implementation of collaborative robots, the SSB designed for this research, and machine learning techniques can help to improve defect detection in composite material manufacturing. These technologies enable real-time quality control which reduces material waste and ensures product reliability. This chapter explores the integration of these technologies into a cohesive defect detection system, implementing machine learning driven automation and thermal imaging-based defect detection in the AFP process.

The traditional approach to defect detection in AFP can rely heavily on post-process inspection methods, including visual inspection, ultrasonic testing, and profilometry. While these techniques are effective, they can be time-consuming and can result in the late identification of defects, requiring costly rework or material disposal. The primary limitation of these traditional methods is their inability to provide real-time defect detection, leading to increased production times and decreased efficiency. In contrast, the real-time defect detection system developed within this PhD research integrated machine learning models, thermal imaging, and cobot based inspections, allowing for continuous in-process monitoring of composite material layups.

A key innovation of this research was the introduction of a multi-robot collaborative system. In this system an industrial AFP robot layed down composite tows while a UR10e cobot equipped with a thermal imaging camera followed in real-time, capturing temperature variation images that indicated potential defects. The SSB acted as the central processing unit, integrating sensor data, running machine learning models, and communicating real-time defect reports to the control system. The SSB’s modular architecture allowed it to interface seamlessly with both the primary AFP robot and the cobot, facilitating flexible and adaptable manufacturing processes.

The machine learning models utilised in this research were trained to classify common AFP defects such as gaps and overlaps using thermal image datasets. Machine learning algorithms were employed to enhance the accuracy and robustness of defect detection. The integration of thermal imaging with machine learning driven analysis presented several benefits. Unlike visual inspection, which is limited by lighting conditions and human error, thermal imaging detected defects based on temperature variations. When composite material was improperly laid, or contained defects, the heat distribution pattern was altered, making it easily identifiable using infrared thermography. By combining this with machine learning based classification, the system could automatically flag defects to minimise material and time wastage.

The cobot’s role in this system was essential, as it provided real-time mobility and adaptive positioning to ensure optimal defect detection. Unlike traditional fixed thermal imaging setups, which can only capture defects from a single angle, the UR10e cobot dynamically adjusted its position based on real-time feedback from the SSB. This allowed for multiple-angle defect analysis, improving the overall accuracy of defect classification. Moreover, the collaborative nature of the cobot ensured that it could work safely alongside human operators, adding an extra layer of flexibility to the manufacturing process.

One of the key challenges addressed in this research was the synchronisation of the AFP robot, cobot, and machine learning processing. Unlike conventional automated manufacturing systems, which operate in sequential steps, this research developed a fully integrated approach where data acquisition, processing, and corrective action occurred simultaneously. The SSB played a pivotal role in managing this integration by coordinating sensor inputs, processing machine learning models, and interfacing with the robot control system. This ensured that defects were detected and rectified before the composite layup process was completed, significantly reducing rework costs and improving overall production efficiency.

Furthermore, this research emphasised the importance of scalability and adaptability in defect detection. The SSB based system was designed to be modular and reconfigurable, meaning it could be easily adapted for different robotic platforms, various types of composite materials, and different manufacturing environments. This made it an ideal solution for industries that require flexible, automated defect detection systems capable of handling a wide range of defect types.

6.2 System Integration

The successful implementation of a real-time defect detection system in AFP required a seamless integration of multiple components, including the SSB, robots, machine learning algorithms, and thermal imaging systems. This integration ensured a fully automated, adaptable, and efficient defect detection system capable of real-time monitoring and corrective action during the AFP process.

6.2.1 System Architecture Overview

The defect detection system comprised of three main subsystems: the AFP System, which served as the primary robotic system responsible for laying composite material; the cobot based inspection system, featuring a UR10e cobot equipped with a thermal imaging camera for real-time defect detection; and the SSB with machine learning processing, which acted as a central processing hub that integrated real-time sensor data, executed machine learning models and coordinated actions between the AFP system and the cobot. The combined multi-robot and SSB system with connections is shown in Figure 6.1. The SSB can be seen attached to the cobot near the base.

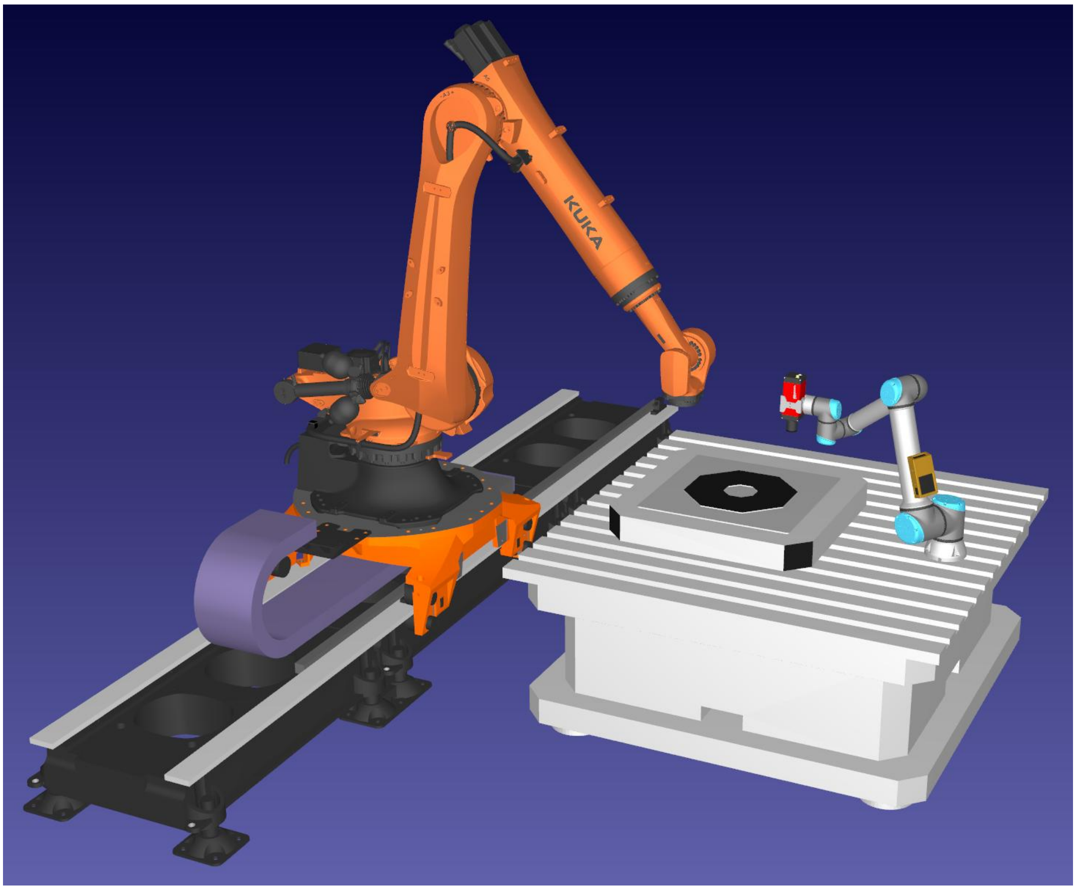

Figure 6.1 – AFP Work Cell

The integration of these components ensured that defect detection occurred in real-time rather than through post-process inspections, significantly improving efficiency, reducing material waste, and minimising rework.

6.2.2 Hardware Integration

The hardware elements of the system worked in unison to achieve highly precise and automated defect detection. The AFP robot served as the primary robotic system responsible for laying down composite tows, while the UR10e cobot moved in synchronisation with it, scanning the composite layup using a thermal imaging camera mounted on the cobot to capture heat distribution patterns across the material. At the core of the system was the SSB, a modular processing unit that collected thermal image data, applied machine learning models for defect detection, and facilitated communication between the AFP system and the cobot, ensuring seamless integration and real-time defect identification.

6.2.3 Software Integration

The ResNet101 deep learning model was ported from the development environment on a Windows laptop PC to the SSB4 based Raspberry Pi 5. To compare the performance with the developed deep learning models, the test data set was evaluated on both the SSB4 and laptop systems. On the Windows laptop PC the test was executed with and without the GPU enabled, as the deep learning function can use the GPU for increased performance. Operating without the GPU showed just the CPU performance on the laptop.

The same ResNet101 deep learning Python program was ported and timed on the Raspberry Pi 5. As was done on the laptop, a run of 50 epochs were set. Training ended early after 29 epochs, lasting just under 10 hours.

Next the trained model was ported to the Hailo-8 accelerator using the Hailo tools. The Hailo development tools only run under the Linux Ubuntu operating system. A separate development system is not needed because the Hailo tools can run on Ubuntu in the WSL (Windows Subsystem for Linux) on a standard Windows computer. The process steps carried out were:

- Convert the trained PyTorch model to the open standard ONNX model file format
- Check if the exported model is compatible with the Hailo tools
- Compile the mode model using the Hailo compiler
- Copy the Hailo-8 specific '.hef' file to the Raspberry Pi file system
- Use the model on the Hailo-8 to run the classification

The Hailo tool set was unable to run the compiled model with the provided test data as there were problems with the Hailo tool set. Compiled models could be benchmarked with performance showing a single pass of the model took 39ms to execute.

The comparison of classification time for the system with different configurations is shown in Table 6.1.

Table 6.1 – ResNet101 Timing Results

System	Average Image Classification (ms)	Correctly Identified	Accuracy
Windows Laptop without GPU	107.8	166/224	74.11%
Windows Laptop with GPU	26.3	167/224	74.55%
SSB4	176.0	171/224	76.34%

The results show that the SSB4 was more than fast enough to process the thermal camera images in real-time given the speed of the UR10e cobot and time required to capture a new image. The variance in accuracy across the three systems are likely due to variations in how each platform handles numerical precision and quantisation. Since different hardware architectures implement floating-point and integer operations differently, small rounding errors and representation limits can accumulate during inference, leading to slight but measurable accuracy variance.

6.3 Testing and Results

The testing phase for the defect detection system focussed on integrating the cobot and the SSB with the machine learning framework. Testing was conducted in a structured manner, progressing from simulations to real-world execution. The objective of the testing process was to evaluate the accuracy, efficiency, and reliability of the integrated system in detecting defects in AFP composite manufacturing. The key focus areas included verifying the synchronisation of robotic components, assessing the performance of machine learning models, and ensuring seamless communication between the cobot, the AFP system, and the SSB processing unit.

6.4 Discussion and Summary

The testing phase successfully validated the effectiveness of the integrated cobot based defect detection system using the SSB with machine learning framework. The combined real-time machine learning processing allowed for accurate and effective defect identification. While some limitations remain, the overall results indicate that this system represents a significant advancement in automated defect detection for composite materials.

Chapter 7

7 Conclusion

The purpose of this thesis was to design, develop and evaluate a real-time defect detection system for AFP composite material creation using thermal imaging and machine learning techniques. The effectiveness of the proposed system, utilising a ResNet101 deep learning model, was proven by demonstrating its ability to detect defects in real-time. Experimental results showed that the system achieved high accuracy in defect detection, which can enhance the quality and performance of the final product while minimising time, cost and material wastage.

7.1 Contributions to Knowledge

Through the research carried out in this thesis, significant developments have been made in the progression of the knowledge about and understanding of AFP defect detection. Specific knowledge has been gained through the analysis of available literature relating to the use of thermal imaging with machine learning to detect defects. Particular contribution has been made to the field of AFP defect detection through the development of a real-time, effective, low-cost solution.

7.1.1 Review and Analysis of Current Literature

A review of existing literature identified current solutions for AFP defect detection and identified areas that were rarely considered or completely lacking in research.

Research indicated the importance of composite materials in various industries primarily due to its strength and lightweight nature. Analysis of automation in production lines identified several advantages with the nervousness around human replacement being the main disadvantage noted. Investigation into AFP defect detection systems showed that thermal imaging had been utilised as a standalone method of defect identification but required intense human interaction. machine learning has been heavily researched and shown to be effective but relies on visual systems and laser profilometry which are not fully successful at recognising defects below the top layer tow. Some expensive AFP defect detection systems have been developed combining machine learning and thermal imaging, but this is an area requiring more research.

The review of available literature successfully identified a key area of AFP defect detection that is not yet fully realised in the form of real-time defect detection. Several proposals offer automated solutions using machine learning and a form of thermal identification but none offer real-time detection which is essential to ensure defects are detected immediately rather than being identified after the complete component has been fabricated.

7.1.2 SSB Development to Enable Real-Time Defect Detection

To enable the real-time defect detection requirement identified through the literature review, the SSB was designed and developed. The SSB was used as proof of concept that a real-time, flexible AFP defect detection system was feasible.

The SSB was designed to be an intelligent smart sensor hub that is easily repositioned enabling the addition of sensors as required at any stage of a manufacturing line. The SSB enables simple sensors to be used in a system providing the functionality of smart connected sensors. The SSB offers a solution for real-time AFP defect detection by providing the capability required to control a cobot and use machine learning to analyse thermal images.

7.1.3 The Development of a Multi-Robot AFP Defect Detection System

Following literature review analysis, a multi-robot system using machine learning and thermography was developed as a new method for AFP defect detection. By integrating a KUKA KR120 industrial robot and UR10e collaborative robot, a highly effective solution for AFP defect detection in composite manufacturing was developed. The thermal imaging system in the UR10e makes rapid identification of defects possible which improves process reliability and ensures product quality.

Testing proved that deep learning models outperformed traditional machine learning models in both accuracy and balance. After an intensive evaluation and analysis of multiple machine learning and deep learning models, the ResNet101 was identified as the optimal model tested for AFP defect detection.

7.2 Limitations and Future Work

While the proposed system demonstrates promising results, several limitations must be acknowledged:

- **Small Dataset** – the dataset used in this study consists of 2214 segmented thermal images. Although sufficient for initial analysis, the small dataset limits the generalisability of the machine learning models. Scaling up the dataset with a larger and more diverse range of defects and material will enhance the robustness and performance of the models.
- **Defect Types** – this study focuses primarily on gap and overlap defects. While these are common in AFP processes, other defect types such as wrinkles, voids, twists, or foreign objects were not extensively evaluated. Including these defect types in future work would provide a more comprehensive assessment of the system's capabilities.
- **Environmental Conditions** – variations in ambient temperature, lighting conditions, and thermal camera calibration could impact the accuracy of defect detection. The experimental setup used controlled conditions, but testing in industrial environments with fluctuating parameters may require additional adjustments.

Addressing the limitations encountered in this thesis would offer several possible areas for future research and development and would provide a stronger foundation for deploying thermal imaging and machine learning systems in real-world AFP manufacturing environments. In addition, interest expressed by industry into the SSB opens up the potential for a future project or related industrial automation research.

Bibliography

[1] P. Thori, P. Sharma, and M. Bhargava, ‘An Approach Of Composite Materials In Industrial Machinery: Advantages, Disadvantages And Applications’, *Int. J. Res. Eng. Technol.*, vol. 02, pp. 350–355, 2013.

[2] R. R. Nagavally, ‘Composite Materials-History , Types , Fabrication Techniques , Advantages , And Applications’, 2016. [Online]. Available: https://api.semanticscholar.org/CorpusID:222334256

[3] Lucintel, ‘Market value of composite materials worldwide in 2019 and 2024 with a forecast for 2031’, Statista. Accessed: Feb. 23, 2025. [Online]. Available: https://www.statista.com/statistics/944471/global-market-value-of-composites/#statisticContainer

[4] Lucintel, ‘Market value of carbon fiber worldwide in 2016 and 2021, with a forecast for 2027’, Statista. [Online]. Available: https://www.statista.com/statistics/947941/global-market-value-carbon-fiber/

[5] E. Asmatulu, A. Alonayni, and M. Alamir, ‘Safety concerns in composite manufacturing and machining’, Mar. 2018, p. 68. doi: 10.1117/12.2296707.

[6] A. Gamage, ‘Research Design, Philosophy, and Quantitative Approaches in Scientific Research Methodology’, *Sch. J. Eng. Technol.*, vol. 13, pp. 91–103, Feb. 2025, doi: 10.36347/sjet.2025.v13i02.004.

[7] H. Taherdoost, ‘What are Different Research Approaches? Comprehensive Review of Qualitative, Quantitative, and Mixed Method Research, Their Applications, Types, and Limitations’, *J. Manag. Sci. Eng. Res.*, Aug. 2022.

[8] R. Attalla, F. N., and A. Gelany, ‘The Technology of Mud-bricks from the Sacred Architecture in Ancient Egypt to the Green House Revolution (Case Study of Luxor City, Egypt)’, vol. 29, pp. 1264–1277, Jan. 2020.

[9] V. Serifi, M. Tarić, D. Jevtic, A. Ristovski, and M. Šahinagić-Isović, ‘Historical Development Of Composite Materials’, *Ann. Univ. Oradea Econ. Sci.*, vol. Vol. XXVII (XVII), Nov. 2018, doi: 10.15660/AUOFMTE.2018-3.3392.

[10] L. M. Seymour, J. Maragh, P. Sabatini, M. D. Tommaso, J. C. Weaver, and A. Masic, ‘Hot mixing: Mechanistic insights into the durability of ancient Roman concrete’, *Sci. Adv.*, vol. 9, no. 1, p. eadd1602, 2023, doi: 10.1126/sciadv.add1602.

[11] ‘Owens Corning Corporation International Directory of Company Histories’, *Encyclopedia.com*. [Online]. Available: https://www.encyclopedia.com/books/politics-and-business-magazines/owens-corning-corporation

[12] D. Edgerton, ‘technologies of war and peace’, *Imperial Magazine*, no. 36, pp. 20–22, spring/summer 2011.

[13] A. T. Bhatt, P. P. Gohil, and V. Chaudhary, ‘Primary Manufacturing Processes for Fiber Reinforced Composites: History, Development & Future Research Trends’, *IOP Conf Ser Mater Sci Eng 330 012107*, 2018.

[14] J. Bošnjak Hordov, N. Bodrožić Ćoko, M. Jurčević, I. Primorac, and S. Nižetić, ‘Investigation of thermal properties of pykrete’, *IOP Conf. Ser. Earth Environ. Sci.*, vol. 1196, p. 012055, Jun. 2023, doi: 10.1088/1755-1315/1196/1/012055.

[15] ‘Composite material | Construction, Strength, Durability | Britannica’. Accessed: Jul. 06, 2025. [Online]. Available: https://www.britannica.com/technology/composite-material

[16] ‘ACMA-Guidelines-and-Recommended-Practices-FRP-Architectural-Products’. American Composites Manufacturers Association, 2021. [Online]. Available: https://compositesuk.co.uk/wp-content/uploads/2021/11/ACMA-Guidelines-and-Recommended-Practices-FRP-Architectural-Products.pdf

[17] M. Egbo, ‘A fundamental review on composite materials and some of their applications in biomedical engineering’, *J. King Saud Univ. - Eng. Sci.*, vol. 33, Jul. 2020, doi: 10.1016/j.jksues.2020.07.007.

[18] H. Farhat, ‘Chapter 3 - Materials and coating technologies’, in *Operation, Maintenance, and Repair of Land-Based Gas Turbines*, H. Farhat, Ed., Elsevier, 2021, pp. 63–87. doi: 10.1016/B978-0-12-821834-1.00007-1.

[19] D. K. Rajak, D. D. Pagar, R. Kumar, and C. I. Pruncu, ‘Recent progress of reinforcement materials: a comprehensive overview of composite materials’, *J. Mater. Res. Technol.*, vol. 8, no. 6, pp. 6354–6374, Nov. 2019, doi: 10.1016/j.jmrt.2019.09.068.

[20] A. Björnsson, M. Jonsson, and K. Johansen, ‘Automation of Composite Manufacturing Using Off-the-shelf Solutions, Three Cases from the Aerospace Industry’, presented at the ICCM20 - The 20th International Conference on Composite Materials, 19-24th July 2015, Copenhagen Denmark, 2015. Accessed: Jun. 21, 2025. [Online]. Available: https://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-120300

[21] Z. Quan *et al.*, ‘Additive manufacturing of multi-directional preforms for composites: Opportunities and challenges’, *Mater. Today*, vol. 255, May 2015, doi: 10.1016/j.mattod.2015.05.001.

[22] R. Protz *et al.*, ‘Manufacture, process simulation, modelling and testing of thick-walled thermoset fibre-polymer composite laminates — A review’, *Compos. Struct.*, vol. 373, p. 119678, Dec. 2025, doi: 10.1016/j.compstruct.2025.119678.

[23] J. Galos, ‘Thin-ply composite laminates: a review’, *Compos. Struct.*, vol. 236, p. 111920, Mar. 2020, doi: 10.1016/j.compstruct.2020.111920.

[24] . P. T., ‘AN APPROACH OF COMPOSITE MATERIALS IN INDUSTRIAL MACHINERY: ADVANTAGES, DISADVANTAGES AND APPLICATIONS’, *Int. J. Res. Eng. Technol.*, vol. 02, no. 12, pp. 350–355, Dec. 2013, doi: 10.15623/ijret.2013.0212060.

[25] K. L. Pickering, M. G. A. Efendy, and T. M. Le, ‘A review of recent developments in natural fibre composites and their mechanical performance’,

Compos. Part Appl. Sci. Manuf., vol. 83, pp. 98–112, Apr. 2016, doi: 10.1016/j.compositesa.2015.08.038.

[26] M. Rajyaguru, 'PROS AND CONS OF CARBON FIBER', Jul. 2019.

[27] A. Gasparetto and L. Scalera, 'From the Unimate to the Delta Robot: The Early Decades of Industrial Robotics: Proceedings of the 2018 HMM IFToMM Symposium on History of Machines and Mechanisms', in *History of Mechanism and Machine Science*, 2019, pp. 284–295. doi: 10.1007/978-3-030-03538-9_23.

[28] B. Singh and N. Sellappan, 'Evolution of Industrial Robots and their Applications', p. 763, May 2008.

[29] *Robotics — Vocabulary*, ISO 8373, Nov. 2021.

[30] J. Edward, W. Wannasuphoprasit, and M. Peshkin, 'Cobots: Robots For Collaboration With Human Operators', Mar. 1999.

[31] C. Parschau and J. Hauge, 'Is automation stealing manufacturing jobs? Evidence from South Africa's apparel industry', *Geoforum*, vol. 115, pp. 120–131, Oct. 2020, doi: 10.1016/j.geoforum.2020.07.002.

[32] J. Frohm, V. Lindström, M. Winroth, and J. Stahre, 'THE INDUSTRY'S VIEW ON AUTOMATION IN MANUFACTURING', *IFAC Proc. Vol.*, vol. 39, no. 4, pp. 453–458, Jan. 2006, doi: 10.3182/20060522-3-FR-2904.00073.

[33] J. Frohm, V. Lindström, M. Winroth, and J. Stahre, 'Levels of Automation in Manufacturing', *Ergon. - Int. J. Ergon. Hum. Factors*, vol. 30, pp. 181–207, Jan. 2008.

[34] A. Asquith and G. Horsman, 'Let the robots do it! – Taking a look at Robotic Process Automation and its potential application in digital forensics', *Forensic Sci. Int. Rep.*, vol. 1, p. 100007, Nov. 2019, doi: 10.1016/j.fsir.2019.100007.

[35] 'Global Process Automation Industry: Growth, Trends and Forecast (2020-2025) - ResearchAndMarkets.com'. Accessed: Jul. 13, 2025. [Online]. Available: https://www.businesswire.com/news/home/20200622005325/en/Global-Process-Automation-Industry-Growth-Trends-and-Forecast-2020-2025---ResearchAndMarkets.com

[36] 'Composite materials global market value 2031', Statista. Accessed: Jul. 19, 2025. [Online]. Available: https://www.statista.com/statistics/944471/global-market-value-of-composites/

[37] A. Levy, J. Tierney, D. Heider, P. Lefébure, and D. Lang, *Simulation And Optimization Of The Thermoplastic Automated Tape Placement (ATP) Process*. 2012.

[38] Department of Management and Engineering, Manufacturing Engineering, Linköping University and A. Björnsson, *Enabling Automation of Composite Manufacturing through the Use of Off-The-Shelf Solutions*. Linköping University Electronic Press, 2014. doi: 10.3384/lic.diva-111936.

[39] N. Izzaty, H. Y. Sastra, and Ilyas, 'The Implementation of Graphene Composites for Automotive: An Industrial Perspective', *IOP Conf. Ser. Mater.*

Sci. Eng., vol. 536, no. 1, p. 012133, Jun. 2019, doi: 10.1088/1757-899X/536/1/012133.

[40] A. Elmarakbi and W. Azoti, *NOVEL COMPOSITE MATERIALS FOR AUTOMOTIVE APPLICATIONS: CONCEPTS AND CHALLENGES FOR ENERGY-EFFICIENT AND SAFE VEHICLES*. 2015.

[41] G. V. Savage, 'Composite Materials Technology in Formula 1 Motor Racing', 2008. [Online]. Available: https://api.semanticscholar.org/CorpusID:17757994

[42] 'Carbon fibre's journey from racetrack to hatchback', *BBC News*, Mar. 10, 2011. Accessed: Jun. 21, 2025. [Online]. Available: https://www.bbc.com/news/technology-12691062

[43] 'BMW i3 - First Mass Produced Car Made With Composites', CompositesLab. Accessed: Jun. 21, 2025. [Online]. Available: https://compositeslab.com/bmw-i3-first-mass-produced-car-made-largely-with-composites/index.html

[44] 'Carbon fiber global market value 2027', Statista. Accessed: Jul. 19, 2025. [Online]. Available: https://www.statista.com/statistics/947941/global-market-value-carbon-fiber/

[45] V. Chauhan, J. Varis, and T. Kärki, 'The Potential of Reusing Technical Plastics', *Procedia Manuf.*, vol. 39, pp. 502–508, Jan. 2019, doi: 10.1016/j.promfg.2020.01.407.

[46] 'Audi Smart Factory: How Drones, AR And VR Gears Are Building Cars', TechTheLead. Accessed: Jun. 21, 2025. [Online]. Available: https://techthelead.com/audi-smart-factory-how-drones-ar-and-vr-gears-help-build-cars/

[47] 'Voith demonstrates Industry 4.0 carbon fibre manufacture with Audi part'. Accessed: Jun. 21, 2025. [Online]. Available: https://www.insidecomposites.com/voith-demonstrates-industry-40-carbon-fibre-manufacture-with-audi-part/

[48] A. Quilter, 'Composites in Aerospace Applications'.

[49] G. P. Thomas, 'Composites Used in the Aerospace Industry', AZoM. Accessed: Jun. 21, 2025. [Online]. Available: https://www.azom.com/article.aspx?ArticleID=8152

[50] D. F. Smith, 'The Use of composites in aerospace: Past, present and future challenges.'.

[51] S. PÂNZARU, 'THE MANAGEMENT OF MODERN TECHNOLOGIES USED IN THE AERONAUTICAL INDUSTRY.', *Rev. Gen. Manag.*, vol. 26, no. 2, pp. 63–71, Jul. 2017.

[52] 'Civil Aerospace'. Accessed: Jun. 21, 2025. [Online]. Available: https://www.rolls-royce.com/products-and-services/civil-aerospace.aspx

[53] ‘Rolls-Royce opens new assembly line for Turboshaft engines’. Accessed: Jun. 21, 2025. [Online]. Available: https://www.rolls-royce.com/media/press-releases-archive/yr-2008/rr-opens-new-assembly.aspx

[54] J. Shury, ‘Rolls Royce Opens New Facility in Bristol’, Composites Today. Accessed: Jun. 21, 2025. [Online]. Available: https://www.compositestoday.com/2020/01/rolls-royce-opens-new-facility-in-bristol/

[55] G. Neşer, ‘Polymer Based Composites in Marine Use: History and Future Trends’, *Procedia Eng.*, vol. 194, pp. 19–24, Jan. 2017, doi: 10.1016/j.proeng.2017.08.111.

[56] A. Caramatescu and C. Mocanu, ‘Review of composite materials applications in marine industry’, vol. 42, pp. 169–174, Dec. 2019, doi: 10.35219/AnnUGalShipBuilding.2019.42.23.

[57] F. Rubino, A. Nisticò, F. Tucci, and P. Carlone, ‘Marine Application of Fiber Reinforced Composites: A Review’, *J. Mar. Sci. Eng.*, vol. 8, no. 1, Art. no. 1, Jan. 2020, doi: 10.3390/jmse8010026.

[58] D.-A. Zisimopoulos, ‘Use of Fiber Reinforced Plastics in Ship Construction: A Study of SOLAS regulation II-2/17 on Alternative Design and Arrangementsfor Fire Safety’, May 10, 2016, *National Technological University of Athens*. doi: 10.26240/HEAL.NTUA.11991.

[59] G. Marsh, ‘50 Years of reinforced plastic boats’, *Reinf. Plast.*, vol. 50, pp. 16–19, Oct. 2006, doi: 10.1016/S0034-3617(06)71125-0.

[60] ‘Automated Ply Placement - Airborne’. Accessed: Jun. 22, 2025. [Online]. Available: https://www.airborne.com/automation/laminates-preforms/, https://www.airborne.com/automation/laminates-preforms/

[61] ‘RAMSSES project hits key milestone for green shipbuilding | Damen’. Accessed: Jun. 22, 2025. [Online]. Available: https://www.damen.com/insights-center/articles/ramsses-project-reaches-milestone-towards-sustainable-composite-ship

[62] P. KOSTOV, ‘RAMSSES’, Innovation and Networks Executive Agency - European Commission. Accessed: Jun. 22, 2025. [Online]. Available: https://wayback.archive-it.org/12090/20190615075620/https://ec.europa.eu/inea/en/horizon-2020/projects/h2020-transport/waterborne/ramsses

[63] ‘Ramsses Project: Ramsses Project’. Accessed: Jun. 22, 2025. [Online]. Available: https://www.ramsses-project.eu/

[64] A. Björnsson, J.-E. Lindback, and K. Johansen, ‘Automated Removal of Prepreg Backing Paper - A Sticky Problem’, presented at the SAE 2013 AeroTech Congress & Exhibition, SAE International, Sep. 2013. doi: 10.4271/2013-01-2289.

[65] A. Bjornsson, J.-E. Lindback, D. Eklund, and M. Jonsson, ‘Low-cost Automation for Prepreg Handling - Two Cases from the Aerospace Industry’,

SAE Int. J. Mater. Manuf., vol. 9, no. 1, pp. 68–74, Sep. 2015, doi: 10.4271/2015-01-2606.

[66] M. Elkington, C. Ward, and A. Sarkytbayev, ‘Automated composite draping: a review’, May 2017. Accessed: Jun. 22, 2025. [Online]. Available: https://www.semanticscholar.org/paper/Automated-composite-draping%3A-a-review-Elkington-Ward/f864e87072a58cbae2a01a0bff371c8f55ea7a7f

[67] L.-P. Ellekilde *et al.*, ‘Design of Automated Robotic System for Draping Prepreg Composite Fabrics’, *Robotica*, vol. 39, pp. 1–16, Apr. 2020, doi: 10.1017/S0263574720000193.

[68] A. Levy, J. Tierney, D. Heider, P. Lefébure, and D. Lang, *Simulation And Optimization Of The Thermoplastic Automated Tape Placement (ATP) Process*. 2012.

[69] P. D. Juarez and E. D. Gregory, ‘In Situ Thermal Inspection of Automated Fiber Placement for manufacturing induced defects’, *Compos. Part B Eng.*, vol. 220, p. 109002, Sep. 2021, doi: 10.1016/j.compositesb.2021.109002.

[70] K. Wu, B. Stewart, and R. Martin, ‘ISAAC – A Testbed for Advanced Composites Research’, Sep. 2014.

[71] S. Sun, Z. Han, H. Fu, H. Jin, J. Dhupia, and Y. Wang, ‘Defect Characteristics and Online Detection Techniques During Manufacturing of FRPs Using Automated Fiber Placement: A Review’, *Polymers*, vol. 12, p. 1337, Jun. 2020, doi: 10.3390/polym12061337.

[72] T. Hudson, F. Palmieri, T. Abbott, J. Seebo, and E. Burke, ‘Design of an Automated Ultrasonic Scanning System for In-situ Composite Cure Monitoring and Defect Detection’, Apr. 2019. doi: 10.33599/nasampe/s.19.1523.

[73] R. Smith, L. Nelson, M. Mienczakowski, and P. Wilcox, ‘Ultrasonic Analytic-Signal Responses From Polymer-Matrix Composite Laminates’, *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. PP, pp. 1–1, Nov. 2017, doi: 10.1109/TUFFC.2017.2774776.

[74] A. Szewieczek, W. Hillger, L. Bühling, and D. Ilse, ‘New Developments and applications for Air Coupled Ultrasonic Imaging Systems’, 2018. Accessed: May 23, 2026. [Online]. Available: https://www.semanticscholar.org/paper/New-Developments-and-applications-for-Air-Coupled-Szewieczek-Hillger/3b8eb1f321dd30a6f20a84185a037ceb63699704

[75] J. Cemenska, T. Rudberg, and M. Henscheid, ‘Automated In-Process Inspection System for AFP Machines’, *SAE Int. J. Aerosp.*, vol. 08, no. 2, pp. 303–309, Sep. 2015, doi: 10.4271/2015-01-2608.

[76] D. H. Nguyen, X. C. Sun, I. Tretiak, M. A. Valverde, and J. Kratz, ‘Automatic process control of an automated fibre placement machine’, *Compos. Part Appl. Sci. Manuf.*, vol. 168, no. 107465, May 2023, doi: 10.1016/j.compositesa.2023.107465.

[77] M. Mussatayev, Q. Yi, M. Fitzgerald, V. K. Maes, P. Wilcox, and R. Hughes, ‘Directional eddy current probe configuration for in-line detection of out-of-

plane wrinkles', *Compos. Part B Eng.*, vol. 268, p. 111048, Jan. 2024, doi: 10.1016/j.compositesb.2023.111048.

[78] S. Pantoji, C. Kassapoglou, and D. Peeters, 'Predicting gaps and overlaps in Automated Fiber Placement composites by measuring sources of manufacturing process variations', *Compos. Part B Eng.*, vol. 291, p. 111891, Feb. 2025, doi: 10.1016/j.compositesb.2024.111891.

[79] R. Harik, C. Saidy, S. Williams, Z. Gurdal, and B. Grimsley, *Automated fiber placement defect identity cards: cause, anticipation, existence, significance, and progression*. 2018.

[80] D. Del Rossi, V. Cadran, P. Thakur, M. Palardy-Sim, M. Lapalme, and L. Lessard, 'Experimental investigation of the effect of half gap/half overlap defects on the strength of composite structures fabricated using automated fibre placement (AFP)', *Compos. Part Appl. Sci. Manuf.*, vol. 150, p. 106610, Nov. 2021, doi: 10.1016/j.compositesa.2021.106610.

[81] T. Sobota, 'Fourier's Law of Heat Conduction', in *Encyclopedia of Thermal Stresses*, R. B. Hetnarski, Ed., Dordrecht: Springer Netherlands, 2014, pp. 1769–1778. doi: 10.1007/978-94-007-2739-7_384.

[82] F. Islam, M. J. Donough, E. Oromiehie, A. W. Phillips, N. A. St John, and B. G. Prusty, 'Modelling the effect of hot gas torch heating on adjacent tows during automated fibre placement consolidation of thermoplastic composites', *J. Thermoplast. Compos. Mater.*, vol. 36, no. 10, pp. 4211–4236, Oct. 2023, doi: 10.1177/08927057221123477.

[83] O. A. Tafreshi, Hoa ,Suong Van, Shadmehri ,Farjad, Hoang ,Duc Minh, and D. and Rosca, 'Heat transfer analysis of automated fiber placement of thermoplastic composites using a hot gas torch', *Adv. Manuf. Polym. Compos. Sci.*, vol. 5, no. 4, pp. 206–223, Oct. 2019, doi: 10.1080/20550340.2019.1686820.

[84] E. D. Gregory and P. D. Juarez, 'In-situ thermography of automated fiber placement parts', presented at the 44th Annual Review of Progress in Quantitative Nondestructive Evaluation, Volume 37, Apr. 2018, p. 060005. doi: 10.1063/1.5031551.

[85] C. Schmidt, B. Denkena, K. Völtzer, and T. Hocke, 'Thermal Image-based Monitoring for the Automated Fiber Placement Process', *Procedia CIRP*, vol. 62, pp. 27–32, Dec. 2017, doi: 10.1016/j.procir.2016.06.058.

[86] C. Schmidt, T. Hocke, and B. Denkena, 'Artificial intelligence for non-destructive testing of CFRP prepreg materials', *Prod. Eng.*, vol. 13, Jul. 2019, doi: 10.1007/s11740-019-00913-3.

[87] P. D. Juarez, K. E. Cramer, and J. P. Seebo, 'Advanced in In Situ Inspection of Automated Fiber Placement Systems', presented at the International Society for Optical Engineering, Bellingham, WA, United States, Apr. 2016. Accessed: Jul. 07, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20160009052

[88] C. Sacco, A. B. Radwan, T. Beatty, and R. Harik, 'Machine Learning Based AFP Inspection: A Tool for Characterization and Integration', presented at the

SAMPE, Diamond Bar, CA, United States, May 2019. Accessed: Jul. 07, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20200002684

[89] V. Daghigh, H. Daghigh, T. E. Lacy, and M. Naraghi, ‘Review of machine learning applications for defect detection in composite materials’, *Mach. Learn. Appl.*, vol. 18, p. 100600, Dec. 2024, doi: 10.1016/j.mlwa.2024.100600.

[90] C. Nelon, O. Myers, and A. Hall, ‘The intersection of damage evaluation of fiber-reinforced composite materials with machine learning: A review’, *J. Compos. Mater.*, vol. 56, no. 9, pp. 1417–1452, Apr. 2022, doi: 10.1177/00219983211037048.

[91] C. Sacco, ‘Machine Learning Methods for Rapid Inspection of Automated Fiber Placement Manufactured Composite Structures’, *Theses Diss.*, Jul. 2019, [Online]. Available: https://scholarcommons.sc.edu/etd/5475

[92] C. Li *et al.*, ‘Intelligent recognition of composite material damage based on deep learning and infrared testing’, *Opt. Express*, vol. 29, pp. 31739–31753, Sep. 2021, doi: 10.1364/OE.435230.

[93] P. Arrabiyeh, M. Bobe, M. Duhovic, M. Eckrich, A. Dlugaj, and D. May, ‘Implementing low budget machine vision to improve fiber alignment in wet fiber placement’, *J. Reinf. Plast. Compos.*, Sep. 2024, doi: 10.1177/07316844241278050.

[94] ‘SIMATIC S7-300 - Proven multiple times!’, siemens.com Global Website. Accessed: Jun. 22, 2025. [Online]. Available: https://www.siemens.com/global/en/products/automation/systems/industrial/plc/simatic-s7-300.html

[95] ‘6ES7312-5BF04-0AB0 | Siemens SIMATIC S7-300 Series PLC CPU for Use with SIMATIC S7-300 Series, Digital Output, 10 (Digital)-Input, Digital | RS’. Accessed: Jun. 22, 2025. [Online]. Available: https://uk.rs-online.com/web/p/plcs-programmable-logic-controllers/7877989?gb=s

[96] ‘6AG40122CA200BX0 | Siemens SIMATIC IPC347G, Industrial Computer, 350W, Intel Core i7 3.6 GHz, 8 MB, 4 Windows | RS’. Accessed: Jun. 22, 2025. [Online]. Available: https://uk.rs-online.com/web/p/industrial-computers/2519132?gb=s

[97] B. A. G. & C. K. Germany Hülshorstweg 20, 33415 Verl, ‘Beckhoff | New Automation Technology’, Beckhoff Automation. Accessed: Jun. 22, 2025. [Online]. Available: https://www.beckhoff.com/en-gb/

[98] ‘WiAutomation | Automation Spares Suppliers’. Accessed: Jun. 22, 2025. [Online]. Available: https://gb.wiautomation.com/search

[99] ‘NVIDIA GeForce RTX 40 Series Family’, NVIDIAGeForce. Accessed: Sep. 07, 2025. [Online]. Available: https://www.nvidia.com/en-gb/geforce/laptops/compare/

[100] ‘KiCad’. Accessed: May 03, 2026. [Online]. Available: https://www.kicad.org/

[101] R. P. Ltd, ‘Powered by Raspberry Pi product catalogue’, Raspberry Pi. Accessed: Apr. 19, 2025. [Online]. Available: https://www.raspberrypi.com/for-industry/powered-by/product-catalogue/?

[102] R. P. Ltd, ‘Powered by Raspberry Pi product catalogue PLC’, Raspberry Pi. Accessed: Apr. 19, 2025. [Online]. Available: https://www.raspberrypi.com/for-industry/powered-by/product-catalogue/?industry=Industrial+control%2Fapplications

[103] R. P. Ltd, ‘Raspberry Pi 5 Lifetime’, Raspberry Pi. Accessed: Apr. 19, 2025. [Online]. Available: https://www.raspberrypi.com/products/raspberry-pi-5/

[104] ‘Enterprise Open Source and Linux’, Ubuntu. Accessed: May 03, 2026. [Online]. Available: https://ubuntu.com/

[105] ‘Regulation - 2024/2847 - EN - EUR-Lex’. Accessed: May 09, 2025. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/2847/oj/eng

[106] ‘The Modbus Organization’. Accessed: Apr. 26, 2025. [Online]. Available: https://modbus.org/

[107] ‘KUKA KR 120 R3100-2’. Accessed: May 07, 2026. [Online]. Available: https://my.kuka.com/s/product/kr-120-r31002/01t58000005qcINAAY?language=en_US&tab=Details

[108] ‘Far infrared thermal sensor array (32x24 RES)’, Melexis. Accessed: May 03, 2026. [Online]. Available: https://www.melexis.com/en/product/MLX90640/far-infrared-thermal-sensor-array

[109] ‘OV5647-70 5MP OIS Camera - Waveshare Wiki’. Accessed: May 03, 2026. [Online]. Available: https://www.waveshare.com/wiki/OV5647-70_5MP_OIS_Camera?srsltid=AfmBOor8rzazmCUNZZ-aPjzuDFs5ruWXBA8vC5z-MijeGn4BBqoTcvkR

[110] ‘Simulator for robot arms and offline programming - RoboDK’. Accessed: May 03, 2026. [Online]. Available: https://robodk.com/

[111] ‘Welcome to Python.org’, Python.org. Accessed: May 03, 2026. [Online]. Available: https://www.python.org/

[112] ‘Simulator for robot arms and offline programming - RoboDK’. Accessed: Apr. 19, 2025. [Online]. Available: https://robodk.com/

[113] M. Dehmer and S. C. Basak, Eds, *Statistical and Machine Learning Approaches for Network Analysis*, 1st edn. Wiley, 2012. doi: 10.1002/9781118346990.